



Configuring Monitored Devices

ScienceLogic version 10.2.0

Table of Contents

Introduction	7
Additional Reading	9
Alibaba Cloud: Aliyun	10
Prerequisites	10
Amazon Web Services	11
Configuring AWS to Report Billing Metrics	12
Filtering EC2 Instances By Tag	13
Automatic SL1 Organization Creation	15
Monitoring Consolidated Billing Accounts	15
ScienceLogic Events and AWS Alarms	15
AMQP: RabbitMQ	17
Prerequisites for Monitoring RabbitMQ	17
Aruba Central	18
Prerequisites for Monitoring Aruba Central	18
Cisco: ACI	19
Prerequisites for Monitoring Cisco ACI	19
Recommended System Values	19
Cisco: AppDynamics	21
Prerequisites for Monitoring Cisco AppDynamics	21
Cisco: CloudCenter	22
Configuration and Discovery for Standard Cisco CloudCenter Deployments	22
Prerequisites for Monitoring Standard CloudCenter Deployments	23
Configuration and Discovery for High-Availability Cisco CloudCenter Deployments	23
Prerequisites for Monitoring High-Availability CloudCenter Deployments	23
Cisco: Contact Center Enterprise	25
Configuring Unified Contact Center Enterprise Monitoring Using SNMP	25
Enabling SNMP in Cisco Unified Contact Center Enterprise	25
Enabling SNMP in Cisco Unified Customer Voice Portal (CVP)	29
Enabling SNMP in Cisco Unified Intelligence Center (CUIC)	32
Enabling SNMP in Cisco Finesse Server	34
Cisco: Cloud Services Platform	37
Prerequisites for Monitoring CSP Clusters	37
Discovering CSP Clusters in the SL1 Classic User Interface	40
Cisco: CUCM Unified Communications Manager	41
Prerequisites for Monitoring CUCM	41
Configuring the ScienceLogic Platform to Monitor CUCM	42
Enabling the CUCM AXL Web Service	45
Configuring a CUCM User Account	46
Configuring Prime License Manager	53
Cisco: ESA	55
Prerequisites for Monitoring Cisco Email Security Appliances	55
Cisco: Hyperflex	56
Prerequisites for Monitoring Cisco HyperFlex	56
Cisco: Meeting Server	57
Prerequisites for Monitoring Cisco Meeting Server	57
Cisco: Meraki [API]	58
Generating a Cisco Meraki API Key	58
Disabling Asynchronous Dynamic Application Collection	61
Re-enabling Asynchronous Dynamic Application Collection	61
Creating Events from Cisco Meraki Emails	62

Formatting Inbound Emails	63
Cisco: Tetration	65
Configuring Cisco Tetration Analytics for Monitoring	65
Cisco: UC Ancillary	67
Prerequisites for Monitoring Ancillary Cisco Unified Communications Devices	67
Cisco: UC VOS Applications	68
Configuring Cisco UC VOS Applications for Monitoring	69
Configuring SNMP for Cisco VOS Applications	69
Creating the User Account for the Platform Administrative Web Services (PAWS) API	70
Configuring Cisco Unity Connection	71
Configuring Cisco Unified Communications Manager IM and Presence	72
Configuring Cisco Prime License Manager	73
Configuring Cisco Prime Collaboration Deployment	74
Configuring Cisco Collaboration Mediation Fulfillment	75
Configuring Hosted Collaboration Solution Intelligent Loader	75
Configuring Cisco Contact Center Express	75
Configuring Cisco Emergency Responder	75
Configuring Cisco SocialMiner	75
Cisco: UCS	77
Prerequisites for Monitoring Cisco UCS Manager	77
Configuring the UCS System	77
Cisco: UCS Director	81
Copying the REST API Access Key for a UCS Director Account	81
Cisco: UCS Standalone Rack Server	83
Prerequisites for Monitoring Cisco UCS Standalone Rack Servers	83
Cisco: Unity Express	84
Prerequisites for Monitoring Cisco Unity Express	84
Cisco: Viptela	85
Prerequisite for Monitoring Cisco Viptela	85
Cisco: Wireless	86
Prerequisites for Monitoring Cisco Wireless LAN Controllers	86
Citrix: Xen	87
Enabling Performance Metrics for XenServer 6.2.0 and Above	87
CouchBase	88
Prerequisites for Monitoring Couchbase	88
Dell EMC: Isilon	89
Prerequisites for Monitoring Dell EMC Isilon	89
Dell EMC: Unity	90
Prerequisites for Monitoring Dell EMC Unity	90
Dell EMC: VMAX and PowerMax Unisphere API	91
Prerequisites for Monitoring Dell EMC VMAX and PowerMax Systems	91
Dell EMC: XtremIO	92
Prerequisites for Monitoring Dell EMC XtremIO	92
Configuring Traps with Dell EMC XtremIO	93
Docker	94
Prerequisites for Monitoring Docker	94
Enabling the Docker API	95
Dynatrace	98
Generating a Dynatrace API Token	98
Filtering Partitions from Host Components	99
ELK: AWS CloudTrail	101
Prerequisites for Monitoring AWS ELK Stacks	101

ELK: Azure Activity Log	102
Prerequisites for Monitoring Azure ELK Stacks	102
EMC: VMAX	103
Prerequisites for Monitoring Dell EMC VMAX	103
EMC: VNX	104
Prerequisites for Monitoring Dell EMC VNX	104
F5 BIG-IP	105
Prerequisites for Monitoring F5 BIG-IP	105
F5: BIG-IP DNS	106
Prerequisites for Monitoring F5 BIG-IP DNS	106
Google Cloud Platform *BETA*	107
Creating a Google Cloud Platform Service Account	107
Enabling Google Cloud Platform APIs	110
Hitachi Data Systems: VSP	112
Prerequisites for Monitoring Hitachi VSP Systems	112
IBM: DataPower	113
Prerequisites for Monitoring IBM DataPower Gateways	113
IBM: Db2	114
Prerequisites for Monitoring IBM Db2	114
Prerequisites for Linux/Unix Users	114
Prerequisites for Windows Users	116
IBM: MQ	118
Prerequisites for Monitoring IBM MQ	118
Installing the IBM MQ PowerShell Snap-In for Monitoring on Windows Servers	118
Configuring the IBM: MQ Queue Discovery Snippet	119
IBM: SVC	120
Prerequisites for Monitoring IBM SVC	120
IBM: WebSphere Application Server	121
Prerequisites for Monitoring IBM WebSphere Application Servers	121
JMX Base Pack *BETA*	123
Prerequisites for Monitoring JMX Resources	123
Kubernetes	124
Prerequisites for Monitoring Kubernetes Clusters	124
Required Permissions for the Service Account Token	124
Linux Base Pack	126
Prerequisites for Monitoring Linux Devices with SSH	126
Configuring Linux Devices to Collect Data	127
Microsoft: Azure	131
Configuring an Azure Active Directory Application	132
Creating an Active Directory Application in the Azure Portal	132
Adding Microsoft Graph APIs Permissions to the Application	134
Generating the Secret Key	136
Locating the Application ID and Tenant ID	137
Locating the Subscription ID	137
Adding Reader Access to the Active Directory Application	138
Setting Up a Proxy Server	140
Microsoft: Azure	141
Configuring an Azure Active Directory Application	142
Creating an Active Directory Application in the Azure Portal	142
Adding Microsoft Graph APIs Permissions to the Application	144
Generating the Secret Key	146
Locating the Application ID and Tenant ID	147

Locating the Subscription ID	147
Adding Reader Access to the Active Directory Application	148
Setting Up a Proxy Server	150
Creating a SOAP/XML Credential for Azure	150
Load-Balancing an Account with Multiple Subscriptions	153
Creating an Azure Credential	153
Testing the Azure Credential Using the Credential Tester Panel	155
Microsoft: Office 365	156
Configuring Office 365 Monitoring	156
Creating an Office 365 Active Directory Application in the Azure Portal	157
Adding API Permissions to the Application	159
Generating the Secret Key	162
Microsoft: SQL Server Enhanced	163
Prerequisites for Monitoring SQL Servers	163
SQL Cluster Monitoring	164
Monitoring SQL Clusters on SL1 8.12.1 or greater.	165
Monitoring SQL CLusters on SL1 8.8.1 to 8.12.0	165
MySQL	167
Prerequisites for Monitoring MySQL	167
NetApp Base Pack	168
Prerequisites for Monitoring NetApp	168
New Relic: APM	170
Prerequisites for Monitoring New Relic Services	170
NGINX: Open Source and Plus	171
Prerequisites for Monitoring NGINX Services	171
Nimble Storage	172
Prerequisites for Monitoring Nimble Storage Arrays	172
OpenStack	173
Configuring OpenStack for Monitoring	173
Prerequisites for Monitoring OpenStack	174
Assigning a Role to a User	174
Adding the User Role to API Policy Endpoints	175
Policy Permissions for Administrators	175
Policy Permissions for Non-Administrator Users	178
Oracle: Database	181
Prerequisites for Monitoring Oracle Database Instances	181
Enabling PEM on a Linux Machine	183
Troubleshooting Discovery Issues	184
Palo Alto	185
Prerequisites for Monitoring Palo Alto Firewalls	185
Pure Storage	186
Generating a Pure Storage API Token	186
Testing TCP Port Connectivity	187
Silver Peak	188
Prerequisites for Monitoring Silver Peak	188
SMI-S: Array	189
Prerequisites for Monitoring SMI-S Providers	189
SoftLayer: Cloud	190
Copying Your SoftLayer API Key	190
VMware: NSX	191
Prerequisites for Monitoring VMware NSX	191
VMware: vSphere Base Pack	192

Prerequisites for Monitoring VMware vCenter Servers	192
Creating a Read-Only User Account for Monitoring	193

Chapter

1

Introduction

Overview

This manual describes the configuration steps required for monitoring third-party products in SL1 using the latest versions of the following PowerPacks:

- Alibaba Cloud: Aliyun
- Apcon
- Amazon Web Services
- AMQP: RabbitMQ
- Aruba Central
- Cisco: ACI
- Cisco: AppDynamics
- Cisco: CloudCenter
- Cisco: Contact Center Enterprise
- Cisco: Cloud Services Platform
- Cisco: CUCM Unified Communications Manager
- Cisco: ESA
- Cisco: Hyperflex
- Cisco: Meeting Server
- Cisco: Meraki [API]
- Cisco: Tetration
- Cisco: UC Ancillary
- Cisco: UC VOS Applications

- Cisco: UCS
- Cisco: UCS Director
- Cisco: UCS Standalone Rack Server
- Cisco: Unity Express
- Cisco: Viptela
- Cisco: Wireless
- Citrix: Xen
- CouchBase
- Dell EMC: Isilon
- Dell EMC: Unity
- Dell EMC: VMAX and PowerMax Unisphere API
- Dell EMC: XtremIO
- Docker
- Dynatrace
- ELK: AWS CloudTrail
- ELK: Azure Activity Log
- EMC: VMAX
- EMC: VNX
- F5 BIG-IP
- F5: BIG-IP DNS
- Google Cloud Platform *BETA*
- Hitachi Data Systems: VSP
- IBM: DataPower
- IBM: Db2
- IBM: MQ
- IBM: SVC
- IBM: Tivoli Storage Manager
- IBM: WebSphere Application Server
- JMX Base Pack *BETA*
- Kubernetes
- Linux Base Pack
- Microsoft: Azure
- Microsoft: Office 365
- Microsoft: SQL Server Enhanced
- MySQL

- NetApp Base Pack
- New Relic: APM
- NGINX: Open Source and Plus
- Nimble Storage (2.3)
- OpenStack
- Oracle: Database
- Palo Alto
- Pure Storage
- Silver Peak
- SMI-S: Array
- SoftLayer: Cloud
- VMware: NSX
- VMware: vSphere Base Pack

Additional Reading

For more information about each of the PowerPacks listed in the previous section, see the appropriate [PowerPack-specific manual](#).

For more information about a PowerPack that is not listed in the previous section, see one of the following manuals:

- Monitoring SNMP-Enabled Devices
- Monitoring Switches, Routers, and Firewalls with SNMP
- Monitoring Video Devices
- Monitoring Windows Systems with PowerShell
- Monitoring Windows Systems with WMI

NOTE: ScienceLogic provides this documentation for the convenience of ScienceLogic customers. Some of the configuration information contained herein pertains to third-party vendor software that is subject to change without notice to ScienceLogic. ScienceLogic makes every attempt to maintain accurate technical information and cannot be held responsible for defects or changes in third-party vendor software. There is no written or implied guarantee that information contained herein will work for all third-party variants. See the End User License Agreement (EULA) for more information.

Chapter

2

Alibaba Cloud: Aliyun

Prerequisites

To configure the SL1 system to monitor Aliyun using the *Alibaba Cloud: Aliyun* PowerPack, you must have the account access key ID and password for the Aliyun service you want to monitor.

NOTE: To properly discover and model your Aliyun service in SL1, the account must have at least Read-Only access to the Aliyun service you want to monitor.

NOTE: For more information about the *Alibaba Cloud: Aliyun* PowerPack, see the *Monitoring Alibaba Cloud* manual.


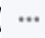
Chapter

3

Amazon Web Services

Overview

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

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

For more information about discovering and monitoring your AWS Infrastructure, watch the video at <https://www.youtube.com/watch?v=ZPqNciWv0Tk>.

The following sections describe several options available for using the *Amazon Web Services PowerPack* to monitor your AWS accounts.

<i>Configuring AWS to Report Billing Metrics</i>	12
<i>Filtering EC2 Instances By Tag</i>	13
<i>Automatic SL1 Organization Creation</i>	15
<i>Monitoring Consolidated Billing Accounts</i>	15
<i>ScienceLogic Events and AWS Alarms</i>	15

NOTE: For more information about the *Amazon Web Services PowerPack*, see the *Monitoring Amazon Web Services* manual.

Configuring AWS to Report Billing Metrics

To use the "AWS: Billing Performance Percent" Dynamic Application, your AWS account must meet the following requirements:

- The user account you supplied in the AWS credential must have permission to view the us-east-1 zone.
- Your AWS account must be configured to export billing metrics to the CloudWatch service.

If your AWS account is not configured to export billing metrics to the CloudWatch service, the "AWS: Billing Performance Percent" Dynamic Application will generate the following event:

```
No billing metrics can be retrieved. Your AWS account is not configured to export billing metrics into CloudWatch.
```

To configure your AWS account to export billing metrics to the CloudWatch service, perform the following steps:

1. Open a browser session and go to aws.amazon.com.
2. Click **[My Account]** and then select *Billing & Cost Management*. If you are not currently logged in to the AWS site, you will be prompted to log in:

amazon
webservices

Sign In or Create an AWS Account

What is your e-mail or mobile number?

E-mail or mobile number:

I am a new user.

I am a returning user and my password is:

Sign in using our secure server

[Forgot your password?](#)

Now Available
Amazon Aurora
Enterprise-class database at 1/10th the cost

[Learn more](#)

Learn more about [AWS Identity and Access Management](#) and [AWS Multi-Factor Authentication](#), features that provide additional security for your AWS Account. View full [AWS Free Usage Tier](#) offer terms.

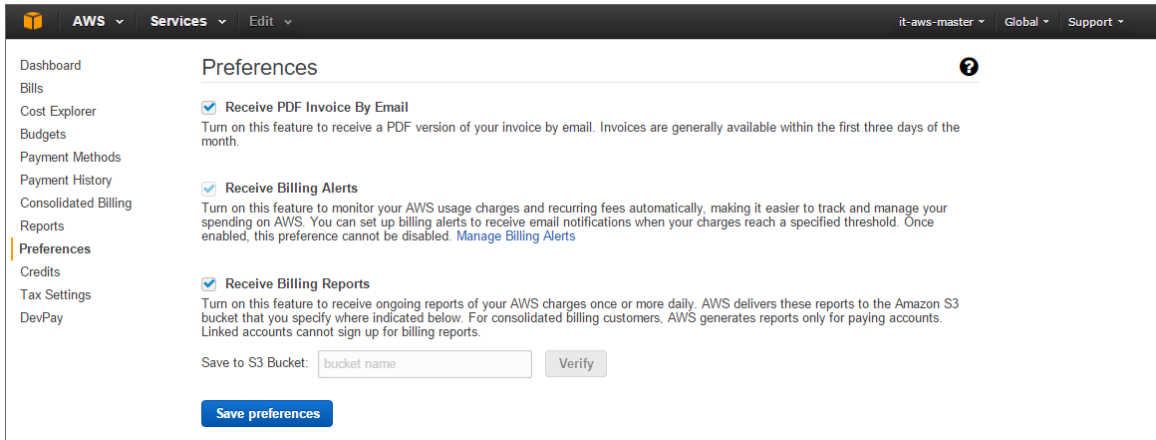
About Amazon.com Sign In

Amazon Web Services uses information from your Amazon.com account to identify you and allow access to Amazon Web Services. Your use of this site is governed by our [Terms of Use](#) and [Privacy Policy](#) linked below.

[Terms of Use](#) [Privacy Policy](#) © 1996-2015, Amazon.com, Inc. or its affiliates

An **amazon.com** company

3. After logging in, the **Billing & Cost Management Dashboard** page appears. In the left navigation bar, click **[Preferences]**. The **Preferences** page appears:



4. Select the **Receive Billing Alerts** checkbox.

CAUTION: If you enable this option, this option cannot be disabled.

5. Click the **[Save Preferences]** button.

Filtering EC2 Instances By Tag

To discover EC2 instances and filter them by tag, you can use the "AWS Credential - Tag Filter" sample credential to enter EC2 tag keys and values.

NOTE: Filtering EC2 instance by tag will apply to **all** accounts discovered.

NOTE: Any EC2 instances that have already been discovered, but do not match the tag filter, will be set to "Unavailable."

To define an AWS credential to discover EC2 instances and filter them by tag:

1. Go to the **Credential Management** page (System > Manage > Credentials).

2. Locate the **AWS Credential - Tag Filter** sample credential and click its wrench icon (🔧). The **Credential Editor** modal page appears:

3. Enter values in the following fields:

Basic Settings

- **Profile Name.** Type a new name for your AWS credential.
- **HTTP Auth User.** Type your AWS access key ID.
- **HTTP Auth Password.** Type your AWS secret access key.

HTTP Headers

- Edit the HTTP header provided:
 - `Tags:<operation>#<EC2-Tag-Key>#<EC2-Tag-Value>`. Type the tag, followed by its operation, tag key, or 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
```

4. Click the **[Save As]** button, and then click **[OK]**.

Automatic SL1 Organization Creation

This feature is only applicable to the two discovery methods that use the Assume Role and automatically discover multiple accounts.

When multiple accounts are discovered, this feature places each account in its own SL1 organization. This feature requires an optional header in the SOAP/XML credential you will create. When this header is present, it will place each account into a new SL1 organization. When this header is not present, each account will be placed in the SL1 organization selected in the discovery session. The name of the organization can be controlled depending on what is provided in the header as follows:

- **OrganizationCreation:NAME:ID**. Autocreates an SL1 organization for accounts using AssumeRole. You can enter one of the following options:
 - **OrganizationCreation:NAME**. The name of the organization will contain the name of the user.
 - **OrganizationCreation:ID**. The name of the organization will contain the ID of the user.
 - **OrganizationCreation:ID:NAME**. The name of the organization will contain both the ID and name of the user, in that order.
 - **OrganizationCreation:NAME:ID**. The name of the organization will contain both the name and ID of the user, in that order.

Monitoring Consolidated Billing Accounts

Consolidated billing is an option provided by Amazon that allows multiple AWS accounts to be billed under a single account. For more information about consolidated billing, see <http://docs.aws.amazon.com/awsaccountbilling/latest/aboutv2/consolidated-billing.html>.

If a consolidated billing account is monitored by SL1, the billing metrics associated with that account include only the consolidated amounts, per service. If you use consolidated billing and want to collect billing metrics per-account, you must discover each account separately. To monitor only the billing metrics for an AWS account, you can create credentials that include only billing permissions.

ScienceLogic Events and AWS Alarms

In addition to SL1 collecting metrics for AWS instances, you can configure CloudWatch to send alarm information to SL1 via API. SL1 can then generate an event for each alarm.

For instructions on how to configure CloudWatch and SL1 to generate events based on CloudWatch alarms, see the [Configuring Inbound CloudWatch Alarms](#) section.

NOTE: All examples shown are for commercial AWS accounts. When AWS Gov is being monitored, the JSON data that refers to ARN will need to be modified from "aws" to "aws-us-gov". For example:
Resource": "arn:aws:iam::<account number>:role/Sciencelogic-Monitor would need to be Resource": "arn:aws:iam-us-gov::<account number>:role/Sciencelogic-Monitor

Chapter

4

AMQP: RabbitMQ

Prerequisites for Monitoring RabbitMQ

To configure SL1 to monitor a RabbitMQ system using the *AMQP: RabbitMQ* PowerPack, you must first have the following information:

- The IP address of the server running the RabbitMQ system
- The username and password for a RabbitMQ user that has read permission to the RabbitMQ API. For information about configuring users in RabbitMQ, see <https://www.rabbitmq.com/management.html>.

NOTE: For more information about the *AMQP: RabbitMQ* PowerPack, see the *Monitoring Monitoring RabbitMQ Systems* manual.

Chapter

5

Aruba Central

Prerequisites for Monitoring Aruba Central

Before you can monitor Aruba Central virtual controllers and their component devices using the *Aruba Central PowerPack*, you must first have the following information:

- Aruba Central username and password
- Aruba Central customer ID
- Aruba Central client ID
- Aruba Central client secret key

You can request these items by registering with Aruba Technical Support.

NOTE: For more information about the *Aruba Central PowerPack*, see the *Monitoring Aruba Central* manual.

Chapter

6

Cisco: ACI

Prerequisites for Monitoring Cisco ACI

To configure the SL1 system to monitor a Cisco ACI system using the *Cisco: ACI PowerPack*, you must first:

- Know the credentials (username and password) for a user account that has access to the API for the Cisco ACI system. The user account must have read-all access.
- Ensure that the APIC in your ACI system supports TLS 1.1 or TLS 1.2. SL1 does not support TLS 1.0.

NOTE: If the credentials for your account have been changed, the PowerPack will not recognize the new credentials. To recognize new credentials, you can either delete or disable the previous administrator account, or delete any cache entries with "1C88582E76AADD40EB8C5E6A6F71B64A_ACI_{host}_{cred_id}_TOKENS".

Recommended System Values

ScienceLogic recommends that you set the following values on your Cisco ACI system:

- **ACI HTTPS Throttle.** 5 requests per second.
- **Web Session Timeout.** 600 seconds or greater.
- **Web Session Idle Timeout.** 600 seconds (default).

NOTE: For more information about the *Cisco: ACI PowerPack*, see the **Monitoring Cisco ACI** manual.

When SL1 performs collection for the ACI cluster, SL1 will create component devices for the components associated with the ACI system and align other Dynamic Applications to those component devices. Some of the Dynamic Applications aligned to the component devices will also be used to create additional component devices.

NOTE: If you delete a Tenant in a monitored device, that component device will still appear in SL1 but the Dynamic Applications aligned to it will stop collecting data, and a message indicating "Failed Availability" will appear in the device log of its child component devices.

You can view all the devices, virtual devices, and component devices in the Cisco ACI system in the following places in the user interface:

- The **Device Investigator** Map page (click **Map** in the **Device Investigator** page) displays a map of a particular device and all of the devices with which it has parent-child relationships. Double-clicking any of the listed devices reloads the page to make the selected device the primary device.
- The **Device Components** page (Devices > Device Components) displays a list of all root devices and component devices discovered by SL1. The **Device Components** page displays all root devices and component devices in an indented view, so you can easily view the hierarchy and relationships between child devices, parent devices, and root devices. To view the component devices associated with a Cisco ACI system, find the Cisco ACI root device and click its plus icon (+).
- The **Component Map** page (Classic Maps > Device Maps > Components) allows you to view devices by root node and view the relationships between root nodes, parent components, and child components in a map. This makes it easy to visualize and manage root nodes and their components. SL1 automatically updates the **Component Map** as new component devices are discovered. The platform also updates each map with the latest status and event information. To view the map for a Cisco ACI device, go to the **Component Map** page and select the map from the list in the left NavBar. To learn more about the **Component Map** page, see the **Views** manual.

Chapter

7

Cisco: AppDynamics

Prerequisites for Monitoring Cisco AppDynamics

Before you can monitor Cisco AppDynamics applications using the *Cisco: AppDynamics PowerPack*, you must first create a user account that is assigned the "Applications and Dashboard Viewer" role in the AppDynamics account portal. This user account must also have sufficient permissions to obtain metrics information from the AppDynamics REST API.

For more information about creating the AppDynamics user account, see <https://docs.appdynamics.com/display/PRO44/Roles+and+Permissions>.

NOTE: For more information about the *Cisco: AppDynamics PowerPack*, see the *Monitoring Cisco AppDynamics* manual.

Chapter

8

Cisco: CloudCenter

Overview

The following sections describe how to configure a Cloud Center Manager for monitoring by SL1 using the Cisco: *CloudCenter PowerPack*:

Configuration and Discovery for Standard Cisco CloudCenter Deployments	22
Prerequisites for Monitoring Standard CloudCenter Deployments	23
Configuration and Discovery for High-Availability Cisco CloudCenter Deployments	23
Prerequisites for Monitoring High-Availability CloudCenter Deployments	23

NOTE: For more information about the Cisco: *CloudCenter PowerPack*, see the *Monitoring Cisco CloudCenter* manual.

Configuration and Discovery for Standard Cisco CloudCenter Deployments

The Cisco: *CloudCenter PowerPack* enables you to discover and collect configuration and performance data about standard or high-availability (HA) CloudCenter deployments and their components. The following sections describe the configuration and discovery steps for monitoring standard (non-HA) CloudCenter deployments.

For information about HA deployments, see the section on [Configuration and Discovery for High-Availability Cisco CloudCenter Deployments](#).

Prerequisites for Monitoring Standard CloudCenter Deployments

To configure the SL1 system to monitor standard (non-HA) Cisco CloudCenter deployments using the Cisco: *CloudCenter* PowerPack, you must first have the following information about the CloudCenter Manager that you want to monitor:

- The IP address of the CloudCenter Manager system
- The username and API key for a Cisco CloudCenter Manager user that has root tenant administration privileges. This account must be an API user, not a GUI user. For information about configuring API users in Cisco CloudCenter Manager, see <http://docs.cloudcenter.cisco.com/display/40API/API+Management+Key>.

Configuration and Discovery for High-Availability Cisco CloudCenter Deployments

The Cisco: *CloudCenter* PowerPack enables you to discover and collect configuration and performance data about standard or high-availability (HA) CloudCenter deployments and their components. The following sections describe the configuration and discovery steps for monitoring HA CloudCenter deployments.

For information about standard (non-HA) deployments, see the section on [Configuration and Discovery for Standard Cisco CloudCenter Deployments](#).

Prerequisites for Monitoring High-Availability CloudCenter Deployments

To configure the SL1 system to monitor HA Cisco CloudCenter deployments using the Cisco: *CloudCenter* PowerPack, you must first have the following information about the CloudCenter components that you want to monitor:

- The IP address or hostname for each of the following components:
 - RabbitMQ
 - RabbitMQ Load Balancer
 - Cisco CloudCenter Manager
 - Cisco CloudCenter Manager Load Balancer
 - CloudCenter PostgreSQL database
 - CloudCenter Orchestrator
 - CloudCenter Orchestrator Load Balancer
 - CloudCenter Health Monitor
 - CloudCenter ELK components

- The username and API key for a Cisco CloudCenter Manager user that has root tenant administration privileges. This account must be an API user, not a GUI user. For information about configuring API users in Cisco CloudCenter Manager, see <http://docs.cloudcenter.cisco.com/display/40API/API+Management+Key>.
- The username and password for a RabbitMQ user that has read permission to the RabbitMQ API. For information about configuring users in RabbitMQ, see <https://www.rabbitmq.com/management.html>.
- The usernames and passwords for Cisco CloudCenter users that have API read permissions for each of the other components in the above list.

Cisco: Contact Center Enterprise

Overview

The following sections describe how to configure Cisco Unified Contact Center Enterprise services for monitoring by SL1 using the *Cisco: Contact Center Enterprise PowerPack*:

<i>Configuring Unified Contact Center Enterprise Monitoring Using SNMP</i>	25
<i>Enabling SNMP in Cisco Unified Contact Center Enterprise</i>	25
<i>Enabling SNMP in Cisco Unified Customer Voice Portal (CVP)</i>	29
<i>Enabling SNMP in Cisco Unified Intelligence Center (CUIC)</i>	32
<i>Enabling SNMP in Cisco Finesse Server</i>	34

<p>NOTE: For more information about the <i>Cisco: Contact Center Enterprise PowerPack</i>, see the <i>Monitoring Cisco Unified Contact Center Enterprise</i> manual.</p>

Configuring Unified Contact Center Enterprise Monitoring Using SNMP

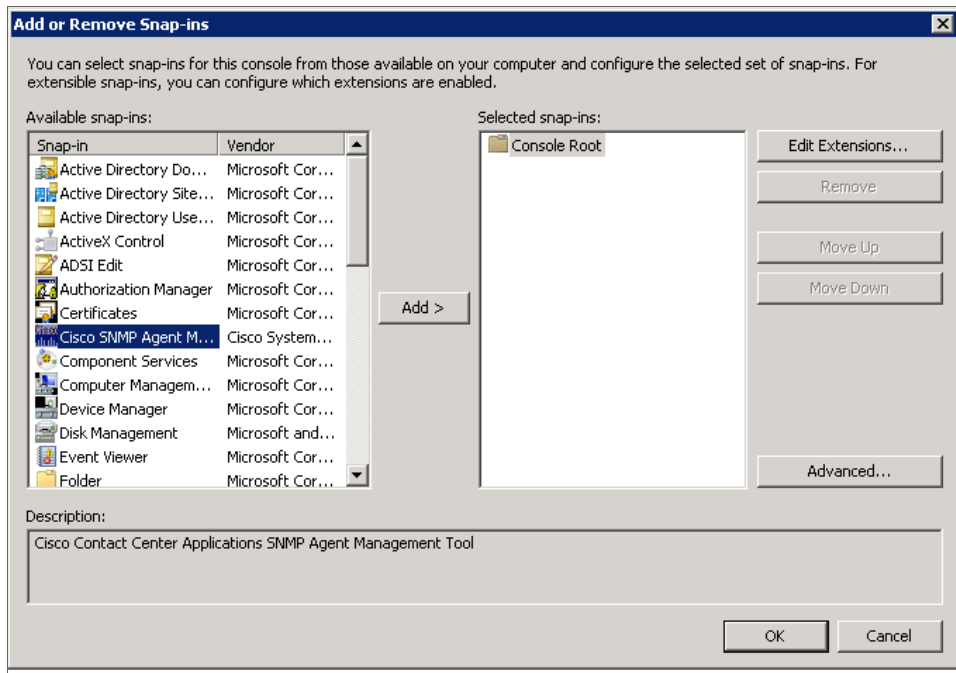
Before you can discover and monitor Cisco Unified Contact Center Enterprise (UCCE) devices in SL1, you must first configure SNMP community strings in each of the UCCE services that you will monitor with SL1.

Enabling SNMP in Cisco Unified Contact Center Enterprise

To enable SNMP in Cisco Unified Contact Center Enterprise, perform the following steps:


1. Log in to the Cisco Unified Contact Center Enterprise Server as an administrator.

2. Open Microsoft Management Console (32-bit).
3. Click **[File]**, then select *Add/Remove Snap-In*. The **Add or Remove Snap-ins** page appears.



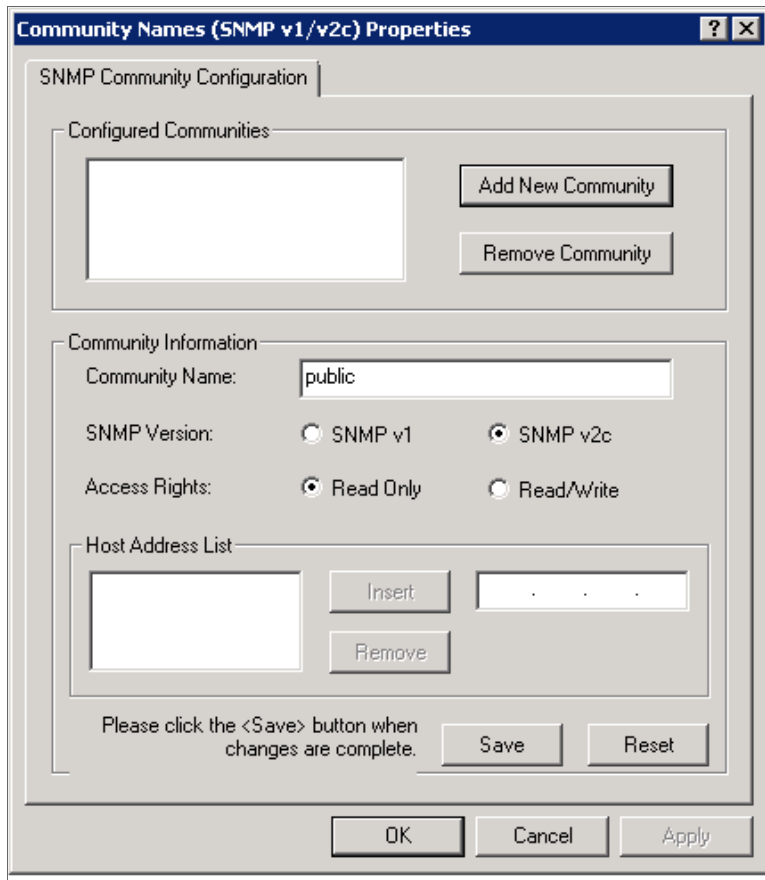
4. In the **Available snap-ins** field, select **Cisco SNMP Agent Management**, then click **[Add >]** to move it to the **Selected snap-ins** field.
5. Click **[OK]**.

6. In the left panel of the Microsoft Management Console, click **Cisco SNMP Agent Management**. Then, in the right panel, right-click **Community Names (SNMP v1, v2c)** and select *Properties*.



7. In the **Community Names (SNMP v1/v2c) Properties** modal page, click the **[Add New Community]** button to enable the fields on the page.


8. Make entries in the following fields:



- **Community Name.** Enter a name for the new community string.
- **SNMP Version.** Select *SNMP v2c*.
- **Access Rights.** Select *Read Only*.

9. Click **[Save]**, and then click **[OK]**.
10. Close the Microsoft Management Console.
11. Open the Microsoft Windows Services console.

- In the Microsoft Windows Services console, select **Cisco Contact Center SNMP Management** from the list of local services, then click the **Restart** hyperlink to restart the service.



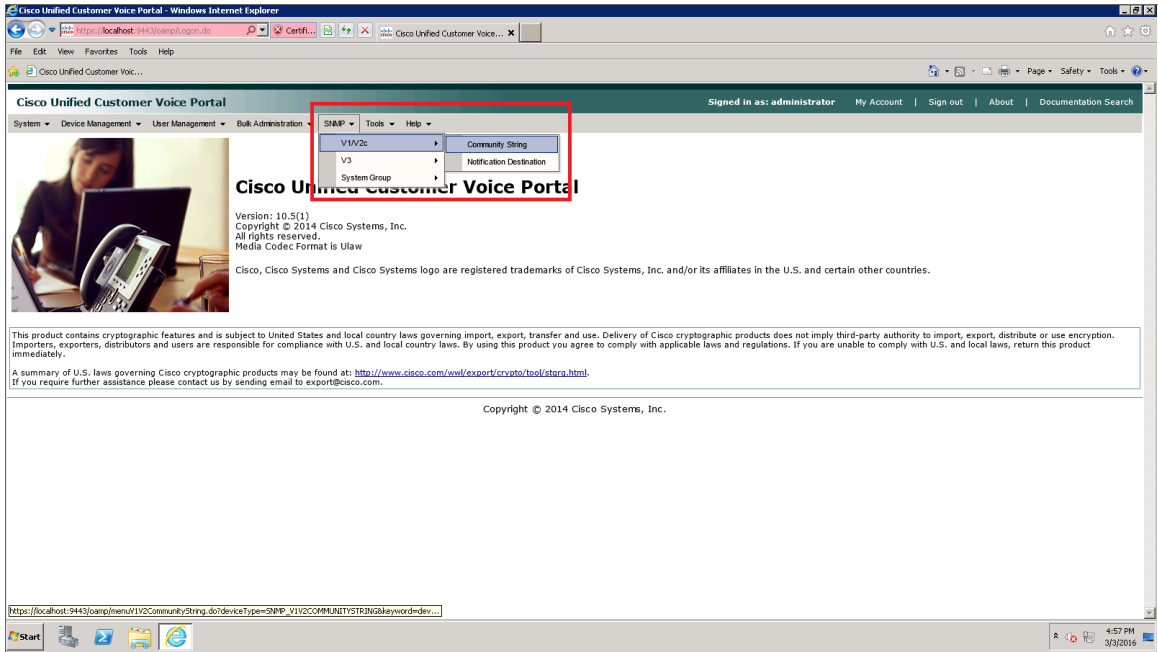
- Close the Microsoft Windows Services console.
- Click the Windows **[Start]** menu, then go to Control Panel > System and Security > Windows Firewall.
- In the left panel, click the **Turn Windows Firewall on or off** hyperlink. The **Customize Settings** page appears.
- Under **Domain network location settings**, select *Turn off Windows Firewall*, then click **[OK]**.
- To enable SNMP in Cisco Unified Contact Center Enterprise Data Server, log in to Cisco Unified Contact Center Enterprise Data Server as an administrator and repeat steps 2-16.

Enabling SNMP in Cisco Unified Customer Voice Portal (CVP)


To enable SNMP in Cisco Unified Customer Voice Portal, perform the following steps:

- Log in to Cisco Unified Customer Voice Portal as an administrator.


2. Click the [SNMP] tab, then select V1/V2c > Community String.



3. On the Find, Add, Delete, Edit V1/V2c Community Strings page, click the [Add New] button.

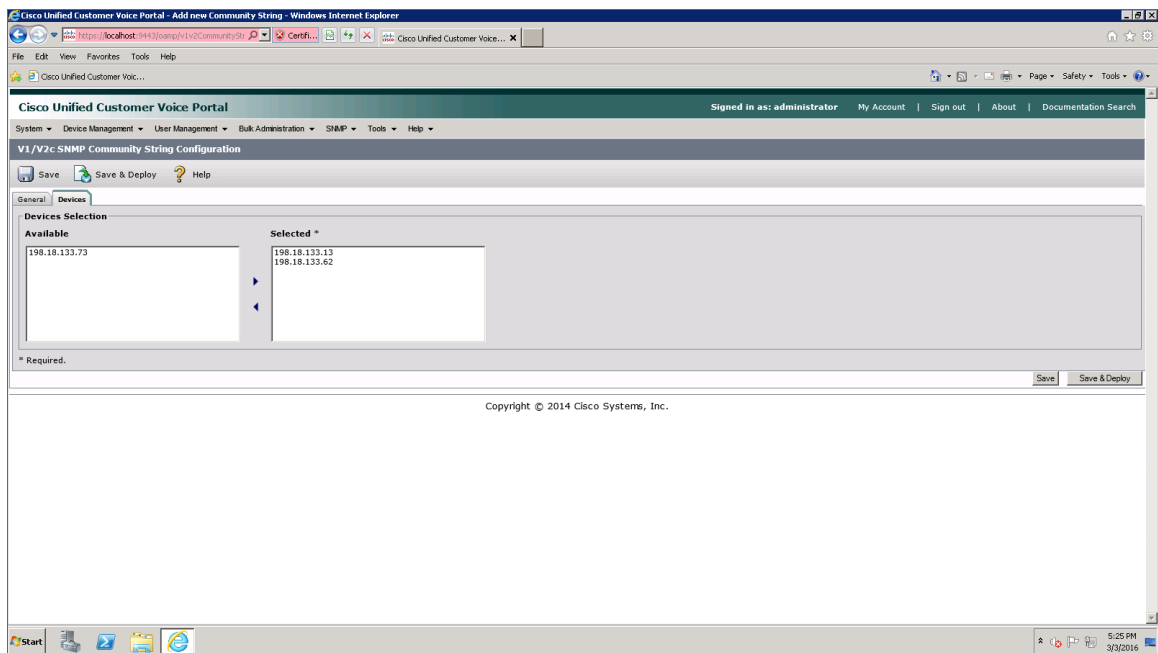


4. The **V1/V2c SNMP Community String Configuration** page appears. Make entries in the following fields:



- **Community String Name.** Enter a name for the new community string.
- **SNMP Version Information.** Select V2C.
- For the other fields on the page, use the default values.

5. Click the **[Devices]** tab.

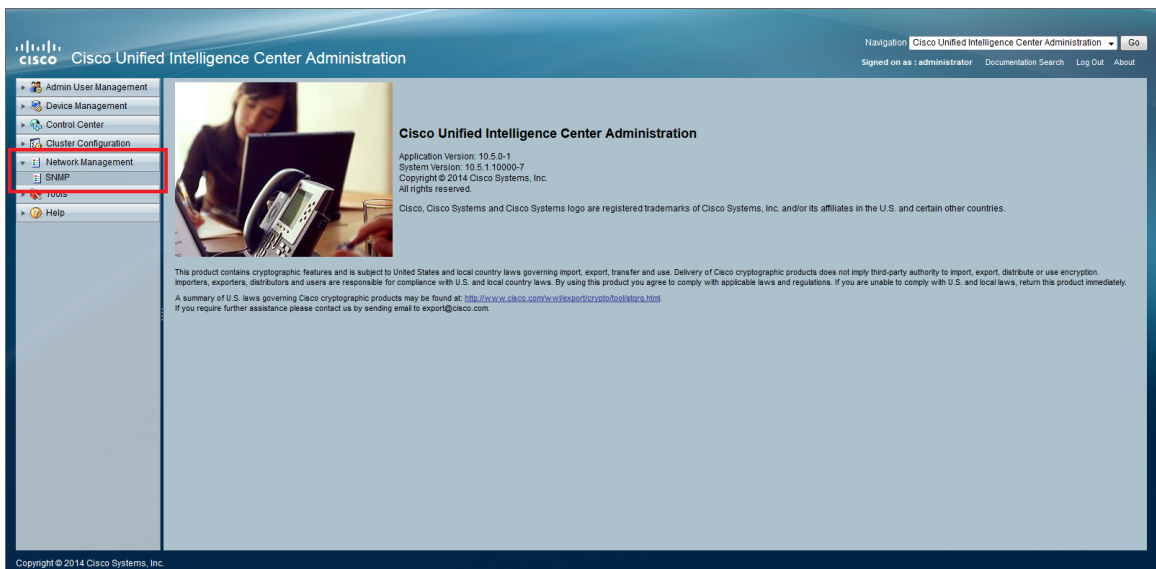


6. Select one or more of the devices in the **Available** field, then click the right-arrow icon to move the selected device(s) to the **Selected** field.
7. Click the **[Save & Deploy]** button. A message confirms that the configuration of the SNMP community string was successfully applied to the selected device(s).

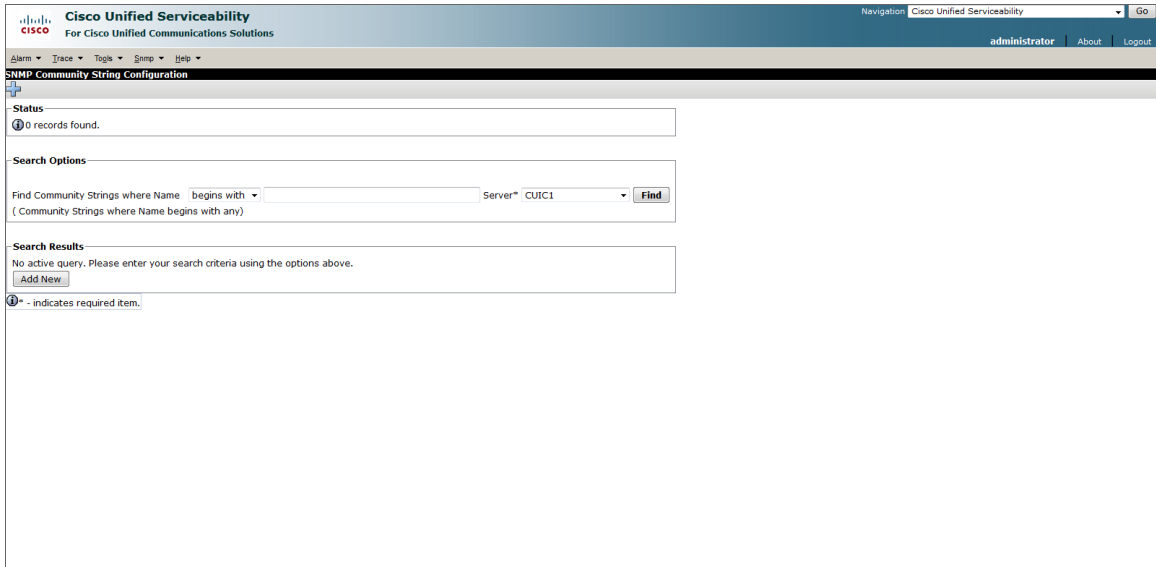
Enabling SNMP in Cisco Unified Intelligence Center (CUIC)

To enable SNMP in Cisco Unified Intelligence Center, perform the following steps:

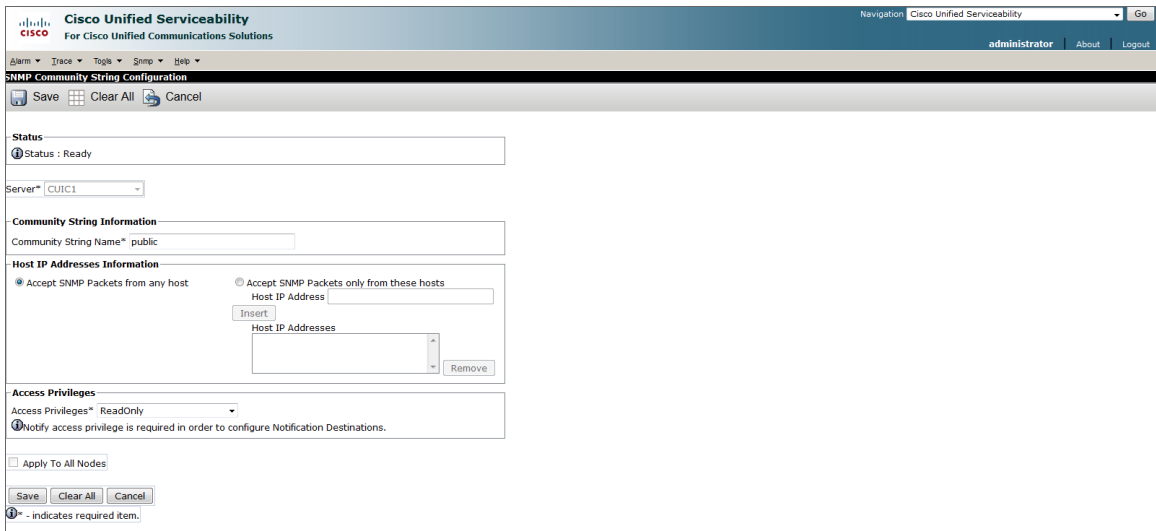
1. Log in to Cisco Unified Intelligence Center as an administrator.
2. In the left panel, click **[Network Management]**, then select **SNMP**.



- On the **SNMP Community String Configuration** page, under **Search Options**, click **[Find]**. The **Search Results** section appears.



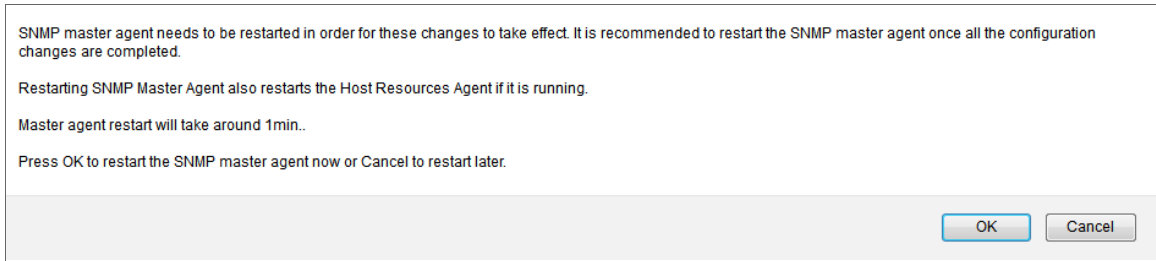
- Under **Search Results**, click **[Add New]**.
- Enter values in the following fields:



- Community String Name.** Enter a name for the new community string.
- Access Privileges.** Select *ReadOnly*.
- For the other fields on the page, use the default values.

- Click **[Save]**.

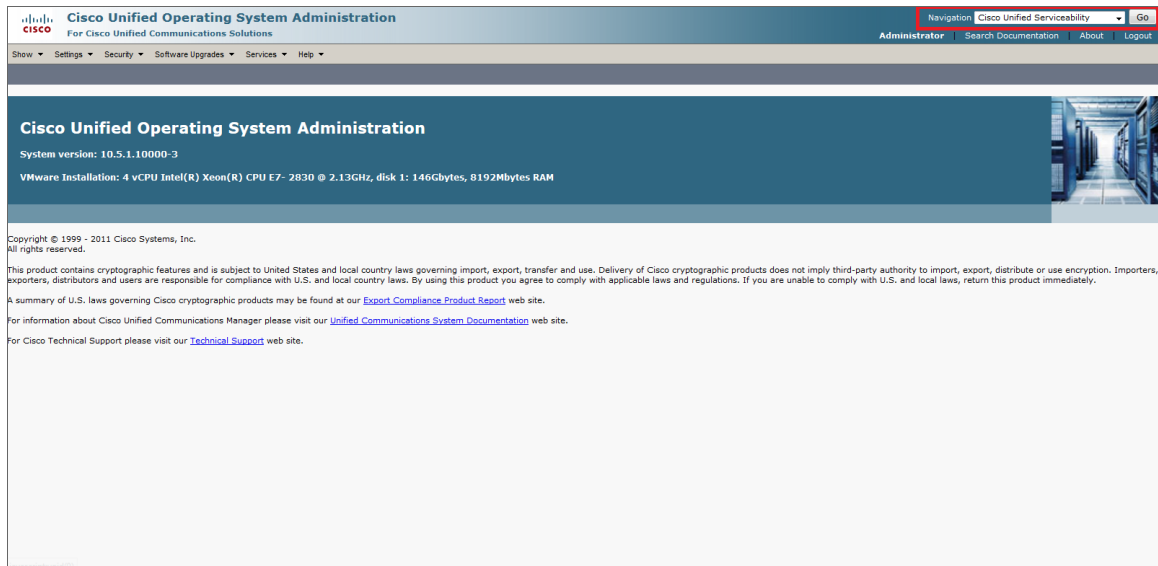
7. Click **[OK]** to restart the SNMP master agent.



Enabling SNMP in Cisco Finesse Server

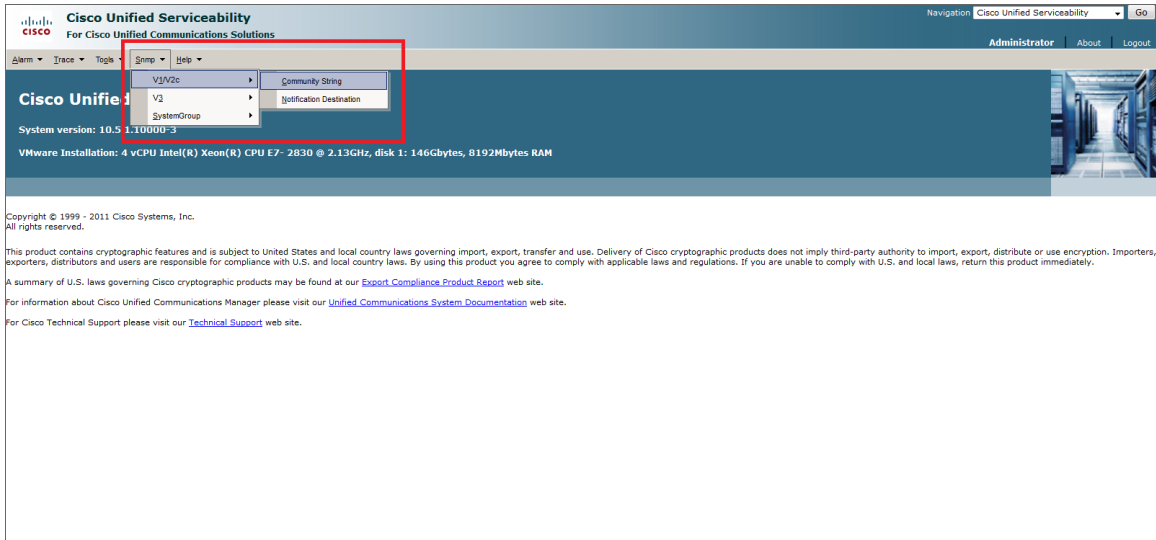
To enable SNMP in Cisco Finesse Server, perform the following steps:

1. Log in to Cisco Unified Operating System Administration as an administrator.
2. In the top-right corner of the page, in the **Navigation** field, select *Cisco Unified Serviceability* and then click **[Go]**.




NOTE: You might be required to enter your login credentials again before proceeding.

3. Click the [SNMP] tab, then select V1/V2c > Community String.



4. On the **SNMP Community String Configuration** page, under **Search Options**, click [Find]. The **Search Results** section appears.



5. Under **Search Results**, click [Add New].

6. Enter values in the following fields:

The screenshot shows the 'SNMP Community String Configuration' page in the Cisco Unified Serviceability interface. The page includes the following fields and controls:

- Status:** A text field containing 'Ready'.
- Server:** A dropdown menu with 'CUIC1' selected.
- Community String Information:** A text field for 'Community String Name*' containing 'public'.
- Host IP Addresses Information:** Two radio buttons: 'Accept SNMP Packets from any host' (selected) and 'Accept SNMP Packets only from these hosts'. Below the second radio button is a 'Host IP Address' text field, an 'Insert' button, a 'Host IP Addresses' list box, and a 'Remove' button.
- Access Privileges:** A dropdown menu with 'ReadOnly' selected. Below it is a message: 'Notify access privilege is required in order to configure Notification Destinations.'
- Apply To All Nodes:** A checkbox that is currently unchecked.
- Buttons:** 'Save', 'Clear All', and 'Cancel' buttons at the bottom.

- **Community String Name.** Enter a name for the new community string.
- **Access Privileges.** Select *ReadOnly*.
- For the other fields on the page, use the default values.

7. Click **[Save]**.

8. Click **[OK]** to restart the SNMP master agent.

The dialog box contains the following text:

SNMP master agent needs to be restarted in order for these changes to take effect. It is recommended to restart the SNMP master agent once all the configuration changes are completed.

Restarting SNMP Master Agent also restarts the Host Resources Agent if it is running.

Master agent restart will take around 1min..

Press OK to restart the SNMP master agent now or Cancel to restart later.

Buttons: OK, Cancel

Chapter

10

Cisco: Cloud Services Platform

Prerequisites for Monitoring CSP Clusters

To configure the SL1 system to monitor CSP clusters using the *Cisco: Cloud Services Platform PowerPack*, you must have the following information about the clusters that you want to monitor:

- Username and password of a user with REST API read access and a role of operator-group or admin-group
- SNMP community string with read privileges and the port set to 161

NOTE: For more information about these requirements, see http://www.cisco.com/c/en/us/td/docs/switches/datacenter/csp_2100/config_guide/b_Cisco_CSP_2100_Config_Guide.html.

Additionally, you must establish a Net-SNMP public community string with the port set to 1610. To do so:

1. Log in to the command line of the CSP device as an administrative user.
2. Run the following commands:


```
netsnmp agent port 1610
netsnmp community public
```

NOTE: For more information about the *Cisco: Cloud Services Platform PowerPack*, see the *Monitoring Cisco Cloud Services Platform 2100* manual.

When you discover your Cloud Services Platform (CSP) cluster with SL1, SL1 auto-aligns a series of Dynamic Applications to discover, configure, and monitor the CSP cluster and all of its associated component devices.

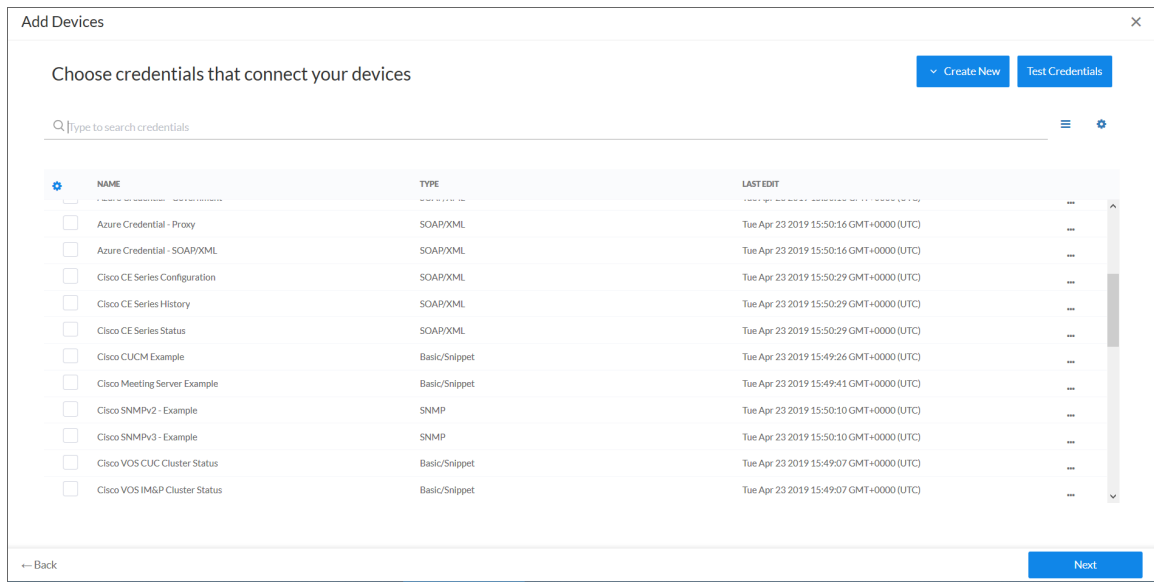
To discover your CSP cluster, perform the following steps:

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]** 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:
 - **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**. Optional. 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.
 - **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:

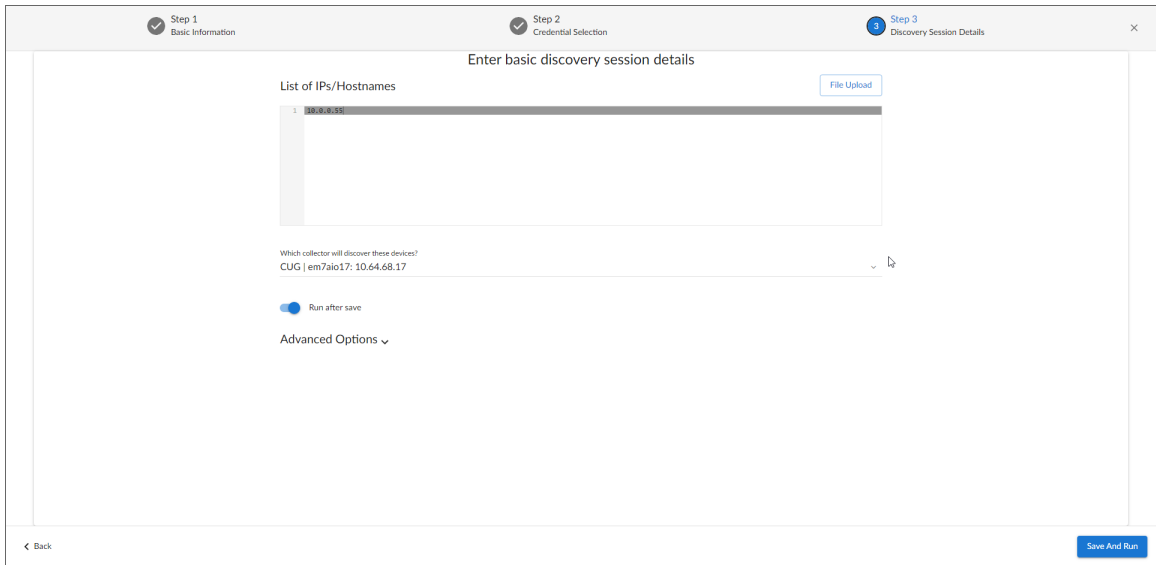


6. On the **Credentials** page, locate and select the two SNMP credentials that you created (one for port 161 and the other for port 1610), and the Basic/Snippet credential and the SSH/Key credential for each of the CSP nodes you want to discover.

NOTE: You must include a minimum of three credentials (one SNMP credential and two Basic/Snippet credentials) for each CSP node with unique credential information.

NOTE: If you are running a Federal Information Processing Standard (FIPS)-compliant installations of the SL1 platform, then you should **not** select an SSH/Key credential.

7. Click **[Next]**. The **Discovery Session Details** page of the **Add Devices** wizard appears:



8. Complete the following fields:

- **List of IPs/Hostnames.** Type the IP address of each CSP node you want to discover.
- **Which collector will monitor these devices?.** Select an existing collector to monitor the discovered devices. Required.
- **Run after save.** Select this option to run this discovery session as soon as you click **[Save and Close]**.

In the **Advanced options** section, click the down arrow icon (▼) to complete the following fields:

- **Discover Non-SNMP.** Enable this setting.
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.

Discovering CSP Clusters in the SL1 Classic User Interface

Chapter

11

Cisco: CUCM Unified Communications Manager

Overview

The following sections describe how to configure a Cisco Unified Communications Manager (CM) system for monitoring by SL1 using the *Cisco: CUCM Unified Communications Manager PowerPack*:

<i>Prerequisites for Monitoring CUCM</i>	41
<i>Configuring the ScienceLogic Platform to Monitor CUCM</i>	42
<i>Enabling the CUCM AXL Web Service</i>	45
<i>Configuring a CUCM User Account</i>	46
<i>Configuring Prime License Manager</i>	53

NOTE: For more information about the *Cisco: CUCM Unified Communications Manager PowerPack*, see the *Monitoring Cisco Unified Communications Manager* manual.

Prerequisites for Monitoring CUCM

During the discovery process, SL1 automatically aligns the IP addresses and hostnames for each node in a Cisco Unified CM cluster via DNS.

If you do not have access to DNS for the Cisco Unified CM systems that you want to monitor with SL1, ensure that you know or have access to the following information about each node:

- IP address
- Hostname

Configuring the ScienceLogic Platform to Monitor CUCM

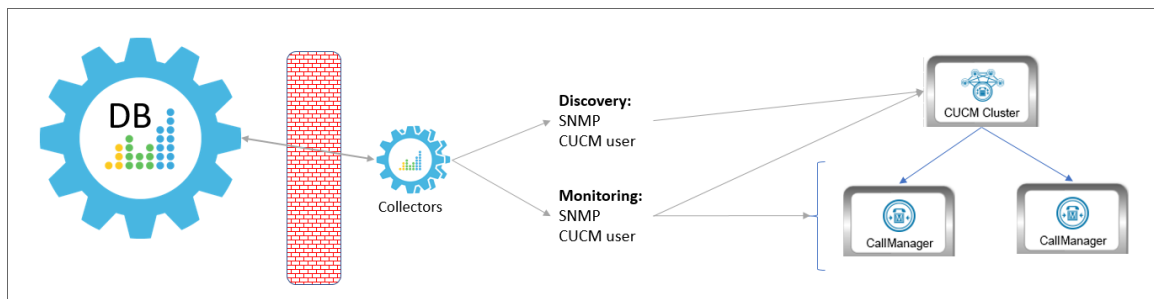
You can choose from several different possible configurations when using SL1 to monitor Cisco Unified CM:

- You can have the ScienceLogic Data Collector either in front of a firewall or behind a firewall.
- You can define the CallManager nodes either by hostname or by IP address in the Cisco Unified CM database.
- In some scenarios, you can also use network address translation (NAT) when defining the CallManagers.

These various methods are described in this section.

Method 1

In the first scenario, the Data Collector sits in front of the firewall and you define the CallManagers by hostname:




In this scenario, you must have the following ports open for the firewall:

Direction	Port	Protocol
ScienceLogic Database Server to the Data Collector	7707	TCP
PhoneHome Collector to the Database Server	7705	TCP

Method 2

In the second scenario, the Data Collector sits in front of the firewall and you define the CallManagers by IP address. This method requires you to *create a host file* that includes the CallManager hostname and IP address:




In this scenario, you must have the following ports open for the firewall:

Direction	Port	Protocol
ScienceLogic Database Server to the Data Collector	7707	TCP
PhoneHome Collector to the Database Server	7705	TCP

Method 3

In the third scenario, the Data Collector sits behind the firewall and you define the CallManagers by hostname:




In this scenario, you must have the following ports open for the firewall:

Direction	Credential	Port	Protocol
ScienceLogic Data Collector to the Cisco Unified CM Cluster and CallManagers	SNMP	161	UDP
	Cisco Unified CM user	8443	TCP

Method 4

In the fourth scenario, the Data Collector sits behind the firewall and you define the CallManagers by hostname, with NAT. This method requires you to [create a host file](#) that includes the CallManager hostname and the IP address the Data Collector can use to access the device:

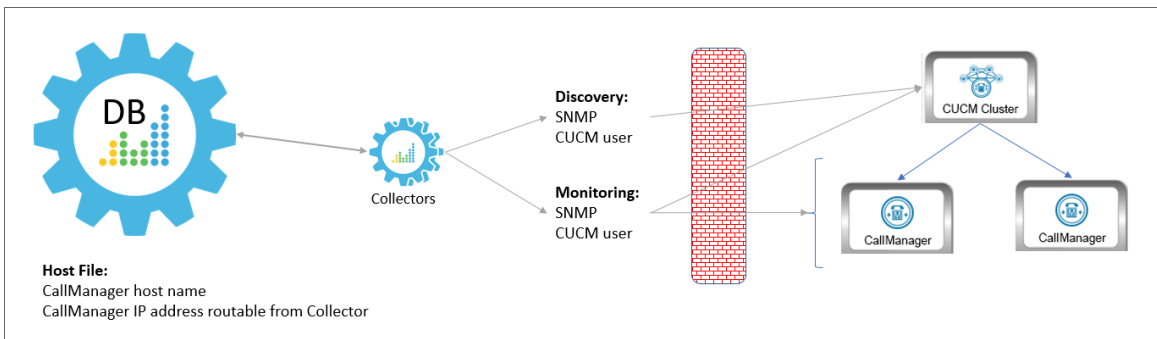


In this scenario, you must have the following ports open for the firewall:

Direction	Credential	Port	Protocol
ScienceLogic Data Collector to the Cisco Unified CM Cluster and CallManagers	SNMP	161	UDP
	Cisco Unified CM user	8443	TCP

Method 5

In the final scenario, the Data Collector sits behind the firewall and you define the CallManagers by IP address, with NAT. This method requires you to [create a host file](#) that includes the CallManager host name and IP address the Data Collector can use to access the device:



NOTE: This method is not supported by versions of the *Cisco: CUCM Unified Communications Manager PowerPack* prior to version 109.

In this scenario, you must have the following ports open for the firewall:

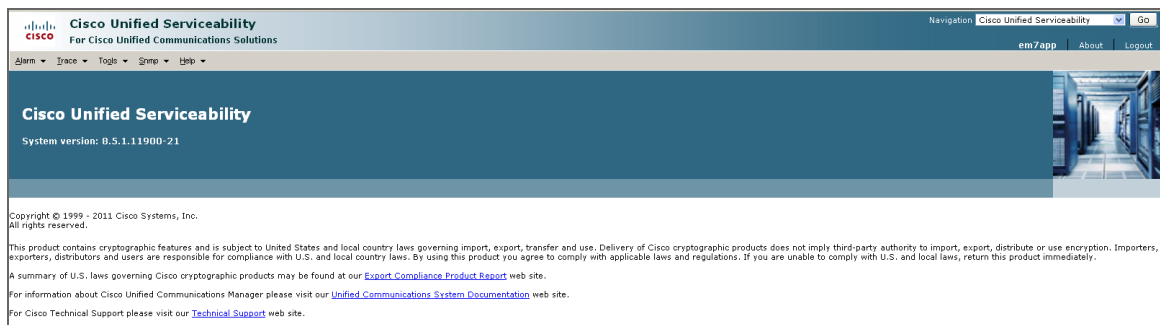
Direction	Credential	Port	Protocol
ScienceLogic Data Collector to the Cisco Unified CM Cluster and CallManagers	SNMP	161	UDP
	Cisco Unified CM user	8443	TCP

Enabling the CUCM AXL Web Service

SL1 can monitor a Cisco Unified CM system by requesting detailed information about the system from the Cisco Unified CM AXL Web Service.

The Cisco Unified CM AXL web service is disabled by default. To enable the AXL web service, perform the following steps:

1. In a browser window, navigate to the following address:
`https://ip-address-of-CM-system:8443/ccmadmin/showHome.do`
2. Log in to the Cisco Unified CM Administration site as an administrator.
3. In the **Navigation** drop-down list at the top-right corner of the page, select *Cisco Unified Serviceability*, and then click the **[Go]** button. The **Cisco Unified Serviceability** page appears:



- In the navigation bar at the top-left of the page, hover over **Tools**, then select **Service Activation**. The **Service Activation** page appears:

The screenshot shows the Cisco Unified Serviceability interface for Service Activation. At the top, there's a navigation bar with 'Alarm', 'Trace', 'Tools', 'Symp', and 'Help'. Below that, the 'Service Activation' section has buttons for 'Save', 'Set to Default', and 'Refresh'. The 'Status' section shows 'Status : Ready'. The 'Select Server' section has a dropdown menu with '192.168.44.22' and a 'Go' button. There's also a checkbox for 'Check All Services'.

The 'CM Services' table is as follows:

Service Name	Activation Status
<input checked="" type="checkbox"/> Cisco CallManager	Activated
<input checked="" type="checkbox"/> Cisco Tftp	Activated
<input checked="" type="checkbox"/> Cisco Messaging Interface	Activated
<input checked="" type="checkbox"/> Cisco Unified Mobile Voice Access Service	Activated
<input checked="" type="checkbox"/> Cisco IP Voice Media Streaming App	Activated
<input checked="" type="checkbox"/> Cisco CTIManager	Activated
<input checked="" type="checkbox"/> Cisco Extension Mobility	Activated
<input checked="" type="checkbox"/> Cisco Extended Functions	Activated
<input checked="" type="checkbox"/> Cisco Dialed Number Analyzer	Activated
<input type="checkbox"/> Cisco DHCP Monitor Service	Deactivated
<input checked="" type="checkbox"/> Cisco Dialed Number Analyzer Server	Activated

The 'CTI Services' table is as follows:

Service Name	Activation Status
<input checked="" type="checkbox"/> Cisco IP Manager Assistant	Activated
<input checked="" type="checkbox"/> Cisco WebDialer Web Service	Activated

- In the **Server** drop-down list, select the Cisco Unified CM server for which you want to enable the AXL web service, and then click the **[Go]** button.
- In the list of services, locate the **Database and Admin Services** section. If the *Activation Status* of the **Cisco AXL Web Service** is "Activated", the AXL web service is already enabled.
- If the *Activation Status* of the **Cisco AXL Web Service** is not "Activated", select the checkbox for the **Cisco AXL Web Service**.
- Click the **[Save]** button at the bottom of the page to save your changes, and then click the **[OK]** button in the pop-up window that appears.

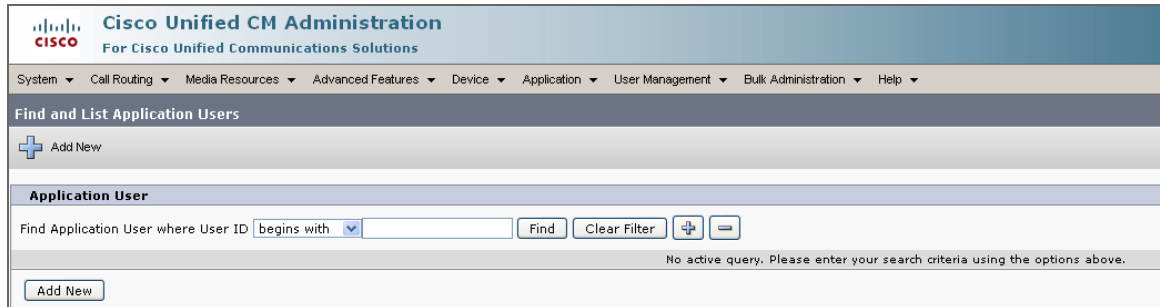
Configuring a CUCM User Account

ScienceLogic recommends that you create a Cisco Unified CM user account that will be used only by SL1 to access the AXL web service. To create a user account in Cisco Unified CM that can access only the AXL web service, perform these two steps:

- Create a user account.
- Create a user group that includes the user account and has permission to access only the AXL web service.

To create a new Cisco Unified CM user group and user account, perform the following steps:

1. In a browser window, navigate to the following address:
`https://ip-address-of-CM-system:8443/ccmadmin/showHome.do`
2. Log in to the Cisco Unified CM Administration site as an administrator.
3. In the navigation bar at the top-left of the page, hover over **User Management**, then select **Application User**. The **Find and List Users** page appears:



4. Click the [+ Add New] button. The **Application User Configuration** page appears:

The screenshot shows the Cisco Unified CM Administration interface for configuring an application user. The page title is "Application User Configuration". At the top, there is a navigation menu with items: System, Call Routing, Media Resources, Advanced Features, Device, Application, User Management, Bulk Administration, and Help. Below the navigation is a "Save" button. The main content area is divided into three sections: "Status", "Application User Information", and "Device Information".

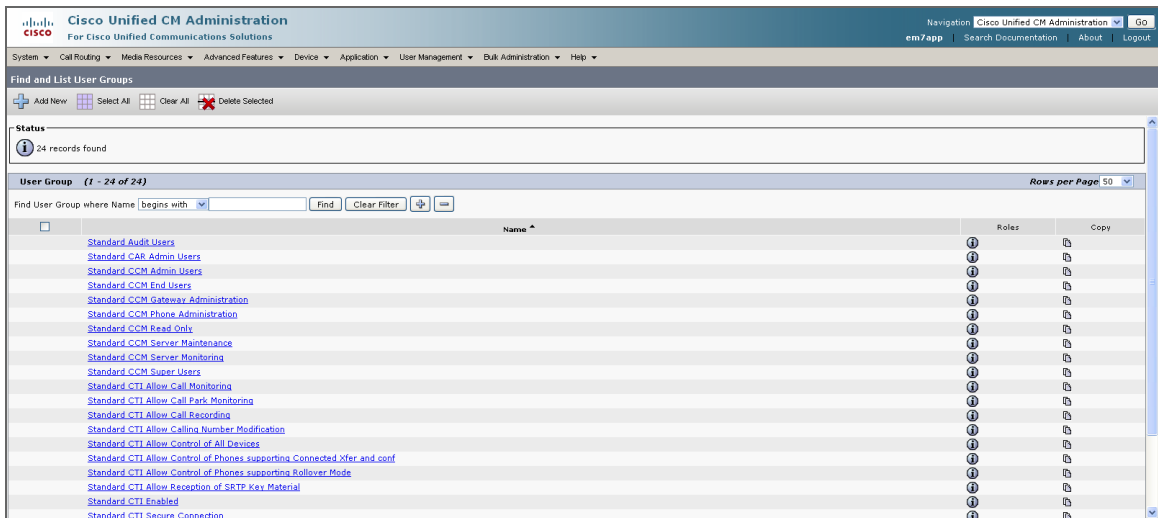
Status
Status: Ready

Application User Information
User ID* [text input]
Password [text input]
Confirm Password [text input]
Digest Credentials [text input]
Confirm Digest Credentials [text input]
Presence Group* [Standard Presence group (dropdown)]
 Accept Presence Subscription
 Accept Out-of-dialog REFER
 Accept Unsolicited Notification
 Accept Replaces Header

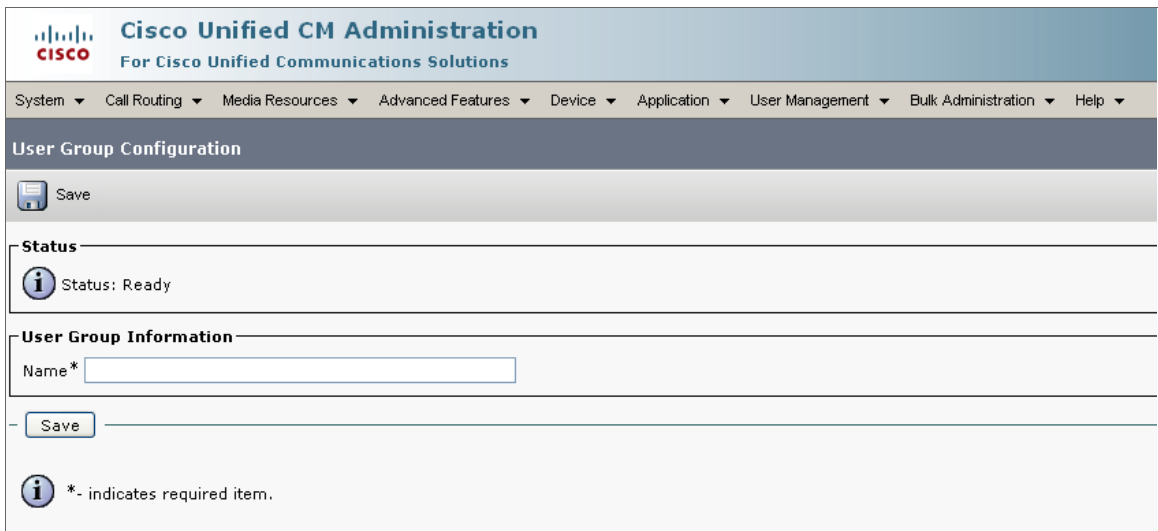
Device Information
Available Devices [Assistant_RP, SEP000F909341F2, SEP001A6C8AC697, SEP04C5A4B0AD9F, SEP44E4D945EF47]
Controlled Devices [empty list]
[Find more Phones] [Find more Route Points]

5. Supply values in the following fields:
 - **User ID**. Type a username for the new user.
 - **Password**. Type a password for the new user.
 - **Confirm Password**. Type the password for the new user again.
6. Click the **[Save]** button.

- In the navigation bar at the top-left of the page, hover over **User Management**, then select **User Group**. The **Find and List User Groups** page appears:

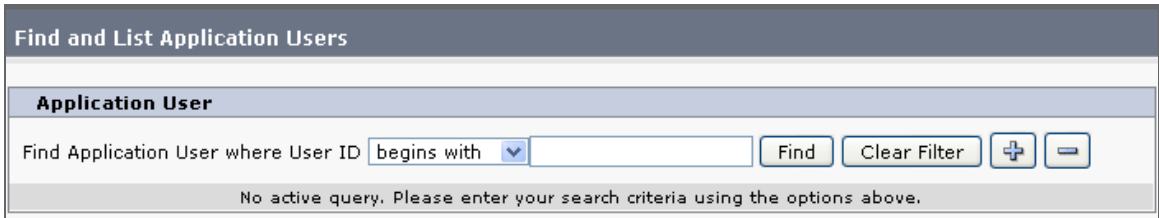


- Click the [+ Add New] button. The **User Group Configuration** page appears:

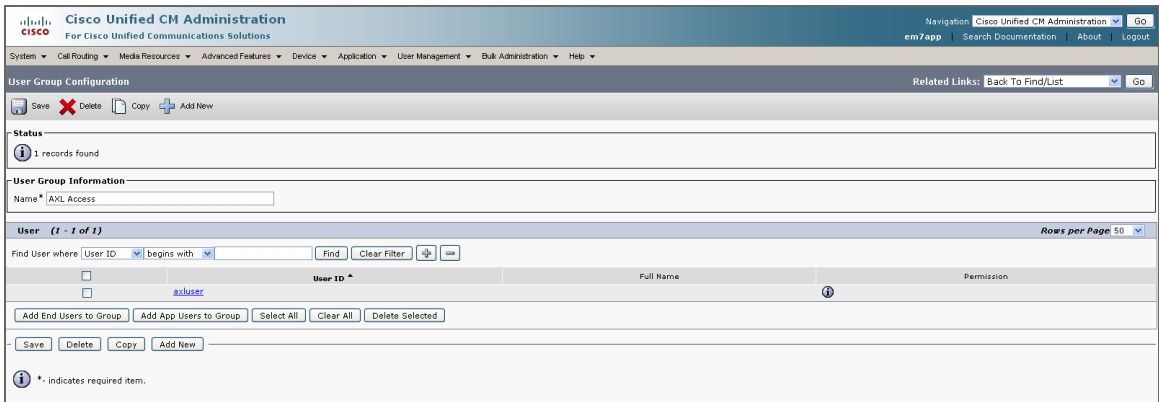


- In the **Name** field, type a name for the user group. For example, you could call the user group "AXL Access".
- Click the **[Save]** button.

- Click the **[Add App Users to Group]** button. The **Find and List Application Users** window appears:



- Click the **[Find]** button. In the list of users, select the checkbox for the user account that you created, then click the **[Add Selected]** button at the bottom of the page.
- The **Find and List Application Users** window closes. In the **User Group Configuration** page, the user account is included in the list of users:



14. In the **Related Links** drop-down list at the top-right hand corner of the page, select *Assign Role to User Group*, and then click the **[Go]** button. The **User Group Configuration** page appears:

The screenshot shows the Cisco Unified CM Administration interface for User Group Configuration. At the top, there is a navigation menu with options: System, Call Routing, Media Resources, Advanced Features, Device, Application, User Management, Bulk Administration, and Help. The main heading is "User Group Configuration". Below this, there is a "Save" button. The "Status" section shows "Status: Ready". The "User Group Information" section shows "Name* AXL Access". The "Role Assignment" section features a "Role" dropdown menu, an "Assign Role to Group" button, and a "Delete Role Assignment" button. At the bottom, there is another "Save" button and three informational messages: "i *- indicates required item.", "i **The role Standard CCM Admin Users must be assigned to a user group to enable its members to logon to CCMAdmin web site", and "i ***The role Standard CCM End Users must be assigned to a user group to enable its members to logon to CCMUser web site".

15. Click the **[Assign Role to Group]** button. The **Find and List Roles** window appears:

The screenshot shows the "Find and List Roles" window. It has a title bar "Find and List Roles" and a "Role" section. Below this, there is a search interface with the text "Find Role where" followed by two dropdown menus: "Name" and "begins with". There is an empty text input field for search criteria, a "Find" button, a "Clear Filter" button, and two small square buttons with "+" and "-" symbols. Below the search fields is a dropdown menu with the text "Select item or enter search text". At the bottom of the window, there is a message: "No active query. Please enter your search criteria using the options above."

16. Click the **[Find]** button. A list of roles appears:

Find and List Roles

Select All Clear All Add Selected Close

Status
 ⓘ 39 records found

Role (1 - 39 of 39) Rows per Page 50

Find Role where Name begins with Find Clear Filter + -
 Select item or enter search text

<input type="checkbox"/>	Name ^	Application	Description	Copy
<input type="checkbox"/>	Standard AXL API Access	Cisco Call Manager AXL Database	Access the AXL APIs	
<input type="checkbox"/>	Standard Admin Rep Tool Admin		Administer CAR	
<input type="checkbox"/>	Standard Audit Log Administration	Cisco Call Manager Serviceability	Serviceability Audit Log Administration	
<input type="checkbox"/>	Standard CCM Admin Users		All users with access to CCM web site	
<input type="checkbox"/>	Standard CCM End Users		Access to CCM User Option Pages	
<input type="checkbox"/>	Standard CCM Feature Management	Cisco Call Manager Administration	Standard CCM Feature Management	
<input type="checkbox"/>	Standard CCM Gateway Management	Cisco Call Manager Administration	Standard CCM Gateway Management	
<input type="checkbox"/>	Standard CCM Phone Management	Cisco Call Manager Administration	Standard CCM Phone Management	
<input type="checkbox"/>	Standard CCM Route Plan Management	Cisco Call Manager Administration	Standard CCM Route Plan Management	
<input type="checkbox"/>	Standard CCM Service Management	Cisco Call Manager Administration	Standard CCM Service Management	
<input type="checkbox"/>	Standard CCM System Management	Cisco Call Manager Administration	Standard CCM System Management	
<input type="checkbox"/>	Standard CCM User Management	Cisco Call Manager Administration	Standard CCM User Management	

17. Select the checkboxes for the following roles:

- *Standard AXL API Access*
- *Standard CCM Admin Users*
- *Standard SERVICEABILITY Read Only*

18. Click the **[Add Selected]** button at the bottom of the page.

19. The **Find and List Roles** window closes. In the **User Group Configuration** page, the **Roles** field includes the *Standard AXL API Access* role:

The screenshot displays the Cisco Unified CM Administration interface for User Group Configuration. The top navigation bar includes System, Call Routing, Media Resources, Advanced Features, Device, Application, User Management, Bulk Administration, and Help. The main content area is titled 'User Group Configuration' and contains a 'Save' button. Below this is a 'Status' section showing 'Status: Ready'. The 'User Group Information' section shows 'Name * AXL Access'. The 'Role Assignment' section features a dropdown menu with three roles: 'Standard AXL API Access', 'Standard CCM Admin Users', and 'Standard SERVICEABILITY Read Only'. To the right of the dropdown are two buttons: 'Assign Role to Group' and 'Delete Role Assignment'. At the bottom left is another 'Save' button. Below the 'Save' button are three information icons with text: '*- indicates required item.', '**The role Standard CCM Admin Users must be assigned to a user group to enable its members to logon to CCMAdmin web site', and '***The role Standard CCM End Users must be assigned to a user group to enable its members to logon to CCMUser web site'.

20. Click the **[Save]** button.

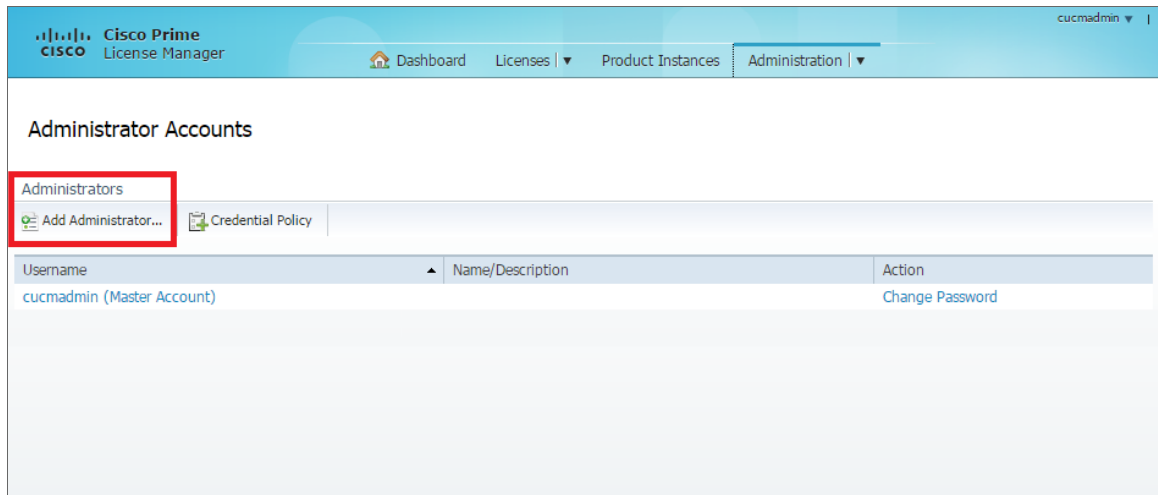
Configuring Prime License Manager

If you want to monitor Cisco Unified CM license information from Cisco Prime License Manager (PLM), you must create an administrator user account that SL1 can use to access PLM.

To create an administrator user in PLM:

1. In a browser window, navigate to the following address:
`https://ip-address-of-plm-server/elm-admin/`
2. Log in to the Cisco PLM site as an administrator.
3. In the **Administration** drop-down menu, select *Administrator Accounts*.

4. Click the **[Add Administrator]** button.



5. In the **Add Administrator Account** modal page, make entries in the following fields:

The screenshot shows a modal window titled 'Add Administrator Account'. At the top, it states '- The minimum password length is 1.' Below this, there are four input fields: 'Name/Description:', '*Username:', '*Password:', and '*Re-enter Password:'. At the bottom right of the modal, there are 'OK' and 'Cancel' buttons.

- **Name/Description.** Type a name or description for the account.
- **Username.** Type the account username.
- **Password.** Type the account password.
- **Re-enter Password.** Type the account password again.

6. Click **[OK]**.

Prerequisites for Monitoring Cisco Email Security Appliances

To configure SL1 to monitor Cisco Email Security Appliances using the *Cisco: ESA PowerPack*, you must first have the following information about the appliance that you want to monitor:

- The appliance's IP address.
- The appliance's SNMP community string.

NOTE: For more information about the *Cisco: ESA PowerPack*, see the *Monitoring Cisco Email Security Appliances* manual.

Chapter

13

Cisco: Hyperflex

NOTE: The *Cisco: Hyperflex PowerPack* supports only HyperFlex API version 2.5 and later.

Prerequisites for Monitoring Cisco HyperFlex

To configure SL1 to monitor Cisco HyperFlex using the *Cisco: Hyperflex PowerPack*, you must have the following information about the HyperFlex data clusters that you want to monitor:

- The Hyperflex Cluster Management IP Address
- SNMP community strings for the voice mailboxes

NOTE: For more information about the *Cisco: Hyperflex PowerPack*, see the *Monitoring Cisco Hyperflex* manual.

Chapter

14

Cisco: Meeting Server

Prerequisites for Monitoring Cisco Meeting Server57

Prerequisites for Monitoring Cisco Meeting Server

To monitor the Cisco Meeting Server, you must be able to access both the Cisco Meeting Server Mainboard Management Processor (MMP) and the Cisco Meeting Server API. Accessing the MMP requires an account with admin access. If you wish to create a new user with admin access, refer to the section "MMP User Account Commands" in the [Cisco Meeting Server MMP Command Line Reference](#) document.

You access the Cisco Meeting Server MMP through SSH, while you access the Cisco Meeting Server API through HTTPS.

- If you can reach both of these through the same IP address, you can typically use a *single Basic/Snippet credential*.
- If the two interfaces have separate IP addresses, or if the API is listening on a port other than 443, you must *create two separate credentials*. In addition, you should include an SNMP credential as part of discovery to correctly classify the device .

NOTE: For more information about the *Cisco: Meeting Server PowerPack*, see the *Monitoring Cisco Meeting Server* manual.

Chapter

15

Cisco: Meraki [API]

Overview

The following sections describe how to configure Cisco Meraki devices for monitoring by SL1 using the *Cisco: Meraki [API] PowerPack* and the Meraki API:

<i>Generating a Cisco Meraki API Key</i>	58
<i>Disabling Asynchronous Dynamic Application Collection</i>	61
<i>Re-enabling Asynchronous Dynamic Application Collection</i>	61
<i>Creating Events from Cisco Meraki Emails</i>	62
<i>Formatting Inbound Emails</i>	63

NOTE: For more information about the *Cisco: Meraki [API] PowerPack*, see the *Monitoring Cisco Meraki (API)* manual.

Generating a Cisco Meraki API Key

To configure Cisco Meraki for monitoring using the Meraki API, you must first generate an API key for a read-only Meraki user. You will then enter this user's API key in the Basic/Snippet credential you create in SL1 to monitor Meraki.

NOTE: If the read-only user has access to multiple organizations, then SL1 can discover all of those organizations with a single discovery session. In this scenario, each organization is created as a separate Cloud Controller in SL1.

However, if you want each Meraki organization to have its own corresponding ScienceLogic organization in SL1, ScienceLogic recommends creating a unique read-only user account and API key for each organization in Meraki. You can then create separate credentials in SL1 for each Meraki organization using those unique API keys, and then use those credentials to run separate discovery sessions for each organization.

To create a read-only user:

1. Log in to the Cisco Meraki web interface.
2. Go to **Organization > Administrators**, and then click the **[Add admin]** button.
3. On the **Create administrator** page, complete the following fields:

The screenshot shows a 'Create administrator' modal window. It contains the following elements:

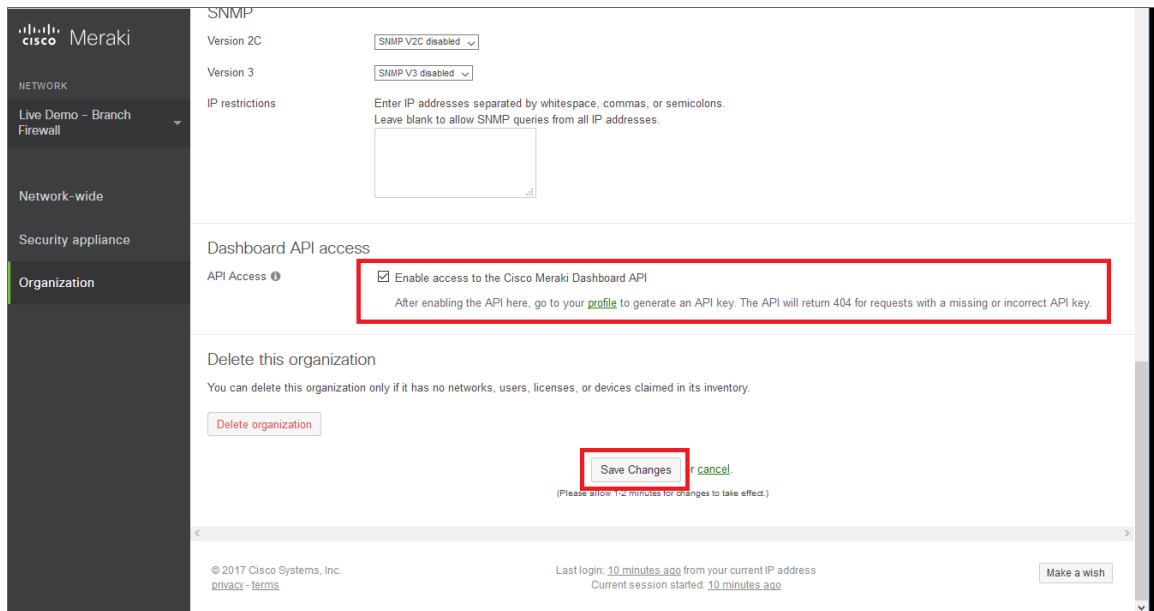
- Name:** A text input field.
- Email:** A text input field.
- Organization access:** A dropdown menu currently set to 'Read-only'.
- Table:** A table with two columns: 'Target' and 'Access'. Below the table is a green link: '+ Add access privileges'.
- Footer:** A 'privacy' link on the left, and 'Close' and 'Create admin' buttons on the right.

- **Name.** Type the user's name.
 - **Email.** Type the user's email address.
 - **Organization access.** Select *Read-only*.
4. Click **[Create admin]**. Cisco Meraki sends an email to the email address provided, describing how the user can complete the registration process. The user must complete those steps before generating the API key.

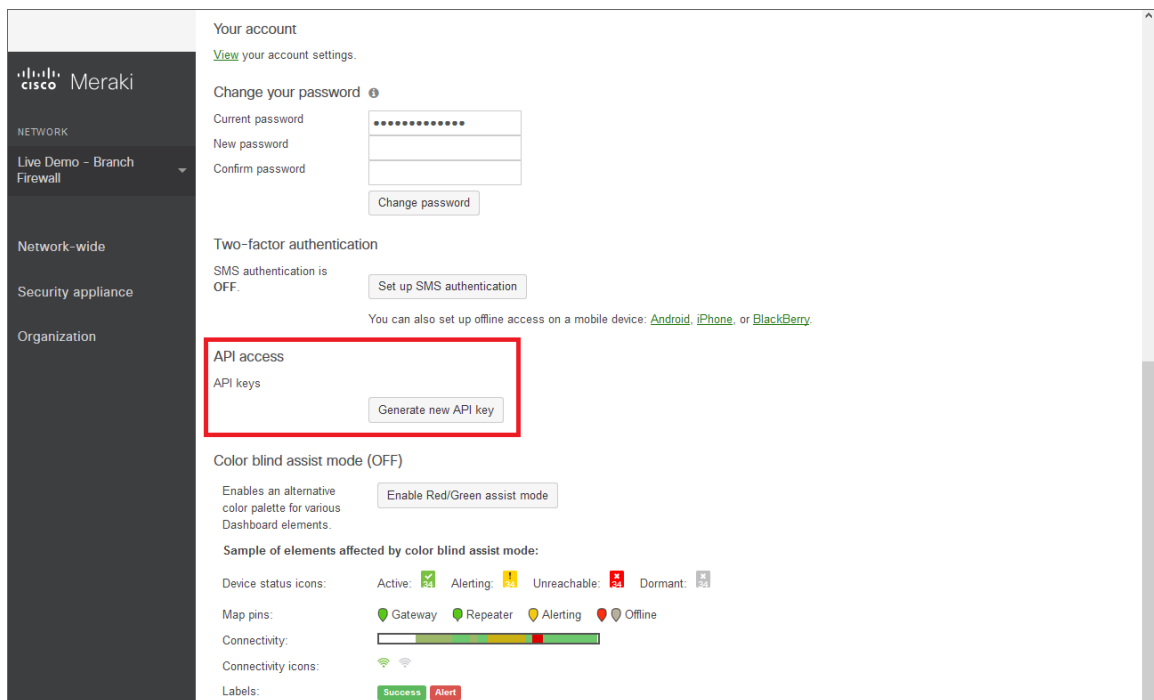
To generate a Cisco Meraki API key for that read-only user:

1. Log in to the Cisco Meraki web interface as the read-only user.

2. Go to **Organization > Settings:**



3. In the **Dashboard API access** section, select the **Enable access to the Cisco Meraki Dashboard API** checkbox.
4. Click the **Save Changes** button.
5. Click the **profile** link in the **Dashboard API access** section.
6. In your user profile, navigate to the **API access** section and click the **Generate new API key** button.



7. In the **API access** section, the API key appears. Copy and save the key value.

NOTE: API keys are visible only to the user that created them.

Disabling Asynchronous Dynamic Application Collection

If the Meraki system you want to monitor consists of more than 200 devices, you must disable the "Data Collection: Async Dynamic App Collection" process before discovering your Meraki system.

NOTE: Disabling asynchronous Dynamic Application collection increases the amount of time it takes the ScienceLogic platform to discover all of the component devices in your Meraki system.

To disable asynchronous Dynamic Application collection:

1. Go to the **Process Manager** page (System > Settings > Admin Processes, or System > Settings > Processes in the SL1 classic user interface).
2. Use the **Process Name** filter field to search for the "Data Collection: Async Dynamic App Collection" process, and then click its wrench icon (🔧). The **Process Editor** page appears.

The screenshot shows the 'Process Editor' window for the process 'Data Collection: Async Dynamic App Collection'. The 'Operating State' dropdown menu is highlighted with a red box and set to 'Disabled'. Other visible settings include 'Program File' (async_dynamic_collect.py), 'Frequency' (Asynchronous), 'Async Throttle' (2), and 'Time Factor (Mins.)' (15). On the right, 'Appliance Types' are listed with checkboxes: All-In-One Server [1] (checked), Database [2] (checked), Administration Portal [3] (unchecked), Customer Portal [4] (unchecked), Data Collection Unit [5] (checked), Message Collection Unit [6] (unchecked), and Integration Server [7] (unchecked). A 'Save' button is at the bottom.

3. In the **Operating State** field, select *Disabled*.
4. Click **[Save]**.

Re-enabling Asynchronous Dynamic Application Collection

If you no longer want to monitor Meraki devices in SL1 and you want to return the system to its original state with asynchronous Dynamic Application collection re-enabled, you must first delete all Meraki devices from the platform. You must then clear the Database Server or Data Collector of any asynchronous processes that are already queued. Failing to do these steps can result in the platform ceasing all data collection until those asynchronous processes are executed.

To re-enable asynchronous Dynamic Application collection:

1. Navigate to the Database Server by typing "<IP address>:8008" into your browser address bar.

2. Log in to the Database Server. The phpMyAdmin browser appears.
3. Select the database from the drop-down **Database** field, and then select the **master_logs** database.
4. In the **master_logs** database, select the **spool_process** table on the left menu, and then click the **[SQL]** tab.
5. Run the following query to clear out the processes on the database:

```
DELETE FROM 'spool_process' WHERE 'proc' = 129 AND 'state' != 0;
```


6. Click **[OK]** at the prompt. Many rows should have been deleted from the table.


If you are using a distributed ScienceLogic system, continue with step 7. Otherwise, go to step 14.

7. In the left menu of the phpMyAdmin browser, select the Data Collector appliance where Meraki devices were discovered.

If the IP address of the Data Collector appears in the upper left-hand corner of the phpMyAdmin browser, go to step 12. Otherwise, if you receive a MySQL error message that your access is denied, continue with step 8.

8. In the Database Server, navigate to the **Master** database and then select the **system_settings_licenses** table.
9. Click **[Browse]** in the upper left-hand side of the page and then identify the Data Collector appliance.
10. Click the **edit** button for the Data Collector:

<input type="checkbox"/>			3	5	SL_ISO1_CU	collector unit: 10.2.8.72	8.5.0	2119	80500002119
--------------------------	---	---	---	---	------------	------------------------------	-------	------	-------------

11. Locate the **db_user** and **db_pass** fields. In those fields, type the same credentials as the Database Server.
12. Click **[Go]**. Wait a few seconds before trying to access the Data Collector in the phpMyAdmin browser. When you do so, the IP address of the Data Collector should appear in the upper left-hand corner of the phpMyAdmin browser.
13. Repeat steps 3-6 on the Data Collector. If successful, many rows should have been deleted from the **spool_process** table.
14. In SL1, go to the **Process Manager** page (System > Settings > Admin Processes, or System > Settings > Processes in the SL1 classic user interface).
15. Use the **Process Name** filter field to search for the "Data Collection: Async Dynamic App Collection" process, and then click its wrench icon (). The **Process Editor** page appears.
16. In the **Operating State** field, select *Enabled*, and then click **[Save]**.

Creating Events from Cisco Meraki Emails

The *Cisco: Meraki [API]* PowerPack includes Event Policies that can generate events in SL1 based on emails that Cisco Meraki sends to SL1.

For SL1 to process events from inbound emails, you must configure your Meraki devices to send email to SL1 using certain formatting rules.

If configured properly, when SL1 domain receives an email with body text that matches a Meraki network component device name and a subject that matches the regular expression (RegEx) pattern of one of the PowerPack's Event Policies, SL1 will generate an event aligned to that network component device.

NOTE: Events from email are always aligned to network devices, even when the email includes references to one or more sub-component devices below the network device.

CAUTION: The email Event Policies included in the *Cisco: Meraki [API] PowerPack* each have an expiry delay setting that specifies the amount of time after which an active event is automatically cleared from SL1 if the event has not reoccurred. However, SL1 clearing an event for reaching its expiry delay setting does not mean that the initial condition that caused the event has been resolved.

Formatting Inbound Emails

Inbound emails must meet the following requirements to be processed as events by SL1:

- The email must be sent to the following address:

```
notify@SL1-domain-name
```

Where "SL1-domain-name" is one of the fully qualified domain names of the Database Server or All-In-One Appliance that is entered in the **Authorized Email Domains** field in the **Email Settings** (System > Settings > Email) page.

- The "from" address used by the external device must be "alerts-noreply@meraki.com" for non-maintenance events, "support-noreply@meraki.com" for maintenance events, or otherwise match an address defined in the **Originator Address** field in an email redirection policy on the **Mailer Redirection** page Events > Inbound Email, or Registry > Events > Inbound Email in the SL1 classic user interface).
- The email subject line must begin with "Alert for" or "Scheduled maintenance for" and match the regular expression (RegEx) pattern of one of the Event Policies included in the *Cisco: Meraki [API] PowerPack*.
- The email body must include the name of a network device monitored by the SL1 system.

The following RegEx patterns are used:

- For scheduled maintenance emails:

```
(Scheduled maintenance for)\s*((network\s|\d\snetworks\s|in\sorganization\s)"([a-zA-Z0-9_-\.\.]+) .*")
```

- For all other emails:

```
(Alert for)\s*([a-zA-Z0-9_-\.\.]+)\s*
```

NOTE: There must be a space between the RegEx pattern and the IP address, hostname, or device ID.

NOTE: The Event Policies included in the *Cisco: Meraki [API] PowerPack* **do not** include RegEx patterns "out of the box". Users can add or modify Event Policy RegEx patterns to best suit their needs.

NOTE: Emails that do not match the RegEx pattern of any Meraki Event Policy will generate a message in the system log. Emails that do not match the name of any component device in SL1 will not generate any events or messages.

NOTE: You can specify how an Event from Email policy will match a RegEx to a device name in the **Behavior Settings** page (System > Settings > Behavior). For more information, see the *Configuring Inbound Email* manual.

Chapter

16

Cisco: Tetration

Configuring Cisco Tetration Analytics for Monitoring

Before you can use SL1 to monitor Cisco Tetration Analytics, you must first generate a Tetration Analytics API key and secret password. You will then use this API key and secret password to create a Basic/Snippet credential that enables SL1 to communicate with and monitor Tetration Analytics clusters.

To configure Cisco Tetration Analytics for monitoring:

1. Log in to the Cisco Tetration Analytics web interface with a **site_admin** or **customer_support** account.
2. Go to **Settings > API Keys**, and then click **[Create API Key]**.
3. Type a **Description** and select the checkbox of the appropriate API key capability.

API Keys

Create API Key

Description

Description (optional)

SW sensor management: API to configure/monitor status of SW sensors

HW sensor management: API to configure/monitor status of HW sensors

Queries on flows: API to query flows in Tetration cluster

At least one capability must be selected.

Create Cancel

4. Click **[Create]**.
5. The API key appears. Copy and save the key value.

NOTE: API keys are visible only to the user that created them.

6. The secret password appears. Copy and save the password value.

WARNING: The secret password value appears only once and cannot be recovered. If you forget or lose the password value, you must generate a new API key with a different password value.

NOTE: For more information about the *Cisco: Tetration PowerPack*, see the **Monitoring Cisco Tetration Analytics** manual.

Chapter

17

Cisco: UC Ancillary

Prerequisites for Monitoring Ancillary Cisco Unified Communications Devices

To configure SL1 to monitor ancillary Cisco Unified Communications (UC) devices using the *Cisco: UC Ancillary PowerPack*, you must have already properly installed and configured the ancillary Cisco UC devices that you want to monitor. You must also note the following information, as appropriate, for each of the ancillary UC devices you want to monitor:

- SNMP community string
- Secure Shell (SSH) username and password to monitor Cisco voice components

NOTE: For more information about the *Cisco: UC Ancillary PowerPack*, see the *Monitoring Cisco Unified Communications Ancillary Devices* manual.


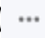
Chapter

18

Cisco: UC VOS Applications

Overview

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

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

The following sections describe how to configure Cisco UC Voice Operating System (VOS) applications for monitoring by SL1 using the *Cisco: UC VOS Applications PowerPack*:

Configuring Cisco UC VOS Applications for Monitoring	69
<i>Configuring SNMP for Cisco VOS Applications</i>	69
<i>Creating the User Account for the Platform Administrative Web Services (PAWS) API</i>	70
<i>Configuring Cisco Unity Connection</i>	71
<i>Configuring Cisco Unified Communications Manager IM and Presence</i>	72
<i>Configuring Cisco Prime License Manager</i>	73
<i>Configuring Cisco Prime Collaboration Deployment</i>	74
<i>Configuring Cisco Collaboration Mediation Fulfillment</i>	75
<i>Configuring Hosted Collaboration Solution Intelligent Loader</i>	75
<i>Configuring Cisco Contact Center Express</i>	75
<i>Configuring Cisco Emergency Responder</i>	75
<i>Configuring Cisco SocialMiner</i>	75

NOTE: For more information about the *Cisco: UC VOS Applications PowerPack*, see the *Monitoring Cisco UC Voice Operating System (VOS) Applications* manual.

Configuring Cisco UC VOS Applications for Monitoring

Before performing the other tasks in this chapter, you must create accounts for the different Cisco VOS applications that you want to monitor in SL1. The following sections describe how to configure the Cisco VOS applications.

Configuring SNMP for Cisco VOS Applications

SL1 uses SNMP to collect information about the following Cisco VOS applications:

- Cisco Contact Center Express (CCX)
- Cisco Unity Connection (CUC) servers
- Cisco Hosted Collaboration Mediation for Fulfillment (HCM-F)
- Cisco HCS Intelligent Loader
- Cisco IM & Presence (IM&P) servers (optional)
- Cisco Emergency Responder
- Cisco Prime Collaboration Deployment (PCD) servers (optional, but recommended)
- Cisco SocialMiner

To configure SNMP for Cisco VOS applications:

1. Log in as an administrative user to the command-line interface of the Cisco VOS application that you want to monitor. You can also use SSH to connect to the application.
2. Run the following command with additional parameters as needed:

```
utils snmp config
```
3. When prompted, add additional SNMP information. The following image displays additional configuration parameters, based on SNMP version (version 1 or 2, or version 3):

```
admin:utils snmp config 1/2c
      utils snmp config 1/2c community-string*
      utils snmp config 1/2c inform*
      utils snmp config 1/2c trap*

admin:utils snmp config 3
      utils snmp config 3 inform*
      utils snmp config 3 trap*
      utils snmp config 3 user*
```

4. For additional SNMP configuration commands and instructions, see the [Cisco Command Line Interface Reference Guide](#).

Creating the User Account for the Platform Administrative Web Services (PAWS) API

To get access to the Platform Administrative Web Services (PAWS) API, you can create a new user account by using the command-line interface on the console of the Cisco VOS application that you want to monitor. You can also use SSH to connect to the application.

You can then use this user account to connect to the following Cisco VOS applications:

- Cisco Contact Center Express (CCX)
- Cisco Unity Connection (CUC) servers
- Cisco Hosted Collaboration Mediation for Fulfillment (HCM-F)
- Cisco HCS Intelligent Loader
- Cisco IM & Presence (IM&P) servers
- Cisco Unified Communications Manager (CUCM)
- Cisco Prime License Manager (PLM)
- Cisco Prime Collaboration Deployment (PCD) servers
- Cisco SocialMiner

To create the PAWS API user account:

1. Log in as an administrative user to the command-line interface of the Cisco VOS application that you want to monitor.
2. To create the new account, run the following command:

```
set account name new_account_username
```
3. The interface prompts you for the privilege level and password for the new account:

```
admin:set account name em7paws

Privilege Levels are:
  Ordinary - Level 0
  Advanced - Level 1

Please enter the privilege level :0
  Please enter the password :*****
    re-enter to confirm :*****
Account successfully created
```

4. Set the privilege level to 0.
5. Type the password, then retype the password to confirm.

6. Newer versions of Cisco Unified Communications products require that new accounts created with the command-line interface must change the password at the first login. This requirement blocks the account from accessing the PAWS API until you change the password. To remove the requirement for this account, run the following command:

```
set password change-at-login disable new_account_username
```

7. To confirm that the user account works with the Cisco PAWS API, log in as an administrator to one of the following addresses:
 - <https://ip-address-of-cisco-application:8443/platform-services/services/ProductService?wsdl>
 - <https://ip-address-of-cisco-application:8443/platform-services/services/ClusterNodesService?wsdl>

NOTE: If you receive a message that the user does not have permission to access a page, then the Cisco VOS application requires a user account like the one you just created to access the PAWS API. You might get this message if you are using Cisco Unified Communications products older than version 9, because those products do not use the PAWS API. In this situation, use the credential setup for non-PAWS API. Also, you cannot use any Dynamic Applications that use the PAWS API, but you can use the SNMP and Application APIs.

Configuring Cisco Unity Connection

You can create a user account for Cisco Unity Connection applications that gives you access to other Cisco APIs such as Administrative XML (AXL), Serviceability, and Real-Time Monitoring Service. You can configure this account using the web-based interfaces for the Cisco applications. This account does not have access to the PAWS API.

NOTE: To create a PAWS API user account for Cisco Unity Connection, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

To create the user account for Cisco Unity Connection:

1. In a browser window, navigate to the following address:

```
https://ip-address-of-cisco-application/cuadmin/home.do
```
2. Navigate to the relevant **Edit Users Basics** page for your version of Cisco Unity Connection (User > Users).
3. Create a new user and complete the fields as needed.
4. Select the role of **Technician** or **System Administrator**.
5. Save the new user account.

6. To confirm that the user account works with the Cisco APIs, log into one of the following addresses:

- `https://ip-address-of-cisco-application:8443/realtimeservice/services/RisPort?wsdl`
- `https://ip-address-of-cisco-application:8443/controlcenterservice/services/ControlCenterServicesPort?wsdl`

7. If you are not prompted for the username and password when testing the addresses, your previous administrative login might still be active. Close the browser and navigate to the addresses again.

Configuring Cisco Unified Communications Manager IM and Presence

You can use the same account for Cisco Unified Communication Manager (CUCM) IM and Presence that you already created for CUCM. If you are creating an account specifically for monitoring IM and Presence, you only need the Standard CCM Server Monitoring Group.

NOTE: Because SL1 does not access the Administrative XML API for IM and Presence, the Standard AXL API Access role is not required.

To create a user account for CUCM IM and Presence:

1. In a browser window, navigate to the Cisco CUCM web interface:

`https://ip-address-of-cisco-cucm/ccmadmin/showHome.do`

2. Navigate to the relevant **User Management** page for your version of Cisco CUCM (User Management > Application User):

The screenshot displays the 'Application User Configuration' interface. At the top, there are navigation buttons: Save, Delete, Copy, and Add New. Below this, the 'Status' section shows 'Status: Ready'. The main section is 'Application User Information', which includes the following fields and options:

- User ID*:
- Password:
- Confirm Password:
- Digest Credentials:
- Confirm Digest Credentials:
- BLF Presence Group*:
- Accept Presence Subscription
- Accept Out-of-dialog REFER
- Accept Unsolicited Notification
- Accept Replaces Header

3. Click the **[Add New]** button and complete the required information for the new user account.

4. In the Permissions Information section, select the **Standard CCM Server Monitoring** and the **Admin-3rd Party API** groups and save the user record.

Permissions Information	
Groups	Admin-3rd Party API Standard CCM Server Monitoring
	View Details
Roles	Standard CCM Admin Users Standard CCMADMIN Read Only Standard SERVICEABILITY
	View Details

NOTE: To create the Level 0 PAWS API user account for CUCM, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#). The discovery process for IM and Presence queries the CUCM servers using this user account to determine the server role (IM and Presence or CUCM). As a result, the PAWS API user account needs to be enabled on the CUCM nodes during discovery for IM and Presence.

Configuring Cisco Prime License Manager

When Cisco Prime License Manager is co-resident with Cisco Unified Communications Manager, this release of the PowerPack cannot monitor Cisco Prime License Manager.

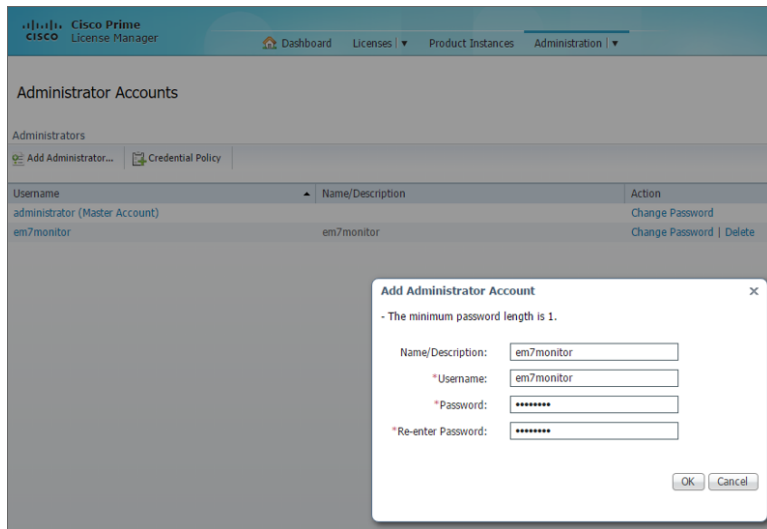
When Cisco Prime License Manager is installed as a standalone system, and is not co-resident with another Cisco product, you can only create administrative users for the application. You can use the existing administrator account or create a new account for monitoring.

To create a user account for Cisco Prime License Manager:

1. In a browser window, navigate to the following address:

`https://ip-address-of-application/elm-admin/faces/main.xhtml`

2. Navigate to the **Administrator Accounts** page for your version of Cisco Prime License Manager (Administration > Administrator Accounts):



3. Click **Add Administrator** and complete the required information.
4. After you create the user account, you can use the following address to confirm that the new account works with the APIs:

`https://ip-address-of-application/elm-admin/faces/license_usage.xhtml?`

NOTE: To create a PAWS API user account for Cisco Prime License Manager, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

Configuring Cisco Prime Collaboration Deployment

To create the PAWS API user account for Cisco Prime Collaboration Deployment, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

Using the PAWS API user account, use SSH to connect to the command-line interface of the application, and then run the following command to get service status:

```
utils service list
```

Configuring Cisco Collaboration Mediation Fulfillment

To create the PAWS API user account for Cisco Collaboration Mediation Fulfillment, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

After you create the account, you can use the following address to confirm that SL1 can monitor Cisco Collaboration Mediation Fulfillment:

```
https://ip-address-of-  
application:8443/controlcenterservice/services/ControlCenterServicesPort?wsdl
```

Configuring Hosted Collaboration Solution Intelligent Loader

Cisco Hosted Collaboration Solution Intelligent Loader requires only a PAWS API user account. To create the PAWS API user account for Cisco Hosted Collaboration Solution Intelligent Loader, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

After you create the account, you can use the following address to confirm that SL1 can monitor Cisco Hosted Collaboration Solution Intelligent Loader:

```
https://ip-address-of-  
application:8443/controlcenterservice/services/ControlCenterServicesPort?wsdl
```

Configuring Cisco Contact Center Express

Cisco Contact Center Express does not let you create additional accounts that can access the Application API. Instead of creating an Application Monitoring user account, you must use the administrative account that was assigned when the product was first installed.

To create the PAWS API user account for Cisco Contact Center Express, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

Configuring Cisco Emergency Responder

You can only use SNMP to monitor the Cisco Emergency Responder. To set up SNMP for the Cisco Emergency Responder, see [Configuring SNMP for Cisco VOS applications](#).

Configuring Cisco SocialMiner

To set up SNMP for Cisco SocialMiner, see [Configuring SNMP for Cisco VOS applications](#).

To create the PAWS API user account for Cisco SocialMiner, see [Creating the User Account for the Platform Administrative Web Services \(PAWS\) API](#).

To use a Social Miner account, make sure that the account has Administrator credentials for API access. You can use an existing SocialMiner administrator account or create a new account for monitoring that has administrator permissions.

NOTE: Because Cisco SocialMiner is a virtual machine that does not support clusters, SL1 creates a cluster for each SocialMiner device during the discovery process. SL1 then uses that cluster to create a component level where it can use the relevant Cisco VOS dynamic applications.

Chapter

19

Cisco: UCS

Overview

The following sections describe how to configure a Cisco Unified Computing System (UCS) Manager for monitoring by SL1 using the *Cisco: UCS PowerPack*:

<i>Prerequisites for Monitoring Cisco UCS Manager</i>	77
<i>Configuring the UCS System</i>	77

NOTE: For more information about the *Cisco: UCS PowerPack*, see the *Monitoring Cisco Unified Computing System (UCS) Manager* manual.

Prerequisites for Monitoring Cisco UCS Manager

To use the Dynamic Component Mapping Dynamic Applications included in the *Cisco: UCS PowerPack*, you must log in to the UCS Manager GUI and create a user account that SL1 can use to access the UCS web service.

Configuring the UCS System

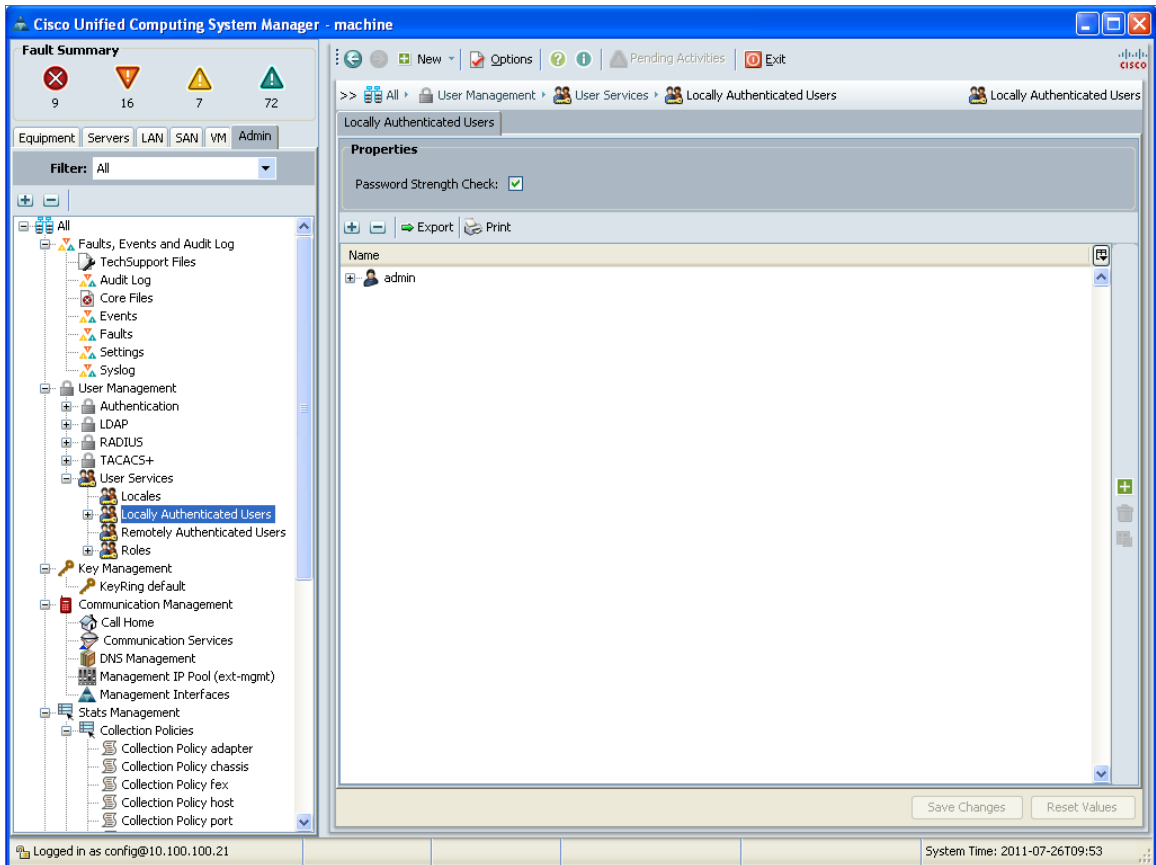
To configure a UCS system for monitoring by SL1, you must:

- Create a user account in UCS that SL1 can use to access the UCS web service
- Enable the CIM XML service

Perform the following steps to complete these tasks:

1. Log in to the UCS Manager GUI as an administrator.

2. At the top of the left pane, click the **[Admin]** tab.
3. In the left pane, go to All > User Management > User Services > Locally Authenticated Users. The **Locally Authenticated Users** page appears in the right pane:



4. Click the green plus icon on the right side of the **Locally Authenticated Users** page. The **Create User** window appears:

Create User

Login ID:

First Name:

Last Name:

Email:

Phone:

Password:

Confirm Password:

Account Status: active inactive

Account Expires:

Roles

- aaa
- admin
- facility-manager
- network
- operations
- read-only
- server-equipment
- server-profile
- server-security
- storage

Locales

OK Cancel

5. Supply values in the following fields:
 - **Login ID.** Enter a username for the user.
 - **Password.** Enter a password for the user.
 - **Confirm Password.** Re-enter the password you entered in the **Password** field.

- **Account Status.** Select *active*.
 - **Account Expires.** Make sure that this checkbox is not selected.
 - **Roles.** To create a read-only user, do not select any checkboxes.
6. Click the **[OK]** button, and then click the **[OK]** button in the confirmation pop-up window.
 7. In the left pane of the UCS Manager GUI, go to All > Communication Management > Communication Services. The **Communication Services** page appears in the right pane.
 8. In the **Admin State** field in the **CIM XML** section, select *Enabled*.

NOTE: Older versions of the UCS software do not include the option to disable the CIM XML service. If the option to enable/disable the CIM XML service does not appear, the service is already enabled.

9. Click the **[Save Changes]** button.

NOTE: When blade servers are replaced in a UCS chassis, and the old blade servers are not properly decommissioned, UCS Manager does not assign new Internal IDs to the new blade servers when they are inserted in the chassis. Instead, UCS Manager assigns an Internal ID of "none" to the new blade servers. This does not cause an error in SL1 if it occurs with only a single blade; however, if more than one blade that you are monitoring is replaced without being decommissioned, multiple blades will have the same Internal ID of "none", which in turn can cause blades to appear under the incorrect chassis or not appear at all in SL1. If this occurs, decommission the affected blades and then reinsert them. For more information, see the section on "Guidelines for Removing and Decommissioning Blade Servers" in the [Cisco UCS documentation](#).

Chapter 20

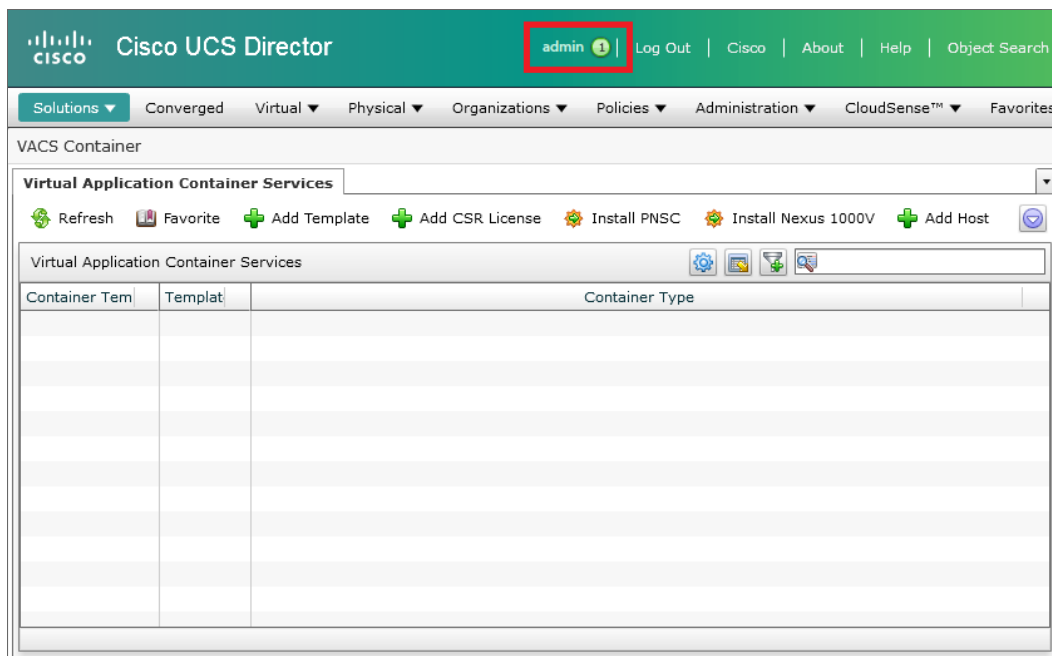
Cisco: UCS Director

Copying the REST API Access Key for a UCS Director Account

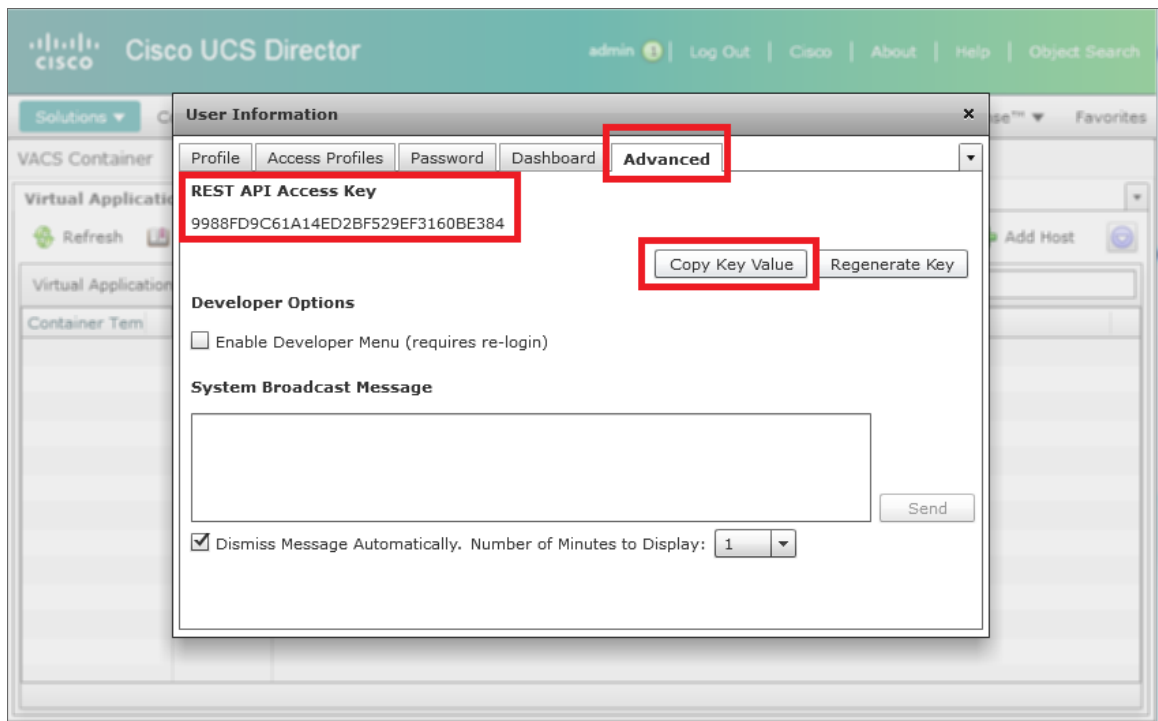
When configuring the Basic/Snippet credential that SL1 uses to discover and monitor UCS Director, you must include the REST API Access Key for a UCS Director administrator user account as the credential password.

To locate and copy the REST API Access Key:

1. Log in to UCS Director as an administrator, and then click the username at the top of the page.



2. The **User Information** modal page appears. Click the **[Advanced]** tab.
3. Click the **[Copy Key Value]** button to copy the REST API Access Key.



NOTE: For more information about the *Cisco: UCS Director PowerPack*, see the *Monitoring Cisco Unified Computing System (UCS) Director* manual.

Chapter

21

Cisco: UCS Standalone Rack Server

Prerequisites for Monitoring Cisco UCS Standalone Rack Servers

In order to monitor Cisco UCS standalone rack servers in SL1 using the *Cisco: UCS Standalone Rack Server PowerPack*, you must know the username and password for a web service user on the rack servers you want to monitor.

NOTE: For more information about the *Cisco: UCS Standalone Rack Server PowerPack*, see the *Monitoring Cisco Unified Computing System (UCS) Standalone Rack Servers* manual.

Chapter

22

Cisco: Unity Express

Prerequisites for Monitoring Cisco Unity Express

To configure the SL1 system to monitor Cisco Unity Express voice mailboxes using the *Cisco: Unity Express PowerPack*, you must first have the following information about the Unity Express voice mailboxes that you want to monitor:

- IP addresses for the voice mailboxes
- SNMP community strings for the voice mailboxes

NOTE: For more information about the *Cisco: Unity Express PowerPack*, see the *Monitoring Cisco Unity Express* manual.

Chapter

23

Cisco: Viptela

Prerequisite for Monitoring Cisco Viptela

To configure the SL1 system to monitor Cisco Viptela resources using the *Cisco: Viptela PowerPack*, you must first know the credentials (username and password) for a user account that has access to the Cisco Viptela system. The user account must have read-all access.

NOTE: For more information about the *Cisco: Viptela PowerPack*, see the *Monitoring Cisco Viptela* manual.

Chapter

24

Cisco: Wireless

Prerequisites for Monitoring Cisco Wireless LAN Controllers

Before you can monitor Cisco wireless LAN controllers using the *Cisco: Wireless PowerPack*, you must have the following information:

- The IP address of the WLC that you want to monitor with SLI
- The settings for an SNMP V2 or SNMP V3 credential that can be used to communicate with the WLC

NOTE: For more information about the *Cisco: Wireless PowerPack*, see the *Monitoring Cisco Wireless LAN Controllers* manual.

Chapter

25

Citrix: Xen

Enabling Performance Metrics for XenServer 6.2.0 and Above

Most performance metrics are disabled by default in Citrix XenServer 6.2.0 and above. Therefore, if you are monitoring XenServer 6.2.0 or above with SL1, you must enable performance metrics on each XenServer host.

NOTE: Performance metrics are enabled by default in XenServer 6.1.0 and below. No additional steps are required to monitor those devices.

To enable performance metrics in XenServer 6.2.0 and above devices:

1. Open the XenServer command line interface.
2. For each XenServer host, enter the following command:

```
xe-enable-all-plugin-metrics true
```

NOTE: For more information about the *Citrix: Xen PowerPack*, see the *Monitoring Citrix XenCenter* manual.

Chapter

26

CouchBase

Prerequisites for Monitoring Couchbase

To configure SL1 to monitor Couchbase servers and component devices using the *CouchBase PowerPack*, you must have the login credentials for a user with administrative access to the Couchbase server.

NOTE: For more information about the *CouchBase PowerPack*, see the *Monitoring Couchbase* manual.

Chapter

27

Dell EMC: Isilon

Prerequisites for Monitoring Dell EMC Isilon

To configure the SL1 system to monitor Dell EMC Isilon storage arrays using the *Dell EMC: Isilon PowerPack*, you must have already installed and configured the storage arrays that you want to monitor.

If you are using a Secure Sockets Layer (SSL) certificate to communicate with the Isilon storage arrays you are monitoring, you must add an Isilon SSL certificate on your SL1 appliance in the following file:

```
/var/lib/em7/content/silo_core_rest/certs.crt
```

NOTE: If you are not using an SSL certificate to communicate with the Isilon storage arrays, then you do not need to add a certificate. For more information about installing an SSL certificate, see the manual *Installing an SSL Certificate*.

Additionally, you should take note of the SNMP community string used by the Isilon storage arrays you want to monitor.

NOTE: For more information about the *Dell EMC: Isilon PowerPack*, see the *Monitoring Dell EMC Isilon* manual.

Chapter

28

Dell EMC: Unity

Prerequisites for Monitoring Dell EMC Unity

Before you can monitor Dell EMC Unity systems using the *Dell EMC: Unity PowerPack*, you must have the following information about the Unisphere REST API:

- Username and password for a user with access to the Unisphere REST API
- IP address for the Unisphere REST API

NOTE: For more information about the *Dell EMC: Unity PowerPack*, see the *Monitoring Dell EMC Unity* manual.

Dell EMC: VMAX and PowerMax Unisphere API

Prerequisites for Monitoring Dell EMC VMAX and PowerMax Systems

Before you can monitor Dell EMC VMAX and PowerMax systems using the *Dell EMC: VMAX and PowerMax Unisphere API* PowerPack, you must have the following information about the Unisphere API that has already been properly configured:

- Username and password for a user with access to the Unisphere REST API
- IP address and port for the Unisphere

NOTE: For more information about the *Dell EMC: VMAX and PowerMax Unisphere API* PowerPack, see the *Monitoring Dell EMC VMAX and PowerMax Unisphere API* manual.

Chapter

30

Dell EMC: XtremIO

Overview

The following sections describe how to configure and discover Dell EMC XtremIO storage devices for monitoring by SL1 using the *Dell EMC: XtremIO PowerPack*:

Prerequisites for Monitoring Dell EMC XtremIO	92
Configuring Traps with Dell EMC XtremIO	93

NOTE: For more information about the *Dell EMC: XtremIO PowerPack*, see the *Monitoring Dell EMC XtremIO* manual.

Prerequisites for Monitoring Dell EMC XtremIO

Before you can monitor Dell EMC XtremIO storage devices in SL1 using the *Dell EMC: XtremIO PowerPack*, you must have already properly installed and configured the XtremIO storage devices that you want to monitor.

In addition, you must create a read-only user in the XtremIO Management Server (XMS) with the following user permissions:

- **User Name:** Type the XMS user's name.
- **Authentication.** Select the **By Password** checkbox.
- **Password:** Type and then confirm the XMS user's password.

You can also configure LDAP authentication for this account.

Finally, take note of the SNMP community string used by the XtremIO storage devices you want to monitor.

For more information about these configuration processes, see the Dell EMC XtremIO documentation.

Configuring Traps with Dell EMC XtremIO

To send alerts to SL1, SNMP traps must be enabled and configured on the Dell EMC XtremIO storage array. When configuring these traps, use the IP address of the ScienceLogic Message Collector, Data Collector, or All-In-One Appliance responsible for monitoring the system as the destination IP.

For more information, see the Dell EMC XtremIO documentation.

Chapter

30

Docker

Overview

The following sections describe how to configure the Docker platform for discovery by SL1 using the *Docker PowerPack*:

Prerequisites for Monitoring Docker	94
Enabling the Docker API	95

NOTE: For more information about the *Docker PowerPack*, see the *Monitoring Docker* manual.

Prerequisites for Monitoring Docker

If you are using Secure Shell (SSH) to monitor Docker or Kubernetes nodes in conjunction with the *Kubernetes PowerPack*, you must install cURL 7.40 or greater on all of the Docker hosts that you want to monitor, prior to discovery. You must then run the following cURL commands on each of those hosts:

- `curl --unix-socket /var/run/docker.sock http://docker/containers/json`
- `curl --unix-socket /var/run/docker.sock http://docker/containers/[container_id]/json`
- `curl --unix-socket /var/run/docker.sock http://docker/containers/[container_id]/stats?stream=0`

If you are using a Basic/Snippet credential, before you can monitor the Docker platform and its component devices in SL1 using the *Docker PowerPack*, you must first follow the instructions in the [Enabling the Docker API](#) section. These steps enable the Dynamic Applications in the *Docker PowerPack* to communicate with and gather data from the Docker API.

NOTE: You do not need to enable the API if you are using SSH to monitor Docker.

WARNING: If you choose to enable the API when monitoring Docker versions through 18.06.1-ce-rc2, be aware that a vulnerability exists. The API endpoints behind the 'docker cp' command are vulnerable to a symlink-exchange attack. (CVE-2018-15664).

Enabling the Docker API

Before you discover Docker components using the *Docker PowerPack*, you must first enable the Docker API. This section describes how to do so for Windows, CentOS, Red Hat Enterprise Linux (RHEL), and Oracle Linux operating systems.

NOTE: If you are using SSH to monitor Docker, skip this section and go to the [Creating an SSH/Key Credential](#) section.

Windows

To enable the Docker API for Windows using the Docker Toolbox:

1. Start Docker Quickstart Terminal.
2. To determine the IP address of the Docker host machine, type the following command:

```
$ docker-machine ip
```
3. Log in to the host machine:

```
$ docker-machine ssh
```
4. Navigate to Boot2Docker:

```
$ cd /var/lib/boot2docker
```
5. Edit the Boot2Docker profile:

```
$ sudo vi profile
```
6. In the profile, change "DOCKER_HOST" to "DOCKER_HOST='-H tcp://0.0.0.0:[port number]'", and set DOCKER_TLS=no.
7. Exit the SSH session, and then restart Docker:

```
$ exit  
$ docker-machine restart
```
8. To verify that the Docker API is accessible, open a browser and navigate to `http://[IP address]:[port number]/version`.

If the Docker API is successfully enabled, the version returns something similar to the following:

```
{"Version":"17.10.0-ce","ApiVersion":"1.33","MinAPIVersion":"1.12","GitCommit":  
"f4fffd25","GoVersion":"go1.8.3","Os":"linux","Arch":"amd64","KernelVersion":  
"4.4.93-boot2docker","BuildTime":"2017-10-17T19:05:23.000000000+00:00"}
```

CentOS

To enable the Docker API for CentOS:

1. Log in to the command-line interface of the server running Docker and navigate to systemd/system:

```
$ cd /etc/systemd/system
```

2. Create a new "docker.service.d" folder, then navigate to that folder:

```
$ mkdir docker.service.d  
$ cd docker.service.d
```

3. Create a new docker.conf file:

```
$ vi docker.conf
```

4. Type the following:

```
INSERT  
[Service]  
ExecStart=  
ExecStart=/usr/bin/dockerd -H tcp://0.0.0.0:[port number] -H  
unix://var/run/docker.sock
```

5. Reload daemon, restart Docker, and open the port on the firewall by typing the following:

```
$ systemctl daemon-reload  
$ systemctl restart docker  
$ firewall-cmd --add-port=[port number]/tcp
```

6. Verify that the Docker API is accessible by typing the following:

```
$ *curl http://localhost:[port number]/version*
```

If the Docker API is successfully enabled, the version returns something similar to the following:

```
{"Version":"17.06.1-ce","ApiVersion":"1.30","MinAPIVersion":"1.12","GitCommit":  
:"874a737","GoVersion":"go1.8.3","Os":"linux","Arch":"amd64","KernelVersion":  
"3.10.0-514.26.2.el7.x86_64","BuildTime":"2017-08-17T23:01:50.155177940+00:00"}
```

RHEL 7 and Oracle Linux 7

To enable the Docker API for RHEL 7 or Oracle Linux 7:

1. Log in to the command-line interface of the server running Docker and navigate to systemd/system:

```
$ cd /etc/systemd/system
```

2. Edit the service.docker file:

```
$ sudo vi docker.service
```


3. Create or edit the file to ensure that it has a [Service] section and a line that starts with "ExecStart=/usr/bin/dockerd". Add "-H tcp://0.0.0.0:[port number] -H unix:///var/run/docker.sock" so that the updated line looks like this:

```
ExecStart=/usr/bin/dockerd -H tcp://0.0.0.0:4243 -H unix:///var/run/docker.sock
```

4. Open the firewall port, if needed, and then reload daemon and restart Docker by typing the following:

```
$ sudo firewall-cmd --add-port=[port number]/tcp
$ sudo firewall-cmd --reload
$ sudo systemctl daemon-reload
$ sudo systemctl restart docker
```

5. Verify that the Docker API is accessible by typing the following:

```
$ curl http://[IP address]:[port number]/version
```

If the Docker API is successfully enabled, the version returns something similar to the following:


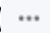
```
{"Version":"17.06.2-ee-4","ApiVersion":"1.30","MinAPIVersion":"1.12","GitCommit":
"dd2c358","GoVersion":"go1.8.3","Os":"linux","Arch":"amd64","KernelVersion":
"3.10.0-514.el7.x86_64","BuildTime":"2017-10-12T16:19:56.386620861+00:00"}
```

Chapter

31

Dynatrace

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

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

Generating a Dynatrace API Token

To configure the SL1 system to monitor Dynatrace resources using the *DynatracePowerPack*, you must first generate a Dynatrace API token.

To do so:

1. Log in to your Dynatrace portal. On the left menu, click **Settings > Integration > Dynatrace API**. The **Dynatrace API** page appears.
2. Click the **[Generate Token]** button.
3. In the blank box that appears, type a token name, and then activate (at a minimum) the "Access problem and event feed, metrics, topology, and RUM JavaScript tag management" permission.
4. Click **[Generate]** to generate the API token.

TIP: You can click the **[Copy]** button next to the generated token to copy the token to your computer's clipboard.

5. The newly generated API token appears in your list of API tokens. Ensure that the **Disable/enable** switch is activated.

6. Optionally, if you want to verify the token, you can use an API tool like Postman or cURL to send a GET request for your Dynatrace environment, and then attach the token to the Api-Token realm for the Authorization HTTP header. For example:

```
curl --request GET \  
  --url https://<Hostname>/e/<Environment-ID>.live.dynatrace.com/api/v1/time \  
  --header 'Authorization: Api-Token <generated API token>' \  

```



NOTE: For more information about the *DynatracePowerPack*, see the **Monitoring Dynatrace** manual.

To configure SL1 to monitor Dynatrace devices, you must first create a SOAP/XML credential. This credential allows the Dynamic Applications in the *Dynatrace PowerPack* to use your Dynatrace user account to retrieve information from the *Dynatrace* environment and component devices.

The PowerPack includes an example SOAP/XML credential (**Dynatrace Credential Example**) that you can edit for your own use.

Filtering Partitions from Host Components

You can filter out partitions from host components in the "Dynatrace: Host Disk Performance" Dynamic Application. To do this, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Locate the "Dynatrace: Host Disk Performance" Dynamic Application and click its wrench icon () .
3. Click on the **[Snippets]** tab.
4. In the **Snippet Editor & Registry** page, click the wrench icon () for the "host_disk_performance" snippet.
5. Edit the `partitions=["/var/lib/docker"]` line to specify the partition(s) you want to filter out. You can specify more than one partition by separating them with commas and enclosing the partitions in quotation marks. Remove the partition if you want to collect data for it.

Dynamic Applications [1729] | Snippet Editor & Registry | Editing Snippet [2121] Guide

Snippet Name	Active State	Required
host_disk_performance	[Enabled]	[Required - Stop Collection]

```

Snippet Code
-- metric in oid_response:
transformed_response = oid_response.get(metric)
else:
    logger.debug("making request {}".format(oid))
    request_path = oid.split('.')[0]
    params = dynatrace_perf.build_params(oid,
                                        relative_time=RELATIVE_TIME,
                                        entities=hosts)
    response = dynatrace_perf.collect_metrics(request_path, params)
    transformed_response = dynatrace_perf.transform_response(response)
    oid_response[metric] = transformed_response
data = {}
for did, info in self.devices.iteritems():
    data[did] = dynatrace_perf.parse_dimensions(oid, transformed_response,
                                              unique_id=info.device.unique_id,
                                              index_key='id',
                                              is_host_disk=True,
                                              partitions=["/var/lib/docker"])

    dynatrace_perf.store_results(data)
except DynatraceError as dyn_err:
    message = "DYNATRACE CLIENT ERROR, reason: {}".format(dyn_err)
    events.generate_event(self.dbc, message)
    logger.exception("DYNATRACE CLIENT ERROR {}".format(dyn_err))
except (KeyError, ValueError, TypeError) as err:
    logger.exception(err)
except Exception as e:
    logger.exception(e)

```

Save Save As

NOTE: The snippet will revert to default values each time the PowerPack is updated. You will need to update the snippet again each time you update the PowerPack.

ELK: AWS CloudTrail

Prerequisites for Monitoring AWS ELK Stacks

To configure SL1 to monitor AWS component devices in ELK stacks using the *ELK: AWS CloudTrail* PowerPack, you must first:

- Install the *Amazon Web Services* PowerPack.
- Create a virtual device in SL1 to represent your AWS service.
- Discover AWS component devices by manually aligning the "AWS Account Discovery" Dynamic Application to the virtual device.
- Ensure that your AWS CloudTrail bucket is properly configured for all read/write events.

NOTE: For more information about the *Amazon Web Services* PowerPack, including how to install the PowerPack and discover AWS devices, see the **Monitoring Amazon Web Services** manual. For more information about the *ELK: AWS CloudTrail* PowerPack, see the **Monitoring Amazon Web Services ELK Stacks** manual.

ELK: Azure Activity Log

Prerequisites for Monitoring Azure ELK Stacks

To configure SL1 to monitor Azure component devices in ELK stacks using the *ELK: Azure Activity Log* PowerPack, you must first:

1. Install the *Microsoft: AzurePowerPack*.
2. Create a virtual device in SL1 to represent your Azure service.
3. Discover Azure component devices by manually aligning the "Microsoft: Azure Account Discovery" Dynamic Application to the virtual device.
4. Ensure that your Azure Activity Log is properly configured for all read/write events.

NOTE: For more information about the *Microsoft: Azure PowerPack*, including how to install the PowerPack and discover Azure devices, see the **Monitoring Microsoft Azure** manual. For more information about the *ELK: Azure Activity Log* PowerPack, see the **Monitoring Microsoft Azure ELK Stacks** manual.

Chapter

34

EMC: VMAX

Prerequisites for Monitoring Dell EMC VMAX

Before you can monitor Dell EMC VMAX systems using the *EMC: VMAX* PowerPack, you must have the following information about an EMC SMI-S Provider that has already been properly installed and configured:

- Username and password for a user with access to the SMI-S Provider
- IP address and port for the SMI-S Provider

NOTE: For more information about the *EMC: VMAX* PowerPack, see the *Monitoring Dell EMC VMAX* manual.

Chapter

35

EMC: VNX

Prerequisites for Monitoring Dell EMC VNX

Before you can monitor Dell EMC VNX storage systems using the *EMC: VNX PowerPack*, you must have the following information about an EMC SMI-S Provider that has already been properly installed and configured:

- Username and password for a user with access to the SMI-S Provider
- IP address and port for the SMI-S Provider

Additionally, statistics logging must be enabled on each Dell EMC VNX storage system that will be monitored. To do so:

1. Log in to Unisphere.
2. Select a Dell EMC VNX storage array from the list, and then click the **[System]** tab.
3. In the **System Management** menu, click **System Properties**.
4. On the **Storage System Properties** dialog box, click the **[General]** tab.
5. Select the *Statistics Logging* checkbox, and then click **[OK]**.

NOTE: For more information about the *EMC: VNX PowerPack*, see the *Monitoring Dell EMC VNX* manual.

Chapter

36

F5 BIG-IP

Prerequisites for Monitoring F5 BIG-IP

Before you can monitor F5 BIG-IP services using the *F5 BIG-IP* PowerPack, you must ensure that SL1 can communicate with BIG-IP using SNMP and you must know the SNMP community string for the BIG-IP system. SL1 can then use the data collected from BIG-IP to create device records for all components managed by BIG-IP.

NOTE: For more information about the *F5 BIG-IP* PowerPack, see the *Monitoring F5 BIG-IP* manual.

Chapter

37

F5: BIG-IP DNS

Prerequisites for Monitoring F5 BIG-IP DNS

Before you can monitor F5 BIG-IP DNS services using the *F5: BIG-IP DNS* PowerPack, you must ensure that SL1 can communicate with BIG-IP DNS using SNMP and you must know the SNMP community string for the BIG-IP DNS system. SL1 can then use the data collected from BIG-IP DNS to create device records for all DNS components.

NOTE: For more information about the *F5: BIG-IP DNS* PowerPack, see the *Monitoring F5 BIG-IP DNS* manual.

Chapter

38

Google Cloud Platform *BETA*

Overview

The following sections describe how to configure Google Cloud Platform resources for monitoring by SL1 using the *Google Cloud Platform *BETA** PowerPack:

Creating a Google Cloud Platform Service Account	107
Enabling Google Cloud Platform APIs	110

NOTE: For more information about the *Google Cloud Platform *BETA** PowerPack, see the *Monitoring Google Cloud Platform* manual.

Creating a Google Cloud Platform Service Account

To monitor Google Cloud Platform (GCP) resources with SL1, you must first create a GCP **service account** for SL1 in the GCP Console. This service account belongs to SL1 instead of an individual end user, and enables SL1 to communicate with Google APIs when monitoring your GCP resources.

This service account's credentials will include a unique email address and a secret JSON key. You will include this email address and key information when you create the SOAP/XML credential that enables SL1 to monitor your GCP resources.

To create a GCP service account:

1. Log in to the GCP Console and go to the **Service accounts** page. If prompted, select a project.
2. Click the **[CREATE SERVICE ACCOUNT]** button.


3. Complete the following fields on the **Create service account** page:

The screenshot shows the Google Cloud Platform interface for creating a service account. The left sidebar lists navigation options: IAM & admin, IAM, Identity & Organization, Organization policies, Quotas, Service accounts (highlighted), Labels, Settings, Privacy & Security, Cryptographic keys, Identity-Aware Proxy, Roles, and Audit Logs. The main content area is titled 'Create service account' and features a progress indicator with three steps: 1. Service account details, 2. Grant this service account access to project (optional), and 3. Grant users access to this service account (optional). The 'Service account details' section contains three input fields: 'Service account name' (with a placeholder 'Display name for this service account'), 'Service account ID' (pre-filled with '@example-project-198515.iam.gserviceaccount.com' and a refresh icon), and 'Service account description' (with a placeholder 'Describe what this service account will do'). At the bottom of the form are 'CREATE' and 'CANCEL' buttons.

- **Service account name.** Type a name for the service account.
- **Service account ID.** This field auto-populates with a service account ID that is based on your **Service account name**.
- **Service account description.** Type a description for the service account.

4. Click **[Create]**. Your service account is created, and the **Service account permissions** page displays.

5. Complete the following fields on the **Service account permissions** page:



- **Role.** Select *Project* > *Viewer*.

NOTE: At a minimum, the service account must have a role of "Project" with "Viewer" permissions for the GCP service that you want to monitor.

6. Click **[Continue]**. The **Grant users access to this service account** page displays.
7. Click **[Create Key]**. The **Create key** pane appears.
8. On the **Create key** pane, select the JSON radio button and then click **[Create]**. The private JSON key is saved to your computer.

Create key (optional)

Download a file that contains the private key. Store the file securely because this key can't be recovered if lost. However, if you are unsure why you need a key, skip this step for now.

Key type

JSON
Recommended

P12
For backward compatibility with code using the P12 format

CREATE **CANCEL**

9. Click **[Close]**, and then click **[Done]**.
10. Open the JSON file that was downloaded to your computer and copy the following information:
 - client_email
 - private_key

TIP: When you copy the private key from the JSON file, it must include the "BEGIN PRIVATE KEY" and "END PRIVATE KEY" lines, including all leading and ending dashes.

If you are discovering GCP resources at the Project level, then you can skip the following steps and continue on to the [Enabling Google Cloud APIs](#) section.

However, *if you are discovering GCP resources at the Organization level*, then you must also do the following:

11. In the GCP Console, go to the **IAM** page and select your organization.
12. Click **[Add]**.
13. Add your service account as a member of the organization, and then add the following mandatory roles:
 - Role > Project > Viewer
 - Role > Resource Manager > Folder Viewer
 - Role > Resource Manager > Organization Viewer
14. When you are finished, click **[Save]**.

Enabling Google Cloud Platform APIs

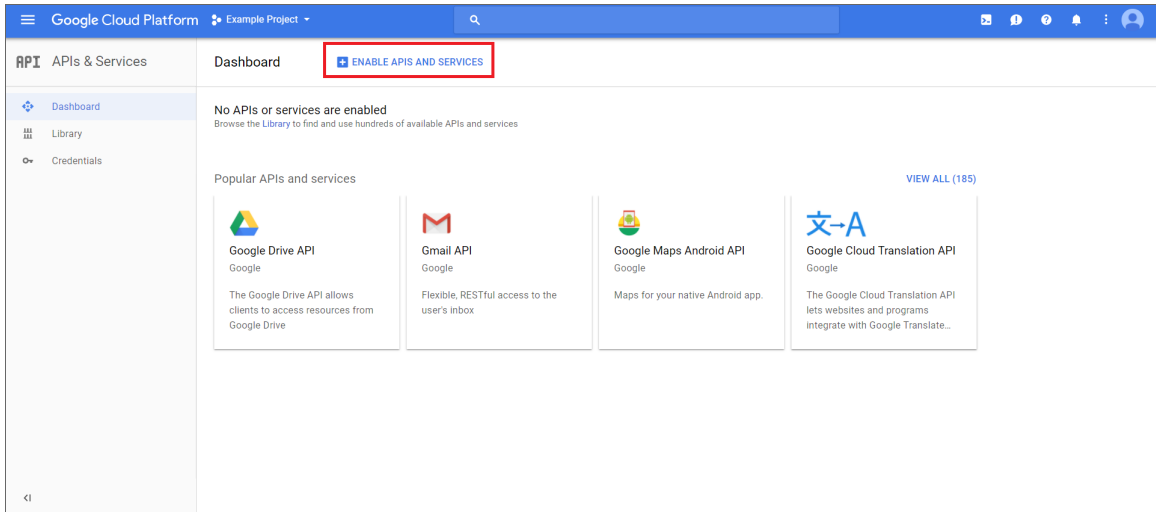
Before SL1 can monitor GCP, you must also enable two APIs in the GCP portal:

- Cloud Resource Manager API
- Compute Engine API

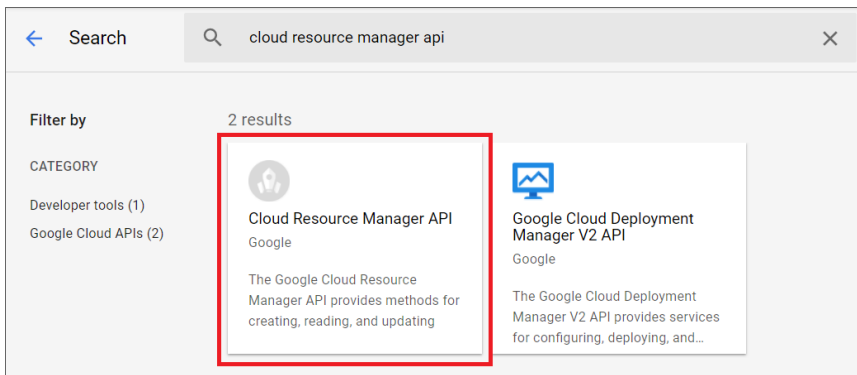
To enable these GCP APIs:

1. Log in to the GCP Console for your project and go to the **API & Services Dashboard** page.

2. Click **[ENABLE APIS AND SERVICES]**. The **API Library** page appears.



3. In the search bar, type "Cloud Resource Manager API". The page will filter search results while you type.
4. Click the **Cloud Resource Manager API** box.



5. On the **Cloud Resource Manager API** page, click the **[Enable]** button.
6. Click **[Dashboard]** on the **API & Services** left menu and then repeat steps 2-5 to enable the **Compute Engine API**.

Hitachi Data Systems: VSP

Prerequisites for Monitoring Hitachi VSP Systems

Before you can monitor Hitachi Virtual Storage Platform (VSP) storage arrays using the *Hitachi Data Systems: VSP PowerPack*, you must have the following information about an Hitachi SMI-S Provider that has already been properly installed and configured:

- IP address and port for the SMI-S Provider
- Username and password for a user with access to the SMI-S Provider

The SMI-S Provider will act as the root device during discovery by SL1.

NOTE: For more information about the *Hitachi Data Systems: VSP PowerPack*, see the *Monitoring Hitachi Data Systems* manual.

Chapter

40

IBM: DataPower

Prerequisites for Monitoring IBM DataPower Gateways 113

Prerequisites for Monitoring IBM DataPower Gateways

Before you can monitor IBM DataPower gateways in SL1 using the *IBM: DataPower PowerPack*, you must first enable SNMP and configure SNMP community strings in each of the DataPower gateways that you will monitor with SL1.


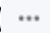
NOTE: For more information about the *IBM: DataPower PowerPack*, see the *Monitoring IBM DataPower Gateway* manual.

Chapter

41

IBM: Db2

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

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

The following sections describe how to configure and discover IBM Db2 databases for monitoring by SL1 using the *IBM: Db2 PowerPack*:

Prerequisites for Monitoring IBM Db2	114
Prerequisites for Linux/Unix Users	114
Prerequisites for Windows Users	116

Prerequisites for Monitoring IBM Db2

To configure the SL1 system to monitor IBM Db2 databases using the *IBM: Db2 PowerPack*, you must first perform the following prerequisites based on your operating system:

Prerequisites for Linux/Unix Users

1. Create a shell session and SSH into the Db2 database you want to monitor.
2. Create a new group to monitor by entering the following command:

```
sudo groupadd <group_name>
```

2. Create a new user for the group you created by entering the following command:

```
sudo useradd -u <user_id> -g <group_name> -m -d /home/<user_name> <user_name>
```

- Set a password for the user you created by entering the following command:

```
sudo passwd <user_name>
```

- Log in with the instance admin user. For example: `su - db2inst1`

- Run the following commands:

```
db2 update database manager configuration using SYSMON_GROUP <group_name>

db2stop

db2start
```

- Connect to your database with the following command:

```
db2 connect to <db_name>
```

- Run the following command to grant the DATAACCESS privilege to the user:

```
db2 "grant DATAACCESS ON DATABASE TO USER <user_name>"
```

- Verify permissions with the following commands:

```
db2 connect to <db_name> user <user_name> using <user_password>

db2 "select SUBSTR(AUTHORITY,1,30), D_USER, D_GROUP, D_PUBLIC, ROLE_USER, ROLE_
GROUP, ROLE_PUBLIC, D_ROLE from table (sysproc.auth_list_authorities_for_authid
(CURRENT_USER, 'U'))"
```

NOTE: Repeat steps 4 - 7 for each Db2 instance.

1	D_USER	D_GROUP	D_PUBLIC	ROLE_USER	ROLE_GROUP	ROLE_PUBLIC	D_ROLE
-----	-----	-----	-----	-----	-----	-----	-----
SYSADM	*	N	*	*	*	*	*
DBADM	N	N	N	N	N	N	*
CREATETAB	N	N	Y	N	N	N	*
BINDADD	N	N	Y	N	N	N	*
CONNECT	N	Y	Y	N	N	N	*
CREATE_NOT_FENCED_ROUTINE	N	N	N	N	N	N	*
SYSCtrl	*	N	*	*	*	*	*
SYSMAINT	*	N	*	*	*	*	*
IMPLICIT_SCHEMA	N	N	Y	N	N	N	*
LOAD	N	N	N	N	N	N	*
CREATE_EXTERNAL_ROUTINE	N	N	N	N	N	N	*
QUIESCE_CONNECT	N	N	N	N	N	N	*
SECADM	N	N	N	N	N	N	*
SYSMON	*	Y	*	*	*	*	*
SQLADM	N	N	N	N	N	N	*
WLMADM	N	N	N	N	N	N	*
EXPLAIN	N	N	N	N	N	N	*
DATAACCESS	Y	N	N	N	N	N	*
ACCESSCTRL	N	N	N	N	N	N	*

NOTE: The user you create will likely need to use KornShell (for Unix systems) or Bash (for Linux systems).

If you are unsure of the shell directory, you can use the command `which ksh` to determine the KornShell directory, or `which bash` to determine the Bash directory.

After you have determined shell directory, run the following commands, replacing `<shell_directory>` with the KornShell or Bash directory:

```
sudo useradd -u <user_id> -g <group_name> -s <shell_directory> -m -d  
/home/<user_name> <user_name>
```

You **should not** use Shell (sh) as the shell for the user. Using Shell for the user shell could result in shell-related errors appearing in the Device Log.

Prerequisites for Windows Users

NOTE: Before performing the steps for the Windows prerequisites, ensure that you have followed the steps in the *Configuring Windows Servers for Monitoring with PowerShell* section of the **Monitoring Windows Systems with PowerShell** manual.

Windows users will need to create a local user and group for the Db2 database. If you have already done so, proceed to [adding the group to the instance database manager](#). To create the user and group, perform the following steps:

1. Click **[Start]** and select **Run**.
2. In the **Run** window, enter `lusrmgr.msc` and click **[OK]**.
3. In the **Local Users and Groups** pane, select the **Users** folder.
4. Click the **Action** menu and select *New User...* Enter the new user's information in the **New User** window and click **[Create]**.
5. In the **Local Users and Groups** pane, select the **Groups** folder.
6. Click the **Action** menu and select *New Group...* Enter the new group's information in the **New Group** window and click **[Create]**.
7. To add the new user to the group, double-click on the group name.
8. Click the **[Add...]** button under the **Members** window and enter the username. Click **[OK]**.

NOTE: You may need to add the user to the Administrators group in order to use PowerShell remoting if you don't have a PowerShell group/policy in place for non-administrative users.

Next, you will need to add the group you created to the instance database manager:

1. Log in to the Db2 database as the instance admin user.
2. Open the Db2 admin shell.
3. Run the following commands:

```
db2 update database manager configuration using SYSMON_GROUP <group_name>
db2stop
db2start
```

Next, you will grant the DATAACCESS privilege to the new user:

1. Log in to the Db2 database as the instance admin user.
2. Open the Db2 admin shell.
3. Run the following commands:

```
db2 connect to <database>
db2 "grant DATAACCESS on database to user <user_name>"
```

NOTE: You will need to grant this access to each database.

NOTE: Perform the steps to add the group to the instance database manager and to grant the DATAACCESS privilege for each Db2 instance that you will monitor.

NOTE: For more information about the *IBM: Db2 PowerPack*, see the *Monitoring IBM Db2* manual.

Chapter

41

IBM: MQ

Prerequisites for Monitoring IBM MQ

To configure the SL1 system to monitor IBM MQ messaging systems using the *IBM: MQ PowerPack*, you must first perform the following:

- [Install the IBM MQ PowerShell Snap-in for Monitoring on Windows Servers](#)
- Give all users the "mgm" group permission

NOTE: For more information about the *IBM: MQ PowerPack*, see the *Monitoring IBM MQ* manual.

Installing the IBM MQ PowerShell Snap-In for Monitoring on Windows Servers

NOTE: Users monitoring MQ on Linux servers do not need to perform these steps.

NOTE: On 64-bit versions of Microsoft Windows, both 32-bit and 64-bit versions of Windows PowerShell are installed. SL1's collection processes using Windows PowerShell will default to using the version of powershell.exe whose folder exists first in the PATH environment variable. Because this will vary from system to system, these steps ensure the WebSphereMQ.dll file is registered for both Windows PowerShell environments.

1. Download the Windows PowerShell library package (mo74.zip) for IBM MQ from the following location:
<https://www.ibm.com/support/pages/mo74-websphere-mq-windows-powershell-library#:~:text=Download%20Description,queue%20managers%20from%20the%20PowerShell>

2. Extract the contents of the zip file to your Windows server, and find the "manual" subfolder from the extracted files (under the mo74_v2.0.1_x86_x64 folder). Create a new folder on your desktop and move the files in the "manual" subfolder to that folder.
3. Register the IBM WebSphere MQ library for use by both 32-bit and 64-bit Windows PowerShell. To do this:

- Start a 32-bit Windows PowerShell console window (this will be the Windows PowerShell (x86) application if running on a 64-bit version of Microsoft Windows) using "Run as administrator", run the following:

```
%WINDIR%\Microsoft.NET\Framework\v4.0.30319\installutil <Directory where  
WebsphereMQ.dll resides>\WebSphereMQ.dll
```

- Start a 64-bit Windows PowerShell console window (this will be the Windows PowerShell application without (x86) in its program name on a 64-bit version of Microsoft Windows) using "Run as administrator" and run the following:

```
%WINDIR%\Microsoft.NET\Framework64\v4.0.30319\installutil <Directory where  
WebsphereMQ.dll resides>\WebSphereMQ.dll
```

4. Open your Windows PowerShell console and add the WebSphere MQ for PowerShell snap-in by running the following command:

```
Add-PSSnapin IBM.PowerShell.WebSphereMQ
```

Configuring the IBM: MQ Queue Discovery Snippet

Chapter

42

IBM: SVC

Prerequisites for Monitoring IBM SVC

To configure the SL1 system to monitor IBM SVC systems using the *IBM: SVC PowerPack*, you must first have the following information about an IBM SMI-S Provider that has already been properly installed and configured:

- The username and password for a user with access to the SMI-S Provider
- IP address and port for the SMI-S Provider

NOTE: For more information about the *IBM: SVC PowerPack*, see the *Monitoring IBM SVC* manual.

IBM: WebSphere Application Server

The following sections describe how to configure and discover IBM WebSphere Application Servers for monitoring by SL1 using the *IBM: WebSphere Application Server PowerPack*:

Prerequisites for Monitoring IBM WebSphere Application Servers 121

Prerequisites for Monitoring IBM WebSphere Application Servers

To configure the SL1 system to monitor IBM WebSphere Application Servers using the *IBM: WebSphere Application Server PowerPack*, you must first set up the following:

- **Performance Monitoring Architecture (PMI)**. PMI is the monitoring structure for the WebSphere Application Server. The performance data provided by the WebSphere PMI helps to monitor and tune the application server performance. To set up PMI, follow the steps here: https://www.ibm.com/support/knowledgecenter/en/SSEQTP8.5.5/com.ibm.websphere.base.doc/ae/tprf_pmi_encoll.html

NOTE: When configuring PMI, it is recommended that you set the status to "All" for each of the application servers you want to monitor.

NOTE: If PMI is disabled on any server, SL1 will continue to show statistics on that server. If the user does not want to see the statistics on the server on which PMI was disabled, they can recursively disable them. SL1 will eventually move that server to **Vanished Devices** and purge it based on the settings that the user has chosen.

- **PerfServlet.** ScienceLogic will use the WebSphere credential that you create to access PMI output through the PerfServlet application. To install PerfServlet, follow the steps here: https://www.ibm.com/support/knowledgecenter/en/SSEQTP8.5.5/com.ibm.websphere.base.doc/ae/tprf_devprfservlet.html
 - After installing, ensure that PerfServlet is mapped to all the WebSphere application servers that you want to monitor
 - To configure the WebSphere credential and access the PerfServlet application, you will need the hostname, default http(s) transport port, and credentials.

NOTE: For more information about the *IBM: WebSphere Application Server PowerPack*, see the **Monitoring IBM WebSphere Application Servers** manual.

Chapter

44

JMX Base Pack *BETA*

Prerequisites for Monitoring JMX Resources

Before you can monitor JMX resources in SL1 using the *JMX Base Pack *BETA** PowerPack, you must have the following information:

- The IP address of the HotSpot, JVM, or OpenJDK system that uses the JMX resources you want to monitor
- The username and password for the system that you want to monitor
- The specific port numbers that you want to monitor

NOTE: For more information about the *JMX Base Pack *BETA** PowerPack, see the *Monitoring Java Management Extensions (JMX)* manual.

Chapter

45

Kubernetes

Prerequisites for Monitoring Kubernetes Clusters

Before you can monitor Kubernetes clusters using the *Kubernetes PowerPack*, you must first do the following:

1. If you will be using Dynamic Applications from the *Linux Base Pack PowerPack*, import and install version 103.
2. Create a Kubernetes service account that SL1 can use to communicate with the Kubernetes API. This service account must have the minimum permissions set in the [Required Permissions for the Service Account Token](#) section.
3. Extract the service account token.
4. Ensure that cURL 7.40 or greater is installed on all Kubernetes nodes that you want to monitor.
5. Configure SSH credentials on the Kubernetes nodes. These credentials must be the same on all nodes, and are used to retrieve data from the underlying Linux OS.

For more information about any of these steps, see <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>.

Required Permissions for the Service Account Token

The minimum required permissions are required for the service account token:

```
...  
apiVersion: rbac.authorization.k8s.io/v1  
kind: ClusterRole  
metadata:  
  name: cluster-limited  
rules:  
- apiGroups:
```

```
- '*'
resources:
- nodes
- pods
- replicationcontrollers
- deployments
- statefulsets
- replicaset
- daemonsets
- cronjobs
- jobs
- componentstatuses
- namespaces
- persistentvolumes
- persistentvolumeclaims
- services
- events
- ingresses
- horizontalpodautoscalers
verbs:
- get
- list
- watch
...
```

NOTE: For more information about the *Kubernetes* PowerPack, see the *Monitoring Kubernetes* manual.

Chapter

46


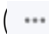
Linux Base Pack

Prerequisites for Monitoring Linux Devices with SSH 126

Configuring Linux Devices to Collect Data 127

NOTE: For more information about the *Linux Base Pack* PowerPack, see the *Monitoring Linux and Solaris* manual.

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

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

Prerequisites for Monitoring Linux Devices with SSH

Before you can monitor Linux devices using the *Linux Base Pack* PowerPack, you must have the following information about the devices that have already been properly configured:

- IP addresses of the devices you want to monitor
- SSH private keys for the devices you want to monitor

Additionally, if you want to collect interface information about your Linux devices, you must install *ifconfig* on those devices.

Configuring Linux Devices to Collect Data

The following tables list the Collection Objects included in those Dynamic Applications and the Linux commands used by each of those objects. You can use these commands to grant or restrict access to certain data types on the user account you will use to monitor your Linux devices.

The following table is a list of configuration and performance Dynamic Applications in the PowerPack:


Dynamic Application	Collection Object	Linux Command
Linux: Configuration Discovery		Determines if a device is a Linux system before discovery in SL1. If the device is not a Linux system, it will not be discovered.
Linux: CPU Configuration	All	cat /proc/cpuinfo/ lscpu
Linux: CPU Cores Performance	All	cat /proc/stat
Linux: CPU Performance	All	cat /proc/stat
Linux: Disk IOPs Performance	All	cat /proc/diskstats
Linux: File System Performance	All	df -kPT
Linux: Hardware Configuration	All	sudo dmidecode -qt 1,2,3
Linux: ICMP Performance	All	cat /proc/net/snmp
Linux: Interface Performance	All	/sbin/ifconfig
Linux: Memory Performance	All	cat /proc/meminfo
Linux: Network Configuration	All	/sbin/ifconfig
Linux: Route Table Configuration	All	netstat -rn
Linux: System Configuration	Kernel Version	cat /proc/sys/kernel/osrelease
	Distribution Genus	cat /etc/os-release
	Host Name	cat /proc/sys/kernel/hostname
	Distribution Release	cat /etc/os-release grep PRETTY_NAME
	AppDynamics Host Name IP Address	hostname=\$(cat /proc/sys/kernel/hostname) && echo \$hostname "<silos:ip>
	AppDynamics Namespace	echo "appdynamics/ns"
	Architecture Type	uname -a

Dynamic Application	Collection Object	Linux Command
	Compiler	cat /proc/version
	Domain Name	cat /proc/sys/kernel/domainname
	Dynatrace Hostname	cat /proc/sys/kernel/hostname
	Dynatrace Namespace	echo "dynatrace/physical/ns"
	New Relic Hostname	cat /proc/sys/kernel/hostname
	New Relic Namespace	echo "newrelic/server/ns"
	Release Date	cat /proc/sys/kernel/version
	SMP Support	cat /proc/sys/kernel/version
	Time Zone	date
	Total Physical Memory (MBytes)	cat /proc/meminfo
	Total Swap Memory (MBytes)	cat /proc/meminfo
Linux: System Load Performance	All	cat /proc/loadavg
Linux: TCP Performance	All	cat /proc/net/snmp
	TCP Ports Listening Cache	netstat -ltn
Linux: TCP Services Configuration	All	netstat -ltn grep tcp
Linux: UDP Performance	All	cat /proc/net/snmp
Linux: UDP Services Configuration	All	netstat -lun grep udp
Linux: Zombie Process	All	ps aux grep Z

The following table is a list of internal collection inventory and performance Dynamic Applications in the PowerPack:

Dynamic Application	Collection Object	Linux Command
Linux: IC Availability	All	Internal Collection that consumes data stored by

Dynamic Application	Collection Object	Linux Command
		the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Detail	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Filesystem Inventory	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Filesystem Performance	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Interface Inventory	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Interface Performance	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Port Performance	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Process Inventory	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: IC Process Performance	All	Internal Collection that consumes data stored by the "Linux: ICDA Cache" Dynamic Application.
Linux: ICDA Cache	Filesystem	df -kPT
	Hardware Config Product Name	cat /sys/devices/virtual/dmi/id/product_name
	Interface	/sbin/ifconfig
	Latency	ping -c1 -W 1 <silosilo:ip>
	Process	ps aux
	Processes CPU Usage	cat /proc/stat
	Processes Memory Usage	free -b
	Software Distribution Release	grep "PRETTY_NAME" /etc/os-release
	Uptime	cat /proc/uptime

NOTE: Linux Base Pack v103 uses a number of standard Linux commands to collect information about a particular device. Most of these commands do not require any specific or elevated permissions to be executed. The PowerPack includes one single command (`dmidecode`) in the "Linux: Hardware Configuration" Dynamic Application which requires root permissions to execute. ScienceLogic recommends configuring a password-less sudo for the user for `dmidecode` as the PowerPack does not support sudo with a password prompt. If the user is not configured correctly the "Linux: Hardware Configuration" Dynamic Application will fail with the following error: `sudo: no tty present and no askpass program specified` You can validate if your configuration is correct by clicking the lightning bolt icon () on the Dynamic Application for the device in question.

Chapter

47

Microsoft: Azure

The following sections describe how to configure Microsoft Azure resources for monitoring by SL1 using the *Microsoft: Azure PowerPack*:

NOTE: The *Microsoft: Azure PowerPack* can monitor Microsoft Azure resources, Microsoft Azure Government resources, and Microsoft Azure resources in Germany and China regions.

Configuring an Azure Active Directory Application	132
<i>Creating an Active Directory Application in the Azure Portal</i>	132
<i>Adding Microsoft Graph APIs Permissions to the Application</i>	134
<i>Generating the Secret Key</i>	136
<i>Locating the Application ID and Tenant ID</i>	137
<i>Locating the Subscription ID</i>	137
<i>Adding Reader Access to the Active Directory Application</i>	138
<i>Setting Up a Proxy Server</i>	140

NOTE: For more information about the *Microsoft: Azure PowerPack*, see the **Monitoring Microsoft Azure** manual.

Configuring an Azure Active Directory Application

To create a SOAP/XML credential that allows SL1 to access Microsoft Azure, you must provide the following information about an Azure application that is already registered with an Azure AD tenant:

- Application ID
- Subscription ID (if monitoring a single subscription)
- Tenant ID
- Secret key

To capture the above information, you must first create (or already have) an application that is registered with Azure Active Directory. The registered application must have Reader access in the subscription. You can then enter the required information about the application when configuring the SOAP/XML credential in SL1. The registered application and the ScienceLogic credential allow SL1 to retrieve information from Microsoft Azure.

TIP: For details on registering an Azure application, see <https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app>.

Creating an Active Directory Application in the Azure Portal


When configuring a SOAP/XML credential in SL1, you must provide the application ID, subscription ID, tenant ID, and secret key of an application that is registered with Azure Active Directory. You will use this registered application to authenticate your Azure account.

NOTE: You must have Service Administrator rights to create an Azure Active Directory application.


To create an application in Azure and register it with Azure Active Directory:

1. Log in to the Azure portal and type "active directory" in the **Search** field at the top of the window.

- From the search results, select *Azure Active Directory*, and then click **App registrations**. The **App registrations** page appears:



- Click the **[New registration]** button.



- When the **Register an application page** appears, enter your application's registration information:
 - Name.** Type a name for the application.
 - Supported account types.** Select *Accounts in this organizational directory only*.
 - Redirect URI (optional).** Select *Web* in the drop-down menu and type a valid URL.

Register an application
PREVIEW

*** Name**
The user-facing display name for this application (this can be changed later).

Sciencelogic Monitoring

Supported account types
Who can use this application or access this API?

Accounts in this organizational directory only (azureteamsciencelogic (Default Directory))

Accounts in any organizational directory

Accounts in any organizational directory and personal Microsoft accounts (e.g. Skype, Xbox, Outlook.com)

[Help me choose...](#)

Redirect URI (optional)
We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

Web

[By proceeding, you agree to the Microsoft Platform Policies](#)

Register

- Click the **[Register]** button. A message appears confirming that your application was added.


Adding Microsoft Graph APIs Permissions to the Application

By default, any new Application has Microsoft Graph API permission. At a minimum, the Microsoft Graph APIs must have permission to directly read data.


To add the Microsoft Graph APIs:

- In the **Search** field of the Azure portal (<https://portal.azure.com>), type "active directory".

2. Click **[App registrations]**, and then click on the name of the Azure Active Directory application you will use to authenticate your Azure account.
3. Click **API Permissions**, and then click **[Add a permission]**. Next, select the **Microsoft Graph** option.




4. In the **Request API permissions** pane, under Select permissions, click the arrow next to **Directory** to open the submenu and select the checkbox for **Directory.Read.all** permission.



5. After you have added the Read directory data, in the **API permissions** page, click the **[Add Permissions]** button.
6. Click **[Grant admin consent for [Directory Name]]**.

7. A pop-up window appears asking if you grant consent for the required permissions for all accounts in your directory. Click **[Yes]**.




Generating the Secret Key

When configuring a SOAP/XML credential for Azure in SL1, you need to provide a secret key for the Azure Active Directory application that you will use to authenticate your account.

To generate a secret key:

1. Log in to the Azure portal at <https://portal.azure.com>, and type "active directory" in the **Search** field at the top of the window.
2. From the search results, select *Azure Active Directory*, and then click **App registrations**.
3. Select the app and then click **[Certificates & secrets]**.
4. In the **Client secrets** pane, click **[+ New client secret]**.




5. In the **Add a client secret** pane, type a name in the **Description** field and select a duration in the **Expires** field.
6. Click **[Add]** to generate the secret key. A new key value displays in the **Client secrets** pane.
7. Copy and save the key value.

Locating the Application ID and Tenant ID

When configuring a SOAP/XML credential for Azure in SL1, you need to provide the Application ID of the Azure Active Directory application you will use to authenticate your Azure account.

To locate the Application ID:

1. Log in to the Azure portal at <https://portal.azure.com>, and type "active directory" in the **Search** field at the top of the window.
2. From the search results, select *Azure Active Directory*, and then click **App registrations**.
3. Click the name of the Active Directory application you will use to authenticate your Azure account. The Application ID and Tenant ID appear in the **Overview** section.



4. Copy and save the values in the corresponding credential fields.


Locating the Subscription ID

If you are monitoring only a single Azure subscription, you must provide the Subscription ID of the Azure Active Directory application you will use to authenticate your account when you configure your SOAP/XML credential for Azure in SL1.

NOTE: If you are monitoring an account with multiple child subscriptions, you can skip this section.

To locate the Subscription ID:

1. In the left pane of the Azure portal (<https://portal.azure.com>), click **[Subscriptions]**.
2. Copy and save the **Subscription ID** of the subscription where you created the Azure Active Directory application you will use to authenticate your account.




Adding Reader Access to the Active Directory Application

To allow ScienceLogic to access your Azure account, you must specify the type of access the user whose information you will use in your SOAP/XML credential has to the Active Directory application used to authenticate your account. Use the **Reader** access role, which is a read-only user that can view everything but cannot make changes.


To specify the access role to the Azure Active Directory application:

1. In the left pane of the Azure Portal (<https://portal.azure.com>), click **[Subscriptions]**.
2. Click the name of your subscription, and then click **[Access control (IAM)]**.


3. In the **Access Control (IAM)** pane, click the **[Add]** button in the **Add a role assignment** section.



4. In the **Add a role assignment** pane, select **Reader** in the **Role** field.



5. In the **Select** field, type the name of the Azure Active Directory application you will use to authenticate your account.



6. Select the application from the search results and click **[Save]**.

Setting Up a Proxy Server

Depending on your needs, you can optionally enable SL1 to connect to Azure through a third-party proxy server such as SQUID. With this configuration, SL1 connects to the proxy server, which then connects to Azure. Azure relays information to the proxy server and SL1 then retrieves that information from the proxy.

NOTE: You can connect to Azure via a proxy server regardless of whether you are monitoring a single subscription or an account with multiple child subscriptions. You can connect to Microsoft Azure, Microsoft Azure Government, and Microsoft Azure Germany and China regions via a proxy server.

NOTE: The *Microsoft: Azure PowerPack* is certified to work with SQUID version 3.5.12 proxy servers.

If you choose to use a proxy server, configure the third-party proxy server based on the third-party documentation. Depending on the type of authentication you require, you might need to specify a user name and password for the proxy server configuration. Also, make a note of the port you opened for the configuration, as this information is needed when creating the SOAP/XML credential.

Chapter

47

Microsoft: Azure

Overview

The following sections describe how to configure Microsoft Azure resources for monitoring by SL1 using the *Microsoft: Azure PowerPack*:

NOTE: The *Microsoft: Azure PowerPack* can monitor Microsoft Azure resources, Microsoft Azure Government resources, and Microsoft Azure resources in Germany and China regions.

Configuring an Azure Active Directory Application	142
<i>Creating an Active Directory Application in the Azure Portal</i>	142
<i>Adding Microsoft Graph APIs Permissions to the Application</i>	144
<i>Generating the Secret Key</i>	146
<i>Locating the Application ID and Tenant ID</i>	147
<i>Locating the Subscription ID</i>	147
<i>Adding Reader Access to the Active Directory Application</i>	148
<i>Setting Up a Proxy Server</i>	150
Creating a SOAP/XML Credential for Azure	150
<i>Load-Balancing an Account with Multiple Subscriptions</i>	153
Creating an Azure Credential	153
<i>Testing the Azure Credential Using the Credential Tester Panel</i>	155

NOTE: For more information about the *Microsoft: Azure PowerPack*, see the **Monitoring Microsoft Azure** manual.

Configuring an Azure Active Directory Application

To create a SOAP/XML credential that allows SL1 to access Microsoft Azure, you must provide the following information about an Azure application that is already registered with an Azure AD tenant:

- Application ID
- Subscription ID (if monitoring a single subscription)
- Tenant ID
- Secret key

To capture the above information, you must first create (or already have) an application that is registered with Azure Active Directory. The registered application must have Reader access in the subscription. You can then enter the required information about the application when configuring the SOAP/XML credential in SL1. The registered application and the ScienceLogic credential allow SL1 to retrieve information from Microsoft Azure.

TIP: For details on registering an Azure application, see <https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-register-app>.

Creating an Active Directory Application in the Azure Portal

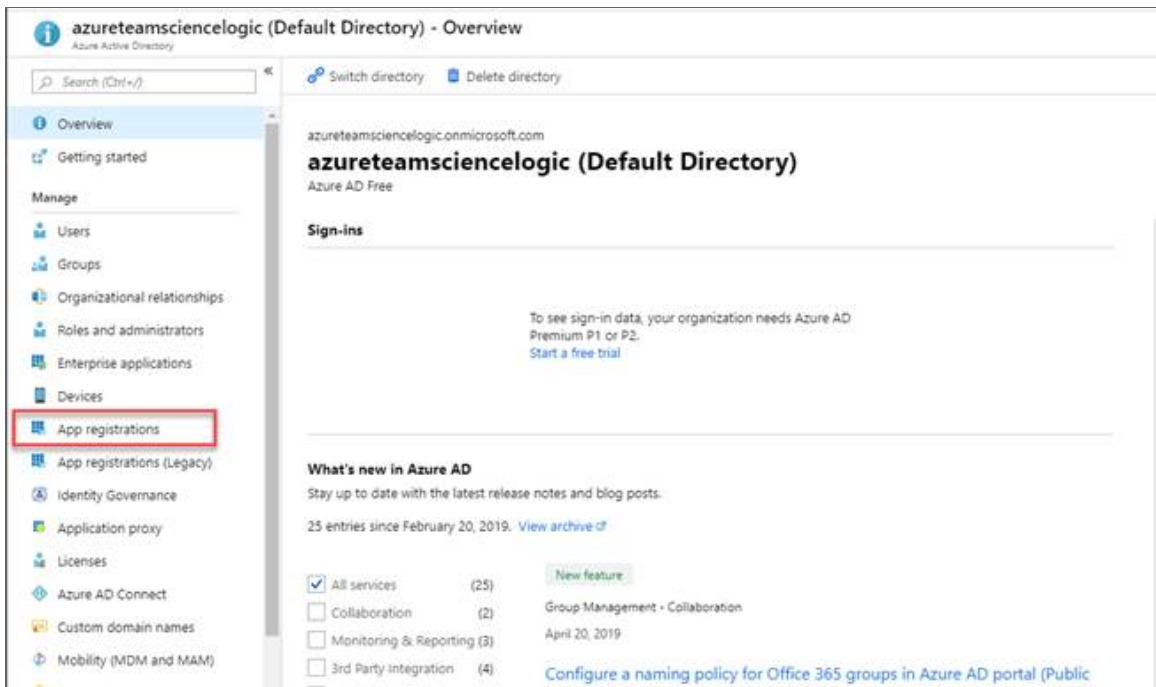
When configuring a SOAP/XML credential in SL1, you must provide the application ID, subscription ID, tenant ID, and secret key of an application that is registered with Azure Active Directory. You will use this registered application to authenticate your Azure account.

NOTE: You must have Service Administrator rights to create an Azure Active Directory application.

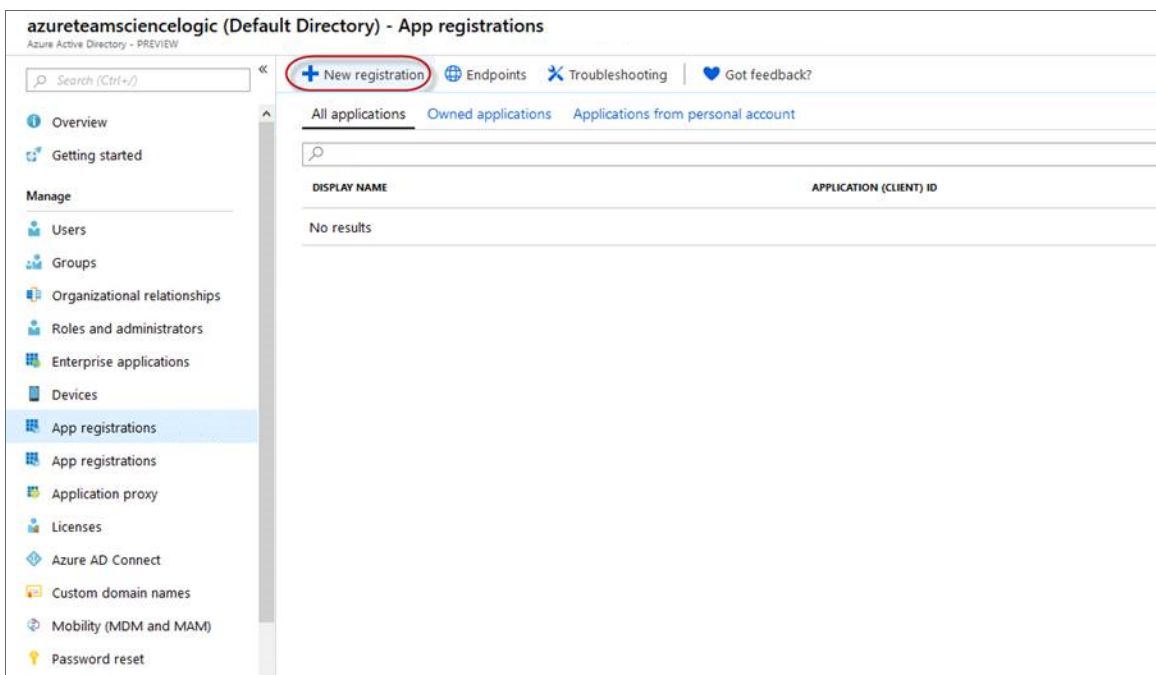
To create an application in Azure and register it with Azure Active Directory:

1. Log in to the Azure portal and type "active directory" in the **Search** field at the top of the window.

- From the search results, select *Azure Active Directory*, and then click **App registrations**. The **App registrations** page appears:



- Click the **[New registration]** button.



- When the **Register an application page** appears, enter your application's registration information:
 - Name.** Type a name for the application.
 - Supported account types.** Select *Accounts in this organizational directory only*.
 - Redirect URI (optional).** Select *Web* in the drop-down menu and type a valid URL.

Register an application
PREVIEW

*** Name**
The user-facing display name for this application (this can be changed later).

Sciencelogic Monitoring

Supported account types
Who can use this application or access this API?

Accounts in this organizational directory only (azureteamsciencelogic (Default Directory))

Accounts in any organizational directory

Accounts in any organizational directory and personal Microsoft accounts (e.g. Skype, Xbox, Outlook.com)

[Help me choose...](#)

Redirect URI (optional)
We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.

Web

[By proceeding, you agree to the Microsoft Platform Policies](#)

Register

- Click the **[Register]** button. A message appears confirming that your application was added.


Adding Microsoft Graph APIs Permissions to the Application

By default, any new Application has Microsoft Graph API permission. At a minimum, the Microsoft Graph APIs must have permission to directly read data.


To add the Microsoft Graph APIs:

- In the **Search** field of the Azure portal (<https://portal.azure.com>), type "active directory".

2. Click **[App registrations]**, and then click on the name of the Azure Active Directory application you will use to authenticate your Azure account.
3. Click **API Permissions**, and then click **[Add a permission]**. Next, select the **Microsoft Graph** option.




4. In the **Request API permissions** pane, under Select permissions, click the arrow next to **Directory** to open the submenu and select the checkbox for **Directory.Read.all** permission.



5. After you have added the Read directory data, in the **API permissions** page, click the **[Add Permissions]** button.
6. Click **[Grant admin consent for [Directory Name]]**.

- A pop-up window appears asking if you grant consent for the required permissions for all accounts in your directory. Click **[Yes]**.




Generating the Secret Key

When configuring a SOAP/XML credential for Azure in SL1, you need to provide a secret key for the Azure Active Directory application that you will use to authenticate your account.

To generate a secret key:

- Log in to the Azure portal at <https://portal.azure.com>, and type "active directory" in the **Search** field at the top of the window.
- From the search results, select *Azure Active Directory*, and then click **App registrations**.
- Select the app and then click **[Certificates & secrets]**.
- In the **Client secrets** pane, click **[+ New client secret]**.




5. In the **Add a client secret** pane, type a name in the **Description** field and select a duration in the **Expires** field.
6. Click **[Add]** to generate the secret key. A new key value displays in the **Client secrets** pane.
7. Copy and save the key value.

Locating the Application ID and Tenant ID

When configuring a SOAP/XML credential for Azure in SL1, you need to provide the Application ID of the Azure Active Directory application you will use to authenticate your Azure account.

To locate the Application ID:

1. Log in to the Azure portal at <https://portal.azure.com>, and type "active directory" in the **Search** field at the top of the window.
2. From the search results, select *Azure Active Directory*, and then click **App registrations**.
3. Click the name of the Active Directory application you will use to authenticate your Azure account. The Application ID and Tenant ID appear in the **Overview** section.



4. Copy and save the values in the corresponding credential fields.


Locating the Subscription ID

If you are monitoring only a single Azure subscription, you must provide the Subscription ID of the Azure Active Directory application you will use to authenticate your account when you configure your SOAP/XML credential for Azure in SL1.

NOTE: If you are monitoring an account with multiple child subscriptions, you can skip this section.

To locate the Subscription ID:

1. In the left pane of the Azure portal (<https://portal.azure.com>), click **[Subscriptions]**.
2. Copy and save the **Subscription ID** of the subscription where you created the Azure Active Directory application you will use to authenticate your account.




Adding Reader Access to the Active Directory Application

To allow ScienceLogic to access your Azure account, you must specify the type of access the user whose information you will use in your SOAP/XML credential has to the Active Directory application used to authenticate your account. Use the **Reader** access role, which is a read-only user that can view everything but cannot make changes.


To specify the access role to the Azure Active Directory application:

1. In the left pane of the Azure Portal (<https://portal.azure.com>), click **[Subscriptions]**.
2. Click the name of your subscription, and then click **[Access control (IAM)]**.

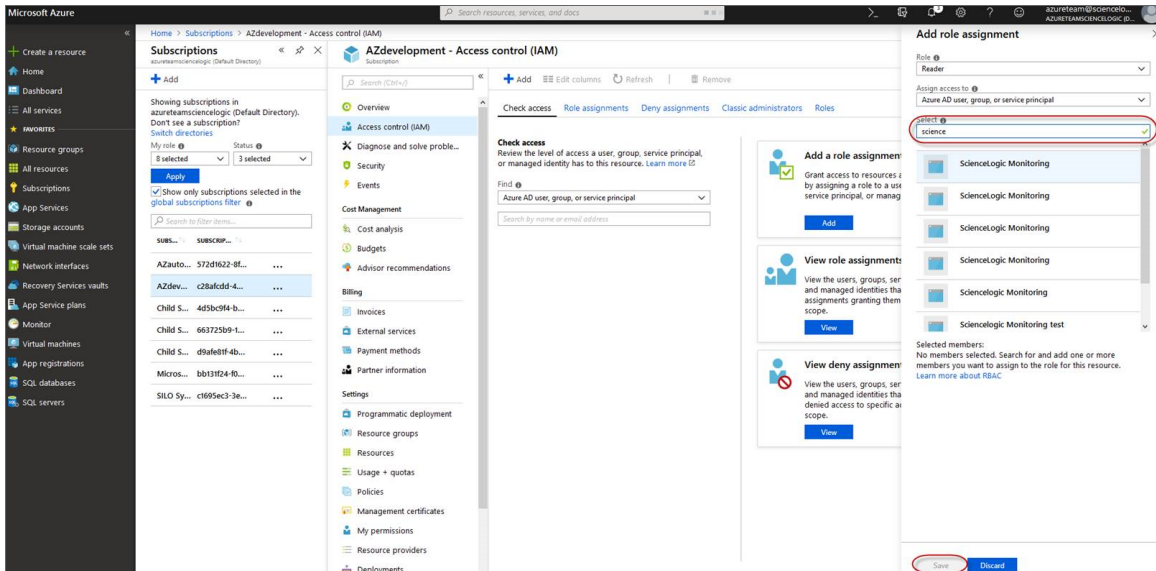
3. In the **Access Control (IAM)** pane, click the **[Add]** button in the **Add a role assignment** section.



4. In the **Add a role assignment** pane, select **Reader** in the **Role** field.



5. In the **Select** field, type the name of the Azure Active Directory application you will use to authenticate your account.



6. Select the application from the search results and click **[Save]**.

Setting Up a Proxy Server

Depending on your needs, you can optionally enable SL1 to connect to Azure through a third-party proxy server such as SQUID. With this configuration, SL1 connects to the proxy server, which then connects to Azure. Azure relays information to the proxy server and SL1 then retrieves that information from the proxy.

NOTE: You can connect to Azure via a proxy server regardless of whether you are monitoring a single subscription or an account with multiple child subscriptions. You can connect to Microsoft Azure, Microsoft Azure Government, and Microsoft Azure Germany and China regions via a proxy server.

NOTE: The *Microsoft: Azure PowerPack* is certified to work with SQUID version 3.5.12 proxy servers.

If you choose to use a proxy server, configure the third-party proxy server based on the third-party documentation. Depending on the type of authentication you require, you might need to specify a user name and password for the proxy server configuration. Also, make a note of the port you opened for the configuration, as this information is needed when creating the SOAP/XML credential.

Creating a SOAP/XML Credential for Azure

After you note the application ID, subscription ID, tenant ID, and secret key of the application (that is registered with Azure Active Directory) that you will use to authenticate your Azure account, you can create a SOAP/XML credential for Azure in SL1. This credential allows the Dynamic Applications in the *Microsoft: Azure PowerPack* to communicate with your Azure subscriptions.

If you want to connect to your Azure account through a third-party proxy server, you must also add the proxy information in the credential. This applies to Microsoft Azure, Microsoft Azure Government, and the Microsoft Azure German and Chinese regions.

The *Microsoft: Azure PowerPack* includes multiple sample credentials you can use as templates for creating SOAP/XML credentials for Azure. They are:

- **Azure Credential - China**, for users who connect to an Azure data center in a Chinese region
- **Azure Credential - Germany**, for users who connect to an Azure data center in a German region (requires a subscription in Germany or Europe)
- **Azure Credential Gov Example**, for users who subscribe to Microsoft Azure Government
- **Azure Credential Proxy Example**, for users who connect to Azure through a third-party proxy server
- **Azure Credential Example**, for all other users.

To create a SOAP/XML credential for Azure:

1. Go to the **Credentials** page (Manage > Credentials).
2. Locate the sample credential you want to use and then click its **[Actions]** icon (☰) and select **Edit**. The **Edit Credential** modal page appears.

The screenshot shows the 'Create Credential' modal page. It features a form with various fields for configuring a credential. On the right side, there is a 'Credential Tester' panel with a 'Test Credential' button. The form includes sections for 'Name', 'All Organizations', 'What organization manages this service?', 'Timeout (ms)', 'Content Encoding', 'Method', 'HTTP Version', 'URL', 'HTTP Auth User', 'HTTP Auth Password', 'Hostname/IP', 'Port (number optional)', 'User', 'Password', 'Embedded Password [%P]', 'Embed Value [%1]', 'Embed Value [%2]', 'Embed Value [%3]', 'Embed Value [%4]', 'HTTP Headers', and 'CURL Options'.

3. Supply values in the following fields:

- **Name**. Type a new name for the Azure credential.
- **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.
- **Timeout (ms)**. Type "120".
- **Content Encoding**. Select *text/xml*.
- **Method**. Select *POST*.

- **HTTP Version.** Select *HTTP/1.1*.
- **URL.** Type the tenant ID in the appropriate place in the URL provided in the sample credential.
- **HTTP Auth User.** Leave this field blank.
- **HTTP Auth Password.** Leave this field blank.

Proxy Settings

- **Hostname/IP.** If you are connecting to Azure via a proxy server, type the server's hostname or IP address. Otherwise, leave this field blank.
- **Port.** If you are connecting to Azure via a proxy server, type the port number you opened when [setting up the proxy server](#). Otherwise, leave this field blank.
- **User.** If you are connecting to Azure via a proxy server using basic authentication, type the server's administrator username. Otherwise, leave this field blank.
- **Password.** If you are connecting to Azure via a proxy server using basic authentication, type the server's administrator password. Otherwise, leave this field blank.

SOAP Options

- **Embedded Password [%P].** Type the secret key for the Azure Active Directory application.
- **Embed Value [%1].** Type the Application ID for the Azure Active Directory application.
- **Embed Value [%2].** Type the Tenant ID for the Azure Active Directory application.
- **Embed Value [%3].** If you are monitoring only a single Azure subscription, type the Subscription ID for the Azure Active Directory application. If you are monitoring multiple subscriptions, leave this field blank.
- **Embed Value [%4].** Leave this field blank. Optionally, you can use this field to add the secret key for the Azure Active Directory application.

HTTP Headers

- **HTTP Headers.** Leave this field blank, unless one of the following scenarios applies to you:
 - If you are using Microsoft Azure Government, this field contains the text "AZGOV".
 - If you are monitoring Microsoft Azure resources in Germany, this field contains the text "AZGER".
 - If you are monitoring Microsoft Azure resources in China, this field contains the text "AZCHINA".
 - If you would like to enable extended logging, enter "LOGGING" in a header field. The log file is located at `/tmp/azure.log`
 - SSL certification verification is enabled by default, but you can disable it in a header field by entering "VERIFY:FALSE".

cURL Options

- **CURL Options.** Do not make any selections in this 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.

Load-Balancing an Account with Multiple Subscriptions

When monitoring an account with multiple child subscriptions, instead of discovering all child subscriptions in a single dynamic component map under their parent account, you can load-balance subscriptions and their components across multiple Data Collectors.

To do this:

- The Collector Group that discovers a group of subscriptions can contain only one Data Collector. You cannot use multiple Data Collectors to discover the Azure components in a single dynamic component map or discover the same device in multiple dynamic component maps.
- To group multiple Azure subscriptions into a single dynamic component map, you need to create a shared credential for that group of subscriptions.
- To create the credential:
 - Perform all of the steps in the section on [Configuring an Azure Active Directory Application](#).
 - Align each subscription in the group with the same application that you registered with Azure AD.
 - In the credential, enter the application ID in the **Embed Value [%1]** field.
 - In the credential, leave the **Embed Value [%3]** field blank.
- During discovery, use this credential to discover the group of subscriptions.
- During discovery, specify the Data Collector you want to use for the group of subscriptions.
- The discovered subscriptions will reside in a common dynamic component map.
- Repeat these steps for each group of subscriptions.

Creating 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 (System > Manage > Credentials).
2. Click the **[Create New]** button and then select *Create Azure 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.
- **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.
- **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.
- **Application ID for Azure AD application**. The Application ID for the Azure Active Directory application.
- **Tenant ID for Azure AD application**. The Tenant ID for the Azure Active Directory application.
- **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.

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 [Testing the Azure Credential](#) section.

Testing the Azure Credential Using the Credential Tester Panel

The *Microsoft: Azure PowerPack* includes a Credential Test for Microsoft Azure. Credential Tests define a series of steps that SL1 can execute on demand to validate whether a credential works as expected.

To test the Azure credential using the Credential Tester panel:

1. After [defining an Azure credential](#), click the **[Save & Test]** button. This activates the Credential Tester fields.
2. 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.
3. Click **[Run Test]** 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".

Chapter

48

Microsoft: Office 365

Overview

The following sections describe how to configure Microsoft Office 365 services for monitoring by SL1 using the *Microsoft: Office 365 PowerPack*:

<i>Configuring Office 365 Monitoring</i>	156
<i>Creating an Office 365 Active Directory Application in the Azure Portal</i>	157
<i>Adding API Permissions to the Application</i>	159
<i>Generating the Secret Key</i>	162

NOTE: For more information about the *Microsoft: Office 365 PowerPack*, see the *Monitoring Microsoft Office 365* manual.

Configuring Office 365 Monitoring

To create a SOAP/XML credential that allows SL1 to access Microsoft Office 365, you must provide the following information about an Office 365 application that is already registered with an Active Directory tenant in Microsoft Azure:

- Application ID
- Tenant ID
- Secret Key

To capture the above information, you must first create or use an existing an Office 365 application that is registered with Azure Active Directory. The application must have access permissions for Office 365 Management APIs and Microsoft Graph APIs. You can then enter the required information about the application when configuring the SOAP/XML credential in SL1. The registered application and the ScienceLogic credential allow SL1 to retrieve information from Office 365.

The following sections describe how to create a registered application, add the appropriate API permissions, and capture the application ID, tenant ID, and secret key.

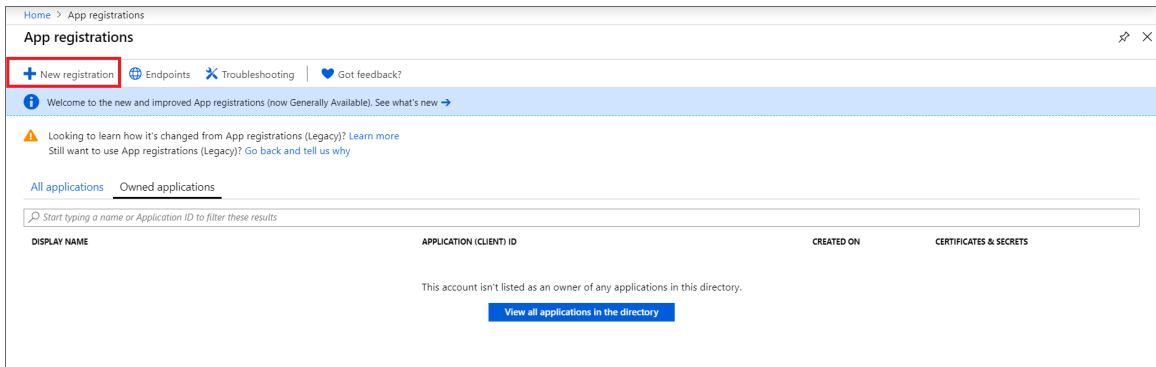
Creating an Office 365 Active Directory Application in the Azure Portal

When configuring a SOAP/XML credential in SL1, you must provide the application ID, tenant ID, and secret key of an Office 365 application that is registered with Azure Active Directory. You use this registered application to authenticate your Office 365 account.

NOTE: You must have Service Administrator rights to create an Active Directory application.

To create an Office 365 application on the Azure portal and register it with Azure Active Directory:

1. Log in to the Azure portal at <https://portal.azure.com> and type "App registrations" in the **Search** field at the top of the window.
2. From the search results, select *App registrations*. The **App registrations** page appears.
3. Click the **[New registration]** button.



4. When the **Register an application page** appears, enter your application's registration information:
 - **Name.** Type a name for the application.
 - **Supported account types.** Select the account types that you want to be supported in your application.
 - **Redirect URI (optional).** Select *Web* in the drop-down menu and type a valid URL. For example: <https://localhost.com>.

Home > App registrations > Register an application

Register an application

*** Name**

The user-facing display name for this application (this can be changed later).

ScienceLogic Monitoring - Office 365 ✓

Supported account types

Who can use this application or access this API?

Accounts in this organizational directory only (azureteamslogic (Default Directory))

Accounts in any organizational directory

Accounts in any organizational directory and personal Microsoft accounts (e.g. Skype, Xbox, Outlook.com)

[Help me choose...](#)

Redirect URI (optional)

We'll return the authentication response to this URI after successfully authenticating the user. Providing this now is optional and it can be changed later, but a value is required for most authentication scenarios.


Web ▼ e.g. <https://myapp.com/auth>

[By proceeding, you agree to the Microsoft Platform Policies](#) [↗](#)

Register

5. Click the **[Register]** button. The **Overview** page for your new application appears.

- On the **Overview** page for your new application, copy and save the values in the *Application (client) ID* and *Directory (tenant) ID* fields. You will need these values when creating your Office 365 credential in SL1.

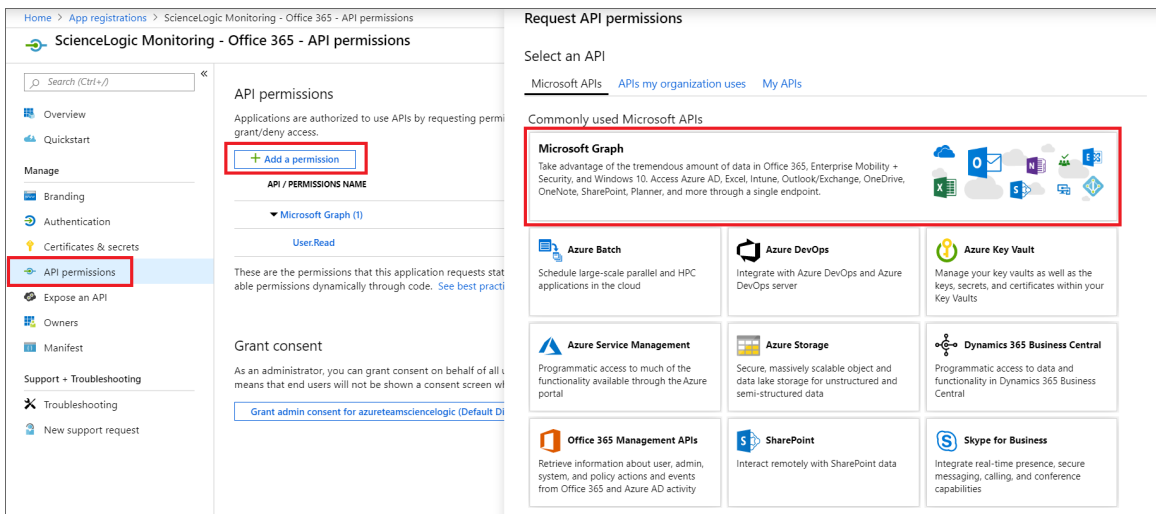


Adding API Permissions to the Application

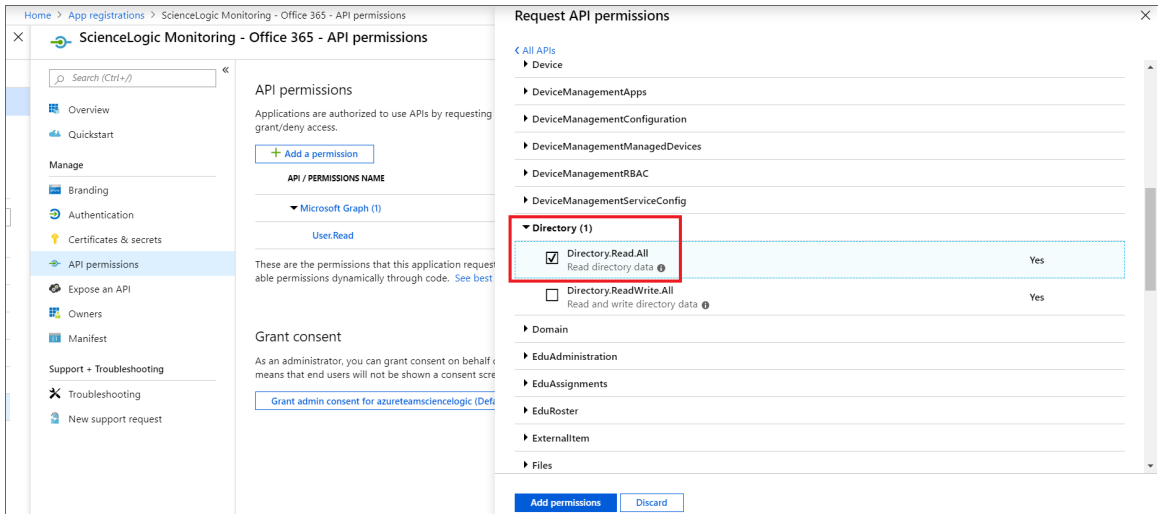
Your Office 365 application must have access permissions for Microsoft Graph APIs and Office 365 Management APIs to be monitored in SL1.

To add API permissions to application:


- From the page for your new application, click **[View API Permissions]**.
- Click **[Add a permission]**, then click the **Microsoft Graph** option.



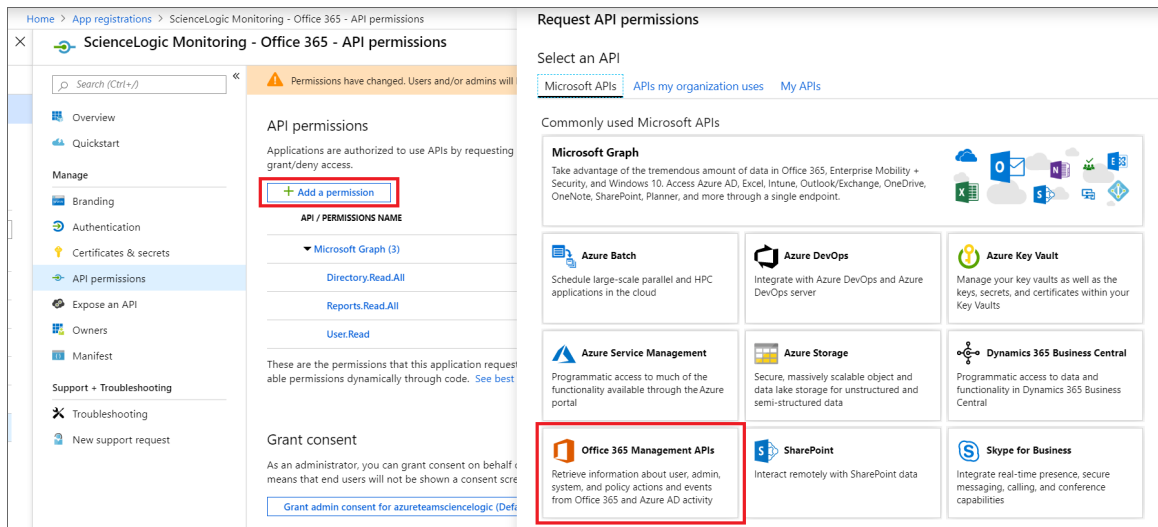
3. In the **Request API permissions** pane, click **Application permissions**.
4. Click the arrow next to **Directory** to open the sub-menu, and then select the checkbox for the *Directory.Read.All* permission.




5. Click the arrow next to **Reports** to open the sub-menu, and then select the checkbox for the *Reports.Read.All* permission.



- Click the **[Add permissions]** button.
- On the **API permissions** page, click **[Add a permission]**, and then click the *Office 365 Management APIs* option.



- In the **Request API permissions** pane, click **Application permissions**.
- Click the arrow next to **ServiceHealth** to open the sub-menu, and then select the checkbox for the *ServiceHealth.Read* permission.



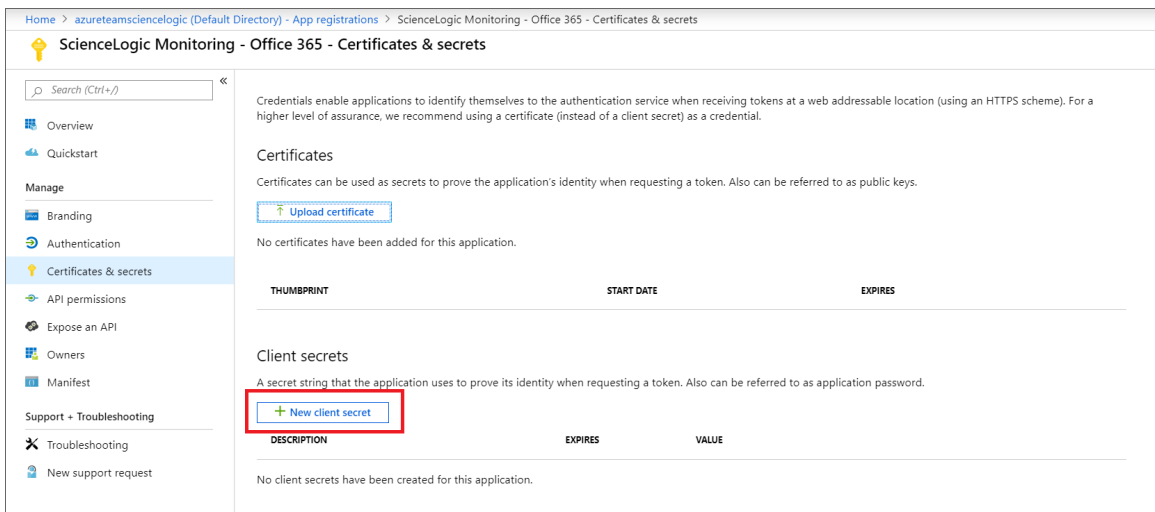
- Click the **[Add permissions]** button.
- On the **API permissions** page, click **[Grant admin consent for [Directory Name]]**.
- A pop-up window appears asking if you grant consent for the required permissions for all accounts in your directory. Click **[Yes]**.

Generating the Secret Key

When configuring a SOAP/XML credential for Office 365 in SL1, you need to provide a secret key for the Office 365 Active Directory application that you will use to authenticate your account.

To generate a secret key:

1. From the Azure portal, type "Active Directory" in the **Search** field at the top of the window.
2. From the search results, select *Azure Active Directory*, and then click **App registrations** on the left pane.
3. Select your Office 365 app from the list.
4. Click [**Certificates & secrets**] on the left pane.
5. In the **Client secrets** pane, click [+ **New client secret**].



6. In the **Add a client secret** pane, type a name in the **Description** field and select a duration in the **Expires** field.
7. Click [**Add**] to generate the secret key. A new key value displays in the **Client secrets** pane.
8. Copy and save the key value.

Microsoft: SQL Server Enhanced

Prerequisites for Monitoring SQL Servers

To configure the SL1 system to monitor SQL servers using the *Microsoft: SQL Server Enhanced PowerPack*, you must first have the following information about the SQL Servers that you want to monitor:

- IP addresses and ports for the SQL Servers
- Username and password for a Windows user account with access to the SQL Servers

The SQL Servers that you monitor must be running PowerShell version 3.0 or later and need to have the SQL Server PowerShell (SQLPS) module installed. This SQLPS module is installed by SQL Server Management Studio. You can also install the SqlServer PowerShell module found here:

<https://www.powershellgallery.com/packages/SqlServer/21.1.18218>

To determine if the proper cmdlets are available for this PowerPack to collect, run `Get-Command Invoke-SqlCmd` to see if the `Invoke-SqlCmd` cmdlet is installed.

In addition, the *Microsoft: SQL Server Enhanced PowerPack* requires the following permissions for the user account used for monitoring:

- SQL 2014 and newer versions require one of the following configurations:
 - The user account has an enabled login on every instance and database to be monitored, with `CONNECT SQL`, `VIEW SERVER STATE`, and `CONNECT ANY DATABASE` permission granted to the login on each instance. The login should have `VIEW DATABASE STATE` permission and `DB_DATAREADER` role granted on the 'master' database, and the `DB_DATAREADER` role granted on the 'msdb' database.
 - The user account has an enabled login on every instance and has the `SYSADMIN` role.

- SQL 2008 to SQL 2012 versions require one of the following configurations:
 - The user account has an enabled login on every instance and database to be monitored, with CONNECT SQL and VIEW SERVER STATE granted to the login on each instance. The login should also have VIEW DATABASE STATE permission and the DB_DATAREADER role granted on the 'master' database, and the DB_DATAREADER role granted on the 'msdb' database. In addition, every database in the instance should have CONNECT access granted to the login.
 - The user account has an enabled login on every instance and has the SYSADMIN role.

ScienceLogic provides a PowerShell script on [the ScienceLogic Support Site](#) that automates the permissions-granting that is required as stated above. The script can be downloaded here: https://portal-cdn.sciencelogic.com/powerpackextras/5819/19047/winrm_configuration_wizardv3.0.zip

After downloading the script, perform the following steps:

1. Copy the winrm_configuration_scriptv3.0.zip file to the Windows server where Microsoft SQL Server is installed and from which you will be collecting data. Unzip the file.
2. Using the credentials for an account that is a member of the Administrator's group, log in to the Windows server you want to monitor. You can log in directly or use Remote Desktop to log in.
3. Right-click on the Windows PowerShell icon and select **Run As Administrator**.
4. At the Windows PowerShell prompt, navigate to the directory where you unzipped the PowerShell script named winrm_configuration_wizard.ps1.
5. At the PowerShell prompt, enter the following to enable execution of the script:

```
Set-ExecutionPolicy -ExecutionPolicy Unrestricted -Scope Process -Force
```

NOTE: The execution policy setting persists only during the current PowerShell session.

6. After the warning text, select Y.
7. To set the required, least-privileged permissions for the user account SL1 will use to monitor all SQL Server instances and databases on the server, run the following script:

```
.\winrm_configuration_wizard.ps1 -user <domain>\<username> -sql_only
```

NOTE: For more information about the *Microsoft: SQL Server Enhanced PowerPack*, see the **Monitoring SQL Servers** manual.



SQL Cluster Monitoring

For SQL Clusters that only include SQL Instances in an Active/Active configuration, follow the steps in the [Discovering SQL Servers](#) section.



For SQL Clusters that include an SQL Instance in an Active/Passive configuration, additional discovery steps are required and listed below.

NOTE: SL1's Active/Passive SQL Instance monitoring leverages the SL1 GUID Component Identifier to allow the SQL Instance component and its child database components to move between SQL Servers during a failover. Adding this GUID Component Identifier on SL1 versions prior to 8.12.1 will create a duplicate SQL Instance component on any already discovered SQL Servers. To prevent this, the GUID Component Identifier is not used by default. The "Enable Active Passive Cluster Failover" threshold in the "Microsoft: SQL Server Discovery" Dynamic Application provides the option to use the GUID Component Identifier when enabled. A value of "0" in the **Threshold Value** disables Active/Passive cluster failover; a value of "1" enables it.

Monitoring SQL Clusters on SL1 8.12.1 or greater.

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Click the wrench icon () for the "Microsoft: SQL Server Discovery" Dynamic Application to open the **Dynamic Applications Properties Editor** page.
3. In the **[Thresholds]** tab, click the wrench icon () for the "Enable Active Passive Cluster Failover" threshold and change the **Threshold Value** to 1.
4. Click **[Save]**.
5. Follow the steps in the [Discovering SQL Servers](#) section on each Windows Server in the cluster.

Monitoring SQL Clusters on SL1 8.8.1 to 8.12.0

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Click the wrench icon () for the "Microsoft: SQL Server Discovery" Dynamic Application to open the **Dynamic Applications Properties Editor** page.
3. In the **[Properties]** tab, change the **Operational State** field to *Disabled*.
4. Click **[Save]**.
5. Follow the steps in the [Discovering SQL Servers](#) section on each Windows Server in the cluster.
6. Go to the **Device Components** page (Registry > Devices > Device Components).
7. Click the wrench icon () for one of the Windows Servers that make up the SQL Cluster to open its **Device Properties** page.
8. In the **[Thresholds]** tab, under **Dynamic App Thresholds | Microsoft: SQL Server Discovery**, change **Enable Active Passive Cluster** to 1.
9. Repeat steps 7 and 8 for each of the Windows Servers that make up the SQL Cluster.
10. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).

11. Click the wrench icon () for the "Microsoft: SQL Server Discovery" Dynamic Application to open the **Dynamic Applications Properties Editor** page.
12. In the **[Properties]** tab, change the **Operational State** field to *Enabled*.
13. Click **[Save]**.

Chapter

50

MySQL

Prerequisites for Monitoring MySQL

To configure the SL1 system to monitor MySQL servers and instances using the MySQL PowerPack, you must first create a read-only MySQL user for each instance to be monitored. For discovery of multiple instances on the same IP address, ScienceLogic recommends creating the same user and password on each instance. The user must have the minimum following privileges:

Privilege	Definition	Level(s)
SELECT	Enables the use of SELECT.	Global, database, table, column.
EXECUTE	Enable the use of statements that execute stored routines (stored procedures and functions). This is necessary for queries on the system database.	

NOTE: For more information about the MySQLPowerPack, see the *Monitoring MySQL* manual.

Chapter

51

NetApp Base Pack

Prerequisites for Monitoring NetApp

Before you discover your NetApp appliances in your SL1 system, you must perform the following configuration tasks on each NetApp Appliance you want to discover:

- Configure a user account on the NetApp device that SL1 will use to connect to the NetApp API. The user account must be assigned a role that includes the following allowed capabilities:
 - login-http-admin
 - api-system-get-*
 - api-aggr-list-info
 - api-lun-list-info
 - api-volume-list-info
 - api-perf-object-get-instances
 - api-storage-shelf-environment-list-info
 - api-net-config-get-active
 - api-vfiler-list-info
 - api-disk-list-info
 - api-snapshot-list-info

NOTE: For Clustered Data ONTAP 8.3 or later, the documentation for customizing the role of a user account is located in the *Clustered Data ONTAP 8.3 System Administration Guide for Cluster Administrators* in the section titled "Customizing an access-control role to restrict user access to specific commands". To view the guide, go to https://library.netapp.com/ecm/ecm_get_file/ECMP1636037. You can download additional NetApp documentation from the NetApp Support Portal at <http://mysupport.netapp.com>.

If you are discovering a Clustered Data ONTAP system, the user account you use for the ScienceLogic credential should be given the built-in "readonly" role and access to the "ontapi" application. For example:

```
security login create [username] -application ontapi -role readonly -vserver  
[clustername]
```

- Determine whether connections to the API on your NetApp device require SSL.
- If you are discovering a NetApp v8 system, you must enable the NetApp multistore license. To do this, execute the following command on your NetApp appliance:

```
options licensed_feature.multistore.enable on
```

NOTE: For more information about the *NetApp Base Pack PowerPack*, see the **Monitoring NetApp Appliances** manual.

Chapter

52

New Relic: APM

Prerequisites for Monitoring New Relic Services

To configure the SL1 system to monitor New Relic services using the *New Relic: APM* PowerPack, you must first have the following information about the New Relic services that you want to monitor *for each account and sub-account*:

- A New Relic REST API key. To generate the REST API key, go to the Account Settings page for your New Relic account.
- The username and password for your New Relic service.
- Insights Query Key. This is optional. Add this to the credential if you want to discover infrastructure groups used for server monitoring. You can generate this from your Insights account.

NOTE: Ensure that you do not have the *New Relic: APM Pro* PowerPack installed before installing the *New Relic: APM* PowerPack. These PowerPacks are not compatible. If you have the *New Relic: APM Pro* PowerPack installed, it will need to be uninstalled prior to installing the *New Relic: APM* PowerPack. The historical data from the *New Relic: APM Pro* PowerPack will be deleted when it is uninstalled.

NOTE: For more information about the *New Relic: APM* PowerPack, see the *Monitoring New Relic* manual.

NGINX: Open Source and Plus

Prerequisites for Monitoring NGINX Services

To configure the SL1 system to monitor NGINX services using the *NGINX: Open Source and Plus* PowerPack, note the following for monitoring the NGINX Open Source Software (OSS):

- The status module must be included when NGINX is instantiated.
- The status stub must be configured in the NGINX configuration.

NOTE: Restart NGINX after editing the configuration.

To learn more about the setup of the status module, see the following NGINX resources:

- Monitoring NGINX (<https://www.nginx.com/blog/monitoring-nginx>)
- Module ngx_http_stub_status_module (http://nginx.org/en/docs/http/ngx_http_stub_status_module.html)

NOTE: For more information about the *NGINX: Open Source and Plus* PowerPack, see the *Monitoring NGINX: Open Source and Plus* manual.

Chapter

54

Nimble Storage

Overview

The following sections describe how to configure Nimble Storage Arrays for monitoring in SL1 using the *Nimble Storage (2.3) PowerPack*:

[Prerequisites for Monitoring Nimble Storage Arrays](#)172

Prerequisites for Monitoring Nimble Storage Arrays

Before you can monitor Nimble Storage Arrays in SL1 using the *Nimble Storage (2.3) PowerPack*, you must have the following:

- Access to TCP port 161 from the SL1 Collector or SL1 All-In-One server
- Nimble Insight SNMP version 2.3 or later

NOTE: For more information about the *Nimble Storage (2.3) PowerPack*, see the ***Monitoring Nimble Storage Arrays*** manual.

Chapter

55

OpenStack

Overview

The following sections describe how to configure OpenStack resources for monitoring by SL1 using the OpenStack PowerPack:

<i>Configuring OpenStack for Monitoring</i>	173
<i>Prerequisites for Monitoring OpenStack</i>	174
<i>Assigning a Role to a User</i>	174
<i>Adding the User Role to API Policy Endpoints</i>	175
<i>Policy Permissions for Administrators</i>	175
<i>Policy Permissions for Non-Administrator Users</i>	178

NOTE: For more information about the OpenStack PowerPack, see the <i>Monitoring OpenStack</i> manual.
--

Configuring OpenStack for Monitoring

To discover OpenStack resources for monitoring by SL1, you must create a SOAP/XML credential that includes authentication information for an OpenStack user.

The user whose information is used in this credential can be either an administrator or a regular (non-administrator) user. Administrator credentials enable SL1 to discover an OpenStack domain and resource pool; regular user credentials enable SL1 to discover only a single project within a specified domain and those components that the user has permissions for in the policy files. The recommended policy edits described in the [Adding the User Role to API Policy Endpoints](#) section will enable non-administrator users to discover resource pools.

The following sections describe how to assign a role to an OpenStack user and then add that user role to the appropriate API policy endpoints.

Prerequisites for Monitoring OpenStack

Before completing the following sections, you must have already created the OpenStack domain and projects you want to monitor, the user whose information you will include in the SOAP/XML credential, and the role you want to assign to that user.

TIP: ScienceLogic recommends that you create a new user role that will be used only for ScienceLogic monitoring and then add this ScienceLogic-specific user role to the policy endpoints described in the [Adding the User Role to API Policy Endpoints](#) section. Having a ScienceLogic-specific user role makes it easier to manage the role's policy permissions without having to change any of your existing user roles.

Assigning a Role to a User

After you have created the user whose information you will include in the SOAP/XML credential, you must assign that user a role in a specific project. This can be done either in the OpenStack portal or using the OpenStackClient command line interface. Both methods are described in this section.

Method 1: OpenStack Portal

To assign the user a role using the OpenStack portal:

1. Log in to the OpenStack portal and navigate to the **Projects** page (Identity > Projects).
2. Locate the project you want to monitor. In the **Actions** column, click **[Manage Members]**.
3. If the user whose information you will include in the SOAP/XML credential does not already appear in the **Project Members** list, locate the user in the **All Users** list and click the plus (+) icon for that user.
4. Locate the user in the **Project Members** list and use the drop-down menu next to the user's name to select the user's role.
5. Click **[Save]**.

Method 2: Command Line Interface

To add a role to the user in a project using the OpenStackClient (OSC) command line interface, SSH into OSC and then use the following command format:

```
openstack role add <role name> --project <project name> user <username>
```

Adding the User Role to API Policy Endpoints

After you have assigned the user a role, you must add the user's assigned role to endpoints in the following OpenStack API policies:

- Keystone (identity services)
- Nova (compute services)
- Neutron (networking services)
- Cinder (block storage services)

For example, to allow a user to list OpenStack projects, Keystone's policy.json file needs to include the following rule:

```
"identity:list_projects" : "role: <user-role>"
```

A rule can also contain multiple roles by using the "or" syntax. For example:

```
"identity:list_projects" : "role: <user-role-1> or role: <user-role-2>"
```

By default, a role can be any of the following:

- admin
- user
- member

Policy Permissions for Administrators

For administrator users, you must update the following policies:

Keystone Policy

An administrator user defined in the SOAP/XML credential needs permissions to the following endpoints defined in /etc/keystone/policy.json:

- identity:get_region
- identity:list_regions
- identity:get_endpoint
- identity:list_endpoints
- identity:get_domain
- identity:list_domains
- identity:get_project
- identity:list_projects

Nova Policy

An administrator user defined in the SOAP/XML credential needs permissions to the following endpoints defined in `/etc/nova/policy.json`:

- `os_compute_api:os-aggregates:index`
- `os_compute_api:os-aggregates:show`
- `os_compute_api:os-extended-server-attributes`
- `os_compute_api:flavors`
- `os_compute_api:os-hosts`
- `os_compute_api:os-hypervisors`
- `os_compute_api:limits`
- `os_compute_api:os-networks`
- `os_compute_api:os-networks:view`
- `os_compute_api:os-networks-associate`
- `os_compute_api:os-security-group-default-rules`
- `os_compute_api:os-security-groups`
- `os_compute_api:os-server-diagnostics`
- `os_compute_api:os-server-groups`
- `os_compute_api:os-server-usage`
- `os_compute_api:servers:detail`
- `os_compute_api:servers:index:get_all_tenants`
- `os_compute_api:servers:detail:get_all_tenants`
- `os_compute_api:servers:show`
- `os_compute_api:servers:show:host_status`
- `os_compute_api:os-services`
- `os_compute_api:os-simple-tenant-usage:show`
- `os_compute_api:os-simple-tenant-usage:list`
- `os_compute_api:os-tenant-networks`
- `os_compute_api:os-virtual-interfaces`
- `os_compute_api:os-volumes`
- `os_compute_api:os-volumes-attachments:show`

Neutron Policy

An administrator user defined in the SOAP/XML credential needs permissions to the following endpoints defined in `/etc/neutron/policy.json`:

- `get_subnet`

- get_subnet:segment_id
- get_subnetpool
- get_address_scope
- get_network
- get_network:router:external
- get_network:segments
- get_network:provider:network_type
- get_network:provider:physical_network
- get_network:provider:segmentation_id
- get_network:queue_id
- get_network_ip_availabilities
- get_network_ip_availability
- get_segment
- get_port
- get_port:queue_id
- get_router
- get_router:distributed
- get_router:ha
- get_dhcp-networks
- get_l3-routers
- get_network_profiles
- get_network_profile
- get_flavors
- get_flavor

Cinder Policy

An administrator user defined in the SOAP/XML credential needs permission to the following endpoint defined in /etc/cinder/policy.json:

- volume_extension:services:index

Policy Permissions for Non-Administrator Users

For regular (non-administrator) users, you must update the following policies:

Keystone Policy

A regular user defined in the SOAP/XML credential needs permissions to the following endpoints defined in `/etc/keystone/policy.json`:

- `identity:get_region`
- `identity:list_regions`
- `identity:get_endpoint`
- `identity:list_endpoints`
- `identity:get_domain`
- `identity:list_domains`
- `identity:get_project`
- `identity:list_projects`

Nova Policy

A regular user defined in the SOAP/XML credential needs permissions to the following endpoints defined in `/etc/nova/policy.json`:

- `os_compute_api:os-aggregates:index`
- `os_compute_api:os-aggregates:show`
- `os_compute_api:os-extended-server-attributes`
- `os_compute_api:flavors`
- `os_compute_api:os-hosts`
- `os_compute_api:os-hypervisors`
- `os_compute_api:limits`
- `os_compute_api:os-networks`
- `os_compute_api:os-networks:view`
- `os_compute_api:os-networks-associate`
- `os_compute_api:os-security-group-default-rules`
- `os_compute_api:os-security-groups`
- `os_compute_api:os-server-diagnostics`
- `os_compute_api:os-server-groups`
- `os_compute_api:os-server-usage`
- `os_compute_api:servers:detail`
- `os_compute_api:servers:index:get_all_tenants`

- os_compute_api:servers:detail:get_all_tenants
- os_compute_api:servers:show
- os_compute_api:servers:show:host_status
- os_compute_api:os-services
- os_compute_api:os-simple-tenant-usage:show
- os_compute_api:os-simple-tenant-usage:list
- os_compute_api:os-tenant-networks
- os_compute_api:os-virtual-interfaces
- os_compute_api:os-volumes
- os_compute_api:os-volumes-attachments:show

Neutron Policy

A regular user defined in the SOAP/XML credential needs permissions to the following endpoints defined in /etc/neutron/policy.json:

- get_subnet
- get_subnet:segment_id
- get_subnetpool
- get_address_scope
- get_network
- get_network:router:external
- get_network:segments
- get_network:provider:network_type
- get_network:provider:physical_network
- get_network:provider:segmentation_id
- get_network:queue_id
- get_network_ip_availabilities
- get_network_ip_availability
- get_segment
- get_port
- get_port:queue_id
- get_router
- get_router:distributed
- get_router:ha
- get_dhcp-networks
- get_l3-routers

- `get_network_profiles`
- `get_network_profile`
- `get_flavors`
- `get_flavor`

Chapter

56

Oracle: Database

Overview

The following sections describe how to configure your Oracle Database instances for monitoring by SL1 using the *Oracle: Database PowerPack*:

<i>Prerequisites for Monitoring Oracle Database Instances</i>	181
<i>Enabling PEM on a Linux Machine</i>	183
<i>Troubleshooting Discovery Issues</i>	184

Prerequisites for Monitoring Oracle Database Instances

To configure the SL1 system to monitor Oracle Database instances using the *Oracle: Database PowerPack*, you must first have the following prerequisites and permissions:

- The Oracle database user must have access the following privileges:
 - sys_privileges: GRANT CREATE SESSION
 - role_privileges: GRANT SELECT_CATALOG_ROLE
 - tan_privileges: GRANT SELECT ON SYS.V_\$DIAG_ALERT_EXT, GRANT SELECT ON SYS.TS\$
- The Oracle database user must have access to the following tables:
 - all_tables
 - dba_data_files
 - dba_free_space
 - dba_registry
 - dba_scheduler_jobs_broken

- dba_scheduler_jobs_failed
- dba_tablespaces
- dba_temp_files
- gv\$sort_segment
- sessions_info
- sys.dba_ind_partitions
- sys.dba_ind_subpartitions
- sys.dba_indexes
- sys.dba_objects
- sys.v_\$database_block_corruption
- v\$archive_dest
- v\$archived_log
- v\$block_change_tracking
- v\$controlfile
- v\$database
- v\$datafile
- v\$datafile_header
- v\$diag_alert_ext
- v\$dispatcher
- v\$latch
- v\$librarycache
- v\$log
- v\$log_history
- v\$logfile
- v\$open_cursor
- v\$parameter
- v\$resource_limit
- v\$rollstat
- v\$rowcache
- v\$session
- v\$sesstat
- v\$statname
- v\$sysstat
- v\$tablespace
- v\$tempfile
- v\$version

- If you are monitoring an RAC system, the user must have access to the following:
 - v\$asm_diskgroup
 - v\$recovery_file_dest
- The Oracle database user must have permission to alter sessions.

Enabling PEM on a Linux Machine

Linux and Unix users can create an SSH/Key credential in order to monitor Oracle Database instances in SL1. The **Private Key (PEM Format)** field may be filled when [creating an SSH/Key credential](#). To enable PEM on a Linux machine, perform the following steps:

1. Create a PEM folder to place the identity keys.

NOTE: ScienceLogic suggests that you create a PEM folder inside the .ssh folder of the user that will use the PEM authentication.

2. Run the following command on your Linux machine to create the SSH key. This command will create public and private keys:

```
ssh-keygen -b 2048 -f identity -t rsa
```

NOTE: The value "identity" in the command above will be the name of the file that is generated. This value can be replaced with any file name.

3. The private key generated from this command is the .pem file needed for the SSH/Key credential. Copy the contents of the file to input into the SL1 credential.
4. Add the generated public key to the `authorized_keys` file that is found in `~/.ssh/authorized_keys` manually or by using the following command:

```
cat identity.pub >> ~/.ssh/authorized_keys
```

5. Restart the SSH service by running the following command:

```
sudo service ssh restart
```

Following the steps above, you can create an SSH/Key credential in SL1 by supplying your Linux server username, Linux server password, and private key. If you would like to create an SSH/Key credential by supplying only your Linux server username and private key, perform the following steps on your Linux machine:

1. Find the `sshd_config` file.
2. Find the `PasswordAuthentication` command line, delete `yes`, and input `no`.
3. Restart the SSH service by running the following command:

```
sudo service ssh restart
```

Troubleshooting Discovery Issues

Oracle: Database v102 includes a `check_oracle` script for Discovery troubleshooting purposes. The script details information about any Discovery/Alignment issues that appear once a Support Escalation is opened.

To run the included `check_oracle` script for the PowerPack:

1. Ensure the script has executable permissions:

```
>sudo chmod 744  
/opt/em7/envs/2BF4A4FD8DC2BA5EDDD565F9CF373156/lib/python2.7/site-  
packages/silo/oracle_db/check_oracle.py
```
2. Execute the script as user `s-em7-core`:

```
>sudo -u s-em7-core  
/opt/em7/envs/2BF4A4FD8DC2BA5EDDD565F9CF373156/lib/python2.7/site-  
packages/silo/oracle_db/check_oracle.py
```
3. When the script has finished, provide the results file:
`/tmp/Oracle_results_10.1.2.3_1622740963`

Chapter

57

Palo Alto

Prerequisites for Monitoring Palo Alto Firewalls

Before you can monitor Palo Alto firewalls in SL1 using the *Palo Alto PowerPack*, you must have the following information:

- SNMP community strings for the devices you want to monitor
- IP addresses for each device you want to monitor
- Username and password for a user with access to the devices you want to monitor

NOTE: The monitored firewalls must be running PAN-OS version 8.0 or later to ensure the proper collection of tunnel performance data.

NOTE: For more information about the *Palo Alto PowerPack*, see the *Monitoring Palo Alto* manual.

Chapter

58

Pure Storage

Overview

The following sections describe how to configure Pure Storage FlashArrays for monitoring by SL1 using the *Pure Storage PowerPack*:

Generating a Pure Storage API Token	186
Testing TCP Port Connectivity	187

NOTE: For more information about the *Pure Storage PowerPack*, see the *Monitoring Pure Storage* manual.

Generating a Pure Storage API Token

The *Pure Storage PowerPack* uses the Pure Storage REST API for collecting configuration and performance data. The Pure API uses port 443; therefore, you must have access to that port. You must also use an API Token, which you can create on the Pure FlashArray and then copy into the [Basic/Snippet credential](#) you create that enables SL1 to discover and monitor the FlashArray.

There are two ways to create the API Token:

- Generate the API token through the Purity user interface (System > Users > Create API Token)
- Generate the API token through the Purity command line interface (`pureadmin create --api-token`)

You can also view existing API tokens in the Purity user interface by navigating to System > Users > API Tokens, clicking the gear icon next to the username, and then selecting *Show API Token*.

After the API Token has been generated, copy and save it for use in the credential.

Testing TCP Port Connectivity

The Pure Storage REST API service runs on TCP port 443 from the primary IP address assigned to the Pure FlashArray. This IP address should be the same one used to access the Purity user interface. To enable SL1 to communicate with the Pure API, your ScienceLogic Data Collector or All-In-One Appliance must have access to TCP port 443.

To test TCP port connectivity, log in to the command line interface of your Data Collector or All-In-One Appliance as the root user and type the following command:

```
nmap -p 443 10.1.1.10
```

If TCP port 443 is open, the following message displays:

```
Starting Nmap 5.51 ( http://nmap.org ) at 2015-10-01 18:42 UTC
Nmap scan report for purestorage-001.mydomain.net (204.110.219.37)
Host is up (0.027s latency).
PORT      STATE SERVICE
443/tcp   open  https

Nmap done: 1 IP address (1 host up) scanned in 0.36 seconds
```

If the port does not appear, or it appears with the state of "filtered", check your firewall settings. If there is a firewall between the ScienceLogic Data Collector or All-In-One Appliance and the Pure Storage REST API, ensure that it can communicate over TCP port 443.

Chapter

59

Silver Peak

Prerequisites for Monitoring Silver Peak

To configure the SL1 system to monitor Silver Peak Unity Orchestrator and edge devices using the *Silver Peak PowerPack*, you must have the following information about the Unity Orchestrator that you want to monitor:

- The IP address or URL of your Orchestrator
- The username and password for the administrator account on your Orchestrator

NOTE: For more information about the *Silver Peak PowerPack*, see the *Monitoring Silver Peak* manual.

Chapter

60

SMI-S: Array

Prerequisites for Monitoring SMI-S Providers

To configure the SL1 system to monitor an SMI-S Provider using the *SMI-S: Array PowerPack*, you must have the following information about the SMI-S Provider that you want to monitor:

- IP address and port for the SMI-S Provider
- Username and password for a user with access to the SMI-S Provider

The SMI-S Provider will act as the root device during discovery by SL1 .

NOTE: For more information about the *SMI-S: Array PowerPack*, see the *Monitoring SMI-S Storage Devices* manual.

Chapter

61

SoftLayer: Cloud

Copying Your SoftLayer API Key

Before you can monitor your SoftLayer account in SL1, you must first generate or retrieve the user-specific API key for your SoftLayer account. SL1 requires this unique API key to communicate with your SoftLayer account.

To generate your SoftLayer API key:

1. Log in to the SoftLayer customer portal and go to the **Users** page (Account > Users).
2. Click the **Generate** link in the **API Key** column for your SoftLayer user. The **Generate** link changes to a **Show** link.
3. Click the **Show** link. Your API key appears.
4. Copy the API key.

To retrieve your SoftLayer API key:

1. Log in to the SoftLayer customer portal.
2. Click your username on the Navigation Pane. The **Edit User Profile** page appears.
3. Locate and copy the API Key.

NOTE: For more information about the *SoftLayer: Cloud* PowerPack, see the **Monitoring SoftLayer** manual.

Chapter

62

VMware: NSX

Prerequisites for Monitoring VMware NSX

To configure SL1 to monitor VMware NSX using the *VMware: NSX PowerPack*, you must first configure an NSX Manager user account to handle RESTful API requests. You will need this account's username and password when creating the Basic/Snippet credential to communicate with the NSX Manager for monitoring.

This API-only user account must be configured in the NSX Manager configuration terminal and have web interface privileges. After you create the user account, perform an API request to give the user account the appropriate role.

NOTE: For more information about the *VMware: NSX PowerPack*, see the *Monitoring VMware NSX* manual.

Chapter

63

VMware: vSphere Base Pack

Overview

The following sections describe how to configure VMware vCenter resources for monitoring by SL1 using the VMware: vSphere Base Pack PowerPack:

Prerequisites for Monitoring VMware vCenter Servers	192
Creating a Read-Only User Account for Monitoring	193

NOTE: For more information about the VMware: vSphere Base Pack PowerPack, see the *Monitoring VMware Systems* manual.

Prerequisites for Monitoring VMware vCenter Servers

Before performing the steps for configuring a vCenter server, you must:

- Have access to a VMware vCenter server that monitors your ESXi and ESX servers.
- Have access to the vCenter server using the vSphere web client.

If the Windows Server that hosts the vCenter server is SNMP-enabled, you must also configure your ESXi or ESX servers for communication using SNMP. To do so, you must:

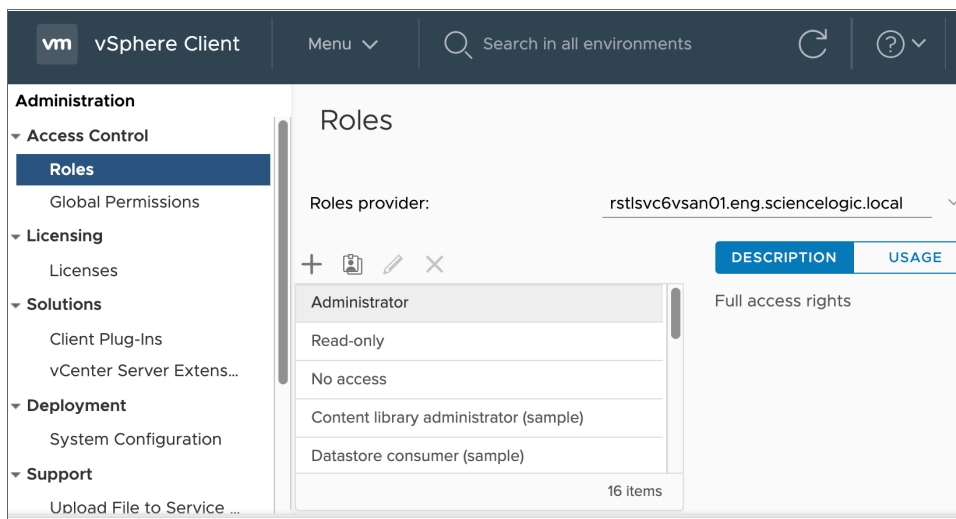
- Configure SNMP community strings, traps, and polling on the ESXi or ESX server. Assign the server at least one SNMP community string. For more information, see VMware's documentation for [Configuring SNMP for ESXi 6.5](#) or [Configuring SNMP for ESXi 6.7](#).

Creating a Read-Only User Account for Monitoring

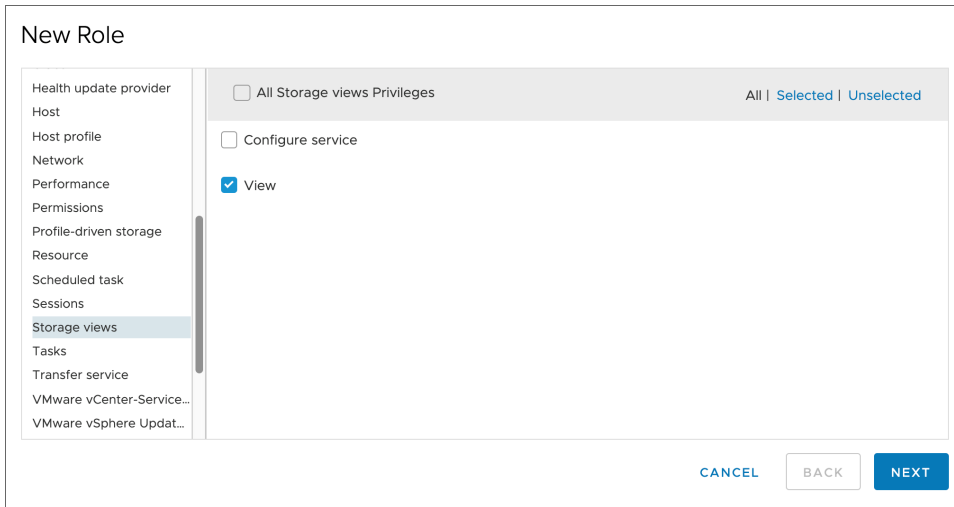
Administrative users are the only default user type that have the level of access SL1 requires to collect data from the VMware vCenter web service. If you do not want to use the username and password of an administrative user in the SOAP/XML credential, you can set up a custom user role with the specific read-only access SL1 requires to the VMware vCenter web service.

To create a custom user role that grants the read-only access SL1 requires, perform the following steps:

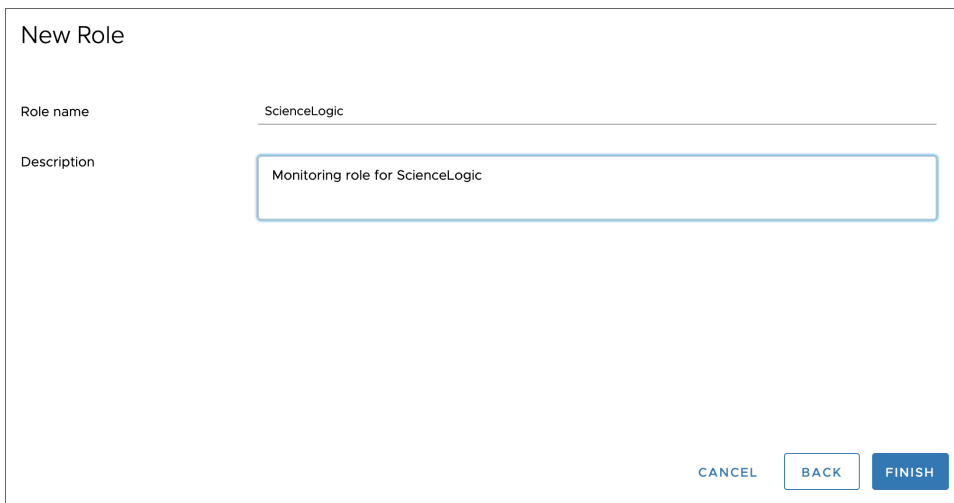
1. Open your vCenter client at `https://<vcenterservername>/ui`
2. Select Menu > Administration from the drop-down.
3. In the menu at the left of the page, click Access Control > Roles. The **Roles** page appears:



4. Click the plus sign (**+**) to add a new Role. The **New Role** page appears:
5. In the pane on the left, select **Storage views** and select the **View** checkbox. Click **[Next]**.




6. In the next screen, enter a name for the role in the **Role name** field. Optionally, you can enter a description in the **Description** field.




7. Click the **[Finish]** button.

To assign the custom role to a user account, perform the following steps:

1. In the vCenter client, select your vCenter server containing the hosts and clusters you are monitoring and click the **[Permissions]** tab.



2. Click the plus sign () to add permissions.
3. Enter values in the following fields:
 - **User**. Select your domain and add the user in the field below.
 - **Role**. Select the role that you just created.
 - **Propagate to children**. Select this checkbox.

A screenshot of a dialog box for adding permissions. It contains the following fields:

- User**: A dropdown menu with 'vsphere.local' selected.
- Role**: A dropdown menu with 'ScienceLogic' selected.
- Propagate to children**: A checkbox that is checked.

At the bottom right of the dialog box, there are two buttons: 'CANCEL' and 'OK'.

4. Click the [OK] button.

© 2003 - 2021, ScienceLogic, Inc.

All rights reserved.

LIMITATION OF LIABILITY AND GENERAL DISCLAIMER

ALL INFORMATION AVAILABLE IN THIS GUIDE IS PROVIDED "AS IS," WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED. SCIENCELOGIC™ AND ITS SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT.

Although ScienceLogic™ has attempted to provide accurate information on this Site, information on this Site may contain inadvertent technical inaccuracies or typographical errors, and ScienceLogic™ assumes no responsibility for the accuracy of the information. Information may be changed or updated without notice. ScienceLogic™ may also make improvements and / or changes in the products or services described in this Site at any time without notice.

Copyrights and Trademarks

ScienceLogic, the ScienceLogic logo, and EM7 are trademarks of ScienceLogic, Inc. in the United States, other countries, or both.

Below is a list of trademarks and service marks that should be credited to ScienceLogic, Inc. The ® and ™ symbols reflect the trademark registration status in the U.S. Patent and Trademark Office and may not be appropriate for materials to be distributed outside the United States.

- ScienceLogic™
- EM7™ and em7™
- Simplify IT™
- Dynamic Application™
- Relational Infrastructure Management™

The absence of a product or service name, slogan or logo from this list does not constitute a waiver of ScienceLogic's trademark or other intellectual property rights concerning that name, slogan, or logo.

Please note that laws concerning use of trademarks or product names vary by country. Always consult a local attorney for additional guidance.

Other

If any provision of this agreement shall be unlawful, void, or for any reason unenforceable, then that provision shall be deemed severable from this agreement and shall not affect the validity and enforceability of any remaining provisions. This is the entire agreement between the parties relating to the matters contained herein.

In the U.S. and other jurisdictions, trademark owners have a duty to police the use of their marks. Therefore, if you become aware of any improper use of ScienceLogic Trademarks, including infringement or counterfeiting by third parties, report them to Science Logic's legal department immediately. Report as much detail as possible about the misuse, including the name of the party, contact information, and copies or photographs of the potential misuse to: legal@sciencelogic.com



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

International: +1-703-354-1010