



Dynamic Application Development

SL1 version 8.9.1

Table of Contents

Overview	7
Introduction	7
What Is A Dynamic Application?	7
Types of Dynamic Applications	8
Elements of a Dynamic Application	9
Optional Features for Dynamic Applications	9
Supplemental Manuals	10
Managing Dynamic Applications	11
Overview	11
Viewing the List of Dynamic Applications	12
Searching and Filtering the List of Dynamic Applications	16
Special Characters	17
Performing Other Tasks in the Dynamic Application Manager Page	21
Viewing the Dynamic Applications Associated with a Device	22
Viewing the Status of a Dynamic Application	24
Found	25
Collect	25
How the ScienceLogic Platform Manages Collect Status	26
Stopping Collection	26
Starting Collection	26
Collection Objects that are Excluded from Maintenance	27
Status of Objects for Deviation	27
Manually Associating a Dynamic Application with a Device	28
Editing the Credential Associated with a Dynamic Application	29
Performing Other Administrative Tasks for an Aligned Dynamic Application	31
Enabling or Disabling Objects	31
Restarting Automatic Maintenance of Collection Objects	32
Editing the Poll Frequency for a Dynamic Application on the Current Device	32
Stopping Data Collection for a Dynamic Application	33
Resetting Statistical Data for a Dynamic Application	33
Resetting Persistent Session Objects for a Dynamic Application	34
Testing Data Collection for a Dynamic Application	35
Removing Data Collected by a Dynamic Application	35
Bulk Un-aligning Dynamic Applications	36
Setting Thresholds for Dynamic Applications	37
Dynamic Applications and Discovery	38
How Does the ScienceLogic Platform Align Dynamic Applications During Discovery?	38
Queuing Discovery from the Dynamic Applications Manager Page	39
System Settings That Affect Dynamic Application Discovery	40
Viewing Dynamic Application Data and Alerts	43
Overview	43
Performance Dynamic Applications	44
Viewing Performance Dynamic Application Data	45
Viewing Details with the Data Table	45
Configuring the Data for the Report	46
Generating a Report from a Performance Graph	46
Defining the Date Range for a Report	47
Configuration Dynamic Applications	47
Selecting Data to View	48
Viewing Data	49

Generating a Report of the Data	49
Viewing Historical Data	50
Editing the Application	51
Journal Dynamic Applications	51
Selecting Data to View	52
Viewing Data	52
Searching & Filtering the List of Data	53
Special Characters	54
Generating a Report of the Data	58
Editing the Application	58
Viewing Alerts Generated by Dynamic Applications	59
Searching the List of Log Entries for a Dynamic Application	60
Viewing Active Events for a Dynamic Application	61
Viewing Details About an Active Event	63
Viewing Cleared Events from a Dynamic Application	64
Dynamic Application Settings	67
Overview	67
Dynamic Application Properties	67
The Abandon Collection Property	73
Creating Dynamic Applications	75
Overview	75
Creating the Container for a Dynamic Application	78
Duplicating an Existing Dynamic Application	79
Creating a Dynamic Application Using the SNMP Walker Tool	80
Collection Objects	81
What is a Collection Object?	81
How Does A Collection Object Get Collected?	81
How Collection Objects are Used by the ScienceLogic Platform	82
Viewing the Collection Objects in a Dynamic Application	83
Creating a Collection Object	84
Basic Settings	86
Collection Settings	89
Grouping & Indexing Settings	89
Configuration Dynamic Application Settings	90
Custom Attribute Settings	93
Standard Deviation Settings	96
Dynamic Component Mapping Settings	97
Creating a Discovery Object	99
Editing a Collection Object	101
Performing Bulk Actions on Collection Objects	101
Caching	102
What is Caching?	102
Caching Responses	103
Consuming Cached Responses	103
Configuring Collection Objects in Cache-Consuming Dynamic Applications	104
SOAP Dynamic Applications	104
WMI Dynamic Applications	104
PowerShell Applications	105
XML Dynamic Applications	105
XSLT Dynamic Applications	105
Collector Affinity for Dynamic Applications that Use Caching	106
Dynamic Component Mapping	107

What is Dynamic Component Mapping?	107
Where in the ScienceLogic Platform are Component Devices Displayed?	108
Configuring a Dynamic Application to Create Component Devices	108
Configuring Collection Objects as Component Identifiers	109
Duplicate MAC Addresses for Component Devices	111
Component Devices, Collector Groups, and Load Balancing	112
Collector Affinity	112
Failover for Collector Groups for Component Devices	113
Merging Component Devices	113
Availability for Component Devices	114
Moving Component Devices Between Root Devices	116
Automatically Aligning Dynamic Applications with Component Devices	116
Automatically Assigning a Device Class to Each Component Device	117
Configuring a Device Class to Match Collected Values	117
Configuring Component Identifiers for the Collection Objects	118
How Does the ScienceLogic Platform Use Component Identifiers to Automatically Assign a Device Class to Each Component Device?	119
Aligning a Default Device Class with a Dynamic Application	120
Defining a Default Device Class in the Device Class Editor	121
Defining a Default Device Class in the Dynamic Application	121
Variables in Dynamic Applications for Component Devices	122
Caching and Component Mapping	123
Dynamic Application Relationships	124
Overview	124
Configuring Collection Objects for Relationships	125
Configuring Identity-Based Relationships	127
Viewing a Map of Component Devices	128
Viewing a Component Map	129
Adding Devices from Another Component Tree to the Component Map	131
Viewing Relationships in the Organization View	132
Viewing an Organizational Map	132
Viewing Relationships in Customized Maps	133
Viewing Relationships for a Single Device	135
Device View Page	137
Presentation Objects for Performance Dynamic Applications	139
What is a Presentation Object?	139
Viewing the Presentation Objects in a Dynamic Application	140
Creating a Presentation Object	141
Presentation Object Formulas	143
How the ScienceLogic Platform Processes Performance Presentation Objects	143
Editing a Presentation Object	144
Deleting a Presentation Object	144
Alerts and Thresholds	146
What is an Alert?	146
What is a Threshold?	146
Viewing the Thresholds in a Dynamic Application	147
Creating a Threshold	148
Editing a Threshold	149
Deleting a Threshold	149
Viewing the Alerts in a Dynamic Application	149
Creating an Alert	151
Alert Formulas	152

Evaluate	152
Operators	153
Dividing Integers	154
Using Quotes in an Alert	154
The result() Function	155
The threshold() Function	156
The active() Function	157
The avg() Function	157
The deviation() Function	158
The find() Function	161
The global() Function	161
The log() Function	162
The prior() Function	163
The round() Function	163
The sum() Function	164
Date & Time Variables	164
The tab_idx Variable	166
Creating an Event Policy for an Alert	167
Editing an Alert	169
Deleting an Alert	169
Validating an Alert	170
Running the Alert Validator Command Line Tool	170
Example	171
Indexing	173
What is Dynamic Application Indexing?	173
How Indexing Affects Configuration Tables	174
How Indexing Affects Performance Graphs	174
How Indexing Affects Alerts	176
The Default Indexing Method	177
Instance Values and Index Creation for SNMP Dynamic Applications	178
Instance Values and Index Creation for Snippet Dynamic Applications	178
Instance Values and Index Creation for Other Dynamic Applications	179
Designating an Index	179
Grouping Dynamic Application Data Using Collection Labels	183
What are Collection Labels and Collection Groups?	183
Viewing the List of Collection Labels	183
Filtering the List of Collection Labels	184
Special Characters	185
Creating a Collection Group	189
Creating a Collection Label	189
What is Normalization?	190
What are Duplicates and How Does the ScienceLogic Platform Manage Them?	193
What is Precedence?	193
Aligning a Presentation Object with a Collection Label	194
Viewing and Managing the List of Presentation Objects Aligned with a Collection Label	195
Viewing and Editing Duplicate Presentation Objects by Collection Label	199
Viewing and Managing the List of Devices Aligned with a Collection Label	201
Editing Duplicate Presentation Objects by Device	204
Editing Duplicate Presentation Objects for a Single Device	206
Editing a Collection Label	207
Deleting a Collection Label	208
Viewing Reports About Collection Labels on a Single Device	209

Viewing Dashboards About Collection Labels	210
--	-----

Introduction

This manual is intended for users who are responsible for developing new Dynamic Applications.

What Is A Dynamic Application?

Dynamic Applications are customizable policies, created for a specific vendor and a specific type of device or system, that tell the ScienceLogic platform:

- What data to collect from devices
- How to present the data that has been collected
- When to generate alerts and events based on the data that has been collected

The ScienceLogic platform includes Dynamic Applications for the most common hardware and software. You can customize these default Dynamic Applications to best work in your environment. You can also create custom Dynamic Applications.

For example, a Dynamic Application that is designed to collect information about file system usage on a device might define that:

- The ScienceLogic platform should collect the total size of each file system and amount of space used for each file system.
- The ScienceLogic platform should use the two collected values to present a graph of "percentage used" values for each file system.
- The ScienceLogic platform should generate an alert when the amount of space used on a file system exceeds a specified percentage of the total size of that file system.

Types of Dynamic Applications

There are many different types of Dynamic Applications that can be created and used in the ScienceLogic platform. The Dynamic Application type defines:

- The *protocol* the ScienceLogic platform should use to collect the data.
- An *archetype*, which defines generally what type of data is being collected.

All Dynamic Application types use one of the following *protocols*:

- **Bulk Snippet**. The ScienceLogic platform will collect data by executing custom code written in the Python programming language. A *single instance of the Dynamic Application* uses custom-written Python code to collect data from *multiple devices*. For a standard Snippet Dynamic Application, the ScienceLogic platform spawns a separate instance of the Dynamic Application for each device or component device. The Bulk Snippet is useful for systems that include a large number of component devices.
- **Database**. The ScienceLogic platform will collect data by connecting to a database and executing a query.
- **Internal Collection**. The ScienceLogic platform will collect data by performing an "Extended Internal Collection". Basic hardware and system data that was previously collected by the platform using an internal SNMP process can now be collected with Internal Collection Dynamic Applications (ICDAs). ICDAs do not use SNMP to provide a way to collect basic hardware and system data from devices that do not support SNMP.
- **PowerShell**. The ScienceLogic platform will collect data from a Windows device by performing PowerShell commands.
- **Snippet**. The ScienceLogic platform will collect data by executing custom code written in the Python programming language.
- **SNMP**. The ScienceLogic platform will collect data by performing SNMP requests.
- **SOAP**. The ScienceLogic platform will collect data by sending HTTP POST requests to a SOAP web service.
- **WMI**. The ScienceLogic platform will collect data by sending requests to a Windows Management Instrumentation service running on an external Windows device.
- **XML**. The ScienceLogic platform will collect data by sending an HTTP GET request for an XML document.
- **XSLT**. Similar to the **SOAP** protocol, the ScienceLogic platform will collect data by sending HTTP POST requests to a SOAP web service. The ScienceLogic platform will also perform XSLT transformations to generate the requests and parse the responses.

All Dynamic Application types use one of the following *archetypes*:

- **Performance**. The ScienceLogic platform will display the collected data in a historical graph. The ScienceLogic platform will also automatically calculate daily normalized data for each graph.
- **Configuration**. The ScienceLogic platform will display the current values of the collected data in tabular format. The ScienceLogic platform will store tables of previously collected data to show when the values have changed.

- **Journal.** The ScienceLogic platform will display collected data in log format. Each log entry can contain multiple collected values and can change over time. The **Journal** archetype is available only when using the **Snippet** protocol.

Each Dynamic Application type has unique configuration options.

Elements of a Dynamic Application

A Dynamic Application can contain the following elements:

- **Properties** that define general settings for the Dynamic Application. For example, the name of the Dynamic Application, type of Dynamic Application (protocol and archetype), and frequency at which the ScienceLogic platform should collect data using the Dynamic Application.
- **Collection Objects** that define each value the ScienceLogic platform should collect, how the platform should collect the value, and any processing the platform must perform before storing the value. For Dynamic Applications with an archetype of **Configuration**, **Collection Objects** also define how each value will be displayed in the table of data.
- **Requests or Snippets** that define how the platform should request data from a device. For Dynamic Application types that use **Requests or Snippets**, each **Collection Object** is associated with the **Request or Snippet** that will populate that **Collection Object**. **Requests** are used in Dynamic Applications that use the SOAP, WMI, or XSLT protocol. **Snippets** are used in Dynamic Applications that use the Snippet protocol.
- **Presentation Objects** that define how the platform should present the collected data (for Dynamic Applications with an archetype of **Performance** or **Journal**).
- **Alerts** that include formulas that the platform will evaluate each time data is collected. Each formula examines the current values for one or more **Collection Objects**. If the formula evaluates to *true*, the platform generates an alert. Additionally, you can align an event policy with each **Alert**, and the platform can trigger that event if the alert evaluates to *true*.
- **Thresholds** that define variables that can be used in **Alerts**. Each **Threshold** defines an allowed range of values for the variable and a default value. The value of the **Threshold** variable can be changed on a per-device basis by users of the Dynamic Application, without having to modify the Dynamic Application.

Optional Features for Dynamic Applications

You can enable the following optional features for a Dynamic Application:

- **Dynamic Component Mapping.** The ScienceLogic platform can use Dynamic Application data that has been collected from a single management system, such as a VMware ESX server, to create a device record for each entity (such as a virtual machine) managed by that single management system.
- **Caching.** When the ScienceLogic platform requests information from a device during Dynamic Application collection, the platform can optionally **cache** the response from the device. If a response is cached, other Dynamic Applications that use the same protocol can use the cached response to populate collection objects. Caching responses reduces the number of requests performed by the platform and speeds up collection.

- **Collector Affinity.** The ScienceLogic platform enables you to specify which Data Collectors are allowed to collect data from a device for a Dynamic Application. This ensures that a group of collection jobs that cannot be split up among multiple Data Collectors—for instance, Dynamic Applications that produce and consume cache—will run on a single Data Collector.
- **Relationships.** Dynamic Applications can be configured to automatically create relationships between devices. For example, the Dynamic Applications in the VMware vSphere and NetApp PowerPacks are configured to create relationships between VMware Datastore component devices and their associated NetApp Volume component devices. Relationships created by Dynamic Applications are used and visualized by the platform in the same manner as relationships created by topology collection, Dynamic Component Mapping, and manually in the user interface.
- **Collection Labels.** When multiple Dynamic Applications collect the same type of performance data from different devices, you can use collection labels to view data from those different Dynamic Applications at the same time. For example, when the ScienceLogic platform displays the CPU utilization, Memory utilization, or Swap utilization for a device, the utilization values come from a presentation object in a Dynamic Application aligned with that device.
- **Custom Attributes.** Dynamic Applications of archetype *configuration* can collect configuration data and use that data to create and/or populate the value of a custom attribute. There are two ways Dynamic Applications can interact with custom attributes. A Dynamic Application can populate the value of an existing custom attribute, or a Dynamic Application can both create the custom attribute and populate its value. For general information on custom attributes, see the manual **Using the ScienceLogic API**.
- **Asset Linking.** The ScienceLogic platform can use data collected by configuration Dynamic Applications to automatically populate fields in asset records.
- **Inventory Linking.** The ScienceLogic platform can use data collected using configuration Dynamic Applications to automatically populate inventories of hardware components and installed software.
- **Standard Deviation Alerting.** The ScienceLogic platform can calculate standard deviation data for values collected by performance Dynamic Applications. The standard deviation data can generate alerts when values deviate by a specified amount.

Supplemental Manuals

This manual covers all the components that are common to all the Dynamic Application protocols. ScienceLogic publishes additional manuals for each Dynamic Application protocol. These manuals describe the unique configuration for each protocol and include examples of Dynamic Applications that use that protocol.

Managing Dynamic Applications

Overview

This chapter will describe how to manage Dynamic Applications. This chapter will include:

- *Viewing the List of Installed Dynamic Applications*
- *Viewing the Dynamic Applications Associated With a Device*
- *Viewing the Status of a Dynamic Application*
- *Maintenance of Collection State*
- *Status of Objects for Deviation Alerting*
- *Manually Associating a Dynamic Application with a Device*
- *Manually Associating a Credential with a Dynamic Application*
- *Performing Other Administrative Tasks for an Aligned Dynamic Application*
- *Setting Thresholds for Dynamic Applications*
- *Dynamic Applications and Discovery*

Viewing the List of Dynamic Applications

The **Dynamic Applications Manager** page (System > Manage > Applications) displays a list of all existing Dynamic Applications. For each Dynamic Application, the page displays the following:

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

Dynamic Application Name	Poll Rate	Type	State	Version	ID	Subscribers	PowerPack	Environment	Collects	Alerts	Events	Thresh	Edited By	Last Edit
1. Not Test ICDA	1 min.	Snippet Configuration	Enabled	1	1502	--	ScienceLogic: ICDA Harc	SYSTEM	1	--	--	--	em7admin	2018-07-20 09:27:05
2. 2-level DCM app cache producer	15 min.	Snippet Configuration	Enabled	0.1	1517	--	2-level DCM app test	SYSTEM	1	--	--	--	em7admin	2018-07-20 09:26:21
3. 2-level DCM app config	1 min.	Snippet Configuration	Enabled	0.1	1520	--	2-level DCM app test	SYSTEM	1	1	1	--	em7admin	2018-07-20 09:26:21
4. 2-level DCM app level1 discovery	1 min.	Snippet Configuration	Enabled	0.1	1518	--	2-level DCM app test	SYSTEM	2	--	--	--	em7admin	2018-07-20 09:26:21
5. 2-level DCM app level2 discovery	1 min.	Snippet Configuration	Enabled	0.1	1519	--	2-level DCM app test	SYSTEM	2	--	--	--	em7admin	2018-07-20 09:26:21
6. Alteon: Configuration	120 min.	SNMP Configuration	Enabled	1.2	1174	--	Alteon Base Pack	n/a	30	5	5	--	em7admin	2018-07-19 23:27:21
7. Alteon: Load Trending	15 min.	SNMP Performance	Enabled	1.1	1173	--	Alteon Base Pack	n/a	6	--	--	--	em7admin	2018-07-19 23:27:21
8. Alteon: Performance	15 min.	SNMP Performance	Enabled	2.7	1175	--	Alteon Base Pack	n/a	6	--	--	--	em7admin	2018-07-19 23:27:22
9. APC: Environmental T/H	5 min.	SNMP Performance	Enabled	6.1	713	--	APC Base Pack	n/a	3	2	2	1	em7admin	2018-07-19 23:24:41
10. APC: Asset Tracking	1440 min.	SNMP Configuration	Enabled	1.1	707	--	APC Base Pack	n/a	11	--	--	--	em7admin	2018-07-19 23:24:41
11. APC: Battery Performance	15 min.	SNMP Performance	Enabled	1.1	708	--	APC Base Pack	n/a	7	--	--	--	em7admin	2018-07-19 23:24:41
12. APC: Battery Status	15 min.	SNMP Configuration	Enabled	1.2	709	--	APC Base Pack	n/a	9	12	12	2	em7admin	2018-07-19 23:24:41
13. APC: Configuration	360 min.	SNMP Configuration	Enabled	1.1	710	--	APC Base Pack	n/a	20	4	4	--	em7admin	2018-07-19 23:24:41
14. APC: Environmental Asset	1440 min.	SNMP Configuration	Enabled	6.1	714	--	APC Base Pack	n/a	12	--	--	--	em7admin	2018-07-19 23:24:41
15. APC: PDU Performance	15 min.	SNMP Performance	Enabled	1.3	712	--	APC Base Pack	n/a	3	--	--	--	em7admin	2018-07-19 23:24:41
16. APC: SmartUPS Power Output	5 min.	SNMP Performance	Enabled	6.1	715	--	APC Base Pack	n/a	2	--	--	--	em7admin	2018-07-19 23:24:42
17. APC: UPS Testing	1440 min.	SNMP Configuration	Enabled	1.2	711	--	APC Base Pack	n/a	6	9	9	--	em7admin	2018-07-19 23:24:41
18. Aruba: Asset	1440 min.	SNMP Configuration	Enabled	2.8	1434	--	Aruba Base Pack	n/a	18	--	--	--	em7admin	2018-07-19 23:28:45
19. Aruba: Voice Counters	5 min.	SNMP Performance	Enabled	2.8	1435	--	Aruba Base Pack	n/a	15	--	--	--	em7admin	2018-07-19 23:28:45






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

- **Dynamic Application Name.** Name of the Dynamic Application, as defined in the **Dynamic Applications Properties Editor** page.
- **Poll Rate.** Frequency, in minutes, at which the ScienceLogic platform will poll all devices that use this Dynamic Application.

NOTE: The **Poll Rate** column displays the default poll frequency for the Dynamic Application, as defined in the **Dynamic Applications Properties Editor** page. You can define a **custom** poll frequency for one or more devices in a device template. The poll frequency defined in the device template overrides the poll frequency defined for the Dynamic Application. Devices to which the device template is applied will use the poll frequency defined in the device template.

- **Type.** Type of Dynamic Application. The choices are:
 - *Bulk Snippet Configuration.* A single instance of the Dynamic Application uses custom-written Python code to collect static configuration data from *multiple devices*. This is useful for systems that include a large number of component devices. For details on creating bulk snippet Dynamic Applications, see the **Snippet Dynamic Application Development Manual**.
 - *Bulk Snippet Performance.* A single instance of the Dynamic Application uses custom-written Python code to collect trendable performance data from *multiple devices*. This is useful for systems that include a large number of component devices. For details on creating bulk snippet Dynamic Applications, see the **Snippet Dynamic Application Development Manual**.
 - *Database Configuration.* The Dynamic Application retrieves configuration data from a database application. The Dynamic Application uses SQL queries. The queried device returns table data. For details on creating database Dynamic Applications, see the **Database Dynamic Application Development Manual**.
 - *Database Performance.* The Dynamic Application retrieves trendable performance data from a database application. The Dynamic Application uses SQL queries. The queried device returns table data. For details on creating database Dynamic Applications, see the **Database Dynamic Application Development Manual**.
 - *Internal Collection Inventory.* The Internal Collection Inventory Dynamic Application (ICDA) retrieves configuration data about filesystems and interface. For filesystem, an ICDA Inventory can retrieve data such as storage size, filesystem type, and storage used. These ICDA's can also collect configuration data about interfaces, such as physical address, operational status, and IP addresses. For details on creating ICDA's, see the **Internal Collection Dynamic Application Development Manual**.
 - *Internal Collection Performance.* The Internal Collection Performance Dynamic Application (ICDA) retrieves data about availability and latency, device information (system description, system uptime, system locale), filesystem performance, and interface performance. For details on creating ICDA's, see the **Internal Collection Dynamic Application Development Manual**.
 - *IT Service.* A special type of Dynamic Application that the ScienceLogic platform uses to monitor IT Services. When you create and edit an IT Service in the **IT Service Editor** page, the ScienceLogic platform will automatically create and maintain a Dynamic Application for that IT Service. Dynamic Applications for IT Services will appear in the **Dynamic Applications Manager** page. However, if you want to edit the settings for an IT Service, you should not edit the Dynamic Application for that IT Service. Instead, use the **IT Service Editor** page to edit IT Services. For details on creating IT Service policies, see the **IT Services Manual**.

- *PowerShell Configuration*. The Dynamic Application uses PowerShell commands to collect static configuration data from a Windows device. For details on creating PowerShell Dynamic Applications, see the manual ***Dynamic Application Development - WMI and PowerShell***. For information on configuring the ScienceLogic platform and external systems to use PowerShell Dynamic Applications, see the manual ***Monitoring Windows Systems***.
- *PowerShell Performance*. The Dynamic Application uses PowerShell commands to collect trendable performance data from a Windows device. For details on creating PowerShell Dynamic Applications, see the manual ***Dynamic Application Development - WMI and PowerShell***. For information on configuring the ScienceLogic platform and external systems to use PowerShell Dynamic Applications, see the manual ***Monitoring Windows Systems***.
- *Snippet Configuration*. The Dynamic Application uses custom-written Python code to collect configuration data from a device. For details on creating snippet Dynamic Applications, see the ***Snippet Dynamic Application Development Manual***.
- *Snippet Journal*. The Dynamic Application uses custom-written Python code to collect data formatted as log entries from a device. For details on creating snippet Dynamic Applications, see the ***Snippet Dynamic Application Development Manual***.
- *Snippet Performance*. The Dynamic Application uses custom-written Python code to collect trendable performance data from a device. For details on creating snippet Dynamic Applications, see the ***Snippet Dynamic Application Development Manual***.
- *SNMP Configuration*. The Dynamic Application uses SNMP to retrieve static, configuration data from devices or applications. For details on creating SNMP Dynamic Applications, see the ***SNMP Dynamic Application Development Manual***.
- *SNMP Performance*. The Dynamic Application uses SNMP to retrieve trendable performance data from devices or applications. For details on creating SNMP Dynamic Applications, see the ***SNMP Dynamic Application Development Manual***.
- *SOAP Configuration*. The Dynamic Application uses XML and SOAP to retrieve static configuration data from a SOAP server. The queried device returns XML data. For details on creating SOAP Dynamic Applications, see the ***XML, SOAP, and XSLT Dynamic Application Development Manual***.
- *SOAP Performance*. The Dynamic Application uses XML and SOAP to retrieve trendable performance data from a SOAP server. The queried device returns XML data. For details on creating SOAP Dynamic Applications, see the ***XML, SOAP, and XSLT Dynamic Application Development Manual***.
- *WMI Configuration*. The Dynamic Application retrieves configuration information from either WMI or WBEM running on a managed device. WMI Dynamic Applications use a query format to request data from a managed device. WBEM Dynamic Applications use wbemcli and HTTP to request data from a managed device. For details on creating WMI Dynamic Applications, see the manual ***Dynamic Application Development - WMI and PowerShell***. For information on configuring the ScienceLogic platform and external systems to use PowerShell Dynamic Applications, see the manual ***Monitoring Windows Systems***.
- *WMI Performance*. The Dynamic Application retrieves trendable performance data from either WMI or WBEM running on a managed device. WMI Dynamic Applications use a query format to request data from a managed device. WBEM Dynamic Applications use wbemcli and HTTP to request data from a managed device.

- **XML Configuration.** The Dynamic Application uses HTTP GET queries. The queried device returns static configuration data in XML format. For details on creating SOAP Dynamic Applications, see the **XML, SOAP, and XSLT Dynamic Application Development Manual**.
 - **XML Performance.** The Dynamic Application uses HTTP GET queries. The queried device returns trendable performance data in XML format. For details on creating SOAP Dynamic Applications, see the **XML, SOAP, and XSLT Dynamic Application Development Manual**.
 - **XSLT Configuration.** The Dynamic Application uses XML and SOAP to retrieve static configuration data from a SOAP server. The requests used to retrieve data are generated by performing an XSLT transformation on an XML document that contains data already collected by the Dynamic Application. The queried device returns XML data, which must be changed to a specific format by performing a second XSLT transformation. For details on creating SOAP Dynamic Applications, see the **XML, SOAP, and XSLT Dynamic Application Development Manual**.
 - **XSLT Performance.** The Dynamic Application uses XML and SOAP to retrieve trendable performance data from a SOAP server. The requests used to retrieve data are generated by performing an XSLT transformation on an XML document that contains data already collected by the Dynamic Application. The queried device returns XML data, which must be changed to a specific format by performing a second XSLT transformation. For details on creating SOAP Dynamic Applications, see the **XML, SOAP, and XSLT Dynamic Application Development Manual**.
- **State.** Specifies whether the Dynamic Application is *Enabled* or *Disabled*.
 - **Version.** Version number to assign to the Dynamic Application. You can customize this value and increment it according to your change-management policies.
 - **ID.** Unique application ID, assigned by the ScienceLogic platform.
 - **Subscribers.** Number of devices that use the Dynamic Application. Clicking on the icon leads to the **Application Subscribers** modal page, where you can view the list of devices and access other pages for each subscriber device. You can also access this page by selecting the wrench icon () for a Dynamic Application and selecting the **[Subscribers]** tab.
 - **PowerPack.** Specifies whether or not the Dynamic Application is included in a PowerPack.
 - **Environment.** The execution environment to which the Dynamic Application is aligned, if it is a snippet or internal collection Dynamic Application. If it is not a snippet or internal collection Dynamic Application, then this column displays "n/a".
 - **Collects.** Number of objects included in the Dynamic Application. Clicking on the icon () leads to the **Collection Objects** page, where you can view the list of collection objects and edit their properties.
 - **Alerts.** Number of custom alerts defined for the Dynamic Application. Clicking on the icon () leads to the **Alert Objects** page, where you can view and edit each alert defined for the Dynamic Application.
 - **Events.** Number of events associated with the Dynamic Application. Clicking on the icon () leads to the **Event Policy Manager** page, where you can view information about each event definition associated with the Dynamic Application definition and edit each event definition.
 - **Thresh.** Number of threshold objects defined for the Dynamic Application. Clicking on the icon () leads to the **Threshold Objects** page, where you can view and edit information about each threshold object defined for the Dynamic Application.

- **Edited By.** Username of the person who created or last edited the Dynamic Application.
- **Last Edit.** Date that the Dynamic Application was created or last edited.

Searching and Filtering the List of Dynamic Applications

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

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

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

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

- **Dynamic Application Name.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching name.
- **Poll Rate.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching polling rate.
- **Type.** You can enter text to match, including *special characters*, , and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching type.
- **State.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching state.
- **Version.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching version number.
- **ID.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching ID number.
- **Subscribers.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching number of subscribers.
- **PowerPack.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching PowerPack.
- **Environment.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching execution environment.
- **Collects.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching number of collection objects.
- **Alerts.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching number of alerts.
- **Events.** You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching number of event policies.

- **Thresh**. You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that have a matching number of thresholds.
- **Edited By**. You can enter text to match, including *special characters*, and the Dynamic Applications Manager page will display only Dynamic Applications that were created or edited by a matching user-name.
- **Last Edited**. Only those Dynamic Applications that match all the previously selected fields and have the specified "last edited" date will be displayed. The choices are:
 - *All*. Display all Dynamic Applications that match the other filters.
 - *Last Minute*. Display only Dynamic Applications that have been modified within the last minute.
 - *Last Hour*. Display only Dynamic Applications that have been modified within the last hour.
 - *Last Day*. Display only Dynamic Applications that have been modified within the last day.
 - *Last Week*. Display only Dynamic Applications that have been modified within the last week.
 - *Last Month*. Display only Dynamic Applications that have been modified within the last month.
 - *Last Year*. Display only Dynamic Applications that have been modified within the last year.

Special Characters

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

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

String and Numeric

- , (comma). Specifies an "OR" operation. Works for string and numeric values. For example:
 "dell, micro" matches all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "AND" operation. Works for string and numeric values. For example:
 "dell & micro" matches all values that contain both the string "dell" AND the string "micro", in any order.

- ! (exclamation point). Specifies a "not" operation. Works for string and numeric values. For example:

"!dell" matches all values that do not contain the string "dell".

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

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

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

"! ^ " would match null values.

"!\$" would match null values.

"!*" would match null values.

"happy, !dell" would match values that contain "happy" OR values that do not contain "dell".

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

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

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

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

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

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

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

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

String

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

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

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

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

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

"! ^ " would match null values.

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

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

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

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

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

"!\$" would match null values.

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

Numeric

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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



Additional Examples

- "aio\$". Matches only text that ends with "aio".
- "^shu". Matches only text that begins with "shu".
- "^silo\$". Matches only the text "silo", with no characters before or after.
- "!silo". Matches only text that does not contain the characters "silo".
- "!^silo". Matches only text that does not start with "silo".
- "!0\$". Matches only text that does not end with "0".
- "!^silo\$". Matches only text that is not the exact text "silo", with no characters before or after.
- "!. Matches null values, typically represented as "--" in most pages.
- " !\$". Matches null values, typically represented as "--" in most pages.
- " !^\$. Matches all text that is not null.
- silo, !aggr". Matches text that contains the characters "silo" and also text that does not contain "aggr".
- "silo, 02, !aggr". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr".
- "silo, 02, !aggr, !01". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr" and also text that does not contain "01".
- "^s*i!*o\$". Matches text that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would match.
- " !^s*i!*o\$". Matches all text that does not contain the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would not match.
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" would match, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" would match, because it does not contain "vol" and it does contain "02".
- "aggr,!vol&02". Matches text that contains "aggr" OR text that does not contain "vol" AND also contains "02".
- "aggr,!vol&!infra". Matches text that contains "aggr" OR text that does not contain "vol" AND does not contain "infra".
- "*". Matches all text.
- " !*". Matches null values, typically represented as "--" in most pages.
- "silo". Matches text that contains "silo".
- " !silo ". Matches text that does not contain "silo".
- " !^silo\$ ". Matches all text except the text "silo", with no characters before or after.
- "-3,7-8,11,24,50-". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50.

- "-3,7-8,11,24,50-,a". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50, and text that includes "a".
- "?n". Matches text that contains any single character and the character "n". For example, this string would match "an", "bn", "cn", "1 n", and "2n".
- "n*SAN". Matches text the contains "n", zero or any number of any characters and then "SAN". For example, the string would match "nSAN", and "nhamburgerSAN".
- "^?n*SAN\$". Matches text that begins with any single character, is following by "n", and then zero or any number of any characters, and ends in "SAN".

Performing Other Tasks in the Dynamic Application Manager Page

From the **Dynamic Applications Manager** page, you can also perform the following tasks:


- **Manually perform dynamic discovery for a Dynamic Application.** In the **Dynamic Applications Manager** page, find the Dynamic Application you want to use for network-wide discovery and select its lightning bolt icon () .
- **Filter the list by application.** Selecting the Filter on this Application icon () sets the **Dynamic Applications Manager** page to display only Dynamic Applications with the same name as the selected Dynamic Application.
- **Change an application's type.** To change the type of one or more Dynamic Applications, select the checkbox () of each Dynamic Application you want to change. In the **Select Action** drop-down menu in the bottom right, scroll to the *Change Type* section and select a new type. Select the **[Go]** button. Each selected Dynamic Application will now be assigned the new type.
- **Change polling interval.** To change the polling interval for one or more Dynamic Application, select the checkbox () of each Dynamic Application you want to change. In the **Select Action** drop-down menu in the bottom right, select a new polling interval for the Dynamic Application. Select the **[Go]** button. Each selected Dynamic Application will now be assigned the new polling interval.
- **Delete an Application.** To delete one or more Dynamic Applications, select the checkbox () of each Dynamic Application you want to change. In the **Select Action** drop-down menu in the bottom right, select *DELETE Application*, and then select the **[Go]** button. Each selected Dynamic Application will be deleted.
- **Clear Application Data.** To clear all historical data, including logs and reports, associated with one or more Dynamic Application, select the checkbox () for each Dynamic Application you want to change. In the **Select Action** drop-down menu in the bottom right, select *CLEAR Application*, and then select the **[Go]** button. The historical data for each selected Dynamic Application will be deleted.
- **Discover Applications using the Select Action drop-down menu.** To start the Dynamic Application discovery process on all managed devices, select the checkbox () for each Dynamic Application you want to use in the discovery process. In the **Select Action** drop-down menu in the bottom right, select *DISCOVER Applications*. Select the **[Go]** button. The ScienceLogic platform starts the Dynamic Application discovery process on all devices in the ScienceLogic platform using the selected Dynamic Applications.

- **Validate and Repair Applications.** For each selected Dynamic Application, ensures that the Dynamic Application includes the objects, thresholds, and alerts that are referenced in the Dynamic Application and in other parts of the ScienceLogic platform. Also ensures that if an event that is associated with the Dynamic Application has an auto-clear event, that auto-clear event is also associated with the Dynamic Application. This tool will troubleshoot and then fix the database tables that reference the Dynamic Application's objects, presentations, thresholds and alerts.
- **Align a Device Dashboard.** To align one or more Dynamic Applications with a Device Dashboard, select the checkbox for each Dynamic Application you want to align with the Device Dashboard. In the **Select Action** drop-down menu in the bottom right, scroll to the *Change Device Dashboard* section and select a Device Dashboard. Select the **[Go]** button. For each device that subscribes to the selected Dynamic Application, the selected device dashboard will appear as an entry in the **Device Dashboard** field in the **Device Summary** page.
- **Align an Execution Environment.** To align one or more Dynamic Applications with an execution environment, select the checkbox for each Dynamic Application you want to align with the environment. In the **Select Action** drop-down menu in the bottom right, scroll to the **Align Collection Environment** section and select an execution environment, and then click the **[Go]** button.

Viewing the Dynamic Applications Associated with a Device

You can view the Dynamic Applications associated with a device from the **Dynamic Application Collections** page. The **Dynamic Application Collections** page also allows you to [associate another Dynamic Application with a device](#), [assign a credential to a Dynamic Application](#), or [perform other administrative tasks for an aligned Dynamic Application](#).

To view the Dynamic Applications associated with a device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to view Dynamic Applications. Select its wrench icon (.
3. In the **Device Administration** panel, select the **[Collections]** tab.

- The **Dynamic Application Collections** page displays a list of all Dynamic Applications aligned with the current device. For each Dynamic Application, the **Dynamic Application Collections** page displays the following read-only information:

Close	Properties	Thresholds	Collections	Monitors	Schedule		
Logs	Toolbox	Interfaces	Relationships	Tickets	Redirects	Notes	Attributes
Device Name	10.2.117.3			Managed Type	Physical Device		
IP Address / ID	10.2.117.3 135			Category	Storage Array		
Class	Pure Storage			Sub-Class	FlashArray Storage System		
Organization	Knights_Pure			Uptime	0 days, 00:00:00		
Collection Mode	Active			Collection Time	2018-07-20 14:07:00		
Description				Group / Collector	CUG-Knights-OL7 10-64-171-139-CU-Knights		
Device Hostname							

Dynamic Application™ Collections							Expand	Actions	Reset	Guide
Dynamic Application	ID	Poll Frequency	Type	Credential	Collector					
+ Pure Storage: Array Capacity Stats	1552	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Array Stats	1543	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Kr					
- Pure Storage: Temperature Stats	1553	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Kr					
Presentation Object *										
+ [Temperature [°C]		1		p_6298	yes	yes	--	--	0	
+ [Temperature [°F]		1		p_6299	yes	yes	--	--	0	
Misc Collection Object *										
[Discovery Object				o_17691	no	yes	--	--		
[Temperature (Labels)				p_17690	yes	yes	--	--		
+ Pure Storage: Array Discovery	1542	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Controller Config	1554	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Drive Config	1550	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Hardware Config	1551	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Hosts & Groups Config	1558	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Message Log Config	1555	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Protection Groups Config	1556	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Volume Discovery	1544	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					
+ Pure Storage: Volume Snapshots Config	1548	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Kr					

[Select Action] [Go]

Save


- **Plus Sign (+)**. Clicking on this icon displays a list of all Presentation Objects included in Dynamic Applications of type "Performance" and "Journal" or a list of all Collection Objects included in Dynamic Applications of type "Configuration". You can click on the plus sign next to each Presentation Object to see all the Collection Objects included in the Presentation Object.
- **Minus Sign (-)**. Collapses a Dynamic Application and hides the display of Presentation Objects and Collection Objects.
- **Dynamic Application**. Name of the Dynamic Application.
- **ID**. Numeric ID for the Dynamic Application.
- **Poll Frequency**. Frequency at which the ScienceLogic platform will query the device to retrieve the data specified in the Dynamic Application. Each Dynamic Application includes a default frequency. From this page (**Dynamic Application Collections**), you can change the poll frequency for a Dynamic Application on the current device. This edited poll frequency will override the default frequency for the Dynamic Application and the poll frequency defined for a Dynamic Application in one or more device templates.
- **Type**. The protocol used by the Dynamic Application (Database [SQL], Internal Collection Inventory or Internal Collection Performance (ICDA), Snippet [Python], SNMP, SOAP, WMI, XML, or XSLT) and the type of data collected by the Dynamic Application (Configuration, Performance, or Journal).


- **Credential.** Name of the credential that the ScienceLogic platform uses to access the device and retrieve the data specified in the Dynamic Application.

NOTE: Cache-consuming Dynamic Applications do not require a credential. If you aligned a cache-consuming Dynamic Application in the **Dynamic Application Alignment** modal page, the **Credential** field displays *N/A* and is grayed out. You do not have to select a credential in the **Dynamic Application Alignment** modal page.

- **Collector.** Name of the specific Data Collector used to collect data from the Dynamic Application.

NOTE: Based on the Dynamic Application's **Collector Affinity** settings, the Dynamic Application might be assigned to a different Data Collector than the Data Collector that is assigned to the device in the Device Properties page (Registry > Devices > wrench icon). In the **Dynamic Application Collections** page, hover your mouse over the **Collector** name for any of the Collection Objects to view a tooltip that explains why the Dynamic Application is assigned to its particular Data Collector.

- **Run Dynamic Application** (). Performs a test run of data collection for the selected Dynamic Application on the current device.

NOTE: If a device is currently unavailable, the lightning-bolt icon () will be grayed out for each Dynamic Application aligned with the device.

- **Checkbox** (). Apply an action from the **Select Action** field to this instance of the Dynamic Application.

5. From this page, you can [associate another Dynamic Application with the device](#), [assign a credential to the Dynamic Application](#), or [perform other administrative tasks for an aligned Dynamic Application](#).

Viewing the Status of a Dynamic Application

For each device, the ScienceLogic platform maintains the collection status for each collection object in each Dynamic Application aligned with that device. The **Dynamic Application Collections** page displays the status of each collection object for a device as represented by two values: **Found** and **Collect**. The **Dynamic Application Collections** page also displays the **Found** and **Collect** values for each presentation object, which are derived from the status of each collection object used by the presentation object.

Found

The **Found** status for a collection object has two possible values:

- Yes. Data has been successfully collected from this device for this object. **Found** is set to Yes the first time data is successfully collected from this device for this object.
- No. Data has never been successfully collected from this device for this object. **No** is the initial value of **Found** for every object when a Dynamic Application is initially aligned with a device.

The **Found** status for a presentation object also has two possible values (Yes and No).

- **If the presentation object uses only one collection object**, the presentation object always has the same default **Found** and default **Collect** values as that collection object.
- **If a presentation object uses multiple collection objects**, the default **Found** value for the presentation object will be Yes only if all the collection objects used by the presentation object have a **Found** value of Yes.

After **Found** is set to Yes for an object, the ScienceLogic platform will never automatically change the value of **Found** for this object.

The value of **Found** is used by the ScienceLogic platform to determine whether icons, tabs, and Navbar links that lead to the **[Performance]** or **[Configs]** page where the collection object is used should be active.

Collect

The **Collect** status for a collection object has two possible values:

- Yes. The ScienceLogic platform will attempt to collect data for this object when collection for this Dynamic Application occurs. Yes is the initial value for **Collect** for every object when a Dynamic Application is initially aligned with a device.
- No. The ScienceLogic platform will not attempt to collect data for this object when collection for this Dynamic Application occurs. The platform might set **Collect** to **No** automatically if no data has been collected.
- If a collection object has a **Collect** value of **No**, all presentation objects that use that collection object will also have a **Collect** value of **No**.

The **Collect** status for a presentation object also has two possible values (Yes and No).

- **If the presentation object uses only one collection object**, the presentation object always has the same default **Found** and default **Collect** values as that collection object.
- **If a presentation object uses multiple collection objects**, the default **Collect** value for the presentation object will be Yes only if **all** the collection objects used by the presentation object have a **Collect** value of Yes. If one or more collection objects used by the presentation object have a **Collect** value of **No**, the presentation object will also have a default **Collect** value of **No**.
- The **Collect** status for a presentation object has no effect upon its collection objects. If you manually change the **Collect** status for a presentation object, the **Collect** status for the collection objects used by the presentation object will not change.

NOTE: Before determining which collection objects defined in a Dynamic Application will be collected, the platform determines whether the Dynamic Application itself should be collected. Dynamic Applications are not collected for devices that are unavailable (because of a failed availability check) or have collection disabled (either manually by a user or because of maintenance scheduled in the platform) regardless of the **Collect** value of the objects.

How the ScienceLogic Platform Manages Collect Status

Stopping Collection

One of the ScienceLogic hourly maintenance tasks checks the last collection time for every collection object being collected from every device. If the last collection time for an object on a device is more than two days ago, collection is stopped for that collection object on that device. The ScienceLogic platform will set the **Collect** status of that object to *No*.

NOTE: If a device is in maintenance mode, is unavailable, or has been manually disabled by a user, the ScienceLogic platform will not automatically set the **Collect** status of objects to *No*. The ScienceLogic platform will automatically set the **Collect** status of objects to *No* only if the device is up and running, but the platform still cannot collect the object.

When the ScienceLogic platform sets the **Collect** status of that object to *No*, the platform generates an event. The event will include the name of the device, the name of the Dynamic Application, the name of the collection object, and the collection object IDs. By default, this event is of severity "notice".

NOTE: For Dynamic Applications that have the **Component Mapping** checkbox selected in the **Dynamic Applications Properties Editor** page, the ScienceLogic platform will never automatically set the **Collect** status to *No* for any of the collection objects in the Dynamic Application.

NOTE: For Dynamic Applications that have the **Caching** fields set to either *Cache Results* or *Consume cached results* in the **Dynamic Applications Properties Editor** page, the ScienceLogic platform will never automatically set the **Collect** status to *No* for any of the collection objects in the Dynamic Application.

Starting Collection

For each object that has the **Collect** status of *No*, the ScienceLogic platform will attempt to re-collect the object once a day. If re-collection is successful, the ScienceLogic platform will automatically set the **Collect** value for that object to *Yes*.

NOTE: If a user manually sets the **Collect** status of a collection object or presentation object to *No*, the ScienceLogic platform will **not** attempt to re-collect the object once a day and will **not** set the **Collect** status to *Yes*.

Collection Objects that are Excluded from Maintenance

The **Collect** status of the following collection objects is never changed automatically:

- Collection objects in Dynamic Applications that have the **Component Mapping** checkbox checked in the **Dynamic Applications Properties Editor** page.
- Collection objects in Dynamic Applications that have the **Caching** fields set to either *Cache Results* or *Consume cached results*, in the **Dynamic Applications Properties Editor** page.
- Collection objects that have the **Disable Object Maintenance** setting enabled.
- Collection objects that have a **Collect** status defined by a user, i.e. collection objects that were manually enabled or disabled by a user.

Status of Objects for Deviation

The ScienceLogic platform allows you to examine the value of an object and trigger an alert if that value falls outside the range of "normal" values for that object at the hour of the day on that day of the week. The deviation function allows you to define such alerts.

To use the deviation function, you must configure the ScienceLogic platform to store and calculate the mean values and standard deviation for an object. You do this by selecting the **Enable Deviation Alerting** field in the **Collection Objects** page. You then specify the minimum and maximum number of weeks to collect deviation data for the object. The ScienceLogic platform must have already collected at least the minimum number of weeks' worth of values for an object before the ScienceLogic platform can evaluate alert formulas that use the deviation function. To use the deviation function, you must specify a minimum value of at least two weeks.

If a Dynamic Application in the **Dynamic Application Collections** page contains one or more alerts that use the deviation function, the **Dynamic Application Collections** page displays the status of the collection objects.

For example, suppose an alert in a Dynamic Application will apply the deviation function to object "o_123". Suppose that you specified that the ScienceLogic platform must collect at least two weeks' worth of deviation data for this object. Suppose that the ScienceLogic platform contains only one weeks' worth of values for object "o_123". In this case, the **Dynamic Application Collections** page will display the following message:

```
Note: object 123 not ready for deviation alerting.
```

When the ScienceLogic platform contains at least two weeks worth of values for object "o_123", the **Dynamic Application Collections** page will display the following message:

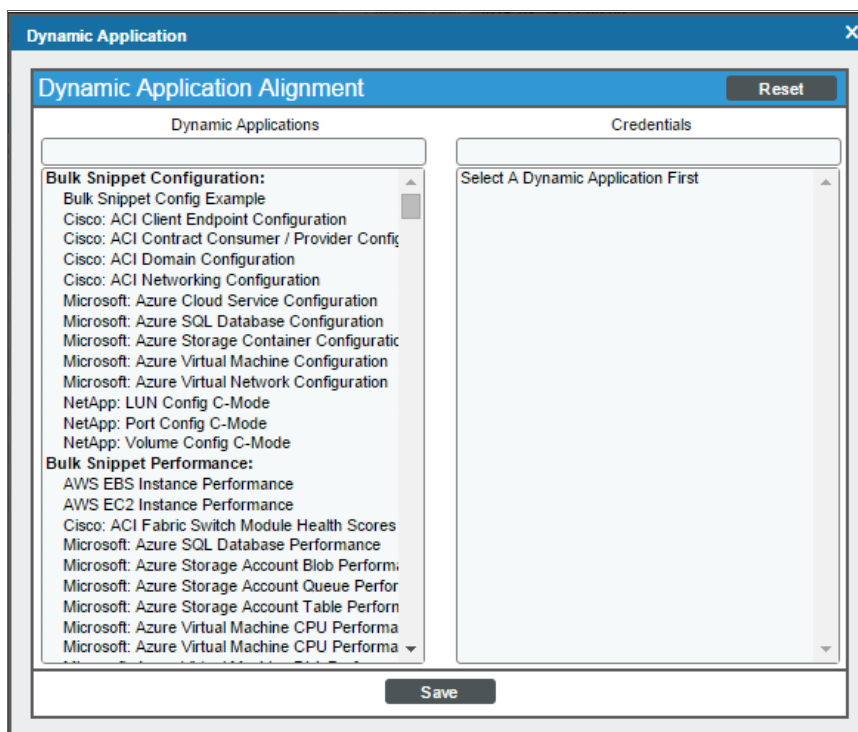
```
All objects ready for deviation alerting.
```

Manually Associating a Dynamic Application with a Device

From the **Dynamic Application Collections** page, you can manually associate a new Dynamic Application with a device.

To manually associate a Dynamic Application with a device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device you want to associate with a Dynamic Application. Click its wrench icon (🔧).
3. In the **Device Administration** panel, click the **[Collections]** tab.
4. In the **Dynamic Application Collections** page, click the **[Actions]** menu and select *Add Dynamic Application*.
5. The **Dynamic Application Alignment** modal page appears. To align a Dynamic Application with a device in this page:



- Select the Dynamic Application you want to align with the device in the **Dynamic Applications** field. You can filter the list of Dynamic Applications using the search field above the **Dynamic Applications** field.

- After selecting a Dynamic Application, you must select a credential. Select a credential in the **Credentials** field. You can filter the list of credentials using the search field above the **Credentials** field.

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

NOTE: Cache-consuming Dynamic Applications **do not** require a credential. If you selected a cache-consuming Dynamic Application in the **Dynamic Application Alignment** modal page, the **Credential** field displays *N/A* and is grayed out. You do not have to select a credential in the **Dynamic Application Alignment** modal page.


6. Click the **[Save]** button in the **Dynamic Application Alignment** modal page to align the Dynamic Application and the credential to the device.
7. The ScienceLogic platform will associate the Dynamic Application with the device and immediately attempt to collect the data specified in the Dynamic Application using the selected credential.
8. After the first, immediate collection, the ScienceLogic platform will collect the data at the frequency defined in the **Polling Frequency** field in the **Application Configuration Editor** page for the Dynamic Application.

Editing the Credential Associated with a Dynamic Application

From the **Dynamic Application Collections** page, you can change the credential associated with a Dynamic Application. This credential will be used by the ScienceLogic platform for this specific Dynamic Application associated with this specific device. For all other devices, the ScienceLogic platform will use the default credential associated with the device, or will use the credential defined in the **Dynamic Application Collections** page for each device.

NOTE: Cache-consuming Dynamic Applications do not require a credential. If you aligned a cache-consuming Dynamic Application with this device (you do this in the **Dynamic Application Alignment** modal page), the **Credential** field displays *N/A* and is grayed out.

To change the credential associated with a Dynamic Application for a device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to define a credential. Select its wrench icon () .
3. In the **Device Administration** panel, select the **[Collections]** tab.
4. In the **Dynamic Application Collections** page, find the Dynamic Application for which you want to change the credential. Select its checkbox (). To apply a credential to multiple Dynamic Applications, select the checkbox for each Dynamic Application.

- From the **Select Action** drop-down list, select the credential from the list of all credentials that you are allowed to use, and then select the **[Go]** button.

The screenshot shows a network management interface with a top navigation bar (Close, Properties, Thresholds, Collections, Monitors, Schedule, Logs, Toolbox, Interfaces, Relationships, Tickets, Redirects, Notes, Attributes). Below the navigation bar is a device information section for a Pure Storage device (IP: 19.2.117.3, Class: Pure Storage, Organization: Knights_Pure). The main area displays a table of Dynamic Application Collections. A dropdown menu is open over the table, showing a list of credentials to select from.

Dynamic Application™ Collections	Dynamic Application	ID	Pol.Frequency	Type	Credential	Collector
+ Pure Storage: Array Capacity Stats		1552	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Array Stats		1543	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Temperature Stats		1553	5 mins	Snippet Performance	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Array Discovery		1542	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Controller Config		1554	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Drive Config		1550	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Hardware Config		1551	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Hosts & Groups Config		1558	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Message Log Config		1555	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Protection Groups Config		1556	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Volume Discovery		1544	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights
+ Pure Storage: Volume Snapshots Config		1548	15 mins	Snippet Configuration	Knights - Pure Storage	10-64-171-139-CU-Knights

The dropdown menu shows the following list of credentials:

- [Select Action]
- Assign Default Credential:
- [Default SNMP Credential]
- _Cisco SNMPV2 - Example
- _Cisco SNMPV3 - Example
- _Cisco_CSP SNMP Port 161 Example
- _Cisco_CSP SNMP Port 1610 Examp
- _Dell EMC Isilon SNMPV2 Example
- _EMT Default V2
- _EMT Default V3
- _IPSLA Example
- _Knights - Dell EMC Isilon SNMPv
- _LifeSize Endpoint SNMP
- _SAC_chris_internal_collection
- _SAC_sanity_std_Deviation
- _SAC_Streaming_Rollups_CBQoS
- _SAC_Streaming_Rollups_Windows
- _SNMP ciscadia V2
- _SNMP Public V1
- [Select Action]

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

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


- You should see your change reflected in the **Credential** column in the **Dynamic Application Collections** page.

Performing Other Administrative Tasks for an Aligned Dynamic Application

You can perform the following other administrative tasks for an aligned Dynamic Application in the **Dynamic Application Collections** page:

- Enable or disable one or more collection objects or presentation objects.
- Stop data collection for the whole Dynamic Application.
- Reset the statistical data that has been stored for standard deviation alerting.
- Reset persistent session objects that have been collected and stored for a Dynamic Application.
- Test collection for a Dynamic Application.
- Remove all data collected using the Dynamic Application and optionally unalign the Dynamic Application from the device.

To perform one of these tasks:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to perform an administrative task. Select its wrench icon (.
3. In the **Device Administration** panel, select the **[Collections]** tab.
4. In the **Dynamic Application Collections** page, find the Dynamic Application for which you want to perform an administrative task. The following sections describe how to perform each task.

Enabling or Disabling Objects

From the **Dynamic Application Collections** page, you can customize the collection performed by the Dynamic Application for the current device. This customization will be used by the ScienceLogic platform only for this specific device. For all other devices, the platform will use the default list of objects from the Dynamic Application's definition or will use the list of objects defined in the **Dynamic Application Collections** page for that device.

NOTE: If a collection object has a **Collect** value of *No*, all presentation objects that use that collection object will also have a **Collect** value of *No*.

To enable or disable collection for one or more objects in a Dynamic Application:

- **To disable collection for one or more collection objects**, unselect the checkbox for each object for which you want to disable collection.
- For each unselected object, the **Collect** column should now display *No*.
- **To enable collection for one or more collection objects**, select the checkbox for each object for which you want to enable collection.
- For each selected object, the **Collect** column should now display *Yes*.
- Select the **[Save]** button.

NOTE: If a user **manually** sets the **Collect** status of a collection object or presentation object to *No*, the platform will **not** attempt to re-collect the object once a day and will **not** automatically set the **Collect** status to *Yes*.

Restarting Automatic Maintenance of Collection Objects

If a user **manually** sets the **Collect** status of a collection object or presentation object, the ScienceLogic platform will **not** automatically change the **Collect** status of that object as described in the [How the ScienceLogic Platform Manages Collect Status](#) section.

If you want the ScienceLogic platform to restart automatic maintenance of the objects in a Dynamic Application, perform the following steps:

1. In the **Dynamic Application Collections** page, select the checkbox () for the Dynamic Application for which you want to restart automatic collection maintenance. To restart automatic collection maintenance for multiple Dynamic Applications, select the checkbox for each Dynamic Application.
2. From the **Select Action** drop-down list, select *Restore System Control of Collection State* and then select the **[Go]** button.
3. Automatic collection maintenance for all objects in the Dynamic Application will now occur. The **Collect** status of the objects in the Dynamic Application will not change immediately.

Editing the Poll Frequency for a Dynamic Application on the Current Device

Poll Frequency is the frequency at which the ScienceLogic platform will query the device to retrieve the data specified in the Dynamic Application. Each Dynamic Application includes a default frequency.

From the **Dynamic Application Collections** page, you can change the poll frequency for a Dynamic Application on the current device. For the current device, the edited poll frequency will override:

- the default frequency for the Dynamic Application.
- the poll frequency defined for a Dynamic Application in one or more device templates.

To edit the poll frequency for a Dynamic Application on the current device:

1. In the **Dynamic Application Collections** page, select the checkbox for the Dynamic Application for which you want to change the poll frequency. To change the poll frequency for multiple Dynamic Applications, select the checkbox for each Dynamic Application.
2. From the **Select Action** drop-down list, select *Poll Frequency* from the list of poll frequencies and then select the **[Go]** button.
3. You should see your change reflected in the **Poll Frequency** column in the **Dynamic Application Collections** page.

Stopping Data Collection for a Dynamic Application

You can stop data collection for a Dynamic Application on the current device. This will affect collection only for this specific device. For all other subscriber devices, the ScienceLogic platform will continue to use this Dynamic Application to collect data.

To stop data collection for a Dynamic Application on this device:

1. Select the checkbox of each Dynamic Application for which you want to stop data collection.
2. From the **Select Action** drop-down list, select the following:
 - **Disable All Collection Objects.** For all collection objects in the selected Dynamic Application(s), the **Collect** value will be set to *No*.
3. Select the **[Go]** button.

NOTE: If a user manually sets the **Collect** status of a collection object or presentation object to *No*, the ScienceLogic platform will not attempt to re-collect the object once a day and will not set the **Collect** status to *Yes*.

Resetting Statistical Data for a Dynamic Application

The ScienceLogic platform allows you to examine the value of an object and trigger an alert if that value falls outside the range of "normal" values for that object at that hour of the day on that day of the week. The deviation function allows you to define such alerts.

To use the deviation function, you must configure the ScienceLogic platform to store and calculate the mean values and standard deviation for an object. You do this by selecting the **Enable Deviation Alerting** field in the **Collection Objects** page. You then specify the minimum and maximum number of weeks to collect deviation data for the object. The ScienceLogic platform must have already collected at least the minimum number of weeks' worth of values for an object before the ScienceLogic platform will evaluate alert formulas that use the deviation function. To use the deviation function, you must specify a minimum value of at least two weeks.

In some cases, you might want to delete all the collected statistics for an object and start over. This is useful if known circumstances change the value of an object, and you no longer want to use the old data to calculate the "normal" ranges. You can do this by "resetting" the statistical data for an object.

For example, suppose you were monitoring bandwidth usage with a standard deviation alert. Suppose your company previously ran on a 09:00 to 17:00 work schedule. Suppose your company has recently added a nightshift to the schedule. In this circumstance, you might want to reset the statistical data to determine the new "normal" usage patterns.

When you reset the statistical data for an object, you are telling the ScienceLogic platform to ignore all previously collected values and to use only values from today onward. When you reset the statistical data for an object, the **Dynamic Application Collections** page will again display a message like:

```
Note: object 123 not ready for deviation alerting.
```

until enough data has been collected to again calculate standard deviation for the object. The ScienceLogic platform will again start collecting the minimum number of weeks of data for the object (as specified in the **Enable Deviation Alerting** field in the **Collection Objects** page) and calculating the "normal" ranges for those objects for each hour at each day of the week.

To delete all current statistical data for an object:

1. In the Dynamic Application, find the object for which you want to reset data.
2. In that Dynamic Application, find the object for which you want to reset data. Select its checkbox .
3. From the **Select Action** drop-down list, select the following option:
 - *Reset Statistical Data*. Removes all previously collected statistical data for the selected object. The ScienceLogic platform will again start collecting the minimum number of weeks of data for the object (as specified in the **Enable Deviation Alerting** field in the **Collection Objects** page) and calculating the "normal" ranges for those objects for each hour at each day of the week.
4. Select the **[Go]** button.
5. The **Dynamic Application Collections** page will display a message like:

```
Note: object 123 not ready for deviation alerting.
```

Resetting Persistent Session Objects for a Dynamic Application

SOAP or XSLT Dynamic Applications can contain a collection object that stores a Session ID. The value for this collection object can be defined as a persistent value. If the ScienceLogic platform has already retrieved and stored a value in the collection object for the Session ID, the ScienceLogic platform will not collect a new value for the collection object until a SOAP fault occurs. You can force the ScienceLogic platform to re-collect a Session ID collection object by deleting the current persistent value.

To delete the current persistent value for a session object:

1. In the Dynamic Application, find the object for which you want to reset data. Select its checkbox .

- From the **Select Action** drop-down list, select the following option:
 - Reset Persistent Session Objects.** Removes the stored value for collection objects of type **SOAP/XSLT Session ID**. **SOAP/XSLT Session ID** objects are persistent across collection periods; the ScienceLogic platform does not collect a **SOAP/XSLT Session ID** object if a collected value is available from a previous poll. After selecting this option, the ScienceLogic platform will delete the existing value for the object and collect a new value during the next collection.
- Select the **[Go]** button.

Testing Data Collection for a Dynamic Application

On a single device, you can perform a test-run of collection with a single Dynamic Application. During this test run, the ScienceLogic platform displays details of each step of the collection process. This information can be very helpful for troubleshooting and debugging.

NOTE: During a test run of a collection with a Dynamic Application, the ScienceLogic platform does not store the collected data or generate alerts. The ScienceLogic platform will continue to collect data and generate alerts using the selected Dynamic Application at the frequency defined in the Dynamic Application.

To execute a test run of collection with a single Dynamic Application:

- Find the Dynamic Application for which you want to test collection and click its lightning bolt icon (⚡).

NOTE: If a device is currently unavailable, the lightning bolt icon (⚡) will be grayed out for each Dynamic Application aligned with the device.

- The ScienceLogic platform displays a **Session Logs** modal page that includes details about each step of the collection process and diagnostic details about alerts in the Dynamic Application. This information can be helpful during troubleshooting.

Removing Data Collected by a Dynamic Application

You can remove the data retrieved with a Dynamic Application from the current device. You have two options for removing Dynamic Application data associated with a device:

- Remove all previously collected data, but continue to collect data at the specified polling frequency.
- Remove all normalized data, but retain all raw collected data and continue to collect data at the specified polling frequency.
- Remove all previously collected data and stop collecting data with this Dynamic Application. This unaligns the device from the Dynamic Application. The device will no longer be a subscriber to the Dynamic Application.


To remove Dynamic Application data associated with a device:

1. In the **Dynamic Application Collections** page, select the checkbox (☑) of the Dynamic Application for which you want to remove data. To remove data for multiple Dynamic Applications, select the checkbox for each Dynamic Application.
2. From the **Select Action** drop-down list, select one of the following options:
 - **Remove Data**. Removes all previously collected data, but data will continue to collect at the specified polling frequency.
 - **Remove Normalized Data**. Removes all normalized data, but all raw collected data is retained and data will continue to collect at the specified polling frequency.
 - **Stop Collection and Remove Data**. Removes all previously collected data and stops collection of data with this Dynamic Application. This "unaligns" the device from the Dynamic Application. The device is no longer considered a subscriber to the Dynamic Application. If you perform this option and later want to subscribe to this Dynamic Application again, you must re-align the device with the Dynamic Application.
3. Select the **[Go]** button.

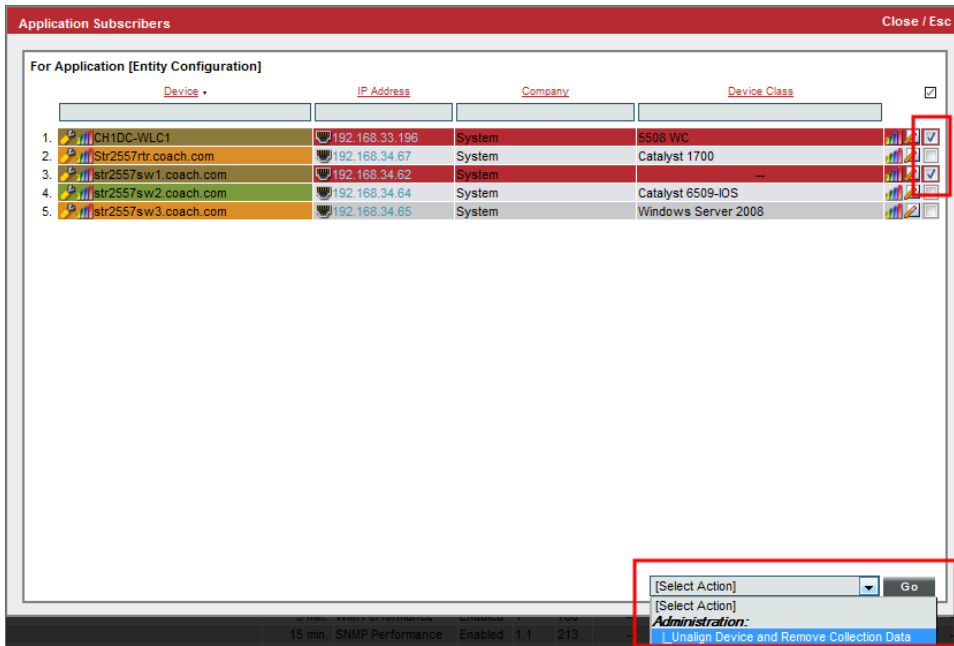
Bulk Un-aligning Dynamic Applications

The **Application Subscribers** page contains a drop-down field in the lower right called **Select Action**. This field allows you to un-align a Dynamic Application from one or more subscriber-devices.

To un-align a Dynamic Application from one or more devices:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. In the **Dynamic Applications Manager** page, find an application with a subscriber icon () in the **Subscribers** column. Select the icon.


3. The **Application Subscribers** page appears.



4. In the **Application Subscribers** page, select the checkbox for each device you want to apply the action to. To select all checkboxes for all devices, select the red checkbox () at the top of the page.
5. In the **Select Action** drop-down list, select one of the following actions.
 - **Unalign Device and Remove Collection Data**. This option unaligns the device from the Dynamic Application and deletes all data collected by the Dynamic Application from the device. The device is no longer considered a subscriber to the Dynamic Application. If you perform this option and later want to subscribe to this Dynamic Application again, you must re-align the device with the Dynamic Application.
6. Select the **[Go]** button to apply the action to all selected devices.

Setting Thresholds for Dynamic Applications

If a Dynamic Application includes one or more **thresholds**, you can change the threshold value on a per-device basis. To change a Dynamic Application threshold for a device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to define a threshold. Select its wrench icon () .
3. In the **Device Administration** panel, select the **[Thresholds]** tab.

- The **Device Thresholds** page displays a list of thresholds defined for each Dynamic Application that is aligned to the device. To change a threshold, move the slider for that threshold or enter a value in the number field for that threshold:

Section	Threshold Name	Current Value	Default Value
Operating System Thresholds	System Latency	100 ms	100 ms
	System Availability	99 %	99 %
Data Retention Thresholds	Device Logs Max	50000 records	50000 records
	Device Logs Age	90 days	90 days
	Bandwidth Data	365 days	365 days
	Daily Rollup Bandwidth Data	365 days	365 days
	Frequent Rollup Bandwidth Data	365 days	365 days
	Hourly Rollup Bandwidth Data	365 days	365 days
	Raw Performance Data	14 days	14 days
	Daily Rollup Performance Data	720 days	720 days
	Frequent Rollup Performance Data	60 days	60 days
	Hourly Rollup Performance Data	365 days	365 days
	Journal Data	60 days	60 days
	Configuration Data	30 records	30 records

- After changing one or more thresholds, select the **[Save]** button to save your changes.

NOTE: Changing a threshold in the **Device Thresholds** page affects only the current device. The threshold values defined in the Dynamic Application remain unchanged.

Dynamic Applications and Discovery

Discovery is the ScienceLogic tool that automatically discovers devices in your network. You supply the discovery tool with a range or list of IP addresses, and the discovery tool determines if a device exists at each IP address. The discovery tool also determines which (if any) Dynamic Applications to align with the device. If the discovery tool finds Dynamic Applications to align with the device, the discovery tool triggers collection for each aligned Dynamic Application.

To learn more about discovery, see the **Discovery and Credentials** manual.

How Does the ScienceLogic Platform Align Dynamic Applications During Discovery?

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

During discovery, the ScienceLogic platform:

- Searches the list of Dynamic Applications.

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

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

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

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

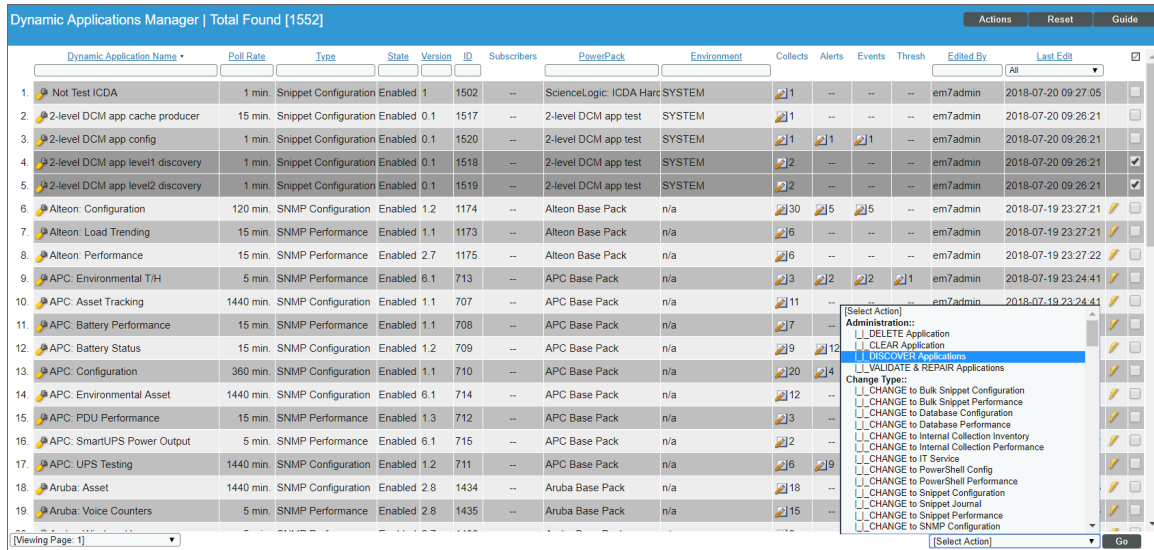
Queuing Discovery from the Dynamic Applications Manager Page

From the **Dynamic Applications Manager** page, you can manually run the Dynamic Application alignment portion of discovery for all devices in the system using one or more selected Dynamic Applications.

To manually queue discovery from the **Dynamic Applications Manager** page:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).

2. In the **Dynamic Applications Manager** page, select the checkbox for each Dynamic Application you want to use for discovery.
3. In the **Select Action** drop-down list, select *Discover Applications*. Select the **[Go]** button.



4. You can also run the Dynamic Application alignment portion of discovery for all devices in the system using a single Dynamic Application. To do this, select the lightning bolt icon (⚡) for that Dynamic Application.

System Settings That Affect Dynamic Application Discovery

Some of the parameters in the **Behavior Settings** page affect how the ScienceLogic platform aligns Dynamic Applications during discovery (discovery, auto-discovery, and re-discovery).

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

1. Go to the **Behavior Settings** page (System > Settings > Behavior).
2. In the **Behavior Settings** page, edit the values in one or more of the following fields:
 - **Initial Discovery Scan Level**. Specifies the data to be gathered during the initial discovery session. The options are:

NOTE: You can override this setting for a single discovery session in the **Discovery Session Editor** modal page.

- *0. Model Device Only*. Discovery tool will discover if device is up and running and if so, collect the make and model of the device. The platform will then generate a device ID for the device, so it can be managed by the platform.

- *1. Initial Population of Apps*. Discovery tool will search for Dynamic Applications to associate with the device. Discovery tool will attempt to collect data for the aligned Dynamic Applications. Discovery will also perform *0. Model Device Only* discovery.
- *2. Discover SSL Certificates*. Discovery tool will search for SSL certificates and retrieve SSL data. Discovery tool will also perform *1. Initial Population of Apps* and *0. Model Device Only*.
- *3. Discover Open Ports*. Discovery tool will search for open ports. Discovery tool will also perform *2. Discover SSL Certificates*, *1. Initial Population of Apps*, and *0. Model Device Only*.

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

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

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

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

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

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

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

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

NOTE: You can override this setting for a single discovery session in the **Discovery Session Editor** modal page.

- *0. Disabled.* The platform will scan only the primary IP address (the one used to communicate with the platform) for open ports.
 - *1. Enabled.* The platform will scan all discovered IP addresses for open ports.
- **Port Scan Timeout.** Length of time, in milliseconds, after which the platform should stop trying to scan an IP address for open ports and begin scanning the next IP address (if applicable). Choices are between 60,000 and 1,800,000 milliseconds.

NOTE: You can override this setting for a single discovery session in the **Discovery Session Editor** modal page.

- **Hostname Precedence.** Specifies which name the platform will use for each discovered device. Choices are:
 - *0. SNMP System Name.* Use the device name specified in the device's MIB.
 - *1. DNS Reverse Lookup.* Use the device name specified in the device's reverse-lookup record.
- **DNS Hostnames.** Specifies which DNS name the platform will use for each discovered device. Choices are:
 - *0. Strip Device Name.* The platform will use the DNS hostname as the device name for each device.
 - *1. Use Full Domain Name (FQD).* The platform will use the fully-qualified domain name as the device name for each device.

Viewing Dynamic Application Data and Alerts

Overview

You can view the data collected from a device by Dynamic Applications in the following places:


- The **Device Performance** page allows you to view detailed reports for the selected device, including the graphs generated for each presentation object in each performance Dynamic Application that is aligned to the device.
- The **Configuration Report** page displays data collected by all configuration Dynamic Applications aligned to a device.
- The **Journal Report** page displays data collected by all journal Dynamic Applications aligned to a device.

This chapter will describe how to view each type of Dynamic Application and how to view alerts generated by Dynamic Applications.

Performance Dynamic Applications

The **Device Performance** page displays performance reports for a device. To view the **Device Performance** page:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Find the device for which you want to view Dynamic Application data. Select its bar graph icon (📊). In the **Device Administration** panel, select the **[Performance]** tab. The **Device Performance** page is displayed:



The **Device Performance** page displays all performance graphs that are available for the device. These graphs include:

- Graphs defined by the Dynamic Applications associated with the device.
- Graphs generated using data collected by other processes, such as availability data, file system usage, interface usage, and data from monitoring policies.

Viewing Performance Dynamic Application Data

To view a performance graph defined in a Dynamic Application:

1. The bottom of the left NavBar of the **Device Performance** page lists all performance Dynamic Applications that include performance graphs. Select the Dynamic Application for which you want to view data.
2. Selecting a performance Dynamic Application will expand the NavBar item and display a list of all presentation objects in that Dynamic Application. Select a presentation object to view the graph for that presentation object.

For each graph associated with a performance Dynamic Application:

- The x-axis of a Performance Dynamic Application graph displays time.
- The y-axis displays the collected value(s).
- Mousing over the graph will display the exact value for each data series at that point on the graph.
- The exact values are displayed in the "Mouse-over" column in the table below the graph.

Viewing Details with the Data Table

At the bottom of each graph is a table that allows you to view details about each data series and overview information about the entire graph.

The data table includes the following:

- **Data Type/Label.** For graphs that include multiple data series on a single graph (when the collection objects used to generate the graph have collected a list of values instead of a single value), each data series has its own row in this table. This column displays the name of the data series and how it is color-coded in the report. Clicking on the check-mark toggles the display of the data series in the graph.
- **Trend.** Allows you to toggle the trendline on and off. The trendline shows average values on each side of the current point. This checkbox is checked (trendline is displayed) by default.
- **Mouse-over.** When you mouse-over the graph, this column displays the exact value for each data series at that time point on the graph.
- **Min.** The column displays the minimum value for the data series in the graph.
- **Max.** This column displays the maximum value for the data series in the graph.
- **Avg.** This column displays the average value for the data series in the graph.
- **Missed Polls.** This column displays the number of times the ScienceLogic platform was unable to collect the data within the time span of the graph.

Configuring the Data for the Report

You can use the **[Options]** menu to determine how you want to configure the data that is displayed in the graph. Your choices are:

- **Default.** The initial graph that is displayed is not normalized and displays every collected value.
- **Normalized.** In the ScienceLogic platform, normalized data does not include polling sessions that were missed or skipped. So for normalized data, null values are not included when calculating maximum values, minimum values, or average values. When you select this option, the ScienceLogic platform normalizes all the data collected in each 24 hour period and displays a single value for each day.
- **Percentile.** Displays percent on the y-axis. This can be applied to normalized or non-normalized graphs.
- **Kiosk.** Displays the report in full-page mode. This is helpful for NOC personnel who need to display graphs on large screens.
- **Detach.** Spawns the graph in a new window.
- **Series Selection.** Each performance graph is limited to displaying eight data series at a time. When you select this option, the **Graph Index Selection** page is displayed, where you can select up to eight data series to display.
- **Edit Current Presentation.** Selecting this option allows you to edit the Presentation Object that was used to generate the graph.

Generating a Report from a Performance Graph

You can use the **[Report]** button to save the report to a file on your local computer. You can select the format to save the file in. Your choices are:

- **HTML with Images.** Saves the graph and a table of all the data in the report, in HTML format.
- **HTML Text Only.** Saves the report as a table of data, in HTML format.
- **HTML Text Only all indexes.** Saves the report as a table of data, in HTML format. In the **Device Performance** page, the report can include up to eight data series (indexes); when you select this option, the HTML report will include all indexes collected by the ScienceLogic platform, even if the number of indexes is greater than eight.
- **CSV.** Saves the data from the report (usually date, time, and value) as comma-separated values.
- **CSV all indexes.** Saves the data from the report (usually date, time, and value) as comma-separated values. In the **Device Performance** page, the report can include up to eight data series (indexes); when you select this option, the CSV report will include all indexes collected by the ScienceLogic platform, even if the number of indexes is greater than eight.
- **Graph Image Only.** Saves only the graph from the report as a .png file.
- **ODS w/Chart Img.** Saves the graph and a table of all the data in the report, in ODS format.
- **ODS Plain.** Saves the table of all the data in the report, in ODS format.
- **XLSX w/ Chart Img.** Saves the graph and a table of all the data in the report, in XLSX (Excel) format.
- **XLSX Plain.** Saves the table of all the data in the report, in XLSX (Excel) format.

- **PDF w/chart Img**. Saves the graph and a table of all the data in the report, in PDF format.
- **PDF Plain**. Saves the table of all the data in the report, in PDF format.

Defining the Date Range for a Report

The fields in the **Date Range Selection** pane allow you to define the time span that will be displayed in the report. When you define a Date Range, only data from the specified time span will be included in the report.

You can use the following fields, menus, and buttons to define the dates and associated data that will be displayed in the report:

- **Start**. Allows you to define the Start date for the report. Data from this date until the specified End date will be included in the report.
- **End**. Allows you to define the end date for the report. Data from the specified Start date until this End date will be included in the report.
- **[Presets]**. Allows you to select some pre-defined time spans for the report.
- **[Set Custom]**. Applies the values from the **Start** and **End** fields to the report.
- **[Range buttons]**. Allows you to zoom in or out for a specific time span. Choices are dependent upon the values selected in the **Date Range Selection** pane and can range from 6 hours to 90 days.
- You can drag the Date Range slider to the left or right, and the Performance graph will display the information for the selected time range.

Configuration Dynamic Applications

The **Configuration Report** page displays data collected from the device by configuration Dynamic Applications. Usually, configuration data contains static information about hardware and configuration settings, such as serial numbers, version numbers, and hardware status.

NOTE: If you select the **Hide Object** checkbox for an object in the **Collection Objects** page (System > Manage > Applications > Create/Edit), the object will not be included in the **Configuration Report** page.

You can define the layout of the **Configuration Report** page in the **Collection Objects** page for the Dynamic Application. In the **Collection Objects** page, you can use the **Group** field and the **Table Alignment** fields to define which objects will be grouped together, and which table each object will appear in.


You can enable change detection for an object in the in the **Collection Objects** page for the Dynamic Application, in the **Change Alerting** field. If an object's value has changed, it will be highlighted in red in the **Configuration Report** page. You can then click on the object's value in the **Configuration Report** page and view a list of historical values for the object.


For more information about configuring the table layout and change detection for a configuration Dynamic Application, see the [Collection Objects](#) chapter.

For objects of type "enum," you can mouseover the object and view all the possible values for the object.

NOTE: The **Configuration Report** page does not display Dynamic Applications that have *Cache Results* selected in the **Caching** field in the **Dynamic Applications Properties Editor** page. Dynamic Applications that cache results are designed to collect data only for other Dynamic Applications and cannot be used to display data.

To view Configuration Dynamic Application information:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Find the device for which you want to view configuration Dynamic Application data. Select its bar graph icon (). In the **Device Administration** panel, select the **[Configs]** tab.
3. The **Device Configuration** page is displayed:



The screenshot shows the 'Configs' tab for a device. The left sidebar lists Dynamic Applications such as 'VMware: Hardware Monitoring Configuration' and 'VMware: HostSystem Configuration'. The main content area displays a 'Configuration Report | VMware: Hardware Monitoring Configuration' with a 'Snap-Shot Date' of 2014-06-16 13:05:00. The report includes several sections:

- Rollup System Health:** A table with one entry: 'VMware Rollup Health State' with a 'Green' health state.
- Processor Health:** A table with 14 entries, all with a 'Green' health state. A note above the table says 'Change detection is disabled. No history available.'.
- Fan Health:** A table with columns for Description, Health State, and RPMs.
- Power Health:** A table with 5 entries, all with a 'Green' health state. Columns include Description, Health State, Units, and Value.

Selecting Data to View

If one or more Dynamic Applications of type "configuration" are associated with the device, the **Configuration Report** page will display that list of Dynamic Applications in the left NavBar.

NOTE: The left navigation bar does not display Dynamic Applications that have *Cache Results* selected in the **Caching** field in the **Dynamic Applications Properties Editor** page. Dynamic Applications that cache results are designed to collect data only for other Dynamic Applications and cannot be used to display data.

Viewing Data

When you select a Dynamic Application in the left NavBar, the right pane displays data collected from the device by the Dynamic Application.

- Some objects may appear in a list at the top of the right pane. These are objects that are not grouped into a table. For each of these values, no values were specified in the **Group** field and the **Table Alignment** field, in the **Collection Objects** page. These are usually objects for which there is only one, non-changing value (like model number, for example).
- Some objects may appear in tables. Tables work best for objects with multiple values, like RAM location. Each row represents one value from each collection object in the group, which all have the same index.
 - Each column heading is the name of an object. Mousing over the column heading displays a description of the object. To edit the description, click on the column heading. The **Collection Objects** page appears, populated with values from the appropriate object. You can edit the value in the *Description* field, and that value will appear when you mouseover the column heading in the **Configuration Report** page.
- Mousing over a value can display the following:
 - If the object is of type "enum", the mouseover text displays the list of all possible values for the object. For example, "0 unknown, 1 disabled, 2 enabled".
 - If change detection has not been enabled, displays the text "Change detection is disabled. No history available".
 - If change detection has been enabled, displays "Click to view change history". If you click, the ScienceLogic platform displays the **Change History** modal page, where you can view all the values collected from the device for the selected object.

Generating a Report of the Data

You can generate a report about the data in the **Configuration Report** page. To do so:

1. In the **Configuration Report** page, in the Navigation Bar (left pane), select the Dynamic Application you want to generate a report from.
2. In the **Configuration Report** page, select the **[Actions]** menu. Select *Print a Report*.
3. The ScienceLogic platform generates an HTML report that contains all the data from the **Configuration Report** page. You can view, print, or save the report.


Viewing Historical Data

By default, the **Configuration Report** page displays data from the latest polling session. However, you can use the **Snap-Shot Selector** page to display data from a previous polling session in the **Configuration Report** page.

The **Snap-Shot Selector** page displays a list of polling sessions where a change was discovered in the configuration data. If none of the data in a Dynamic Application changes from one polling session to the next, then the ScienceLogic platform does not include an entry in the **Snap-Shot Selector** page.

To display data from a previous polling session in the **Configuration Report** page:

1. In the **Configuration Report** page, in the Navigation Bar (left pane), select the Dynamic Application for which you want to view historical data.
2. When the data is displayed in the right pane, select the **[Snap-Shots]** button.
3. The **Snap-Shot Selector** modal page appears. This page displays a calendar interface, in which you can select a date for which you want to view a list of Snap-Shots.



4. To select a date for a Snap-Shot, scroll through the calendar until you find the month that you are interested in. Click on the date you are interested in.
5. The pane to the right will display a list of all available Snap-Shots for the selected date. Each Snap-Shot is labeled with a date and time stamp and specifies how many objects had changed values. To select a Snap-Shot, click on it and select the **[View Snapshot]** button.

NOTE: If the pane to the right does not display one or more available Snap-Shots, this means that the ScienceLogic platform did not detect any changes to the objects on the selected date.

6. The data from the selected Snap-Shot is loaded and displayed in the **Configuration Report** page.

Editing the Application

From the **Configuration Report** page, you can edit the properties of a Dynamic Application. When you do so, you change the behavior of the Dynamic Application for all subscriber devices, not just the current device.

To edit a Dynamic Application from the **Configuration Report** page:


1. In the **Configuration Report** page, in the Navigation Bar (left pane), select the Dynamic Application you want to view and edit.
2. When the data from the Dynamic Application is displayed in the right pane, select the **[Actions]** menu and choose *Edit This Application*.
3. The **Collection Objects** page appears. In this page, you can edit how the ScienceLogic platform retrieves values for an object and how those values are displayed in the **Configuration Report** page. You can also access all the other tabs in the Dynamic Applications panel for the Dynamic Application.

For more information about editing collection objects, see the [Collection Objects](#) chapter.

Journal Dynamic Applications

The **Journal View** page displays journal entry information collected from the device by Dynamic Applications. All information from Dynamic Applications of type journal is included in the **Journal View** page. Journal Dynamic Applications store information in log format; for example, telephone call records or access logs.

To view journal Dynamic Application information:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Find the device for which you want to view journal Dynamic Application data. Select its bar graph icon (). In the **Device Reports** panel, select the **[Journals]** tab.

3. The **Journal View** page is displayed:

Date	Health	State	Collected On
1. 2014-06-14 15:00:00	informational message: [RESOLVED] Elevated rates of connection timeouts	Open	2014-06-15 05:05:09

Selecting Data to View

If one or more Dynamic Applications of type "journal" are associated with the device, the **Journal View** page will display that list of Dynamic Applications in the left NavBar.

When you select a Dynamic Application in the left NavBar, the right pane displays data collected from the device by the Dynamic Application.

Viewing Data

The **Journal View** page arranges collected journal entries in tabular format.

- The table contains a row for each journal entry.
- The table contains a column for each presentation object, plus the **State** and **Collected On** columns. Presentation objects define the text to display in each row in the column, including which collection values will be displayed. Presentation objects are defined in the **Presentation Objects** page for the Dynamic Application.

The **Journal View** page displays the following about each journal entry:

TIP: To sort by descending order, click the column heading again. To sort a column that contains presentation objects, sorting must be enabled in the **Presentation Objects** page (System > Manage > Applications > Create/Edit). Date and time column sorts by descending order on the first click; to sort by ascending order, click the column heading again.

- **Presentation Objects.** One or more columns in the table of journal entries will be presentation objects defined in the Dynamic Application. The values in this column can be based on one or more collection objects, and can be a text string, a number, or a time and date value.
- **State.** Specifies the current state of the journal entry. Journal entries can have one of the following states:
 - Open
 - Closed
 - Abandoned
 - Error
 - Reopened
- **Collected On.** Specifies the last time the journal entry was updated.

Searching & Filtering the List of Data

You can filter the list on the **Journal View** page by one or more parameters. Only journal entries that meet all the filter criteria will be displayed in the **Journal View** page.

To filter by parameter, enter text into the desired filter-while-you-type field. The **Journal View** page searches for journal entries that match the text, including partial matches. By default, the cursor is placed in the left-most filter-while-you-type field. You can use the <Tab> key or your mouse to move your cursor through the fields. The list is dynamically updated as you type. Text matches are not case-sensitive.

You can also use *special characters* to filter each parameter.

Filter the list by one or more of the following parameters:

- **Presentation Objects.** Each presentation object column has a filter. For columns that contain a text string or a numeric value, you can enter text to match, including special characters, and the **Journal View** page will display only journal entries that have a matching value for that presentation object. For each journal entry, the value that is matched for a presentation object is the value of the first collection object that appears in the presentation object text. For columns that contain a time and date, you can select a time span, and the **Journal View** page will display only journal entries that have a time and date value within the selected time span. Choices are:
 - *All.* Display all journal entries that match the other filters.
 - *Last Minute.* Display only journal entries that have been created within the last minute.
 - *Last Hour.* Display only journal entries that have been created within the last hour.

- *Last Day*. Display only journal entries that have been created within the last day.
 - *Last Week*. Display only journal entries that have been created within the last week.
 - *Last Month*. Display only journal entries that have been created within the last month.
 - *Last Year*. Display only journal entries that have been created within the last year.
- **State**. You can enter text to match, including special characters, and the **Journal View** page will display only journal entries that have a matching state. Journal entries can have one of the following states:
 - Open
 - Closed
 - Abandoned
 - Error
 - Reopened
- **Collected On**. You can select a time span, and the **Journal View** page will display only journal entries that have been updated within that time period. Choices are:
 - *All*. Display all journal entries that match the other filters.
 - *Last Minute*. Display only journal entries that have been created within the last minute.
 - *Last Hour*. Display only journal entries that have been created within the last hour.
 - *Last Day*. Display only journal entries that have been created within the last day.
 - *Last Week*. Display only journal entries that have been created within the last week.
 - *Last Month*. Display only journal entries that have been created within the last month.
 - *Last Year*. Display only journal entries that have been created within the last year.

Special Characters

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

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

String and Numeric

- , (comma). Specifies an "OR" operation. Works for string and numeric values. For example:
 "dell, micro" matches all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "AND " operation. Works for string and numeric values. For example:
 "dell & micro" matches all values that contain both the string "dell" AND the string "micro", in any order.

- ! (exclamation point). Specifies a "not" operation. Works for string and numeric values. For example:

"!dell" matches all values that do not contain the string "dell".

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

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

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

"! ^ " would match null values.

"!\$" would match null values.

"!*" would match null values.

"happy, !dell" would match values that contain "happy" OR values that do not contain "dell".

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

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

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

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

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

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

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

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

String

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

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

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

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

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

"! ^ " would match null values.

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

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

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

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

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

"!\$" would match null values.

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

Numeric

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Additional Examples


- "aio\$". Matches only text that ends with "aio".
- "^shu". Matches only text that begins with "shu".
- "^silo\$". Matches only the text "silo", with no characters before or after.
- "!silo". Matches only text that does not contain the characters "silo".
- "!^silo". Matches only text that does not start with "silo".
- "!O\$". Matches only text that does not end with "O".
- "!^silo\$". Matches only text that is not the exact text "silo", with no characters before or after.
- "!. Matches null values, typically represented as "--" in most pages.
- "!. Matches null values, typically represented as "--" in most pages.
- "!. Matches all text that is not null.
- silo, !aggr". Matches text that contains the characters "silo" and also text that does not contain "aggr".
- "silo, 02, !aggr". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr".
- "silo, 02, !aggr, !01". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr" and also text that does not contain "01".
- "^s*i!*o\$". Matches text that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXlXo" would match.
- "!^s*i!*o\$". Matches all text that does not contain the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXlXo" would not match.
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" would match, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" would match, because it does not contain "vol" and it does contain "02".
- "aggr,!vol&02". Matches text that contains "aggr" OR text that does not contain "vol" AND also contains "02".
- "aggr,!vol&!infra". Matches text that contains "aggr" OR text that does not contain "vol" AND does not contain "infra".
- "*". Matches all text.
- "!*". Matches null values, typically represented as "--" in most pages.
- "silo". Matches text that contains "silo".
- "!silo". Matches text that does not contain "silo".
- "!^silo\$". Matches all text except the text "silo", with no characters before or after.
- "-3,7-8,11,24,50-". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50.

- "-3,7-8,11,24,50-,a". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50, and text that includes "a".
- "?n". Matches text that contains any single character and the character "n". For example, this string would match "an", "bn", "cn", "1 n", and "2n".
- "n*SAN". Matches text the contains "n", zero or any number of any characters and then "SAN". For example, the string would match "nSAN", and "nhamburgerSAN".
- "^?n*SAN\$". Matches text that begins with any single character, is following by "n", and then zero or any number of any characters, and ends in "SAN".

Generating a Report of the Data

You can generate a report about the data in the **Journal View** page.


To generate a report about the data in the **Journal View** page:

1. Go to Registry > Devices > Device Manager. Find the device for which you want to generate a report. Select its bar graph icon (). Select the **[Journals]** tab.
2. In the **Journal View** page, in the left NavBar, select the Dynamic Application you want to generate a report from.
3. You can filter the journal entries to include in the report. Using the search filters at the top of the table of journal entries, filter the list of journal entries so that only the journal entries you want to include on the report are displayed.
4. In the **Journal View** page, select the **[Actions]** menu. Select **Generate Report**.
5. The **Export current view as a report** page is displayed. Select the output format for the report, optionally select if the ScienceLogic platform must force the browser to save the file to disk, and then select the **[Generate]** button.

Editing the Application

From the **Journal View** page, you can edit the properties of a Dynamic Application. When you do so, you change the behavior of the Dynamic Application for all subscriber devices, not just the current device.

To edit a Dynamic Application from the **Journal View** page:

1. Go to Registry > Devices > Device Manager. Find the device for which you want to view data. Select its bar graph icon (). Select the **[Journals]** tab.
2. In the **Journal View** page, in the left NavBar, select the Dynamic Application you want to view and edit.
3. When the data from the Dynamic Application is displayed in the right pane, select the **[Actions]** menu and choose **Edit This Application**.
4. The **Collection Objects** page appears. In this page, you can edit how the ScienceLogic platform retrieves values for an object. You can also access all the other tabs in the Dynamic Applications panel for the Dynamic Application.

Viewing Alerts Generated by Dynamic Applications

You can view alerts generated by a Dynamic Application in the **Device Logs & Messages** page. To view alerts on a device generated by a Dynamic Application:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to view alerts generated by a Dynamic Application. Select its bar graph icon (📊).
3. In the **Device Reports Panel**, select the **[Logs]** tab. The **Device Logs & Messages** page displays the following for each log message:

The screenshot shows the 'Device Logs & Messages' page for device 'em7_ap_203'. It includes a navigation menu at the top with tabs like Close, Summary, Performance, Topology, Configs, Interfaces, and Organization. Below this is a detailed view of the device's configuration, including IP address (10.0.9.203), class (ScienceLogic, Inc.), and system information. The main area displays a list of 315 messages found, with a search bar and filters. The log messages are listed in a table with columns for Date Time, Source, Event ID, and Severity. The messages include various alerts such as 'Device Now Available', 'Network Latency below threshold', and 'Device Failed Availability Check: UDP - SNMP'.

Date Time	Source	Event ID	Severity
2014-10-06 20:17:21	Internal	111284	--
2014-10-06 20:17:21	Internal	111285	--
2014-10-06 20:12:16	Internal	111273	--
2014-10-06 20:12:16	Internal	111274	--
2014-10-06 20:07:16	Internal	111273	--
2014-10-06 20:07:16	Internal	111274	--
2014-10-06 20:05:18	Internal	111125	--
2014-10-06 20:02:20	Internal	111273	--
2014-10-06 20:02:20	Internal	111274	--
2014-10-06 20:01:13	Internal	111125	--
2014-10-06 20:00:58	Internal	--	--
2014-10-06 20:00:36	Internal	--	--
2014-10-06 20:00:33	Internal	--	--
2014-10-06 20:00:21	Internal	--	--
2014-10-06 19:57:14	Internal	111264	--
2014-10-06 19:57:14	Internal	111265	--
2014-10-06 19:52:15	Internal	111178	--
2014-10-06 19:52:15	Internal	111179	--
2014-10-06 19:47:16	Internal	111178	--
2014-10-06 19:47:16	Internal	111179	--
2014-10-06 19:42:18	Internal	111178	--
2014-10-06 19:42:18	Internal	111179	--
2014-10-06 19:37:17	Internal	111178	--
2014-10-06 19:37:17	Internal	111179	--
2014-10-06 19:32:15	Internal	111178	--
2014-10-06 19:32:15	Internal	111179	--
2014-10-06 19:27:16	Internal	111178	--
2014-10-06 19:27:16	Internal	111179	--
2014-10-06 19:22:21	Internal	111178	--

- **Date Time.** The date and time the entry was made in the log.
- **Source.** The entity or process that generated the message. If generated by a Dynamic Application, the source will be *Dynamic*.
- **Event ID.** If an event was created, a unique event ID, generated by the ScienceLogic platform. If the log entry is not associated with an event, no ID appears in this column.

- **Priority.** This column is not applicable to alerts generated by Dynamic Applications.
- **Message.** Text of the log entry, color-coded to match event severity (if applicable).

Searching the List of Log Entries for a Dynamic Application

The search fields at the top of the **Device Logs & Messages** pane allows you to search log entries for messages generated by a Dynamic Application associated with a device. To search the list of log entries for a Dynamic Application:

1. In the **Device Logs & Messages** pane, select the **[Search All Messages]** drop-down list.
2. Select *[Search Dynamic]*.
3. In the **drop-down list** search bar, you can optionally define the parameters for the search. Choices are:
 - *Where Message is like.* Search the message portion of the log entry.
 - *Where Date received is like.* Search the date portion of the log entry.
 - *Where Priority is like.* This selection is not applicable to alerts generated by Dynamic Applications.
4. In the regular expression field, you manually enter the text to search for. You can use the following special characters in this field:
 - * Match zero or more characters preceding the asterisk. For example:
"dell*" would match "dell", "dell2650", "dell7250" and "dell1700N".
"*dell*" would match "mydell", "dell", "dell2650", "dell7250" and "dell1700N".
 - % Match zero or more characters preceding the percent symbol. This special character behaves in the same way as the asterisk.

- When you select the **[Search]** button, the **Device Logs & Messages** page will be refreshed and will display only log entries that match the search parameters.

The screenshot displays the 'Device Logs & Messages' page for a device named 'em7_ag_203'. The search bar is set to 'where Message is like' and contains the text 'Search All Messages'. The search results show a list of 29 log entries, each with a date and time, a search term, and a message description. The messages include various system alerts such as 'Device Now Available', 'Network Latency below threshold', and 'Device Failed Availability Check: UDP - SNMP'. The page also features a navigation menu at the top with tabs like Close, Summary, Performance, Topology, Configs, Interfaces, and Organization.



Viewing Active Events for a Dynamic Application

Events are messages that are triggered when a specific condition is met. For example, an event can signal that a server has gone down, that a device's hard drives are getting too full, or simply display the status of a device. You can view events generated by Dynamic Applications in the **Viewing Events** page. You can view active events, as well as events that have been cleared.

To view active events generated by a Dynamic Application:

- Go to the **Device Manager** page (Registry > Devices > Device Manager).
- In the **Device Manager** page, find the device for which you want to view events generated by a Dynamic Application. Select its bar graph icon (📊).

- In the **Device Reports Panel**, select the **[Events]** tab. The **Viewing Events** page displays the following for each event message:

Event Message Severity	Acknowledged	Age / Elapse	Ticket	Last Detected	EID	Source	Count	Del
System or agent has recently restarted	<input checked="" type="checkbox"/>	6 mins 30 secs	--	2014-10-06 20:25:07	111288	Internal	2	<input type="checkbox"/>
Device Now Available	<input checked="" type="checkbox"/>	11 mins 17 secs	--	2014-10-06 20:17:21	111284	Internal	1	<input type="checkbox"/>
Network Latency below threshold	<input checked="" type="checkbox"/>	11 mins 17 secs	--	2014-10-06 20:17:21	111285	Internal	1	<input type="checkbox"/>
Database connection problem: Dynamic App: 894. Could not connect to database using credential E	<input checked="" type="checkbox"/>	2 hrs 13 mins	--	2014-10-06 20:27:14	111125	Internal	24	<input type="checkbox"/>

Legend: 2 Healthy, 1 Minor, 1 Major

- **Event Message | Severity.** Message generated by event, as defined in the **Event Policy Editor** page. The message is color-coded for severity.
 - **Critical.** Critical events are those that require immediate attention.
 - **Major.** Major events are those that require immediate investigation.
 - **Minor.** Minor events are those that need to be investigated before problems become severe.
 - **Notice.** Notice events are those that require attention but are not problem-related.
 - **Healthy.** Healthy events are those that are not urgent.
- **Acknowledged.** Specifies whether anyone has acknowledged this event.
 - *Red check.* Event has not been acknowledged.
 - *Gray check with name.* Event has been acknowledged.
- **Age / Elapse.** Number of days, hours, and minutes since the last occurrence of the event.
- **Ticket.** Ticket ID associated with this event, if applicable.
- **Last Detected.** Date and time of last occurrence of the event.

- **EID**. Unique ID for the event, generated by the ScienceLogic platform. By default, this column appears only if specified in your Account Settings or if you select the **[Advanced]** button.
- **Source**. Source of the log message that triggers the event, as defined in the **Event Policy Editor** page. If the event was generated by a Dynamic Application, the **Source** will be *Dynamic*.
- **Count**. Number of times this event has occurred.
- **External Ticket**. The numeric ID associated with a ticket from an external ticketing system (that is, a ticket that was not created in the ScienceLogic platform). If this field displays a value, you can click on that value to spawn a new window and view the external ticket.

Viewing Details About an Active Event

You can view details about an active event in the **Event Information** page. The **Event Information** page displays more detailed data about a single active event. To view details about an active event:

1. In the **Viewing Active Events** pane, find the event for which you want to view detailed information.
2. Select its information icon (i).
3. The **Event Information** page displays the following for each active event:




- **Event ID**. Unique ID for the event, generated by the ScienceLogic platform.
- **Event Message**. Message generated by event, as defined in the **Event Policy Editor** page.

- **Severity.** Severity of the event. Choices are:
 - *Critical*
 - *Major*
 - *Minor*
 - *Notice*
 - *Healthy*
- **For Element.** Name of the element associated with the event.
- **First Occurrence.** Number of days and hours since the first occurrence of the event. Date and time of first occurrence of the event.
- **Last Occurrence.** Number of days and hours since the last occurrence of the event. Date and time of last occurrence of the event.
- **Occurrence Count.** Number of times this event has occurred on this entity.
- **Acknowledged On.** Date and time the event was acknowledged.
- **Acknowledged By.** Username of user who acknowledged the event.
- **Policy Name/ID.** Name of the event policy, as defined in the **Event Policy Editor** page and policy ID.
- **Policy Type.** Source of the log message that triggers the event, as defined in the **Event Policy Editor** page. If the event was generated by a Dynamic Application, the **Policy Type** will be *Dynamic*.
- **Ticket Description.** Description field from the associated ticket, if applicable.
- **Probable Cause & Resolution Text.** This pane displays additional information about the event, as defined in the **Event Policy Editor** page.


Viewing Cleared Events from a Dynamic Application

You can view cleared events from a Dynamic Application from the **Viewing Cleared Events** page. A cleared event is an event that is no longer occurring, and has been manually cleared by a user, or has expired.

To view cleared events from a Dynamic Application:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. In the **Device Manager** page, find the device for which you want to view events generated by a Dynamic Application. Select its bar graph icon ().

- In the **Device Reports Panel**, select the **[Events]** tab. In the **Viewing Events** page, select the **[Cleared]** button. The **Viewing Cleared Events** page displays the following information:

Close	Summary	Performance	Topology	Configs	Log Files	Interfaces				
Logs	Events	Tickets	Configuration	Configuration	Configuration	TCP Ports	Organization			
Device Name	em7_ap_203	Managed Type	Physical Device							
IP Address / ID	10.0.9.203 246	Category	System:EM7							
Class	ScienceLogic, Inc.	Sub-Class	EM7 Admin Portal							
Organization	System	Uptime	0 days, 00:00:42							
Collection Mode	Active	Collection Time	2014-10-06 20:25:00							
Description	ScienceLogic EM7 G3 - Administration Portal	Group / Collector	CUG MOSS_Patch_AIO							
Device Hostname										
Viewing Cleared Events							Active	Stats	Reset	Guide
Event Message Severity		Acknowledged	Ticket	Cleared By	Cleared Date	Time Since Cleared	EID	Source	Count	
Device Failed Availability Check: UDP - SNMP		--	--	auto-clear	2014-10-06 20:17:29	12 mins 26 secs	111273	Internal	3	
Device Failed Availability Check: UDP - SNMP		--	--	auto-clear	2014-10-06 19:57:19	32 mins 36 secs	111178	Internal	15	
Device Failed Availability Check: UDP - SNMP		--	--	auto-clear	2014-10-03 00:52:31	3 days 19 hrs	104687	Internal	6	
Network latency exceeded threshold: No Response		--	--	auto-clear	2014-10-06 20:17:29	12 mins 26 secs	111274	Internal	3	
Network latency exceeded threshold: No Response		--	--	auto-clear	2014-10-06 19:57:19	32 mins 36 secs	111179	Internal	15	
Network latency exceeded threshold: No Response		--	--	auto-clear	2014-10-03 00:52:31	3 days 19 hrs	104688	Internal	6	
System or agent has recently restarted		--	--	expiration	2014-09-27 01:26:14	1 wk 2 days	95309	Internal	3	
Added dynamic application for device: Net-SNMP: CPU		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79078	Internal	1	
Added dynamic application for device: Net-SNMP: Physical Memory		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79079	Internal	1	
Added dynamic application for device: Net-SNMP: Swap		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79080	Internal	1	
Added dynamic application for device: EM7: Asset Information		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79081	Internal	1	
Added dynamic application for device: Host Resource: CPU Config		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79082	Internal	1	
Added dynamic application for device: Host Resource: Memory Config		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79083	Internal	1	
Added dynamic application for device: Support: File System		--	--	expiration	2014-09-16 20:19:09	2 wks 6 days	79084	Internal	1	
Device Now Available		--	--	auto-clear	2014-10-06 20:02:29	27 mins 26 secs	111264	Internal	1	
Network Latency below threshold		--	--	auto-clear	2014-10-06 20:02:29	27 mins 26 secs	111265	Internal	1	
Device Now Available		--	--	expiration	2014-10-03 01:08:22	3 days 19 hrs	105319	Internal	1	
Network Latency below threshold		--	--	expiration	2014-10-03 01:08:21	3 days 19 hrs	105320	Internal	1	



- **Event Message | Severity.** Message generated by event, as defined in the **Event Policy Editor** page. The message is color-coded for severity.
 - **Critical.** Critical events are those that require immediate attention.
 - **Major.** Major events are those that require immediate investigation.
 - **Minor.** Minor events are those that need to be investigated before problems become severe.
 - **Notice.** Notice events are those that require attention but are not problem-related.
 - **Healthy.** Healthy events are those that are not urgent.
- **Acknowledged.** Specifies whether anyone has acknowledged this event.
 - *Gray Dash.* Event has not been acknowledged.
 - *Gray check with name.* Event has been acknowledged.
- **Ticket.** Ticket ID associated with this event, if applicable.
- **Cleared By.** Username of the person who cleared the ticket.
- **Cleared Date.** Date and time on which event was cleared.

- **Time Since Cleared.** Number of days, hours, and minutes since the event was cleared.
- **EID.** Unique ID for the event, generated by the ScienceLogic platform. By default, this column appears only if specified in your Account Settings or if you select the **[Advanced]** button.
- **Source.** Source of the log message that triggers the event, as defined in the **Event Policy Editor** page. If the event was generated by a Dynamic Application, the **Source** will be *Dynamic*.
- **Count.** Number of times this event has occurred.

Dynamic Application Settings

Overview

Every Dynamic Application has a set of properties specific to that Dynamic Application. This chapter will describe how to view and edit the properties for a Dynamic Application. There are also some global settings that affect the operation of Dynamic Applications; these global settings are described in the [Managing Dynamic Applications](#) chapter.

Dynamic Application Properties

Every Dynamic Application has a set of properties that define the general settings for that Dynamic Application. To view and edit the settings for a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).

2. Find the Dynamic Application you want to view or edit the properties for. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed:

The screenshot shows the 'Dynamic Applications [1011] | Properties Editor' window. It features a top navigation bar with 'Guide' and 'Reset' buttons. The main area is divided into several sections:

- Configuration Fields:**
 - Application Name:** Microsoft: Azure Account Discovery
 - Application Type:** [Snippet Configuration]
 - Execution Environment:** [-- Default: Microsoft: Azure]
 - Caching:** [No caching]
 - Device Dashboard:** None
 - Version Number:** [Version 1.2]
 - Operational State:** [Enabled]
 - Collector Affinity:** [Default]
 - Poll Frequency:** [Every 15 Minutes]
 - Abandon Collection:** [Default]
 - Context:** (empty field)
 - Null Row Option:** [Hide row]
 - Null Column Option:** [-- values]
- Advanced Settings:**
 - Disable Rollup of Data:**
 - Component Mapping:**
 - Buttons:** Save, Save As
- Description:** This dynamic application discovers and models Azure accounts.
- Release Notes & Change Log:** A rich text editor showing:
 - Version 1.2:** 1. Added support for Azure Government subscriptions.
 - Version 1.1:** 1. Added availability collection object. 2. Added multiple subscription support. 3. Updated ID collection object to be visible. 4. Removed self monitoring process.
 - Version 1.0:** 1. GA release of Microsoft: Azure Account Discovery dynamic application.

3. A Dynamic Application has the following properties:

- **Application Name.** The name of the Dynamic Application.
- **Application Type.** Specifies the Dynamic Application type, which comprises a *protocol* and an *archetype*. For a description of protocols and archetypes, see the **Types of Dynamic Application** section in the [Overview](#) chapter. The following options are available in the **Application Type** field:
 - Bulk Snippet Configuration
 - Bulk Snippet Performance
 - Database Configuration
 - Database Performance
 - IT Service
 - PowerShell Configuration
 - PowerShell Performance
 - SNMP Configuration
 - SNMP Performance

- SOAP Configuration
- SOAP Performance
- Snippet Configuration
- Snippet Journal
- Snippet Performance
- WMI Configuration
- WMI Performance
- XML Configuration
- XML Performance
- XSLT Configuration
- XSLT Performance

NOTE: The **Application Type** field cannot be changed after the Dynamic Application is created. When editing an existing Dynamic Application, the **Application Type** field is view-only.

- **Execution Environment.** The execution environment to which the Dynamic Application is aligned. This field appears only for snippet and internal collection Dynamic Applications. An execution environment contains the supporting modules, code, scripts, directories, and files (packaged in ScienceLogic Libraries) for the snippet. An execution environment includes its own installation directories, doesn't share libraries with other environments, and allows granular control of dependencies and versions. The default execution environment is *System*. For more information, see the **ScienceLogic Libraries and Execution Environments** manual.
- **Caching.** This field is disabled for Dynamic Applications that use Bulk Snippet, Database, or Internal Collection (ICDA) type. When the ScienceLogic platform requests information from a device during Dynamic Application collection, the platform can optionally **cache** the response from the device. If a response is cached, other Dynamic Applications that use the same protocol can retrieve collection object values from the cached response. When you cache responses, you reduce the number of requests performed by the platform and speed up collection. Typically, Dynamic Applications aligned to component devices retrieve values from a cached repository. The cached repository is created with a Dynamic Application aligned with the root device. Choices are:
 - *No caching.* Specifies that this Dynamic Application does not use the caching feature. The requests that are used to populate the collection objects in this Dynamic Application are defined within this application and the responses to those requests are not available to other Dynamic Applications.
 - *Cache Results.* Specifies that the ScienceLogic platform should cache the responses from the requests defined in this Dynamic Application. Other Dynamic Applications that use the same protocol or Snippet Dynamic Applications can retrieve data from the cached responses associated with this Dynamic Application. When this option is selected, the collection objects defined in this Dynamic Application are no longer collected at a regular frequency. If you select this option, you must not define collection objects that are used to generate configuration tables or performance graphs.

- *Consume Cached Results*. Specifies that the collection objects defined in this Dynamic Application are populated with values from other Dynamic Applications that use the same protocol and have *Cache Results* enabled.
 - In a SOAP or WMI Dynamic Application that uses the *Consume Cached Results* option, you cannot define requests.
 - For XSLT Dynamic Applications that use the *Consume Cached Results* option, you must still define the XSLT Parser Code that will be used to transform the cached responses.
 - For Snippet Dynamic Applications that use the *Consume Cached Results* option, you must still define snippet code that processes the cached responses; however, Snippet Dynamic Applications that use the *Consume Cached Results* option can consume cached responses from any Dynamic Application type that can cache responses.
 - Dynamic Applications that use the *Consume Cached Results* option can use cached responses from multiple different Dynamic Applications.
 - If a cached response is unavailable when the ScienceLogic platform performs collection for a "Consume Cached Results" Dynamic Application, the platform will automatically attempt to execute the appropriate request to re-populate the cache.

NOTE: If a Dynamic Application includes *Consume Cached Results* and is aligned with a device that is **not** a component device, the Dynamic Application that provides the results to be consumed (includes *Cache Results*) must be aligned with the same device. If a Dynamic Application includes *Consume Cached Results* and is aligned with device that is a **component device**, the related Dynamic Application that provides the results to be consumed (includes *Cache Results*) must be aligned with the root device for that component.

- **Device Dashboard**. Select a device dashboard from a list of all device dashboards in the ScienceLogic platform. The selected device dashboard will be associated with all devices that subscribe to this Dynamic Application and will appear as an option in the **Device Summary** page. For more information about device dashboards, see the **Dashboards** manual.
- **Version Number**. Specifies the current revision level of the Dynamic Application as specified by the developer of the Dynamic Application.
 - **Operational State**. Specifies whether the ScienceLogic platform will collect data from devices using the Dynamic Application. The **Operational State** also specifies whether the ScienceLogic platform will automatically align the Dynamic Application to devices during discovery, re-discovery, and nightly auto-discovery. The following options are available:
 - *Enabled*. The platform will collect data from devices using the Dynamic Application and will automatically align this application to devices during discovery, re-discovery, and nightly auto-discovery.
 - *Disabled*. The platform will not collect data from devices using the Dynamic Application and will not align this application to devices during discovery, re-discovery, and nightly auto-discovery. Selecting *Disabled* does not remove existing device alignments for the Dynamic Application and does not prevent a user from manually aligning the Dynamic Application with a device.

- **Collector Affinity.** Specifies which Data Collectors are allowed to run collection jobs for this Dynamic Application. Choices are:
 - *Default.* If the Dynamic Application is auto-aligned to a component device during discovery, then the Data Collector assigned to the root device will collect data for this Dynamic Application as well. For devices that are not component devices, the Data Collector assigned to the device running the Dynamic Application will collect data for the Dynamic Application.
 - *Root Device Collector.* The Data Collector assigned to the root device will collect data for the Dynamic Application. This guarantees that Dynamic Applications for an entire Device Component Map (DCM) tree will be collected by a single Data Collector. You might select this option if:
 - The Dynamic Application has a cache dependency with one or more other Dynamic Applications.
 - You are unable to collect data for devices and Dynamic Applications within the same Device Component Map on multiple Data Collectors in a collector group.
 - If the Dynamic Application will consume cache produced by a Dynamic Application aligned to a non-root device (for instance, a cluster device).
 - *Assigned Collector.* The Dynamic Application will use the Data Collector assigned to the device running the Dynamic Application. This allows Dynamic Applications that are auto-aligned to component devices during discovery to run on multiple Data Collectors. You might select this option if:
 - The Dynamic Application has no cache dependencies with any other Dynamic Applications.
 - You want the Dynamic Application to be able to make parallel data requests across multiple Data Collectors in a collector group.
 - The Dynamic Application can be aligned using mechanisms other than auto-alignment during discovery (for instance, manual alignment or alignment via Device Class Templates or Run Book Actions).
- **Poll Frequency.** Specifies how often the platform will use this Dynamic Application to collect data from the device(s) the Dynamic Application is aligned with. Choices range from *Every 1 Minute* to *Every 24 Hours*.

NOTE: You can define a custom poll frequency for a Dynamic Application aligned with one or more devices in a device template. The poll frequency defined in the device template overrides the poll frequency defined for the Dynamic Application. Devices to which the device template has been applied will use the poll frequency defined in the device template. You can override both the **Poll Frequency** field and on a per-device basis from the **Dynamic Application Collections** page.

- **Maximum Devices.** Appears only for Dynamic Applications of type Bulk Snippet Configuration and Bulk Snippet Performance. Specifies the maximum number of devices that a single instance of the Dynamic Application can collect data from. If the number of discovered devices exceeds the value in this field, the platform runs additional instances of the Dynamic Application. Enter a number between 1 and 500.

- **Abandon Collection**. Specifies how many collection objects must timeout in one poll period before the platform stops trying to collect data for the Dynamic Application from a specific device. For more information about how this field works, see the [Abandon Collection Property](#) section.
- **Context**. Optional field that is used for Dynamic Applications that use the SNMP protocol. For SNMPv3, specifies the snmp ID of the device from which you want to retrieve data. This helps differentiate between multiple instances of MIBs and OIDs.
- **Disable Rollup of Data**. If you select this checkbox, the ScienceLogic platform does not roll up the collected data from this Dynamic Application. Selecting this checkbox decreases the workload on the ScienceLogic system. The ScienceLogic platform displays rolled-up data from Dynamic Applications in multiple pages throughout the user interface, including, but not limited to, dashboard widgets, reports, and device summary/performance pages. If you select this checkbox, those pages that depend on rolled-up data will no longer display data for this Dynamic Application.
- **Description**. Optional field that describes the purpose and function of the Dynamic Application.
- **Release Notes & Change Log**. Optional field that specifies any additional information the developer wants to provide with the Dynamic Application.

For Dynamic Applications that use the Snippet, SOAP, WMI, XML, or XSLT protocols, the following additional fields appear:

- **Caching**. For information about this field, see the [Caching](#) chapter. If you do not want to use the caching feature, select *No Caching* in this field.

NOTE: The **Caching** field does not appear for Bulk Snippet Dynamic Applications.

- **Component Mapping**. If you select this checkbox, the ScienceLogic platform will use data collected by the Dynamic Application to create Component Devices. For more information about this checkbox, see the [Dynamic Component Mapping](#) chapter.

For Dynamic Applications of type "configuration", the following additional fields appear:

- **Null Row Option**. This field is active only for Dynamic Applications of type "configuration". If this Dynamic Application does not collect any values for a row in a table of configuration data that previously returned data, you can specify how the ScienceLogic platform will display the row in the **Configuration Report** page. Choices are:
 - *--values*. The platform will display the "--" (dash dash) characters to signify that no data was collected for the current snapshot.
 - *Last collected*. The platform will display the last successfully collected values for the row. The platform will not indicate that it has collected no values for the row.
 - *Last collected with indicator*.
 - *Hide Row*. The platform will not display the row that contains no values.

- **Null Column Option**. This field is active only for Dynamic Applications of type "configuration". If this Dynamic Application does not collect a value for a table cell for which data was previously collected, you can specify what the platform will display in the table cell in the **Configuration Report** page. Choices are:
 - *--values*. The platform will display the "--" (dash dash) characters to signify that no data was collected for the current snapshot.
 - *Last collected*. The platform will display the last successfully collected value for the table cell. The platform will not indicate that it has collected no value for the table cell.
 - *Last collected with indicator*. The platform will display the last successfully collected value for the table cell. The value in the table cell is highlighted, to indicate that this is not the latest value, but only the latest successfully collected value.

4. To edit one or more of the properties for a Dynamic Application, change the value for that property, and then select the **[Save]** button.

NOTE: If you modify a Dynamic Application that is included in an imported PowerPack, for example, Dynamic Applications provided by ScienceLogic, your changes will be overwritten if the PowerPack is reimported and reinstalled.

The Abandon Collection Property

The **Abandon Collection** field in the **Dynamic Applications Properties Editor** specifies how many collection objects must timeout before the ScienceLogic platform stops trying to collect data for the Dynamic Application from a specific device. At the next scheduled collection interval, the ScienceLogic platform will try again to collect data for this Dynamic Application from the device.

The following options are available in the **Abandon Collection** field:

- **Default**. Specifies a threshold of two collection objects.
- **After 1 Timed-Out Objects to After X Timed-Out Objects**. Specifies how many collection objects must time out before this collection session is abandoned on the device. Options will be displayed for each number between 1 and the number of collection objects defined in the Dynamic Application. For example, if three collection objects are defined in a Dynamic Application, the following options are available:
 1. *Default*
 2. *After 1 Timed-Out Objects*
 3. *After 2 Timed-Out Objects*
 4. *After 3 Timed-Out Objects*

During a polling session on a device, the frequency of which is controlled by the **Poll Frequency** field, the Dynamic Application collection process connects to the aligned device using the aligned credential and then requests each collection object. If the ScienceLogic platform cannot retrieve a value for a collection object before the timeout limit (defined in the credential), the collection object is considered timed-out. In a single polling session on a single device, when the number of collection objects that are timed-out reaches the number specified in the **Abandon Collection** field, the ScienceLogic platform abandons collection of the Dynamic Application on the device for the current polling session. When collection is abandoned:

- The ScienceLogic platform will not attempt to collect any of the other collection objects in the Dynamic Application from that device in this polling session.
- The ScienceLogic platform can generate an event with the message "Dynamic Application stopped collection - abandon threshold crossed". The event will include the Dynamic Application ID, device ID, and the collection object IDs that timed out.

NOTE: The event "Dynamic Application Stopped collection - abandon threshold crossed" is disabled by default. To enable this event, go to the **Event Category Manager** page and edit the definition for this event.

- Subsequent polling sessions are not affected. During the next polling session, the collection process for the Dynamic Application will attempt to collect all the collection objects
- If a device is unavailable or consistently abandons collection for more than two days, the ScienceLogic platform will set the **Collect** status of all objects for all the Dynamic Applications aligned with the device to *No*.

NOTE: For all objects except those retrieved from a database, the timeout limit is specified in the credential. For database objects, the timeout limit is specified internally by the ScienceLogic platform.

Creating Dynamic Applications

Overview

To create a new Dynamic Application, you must perform the following general steps:

1. Determine what data you want to collect from a monitored device that would help you manage that device.
2. Determine what management interfaces are available on the monitored device and how the data that you need is made available - through an SNMP agent, WMI agent, web server, database, or another method. You can then choose the appropriate protocol type for your Dynamic Application. The same information might be available through multiple management interfaces on the monitored device; you will need to determine which management interface you will use based on ease of implementation, security, and other factors.
3. Determine which data points you want to collect, and how you want that data displayed in the ScienceLogic platform. You can then:
 - Choose the archetype of your Dynamic Application.
 - List the requests, collection objects, and presentation objects you will need to create.
4. Determine the other features that you will use for this Dynamic Application. In addition to collecting and displaying data, the ScienceLogic platform can:
 - Automatically align your Dynamic Application to devices during discovery.
 - Generate alerts when the collected data meets specified thresholds.
 - Create component devices using the collected data.
 - Cache the results of requests defined in Snippet, SOAP, WMI, XML, and XSLT Dynamic Applications for use by other Dynamic Applications.

NOTE: If you configure your Dynamic Application to cache the results, you cannot include collection objects (other than a discovery object) or presentation objects in the same Dynamic Application. You must create a second Dynamic Application to use the cached results to display data in the ScienceLogic platform. If you configure your Dynamic Application to cache the results, skip steps six and seven of these instructions.

- Use the performance graphs defined in a Dynamic Application to represent CPU, Memory, or Swap utilization for subscriber devices.
 - Use the collected data to automatically populate asset records.
 - Use the collected data to automatically populate inventory lists for hardware components and installed software.
 - Use the collected data to create and/or populate custom attributes.
5. Create the container for the Dynamic Application. During this step, you will use the protocol and archetype you chose in the previous steps to define the Dynamic Application type. You will also define the basic properties for your Dynamic Application, such as how often the ScienceLogic platform will perform collection. There are three methods for creating a container for your Dynamic Application:
 - Creating a new Dynamic Application from scratch.
 - Creating a new Dynamic Application by duplicating an existing Dynamic Application using the **[Save As]** button. This option is useful if you are creating multiple, similar Dynamic Applications.
 - If you are creating an SNMP Dynamic Application, use the SNMP Walker Tool to create the Dynamic Application and populate the collection objects.
 6. If you are creating a Snippet, Bulk Snippet, SOAP, WMI, or XSLT Dynamic Application, define the requests or snippets that will be used to populate the collection objects in your Dynamic Application. If one or more other Dynamic Applications are already configured to cache all the data you need to populate the collection objects in this Dynamic Application, you can configure your Dynamic Application to consume cached requests.
 7. Create the collection objects for the Dynamic Application.
 8. If you are creating a performance or journal Dynamic Application, create the presentation objects for the Dynamic Application.
 9. At this stage of development, you might want to test the basic functionality of your Dynamic Application before adding additional features such as discovery objects or alerts. To test your Dynamic Application:
 - If required, create the appropriate credential for your Dynamic Application.
 - Discover a device from which the Dynamic Application can collect data.
 - In the **Dynamic Application Collections** page for the device, manually align the Dynamic Application.

TIP: To speed up your testing, you might want to change the **Poll Frequency** setting in the **Dynamic Applications Properties Editor** to *Every 1 minute*. When your Dynamic Application is complete, you can change this setting back to the desired value.

10. Implement the other features you want to use for your Dynamic Application. After implementing each feature, you might want to re-test your Dynamic Application. To implement additional features:
 - To automatically align your Dynamic Application to devices during discovery, create one or more [discovery objects](#).
 - To generate alerts when the collected data meets specified thresholds, define [thresholds and alerts](#).
 - To create component devices using the collected data, enable [Dynamic Component Mapping](#).
 - To cache the results of requests for use by other Dynamic Applications, enable [caching](#).
 - To use the performance graphs defined in your Dynamic Application to represent CPU, Memory, or Swap utilization for subscriber devices, define a vitals link for one or more [presentation objects](#).
 - To use the data collected by your Dynamic Application to automatically populate asset records, define an asset link for one or more [collection objects](#).
 - To use the data collected by your Dynamic Application to automatically populate inventory lists for hardware components and installed software, define an inventory link for one or more [collection objects](#).
 - To use data collected by your Dynamic Application to automatically create and/or populate custom attributes, align a custom attribute with one or more [collection objects](#).

Creating the Container for a Dynamic Application

If you are creating a new Dynamic Application from scratch, you must create a container for your Dynamic Application. To do this:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Select the **[Actions]** button, and then select *Create New Dynamic Application*. The **Dynamic Applications Create New Application** page appears in a new window:

The screenshot shows the 'Dynamic Applications | Create New Application' page. The top navigation bar includes 'Guide' and 'Reset' buttons. The main configuration area is divided into several columns of settings:

- Application Name:** [Application]
- Application Type:** [SNMP Performance]
- Caching:** [No caching]
- Device Dashboard:** [None]
- Version Number:** [Version 1.0]
- Operational State:** [Enabled]
- Collector Affinity:** [Default]
- Poll Frequency:** [Every 15 Minutes]
- Maximum Devices:** [50]
- Abandon Collection:** [Default]
- Context:** [Empty field]
- Null Row Option:** [-- values]
- Null Column Option:** [-- values]
- Disable Rollup of Data:** [Unchecked checkbox]
- Component Mapping:** [Unchecked checkbox]

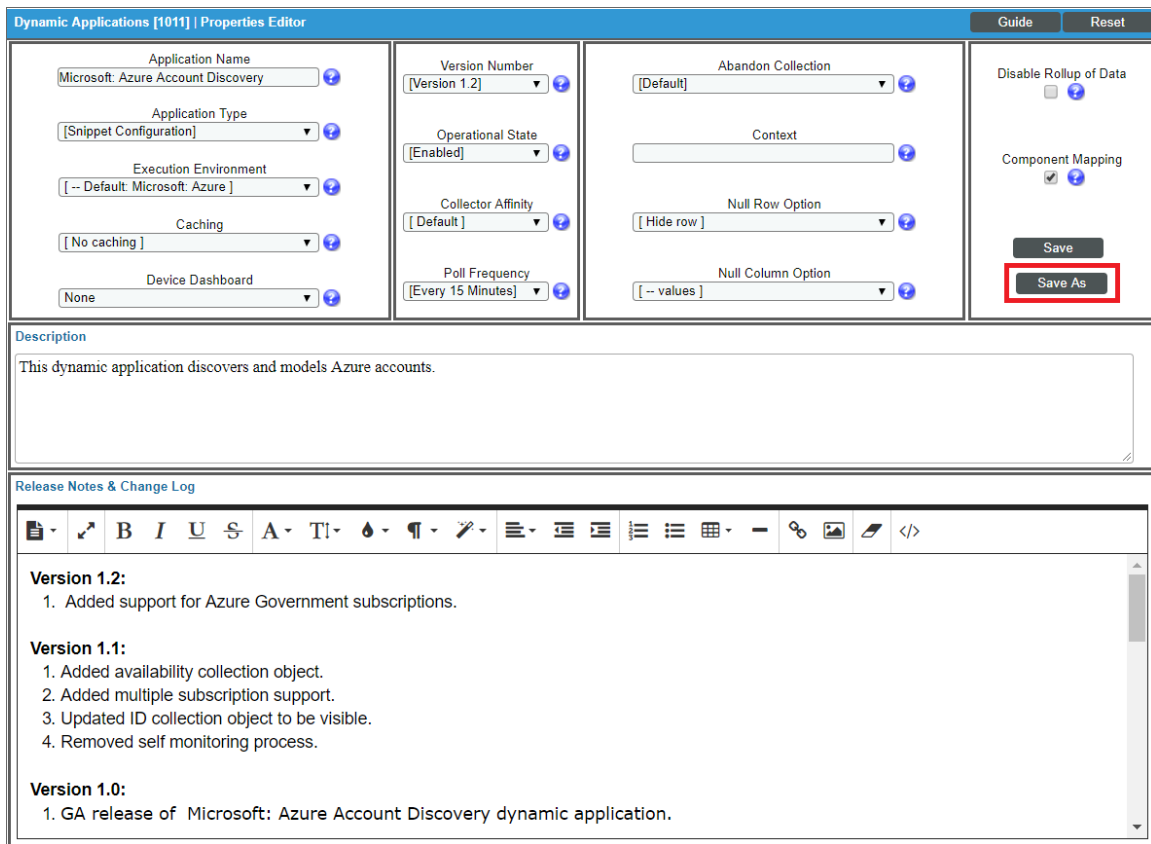
Below the configuration area is a 'Description' text area. At the bottom is a 'Release Notes & Change Log' section with a rich text editor toolbar and a text area containing 'Start typing ...'.

3. Enter a value in each field in the **Dynamic Applications Create New Application** page. For a description of each field, see the [Dynamic Application Settings](#) chapter.
4. Select the **[Save]** button to create your new Dynamic Application.

Duplicating an Existing Dynamic Application

If you are creating a new Dynamic Application that is similar to an existing Dynamic Application, you can duplicate the existing Dynamic Application. To do this:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to duplicate. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Enter a new name for your Dynamic Application in the **Application Name** field. You can also change the values in the other fields on this page. For a description of each field on this page, see the [Dynamic Application Settings](#) chapter.
4. Select the **[Save As]** button to duplicate the Dynamic Application:



When a Dynamic Application is duplicated:

- A new Dynamic Application is created using the properties defined in this page.
- New instances of the collection objects, presentation objects, thresholds, and alerts in the existing Dynamic Application are created and associated with the new Dynamic Application.
- New instances of the event policies associated with the existing Dynamic Application are created and associated with the alerts in the new Dynamic Application.
- For Dynamic Applications that have the **Component Mapping** checkbox enabled, a new instance of the associated component device class is created and associated with the new Dynamic Application.
- For Dynamic Applications that have the **Component Mapping** checkbox enabled, the Dynamic Applications aligned in the **Dynamic App Component Alignment** page will also be aligned with the new Dynamic Application. The ScienceLogic platform does not create new instances of the Dynamic Applications aligned in the **Dynamic App Component Alignment** page.
- The new Dynamic Application is not automatically included in a PowerPack.

Creating a Dynamic Application Using the SNMP Walker Tool

The easiest way to create an SNMP Dynamic Application is to use the SNMP Walker Tool to determine what information is available from the SNMP agent on a device. You can then select the SNMP OIDS directly from the SNMP Walker Tool and add them to a new or existing Dynamic Application. For information on how to use the SNMP Walker Tool to create an SNMP Dynamic Application, see the **SNMP Dynamic Application Development** manual.

Chapter

6

Collection Objects

What is a Collection Object?

Each **Collection Object** in a Dynamic Application specifies a data point that the ScienceLogic platform will attempt to collect. A Dynamic Application must have at least one Collection Object. Each Dynamic Application protocol uses a different method to specify how the ScienceLogic platform should collect each Collection Object.

How Does A Collection Object Get Collected?

Dynamic Applications that use the Snippet, SOAP, WMI, XSLT, or PowerShell protocols include one or more **requests** that define how the ScienceLogic platform should request data from a device. For Snippet Dynamic Applications, the requests are referred to as **snippets**. Each request specifies an operation that will gather a response from the device.

Each collection object in a Snippet, SOAP, WMI, XSLT, or PowerShell Dynamic Application is associated with a request. The collection object specifies how to parse a data point from the response. A single request can be used to populate multiple collection objects.

For example, the default "Windows Interface" WMI Dynamic Application includes one request and will return one response. The request includes a WMI query that retrieves six properties. The Dynamic Application includes six collection objects, one collection object for each property included in the response.

Collection objects in Dynamic Applications that use the Database, SNMP, XML, or PowerShell protocols do not use separate requests. For Database, SNMP, XML, or PowerShell Dynamic Applications, the collection object itself defines the method of collection:

- A **Database** collection object specifies a database query that the ScienceLogic platform must execute to populate the collection object.

- An **SNMP** collection object specifies the SNMP OID that the ScienceLogic platform must retrieve to populate the collection object.
- An **XML** collection object specifies how to parse a value from an XML document to populate the collection object. The XML document to parse is specified in the SOAP/XML credential aligned with the Dynamic Application.
- A **PowerShell** collection object specifies the PowerShell command that the ScienceLogic platform must execute to populate the collection object.

This chapter describes how to define the collection object properties that are common to all Dynamic Application protocol types. For information on how to define requests and define the collection object fields specific to each Dynamic Application protocol, see the supplementary manual for each Dynamic Application protocol type.

How Collection Objects are Used by the ScienceLogic Platform

The data collected for each Collection Object is displayed differently for each Dynamic Application archetype:

- For the **Configuration** archetype, each collection object represents a column in a table of collected data. The value(s) collected by the ScienceLogic platform for the collection object are displayed in that column. For example, the default "Host Resource: Software" Dynamic Application includes two collection objects called "Install Date" and "Software Title". The table of data looks like this:

The screenshot shows a 'Configuration Report | APC: Battery Status' window. It includes a 'Snap-Shot Date [2013-05-10 14:00:00]' and a 'Snap-Shots' button. The main table displays the following data:

UPS Basic Battery Status	batteryNormal
Elapsed Time on Battery	66000
Date Batteries Last Replaced	08/01/08
Percent Battery Remaining	5748
Battery Run Time Remaining	60000
UPS Battery Replacement Message	noBatteryNeedsReplacing
Number of External Battery Packs on UPS	0
Number of Defective External Battery Packs	0

NOTE: A Dynamic Application can define multiple tables.

- For the **Performance** archetype, collection objects are used in **Presentation Objects**. Each presentation object specifies a formula that must include one or more collection objects. The formula in a presentation object is used to generate a time series graph of data. For more information on presentation objects for Dynamic Applications with an archetype of Performance, see the [Presentation Objects](#) chapter.
- For the **Journal** archetype, collection objects are attributes of a **Journal Entry**. The collection objects are used in **Presentation Objects** to define how the collection objects will be formatted in the log of journal entries. For information on collection objects in Dynamic Applications with an archetype of Journal, see the [Snippet Dynamic Application Development](#) manual.

Collection Objects are also used to define alerts in Dynamic Applications. Each alert specifies a formula that must include one or more collection objects. The ScienceLogic platform examines alert formulas to evaluate whether an alert should be generated based on the collected data. For more information on alerts for Dynamic Applications, see the [Alerts and Thresholds](#) chapter.

Viewing the Collection Objects in a Dynamic Application

To view the collection objects in a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to view the collection objects for. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed. The **Collection Object Registry** pane at the bottom of the page displays the following information about each collection object:


Collection Object Registry												
	Object Name	Class Type	Class ID	SNMP OID	Group	ID	Asset Link	Change Alerting	Align	Edit Date		<input type="checkbox"/>
1.	🔧 Pccm Gateway DChannel Status	Config Enum	11	.1.3.6.1.4.1.9.9.156.1.3.1.1.13	--	o_614	--	Disabled	Left	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>
2.	🔧 Pccm Gateway Status Reason	Config Enum	11	.1.3.6.1.4.1.9.9.156.1.3.1.1.10	--	o_611	--	Disabled	Left	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>
3.	🔧 Pccm Gateway Time Last Registered	Config Date & Time	19	.1.3.6.1.4.1.9.9.156.1.3.1.1.12	--	o_613	--	Disabled	Left	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>
4.	🔧 Pccm Gateway Time Last Status Update	Config Date & Time	19	.1.3.6.1.4.1.9.9.156.1.3.1.1.11	--	o_612	--	Disabled	Left	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>
5.	🔧 Discovery	Discovery	100	.1.3.6.1.4.1.9.9.156.1.3.1.1.1	--	o_610	--	--	--	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>
6.	🔧 Gateway Status	Config Enum	11	.1.3.6.1.4.1.9.9.156.1.3.1.1.5	--	o_609	--	Disabled	Left	2013-05-07 22:04:21	🔍 📄 🔄	<input type="checkbox"/>


- **Object Name.** The name of the collection object.
- **Class Type** and **Class ID.** The data type of the collection object. See the [Creating a Collection Object](#) section for a description of the possible class types.
- **Protocol Specific Collection Method.** The label displayed for this column is different for each Dynamic Application protocol. The values displayed in this column specify how the ScienceLogic platform will request the collection object or parse the collection object from a separately defined request.
- **Group.** The group number for the collection object.
- **ID.** The unique ID assigned to the collection object by the ScienceLogic platform. You can use this ID to reference the collection object in presentation formulas and alert formulas. The unique ID will always start with "o_" (lowercase "oh", underscore).
- **Edit Date.** The last time a user edited the collection object.

For collection objects in configuration Dynamic Applications, the following additional information is displayed:

- **Asset Link.** The asset record field that is automatically populated with values collected for the collection object.

- **Change Alerting**. Specifies whether change alerting is enabled for the collection object. Possible values are:
 - *Enabled*. Specifies that if a collected value for the collection object changes, the ScienceLogic platform will trigger an event.
 - *Disabled*. Specifies that if a collected value for the collection object changes, the ScienceLogic platform will not trigger an event.
- **Align**. Specifies how the collected values for the collection object will be justified in the table of configuration data. Possible values are:
 - *Left*
 - *Center*
 - *Right*


To edit a collection object, select its wrench icon ()

For SNMP Dynamic Applications, an information icon () is displayed for each collection object. Select this icon to view information about the SNMP OID defined for a collection object. The information displayed is defined in the SNMP MIB for that SNMP OID.

Creating a Collection Object

NOTE: This section does not apply to Dynamic Applications with an archetype of Journal. For information on collection objects in Dynamic Applications with an archetype of Journal, see the *Snippet Dynamic Application Development* manual.

To add a Collection Object to a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to add a collection object to. Select its wrench icon () . The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed. The fields that appear on this page are different for each Dynamic Application type.

For example, the **Dynamic Applications | Collections Objects** page for an SNMP Configuration Dynamic Application looks like this:

The **Dynamic Applications | Collections Objects** page for an SNMP Performance Dynamic Application looks like this:

4. Enter values in the fields on this page. This chapter describes the fields on this page in separate sections. Depending on the type, purpose, and settings of your Dynamic Application, you will need to define values listed in different sections. The sections are:
 - **Basic Settings.** This section describes the fields that specify the name, description, and data type of the collection object. You must supply values in these fields for all collection objects.
 - **Collection Settings.** This section describes the fields that specify how the ScienceLogic platform should collect values for the collection object. These fields are different for each Dynamic Application protocol. You must supply values in these fields for all collection objects.
 - **Grouping & Indexing Settings.** This section describes the fields that specify how the ScienceLogic platform should group collection objects. These fields are optional; however, collection objects in most Dynamic Applications have values specified in these fields. The **Usage Type** field also defines how the data collected by a Dynamic Application is used to build device relationships.

- **Configuration Dynamic Application Settings.** This section describes the fields that specify how the ScienceLogic platform should use and display the collected data for collection objects in configuration Dynamic Applications. These fields appear only for configuration Dynamic Applications. These fields are optional.
- **Standard Deviation Settings.** This section describes the fields that specify how the ScienceLogic platform should calculate standard deviation data for a collection object. These fields appear only for performance Dynamic Applications. These fields are optional. Supply values in these fields only if you are using this collection object in the standard deviation function in an alert formula.
- **Dynamic Component Mapping Settings.** This section describes the fields that specify how the ScienceLogic platform should create component devices using the collection object. These fields appear only for Dynamic Applications that have the **Component Mapping** checkbox checked in the **Dynamic Applications Properties Editor** page. These fields are optional. Supply values in these fields only if you are using this collection object to create a component device.

5. Select the **[Save]** button to save the collection object.

Basic Settings

To define the basic settings for a collection object, supply a value in the following fields:

- **Object Name.** The name of the collection object. This name is used to label the collection object in the following places in the ScienceLogic platform:
 - In the **Collection Object Registry** pane at the bottom of this page.
 - In the list of collection objects in the **Dynamic Applications Presentation Objects** page.
 - In the list of collection objects in the **Dynamic Applications Alert Objects** page.
 - For configuration Dynamic Applications, in the **Dynamic Application Collections** page (the **[Collections]** tab in the **Device Administration** panel). This page displays the collection status of each collection object.
 - For configuration Dynamic Applications, in the **Configuration Report** page (the **[Configs]** tab in the **Device Reports** panel). The **Object Name** appears as the column heading for the collection object.

- **Class Type.** The type of data that will be collected for this collection object. The data type defines how the ScienceLogic platform will process and display the collected values. For configuration Dynamic Applications, the choices are:
 - [10] *Config Character.* Collection object will be populated with strings that the ScienceLogic platform should display with no text formatting applied.
 - [11] *Config Enum.* Collection object will be populated with numeric values that map to a numeric label. You must define the label value for each of the possible numeric values in the **Enum Values** field.
 - [12] *Config Timeticks.* Collection object will be populated with numeric values that represent timeticks.
 - [14] *Config Gauge.* Collection object will be populated with numeric values that can increase and decrease with each polling session.
 - [15] *Config Numeric.* Collection object will be populated with numeric values.
 - [16] *Config Time Stamp.* Collection object will be populated with timestamps.
 - [17] *Config MAC Address.* Collection object will be populated with MAC addresses.
 - [18] *Config IP Address.* The values collected by this collection object are IP addresses. If the collected value is a hexadecimal string, the ScienceLogic platform will convert the string to the standard dotted decimal format and store that dotted decimal value for the collection object.
 - [19] *Config Date & Time.* Collection object will be populated with date and time values in the format "Year-Month-Day,Hour:Minute:Second" or in Hex notation. The ScienceLogic platform will convert the collected value to a human-readable date & time string before storing the value.
 - [100] *Discovery.* For information about this class type, see the [Creating a Discovery Object](#) section.
 - [101] *Label (Hourly Polled).* Collection object will be populated with strings that will be collected only once an hour and for which only the newest collected value will be stored.
 - [104] *Label (Always Polled).* Collection object will be populated with strings for which only the newest collected value will be stored.
 - [105] *Index.* Collection object will be populated with the indexes associated with the list of values, not the values themselves.
 - [108] *Label (Config Group).* Collection objects of this type are not collected. The **Object Name** defined for this collection object will be used to label the table of data for the group selected in the **Group** field for this collection object.

For performance Dynamic Applications, the choices are:

- [1] *Performance Counter.* Collection object will be populated with counter values. When the ScienceLogic platform evaluates a presentation formula or alert formula that contains a collection object of this type, the value for the collection object is calculated by subtracting the previous collected value for the collection object from the current collected value.
- [2] *Performance Percent.* Collection object will be populated with percentage values (i.e. values from 0 - 100).
- [3] *Performance Kilobyte.* Collection object will be populated with values that represent size in kilobytes.

- [4] *Performance Gauge*. Collection object will be populated with numeric values that can increase or decrease with each polling session. The values collected are the actual values that should be used when the ScienceLogic platform evaluates presentation formulas and alert formulas.
 - [100] *Discovery*. For information about this class type, see the [Creating a Discovery Object](#) section.
 - [101] *Label (Hourly Polled)*. Collection object will be populated with strings that will be collected only once an hour. Collection objects of this type are used to label the lines on a performance graph.
 - [104] *Label (Always Polled)*. Collection object will be populated with strings. Collection objects of this type are used to label the lines on a performance graph.
 - [105] *Index*. Collection object will be populated with the indexes associated with the list of values, not the values themselves.
 - [110] *Label IP (Hourly Polled)*. The object value for the label (IP address) will be retrieved only hourly, not at the **Poll Frequency** defined for the Dynamic Application. In the **String Type** field, specify the format of the collected string, to allow the ScienceLogic platform to convert the string to dotted decimal notation.
 - [111] *Label IP (Always Polled)*. The object value for the label (IP address) will be at the **Poll Frequency** defined for the Dynamic Application. In the **String Type** field, specify the format of the collected string, to allow the ScienceLogic platform to convert the string to dotted decimal notation.
- **String Type**. Specifies whether a collected string is in hex notation or standard text. This field helps the ScienceLogic platform convert collected strings to the appropriate format. Choices are:
 - *Standard*. Collected string is a standard text string.
 - *Hex*. Collected string is in hex notation.
 - **Description**. A description of the collection object. This field is optional.
 - **Formula**. If you require the ScienceLogic platform to examine and manipulate the object's value before storing the value in the platform, you can enter a formula in this field. Instead of storing the collected values for this collection object, the platform stores the result of this formula. The formula field can include any arithmetic Python statement and any Python time statement that returns a single value. That returned, single value will be stored instead of the collected value. You can use the following variables in the formula field:
 - **%V**. Collected value for the object.
 - **%P**. Collected value for the object from the last successful poll.
 - **%L**. Date and time of the last successful poll of the object, in DateTime format, including seconds.
 - **%C**. Date and time of the current poll of the object, in DateTime format, including seconds.
 - **%E**. Number of seconds, in integer format, that have elapsed since last poll of object.
 - **Disable Object Maintenance**. When you select this checkbox, the ScienceLogic platform will never automatically disable collection of this object on an aligned device. For details on how the ScienceLogic platform manages collection status, see the chapter on [Managing Dynamic Applications](#).

Collection Settings

The fields that define how the ScienceLogic platform should collect data for a collection object are different for each Dynamic Application protocol.

In general, one field specifies how the ScienceLogic platform should request the collection object or parse the collection object from a separately defined request. For example, for the SNMP protocol, this field is called **SNMP OID** and specifies the SNMP OID that the ScienceLogic platform must request to obtain a value for this collection object. This field has one of the following labels:

- *SNMP OID*
- *SQL Query*
- *SOAP Tags*
- *Snippet Arguments*
- *WMI Request Arguments*
- *PowerShell Arguments*
- *XML Tags*
- *XSLT Tags*

For SOAP, Snippet, WMI, and XSLT Dynamic Applications, a second field specifies the request or snippet that will return a value for the collection object. This field has one of the following labels:

- *SOAP Request*
- *Snippet*
- *WMI Request*
- *PowerShell Request*
- *XSLT Request*

For more information about these fields, see the supplementary manual for the protocol you are working with.

Grouping & Indexing Settings

The collection objects in a Dynamic Application can be divided into *groups*. When collection objects are in the same group:

- For configuration Dynamic Applications, the collection objects will appear in the same table of configuration data.
- The collected values for the collection objects will share the same **indexes**. When a list of values is collected for a collection object, indexes are used to maintain an association between values collected at different times for the same collection object and maintain associations between collection objects in the same Dynamic Application. For more information, see the [Indexing](#) chapter. Typically, you should group collection objects together if:
 - The collection objects will be used in the same presentation object formula.

- The collection objects will be used in the same alert formula.

The following fields define the grouping and indexing settings:

- **Group Number.** The group for this collection object. The options are *No Group* or a group number from 1 to 50.
- **Usage Type.** If *Group Index* is selected in this field, the ScienceLogic platform will use the collected values for this collection object to assign indexes for this **Group Number**. To use the default indexing method for the collection objects in this group, select *Standard* in this field. See the [Indexing](#) chapter for information on how to use this field. The **Usage Type** field also defines how the data collected by a Dynamic Application is used to build device relationships. For more information about how to use the **Usage Type** field to define relationships, see the [Dynamic Application Relationships](#) chapter.

Configuration Dynamic Application Settings

For collection objects in configuration Dynamic Applications, the following additional fields are displayed:

- **Custom Attribute.** This field appears only for collection objects in configuration Dynamic Applications. For details, see the section on [Custom Attribute Settings](#).
- **Asset/Form Link.** The ScienceLogic platform can automatically populate the asset record for a device using data collected by configuration Dynamic Applications. In these two drop-down lists, you can select a standard asset field and/or a customized asset field that appears in the **Asset Properties** page or the other pages in the **Asset Administration** panel. The selected asset field will be automatically populated with the value collected for this collection object. To populate a standard asset field, select the name of the field from the first drop down list. Choices are:
 - *Make*
 - *Model*
 - *Serial Number*
 - *Operating System*
 - *Host ID / SID*
 - *OS System Name*
 - *CPU Count*
 - *CPU Type/Make*
 - *CPU Speed*
 - *Firmware / BIOS Name*
 - *Installed Memory*
 - *Asset Tag*
 - *Disk Array Size*
 - *Disk Count*
 - *Disk Size*

To populate a custom asset field, select the name of a field in the second drop-down list. The choices are all fields included in the Application Form for "Embedded Asset Form Fields". To view or configure this form, go to System > Customize > Form Fields.

- **Inventory Link.** The ScienceLogic platform can automatically populate inventory lists of hardware components and installed software using data collected by configuration Dynamic Applications. In the first drop-down list, you can select the type of inventory record to populate. In the second drop-down list, select the field in that record to populate. If you select *Disabled* in the first drop-down list, the second drop-down list is not displayed. Choices are:
 - *CPU.* The collection object will populate a hardware inventory record for a CPU. If you select this option, the following field choices appear in the second drop-down list:
 - *CPU Index ID*
 - *CPU Description*

- *Disk*. The collection object will populate a hardware inventory record for a physical disk. If you select this option, the following field choices appear in the second drop-down list:
 - *Disk Index ID*
 - *Disk Name*
 - *Disk Capacity Kb*
 - *Disk Read/Write Mode*
 - *Media Description*

- *Memory*. The collection object will populate a hardware inventory record for installed memory. If you select this option, the following field choices appear in the second drop-down list:
 - *Memory Index ID*
 - *Memory Name*
 - *Total Physical Memory Kb*
 - *Physical Memory in Use Kb*
 - *Physical Memory Percent in Use*

- *Virtual Memory*. The collection object will populate a hardware inventory record for virtual memory. If you select this option, the following field choices appear in the second drop-down list:
 - *Swap Index ID*
 - *Swap Description*
 - *Total Swap Memory Kb*
 - *Swap Memory in Use Kb*
 - *Swap Memory Percent in Use*

- *Hardware Components*. The collection object will populate a hardware inventory record for any other type of hardware component. If you select this option, the following field choices appear in the second drop-down list:
 - *Component Index ID*
 - *Component Description*
 - *Component Type*

- *Software Titles*. The collection object will populate a software inventory record. If you select this option, the following field choices appear in the second drop-down list:
 - *Install Date*
 - *Software Title*

- **Change Alerting.** This field enables or disables change detection. If the collected value of the collection object changes, the ScienceLogic platform can automatically trigger an event. You can view the change history for the collection object in the **Configuration Report** page for each subscriber device. The choices for this field are:
 - *Enabled.* Specifies that if the value of this collection object changes, the ScienceLogic platform will trigger an event.
 - *Disabled.* Specifies that if the value of this collection object changes, the ScienceLogic platform will not trigger an event.
- **Table Alignment.** These two drop-down lists define how the values collected for this collection object will be displayed in the table of configuration data in the **Configuration Report** page. The first drop-down list defines how the value will be justified in the table cell. The choices in the first drop-down list are:
 - *Left*
 - *Center*
 - *Right*
- **Hide Object.** If you select this checkbox, the object will not appear in the **Configuration Report** page for each subscriber device.
- **Enum Values.** Applies only to collection objects with a class type `[11] Config Enum`. In this field you can map each possible collected value to a descriptive string that will be displayed in the ScienceLogic platform. Use the following format:

```
#number:value#number:value#number:value
```

Where **number** is the collected value and **value** is the descriptive string. You must include a "#" as the first character in this field and as a separator between number:value pairs.

For example, suppose a `[11] Config Enum` object returns an integer that represents a status condition. The possible returned values and the conditions they represent might be:

```
3 = healthy4 = minor5 = major6 = critical
```

You would supply the following string in the **Enum Values** field:

```
#3:healthy#4:minor# 5:major#6:critical
```

Custom Attribute Settings

In Dynamic Applications of archetype *configuration*, you can:

- Use a collection object to populate the value of an existing custom attribute.
- Use a pair of collection objects to create a custom attribute and provide a value for that custom attribute. You must define a collection object to define the name of the custom attribute; this causes the ScienceLogic system to create a custom attribute with the name from the collection object. You must also define a second collection object to populate the value of the custom attribute.

NOTE: For details on creating and managing custom attributes, see the manual *Using the ScienceLogic API*.

The following fields in the **Collection Objects** page allow you to use one or more collection objects to define and/or populate a custom attribute:

The screenshot shows a configuration form for a collection object. The 'Object Name' field is circled in red and contains the text 'dyn_val'. Below it, the 'XML Tags' field contains the text '<tags><tag><shape>'. The 'Class Type' dropdown is set to '[10 Config Character]'. The 'String Type' dropdown is set to '[Standard]'. The 'Custom Attribute' dropdown is circled in red and set to 'Dynamic Value'. The 'dyn_key' dropdown is also circled in red and set to 'dyn_key'. The 'Group / Usage Type' dropdown is set to '[Group 1]' and '[Standard]'. The 'Asset / Form Link' dropdown is set to '[None]'. The 'Inventory Link' dropdown is set to '[Disabled]'. The 'Change Alerting' dropdown is set to '[Disabled]'. The 'Table Alignment' dropdown is set to '[Left]'. The 'Hide Object' checkbox is unchecked. At the bottom, there is a 'Save' button and a checkbox labeled 'Disable Object Maintenance'.

- **Align to Custom Attribute.** Specify the custom attribute to associate with this collection object. The custom attribute will be populated with a value from a collection object. Choices are:
 - *None.* This collection object is not associated with a custom attribute.
 - *Static.* This collection object is associated with a specific custom attribute.
 - **Static Name.** If you selected *Static* in the **Custom Attribute** field, the *Static Name* field appears. In this field, specify the name of the custom attribute that you want to populate with the value of the collection object. You can select from a list of existing custom attributes.
 - If the list does not include the custom attribute you want to align with the collection, select the plus-sign icon (+). The icon clears the field and allows you to manually enter a value.
 - If you manually specify a custom attribute, the platform will search for a custom attribute with a matching name and populate the custom attribute with the value of this collection object. If the platform does not find a custom attribute with a matching name and therefore creates the custom attribute, the new custom attribute will be an extended custom attribute, for devices. The data type will be integer (for numeric values) or string (for all other value types).
 - *Dynamic Name.* You can use a pair of collection objects to populate the name and value of a custom attribute. You must define each collection object separately. When you select *Dynamic Name* in the **Custom Attribute** field, the name of the custom attribute is populated with the value of the collection object. If the platform does not find a custom attribute with a matching name, the platform will create the custom attribute. If the platform does not find a custom attribute with a matching name and therefore creates the custom attribute, the new custom attribute will be an extended custom attribute, for devices. The data type will be integer (for numeric values) or string (for all other value types).

NOTE: If you select *Dynamic Name* in the **Custom Attribute** field, you must create a second collection object that will populate the value of the custom attribute.

NOTE: Names for custom attributes must conform to XML naming standards. The attribute name can contain any combination of alphanumeric characters, a period, a dash, a combining character or an extending character. If a collected value for an attribute name does not conform to XML standards, the platform will replace non-valid characters with an underscore + the hexadecimal value of the illegal character + an underscore. So "serial number" would be replaced with "serial_X20_number". The attribute label will use the original, non-converted value ("serial number").

- **Dynamic Value.** The value of the custom attribute selected in the *Dynamic Name* field is populated with the value of the collection object.
- **Dynamic Name.** If you selected *Dynamic Value* in the **Custom Attribute** field, the *Dynamic Name* field appears. Select from the list of collection objects that have a **Custom Attribute** value of *Dynamic Name*.

NOTE: The collection object assigned to the *Dynamic Value* is added to the same **Group** as the collection object assigned to the associated *Dynamic Name*. If the collection object for *Dynamic Name* is not assigned to a **Group**, you will be prompted to select a **Group** for the both the collection object for *Dynamic Name* and the collection object for *Dynamic Value*.

NOTE: Each group can contain only one collection object that is assigned to a *Dynamic Value* and only one collection object that is assigned to a *Dynamic Name*. The group can contain other collection objects, but should not contain more than one collection object assigned to a *Dynamic Value* and not more than one collection object assigned to a *Dynamic Name*.

Standard Deviation Settings

To use a collection object with the standard deviation function (deviation) in an alert formula, you must enable standard deviation alerting for the collection object. To enable standard deviation alerting for a collection object, supply a value in the following fields:

NOTE: These fields appear only for collection objects in performance Dynamic Applications that collect numeric values.

The screenshot shows a configuration form for a collection object. The left-hand side contains several input fields: 'Object Name', 'SNMP OID', 'Class Type' (a dropdown menu currently showing '[1] Performance Counter'), and 'Group / Usage Type' (two dropdown menus showing '[No Group]' and '[Standard]'). Below these is the 'Enable Deviation Alerting' section, which is highlighted with a red oval. It consists of a checkbox labeled 'Enable Deviation Alerting', followed by 'max weeks data' and 'min weeks data' input fields. The right-hand side of the form has two large text areas labeled 'Description' and 'Formula'. At the bottom center is a 'Save' button, and at the bottom right is a checkbox labeled 'Disable Object Maintenance'.

- **Enable Deviation Alerting.** If selected, the ScienceLogic platform will calculate and store standard deviation data for the collection object. To use a collection object with the standard deviation function in an alert formula, this checkbox must be checked.
- **Max Weeks Data.** If you selected **Enable Deviation Alerting**, you must specify a value in this field. This field specifies the maximum number of weeks of data that the ScienceLogic platform should use when evaluating a standard deviation function that uses this collection object. The default value is "4". This means that when the ScienceLogic platform evaluates an alert formula that contains a standard deviation function that uses this collection object, the platform will use the last four weeks of data for the collection object to calculate the standard deviation. Any data older than the last 4 weeks is not included in calculations for standard deviation.
- **Min Weeks Data.** If you selected **Enable Deviation Alerting**, you must specify a value in this field. This field specifies the minimum number of weeks of data required for the ScienceLogic platform to evaluate a standard deviation function that uses this collection object. The default value is "2". This means that the platform requires at least the data from the last two weeks from today to calculate standard deviation for this collection object.

NOTE: If you enable deviation alerting for an existing collection object and if the ScienceLogic platform has already collected data for that collection object, the platform will use the already-collected data to "backfill" and calculate mean and standard deviation using the historical data. For example, suppose a Dynamic Application has already been collecting data for 12 weeks. Suppose during the twelfth week of collection, you enable deviation alerting for a collection object for the first time. Suppose you set **Max Weeks Data** to "6" and **Min Weeks Data** to "4". The first time the platform performs calculations for standard deviation for this collection object (once every hour), the platform would then "backfill" and include the last six weeks of historical data when calculating standard deviation for the collection object. If the platform has already collected at least the Min Weeks Data, the platform will use the historical data to calculate standard deviation.

NOTE: The backfill functionality works only during the very first time the ScienceLogic platform performs calculations for standard deviation for a collection object. If you change the settings for **Min Weeks Data** and **Max Weeks Data** so that the new value of **Min Weeks Data** exceeds the previous value for **Max Weeks Data**, the ScienceLogic platform will require you to collect additional data to meet the new requirement. For example, suppose you initially set **Min Weeks Data** to "2" (two weeks) and **Max Weeks Data** to "4" (four weeks). Suppose that 9 weeks later, you edit **Min Weeks Data** to "5" (five weeks) and **Max Weeks Data** to "7". Because of the previous **Max Weeks Data** of "4", the ScienceLogic platform will have retained only four weeks of calculated standard deviation data. After editing the values, you will have to wait one additional week before the ScienceLogic platform will have enough data (five weeks, as specified by **Min Weeks Data**) to generate standard deviation alerts. The ScienceLogic platform will not use historical data to backfill the additional week.

Dynamic Component Mapping Settings

For Dynamic Applications that have the **Component Mapping** checkbox selected in the **Dynamic Applications Properties Editor** page, the **Component Identifiers** field is displayed in the **Collection Objects** page:

The screenshot shows a configuration form for a collection object. On the left side, there are several dropdown menus and checkboxes: Object Name, Snippet Arguments, Class Type ([10] Config Character), String Type ([Standard]), Snippet (Discovery), Group / Usage Type (No Group), Asset / Form Link (None), Inventory Link (Disabled), Change Alerting (Disabled), Table Alignment (Left), and Hide Object. On the right side, there are text areas for Description and Formula, and a list box for Component Identifiers. The Component Identifiers list is circled in red and contains: Availability, Class Identifier 1, Class Identifier 2, Distinguished Name (%C), GUID (%G), and MAC Address. At the bottom, there is a Save button and a checkbox for Disable Object Maintenance.

In this field, select one or more identifiers that the ScienceLogic platform can use to create component devices. The values collected for this collection object will populate the selected identifier for the component devices. Collection objects can define the following component identifiers:

- *Availability*. The availability of the component devices. The values collected using this collection object will be used to generate availability reports and graphs for each component device.
- *Class Identifier 1*. The class type of the component devices, typically the vendor.
- *Class Identifier 2*. The class sub-type of the component devices, typically the model.
- *Device Name (%N)*. The name of the component devices in the ScienceLogic platform.
- *Distinguished Name (%C)*. The name that the system that monitors the component devices (the root device) uses to identify the component devices.
- *GUID (%G)*. The globally unique identifier the ScienceLogic platform will use for the component devices. This value must be unique for all component devices in the ScienceLogic platform. This value is used to determine when a component device should be associated with a different root device.
- *MAC Address*. The MAC Address of the component device.
- *Organization*. During discovery, if the value of the collection object exactly matches an organization name in the ScienceLogic platform, the platform will align the component device with that organization. During subsequent collections, if the value of this collection object changes and exactly matches an organization name in the ScienceLogic platform, the platform aligns the new organization with the component device. If the value of this collection object does not exactly match any organizations in the ScienceLogic platform, the platform will create an event and align the event with the device aligned with the Dynamic Application.
- *Unique Identifier (%U)*. The unique identifier the ScienceLogic platform will use for the component devices. This value must be unique for all component devices associated with the same root device.
- *UUID*. The universally unique identifier for the component devices.
- *GUID*. Globally unique identifier. This value will be unique within the ScienceLogic platform. No two component devices can have the same GUID. This component identifier protects the integrity of a component device even if it is aligned with a different root device (for example, with a vmotion event).

NOTE: Each of these options can be defined by only one collection object in a Dynamic Application. If a component identifier is already defined by a collection object in this Dynamic Application, that component identifier will not be displayed in the list.

If the Dynamic Application is aligned with a component device class, the ScienceLogic platform will create a component device when the collection objects that map to the *Device Name* and *Unique Identifier* of a component device are successfully collected from a device.

For more information about configuring a Dynamic Application to use Dynamic Component Mapping, see the [Dynamic Component Mapping](#) chapter.

Creating a Discovery Object

A **discovery object** is a type of collection object that allows the ScienceLogic platform to automatically align a Dynamic Application to a device during discovery.

The ScienceLogic platform attempts to automatically align Dynamic Applications to devices under the following circumstances:

- When a discovery session that has an **Initial Scan Level** of 1. *Initial Population of Apps* or higher is executed, either manually or on a schedule, the ScienceLogic platform will attempt to automatically align Dynamic Applications using the credentials specified in the discovery session.
- For devices that have the **Auto-Update** and **Dynamic Discovery** checkboxes enabled on the **Device Properties** page, the ScienceLogic platform will attempt to automatically align Dynamic Applications once a day using the credentials aligned with the device.
- When a user selects the lightning bolt icon (⚡) for a Dynamic Application in the **Dynamic Applications Manager** page, the ScienceLogic platform will attempt to automatically align that single Dynamic Application with all devices using the credentials aligned with each device.

Every discovery object specifies a condition that can evaluate to *true* or *false*. When the ScienceLogic platform attempts to automatically align a Dynamic Application to a device, every available credential is used to try and collect every discovery object in the Dynamic Application. If **all** the discovery objects in a Dynamic Application evaluate to *true* for a device, the ScienceLogic platform will align the Dynamic Application with the device.

A discovery object can specify one of the following conditions:

- If the device returns a value for the discovery object, evaluate to *true*. Otherwise, evaluate to *false*.
- If the device returns a value for the discovery object and that value matches specific criteria, evaluate to *true*. Otherwise, evaluate to *false*.
- If the device **does not** return a value for the discovery object, evaluate to *true*. Otherwise, evaluate to *false*.

To add a discovery object to a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to add a collection object to. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed. Enter values in the following fields:
 - **Object Name**. The name of the discovery object. If you are defining only one discovery object for this Dynamic Application, best practice is to enter "Discovery" in this field.
 - **Protocol Specific Collection Method**. The label displayed for these fields is different for each Dynamic Application protocol. Specify how the ScienceLogic platform should collect the discovery object.
 - **Class Type**. Select *[100] Discovery* in this from the drop-down list.

4. Select the **[Save]** button to save the collection object. Because this collection object has been defined with a **Class Type** of *[100] Discovery*, additional fields that are specific to discovery objects are displayed:

The screenshot shows a configuration form for a discovery object. The fields are as follows:

- Object Name: Presence
- SNMP OID: .1.3.6.1.4.1.318.1.1.1.2.1
- Class Type: [100 Discovery]
- Tabular:
- Alignment Condition: [Align if OID is present]
- Validity Check: Where: is > Result

A red warning message is displayed below the Validity Check field: "Validity Check is an optional setting that is used to validate if a discovery object is reliably reporting data. Some agents or clients may respond to a query, even if the related feature or service is not operational."

At the bottom of the form are three buttons: Save, Save As, and Disable Object Maintenance.



5. Enter values in the following fields:

- **Tabular**. If the value specified in the **Protocol Specific Collection Method** fields will return tabular data (a list of values) and you want to match the returned values to specific criteria, select this checkbox. You must then specify a value in the **Validity Check** field. The ScienceLogic platform will iterate through the returned tabular data and search for the value specified in the **Validity Check** field.
- **Alignment Condition**. Specifies how this discovery object should be evaluated. Choices are:
 - *Align if OID is present*. Specifies that if the device returns a value for the discovery object, evaluate to *true*. Optionally, in the **Validity Check** field you can specify the criteria that the returned value must meet.
 - *Align if OID is NOT present*. Specifies that if the device does not return a value for the discovery object, evaluate to *true*.
- **Validity Check**. If you selected *Align if OID is present* in the **Alignment Condition** field, this field is displayed. In this field, you can specify criteria that the returned value must meet for the discovery object to evaluate to *true*. To use this field, enter a value or regular expression in the text area, then select from the drop-down list of operators to specify how that value should be compared to the returned value. The choices for the comparison operator are:
 - *is > Result*. Specifies that the discovery object evaluates to *true* if the specified value is greater than the returned value.
 - *is == Result*. Specifies that the discovery object evaluates to *true* if the specified value is equal to the returned value.
 - *is < Result*. Specifies that the discovery object evaluates to *true* if the specified value is less than the returned value.
 - *is != Result*. Specifies that the discovery object evaluates to *true* if the specified value is not equal to the returned value.
 - *is LIKE Result*. Specifies that the discovery object evaluates to *true* if the returned value contains the specified value.

- *matches REGEX*. Specifies that the discovery object evaluates to *true* if the returned value matches the specified regular expression.
6. Select the **[Save]** button to save the discovery object.

Editing a Collection Object

To edit a Collection Object in a Dynamic Application, perform the following steps:


1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to edit. Select its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed.
4. In the **Collection Object Registry** pane at the bottom of the page, locate the collection object you want to edit. Select the wrench icon () for the collection object you want to edit.
5. The fields in the top pane will be populated with the saved values for the selected collection object. Edit the values in one or more fields. For a description of each field, see the [Creating a Collection Object](#) section.
6. Select the **[Save]** button to save your changes to the collection object. Select the **[Save As]** button to save the changes as a new collection object.

Performing Bulk Actions on Collection Objects

You can perform the following bulk actions on the collection objects in a Dynamic Application:

- Delete collection objects from the ScienceLogic platform.
- Change the **group** value.
- Enable or disable **Change Alerting** for collection objects in configuration Dynamic Applications.

To perform a bulk action on the collection objects in a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to edit. Select its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed.
4. In the **Collection Object Registry** pane at the bottom of the page, locate the collection object(s) you want to perform the action on. Select the checkbox for each collection object you want to perform the action on.
5. In the **Select Action** drop-down list, select the action you want to perform, and then select the **[Go]** button. The selected collection objects will have the selected action applied to them.

Chapter

7

Caching

What is Caching?

When the ScienceLogic platform requests information from a device during Dynamic Application collection, the platform can **cache** the response from the device. If a response is cached, other Dynamic Applications that use the same protocol can use the cached response to populate collection objects. Caching responses reduces the number of requests performed by the ScienceLogic platform and speeds up collection. Typically, Dynamic Applications aligned to component devices must use cached responses from Dynamic Applications aligned with the root device (the device that "manages" the component device).

Caching is configured in the **Caching** drop-down list in the general properties for a Dynamic Application. The following options are available:

- *No caching*. Specifies that this Dynamic Application does not use the caching feature. The requests that populate the collection objects in this Dynamic Application are defined within this Dynamic Application. The responses to those requests are not available to other Dynamic Applications.
- *Cache Results*. The ScienceLogic platform will cache all responses retrieved by the Dynamic Application. See the [Caching Responses](#) section for more information.
- *Consume Cached Results*. The ScienceLogic platform will use the cached responses from other Dynamic Applications to populate the collection objects defined in this Dynamic Application. See the [Consuming Cached Responses](#) section for more information.

NOTE: All types of Dynamic Applications except Bulk Snippet, Database, and ICDA can both cache results and consume cached results.

Caching Responses

Dynamic Applications that use the following protocols can cache responses:

- SNMP
- Snippet
- SOAP
- WMI
- PowerShell
- XML
- XSLT

NOTE: Bulk Snippet Dynamic Applications cannot cache responses.

For XSLT Dynamic Applications, the untransformed response is cached, not the transformed response.

For Dynamic Applications that cache responses, the collection objects defined in this Dynamic Application are no longer collected at a regular frequency. You must not define collection objects that are used to generate configuration tables or performance graphs in a Dynamic Application that caches responses. Typically, Dynamic Applications that cache results will contain only a discovery object.

For information about caching responses in Snippet Dynamic Applications, see the *Snippet Dynamic Application Development* manual.

Consuming Cached Responses

Dynamic Applications that use the following protocols can consume cached responses:

- SNMP
- Snippet
- SOAP
- WMI
- PowerShell
- XML
- XSLT

NOTE: Bulk Snippet Dynamic Applications cannot consume cached responses.

Snippet Dynamic Applications can consume cached results from any Dynamic Application that caches responses. For more information about using cached responses in snippet code, see the **Snippet Dynamic Application Development** manual. All other Dynamic Applications that consume cached responses can do so only from Dynamic Applications that use the same protocol. For example, SOAP Dynamic Applications can consume cached responses only from other SOAP Dynamic Applications. A SOAP Dynamic Application cannot consume a cached response from an XSLT Dynamic Application.

For Dynamic Applications that consume cached responses:

- All collection objects in the Dynamic Application must be populated using cached responses from other Dynamic Applications. If a Dynamic Application consumes cached responses, you cannot define additional requests in that Dynamic Application.
- If the Dynamic Application is aligned to a non-component device, all the other Dynamic Applications that cache requests used by the Dynamic Application must be aligned to the same device.
- If the Dynamic Application is aligned to a component device, all the other Dynamic Applications that cache requests used by the Dynamic Application must be aligned to the same device or the root device.
- If a cached response is unavailable when the ScienceLogic platform performs collection for a Dynamic Application that consumes cached responses, the platform will automatically attempt to execute the appropriate request to re-populate the cache.
- For XSLT Dynamic Applications that use the *Consume Cached Results* option, XSLT Parser Code must still be defined in the Dynamic Application. The XSLT Parser Code will be used to transform the cached responses.

Configuring Collection Objects in Cache-Consuming Dynamic Applications

The following sections describe how to configure collection objects to use responses that are cached by other Dynamic Applications in SOAP, WMI, XML, and XSLT Dynamic Applications.

SOAP Dynamic Applications

When a SOAP Dynamic Application is configured to *Consume Cached Results*:

- The **[Requests]** tab will not be displayed.
- In the **Dynamic Applications | Collections Objects** page, the **SOAP Request** field will display a list of all requests defined in SOAP Dynamic Applications that cache responses. Select the cached response that the ScienceLogic platform should use to populate the collection object.

WMI Dynamic Applications

When a WMI Dynamic Application is configured to *Consume Cached Results*:

- The **[Requests]** tab will not be displayed.
- In the **Dynamic Applications | Collections Objects** page, the **WMI Request** field will display a list of all requests defined in WMI Dynamic Applications that cache responses. Select the cached response that the ScienceLogic platform should use to populate the collection object.

PowerShell Applications

When a PowerShell Dynamic Application is configured to *Consume Cached Results*:

- The **[PowerShell]** tab will not be displayed.
- In the **Dynamic Applications | Collections Objects** page, the **PowerShell Request** field will display a list of all requests defined in PowerShell Dynamic Applications that cache responses. Select the cached response that the ScienceLogic platform should use to populate the collection object.

XML Dynamic Applications

When an XML Dynamic Application is configured to *Consume Cached Results*, an additional field is displayed in the **Dynamic Applications | Collections Objects** page:

- **XML Document**. This field lists all XML Dynamic Applications that cache responses. Select the XML Dynamic Application that requests the XML Document that contains this collection object.

NOTE: For the XML Dynamic Application that caches responses, the URL of the requested XML Document is defined in the credential that is aligned to that Dynamic Application.

XSLT Dynamic Applications

When an XSLT Dynamic Application is configured to *Consume Cached Results*, you must still define XSLT Requests for the Dynamic Application in the **Dynamic Applications Request Editor and Registry** page (the **[Requests]** tab). To define an XSLT Request for an XSLT Dynamic Application that is configured to *Consume Cached Results*:

1. Go to the **Dynamic Applications Request Editor and Registry** page (the **[Requests]** tab) for the Dynamic Application.
2. Supply values in the **Request Name**, **Execution Sequence**, and **Active State** fields.
3. The **XSLT Request Code** field in the **Dynamic Applications Request Editor and Registry** page (the **[Requests]** tab) is renamed **Cached XSLT Request**. The **Cached XSLT Request** field displays a list of all requests defined in XSLT Dynamic Applications that cache responses. Select the cached response that the ScienceLogic platform should transform using the **XSLT Parser Code** defined in this request.
4. Define the XSLT document that the ScienceLogic platform should use to transform the cached response selected in the **Cached XSLT Request** field. The output of the transformation should be in the standard format from which the ScienceLogic platform parses collection objects.

For an XSLT Dynamic Application that is configured to *Consume Cached Results*, the **XSLT Request** field in the **Dynamic Applications | Collections Objects** page will contain only the XSLT Requests defined in the same Dynamic Application.

Collector Affinity for Dynamic Applications that Use Caching

Collector Affinity specifies which Data Collectors are allowed to run collection jobs for Dynamic Applications aligned to component devices. You can define Collector Affinity for each Dynamic Application.

Dynamic Applications [1011] | Properties Editor Guide Reset

Application Name Microsoft: Azure Account Discovery	Version Number [Version 1.2]	Abandon Collection [Default]	Disable Rollup of Data <input type="checkbox"/>
Application Type [Snippet Configuration]	Operational State [Enabled]	Context []	Component Mapping <input checked="" type="checkbox"/>
Execution Environment [-- Default: Microsoft: Azure]	Collector Affinity [Default]	Null Row Option [Hide row]	Save Save As
Caching [No caching]	Poll Frequency [Every 15 Minutes]	Null Column Option [-- values]	
Device Dashboard None			

Description

This dynamic application discovers and models Azure accounts.

Release Notes & Change Log

Version 1.2:

1. Added support for Azure Government subscriptions.

Version 1.1:

1. Added availability collection object.
2. Added multiple subscription support.
3. Updated ID collection object to be visible.
4. Removed self monitoring process.

Version 1.0:

1. GA release of Microsoft: Azure Account Discovery dynamic application.

For Dynamic Applications that create cache or consume cache, you should select the following option for **Collector Affinity**:

- **Root Device Collector.** The Data Collector assigned to the root device will collect data for the Dynamic Application. This guarantees that Dynamic Applications for an entire DCM tree will be collected by a single Data Collector. This option ensure that caching Dynamic Applications can access the required cache.

Dynamic Component Mapping

What is Dynamic Component Mapping?

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

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

- In the ScienceLogic platform, a managed entity is called a **component device**. A component device is an entity that runs under the control of a management device.
- In the ScienceLogic platform, the **root device** is the physical device that manages one or more component devices.
- In the ScienceLogic platform, a **parent device** is a device that has associated entities modeled as component devices. A parent device can be either a root device or another component device.
- In the ScienceLogic platform, a **component identifier** is a collection object that the platform uses to populate information about a component device, such as the device name.

For example, in a Cisco UCS system, the ScienceLogic platform might discover a physical server that hosts the UCS manager. The ScienceLogic platform might discover a chassis as a component device. The chassis is a child device to the physical server; the physical server is the root device. The platform might also discover a blade as a component device that is part of the chassis. The blade is a child device to the chassis. The chassis is the parent device.

Where in the ScienceLogic Platform are Component Devices Displayed?

You can view root devices and their associated component devices in the following places in the ScienceLogic platform:

- The **Device Components** page (Registry > Devices > Device Components) displays a list of all root devices and component devices discovered by the ScienceLogic platform. 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.
- The **Component Map** page (Views > 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. The ScienceLogic platform automatically updates the **Component Map** as new component devices are discovered. The platform also updates each map with the latest status and event information.

Configuring a Dynamic Application to Create Component Devices

Dynamic Applications that use the following protocols can be configured to create component devices:

- SNMP
- Snippet
- SOAP
- WMI
- PowerShell
- XML
- XSLT

For the ScienceLogic platform to create component devices using your Dynamic Application, you must perform the following steps to configure Dynamic Component Mapping:

1. In the **Dynamic Applications Properties Editor** page, check the **Component Mapping** checkbox.
2. In the **Dynamic Applications | Collections Objects** page, configure a collection object that maps to the *Device Name* component identifier.
3. In the **Dynamic Applications | Collections Objects** page, configure a collection object that maps to the *Unique Identifier* component identifier.
4. Create a Device Class with a **Device Type** of *Component* to align with the Dynamic Application.

If a Dynamic Application is configured to create component devices, the ScienceLogic platform will create a component device when the collection objects that map to the *Device Name* and *Unique Identifier* of a component device are successfully collected from a device.

For information on configuring device classes for component devices, see the section on [aligning device classes with component devices](#).

You can also perform the following additional configuration steps for Dynamic Applications that create component devices:

- In the **Dynamic Applications | Collections Objects** page, configure collection objects that map to the following component identifiers:
 - The availability of the component devices. The values collected using this collection object will be used to generate availability reports and graphs for each component device.
 - The manufacturer of the component devices.
 - The model of the component devices.
 - The names the management system uses to refer to the devices, known as the *Distinguished Name*.
 - The globally unique identifier of the component devices. This value is used to determine when a component device should be associated with a different root device.
 - The MAC Address of the component devices.
 - The UUID of the component devices.
- In the **Dynamic Application Component Alignment** page, define a list of Dynamic Applications the ScienceLogic platform should automatically align to component devices created using this Dynamic Application.

Configuring Collection Objects as Component Identifiers

To configure a collection object to map to a component identifier:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application that you want to configure. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Collections]** tab. The **Dynamic Applications | Collections Objects** page is displayed.
4. Select the wrench icon (🔧) for the collection object you want to edit.
5. For Dynamic Applications that have the **Component Mapping** checkbox selected in the **Dynamic Applications Properties Editor** page, the **Component Identifiers** field is displayed. In this field, select one or more identifiers that the ScienceLogic platform can use to create and monitor component devices. The values collected for this collection object will populate the selected identifier for the component devices. Collection objects can define the following component identifiers:

- *Availability*. The availability of the component devices. The ScienceLogic platform uses *Availability* collection objects to generate availability reports and graphs for each component device. For information about how the platform calculates availability using this component identifier, see the [Availability for Component Devices](#) section.
- *Class Identifier 1*. The class type of the component devices, typically the vendor. This field can be used to align a discovered component device with an existing device class.
- *Class Identifier 2*. The class sub-type of the component devices, typically the model. This field can be used to align a discovered component device with an existing device class.
- *Device Name (%N)*. The name of the component devices in the ScienceLogic platform.
- *Distinguished Name (%C)*. The name that the system that monitors the component devices (the root device) uses to identify the component devices. Using this component identifier is optional; this name is not used as a unique identifier by the platform. This component identifier is useful if an additional substitution character is required in your Dynamic Application.
- *GUID (%G)*. The globally unique identifier the ScienceLogic platform will use for the component devices. This value must be unique for all component devices in the ScienceLogic platform. This value is used to determine when a component device should be associated with a different root device. For information about how the ScienceLogic platform moves component devices using this component identifier, see the [Moving Component Devices Between Root Devices](#) section.
- *MAC Address*. The MAC Address of the component device. This value is used to match component devices to devices that are discovered using the discovery tool. If a component device is matched to a physical device, the ScienceLogic platform discovers a component device directly using the discovery tool. For more information, see the [Merging Component Devices with Physical Devices](#) section.
- *Organization*. During discovery, if the value of the collection object exactly matches an organization name in the ScienceLogic platform, the platform will align the component device with that organization. During subsequent collections, if the value of this collection object changes and exactly matches an organization name in the ScienceLogic platform, the platform aligns the new organization with the component device. If the value of this collection object does not exactly match any organizations in the ScienceLogic platform, the platform will create an event and align the event with the device aligned with the Dynamic Application.
- *Unique Identifier (%U)*. The unique identifier for the component devices. This value must be unique for all component devices associated with the same root device.
- *UUID*. The universally unique identifier for the component devices.
- *GUID*. Globally unique identifier. This value will be unique within the ScienceLogic platform. No two component devices can have the same GUID. This component identifier protects the integrity of a component device even if it is aligned with a different root device (for example, with a vmotion event).

NOTE: Each of these options can be defined by only one collection object in a Dynamic Application. If a component identifier is already defined by a collection object in this Dynamic Application, that component identifier will not be displayed in the list. For example, if you have already assigned a collection object to be the device name, the *Device Name (%N)* entry will not appear in the **Component Identifiers** list for other collection objects in the Dynamic Application.

6. Select the **[Save]** button to save your changes.

Duplicate MAC Addresses for Component Devices

The ScienceLogic platform handles duplicate MACs for component devices differently than duplicate MACs for physical devices. When a component device is assigned a MAC address, the ScienceLogic platform does not enforce uniqueness and will allow a component device to be created with the same MAC address as existing physical devices and/or existing component devices.

Unlike how the ScienceLogic platform discovers physical devices, the platform uses Dynamic Applications to retrieve data from a management device and "discover" each entity managed by that management device as a component device. The ScienceLogic platform then uses that retrieved data to create a device for each managed entity. In some cases, the managed entities are nested. In the ScienceLogic platform, physical devices are identified by IP address and MAC address. In the ScienceLogic platform, component devices are identified by a device name, a unique identifier, and a device class. A Dynamic Application that creates component devices can assign a MAC address to each component device, but is not required to.

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

For example, in a Cisco UCS system, the platform might discover a physical server that hosts the UCS manager. This physical server is the **root device**. The ScienceLogic platform might discover a chassis on the root device. The chassis is a **component device**. The chassis is a child device to the physical server. The platform might also discover a blade as a component device that is part of the chassis. The blade is a child device to the chassis. The chassis is the **parent device**.

The ScienceLogic platform does not automatically combine new component devices with any existing device record using the MAC address of the new component device. A component device can be combined with an existing device record under the following conditions:

- Dynamic Applications that create component devices can assign a globally unique identifier (GUID) to each component device. When the ScienceLogic platform performs collection for a Dynamic Application, and the Dynamic Application includes a collection object with a GUID component identifier, the platform compares the collected values for that collection object with all GUID values for all component devices discovered in the system. If a newly collected value matches a GUID value for an existing component device, the device from which the ScienceLogic platform collected the new value will become the parent of the existing component device. The existing component device will no longer be associated with its previous parent device. No new component device will be created.

- You can merge a physical device and a component device. You can do this in the **[Actions]** menu in the **Device Properties** page (Registry > Devices > wrench icon) for either the physical device or the component device. When you merge a physical device and a component device, the device record for the component device is no longer displayed in the user interface; the device record for the physical device is displayed in user interface pages that previously displayed the component device. For example, the physical device is displayed in lieu of the component device in the **Device Components** page and the **Component Map** page. All existing and future data for both devices will be associated with the physical device. You can unmerge a component device from a physical device in the **[Actions]** menu in the **Device Properties** page for the physical device (Registry > Devices > wrench icon).

Component Devices, Collector Groups, and Load Balancing

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

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

If any of the following circumstances occur, the platform will redistribute all the devices in the Collector Group:

- A new Data Collector is added to a Collector Group.
- Failover or failback occurs within a Collector Group (if failover is enabled).
- A user clicks the lightning bolt icon (⚡) for a Collector Group to manually force redistribution.

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

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

Collector Affinity

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

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

- **Root Device Collector.** The Data Collector assigned to the root device will collect data for the Dynamic Application. This guarantees that Dynamic Applications for an entire DCM tree will be collected by a single Data Collector. You might select this option if:
 - The Dynamic Application has a cache dependency with one or more other Dynamic Applications.
 - You are unable to collect data for devices and Dynamic Applications within the same Device Component Map on multiple Data Collectors in a collector group.
 - If the Dynamic Application will consume cache produced by a Dynamic Application aligned to a non-root device (for instance, a cluster device).

- **Assigned Collector.** The Dynamic Application will use the Data Collector assigned to the device running the Dynamic Application. This allows Dynamic Applications that are auto-aligned to component devices during discovery to run on multiple Data Collectors. You might select this option if:
 - The Dynamic Application has no cache dependencies with any other Dynamic Applications.
 - You want the Dynamic Application to be able to make parallel data requests across multiple Data Collectors in a collector group.
 - The Dynamic Application can be aligned using mechanisms other than auto-alignment during discovery (for instance, manual alignment or alignment via Device Class Templates or Run Book Actions).

Failover for Collector Groups for Component Devices

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

Merging Component Devices

The ScienceLogic platform does not automatically combine new component devices with any existing device record. A component device can be combined with an existing device record under the following conditions:

- Dynamic Applications that create component devices can assign a globally unique identifier (GUID) to each component device. When the ScienceLogic platform performs collection for a Dynamic Application, and the Dynamic Application includes a collection object with a GUID component identifier, the platform compares the collected values for that collection object with all GUID values for all component devices discovered in the system. If a newly collected value matches a GUID value for an existing component device, the device from which the platform collected the new value will become the parent of the existing component device. The existing component device will no longer be associated with its previous parent device. No new component device will be created.

- You can merge a physical device and a component device. You can do this two ways:
 - From the **[Actions]** menu in the **Device Properties** page (Registry > Devices > wrench icon) for either the physical device or the component device.
 - From the **[Actions]** menu in the **Device Manager** page (Registry > Devices > Device Manager), select *Merge Devices* to merge devices in bulk.

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

If you manually merge a component device with a physical device, the ScienceLogic platform allows data for the merged component device and data from the physical device to be collected on different Data Collectors. Data that was aligned with the component device can be collected on the Data Collector for its root device. If necessary, data aligned with the physical device can be collected on a different Data Collector.

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

- You can unmerge a component device from a physical device. You can do this in two ways:
 - From the **[Actions]** menu in the **Device Properties** page (Registry > Devices > wrench icon) for either the physical device or the component device.
 - From the **[Actions]** menu in the **Device Manager** page (Registry > Devices > Device Manager), select *Unmerge Devices* to unmerge devices in bulk.

Availability for Component Devices

For non-component devices, the ScienceLogic platform checks availability every 5 minutes by performing a ping request, SNMP request, or TCP connection to the admin primary IP address for the device. The result of the availability check is either *available* or *unavailable*. The ScienceLogic platform uses availability data to generate events when a device is unavailable and to generate availability reports and graphs.

In the **Dynamic Applications | Collections Objects** page, you can configure a collection object to represent the availability of the component devices created by a Dynamic Application. The ScienceLogic platform processes and stores the availability values returned by the availability collection object in the same way that the platform processes and stores the values collected by the availability check for physical devices.

When you configure a collection object to represent the availability of component devices:

- When the ScienceLogic platform performs collection for this Dynamic Application, if the platform can collect a value for a component device using the collection object and the value is not 0 (zero) or "false", the component device is considered "available".

- When the ScienceLogic platform performs collection for this Dynamic Application, if the platform cannot collect a value for a component device using the collection object or the platform collects a value that is 0 (zero) or "false", the component device is considered "unavailable".
- If the collection objects aligned with the *Device Name* and *Unique Identifier* component identifiers return lists of values, the ScienceLogic platform will create multiple component devices. Each component device will be associated with an index, i.e. a location in the list of values. To monitor the availability of all the component devices in the list, the collection object aligned with the *Availability* component identifier should return a list of values with a value at each index associated with a component device. The ScienceLogic platform will consider a component device "unavailable" when the list of values returned by the collection object aligned with the *Availability* component identifier does not include a value, or returns a value of 0 (zero) or false, at the index for that component device. For more information about Dynamic Application indexing, see the [Indexing](#) chapter.
- If you align a collection object with the *Availability* component identifier, the ScienceLogic platform will create a system availability graph in the **Device Performance** page of each component device created using the Dynamic Application.
- If you align a collection object with the *Availability* component identifier and the ScienceLogic platform determines that a component device is unavailable, the platform will generate an event and supply the value "Unavailable" in the **Collection State** column in the **Device Components** page.

The following rules apply to the availability state for component devices:

- Component devices can use a Component Identifier to monitor availability. However, in a tree of component devices, some component devices might have a component identifier for availability and others might not. For example, suppose a component device has a component identifier for availability, and the ScienceLogic platform considers that component device "unavailable". All the descendents of that component device that do not have their own component identifier for availability will be considered unavailable. As soon as the ScienceLogic platform finds a descendent with its own component identifier for availability, the platform stops checking that descendent and its descendents for availability. Component devices without their own component identifier for availability inherit their availability from their nearest ancestor that has a component identifier for availability.
- For trees that include merged devices, i.e. those that include both hardware devices and component devices, the ScienceLogic platform skips over the hardware devices and allows them to use a network-based protocol to determine availability. For example, suppose you have a tree like this:
 - Grandparent device is a component device with a component identifier for availability. The platform has determined that the grandparent device is unavailable.
 - Child device is a hardware device that uses ICMP and ping to determine availability. When the platform evaluates the grandparent's component identifier, the platform skips over this device. ICMP and ping determine the availability of this device.
 - Grandchild device is a component device that does not have its own component identifier for availability. When the platform evaluates the grandparent's component identifier, the platform assigns the grandparent's availability state to this grandchild device.
- If all the hosts in a cluster are powered off or unavailable, both the hardware-based hosts and the associated component devices will have an availability value of "Unavailable". When at least one host in the cluster becomes available, some or all of the associated component devices will also become available.

For instructions on how to align a collection object to a component identifier, see the [Configuring Collection Objects as Component Identifiers](#) section.

Moving Component Devices Between Root Devices

Depending on the type of management system that you want to monitor using Dynamic Component Mapping, it might be possible for a component device to move from one root device to another root device. For example, if a Dynamic Application with Dynamic Component Mapping configured is aligned with multiple VMware ESX servers that are root devices with guest VMs discovered as component devices, a vMotion event might move a guest VM from one VMware ESX server to another VMware ESX server.

If your Dynamic Applications will create component devices that can move from one root device to another root device, you can configure a collection object that maps to a globally unique identifier (GUID) for each component device. The value of the GUID component identifier must be unique for every component device in the ScienceLogic platform.


When the ScienceLogic platform performs collection for a Dynamic Application, and the Dynamic Application includes a collection object with a GUID component identifier, the platform compares the collected values for that collection object with all GUID values for all component devices discovered in the system. If a newly collected value matches a GUID value for an existing component device, the device from which the ScienceLogic platform collected the new value will become the parent of the existing component device. The existing component device will no longer be associated with its previous parent device.

Automatically Aligning Dynamic Applications with Component Devices

In a Dynamic Application that has Dynamic Component Mapping enabled, you can define a list of Dynamic Applications that the ScienceLogic platform will automatically align to the component devices that are discovered by the Dynamic Application. The Dynamic Applications in the list are automatically aligned to a component device when the component device is created. Every time collection is performed for the Dynamic Application that created a component device, any new Dynamic Applications that have been added to the list of Dynamic Applications are aligned to the component device.

This feature is useful when you want to nest component devices, that is, discover additional component devices that have a component device as the parent. To nest component devices, you must specify that the ScienceLogic platform should immediately align Dynamic Applications that discover additional component devices upon creation of a component device.

To define a list of Dynamic Applications that you want the ScienceLogic platform to automatically align to component devices:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application that has Dynamic Component Mapping enabled. Select its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.

3. Select the **[Component]** tab. The **Dynamic Application Component Alignment** page is displayed.

NOTE: The **[Component]** tab and the **Dynamic Application Component Alignment** page do not appear for Dynamic Applications that do not have component mapping enabled in the **Dynamic Applications Properties Editor** page.

4. The **Dynamic Application Component Alignment** page displays two lists of Dynamic Applications:
 - **Unaligned Dynamic Apps.** Displays a list of all Dynamic Applications in the ScienceLogic platform that are not automatically aligned to the component devices discovered by this Dynamic Application.
 - **Aligned Dynamic Apps.** Displays a list of Dynamic Applications that the ScienceLogic platform will automatically align to the component devices discovered by this Dynamic Application.

To move a Dynamic Application from the **Unaligned Dynamic Apps** list to the **Aligned Dynamic Apps** list, select the Dynamic Application and then select the right-arrow button.

To move a Dynamic Application from the **Aligned Dynamic Apps** list to the **Unaligned Dynamic Apps** list, select the Dynamic Application and then select the left-arrow button.

5. Select the **[Save]** button to save your changes.

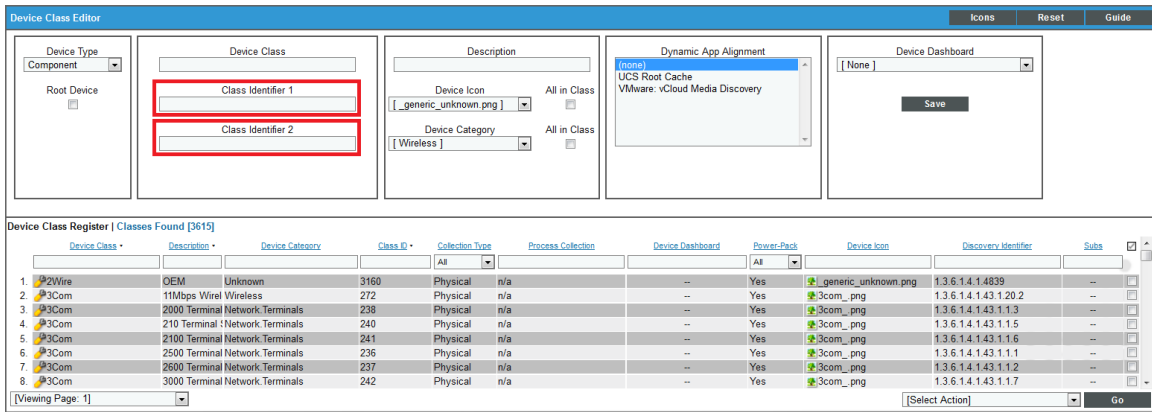
Automatically Assigning a Device Class to Each Component Device

The ScienceLogic platform can automatically align a device class with each discovered component device. If you do not want to automatically assign all component devices to the same device class, you can use Component Identifiers (variables) to align component devices with device classes. The value of two collection objects will determine the device class for each component device.

To customize the alignment between device classes and component devices, perform the steps in the following two section.

Configuring a Device Class to Match Collected Values

1. Go to the **Device Class Editor** page (System > Customize > Device Classes).



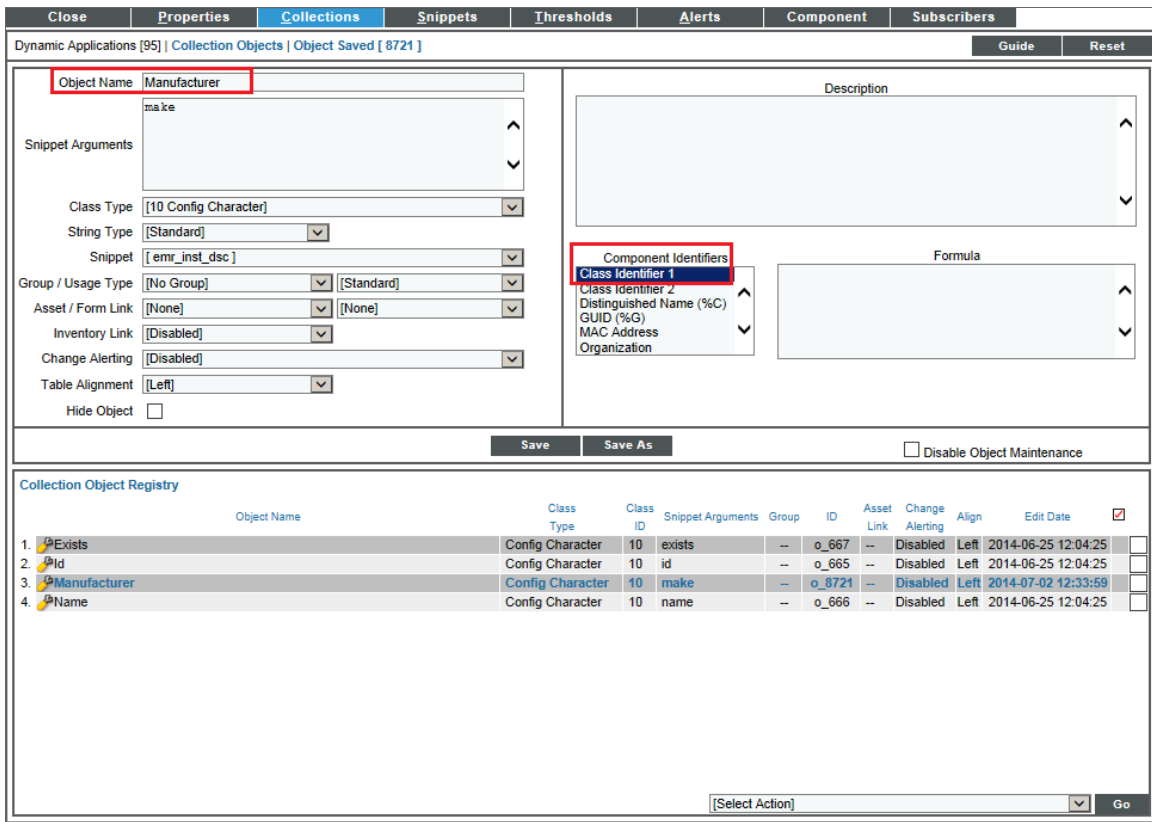
2. When you create or edit a device class for component devices, you must define these three values:

- **Device Type.** Select "Component".
- **Class Identifier 1.** Enter a value collected by one of the collection objects in the Dynamic Application that creates the component devices. Typically, this value is the manufacturer of the device.
- **Class Identifier 2.** Enter a value collected by one of the collection objects in the Dynamic Application that creates the component devices. Typically, this value is the model of the device. If you want to assign this device class based only on the **Class Identifier 1**, e.g. if you are creating a "fall back" device class for components that do not match specific **Class Identifier 1** and **Class Identifier 2** combinations, do not enter a value in this field.

NOTE: The combination of the **Class Identifier 1** and **Class Identifier 2** values should be unique to devices that subscribe to the device class.

Configuring Component Identifiers for the Collection Objects

When you create a Dynamic Application, you can [map the value of a collection object to a component identifier](#). The value collected for the collection object (from each device) will populate the component identifier (for that device).



Two of these component identifiers help the platform automatically align a device class with each component device.

- **Class Identifier 1.** For a component device to be assigned to a device class based on component identifiers, the value(s) collected by this collection object must match the value you entered in the **Class Identifier 1** field for that device class.
- **Class Identifier 2.** For a component device to be assigned to a device class based on component identifiers, either the value(s) collected by this collection object must match the value you entered in the **Class Identifier 2** field for that device class or the device class must not include a value for the **Class Identifier 2** field.

How Does the ScienceLogic Platform Use Component Identifiers to Automatically Assign a Device Class to Each Component Device?

The ScienceLogic platform uses the following procedure to automatically assign a device class to each new component device:

1. If the values for component identifiers **Class Identifier 1** and **Class Identifier 2** exactly match the values specified in the device class fields **Class Identifier 1** and **Class Identifier 2** in an existing component device class, the platform will align the component device with that device class. During subsequent collections, if the value of the one or both of the collection objects changes and exactly matches a device class name and device class description, the platform will change the device class to match the values of **Class Identifier 1** and **Class Identifier 2**.
2. If the value for component identifier **Class Identifier 1** exactly matches the device class field **Class Identifier 1** in a component device class definition, but the component identifier **Class Identifier 2** either does not match a device class description or is not defined, the platform will assign the component device to the device class with the matching value in the field **Class Identifier 1** and no value defined for **Class Identifier 2**.

NOTE: ScienceLogic recommends that you create such a device class as a "fallback" for the specified **Class Identifier 1**. The platform can then use this device class for any component devices that have an unrecognized **Class Identifier 2**.

3. If the Dynamic Application does not include component identifiers **Class Identifier 1** and **Class Identifier 2** OR the values for **Class Identifier 1** and **Class Identifier 2** do not match the fields in an existing component device class, the platform aligns the component device with the device class associated with the Dynamic Application. You can define this device class in the **Device Class Editor** (System > Customize > Device Classes > edit a device class > Dynamic App Alignment field) or in the Dynamic Application **Editor** window, in the **Components** tab (System > Manage > Applications > edit a Dynamic Application > Components tab).
4. If the Dynamic Application does not include component identifiers **Class Identifier 1** and **Class Identifier 2** OR the values for **Class Identifier 1** and **Class Identifier 2** do not match the fields in an existing device class, AND the Dynamic Application is not aligned with a device class (either in the **Device Class Editor** page or in the **Components** tab), the platform automatically assigns the component device to the device class "Generic: Component".

Aligning a Default Device Class with a Dynamic Application

You can align a device class with a Dynamic Application that creates component devices. If you choose not to use the component identifiers **Class Identifier 1** and **Class Identifier 2** in your Dynamic Application to match a device class, the ScienceLogic platform will automatically align each newly discovered component device with the default device class that you aligned with the Dynamic Application.

There are two ways to align a default device class with a Dynamic Application:

- In the **Device Class Editor** page (System > Customize > Device Classes), in the **Dynamic app Alignment** field.
- In the **Dynamic Application Editor** window, in the **Components** tab (System > Manage > Applications > Edit a Dynamic Application > Components tab).

Defining a Default Device Class in the Device Class Editor

To define a default device class for a Dynamic Application from the **Device Class Editor** page:

1. Go to the **Device Class Editor** page (System > Customize > Device Classes).
2. Create or edit the device class you want to align with the Dynamic Application.

The screenshot shows the 'Device Class Editor' window. It has several sections: 'Device Type' (Component), 'Device Class' (empty), 'Description' (empty), 'Device Icon' (generic_unknown.png), 'Device Category' (Wireless), and 'Dynamic App Alignment' (highlighted with a red box, showing options: (None), UCS Root Cache, VMware vCloud Media Discovery). There is also a 'Device Dashboard' dropdown set to '[None]' and a 'Save' button.

Below the editor is a table titled 'Device Class Register | Classes Found [3615]':

Device Class	Description	Device Category	Class ID	Collection Type	Process Collection	Device Dashboard	Power-Pack	Device Icon	Discovery Identifier	Subs
1. P2Wire	OEM	Unknown	3160	Physical	n/a	--	Yes	generic_unknown.png	1.3.6.1.4.1.4839	--
2. P3Com	11Mbps Wirel	Wireless	272	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.20.2	--
3. P3Com	2000 Terminal Network	Terminals	238	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.3	--
4. P3Com	210 Terminal Network	Terminals	240	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.5	--
5. P3Com	2100 Terminal Network	Terminals	241	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.6	--
6. P3Com	2500 Terminal Network	Terminals	236	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.1	--
7. P3Com	2600 Terminal Network	Terminals	237	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.2	--
8. P3Com	3000 Terminal Network	Terminals	242	Physical	n/a	--	Yes	3com_.png	1.3.6.1.4.1.43.1.1.7	--

3. In the field **Dynamic App Alignment**, select a Dynamic Application.
4. Select the **[Save]** button.

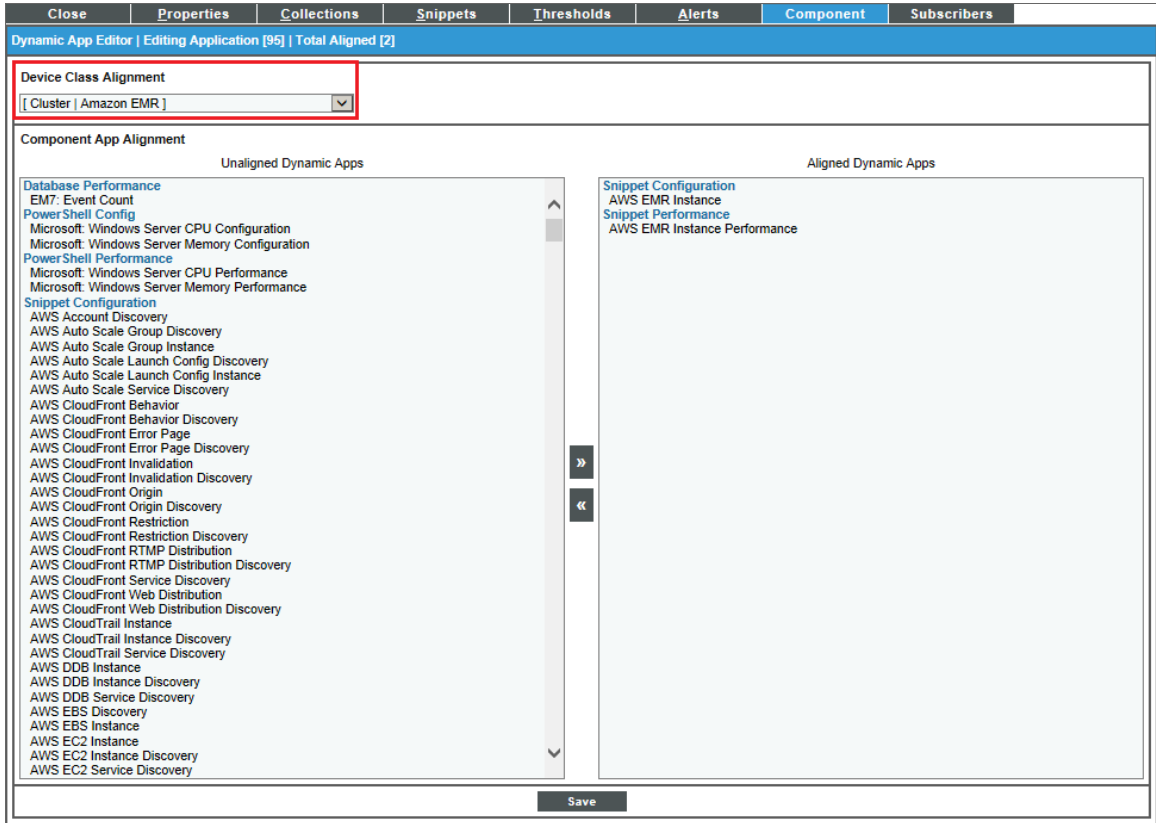
Defining a Default Device Class in the Dynamic Application

To define a default device class for a Dynamic Application from the **Dynamic Application Editor** window:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application that you want to associate a device class with. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.

3. Select the **[Component]** tab. The **Dynamic Application Component Alignment** page is displayed.

NOTE: The **[Component]** tab does not appear for Dynamic Applications that do not have dynamic component mapping enabled.



4. In the **Device Class Alignment** field, select a device class.
5. Select the **[Save]** button.

Variables in Dynamic Applications for Component Devices

If Dynamic Applications are aligned to component devices, the ScienceLogic platform uses the following rules when performing collection:

- If the %D substitution is used in the credential aligned with the Dynamic Application, the ScienceLogic platform substitutes the IP address of the root device.
- If the Dynamic Application uses cached responses from other Dynamic Applications, the Dynamic Application that caches the response can be aligned to either the component device or the root device.

In SOAP and XSLT Dynamic Applications that are aligned to component devices, you can include the following substitution characters in the **SOAP Request Code**, **XSLT Request Code**, and **XSLT Parser Code** fields:

- **%N**. The platform will substitute the value of the *Device Name* component identifier for the component device for which collection is being performed.
- **%U**. The platform will substitute the value of the *Unique Identifier* component identifier for the component device for which collection is being performed.
- **%C**. The platform will substitute the value of the *Distinguished Name* component identifier for the component device for which collection is being performed.
- **%G**. The platform will substitute the value of the *GUID* component identifier for the component device for which collection is being performed.

Caching and Component Mapping

Although the caching and dynamic component mapping features can be used separately, the caching feature is useful to collect performance and configuration data for component devices. For example, the following set of Dynamic Applications might be used to discover and monitor component devices:

- A Dynamic Application that requests all the information needed for all other Dynamic Applications in this set. This Dynamic Application caches the responses and is aligned to the root device.
- A Dynamic Application that creates component devices. This Dynamic Application consumes cache and is aligned to the root device.
- One or more Dynamic Applications that collect performance and/or configuration data for each component device. These Dynamic Applications consume the cached responses from the first Dynamic Application and are aligned to each component device.

To associate performance or configuration data with only the appropriate component device (that is, the component device that generated the performance or configuration data), you would use substitution characters in your requests. When the request is executed, the ScienceLogic platform replaces the substitution character with a component identifier for the component device for which collection is currently being performed.

Dynamic Application Relationships

Overview

Dynamic Applications can be configured to **automatically create relationships between devices**. For example, the Dynamic Applications in the VMware vSphere and NetApp PowerPacks are configured to create relationships between VMware Datastore component devices and their associated NetApp Volume component devices. Relationships created by Dynamic Applications are used and visualized by the platform in the same manner as relationships created by topology collection, Dynamic Component Mapping, and manually in the user interface.

There are three methods for configuring Dynamic Applications to automatically create relationships:

- A single configuration Dynamic Application that collects the unique identifier (UID) for a component device. When the ScienceLogic platform polls a subscriber device and collects a UID value, and that UID value matches the UID for an existing component device in the same component tree as the subscriber device, the ScienceLogic platform creates a relationship between the subscriber device and the component device.
- A single configuration Dynamic Application that collects the globally unique identifier (GUID) for a component device, typically a component device that is in a different component tree than the subscriber device. When the ScienceLogic platform polls a subscriber device and collects a GUID value, and that GUID matches the GUID of an existing component device that is in a different component tree than the subscriber device, the ScienceLogic platform creates a relationship between the subscriber device and the component device.
- A pair of configuration Dynamic Applications collect the same data. One of these Dynamic Applications is configured to define a unique identity for subscriber devices and also references the relationship. The other Dynamic Application defines the relationship and also references the unique identity. If each Dynamic Application collects the same value for unique identity and the same value for relationship, a relationship is created between the two subscriber devices. This type of relationship is described in the section [Configuring Identity-based Relationships](#).

Configuring Collection Objects for Relationships

To configure a collection object for the creation of relationships:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the configuration Dynamic Application that you want to configure. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the [Collections] tab. The **Dynamic Applications | Collections Objects** page is displayed.

Dynamic Applications [992] | Collection Objects

Object Name:

Snippet Arguments:

Class Type: [10] Config Character

String Type: [Standard]

Snippet: snippet out

Group / Usage Type: [Standard] (dropdown menu open showing: [Standard], Group Index, Identity Namespace, Relationship Namespace, GUID Relationship, Unique Identifier Relationship, GUID + Index, Unique Identifier + Index)

Asset / Form Link: [None]

Inventory Link: [Disabled]

Change Alerting: [Disabled]

Table Alignment: [Left]

Hide Object:

Save Disable Object Maintenance

Collection Object Registry

Object Name	Class Type	Class ID	Snippet Arguments	Group	ID	Asset Link	Change Alerting	Align	Edit Date	
1. Test	Config Character	10		--	o_9497	--	Disabled	Left	2014-06-10 08:51:29	<input type="checkbox"/>

[Select Action]

4. Select the wrench icon (🔧) for the collection object you want to edit.
5. The **Usage Type** field specifies how this collection object is used to build relationships. The options in this field are:
 - *Standard*. The collection object is not used to create relationships or designated as an index for the specified collection object group.

- *Group Index*. The collection object is configured as the index for the selected collection object **Group**. For more information about indexing, see the [Indexing](#) chapter. Additionally, if other collection objects in the same group are configured with a **Usage Type** of *Identity Namespace* or *Relationship Namespace*, the ScienceLogic platform uses the collection object with a **Usage Type** of *Group Index* to define the unique identifier and create relationships between the subscribers of two related Dynamic Applications. See the [Configuring Identity-based Relationships](#) section for a description of how to implement this type of relationship.
- *Identity Namespace*. The ScienceLogic platform uses the collection object to create relationships between the subscribers of two related Dynamic Applications. This collection object defines the identity portion of the relationship. If the ScienceLogic platform collects this collection object from a device, the ScienceLogic platform will make that device the child in a parent-child relationship. See the [Configuring Identity-based Relationships](#) section for a description of how to implement this type of relationship.
- *Relationship Namespace*. The ScienceLogic platform uses this collection object to create relationships between the subscribers of two related Dynamic Applications. If the ScienceLogic platform collects this collection object from a device, the ScienceLogic platform will make that device the parent in a parent-child relationship. See the [Configuring Identity-based Relationships](#) section for a description of how to implement this type of relationship.

NOTE: **Usage Type** of *Identity Namespace* defines the **child** portion of an identity-based relationships. **Usage Type** of *Relationship Namespace* defines the **parent** portion of an identity-based relationships.

- *GUID Relationship*. When the ScienceLogic platform collects a value for this collection object, and that value matches the GUID for a component device in the system, the ScienceLogic platform creates a relationship between the subscriber device and the component device. The subscriber device will be the parent device and the component device will be the child device. The relationship will be labeled with the name of this collection object. For information about globally unique identifiers for component devices,
- *Unique Identifier Relationship*. When the ScienceLogic platform collects a value for this collection object, and that value matches the unique identifier for a component device in the same component tree as the subscriber device, the ScienceLogic platform creates a relationship between the subscriber device and the component device. The subscriber device will be the parent device and the component device will be the child device. The relationship will be labeled with the name of this collection object.
- *GUID Relationship + Index*. The ScienceLogic platform uses this collection object to create relationships in the same manner as the *GUID Relationship* option. When the ScienceLogic platform collects a value for this collection object, and that value matches the GUID for a component device in the system, the ScienceLogic platform creates a relationship between the subscriber device and the component device. The subscriber device will be the parent device and the component device will be the child device. The relationship will be labeled with the name of this collection object. Additionally, the collection object is configured as the index for the specified collection object group.

- *Unique Identifier + Index*. The ScienceLogic platform uses this collection object to create relationships in the same manner as the Unique Identifier Relationship option. When the ScienceLogic platform collects a value for this collection object, and that value matches the unique identifier for a component device in the same component tree as the subscriber device, the ScienceLogic platform creates a relationship between the subscriber device and the component device. The subscriber device will be the parent device and the component device will be the child device. The relationship will be labeled with the name of this collection object. Additionally, the collection object is configured as the index for the specified collection object group.

4. Click **[Save]**.

Configuring Identity-Based Relationships

An identity-based relationship is created using a pair of Dynamic Applications:


- A Dynamic Application that defines a unique identity for subscriber devices. This Dynamic Application defines:
 - a unique ID for the device, using a collection object that usually has a **Usage Type** of *Group Index*.
 - the identity side of the relationship, using a collection object that usually has a **Usage Type** of *Identity Namespace*. The devices that subscribe to this Dynamic Application will be children in the parent-child relationship that is created using the unique identity.

NOTE: This identifier must be unique to this relationship, meaning no two subscribers of the same Dynamic Applications should collect the same identifier for the same relationship. However, both Dynamic Applications in the pair must collect the same identifier value for a relationship to be created.

- A Dynamic Application that defines the relationship for subscriber devices. This Dynamic Application defines:
 - a unique ID for the device, using a collection object that usually has a **Usage Type** of *Group Index*.
 - a relationship, using a collection object that usually has a **Usage Type** of *Relationship Namespace*. The devices that subscribe to this Dynamic Application will be parents in the parent-child relationship that is created using the unique identity.

A relationship is created when:

- both Dynamic Applications in a pair collect identical values for the unique identifier (the collection object with a **Usage Type** of *Group Index*).
- the value for the collection object that has a **Usage Type** of *Identity Namespace* is identical to the value for the collection object that has a **Usage Type** of *Relationship Namespace*.



To configure a pair of Dynamic Applications to create identity-based relationships, go to the Dynamic Applications Collection Objects page:

- In both Dynamic Applications, select the collection object that collects the identifier for the relationship. In the **Usage Type** field, select *Group Index*. Select a **Group** for the collection object. The Group should be the same in both Dynamic Applications.
- In the Dynamic Application that defines a unique identity for subscriber devices, select the collection object that collects the namespace for the relationship. In the **Usage Type** field, select *Identity Namespace*. For this collection object, select the same **Group** as the *Group Index* collection object.
- In the Dynamic Application that defines the relationship, select the collection object that collects the namespace for the relationship. In the **Usage Type** field, select one of the values that includes the text "Relationship". For this collection object, select the same **Group** as the *Group Index* collection object.

NOTE: The *Identity Namespace* and *Relationship Namespace* objects must return a value for each of the indexes returned by the *Group Index* object. Typically, the values at each index returned by the *Identity Namespace* and *Relationship Namespace* objects are equal.

Viewing a Map of Component Devices

The **Component Map** page allows you to view devices by root node and also view the relationships between root nodes, parent components, and child components. This map makes it easy to visualize and manage root nodes and their components.

NOTE: User accounts of type "user" can view only root nodes and device components that belong to their organization(s).

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

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

For example, in a Cisco UCS system, the ScienceLogic platform might discover a physical server that hosts the UCS manager. The platform might discover a chassis as a component device. The chassis is a child device to the physical server; the physical server is the root device. The platform might also discover a blade as a component device that is part of the chassis. The blade is a child device to the chassis. The chassis is the parent device.

The ScienceLogic platform automatically updates the Component Map as new component devices are discovered. The platform also updates each map with the latest status and event information.

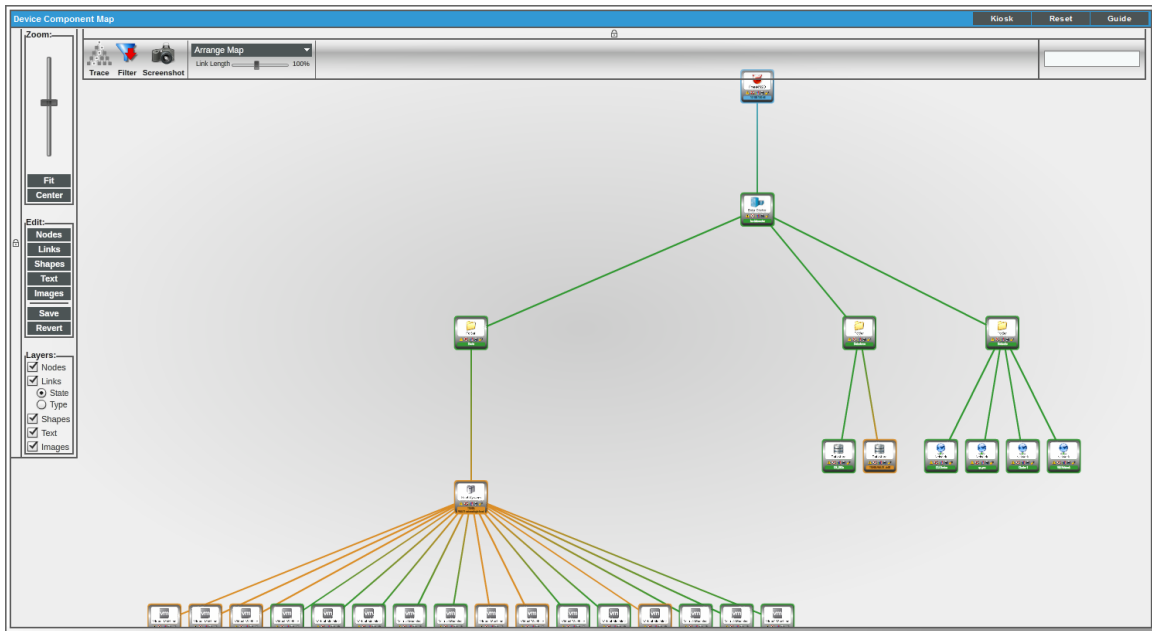
NOTE: For a user of type *User*, you can view only root nodes and device components that belong to your organization(s).

Viewing a Component Map

You can view a component map from the **Component Map** page. To view a Component Map:

1. Go to the **[Views]** tab and then go to the left NavBar.


- Expand the links for Device Maps > Components and then select a root node. The **Device Component Map** page appears and displays the map of the root node and its component devices:



Adding Devices from Another Component Tree to the Component Map

You can add a device from another component tree to the map of the current component device.

For example, suppose you have a UCS system. Suppose you are running an ESX server and VMs on one of the blades of the UCS system:



You can add the device where the ESX server resides to the map of the UCS system. To do this:

1. Go to the **Component Map** page (Views > Device Maps > Components).
2. Expand the list of maps. Find the component map to which you want to add a device.
3. Select the **[Nodes]** button. Select the **Add Dev** icon (in the upper left).
4. In the **Device Browser** modal page, select the root node you want to add to the component map. Select the **[Add To Map]** button.

NOTE: When you complete these steps, the ScienceLogic system will automatically add all the child devices for the newly added node to the map.

5. Select the **[Links]** button. Click on the parent device and then click on the newly added child device.
6. The link you created is an **Event Correlation Override** link. If you want to manually define a parent device and child device for two devices that do not share a Layer-2 link, a Layer-3 link, a CDP link, an LLDP link, or a VMware relationship, you can create an **Event Correlation Override** link between the devices. Additionally, you can then define an event hierarchy for these devices.

7. Select the **[Save]** button.
8. Select the **[Reset]** button to view the changes to the map.

Viewing Relationships in the Organization View

The **Organizational Map** page allows you to view devices by organization. This makes it easy to visualize and manage devices in organizations. If devices in the organization include relationships created by Dynamic Applications, these relationships are displayed in the map. If devices in this organization include CDP, LLDP, layer-2, or layer-3 relationships, they are included in the map.

All elements, policies, events, tickets, and users in the ScienceLogic platform are associated with an organization. An organization is a group for managing elements and user accounts. You can define organizations by geographic area, departments, types of devices, or any structure that works best for your needs. The minimum characteristics of an organization consist of:

- A unique name
- Users who are members of the organization
- Elements, such as devices, associated with the organization

The ScienceLogic platform automatically updates the Organizational Maps as new devices are added to organizations and as new organizations are defined. The ScienceLogic platform also updates each map with the latest status and event information.


To select an Organizational Map, go to the **[Views]** tab and then go to the left NavBar. Expand the links for Device Maps > Organizational and then select an organization. The **Organizational Map** page appears and displays a map of the selected organization and its devices.

Viewing an Organizational Map

You can view an organizational map from the **Organizational Map** page. To view an organizational map:

1. Go to the **[Views]** tab and then go to the left NavBar.

- Expand the links for Device Maps > Organizational, and then select an organization. The **Organizational Map** page appears and displays a map of the organization and its member devices:



- The blank field in the upper right allows you to find one or more devices in the map. You can enter a string, and the ScienceLogic platform will highlight only the devices that have a device name that matches the string.
 - Each member device appears as an icon in the map.
 - The Organizational Map page for an organization also displays the relationships between devices in the organization. This includes:
 - Layer-2 devices and their clients
 - Layer-3 devices and layer-2 devices
 - Component devices and their parents, e.g. virtual machines and their hypervisors
 - Network devices that use CDP (Cisco Discovery Protocol) and devices that are specified as neighbors in the CDP tables
 - Network devices that use LLDP (Link Layer Discovery Protocol) and devices that are specified as neighbors in the CDP tables
 - Device relationships created with Dynamic Applications
 - Manually created parent-child relationships that affect event correlation
- When the map appears, you can view and reposition the components. The map can be edited and rearranged using drag-and-drop features. Devices can be repositioned for easier reading, if necessary.
 - Mousing over a device displays its name, IP address, device type, and device category.

Viewing Relationships in Customized Maps

A customized map allows you to view the devices and links that are most important to you.

When you create a customized map, you are also creating a new device group (which appears in the **Device Groups** page). You can add devices and other sub-device groups to the new map, just as you would to a standard device group.

If the ScienceLogic platform has information about the relationships between the devices in a customized map, the platform automatically includes the appropriate links in the customized map. Customized maps display every type of relationship data, which includes:

- Layer-2 devices and their clients
- Layer-3 devices and layer-2 devices
- Component devices and their parents, e.g. virtual machines and their hypervisors
- Network devices that use CDP (Cisco Discovery Protocol) and devices that are specified as neighbors in the CDP tables
- Network devices that use LLDP (Link Layer Discovery Protocol) and devices that are specified as neighbors in the LLDP tables
- Device relationships created with Dynamic Applications
- Manually created parent-child relationships that affect event correlation

Customized maps appear in the following sections in the left NavBar:

- **Views > My Customized Maps**. Personalized maps that you create.
- **Views > User Customized Maps**. If you are a user of type "administrator", you can navigate to the maps in this section to view and edit all customized maps in the ScienceLogic platform, even if the device group associated with the map was defined with the field **Shared (visible to all users)** set to *no*.
- **Views > Shared Customized Maps**. If a customized map or device group is defined as "shared", you can view the maps in this section. The maps in **Shared Customized Maps** require the same Access Hooks and Access Keys as device groups. Depending upon the Access Keys assigned to your account, you might be able to edit **Shared Customized Maps** created by other users. To learn more about Access Hooks and Access Keys, see the manual **Access Permissions**.

NOTE: If you create a device group from the **Device Groups** page and set the **Visibility** field to include *Maps/Views*, the device group will appear as a map in the **Custom Device Group Map** page (Views > My Customized Maps). If you set the **Shared** field to *yes*, the device group will appear for view by other users as a map in the **Shared Customized Maps** page (Views > Shared Customized Maps).

To learn more about Customize Maps, see the manual **Views**.

Viewing Relationships for a Single Device

You can view all links for a single device in the **Device Relationships** page, in the **Device Properties** panel. To view all links for a single device:

1. Go to the **Device Relationships** page (Registry > Networks > Device Relationships) and click the Device Properties icon (🔧) for the device you want to see relationships. If a link has been defined on a device, you can also go to the **Device Manager** page (Registry > Devices > Device Manager), click the wrench icon for a device (🔧) and click the **[Relationships]** tab in the **Device Properties** pane.

Child	Child IP	Child Interface	Child Phys Addr	Child If Manufacturer	Parent	Parent Interface	Parent If Alias	Parent Phys Addr	Parent If Manufacturer	Type
Technology Device 3	10.40.40.6	HP Network T 00:09:97:c0:42:99		NortelNetw	Technology Switch B	Fa0/10	--	00:06:b1:1b:9d:2c	Sonicwall	Layer 2
Technology Device 4	10.40.40.7	HP Network T 00:0e:08:ab:65:51		SapuraTech	Technology Switch B	Fa0/11	--	00:06:b1:1b:9d:2d	Sonicwall	Layer 2
Technology Switch B	--	Fa0/12	00:06:b1:1b:9d:2d	Sonicwall	Technology Switch A	Vt.1	--	00:50:00:01:3c:33	Hendberg/9	Layer 2

- The **Device Relationships** page appears. The left pane of the **Device Relationships** page displays links to parent devices. The right pane of the **Device Relationships** page displays links to child devices. For each relationship, the **Device Relationships** page displays the following information:

- **Type of relationship.** Possible values are:
 - *Layer 2.* Layer-2 devices and their clients.
 - *Layer 3.* Layer-3 devices and layer-2 devices.
 - *VMware.* Hypervisors and their virtual machines.
 - *CDP.* Network devices that use CDP (Cisco Discovery Protocol) and devices that are specified as neighbors in CDP tables.
 - *LLDP.* Network devices that use LLDP (Link Layer Discovery Protocol) and devices that are specified as neighbors in LLDP tables.
 - *Event Correlation.* Relationships defined manually by users through the user interface.
 - *Component Mapping.* Relationships defined using Dynamic Applications.
- **Child Interface.** Name of the interface through which the child device communicates with the parent device and a link to the **Interfaces Found** page for the child interface.
- **Parent Device.** The name of the parent device and a link to the **Device Properties** page for the parent device.

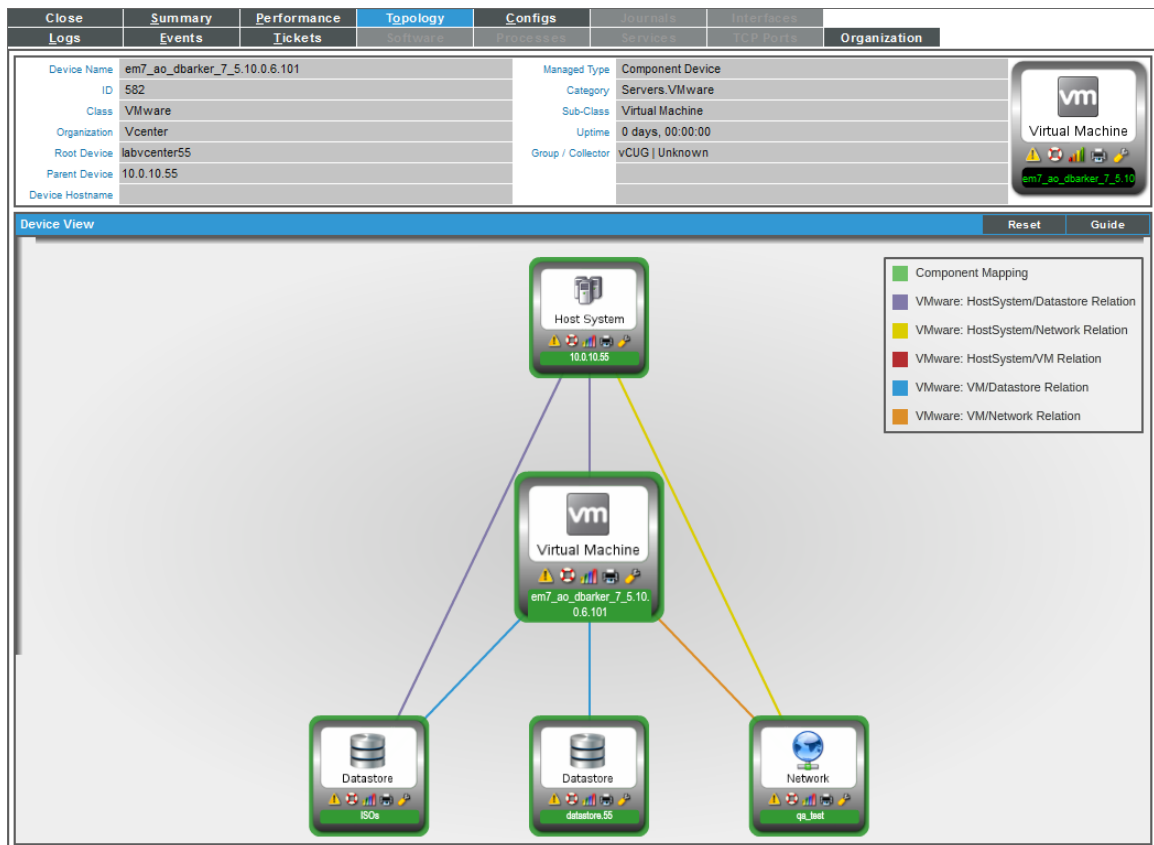
- **Parent Interface.** The name of the interface through which the parent device communicates with the child device and a link to the **Interfaces Found** page for the parent interface.

NOTE: Clicking on a device reloads the **Device Relationships** page and makes the selected device the primary device.

Device View Page

You can also view information about relationships for a single device in the Device View page.

The **Device View** page appears when a user clicks the **Topology** tab in the Device Reports panel. The **Device View** page displays a map of the device and all of the devices with which the device has relationships.



These relationships include:

- Layer-2 devices and their clients
- Layer-3 devices and Layer-2 devices
- Component devices and their parent devices. For example, virtual machines and their hypervisors and their virtual machines.

- Network devices that use CDP (Cisco Delivery Protocol) and devices that are specified as neighbors in CDP tables
- Links between network devices that use CDP (Cisco Discovery Protocol) and devices that are specified as neighbors in CDP tables
- Network devices that use LLDP (Link Layer Delivery Protocol) and devices that are specified as neighbors in LLDP tables
- Links between network devices that use LLDP (Link Layer Discovery Protocol) and devices that are specified as neighbors in LLDP tables
- Device relationships between root devices, parent devices, and component devices (Component Mapping)
- Device relationships created with Dynamic Applications
- Manually created parent-child relationships that affect event correlation

NOTE: Double-clicking on a device reloads the **Device View** page and makes the selected device the primary device.

For details on the toolbars that appear in this page, see the **Views** manual.

Presentation Objects for Performance Dynamic Applications

What is a Presentation Object?

Presentation Objects define how the ScienceLogic platform should present data for Dynamic Applications with an archetype of **Performance** or **Journal**. For performance Dynamic Applications, a presentation object defines how the ScienceLogic platform uses the collected values for the collection objects to generate graphs.


For example, a Dynamic Application that is designed to collect information about file system usage on a device might define:






- Two collection objects, one that collects the total size of each file system and one that collects the amount of space used on each file system.
- A presentation object that defines a graph. This graph displays the percentage of space used for each file system. The presentation object would include a formula that uses both collection objects.

NOTE: This chapter describes presentation objects in performance Dynamic Applications. This chapter does not apply to journal Dynamic Applications with an archetype of Journal. For information on presentation objects in journal Dynamic Applications, see the **Snippet Dynamic Application Development** manual.

Viewing the Presentation Objects in a Dynamic Application

To view the presentation objects in a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to view the presentation objects for. Click its wrench icon () . The **Dynamic Applications Properties Editor** page is displayed.
3. Click the **[Presentations]** tab. The **Dynamic Applications Presentation Objects** page is displayed. The **Presentation Object Registry** pane at the bottom of the page displays a list of defined reports.
4. For each report in a Dynamic Application of type Performance, the **Presentation Object Registry** displays the following:

	Report Name	State	Abbreviation Suffix	Group	Label	Precedence	Show as Percent	ID	Date Edit
1.	 ARP Entries	Enabled	--	--	--	--	No	pres_14	2014-06-11 14:11:01
2.	 Concurrent Connections Per Port	Enabled	o	--	--	--	No	pres_11	2014-06-11 14:11:01
3.	 Connection Rate Per Virtual Server	Enabled	--	--	--	--	No	pres_12	2014-06-11 14:11:01
4.	 Port Bindings	Enabled	--	--	--	--	No	pres_13	2014-06-11 14:11:01
5.	 Real Server Current Sessions	Enabled	--	--	--	--	No	pres_15	2014-06-11 14:11:01

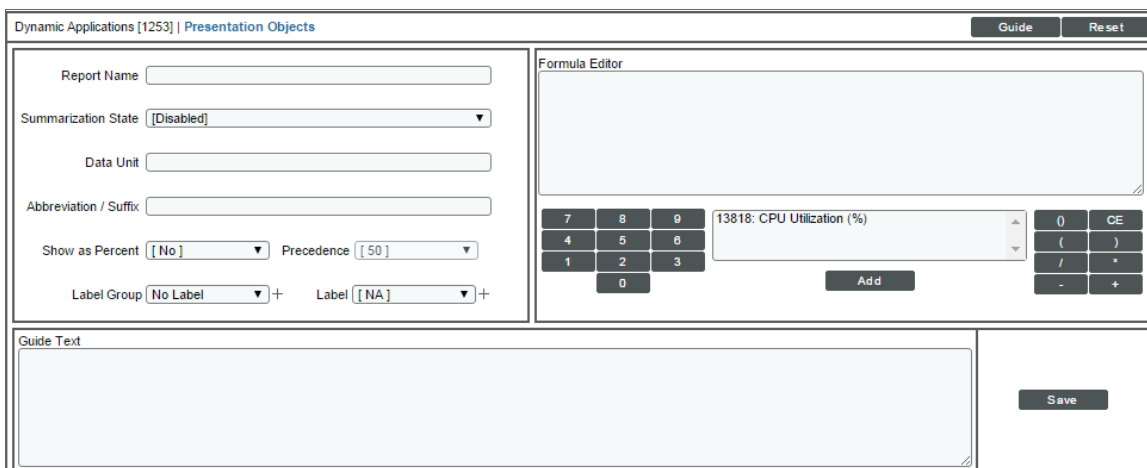
- **Report Name.** The name of the presentation object.
- **State.** The state of the presentation object. Possible values are:
 - *Enabled.* The platform will generate graphs for this presentation object.
 - *Disabled.* The platform will not generate graphs for this presentation object.
- **Abbreviation Suffix.** Abbreviation for the unit of measure to display in the report.
- **Label Group.** Name of the Collection Group aligned with the presentation object. For more information on Collection Labels, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Label.** Name of the Collection Label aligned with the presentation object. The **Vitals** group is available as a Collection Label and preserves the previous behavior of **Vitals Link**. For more information on Collection Labels, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Precedence.** Weight assigned to this presentation object. For more information on Collection Labels and Precedence, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Show as Percent.** Specifies whether the resulting graph displays percent values. If so, the y-axis of the graph will range from 0 - 100 percent.
 - *Enabled.* The graph displays percent values.
 - *Disabled.* The graph does not display percent values.

- **ID**. The unique ID assigned to the report. The ScienceLogic platform automatically generates and assigns this identifier.
- **Date Edit**. Date and time the report definition was last edited.

Creating a Presentation Object

To add a presentation object to a performance Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the performance Dynamic Application you want to add a presentation object to. Click its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Click the **[Presentations]** tab. The **Dynamic Applications Presentation Objects** page is displayed.
4. In the **Dynamic Applications Presentation Objects** page, click the **[Reset]** button to clear values from the fields in the top pane, the **Formula Editor** pane, and the **Guide Text** pane.
5. Supply values in the following fields:

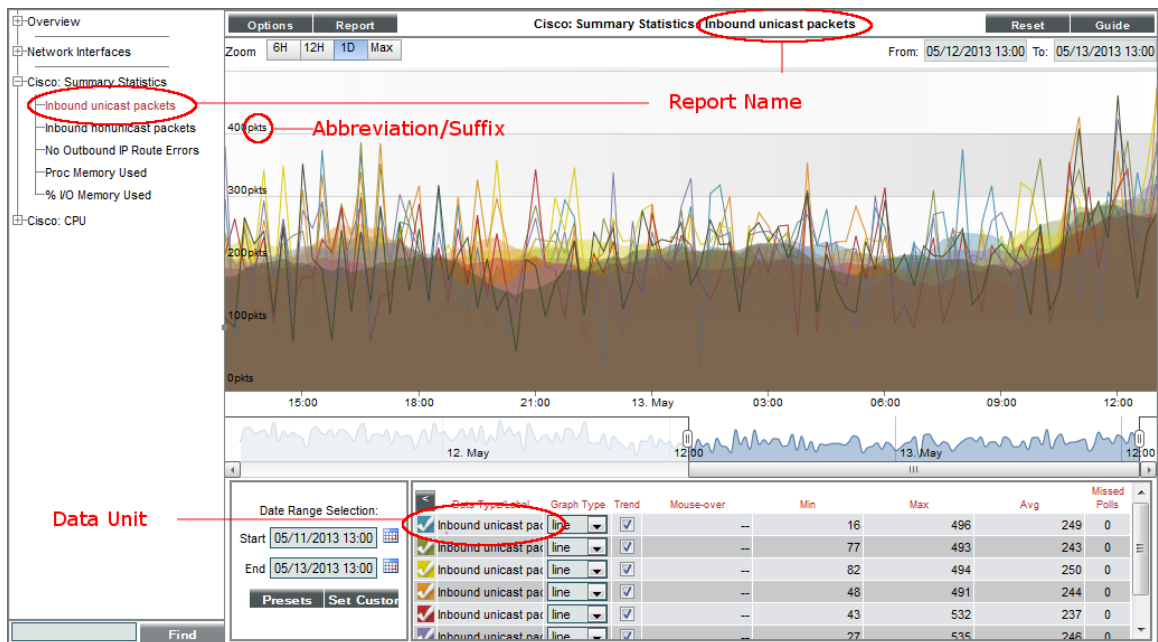


- **Report Name**. Name of the report. Can be any combination of alpha-numeric characters. This name will appear as the name of the report.
- **Summarization State**. Specifies whether the ScienceLogic platform should include hourly and daily summary statistics for the presentation object in the graph. Choices are:
 - *Enabled*. The platform will include summary statistics for the presentation object in the graph.
 - *Disabled*. The platform will not include summary statistics for the presentation object in the graph.
- **Data Unit**. Units used in the graph. Can be any combination of alphanumeric characters, up to 25 characters in length.
- **Abbreviation/Suffix**. Abbreviation for units used in the graph. Can be any combination of alphanumeric characters, up to 16 characters in length.

- **Show as Percent**. Specifies whether the resulting graph displays percent values. If so, the y-axis of the graph will range from 0 – 100 percent.
 - *Enabled*. The graph displays percent values.
 - *Disabled*. The graph does not display percent values.
- **Precedence**. Weight to assign to this presentation object. For more information on Collection Labels and Precedence, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Label Group**. Name of the Collection Group to align with the presentation object. The **Vitals** group is available as a Collection Label and preserves the previous behavior of **Vitals Link**. You can click the plus-sign (+) to create a new Collection Group. For more information on Collection Labels, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Label**. Name of the Collection Label to align with the presentation object. You can click the plus-sign (+) to create a new Collection Label. For more information on Collection Labels, see the chapter [Grouping Dynamic Application Data Using Collection Labels](#).
- **Formula Editor**. The formula the ScienceLogic platform will use to generate graphs for the presentation object. For a full description of this field, see the [Presentation Object Formulas](#) section.
- **Guide Text**. Enter optional guide text for the presentation object.

6. Click the **[Save]** button to save the presentation object.

The following picture shows where the **Report Name**, **Data Unit**, and **Abbreviation/Suffix** values are displayed for a presentation object:



Presentation Object Formulas

The **Formula Editor** allows users to define which object value(s) that will be used to generate the graph for this presentation object and the mathematical manipulations that must be performed on those object values. The result of a defined formula must be a single numeric value.

A presentation object formula can include:

- The ID for one or more collection objects in the same Dynamic Application. All collection object IDs start with "o_" (lowercase "oh", underscore).
- The four main arithmetic operators (+, -, *, /).
- Parentheses, which set precedence for arithmetic operators.
- Braces ({ and }), which can be used to specify that an object is optional, and the formula should still be evaluated even if a value for the optional object has not been collected.
- Any PHP math function that evaluates to a numeric value. For information about PHP math functions, see <http://www.php.net/math>.

The operators and PHP math functions may be used in any combination and with any number of collection object IDs, provided that the result of the expression always equates to a single numeric value.

The following examples are all valid presentation formulas:

```
o_123
abs(o_123)
o_123 / o_124 * 100
(o_123 + o_124) * o_125
o_123 + {o_124}
```

The scrolling list below the **Formula Editor** contains a list of all objects in the Dynamic Application. To insert the ID for a collection object in to the formula, select the collection object from the list, and then select the **[Add]** button.

The **Guide Text** pane allows users to include descriptive text with each graph. When a user views the graph in the **Device Performance** page, the **[Guide]** button will display the text from this pane.

How the ScienceLogic Platform Processes Performance Presentation Objects

The ScienceLogic platform does not store the result of a presentation formula. Every time a user views the graph associated with a presentation object, the platform recalculates the formula using the stored values for the collection objects and generates the graph.

For each poll period in the x-axis of the graph, the ScienceLogic platform calculates the presentation formula and plots the value for that poll period.

If one or more of the collection objects in the presentation formula are not available for a poll period, the graph does not display a value for that poll period.



If one of the collection objects used in a presentation object collects a list of values, the ScienceLogic platform calculates the presentation formula for each entry in the list. The graph will include one line for each entry in the list. For example, suppose the following presentation object formula is used to generate a graph:

$$o_123 / o_124 * 100$$

Suppose that for each poll period for a device, the collection objects (o_123 and o_124) return a list of two values. When the ScienceLogic platform generates the graph for this device using this presentation object, the graph will include two graph lines. The data point for each poll period on the first graph line will be calculated by evaluating the presentation formula for the first value from the list of values collected for o_123 at that poll period and the first value from the list of values collected for o_124 at that poll period. The second graph line will use the second values in the list of collected values.


Editing a Presentation Object

To edit a presentation object in a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to edit. Click its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.
3. Click the **[Presentations]** tab. The **Dynamic Applications Presentation Objects** page is displayed.
4. In the **Presentation Object Registry** pane at the bottom of the page, locate the presentation object you want to edit. Click the wrench icon () for the presentation object you want to edit.
5. The fields in the top pane will be populated with the saved values for the selected presentation object. Edit the values in one or more fields. For a description of each field, see the [Creating a Presentation Object](#) section.
6. Click the **[Save]** button to save your changes to the presentation object. Click the **[Save As]** button to save the changes as a new presentation object.

Deleting a Presentation Object

To delete a presentation object from a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to edit. Click its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.

3. Click the **[Presentations]** tab. The **Dynamic Applications Presentation Objects** page is displayed.
4. In the **Presentation Object Registry** pane at the bottom of the page, locate the presentation object you want to delete. Click its bomb icon (💣). The presentation object is deleted from the ScienceLogic platform. Any collected data that the presentation object used will remain in the ScienceLogic platform. However, you will no longer be able to view any graphs generated using the presentation object.

Alerts and Thresholds

What is an Alert?


An **alert** defines a formula that the ScienceLogic platform will evaluate each time data is collected. Each formula includes **collection objects**. When the ScienceLogic platform evaluates the alert formula, the platform will substitute in the most recently collected values for each collection object. If the formula evaluates to *true*, the platform generates an alert. When the platform generates an alert, the platform makes an entry in the device log. Additionally, you can define an event policy that will trigger an event if a specific alert is generated.





What is a Threshold?

A **threshold** defines a variable that can be used in one or more **alerts**. Each threshold defines a range of values for the variable and a default value. Users can change the value of the threshold on a per-device basis without having to modify the Dynamic Application.

Viewing the Thresholds in a Dynamic Application

To view the thresholds in a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to view the thresholds for. Select its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Thresholds]** tab. The **Dynamic Applications Threshold Objects** page is displayed. The **Thresholds Object Registry** pane at the bottom of the page displays the following information about each threshold:

Threshold Object Registry										
	Name	Override	Type	Numeric Range High	Numeric Range Low	Threshold Unit	Threshold Value	ID	Date	Edit
1.	 CPU Utilization High	Enabled	Percentage	100	0		90	t_219	2013-04-12 17:43:27	
2.	 Test Threshold 1	Enabled	Decimal	100.000	0.500	5	90.000	t_218	2013-04-12 17:41:56	

- **Name**. The name of the threshold. This name will be used to label the threshold in the **[Thresholds]** tab in the **Device Administration** panel.
- **Override**. Defines whether this threshold can be overridden on a per-device basis. Possible values are:
 - *Enabled*. When the Dynamic Application that contains this threshold is aligned with a device, a user will be able to set a new value for this threshold in the **[Thresholds]** tab in the **Device Administration** panel.
 - *Disabled*. Users cannot set a new value for this threshold on a per-device basis.
- **Type**. Specifies whether the threshold will be a *Percentage*, *Integer*, or *Decimal*.
- **Numeric Range High**. When a user overrides this threshold for a device, this is the maximum value the threshold can be set to.
- **Numeric Range Low**. When a user overrides this threshold for a device, this is the minimum value the threshold can be set to.
- **Threshold Unit**. Specifies the unit of measure for the threshold value.
- **Threshold Value**. The default value for the threshold. If a user does not define an override for a device, the ScienceLogic platform will use this threshold value for that device.
- **ID**. The unique ID assigned to the threshold by the ScienceLogic platform. This ID is used to reference the threshold in alert formulas. The unique ID will always start with "t_".
- **Date Edit**. The last time a user edited this threshold.

Creating a Threshold



To add a threshold to a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to add a threshold to. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Thresholds]** tab. The **Dynamic Applications Threshold Objects** page is displayed.
4. Supply values in the following fields:

- **Threshold Name.** The name of the threshold. This name will be used to label the threshold in the **[Thresholds]** tab in the **Device Administration** panel.
 - **Override Threshold Value.** Defines whether this threshold can be overridden on a per-device basis. When disabled, this threshold does not appear on the **Device Thresholds** page for each subscriber device and cannot be edited for each subscriber device. Choices are:
 - *Enabled.* Threshold appears on the **Device Thresholds** page for each subscriber device and can be edited for each subscriber device.
 - *Disabled.* Threshold does not appear on the **Device Thresholds** page for each subscriber device.
 - **Numeric Range High.** Specifies the high end of the possible values. In the **Device Thresholds** page for each subscriber device, this number will appear at the high end of the slider. When a user overrides this threshold for a device (in the **Device Thresholds** page), this is the maximum value the threshold can be set to.
 - **Numeric Range Low.** Specifies the low end of the possible values. In the **Device Thresholds** page for each subscriber device, this number will appear at the low end of the slider. When a user overrides this threshold for a device (in the **Device Thresholds** page), this is the minimum value the threshold can be set to.
 - **Threshold Type.** Specifies whether the threshold will be a *Percentage*, *Integer*, or *Decimal*.
 - **Threshold Unit.** Specifies the unit of measure for the threshold value.
 - **Threshold Value.** Assigns a default value to the threshold object. If a user does not override the threshold in the **Device Thresholds** page for a subscriber device, the ScienceLogic platform will use this threshold value for the subscriber device.
5. Select the **[Save]** button to save the threshold.



Editing a Threshold

To edit an already-defined threshold:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. In the **Dynamic Applications Manager** page, find the Dynamic Application for which you want to edit a threshold. Select its wrench icon () .
3. Select the **[Thresholds]** tab for the Dynamic Application.
4. In the **Threshold Objects** page, find the threshold in the **Threshold Object Registry** pane. Select its wrench icon () .
5. The fields in the top pane are populated with values from the selected threshold. You can edit the value of one or more fields. For a description of each field, see the [Creating a Threshold](#) section.
6. Select the **[Save]** button to save your changes to the threshold.


Deleting a Threshold

You can delete a threshold from a Dynamic Application. To do this:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. In the **Dynamic Applications Manager** page, find the Dynamic Application for which you want to delete a threshold. Select its wrench icon () .
3. Select the **[Thresholds]** tab for the Dynamic Application.
4. In the **Threshold Objects** page, find the threshold in the **Threshold Object Registry** pane. Select its bomb icon () . The threshold will be deleted from the ScienceLogic platform. You must manually remove the threshold from all alert formulas in which that threshold appears.

Viewing the Alerts in a Dynamic Application

To view the alerts in a Dynamic Application:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to view the alerts for. Select its wrench icon () . The **Dynamic Applications Properties Editor** page is displayed.

- Select the **[Alerts]** tab. The **Dynamic Applications Alert Objects** page is displayed. The **Alert Object Registry** pane at the bottom of the page displays the following information about each alert:

Alert Object Registry							
	Policy Name	Formula	State	Maintain	Events	ID	Edit Date
1.	CCM Down	result(o_2062, label = 'o_2061') == 3	Enabled	Yes	Yes	a_309	2013-03-11 11:22:47
2.	CCM Unknown	result(o_2062, label = 'o_2061') == 1	Enabled	Yes	Yes	a_311	2013-03-11 11:22:47
3.	CCM Up	result(o_2062, label = 'o_2061') == 2 and (active(a_311) or active(a_309))	Enabled	No	Yes	a_310	2013-03-11 11:22:47

- **Policy Name.** The name of the alert.
- **Formula.** The formula that the ScienceLogic platform will evaluate to determine whether to generate the alert. An alert formula can contain references to collection objects ("o_" followed by a numeric value), threshold objects ("t_" followed by a numeric value), and/or other alerts ("a_", followed by a numeric value).
- **State.** Specifies whether the ScienceLogic platform will evaluate the alert when new data is collected for the Dynamic Application. Possible values are:
 - *Enabled.* The platform will evaluate the alert when new data is collected.
 - *Disabled.* The platform will not evaluate the alert when new data is collected.
- **Maintain.** Specifies whether the ScienceLogic platform will track the status of this alert (active or inactive) for use as a condition in other alert formulas. Possible values are:
 - *Yes.* The platform tracks the status of this alert. The status of this alert can be used as a condition in other alert formulas.
 - *No.* The platform does not track the status of this alert. The status of this alert cannot be used as a condition in other alert formulas.
- **Events.** Specifies whether an event policy has been created for this alert. Possible values are:
 - *Yes.* An event policy has been created for this alert. Selecting the event icon () will open the **Event Policy Editor**, where you can edit the associated event.
 - *No.* An event policy has not been created for this alert. Selecting the gray event icon () will open the **Event Policy Editor**, where you can define an event policy for this alert.
- **ID.** The unique ID assigned to the alert by the ScienceLogic platform. This ID is used to reference the alert in other alert formulas. The unique ID will always start with "a_".
- **Date Edit.** The last time a user edited this alert.

Creating an Alert

To add an alert to a Dynamic Application, perform the following steps:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to add an alert to. Select its wrench icon (🔧). The **Dynamic Applications Properties Editor** page is displayed.
3. Select the **[Alerts]** tab. The **Dynamic Applications Alert Objects** page is displayed.
4. Supply values in the following fields:

Policy Name	<input type="text"/>
Active State	[Disabled]
Log Message	<input type="text"/>
Maintain State	[No]
Trigger Alert	[None]

Save

Formula Editor

7	8	9	o_2060: IP Address	r()	()	CE
4	5	6	o_2061: Name	*	=	==
1	2	3	o_2062: Status	/	*	>=
0				-	+	<=

Add

- **Policy Name.** Enter a name for the alert.
- **Active State.** Specifies whether the ScienceLogic platform will evaluate the alert when new data is collected for the Dynamic Application. Possible values are:
 - *Enabled.* The platform will evaluate the alert when new data is collected.
 - *Disabled.* The platform will not evaluate the alert when new data is collected.
- **Log Message.** Enter the log message associated with the alert. When the ScienceLogic platform generates this alert, this message will appear in the device logs. You can use the following substitution characters in this field. The substitution characters are populated by including certain functions in the **Formula Editor** for the alert:
 - *%V.* The value returned by the result() function.
 - *%L.* The value returned by the "label" argument in the result() function.
 - *%T.* The value returned by the threshold() function.
- **Maintain State.** Specifies whether the ScienceLogic platform will track the status of this alert (active or inactive) for use as a condition in other alert formulas. Possible values are:
 - *Yes.* The platform tracks the status of this alert. Select this option if you want to use the status of this alert as a condition in other alert formulas.
 - *No.* The platform does not track the status of this alert. Select this option if you do not want to use the status of this alert as a condition in other alert formulas.

- **Trigger Alert.** This is a deprecated field. Leave this field set to the default value.
- **Formula Editor.** Enter the formula that the ScienceLogic platform will evaluate to determine whether to generate the alert. For a full description of this field, see the [Alert Formulas](#) section.

5. Select the **[Save]** button to save the alert.

Alert Formulas

The **Formula Editor** field allows you to define the formula that the ScienceLogic platform will evaluate for an alert. The result of a defined alert formula must be either *True* or *False*. If the formula evaluates to *True*, the platform generates an alert.

An alert formula can include:

- The ID for one or more collection objects in the same Dynamic Application. All collection object IDs start with "o_" (lowercase "oh", underscore). However, do not use a collection object of type *Label: Hourly Polled* in an alert formula.
- The ID for one or more thresholds in the same Dynamic Application. All threshold IDs start with "t_".
- Arithmetic operators (+, -, *, /), comparison operators (>, <, !=, ==), and/or boolean operators ("and", "or", "not").

NOTE: The "==" operator is used to test equality. The single "=" assignment operator is not supported, except when used inside some functions).

- The result() and/or threshold() functions. The results of these functions can be included in the log message for the alert.
- One or more ScienceLogic specific functions described in this section.
- One or more ScienceLogic specific date & time variables or the tab_idx variable described in this section.
- Any standard Python function.

The operators, functions, and variables may be used in any combination and with any number of collection object IDs and threshold IDs, provided that the result of the expression always equates to *True* or *False*.

The scrolling list below the **Formula Editor** contains a list of all collection objects and thresholds in the Dynamic Application. To insert the ID for a collection object or threshold in to the formula, select the collection object or threshold from the list, then select the **[Add]** button.

Evaluate

To evaluate the value of a collection object, enter an expression in the **Formula Editor** that compares the value of the object to something else, where "something else" can be another collection object, a threshold object, a simple number, a string, the result of a function, or a combination of objects and numbers.

For example, suppose you want to generate an alert based on the value of collection object "o_15160", which contains the current temperature of a hardware component in a Dell server. You could select object "o_15160: Temp. Reading" from the scrolling list. When the object appears in the **Formula Editor**, you could enter the following formula:

```
o_15160 > 50
```

The formula evaluates to *True* if the value returned by object o_15160 is greater than 50.

NOTE: The ScienceLogic platform assigns each object an object ID. This is different than the object number assigned to the object in the MIB. The ScienceLogic platform requires that you use the object ID when referring to an object.

If one or more of the collection objects used in an alert formula collect a list of values, the ScienceLogic platform evaluates the alert formula for each entry in the list. An alert can be generated for each entry in the list. For example, suppose the following alert formula is evaluated:

```
o_123 / o_124 * 100 > 90
```

Suppose that for each poll period for a device, the collection objects (o_123 and o_124) return a list of two values. The ScienceLogic platform will evaluate the alert formula twice for this device. For the first evaluation, the first value from the list of values collected for o_123 at that poll period and the first value from the list of values collected for o_124 at that poll period will be substituted into the alert formula. For the second evaluation, the second values in the list of collected values will be substituted in to the alert formula.

If you compare a collection object to a string or an integer, be careful to use quotation marks consistently. The collection object and the comparison string or comparison integer must use the same quotation marks. Either both the collection object and the comparison value must be surrounded by single quotation marks or neither the collection object nor the comparison value should be surrounded by single quotation marks. For example:

- o_1346 == '5' is incorrect syntax (notice no single quotation marks around o_1346 but single quotation marks around 5)
- 'o_1346' == '5' is correct
- o_1346 == 5 is correct

For more details, see the section on [Using Quotes in an Alert](#).

Operators

The **Formula Editor** accepts only arithmetic operators (+, -, *, /), comparison operators (>, <, !=, ==) and Boolean operators (and, or, not). Parentheses are used to group and set precedence for operators.

The comparison operators are those commonly used in many modern programming languages. It is important to note that the '==' operator is used to test equality.

The single '=' assignment operator is not supported (except when passing optional arguments to the result() function, see below).

The operators may be used in any combination and with any number of object IDs, provided that the result of the expression always equates to TRUE or FALSE.

The following examples are all valid formulas:

`o_123 / 100 > 3`

`o_123 > 3 and o_123 < 7`

`o_123 == 5`

`o_123 == 2 or o_123 == 4`

Valid examples using multiple objects:

`o_123 == o_124`

`o_123 / o_124 > 2.5`

`(o_123 + o_124) > (o_125 + 6)`

Dividing Integers

When you divide two integers, sometimes the result is a number that includes a decimal point. For example, if you divide $3/2$, the result is 1.5. A value with a decimal point is called a floating-point number.

In the ScienceLogic platform, **only if the divisor is a float can the resulting quotient be a float**. If an integer is divided by another integer, the result is an integer. For example, if "3/2" is evaluated in an alert formula, the result will be 1, not 1.5. If you are dividing two numbers in the ScienceLogic platform, you can use the python `float()` function to ensure that the divisor is converted to a floating-point number. This will ensure that the resulting quotient is also a floating-point number

For example, suppose your Dynamic Application includes two objects, `o_123` and `o_456`. Suppose you want to divide the value of `o_123` by the value of `o_456`. You should use the `float()` function like this:

```
o_123/float(o_456)
```

Because the divisor is a floating-point number, the resulting quotient can be a floating-point number.

Using Quotes in an Alert

Some alert functions accept either a string or a number as an argument. Some alert functions also require quotes around the arguments. When using quotes with alert functions, follow these guidelines:

- If you want to perform string operations on the object (such as `find()`), put the collection object ID in quotes.
- If you want to perform numeric operations on the object, do not put the collection object ID in quotes.
- You can use single quotes or double quotes – either work, but they must be paired (you cannot start with ' and end with ").
- For consistency, ScienceLogic recommends that you always use single quotes around strings in formulas.

- Take care if you are copying and pasting text from a text editor – the Python interpreter that evaluates alert formulas requires specific quote characters. These are valid quote characters: `'`, these are not: `''` `""`.

The `result()` Function

The **`result()`** function passes a value for use in the alert message or event message associated with the alert. The passed value is available in the `%V` substitution. The expression defined in the `result()` function can be any combination of object IDs, numeric values, and strings allowed in the normal alert formula logic. The syntax is:

```
result(expression)
```

For example, if temperature is provided via object ID `o_126`, the following formula would return the temperature value for use in the alert message or the event message associated with the alert:

```
result(o_126) > 35
```

If the temperature is returned in Celsius, and you want it displayed in the alert message or event message in Fahrenheit, you could use the following formula:

```
result((o_126 * 1.8) + 32) > 120
```

The expression passed to the `result()` function can also refer to multiple objects. For example, suppose you wanted to determine the total number of IP DHCP addresses in a subnet. You could add together the following example objects:

- `o_5678` = number of addresses in use.
- `o_8910` = number of addresses free.
- `o_1112` = number of addresses pending.

The `result()` function would look like this:

```
result(o_5678 + o_8910 + o_1112)
```

There are two optional arguments to the `result` function: `enums` and `label`. These provide for further descriptive text to be substituted into the alert message or event message associated with the alert. The syntax is:

```
result((expression), enums={number:'value',...}, label='label object')
```

The **`enums`** argument specifies a list of substitution characters that the ScienceLogic platform should use to translate the value passed by the `result()` function. In general, this argument is used only if the expression refers to a collection object that is of type *Config Enum*. This argument must be contained within curly braces `{}` and is formatted with the value and substitution characters separated by a colon, and each element separated by commas.

For example, suppose you want to monitor a status object, `o_987`. Suppose this status object is of type *Config Enum* and can return integer values 3, 4, 5, or 6. These integer values map to the following states:

- 3 = OK
- 4 = nonCritical

- 5 = critical
- 6 = nonRecoverable

If collection object `o_987` is passed using the `result()` function, the integer value of `o_987` is stored in the `%V` variable. By using the `enums` argument, you can pass the string value instead:

```
result(o_987, enums = {3:'OK',4:'nonCritical', 5:'critical', 6:'nonRecoverable'})
```

The **label** argument passes an additional value for use in the alert message or event message associated with the alert. The passed label value is available in the `%L` substitution variable.

For example, suppose you have multiple temperature probes and you want to capture which temperature probe is generating the alert. If collection object `o_5678` contains the name of the temperature probe, you could supply the following value in the label argument (again using a conversion from Celsius to Fahrenheit):

```
result(((1.8 * o_1234)+32), label='o_5678')
```

The value of the label argument can now be used in an event definition using the `%L` substitution.

NOTE: The value of the label argument must be in single quotes.

The `threshold()` Function

Like the **result()** function, the **threshold()** function passes a value for use in the alert message or event message associated with the alert. The passed value is available in the `%T` substitution variable. The syntax is:

```
threshold(expression)
```

The `threshold()` function is intended to make the value of a threshold that triggered an alert visible in the corresponding log message and/or event.

As with the `result()` function, the `threshold()` function can contain a combination of object IDs and numeric values.

For example, suppose you are monitoring temperature. The object `o_1234` contains temperature in Celsius. The threshold `t_99` contains the maximum allowable temperature. You could enter the following formula to convert the temperature values to Fahrenheit and generate an alert when the temperature reading exceeded the allowable maximum:

```
result((1.8 * o_1234)+32) > threshold((1.8 * o_t_99)+32)
```

If an alert used the `result()` function with the label argument and the `threshold()` function, the alert message might look like this:

```
%L reporting temperature of %V F, threshold is %T F
```

The values substituted in to replace `%L`, `%V` and `%T` are the actual values captured in the alert formula definition by the `result()` and `threshold()` functions. If the alert is generated, the alert message might look like this:

```
Backplane-probe reporting temperature of 130 F, threshold is 120 F
```

The items in bold are the substituted values.

The active() Function

The **active()** function checks the state of another alert in the same Dynamic Application. If the alert is active, the active() function returns the value TRUE. The active() function can be used in combination with other expressions. When used in combination with other expressions, if the entire expression resolves to TRUE, a new alert is generated. The syntax is:

```
active(alert_ID)
```

To use the active() formula to check the state of an alert, the alert being checked must have the **Maintain State** set to Yes.

For example, suppose you are monitoring a circuit for failed authentications. You could define two alerts:

The first alert is assigned ID a_175 by the ScienceLogic platform. In the first alert, the object o_16060 contains the number of failed authentications. The object o_16056 contains the name of the circuit being monitored for failed authentication.

```
result(o_16060, label='o_16056') > threshold(13)
```

In this expression, a high-end threshold is defined. In this expression, if more than 13 failed authentications occur on the circuit specified in o_16056, an alert is triggered. This alert defines the problem. For this alert, the **Maintain State** drop-down list is set to Yes.

The second alert is assigned ID a_176 by the ScienceLogic platform. In the second alert, we define a low-end threshold for failed authentications.

```
result(o_16060, label='o_16056') < threshold(7) and active(a_175)
```

In this expression, if less than seven failed authentications occur on the circuit specified in o_16056, AND alert a_175 is still active, a new alert is triggered. The second alert defines the "healthy" criteria.

Suppose we had defined a critical event policy based on alert a_175. This event would warn that too many failed authentications occurred on circuit o_16056.

Now suppose we wanted to clear that critical event when failed authentications fall back to an acceptable number. We could create a healthy event policy based on alert a_176. We could define this event to auto-clear the previous critical event. The new event would inform you that failed authentications were back within acceptable levels.

The avg() Function

The **avg()** function calculates and returns the average of all values returned by a collection object when that collection object returns a list of values. When a collection object returns a single value, the avg() function will return that single value. The syntax is:

```
avg(object_ID)
```

For example, suppose a device has three network interfaces. Now suppose the object o_1234 contains the value for "octets in".

On our example device, we would have three separate instances of object o_1234. We could use the avg() function to calculate and return the average of ALL instances of object o_1234. If that average is greater than 1,000,000 octets, an alert is triggered:

```
avg(o_1234) > 1000000
```

The deviation() Function

The **deviation()** function allows you to define alerts that are triggered when an object has a value that is outside the "normal" range for the current hour on the current day of the week. The syntax is:

```
deviation(object_ID)
```

The deviation() function allows you to compare the current value of a collection object to a specified number of standard deviations outside the "normal" value for the object. You can specify that you want to trigger an alert if the current value falls outside the specified number of standard deviations from "normal".

NOTE: To use the deviation() function, you must configure the ScienceLogic platform to store and calculate the mean values and standard deviations for a collection object. You do this by selecting the **Enable Deviation Alerting** field in the **Collection Objects** page. You then specify the minimum and maximum number of weeks to collect deviation data for the object. **The ScienceLogic platform must have already collected at least the minimum number of weeks' worth of data for an object before the platform will evaluate alerts that use the deviation() function. To use the deviation() function, you must specify a minimum value of at least two weeks. The ScienceLogic platform evaluates the deviation() function against the aggregated values collected for multiple weeks. The number of weeks used during the evaluation depends on the amount of data available and is at least the minimum number of weeks specified for the collection object.**

After the ScienceLogic platform has collected the minimum number of weeks' worth of data for an object, the platform calculates the mean value for that object at every hour of every day and then calculates the standard deviations from that mean value.

The deviation function uses the following formula to examine the value of an object:

$$\frac{\text{(current value of a collection object - mean value for current hour on current day of week)}}{\text{(standard deviation for current hour at current day of week)}}$$

NOTE: The deviation() function converts the value of *(current value of a collection object - mean value for current hour on current day of week)* to a positive value. Therefore, you should always compare the results of the deviation function to zero ("0") or to a positive number.

The `deviation()` function sends data to the collector in the form of upserts to calculate the standard deviation. In cases where a collector failover might happen, or if a device moves collector groups, the standard deviation data is re-upserted on the new collector.

Also, when enabling deviation alerting on a collector, if you set the raw data retention threshold and the standard deviation minimum weeks alerting window, any time that the deviation alerting window is higher than the raw data retention threshold, deviation alerting will not be available after a collector failover. Deviation alerting will be available only when the ScienceLogic platform has gathered enough new raw data to meet the minimum weeks specified by standard deviation.

Some possible uses for the deviation function are:

- Determining if an application is functioning properly. For example, if a log file for an application begins to grow at a rate outside the "normal" range, you can trigger an alert to determine if there is a problem with the application.
- Monitoring security. For example, if bandwidth usage exceeds the normal activity, you can trigger an alert that indicates that your network might have been compromised.

Example 1

For example:

- Suppose the platform has already collected four weeks' worth of data for an object (`o_123`).
- Suppose that for the last four weeks, every Monday between 09:00 and 10:00, the mean value for `o_123` is "50".
- Suppose that 68% of all values on Monday between 09:00 and 10:00 fall within +/- "10" of the mean value (that is, between 40 and 60).
- For `o_123` on Mondays between 09:00 and 10:00, a standard deviation of "1" is +/- "10".

Suppose we specify the following alert formula:

```
deviation(o_123) > 1
```

This alert formula specifies that if `o_123` collected a value outside "1" standard deviation of the "normal" value, the platform should trigger an alert.

If on Monday at 09:30, the value of `o_123` is "45", the deviation function performs the following calculation:

```
(45-50)/10 > 1
```

This evaluates to FALSE, therefore the platform would not trigger an alert.

Example 2

For another example:

- Suppose the ScienceLogic platform has already collected four weeks' worth of data for a collection object (o_789).
- Suppose that on Monday, between 09:00 and 10:00, the value of o_789 is always "50" and has not varied **at all during the entire four weeks.**
- The mean value of o_789 on Mondays, between 09:00 and 10:00 is "50".
- 68% of all values on Mondays between 09:00 and 10:00 will fall within +/- "0" of the mean.
- For o_789 on Monday between 09:00 and 10:00, a standard deviation of "1" is +/- "0".

Suppose on Monday at 09:15, the value of o_789 is "50". If this is the expected behavior, we could specify:

```
deviation(o_789) >1
```

The deviation function performs the following calculation:

```
(50-50)/0 > 1
```

This evaluates to FALSE, therefore the platform would not trigger an alert.

Now suppose that on Monday at 09:30, the value of o_789 is "45", the deviation function performs the following calculation:

```
(45-50)/0 > 1
```

In this case, the deviation function would return "infinity". The formula will evaluate to TRUE, and trigger an alert. If we did not expect to repeatedly collect the same value every Monday between 09:00 and 10:00, this formula might be appropriate to monitor object o_789.

NOTE: In the alert editor, you can specify infinity as `float("inf")`

If you do expect the first standard deviation to be "+/- 0", that is, if you expect the value of an object to never fluctuate, you can test for this with the following:

```
deviation(o_789) == float("inf")
```

Because $(45-50)/0$ does equal infinity, this will evaluate to TRUE and trigger an alert. We can define this alert to provide details about a standard deviation of "0".

The 68-95-99.7 rule

Statistically, 99.7% of all values lie within three standard deviations from the mean.

In theory, if an alert formula uses a standard deviation of "1" as the threshold, such as the one in Example 1, the alert will trigger for 32% of the collected values.

Using the 68-95-99.7 rule, the following list indicates the percentage of collected values that will trigger an alert for different standard deviation thresholds:

- **1 (one) standard deviation. 32%** of collected values will fall outside this range. Therefore, an alert that compares an object's value to a value of :1 (one) standard deviation will be triggered 32% of the time.

For example:

```
deviation(o_123) > 1
```

will trigger an alert on 32% of the values of o_123.

- **2 (two) standard deviations. 5%** of collected values will fall outside this range.
- **3 (three) standard deviations. 0.3%** of collected values will fall outside this range.

The find() Function

The **find()** function searches a collection object for a specific alpha-numeric string. The syntax is:

```
object_ID.find('alpha-numeric string')
```

If the collection object contains the string, the find() function returns the location of the string in the collection object, where 0 is the first character of the collection object. If the collection object does not contain the string, the find() function returns -1 (negative one).

For example, suppose you want to trigger an alert if inbound bandwidth-usage drops below a certain level. Suppose collection object o_4567 collects "inbound octets". Suppose collection object o_4568 returns the type of interface. Now suppose you don't want to trigger an alert if this condition occurs on a loopback interface. You could write an alert like this:

```
result(o_4567) < 500 and o_4568.find('loopback') = -1
```

This logic says "if object o_4557 is less than 500 and object o_4558 does not contain the text "loopback" (that is, the find function cannot find the string "loopback", so returns "-1"), trigger an alert.

The global() Function

The **global()** function returns the collected value (or, for list objects that are not collected at each collection, the last collected value) of a collection object regardless of the group or index that is currently being evaluated for the alert formula.

The syntax is:

```
global(object_ID)
```

When you include multiple collection objects in an alert, those collection objects must be in the same group. If a list or table of values is returned by a collection object, the platform evaluates the alert for each index in that group.

You can use the global() function to include a collection object that returns a single value in an alert formula that also includes collection objects that return lists.

For example:

Suppose that you have a group, **Group 1**, that contains the following two collection objects:

- **o_123**. Contains file system utilization in percent.
- **o_124**. Contains name of file system.

Suppose that you have an additional collection object that is not aligned with any group:

- **o_125**. Contains a single value - the file system that is used for data storage for the primary application.

Suppose you want to trigger an event when a file system is over 80% utilized:

- Suppose that you want to trigger a Major event for all file systems except the one used for data storage for the primary application.
- Suppose you want to trigger a Critical event for the file system used for data storage for the primary application.

To do this, you can create two alerts:

- An alert for the file system for data storage for the primary application
- An alert for all other file systems.

In the alert, the object **o_125** will be compared to **o_124** to determine whether this is the file system for data storage for the primary application.

Because **o_125** is a single value in a different group/index than the other collection objects, we use the **global()** function to return the value for **o_125**.

The two alerts would look like this:

```
o_123 > 80 and 'o_124'.find('global(o_125)') > -1
```

```
o_123 > 80 and 'o_124'.find('global(o_125)') = -1
```

The first alert will trigger only for the file system for data storage (file system name is the same name as contained in **o_125**).

The second alert will trigger for all other file systems (file system name is not the same as in **o_125**).

NOTE: If the collection object specified in the **global()** function is a string, the ScienceLogic system will automatically substitute the value surrounded by single quote characters.

The log() Function

The **log()** function calculates and returns the logarithm of a specified number to the specified base. If base is not specified, the log function uses base e to return the natural logarithm. The syntax is:

```
log(number, base)
```

For example:

```
log(o_123, 10)
```

This example would calculate the base-10 logarithm of the value of collection object `o_123`.

The `prior()` Function

The `prior()` function returns the previous value of a collection object (from the previous polling session). The syntax is:

```
prior(object_ID)
```

For example:

```
o_123 > prior(o_123)
```

This example evaluates to true if the value of object `o_123` is now greater than it was during the last polling session.

Suppose you have defined a polling frequency of five minutes for a Dynamic Application. Every five minutes, the ScienceLogic platform will retrieve the value from object `o_123`. The example above could be used to trigger an alert if the value of object `o_123` is now greater than it was five minutes ago.

The `round()` Function

The `round()` function rounds the value of a collection object to a specified number of digits. The `round()` function returns a floating point value.

The syntax is

```
round(object ID, number_of_digits)
```

- The `round()` function rounds values to the closest multiple of 10 to the negative n.
 - `round(0.5)` returns 1.0
 - `round(-0.5)` returns -1.0
- If you omit the optional `number_of_digits` argument, the function defaults to a whole number and ".0:", like "2.0"

For example, suppose object `o_567` contains the value "4.688".

```
round(o_567, 2)
```

would return the value

```
4.67
```

The sum() Function

The **sum()** function calculates and returns the sum of all values returned by a collection object when that collection object returns a list of values. When a collection object returns a single value, the **sum()** function will return that single value. The syntax is:

```
sum(object_ID)
```

For example, suppose a device has three network interfaces. Now suppose the object `o_1234` contains the value for "octets in".

On our example device, we would have three separate instances of object `o_1234`. We could use the **sum()** function to calculate and return the sum of ALL instances of object `o_1234`. If that sum is greater than 1,000,000 octets, an alert is triggered:

```
sum(o_1234) > 1000000
```

Date & Time Variables

You can use the time and date variables to define alerts that occur only during specified dates or times.

Values starting with the word "local" refer to the local date and time, specified in **System Timezone** field in the **Behavior Settings** page (System > Settings Behavior).

For every "local" value there is a "utc" equivalent that uses values based on UTC.

The following is a list of time and date variables that can be used in alerts:

- **localyear**. The year, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be any four-digit year. For example, 2007.
- **localmonth**. The month, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 01 – 12, with 01 being January and 12 being December.
- **localmonthday**. The day of the month, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Value can be 01 – 31.
- **localhour**. The hour, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 0 – 23, with 0 being midnight and 23 being 11 PM.
- **localminute**. The minutes, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 0 -59, with 0 being 0 minutes after the hour and 59 being 59 minutes after the hour.
- **localsecond**. The seconds, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 0 – 61. 60 is used for leap seconds; 61 is used for double leap-seconds.
- **localweekday**. Day of the week, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 0 – 6, with 0 being Monday and 6 being Sunday.

- **localyear**. Day of the year, in local date and time, as specified in the **System Timezone** field in the **Behavior Settings** page. Values can be 1 – 366.
- **utcyear**. The year, in UTC date and time. Values can be any four-digit year. For example, 2007.
- **utcmonth**. The month, in UTC date and time. Values can be 01 – 12, with 01 being January and 12 being December.
- **utcmonthday**. The day of the month, in UTC date and time. Value can be 01 – 31.
- **utcheour**. The hour, in UTC date and time. Values can be 0 – 23, with 0 being midnight and 23 being 11 PM.
- **utcminute**. The minutes, in UTC date and time. Values can be 0 -59, with 0 being 0 minutes after the hour and 59 being 59 minutes after the hour.
- **utcsecond**. The seconds, in UTC date and time. Values can be 0 – 61. 60 is used for leap seconds; 61 is used for double leap-seconds.
- **utcweekday**. Day of the week, in UTC date and time. Values can be 0 – 6, with 0 being Monday and 6 being Sunday.
- **utcyearchday**. Day of the year, in UTC date and time. Values can be 1 – 366.

Example 1

For example, you could define an alert that is triggered only when both the following are TRUE:

- the value of `o_999` falls to zero
- the condition occurs during business hours (between 8AM and 6 PM)

To define the alert, you could use the `localhour` variable:

```
result(o_999) == 0 and localhour >= 8 and localhour <= 18
```

Example 2

You could also define an alert that excludes weekend days. The alert would be triggered only when both the following are TRUE:

- the value of `o_999` falls to zero
- the condition occurs during weekdays

To define the alert, you could use the `localweekday` variable. The days of the week are numbered from 0 to 6, with Saturday being 5 and Sunday being 6. To exclude weekends you could define the alert to trigger only when the “localweekday” number is less than 5

```
result(o_999) == 0 and localweekday < 5
```

Example 3

You could also define an alert that is triggered only on the first day of the month. The alert would be triggered only when both the following are TRUE:

- the value of `o_999` falls to zero
- the condition occurs when the day of the month is "1"

To define the alert, you could use the `localmonthday` variable:

```
result(o_999) == 0 and localmonthday == 1
```

The `tab_idx` Variable

The `tab_idx` variable allows you to apply an alert only to specific indexes when a collection object returns a list of values. For SNMP Dynamic Applications, the `tab_idx` variable returns the SNMP index that corresponds to the collection object values that are currently being evaluated. For other Dynamic Applications, the `tab_idx` variable returns the location in the list of values that corresponds to the collection object values that are currently being evaluated.

For example:

- Suppose that in a list of values representing device state, there may be six potential indexes in the list of values, but only 1, 3, 4, and 6 are used.
- Suppose that the object `o_999` represents the device state, with 0 being healthy.

You could define the alert to check when object `o_999` does not equal zero (is not healthy). But you could limit the alert to check only indexes 1, 3, 4, and 6. For an SNMP Configuration or SNMP Performance Dynamic Application, the alert formula would be:

```
o_999 != 0 and tab_idx in ['.1','.3','.4','.6']
```

For all other types of Dynamic Applications, the alert formula would be:

```
o_999 != 0 and tab_idx in ['1','3','4','6']
```

Notice the syntax for the `tab_idx` variable. You must use the syntax as it appears in the example:



- Include the text "in" after the `tab_idx` variable
- Surround the list of index numbers with square brackets
- Surround each index with single-quotation marks (')
- For SNMP Configuration and SNMP Performance Dynamic Applications, precede each index with a period (.)
- Separate the list of index numbers with commas (,)

Creating an Event Policy for an Alert

When the ScienceLogic platform generates an alert for a device, the log message associated with that alert is inserted into the device logs. Optionally, you can associate an alert with an event policy. When the ScienceLogic platform generates an alert that is associated with an event policy and all the conditions defined in the event policy are met, the platform will generate an event. In addition to appearing in the device logs, the event will appear:

- In the **Event Console** page (the **[Events]** tab).
- In the **[Summary]** tab in the **Device Reports** panel.
- In the **[Events]** tab in the **Device Reports** panel.

To create an event policy for an alert:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the Dynamic Application you want to add an event policy for. Click its wrench icon (). The **Dynamic Applications Properties Editor** page is displayed.
3. Click the **[Alerts]** tab. The **Dynamic Applications Alert Objects** page is displayed.
4. In the **Dynamic Applications Alert Objects** page, find the alert in the **Alert Object Registry** pane. Click its gray event icon (.

- The **Event Policy Editor** page is displayed. The following fields in the **Event Policy Editor** are populated with information about the alert:

The screenshot shows the 'Event Policy Editor | Create New Event Policy' window. It has three tabs: 'Policy', 'Advanced', and 'Suppressions'. The 'Policy' tab is active. On the left, there are three dropdown menus: 'Event Source' (set to 'Dynamic'), 'Operational State' (set to '[Enabled]'), and 'Event Severity' (set to '[Major]'). Below these is a checkbox for 'Use Modifier'. On the right, there are two text input fields: 'Policy Name' and 'Event Message'. Below these is a large text area for 'Policy Description' with a rich text editor toolbar. At the bottom center is a 'Save' button.

- **Event Source.** This field is set to *Dynamic*, which tells the platform that this event policy matches alerts generated using Dynamic Applications.
 - **Policy Name.** This field is populated with the value you specified in the **Policy Name** field for the alert.
 - **Link-Alert.** This field appears on the **[Advanced]** tab in the **Event Policy Editor**. This field is populated with a link to the alert, which tells the platform that the alert you selected can trigger this event policy.
- Before you save the event policy, you must supply a value in the following fields:
 - **Operational State.** Specifies whether the ScienceLogic platform should trigger this event if the associated alert is generated and all the conditions defined in the event policy are met. Choices are:
 - *Enabled.* The platform will trigger this event if the associated alert is generated and all the conditions defined in the event policy are met.
 - *Disabled.* The platform will never trigger this event. Selecting this option does not stop the ScienceLogic platform from generating the associated alert.

- **Event Severity.** The severity value to associate with events created with this event policy. Choices are:
 - *Healthy*
 - *Notice*
 - *Minor*
 - *Major*
 - *Critical*
- **Event Message.** The event message that the ScienceLogic platform will associate with instances of this event. If you want to use the message generated by the alert that triggers this event, enter "%M" in this field.



7. You can optionally specify values in the additional fields, including the **Policy Description** field and the fields that appear in the **[Advanced]** tab. For more information about the **Event Policy Editor**, see the **Events** manual or click the **[Guide]** button. If you choose to make further changes to your event policy, you must not change the values in the following fields:

- The **Event Source** field in the **[Policy]** tab.
- The **Link-Alert** field in the **[Advanced]** tab.

8. Click the **[Save]** button to save the event policy.

Editing an Alert

To edit an already-defined alert:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. In the **Dynamic Applications Manager** page, find the Dynamic Application for which you want to edit an alert. Click its wrench icon ()
3. Click the **[Alerts]** tab for the Dynamic Application.
4. In the **Dynamic Applications Alert Objects** page, find the alert in the **Alert Object Registry** pane. Click its wrench icon ()
5. The fields in the top pane are populated with values from the selected alert. You can edit the value of one or more fields. For a description of each field, see the [Creating an Alert](#) section.
6. Click the **[Save]** button to save your changes to the alert.

Deleting an Alert

You can delete an alert from a Dynamic Application. To do this:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).

2. In the **Dynamic Applications Manager** page, find the Dynamic Application for which you want to delete an alert. Click its wrench icon (🔧).
3. Click the **[Alerts]** tab for the Dynamic Application.
4. In the **Dynamic Applications Alert Objects** page, find the alert in the **Alert Object Registry** pane. Click its bomb icon (💣). The alert will be deleted from the ScienceLogic platform. If an event policy was associated with the alert, you must manually delete the event in the **Event Policy Manager** page (Registry > Events > Event Manager).

Validating an Alert

After you have created an alert formula, you can use a command line interface (CLI) tool to validate that the alert works as intended. This tool enables Dynamic Application developers to validate alerts by directly inserting data into the Alerting module without using a full ScienceLogic system.

Running the Alert Validator Command Line Tool

The `alert_validator.py` Python CLI tool enables Dynamic Application developers to validate the alert formulas that are included in Dynamic Applications using shell sessions at the command line. This tool is located in the following folder:

```
/opt/em7/backend/alert_validator.py
```

NOTE: The `alert_validator.py` tool can be run only from a Data Collector or All-In-One Appliance.

The `alert_validator.py` tool checks an alert formula to validate that the following are all true:

- All collection object and device threshold object references are structured correctly and are valid for the specified Dynamic Application.
- Any collection objects that are referenced outside of `global()`, `sum()`, and `avg()` all belong to the same group of metrics.
- The following functions are structured correctly and contain a valid collection object reference:
 - `sum`
 - `avg`
 - `global`
 - `prior`
 - `changed`
 - `deviation`
- The `active()` function alert reference is structured correctly and is valid for the specified Dynamic Application, with an alert formula that belongs to the same group as the alert being tested.

- The result() and threshold() functions are used correctly (including optional result arguments).
- The alert formula contains no syntax errors.

To run the alert_validator.py tool, you must first run a command to create JSON files that capture data about the Dynamic Application. This command uses the following structure, and **must** always end with the device ID and Dynamic Application ID:

```
sudo -u s-em7-core python /opt/em7/backend/alert_validator.py --capture-data [Device ID Dynamic Application ID]
```

NOTE: Replace [Device ID Dynamic Application ID] with the Device ID and Dynamic Application ID, respectively. Do not include the brackets. The Device ID and Dynamic Application ID should be separated by a space and no other characters.

When you run this command, a folder containing all of the generated JSON files is created in the /tmp/ folder. The folder name includes the Dynamic Application ID.

After the folder with the JSON files has been created, you can then run a separate command to evaluate the alert formulas. This command uses the following structure:

```
sudo -u s-em7-core python /opt/em7/backend/alert_validator.py -j AlertInternalState.json -j DynamicAppAlert_[Alert ID].json -j DynamicAppThreshold_[Threshold ID].json -j DynamicAppObjects_[Object ID].json -j DynamicAppObjectSchema_[Dynamic Application ID].json
```

NOTE: Replace [Alert ID], [Threshold ID], [Object ID], and Dynamic Application ID with the Alert ID, Threshold ID, Object ID, and Dynamic Application ID, respectively. Do not include the brackets.

TIP: ScienceLogic recommends running the command from the folder in which the JSON files are stored so you do not need to include the path for each JSON file.

If the alert validation is successful, a detailed report will display indicating that the specified alerts were successfully triggered.

If the alert validation fails, a message will display that specifies the reason why it failed.

Example

The following example creates JSON files to capture data for device ID 3, Dynamic Application ID 1429:

```
sudo -u s-em7-core python /opt/em7/backend/alert_validator.py --capture-data 3 1429
```

After the JSON files have been created, you can then use them to evaluate the alert formulas. For example:

```
sudo -u s-em7-core python /opt/em7/backend/alert_validator.py -j  
AlertInternalState.json -j DynamicAppObjects_3.json -j DynamicAppThreshold_651.json  
-j DynamicAppAlert_2111.json -j DynamicAppObjectSchema_1429.json
```

What is Dynamic Application Indexing?

When the ScienceLogic platform performs collection for a collection object, the response might contain a list of values for that collection object instead of a single value.

For example, suppose a collection object collects the amount of space used in a file system. For that single collection object, a device might respond with a list of values. Each value in the list is the amount of space used for each of the file systems on that device.

When a collection object returns multiple values, the ScienceLogic platform uses **Indexing** to track each of the values. Each of the values in the list is given an **index**, which is a numerical value that allows the platform to:

- **Maintain associations between collection objects in the same Dynamic Application.** For example, suppose a collection object collects the amount of space used in a file system and a device responds with multiple values. Suppose a second collection object returns the names of those file systems. The lists of values for each collection object should be assigned the same index, so that the values for space used are associated with the correct file system name in graphs and reports. The first value in the list of amount of space used values is associated with the first value in the list of names, etc.
- **Maintain an association between values collected at different times for the same collection object.** Every time the ScienceLogic platform collects a list of values for a collection object, each value that is associated with a value in the previously collected list should have the same index value as the previous value. For example, suppose a collection object collects the amount of space used in a file system and a device responds with multiple values. Suppose that during the first collection, the amount of space used for the file system called "C:\\" is assigned index "1". The next time the platform collects the list of values, the amount of space used for the file system called "C:\\" should again be associated with index "1". When the platform generates a graph of data, each line on the graph is drawn using values with the same index.

Index values are unique and maintain association only for each group of collection objects in each Dynamic Application for each device. For example:

- In a single Dynamic Application, a collection object in group 1 can have the same index values as a collection object in group 2, but the collected values for these collection objects are not associated with each other.
- In a single Dynamic Application, for each group of collection objects, the platform maintains separate sets of indexes for each aligned device. A collection object in group 1 might have different index values for device A and for device B.

The ScienceLogic platform has a default method to assign index values. For some Dynamic Applications, the default method is not sufficient and might create incorrect associations between values. Typically, this is caused when the list of values for a collection object changes order from one collection to the next. If the default indexing method is not sufficient, you can designate a collection object that will control the indexing behavior.

NOTE: Index values are internal to the ScienceLogic platform and are typically not displayed to users. The platform displays index values as the labels for performance graphs only when collected labels are not available.

How Indexing Affects Configuration Tables

For each group of collection objects defined in a configuration Dynamic Application, the ScienceLogic platform displays a table in the configuration report for that Dynamic Application:

- Each collection object is assigned its own column in the table.
- Each index is assigned its own row in the table.

If a given collection object does not have a collected value for an index, an empty field will be displayed for that row and column.

How Indexing Affects Performance Graphs

For each presentation object defined in a Dynamic Application, the ScienceLogic platform generates a graph that includes one line for each index.

For example, suppose a presentation formula in a Dynamic Application includes a single collection object:

`o_123`

Suppose that when this Dynamic Application is aligned to a device, the ScienceLogic platform collects a list of three values for `o_123`.

- The graph for this presentation object will display three lines, one for each index.
- When new data is collected, the new poll time will be included on the x-axis of the graph.

- For each of the three lines, the data point (y-axis) for the new poll time will be the collected value for that index.

For example, suppose the ScienceLogic platform assigns indexes 0, 1, and 2 to the lists of collected values for o_123. Suppose the platform then performs three polls at 00:00, 00:05, and 00:10. Suppose the indexed collected values look like this:

Index	00:00	00:05	00:10
0	6	8	7
1	15	30	20
2	50	60	55

When the platform generates the performance graph:

- One graph line will include the values 6, 8, and 7.
- One graph line will include the values 15, 30, and 20.
- One graph line will include the values 50, 60, and 55.

Now suppose that during the 00:15 poll, a list of three values is returned for the collection object, but the values are assigned the indexes 1, 2, and 3:

Index	00:00	00:05	00:10	00:15
0	6	8	7	-
1	15	30	20	9
2	50	60	55	25
3	-	-	-	70

The graph will now include four lines:

- One graph line will include the values 6, 8, and 7. This line will not include a data point for the 00:15 poll.
- One graph line will include the values 15, 30, 20, and 9.
- One graph line will include the values 50, 60, 55, and 25.
- One graph line will not include data points for the 00:00, 00:05, and 00:10 polls, and will include the value 70 for the 00:15 poll.

For presentation formulas that contain multiple collection objects, the ScienceLogic platform evaluates the formula for the values at each index. A graph line is drawn for each index for which there are enough collected values to calculate a result. For example, suppose a presentation object has the following formula:

$$o_{123} + o_{321}$$

If the ScienceLogic platform collects the following lists of values for o_123 and o_321 and assigns the following index values:

Index	o_123	o_321
0	100	-
1	1000	750
2	100	30
3	-	80

The ScienceLogic platform will evaluate the presentation formula for each of the four indexes. However, there are only enough collected values to calculate a result for two indexes (1 and 2). Therefore, two graph lines will appear on the graph.

Suppose the presentation formula was changed to make o_123 optional:

`o_123 + {o_321}`

For the same set of collected values, the graph would include three graph lines, which would correspond to indexes 0, 1, and 2. Index 3 would not be included, because object o_123 is not optional and does not include a value at index 3.

The indexes assigned to collected values are also used to label the lines in the graph key that appears in the lower right of the **Performance** page. When a graph line is drawn for an index, the ScienceLogic platform looks for a collected value that matches the following criteria:

- Was collected for a collection object in the Dynamic Application that has a **Class Type** of *101 Label (Hourly Polled)* or *104 Label (Always Polled)*.
- Has the same group value as the collection objects that were used to draw the graph line.
- Has the same index value as the values used to draw the graph line.

If the ScienceLogic platform finds a collected value that matches the criteria, that value is displayed in the graph key for that graph line. If the platform does not find a collected value that matches the criteria, the graph line is labeled "Index X", where X is the assigned index for the values used to generate the graph line.

NOTE: The poll time for the collected value is not considered when the platform chooses a label for a graph line. For collection objects that have a **Class Type** of *101 Label (Hourly Polled)* or *104 Label (Always Polled)*, the platform stores only the last collected value.

How Indexing Affects Alerts

When the ScienceLogic platform evaluates an alert, the formula is evaluated separately for each index assigned to the collected values.

For example, suppose an alert formula in a Dynamic Application uses one collection object:

$$o_{123} > 0$$

Suppose that when this Dynamic Application is aligned to a device, the ScienceLogic platform collects a list of three values for o_{123} . The platform will evaluate the alert formula three times, once for each collected value.

Suppose another alert formula uses two collection objects:

$$(o_{123} / 100) * o_{321} > 90$$

Suppose the ScienceLogic platform collects the following lists of values for o_{123} and o_{321} and assigns the following index values:

Index	o_{123}	o_{321}
0	100	80
1	1000	750
2	100	30

During each poll period, the ScienceLogic platform will evaluate the alert formula three times, once for each pair of values for each index.

Now suppose that for a different device, the ScienceLogic platform collects a list of three values for each collection object, but the indexes assigned to the collected values do not match:

Index	o_{123}	o_{321}
0	200	-
1	300	160
2	100	30
3	-	80

For this set of collected values, the alert formula will be evaluated four times, once for each index assigned to the collected values. However, the evaluation for indexes 0 and 3 will fail because not all the collection objects have a value for that index.

The Default Indexing Method

For each group of collection objects in each Dynamic Application associated with a single device, the ScienceLogic platform maintains a list of index values. When a list of values is collected for a collection object, the platform uses the following information about each value in the list to determine the index for that value:

- The ID for the device the data was collected from.
- The ID for the Dynamic Application the data was collected using.

- The ID for the group the collection object is included in.
- Where in the list of values the value appears, called the *instance*.

Each index in the list of indexes maintained by the ScienceLogic platform has the same four attributes (device ID, Dynamic Application ID, group ID, and instance). During collection, the ScienceLogic platform attempts to match each value in the list of collected values with an index in the list of indexes. If a match is found, i.e. the device ID, Dynamic Application ID, group ID, and instance are all identical, the platform associates the matching index with the collected value. If a match is not found, the platform creates a new record in the list of indexes for that device ID, Dynamic Application ID, group ID, and instance and assigns a new index value.

Instance Values and Index Creation for SNMP Dynamic Applications

For collection objects in Dynamic Applications that use the SNMP protocol, the *instance* of a value in a list of collected values is the SNMP index associated with the value. For example, suppose the SNMP OID used in a collection object is ".1.3.6.1.2.1.25.2.3.1.6". This OID represents the amount of storage space used (as reported by the Host Resources MIB). Suppose a device responds to this OID in the following way:

```
HOST-RESOURCES-MIB::hrStorageUsed.1 = INTEGER: 6097420
HOST-RESOURCES-MIB::hrStorageUsed.3 = INTEGER: 6097420
HOST-RESOURCES-MIB::hrStorageUsed.6 = INTEGER: 472112
HOST-RESOURCES-MIB::hrStorageUsed.7 = INTEGER: 1984116
```

The SNMP indexes for the returned values are ".1", ".3", ".6", and ".7". These SNMP indexes are used as the instance for each value when the ScienceLogic platform assigns an index.

If the ScienceLogic platform needs to create a new index record for an SNMP collection object, that is, the device ID, Dynamic Application ID, group ID, and instance do not match an existing index record, the platform will first attempt to use the numeric equivalent of the instance as the index. In the example above, the platform will attempt to create indexes 1, 3, 6, and 7.

If the ScienceLogic platform tries to add an index record and determines that the index already exists for the same device ID, Dynamic Application ID, and group ID, the platform will create the new record with an index equal to one more than the current highest index value for that device ID, Dynamic Application ID, and group ID.

In the example above, suppose that the platform has already created index records for an instance other than ".1", ".3", ".6", and ".7" for the same device ID, Dynamic Application ID, and group ID. Suppose that it is the only instance that has an existing index record. Suppose that instance has been assigned index 7. When the platform attempts to create index records for the four instances in the example above, index records for instances ".1", ".3", and ".6" are successfully created with indexes 1, 3, and 6. Because index 7 is already taken for this device ID, Dynamic Application ID, and group ID, the index record for instance ".7" is created with index 8, which is one more than the current highest index value (7).

Instance Values and Index Creation for Snippet Dynamic Applications

For collection objects in Snippet Dynamic Applications, the *instance* of a value in a list of collected values is defined by the snippet code that collects that collection object. Collection objects are passed back to the ScienceLogic platform by snippet code in a tuple. The first value in the tuple is the *instance*, the second value in the tuple is the collected value.

The ScienceLogic platform creates index records for Snippet Dynamic Application collection objects with a method similar to that used for SNMP Dynamic Applications:

- If the ScienceLogic platform needs to create a new index record for a snippet collection object, the platform will first attempt to use the value of the instance as the value of the index.
- If the ScienceLogic platform tries to add a record for an index value that already exists for the same device ID, Dynamic Application ID, and group ID, the platform will create the new record with an index equal to one more than the current highest index value for that device ID, Dynamic Application ID, and group ID.

Instance Values and Index Creation for Other Dynamic Applications

For collection objects in Dynamic Applications that use a protocol of Database, SOAP, WMI, XML, and XSLT, the **instance** of a value in a list of collected values is the location at which that value appears in the response. The first value in the list is instance "0", the second value in the list is instance "1", and so on.

For example, suppose a collection object for a Database Dynamic Application runs a query that selects the column "speed" from a database table. Suppose the following list of data is returned:

speed
1000
500
1200
700

The instance values would be:

speed	instance
1000	0
500	1
1200	2
700	3

To create a new index record for a Database, SOAP, WMI, XML, or XSLT collection object, the ScienceLogic platform will create the new record with an index equal to one more than the current highest index value for that device ID, Dynamic Application ID, and group ID. The first created index for a device ID, Dynamic Application ID, and group ID is 0.

Designating an Index

In cases where the order or size of the list of values for a collection object might change, the default method of assigning indexes is not sufficient to maintain the associations between collected values.

For example, consider the initial instance values that were assigned to the list of "speed" values in the previous example:

speed	instance
1000	0
500	1
1200	2
700	3

Now suppose that the list of values changes order during the next poll:

speed
1000
500
700
1200

The "700" and "1200" values are now matched to different instance values:

speed	instance
1000	0
500	1
700	2
1200	3

This means that when the ScienceLogic platform stores the new values "700" and "1200", they are associated with different index values than the previously stored data. If these values are used in a performance graph, the new value of "700" will appear in the graph line for index 2 (which has "1200" as the previous data points), and the new value of "1200" will appear in the graph line for index 3 (which has "700" as the previous data points). If these values are used in a configuration table, the ScienceLogic platform will detect a change from "700" to "1200" for index 2 and a change from "1200" to "700" for index 3.

When the default method of assigning indexes is not sufficient, you can designate a collection object in a group of collection objects as the index. When you designate a collection object as an index, the ScienceLogic platform still uses the same information about each value in the list to determine the index for that value:

- The ID for the device the data was collected from.
- The ID for the Dynamic Application the data was collected using.
- The ID for the group the collection object is included in.
- The **instance** associated with the value.

When you designate a collection object as an index, the ScienceLogic platform uses the value and location of the "index" collection object to define the **instance**. The collected values for the "index" collection object are recorded as the instances. For the collected values for the other collection objects in the same group, the instance is the collected value for the "index" collection object that appears at the same location in the list for that poll period.

For example, suppose that a group contains two collection objects:

- A collection object that collects the storage space used as reported by the Host Resources MIB (SNMP OID .1.3.6.1.2.1.25.2.3.1.6)
- A collection object that collects the storage description as reported by the Host Resources MIB (SNMP OID .1.3.6.1.2.1.25.2.3.1.3).

Suppose that the collection object that collects the storage description is designated as the index for the group. Suppose the first time this Dynamic Application collects data from "Device A", the device responds to the storage description OID in the following way:

```
HOST-RESOURCES-MIB::hrStorageDescr.1 = STRING: C:\
HOST-RESOURCES-MIB::hrStorageDescr.3 = STRING: D:\
HOST-RESOURCES-MIB::hrStorageDescr.6 = STRING: Virtual Memory
HOST-RESOURCES-MIB::hrStorageDescr.7 = STRING: Physical Memory
```

For this group of collection objects in this Dynamic Application for this device, the ScienceLogic platform will create the following record of instances to internal indexes:

instance	index
C:\	1
D:\	3
Virtual Memory	6
Physical Memory	7

NOTE: The ScienceLogic platform will still attempt to use the numeric equivalent of the SNMP index as the index.

For the same first poll of the device, suppose the device responds to the storage used OID in the following way:

```
HOST-RESOURCES-MIB::hrStorageUsed.1 = INTEGER: 6097420
HOST-RESOURCES-MIB::hrStorageUsed.3 = INTEGER: 0
HOST-RESOURCES-MIB::hrStorageUsed.6 = INTEGER: 24863
HOST-RESOURCES-MIB::hrStorageUsed.7 = INTEGER: 27257
```

To determine the internal index to associate with the value at SNMP index ".1", 6097420:

1. The ScienceLogic platform determines the value for the designated index (the storage description) at the same location (SNMP index ".1"). This value is "C:\".
2. The ScienceLogic platform uses "C:\" as the instance value to match against the record of internal indexes for this group, Dynamic Application, and device.

- The instance value of "C:\\" matches to internal index 1. When the ScienceLogic platform stores the value "6097420", the platform associates the value with index 1.

The ScienceLogic platform repeats this process for the other values returned for the storage used OID. For the first poll of the device, the following table shows the values returned for the storage used OID, the instance value the platform used, and the internal index that is associated with the stored values:

collected value	instance	index
6097420	C:\	1
0	D:\	3
24863	Virtual Memory	6
27257	Physical Memory	7

Every time the ScienceLogic platform uses this Dynamic Application to collect data from the device, the platform determines the instance values using this method. Suppose that during the next poll, the device responds to the storage description OID in the following way:

```
HOST-RESOURCES-MIB::hrStorageDescr.1 = STRING: C:\
HOST-RESOURCES-MIB::hrStorageDescr.6 = STRING: Virtual Memory
HOST-RESOURCES-MIB::hrStorageDescr.7 = STRING: Physical Memory
HOST-RESOURCES-MIB::hrStorageDescr.8 = STRING: D:\
```

And responds to the storage used OID in the following way:

```
HOST-RESOURCES-MIB::hrStorageUsed.1 = INTEGER: 6097420
HOST-RESOURCES-MIB::hrStorageUsed.6 = INTEGER: 24863
HOST-RESOURCES-MIB::hrStorageUsed.7 = INTEGER: 27257
HOST-RESOURCES-MIB::hrStorageUsed.8 = INTEGER: 0
```

Notice that the values for the "D:\" drive have moved from SNMP index ".3" to SNMP index ".8".

To determine the internal index to associate with the values at SNMP index ".8":

- The ScienceLogic platform determines the value for the designated index (the storage description) at SNMP index ".8". This value is "D:\".
- The ScienceLogic platform uses "D:\" as the instance value to match against the record of internal indexes for this group, Dynamic Application, and device.
- The instance value of "D:\" matches to internal index 3. When the ScienceLogic platform stores the value "D:\" and "0", the platform associates the values with index 3.

Therefore, the storage description and storage used values for the "D:\" drive are associated with the same internal index as was associated with the previous values for the storage description and storage used.

Grouping Dynamic Application Data Using Collection Labels

What are Collection Labels and Collection Groups?

Collection Labels and Collection Groups allow you to group and view data from multiple performance Dynamic Applications in a single dashboard widget.

For example:

- Suppose you monitor phone systems from multiple vendors.
- Suppose you want to create a dashboard that displays the ten phone systems that drop the most calls.
- You could create a Collection Group called "Dropped Calls".
- You could create two Collection Labels: "Average Dropped Calls", and "Raw Dropped Calls".
- For each vendor, you could edit the appropriate performance Dynamic Application and align a collected value with "Average Dropped Calls" and align another collected value with "Raw Dropped Calls".
- You could then create a dashboard that displays the ten phone systems with the highest values for "Raw Dropped Calls" and also displays the ten phone systems with the highest values for "Average Dropped Calls".

Viewing the List of Collection Labels

The **Collection Labels** page (System > Manage > Collection Labels) displays a list of all the existing Collection Labels. By Default, the ScienceLogic platform includes the following Collection Groups:

- **Vitals.** Includes the Collection Labels "CPU", "Memory", and "Swap".
- **Video Performance.** Includes Collection Labels for common performance metrics associated with video endpoint devices.

The **Collection Labels** page displays the following about each existing Collection Label:

Collection Labels Collection Labels Found [17]							Reset	Guide
Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates		
1. In Use	In Use	Video Performance	No	2	--	--		
2. Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--		
3. Max Jitter	Max Jitter	Video Performance	No	2	--	--		
4. Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--		
5. Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--		
6. Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--		
7. Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--		
8. Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--		
9. Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--		
10. Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--		
11. Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--		
12. Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--		
13. Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--		
14. Usage	Usage	Video Performance	No	2	--	--		
15. CPU		Vitals	No	37	9	--		
16. Memory		Vitals	No	17	--	--		
17. Swap		Vitals	No	6	--	--		

- **Label Name.** Name of the Collection Label.
- **Label Description.** Description of the Collection Label. This field is optional.
- **Group Name.** Collection Group that contains this Collection Label.
- **Frequent Data.** Specifies whether *frequently rolled up data* is calculated for the Collection Label.
- **Aligned Presentations.** Presentation Objects aligned with this Collection Label.
- **Aligned Devices.** Devices that currently populate the Collection Label.
- **Duplicates.** Number of devices for which two or more Presentation Objects are aligned with the same Collection Label.

Filtering the List of Collection Labels

You can filter the list of Collection Labels on the **Collection Labels** page by one or more parameters. Only Collection Labels that meet all the filter criteria will be displayed in the **Collection Labels** page.

To filter by parameter, enter text into the desired filter-while-you-type field. The **Collection Labels** page searches for Collection Labels that match the text, including partial matches. By default, the cursor is placed in the left-most filter-while-you-type field. You can use the <Tab> key or your mouse to move your cursor through the fields. The list is dynamically updated as you type. Text matches are not case-sensitive.

You can also use *special characters* to filter each parameter.

Filter by one or more of the following parameters:

- **Label Name.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching label name.
- **Label Description.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching label description.

- **Group Name.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching group name.
- **Frequent Data.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that have a matching value in the **Frequent Data** field.
- **Aligned Presentations.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching number of presentations.
- **Aligned Devices.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching number of aligned devices.
- **Duplicates.** You can enter text to match, including special characters (comma, ampersand, and exclamation mark), and the **Collection Labels** page will display only Collection Labels that are associated with a matching number of duplicates.

Special Characters

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

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

String and Numeric

- , (comma). Specifies an "OR" operation. Works for string and numeric values. For example:
"dell, micro" matches all values that contain the string "dell" OR the string "micro".
- & (ampersand). Specifies an "AND " operation. Works for string and numeric values. For example:
"dell & micro" matches all values that contain both the string "dell" AND the string "micro", in any order.

- ! (exclamation point). Specifies a "not" operation. Works for string and numeric values. For example:

"!dell" matches all values that do not contain the string "dell".

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

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

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

"! ^ " would match null values.

"!\$" would match null values.

"!*" would match null values.

"happy, !dell" would match values that contain "happy" OR values that do not contain "dell".

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

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

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

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

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

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

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

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

String

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

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

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

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

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

"! ^ " would match null values.

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

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

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

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

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

"!\$" would match null values.

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

Numeric

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Additional Examples

- "aio\$". Matches only text that ends with "aio".
- "^shu". Matches only text that begins with "shu".
- "^silo\$". Matches only the text "silo", with no characters before or after.
- "!silo". Matches only text that does not contain the characters "silo".
- "!^silo". Matches only text that does not start with "silo".
- "!0\$". Matches only text that does not end with "0".
- "!^silo\$". Matches only text that is not the exact text "silo", with no characters before or after.
- "!. Matches null values, typically represented as "--" in most pages.
- "!"\$. Matches null values, typically represented as "--" in most pages.
- "!. Matches all text that is not null.
- silo, !aggr". Matches text that contains the characters "silo" and also text that does not contain "aggr".
- "silo, 02, !aggr". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr".
- "silo, 02, !aggr, !01". Matches text that contains "silo" and also text that contains "02" and also text that does not contain "aggr" and also text that does not contain "01".
- "^s*i*i*o\$". Matches text that contains the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would match.
- "!^s*i*i*o\$". Matches all text that does not contain the letter "s", "i", "l", "o", in that order. Other letters might lie between these letters. For example "sXiXo" would not match.
- "!vol&!silo". Matches text that does not contain "vol" AND also does not contain "silo". For example, "volume" would match, because it contains "vol" but not "silo".
- "!vol&02". Matches text that does not contain "vol" AND also contains "02". For example, "happy02" would match, because it does not contain "vol" and it does contain "02".
- "aggr,!vol&02". Matches text that contains "aggr" OR text that does not contain "vol" AND also contains "02".
- "aggr,!vol&!infra". Matches text that contains "aggr" OR text that does not contain "vol" AND does not contain "infra".
- "*". Matches all text.
- "!*". Matches null values, typically represented as "--" in most pages.
- "silo". Matches text that contains "silo".
- "!silo". Matches text that does not contain "silo".
- "!^silo\$". Matches all text except the text "silo", with no characters before or after.
- "-3,7-8,11,24,50-". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50.

- "-3,7-8,11,24,50-,a". Matches numbers 1, 2, 3, 7, 8, 11, 24, 50, and all numbers greater than 50, and text that includes "a".
- "?n". Matches text that contains any single character and the character "n". For example, this string would match "an", "bn", "cn", "1n", and "2n".
- "n*SAN". Matches text the contains "n", zero or any number of any characters and then "SAN". For example, the string would match "nSAN", and "nhamburgerSAN".
- "^?n*SAN\$". Matches text that begins with any single character, is following by "n", and then zero or any number of any characters, and ends in "SAN".

Creating a Collection Group

You cannot create a Collection Group separately from creating a Collection Label. When you [create a Collection Label](#), you can specify a new Collection Group or specify an existing Collection Group. If you specify a new Collection Group, the ScienceLogic platform saves the new Collection Group when it saves the new Collection Label.

Creating a Collection Label

You can create a new Collection Label from the **Collection Labels** page (System > Manage > Collection Labels). To do so:

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).
2. Select the green plus-sign in the lower left of the page.

	Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates
1	In Use	In Use	Video Performance	No	2	--	--
2	Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--
3	Max Jitter	Max Jitter	Video Performance	No	2	--	--
4	Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--
5	Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--
6	Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--
7	Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--
8	Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--
9	Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--
10	Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--
11	Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--
12	Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--
13	Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--
14	Usage	Usage	Video Performance	No	2	--	--
15	CPU	Vitals	Vitals	No	37	9	--
16	Memory	Vitals	Vitals	No	17	--	--
17	Swap	Vitals	Vitals	No	6	--	--

3. Enter values in the following columns:
 - **Label Name**. Name of the Collection Label. This field is required.
 - **Label Description**. Description of the Collection Label. This field is optional.

- **Group Name.** Collection Group to align with the Collection Label. You can select from a list of existing Collection Groups or enter the name of a new Collection Group. This field is required.
- **Frequent Data.** Specifies whether *frequently rolled up data* is calculated for the Collection Label. If the Collection Label will include data that is collected every five minutes or more frequently, and you require that dashboard data be updated every 15 minutes or 20 minutes, select Yes in this field. This data is available immediately for use in a collection label.
- **Save icon** (📌). Select this icon to save your new Collection Label.

4. The new Collection Label appears in the page.

What is Normalization?

Normalization and roll-up are the processes by which the ScienceLogic platform manages collected performance data for display and storage.

- **Raw data** is the data exactly as it was collected from a device or application.
- **Normalized** and **rolled up** data is data for which the platform has performed calculations, usually averaging raw data over a period of time.

Dynamic Applications can collect raw performance data from a device at the following intervals:

- 1 minute
- 2 minutes
- 3 minutes
- 5 minutes
- 10 minutes
- 15 minutes
- 30 minutes
- 1 hour
- 2 hours
- 6 hours
- 12 hours
- 24 hours

For performance Dynamic Applications, you specify this interval in the **Poll Frequency** field, in the **Properties Editor** page (System > Manage > Applications > Create or 🛠️)

The platform **rolls up** data so that reports with a larger timespan do not become difficult to view and to save storage space on the ScienceLogic database. When the platform rolls up data, the platform groups data into larger sets and calculates the average value for the larger set.

There are two types of roll up:

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

The platform rolls up raw performance data as follows:

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

Before the platform normalizes date, EM7 **transforms** the data. To transform data, the platform:

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

NOTE: During the data transform steps, the platform does not directly roll up the raw data in the database tables.

When the platform rolls up data, the platform must **normalize** that data. To normalize data, the platform:

- groups and orders the data
- determines the sample size
- calculates count
- determines the maximum value
- determines the minimum value
- calculates the mean value
- calculates the average value
- calculates the sum
- determines the standard deviation

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

Example

For example, suppose that **every five minutes**, the platform collects data about file system usage on the device named **my_device**. When the platform normalizes and rolls up the collected data for file system usage for **my_device**, the platform will:

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

What are Duplicates and How Does the ScienceLogic Platform Manage Them?

Multiple presentation objects can be aligned with a single Collection Label. For example, suppose that a Dynamic Application includes a presentation object for "memory used", and another Dynamic Application includes a presentation object for "memory usage". Suppose that both of these presentation objects are aligned with the Collection Label named "Memory".

Suppose that one of the devices monitored by the ScienceLogic platform subscribes to both of those Dynamic Applications (for example, a Dynamic Application that monitors OEM hardware and a Dynamic Application that monitors the operating system). For that device, the ScienceLogic platform will collect values for both presentation objects that are aligned with the Collection Label named "Memory".

When this situation arises, the platform uses precedence and some internal rules to assign a single presentation object to the Collection Label for that device. However, you can manually assign a different presentation object to the Collection Label after discovery.

If a device has a duplicate, the platform uses the following rules to determine which presentation object to use for that Collection Label for that device:

- If a manually defined Collection Label-presentation object pair exists, use that pair.
- If the platform cannot find a manually defined Collection Label-presentation object pair, use the pair with the lowest **precedence** value.
- If the platform finds more than one Collection Label-presentation object pair with the same precedence value, the platform will create a pair using the presentation object with the lowest presentation ID.

What is Precedence?

The ScienceLogic platform performs discovery (during initial discovery and during nightly updates) and aligns Dynamic Applications with devices. During discovery, the platform will also align Collection Labels with devices. For devices with **duplicates**, the platform evaluates **precedence** to automatically align a single presentation object with each Collection Label. For devices with duplicates, the platform assigns the Collection Label-presentation object pair with the lowest precedence value.

The ScienceLogic platform evaluates precedence:

- During nightly update discovery.

NOTE: If you have manually defined a Collection Label-presentation object pair for one or more devices, nightly update discovery will not change the Collection Label-presentation object pair.

- When a Dynamic Application is manually aligned with a device in the **Dynamic Application Collections** page
- When devices are manually merged.

Aligning a Presentation Object with a Collection Label

You can align one or more presentation objects with a collection label. This allows the ScienceLogic platform to compare and display reports on data from multiple performance Dynamic Applications.

To align a presentation object with a collection label:

1. Go to the **Dynamic Applications Manager** page (System > Manage > Applications).
2. Find the performance Dynamic Application that contains the presentation object you are interested in. Select the wrench icon (🔧) for that Dynamic Application.
3. In the Dynamic Application panel, select the **Presentations** tab.
4. In the **Presentation Objects** page, go to the **Presentation Object Registry** pane and find the presentation object you want to align with a Collection Label. Select the wrench icon (🔧) for that presentation object.

The screenshot shows the 'Presentations' tab for a Dynamic Application. The 'Presentation Objects' section is active, displaying a 'Receive Audio Jitter' report configuration. The 'Label Group' is set to 'Audio Jitter' and the 'Label' is 'Audio Jitter - Inc'. The 'Precedence' is set to 50. Below this, the 'Presentation Object Registry' table lists various presentation objects, with 'Receive Audio Jitter' highlighted in red.

	Report Name	State	Abbreviation Suffix	Group	Label	Precedence	Show as Percent	ID	Date Edit
1	Receive Audio Jitter	Enabled	ms	--	--	--	No	pres_2283	2014-06-25 12:06:29
2	Receive Audio Packet Loss	Enabled	pkts	--	--	--	No	pres_2284	2014-06-25 12:06:29
3	Receive Audio Rate	Enabled	kbps	--	--	--	No	pres_2285	2014-06-25 12:06:29
4	Receive Packet Loss Percent	Enabled	--	--	--	--	No	pres_2292	2014-06-25 12:06:29
5	Receive Video Frame Rate	Enabled	frps	--	--	--	No	pres_2286	2014-06-25 12:06:29
6	Receive Video Jitter	Enabled	ms	--	--	--	No	pres_2287	2014-06-25 12:06:29
7	Receive Video Packet Loss	Enabled	pkts	--	--	--	No	pres_2288	2014-06-25 12:06:29
8	Receive Video Rate	Enabled	kbps	--	--	--	No	pres_2289	2014-06-25 12:06:29
9	Receive Video Rate Used	Enabled	kbps	--	--	--	No	pres_2290	2014-06-25 12:06:29
10	Transmit Audio Jitter	Enabled	ms	--	--	--	No	pres_2282	2014-06-25 12:06:29
11	Transmit Audio Packet Loss	Enabled	pkts	--	--	--	No	pres_2280	2014-06-25 12:06:29
12	Transmit Audio Rate	Enabled	kbps	--	--	--	No	pres_2275	2014-06-25 12:06:29
13	Transmit Packet Loss Percent	Enabled	--	--	--	--	No	pres_2291	2014-06-25 12:06:29
14	Transmit Video Frame Rate	Enabled	frps	--	--	--	No	pres_2278	2014-06-25 12:06:29


- The top pane is populated with values from the selected presentation object. Select values for the following fields:
 - Precedence.** Set the global precedence for this Collection Label-presentation object pair. For more information, see the section on [Precedence](#).
 - Label Group.** Select from a list of existing Collection Groups or click on the plus-sign icon (+) and enter the value for a new Collection Group. The current presentation object will be a member of the specified Collection Group.
 - Label.** Select from a list of existing Collection Labels or click on the plus-sign icon (+) and enter the value for a new Collection Label. The current presentation object will be aligned with the specified Collection Label.
- When you generate reports on the selected Collection Label, this presentation object will be included in the report.


Viewing and Managing the List of Presentation Objects Aligned with a Collection Label

From the **Collection Labels** page, you can view information about each Collection Label. For each Collection Label, you can view a list of presentation objects aligned with that Collection Label. To view this information:

- Go to the **Collection Labels** page (System > Manage > Collection Labels).

	Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates
1.	In Use	In Use	Video Performance	No	2	--	--
2.	Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--
3.	Max Jitter	Max Jitter	Video Performance	No	2	--	--
4.	Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--
5.	Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--
6.	Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--
7.	Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--
8.	Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--
9.	Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--
10.	Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--
11.	Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--
12.	Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--
13.	Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--
14.	Usage	Usage	Video Performance	No	2	--	--
15.	CPU		Vitals	No	37	9	--
16.	Memory		Vitals	No	17	--	--
17.	Swap		Vitals	No	6	--	--

- Find the Collection Label you are interested in. In the **Aligned Presentations** column, select the pencil icon (). The **Aligned Presentations** modal page appears:




	Dynamic Application Name	Presentation Name	Number Devices	Precedence	
1.	NetApp: System Stats 7-Mode	Average Processor Busy	--	0	<input type="checkbox"/>
2.	AWS RDS Instance Performance	CPU Utilization	--	50	<input type="checkbox"/>
3.	Cisco: CPU	CPU 5 minute average percent	--	50	<input type="checkbox"/>
4.	Cisco: CPU Old	CPU Percentage	--	50	<input type="checkbox"/>
5.	Cisco: Nexus: CPU Performance	CPU Utilization	--	50	<input type="checkbox"/>
6.	Cisco: PIX CPU	CPU Percentage	--	50	<input type="checkbox"/>
7.	Cisco: PIX CPU	CPU Percentage - Old	--	50	<input type="checkbox"/>
8.	Cloudkick: CPU	I/O Wait	--	50	<input type="checkbox"/>
9.	Cloudkick: CPU	Steal	--	50	<input type="checkbox"/>
10.	Cloudkick: CPU	User	--	50	<input type="checkbox"/>
11.	Cloudkick: CPU	CPU Idle	--	50	<input type="checkbox"/>
12.	Cloudkick: CPU	IRQ	--	50	<input type="checkbox"/>
13.	Cloudkick: CPU	System	--	50	<input type="checkbox"/>
14.	Cloudkick: CPU	CPU Utilization	--	50	<input type="checkbox"/>
15.	Dell Switch: CPU Collection (Via Telnet)	CPU 5 minutes	--	50	<input type="checkbox"/>
16.	EMC Storage Processor Performance	Percent Busy	--	50	<input type="checkbox"/>
17.	Force 10: CPU Utilization	chStackUnitCpuUtil5Min	--	50	<input type="checkbox"/>
18.	Fortinet: CPU	CPU Percentage	--	50	<input type="checkbox"/>
19.	Foundry: CPU	CPU Percentage	--	50	<input type="checkbox"/>
20.	Host Resource: CPU	CPU Average	--	50	<input type="checkbox"/>
21.	HP UX: CPU	Overall CPU	--	50	<input type="checkbox"/>
22.	Juniper: CPU	CPU Percentage	--	50	<input type="checkbox"/>
23.	Microsoft HyperV Guest CPU Load	LoadPercentage	--	50	<input type="checkbox"/>
24.	Microsoft: Azure Virtual Machine CPU Performance	CPU Utilization	9	50	<input type="checkbox"/>
25.	Microsoft: Windows Server CPU Performance	CPU Utilization	--	50	<input type="checkbox"/>
26.	Net-SNMP: CPU	Overall CPU	1	50	<input type="checkbox"/>
27.	NetApp: Svsstem Stats C-Mode	Average Processor Busy	--	50	<input type="checkbox"/>

- In the **Aligned Presentations** modal page, you can view information about the presentation objects aligned with the current Collection Label and perform actions to manage those presentation objects. You can also [unalign a presentation object](#) from a Collection Label and [change the precedence](#) for one or more Collection Label-presentation object pairs.

To globally unalign a presentation object from a Collection Label:

1. In the **Aligned Presentations** modal page, find the presentation object that you want to unalign from the Collection Label and select its checkbox.




2. From the **Select Action** field in the lower right, select *Unalign from Label*. Select the Go button.
3. The selected presentation object will no longer be associated with the Collection Label.

For each Collection Label-presentation object pair, you can define precedence. For example, suppose that both the "Cisco: CPU" Dynamic Application and the "Host Resource: CPU" include a presentation object that is aligned with the **CPU** Collection Label. You can define precedence to specify priority for each presentation object associated with a Collection Label.

Collection Group / Collection Label	Presentation Object	Dynamic Application
Vitals / CPU	CPU Average	Host Resource: CPU
Vitals / CPI	CPU 5 minutes average percent	Cisco: CPU

To set the precedence for the Collection Label (in our example, "CPU"):

1. The **Aligned Presentations** modal page displays all the presentation objects associated with the selected Collection Label. By default, each presentation object has a precedence of 50.



2. In the **Aligned Presentations** modal page, you can edit precedence in two ways:
 - In the **Precedence** column, use the up arrow and down arrow to change the value for a single presentation object. Repeat for each presentation object for which you want to edit precedence.
 - Select the checkbox of one or more presentation objects. In the **Select Action** field, select *Change Precedence* and a value. Select the **[Go]** button. Each selected presentation object will be assigned the new (and identical) precedence value.
3. Repeat steps 2–4 for each Presentation Object for which you want to edit the precedence value.

NOTE: The precedence values you define in the **Aligned Presentations** modal page override the precedence value you set per presentation object in the **Presentation Objects** page.

Viewing and Editing Duplicate Presentation Objects by Collection Label

You can view a list of devices where duplicates occur, view how the ScienceLogic platform assigned the Collection Label-presentation object pair, and edit the Collection Label-presentation object pair for one or more devices. When you manually define a Collection Label-presentation object pair for a device, the platform will not edit or change that pair.

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).

Collection Labels | Collection Labels Found [17] Reset Guide

	Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates
1	In Use	In Use	Video Performance	No	2	--	--
2	Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--
3	Max Jitter	Max Jitter	Video Performance	No	2	--	--
4	Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--
5	Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--
6	Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--
7	Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--
8	Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--
9	Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--
10	Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--
11	Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--
12	Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--
13	Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--
14	Usage	Usage	Video Performance	No	2	--	--
15	CPU		Vitals	No	37	10	1
16	Memory		Vitals	No	17	1	--
17	Swap		Vitals	No	6	1	--
					--	--	--

- Find the Collection Label you are interested in. In the **Duplicates** column, select the pencil icon (✎). The **Duplicates** modal page appears.

DID	Device Name	Organization	Device IP Address	Device Class Sub-class	Dynamic Application Name	Presentation Name	Aligned
677	em7_ao	System	10.100.100.7	ScienceLogic, Inc. OEM	Net-SNMP: CPU	Overall CPU	Yes <input checked="" type="radio"/>
677	em7_ao	System	10.100.100.7	ScienceLogic, Inc. OEM	Host Resource: CPU	CPU Average	No <input type="radio"/>

[Select Action]

- In the **Duplicates** modal page, you can view a list of devices for which there are multiple possible Collection Label-presentation object pairs. You can view which pair is currently assigned to the device.
- To change the pair for a device, click on the pair's radio button.
- Repeat step #4 for each device on which you want to edit the duplicate.
- In the **Select Action** field (in the lower right), select *Align Presentation for Device*. Select the **[Go]** button.
- Each edited device will now use the selected Collection Label-presentation object pair.

Viewing and Managing the List of Devices Aligned with a Collection Label

From the **Collection Labels** page, you can view information about each Collection Label. For each Collection Label, you can view a list of devices from which the ScienceLogic platform is collecting values. To view this information:

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).

Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates
In Use	In Use	Video Performance	No	2	--	--
Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--
Max Jitter	Max Jitter	Video Performance	No	2	--	--
Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--
Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--
Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--
Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--
Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--
Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--
Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--
Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--
Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--
Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--
Usage	Usage	Video Performance	No	2	--	--
CPU		Vitals	No	37	10	--
Memory		Vitals	No	17	1	--
Swap		Vitals	No	6	1	--

2. Find the Collection Label you are interested in. In the **Aligned Devices** column, select the pencil icon (✎).

Device Name	Organization	Device IP Address	Device Class Sub-class	Dynamic App Name	Aligned Presentation
2008r2-2	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
Azure-Team-DC1	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
Azure-Team-DC2	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
azureteam-vm2	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
em7_ao	System	10.100.100.7	ScienceLogic, Inc. OEM	Net-SNMP: CPU	Overall CPU
vm-2008r1-tmp	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
vm-temp-006	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
vm-temp-201	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
vm-tmp-1	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization
vm-tmp-100	Azure	--	Microsoft Azure Virtual Mac	Microsoft: Azure Virtual Mac	CPU Utilization

3. In the **Aligned Devices** modal page, you can view information about the devices that are aligned with the current Collection Label and perform actions to manage those devices.

For devices that include duplicates, you can reset the presentation object for one or more devices. When you manually define a Collection Label-presentation object pair for a device, the ScienceLogic platform will not edit or change that pair.

1. In the **Aligned Devices** modal page, select the checkbox for one or more devices for which you want to change the Collection Label-presentation object pair.

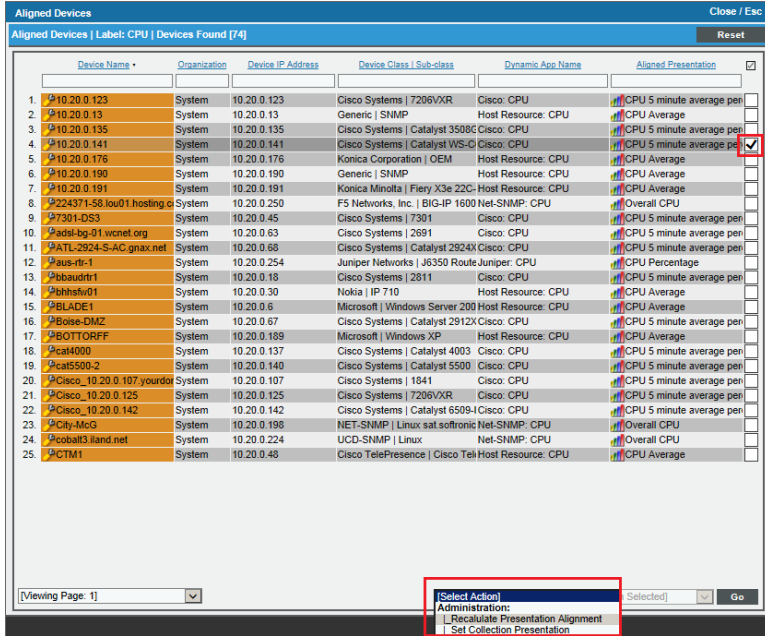
Device Name	Organization	Device IP Address	Device Class Sub-class	Dynamic App Name	Aligned Presentation
10.20.0.123	System	10.20.0.123	Cisco Systems 7206VXR	Cisco: CPU	CPU 5 minute average per
10.20.0.13	System	10.20.0.13	Generic SNMP	Host Resource: CPU	CPU Average
10.20.0.135	System	10.20.0.135	Cisco Systems Catalyst 3508C	Cisco: CPU	CPU 5 minute average per
10.20.0.141	System	10.20.0.141	Cisco Systems Catalyst WS-C	Cisco: CPU	CPU 5 minute average per
10.20.0.176	System	10.20.0.176	Konica Corporation OEM	Host Resource: CPU	CPU Average
10.20.0.190	System	10.20.0.190	Generic SNMP	Host Resource: CPU	CPU Average
10.20.0.191	System	10.20.0.191	Konica Minolta Fiery X3e 22C	Host Resource: CPU	CPU Average
224371-56.lou01.hosting.c	System	10.20.0.250	F5 Networks, Inc. BIG-IP 1600	Net-SNMP: CPU	Overall CPU
7301-DS3	System	10.20.0.45	Cisco Systems 7301	Cisco: CPU	CPU 5 minute average per
addl-bg-01.wcnet.org	System	10.20.0.63	Cisco Systems 2691	Cisco: CPU	CPU 5 minute average per
ATL-2924-S-AC.gnax.net	System	10.20.0.68	Cisco Systems Catalyst 2924X	Cisco: CPU	CPU 5 minute average per
aus-ri-1	System	10.20.0.254	Juniper Networks J6350 Route	Juniper: CPU	CPU Percentage
abaudrfr1	System	10.20.0.18	Cisco Systems 2611	Cisco: CPU	CPU 5 minute average per
bbhshv01	System	10.20.0.30	Nokia IP 710	Host Resource: CPU	CPU Average
BLADE1	System	10.20.0.6	Microsoft Windows Server 200	Host Resource: CPU	CPU Average
Boise-DMZ	System	10.20.0.67	Cisco Systems Catalyst 2912X	Cisco: CPU	CPU 5 minute average per
BOTTORFF	System	10.20.0.189	Microsoft Windows XP	Host Resource: CPU	CPU Average
pcat4000	System	10.20.0.137	Cisco Systems Catalyst 4003	Cisco: CPU	CPU 5 minute average per
pcat5500-2	System	10.20.0.140	Cisco Systems Catalyst 5500	Cisco: CPU	CPU 5 minute average per
Cisco_10.20.0.107.yourdo	System	10.20.0.107	Cisco Systems 1841	Cisco: CPU	CPU 5 minute average per
Cisco_10.20.0.125	System	10.20.0.125	Cisco Systems 7206VXR	Cisco: CPU	CPU 5 minute average per
Cisco_10.20.0.142	System	10.20.0.142	Cisco Systems Catalyst 6509-I	Cisco: CPU	CPU 5 minute average per
City-MCG	System	10.20.0.198	NET-SNMP Linux sat softronic	Net-SNMP: CPU	Overall CPU
acobbals.liand.net	System	10.20.0.224	UCD-SNMP Linux	Net-SNMP: CPU	Overall CPU
ACTM1	System	10.20.0.48	Cisco TelePresence Cisco Tel	Host Resource: CPU	CPU Average

2. In the menus in the lower right, select **Set Collection Presentation** and then select the presentation object. Select the **[Go]** button.

For devices that include duplicates, you can clear all current settings, including manual settings. The ScienceLogic platform will then automatically evaluate the precedence for each possible presentation object and assign the Collection Label-presentation object pair with the lowest precedence.

To clear the current Collection Label-presentation object pair for one or more devices:

1. In the **Aligned Devices** modal page, select the checkbox for one or more devices for which you want to clear the aligned presentation object.



2. In the menus in the lower right, select **Recalculate Presentation Alignment**. Select the **[Go]** button.
3. The ScienceLogic platform will evaluate the precedence of each possible presentation object and assign the presentation object with the lowest precedence.

Editing Duplicate Presentation Objects by Device

You can view a list of devices where duplicates occur, view how the ScienceLogic platform assigned the Collection Label-presentation object pair, and edit the Collection Label-presentation object pair for one or more selected devices. When you manually define a Collection Label-presentation object pair for a device, the platform will not edit or change that pair:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Select the checkbox for each device you are interested in.

Device Name	Device Hostname	IP Address	Device Category	Device Class / Sub-class	DID	Organization	Current State	Collection Group	Collection State	SNMP Capable	SNMP Access
10.100.100.40	--	10.100.100.40	Pingable	Ping ICMP	274	System	Healthy	CUG	Active	--	--
10.100.100.48	--	10.100.100.48	Pingable	FreeBSD ICMP	294	System	Notice	CUG	Active	--	--
10.100.100.10	--	--	Servers.VMware	Virtual Machine	348	System	Healthy	CUG	Active	--	--
10.100.100.8	--	--	Servers.VMware	Virtual Machine	353	System	Healthy	CUG	Active	--	--
10.100.100.11	--	--	Servers.VMware	Virtual Machine	349	System	Healthy	CUG	Active	--	--
Cluster_1	--	--	Virtual.Infrastr	Network	341	System	Healthy	CUG	Active	--	--
CU-Device	--	--	Virtual.Infrastr	Network	339	System	Healthy	CUG	Active	--	--
Databases	--	--	Virtual.Infrastr	Folder	334	System	Healthy	CUG	Active	--	--
902_svc_PRODUCION_1002	--	--	Servers.VMware	Virtual Machine	348	System	Healthy	CUG	Active	--	--
em7_7302_test1	--	192.168.33.51	System.EM7	ScienceLogic, Inc. EM7 Data Collector	299	System	Healthy	CUG	Active	EM7 Default V2	V2
em7_7302_test2	--	192.168.33.52	System.EM7	ScienceLogic, Inc. EM7 Database	288	System	Healthy	CUG	Active	EM7 Default V2	V2
em7_7302_test3	--	192.168.33.50	System.EM7	ScienceLogic, Inc. EM7 Database	287	System	Minor	CUG	Active	EM7 Default V2	V2
GW_CU_100_15	--	--	Servers.VMware	Virtual Machine	358	System	Healthy	CUG	Active	--	--
SQLSRV_100_13	--	--	Servers.VMware	Virtual Machine	354	System	Healthy	CUG	Active	--	--
ha-datascenter	--	--	Virtual.Infrastr	Datacenter	332	System	Healthy	CUG	Active	--	--
Hosts	--	--	Virtual.Infrastr	VMware Folder	333	System	Healthy	CUG	Active	--	--
Hughes_AID_10_100_100_9	--	--	Servers.VMware	Virtual Machine	344	System	Major	CUG	Active	--	--
SQL_100_40	--	--	Servers.VMware	Virtual Machine	350	System	Healthy	CUG	Active	--	--
Networks	--	--	Virtual.Infrastr	Folder	335	System	Healthy	CUG	Active	--	--
OS-ISOs	--	--	Virtual.Infrastr	Datastore	337	System	Healthy	CUG	Active	--	--
qa_pm	--	--	Virtual.Infrastr	Network	340	System	Healthy	CUG	Active	--	--
SUSE_10_100_35	--	--	Servers.VMware	Virtual Machine	345	System	Major	CUG	Unavailable	--	--
SUSE_10_100_38	--	--	Servers.VMware	Virtual Machine	351	System	Major	CUG	Unavailable	--	--
SUSE_11_100_30	--	--	Servers.VMware	Virtual Machine	343	System	Major	CUG	Unavailable	--	--
SUSE_11_NS	--	--	Servers.VMware	Virtual Machine	352	System	Major	CUG	Unavailable	--	--
TestVM15	--	--	Virtual	KVM Virtual Machine	331	System	Healthy	CUG	Active	--	--
TRAIN-VM-01.sciencelogic.io	--	--	Servers.VMware	Host Server	336	System	Major	CUG	Active	--	--
TRAIN-VM-01_d801	--	--	Virtual.Infrastr	Datastore	335	System	Major	CUG	Active	--	--
UCSPE_10.100.100.21 (2)	--	--	Servers.VMware	Virtual Machine	355	System	Major	CUG	Unavailable	--	--
VM Network	--	--	Virtual.Infrastr	Network	342	System	Healthy	CUG	Active	--	--
WIN2012-21.DDCS.LOCAL	--	10.100.100.21	Servers	Microsoft Windows Server 2012 R2	325	System	Healthy	CUG	Active	--	--
WIN2012-22.DDCS.LOCAL	--	10.100.100.22	Servers	Microsoft Windows Server 2012 R2	327	System	Healthy	CUG	Active	--	--
WIN2012-23.DDCS.LOCAL	--	10.100.100.23	Servers	Microsoft Windows Server 2012 R2	329	System	Major	CUG	Active	--	--
WIN2012_100_21	--	--	Servers.VMware	Virtual Machine	347	System	Healthy	CUG	Active	--	--
WIN2012_100_22	--	--	Servers.VMware	Virtual Machine	350	System	Healthy	CUG	Active	--	--
WIN2012_100_23	--	--	Servers.VMware	Virtual Machine	357	System	Healthy	CUG	Active	--	--

3. If you want to view a list of duplicates for all possible devices, select the red check-box in the top row of the page. This selects all devices.

- In the **Select Action** field (lower right), select **FIND Collection Label Duplicates**. Select the **[Go]** button.

The screenshot shows a window titled "Current Duplicates" with a "Close / Esc" button in the top right. Below the title bar, there is a header area with "Current Duplicates | Label: [Vitals: CPU] | Records Found [2]" and a "Reset" button. The main area contains a table with the following columns: DID, Device Name, Organization, Device IP Address, Device Class | Sub-class, Dynamic Application Name, Presentation Name, and Aligned. Two rows of data are visible, both for device "CTM1" with IP "10.20.0.48". The first row is for "Net-SNMP: CPU" with "Overall CPU" presentation and is not aligned. The second row is for "Host Resource: CPU" with "CPU Average" presentation and is aligned. At the bottom right, there is a "[Select Action]" dropdown menu and a "Go" button.

DID	Device Name	Organization	Device IP Address	Device Class Sub-class	Dynamic Application Name	Presentation Name	Aligned
28	CTM1	System	10.20.0.48	Cisco TelePresence Cisco	Net-SNMP: CPU	Overall CPU	No
28	CTM1	System	10.20.0.48	Cisco TelePresence Cisco	Host Resource: CPU	CPU Average	Yes

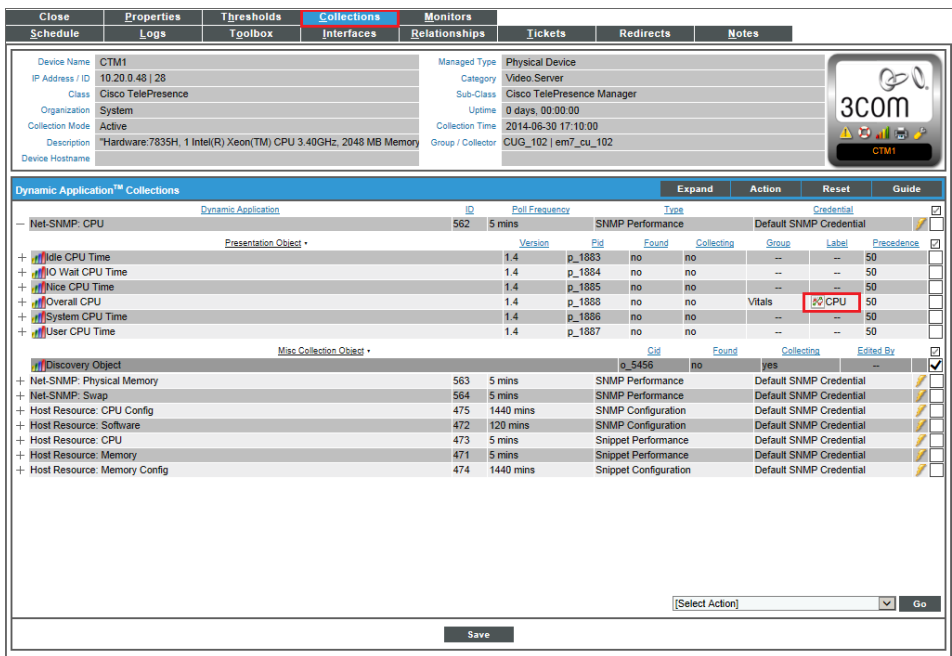
- The **Current Duplicates** page is displayed. For each device, you can edit the presentation object that is aligned with a Collection Label.
 - To select a Collection Label, use the drop-down list in the upper left.
 - To change the aligned presentation object for one or more devices:
 - Click on the radio button for the desired presentation object for the device.
 - For each additional device you want to edit, click on the radio button for the desired presentation object.
 - In the **Select Action** menu (lower right), select **Align Presentation for Device**. Select the **[Go]** button.

Editing Duplicate Presentation Objects for a Single Device

You can edit the Collection Label-presentation object pair for a single device. If a single device includes duplicate Collection Label-presentation object pairs, you can specify which one the ScienceLogic platform should use for that device.

To edit the Collection Label-presentation object pairs for a single device:

1. Go to the **Device Manager** page (Registry > Devices > Device Manager).
2. Find the device you want to edit. Select its wrench icon (🔧).
3. Select the **[Collections]** tab. In the **Dynamic Application Collections** page, click on the plus signs (+) to expand each Dynamic Application.



4. You will notice that some presentation objects include the chart icon in the **Label** column. These presentation objects are duplicates that are not currently aligned with a Collection Label. If you want to align one of these presentation objects with the Collection Label (instead of the current alignment), click on the chart icon.
5. You will be prompted before the ScienceLogic platform aligns the presentation object with the Collection Label. After approving, you will notice that a new presentation object now displays a chart icon in its **Label** column. This is because this presentation object is no longer associated with a Collection Label.

Editing a Collection Label

You can edit a Collection Label from the **Collection Labels** page (System > Manage > Collection Labels). To do so:

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).

	Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates
1	In Use	In Use	Video Performance	No	2	--	--
2	Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--
3	Max Jitter	Max Jitter	Video Performance	No	2	--	--
4	Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--
5	Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--
6	Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--
7	Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--
8	Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--
9	Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--
10	Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--
11	Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--
12	Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--
13	Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--
14	Usage	Usage	Video Performance	No	2	--	--
15	CPU		Vitals	No	37	9	--
16	Memory		Vitals	No	17	--	--
17	Swap		Vitals	No	6	--	--

2. Find the Collection Label you want to edit. Select its wrench icon (🔧).
3. You can edit one or more of the following:


- **Label Name**. Name of the Collection Label. This field is required.
- **Label Description**. Description of the Collection Label. This field is optional.
- **Group Name**. Collection Group to align with the Collection Label. You can select from a list of existing Collection Groups or enter the name of a new Collection Group. This field is required.
- **Frequent Data**. Specifies whether *frequently rolled up data* is calculated for the Collection Label. If the Collection Label will include data that is collected every five minutes or more frequently, and you require that dashboard data be updated every 15 minutes or 20 minutes, select Yes in this field. This data is available immediately for use in a collection label.
- **Save icon** (📌). Select this icon to save your changes.
















Deleting a Collection Label

You can delete a Collection Label from the **Collection Labels** page (System > Manage > Collection Labels) only if the Collection Label has no **Aligned Presentations**. To delete a Collection Label:

NOTE: You can delete a Collection Label only if no presentation objects are aligned with that label.

1. Go to the **Collection Labels** page (System > Manage > Collection Labels).



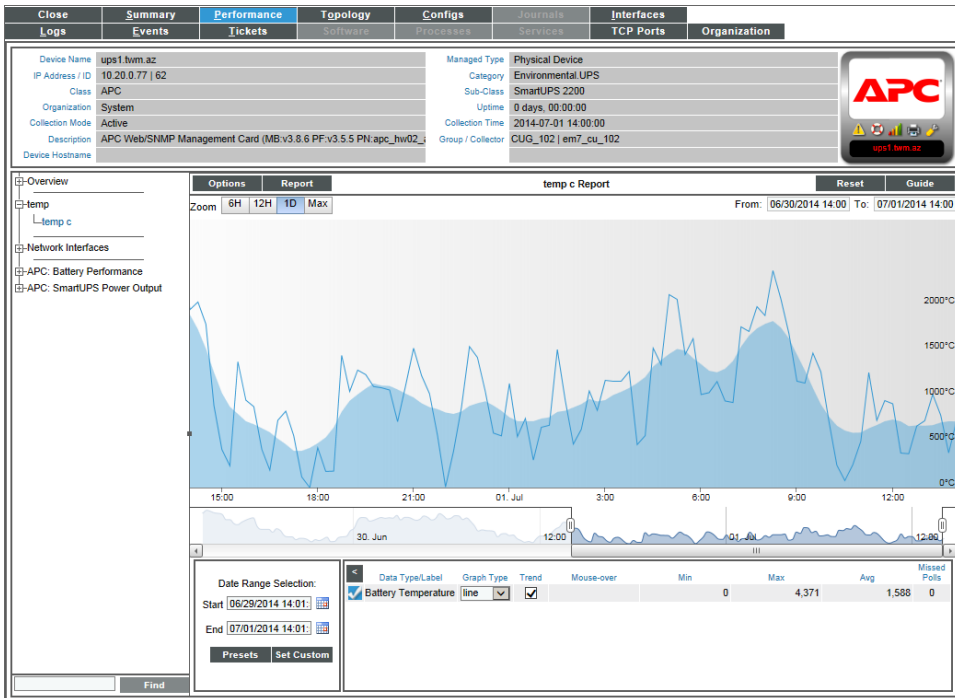
Collection Labels Collection Labels Found [18]								Reset	Guide
Label Name	Label Description	Group Name	Frequent Data	Aligned Presentations	Aligned Devices	Duplicates			
1 Example	Example	Video Performance	No	--	--	--			
2 In Use	In Use	Video Performance	No	2	--	--			
3 Max % Packet Loss	Max % Packet Loss	Video Performance	No	2	--	--			
4 Max Jitter	Max Jitter	Video Performance	No	2	--	--			
5 Rx Audio Jitter	Receive Audio Jitter	Video Performance	No	2	--	--			
6 Rx Audio Pkts Lost	Receive Audio Packets Lost	Video Performance	No	2	--	--			
7 Rx Total BW	Total Incoming BW	Video Performance	No	2	--	--			
8 Rx Video Jitter	Rx Video Jitter	Video Performance	No	2	--	--			
9 Rx Video Pkts Lost	Rx Video packets lost	Video Performance	No	2	--	--			
10 Tx Audio Jitter	Transmit Audio Jitter	Video Performance	No	2	--	--			
11 Tx Audio Pkts Lost	Transmit Audio Packets Lost	Video Performance	No	2	--	--			
12 Tx Total BW	Total Outgoing Bandwidth	Video Performance	No	2	--	--			
13 Tx Video Jitter	Outgoing Video Jitter	Video Performance	No	2	--	--			
14 Tx Video Pkts Lost	Transmit Video Packets Lost	Video Performance	No	2	--	--			
15 Usage	Usage	Video Performance	No	2	--	--			
16 CPU		Vitals	No	37	10	--			
17 Memory		Vitals	No	17	1	--			
18 Swap		Vitals	No	6	1	--			

2. Find the Collection Label you want to delete.
3. Select its bomb icon (🧨).
4. The Collection Label will be deleted from the ScienceLogic platform.

Viewing Reports About Collection Labels on a Single Device

For each device in the ScienceLogic platform, the **Device Performance** page displays time-series graphs about the data collected from that device.

If a device subscribes to a Dynamic Application that includes Collection Labels, the platform will display the Collection Group in the left pane of the **Device Performance** page. You can expand the Collection Group and select a Collection Label.



The graph for a Collection Label displays collected values on the Y-axis and time on the X-axis.

Viewing Dashboards About Collection Labels

You can use the following dashboard widgets to include data associated with Collection Labels in a dashboard:

- Multi-Series Performance Widget
- Leaderboard / Top-N Widget
- Gauge / Meter

For details on each widget, see the **Dashboards** manual.

© 2003 - 2018, 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