



User's Guide

Microsoft® Systems Center Operations Manager Management Pack for InfraStruxure® Central

SFSCOM

Contents

- Introduction 1**
- Integration Concepts 1**
 - Inventory Integration (Service Model) 1
 - Monitoring Integration (Health Model) 2
 - Console Integration 4
- Initial Setup 5**
 - Installation 5
 - Changing the Default Management Pack Behavior Using an Overrides Management Pack 6
 - Registering InfraStruxure Central Servers on Proxy Health Services 7
 - Registry Use in Alarm Collection 8
 - Creating Run As Accounts and Assigning Them to the Run As Profile 9
 - Partitioning Large InfraStruxure Central Servers across Multiple Health Services 9
 - Setting the Environment Variable ISXCU! 10
- Topics Related to Operations 11**
 - Inventory Discovery 11
 - Inventory Discovery Segments and Access to the %TEMP% Folder 11
 - Warning Events on Initial Inventory Discovery 12
 - International Characters in Inventoried Properties 12
 - Drilling Down on Distributed Application (DA) Diagrams 12
 - Special Parameters for Monitor Overrides 12
 - Optimizations to Reduce Impact and the Effect of Overrides . . . 13
 - Script Events Generated by the Management Pack 16
- Appendices 17**
 - Appendix 1 – Included Sensor Types 17

Introduction

The APC InfraStruxure Central server is a network appliance that collects data from and monitors status of devices that support data center critical infrastructure.

The Microsoft System Center Operations Manager Management Pack provides a tight integration between the physical infrastructure space for which InfraStruxure Central provides unprecedented visibility and control and the systems infrastructure space for which Operations Manager provides the same type of solution.

Events from the physical infrastructure space can be viewed side-by-side with the system events to which they may be related, providing an ability to predict, respond to, and prevent issues in the physical infrastructure from affecting system and application availability.

The Management Pack integrates with InfraStruxure Central version 5.1 and higher through the new Web Services interface.

Integration Concepts

Inventory Integration (Service Model)

The Management Pack integrates the inventory of InfraStruxure Central servers in your environment and the inventory that they host using a two-part discovery process. The first part discovers all InfraStruxure Central servers that are registered on a given Health Service and is called the Registration Discovery. The second part discovers the inventory of each InfraStruxure Central server.

In this document, we will make frequent references to “Health Service”. The initial discovery is actually targeted at computers running Windows Server (the `Microsoft.Windows.Server.Computer.Operations Manager` class type), but it is the Health Service running on that server that is the primary component that performs the functions dictated by this Management Pack.

The inventory discovered is defined as follows:

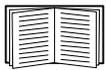
- For each Health Service for which the Registration Discovery is enabled:
 - An object representing each registered InfraStruxure Central server. Each object becomes the target of several monitoring rules, the host for all devices, and the container for all device groups.
 - An object representing each registered InfraStruxure Central server (key properties only) as a Component Group for the By Device Group Dynamic Distributed Application.
 - An object representing each registered InfraStruxure Central server (key properties only) as a Component Group for the By Device Type Dynamic Distributed Application.
- For each discovered InfraStruxure Central server:
 - Objects representing every device group, hosted device, and sensor.
 - Dynamically-created device groups that are derived from the system group type.



Note: Device groups can be used for overrides, but cannot be edited through the user interface.

- For the management group:
 - A group that contains all InfraStruxure Central servers registered on all Health Services.
 - A distributed application root object that represents the hierarchy by device groups: the root object, which contains component groups for each InfraStruxure Central server. Each component group contains device groups discovered for that InfraStruxure Central server.
 - A distributed application root object that represents the hierarchy by device types: the root object, which contains component groups for all InfraStruxure Central servers registered on all Health Services, which contain component groups representing each distinct device type for which at least one device was discovered for that InfraStruxure Central server, which contain the devices discovered for that InfraStruxure Central servers that are of that particular type.
 - The dynamic distributed applications cannot be edited using the Distributed Application Designer.

The sensor types discovered by the Management Pack represent a subset of sensor types that may be found on an InfraStruxure Central server.



See “Appendix 1 – Included Sensor Types” on page 17 for a complete list of the included sensor types.

For any workflow that targets sensors instead of devices, only sensor types that are included in the Management Pack will be monitored by the sensor rule. Sensor types not included in the Management Pack will be covered by the device rule. For instance, there are two rules for the overall alarm health state of an entity (device or sensor). For sensor types that are discovered, their overall alarm health state will be monitored by the unit monitor that targets the sensors.

For sensor types that are not discovered, their overall alarm health state will be reflected in the unit monitor that targets the device that hosts that sensor. This applies to any monitor or rule that monitors a trait of a sensor that is also reflected in the device (alarm state, overall state, etc.) Some sensor traits (e.g. data points) are not reflected in the device and thus are not available for sensor types not included in the Management Pack.

Monitoring Integration (Health Model)

The Management Pack defines several integration points of the health model of an InfraStruxure Central server. In general, all monitors and rules are configured to run every 5 minutes.

- For Devices
 - Each device is monitored for its overall device state. This is an attribute of the device itself whose value is mapped to one of the three Operations Manager health states, Healthy, Warning, and Critical.

- Each device is monitored for its overall alarm health state. This is an indicator of the overall state of the active alarms for a device. This monitor will generally follow the health state of the most recent, most severe, active alarm for a device. This is an overall health monitor, however, and does not map every alarm to a monitor state transition. See the Product Knowledge article defined for the monitor in the Management Pack for more details.
- Each device is also targeted by a pair of rules that collect alarms that have been generated since the last collection and perform these actions:
 - Insert them into the Operations Manager operational database and data warehouse for display in event views and reports.
 - Optionally generate an alert for each alarm. This rule is disabled by default.
- For Sensors
 - Each sensor is monitored for its overall sensor state. This is an attribute of the sensor itself whose value is mapped to one of the three Operations Manager health states, Healthy, Warning, and Critical.
 - Each sensor is monitored for its overall alarm health state. This is an indicator of the overall state of the active alarms for a sensor. This monitor will generally follow the health state of the most recent, most severe, active alarm for a sensor. This is an overall health monitor, however, and does not map every alarm to a monitor state transition. See the Product Knowledge article defined for the monitor in the Management Pack for more details.
 - Each sensor is also targeted by a pair of rules that collect alarms that have been generated since the last collection and:
 - Inserts them into the Operations Manager operational database and data warehouse for display in event views and reports.
 - Optionally generates an alert for each alarm. This rule is disabled by default.
 - All sensor types except state sensors are also targeted by a rule that collects the data point for that sensor and inserts it into the operational database and data warehouse for display in performance views or reports.
 - Each state sensor is targeted by a monitor that reads the sensor data point and maps it to a health state. There are two variants of this monitor:
 - **Pre-Defined:** This monitor will only have an active health state for sensor types whose data point represents a state that is well-defined and has a fixed mapping of its state values to Operations Manager health states.
 - **User-Defined:** This monitor will only have an active health state for sensor types whose data point represents a state that is undefined outside of a user-provided context. The user can establish regular expression patterns to determine which states should be considered Healthy, Warning, and Normal. An example of such a state sensor is a generic switch sensor. Depending on the switch and the application, a state of open, closed, etc. could be normal or cause for warning or a critical condition. This monitor is disabled by default, as no meaningful state can be established as the default.

- Dependency Monitors
 - Each InfraStruxure Central server has a dependency monitor that gathers the state of the devices it hosts and another dependency monitor that gathers the state of the device groups it contains.
 - Each device group has a dependency monitor that gathers the state of the devices in the group.
 - Each device has a dependency monitor that gathers the state of the sensors hosted by the device.
 - The management-group-wide group that contains all InfraStruxure Central servers has a dependency monitor that gathers the state of each InfraStruxure Central server.
 - Both dynamic DAs have dependency monitors that gathers the health of the component groups at lower levels of the hierarchy to their parent, up to the root object.

Console Integration

The Management Pack defines several views and tasks to provide the required window into the InfraStruxure Central server's inventory state.

- Views
 - A state view that shows the state of all registered InfraStruxure Central servers on all Health Services is provided. State detail columns are provided for the hosted device state and the contained device group state.
 - A state view that shows the state of all devices discovered on all InfraStruxure Central servers on all Health Services is provided. State detail columns are provided for the hosted sensor state (for state-bearing sensors only).
 - Three diagram views are provided with the All InfraStruxure Centrals group, the root object of the By Device Group DA, and the root object of the By Device Type DA as their root object, respectively.
 - A performance view folder is provided with pre-defined performance views for each data-bearing sensor class type.
 - An environmental dashboard view is provided that displays a 2x2 matrix of the most common performance views.

- Tasks
 - For an InfraStruxure Central server
 - Tasks to launch the InfraStruxure Central client or a web browser to the InfraStruxure Central server.
 - For a device:
 - Tasks to launch the InfraStruxure Central client, a Web browser to the InfraStruxure Central server, or a Web browser to the device itself. A special proxy URL is automatically selected to enable devices on the private network of an InfraStruxure Central server to be browsed.

– For a sensor:

Tasks to launch the InfraStruxure Central client, a Web browser to the InfraStruxure Central server, or a Web browser to the device that hosts the sensor. A special proxy URL is automatically selected to enable devices on the private network of an InfraStruxure Central server to be contacted.

- Images

– Wherever possible, images have been provided for the class types discovered by the Management Pack. This includes:

- Images for each device type
- Images for each sensor type
- Images for the InfraStruxure Central server
- Images for device groups
- Images for the dynamic DA Component Groups
- Images for the All ISX Centrals group and DA root objects

Initial Setup

Installation

To install the Management Pack, install the .msi package downloaded from the APC Web site (<http://www.apc.com/tools/download>) into an appropriate directory. The installation package consists of four files: this document, the sealed Management Pack (.mp), the optional Overrides Management Pack (.xml), and the registration utility (.exe).

Only the sealed Management Pack is required for import. Perform the installation on a system running the Operations Manager console. The Management Pack references only built-in Operations Manager Management Packs, but requires at least Operations Manager 2007 SP1.

The registration utility is run on each Health Service that will act as a proxy between Operations Manager and InfraStruxure Central. The registration utility is a tool that provides a user interface to set certain registry keys that are used by the Management Pack to perform the initial discovery of the InfraStruxure Central servers that should be managed from that Health Service. Therefore, the installation should also be performed on the systems that will act as proxy, unless you intend to manage these registry keys by some other means.



See “Registering InfraStruxure Central Servers on Proxy Health Services” on page 7 for more information.

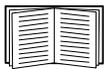
Changing the Default Management Pack Behavior Using an Overrides Management Pack

The Management Pack installs in a dormant state and with a minimum monitoring configuration enabled by default. An optional unsealed Overrides Management Pack is included in the installation package; it enables everything defined in the Management Pack. Activating the Overrides is acceptable in a test and evaluation environment; however, it is imperative that the activities of the Management Pack are completely understood before using this configuration in a production environment. Follow this procedure to control the introduction of this Management Pack into your environment:

1. Import this Management Pack into a test environment.
2. Create a custom Management Pack for overrides or use the included optional Overrides Management Pack.
3. Create a group for the entities targeted by each set of rules you wish to enable or disable.
4. Create overrides that target the aforementioned groups.
5. Export this custom Management Pack.
6. Import the sealed Management Pack with the custom Management Pack into the production environment.
7. Populate the groups as required.

Specifically, the following components are disabled by default:

- **The InfraStruxure Central Registration discovery.** This discovery targets computers running Windows Server (i.e. the `Microsoft.Windows.Server.Computer Operations Manager` class) and looks for the registry keys that are written by the registration utility. This rule is disabled by default, which renders the Management Pack dormant on import. It is recommended that you create a group containing the computers running Windows Server in the overrides Management Pack and enable the rule for that group.
- **The various data collection rules.** These rules collect data points from various sensor types discovered and insert them into the operational database and data warehouse. Because they can generate large amounts of data depending on the environment, they are disabled by default. It is recommended that you create a group containing sensors in the overrides Management Pack and enable the desired rules for that group.



See “Optimizations to Reduce Impact and the Effect of Overrides” on page 13 for an important discussion regarding the ramifications of certain overrides for these rules.

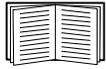
- **The alarm-to-alert mapping rules.** These rules collect all alarms that are raised on an InfraStruxure Central server and map them to alerts in the Operations Manager console. Because these rules can generate a high volume of alerts depending on the environment and because these alerts must be closed manually, the rules are disabled by default. There are two rules: one for devices and one for sensors. Two other rules that mirror these rules collect the alarms and generate Operations Manager events for every alarm raised on an InfraStruxure Central server. These rules are enabled by default; it is only the alert rules that are disabled by default. It is strongly recommended that your create a group containing devices and/or sensors (or multiple groups) in the overrides Management Pack and enable the desired rules for those groups.



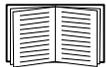
See “Optimizations to Reduce Impact and the Effect of Overrides” on page 13 for an important discussion regarding the ramifications of certain overrides for these rules.

Another important override to consider is the **TimeoutSeconds** parameter defined for rules and monitors. The default value for all workflows is 15 minutes (900 seconds). As the device count of an InfraStruxure Central server grows, the time required by these scripts increases and may exceed 15 minutes.

You will be alerted to this condition by the standard Operations Manager alerts. When notified, you will need to override the TimeoutSeconds parameter. This condition is not affected by the number of devices visible to the Health Server, but is dependent on the total number of devices managed by the InfraStruxure Central server.



See “Partitioning Large InfraStruxure Central Servers across Multiple Health Services” on page 9 for more information.



See “Optimizations to Reduce Impact and the Effect of Overrides” on page 13 for an important discussion regarding the ramifications of certain overrides for these rules.

Registering InfraStruxure Central Servers on Proxy Health Services

The initial discovery that generates the first instances of class types defined in this Management Pack is the registration discovery. This discovery is targeted to computers running Windows Server (i.e. the `Microsoft.Windows.Server.Computer` Operations Manager class) and runs every six hours. It discovers the registry keys that are written by the registration tool and maps them to instances of the corresponding ISX Central class types.

It also defines the management relationship between the Health Service and any InfraStruxure Central servers. This ensures that all subsequent activities related to the InfraStruxure Central server are processed by the same Health Service that discovered it. The registration discovery also discovers component groups that mirror the InfraStruxure Central server itself, which are subsequently linked to the two Dynamic Distributed Application root objects by a separate discovery rule that creates containment relationships between them.

Any InfraStruxure Central server that is registered on a Health Service will be managed from that Health Service. All transactions with the InfraStruxure Central web services interface will be conducted from that Health Service. You must register the InfraStruxure Central servers that will be monitored on the Health Services that will be acting as proxy before you create Run As accounts and assign them to the Run As profile defined in the Management Pack.

The first discovery run on a Health Service will run almost immediately after the Run As account is set. If the InfraStruxure Central servers are already registered (i.e., present in the registry), the initial discovery will discover them. Otherwise, the initial discovery will not discover any registered InfraStruxure Central servers until the next run six hours later.



Note: When using a Health Service with an InfraStruxure Central server, you must enable the “Agent Proxy” option under **Agent Properties > Security** on the Health Service.

To use the registration utility, start it from the Health Service that will be acting as the proxy. A grid will be displayed that shows the InfraStruxure Central servers currently registered on that Health Service.

- To add a new InfraStruxure Central registration, click Add. You will be asked to provide a name and fully-qualified domain name (FQDN) for the InfraStruxure Central server. The name can be any descriptive name. The fully-qualified domain name should be the IP address or FQDN of the InfraStruxure Central server. Do not use a URL here. The registry is updated immediately.
- InfraStruxure Central registrations on a Health Service must contain unique FQDNs as this is a key field in Operations Manager. We also recommend that the name be descriptive and clear, as it will be used in the Operations Manager object paths.
- To remove an InfraStruxure Central registration, highlight the row of the registration and click Remove. The registry is updated immediately.
- To set the collection date for an InfraStruxure Central server, highlight the row of the registration and click Advanced.



See “Registry Use in Alarm Collection” on page 8 for more information.

All registry changes are made in the key:

```
HKEY_LOCAL_MACHINE\Software\APC\MsScOpsMgrMpForIsxc
```

- The registered names appear as the REG_SZ values Name0...Name<n>
- The registered FQDNs appear as the REG_SZ values FQDN0...FQDN<n>
- These values must be consecutively numbered from 0.

Registry Use in Alarm Collection

The four rules that collect alarms from an InfraStruxure Central server use the registry of the Health Service to manage the collection process by storing the date and time of the last collection in the registry. This date and time, expressed in the XML date and time format, is used as the starting point for the next collection. This information is stored under the key:

```
HKEY_LOCAL_MACHINE\Software\APC\MsScOpsMgrMpForIsxc\AlarmCollection
```

- The last collection date appears as a REG_SZ value named for the FQDN of the InfraStruxure Central server it represents. The value is the XML-formatted date and time value.

Under normal circumstances, you can accept the default behavior of the Management Pack, which is to set the last collection date and time to the current date and time when the collection rules first run. This prevents historical alarms from being brought into Operations Manager, but allows future alarms to be imported.

There are certain circumstances under which you may need to change this value for an InfraStruxure Central server:

- With the registration of an InfraStruxure Central server, you wish to specify a date and time from which historical alarms should be collected on the first run.
- You experience an issue related to the configuration sets that result from a certain combination of overrides, as described in the “Optimizations to Reduce Impact and the Effect of Overrides” on page 13, and wish to reset the collection to a prior date and time after correcting the issue.

In these cases, you may use the registration tool to specify the date and time from which the Management Pack should collect alarms the next time it does a collection:

1. Start the registration tool.
2. Click on the row containing the InfraStruxure Central server for which you wish to change the last collection date and time.
3. Click the **Advanced** button.
4. Specify the date and time the Management Pack should use the next time it collects alarms.

Creating Run As Accounts and Assigning Them to the Run As Profile

After registering the InfraStruxure Central servers that will be monitored by the selected Health Services, you must create the Run As accounts required to access the InfraStruxure Central servers and assign them to the Run As Profile that is defined in the Management Pack. Each Health Service can be assigned a different Run As account. You will need multiple Run As accounts in the following cases:

- If you are partitioning an InfraStruxure Central server across multiple Health Services (discussed below).
- If multiple InfraStruxure Central servers do not share a common login account. In this case, multiple Health Services may be required to act as proxy for the different InfraStruxure Central servers. This is because only one Run As account can be assigned to the Run As profile for a particular Health Service at one time.

To assign a Run As account to a health service's Run As profile, you must select the "Simple Authentication Run As" option when creating the account.

Partitioning Large InfraStruxure Central Servers across Multiple Health Services

Because there are inherent limits to the number of concurrent workflows that Operations Manager can run, the Management Pack has been designed to require a minimum number of transactions with the InfraStruxure Central server. This is called "cook down" and is covered in the section "Optimizations to Reduce Impact and the Effect of Overrides" on page 13.

Cook down solves the problem of concurrent workflows, but can introduce other issues, namely capacity and outstanding data in a workflow. To mitigate these issues, the Management Pack allows multiple Health Services to manage a single InfraStruxure Central server by utilizing device groups and restricted user accounts.

As an example, consider an environment that has a single InfraStruxure Central server that manages 500 devices. In testing, you discover that the activity from this InfraStruxure Central server overwhelms the Health Service. You can partition this InfraStruxure Central server to be managed by multiple Health Services by using the following method:

- Create device groups on the InfraStruxure Central server and split the device inventory among the device groups. For example, Group 1 through Group 4 with 125 devices each.
- Create four restricted logins on the InfraStruxure Central server, User1 through User4, and assign them rights to only Group 1 through Group 4, respectively.
- Use the registration utility to register the InfraStruxure Central server on four different servers; for example SERVER1, SERVER2, SERVER3, and SERVER4. While the name and FQDN can be identical across Health Services (the proxy Health Service is a key field), it is recommended that you appropriately tag the InfraStruxure Central server name in a way that uniquely identifies the partition (e.g. a number)

- Create four Simple Authentication Run As accounts in Operations Manager, one for each user account.
- Assign the Run As profile that installs with the Management Pack to the four user accounts, using a different user name for each of the four Health Services on which you would like to partition the load. For example:
 - On SERVER1, assign User1 to the APC ISX Central Login Run As Profile
 - On SERVER2, assign User2 to the APC ISX Central Login Run As Profile
 - On SERVER3, assign User3 to the APC ISX Central Login Run As Profile
 - On SERVER4, assign User4 to the APC ISX Central Login Run As Profile

Additional Steps for Operations Manager 2007 R2. The following steps should be executed on Operations Manager 2007 R2 platforms only.

- Ensure that each Run As account is distributed to the Windows Server computer on which it is intended to be used. This is configured on the last tab of the Run As account properties.
- Configure the Run As profile and assign each Run As account to the Windows Server computer on which it is intended to be used.
- After the InfraStruxure Central objects registered on each Windows Server are discovered, configure the Run As Profile and additionally assign each Run As Account to the InfraStruxure Central object to which it applies. The InfraStruxure Central objects should be uniquely tagged, per the suggestion in the third step of this process.

Operations Manager 2007 R2 Example.

- Assign User1 to SERVER1
- Assign User2 to SERVER2
- Assign User3 to SERVER3
- Assign User4 to SERVER4

Assuming that the InfraStruxure Central server is named "central01" and that it is registered as "central01/1" on SERVER1, "central01/2" on SERVER2, etc., after the four InfraStruxure Central objects are discovered:

- Assign User1 to central01/1
- Assign User2 to central02/2
- Assign User3 to central03/3
- Assign User4 to central04/4



Note: With Operations Manager 2007 R2 and partitioning, the Run As Profile configuration must be kept in sync with the registered InfraStruxure Central servers and their partitions

Setting the Environment Variable ISXCUI

In order to use certain console tasks defined in the Management Pack, register the install location of the InfraStruxure Central client application using a system environment variable. The variable should be named ISXCUI and should be set to the full pathname of the folder that contains the InfraStruxure Central Client. Do not use the full pathname of the client executable. You must use the full pathname of the folder that contains that executable to enable the “Launch ISX Central” console tasks to locate the client executable.

Topics Related to Operations

Inventory Discovery

The second phase of the discovery, the inventory discovery, is targeted to each InfraStruxure Central server that is discovered by the registration discovery. This discovery is scheduled to run every 30 minutes, with the first run for any given InfraStruxure Central server occurring shortly after the InfraStruxure Central server's registration is discovered. It discovers all the device groups, device group membership relationships, devices, and sensors visible to the discovery (visibility is determined by the rights granted to the InfraStruxure Central server login represented by the Run As account assigned to the Run As Profile for this Health Service).

Inventory Discovery Segments and Access to the %TEMP% Folder

The maximum size of a single discovery data item in Operations Manager is 4 MB. This is a security measure built into Operations Manager. The discovery data item is an XML document and as such has the overhead associated with the text associated with starting and ending tags, etc. The result is that with larger InfraStruxure Central server inventories, a single discovery data item cannot represent the entire inventory. To overcome this, the discovery is divided into segments. There are 21 segments that are defined with the base Management Pack, which will accommodate even the largest InfraStruxure Central server.

The operation of these segments is essentially identical, with the exception of segment 0. Only segment 0 actually contacts the InfraStruxure Central server to retrieve the visible inventory through the Web Services interface. Segment 0 writes the entire inventory to an intermediate file it places in the temporary directory defined by the %TEMP% variable. If the Health Service is running as Local System, this will be %SystemRoot%\temp. If the Health Service is running as any other account, this will be that account's temporary directory.

This file includes directives for the class instances and properties for device groups, device group members, devices, and sensors. It also includes segment directives, with the directive to start segment 0 being the first line in the file. As segment 0 writes the contents of this file to disk, it estimates the contribution each directive will have on the size of the resultant discovery data item.

When this estimate approaches 4M, it writes a directive to start the next segment in order and resets the estimate. The net effect is that the intermediate file will be split into segments that roughly divide the discovery into segments that should be well under the 4M limit.

All segments read this master intermediate file and create their own copy, with only the directives specific to that segment being included. This includes segment 0, which proceeds to this step after it completes the inventory and writes the master intermediate file. All segments, including segment 0, then process this file and create their discovery data item.

Segment 0 will re-write the master intermediate file every time it runs. It attempts to minimize the amount of time the file is in an inconsistent state by writing the new file using a temporary file name and then deleting the original and immediately renaming the temporary file. All segments use this two-step process of splitting the file first and then processing it to minimize the amount of time they have the master intermediate file open. There are retries built into the process.

Warning Events on Initial Inventory Discovery

Since there are no mechanisms for inter-workflow synchronization, there is one race condition that occurs on the first discovery cycle for a new InfraStruxure Central server on a Health Service. This also occurs on the first discovery after the temporary directory is cleared, if that happens. All segments except segment 0 will attempt to open the master intermediate file. Since it may take segment 0 much longer to create this file than is prudent for the other segments to wait, they will log a 2724 event in the Operations Manager log and exit without creating a discovery data item.

This will also cause event 21405 to be logged by Operations Manager itself, indicating that the module was expected to create a data item but did not. The discovery process must terminate abnormally in this case. On the initial discovery, the discovery could return an empty discovery, but if this condition is occurring because the temporary directory has been cleared, an empty discovery would cause all objects for segments other than 0 to be deleted until the next discovery cycle.

Another consequence of this behavior is that the entire inventory for larger InfraStruxure Central servers will not be discovered until the second discovery cycle. The initial discovery cycle will discover the inventory that fits into the initial 4M discovery data item (segment 0). The remainder of the segments will terminate and not produce discovery data until the second cycle.

This behavior is a consequence of the discovery data item size limit and the lack of synchronization facilities between workflows.

International Characters in Inventoried Properties

Because the InfraStruxure Central interface uses Web Services that are accessed from Operations Manager using a script, the output of the script is available in ANSI characters only. The output is converted to the UTF-8 encoding upon being picked up from the script by Operations Manager. While international characters will import properly, some character sets that are double-byte only (UTF-16, etc.) will not import correctly. The data will import, but will be represented as question mark (?) characters in Operations Manager. There is no work-around for this limitation.

Drilling Down on Distributed Application (DA) Diagrams

The Management Pack installs with several pre-defined views, among them two special diagram views for the two dynamic Distributed Applications that the Management Pack creates. One DA is the hierarchy of InfraStruxure Central servers by device group; the other is the hierarchy of InfraStruxure Central servers by device type. For both of these diagram views, options have been set that streamline the display of these DA within the view.

The consequence of this is that the diagram view's tree ends at the device group or device type level. These are leaf nodes that cannot be expanded to see the devices below them. To further investigate the devices and sensors that are contributing to the overall health state of the leaf node, right click on the leaf node and select Open Diagram View. A new diagram view will open with the device group or device type as its root object. The devices and sensors subordinate to the root object in the hierarchy will be visible in this diagram view.

Special Parameters for Monitor Overrides

Where possible, monitors have been equipped with two parameters that you may override to finely tune the health state into which a given monitor enters due to the various states of the underlying InfraStruxure Central server: UpgradeWarning and DowngradeCritical.

As their name implies, these two Boolean (true/false) values can upgrade warning conditions to critical conditions or downgrade critical conditions to warning conditions.

While Operations Manager provides the ability to override the severity of an alert generated from a monitor, the severity of the underlying monitor is usually fixed. This difference in severity can cause confusion in some cases when investigating the cause of an alert. The alert will show one severity, but the underlying monitor states in the Health Explorer will show another severity. These parameters allow the state of the underlying monitor to be changed. While it is still possible to override the severity of the alert, the alert severity can be left to follow the monitor severity in most cases

Most monitors defined in the Management Pack use a pre-defined mapping to reflect the underlying entity state in the Operations Manager monitor health state. Given these two parameters, you can change this behavior. The net effect of these parameters is given in the following matrix:

UpgradeWarning	DowngradeCritical	Net Effect
False (Default)	False (Default)	The health state of the monitor will be the health state mapped by the Management Pack from the state of the underlying device or sensor in InfraStruxure Central.
True	False (Default)	All abnormal states of the underlying device or sensor in InfraStruxure Central will be represented as Critical in the health state of the monitor.
False (Default)	True	All abnormal states of the underlying device or sensor in InfraStruxure Central will be represented as Warning in the health state of the monitor.
True	True	Warning states will be Critical and Critical states will be Warning. Use this option only under unique conditions where the configuration warrants it.

Optimizations to Reduce Impact and the Effect of Overrides

The scripts defined in this Management Pack that collect device and sensor state and data take advantage of Operations Manager’s ability to “cook down” a script. This applies to every script that is used in a workflow targeted to a device or sensor, which includes the following:

- Device Alarm Health Monitor
- Collection Rule for Device Alarms
- Alerting Rule for Device Alarms
- Device State Health Monitor
- Sensor Alarm Health Monitor
- Collection Rule for Sensor Alarms
- Alerting Rule for Sensor Alarms
- Sensor State Health Monitor
- Sensor State Monitor (Pre-Defined)
- Sensor State Monitor (User-Defined)

This design prevents Operations Manager from requiring a separate invocation of any particular script for each instance of a device or sensor. If this were not done in large environments, the Health Service could be required to run literally thousands of scripts each polling cycle. This would be an undue strain on the InfraStruxure Central server, the Health Service, and the entire Operations Manager Management Group.

To understand how this design mitigates this risk, it is necessary to understand the anatomy of these workflows. Consider the Device State Monitor workflow as an example:

- The Device State Health Monitor targets an InfraStruxure Central device and has three states: Healthy, Warning, and Critical. Therefore, for each device, three independent workflows will be required:
 - Check the Device State, Filter for Healthy State, Set the State.
 - Check the Device State, Filter for Warning State, Set the State.
 - Check the Device State, Filter for Critical State, Set the State.
- Since Operations Manager realizes that the “Check the Device State” component (module) of the workflows are all the same, it will run this only once. This is what is known as “cooking down” the script. The output from that one invocation of the script will be fed to each filter (condition detection) and only the one that matches the specified criteria will continue the workflow and set the state accordingly. This is the basic theory of operation of any monitor.
- If the “Check the Device State” script requires the InfraStruxure Central Device Element ID as a parameter, the extent to which Operations Manager can cook down the script ends there. Although each device has three workflows that use the same invocation of the script, Operations Manager will still be required to run the script once for each device. This is the common design pattern of a scripted monitor: that the script runs against a specific instance of the target and therefore must be run once for each target.

- In this Management Pack, the aforementioned scripts do not require the target ID to be passed. They only require the endpoint of the InfraStruxure Central server. They all contact the InfraStruxure Central server and return a data set that represents all targets on that InfraStruxure Central server. For the Device State Health Monitor, for instance, the script returns the Device State for all devices on an InfraStruxure Central server. Because the data set contains State for multiple devices, each data point also contains the device element ID. This causes Operations Manager to feed each data point in the data set to each workflow in the pipeline (in this case, the three filters for each device).
- To compensate for the lack of specificity in the data source module, the filters for each health state are more complex than usual. Instead of evaluating only the state returned, they evaluate that the device element ID matches the target of any given workflow and the state matches the state of that workflow (plus some other criteria that are not directly related to cook down). Although this step only shifts the burden to the filter, it profoundly reduces the overall impact of the monitor:
 - Only one set of transactions is made with the InfraStruxure Central server.
 - Only one script is run, which requires a separate process and is generally the most resource-intensive type of operation in Operations Manager.
 - The bulk of the burden is borne by other module types, such as filters, that are implemented in native or managed code (.dlls) and that are tightly integrated and optimized internal to the Health Service.

The underlying data source module is not specific in its configuration. It is only passed a few parameters that are not directly related to the target; specifically: (1) the endpoint of the InfraStruxure Central server, (2) the interval between script runs, (3) the synchronization time of the runs, and (4) the script timeout in seconds.

As with any workflow in Operations Manager, the Management Pack allows the Operations Manager administrator to override values (2) through (4) to fit any particular need. These overrides carry with them special potential consequences that must be understood and taken into account when you specify the overrides. Consider the following example:

- The default configuration sets the Device State Health Monitor to run every 5 minutes with no synchronization time and a 15 minute timeout.
- You have three different groups of devices that you have created, High Importance, Medium Importance, and Low Importance.
- You override the monitor interval, specifying every 5 minutes for High Importance devices, every 15 minutes for Medium Importance devices, and every hour for Low Importance Devices.
- Because of the design of the system as described above, the ability for Operations Manager to cook down the “Check the Device State” script has been reduced to one run for each group. This is because one of the overridden values is a configuration parameter of the script. As such, Operations Manager must run the script once for each unique set of configuration parameters.
- Although logically you have separated the devices into three groups, the script is designed to collect the device state for all devices. The overall operation of the monitor will still be intact, however:
 - For the High Importance devices, the script will return the state for all devices every 5 minutes, but the filters will ensure that only those device element IDs actually targeted in this group are affected.
 - The same is true for the Medium Importance devices and Low Importance devices.

- Therefore, although you have separated the devices and have alleviated Operations Manager from evaluating the monitor on the Medium and Low Importance devices so frequently, the script that runs every 5 minutes still does the same amount of work. In fact, since you are running one script every 5 minutes, two every 15 minutes (the 5 and 15), and three every 60 minutes (the 5, 15, and 60), you have actually increased the number of transactions that run against the InfraStruxure Central server.
- For most workflows, this number of transactions will be acceptable. The benefits of cooked down scripts far outweigh the overhead incurred when the cook down is counteracted by varying configuration sets caused by overrides.



Caution: If overrides are specified for very small groups (or even individual devices or sensors) and with even slightly different values, the script could be forced to run, collecting all data points, but only being used by a single target. **Take care in specifying overrides in this Management Pack.**



Warning: Four of the rules provided in this Management Pack collect the alarms from the InfraStruxure Central server for devices (two rules) and sensors (two rules) and create Operations Manager events based on the content of the alarms as well as optionally generate alerts for these events. These four rules actually share a common data source that assembles a data set of all alarms that have occurred on a given InfraStruxure Central server since the last collection. **These four rules must only be overridden together, for all devices/sensors on a given InfraStruxure Central server, and using the same exact values.**

If multiple configuration sets are created by overrides that do not follow these criteria, the last collection time will be set in the registry of the Health Service computer by the first such configuration set to run. The remaining configuration sets will run, but will incorrectly determine the date at which they last ran. These rules are designed to be overridden only under the conditions listed above.

In summary, for all of the workflows referenced in the beginning of this section, the cook down design pattern must be taken into account when specifying overrides for the values that affect the data source script: TimeoutSeconds, IntervalSeconds, and SyncTime.

Script Events Generated by the Management Pack

Events 2720 and 2721 are logged when the registration discovery rule runs on a Health Service. Event 2720 marks the start of the discovery; event 2721 marks the end. When this discovery completes, all InfraStruxure Central servers registered on the Health Service should be discovered, though they will not have any inventory or health state.

Events 2722 and 2723 are logged when segment 0 of the inventory discovery rule runs on a Health Service. Event 2722 marks the start of the discovery; event 2723 marks the end. When this discovery completes, the inventory for the InfraStruxure Central server should be at least partially present in Operations Manager.

Event 2724 is logged by all scripts and signifies a general exception that the script encountered. This event will be logged with the script name, an internal function name, a descriptive message, and the contents of the exception.

Event 2729 is similar to 2724 but is logged specifically when an exception is encountered during a transaction with an InfraStruxure Central server's Web Services interface. The same content is logged with the message as with event 2724, but the message may also include the XML involved in the transaction.

Appendices

Appendix 1 – Included Sensor Types

ISXCSensorType	Operations Manager ClassType
AIRFLOW	com.apcc.ISXC.ISXCSensor.Airflow
ALARM_STATUS (AC)	com.apcc.ISXC.ISXCSensor.State
ALARM_STATUS (ENVIRONMENT)	com.apcc.ISXC.ISXCSensor.State
ANALOG	com.apcc.ISXC.ISXCSensor.Generic
BATTERY_REPLACEMENT	com.apcc.ISXC.ISXCSensor.State
CAMERA_MOTION	com.apcc.ISXC.ISXCSensor.State
COOLANT_LEVEL	com.apcc.ISXC.ISXCSensor.State
COOLING_DEMAND	com.apcc.ISXC.ISXCSensor.Power
COOLING_OUTPUT	com.apcc.ISXC.ISXCSensor.Power
DOOR_SWITCH	com.apcc.ISXC.ISXCSensor.State
EXTERNAL_CONTACT	com.apcc.ISXC.ISXCSensor.State
EXTERNAL_HUMIDITY	com.apcc.ISXC.ISXCSensor.Humidity
EXTERNAL_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
FRONT_DOOR_HANDLE_STATUS	com.apcc.ISXC.ISXCSensor.State
FRONT_DOOR_LOCK_STATUS	com.apcc.ISXC.ISXCSensor.State
FRONT_DOOR_STATUS	com.apcc.ISXC.ISXCSensor.State
GENERATOR_MODE	com.apcc.ISXC.ISXCSensor.State
GENERATOR_STATUS	com.apcc.ISXC.ISXCSensor.State
INPUT_CURRENT	com.apcc.ISXC.ISXCSensor.Current
INPUT_FLUID_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
INPUT_VOLTAGE	com.apcc.ISXC.ISXCSensor.Voltage
INTEGRATED_HUMIDITY	com.apcc.ISXC.ISXCSensor.Humidity
INTEGRATED_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
LOWER_INLET_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
MAIN_INPUT_BREAKER	com.apcc.ISXC.ISXCSensor.State
MIDDLE_INLET_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
OIL_LEVEL	com.apcc.ISXC.ISXCSensor.State
ON_BATTERY_STATE	com.apcc.ISXC.ISXCSensor.State
ON_BYPASS_STATE	com.apcc.ISXC.ISXCSensor.State
OPERATING_MODE	com.apcc.ISXC.ISXCSensor.State
OUTPUT_CURRENT	com.apcc.ISXC.ISXCSensor.Current
OUTPUT_CURRENT_TOTAL	com.apcc.ISXC.ISXCSensor.Current
OUTPUT_LOAD_VA	com.apcc.ISXC.ISXCSensor.ApparentPower
OUTPUT_LOAD_WATTS	com.apcc.ISXC.ISXCSensor.Load
OUTPUT_POWER_TOTAL_VA	com.apcc.ISXC.ISXCSensor.ApparentPower
OUTPUT_POWER_TOTAL_WATTS	com.apcc.ISXC.ISXCSensor.Power
OUTPUT_VOLTAGE	com.apcc.ISXC.ISXCSensor.Voltage
Q1_INPUT_BREAKER	com.apcc.ISXC.ISXCSensor.State
Q2_OUTPUT_BREAKER	com.apcc.ISXC.ISXCSensor.State
Q3_BYPASS_BREAKER	com.apcc.ISXC.ISXCSensor.State

RACK_INLET_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
REAR_DOOR_HANDLE_STATUS	com.apcc.ISXC.ISXCSensor.State
REAR_DOOR_LOCK_STATUS	com.apcc.ISXC.ISXCSensor.State
REAR_DOOR_STATUS	com.apcc.ISXC.ISXCSensor.State
RUN_TIME_REMAINING	com.apcc.ISXC.ISXCSensor.Time
SENSOR_STATE	com.apcc.ISXC.ISXCSensor.State
SUPPLEMENTAL_DEVICE_STATUS	com.apcc.ISXC.ISXCSensor.State
SYSTEM_OUTPUT_CURRENT	com.apcc.ISXC.ISXCSensor.Current
SYSTEM_OUTPUT_POWER_TOTAL	com.apcc.ISXC.ISXCSensor.Power
SYSTEM_OUTPUT_POWER_USAGE_TOTAL	com.apcc.ISXC.ISXCSensor.Power
TIME_ON_BATTERY	com.apcc.ISXC.ISXCSensor.Time
UPPER_INLET_TEMPERATURE	com.apcc.ISXC.ISXCSensor.Temperature
UTILITY_FAILURE_STATE	com.apcc.ISXC.ISXCSensor.State
VOLUME_AIRFLOW	com.apcc.ISXC.ISXCSensor.Airflow

APC Worldwide Customer Support

Customer support for this or any other APC product is available at no charge in any of the following ways:

- Visit the APC Web site to access documents in the APC Knowledge Base and to submit customer support requests.
 - **www.apc.com** (Corporate Headquarters)
Connect to localized APC Web sites for specific countries, each of which provides customer support information.
 - **www.apc.com/support/**
Global support searching APC Knowledge Base and using e-support.
- Contact the APC Customer Support Center by telephone or e-mail.
 - Local, country-specific centers: go to **www.apc.com/support/contact** for contact information.

For information on how to obtain local customer support, contact the APC representative or other distributors from whom you purchased your APC product.

© 2010 APC by Schneider Electric. APC, the APC logo, and InfraStruXure are owned by Schneider Electric Industries S.A.S., American Power Conversion Corporation, or their affiliated companies. All other trademarks are property of their respective owners.

