Operations Smart Plug-ins for Infrastructure

Software Version: 12.04
Operations Manager for Windows®, HP-UX, Linux, and Solaris operating systems

Concepts Guide

Document Release Date: August 2017
Software Release Date: August 2017

Hewlett Packard Enterprise

Legal Notices

Warranty

The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

The information contained herein is subject to change without notice.

Restricted Rights Legend

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Copyright Notice

Trademark Notices

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Microsoft® and Windows® are U.S. registered trademarks of the Microsoft group of companies.

UNIX® is a registered trademark of The Open Group.

Java is a registered trademark of Oracle and/or its affiliates.

Documentation Updates

To check for recent updates or to verify that you are using the most recent edition of a document, go to: https://softwaresupport.hpe.com/.

This site requires that you register for an HPE Passport and to sign in. To register for an HPE Passport ID, click **Register** on the HPE Software Support site or click **Create an Account** on the HPE Passport login page.

You will also receive updated or new editions if you subscribe to the appropriate product support service. Contact your HPE sales representative for details.

Support

Visit the HPE Software Support site at: https://softwaresupport.hpe.com/.

Most of the support areas require that you register as an HPE Passport user and to sign in. Many also require a support contract. To register for an HPE Passport ID, click **Register** on the HPE Support site or click **Create an Account** on the HPE Passport login page.

To find more information about access levels, go to: https://softwaresupport.hpe.com/web/softwaresupport/access-levels.

HPE Software Solutions Now accesses the HPE Software Solution and Integration Portal website. This site enables you to explore HPE Product Solutions to meet your business needs, includes a full list of Integrations between HPE Products, as well as a listing of ITIL Processes. The URL for this website is https://softwaresupport.hpe.com/km/KM01702731.

Contents

Chapter 1: Introduction

This chapter provides a broad overview of Infrastructure Management and how the Operations Smart Plug-ins for Infrastructure (Infrastructure SPIs) can be effectively used to manage the enterprise wide infrastructure. It introduces concepts that will help you manage your infrastructure resources and processes, monitor applications, and monitor your systems.

What is Infrastructure Management

The dependency of enterprise on infrastructure has made it really imperative to look for ways and means to manage IT infrastructure. Infrastructure management not only helps to maintain and optimize business-critical systems but also assures availability of resources in a complex and decentralized IT infrastructure setup.

Logical representation of Infrastructure management

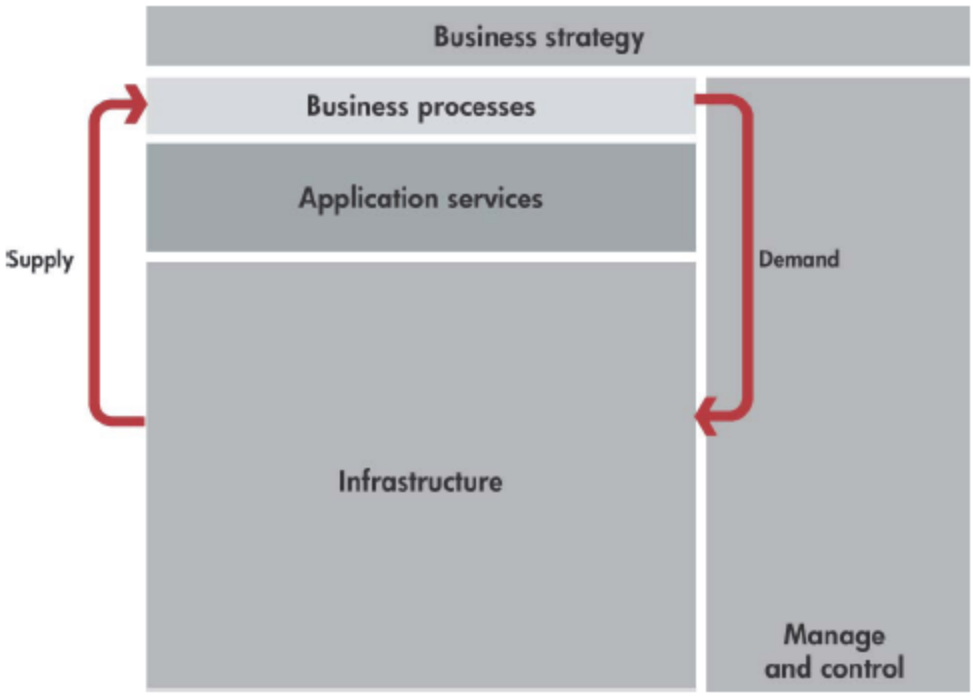

Infrastructure management helps you to monitor, analyze, and optimize the usage of distributed operating systems, application and storage servers, clusters, and virtual machines. It helps to predict

infrastructure resource utilization so that IT infrastructure issues can be either avoided or resolved before critical business availability is impacted. It also ensures optimal system performance and availability across the entire setup.

The challenges of managing disparate infrastructure environments within a global organization are to coordinate the flow of critical information to resolve problems quickly, and to reduce the downtime and costs.

Systems Infrastructure Management

Systems infrastructure is the foundation or base infrastructure that is integral to an enterprise. It includes CPU, operating system, disk, memory, and network resources that need to be continuously monitored to ensure availability, performance, security and smooth functioning of underlying physical systems.

The system downtime can affect the quality of service you offer to the customers. For instance, a CPU bottleneck on the central web server could mean slow response to the customer accessing the server through a client application. This can directly encumber the customer satisfaction for your products and services. Such a scenario can be avoided by continuous monitoring of the systems infrastructure.

Systems infrastructure management enables you to achieve greater efficiency and productivity. It helps you to correlate, identify, and correct root cause of infrastructure faults and performance degradations. By analyzing the trend and performance of base infrastructure you can determine and plan for future requirements.

Virtual Infrastructure Management

Virtualization enables dividing the computer resources into multiple execution environments. It abstracts the physical hardware layer to enhance IT resource utilization. Virtual machines are used to consolidate the workloads of several under-utilized servers to fewer machines for effective utilization of hardware. It helps to reduce environmental costs and aids easier management and administration of the server infrastructure. Virtual machines are also used to run multiple operating systems simultaneously. These operating systems can be different versions, or even entirely different systems, which can be on hot standby. Virtualization enables existing operating systems to run on shared memory multiprocessors. Since virtual machines are logical entities and are separate from the physical resources they use, the host environment is able to dynamically assign the resources among them.

Virtualization infrastructure management maximizes resource utilization by providing monitoring services that allow you to visualize and manage your virtual infrastructure. The benefits of managing a

virtual infrastructure are lower management costs, centralized management of heterogeneous resources, improved performance, and increased availability with greater visibility into your virtual systems.

Cluster Infrastructure Management

A cluster is a group of systems grouped together over a network, to improve performance and availability of systems over that provided by a single system. The distributed software installed on the networked computers turns them into a distributed system and presents the user with a single-system image. Clusters such as Serviceguard (for HP-UX and Linux) and Microsoft Cluster Server (MSCS), are often used to manage servers for high availability.

There are various ways of clustering the systems, depending upon the purpose. For example, the *High-availability (HA)* clusters are created to ensure the service availability especially for business critical applications and services. The HA clusters have redundant nodes. If a server with a particular application crashes, the application is immediately restarted on another system without administrative intervention. This redundancy provides high availability of services by eliminating single points of failure. Another category of clusters is *Load-balancing.* Load balancing clusters share the workload among the systems that are members of a cluster, and function as one single virtual computer.

Cluster infrastructure monitoring maximizes resource availability and system performance by providing monitoring services that allow you to visualize and manage all the nodes in your cluster. The benefits of managing a cluster infrastructure are centralized management, improved performance, and increased availability of cluster nodes and resource groups.

Operations Smart Plug-ins for Infrastructure

The Operations Smart Plug-ins for Infrastructure (Infrastructure SPIs) forms a software suite that integrates fully with the Operations Manager (OM) and extends OMs management scope to include distributed enterprise-wide base infrastructure including systems, high-availability clusters (HA clusters), and virtual infrastructure.

The Infrastructure SPIs provide pre-defined management policies to enable you to quickly gain control of the essential elements of your IT infrastructure. It enables relating the cross domain IT infrastructure events with relevant applications and maps them into a hierarchical service map. The map view displays the real-time status of your infrastructure environment and helps to identify the root-cause of alarms reported on operating systems, associated software services, and, in addition, essential hardware elements such as CPU, memory, swap space and so on.

The Infrastructure SPIs can be used to monitor and manage the functionality of the operating systems and associated software and hardware. The operating systems and related infrastructure can be in a clustered and virtualized environment.

The Infrastructure SPIs are integrated with other OM products, such as Operations Agent, Performance Agent, Reporter, and Performance Manager.

Infrastructure SPIs Documentation Set

The following documentation map lists the documents that are required for understanding the product.

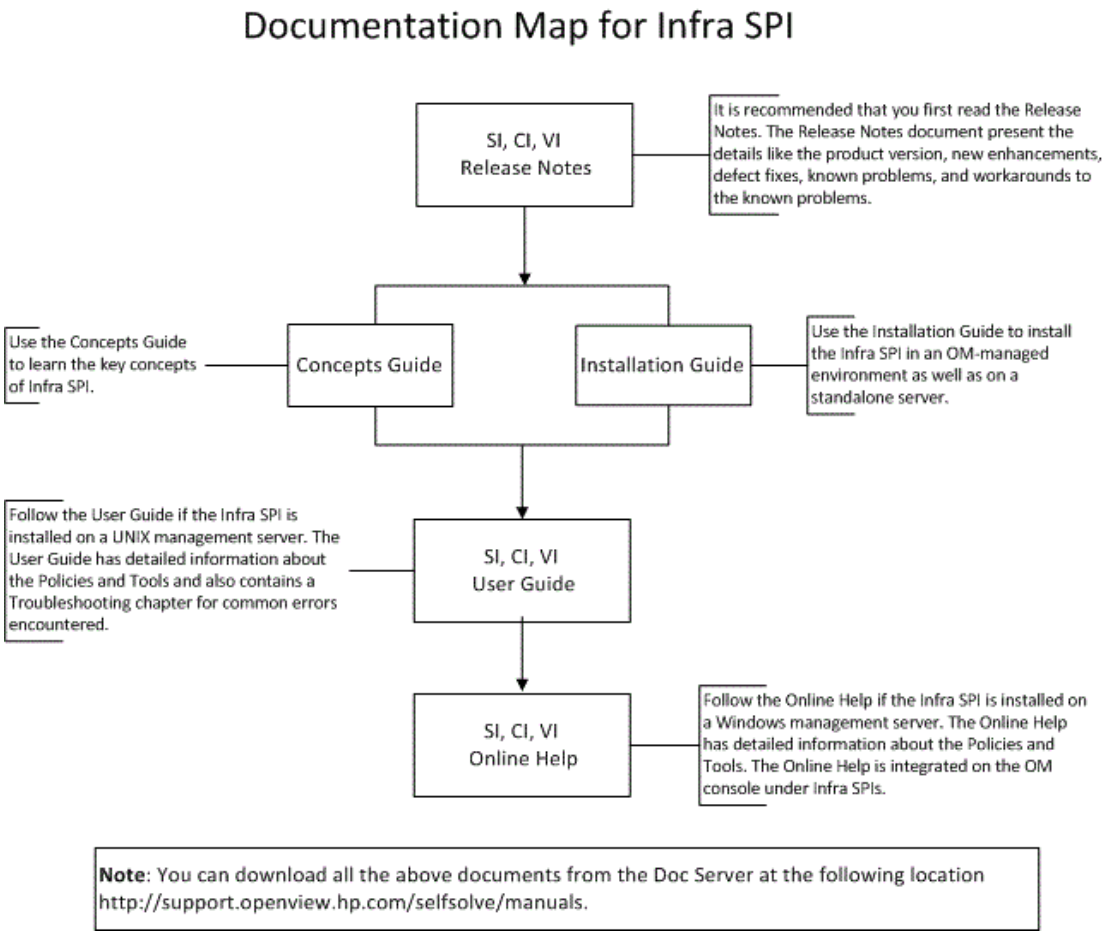

Related Documentation

The Infrastructure SPIs are integrated with other OM products, such as Operations Agent, Performance Agent, Reporter, and Performance Manager.

Before you start the installation of the infrastructure SPIs, you must plan the infrastructure. The Operations Manager (OM) presents you with the framework to monitor and manage multiple systems through a single, interactive console. The Operations agent deployed on individual nodes helps you gather vital information to facilitate the monitoring process.

To install and deploy the Operations agent on the nodes, see the following Operations Agent documentation:

- Operations Agent Installation Guide
- Operations Agent Deployment Guide
- Operation Agent Release Notes
- Performance Agent

You can integrate with Reporter to create reports in multiple formats from the data collected by the Operations Agent. See the following documentation:

- Reporter Concepts Guide
- Reporter Release Notes
- Reporter Installation Guide

Integrate with Performance Manager to view and analyze the available data in the form of charts and graphs. See the following documentation:

- Performance Manager Installation, Upgrade, and Migration Guide
- Performance Release Notes

Chapter 2: Infrastructure SPIs Architecture

The Infrastructure SPIs help you to increase the infrastructure availability and performance, visualize the capacity shortages and trends, and lower the overall operational maintenance cost across your entire environment. They offer common and unified model for managing infrastructure problem on single systems, cluster environments, and virtualized setups.

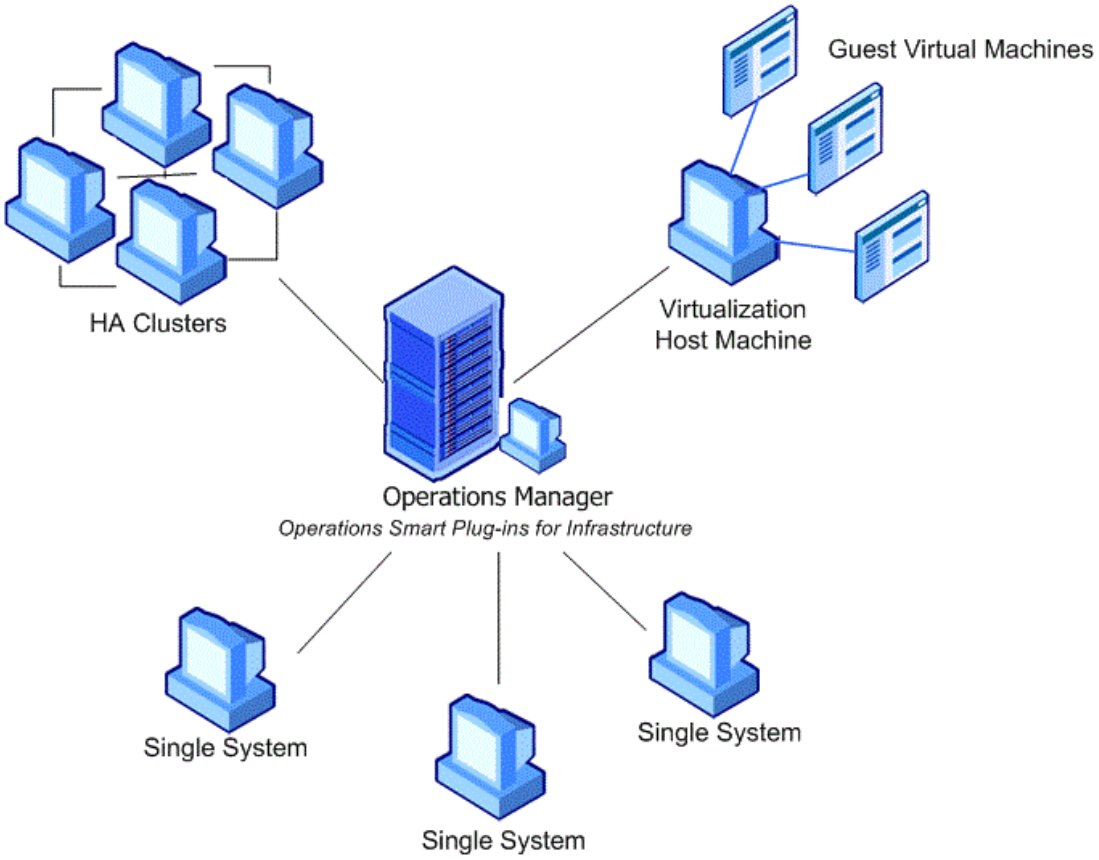

Operations Smart Plug-ins for Infrastructure is a software suite comprising of three SPIs:

- Operations Smart Plug-in for Systems Infrastructure
- Operations Smart Plug-in for Virtualization Infrastructure
- Operations Smart Plug-in for Cluster Infrastructure

To find an infrastructure-specific policy, such as for a cluster, you can use the policy folder. For example:

On OM for Windows console:

Policy management → Policy groups → Infrastructure Management → en → Cluster Infrastructure

On OM for UNIX/Linux console:

Policy Bank → Infrastructure Management → en → Cluster Infrastructure

Smart Plug-in for Systems Infrastructure

The Smart Plug-in for Systems Infrastructure (Systems Infrastructure SPI) helps you monitor enterprise wide single systems running Microsoft Windows or enterprise Linux distributions. It sends out alerts to the OM console for performance, capacity utilization, availability, and security of the monitored systems. The Systems Infrastructure SPI discovery policy gathers service information from the managed nodes such as hardware resources, operating system attributes, and applications and adds this information to the OM Services area.

If the managed node is a cluster node, the discovery policy initiates Cluster Infrastructure SPI discovery. The Cluster Infrastructure SPI discovers the clusters, cluster nodes, and resource groups.

Smart Plug-in for Virtualization Infrastructure

The Smart Plug-in for Virtualization Infrastructure monitors performance, capacity and availability aspects of the virtual resources.

The virtual infrastructure consists of the following components:

Host Machines are the physical machines allow sharing of the machine resources between different virtual machines.

Guest Machines are the virtual machines that run on the host machines abstracting the details of the underlying hardware or operating system.

Smart Plug-in for Cluster Infrastructure

The Smart Plug-in for Cluster Infrastructure helps you monitor high availability (HA) cluster infrastructure on the network. It sends out alerts to the OM console for performance and availability of the clustered nodes. The availability of clustered nodes can be affected due to downtime. Downtime may be planned due to maintenance or routine operations such as upgrade, space management or

system reconfiguration or unplanned due to power outage, human error, data corruption, and software or hardware errors.

The HA cluster infrastructure consists of the following components:

Cluster service controls cluster activity, communication between the cluster servers, and operations in case of a failure.

Cluster nodes are specially linked servers running the cluster service.

Cluster resource group is a group of cluster resources that are managed as a unit of failover.

Operations Manager i

Operations Manager i (referred to in this guide as OMi) is an event and performance manager that helps you to monitor heterogeneous network from the perspective of its operational infrastructure and, at the same time, from the perspective of business services, in a large and complex environment. It is responsible for gathering information about infrastructure health, systems, and applications to help identify and resolve specific operations problems.

OMi helps you achieve the following:

- Create a superior view of infrastructure health across systems and networks.
- Consolidate events into one central console and correlate the IT infrastructure events and end user management events.
- Automatically determine with topology-based event correlation the root cause and event relationship.
- Effectively make use of the advanced service health for the IT infrastructure.

To achieve this, OMi requires the Business Service Management (BSM) platform and Operations Manager. OM forwards the OMi ready events that occur in your managed environment to the OMi where they are evaluated. The image below is a logical representation of the communication between OM and OMi.

Logical representation of communication between OM and OMi

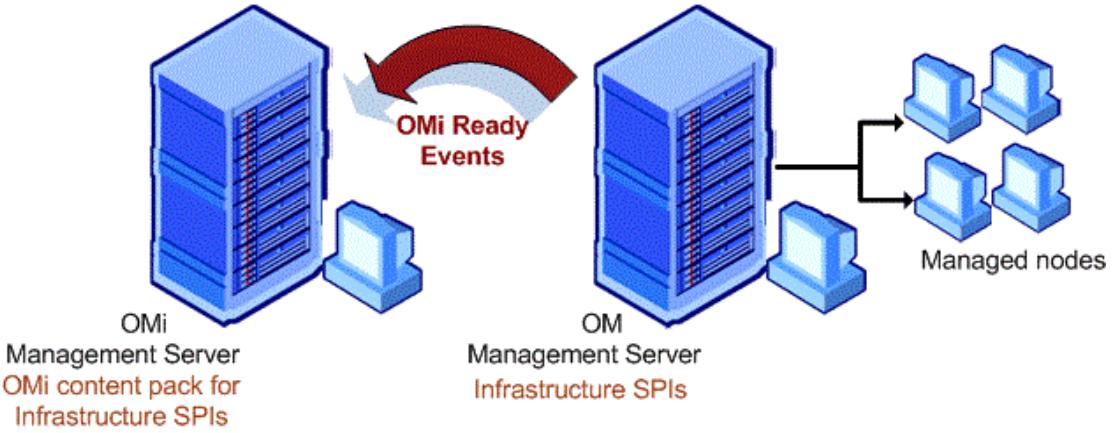

OMi uses content packs to exchange customized, OMi-related data between instances of the OMi management server. A content pack contains a complete snapshot of all (or any part of) the OMi rules, tools, mappings, assignments, and menu options that you define and configure to help users manage your IT environment with OMi. You can specify the content you want to update when synchronizing the topology between OMi and Operations Manager.

Using the Infrastructure SPIs, you can manage your infrastructure in OM and monitor the availability, use, and performance of the OM as managed nodes. Infrastructure SPIs include pre-configured content pack that enables OMi to analyze events related to infrastructure.

How Infrastructure SPIs display Alerts and Ancillary Information

Infrastructure SPIs displays information that helps to analyze, isolate and resolve infrastructure issues. It displays information in various forms:

Message alerts are displayed in the OM message browser. Using the measurement threshold policy settings and the collected values for each targeted metric that the collector/analyzer has gathered, the Infrastructure SPIs forward appropriate messages to the OM console, where they are displayed with a color-coded severity level.

Instructions text offers problem resolution suggestion. It is available within the generated message Properties sheet. To view the text, right-click the message, select **Properties**, and choose the **Instructions** tab.

Reports provide a picture of system, cluster, or virtualization resources. You can use reports to observe performance and usage trends.

Graphs help you to look at usage trends, compare performance between systems, and analyze the metric collected data. A set of preconfigured graphs are provided with the System Infrastructure SPI and Virtualization Infrastructure SPI.

Annotations are additional notes in the OM messages. Infrastructure SPI messages contains annotations providing status and the output of the automatic actions that run on the managed node.

Monitored Aspects

The Infrastructure SPI adds to the monitoring capabilities of OM by monitoring Infrastructure data that is targeted and gathered according to rules and schedule specifications contained within policies.

As the usage of IT resources changes, and functionality evolves, the amount of disk space, processing power, memory, and other parameters also change. It is essential to understand the current demands, and how they will change over time. Monitoring these aspects over a period of time is beneficial in understanding the impact on IT resource utilization. Infrastructure management analyzes current and historical performance to accurately predict future resource capacity needs.

Capacity monitoring

Capacity monitoring helps to deliver performance at the required service level and cost. It ensures that the capacity of the IT infrastructure corresponds to the evolving demands of the business. It helps identify the under-utilized and over-utilized resources.

Performance monitoring

Performance monitoring helps to preempt performance disruption and identify infrastructure issues that can threaten service quality. The collected performance data is used to correlate events across the entire infrastructure of servers, operating systems, network devices, and applications in order to prevent or identify the root cause of a developing performance issue.

Availability monitoring

Availability monitoring helps to ensure adequate availability of resources. It is important to identify unacceptable resource availability levels. The current load on IT infrastructure is computed and

compared with threshold levels to see if there is any shortfall in resource availability.

Security

Security monitoring helps to identify security issues and vulnerabilities across heterogeneous operating environments so that corrective steps can be initiated in a timely manner. This is necessary to ensure the continuity of service and safety of information.

Chapter 3: Key Concepts

Setting Thresholds

Most policies have multiple threshold levels. It is essential to set the threshold values carefully for each level because the message alert on the OM for Windows console depends upon the set threshold value.

Infrastructure SPIs enable you to set and modify policy threshold values in the following ways:

Customizing Threshold Values using Script Parameters

If the default values do not suit your particular environment, you can change the type of threshold as well as the defined threshold limits. The Script-Parameter tab in a policy displays the parameter names for the thresholds. These parameter names are case-sensitive.

To modify thresholds for Measurement Threshold type policies:

1. Start the Operations GUI and use the Console Tree to browse to the **Infrastructure Management** policy group/ policy bank.
2. Locate the policy you want to modify by expanding the appropriate policy group Virtualization Infrastructure, Systems Infrastructure, or Cluster Infrastructure and their appropriate subordinate groups.
3. Open the **Policy Edit** window.
4. Select the **Script-Parameters** tab and set the new threshold value for the thresholds as appropriate. In the absence of the Script-Parameters tab in the policy, you can use the Threshold levels tab to set threshold values.

   Example of modifying threshold for Measurement Threshold policy

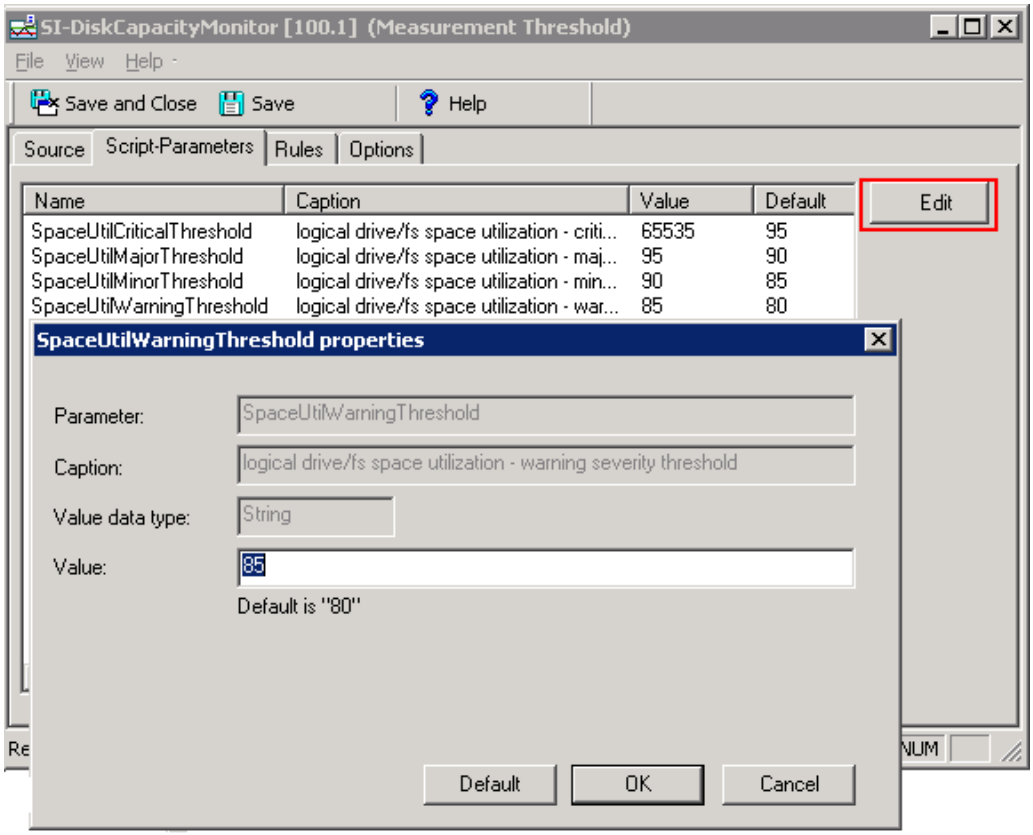

5. Click **OK** to save the changes.

   Redeploy the policy to the appropriate managed nodes.

Customizing Threshold Values using Threshold Overrides

The threshold overrides concept provides additional flexibility to the Infrastructure SPIs by controlling how a set of threshold values is applied to the target managed nodes. The *ThresholdOverrides* policy enables you to override the thresholds for multiple polices on a managed node.

You can set the threshold value across multiple policies in one step, by creating a list of all threshold parameter names and values, and override the thresholds across policies on a managed node. You can use this list of threshold parameters to standardize the setting on multiple managed nodes. In case you want to change values for a node, you can modify the values in the list and deploy the *ThresholdOverrides* policy on that particular managed node.

Example of threshold overrides policy

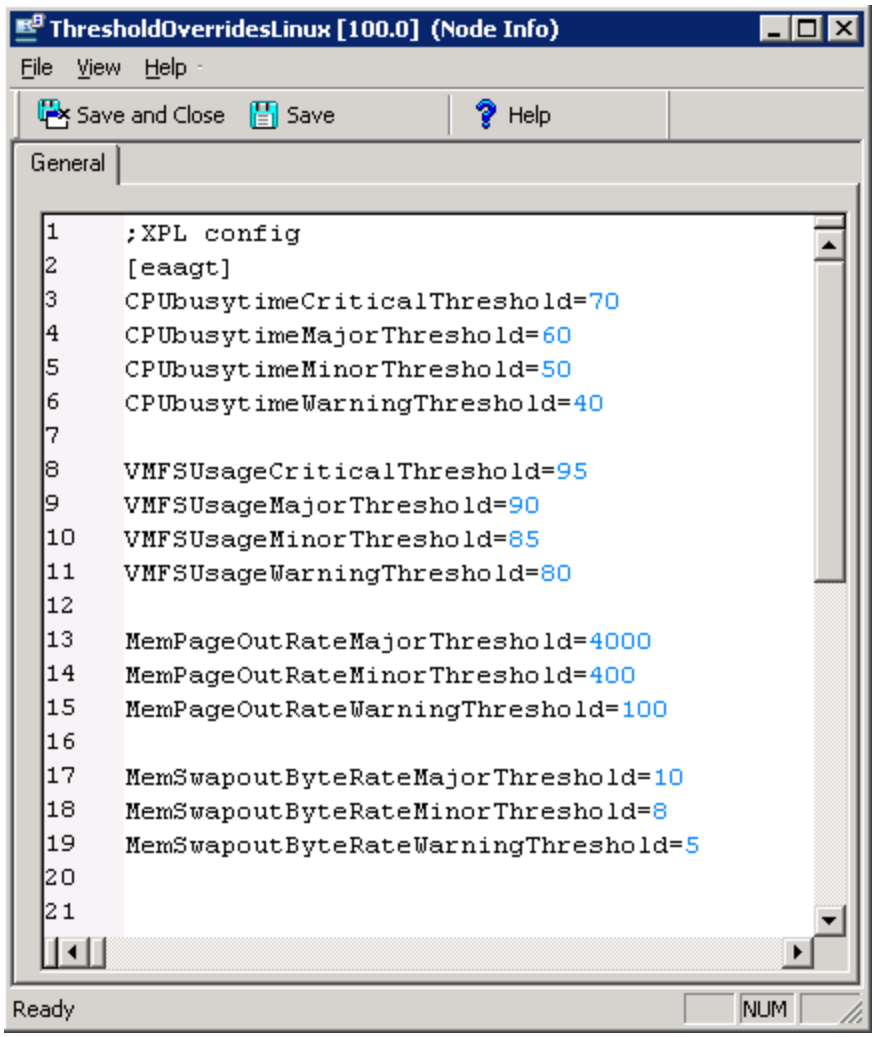

These steps help you to override the threshold settings for the policies on a particular managed node. You can create copies of this policy to set different sets of values on other managed nodes. After specifying the overriding thresholds, make sure that you deploy the policy to the managed nodes.

> **Note:** ThresholdOverrides policy is of the type Node Info. These polices do not generate and send messages to the OM console.

Additionally, you can modify the threshold values in a policy directly by using XPL configuration settings. To view and change the settings for policy thresholds in the XPL, use the following commands:

- View the XPL configuration settings namespace:

  ```
  ovconfget eaagt
  ```

- Change the threshold values:

```
ovconfchg -ns eaagt -set <threshold name> <overriding threshold value>
```

Example:

```
ovconfchg -ns eaagt -set VMFSUsageCriticalThreshold 91
              -set VMFSUsageMajorThreshold 86
              -set VMFSUsageMinorThreshold 81
              -set VMFSUsageWarningThreshold 76
```

Adaptive Thresholds

The adaptive threshold concept is used to determine optimal threshold values by using the historical records for performance characteristics and usage patterns of infrastructure resources, instead of using fixed threshold values specified in the policies. The policies that use adaptive thresholds calculate the average value of the metric during the last hour. The average value is then compared to the data collected in the previous 4 weeks during the same hour or day interval. If there is a significant difference in the value, the policy generates an alert.

Constant threshold values set in the policies are ideal for a specific scenario, but not for all scenarios. It is necessary to change the threshold values according to the type of environment for improved performance of the infrastructure resources. Distributed system environments generally follow predictable trends over time. Adaptive threshold helps to automatically calculate the threshold values according to available performance data for previous weeks.

When policies that use adaptive threshold are deployed on managed nodes, the adaptive threshold script establishes a baseline from the historic samples. The policy stores the value of the metric collected from the previous 4 weeks for the same day and same time slot known as **Historic Data or Baseline Data**. The sampled data is collected from the Embedded Performance Component or Performance Agent. These samples help identify previous trends in infrastructure performance. Based on these trends, the threshold range is automatically calculated.

Current Data is the data collected over the last hour period. Comparing the current performance data with the historical data enables to detect deviations from normal behavior. An alert is generated when an abnormal behavior is detected.

Standard deviation is calculated by comparing the value of current data against the average value of the historic data. If the standard deviation value is high or low it means the metric value varies historically.

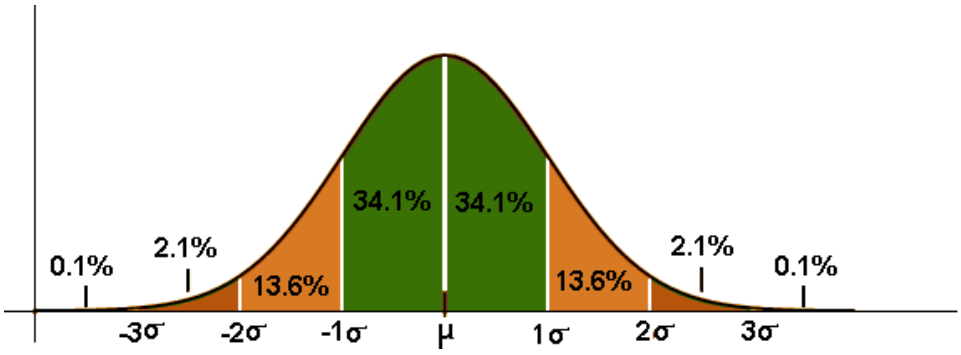

The script compares the current data average with the standard deviation levels, to determine if the current data exceeds or precedes either one standard deviation or two standard deviations of the baseline data.

If the current value exceeds the one standard deviation, but not two standard deviations, then the policy generates warning severity message. Similarly, if the current value exceeds two standard deviations, the policy generates a major severity message. The number of standard deviations and messages severities are configurable using the policy script parameters.

You can pre-define the values for MajorDeviation, WarningDeviation, and MinorDeviation in the policy as part of the script parameters. The parameters display the number of standard deviations from normal, at which the policy generates a major, warning, or minor severity message respectively.

The infrastructure resources that have sporadic usage at intervals greater than one month may not be suitable candidates for implementation of adaptive thresholds. For example, if you have an analysis server that really works hard only during the quarters last week, then the policy that use adaptive threshold will not be suitable. This is because it will compare this activity against the activity during the previous 4 weeks that are not representative of the activity during the last week.

Adaptive Threshold is calculated as follows:

```
Range for normal behavior = Historic Average ± N * Historic Standard Deviation
```

In this instance, Historic Average is the average of the Historic Data and N is the value of the Major, Minor, or Warning Deviation.

For example:

Consider the average for memory utilization of a system during the last hour is 60% and the historic data over the last 4 weeks on the same day and hour averages 40% with a standard deviation of 5. Next, consider that the policy is configured with a value of 3.5 for the parameter MajorDeviation.

As per the Adaptive Threshold calculation:

Range for normal behavior = `40 ± 3.5 * 5 = 57.5-22.5`

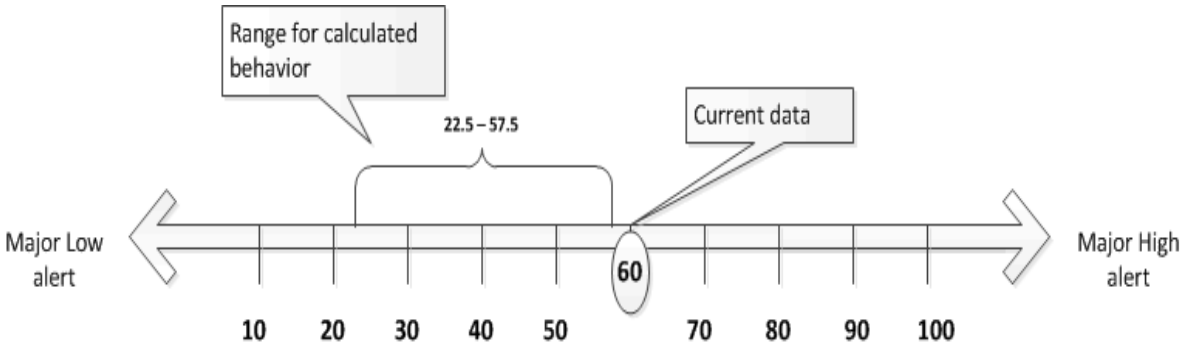

Since standard deviation is 5 and the current value is 60%, it means there was a variation in the current value compared to the average historic data. On a scale of 100, the current data of 60% do not lie in the range of 57.5-22.5%. As a result the policy sends a major high alert. The policy with adaptive threshold perform statistical analysis and is unaware of the nature of the metric or of any environmental changes that may justify the current metric behavior.

Cutoff Threshold

Adaptive threshold triggers alarms when the value observed has significant variance compared to baseline data. For under-utilized systems the standard values observed for statistics under consideration are low. Even a small variance from these historical data seems as a large deviation in comparison and hence adaptive threshold triggers alarms.

For example, the CPU utilization of a system has been recorded in the range of 0.5-4% and the current value is 7%.

Example of CPU utilization graph with data of a under-utilized system

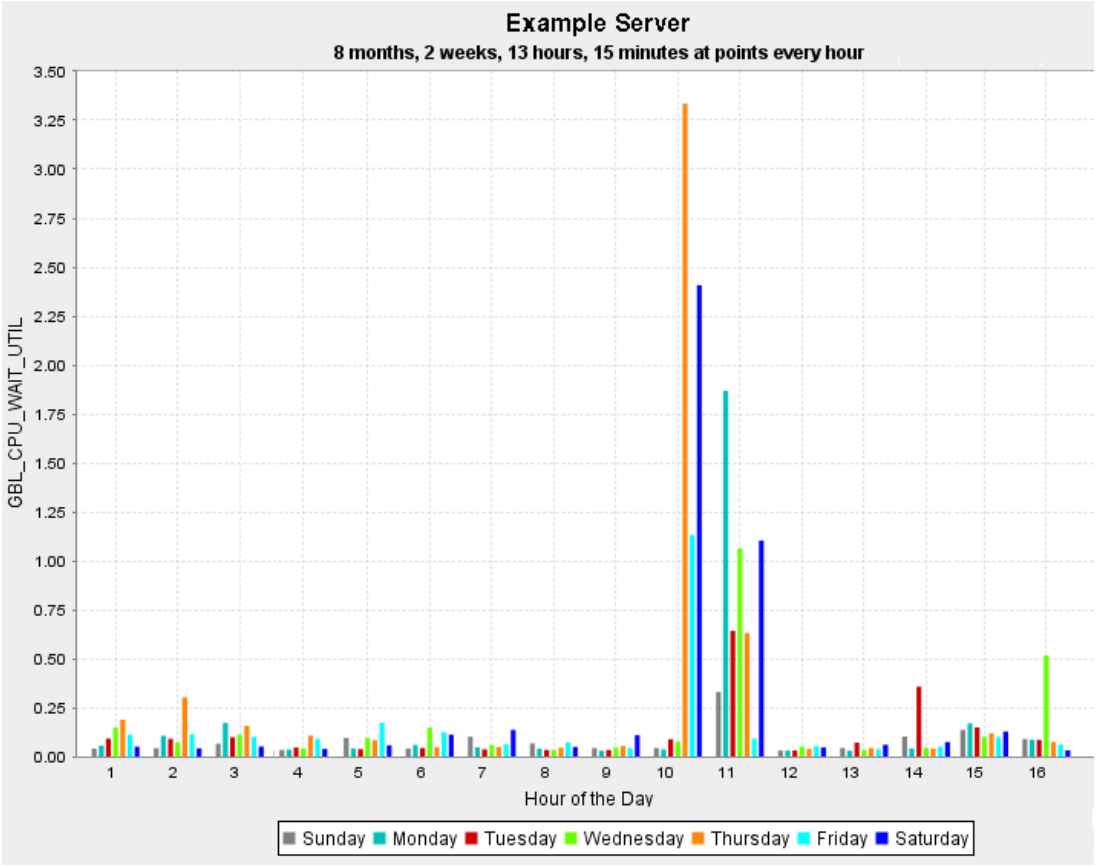

At this level, you may not want to receive an alert for CPU utilization higher than normal. However, if the CPU utilization of the system is in 50% there is cause for concern, and you may want to receive alerts.

To avoid getting adaptive threshold alarms for under-utilized systems, you can define cutoff thresholds for the parameters. The cutoff threshold is used to determine whether the alarm should be raised or not.

The cutoff value (say ‘x’) can be set in the following two ways:

- 'x'
- ‘>x'

  where ‘x’ is a numeric value.

The above setting impacts the AT policy processing, as the CutOff value is verified against the current data sample average and then it is decided whether or not to continue with further processing.

The CutOff value is checked for the following conditions:

- When cutoff value is set to 'x', continue processing only if current data sample average is greater than x.
- When cutoff value is set to '>x', continue processing only if current data sample average is less than x.

For example, the cutoff threshold is set at 50%, the current data sample average value is 7%, and the CPU utilization of a system has been recorded in the range of 0.5-4%. At this level, although the CPU utilization is higher, you will not receive an alert till the utilization levels reach 50% or above.

By default, the cutoff parameter has no value defined. You can decide to assign a value depending upon the system usage trends. The table below contains recommended cutoff values for adaptive threshold policies.

Recommended cutoff values for adaptive threshold policies

Policy Name	Threshold Parameter Name	Recommended Value
SI-SwapUtilization-AT	SwapUtilCutOff	10
SI-PerCPUUtilization-AT	CPUUtilCutOff	50
SI-PerNetifInbyteBaseline-AT	ByNetifInByteCutOff	1000 (~1 KB)
SI-PerDiskUtilization-AT	DiskUtilCutOff	30
SI-MemoryUtilization-AT	MemUtilCutOff	60
SI-PerNetifOutbyteBaseline-AT	ByNetifOutByteCutOff	1000 (~1 KB)
VI-VMCPUEntitlementUtilizationMonitor-AT	CPUEntlUtilCutOff	100
VI-IBMLPARFrameCPUUtilMonitor-AT	*LPARFrameCPUUtilCutOff*	50
VI-HPVMGuestCPUEntlUtilMonitor-AT	*CPUEntlUtilCutOff*	80
VI-IBMLPARCPUEntlUtilMonitor-AT	*CPUEntlUtilCutOff*	80
	CPUEntlUtilCutOff	80
VI-MSHyperVGuestCPUEntlUtilMonitor-AT	*CPUEntlUtilCutOff*	80
VI-OracleSolarisZoneCPUEntlUtilMonitor-AT	*CPUEntlUtilCutOff*	80

Recommended cutoff values for adaptive threshold policies, continued

Policy Name	Threshold Parameter Name	Recommended Value
VI-IBMLPARMemoryEntlUtilMonitor-AT	*MEMEntlUtilCutOff*	80
	MEMEntlUtilCutOff	80
VI-OracleSolarisMemoryEntlUtilMonitor-AT	*MEMEntlUtilCutOff*	80
VI-OracleSolarisZoneSwapUtilMonitor-AT	*SwapUtilCutOff*	50
VI-LinuxVirtDiskPhysByteRateBaseline-AT	DiskPhysbyteCutOff	1000*the number of VMs running on KVM or Xen hosts
VI-LinuxVirtNetByteRateBaseline-AT	NetbyteRateCutOff	1000*the number of VMs running on KVM or Xen hosts.
VI-LinuxVirtGuestCPUTotalUtilMonitor-AT	CPUTotUtilCutOff	50
VI-LinuxVirtVMMemoryUsage-AT	MemUsageCutOff	75

You can set the cutoff values for adaptive threshold policies by using threshold overrides policy. To know more about threshold overrides, see "Adaptive Thresholds" on page 18.

Remote Monitoring

In a network topology, there are clients that are directly monitored by the server. In such a setup it is common that the client in turn hosts or interacts with other servers, disks, or machines known as remote entities. The server interacts with these remote entities indirectly through the clients. Typically, these remote entities are agent-less nodes that are monitored through proxy connections. The remote entities are network addressable entities like cluster nodes, cluster resource groups, virtual machines, or remote drives.

Infrastructure SPIs help you to monitor the remote infrastructure. These are added as managed nodes or nodes monitored as SNMP/message-allowed nodes.

The remote drives mounted on *individual systems* are monitored through the Systems Infrastructure SPI policies. The remote drive space utilization monitor policies monitor space utilization on file shares provided by another system.

In a virtualized environment it is important to monitor virtual machine performance and availability remotely without having the need for running the Operations Agent on them. The Virtualization Infrastructure SPI monitors virtual machines remotely.

In a clustered environment the Cluster Infrastructure policies monitor the nodes and resource groups of a cluster remotely.

When Systems Infrastructure SPI identifies a managed node as a virtualized server or a cluster system, it initiates virtualization discovery policy or cluster discovery policy, as appropriate. To enable assignment of alerts to the appropriate entity, such as the nodes that do not have Operations Agent running on them, can be added as message-allowed nodes in the OM node bank. The Infrastructure SPIs automatically add nodes in the OM node bank in some cases and in some cases it is left to be done manually. To ascertain whether nodes are auto-added in node-bank by Infrastructure SPI or to add nodes manually, refer to the tables below:

Scenarios when nodes are added automatically in the OM node bank

Node Auto-addition Scenarios	Node Type
Cluster Nodes	Message allowed/Managed node
Cluster Resource Groups	Message allowed/Virtual node/Managed node

Scenarios when nodes are not added automatically in the OM node bank

Scenarios when node is not auto-added	Node Type (can be added as)
NFS/Samba (CIFS) share providers	Message allowed/Managed node
Hyper-V guest virtual machines	Message allowed/Managed node

To ensure that the alerts are assigned to the remote system rather than the host from which the monitoring is done, many Infrastructure SPI policies (that are capable of doing remote monitoring) contain a parameter setting *AssignMessageToRemoteHost* that can be set to 0 or 1 depending on the need.

Example of AssignMessageToRemoteHost parameter in a policy

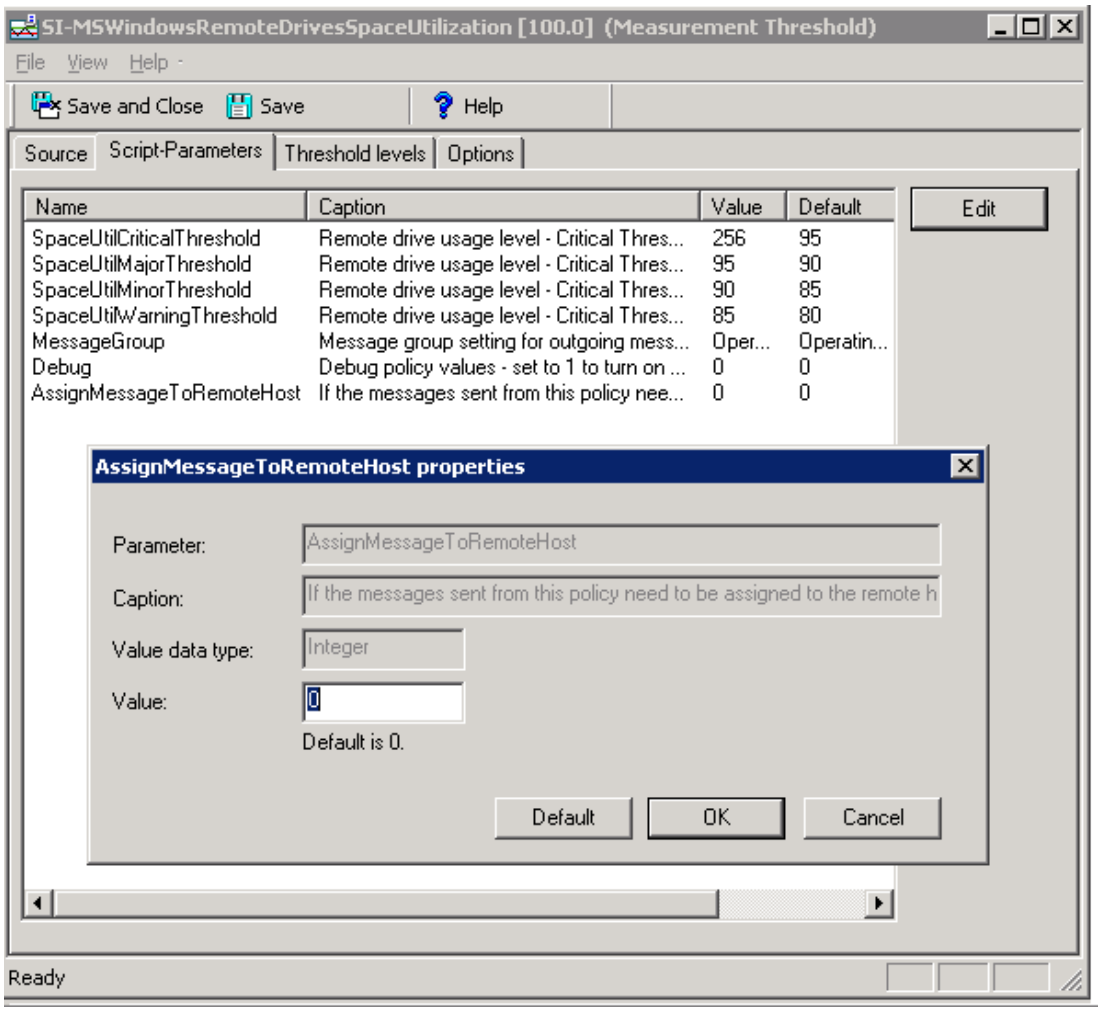

You can set the value to 1 to display the primary node of the alert message as the remote host. By default, the messages are assigned to the managed node from which the message is sent out.

Dialogue with Infrastructure SPI Expert

This section presents some possible scenarios of the Infrastructure SPIs.

In this example, Newbie, a recently employed administrator, has been given the responsibility of installing and deploying the Infrastructure SPIs. The new administrator does not have an idea about Infrastructure SPIs or earlier version of Virtualization Infrastructure SPI. The new administrator approaches InfraSPI Expert, a senior administrator and a power-user of operations management features and functionality, for assistance in understanding the product. They have the following conversations.

Newbie: One of our servers lately seems to be slowing down and response times of the database instance running on the system are getting erratic.

InfraSPI Expert: Did you monitor the CPU and memory utilization for any bottlenecks on the system?

Newbie: No.

InfraSPI Expert: Deploy the Systems Infrastructure SPIs CPU bottleneck diagnosis and memory bottleneck diagnosis policies on the node. These policies report back the top 10 CPU and memory hogging processes respectively. This should help you identify the cause of the problem.

Newbie: Okay.

Newbie: I deployed the policies and now I know the reason for our server slowing down. The Systems Infrastructure SPI sends alert messages to the OM console and I was quickly able to see the root cause of the problem. The messages display the top 10 processes hogging CPU and memory utilization.

InfraSPI Expert: Great going, so what was the issue?

Newbie: There was a rogue application running that was sporadically increasing CPU consumption as well as memory usage. I have fixed the rogue application.

InfraSPI Expert: Okay.

For information about the policies provided by Systems Infrastructure SPI, see *Operations Smart Plug-in for Systems Infrastructure User Guide*.

Newbie: This was a great way of identifying the problem. I have a question though, what do I do for systems that are perennially busy or the ones that are idle?

InfraSPI Expert: For systems that are highly utilized or under-utilized you can deploy the Systems Infrastructure SPIs CPU utilization and memory utilization monitor policies on the node. These are adaptive threshold policies and will give you accurate results.

Newbie: I remember reading about adaptive threshold policies. These policies learn from the performance characteristics and patterns of the previously collected data, and statistically determine if the current utilization is normal or not. The thresholds are automatically calculated according to historic data.

InfraSPI Expert: Yes, that way you get alerts only for actual problem scenarios. An easy way to identify auto threshold determination policies is to look for AT appended at the end of policy name, for example, VI-VMwareVMMemoryUsage-AT.

For more information on how adaptive thresholds are calculated, see "Key Concepts" on page 15.

Newbie: Okay. While going through many policies, I noticed the critical threshold parameter has a strange value “65535” assigned.

InfraSPI Expert: Yes, this value is deliberately assigned to mask the critical threshold parameter. This avoids system generating critical alert messages. We can manually define the critical threshold level if required.

Newbie: But why 65535?

InfraSPI Expert: Just a number! For masking, it can be any number above 100.

Newbie: What about machines which hang or get stuck in a transient stage?

InfraSPI Expert: Good you asked. Virtualization Infrastructure SPI sends out an alert, if any virtual machine is stuck in a transient stage such as starting, snapshots, migrating, saving, and stopping for more than 30 minutes. This policy is very useful to identify any state transitions or issues (if any) with the running of virtual machine.

Newbie: Another thing, what about the planned outages that occur every night for our test virtual machines? This is intentional and I don’t want messages about VM state changes for the scheduled outages. How can I avoid getting alerts for these servers?

InfraSPI Expert: Good question. For nodes that undergo planned outages at predefined time every day, like our test virtual machines that are shutdown at 21:00:00 hours and brought back up at 05:00:00 hours next morning, you can avoid getting state change alerts messages from VI-StateMonitor policy by setting the **AlertOnPlannedOutage** parameter to true. When the policy is deployed with this setting, the VI-StateMonitor policy will not generate VM suspension alert for the monitored node during the specified time duration. For information about the policies provided for monitoring virtual infrastructure, see *Smart Plug-in for Virtualization Infrastructure SPI User Guide.*

Newbie: We have a clustered web server where we want to ensure high availability of the server. We just cannot afford to let it go down. Is there a policy which can help me monitor and raise alerts on critical service levels watermarks like quorum breach or single point of failure situation?

InfraSPI Expert: Yes, Cluster Infrastructure SPI provides the CI-ClusterMonitor policy. This policy monitors for conditions such as:

- Cluster is down or offline.
- Majority of the nodes are down and cluster quorum is not maintained. If (n/2 + 1) nodes are not in active state. For example, if there are six nodes in your cluster, less than 4 nodes are active, the policy will raise an alert.
- There is only 1 active node in the whole cluster, thus being a single point of failure (SPOF) and a potential risk to cluster availability.

The policy will send out alerts for all the above conditions, thus ensuring that any breach of cluster availability watermarks are brought to notice. For information about the policies provided for monitoring cluster infrastructure, see *Smart Plug-in for Cluster Infrastructure SPI User Guide.*

Newbie: Great, thank you for all the useful information!

Send documentation feedback

If you have comments about this document, you can contact the documentation team by email. If an email client is configured on this system, click the link above and an email window opens with the following information in the subject line:

Feedback on Concepts Guide (Operations Smart Plug-ins for Infrastructure 12.04)

Just add your feedback to the email and click send.

If no email client is available, copy the information above to a new message in a web mail client, and send your feedback to docfeedback@hpe.com.

We appreciate your feedback!