HP Operations Smart Plug-ins for Infrastructure

for HP Operations Manager for Windows®, UNIX, and Linux operating systems

Software Version:1.00

Concepts Guide

Document Release Date: September 2009
Software Release Date: September 2009

Legal Notices

Warranty

The only warranties for HP 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. HP 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 HP 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 Notices

Trademark Notices

Microsoft® and Windows® are U.S. registered trademarks of Microsoft Corporation.

UNIX® is a registered trademark of The Open Group.

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

Documentation Updates

The title page of this document contains the following identifying information:

- Software Version number, which indicates the software version.
- Document Release Date, which changes each time the document is updated.
- Software Release Date, which indicates the release date of this version of the software.

To check for recent updates or to verify that you are using the most recent edition of a document, go to:

http://h20230.www2.hp.com/selfsolve/manuals

This site requires that you register for an HP Passport and sign in. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

Or click the **New users - please register** link on the HP Passport login page.

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

Support

Visit the HP Software Support Online web site at:

www.hp.com/go/hpsoftwaresupport

This web site provides contact information and details about the products, services, and support that HP Software offers.

HP Software online support provides customer self-solve capabilities. It provides a fast and efficient way to access interactive technical support tools needed to manage your business. As a valued support customer, you can benefit by using the support web site to:

- Search for knowledge documents of interest
- Submit and track support cases and enhancement requests
- Download software patches
- Manage support contracts
- Look up HP support contacts
- Review information about available services
- Enter into discussions with other software customers
- Research and register for software training

Most of the support areas require that you register as an HP Passport user and sign in. Many also require a support contract. To register for an HP Passport ID, go to:

http://h20229.www2.hp.com/passport-registration.html

To find more information about access levels, go to:

http://h20230.www2.hp.com/new_access_levels.jsp

Contents

1 Introduction

This chapter provides a broad overview of infrastructure management and how HP 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, it assures availability of resources in a complex and decentralized IT infrastructure setup.

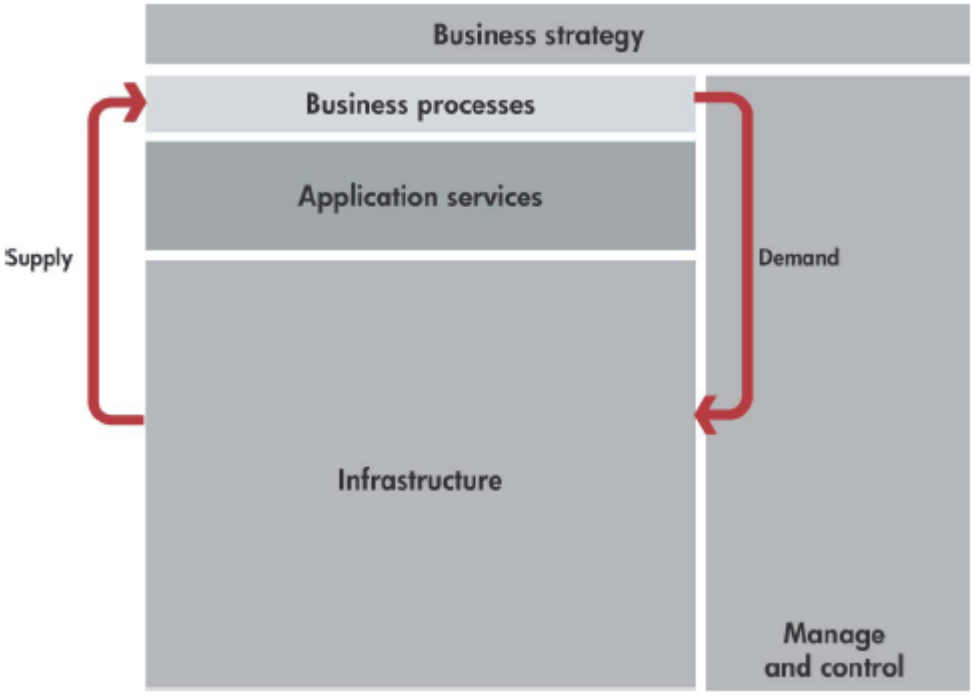

Infrastructure management helps you to monitor, analyze, and optimize the usage of distributed operating systems, application and storage servers, clusters, and virtual machines. It helps to predict infrastructure resource utilization so that IT infrastructure issues can be either avoided or resolved 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 resource 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), TruClusters (for Tru64 UNIX®) 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 specially 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. The 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 a member of the 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.

What are Infrastructure SPIs

The Infrastructure SPIs form a software suite that integrates fully with HP Operations Manager (HPOM) and extends HPOM's management scope to include distributed enterprise-wide base infrastructures of Microsoft Windows and Linux systems. It relates the cross domain IT infrastructure events with relevant applications in the HPOM console, and maps them into a hierarchical service map while eliminating redundant infrastructure and processes.

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 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 are integrated with other HPOM products, such as HP Performance Agent, HP Reporter, and HP Performance Manager.

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.

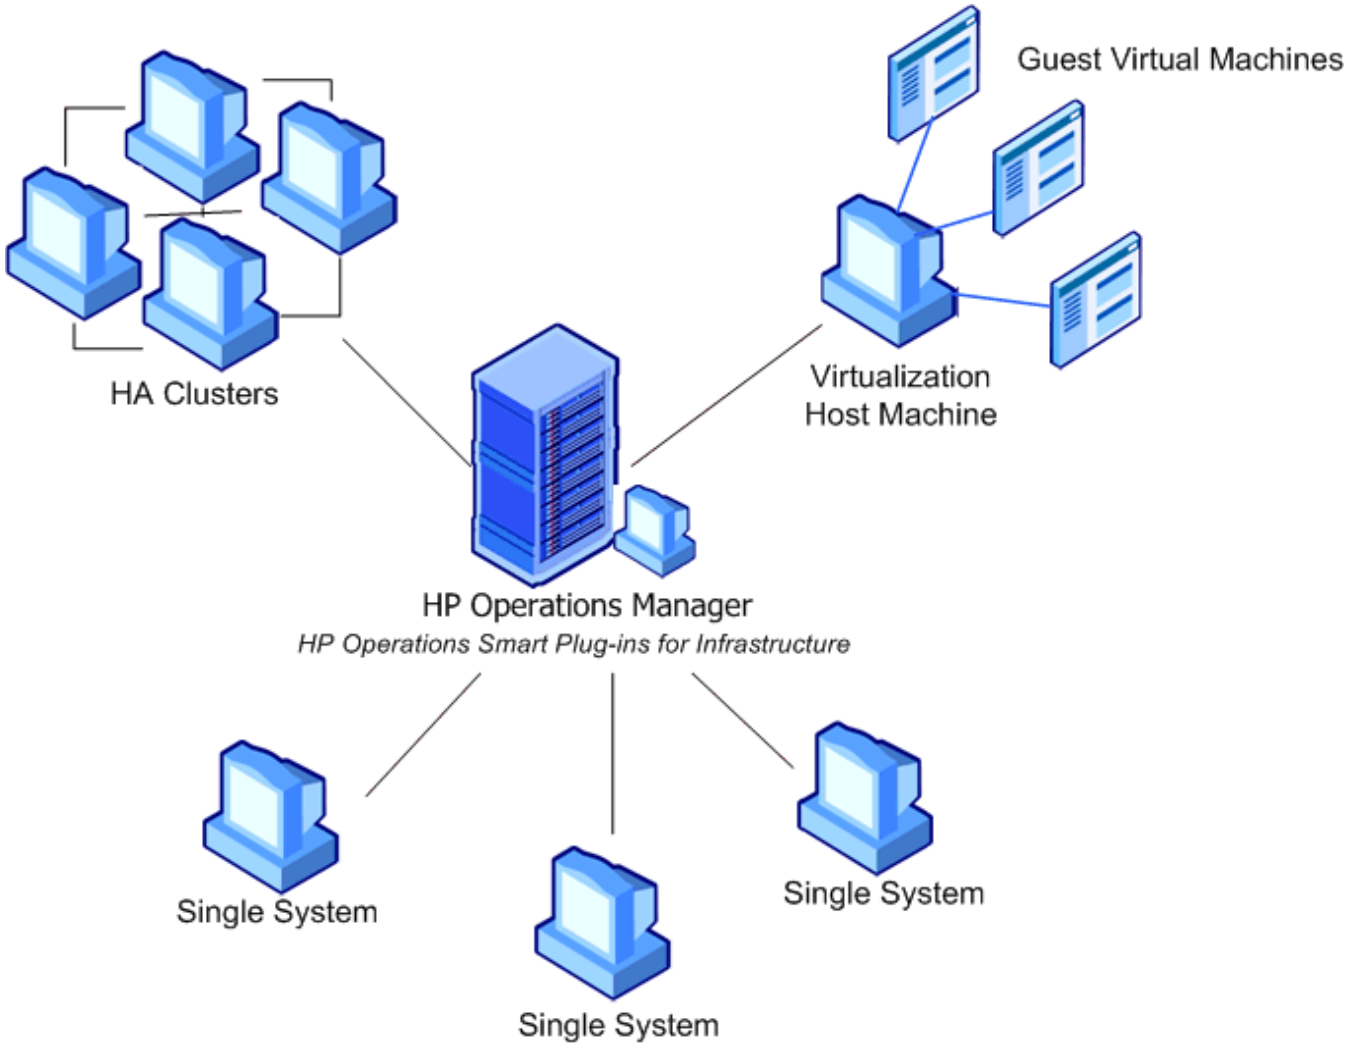

HP Operations Smart Plug-ins for Infrastructure is a software suite comprising of three SPIs: *Smart Plug-in for Systems Infrastructure*, *Smart Plug-in for Virtualization Infrastructure*, and *Smart Plug-in for Cluster Infrastructure*.

To easily find a infrastructure-specific policy, such as one for cluster, you can use the policy folder: **Policy management** → **Policy groups** → **SPI for Infrastructure** → **en** → **Cluster Infrastructure** on HPOM for Windows console, and **Policy Bank** → **SPI for Infrastructure** → **en** → **Cluster Infrastructure** on HPOM for UNIX/Linux console.

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 HPOM 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 HPOM Services area.

If the managed node is VMware vMA or Microsoft Hyper-V, the discovery policy initiates Virtualization Infrastructure SPI discovery. The Virtualization Infrastructure SPI discovers the host machines, virtual machines (guest machines), and the resource pools.

Similarly, 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.

vMA is VMware's vSphere management assistant that includes prepackaged software to manage ESX/ESXi and vCenter Server systems.

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 HPOM 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.

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 HPOM 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 HPOM console, where they are displayed with a color-coded severity level.

Instructions text offers problem resolution suggestion. It is available within the generated message's 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 HPOM 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 HPOM 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.

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 HPOM 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 HP Operations GUI and use the Console Tree to browse to the **SPI for Infrastructure** 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.

Figure 1 Example of modifying threshold for Measurement Threshold policy

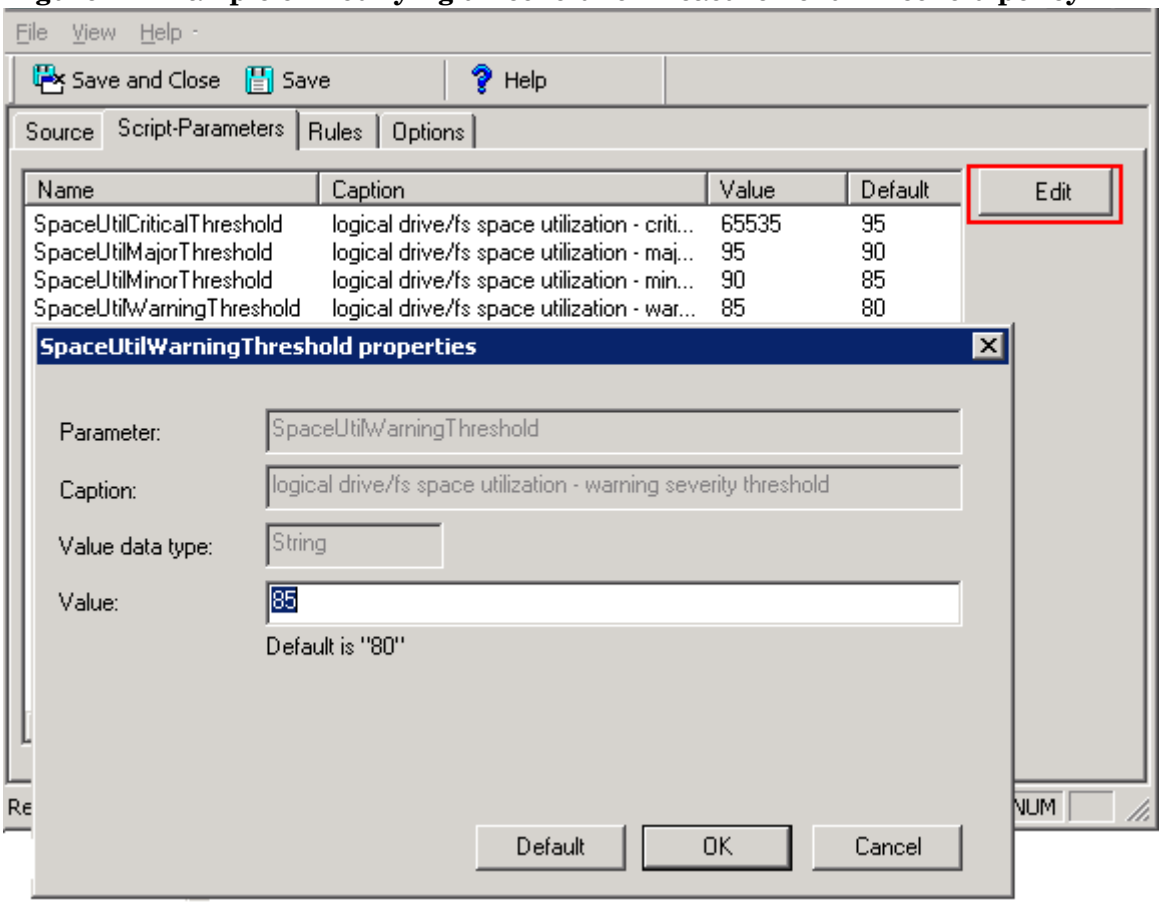

5 Click **OK** to save the changes.

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.

Figure 2 Example of threshold overrides policy

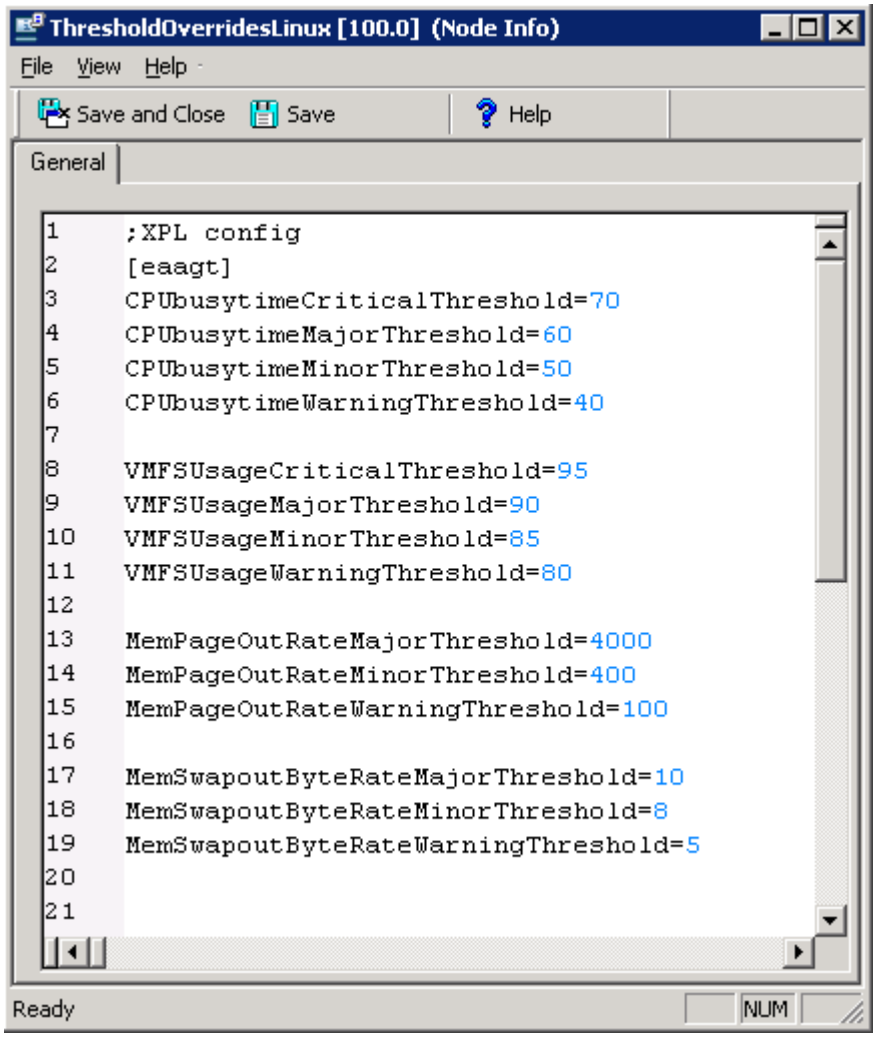

These steps help you to override the threshold settings for the policies on a particular managed node. You can create copies of this policy to set different sets of values on other managed nodes. After specifying the overriding thresholds, make sure that you deploy the policy to the managed nodes.

ThresholdOverrides policy is of the type Node Info. These polices do not generate and send messages to the HPOM 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 threshold values using the historical records for performance characteristics and usage patterns of infrastructure resources, instead of using fixed threshold values specified in the policies.

Constant threshold values set in the policies may be ideal for one situation but not for all situations. Therefore, it may be necessary to change the threshold values according to the environment setup for infrastructure performance. 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 days.

When policies that use adaptive threshold are deployed on managed nodes, the adaptive threshold script establishes a baseline from the historic samples. The sampled data is collected from the HP Embedded Performance Component or HP Performance Agent. These samples help identify previous trends in infrastructure performance. Based on these trends the threshold range is automatically calculated. Comparing the current performance data with the adaptive thresholds indicates whether the current infrastructure resource utilization is normal or not. An alert is generated when abnormal behavior is detected.

Standard deviation is calculated by subtracting each individual sample from the average of the entire sample set, squaring the result, and adding each square to an aggregate counter. The aggregate count is divided by the number of samples minus one. Standard deviations have a trust factor denoting the variance level that goes up with the number of points in the statistical sample and the standard deviation compared to the maximum value. This trust factor dictates if an alert is generated when abnormal behavior is detected or not. If a value has a trust factor of 40% or lower, the alert is not processed. This eliminates false alarms when there is insufficient data.

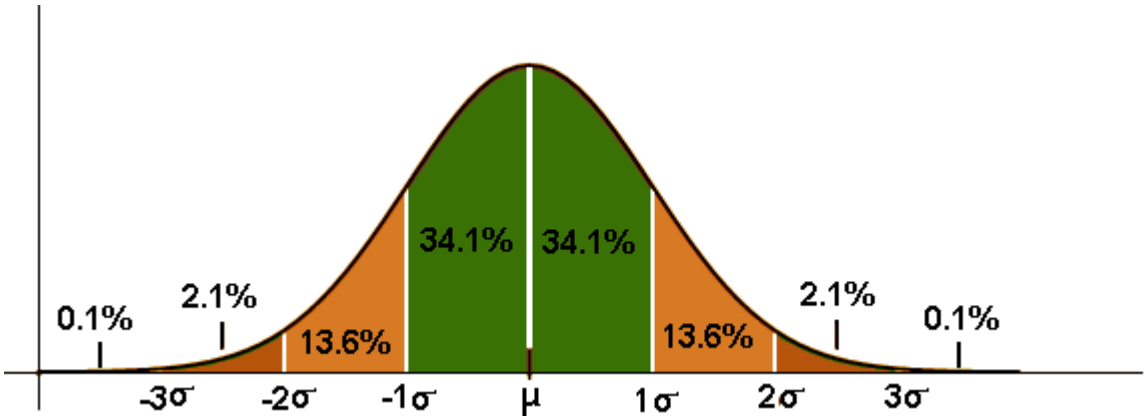

The script compares the current data average with the standard deviation levels, to determine if the current data exceeds or precedes either one standard deviation or two standard deviations of the baseline data.

The data within one standard deviation, represented in the diagram from -1σ to μ and μ to 1σ, is not processed due to t he *trust* factor. 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.

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 only really works hard during quarter end processing then the baseline data will not be helpful in calculating an adaptive threshold.

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%.

Figure 3 Example of CPU utilization graph with data of a under-utilized system

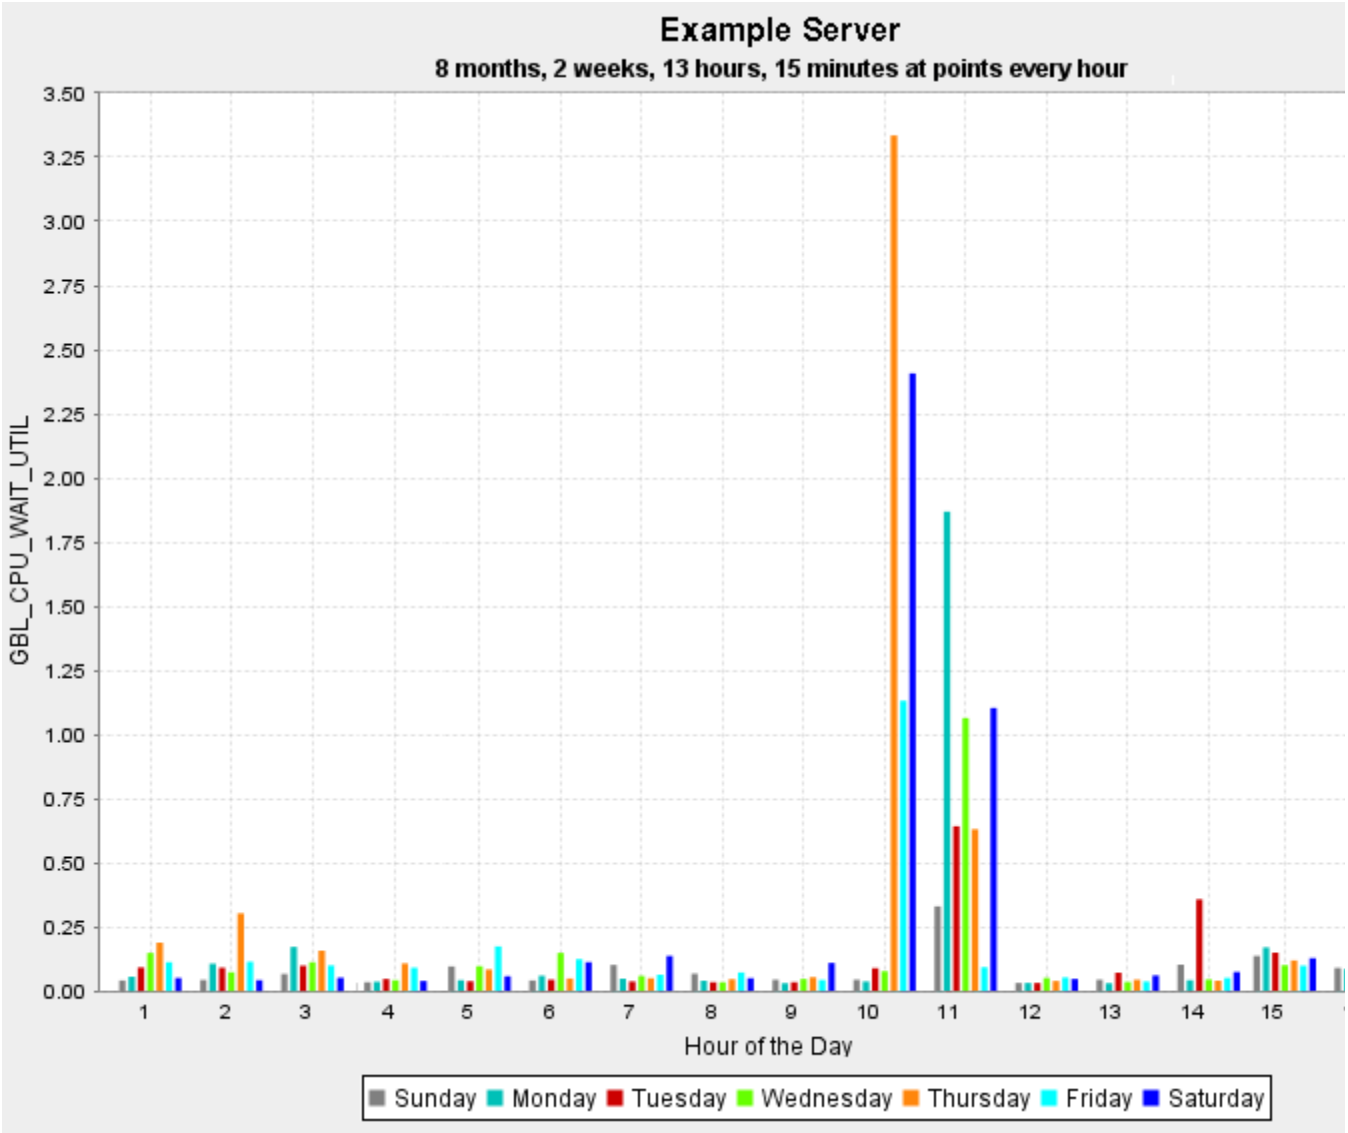

At this level you may not want to receive an alert for CPU utilization higher than normal. However, if the CPU utilization of the system is in 50% there is cause for concern, and you may want to receive alerts.

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.

For example, the CPU utilization of a system has been recorded in the range of 0.5-4%, the current value is 7%, and the cutoff threshold is set at 50%. At this level, although the CPU utilization is higher than the recorded data, 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.

Table 1 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-VMwareVMMemoryUsage-AT	MemUsageCutOff	75
VI-VMCpuEntitlementUtilizationMonitor-AT	CPUEntlUtilCutOff	100
VI-VMwareHostDiskUtilization-AT	HostDiskUtilCutOff	40
VI-VMwareNetifOutbyteBaseline-AT	HostNetifOutbyteCutOff	1000*the number of VMs running on ESX/ ESXi hosts
VI-VMwareNetifInbyteBaseline-AT	HostNetifInbyteCutOff	1000*the number of VMs running ESX/ ESXi hosts

You can set the cutoff values for adaptive threshold policies by using threshold overrides policy. To know more about threshold overrides, see Customizing Threshold Values using Threshold Overrides.

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 adressable 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 HP Operations Agent on them. The Virtualization Infrastructure SPI monitors virtual machines remotely. In the case of VMware monitoring the Virtualization Infrastructure SPI monitors remote ESX and ESXi hosts from the vMA.

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 HP Operations Agent running on them, can be added as message-allowed nodes or SNMP nodes in the HPOM node bank. The Infrastructure SPIs automatically add nodes in the HPOM 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:

Table 2 Scenarios when nodes are added automatically in the HPOM node bank

Node Auto-addition Scenarios	Node Type
Cluster Nodes	SNMP/Message allowed/Managed node
Cluster Resource Groups	SNMP/Message allowed/Virtual node/Managed node
ESX hosts registered to vMA	SNMP/Message allowed/Managed node
ESXi hosts registered to vMA	SNMP/Message allowed node
vCenter hosts registered to vMA	SNMP/Message allowed/Managed node

Table 3 Scenarios when nodes are not added automatically in the HPOM node bank

Scenarios when node is not auto-added	Node Type (can be added as)
NFS/Samba (CIFS) share providers	SNMP/Message allowed/Managed node
vCenter hosts managing ESX/ESXi hosts registered to vMA but such vCenter hosts are not registered with vMA	SNMP/Message allowed/Managed node
Hyper-V guest virtual machines	SNMP/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.

Figure 4 Example of AssignMessageToRemoteHost parameter in a policy

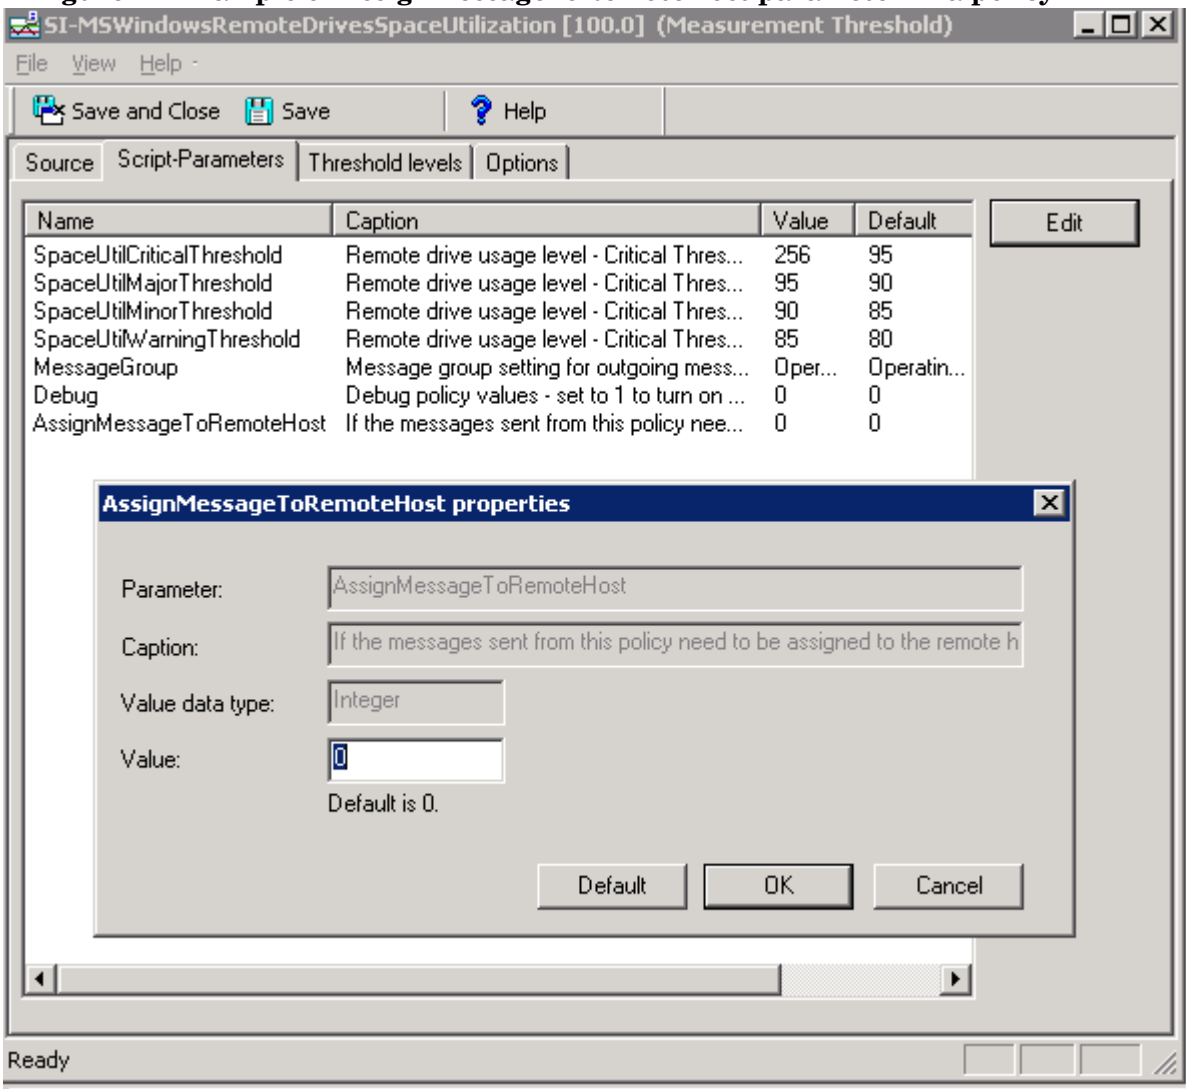

You can set the value to 1 to display the primary node of the alert message as the remote host. By default, the messages are assigned to the managed node from which the message is sent out.

4 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. She 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 SPI's 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 HPOM 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 *HP 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 SPI's 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, SI-SwapUtilization-AT and VI-VMwareHostDiskUtilization-AT.

For more information on how adaptive thresholds are calculated, see Adaptive Thresholds.

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 out 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: Also I noticed the node grouping for Infrastructure SPI managed nodes is according to the operating system. I wonder what is the reason for this.

InfraSPI Expert: Grouping the managed nodes according to operating system and platforms has many benefits. It aids better classification of the nodes, and allows easy assignment and deployment of policies. You can assign policies to the group at one go instead of individually assigning or deploying them. Another reason for grouping the nodes like this is to separate them from the Operating System SPIs (HP Operations Smart Plug-in for Microsoft Windows Operating System and HP Operations Smart Plug-in for UNIX Operating Systems). Let me show you how the node groups appear on HPOM for Windows and UNIX/ Linux.

The node groups in HPOM for UNIX/ Linux appear like this:

Label	Name	Nodes
InfraSPI Managed Nodes - MS Windows	InfraSPI Managed Nodes - MS Windows	2
InfraSPI Managed Nodes - Red Hat Enterprise Linux	InfraSPI Managed Nodes - Red Hat Enterprise Linux	1
InfraSPI Managed Nodes - VMware vMA	InfraSPI Managed Nodes - VMware vMA	1
linux	linux	2

On the other hand, the node groups in HPOM for Windows are nested and appear like this:

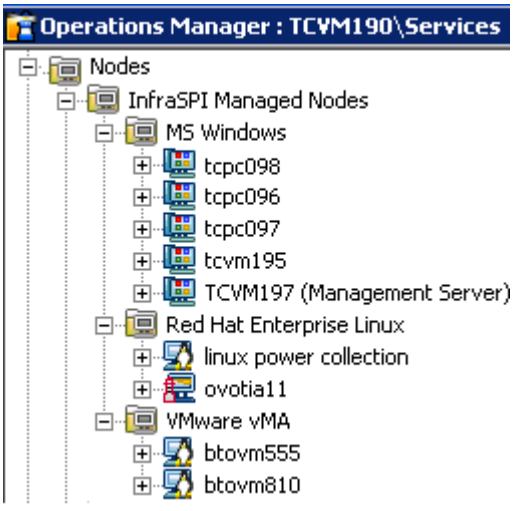

Newbie: Okay. I was going through the release notes and found that Virtualization Infrastructure is version 1.5 and Infrastructure SPI is 1.0. What has changed from Virtualization Infrastructure SPI version 1.0 to 1.5?

InfraSPI Expert: Virtualization Infrastructure SPI version 1.0 was an independent standalone SPI, used to monitor the ESX servers in the data center. Now it must be upgraded to the current version of Virtualization Infrastructure SPI (version 1.50) that is a part of Infrastructure SPIs 1.00. The current version of Virtualization Infrastructure SPI monitors VMware ESX/ ESXi servers, Microsoft Hyper-V servers, and the virtual machines hosted on these servers. For more information about the changes in VISPI and its architecture, see *HP Operations Smart Plug-ins for Infrastructure Installation Guide*. The guide covers upgradation from Virtualization Infrastructure SPI to Infrastructure SPIs.

Newbie: Okay! How can I use VM state monitor policy?

InfraSPI Expert: VI-StateMonitor policy reports the state of host server and guest virtual machines configured on them. Basically the policy monitors and reports any state changes for the virtual machines.

Newbie: What about any machines which are hanged or 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, snapshoting, 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 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 information!

We appreciate your feedback!

If an email client is configured on this system, by default an email window opens when you click on the bookmark “Comments”.

In case you do not have the email client configured, copy the information below to a web mail client, and send this email to **docfeedback@hp.com**

Product name:

Document title:

Version number:

Feedback: