

Foglight® for Infrastructure 7.1.0
User and Reference Guide



© 2023 Quest Software Inc.

ALL RIGHTS RESERVED.

This guide contains proprietary information protected by copyright. The software described in this guide is furnished under a software license or nondisclosure agreement. This software may be used or copied only in accordance with the terms of the applicable agreement. No part of this guide may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's personal use without the written permission of Quest Software Inc.

The information in this document is provided in connection with Quest Software products. No license, express or implied, by estoppel or otherwise, to any intellectual property right is granted by this document or in connection with the sale of Quest Software products. EXCEPT AS SET FORTH IN THE TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST SOFTWARE ASSUMES NO LIABILITY WHATSOEVER AND DISCLAIMS ANY EXPRESS, IMPLIED OR STATUTORY WARRANTY RELATING TO ITS PRODUCTS INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. IN NO EVENT SHALL QUEST SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, CONSEQUENTIAL, PUNITIVE, SPECIAL OR INCIDENTAL DAMAGES (INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOSS OF PROFITS, BUSINESS INTERRUPTION OR LOSS OF INFORMATION) ARISING OUT OF THE USE OR INABILITY TO USE THIS DOCUMENT, EVEN IF QUEST SOFTWARE HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest Software makes no representations or warranties with respect to the accuracy or completeness of the contents of this document and reserves the right to make changes to specifications and product descriptions at any time without notice. Quest Software does not make any commitment to update the information contained in this document.

If you have any questions regarding your potential use of this material, contact:

Quest Software Inc.
Attn: LEGAL Dept.
4 Polaris Way
Aliso Viejo, CA 92656

Refer to our website (<https://www.quest.com>) for regional and international office information.

Patents

Quest Software is proud of our advanced technology. Patents and pending patents may apply to this product. For the most current information about applicable patents for this product, please visit our website at <https://www.quest.com/legal>.

Trademarks

Quest, the Quest logo, and Where next meets now are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. "Apache HTTP Server", Apache, "Apache Tomcat" and "Tomcat" are trademarks of the Apache Software Foundation. Google is a registered trademark of Google Inc. Android, Chrome, Google Play, and Nexus are trademarks of Google Inc. Red Hat, JBoss, the JBoss logo, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries. CentOS is a trademark of Red Hat, Inc. in the U.S. and other countries. Fedora and the Infinity design logo are trademarks of Red Hat, Inc. Microsoft, .NET, Active Directory, Internet Explorer, Hyper-V, Office 365, SharePoint, Silverlight, SQL Server, Visual Basic, Windows, Windows Vista and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. AIX, IBM, PowerPC, PowerVM, and WebSphere are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Java, Oracle, Oracle Solaris, PeopleSoft, Siebel, Sun, WebLogic, and ZFS are trademarks or registered trademarks of Oracle and/or its affiliates in the United States and other countries. SPARC is a registered trademark of SPARC International, Inc. in the United States and other countries. Products bearing the SPARC trademarks are based on an architecture developed by Oracle Corporation. OpenLDAP is a registered trademark of the OpenLDAP Foundation. HP is a registered trademark that belongs to Hewlett-Packard Development Company, L.P. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. MySQL is a registered trademark of MySQL AB in the United States, the European Union and other countries. Novell and eDirectory are registered trademarks of Novell, Inc., in the United States and other countries. VMware, ESX, ESXi, vSphere, vCenter, vMotion, and vCloud Director are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Sybase is a registered trademark of Sybase, Inc. The X Window System and UNIX are registered trademarks of The Open Group. Mozilla and Firefox are registered trademarks of the Mozilla Foundation. "Eclipse", "Eclipse Foundation Member", "EclipseCon", "Eclipse Summit", "Built on Eclipse", "Eclipse Ready", "Eclipse Incubation", and "Eclipse Proposals" are trademarks of Eclipse Foundation, Inc. IOS is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. Apple, iPad, iPhone, Mac OS, Safari, Swift, and Xcode are trademarks of Apple Inc., registered in the U.S. and other countries. Ubuntu is a registered trademark of Canonical Ltd. Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. OpenSUSE, SUSE, and YAST are registered trademarks of SUSE LCC in the United States and other countries. Citrix, AppFlow, NetScaler, XenApp, and XenDesktop are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. AlertSite and DéjàClick are either trademarks or registered trademarks of Boca Internet Technologies, Inc. Samsung, Galaxy S, and Galaxy Note are registered trademarks of Samsung Electronics America, Inc. and/or its related entities. MOTOROLA is a registered trademark of Motorola Trademark Holdings, LLC. The Trademark BlackBerry Bold is owned by Research In Motion Limited and is registered in the United States and may be pending or registered in other countries. Quest is not endorsed, sponsored, affiliated with or otherwise authorized by Research In Motion Limited. Ixia and the Ixia four-petal logo are registered trademarks or trademarks of Ixia. Opera, Opera Mini, and the O logo are trademarks of Opera Software ASA. Tevtron, the Tevtron logo, and CitraTest are registered trademarks of Tevtron, LLC. PostgreSQL is a registered trademark of the PostgreSQL Global Development Group. MariaDB is a trademark or registered trademark of MariaDB Corporation Ab in the European Union and United States of America and/or other countries. Vormetric is a registered trademark of Vormetric, Inc. Intel, Itanium, Pentium, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Debian is a registered trademark of Software in the Public Interest, Inc. OpenStack is a trademark of the OpenStack Foundation. Amazon Web Services, the "Powered by Amazon Web Services" logo, and "Amazon RDS" are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Infobright, Infobright Community Edition and Infobright Enterprise Edition are trademarks of Infobright Inc. POLYCOM®, RealPresence® Collaboration Server, and RMX® are registered trademarks of Polycom, Inc. All other trademarks and registered trademarks are

property of their respective owners.

Legend

- **WARNING:** A WARNING icon indicates a potential for property damage, personal injury, or death.

- ⚠ **CAUTION:** A CAUTION icon indicates potential damage to hardware or loss of data if instructions are not followed.

- ⓘ **IMPORTANT NOTE, NOTE, TIP, MOBILE, or VIDEO:** An information icon indicates supporting information.

Foglight for Infrastructure User and Reference Guide
Updated - 2023
Foglight Version - 7.1.0
Cartridge Version - 7.1.0

Contents

Using Foglight for Infrastructure	7
Host availability alerting	7
Virtual hosts and reported metrics	8
Exploring the Infrastructure Environment dashboard	8
Accessing the Infrastructure dashboard	9
Selecting a service	9
Running a report for the Infrastructure Environment	9
Exploring the Monitoring tab	10
Exploring monitored hosts	14
Exploring the FAQ Question Viewer	16
Exploring the Administration tab	18
Using Foglight for Infrastructure agents	29
Configuring Multiple Instances for Monitoring using Silent Installation	30
Using the PowerVM HMC agent	32
Using the UnixAgentPlus	32
Monitoring remote hosts	32
Deploying Foglight for Infrastructure agents	33
Creating agent instances	33
Defining credentials	33
Activating the agent	34
Monitoring the infrastructure	34
Adding a monitored host	34
Adding multiple monitored hosts	43
About the WindowsAgent	49
Supported platforms	50
Agent properties	50
About the UnixAgentPlus	55
Supported platforms	56
Agent properties	56
About the UnixAgent	62
Supported platforms	62
Agent properties	62
About the MultiHostProcessMonitorAgent	68
Supported platforms	68
Agent properties	68
Monitoring log files with Foglight Log Monitor	73
Configuring monitoring agents	73
Investigating log records	77
Configuring agent properties	79
Configuring File Log Monitor agent properties	79
Configuring Windows Event Log Monitor agent properties	85
Configuring connections to remote Windows platforms	89
Foglight Log Monitor command shell types	89

Foglight Log Monitor remote command access permissions	90
Foglight Log Monitor remote connection requirements	90
Monitoring IBM PowerVM environments	92
Before you begin	92
Introducing the PowerVM infrastructure	92
Managing PowerVM HMC agents	93
Configuring HMC user accounts	93
Configuring PowerVM HMC agent credentials	93
Creating PowerVM HMC agents	94
Reviewing and editing PowerVM HMC agent properties	95
Monitoring your PowerVM environment	97
Exploring the Infrastructure Environment dashboard	97
Investigating the collective use of all HMC resources	103
Viewing HMC details	105
Identifying top consumers of managed server resources	106
Viewing individual managed server details	108
Identifying top consumers of PowerVM partition resources	110
Viewing individual PowerVM partition details	111
Identifying top consumers of PowerVM virtual I/O server resources	114
Viewing individual PowerVM VIOS details	115
Investigating additional managed server, partition, and VIOS details	117
Reviewing frequently asked questions	139
Advanced system configuration and troubleshooting	145
Advanced system configuration for WinRM	145
Adding a non-administrative user to user groups	145
Setting WinRM RootSDDL for a non-administrative user	146
Granting permission to the namespace	147
Granting permission to the service	147
Additional WinRM configuration in FIPS-compliant mode	148
Configuring default local user credentials for Infrastructure Agents	149
Reference	153
Foglight for Infrastructure views	153
Quick View	153
Summary view	154
Resource Utilizations view	156
Host Monitor views	158
Processes views	162
Foglight Log Monitor views	164
Rules	168
Metrics	168
AIX metrics	168
HP-UX metrics	169
Linux metrics	170
Solaris metrics	171
Windows metrics	172

AIXHostDetails topology object	172
AIXMemoryDetails topology object	173
AIXPhysicalCPUUsage topology object	174
AIXUsedMemoryDetails topology object	174
CPUCounts topology object	174
Host topology object	175
HostCPUs topology object	177
HostNetwork topology object	178
HostProcess topology object	180
HostProcessInstance topology object	182
HostService topology object	184
HostTopProcessEntry topology object	185
LPARNetworkInterfaceDetails topology object	185
LPARPhysicalDiskDetails topology object	185
LogicalDisk topology object	186
MSCluster topology object	188
MSClusterNode topology object	189
MSClusterResource topology object	189
MSClusterResourceGroup topology object	189
MShostedResourceGroup topology object	190
Memory topology object	190
NetworkInterface topology object	192
NetworkInterfaceDetails topology object	194
NixHostProcessInstance topology object	195
OperatingSystem topology object	195
PhysicalDisk topology object	196
PowerVMVIOStopology object	198
PowerVMVIOStopology object	199
PowerVMVIOStopology object	200
PowerVMVIOStopology object	200
PowerVMVIOStopology object	201
Processor topology object	201
WinHostProcessInstance topology object	203
Appendix: Building regular expressions in Foglight	204
What is a regular expression?	204
Where can I find regular expressions?	204
Regular expression basics	205
Building a simple pattern	206
Building a pattern that matches a specific character	207
Building a pattern that matches multiple characters	207
Using advanced quantifiers	207
Using special characters and regular expression flags	208
Grouping elements in a pattern	208
Additional information	209
About Us	210
Technical support resources	210

Using Foglight for Infrastructure

This *User and Reference Guide* describes the dashboards included with Foglight for Infrastructure and contains instructions for configuring Foglight for Infrastructure Agents, used by the cartridge to monitor the health of your system infrastructure. It also provides reference information about views, rules, and metrics that are included with Foglight for Infrastructure.

This guide is intended for any user who wants to know more about the Agent properties and dashboards included with Foglight for Infrastructure. It is also meant for users who want to learn about the rules that are included with Foglight for Infrastructure, and the data collected by monitoring agents.

Foglight Log Monitor is an Infrastructure utility shipped with Foglight. Foglight Log Monitor allows you to monitor the contents of the log files on your system and to identify activities that can lead to performance problems.

Better management of your system can be achieved when you are alerted of potential problems before the system stability is affected. This ensures consistent performance of your system at established service levels. This guide describes the dashboards included with Foglight Log Monitor.

Foglight for Infrastructure monitors your infrastructure environment and helps you analyze performance issues affecting various hosts. It may be used by itself to monitor physical hosts, or in combination with other agents to monitor virtual hosts. It contains the following major components:

- Infrastructure Environment dashboard: provides a visual representation of the hosts being monitored in your environment. For more information, see [Exploring the Infrastructure Environment dashboard](#).
- Foglight for Infrastructure agents: collect information about the monitored hosts. Agents communicate with the Foglight Management Server using the Foglight Agent Manager. For more information, see: [Using Foglight for Infrastructure agents](#), [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

For more information about Foglight for Infrastructure, see the following topics:

- [Host availability alerting](#)
- [Virtual hosts and reported metrics](#)

Host availability alerting

Foglight for Infrastructure 5.8.5.3 has changed the way it alerts users about host availability. It is now able to identify and report the following three states for a monitored host:

- **MONITORED.** This is the default state. It is set on the *Host* object for each successful collection topology submission.
- **UNMONITORED.** It is set on the *Host* object when the agent fails to authenticate, or is otherwise unable to connect to the monitored host.
- **UNAVAILABLE.** It is set on the *Host* object when agent is configured to validate host availability with `ping` (by setting the **Use ping to validate host availability** property for that agent to “True”), and a `ping` command is unsuccessful. For details about configuring agent properties, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

i | **NOTE:** When the **Use ping to validate host availability** property is enabled on a UNIX® platform, the `sudoer` file needs to be configured to allow the ICMP process to run with `NOPASSWD`.

In order to allow for accurate host reporting, the agent is able to submit an unmonitored collection only if it had previously connected to the host during its current activation cycle. When the agent connects to the monitored host, a cached host instance is retained (and refreshed each successful collection cycle) as a reference, should the host become unavailable at a future date.

Virtual hosts and reported metrics

Several Foglight components, including Foglight for Virtualization components (Foglight for VMware, Foglight for Hyper-V, Foglight for OpenStack) and Foglight for Infrastructure report metrics for CPU, Memory, Network, and Disk on a monitored host. However, where both virtualization monitoring agents and Infrastructure HostAgents are set up to collect metrics, those reported metrics would not agree. Virtualization monitoring agent metrics are the more accurate values for monitoring purposes as they come from the virtual hypervisor, for example, an ESX host via vCenter®. vCenter accounts for virtualization and its reported metrics should be considered correct.

There are a number of reasons why the reported infrastructure metrics for a virtual host are not as accurate as those of the virtualization cartridge. A guest OS can only identify the virtual hardware resources allocated to the virtual machine on which it is installed and not the actual total physical hardware resources available to the system. Also, the hypervisor adds overhead which manifests in the reported metrics as increased utilization. For more information, see [CPU usage of an application differs in virtual machines compared to physical machines \(2032\)](#).

If, in addition to the virtualization monitoring agent, you set up an Infrastructure agent to collect core metrics, you run the risk of having two agents reporting different values to the same host topology object. This can create frequent topology changes which impacts Management Server performance.

It is recommended that you disable core IC metric collection (CPU, Memory, Network, Disk) for a guest OS system in favour of the virtualization monitoring agent. You can still use the IC agent to collect process and Windows Service data, if desired.

The following virtualization cartridges are recommended for monitoring virtual hosts:

- [Foglight for VMware](#)
- [Foglight for OpenStack](#)
- [Foglight for Hyper-V](#)

Exploring the Infrastructure Environment dashboard

The Infrastructure Environment dashboard helps you monitor, analyze, and investigate the performance of your infrastructure environment.

It allows you to determine potential bottlenecks in your system's performance.

For more details, see these topics:

- [Accessing the Infrastructure dashboard](#)
- [Selecting a service](#)
- [Running a report for the Infrastructure Environment](#)
- [Exploring the Monitoring tab](#)
- [Exploring monitored hosts](#)
- [Exploring the FAQ Question Viewer](#)
- [Exploring the Administration tab](#)

Accessing the Infrastructure dashboard

To access the Infrastructure dashboard:

- On the navigation panel, under Dashboards, click **Infrastructure**.
The Infrastructure Environment dashboard appears, displaying the Windows® service (by default).

The dashboard contains the following components:

- **Select a Service** drop-down list: allows you to select the service that you want to monitor. For more information, see [Selecting a service](#).
- **Reports** button : allows you to run a report for the Infrastructure Environment dashboard. For more information, see [Running a report for the Infrastructure Environment](#).
- **Add OS Monitor**: provides access to the *Add Monitored Host* and *Add Monitored Host - List* wizards. For more information about these wizards, see [Adding a monitored host](#) and [Adding multiple monitored hosts](#).
 - **NOTE:** If Foglight for SNMP is installed on your server, the **SNMP Hosts** button appears in the **Add OS Monitor** section. For more information about this functionality, see the Foglight for SNMP product documentation.
 - **IMPORTANT:** The **Add OS Monitor** section is hidden if the server is a federation server.
- **Monitoring** tab: displays the domains associated with the selected service; allows you to monitor the objects associated with these domains. For more information, see [Exploring the Monitoring tab](#).
- **FAQs** tab (FAQ Question Viewer): displays a list of questions that helps you investigate performance problems with the objects defined in the selected service/domain. For more information, see [Exploring the FAQ Question Viewer](#).
- **Administration** tab: allows you to edit the agent properties for a selected agent or agents. For more information, see [Exploring the Administration tab](#).

Selecting a service

To filter by service:

- 1 On the navigation panel, under **Dashboards**, click **Infrastructure**.
The Infrastructure Environment dashboard appears.
- 2 From the Select a Service list, select the service for which you want to see the infrastructure information.
The Monitoring tab, the FAQ Question Viewer, and the Administration tab are automatically updated to display the information specific to the selected service.

- **NOTE:** For detailed information about monitoring **PowerVM** services, see the [Monitoring IBM PowerVM environments](#).

Running a report for the Infrastructure Environment

Reports are a convenient way to share data about your monitored environment with others in your organization.

To run the report associated with the Infrastructure Environment dashboard:

- 1 In the Infrastructure Environment dashboard, click **Reports** Reports on the upper-right corner.
 - 2 On the list that appears, click **Hosts**. This is the report associated with the Infrastructure Environment dashboard.
- The Hosts wizard appears.
- 3 On the Set Input Parameters page, select the input parameters for the report from the *Time Range* and the *Service* lists, then click **Next**.
 - 4 On the Set Properties page, enter a name for the report, select a format for it, enter the email addresses of the people who should receive this report, schedule the report delivery (optionally) and the number of copies to be retained, then click **Next**.

If you selected to schedule the report, the Select Schedule page appears. Continue with Step 5.

If you did not select to schedule the report, the Summary page appears. Continue with Step 6.

- 5 On the Select Schedule page, select a schedule type from the list of available options, or click **New Schedule** to define a custom schedule, then click **Next**.
- 6 On the Summary page, review the settings defined, then click **Finish**.

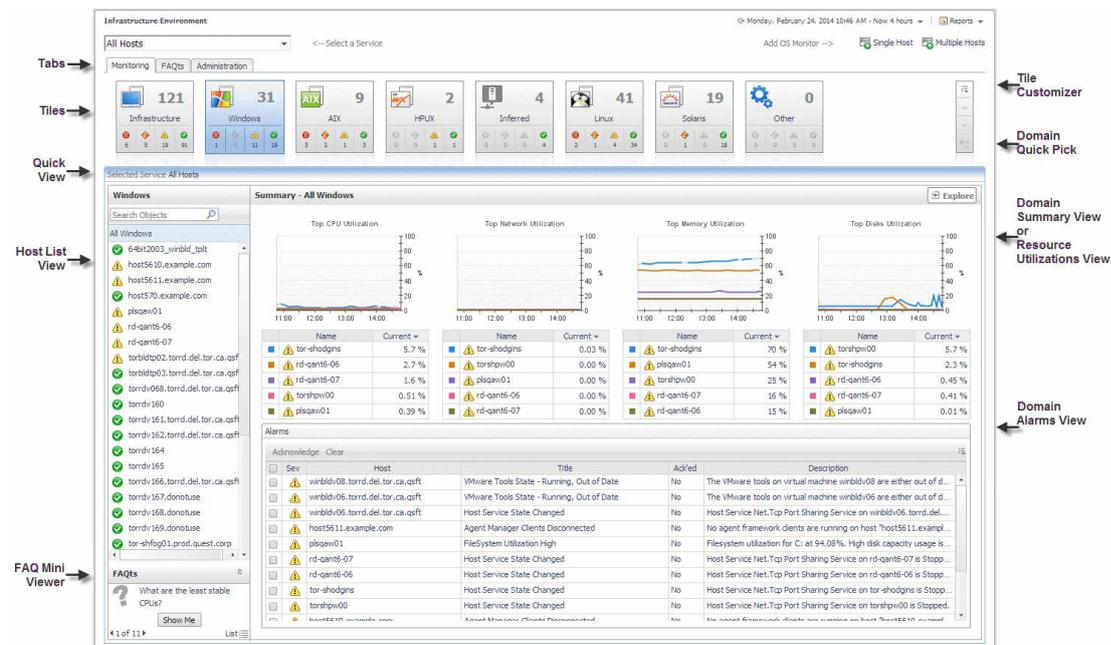
The report is generated and delivered to the recipients indicated in the report settings.

Exploring the Monitoring tab

The Monitoring tab contains several tiles (at the top of the tab), which represent the domains associated with the selected service. When you click a tile, the Quick view (at the bottom of the tab) is refreshed to display the list of objects associated with the selected domain and a summary of the domain's host status.

NOTE: The following figure illustrates the information displayed in the Monitoring tab, for the Windows® service and the domains included in this service, as shown in the tiles. This information is just an example; it may vary depending on the selected service, selected domain, and the monitored hosts. When you install Foglight for Infrastructure, the Windows service is displayed by default.

Figure 1. Infrastructure Environment - Monitoring tab



Tiles

To customize the tiles section:

- 1 Click the **Tile Customizer**  icon.
A dialog box appears, displaying all the tiles available for the domain viewer.
- 2 Select the tile that you want to display by clicking its icon once; click the icon again to toggle the selection.
i | **NOTE:** The currently displayed tile cannot be removed from the list of the objects associated with a domain.
- 3 Click **Apply**.
The tiles area is updated to display the selected tiles.

To quickly find a tile and display the information for that domain:

- 1 Click the **Quick Pick**  icon.
A dialog box appears, displaying the list of all tiles displayed for that domain.
- 2 Select the tile that you want to display by clicking its icon once.
The Quick view is updated to display the information for the selected domain.

Quick View

The Quick View consists of several views:

- **Host List** view: displays the list of objects associated with the selected domain.
- **FAQts** view (**FAQ Mini Viewer**): contains the list of questions relevant to the selected domain.
- **Domain Summary** view: displays the utilization summary for all hosts monitored in the selected domain.
- **Resource Utilizations** view: displays the resource utilization values for the selected monitored host.

Domain Summary view

The *Domain Summary* view is the default view displayed when you select a tile in the Monitoring tab.

- i** | **TIP:** If the *Resource Utilizations* view is displayed for a selected monitored host, you can return to displaying the *Domain Summary* view by clicking **All <Domain_Name>** in the *Host List* view.

At the top of the *Domain Summary* view, you can see information about the top resource users (that is, *Top CPU Utilization*, *Top Network Utilization*, *Top Memory Utilization*, and *Top Disks Utilization*). The number of top utilization indicators to be displayed in these charts and tables is defined by settings the monitoring agent properties. For more information, see [Creating agent instances](#).

Clicking any data series on the charts allows you to drill down into the Metric Analyzer dashboard for the associated metric and view the metrics collected for that topology object.

At the bottom of the view, you can see an embedded *Alarms* view, which displays the list of alarms fired for all hosts in the selected domain. To drill down into the *Alarm* view for a specific host, click the alarm status in the alarms table. This *Alarm* view shows the history of key metrics used in the rules that produced the alarm. You can click the various tabs to review additional diagnostic details, if required.

From the *Domain Summary* view, you can drill down into the All Hosts view by clicking the **Explore**  button in the upper-right corner of the view.

Resource Utilizations view

To display the Resource Utilizations view:

- In the *Host List* view, click the name of the host.
 - If the host is not being monitored, the *Resource Utilizations* view displays a warning message. You can add the selected host to the list of monitored hosts by clicking **Configure Host Monitoring** . The Add Monitored Host wizard appears. For more information, see [Adding a monitored host](#).
 - If the host is being monitored, the *Resource Utilizations* view shows the *CPU*, *Memory*, *Network*, and *Storage* utilization charts for the selected host, and the *Other Alarms* table (which displays all alarms on the selected host, other than CPU, Memory, Network, and Storage alarms).

Figure 2. Resource Utilizations view

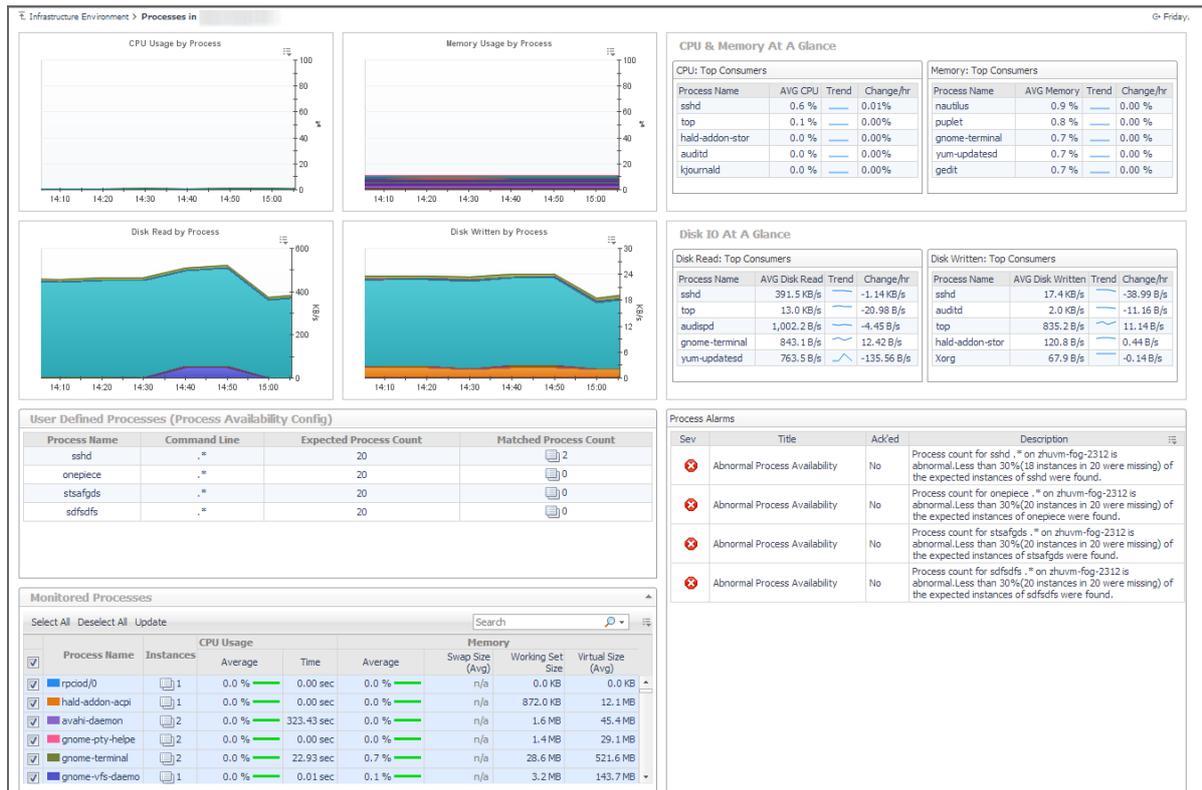


Clicking any data series on the charts allows you to drill down into the Metric Analyzer dashboard for the associated metric and view the metrics collected for that topology object.

From the *CPU*, *Memory*, and *Disk* views you can drill down into the *Processes* dashboard, by clicking the *CPU Usage by Process*, *Memory Usage by Process*, and *Disk IO by Process* links respectively. The *Processes* dashboard provides a summary of CPU, memory, and disk usage by process, as well as a table of process details. If a host does not have process data (for example, because it is monitored by a virtual agent which does not collect process data), then it is possible to create an agent to monitor processes directly from this drill-down view.

- **NOTE:** The drill-down mechanism may change in a future release when tab-based views are expected to become available under *Explore*. The *Processes* dashboard is moved from Foglight for Guest Process Investigation to Foglight for Infrastructure. In Foglight for Virtualization, Enterprise Edition, Foglight for Infrastructure is pre-installed and contributes a *Processes* tab to the **VMware Explorer**. This replaces the tab provided by Foglight for Guest Process Investigation in previous releases. To access this tab, on the navigation panel, under *Dashboards*, click **VMware > VMware Explorer**, and in the *Virtual Infrastructure* view, select a monitored ESX® host or a virtual machine. This feature allows you to monitor guest processes on virtual machines, and drill down into the processes running within virtual machines, to better understand internal performance and resource utilization. System administrators can use the Guest Processes Investigation views to quickly determine which processes are impacting VM or physical host performance and which VMs and physical hosts are utilizing resources ineffectively. If a host is not configured for monitoring, click the **Add Monitoring** button in the *Processes* tab to open the configuration wizard. For more information, see [Adding a monitored host](#).

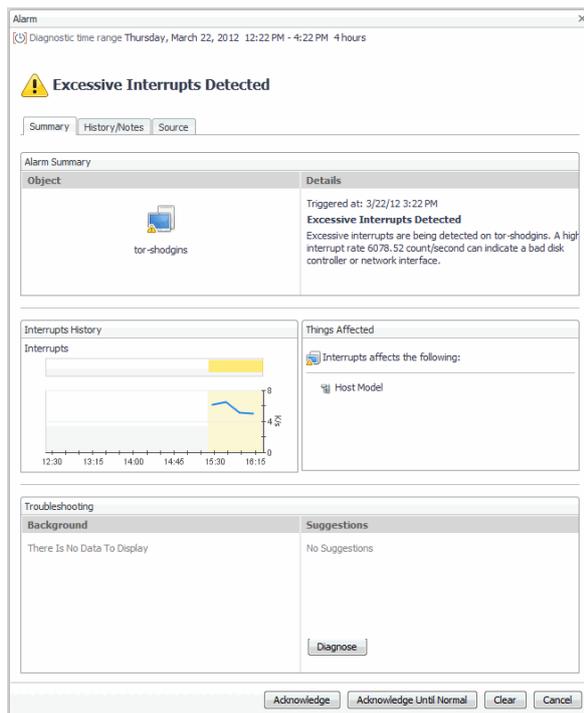
Figure 3. Processes dashboard



From the *Resource Utilizations* view you can drill down into the selected Host view by clicking the **Explore**  button in the upper-right corner of the view. For more information, see [Exploring monitored hosts](#).

From the *Resource Utilizations* view for a monitored host you can drill down into more detailed alarm views, by clicking a specific alarm icon.

Figure 4. Alarm view



The *Alarm* view shows the history of key metrics used in the rules that produced the alarm. You can click the various tabs to review additional diagnostic details, if required.

FAQ Mini Viewer

To see the answer to frequently asked questions relevant to a selected domain:

- 1 In the FAQ Mini Viewer, click the left or right arrow to navigate through the list of questions associated with the selected domain.

The question is displayed in the FAQ Mini Viewer.

- 2 To see the answer to the selected question, click **Show Me**.

A dialog box appears, displaying the answer to the selected question.

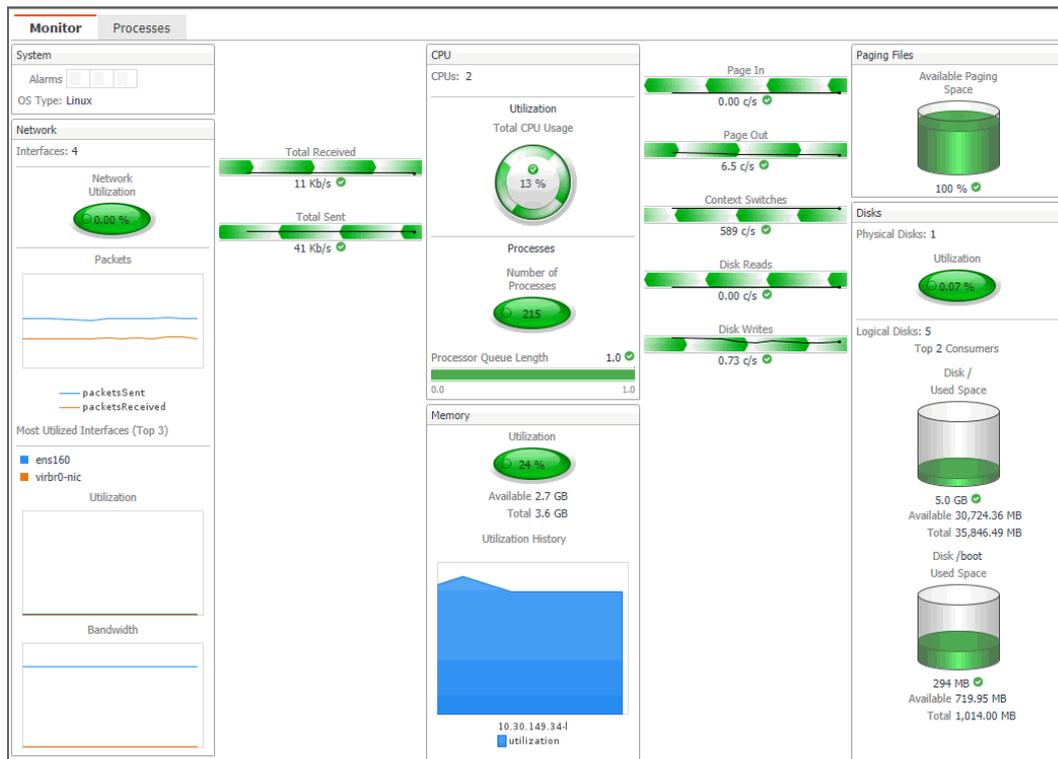
Exploring monitored hosts

When you select a monitored host and display its resource utilization views in the Quick View, you can drill down to a detailed view to explore that host's activity. For example, selecting a Windows® host and choosing **Explore** displays a monitoring dashboard that helps you understand the state of the host's resources and how they affect your monitored system as a whole. Along with displaying the system, network, CPU, memory, disk usage metrics, and any related alarms, this intuitive dashboard connects these visual elements with a series of graphical flows that illustrate how quickly the hosts transmits and processes data in real time. For example, you can review the rates of incoming and outgoing data and how they affect your network resources.

i **TIP:** The value of Processor Queue Length metric is color-coded to indicate the threshold severity. For generic information about metrics' thresholds, see "Working with Metric Thresholds" in the Administration and Configuration Help.

TIP: The value of Available Paging Space metric is color-coded to indicate the threshold severity. For generic information about metrics' thresholds, see "Working with Metric Thresholds" in the Administration and Configuration Help.

Figure 5. Exploring hosts

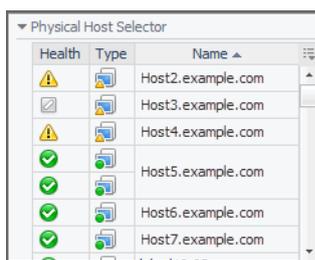


After exploring the hosts, there are two tabs displayed on the top-left corner:

- The Monitor tab. For complete information about the Monitor views and metrics appearing on this dashboard, see [Host Monitor views](#) on page 158.
- The Process tab. For complete information about the Monitor views and metrics appearing on this dashboard, see [Processes views](#) on page 162.

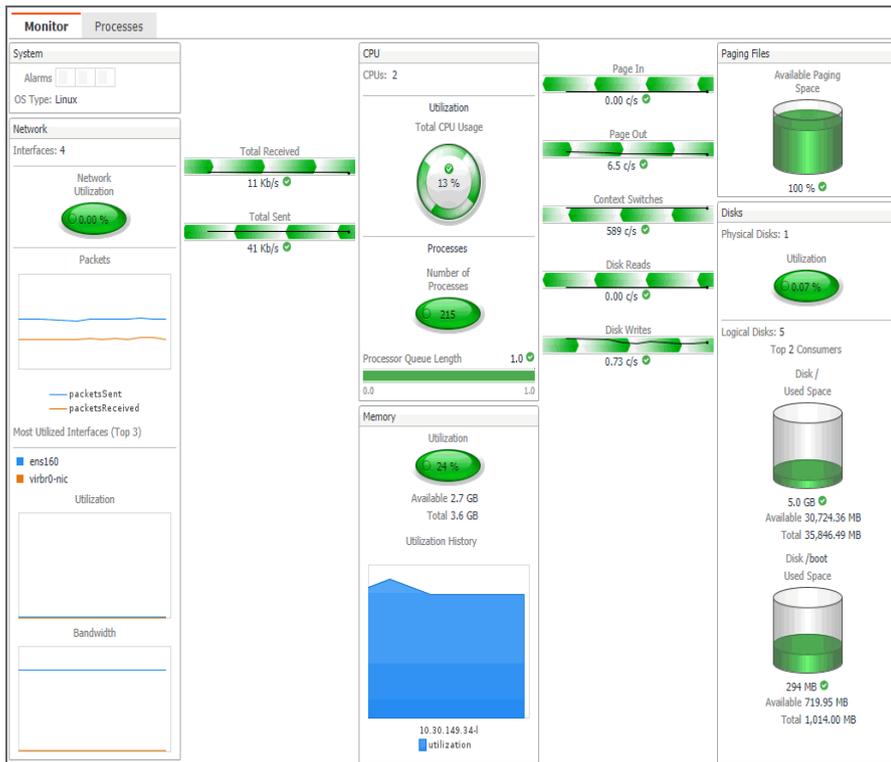
This dashboard also provides the **Physical Host Selector** on the navigation panel, which allows you to quickly drill down on a different host and investigate its state, without returning to the Quick View.

Figure 6. Physical Host Selector



To drill down on a host:

- 1 On the Infrastructure Environment dashboard, click the tile reflecting the domain of the host that you want to drill down to. For example, to drill down on a Windows® host, click the **Windows** tile.
- 2 In the Quick View, in the left pane, select the host.
The Quick View refreshes, showing the **Resource Utilizations** view on the right.
- 3 In the top-right corner of the **Resource Utilizations** view, click **Explore**.
The host details appear in the display area.



TIP: These views also appear when you drill down on a monitored host from the Hosts dashboard. For more information about this dashboard, see the [Host Monitor views](#).

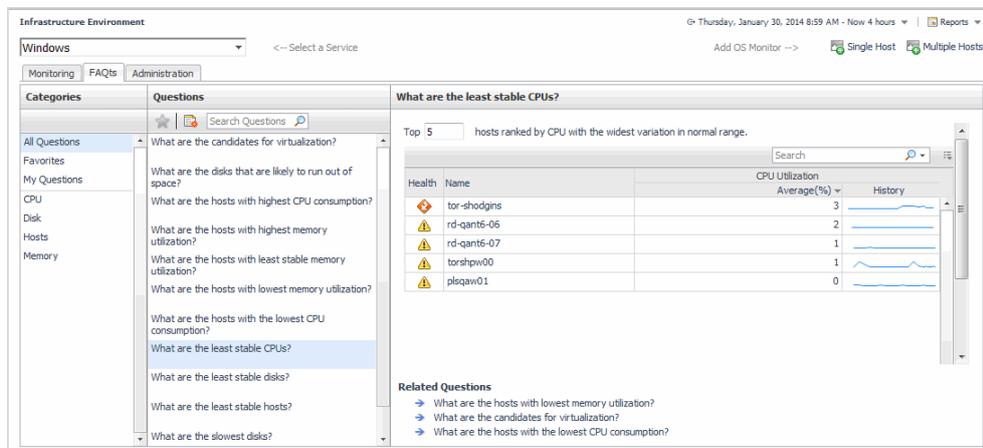
- Optional.** Drill down on a different host. On the navigation panel, under **Physical Host Selector**, click a host.

The display area refreshes, showing the monitoring statistics about the selected host.

Exploring the FAQ Question Viewer

The FAQ Question Viewer allows you to browse your infrastructure environment by selecting a starting point. A set of relevant questions is available in this interface, to help you narrow down the list of objects of interest. The number of questions may change depending on the selected service and the selected domain.

Figure 7. FAQ Question Viewer



This tab is made up of the following embedded views:

- [Categories](#)
- [Questions](#)
- [Answers](#)

Categories

This view lists the categories for which questions can be answered by Foglight for Infrastructure.

Click a category in the list to select it.

Questions

This view lists the questions, for the category selected in the [Categories](#), that can be answered by Foglight for Infrastructure. A set of default questions are included with Foglight for Infrastructure.

Click a question in the list to select it.

To mark a question as your favorite:

- Select the question that you want to mark as your favorite in the Questions view, and click **Mark as Favorite** .

i | **NOTE:** To see all of your favorites, click **Favorites** in the *Categories* view.

To run a report based on selected questions:

- 1 Click **Create Report** .

The **Create Report** wizard appears.

- 2 Select the check boxes for the questions you want to add to your report and click **Finish**.

The report is displayed in the My Report dashboard.

i | **NOTE:** You can run or schedule a report. For more information, see the “Generating a Report” and “Scheduling a Report” topics in the *Foglight User Guide*.

To search for a question:

- If the list of questions is long and you want to narrow it down, search for a particular text string using the **Search Questions** field.

Answers

This view provides an answer to the question selected in the [Questions](#) view. The answer appears on the right side of the dashboard, and it corresponds to the time interval specified in the Time Range (top right corner of the dashboard).

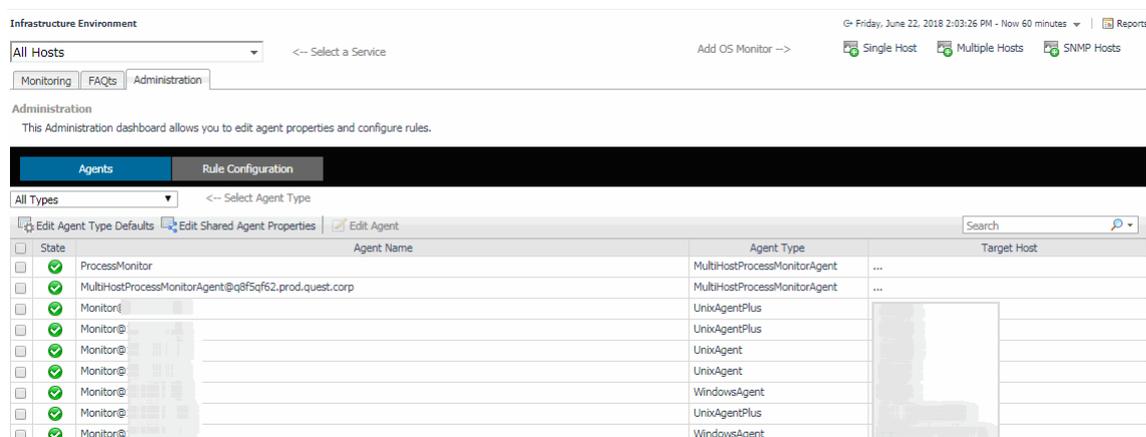
When no objects in your environment match the selected question and the selected time range, no data is displayed in the Answers table. Extend the time interval by selecting a different option from the Time Range, or select a different question.

To view detailed information about one of the monitored hosts, click the hosts's name in the table. This allows you to drill down to the **Infrastructure Environment > Host_Name** dashboard.

Exploring the Administration tab

The Administration tab allows you to manage infrastructure agent instances, you can edit the default shareable and agent properties, configure agent properties that apply only to a specific agent instance. This tab also contains the link of infrastructure rules that you can manage them and create custom rule for specific topology object.

Figure 8. Infrastructure Environment - Administration tab



The Administration tab consists of the following two embedded views:

- [Agents](#)
- [Rule Configuration](#)

Agents

The Agents view shows all the infrastructure agents monitored in the Infrastructure environment. The Agents table refreshes when a different agent type is selected from the **Selected Agent Type** drop-down list.

The Agents table includes the following columns:

- **State:** icon indicating the state of the agent
- **Agent Name:** the name of the monitoring agent
- **Agent Type:** the type of the monitoring agent
- **Target Host:** the name of the host monitored by the agent

The Administration view's toolbar includes the following buttons:

- **Edit Agent Type Defaults:** allows you to edit the default agent type properties.
- **Edit Shared Agent Properties:** allows you to edit the shared agent type properties.
- **Edit Agent:** allows you to edit the agent properties for the selected agent or agents.

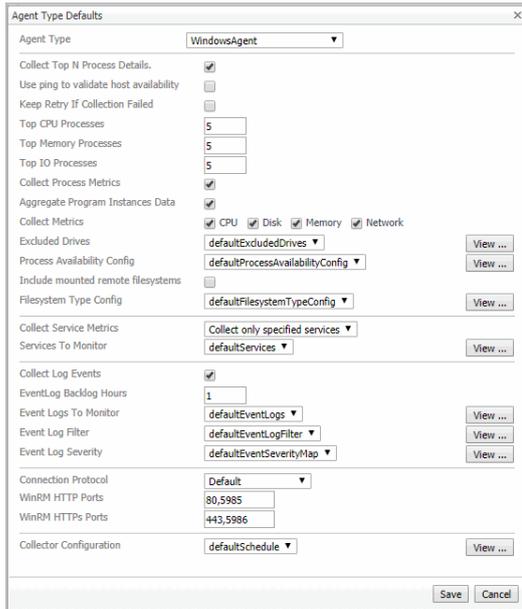
To edit the default agent type properties:

- 1 Click **Edit Agent Type Defaults**.

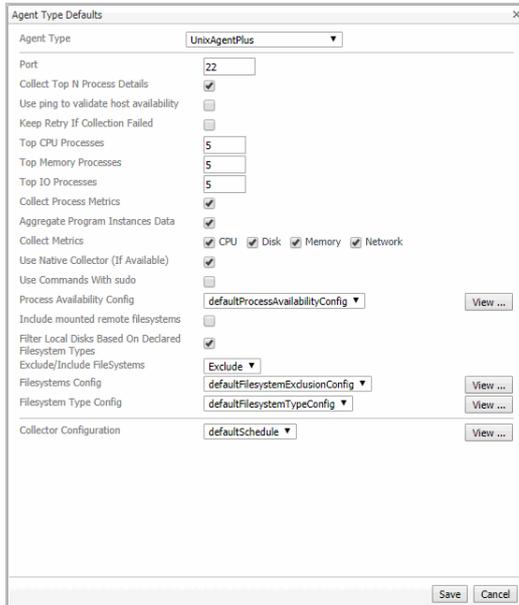
The Agent Type Defaults dialog box appears.

- 2 From the Agent Type list, select whether you want to change the default properties for a Windows, UnixAgentPlus, Unix, or MultiHostProcessMonitor agent type.

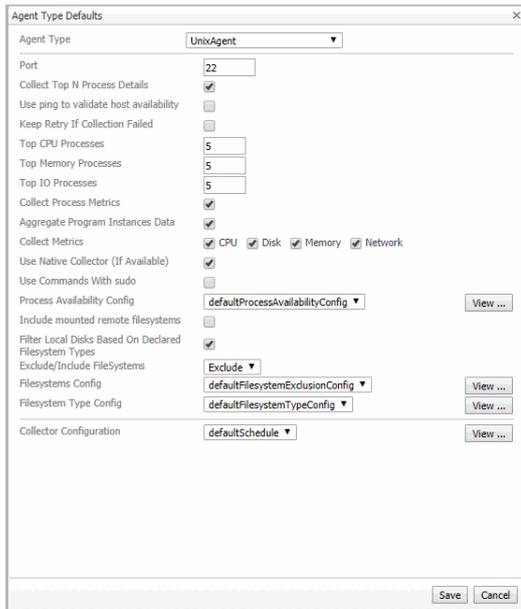
The dialog box refreshes to display the properties specific to the selected agent type. The following illustration shows the default properties for a WindowsAgent.



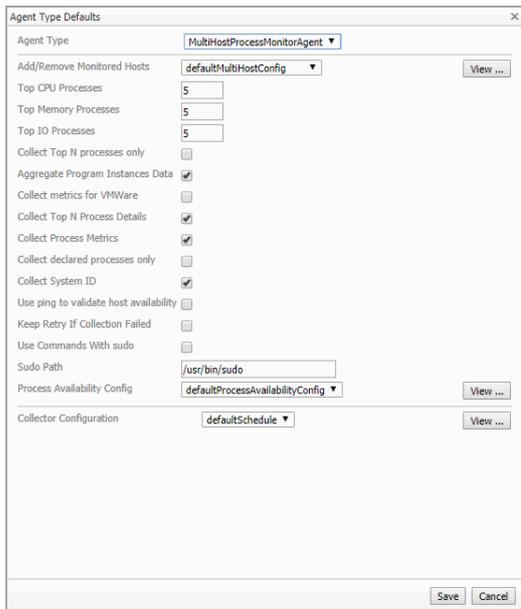
The following illustration shows the default properties for a UnixAgentPlus.



The following illustration shows the default properties for a UnixAgent.



The following illustration shows the default properties for a MultiHostProcessMonitorAgent.



For detailed information about Foglight for Infrastructure agents and their properties, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

i | NOTE: UnixAgentPlus only supports the Linux and Solaris platforms.

i | NOTE: MultiHostProcessMonitorAgent includes both the Windows and Linux® platforms.

- Update the agent's properties, by selecting other properties from the drop-down lists, by creating new shared property lists, or by editing the default existing ones.

Clicking the **View** button on the right side of a default property opens the Agent Property List dialog box, which lists the property rows. From here, you can click:

- Edit:** allows you to edit the currently selected property list, by adding and removing rows.

- **Copy:** allows you to edit the displayed properties and save the copy as a new shared property list. When saved, the copy becomes the selected item in the drop-down list.

i | IMPORTANT: When you click Save in the Agent Property List dialog box, nothing is persisted on the server. Changes (new copy or edits) are saved after [Step 4](#).

4 Click **Save** to save the changes made to the default properties.

i | IMPORTANT: Changes (new copy or edits) are now saved to the server, provided that the copy or edited version is the one selected in the drop-down to be used by the currently-selected agent. Changes that involve default properties that are not used by this agent are discarded.

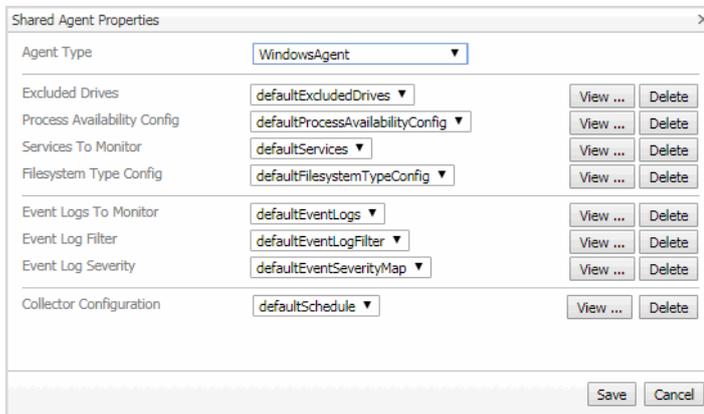
To edit the shared agent type properties:

1 Click **Edit Shared Agent Properties**.

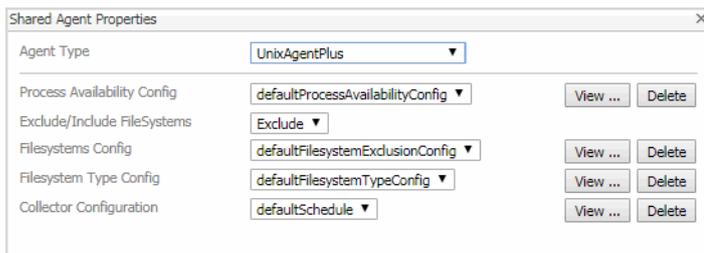
The Shared Properties dialog box appears.

2 From the Agent Type list, select whether you want to change the properties shared by Windows, UnixAgentPlus, UNIX, or MultiHostProcessMonitorAgent agents.

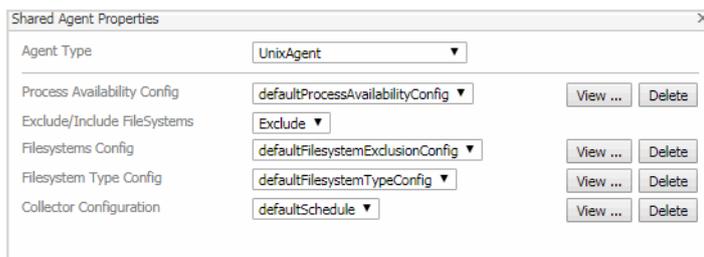
The dialog box refreshes to display the properties specific to the selected agent type. The following illustration shows the shared properties for Windows agents.



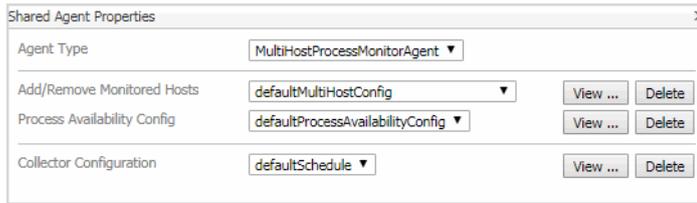
The following illustration shows the shared properties for UnixAgentPlus agents.



The following illustration shows the shared properties for Unix agents.



The following illustration shows the shared properties for MultiHostProcessMonitor agents.



For detailed information about Foglight for Infrastructure agents and their properties, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

- 3 Update the agent's properties, by selecting other properties from the drop-down lists, by creating new lists, or by editing the existing ones.

Clicking the **View** button on the right side of a shared property opens the Agent Property List dialog box, which lists the property rows. From here, you can click:

- **Edit:** allows you to edit the currently selected property list, by adding and removing rows.
- **Copy:** allows you to edit the displayed properties and save the copy as a new property list. When saved, the copy becomes the selected item in the drop-down list.

i | **IMPORTANT:** When you click Save in the Agent Property List dialog box, nothing is persisted on the server. Changes (new copy or edits) are saved after [Step 4](#).

Clicking the **Delete** button on the right side of a shared property deletes the property list currently selected in the drop-down. Default property lists cannot be deleted from the system.

- 4 Click **Save** to save the changes made to the shared properties.

i | **IMPORTANT:** Changes (new copy or edits) are now saved to the server, provided that the copy or edited version is the one selected in the drop-down to be used by the currently-selected agent. Changes that involve shared properties that are not used by this agent are discarded.

To edit the agent properties for a single agent:

- 1 Select the check box for that agent on the Agents table.

- 2 Click **Edit Agent**.

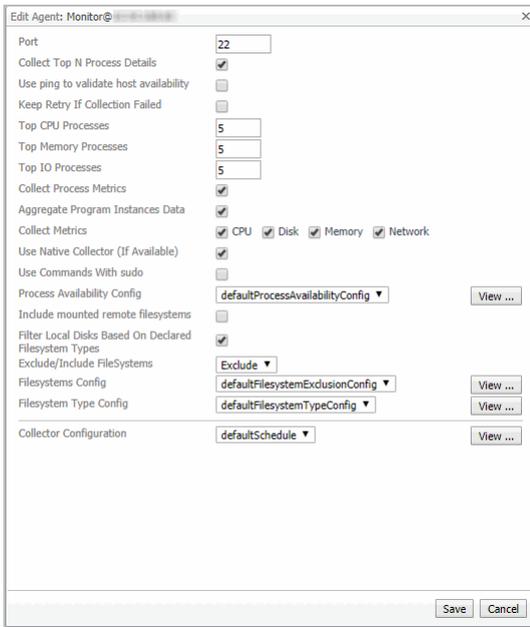
The Edit Agent dialog box appears.

- 3 Update the agent's properties, then click **Save**.

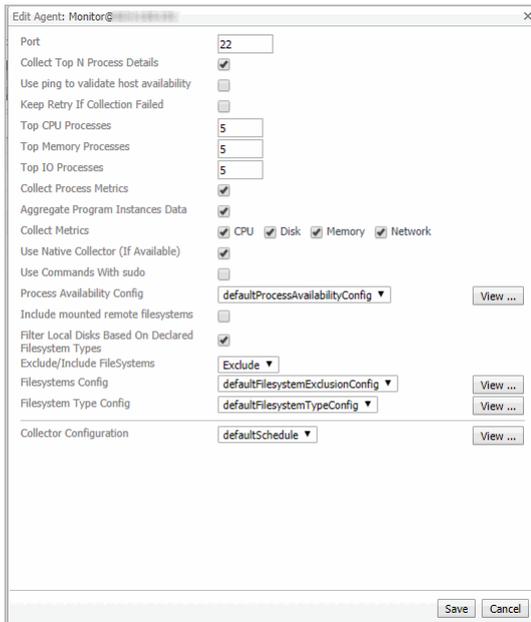
The list of properties may vary depending on the agent type. The following illustration shows the properties for a WindowsAgent.



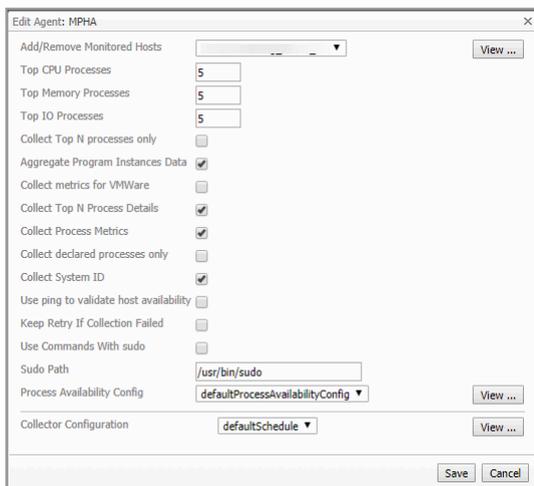
The following illustration shows the properties for a UnixAgentPlus.



The following illustration shows the properties for a UnixAgent.



The following illustration shows the properties for a MultiHostProcessMonitorAgent.



For detailed information about Foglight for Infrastructure agents and their properties, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

To edit the agent properties for multiple agents:

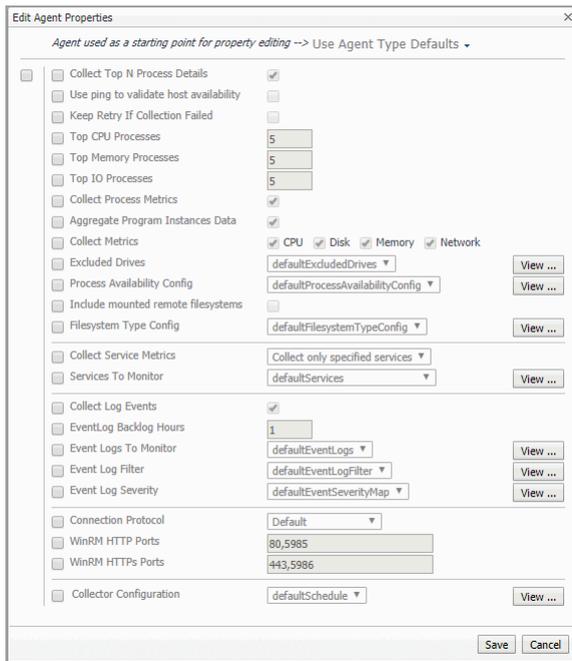
- 1 Select the check boxes for these agents on the Agents table.

i | NOTE: The selected agents must be monitored by the same Foglight for Infrastructure agent type.

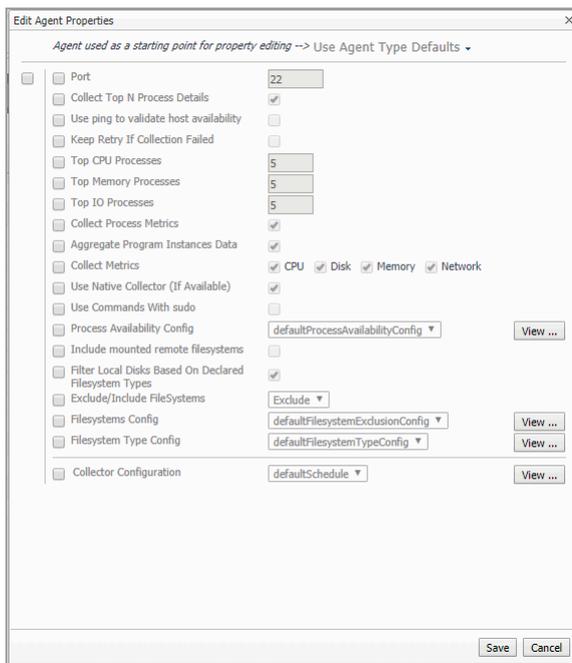
- 2 Click **Edit Agent**.

i | NOTE: If the agents selected are monitored by a mix of Windows® and UNIX® agents, a warning message is displayed.

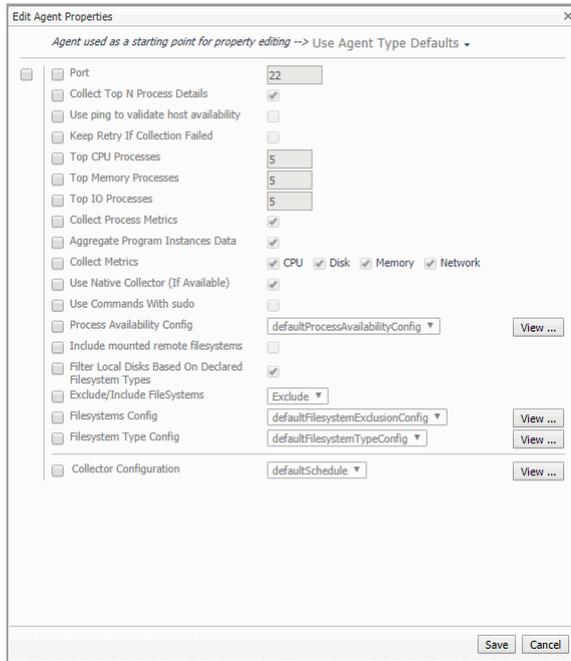
The Edit Agent Properties dialog box appears. The list of properties may vary depending on the agent type. The following illustration shows the properties for WindowsAgent.



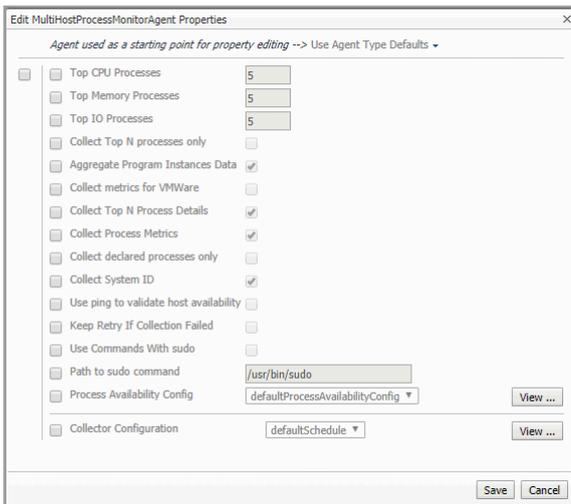
The following illustration shows the properties for UnixAgentPlus.



The following illustration shows the properties for UnixAgent.



The following illustration shows the properties for MultiHostProcessMonitorAgent.



For detailed information about Foglight for Infrastructure agents and their properties, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

- 3 Select the monitoring agent that should be used as a starting point for editing the properties for this agent type. From the drop-down list, select **Use Agent Type Defaults** or one of the available host names.

The Edit Agent Properties dialog box refreshes to show the properties for the selected agent.

- 4 Update the agent's properties, as needed. Select the check boxes for the fields that need to be edited, then select the options you want from the drop-down lists.

Clicking the **View** button on the right side of a property opens a pop-up which lists the property rows.

- 5 Click **Save** to save the changes made to the agent properties.

Rule Configuration

The Rule Configuration view of the Infrastructure Environment dashboard contains links to rules and alarms tasks that you can use to manage Infrastructure rules and alarms.

i | **NOTE:** This Rule Configuration view is only available in the Foglight Management Server 5.9.2 or later.

The Rule Configuration view consists of the following:

- [Rules view](#)
- [Enabling/Disabling rule\(s\)](#)
- [Adding a custom rule](#)
- [Removing custom rule\(s\)](#)

Rules view

By default, the following columns are displayed in the *Rules* view:

- **Enabled:** Indicates if the rule is enabled  or disabled . You can sort the list of rules by state, by clicking the Enabled column.
- **Rule:** Contains the rule name. Click the rule name to start the workflow for viewing and editing rule details.
- **Fatal** , **Critical** , and **Warning**  thresholds (multiple-severity rules only):
 - For expressions that include one registry variable, these columns contain the current value of that variable. Click the value to edit it.
 - For expressions that include multiple registry variables, the column contains an icon . Clicking that icon shows the list of referenced registry variables and their values. Click a value to edit it.
 - For expressions that do not include any registry variables, this column contains an icon . Clicking that icon navigates to the **Edit Rule** dashboard.
 - For rule states that do not have a conditional expression defined, this column is empty.
- **Alarms:** Contains the number of alarms (multiple-severity rules only) generated by the rule. Clicking that column shows a list of alarms indicating for each alarm its severity, when the alarm was generated, and the alarm message.
- **Applies to:** Shows the object name that is applied to this custom rule.
- **Description:** Contains the rule description.

Enabling/Disabling rule(s)

The *Rule Configuration* dashboard shows a list of existing rules and a set of rule management commands at the top of the list. Use the **Enable Rule** and **Disable Rule** buttons to activate or deactivate one or multiple rules at once.

To enable a rule:

- 1 Log in to the Foglight browser interface.
- 2 Ensure that the navigation panel is open.

To open the navigation panel, click the right-facing arrow  on the left.
- 3 On the navigation panel, under *Dashboards*, click **Infrastructure**.

The **Infrastructure** dashboard opens.
- 4 Click **Administration > Rule Configuration**.

The **Rule Configuration** dashboard opens.
- 5 On the *Rules* list, select one or more check boxes in the left-most column, and then click **Enable Rule**.

The *Enable Rules* dialog box opens.

- 6 In the *Enable Rules* dialog box, click **Yes**.

The *Rules* list refreshes with the rules' status updated automatically.

To disable a rule:

- 1 Log in to the Foglight browser interface.

- 2 Ensure that the navigation panel is open.

To open the navigation panel, click the right-facing arrow  on the left.

- 3 On the navigation panel, under *Dashboards*, click **Infrastructure**.

The **Infrastructure** dashboard opens.

- 4 Click **Administration > Rule Configuration**.

The **Rule Configuration** dashboard opens.

- 5 On the *Rules* list, select one or more check boxes in the left-most column, and then click **Disable Rule**.

The *Disable Rules* dialog box opens.

- 6 In the *Disable Rules* dialog box, click **Yes**.

The *Rules* list refreshes with the rules' status updated automatically.

Adding a custom rule

The *Rule Configuration* dashboard shows a list of existing rules and a set of rule management commands at the top of the list. Use the **Add Custom Rule** button to create a new rule as needed.

To customize a rule:

- 1 Log in to the Foglight browser interface.

- 2 Ensure that the navigation panel is open.

To open the navigation panel, click the right-facing arrow  on the left.

- 3 On the navigation panel, under *Dashboards*, click **Infrastructure**.

The **Infrastructure** dashboard opens.

- 4 Click **Administration > Rule Configuration**.

The **Rule Configuration** dashboard opens.

- 5 Click **Add Custom Rule** on the *Rules* table.

The *Create Custom Rule* dialog box opens.

- 6 In the *Create Custom Rule* dialog box, specify the following:

- a Alarm Type:

- a Type the name of custom rule in the *Name* field.

- b Select an *Object Type*, and then select a metric from the *Metric* drop-down list. The value of *Metric* varies from the *Object Type*.

- c Select either *Threshold* or *% Change*, and then specify the following values as needed.

- *Threshold*: Specify *Condition*, *Time Period*, *Severity*, and then specify whether or not fire actions if the specified data attempts are reached. The value of *Condition* cannot be negative.
- *% Change*: Specify *Condition*, *Time Period*, and *Severity Label*. The value of *Condition* cannot be negative.

- b (Optional) Scope: Choose the objects to which you want to apply this rule. If no objects are selected in this step, the custom rule will apply to all objects which type is the *Object Type* specified in [Step 6](#).
 - c (Optional) Notifications: Click **Add New**, then the *Edit Notification Config - Dialog* box appears. In this dialog box, type the *E-mail Address* and *Description* as needed, and then click **Add**.
- 7 Click **Save**.

The *Rules* table refreshes automatically to show the newly added rule.

Removing custom rule(s)

The *Rule Configuration* dashboard shows a list of existing rules and a set of rule management commands at the top of the list. Use the **Remove Custom Rule** button to delete existing custom rule(s) as needed.

To remove a custom rule:

- 1 Log in to the Foglight browser interface.
- 2 Ensure that the navigation panel is open.

To open the navigation panel, click the right-facing arrow  on the left.
- 3 On the navigation panel, under *Dashboards*, click **Infrastructure**.

The **Infrastructure** dashboard opens.
- 4 Click **Administration > Rule Configuration**.

The **Rule Configuration** dashboard opens.
- 5 Click **Remove Custom Rule** on the *Rules* table.

The *Remove* dialog box opens.
- 6 Click **Yes**.

The *Rules* table refreshes automatically and removes the selected rule.

Using Foglight for Infrastructure agents

Foglight for Infrastructure uses agents to collect information about the monitored hosts. Agents communicate with the Foglight Management Server using the Foglight Agent Manager.

For more information, see these topics:

- [Configuring Multiple Instances for Monitoring using Silent Installation](#)
- [Using the UnixAgentPlus](#)
- [Monitoring remote hosts](#)
- [Deploying Foglight for Infrastructure agents](#)
- [Creating agent instances](#)
- [Defining credentials](#)
- [Activating the agent](#)
- [Monitoring the infrastructure](#)
- [Adding a monitored host](#)
- [Adding multiple monitored hosts](#)

Configuring Multiple Instances for Monitoring using Silent Installation

Foglight for Infrastructure silent installation uses scripts and a CSV file to setup multiple Infrastructure agents on a given Foglight Agent Manager (FglAM). The CSV file defines the instances to be created. For more information, see [Contents of the Input CSV File](#) on page 31.

i | **NOTE:** Foglight for Infrastructure silent installation only supports the following agent types: WindowsAgent, UnixAgent, and UnxiAgentPlus.

Prerequisite

- A running Foglight Management Server
- At least one FglAM process running on the selected Foglight Management Server
- An Infrastructure cartridge installed

To add multiple agents using the silent installation through a command line interface:

- 1 Navigate to **Administration > Cartridges > Components for downloads**.
- 2 Download the Infrastructure_Command_Line_Installer file to a temporary folder.

The Infrastructure_Command_Line_Installer.zip file contains the following files:

- README.txt
- infrastructure_agents_silent_install.groovy

A groovy script file that runs the command line installation. This file should be extracted to the <FMS_HOME>/bin/ directory.

- infrastructure_agents_input_template.csv

A template file that serves as the basis for the input CSV file. This file can be extracted to any folder of your choice, provided that the instance file name points to the selected path.

- 3 Extract the Infrastructure_Command_Line_Installer.zip file to a temporary folder.
- 4 Extract infrastructure_agents_silent_install.groovy to the <FMS_HOME>/bin/ directory.
- 5 Go to the command line and run the command: <FMS_HOME>/bin/fglcmd -srv <fms_host_name> -port <fms_url_port> -usr <fms_user_name> -pwd <fms_user_password> -cmd script:run -f infrastructure_agents_silent_install.groovy fglam_name <fglam_name> instances_file_name <instances_file_name> lockbox_name <lockbox_name> lockbox_password <lockbox_password>

The descriptions of the flags and parameters in this command are as follows:

- <FMS_HOME>: The Foglight Management Server installation directory.
- <fms_host_name>: The host name where the Foglight Management Server is installed.
- <fms_url_port>: The Foglight Management Server port.
- <fms_user_name>: The user name used for connecting to the Foglight Management Server.
- <fms_user_password>: The password of the specified user.
- <fglam_name>: The name of the selected FglAM to add the new agents.
- <instances_file_name>: The full path and the name of the input CSV file.

i | **NOTE:** If the full path contains spaces, instances_file_name must be quoted.

- <lockbox_name>: Optional - Specifies the name of an existing lockbox that would be used.

- <lockbox_password>: Optional - The selected lockbox's password.

Contents of the Input CSV File

The CSV file contains data regarding a list of instances to monitor.

Table 1. Input CSV File

Name	Description
Agent Name (Optional)	The required Infrastructure agent name. If no value is entered in this field, a default value from the template will be imported as follows: <i>Monitor@ + hostName</i> . For example: for the host name <i>MyHost</i> , an agent will be created with the name <i>Monitor@MyHost</i>
Host Name	The target host name or IP address to be monitored.
Host Name Override (Optional)	The host name to be used to store this host's data in the Foglight data model.
OS Type	The operating system of the monitored host. The possible value includes: <i>Windows, Linux, Solaris, AIX, and HP-UX</i> .
Virtual/Guest (Optional)	Specifies whether the monitored host is a virtual machine. The possible values are either <i>True</i> or <i>False</i> . If no value is entered in this field, the default value <i>False</i> will be used. If this field is set to <i>True</i> , it will not collect metrics for CPU, Memory, Network, and Disk.
SSH Port (Optional)	SSH port the agent is connected to. If no value is entered in this field, it will use the default value 22.
OS Domain (Optional)	The domain name of your Windows environment (if configured in your environment).
OS User Name	The user name required for connecting to the operating system.
OS Password	The password for the user of the operating system specified above.

Contents of the Installation Command Line Output File

Table 2. Command Line Output File

Name	Description
Status	A CSV file that specifies the monitoring output of all given instances from the input CSV file: <i>original file name + output.csv</i> .
Agent Name	The name of the monitored Infrastructure agent.
Agent Type	The agent type of the monitored Infrastructure agent.
Target Host	The target host name or IP address monitored by Infrastructure agent.
Creation Status	Specifies whether the Infrastructure agent was created successfully or failed. The possible values are either <i>Successful</i> or <i>Failed</i> .
Credential Error Message	An error message in case the credential creation failed.
Creation Error Message	An error message in case the agent creation failed.

Using the PowerVM HMC agent

Foglight for PowerVM allows you to monitor IBM® PowerVM® environments. Foglight alerts you about infrastructure problems as soon as they develop, enabling you to resolve issues pro actively before end users are affected. Early intervention ensures consistent application performance at established service levels. Foglight for PowerVM monitors the health of your virtual system by tracking the levels of resource utilization such as CPU, network, and memory consumption of individual objects in your integrated environment.

To monitor a collection of PowerVM® servers, you need a running instance of the PowerVM HMC Agent. This agent is provided with Foglight for PowerVM.

For detailed information about monitoring a PowerVM infrastructure, see [Monitoring IBM PowerVM environments](#).

Using the UnixAgentPlus

The UnixAgentPlus is fully supported for use in monitoring Linux® and Oracle Solaris® systems.

While a UnixAgentPlus collects all the metrics collected by a UnixAgent, it also:

- collects LogicalDisk.percentNodesRemaining as an additional metric.
- returns metrics for SAN mounted disks.
- provides the ability to specify which declared processes are reported when the Collect declared processes only property is set to true.

i | **NOTE:** If you need to monitor AIX or HP-UX platforms, then UnixAgent must be used.

Monitoring remote hosts

Unlike the previous generation of system agents, Foglight for Infrastructure agents can be set up to monitor hosts remotely. A single Foglight Agent Manager can host many Foglight for Infrastructure agents, remotely monitoring different hosts.

Foglight for Infrastructure agents are best deployed on a Foglight Agent Manager that is closely located to the monitored host, for example on the same subnet. This is especially important for WindowsAgents because the communication interfaces they depend on are not well-suited to a high-latency environment, performing many remote operations to gather data. UnixAgentPlus and UnixAgents are also affected. A high latency between the Agent Manager and the monitored host can result in collections being skipped, and observed as gaps in data.

i | **NOTE:** Monitoring remote Linux® and UNIX® hosts is possible using the SSH (secure shell) protocol. The Agent Manager does not support Telnet as a remote monitoring protocol because of potential security issues. For more information, see the *Agent Manager Guide*.

Foglight for Infrastructure agents can still be used locally in situations where remote operation is not possible.

Preventing memory leaks in Windows Server® 2008 R2 and Windows 7

Querying the `Win32_Service` WMI class can result in memory leaks when the agent monitors a Windows® Server 2008 R2 or Windows 7 system. If you see that the amount of memory used by the `Wmiprvse.exe` process keeps increasing, and the system performance decreasing, this is likely caused by this known Microsoft® issue. To prevent memory leaks, install the Microsoft hotfix #981314 available from <http://support.microsoft.com/kb/981314>.

Deploying Foglight for Infrastructure agents

When you install Foglight for Infrastructure on the Foglight Management Server, the installation process makes its agent package (*HostAgents.gar*) available for deployment.

Before creating Foglight for Infrastructure agents for the hosts that you want to monitor, you must have the Foglight for Infrastructure cartridge file installed on the Foglight Management Server, and the agent package deployed to one or more Foglight Agent Managers used to manage the remote agents.

Creating agent instances

After a successful agent deployment, you can proceed to create agent instances using the agent types included in the newly-deployed infrastructure agent package: *UnixAgentPlus*, *UnixAgent*, *WindowsAgent*, or *MultiHostProcessMonitorAgent*.

i | **NOTE:** *UnixAgentPlus* includes the Linux and Oracle Solaris agent types.

i | **NOTE:** *UnixAgent* includes the AIX, HP-UX, Linux, and Oracle Solaris agent types.

i | **NOTE:** *MultiHostProcessMonitorAgent* includes both the Windows and Linux platforms.

The recommended method for creating Foglight for Infrastructure agents for the hosts that you want to monitor is by using the **Add Monitored Host** wizard or the **Add Monitored Host - List** wizard. For detailed information about these wizards, see sections [Adding a monitored host](#) and [Adding multiple monitored hosts](#).

For generic information about creating agent instances and editing properties using the Agent Status dashboard, see “Create agent instances on monitored hosts” in the *Administration and Configuration Help*. Edit the agent instance properties, as necessary for each monitored host. For generic information about editing agent properties, see “Edit instance-specific agent properties” in the *Administration and Configuration Help*.

For detailed information about Foglight for Infrastructure agents, see [About the WindowsAgent](#), [About the UnixAgentPlus](#), [About the UnixAgent](#), and [About the MultiHostProcessMonitorAgent](#).

Defining credentials

Credentials are security data that provide the Foglight for Infrastructure agents with the permission to monitor system resources, such as a host or a range of hosts.

i | **NOTE:** Credentials are needed for any remote agent. They are not needed when the Foglight Agent Manager is physically located on the host being monitored.

If you create agent instances using the **Add Monitored Host** wizard or the **Add Monitored Host - List** wizard, the credentials setup is part of the wizard’s flow. For details, see sections [Adding a monitored host](#) and [Adding multiple monitored hosts](#).

If you create agent instances using the Agent Status dashboard, you need to set up the credentials required for your Foglight for Infrastructure Agent in order to start collecting data. The credentials provided for monitoring should be those of a local user account on the remote system. The local user account should also have a local home directory. If there is a failure in the remote storage or remote user services, a local account is accessible to the agent, allowing it to detect and report the problem. If a network account, or an account with a remotely mounted home directory is used, this may prevent the agent from accessing the remote system when a problem occurs. For detailed information about setting up credentials in Foglight, see the “Managing Credentials” topic in the *Foglight Administration and Configuration Help*.

Activating the agent

After you create the agent instance and define the proper credentials, you can enable the agent to start the data collection:

- 1 On the navigation panel, navigate to Dashboards > Administration > Agents > Agent Status.
- 2 Select the agent that you want to activate, and click **Activate**.

i | **TIP:** The activate operation fails if the credentials are not properly defined. Click **Get Log** to see the list of errors and warnings.

i | **NOTE:** You can also create/activate an agent when adding a monitored host. For more information, see [Adding a monitored host](#).

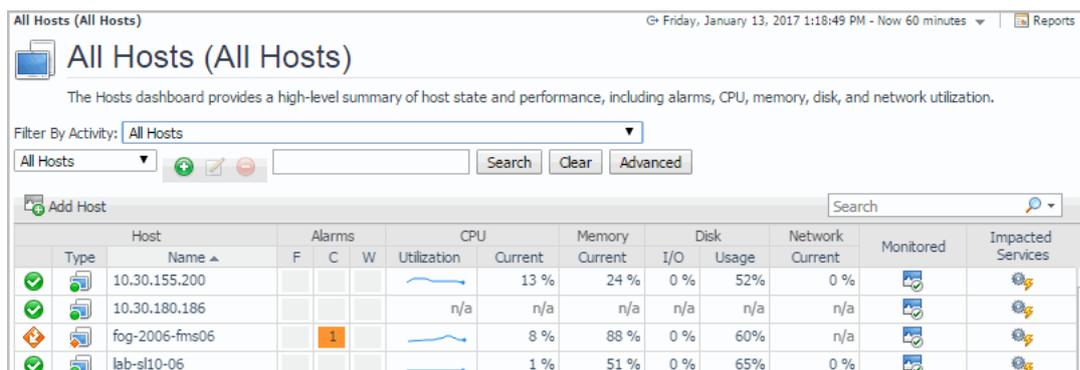
Monitoring the infrastructure

After activating the agent, you can start monitoring your infrastructure environment. If your agent is configured correctly, the Hosts dashboard should now display the data being collected.

To access the Hosts dashboard:

- On the navigation panel, under Dashboards, click **Infrastructure > Hosts > Hosts**.

Figure 9. Hosts dashboard



i | **NOTE:** The *Add Host* icon and the *Monitored* column are available only if Foglight for Infrastructure is installed.

The *Monitored* column shows whether the host is already configured for monitoring , or is not monitored yet (in which case, you can click the icon to add a monitored host).

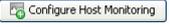
Once sufficient data has been collected, the *CPU (Utilization, Current)*, *Memory Current*, *Disk (I/O, Usage)*, and *Network Current* columns and the *Utilization* sparklines display the resource utilization for each monitored host. Hover over each of these columns to see the resource utilization graph for a host.

Adding a monitored host

The Add Monitored Host wizard allows you to add different types of hosts to the list of components being monitored in your infrastructure environment.

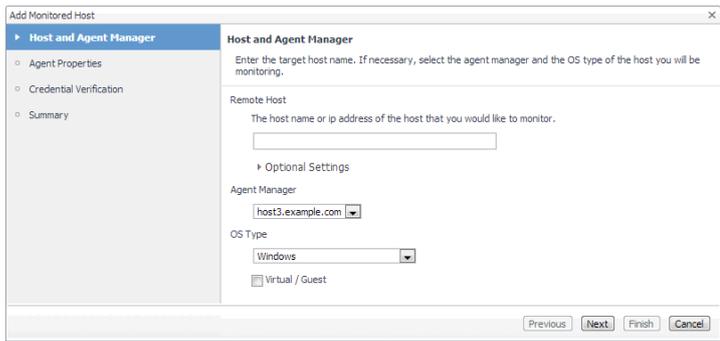
You can launch this wizard from several places in the Foglight browser interface:

- From the Infrastructure Environment dashboard, by clicking **Single Host** in the upper-right corner.

- From the Infrastructure Environment dashboard > **Monitoring tab > Resource Utilizations** view. If a host is not being monitored, the Resource Utilizations view displays a warning message. You can add the selected host to the list of monitored hosts by clicking **Configure Host Monitoring** .
- From the Hosts dashboard, by clicking **Add Host** , or (for a known but unmonitored host) by clicking the icon in the Monitored column.

To add a host using the Add Monitored Hosts wizard:

- 1 In the **Add Monitored Host** wizard, on the **Host and Agent Manager** page, specify the host that you want to monitor.



- **Remote Host:** Type the host name or the IP address of the host you want to monitor. If you clicked the Monitored icon for an unmonitored host, this box already contains the host name.
 - **NOTE:** This value must resolve to a real host name or IP address in order to pass the credential validation test in [Step 3](#).
- **Optional Settings:** Expand this area and type the host name that you want to override the previously specified host. This step is optional.
- **Agent Manager:** Select the agent manager that you want to use to collect the information about the remote host. The list includes only the Foglight Agent Managers to which the Foglight for Infrastructure agent package is deployed.
- **OS Type:** Select the OS type of the remote host that you want to monitor: **Windows**, ***nix (HP-UX, AIX)**, **Linux**, or **Solaris**.
 - **NOTE:** The wizard uses the specified OS type to determine what agent type to create, WindowsAgent or UnixAgent. This box appears disabled if the wizard is invoked by selecting a host when its OS type is known.
- **Virtual / Guest:** Ensure this check box is disabled.
 - **NOTE:** Ability to monitor virtual hosts is currently implemented only in Foglight™ for Virtualization, Enterprise Edition. For information about how to configure a virtual agent, refer to the Foglight™ for Virtualization, Enterprise Edition documentation.

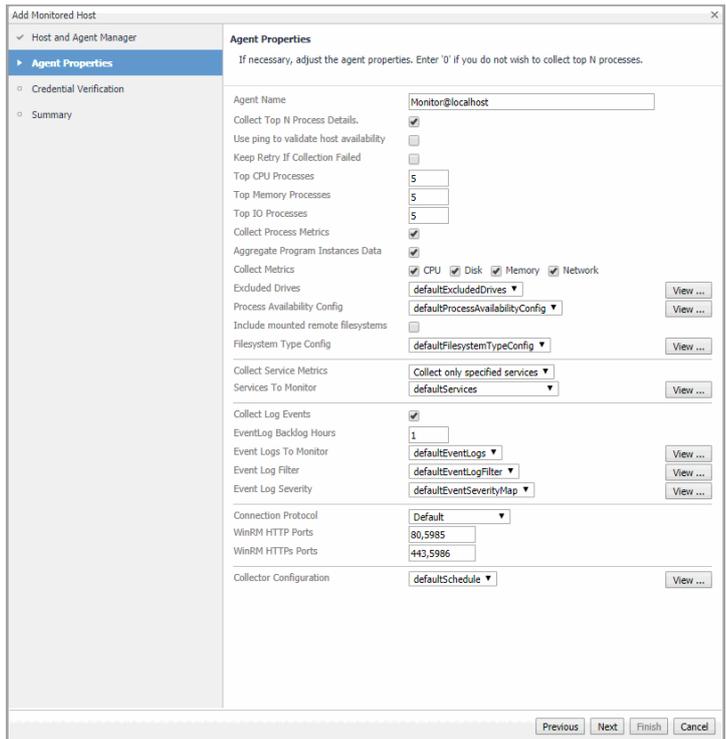
- 2 Click **Next**.

The host name is validated.

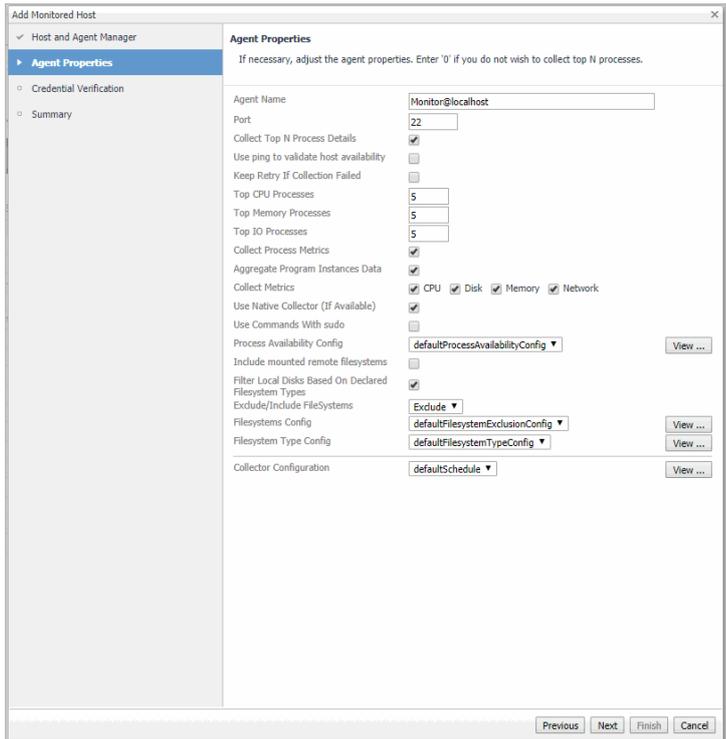
If the host name does not resolve to a real host name, the message `Unable to locate host using the provided host name` appears in the **Validation Result** dialog box. Click **Cancel** and return to [Step 1](#) to define a valid host name or IP address.

If the host name passes the validation test, the **Add Monitored Host** wizard refreshes, showing the **Agent Properties** page. The layout of the page is different for WindowsAgents and UnixAgents.

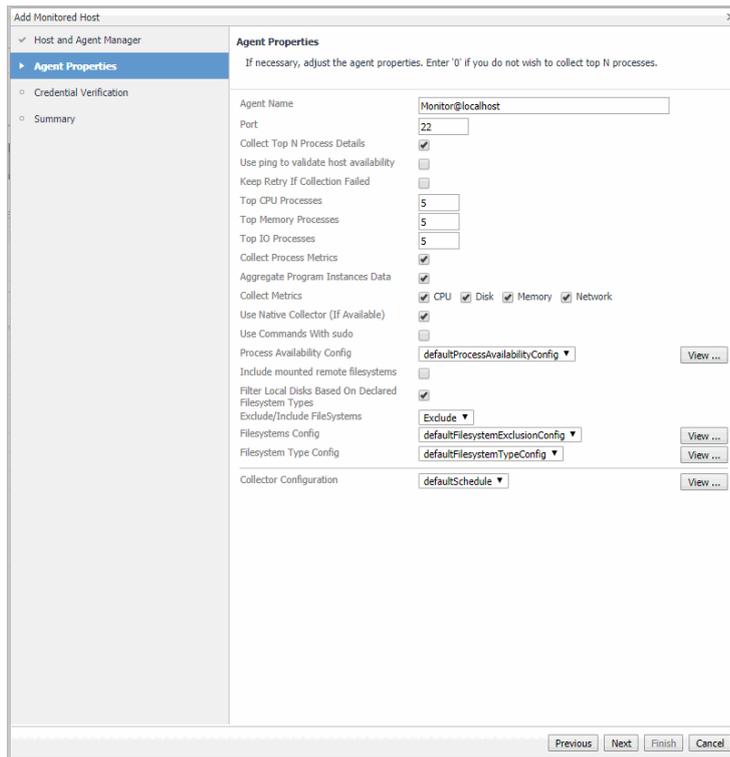
- WindowsAgent



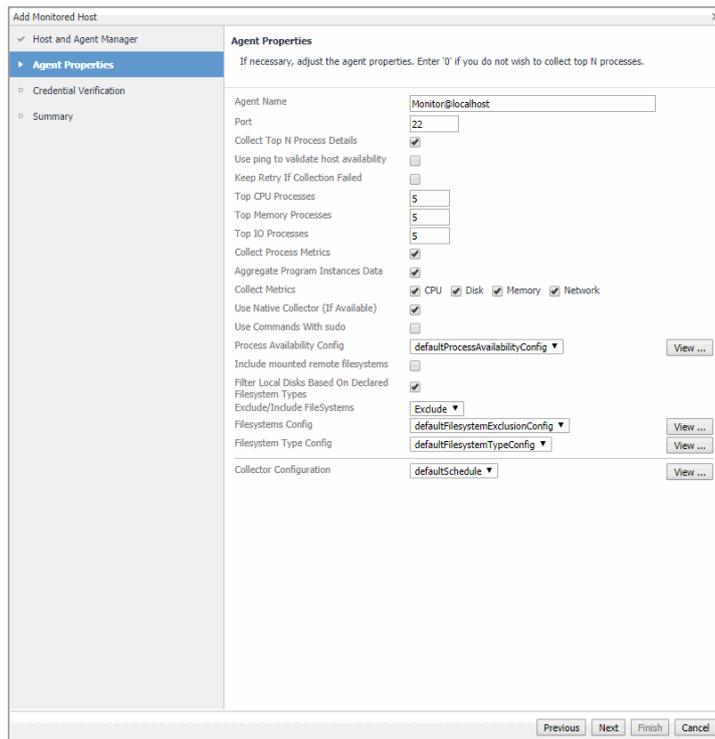
- *nix (HP-UX, AIX) Agent



- LinuxAgent



- Solaris Agent



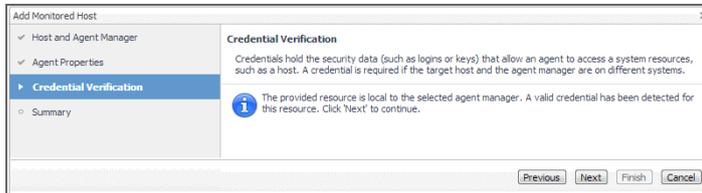
i | **NOTE:** If you selected the Virtual/Guest check box in [Step 1](#), the Collect Metrics check boxes are disabled, preventing the agent from collecting these type of metrics.

- 3 On the **Agent Properties** page, specify the properties of the monitoring agents. For complete information, see [Creating agent instances](#). Click **Next**.

The wizard checks whether the Agent Manager has any credentials configured for this host, and displays the findings on the Credential Verification page. Matching credentials are those with a resource mapping that matches the host to be monitored.

NOTE: Credentials are security data that provide the Infrastructure agent with the permission to monitor system resources, such as a host or a range of hosts.

If the selected host name matches the host on which the Agent Manager is running, that is indicated on the **Credential Verification** page.



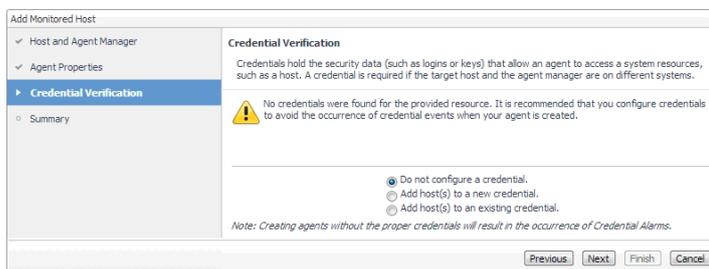
If any matching credentials are found, they are displayed in table format. The result of the validation is displayed in the **Test** column.

- If the test fails, click the **FAILED** link to open a dialog box that displays additional information about the issues encountered during the validation, and helps you troubleshoot the root of the problem.
- If the test is successful, the **SUCCESS** link appears in the **Test** column.



TIP: When several credentials are available for a resource, the list of credentials is ordered with the most applicable credentials at the top. The ordering is based on a relative ordering value assigned to each credential.

If no matching credentials are found, that is indicated on the **Credential Verification** page.



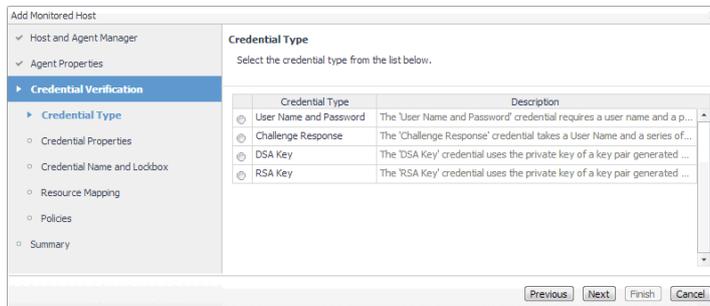
4 On the **Credentials Verification** page, select one of the following options:

- **Do not configure a credential:** Select this option if you want to configure the credential for this resource at a later time. Local credentials for Windows® and UNIX® are set up by default when Foglight for Infrastructure is installed. Click **Next** and continue with [Step 7](#).
- **Add host(s) to a new credential:** Select this option if you want to add the host to a new credential. This option is suitable if none of the existing credentials have the connection details needed to access the new host. Click **Next** and continue with [Step 5](#).

- **Add host(s) to an existing credential:** Select this option if you want to add the host to an existing credential. This option is suitable if an existing credential has the security data needed to access the new host, but you need to edit its resource mappings to include this host. Click **Next** and continue with [Step 6](#).
- If the selected host name matches the host on which the Agent Manager is running, click **Next** and continue with [Step 7](#).

5 Create a new credential.

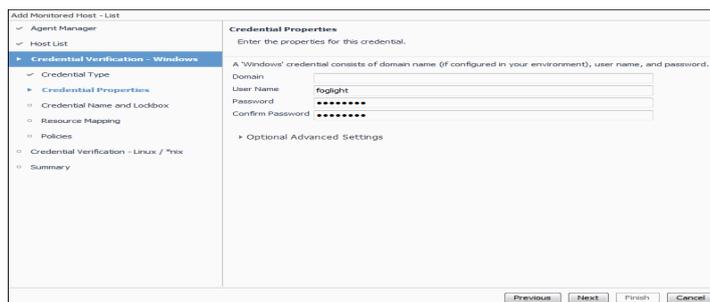
- a On the **Credential Type** page that appears, select the credential type from the available options.



TIP: The list of credential types varies depending on the agent type.

- b Click **Next**.

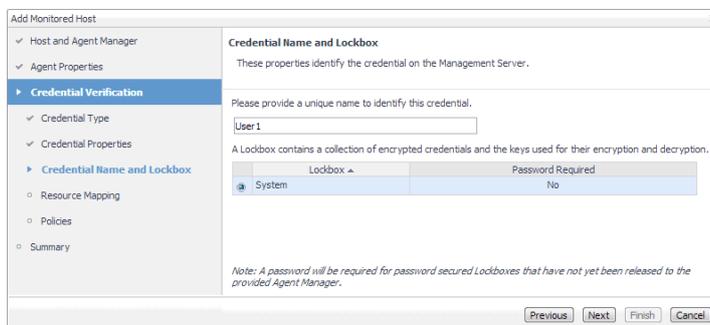
The **Credential Properties** page appears.



The level of required information depends on the selected credential type. For example, the **User Name and Password** type needs a user name and a password, while the **Challenge Response** type needs a user name along with a question/response pair.

- c On the **Credential Properties** page, type the required properties, and click **Next**.

The **Credential Name and Lockbox** page appears.



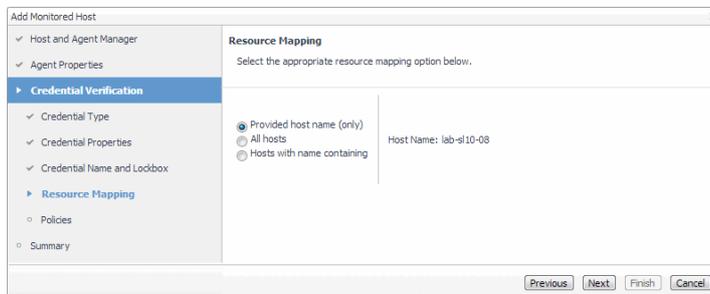
- d On the **Credential Name and Lockbox** page, provide a name to identify the credential, and select a lockbox in which you want to keep the credential. A lockbox can be used to group credentials for

access and/or security. In smaller Foglight installations, using the default **System** lockbox should be sufficient.

- NOTE:** If a lockbox is password protected and is not released to the target Foglight Agent Manager, you can provide the lockbox password on the last page of the wizard.

Click **Next**.

The **Resource Mapping** page appears.



- e On the **Resource Mapping** page, indicate which hosts you want to associate with this credential. You can either select the host that you are about to start monitoring, all monitored hosts, or hosts whose name contains a specific text string.

Click **Next**.

The **Policies** page appears.



- f Optional. On the **Policies** page, define one or more policies for this credential. A policy defines the number of times a credential can be used, the number of allowed authentication failures, the time range during which the credential is valid, or the length of time the credential data can be cached on the client. For example, you can specify the number of times the credential can be used, or the time period during which it can be used. For complete information about the available credential policies, see the *Administration and Configuration Help*.

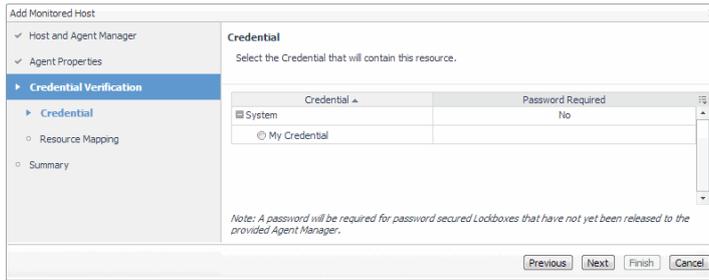
Click **Next**.

The **Summary** page appears.

- g Continue with [Step 7](#).

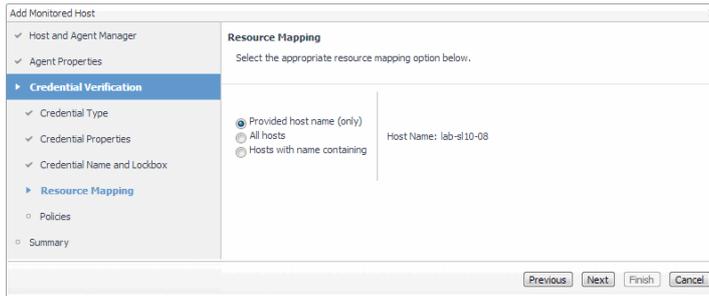
6 Use an existing credential.

- a On the **Credential** page that appears, select an existing credential to contain this host.



b Click **Next**.

The **Resource Mapping** page appears.



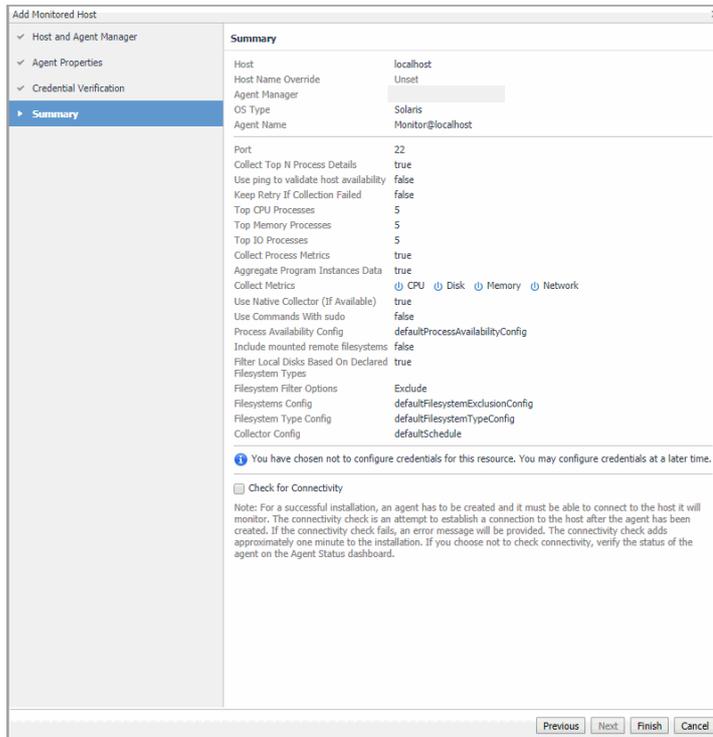
c On the **Resource Mapping** page, indicate which hosts you want to associate with this credential. You can either select the host that you are about to start monitoring, all monitored hosts, or hosts whose name contains a specific text string.

Click **Next**.

The **Summary** page appears.

d Continue with **Step 7**.

7 On the **Summary** page that appears, review the information provided about the host and the monitoring agent.



NOTE: A password field appears on this page in the case where you are using a password-secured lockbox that has not been released to the Agent Manager yet.

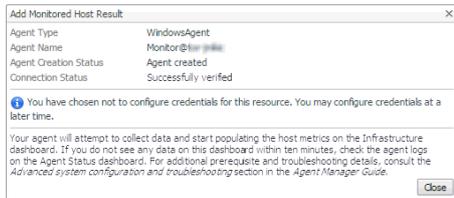
8 To perform a connectivity check, select **Check for Connectivity**.

For a successful installation, an agent has to be created and it must be able to connect to the host it will monitor. The connectivity check is an attempt to establish a connection to the host after the agent has been created. If the connectivity check fails, an error message is provided. The connectivity check adds approximately one minute to the installation. If you choose not to check connectivity, you can verify the agent status on the Agent Status dashboard.

9 Click **Finish**.

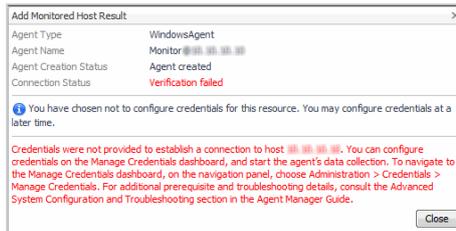
The new host is added to the Hosts dashboard after a short delay. The monitoring agent is created.

- If the operation is successful, the **Add Monitored Host** dialog box appears.



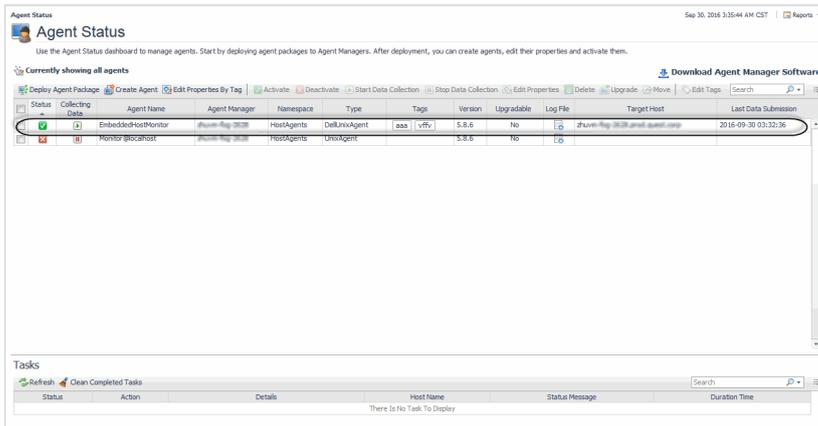
Review the information and close the dialog box.

- If you performed a connectivity check and verification fails, an error message is displayed to help you better understand and resolve the problem.



Review the error message. Close the dialog box, and perform the necessary steps, as instructed.

The agent instances created to monitor the new host appear on the Agent Status dashboard.



NOTE: The Agent Status dashboard does not refresh with the latest state unless an action is performed on the page, or the **Refresh** button on the page is clicked. Agents created using the **Add Monitored Host** wizard appear on the Agent Status dashboard only after clicking the Refresh button.

Adding multiple monitored hosts

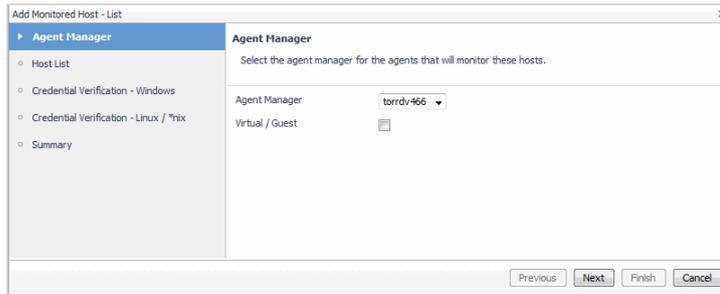
The **Add Monitored Host - List** wizard allows you to start monitoring multiple hosts in your infrastructure environment.

You can launch this wizard from the Infrastructure Environment dashboard, by clicking **Multiple Hosts** in the upper-right corner.

To add multiple monitored hosts:

- 1 On the Infrastructure Environment dashboard, in the upper-right corner, click **Multiple Hosts**.

The **Add Monitored Host - List** wizard appears with the **Agent Manager** page open.

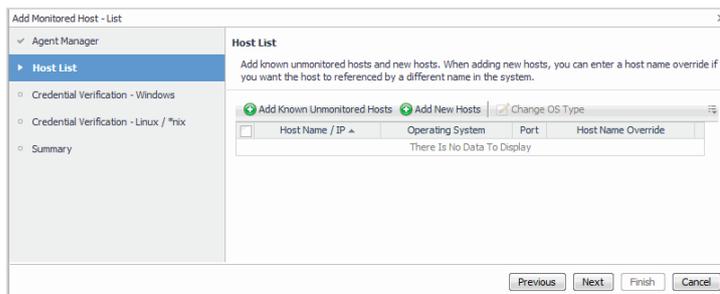


- **Agent Manager:** Select the agent manager that you want to use to collect the information about the hosts that you want to monitor. The list shows the Foglight Agent Managers to which the Foglight for Infrastructure agent package is deployed.
- **Virtual / Guest:** Ensure this check box is disabled.

NOTE: Ability to monitor virtual hosts is currently implemented only in Foglight™ for Virtualization, Enterprise Edition. For information about how to configure a virtual agent, refer to the Foglight™ for Virtualization, Enterprise Edition documentation.

Click **Next**.

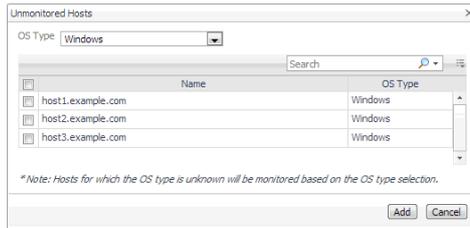
The **Add Monitored Host - List** wizard refreshes, showing the **Host List** page.



- 2 Specify the hosts that you want to monitor. You have two options, that you can combine, depending on which hosts you want to start monitoring:

- **Start monitoring existing unmonitored hosts:** Choose this option if you want to start monitoring one or more hosts that already exist in your environment, but are not currently monitored. On the **Host List** page, click **Add Known Unmonitored Hosts**.

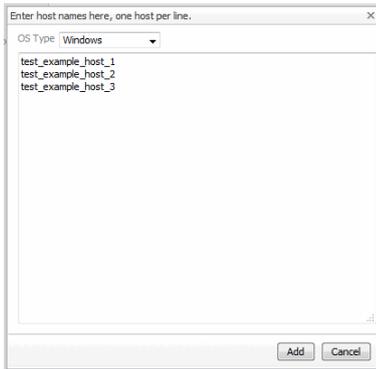
The **Unmonitored Hosts** dialog box appears.



In the **Unmonitored Hosts** dialog box, select one or more hosts that you want to start monitoring, and click **Add**.

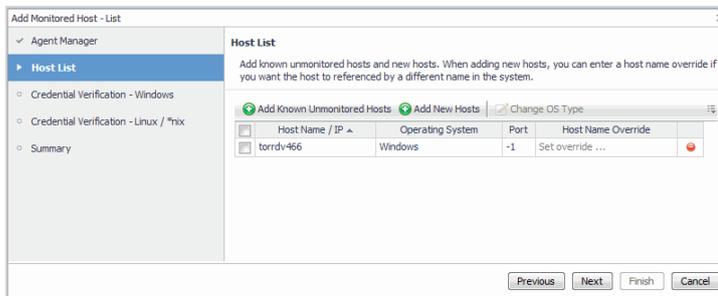
- **Start monitoring new hosts.** Choose this option if you want to start monitoring new hosts, that are not currently showing as unmonitored hosts in your environment. Click **Add New Hosts**.

The **Enter host names here, one host per line** dialog box appears.



In the **Enter host names here, one host per line** dialog box, type a list of one or more host addresses, one host per line. To specify an override name for a host, type a comma immediately after the host name, followed by the override name. Click **Add**.

The **Host List** page refreshes, showing the list of hosts that you selected for monitoring.

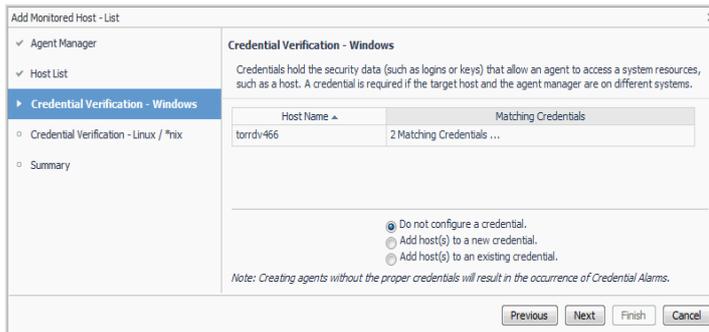


- 3 Optional. Specify host name overrides for one or more hosts. To specify a name override for a host, in the row containing the host, in the **Host Name Override** column, click **Set override**. In the dwell that appears, type the name and click **Save**.

- 4 Click **Next**.

The wizard checks whether the selected Agent Manager has any credentials configured for the selected Windows® hosts, and displays the findings on the **Credential Verification - Windows** page that appears. A matching credential has a resource mapping that resolves to a host name.

i | **NOTE:** Credentials are security data that provide the Foglight for Infrastructure agents with the permission to monitor system resources, such as a host or a range of hosts.



5 On the **Credential Verification - Windows** page, select one of the following options:

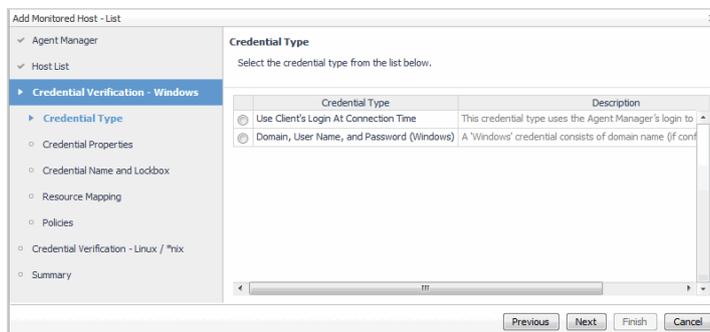
- **Do not configure a credential:** Select this option if you want to configure the credential for these hosts at a later time. Local credentials for Windows® and UNIX® are set up by default when Foglight for Infrastructure is installed. Click **Next**.

The wizard checks whether the selected Agent Manager has any credentials configured for the selected Unix hosts. Continue with [Step 8](#).

- **Add host(s) to a new credential:** Select this option if you want to add the hosts to a new credential. This option is suitable if none of the existing credentials have the connection details needed to access the new host. Click **Next** and continue with [Step 6](#).
- **Add host(s) to an existing credential:** Select this option if you want to add the host to an existing credential. This option is suitable if an existing credential has the security data needed to access the new host, but you need to edit its resource mappings to include this host. Click **Next** and continue with [Step 7](#).

6 Create a new credential.

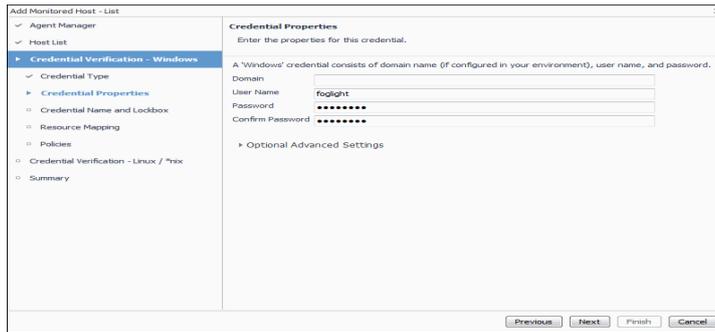
- On the **Credential Type** page that appears, select the credential type.



i | TIP: The list of credential types varies depending on the agent type.

- Click **Next**.

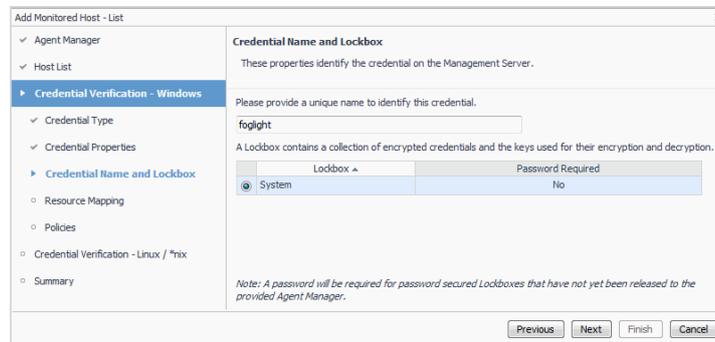
The **Credential Properties** page appears.



The level of required information depends on the selected credential type. For example, the **User Name and Password** type needs a user name and a password, while the **Challenge Response** type needs a user name along with a question/response pair.

- c On the **Credential Properties** page, type the required properties, and click **Next**.

The **Credential Name and Lockbox** page appears.

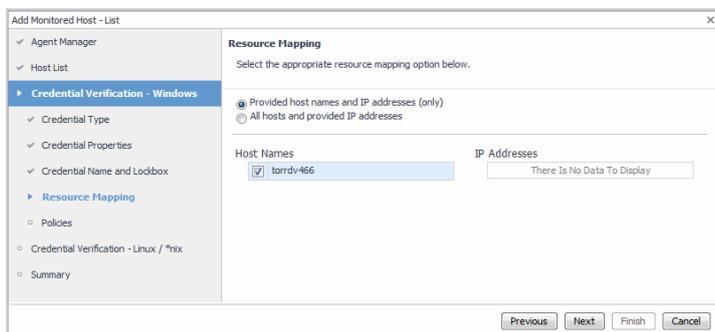


- d On the **Credential Name and Lockbox** page, provide a name to identify the credential, and select a lockbox in which you want to keep the credential. A lockbox can be used to group credentials for access and/or security. In smaller Foglight installations, using the default **System** lockbox should be sufficient.

NOTE: If a lockbox is password protected and is not released to the target Foglight Agent Manager, you can provide the lockbox password on the last page of the wizard.

Click **Next**.

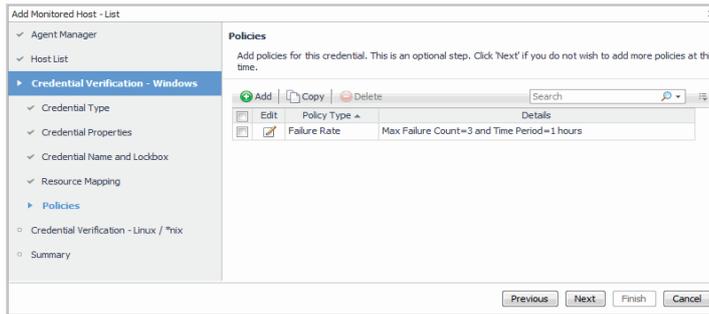
The **Resource Mapping** page appears.



- e On the **Resource Mapping** page, indicate which hosts you want to associate with this credential. You can either select all hosts in the list, or all monitored hosts.

Click **Next**.

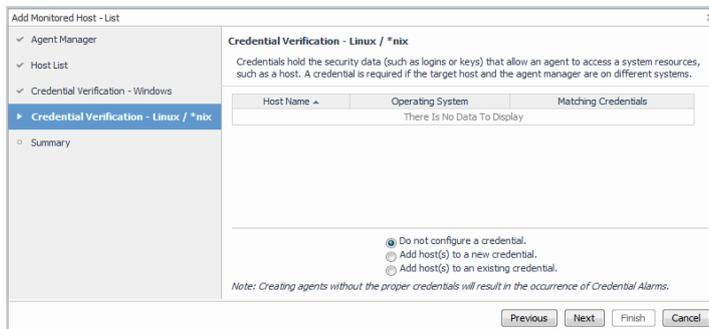
The **Policies** page appears.



- f Optional. On the **Policies** page, define one or more policies for this credential. A policy defines the number of times a credential can be used, the number of allowed authentication failures, the time range during which the credential is valid, or the length of time the credential data can be cached on the client. For example, you can specify the number of times the credential can be used, or the time period during which it can be used. For complete information about the available credential policies, see the *Administration and Configuration Help*.

Click **Next**.

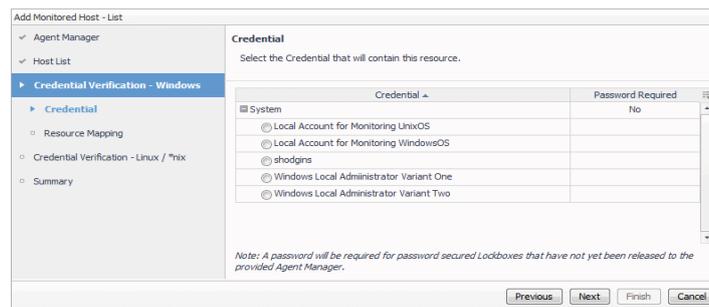
The **Credential Verification - Linux / *nix** page appears.



- g Continue with [Step 8](#).

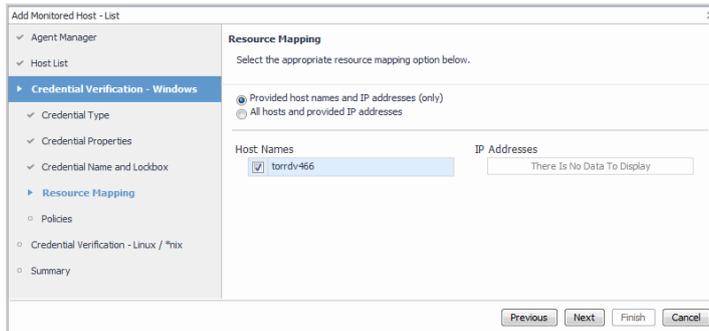
7 Use an existing credential.

- a On the **Credential** page that appears, select an existing credential to contain this host.



- b Click **Next**.

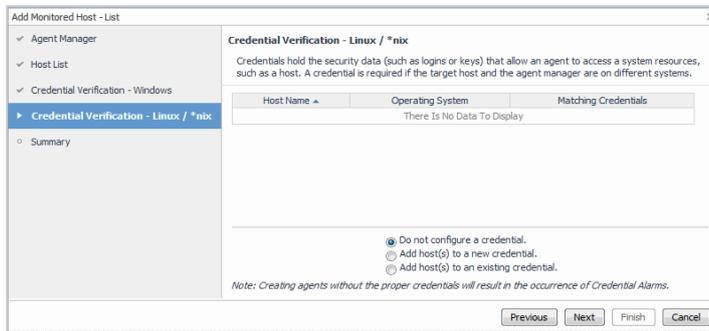
The **Resource Mapping** page appears.



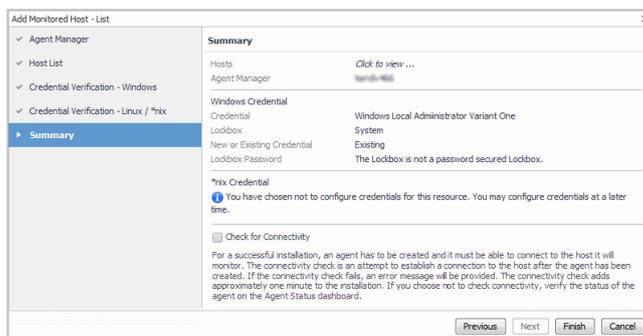
- c On the **Resource Mapping** page, indicate which hosts you want to associate with this credential. You can either select the host that you are about to start monitoring, all monitored hosts, or hosts whose name contains a specific text string.

Click **Next**.

The **Credential Verification - Linux / *nix** page appears.



- 8 On the **Credential Verification - Linux / *nix** page that appears, start specifying the credentials for your UNIX® hosts. For instructions, refer to [Step 5](#).
- 9 On the **Summary** page, review the information provided about the hosts you want to monitor and their credentials.



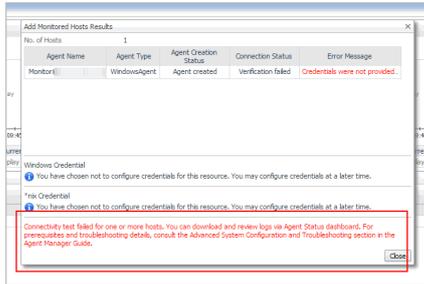
- 10 To perform a connectivity check, select **Check for Connectivity**.

For a successful installation, an agent has to be created and it must be able to connect to the host it will monitor. The connectivity check is an attempt to establish a connection to the host after the agent has been created. If the connectivity check fails, an error message is provided. The connectivity check adds approximately one minute to the installation. If you choose not to check connectivity, you can verify the agent status on the Agent Status dashboard.

- 11 Click **Finish**.

The new hosts are added to the Hosts dashboard after a short delay. The monitoring agents are created.

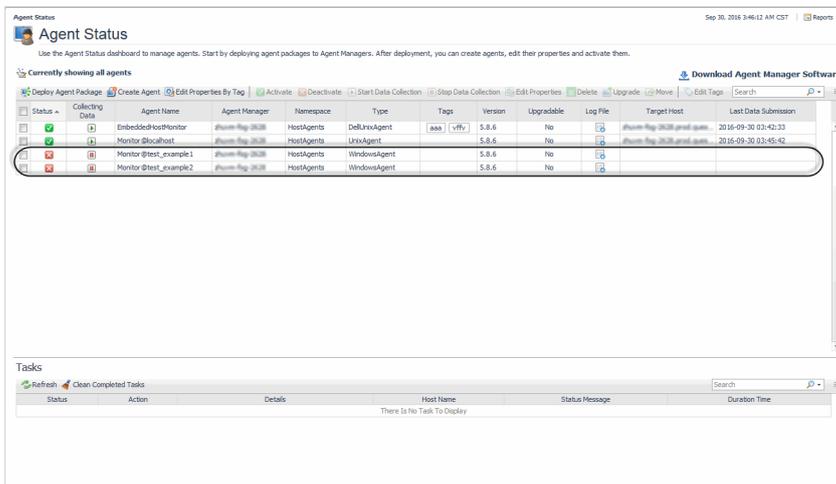
If you performed a connectivity check and verification fails for one or more agents, an error message is displayed to help you better understand and resolve the problem.



Review the information in the **Add Monitored Host** dialog box. If applicable, review the error message and perform the necessary steps, as instructed.

12 Close the **Add Monitored Host** dialog box.

The agent instances created to monitor the new host appear on the Agent Status dashboard.



NOTE: The Agent Status dashboard does not refresh with the latest state unless an action is performed on the page, or the **Refresh** button on the page is clicked. Agents created using the **Add Monitored Host** wizard appear on the Agent Status dashboard only after clicking the **Refresh** button.

About the WindowsAgent

The *WindowsAgent* monitors Windows® systems and collects the following information:

- CPU utilization and performance
- Memory and network performance
- Physical disk devices on a system
- Paging files (such as available paging space)
- Processes running or waiting to run on a monitored host
- Services configured on a monitored host
- Windows event logs
- One or more network interface cards for network traffic data (such as the number of packets, throughput, errors, utilization, and collisions)

i | **NOTE:** Services and event logs are collected for use by other cartridges, and are not directly visible in the Infrastructure dashboard.

There are views, rules, and data associated with this agent. For more information, see [Reference](#).

Windows Management Instrumentation (WMI) and Windows Remote Management (WinRM) are two different mechanisms that monitoring agents can use to establish remote connections. In most scenarios, only one of these mechanisms needs to be configured. Due to limitations in scalability for WMI that cannot be overcome, the preferred mechanism for monitoring Windows resources is via WinRM. For platform-specific information for configuring the Foglight Agent Manager on Windows, when using WMI or WinRM for remote monitoring access, see sections “Configuring Windows Management Instrumentation (WMI)” and “Configuring Windows Remote Management (WinRM)” in the *Foglight Agent Manager Guide*.

i | **NOTE:** When the Infrastructure WindowsAgent uses a WMI mechanism for remote connection, and non-administrative users for credentials, the following data cannot be collected with this type of credential:

- Services data collection (*Win32_Service* class)
- Event log data collection (*Win32_NTLogEvent* class)
- Cluster data collection (*MSCluster_** classes)

For more details, see these topics:

- [Supported platforms](#)
- [Agent properties](#)

Supported platforms

For a list of platforms supported for the WindowsAgent, see the “System Requirements” section in the *Foglight for Infrastructure Release Notes*.

Agent properties

When an agent connects to the Foglight Management Server, it is provided with a set of properties that it uses to configure its correct running state. For more information about working with agent properties, see [Creating agent instances](#).

The WindowsAgent is shipped with default properties that can be modified to suit your system requirements. The properties specific to the WindowsAgent are illustrated in the following screenshot.

Figure 10. WindowsAgent properties

You can configure the following settings for this agent:

- **Properties:**
 - **Host:** host name or IP address.
 - **Host name override:** host name to be used to store this host's data in the Foglight data model.
 - **Top CPU Processes:** number of top CPU processes to be monitored. Default value = 5.
 - **Top Memory Processes:** number of top memory processes to be monitored. Default value = 5.
 - **Top IO Processes:** number of top IO processes to be monitored. Default value = 5.
 - **Keep Retry If Collection Failed:** Default value = *False*. When set to *True*, the agent keeps trying to collect data when failed to connect to the target host during startup.
 - **Collect Top N processes only:** Default value = *False*. When set to *True*, the agent collects data for Top CPU Processes, Top Memory Processes, and Top IO Processes only.
 - **Collect process metrics:** Default value = *True*. When set to *True*, the agent collects process metrics.
 - **Collect CPU metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system's CPUs.
 - **Collect disk metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system's disks.
 - **Include mounted remote filesystems:** Default value = *False*. When set to *True*, the agent collects metrics about remotely mounted disks.
 - **Collect memory metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system's memory.

- **Collect network metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the network.

i | **TIP:** If you are collecting basic host metrics using Foglight for VMware and/or Foglight for Hyper-V, you may need to set the **Collect CPU/disk/memory/network metrics** options to *False*, to prevent different or conflicting values from being reported. For Foglight for VMware, consider setting all four flags to *False*. In the case of Foglight for Hyper-V, consider setting CPU to *False*, and possibly memory as well. In some cases, Foglight for Hyper-V can collect memory data more accurately than Foglight for Infrastructure, however that depends on the version of Foglight for Hyper-V.

- **Collect System ID:** Default value = *True*. This property indicates to the agent whether or not to collect a unique system ID from this system. This is not always desirable when monitoring Hyper-V® systems, as some Hyper-V systems use the same ID for multiple systems and are not unique.
- **Collect cluster information:** Default value = *False*. This property indicates to the agent whether or not to collect additional metrics about Windows® clusters.
- **Treat cluster resource groups as virtual host:** Default value = *False*. This property indicates to the agent whether or not hosts in a Windows cluster are be treated as virtual machines or distinct physical hosts.
- **Collect service metrics:** When set to the default value, *Collect only specified services*, the agent collects metrics only for the services specified in the *Services to Monitor* list. When set to *Collect all services*, all service metrics are collected. When set to *Do not collect services*, no service metrics are collected.
- **Services collect type:** When set to the default value, *Collected services by display name*, the agent collects services by service display name. When set to *Collected services by name*, the agent collects services by service name.
- **Submit changed services only:** Default value = *False*. When set to the default value, the agent submits all collected services. When set to *True*, the agent only submits the changed services.
- **Collect log events:** Default value = *True*. When set to *True*, the agent collects log events information.
- **Report only aggregate process metrics:** Default value = *True*. When set to *True*, the agent collects metrics from all processes, aggregates the information, and presents it in a unified report. When set to *False*, the agent still collects aggregate data, but it also includes details about every process. This could result in a lot of data sent to the Management Server and may have a performance impact.
- **Collect Top N Process Details:** Default value = *True*. When set to *True*, the agent collects data for the Top CPU Processes, Top Memory Processes, and Top IO Processes. Details about these top processes are accessible from the Infrastructure Environment dashboard (for example, to see the top CPU processes, in the Monitoring tab, select a host on the Quick view, click the **Explore** button in the Resource Utilizations view, and click any metric indicator in the CPU area; the top CPU consumers are displayed in the CPU Details dashboard.)

When set to *False*, some extra details (owning username/domain) are unavailable for processes reported in the various “Top N” collections. Gathering this information can be expensive if the connection to the remote machine is slow or the “Top N” collections are configured to be very large.

- **Use ping to validate host availability:** Default value = *False*. When set to *True*, the agent is configured to use `ping` to detect if the monitored host is unavailable. If the agent fails to make a connection to the monitored host, and this property is set to *True*, the agent sends a ping command to the host. If the host does not respond, the *Host.monitored* observation is set to UNAVAILABLE (for more details, see [Host availability alerting](#)).
- **Excluded Drives:** A list of drives that are excluded from monitoring. You can modify, clone, and delete lists of excluded drives, as necessary. The list contains two columns: **Host Name Expression** and **Drive Name Expression**. An entry in the list consists of two regular expressions that together identify one or more drives that are excluded from monitoring. For example:

Table 3. Example of drives excluded from monitoring

Column	Value	Meaning
Host Name Expression	tor.*	All hosts starting with “tor” such as tor.test.com or tor.prod.com
Drive Name Expression	[ST]:.*	Drive letters starting with either S or T, followed by a colon ‘:’, followed by any other characters (typically a backslash ‘\’), such as S:\ and T:\

i | **IMPORTANT:** Failing to provide a valid regular expression can result in data loss. A list containing invalid expressions (for example, [ST]:* instead of [ST]:.*) causes the loss of all OS metrics collected from the hosts associated with that list.

i | **NOTE:** For complete information about the regular expressions syntax, see [Appendix: Building regular expressions in Foglight](#).

Click **Edit** to modify the entries in the list. In the **HostAgents - WindowsAgent - ListName** dialog box that appears, you have several options: add or delete rows, edit the existing entries, select drives to be excluded, save or revert changes. To exclude one or more drives from monitoring, add table rows entries and populate them with regular expressions, as required.

Click **Clone** to clone the selected exclude list. In the **Clone ListName** dialog box that appears, enter a name for the new list, and click **OK** to save the it. The new exclude list is added to the drop-down list.

Click **Delete** to delete the selected exclude list. In the **Delete ListName** dialog box that appears, confirm the deletion by clicking **Yes**. The exclude list is deleted from the drop-down list.

- **Process Availability Config:** A list of monitored processes and their expected instance counts. The list contains three columns: **Process Name**, **Command Line**, and **Expected Process Count**, and can be edited, as required. The agent compares the number of actual processes with the number of expected processes, found in this list. Results are displayed in the **Processes > User Defined Processes (Process Availability Config)** view (for details, see [User Defined Processes \(Process Availability Config\)](#)).

i | **TIP:** This feature replaces the AppMonitor capability provided by the legacy OS cartridge.

The combination of Process Name, Command Line, and Expected Process Count consists of the expression that identifies the query criteria for filtering out the process instances. The following samples demonstrate how to find the values of Process Name and Command Line on Solaris and Windows platforms:

i | **TIP:** So far, we only support exact string matching for Process Name, while regular expression is supported for Command Line.

- Windows: Open the Task Manager, and click Process.

i | **TIP:** The Command Line column is hidden by default. In the Task Manager, click View > Select Columns... > Command Line to show this column.

```
jusched.exe *32 "C:\Program Files (x86)\Common Files\Java\Java Update\jusched.exe"
```

In the sample above:

- Process Name is jusched.exe
- Command Line can be either C:\Program Files (x86)\Common Files\Java\Java Update\jusched.exe or some keywords (for example, jusched)

- **Filesystem Type Config: A list of filesystem types.**

- **Event Log:**

- **Hours of backlog events to read:** controls how much historical data is collected by the agent when it is first activated.
- **Event Logs to Monitor:** provides the list of event logs that users want to monitor. Windows® automatically maintains a few different event logs. Users can modify, clone, and delete these event logs, as necessary.

Click **Edit** to modify the log selected from the drop-down list. In the **HostAgents - WindowsAgent - LogName** dialog box, you have several options: add or delete rows, modify fields, select logs to be monitored, save or revert changes. The *Event Log Name* is the name of the log to monitor; this column is pre-populated with a few default values; users who want to monitor other logs can add new rows. The *Monitor* check boxes instruct the agent whether to monitor that log or not.

Click **Clone** to clone the log selected from the drop-down list. In the **Clone LogName** dialog box, enter a name for the new event log, and click **OK** to save the it. The new event log is added to the drop-down list.

Click **Delete** to delete the log selected from the drop-down list. In the **Delete LogName** dialog box, confirm the deletion by clicking **Yes**. The event log is deleted from the drop-down list.

- **Event Log Filter:** provides the ability to filter out the events that should not be monitored (for example, messages that are present in the event log but are irrelevant to users, therefore users should not be alerted of them). Windows provides a *default* log filter. Users can modify, clone, and delete this filter, as necessary.

Each entry in the list indicates the event type, whether to include or exclude it, the event source, category, event ID and description, and event throttle count and its duration.

Event Throttle Count: When set, this ensures that one event in every count (the event log entry occurrence that the filter applies to) is submitted to the Management Server. When If the count is one or less, then every event log entry occurrence is submitted and no throttling is done. The default is zero '0'.

Event Throttle Duration (seconds): This value represents the duration in seconds for the throttle count to be applied. When set, the throttle count is applied within a duration. After the duration expires, the throttling restarts from the beginning regardless of the current throttle state. If the count is one or less, then only one event log entry the filter matches is submitted within the specified duration. If the count is larger than one, then only one in every count (the event log entry occurrence that the filter matches) is submitted, and the agent starts counting pattern matches from zero after the duration. The default value is zero '0', which means the duration is not applied.

i | **NOTE:** The **Event Throttle Count** and **Event Throttle Duration (seconds)** properties only apply to INCLUDE-type filters (and not EXCLUDE-type filters), since throttling is necessary only if a message is being included (and submitted).

Click **Edit** to modify the filter selected from the drop-down list. In the **HostAgents - WindowsAgent - FilterName** dialog box, you have several options: add or delete rows, modify fields, save or revert changes.

Click **Clone** to clone the filter selected from the drop-down list. In the **Clone FilterName** dialog box, enter a name for the new event log, and click **OK** to save the it. The new filter is added to the drop-down list.

Click **Delete** to delete the filter selected from the drop-down list. In the **Delete FilterName** dialog box, confirm the deletion by clicking **Yes**. The filter is deleted from the drop-down list.

- **Event Log Severity:** allows users to map the severity levels that are defined in Foglight to the severity levels (or Type) that are defined in the Windows event log.
- **Monitored Services:**
 - **Services to Monitor:** provides lists of services that users want to monitor. Windows provides a *default* list of services. Users can modify, clone, and delete lists of services, as necessary.

Click **Edit** to modify the selected service list. In the **HostAgents - WindowsAgent - ListName** dialog box, you have several options: add or delete rows, modify fields, select services to be monitored, save or revert changes. The *Service Name* is the Display Name of the service to monitor; this column is pre-populated with a few default values; users who want to monitor other

services can add new rows. The *Monitor* check boxes instruct the agent whether to monitor that service or not. The check boxes in the *Restart if Stopped* column allow you to restart stopped Windows services, as required.

Click **Clone** to clone the selected service list. In the **Clone ListName** dialog box, enter a name for the new service list, and click **OK** to save it. The new list is added to the monitored services drop-down list.

Click **Delete** to delete the selected service list from the monitored services drop-down list. In the **Delete ListName** dialog box, confirm the deletion by clicking **Yes**. The service list is deleted from the monitored services drop-down list.

- Services collect type: provides the following two options for the service queries.
 - Collected services by display name (default option): Collect the monitored services based on the Display Name.
 - Collected services by name: Collect the monitored services based on the Service Name.
- **Connection Config:**
 - **Connection Protocol:** Provides four connection protocols: *WMI*, *Default*, *WinRM with HTTP*, and *WinRM with HTTPS*. When *Default* is selected, WindowsAgent will try to set up the connection using WinRM with different ports(80, 443, 5985, 5986) first. If failed, WindowsAgent will set up the connection to the monitored host using WMI.
 - **WinRM HTTP Ports:** Defines the port number when using the *WinRM with HTTP* protocol.
 - **WinRM HTTPS Ports:** Defines the port number when using the *WinRM with HTTPS* protocol.
- **Data Collection Scheduler:**
 - **Collector Config:** defines how quickly the agent collects data. Windows provides a defaultSchedule configuration. Users can modify, clone, and delete configurations, as necessary.

Click **Edit** to modify the configuration selected from the drop-down list. In the **HostAgents - WindowsAgent - ConfigurationName** dialog box, you have several options: add or remove rows (that is, collectors and their settings), modify fields, and save or revert changes.

Click **Clone** to clone the configuration selected from the drop-down list. In the **Clone ConfigurationName** dialog box, enter a name for the new configuration, and click **OK** to save the it. The new configuration is added to the drop-down list.

Click **Delete** to delete the configuration selected from the drop-down list. In the **Delete ConfigurationName** dialog box, confirm the deletion by clicking **Yes**. The configuration is deleted from the drop-down list.

About the UnixAgentPlus

The UnixAgentPlus monitors Linux® and Oracle Solaris® platforms and collects the following information:

- CPU utilization and performance
- Memory and network performance
- Physical disk devices on a system
- Paging files (such as available paging space)
- Processes running or waiting to run on a monitored host
- One or more network interface cards for network traffic data (such as the number of packets, throughput, errors, utilization, and collisions).

- i** **NOTE:** The *UnixAgent* monitoring Solaris platforms treated ZFS® pools (retrieved by executing the `zpool list` command) as *LogicalDisks*, however the *UnixAgentPlus* treats them as *PhysicalDisk*. Due to this change, when switching from *UnixAgent* to *UnixAgentPlus*, you may notice that certain filesystem-related alarms that were raised for *LogicalDisks* under *UnixAgent* are now raised for *PhysicalDisks* under *UnixAgentPlus*. Additionally, charts for *ZFS Pools* that used to be populated for *LogicalDisks* under *UnixAgent* are populated for *PhysicalDisks* under *UnixAgentPlus*.

For more details, see the following:

- [Supported platforms](#)
- [Agent properties](#)

Supported platforms

For a list of platforms supported for the *UnixAgentPlus*, see “System Requirements” in the *Foglight for Infrastructure Release Notes*.

- i** **NOTE:** The infrastructure UNIX® agents do not execute shell commands on the monitored system. Therefore, there are no special requirements or restrictions regarding the shells on which they are installed. The agents for the AIX® and HP-UX platforms include an optional native executable to gather certain metrics.

Agent properties

When an agent connects to the Foglight Management Server, it is provided with a set of properties that it uses to configure its correct running state. For more information about working with agent properties, see [Creating agent instances](#).

The *UnixAgentPlus* is shipped with default properties that can be modified to suit your system requirements. The properties specific to the *UnixAgentPlus* are illustrated in the following screenshot.

Figure 11. UnixAgentPlus properties

You can configure the following settings for this agent:

- **Properties:**

- **Host:** host name or IP address.
- **Host name override:** host name to be used to store this host's data in the Foglight data model.
- **Port:** SSH port on which the agent connects. Default value = 22.
- **Top CPU Processes:** number of top CPU processes to be monitored. Default value = 5.
- **Top Memory Processes:** number of top memory processes to be monitored. Default value = 5.
- **Top IO Processes:** number of top IO processes to be monitored. Default value = 5.
- **Keep Retry If Collection Failed:** Default value = *False*. When set to *True*, the agent keeps trying to collect data when failed to connect to the target host during startup.
- **Collect Top N processes only:** Default value = *False*. When set to *True*, the agent collects data for Top CPU Processes, Top Memory Processes, and Top IO Processes only.
- **Use native collector (if available):** Default value = *True*.

i NOTE: The AIX® and HP-UX agents deploy a simple binary program to the remote system and execute it to gather a better collection of metrics. It runs in the same user account under which the Agent Manager connects, and is deleted from the system when the agent is deleted. This program is not required for the agent to operate, but some data become unavailable if it is turned off. HP-UX in particular requires the native binary to collect many of the standard metrics.

- **Aggregate data for all instances of a program:** Default value = *True*. When set to *True*, the agent collects data from all the instances of a program (for example all Oracle® instances), aggregates the information, and presents it in a unified report.

When set to *False*, the agent still collects aggregate data, but it also includes details about every process. This could result in a lot of data sent to the Management Server and may have a performance impact.

- **Collect Top N Process Details:** Default value = *True*. When set to *True*, the agent collects data for the Top CPU Processes, Top Memory Processes, and Top IO Processes. Details about these top processes are accessible from the Infrastructure Environment dashboard (for example, to see the top CPU processes, in the Monitoring tab, select a host on the Quick view, click the **Explore** button in the Resource Utilizations view, and click any metric indicator in the CPU area; the top CPU consumers are displayed in the CPU Details dashboard.)

When set to *False*, some extra details (owning username/domain) are unavailable for processes reported in the various “Top N” collections. Gathering this information can be expensive if the connection to the remote machine is slow or the “Top N” collections are configured to be very large.

- **Collect process metrics:** Default value = *True*. When set to *True*, the agent collects process metrics.
- **Collect declared processes only:** Default value = *False*. When set to *False*, the agent collects metrics for all processes that matches the values defined in Process Availability Config. An alarm will be raised when the actual process count is lower than Expected Process Count.

When the Collect declared processes only property is set to *true*, only processes declared under Process Availability Config will be reported.

- **Collect CPU metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s CPUs.
- **Collect disk metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s disks.
- **Include filesystems mounted from memory:** Default value = *False*. This property indicates to the agent whether or not to collect information about RAM disks. This information is typically collected when monitoring Linux® and Solaris® platforms, and not collected for HP-UX and AIX® platforms.
- **Include mounted remote filesystems:** Default value = *False*. When set to *True*, the agent collects metrics about remotely mounted disks.
- **Include SAN mounted disks:** Default value = *False*. When set to *True*, the agent collects metrics about SAN mounted disks.
- **Include Unmounted disks:** Default value = *True*. When set to *True*, the agent collects metrics about unmounted LogicalDisks.
- **Collect memory metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s memory.
- **Collect network metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the network.

i | **TIP:** If you are collecting basic host metrics using Foglight for VMware, you may need to set the **Collect CPU/disk/memory/network metrics** options to *False*, to prevent Foglight for VMware and Foglight for Infrastructure from reporting different or conflicting values. For Foglight for VMware, consider setting all four flags to *False*.

- **Collect System ID:** Default value = *True*. This property indicates to the agent whether or not to collect a unique system ID from this system. This is not always desirable when monitoring Hyper-V® systems, as some Hyper-V systems use the same ID for multiple systems and are not unique.
- **Collect Hypervisor metrics:** Default value = *False*. This property indicates to the agent whether or not to collect additional metrics from hypervisor systems (for example, Solaris global Zone, AIX® LPAR, and so on).
- **Use ping to validate host availability:** Default value = *False*. When set to *True*, the agent is configured to use `ping` to detect if the monitored host is unavailable. If the agent fails to make a connection to the monitored host, and this property is set to *True*, the agent sends a ping command to the host. If the host does not respond, the *Host.monitored* observation is set to UNAVAILABLE (for more details, see [Host availability alerting](#)).

i | **NOTE:** When the **Use ping to validate host availability** property is enabled on a UNIX® platform, the `sudoer` file needs to be configured to allow the ICMP process to run with NOPASSWD. For details, see [Configuring secure launcher permissions using sudo](#).

- **Use commands with sudo:** Default value = *False*. When set to *False*, the agent does not use commands that require sudo, and does not collect metrics that require root permissions. For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).
- **Filter local disks based on declared filesystem types:** Default value = *True*. When set to *True*, the agent enables the local filesystem type filtering.
- **Path to sudo command:** The path to the sudo executable.
- **Process Availability Config:** A list of monitored processes and their expected instance counts. The list contains three columns: Process Name, Command Line, and Expected Process Count, and can be edited, as required. When the Collect declared processes only property is set to true, only processes declared under Process Availability Config will be reported.

The agent compares the number of actual processes with the number of expected processes found in this list. Results are displayed in the **Processes > Processes > User Defined Processes (Process Availability Config)** view (for details, see [User Defined Processes \(Process Availability Config\)](#)).

i | TIP: This feature replaces the AppMonitor capability provided by the legacy OS cartridge.

The combination of Process Name, Command Line, and Expected Process Count consists of the expression that identifies the query criteria for filtering out the process instances. The following samples demonstrate how to find the values of Process Name and Command Line on Solaris:

i | TIP: So far, we only support exact string matching for Process Name, while regular expression is supported for Command Line.

- **Solaris: Execute the “/usr/bin/ps -e -o uid,pid,ppid,vsz,rss,time,pcpu,sid,s,user,comm,args” command.** Then you will get the following process details.

```
16801 21353 21351 249104 159096 01:08 0.0 2665 S murex
/usr/local/java/jdk1.7.0_51/bin/java
/usr/local/java/jdk1.7.0_51/bin/java -Dname=RTZSP_DEALGEN -
Djava.library.path=/
```

In the sample above:

- Process Name is /usr/local/java/jdk1.7.0_51/bin/java
- Command Line can be either usr/local/java/jdk1.7.0_51/bin/java - Dname=RTZSP_DEALGEN -Djava.library.path=/ or some keywords (for example, RTZSP_DEALGEN)

- **Filesystem Type Config: A list of filesystem types.**
- **Exclude/Include FileSystems:** The type of FileSystems list to be used for monitoring.
 - **Exclude** (default) indicates that the file systems listed in the FileSystems list should be excluded from monitoring.
 - **Include** performs system monitoring on the file system that you are defining.
- **Filesystem Config:** A list of file systems that are excluded from monitoring (if the **Exclude/Include FileSystems** property is set to **Exclude**) or included in the monitoring (if the **Exclude/Include FileSystems** property is set to **Include**). You can modify, clone, and delete lists of excluded/included file systems, as necessary. The list contains three columns: **MountPoint regular expression**, **Remote host name regular expression**, and **Monitored host regular expression**. An entry in the list consists of three regular expressions that together identify one or more file systems that are excluded from/ included in the monitoring. For example:

Table 4. Example of file systems list

Column	Value	Meaning
MountPoint regular expression	/workspace	If set, the file systems with the matching mount point should be excluded/included. In this example: All file systems located in /workspace.
Remote host name regular expression	tor.*	This is related to the file systems which are remotely mounted. If set, the remote file systems with the matching remote host should be excluded/included. In this example: All remote hosts starting with “tor” such as tor.test.com or tor.prod.com.
Monitored host regular expression	tor.*	This refers to the hostname collected by the OS. If set, all the OS-collected metrics for the matching host should be excluded/included. In this example: All hosts starting with “tor” such as tor.test.com or tor.prod.com.

i | **NOTE:** For complete information about the regular expressions syntax, see [Appendix: Building regular expressions in Foglight](#).

Click **Edit** to modify the entries in the list. In the **HostAgents - UnixAgentPlus - ListName** dialog box that appears, you have several options: add or delete rows, edit the existing entries, select file systems to be excluded/included, save or revert changes. To exclude/include one or more file systems from monitoring, add table rows entries and populate them with regular expressions, as required.

Click **Clone** to clone the selected exclude/include list. In the **Clone ListName** dialog box that appears, enter a name for the new list, and click **OK** to save the it. The new exclude/include list is added to the drop-down list.

Click **Delete** to delete the selected exclude/include list. In the **Delete ListName** dialog box that appears, confirm the deletion by clicking **Yes**. The exclude/include list is deleted from the drop-down list.

- **Data Collection Scheduler:**

- **Collector Config:** Defines how quickly the agent collects data. UNIX® provides a defaultSchedule configuration. Users can modify, clone, and delete configurations, as necessary.

Click **Edit** to modify the configuration selected from the drop-down list. In the **HostAgents - UnixAgentPlus - ConfigurationName** dialog box, you have several options: add or remove rows (that is, collectors and their settings), modify fields, and save or revert changes.

Click **Clone** to clone the configuration selected from the drop-down list. In the **Clone ConfigurationName** dialog box, enter a name for the new configuration, and click **OK** to save the it. The new configuration is added to the drop-down list.

Click **Delete** to delete the configuration selected from the drop-down list. In the **Delete ConfigurationName** dialog box, confirm the deletion by clicking **Yes**. The configuration is deleted from the drop-down list.

Configuring secure launcher permissions using sudo

Some UnixAgentPlus can function without root privileges, but certain metrics can only be collected by commands which must be run as root. In order to give these agents the required access, Foglight Agent Manager is configured to launch these agents using a tool such as sudo that allows privilege escalation (without a password).

To this effect, the sudo configuration file (*/etc/sudoers*) must be configured so that password prompts are not required for a number of executables. The commands requiring elevated privileges differ by platform. The following commands must be configured for this version of Foglight for Infrastructure.

Table 5. Sudo commands that must be configured

Agent Type	Command	Notes
Linux®	<code>/usr/bin/find, /bin/cat</code>	Used to read IO statistics from the <code>/proc</code> filesystem.
	<code>/sbin/ethtool</code> or <code>/usr/sbin/ethtool</code> (depending on distribution)	Used to determine the network card bandwidth; <code>ethtool</code> is favoured if it is found.
	<code>/sbin/mii-tool</code> or <code>/usr/sbin/mii-tool</code> (depending on distribution)	
Oracle Solaris®	N/A	Oracle Solaris agents do not require sudo access.

The following is an example of how to configure the `/etc/sudoers` file to allow the user `foglight` to execute Linux® commands without being prompted for a password:

```
foglight ALL = NOPASSWD: /usr/bin/find, /bin/cat, /sbin/ethtool
```

In addition, the `requiretty` flag must **not** be set in `/etc/sudoers` for the user, since Foglight for Infrastructure agents use non-interactive shells.

The following is an example of how to unset the `requiretty` flag for a single user named `foglight`, so that this user can run `sudo` commands remotely:

```
Defaults:foglight !requiretty
```

i | **NOTE:** If `requiretty` flag is set, `sudo` can run only when the user is logged in to a real `tty`. When this flag is set, `sudo` can only be run from a login session and not via other means, such as `cron` or `cgi-bin` scripts. This flag is off (unset) by default.

Controlling log usage for sudo access

Using commands with `sudo` access can result in increased logging. Sudo provides the following levels of logging, each resulting in the capture of a specific type of information:

- Log all commands that run with `sudo`.
- Log only commands that run with `sudo` and cause errors.
- Log all command-line output produced by the commands that run with `sudo`.
- Log all command-line input provided to the commands that run with `sudo`.

Depending on the user's `sudo` and `syslog.conf` configuration, `sudo` use may result in excess logging. To minimize the amount of log messages, ensure that `sudo` does not make use of the `LOG_INPUT` or `LOG_OUTPUT` tags for the commands that the UnixAgent runs. Depending on the existing monitored hosts' configuration, any lines added to the `/etc/sudoers` file for Foglight monitoring may have to include `NOLOG_OUTPUT` or `NOLOG_INPUT` to override the default configuration. For example, for a user named `foglight` connecting to a monitored host, the following lines are required:

Linux

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,
  /bin/cat, [/sbin/ethtool|/usr/sbin/ethtool|sbin/mii-tool|/usr/sbin/mii-tool]
```

The last argument in this syntax depends on the type and location of the tool, `ethtool` or `mii-tool`, used to determine the network card bandwidth. If you are unsure which tool your system uses, you can specify all of them:

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,
  /bin/cat, /sbin/ethtool, /usr/sbin/ethtool, /sbin/mii-tool, /usr/sbin/mii-tool
```

Oracle Solaris

N/A

About the UnixAgent

The *UnixAgent* monitors Linux®, Oracle Solaris®, HP-UX, or AIX® systems and collects the following information:

- CPU utilization and performance
- Memory and network performance
- Physical disk devices on a system
- Paging files (such as available paging space)
- Processes running or waiting to run on a monitored host
- One or more network interface cards for network traffic data (such as the number of packets, throughput, errors, utilization, and collisions)

i | **NOTE:** The *UnixAgent* monitoring Solaris platforms treated ZFS® pools (retrieved by executing the `zpool list` command) as *LogicalDisks*, however the *UnixAgentPlus* treats them as *PhysicalDisk*. Due to this change, when switching from *UnixAgent* to *UnixAgentPlus*, you may notice that certain filesystem-related alarms that were raised for *LogicalDisks* under *UnixAgent* are now raised for *PhysicalDisks* under *UnixAgentPlus*. Additionally, charts for *ZFS Pools* that used to be populated for *LogicalDisks* under *UnixAgent* are populated for *PhysicalDisks* under *UnixAgentPlus*.

i | **IMPORTANT:** Using the native collector to monitor an HP-UX system requires that the monitoring *UnixAgent*'s account belong to the `sys` and `bin` groups on the monitored system. Failing to add the account to these groups prevents the agent from collecting some logical disk metrics.

There are views, rules, and data associated with this agent. For more information, see [Reference](#).

For more details, see these topics:

- [Supported platforms](#)
- [Agent properties](#)

Supported platforms

For a list of platforms supported for the *UnixAgent*, see “System Requirements” in the *Foglight for Infrastructure Release Notes*.

i | **NOTE:** The infrastructure UNIX® agents do not execute shell commands on the monitored system. Therefore, there are no special requirements or restrictions regarding the shells on which they are installed. The agents for the AIX® and HP-UX platforms include an optional native executable to gather certain metrics.

Agent properties

When an agent connects to the Foglight Management Server, it is provided with a set of properties that it uses to configure its correct running state. For more information about working with agent properties, see [Creating agent instances](#).

The *UnixAgent* is shipped with default properties that can be modified to suit your system requirements. The properties specific to the *UnixAgent* are illustrated in the following screenshot.

Figure 12. UnixAgent properties

You can configure the following settings for this agent:

- **Properties:**

- **Host:** host name or IP address.
- **Host name override:** host name to be used to store this host's data in the Foglight data model.
- **Port:** SSH port on which the agent connects. Default value = 22.
- **Top CPU Processes:** number of top CPU processes to be monitored. Default value = 5.
- **Top Memory Processes:** number of top memory processes to be monitored. Default value = 5.
- **Top IO Processes:** number of top IO processes to be monitored. Default value = 5.
- **Keep Retry If Collection Failed:** Default value = *False*. When set to *True*, the agent keeps trying to collect data when failed to connect to the target host during startup.
- **Collect Top N processes only:** Default value = *False*. When set to *True*, the agent collects data for Top CPU Processes, Top Memory Processes, and Top IO Processes only.
- **Use native collector (if available):** Default value = *True*.

i **NOTE:** The AIX® and HP-UX agents deploy a simple binary program to the remote system and execute it to gather a better collection of metrics. It runs in the same user account under which the Agent Manager connects, and is deleted from the system when the agent is deleted. This program is not required for the agent to operate, but some data become unavailable if it is turned off. HP-UX in particular requires the native binary to collect many of the standard metrics.

i **IMPORTANT:** Using the native collector to monitor an HP-UX system requires that the monitoring UnixAgent's account belong to the `sys` and `bin` groups on the monitored system. Failing to add the account to these groups prevents the agent from collecting some logical disk metrics.

- **Aggregate data for all instances of a program:** Default value = *True*. When set to *True*, the agent collects data from all the instances of a program (for example all Oracle® instances), aggregates the information, and presents it in a unified report.

When set to *False*, the agent still collects aggregate data, but it also includes details about every process. This could result in a lot of data sent to the Management Server and may have a performance impact.

- **Collect Top N Process Details:** Default value = *True*. When set to *True*, the agent collects data for the Top CPU Processes, Top Memory Processes, and Top IO Processes. Details about these top processes are accessible from the Infrastructure Environment dashboard (for example, to see the top CPU processes, in the Monitoring tab, select a host on the Quick view, click the **Explore** button in the Resource Utilizations view, and click any metric indicator in the CPU area; the top CPU consumers are displayed in the CPU Details dashboard.)

When set to *False*, some extra details (owning username/domain) are unavailable for processes reported in the various “Top N” collections. Gathering this information can be expensive if the connection to the remote machine is slow or the “Top N” collections are configured to be very large.

- **Collect process metrics:** Default value = *True*. When set to *True*, the agent collects process metrics.
- **Collect CPU metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s CPUs.
- **Collect disk metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s disks.
- **Include filesystems mounted from memory:** Default value = *False*. This property indicates to the agent whether or not to collect information about RAM disks. This information is typically collected when monitoring Linux® and Solaris® platforms, and not collected for HP-UX and AIX® platforms.
- **Include mounted remote filesystems:** Default value = *False*. When set to *True*, the agent collects metrics about remotely mounted disks.
- **Collect memory metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the system’s memory.
- **Collect network metrics:** Default value = *True*. When set to *True*, the agent collects performance metrics about the network.

i | **TIP:** If you are collecting basic host metrics using Foglight for VMware, you may need to set the **Collect CPU/disk/memory/network metrics** options to *False*, to prevent Foglight for VMware and Foglight for Infrastructure from reporting different or conflicting values. For Foglight for VMware, consider setting all four flags to *False*.

- **Collect System ID:** Default value = *True*. This property indicates to the agent whether or not to collect a unique system ID from this system. This is not always desirable when monitoring Hyper-V® systems, as some Hyper-V systems use the same ID for multiple systems and are not unique.
- **Collect Hypervisor metrics:** Default value = *False*. This property indicates to the agent whether or not to collect additional metrics from hypervisor systems (for example, Solaris global Zone, AIX® LPAR, and so on).
- **Use ping to validate host availability:** Default value = *False*. When set to *True*, the agent is configured to use `ping` to detect if the monitored host is unavailable. If the agent fails to make a connection to the monitored host, and this property is set to *True*, the agent sends a ping command to the host. If the host does not respond, the *Host.monitored* observation is set to UNAVAILABLE (for more details, see [Host availability alerting](#)).

i | **NOTE:** When the **Use ping to validate host availability** property is enabled on a UNIX® platform, the `sudoer` file needs to be configured to allow the ICMP process to run with NOPASSWD. For details, see [Configuring secure launcher permissions using sudo](#).

- **Use commands with sudo:** Default value = *False*. When set to *False*, the agent does not use commands that require `sudo`, and does not collect metrics that require root permissions. For more

information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).

- **Filter local disks based on declared filesystem types:** Default value = *True*. When set to *True*, the agent enables the local filesystem type filtering.
- **Path to sudo command:** The path to the sudo executable.
- **Process Availability Config:** A list of monitored processes and their expected instance counts. The list contains three columns: **Process Name**, **Command Line**, and **Expected Process Count**, and can be edited, as required. The agent compares the number of actual processes with the number of expected processes, found in this list. Results are displayed in the **Processes > User Defined Processes (Process Availability Config)** view (for details, see [User Defined Processes \(Process Availability Config\)](#)).

i | **TIP:** This feature replaces the AppMonitor capability provided by the legacy OS cartridge.

The combination of Process Name, Command Line, and Expected Process Count consists of the expression that identifies the query criteria for filtering out the process instances. The following samples demonstrate how to find the values of Process Name and Command Line on Solaris:

i | **TIP:** So far, we only support exact string matching for Process Name, while regular expression is supported for Command Line.

- **Solaris:** Execute the `"/usr/bin/ps -e -o uid,pid,ppid,vsz,rss,time,pcpu,sid,s,user,comm,args"` command. Then you will get the following process details.

```
16801 21353 21351 249104 159096 01:08 0.0 2665 S murex
/usr/local/java/jdk1.7.0_51/bin/java
/usr/local/java/jdk1.7.0_51/bin/java -Dname=RTZSP_DEALGEN -
Djava.library.path=/
```

In the sample above:

- Process Name is `/usr/local/java/jdk1.7.0_51/bin/java`
 - Command Line can be either `usr/local/java/jdk1.7.0_51/bin/java -Dname=RTZSP_DEALGEN -Djava.library.path=/` or some keywords (for example, `RTZSP_DEALGEN`)
- **Filesystem Type Config: A list of filesystem types.**
 - **Exclude/Include FileSystems:** The type of FileSystems list to be used for monitoring.
 - **Exclude** (default) indicates that the file systems listed in the FileSystems list should be excluded from monitoring.
 - **Include** performs system monitoring on the file system that you are defining.
 - **Filesystem Config:** A list of file systems that are excluded from monitoring (if the **Exclude/Include FileSystems** property is set to **Exclude**) or included in the monitoring (if the **Exclude/Include FileSystems** property is set to **Include**). You can modify, clone, and delete lists of excluded/included file systems, as necessary. The list contains three columns: **MountPoint regular expression**, **Remote host name regular expression**, and **Monitored host regular expression**. An entry in the list consists of three regular expressions that together identify one or more file systems that are excluded from/ included in the monitoring. For example:

Table 6. Example of file systems list

Column	Value	Meaning
MountPoint regular expression	/workspace	If set, the file systems with the matching mount point should be excluded/included. In this example: All file systems located in /workspace.
Remote host name regular expression	tor.*	This is related to the file systems which are remotely mounted. If set, the remote file systems with the matching remote host should be excluded/included. In this example: All remote hosts starting with “tor” such as tor.test.com or tor.prod.com.
Monitored host regular expression	tor.*	This refers to the hostname collected by the OS. If set, all the OS-collected metrics for the matching host should be excluded/included. In this example: All hosts starting with “tor” such as tor.test.com or tor.prod.com.

i | **NOTE:** For complete information about the regular expressions syntax, see [Appendix: Building regular expressions in Foglight](#).

Click **Edit** to modify the entries in the list. In the **HostAgents - UnixAgentPlus - ListName** dialog box that appears, you have several options: add or delete rows, edit the existing entries, select file systems to be excluded/included, save or revert changes. To exclude/include one or more file systems from monitoring, add table rows entries and populate them with regular expressions, as required.

Click **Clone** to clone the selected exclude/include list. In the **Clone ListName** dialog box that appears, enter a name for the new list, and click **OK** to save the it. The new exclude/include list is added to the drop-down list.

Click **Delete** to delete the selected exclude/include list. In the **Delete ListName** dialog box that appears, confirm the deletion by clicking **Yes**. The exclude/include list is deleted from the drop-down list.

- **Data Collection Scheduler:**

- **Collector Config:** defines how quickly the agent collects data. UNIX® provides a defaultSchedule configuration. Users can modify, clone, and delete configurations, as necessary.

Click **Edit** to modify the configuration selected from the drop-down list. In the **HostAgents - UnixAgent - ConfigurationName** dialog box, you have several options: add or remove rows (that is, collectors and their settings), modify fields, and save or revert changes.

Click **Clone** to clone the configuration selected from the drop-down list. In the **Clone ConfigurationName** dialog box, enter a name for the new configuration, and click **OK** to save the it. The new configuration is added to the drop-down list.

Click **Delete** to delete the configuration selected from the drop-down list. In the **Delete ConfigurationName** dialog box, confirm the deletion by clicking **Yes**. The configuration is deleted from the drop-down list.

Configuring secure launcher permissions using sudo

Some UnixAgents can function without root privileges, but certain metrics can only be collected by commands which must be run as root. In order to give these agents the required access, Foglight Agent Manager is configured to launch these agents using a tool such as sudo that allows privilege escalation (without a password).

To this effect, the sudo configuration file (*/etc/sudoers*) must be configured so that password prompts are not required for a number of executables. The commands requiring elevated privileges differ by platform. The following commands must be configured for this version of Foglight for Infrastructure.

Table 7. Sudo commands that must be configured

Agent Type	Command	Notes
AIX®	<code>/usr/bin/svmon</code>	Some AIX versions require elevated permissions to access virtual memory information.
Linux®	<code>/usr/bin/find, /bin/cat</code>	Used to read IO statistics from the <code>/proc</code> filesystem.
	<code>/sbin/ethtool</code> or <code>/usr/sbin/ethtool</code> (depending on distribution)	Used to determine the network card bandwidth; <code>ethtool</code> is favoured if it is found.
	<code>/sbin/mii-tool</code> or <code>/usr/sbin/mii-tool</code> (depending on distribution)	
Oracle Solaris®	N/A	Oracle Solaris agents do not require sudo access.
HP-UX	N/A	HP-UX agents do not require sudo access.

The following is an example of how to configure the `/etc/sudoers` file to allow the user `foglight` to execute Linux® commands without being prompted for a password:

```
foglight ALL = NOPASSWD: /usr/bin/find, /bin/cat, /sbin/ethtool
```

In addition, the `requiretty` flag must **not** be set in `/etc/sudoers` for the user, since Foglight for Infrastructure agents use non-interactive shells.

The following is an example of how to unset the `requiretty` flag for a single user named `foglight`, so that this user can run sudo commands remotely:

```
Defaults:foglight !requiretty
```

i | **NOTE:** If `requiretty` flag is set, sudo can run only when the user is logged in to a real `tty`. When this flag is set, sudo can only be run from a login session and not via other means, such as `cron` or `cgi-bin` scripts. This flag is off (unset) by default.

Controlling log usage for sudo access

Using commands with sudo access can result in increased logging. Sudo provides the following levels of logging, each resulting in the capture of a specific type of information:

- Log all commands that run with sudo.
- Log only commands that run with sudo and cause errors.
- Log all command-line output produced by the commands that run with sudo.
- Log all command-line input provided to the commands that run with sudo.

Depending on the user's sudo and `syslog.conf` configuration, sudo use may result in excess logging. To minimize the amount of log messages, ensure that sudo does not make use of the `LOG_INPUT` or `LOG_OUTPUT` tags for the commands that the UnixAgent runs. Depending on the existing monitored hosts' configuration, any lines added to the `/etc/sudoers` file for Foglight monitoring may have to include `NOLOG_OUTPUT` or `NOLOG_INPUT` to override the default configuration. For example, for a user named `foglight` connecting to a monitored host, the following lines are required:

AIX

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/svmon
```

Linux

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,
```

```
/bin/cat, [/sbin/ethtool|/usr/sbin/ethtool|sbin/mii-tool|/usr/sbin/mii-tool]
```

The last argument in this syntax depends on the type and location of the tool, *ethtool* or *mii-tool*, used to determine the network card bandwidth. If you are unsure which tool your system uses, you can specify all of them:

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,  
/bin/cat, /sbin/ethtool, /usr/sbin/ethtool, /sbin/mii-tool, /usr/sbin/mii-tool
```

Oracle Solaris

N/A

HP-UX

N/A

For complete information about sudo and */etc/sudoers*, refer to the sudo and */etc/sudoers* man pages.

About the MultiHostProcessMonitorAgent

The MultiHostProcessMonitorAgent monitors both Windows® and Linux® systems and collects the following information:

- Processes running or waiting to run on a monitored host

There are views, rules, and data associated with this agent. For more information, see [Reference](#)

For more details, see these topics:

- [Supported platforms](#)
- [Agent properties](#)

Supported platforms

For a list of platforms supported for the MultiHostProcessMonitorAgent, see “System Requirements” in the *Foglight for Infrastructure* Release Notes.

- i** | **NOTE:** The infrastructure UNIX® agents do not execute shell commands on the monitored system. Therefore, there are no special requirements or restrictions regarding the shells on which they are installed.

Agent properties

When an agent connects to the Foglight Management Server, it is provided with a set of properties that it uses to configure its correct running state. For more information about working with agent properties, see [Creating agent instances](#).

The MultiHostProcessMonitorAgent is shipped with default properties that can be modified to suit your system requirements. The properties specific to the MultiHostProcessMonitorAgent are illustrated in the following screenshot.

Figure 13. MultiHostProcessMonitorAgent properties

You can configure the following settings for this agent:

- Properties:
 - Multiple Hosts Config: The list of hosts to be monitored. Click Edit to modify the following Multiple Host Config properties
 - OS Type: type of the operating system (Unix® or Windows®).
 - Host: host name or IP address.
 - Host name override: host name to be used to store this host's data in the Foglight data model.
 - SSH Port: SSH port on which the agent connects. Default value = 22. Applicable to the Linux® platform only.
 - Top CPU Processes: number of top CPU processes to be monitored. Default value = 5.
 - Top Memory Processes: number of top memory processes to be monitored. Default value = 5.
 - Top IO Processes: number of top IO processes to be monitored. Default value = 5.
 - Keep Retry If Collection Failed: Default value = *False*. When set to *True*, the agent keeps trying to collect data when failed to connect to the target host during startup.
 - Collect Top N processes only: Default value = *False*. When set to *True*, the agent collects data for Top CPU Processes, Top Memory Processes, and Top IO Processes only.
 - Aggregate data for all instances of a program: Default value = *True*. When set to *True*, the agent collects data from all the instances of a program (for example all Oracle® instances), aggregates the information, and presents it in a unified report.
When set to *False*, the agent still collects aggregate data, but it also includes details about every process. This could result in a lot of data sent to the Management Server and may have a performance impact.
 - Collect additional metrics for VMWare: Default value = *True*. When set to *True*, the agent collects the following data:
 - Host: *availablePagingSpace, runQueueLength, contextSwitches*
 - CPU: *totalHz, percentUserTime*
 - Memory: *capacity, consumed, pageInRate, pageOutRate, and utilization*
 - Collect Top N Process Details: Default value = *True*. When set to *True*, the agent collects data for the Top CPU Processes, Top Memory Processes, and Top IO Processes. Details about these top processes are accessible from the Infrastructure Environment dashboard (for example, to see the top CPU processes, in the Monitoring tab, select a host on the Quick view, click the Explore button

in the Resource Utilizations view, and click any metric indicator in the CPU area; the top CPU consumers are displayed in the CPU Details dashboard).

- Collect process metrics: Default value = True. When set to True, the agent collects process metrics.
- Collect declared processes only: Default value = False. When set to False, the agent collects metrics for all processes that matches the values defined in Process Availability Config. An alarm will be raised when the actual process count is lower than Expected Process Count.

When the Collect declared processes only property is set to true, only processes declared under Process Availability Config will be reported.

- Collect System ID: Default value = True. This property indicates to the agent whether or not to collect a unique system ID from this system. This is not always desirable when monitoring Hyper-V® systems, as some Hyper-V systems use the same ID for multiple systems and are not unique.
- **Use ping to validate host availability:** Default value = *False*. When set to *True*, the agent is configured to use `ping` to detect if the monitored host is unavailable. If the agent fails to make a connection to the monitored host, and this property is set to *True*, the agent sends a ping command to the host. If the host does not respond, the *Host.monitored* observation is set to UNAVAILABLE (for more details, see [Host availability alerting](#)).

i | **NOTE:** When the **Use ping to validate host availability** property is enabled on a UNIX® platform, the `sudoer` file needs to be configured to allow the ICMP process to run with NOPASSWD. For details, see [Configuring secure launcher permissions using sudo](#).

- Use commands with sudo: Default value = False. Applicable to the Linux® platform only. When set to False, the agent does not use commands that require sudo, and does not collect metrics that require root permissions. For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).
- Path to sudo command: The path to the sudo executable. Applicable to the Linux® platform only.

Process Availability Config: A list of monitored processes and their expected instance counts. The list contains three columns: Process Name, Command Line, and Expected Process Count, and can be edited, as required. The agent compares the number of actual processes with the number of expected processes found in this list. Results are displayed in the **Processes > Processes > User Defined Processes (Process Availability Config)** view (for details, see [User Defined Processes \(Process Availability Config\)](#)).

i | **TIP:** This feature replaces the AppMonitor capability provided by the legacy OS cartridge.

The combination of Process Name, Command Line, and Expected Process Count consists of the expression that identifies the query criteria for filtering out the process instances. The following samples demonstrate how to find the values of Process Name and Command Line on Solaris:

i | **TIP:** So far, we only support exact string matching for Process Name, while regular expression is supported for Command Line.

- **Solaris: Execute the** `"/usr/bin/ps -e -o uid,pid,ppid,vsz,rss,time,pcpu,sid,s,user,comm,args"` command. Then you will get the following process details.

```
16801 21353 21351 249104 159096 01:08 0.0 2665 S murex
/usr/local/java/jdk1.7.0_51/bin/java
/usr/local/java/jdk1.7.0_51/bin/java -Dname=RTZSP_DEALGEN -
Djava.library.path=/
```

In the sample above:

- Process Name is `/usr/local/java/jdk1.7.0_51/bin/java`
- Command Line can be either `usr/local/java/jdk1.7.0_51/bin/java -Dname=RTZSP_DEALGEN -Djava.library.path=/` or some keywords (for example, `RTZSP_DEALGEN`)

- Data Collection Scheduler

- **Collector Config:** defines how quickly the agent collects data. Both Windows and Linux® provide a defaultSchedule configuration. Users can modify, clone, and delete configurations, as necessary.

Configuring secure launcher permissions using sudo

Some UnixAgents can function without root privileges, but certain metrics can only be collected by commands which must be run as root. In order to give these agents the required access, Foglight Agent Manager is configured to launch these agents using a tool such as sudo that allows privilege escalation (without a password).

To this effect, the sudo configuration file (*/etc/sudoers*) must be configured so that password prompts are not required for a number of executables. The commands requiring elevated privileges differ by platform. The following commands must be configured for this version of Foglight for Infrastructure.

Table 8. Sudo commands that must be configured

Agent Type	Command	Notes
Linux®	<code>/usr/bin/find, /bin/cat</code>	Used to read IO statistics from the <i>/proc</i> filesystem.
	<code>/sbin/ethtool or /usr/sbin/ethtool (depending on distribution)</code>	Used to determine the network card bandwidth; <i>ethtool</i> is favoured if it is found.
	<code>/sbin/mii-tool or /usr/sbin/mii- tool (depending on distribution)</code>	

The following is an example of how to configure the */etc/sudoers* file to allow the user `foglight` to execute Linux® commands without being prompted for a password:

```
foglight ALL = NOPASSWD: /usr/bin/find, /bin/cat, /sbin/ethtool
```

In addition, the `requiretty` flag must **not** be set in */etc/sudoers* for the user, since Foglight for Infrastructure agents use non-interactive shells.

The following is an example of how to unset the `requiretty` flag for a single user named `foglight`, so that this user can run `sudo` commands remotely:

```
Defaults:foglight !requiretty
```

i | **NOTE:** If `requiretty` flag is set, `sudo` can run only when the user is logged in to a real `tty`. When this flag is set, `sudo` can only be run from a login session and not via other means, such as `cron` or `cgi-bin` scripts. This flag is off (unset) by default.

Controlling log usage for sudo access

Using commands with `sudo` access can result in increased logging. Sudo provides the following levels of logging, each resulting in the capture of a specific type of information:

- Log all commands that run with `sudo`.
- Log only commands that run with `sudo` and cause errors.
- Log all command-line output produced by the commands that run with `sudo`.
- Log all command-line input provided to the commands that run with `sudo`.

Depending on the user's `sudo` and *syslog.conf* configuration, `sudo` use may result in excess logging. To minimize the amount of log messages, ensure that `sudo` does not make use of the `LOG_INPUT` or `LOG_OUTPUT` tags for the commands that the UnixAgent runs. Depending on the existing monitored hosts' configuration, any lines added to the */etc/sudoers* file for Foglight monitoring may have to include `NOLOG_OUTPUT` or `NOLOG_INPUT` to override the default configuration. For example, for a user named `foglight` connecting to a monitored host, the following lines are required:

Linux

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,
```

```
/bin/cat, [/sbin/ethtool|/usr/sbin/ethtool|sbin/mii-tool|/usr/sbin/mii-tool]
```

The last argument in this syntax depends on the type and location of the tool, *ethtool* or *mii-tool*, used to determine the network card bandwidth. If you are unsure which tool your system uses, you can specify all of them:

```
foglight ALL = NOLOG_INPUT: ALL, NOLOG_OUTPUT: ALL, NOPASSWD: /usr/bin/find|,  
/bin/cat, /sbin/ethtool, /usr/sbin/ethtool, /sbin/mii-tool, /usr/sbin/mii-tool
```

Monitoring log files with Foglight Log Monitor

Foglight for Infrastructure relies on the File Log Monitor and Windows Event Log Monitor agents to collect data. These agents collect desired information from selected logs. A log file consists of one or more entries, or log records. Depending on the format of a monitored log file a log record can span multiple lines. The collected information is visualized on the Log Monitor dashboard.

Start by ensuring that Foglight for Infrastructure is installed on the Management Server, and that the agent package is deployed. For installation instructions, see the *Foglight for Infrastructure Release Notes*.

Next, configure the File Log Monitor and Windows Event Log Monitor agents for data collection. For more information, see [Configuring monitoring agents](#), [Configuring agent properties](#), and [Configuring connections to remote Windows platforms](#).

When your monitoring agent instances are configured and are collecting data, navigate to the Log Monitor dashboard. This dashboard allows you to look at individual log records, and observe their growth rate over the selected time range. For more information, see [Investigating log records](#).

Configuring monitoring agents

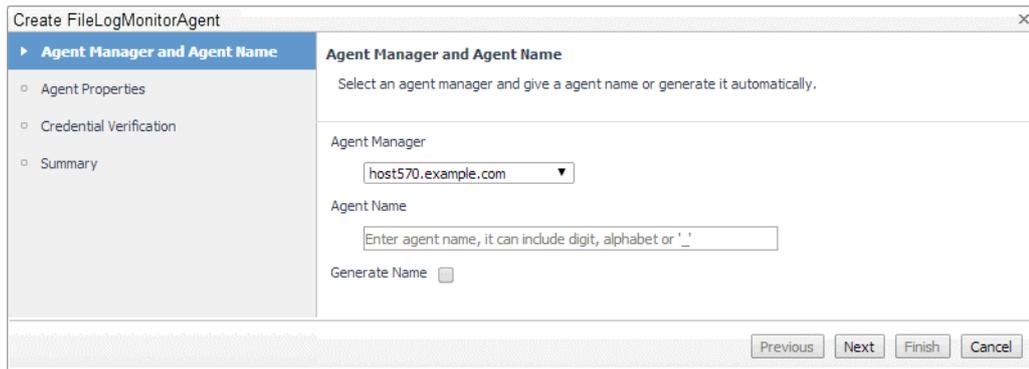
Foglight for Infrastructure uses Quest File Log Monitor Agent and Quest Windows Log Monitor Agent instances to collect information from monitored hosts. When Foglight for Infrastructure is installed on the Management Server and the Host Agents package is deployed to a desired FoglightAgent Manager host, you can create these agents and configure them for data collection.

While the Quest File Log Monitor Agent collects information from selected text files, Quest Windows Event Log Monitor Agent collects information from Windows Event Log files. Both agents can look for the text patterns you specify in the monitored logs.

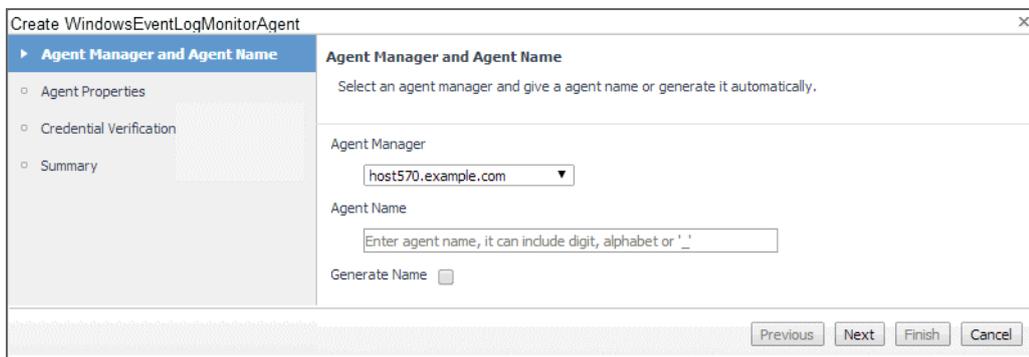
To create an agent instance, activate it, and start its data collection, use the appropriate wizard (**Create FileLogMonitorAgent** or **Create WindowsEventLogMonitorAgent**), accessible from the Log Monitor dashboard.

To create a monitoring agent instance, activate it, and start its data collection:

- 1 Log in to the Foglight browser interface.
- 2 On the navigation panel, under **Dashboards**, click **Log Monitor**.
 - To create a Quest File Log Monitor Agent instance, on the Log Monitor dashboard, in the top-right corner, click **File Log Monitor** to launch the **Create FileLogMonitorAgent** wizard.



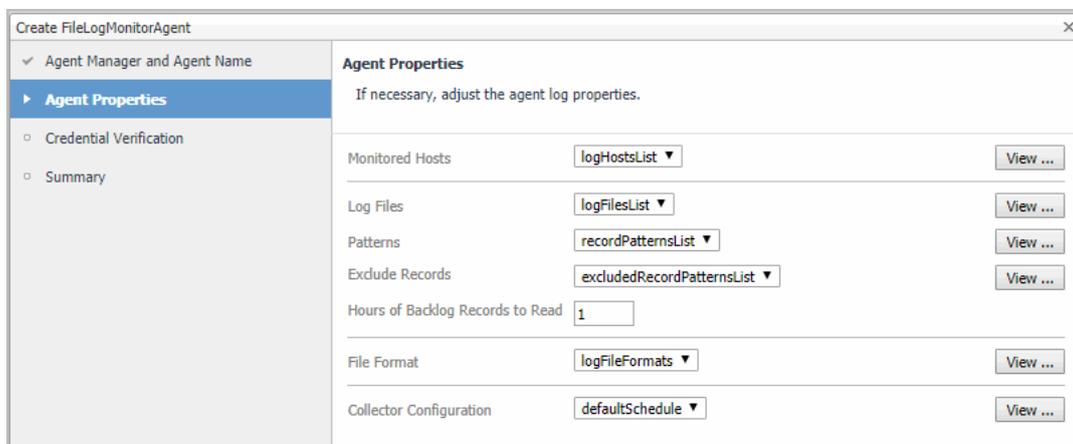
- To create a Quest Windows Event Log Monitor Agent instance, on the Log Monitor dashboard, in the top-right corner, click **Windows Event Log Monitor** to launch the **Create WindowsEventLogMonitorAgent** wizard.



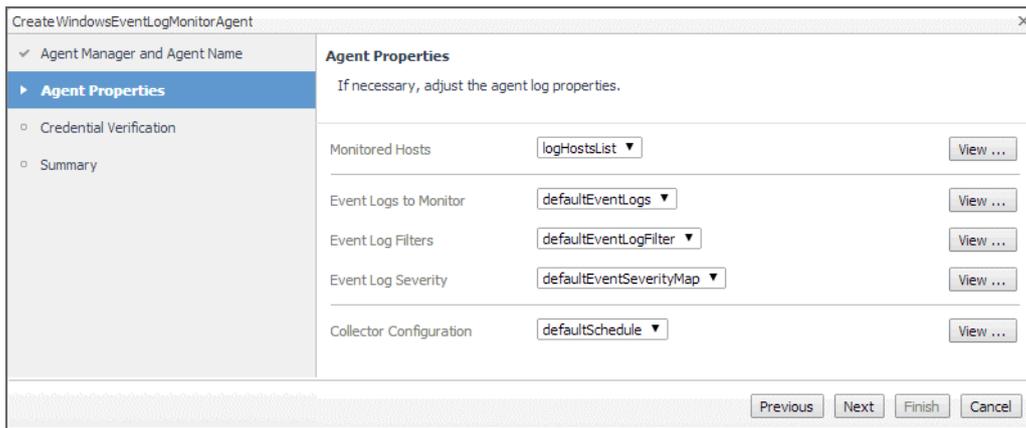
- 3 Select the host where you want to run the agent instance, and specify the agent name.
 - a On the **Agent Manager and Agent Name** page in the wizard, click **Agent Manager**, and from the list that appears, select a host running an Agent Manager instance that you want to use to manage the agent instance that you are about to create.
 - b Specify the name you want to assign to this agent instance. In the **Agent Name** box, type the agent name. Optionally, select the **Generate Name** check box to have the wizard generate the name automatically.
 - c Click **Next**.

Quest File Log Monitor and Windows Event Log Monitor agents each come with a different set of agent properties, so the contents of this page are different, reflecting the type of the agent instance that you are creating.

File Log Monitor Agent properties



Windows Event Log Monitor Agent properties

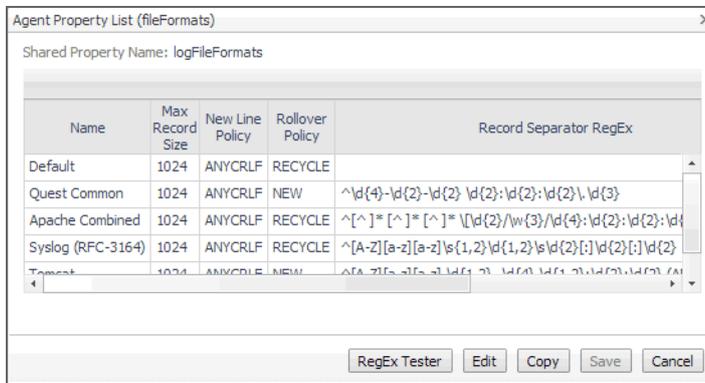


- 4 Configure the agent properties, as required.

For more information about specific agent properties, see the following:

- [Monitored Hosts](#)
 - [File Formats](#)
 - [Log Files](#)
 - [Data Collection Scheduler](#)
- File Log Monitor Agent properties**
- [Monitored Hosts](#)
 - [Event Logs](#)
 - [Data Collection Scheduler](#)
- Windows Event Log Monitor Agent properties**

- a To view or edit a property, click **View** on the right of the list property name.



In the dialog box that appears, click **Edit**, and make changes to the existing entries, or add new ones, as required.

- 5 Review the information on the **Credential Verification** page, and make any changes, if required.
 - If the page indicates that you have a valid set of credentials in place, you do not need to make any changes. This is typically the case if the hosts whose files you are about to start monitoring are already monitored by other Foglight for Infrastructure agents (a WindowsAgent or a UnixAgent).
 - If the page indicates that the Agent Manager does not have the credentials needed to access the monitored hosts, click **Manage Credentials**, and create a new credential, as required. For more information, see “Controlling System Access with Credentials” in the *Administration and Configuration Help*.

Click **Next**.

The **Summary** page reflects the newly configured settings, including agent properties. Because File Log Monitor and Windows Event Log Monitor agents each come with a different set of agent properties, the contents of this page are different, reflecting the type of the agent instance that you are creating.

- 6 Review the information on the **Summary** page, and click **Finish** to start collecting data.

Investigating log records

The Log Monitor dashboard displays the amount of log records over the monitored time range, lists the monitored log records, along with individual record details, such as its source, message, and severity, among others. Use this dashboard on a daily basis to review the monitored log records and identify potential signs that can lead to performance bottlenecks. For example, a significant increase in the number of log records can help you predict and prevent potential system-level issues.

To access this dashboard, from the navigation panel, choose **Dashboards > Log Monitor**.

The screenshot displays the Foglight Log Monitor dashboard. At the top, it shows the current time as Tuesday, November 4, 2014 12:20 PM and indicates 1 alarm. The dashboard is for host 'torrdv466' and includes monitors for 'File Log Monitor' and 'Windows Event Log Monitor'. A 'File Selector' section allows filtering by directory ('All') and lists event log files like 'Application' and 'System'. A 'New Records Found' chart shows a single data point at 12:25 with a count of 1. Below the chart is a table of log records:

TimeStamp	Tag	Severity	Log Record
11/04/14 03:21:02 PM	n/a	Informational	Microsoft Security Client successfully applied security policy: "Antimalware Policy".

The 'Record Details' section provides further information for the selected record:

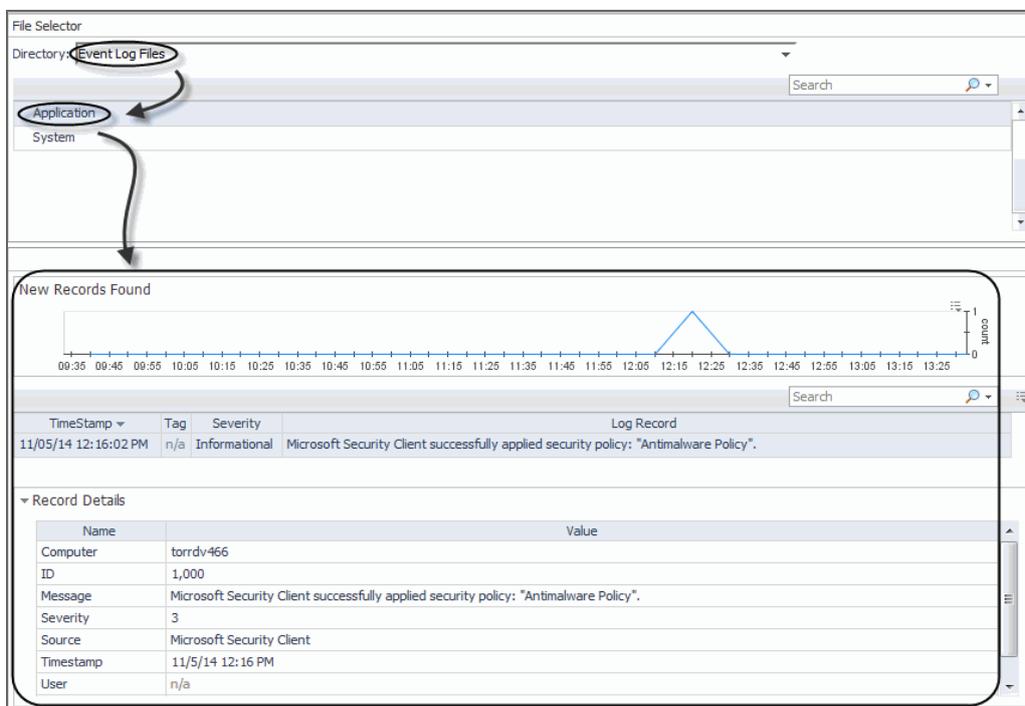
Name	Value
Computer	torrdv466
ID	1,000
Message	Microsoft Security Client successfully applied security policy: "Antimalware Policy".
Severity	3
Source	Microsoft Security Client
Timestamp	11/4/14 3:21 PM
User	n/a

Start by choosing the host containing the log files whose records you want to review, using the **File Selector** view. From here, for file logs, you can select the directory and the log file name.



For Windows Event Logs, select **Event Log Files**, and choose the name of the Windows Event Log (for example, **Application**).

NOTE: Log names are specified in the Windows Event Log Monitor agent properties. For more information, see [Event Logs](#).



For complete information about the data appearing in these views, see [Foglight Log Monitor views](#).

Configuring agent properties

The File Log Monitor and Windows Event Log Monitor agents collect data from log files and send it to the Management Server.

The monitoring agents look for specific records given a text pattern in the monitored log files. They support regular expressions, which allows you to search for desired text patterns. The agents support PCRE (Perl Compatible Regular Expressions). For details about the PCRE syntax, visit <http://perldoc.perl.org/perlre.html>.

To monitor log records on a host, an agent establishes an SSH connection to Unix hosts, or a WMI or WinRM connection to Windows® hosts. Once this connection is in place, an agent executable is uploaded to the remote host (into %TEMP% on Windows and /tmp on UNIX®). The name of the executable name is unique to the agent instance. The executable then starts locating the monitored log files on the host and extracts the desired records, as specified in the agent properties. The agent sends back the extracted records to the agent for processing. The agent keeps track of the current location in all of the scanned log files so that the same information is only scanned once.

The monitoring agents use the same set of credentials as the other Foglight for Infrastructure agents, and the same credential purposes. For more information, see the [Using Foglight for Infrastructure](#).

When an agent connects to Foglight, it is provided with sets of properties that it uses to configure its correct running state. The File Log Monitor and Windows Event Log Monitor agents are provided with list properties. List properties can be shared amongst multiple agent instances. Any changes you make to a list property affects all agent instances that are associated with that list.

Default versions of these properties are installed with Foglight for Infrastructure, and configured for your agent instances when you run the configuration wizards. If you need to make changes to any list properties after configuring your monitoring agents, you can do so (see [Configuring File Log Monitor agent properties](#) and [Configuring Windows Event Log Monitor agent properties](#)). However, keep in mind that any changes you make to a list property can affect other agent instances. For more information about LogMonitor remote monitoring, see [Configuring connections to remote Windows platforms](#).

For complete information about working with agent properties, see the *Administration and Configuration Help*.

To configure agent properties:

- 1 Log in to the Foglight browser interface.
- 2 Ensure that the navigation panel is open.
To open the navigation panel, click the right-facing arrow on the left.
- 3 Open the Agent Status dashboard and navigate to the agent properties.
 - a On the navigation panel, under Dashboards, select **Administration > Agents > Agent Status**.
 - i** **IMPORTANT:** Another way of editing agent properties is through the Agent Properties dashboard. The properties you specify on this dashboard apply to all instances of the selected agent type, excluding any of the existing agents. To be certain that you are editing the properties of a particular agent instance and prevent overwriting of properties of another agent instance, use the Agent Status dashboard instead of the Agent Properties dashboard.
 - b On the Agent Status dashboard, select the instance of the File Log Monitor or Windows Event Log Monitor agent whose properties you want to modify and click **Edit Properties**.
 - c Indicate that you want to edit the properties of the selected agent instance.
A list of agent properties appears in the display area.

Configuring File Log Monitor agent properties

The File Log Monitor Agent collects information from selected text-based log files. A log file consists of one or more records; a record can span multiple lines, depending on the format of the log files.

The agent includes a set of properties that you can use to define the location, name, and structure of the log files it monitors. It includes the following groups of agent properties:

- [Monitored Hosts](#)
- [Log Files](#)
- [File Formats](#)
- [Record Transformations](#)
- [Data Collection Scheduler](#)

For a configuration example, see [FileLogMonitor configuration example](#).

Monitored Hosts

The **Monitored Host** properties specify the hosts whose log files you want to monitor with this agent.



- **Hosts:** A list specifying the hosts monitored by the agent instance. Typically you want a cloned list that is associated with a specific agent instance. Each entry in the list includes the following columns:
 - **Host:** The name of the monitored host or its IP address.
 - **Host name override:** The host name under which this host's data is stored in the data model. This property is optional.
 - **Host Type:** Windows or Unix. This property determines how the agent connects to the host: using SSH (Unix hosts), or using WMI or WinRM (Windows hosts).
 - **SSH Port:** The port number used for secure connections, if applicable. For Unix and Linux hosts, this value is typically set to 22. For Windows hosts, this is not applicable, and -1 should be specified (meaning not applicable). This property is optional.
 - **Operation Timeout:** The maximum amount of time in seconds given to the agent for each phase of a collection attempt. This includes uploading the native executable, scanning for log entries, and retrieving log content.
 - **Collect System ID:** This property indicates to the agent whether or not to collect a unique system ID from this system. This is not desirable when monitoring Hyper-V systems, as some Hyper-V systems use the same ID for multiple systems, preventing them from being unique.
 - **Remote Collector Executable:** The name of the agent native executable on the remote monitored host. This property is optional. If not specified, a random name is used. Configure this property only if you need to set a specific name for the executable so that you can write a sudo rule for it, or to have it uploaded to a non-default directory. In that case, provide a complete a full path name along with the file name.
 - **TIP:** By default, the executable is created on the monitored host in the `%TEMP%` directory (Windows) or `/tmp` (Unix).
 - **Secure Launcher:** The name and path to the sudo that enables the agent to launch on Unix and Linux machines, for example: `/usr/bin/sudo`. This property is optional.

Log Files

The **Log Files** properties allow you to specify the monitored log files on each host the agent instance connects to, and the type of log records that you want to scan.

Log Files				
Log Files	logFilesList ▼	Edit	Clone	Delete
Patterns	recordPatternsList ▼	Edit	Clone	Delete
Exclude Records	excludedRecordPatternsList ▼	Edit	Clone	Delete
Hours of Backlog Records to Read	<input type="text" value="1"/>			

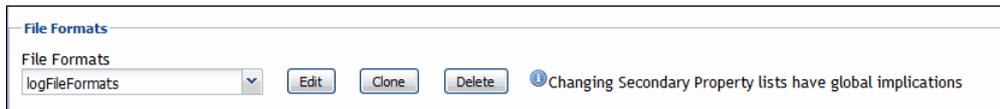
- **Log Files:** A list specifying the log files monitored by this agent. If the list is shared between agent instances, or if the agent instance is configured to connect to multiple hosts, the log file locations specified in this list are checked on every host the agent connects to. This is useful in situations when you want to scan a standard log file, for example, `/var/log/messages`, across multiple hosts. To do that, create one agent instance with its own *Hosts* list, (see [Monitored Hosts](#)), and a single row in this list.

Each entry in the list includes the following columns:

- **Directory:** The directory containing the log files that you want to monitor.
- **Filename Pattern:** A regular expression that specifies which log files to monitor.
 - **TIP:** The agent supports PCRE (Perl Compatible Regular Expressions). For details about the PCRE syntax, visit <http://perldoc.perl.org/perlre.html>.
- **File Format Name:** The name of the file format the log file uses. File format definitions are specified in the [File Formats](#) properties. The value you provide in this column must match an existing file format.
- **Patterns:**
 - **RegEx Match Pattern:** A regular expression that the agent uses to look for specific text in the monitored log files.
 - **TIP:** The agent supports PCRE (Perl Compatible Regular Expressions). For details about the PCRE syntax, visit <http://perldoc.perl.org/perlre.html>.
 - **Match Severity:** The severity associated with log records that match the specified regular expression, in the monitored log file. There are five available severities that you can choose from: `Warning`, `Critical`, `Fatal`, `Debug`, and `Informational`.
 - **NOTE:** The `Critical` severity is assigned to any log record with the `Error` severity.
 - **Tags:** One or more comma-separated tags that you want to add to log records that match the specified regular expression, in the monitored log file. This property is optional. Tags are useful because they can help you quickly locate records with a desired tag. If set, tags are reported along with any record that matches the specified regular expression. For example, the tag `security, auth` can be applied to any records that match the regular expression `.*login failed.*`. This allows the agent to identify all records (regardless of file name, host, agent or content) that relate to either security or authorization, and to display them on the Log Monitor dashboard.
- **Exclude Records:** Enable/disable to trigger an alarm for agents that are specified in the pattern.
 - **RegEx Record Exclude Pattern:** The regular expression of the specific text that the agent uses to exclude records from the monitored log files.
 - **RegEx File Exclude Pattern:** The regular expression of the specific filename or directory that the agent uses to exclude records from the monitored log files.
 - **Exclude Enable:** Sets to true to enable excluding records from the log files, and vice versa.
- **Hours of Backlog Records to Read:** Specifies the time duration of scanning the monitored log files.

File Formats

The **File Formats** properties allow you to specify the format of the log files that you want to monitor.



- **File Formats:** A list describing the structure of contents in the monitored log files. Any file that you monitor must have its format specified in this list. In most cases, all agent instances refer to one global list of file formats. The global list is pre-populated to describe some common log file formats.

Each entry in the list includes the following columns:

- **Name:** The name of the file format.
- **Max Record Size:** The maximum length of a record entry in bytes to use for pattern matching and submission to the Management Server. If a record is larger in size, it is truncated. The pattern is still applied to the entire record when searching for a match. This property is optional, and defaults to 1024 if not specified.
- **New Line Policy:** The character or sequence of characters used to signify the end of a text line (EOL) in the file format. The following values are available:

CR: The carriage return character (`\r`), ASCII code `0x0d`. This is common on Mac OS systems up to version 9.

LF: The line feed character (`\n`), ASCII code `0x0a`. This is common on Unix and Unix-like systems (including Mac OS X systems).

CRLF: A carriage return character (`\r`, ASCII code `0x0d`), followed by a line feed (`\n`, ASCII code `0x0a`). This is common on Windows systems.

ANYCRLF: A carriage return character (`\r`, ASCII code `0x0d`), a line feed (`\n`, ASCII code `0x0a`), or both. This is common on mixed platform log files.

ANY: Any Unicode new line sequence, including **CR**, **LF**, **CRLF**, plus the single characters VT (vertical tab, `U+000B`), FF (form feed, `U+000C`), NEL (next line, `U+0085`), LS (line separator, `U+2028`), and PS (paragraph separator, `U+2029`).

- **Rollover Policy:** Indicates to the agent the way the contents of the log files using this format are rotated when they reach their maximum size.
 - RECYCLE:** The name of the active log file remains the same, while older logs are renamed by appending a '1' to the name and increasing that number each time a new log file is created. For example, your system always writes log records to the same file, *abc.log*. When that file reaches its maximum size, the system renames it to *abc1.log*, and creates a new *abc.log* file for storing new logs. When that *abc.log* file becomes full, the system renames it to *abc.log2* file, and creates a new *abc.log* file, and so on.
 - NEW:** Each time a new log file is created, the number in the file name of the active log is increased by '1'. For example, your system starts writing log records to *abc.log*. When that file reaches its maximum size, the system creates a new log file, *abc1.log*, and continues to write log records to that file. When *abc1.log* file becomes full, the system creates an *abc2.log* file, and so on.
- **Record Separator RegEx:** A regular expression indicating when a log record ends and a new one starts. This property is optional. The default is an empty string which indicates that each record is a single line.
 - **TIP:** The agent supports PCRE (Perl Compatible Regular Expressions). For details about the PCRE syntax, visit <http://perldoc.perl.org/perlre.html>
- **Character Set:** The character encoding used in the log file using this format. The character set must be installed on the remote host, and available through the `iconv` utility. The native character set is translated to UTF-8 when a log record is reported. This property is optional. If not specified, it defaults to `UTF-8`.
- **Maximum Match Count:** The maximum number of records the agent can read during one collection interval. If the agent reaches the number of records before the end of the file, it continues to read the remainder of the file during the next collection interval, and so on.

Setting this value to a reasonable number of records (for example, 200) allows you to control the amount of time and resources the agent spends to read monitored logs during a single collection interval, and to prevent bottlenecks during data collection. If you do not want to specify a limit, type -1.

- **Max Processing Time(s):** The amount time in seconds the agent can spend on reading all log files in one collection cycle while being connected to the remote host.

This value should be equal to or less than the **Operation Timeout** value in the **Hosts** list. For more information, see [Monitored Hosts](#).

Setting this value to a reasonable number of seconds allows you to control the amount of time the agent spends to read monitored logs during a single collection interval, and to prevent bottlenecks during data collection. Once surpassed, the collector stops, and continues from that same point in the next collection cycle.

Record Transformations

The **Record Transformations** properties allow you to transform any log message before it is sent to the Management Server. This could be used to add extra information or to remove sensitive information from a log record.



- **Record Transformations:** A list of record transformations that the agents must use in conjunction with the match patterns to convert any log messages. When no transformation is specified, the log record is transmitted to the Management Server without changes.

Each entry in the list includes the following columns:

- **RegEx Record Transformation Pattern:** A regular expression that the agent uses to look for specific text in the collected log record.
- **Record Transformation:** The replacement text that the agent uses in the log record to be transmitted to the Management Server.

Data Collection Scheduler

The **Datacenter Collection Scheduler** agent properties specify the data frequency settings the agent uses to read monitored log files.



- **Collector Config:** A list containing the data collectors the agent uses. Each entry in the list includes the following columns:
 - **Collector Name:** The name of the collector the agent uses to gather data.
 - **Default Collection Interval:** The number of milliseconds, seconds, minutes, hours, or days during which the agent collects data.
 - **Time Unit:** The time unit associated with the **Default Collection Interval**.
 - **Fast-Mode Collection Interval:** The number of milliseconds, seconds, minutes, hours, or days during which the agent collects data when working in the fast collection mode.
 - **Fast-Mode Time Unit:** The time unit associated with the **Fast-Mode Collection Interval**.
 - **Fast-Mode Max Count:** The maximum number of the times the agent can stay in fast collection mode.

FileLogMonitor configuration example

Scan example

This example provides the configuration settings for monitoring the FglAM log files on a UNIX® system for `WARN` and `ERROR` records. The FglAM log files are located in the `/home/user/FglAM/state/default/logs` folder. FglAM log records have a date at the beginning of each record that look like this:

```
2014-11-17 13:37:22.315 ...
```

This format can be set as the regular expression for the record separator.

Table 1. Configuration settings for the Scan example

Group Name	Property Name	Property Details	Value set for this example		
Monitored Hosts	Hosts	Host	host.domain.com		
		Host name override	(optional)		
		Host Type	UNIX		
		SSH Port	22		
		Operation Timeout	60000		
		Collect System ID			
		Remote Collector Executable	(optional)		
		Secure Launcher	(optional)		
		Log Files	Log Files	Directory	/home/user/FglAM/state/default/logs
				Filename Pattern	FglAM_*.log\$
File Format Name					
Patterns	RegEx Match Patterns		WARN ERROR		
	Match Severity				
	Tags				
	Exclude Records		RegEx Record Exclude Pattern WARN ERROR		
	RegEx File Exclude Pattern		C:\temp.log D:\temp.txt C:\apache\logs\FglAM*.log Test*.log		
Exclude Enable					
File Formats	File Formats	Name			
		Max Record Size	1024		
		New Line Policy	ANYCRLF		
		Rollover Policy	NEW		
		Record Separator RegEx	^\d{4}-\d{2}-\d{2} \d{2}:\d{2}:\d{2}\.\d{3}		
		Character Set	UTF-8		
		Maximum Match Count	200		
		Max Processing Time(s)	120		
Record Transformations	Record Transformations	RegEx Record Transformation Pattern	(optional)		
		Record Transformation	(optional)		
Data Collection Scheduler	Collector Config	Collector Name	(default)		

Table 1. Configuration settings for the Scan example

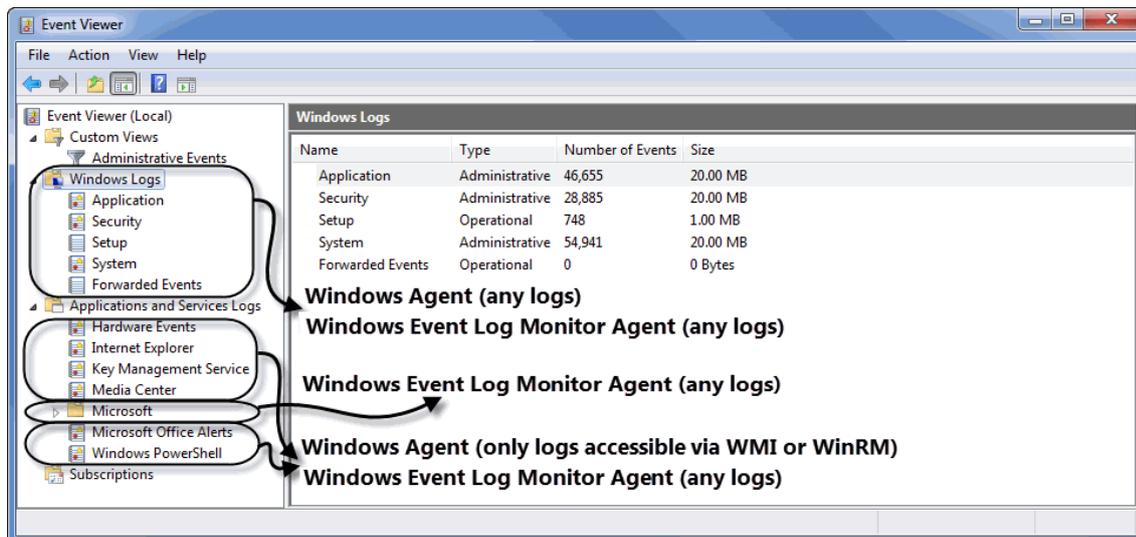
Group Name	Property Name	Property Details	Value set for this example
		Default Collection Interval	(default)
		Time Unit	(default)
		Fast-Mode Collection Interval	(default)
		Fast-Mode Time Unit	(default)
		Fast-Mode Max Count	(default)

This example only shows one scan, but the scan can be performed multiple times at regular intervals since more records can be added to the log files over time.

Configuring Windows Event Log Monitor agent properties

The Windows Event Log Monitor Agent collects information from Windows Event Log files. This agent expands the functionality of the Windows Agent, also included with Foglight for Infrastructure. However, while the Windows Agent can collect information only from Windows Logs and those Application and Service Logs that are accessible through WMI or WinRM, the Windows Event Log Monitor Agent can monitor both Windows Logs and any Application and Service Logs, including the Microsoft Application and Services Logs, available on some newer Microsoft Windows versions.

For more information about the Windows Agent, see the *Managing Infrastructure User and Reference Guide*.



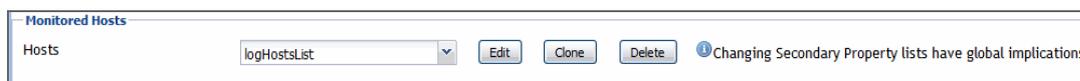
This agent includes the following groups of agent properties:

- [Monitored Hosts](#)
- [Event Logs](#)
- [Record Transformations](#)
- [Data Collection Scheduler](#)

For a configuration example, see [WindowsEventLogMonitor configuration example](#).

Monitored Hosts

The **Monitored Host** properties specify the hosts whose log files you want to monitor with this the agent.



- **Hosts:** A list specifying the hosts monitored by the agent instance. Typically you want a cloned list that is associated with a specific agent instance. Each entry in the list includes the following columns:
 - **Host:** The name of the monitored host or its IP address.
 - **Host name override:** The host name under which this host's data is stored in the data model. This property is optional.
 - **Network Operation Timeout (seconds):** The maximum amount of time in seconds given to the agent for each phase of a collection attempt. This includes uploading the native executable, scanning for log entries, and retrieving log content.
 - **Collect System ID:** This property indicates to the agent whether or not to collect a unique system ID from this system. This is not desirable when monitoring Hyper-V systems, as some Hyper-V systems use the same ID for multiple systems, preventing them from being unique.
 - **Remote Collector Executable:** The name of the agent native executable on the remote monitored host. This property is optional. If not specified, a random name is used. Configure this property only if you need to set a specific name for the executable so that you can write a sudo rule for it, or to have it uploaded to a non-default directory. In that case, provide a complete a full path name along with the file name.

TIP: By default, the executable is created on the monitored host in the %TEMP% directory.

- **Maximum Record Match Count Per Log File:** The maximum number of records the agent reads per log file. Setting this value to a reasonable number of records (for example, 200) allows you to control the amount of time and resources the agent spends to read monitored logs during a single collection interval, and to prevent bottlenecks during data collection. If you do not want to specify a limit, type -1.
- **Backlog of Events (seconds):** The length of time in the past to start collecting data from until the present moment, if not already processed. Specifying a reasonable amount of time using this property (for example, 3,600 seconds or one hour) allows you to bring in historical data, providing a point of reference for future collections.
- **Max Logs Processing Time (seconds):** The amount of time in seconds given to the agent for a data collection attempt. Setting this value to a reasonable number of seconds (for example, 120) allows you to control the amount of time the agent spends to read monitored logs during a single collection interval, and to prevent bottlenecks during data collection.

This value should be equal to or less than the **Network Operation Timeout (seconds)** value in the **Hosts** list.

Event Logs

The **Event Logs** properties allow you to specify the Windows Event Logs and the type of records you want to monitor.



- **Event Logs to Monitor:** A list specifying the types of Windows Event Logs monitored by this agent. Each entry in the list includes the following column:

- **Event Log Name:** The name of the Windows Event Log whose files you want to monitor with this agent. This value must be identical to the **Full Name** value, seen in the **Log Properties** dialog box. To find out an event log name, in the Event Viewer, right-click a Windows Log, choose **Properties** from the shortcut menu, and look for the **Full Name** property in the dialog box that appears.
- **Event Log Filters:** A list specifying the types of Windows Event Log entries that you want to monitor with this agent. Using this list you can configure the agent to include and exclude specific entries from its collections using `Include` and `Exclude` commands, as required. By default, the agent does not collect any data unless configured to do so. For example, instructing the agent to `exclude` certain log types from its collections, without specifying which types of log items to `include`, prevents the agent from collecting any data from Windows Event Logs.

i | **NOTE:** For log filters that signify “any” (that is, you do not want to filter on this property), the **User**, **Source**, **Category**, and **Event Description** properties should all have a value of “*”, and the **EventID** property should have a value of “0”.

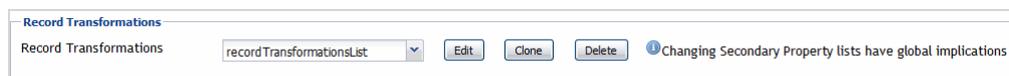
Each entry in the list includes the following columns:

- **Include/Exclude:** Instructs the agent to include or exclude certain types of logs from its collections.
 - **Type:** The Windows severity level: * (All) , Information, Warning, Error, Success Audit, or Failure Audit. For more information about these levels, see your Windows documentation.
 - **User:** The name of the user who initiated the Windows Event.
 - **Source:** The name of the Windows Event Log to which the event is published.
 - **Category:** The category of logs to `include` or `exclude` from agent collections.
 - **EventID:** The Windows Event ID. For more information about Windows Event IDs, see your Windows documentation.
 - **Event Description:** The description of the Windows event.
 - **Tags:** The tag associated with the record, as specified in the agent properties.
 - **Event Throttle Count:** When set, this ensures that one event in every count (the event log entry occurrence that the filter applies to) is submitted to the Management Server. If the count is one or less, then every event log entry occurrence is submitted and no throttling is done. The default is zero '0'.
 - **Event Throttle Duration (seconds):** This value represents the duration in seconds for the throttle count to be applied. When set, the throttle count is applied within a duration. After the duration expires, the throttling restarts from the beginning regardless of the current throttle state. If the count is one or less, then only one event log entry the filter matches is submitted within the specified duration. If the count is larger than one, then only one in every count (the event log entry occurrence that the filter matches) is submitted, and the agent starts counting pattern matches from zero after the duration. The default value is zero '0', which means the duration is not applied.
- **i** | **NOTE:** The **Event Throttle Count** and **Event Throttle Duration (seconds)** properties only apply to `INCLUDE`-type filters (and not `EXCLUDE`-type filters), since throttling is necessary only if a message is being included (and submitted).
- **Event Log Severity:** A mapping that specifies how Windows Event Log severities relate to Foglight severity levels. A default agent installation includes a complete mapping. You can make changes to it when configuring Windows Event Log Monitor Agent instances, if required.

Event Log Severity	Foglight Severity
Error	Critical
Warning	Warning
Information	Informational
Success Audit	Informational
Failure Audit	Warning

Record Transformations

The **Record Transformations** properties allow you to transform any log message before it is sent to the Management Server. This could be used to add extra information or to remove sensitive information from a log record.



- **Record Transformations:** A list of record transformations that the agents must use in conjunction with the match patterns to convert any log messages. When no transformation is specified, the log record is transmitted to the Management Server without changes.

Each entry in the list includes the following columns:

- **RegEx Record Transformation Pattern:** A regular expression that the agent uses to look for specific text in the collected log record.
- **Record Transformation:** The replacement text that the agent uses in the log record to be transmitted to the Management Server.

Data Collection Scheduler

The **Datacenter Collection Scheduler** agent properties specify the data frequency settings the agent uses to read monitored log files.



- **Collector Config:** A list containing the data collectors the agent uses. Each entry in the list includes the following columns:
 - **Collector Name:** The name of the collector the agent uses to gather data.
 - **Default Collection Interval:** The number of milliseconds, seconds, minutes, hours, or days during which the agent collects data.
 - **Time Unit:** The time unit associated with the **Default Collection Interval**.
 - **Fast-Mode Collection Interval:** The number of milliseconds, seconds, minutes, hours, or days during which the agent collects data when working in the fast collection mode.
 - **Fast-Mode Time Unit:** The time unit associated with the **Fast-Mode Collection Interval**.
 - **Fast-Mode Max Count:** The maximum number of the times the agent can stay in fast collection mode.

WindowsEventLogMonitor configuration example

Scan example

This example provides the configuration settings for monitoring the “System” and “Application” Windows® event log files. Any records with a source value of `Perflib` are excluded from the monitoring, and only records that are of type `Warning` are included in the monitoring.

Table 2. Configuration settings for the Scan example

Group Name	Property Name	Property Details	Value set for this example
Monitored Hosts	Hosts	Host	host.domain.com
		Host name override	(optional)

Table 2. Configuration settings for the Scan example

Group Name	Property Name	Property Details	Value set for this example	
		Network Operation Timeout (seconds)	120	
		Collect System ID		
		Remote Collector Executable	(optional)	
		Maximum Record Match Count Per Log File	200	
		Backlog of Events (seconds)	3600	
		Max Logs Processing Time (seconds)	120	
Event Logs	Event Logs to Monitor	Event Log Name	System Application	
		Event Log Filters	Include/Exclude INCLUDE Type WARNING User * Source * Category * EventID 0 Event Description * Tags Event Throttle Count Event Throttle Duration (seconds)	
		Event Log Severity		
	Record Transformations	Record Transformations	RegEx Record Transformation Pattern	(optional)
			Record Transformation	(optional)
	Data Collection Scheduler	Collector Config	Collector Name	(default)
			Default Collection Interval	(default)
			Time Unit	(default)
			Fast-Mode Collection Interval	(default)
			Fast-Mode Time Unit	(default)
			Fast-Mode Max Count	(default)

This example only shows one scan, but the scan can be performed multiple times at regular intervals since more records can be added to the log files over time.

Configuring connections to remote Windows platforms

Foglight Log Monitor command shell types

Foglight Log Monitor requires that a Windows® command shell connection be established to execute Windows commands on remote machines. There are two types of command shell connections that can be established to

execute remote commands: *WinRMCommandShell* and *DCOMWindowsCommandShell*. You need to setup the remote machine based on the type of command shell connection you need to establish.

To execute Windows commands on a local machine, a *LocalWindowsCommandShell* may be used, if local user credentials are provided.

The Foglight Log Monitor command shells are described in the following sections.

WinRMCommandShell

Uses Windows Remote Management (WinRM) to execute remote commands. For configuration information, see section “Configuring Windows Remote Management (WinRM)” in the *Foglight Agent Manager Guide*.

i | **NOTE:** *WinRmCommandShell* connections are attempted before *DCOMWindowsCommandShell*.

DCOMWindowsCommandShell

This command shell type executes commands remotely using Windows Management Instrumentation (WMI). WinShell must be setup as well.

For configuration information, see sections “Configuring Windows Management Instrumentation (WMI)” and “Configuring Registry Settings for WinShell Access through DCOM” in the *Foglight Agent Manager Guide*.

LocalWindowsCommandShell

This command shell type is for local command execution. No setup is required for executing commands on a local machine.

Foglight Log Monitor remote command access permissions

The Foglight for Infrastructure *WindowsAgent* can use the WMI mechanism to establish remote connections for monitoring Windows resources. In this case it can collect data only from specific event logs, but not all (for details, see [About the WindowsAgent](#)).

To monitor event logs within the “Applications and Services” category, you must use the LogMonitor agents (*FileLogMonitorAgent* or *WindowsEventLogMonitorAgent*).

Foglight LogMonitor copies an executable to the remote machine and runs this executable, which outputs the collected data and then Foglight Agent Manager processes it. The executable uses Windows native APIs to obtain the relevant data from the Windows Event Logs. To copy and run the executable on the remote machine, access to the Windows command prompt is required. If DCOM is used, an extra setup step is required (for details, see “Configuring Registry Settings for WinShell Access through DCOM” in the *Foglight Agent Manager Guide*). There are no extra setup steps required if WinRM is used.

Foglight Log Monitor remote connection requirements

The remote monitoring of Windows® and UNIX® hosts has unique requirements, as presented in the *Foglight Agent Manager Guide*. For example, the following log entry indicates that the Remote Connection failed.

```
2015-06-02 11:05:44.286 ECHO <HostAgents/5.7.2/FileLogMonitorAgent/LogMonitor-IIRWin_Webserver-agent> WARN [Quartz[0]-1228]
com.quest.foglight.infrastructure.actions.logmonitor.file.FileLogScanAction - Could
```

not execute data collection commands for File Log Scan Action
[Host=host.example.com, HostType=WINDOWS, Directory=D:\Program
Files(x86)\FglAM\state\default\logs, Filename=temp.log]. It will be skipped in this
collection period.

com.quest.glue.api.services.RemoteConnectionException: a shell connection could not
be established

Monitoring IBM PowerVM environments

Foglight™ for PowerVM allows you to monitor IBM® PowerVM® environments. Foglight alerts you about infrastructure problems when they develop, enabling you to resolve issues pro actively before end users are affected. Early intervention ensures consistent application performance at established service levels. Foglight for PowerVM monitors the health of your virtual system by tracking the levels of resource utilization such as processor, network, and memory consumption of individual objects in your integrated environment.

- [Before you begin](#)
- [Managing PowerVM HMC agents](#)
- [Monitoring your PowerVM environment](#)

Before you begin

Ensure that Foglight for Infrastructure is installed on the Management Server. For installation instructions, see the *Foglight for Infrastructure Release Notes*.

- To monitor PowerVM® servers, you need a running instance of the PowerVM HMC Agent. This agent is provided with Foglight for Infrastructure.
- Collecting OS-level data from hosts in your PowerVM environment requires running instances of Foglight for Infrastructure Windows or Unix agents. These agents are also provided with Foglight for Infrastructure.

Introducing the PowerVM infrastructure

When you deploy Foglight™ for PowerVM and set up the monitoring agents for data collection, the Infrastructure Environment dashboard enables you to review the performance of your environment at a glance. To ensure consistent application performance, use this dashboard to find out more details about specific components, and to look for potential performance degradation indicators, such as high processor loads or low network utilization.

Foglight for PowerVM focuses on the following components in your monitored PowerVM® environment:

- *HMC* (Hardware Management Console) manages a groups of PowerVM servers that are running in your integrated environment.
- *Managed Server* is a collection of physical components managed by the PowerVM Hypervisor software layer. PowerVM Hypervisor divides physical system resources into isolated logical partitions. An HMC typically includes multiple logical partitions.
- *Partition* encapsulates an operating system and application components that are running in a PowerVM Hypervisor. It has a dedicated share of the system's physical resources.
- *VIOS* provides a mechanism that enables logical partitions to share physical resources of a managed PowerVM server.

Managing PowerVM HMC agents

Foglight™ for PowerVM relies on PowerVM HMC Agents to collect data from the monitored PowerVM® infrastructure. Foglight for Infrastructure features are provided with Foglight for Infrastructure.

When you install Foglight for Infrastructure, you can create PowerVM HMC Agents to collect performance information about your PowerVM environment. To connect to the monitored PowerVM system, the monitoring agents must have sufficient permissions. For complete information, see the following sections:

- [Configuring HMC user accounts](#)
- [Configuring PowerVM HMC agent credentials](#)
- [Creating PowerVM HMC agents](#)
- [Reviewing and editing PowerVM HMC agent properties](#)

Configuring HMC user accounts

To connect to the monitored Hardware Management Console (HMC), each PowerVM HMC Agent requires a valid user account. An HMC account requires a user name and password to connect to the HMC. You specify this information by creating a Foglight credential of the *User Name and Password* type (see [Configuring PowerVM HMC agent credentials](#)).

Each HMC user account has a role that grants access to specific parts of the HMC. Each PowerVM HMC Agent instance must be associated with a valid HMC account. To ensure that the HMC account provides access to all elements of your monitored environment, including Virtual I/O Servers (VIOS), the account must be granted a role with sufficient permissions.

To create an HMC role for the PowerVM HMC Agent:

- 1 Create a custom HMC role based on the `hmcoperator` role.
- 2 Remove any unnecessary permissions from the newly created role, as required. For example, your organization may require that such role should grant permissions for changing resources or restarting hosts.
- 3 Ensure that the *Issue virtual I/O server command* permission is granted to the newly created HMC role. This permission enables the agent to issue the `viosvr cmd` command. Removing it causes the HSCL350B error, indicating that the user account does not have appropriate permissions. For more information about this error, visit:

<http://www-01.ibm.com/support/knowledgecenter/P8ESS/p8eai/HSCL350B.htm>

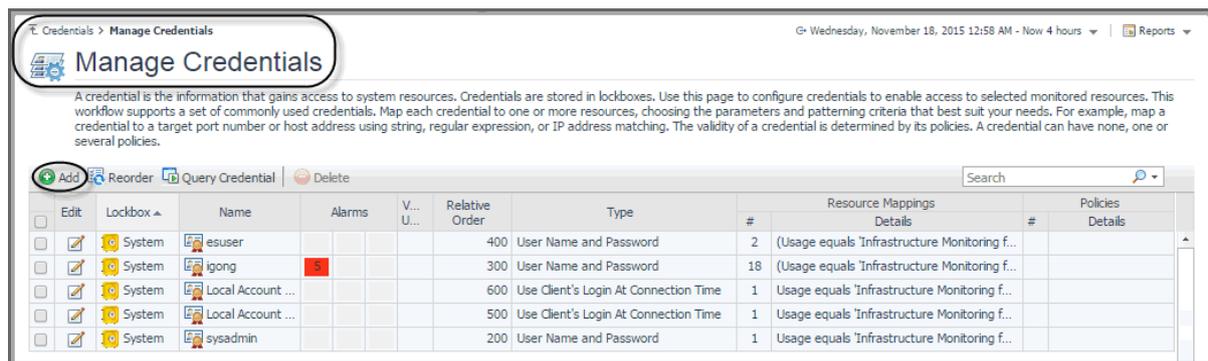
For more information about HMC roles, see your HMC documentation.

Configuring PowerVM HMC agent credentials

A valid user account is required to monitor an HMC. When you configure an HMC account (as described in [Configuring HMC user accounts](#)), you must ensure the HMC account information is provided to a Foglight credential:

You can create a credential using the Manage Credentials dashboard. To access this dashboard, on the navigation panel, under **Dashboards**, click **Credentials**. On the Credentials page that appears, click **Manage Credentials**. To create a credential using this dashboard, click **Add**.

Figure 1. Creating a credential



The credential must include the following information in order enable your PowerVM HMC Agent to access the monitored HMC environment:

- **Type:** User Name and Password.
- **User Name:** The name of the HMC user account.
- **Password:** The password associated with the HMC account.
- **Usage:** Infrastructure Monitoring for PowerVM.
- **Target Host Name** or **Target Host Name:** A mapping to the name or IP address of the host on which HMC is running.

After creating the credential, you must ensure that the lockbox in which the credential is created is released to the Foglight Agent Manager associated with the PowerVM HMC Agent instance that you are about to create (see [Creating PowerVM HMC agents](#)). You can release a credential to a desired Agent Manager using the Manage Lockboxes dashboard. To access this dashboard, on the navigation panel, under **Dashboards**, click **Credentials**. On the Credentials page, click **Manage Lockboxes**.

Figure 2. Releasing a lockbox to an Agent Manager



For complete information about managing Foglight credentials and lockboxes, see the Foglight *Administration and Configuration Help*.

Creating PowerVM HMC agents

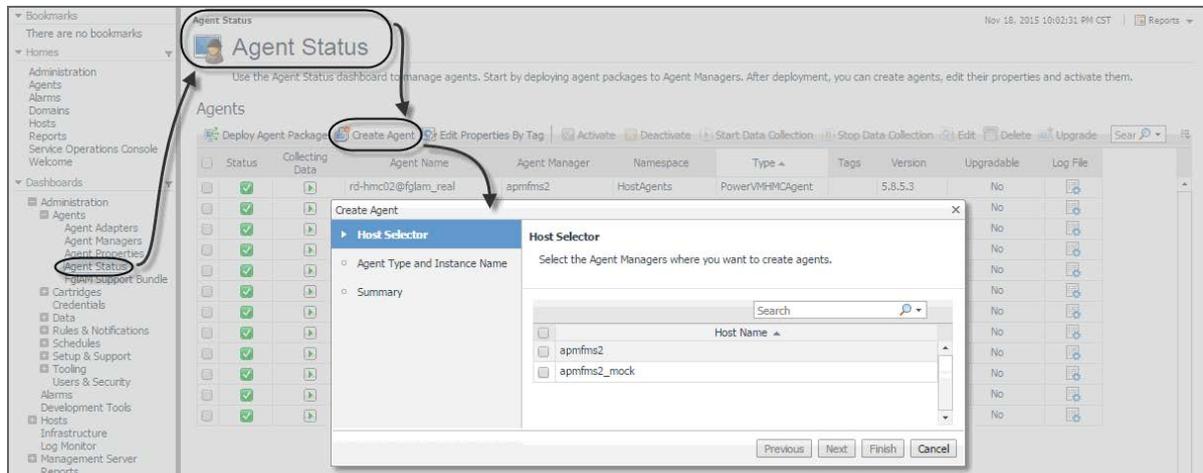
When you install Foglight for Infrastructure, you can create PowerVM HMC Agents to collect performance information about your PowerVM environment. PowerVM HMC Agents collect performance information about your PowerVM environment. Each PowerVM HMC Agent monitors an HMC instance. You can create one or more PowerVM HMC Agent instances, as required.

Before you create a PowerVM HMC Agent instance, you must ensure that Foglight for Infrastructure is installed on the Management Server, and that the Host Agents package containing the PowerVM HMC Agent is deployed to the desired Agent Manager. For more information about installing Foglight for Infrastructure on the Management Server and deploying the Host Agents package, see your Foglight for Infrastructure documentation. For general

instructions on how to install a cartridge or deploy an agent package, see the Foglight *Administration and Configuration Help*.

Create PowerVM HMC Agent instance by running the **Create Agent** wizard. This wizard is accessible from the Agent Status dashboard. To access this dashboard, on the navigation panel, under **Dashboards**, choose **Administration > Agents > Agent Status**. For complete information about the Agent Status dashboard and the **Create Agent** wizard, see the Foglight *Administration and Configuration Help*.

Figure 3. Accessing the Create Agent wizard from the Agent Status dashboard



Reviewing and editing PowerVM HMC agent properties

PowerVM HMC Agents collect data from your monitored PowerVM[®] infrastructure and send it to the Management Server. They keep track of resource utilization metrics and alerts you when applicable pre-defined thresholds are reached.

When an agent connects to Foglight[™], it is provided with sets of properties that it uses to configure its correct running state. Each agent is provided with a combination of two types of properties: agent properties and shareable properties. Default versions of these properties are installed with Foglight for PowerVM. You can review and edit PowerVM HMC Agent properties using the Agent Status dashboard.

To view and edit PowerVM HMC Agent properties:

- 1 Log in to the Foglight browser interface.
- 2 On the navigation panel, under **Dashboards**, choose **Administration > Agents > Agent Status**.
- 3 On the Agent Status dashboard, select a PowerVM HMC Agent instance whose properties you want to review or edit.

i **TIP:** Your monitoring environment typically includes a number of monitoring agent instances, such as PowerVM HMC Agents, along with agents of different types. To locate a PowerVM HMC Agent, look for `PowerVMHMCAGENT` in the **Type** column.

- 4 Click **Edit**, and from the menu that appears, choose **Edit Properties**.

The Edit Properties dashboard appears in the display area.

- 5 Click **Modify the private properties for this agent**.

The PowerVM HMC Agent properties appear enabled in the display area.

Figure 4. PowerVM HMC Agent properties

Configuring HMC Settings

The **HMC Settings** specify general settings the agent needs to connect to the monitored environment.

Figure 5. HMC Settings

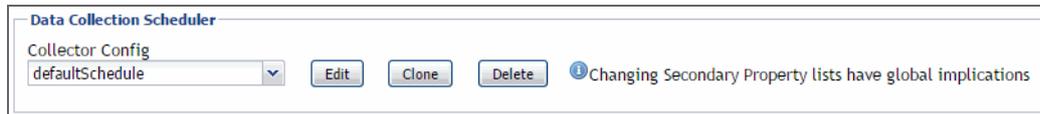
- **Primary HMC:** The host name or IP address of the primary HMC the PowerVM Agent connects to for collecting data.
- **Secondary HMC:** The host name or IP address of the secondary HMC the PowerVM Agent connects to for collecting data when the primary HMC is not connectable
- **Collect All Virtual I/O Servers:** Indicates if the agent collects information about all monitored PowerVM virtual I/O servers (**True**), or only about those listed under Specified Virtual I/O Servers (**False**).
- **Only collect the specified Virtual I/O Servers:** Indicates if the agent only collects information about PowerVM virtual I/O servers listed under Specified Virtual I/O Servers (**True**), or about all monitored PowerVM virtual I/O servers (**False**). If set to **True**, it overrides the **Collect All Virtual I/O Servers** property.
- **Specified Virtual I/O Servers:** A list indicating which PowerVM virtual I/O servers to monitor. The settings in this list take place when **Only collect the specified Virtual I/O Servers** is set to **True**.
 - **Virtual I/O Server Partition Name:** The name of the PowerVM virtual I/O server partition.
 - **Enable:** Indicates whether to monitor the listed PowerVM VIOS. If selected, the PowerVM VIOS is monitored. If not selected, the PowerVM VIOS is not monitored.
- **Collect Virtual IO Server in Multiple Threads Mode:** Indicates if the data collection runs with a single thread (**False**), or multiple threads (**True**). Running a multi-thread data collection allows the data collection process to run faster, but can cause a higher load on the PowerVM Virtual I/O Server. Since a VIOS has a

role of an I/O bus in your monitored system, selecting this option is configurable, allowing you to select a desired type of collection, without compromising your overall system performance.

Configuring Data Collection Scheduler properties

The Data Collection Scheduled properties allow you to adjust the frequency at which the PowerVM HMC Agent collects data from the monitored system.

Figure 6. Data Collection Scheduler properties



- **Collector Config:** A list identifying the data collectors the agent uses. Each entry in the list includes the following columns, allowing you to adjust the data collection settings for each individual collector:
- **Collector Name:** The name of the collector: Essential Collection, Resource Collection (Resource, Disk, Network), and Inventory Collection.
- **Default Collection Interval:** The length of the default collection interval.
- **Time Unit:** The time unit for measuring the default collection interval: milliseconds, seconds, minutes, hours, or days.
- **Fast-Mode Collection Interval:** The length of the collection interval when the agent is running in fast mode.
- **Fast-Mode Time Unit:** The time unit of the collection interval when the agent is running in fast mode.
- **Fast-Mode Max Count:** The maximum count of entries when the agent is running in fast mode.

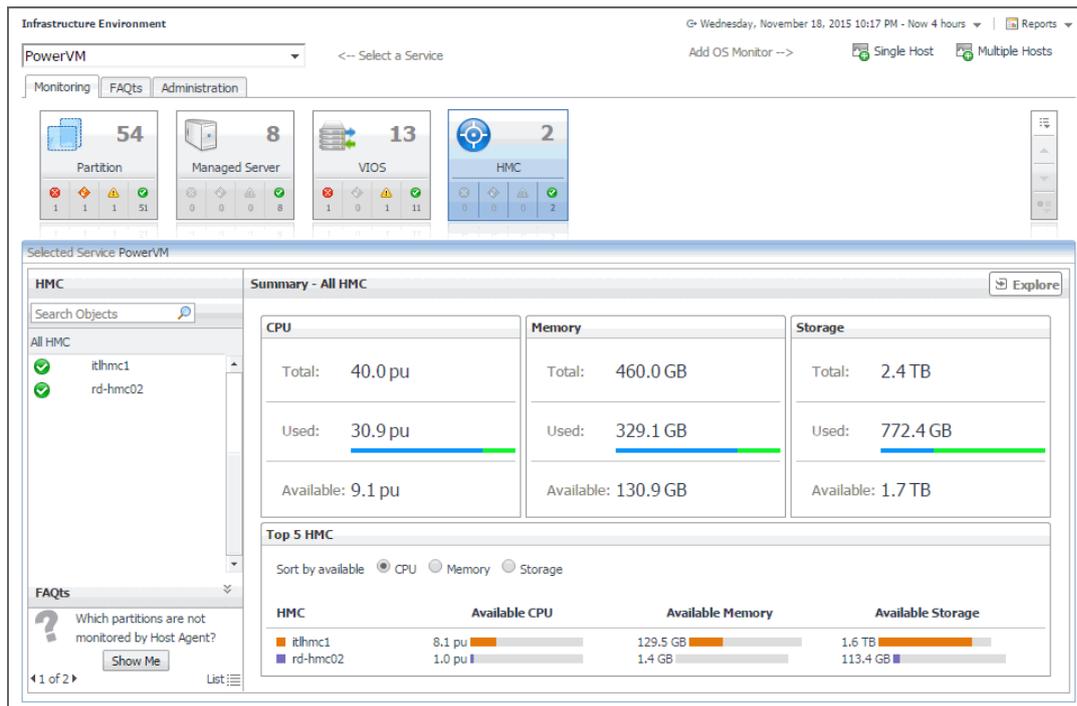
Monitoring your PowerVM environment

You can monitor the performance of PowerVM® servers, PowerVM partitions, and PowerVM Virtual I/O servers on the Monitoring tab of the Infrastructure Environment dashboard. The information appearing on this dashboard can tell you how well the selected monitored objects are performing given the current workload, and to discover potential resource-level issues, typically indicated by higher than usual processor, memory, and disk storage usage. This allows you to pro-actively reallocate resources where they are most needed, and to prevent service interruptions before they happen.

Exploring the Infrastructure Environment dashboard

A typical PowerVM® environment includes one or more HMC instances, managed PowerVM servers, PowerVM partitions, and PowerVM Virtual I/O server (VIOS) instances. You can view the state of these components on the Infrastructure Environment dashboard when you select the PowerVM service.

Figure 7. Infrastructure Environment dashboard



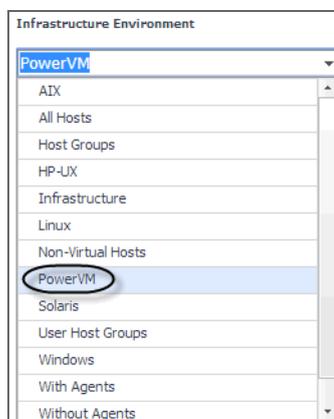
The Infrastructure Environment dashboard provides a set of tabs, each displaying a different aspect of your monitored system.

The **Monitoring** tab displays main components of your monitored PowerVM infrastructure such as HMC instances, managed servers, partitions, and VIOS instances. When you navigate to the Infrastructure Environment dashboard and select the **PowerVM** service for the first time, the **Monitoring** tab appears open.

To view PowerVM objects on the Infrastructure Environment dashboard:

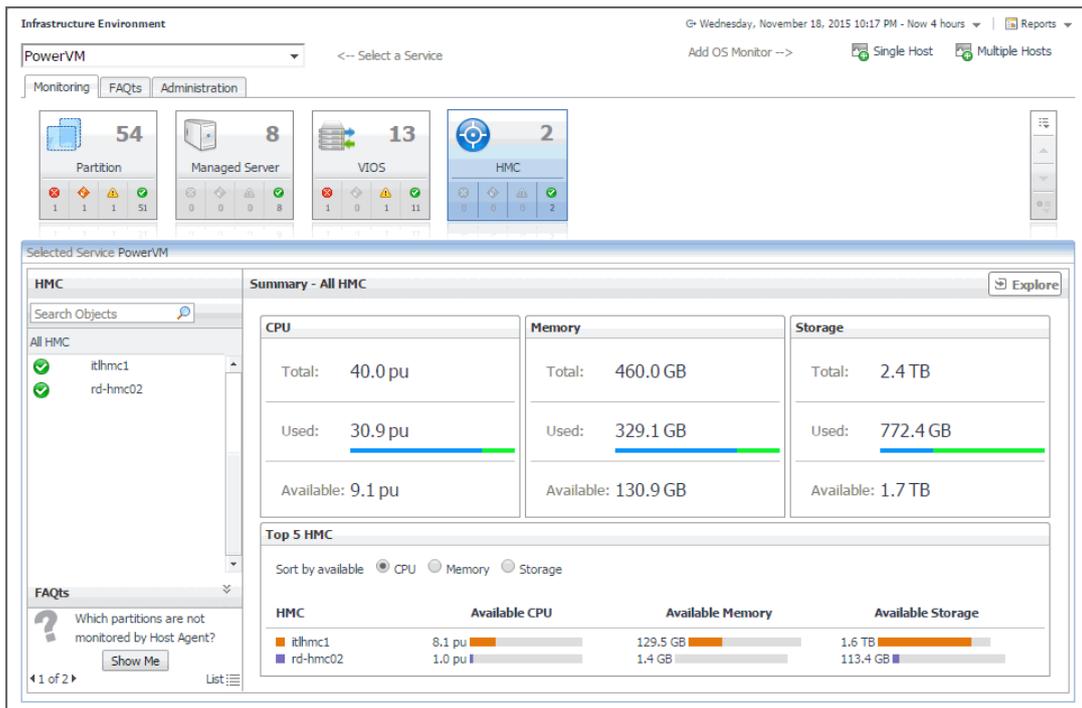
- 1 Log in to the Foglight™ browser interface.
- 2 On the navigation panel, under **Dashboards**, click **Infrastructure**.
The Infrastructure Environment dashboard appears in the display area.
- 3 On the Infrastructure Environment dashboard, in the top-left corner, click the service box, and from the list that appears, select **Power VM**.

Figure 8. Infrastructure Environment PowerVM service



The Infrastructure Environment dashboard refreshes, showing your monitored PowerVM objects.

Figure 9. Infrastructure Environment dashboard

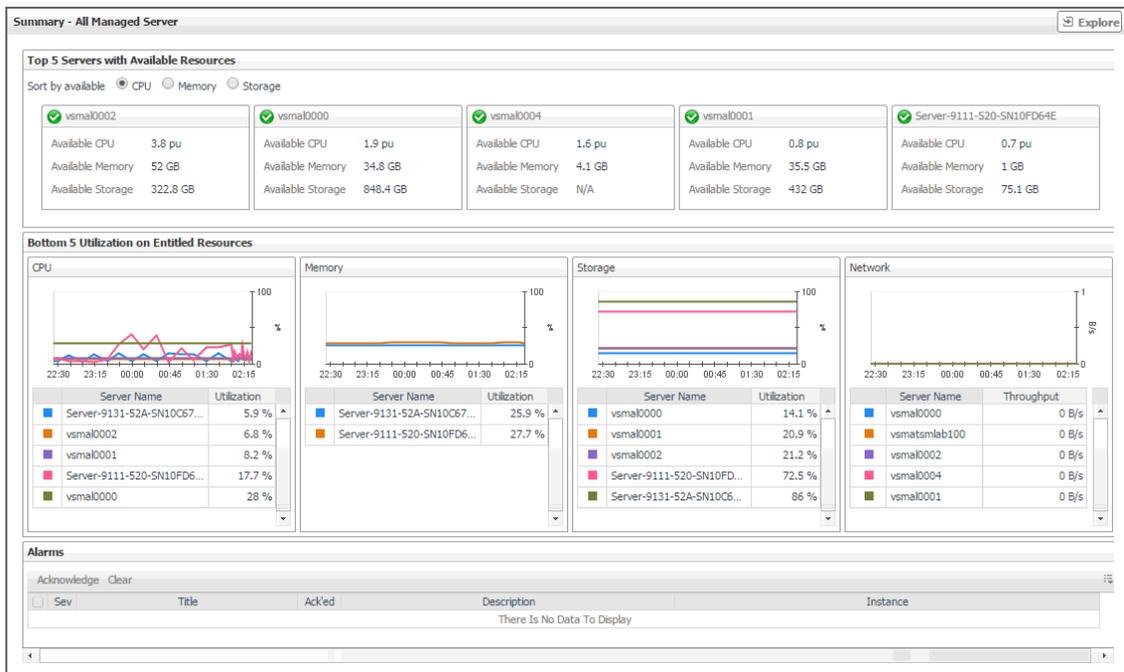


- 4 On the Infrastructure Environment dashboard, on the **Monitoring** tab, click a tile representing the object type that you want to investigate.

For example, to investigate managed servers, click the **Managed Server** tile.

The **Selected Service PowerVM** view refreshes, showing the performance summary of all managed servers in the **Summary - All Managed Server** view on the right.

Figure 10. Summary - All Managed Server view

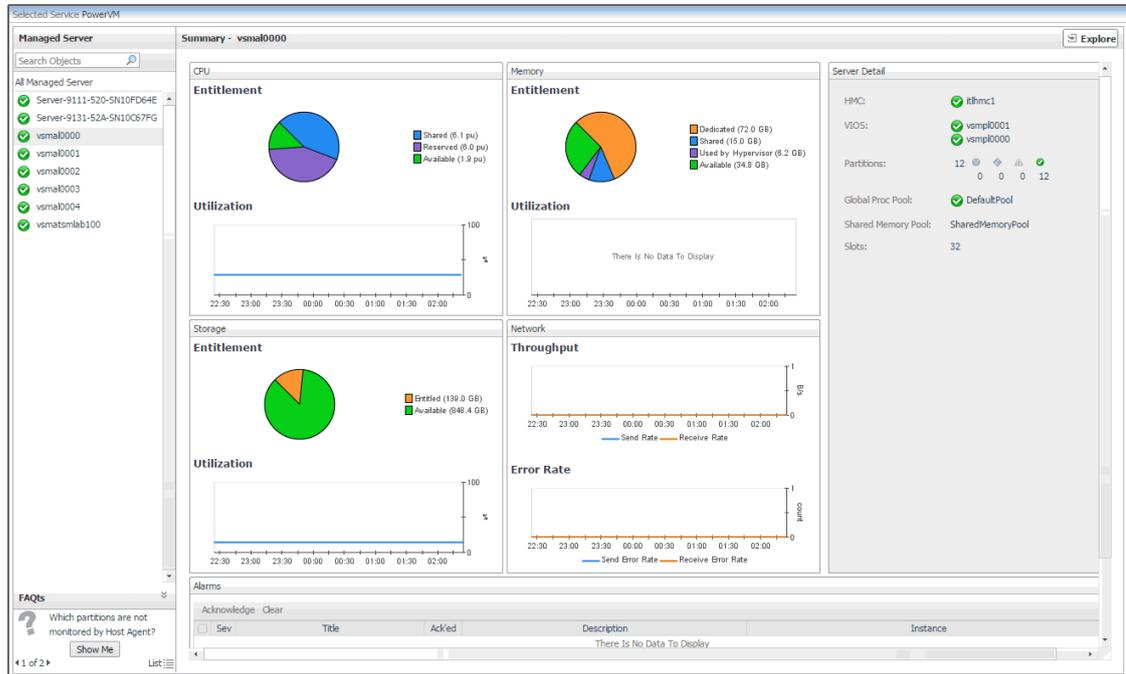


- 5 In the **Managed Server** view on the left, select an object instance from the list.

For example, if you selected **All Managed Server**, click a managed server in the list.

The Selected Service PowerVM refreshes, showing the managed server summary view on the right.

Figure 11. Managed server summary view

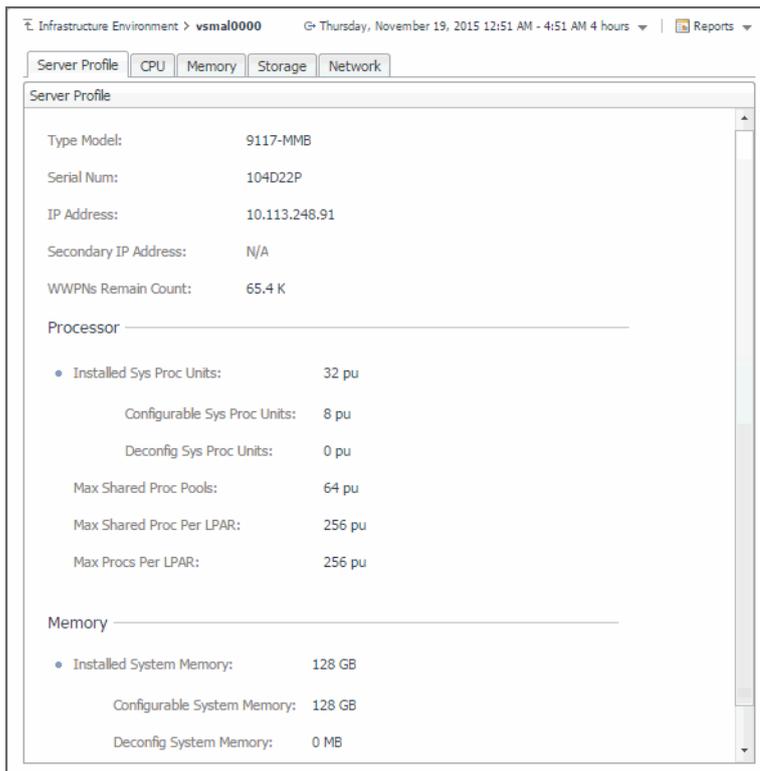


This view displays the resource utilization for the selected object.

- 6 To display more details about the selected managed server, partition, or VIOS, in the top-right corner, click **Explore**.

The display area refreshes.

Figure 12. Managed Server Details view



This view provides a set of tabs, providing physical system configuration details along with processor, memory, disk, and network utilization details. For more details about the information appearing on these tabs, see [Investigating additional managed server, partition, and VIOS details](#).

TIP: To return to the Infrastructure Environment dashboard, use the bread crumb trail in the top-left corner.

Selecting monitored objects

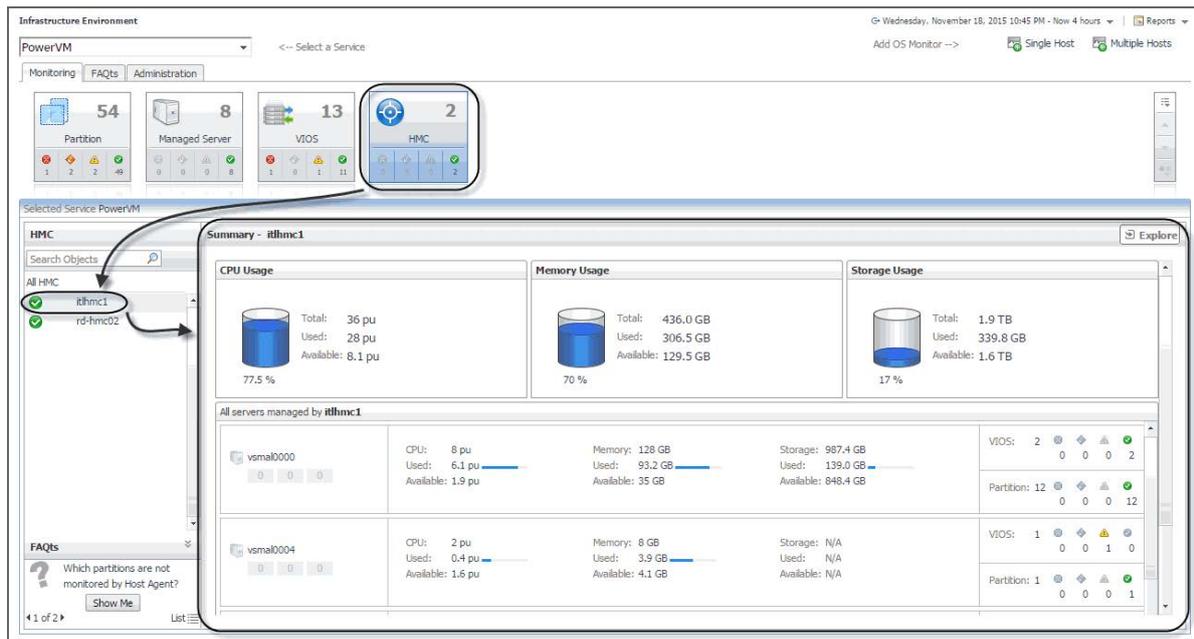
A set of tiles along the top of this tab gives you a quick overview of the monitored objects: HMC instances, managed servers, partitions, and VIOS instances. Each tile represents a collection of a specific object type, shows the object count, and the count of objects in each alarm state (Normal, Warning, Critical, and Fatal).

Figure 13. Tiles representing monitored objects



The **Selected Service PowerVM** view appearing immediately below the tiles allows you to select a specific instance of the tile selection. From here, you can drill down on a desired object instance, and review the metrics associated with the selected object.

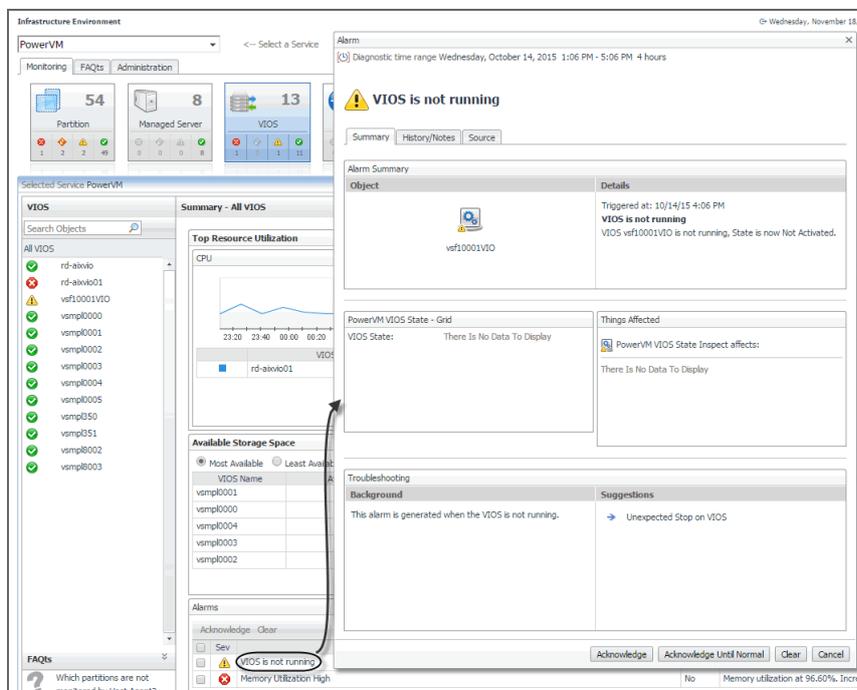
Figure 14. Selecting monitored objects



Observing alarms

If any alarms are generated against certain types of monitored objects, they appear along the bottom of the summary view on the right. Drill down on an alarm to find out what triggered it. The information appearing in the alarm dialog box can help you understand the cause of the problem, and suggest the steps you can take to further investigate the related issue.

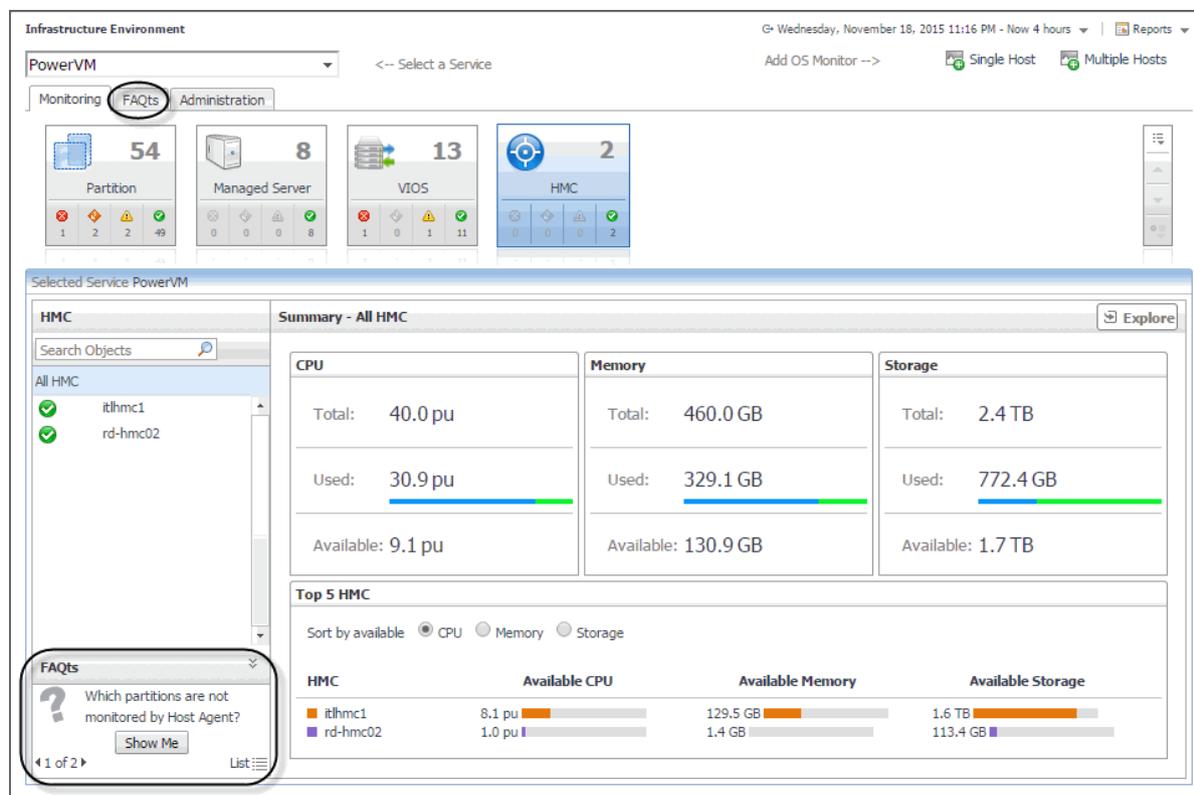
Figure 15. Drilling down on an alarm



Reviewing FAQs

The bottom-left corner of the Monitoring tab allows you to navigate through commonly asked questions associated with the selected object type, and to review their answers. You can also review these questions using the **FAQs** tab. For more information, see [Reviewing frequently asked questions on page 139](#).

Figure 16. FAQs view



Investigating the collective use of all HMC resources

Your monitored PowerVM[®] servers are managed by the one or a pair of HMCs (Hardware Management Console). You can monitor the performance of HMC instances when you select the **HMC** tile on the Infrastructure Environment dashboard.

The **Summary - All HMC** view displays the levels of total, used, and available processor, memory, and disk resources for all HMCs in your integrated infrastructure. This view appears in the Selected **Service PowerVM** view when you select **All HMCs** in the **HMC** view on the left. It identifies the elements with the highest available levels of these resources.

Figure 17. Summary - All HMC view

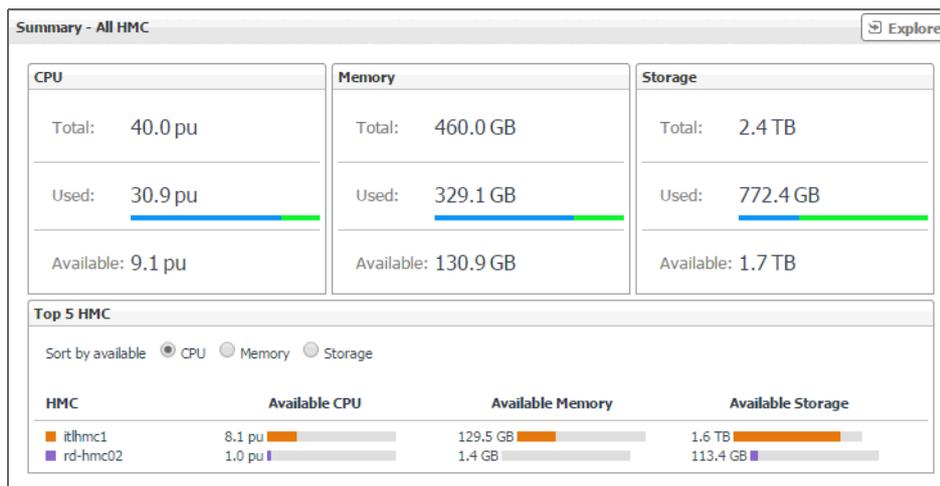


Table 1. Summary - All HMC view

CPU

- Total** The total number of processing units that are allocated to all managed servers belonging to all HMC instances in your monitored environment.
- Used** The number of processing units that are allocated to all managed servers belonging to all HMC instances in your monitored environment, that are currently in use.
- Available** The number of processing units that are allocated to all managed servers belonging to all HMC instances in your monitored environment, that are currently available for use.

Memory

- Total** The total amount of memory in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment.
- Used** The amount of memory in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment, that is currently in use.
- Available** The amount of memory in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment, that is currently available for use.

Storage

- Total** The total amount of disk storage in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment.
- Used** The amount of disk storage in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment, that is currently in use.
- Available** The amount of disk storage in GB that is allocated to all managed servers belonging to all HMC instances in your monitored environment, that is currently available for use.

Top 5 HMC

- The HMC instances that consume the highest amounts of processor, memory, and disk resources.
- Sort by available** Select **processor**, **Memory**, or **Storage**, to sort the list.
- HMC** The name of the HMC instance.
- Available CPU** The number of processing units that are allocated to all managed servers belonging to the HMC instance, that are currently available for use.
- Available Memory** The amount of memory in GB that is allocated to all managed servers belonging to the HMC instance, that is currently available for use.
- Available Storage** The amount of disk storage in GB that is allocated to all managed servers belonging to the HMC instance, that is currently available for use.

Viewing HMC details

The HMC **Summary** view displays the resource utilization for all managed servers associated with the selected HMC instance. This view shows the levels of processor, memory, and disk usage for the selected HMC. The view appears in the **Selected Service PowerVM** view when you select the **HMC** tile, and then click an HMC instance in the **HMC** view.

Figure 18. HMC summary view

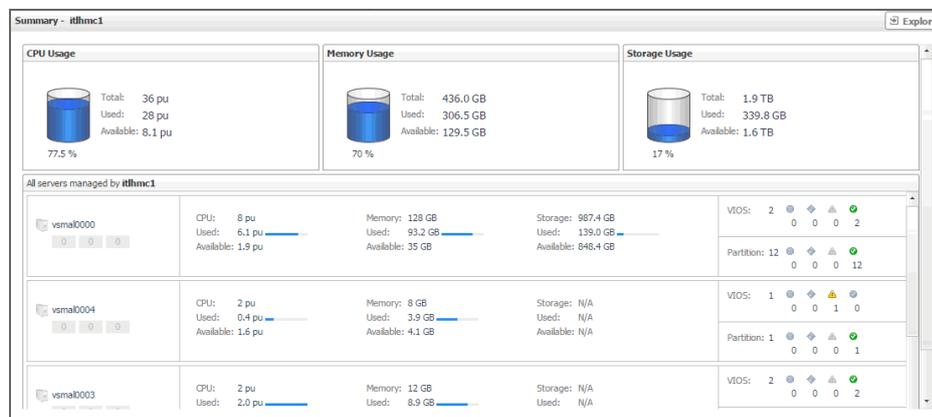


Table 2. HMC view

CPU Usage

- Total** The total number of processing units that are allocated to all managed servers belonging to the selected HMC.
- Used** The number of processing units that are allocated to all managed servers belonging to the selected HMC, that are currently in use.
- Available** The number of processing units that are allocated to all managed servers belonging to the selected HMC, that are currently available for use.
- Usage percentage** The percentage of processor resources that are allocated to all managed servers belonging to the selected HMC, that are currently in use.

Memory Usage

- Total** The total amount of memory, in GB, that is allocated to all managed servers belonging to the selected HMC.
- Used** The amount of memory, in GB, that is allocated to all managed servers belonging to the selected HMC, that is currently in use.
- Available** The amount of memory, in GB, that is allocated to all managed servers belonging to the selected HMC, that is currently available for use.
- Usage percentage** The percentage of memory resources that is allocated to all managed servers belonging to the selected HMC, that is currently in use.

Storage Usage

- Total** The total amount of disk storage, in GB, that is allocated to all managed servers belonging to the selected HMC.
- Used** The amount of disk storage, in GB, that is allocated to all managed servers belonging to the selected HMC, that is currently in use.
- Available** The amount of disk storage, in GB, that is allocated to all managed servers belonging to the selected HMC, that is currently available for use.
- Usage percentage** The percentage of disk storage resources that is allocated to all managed servers belonging to the selected HMC, that is currently in use.

All servers managed by HMC

Table 2. HMC view

Managed server	The name of the managed server.
Alarm counts	The counts of alarms generated against the managed server, given for each severity state: Warning, Critical, and Fatal.
CPU	
Used	The number of processing units that are allocated to the managed server, and are currently in use.
Available	The number of processing units that are allocated to the managed server, and are currently available for use.
Memory	
Used	The amount of memory in GB that is allocated to the managed server, and is currently in use.
Available	The amount of memory in GB that is allocated to the managed server, and is currently available for use.
Storage	
Used	The amount of disk space in GB that is allocated to the managed server, and is currently in use.
Available	The amount of disk space in GB that is allocated to the managed server, and is currently available for use.
VIOS	
Count	The number of PowerVM virtual I/O servers that manage the use of the managed server's resources by PowerVM partitions.
Alarm counts	The counts of alarms generated against the associated VIOS instances, given for each severity state: Warning, Critical, and Fatal.
Partition	
Count	The number of PowerVM partitions that use the managed server's resources.
Alarm counts	The counts of alarms generated against the associated PowerVM partitions, given for each severity state: Warning, Critical, and Fatal.

Identifying top consumers of managed server resources

A managed server is a physical system managed by the PowerVM Hypervisor software layer. PowerVM Hypervisor divides physical system resources into isolated logical partitions. A managed server typically has multiple logical PowerVM partitions. You can monitor the performance of managed servers when you select the **Managed Server** tile on the Infrastructure Environment dashboard.

The **Summary - All Managed Servers** view identifies the top five servers with available processor, memory, and disk resources for each of these categories. This view appears in the **Selected Service PowerVM** view when you select **All Managed Server** in the **Managed Server** view on the left. It recognizes the five managed servers with the lowest utilization of processor, memory, disk, and network resources that are allocated to them. Use it to quickly locate a managed server with available resources that can be allocated to new or existing logical partitions, when needed.

Figure 19. Summary - All Managed Server view

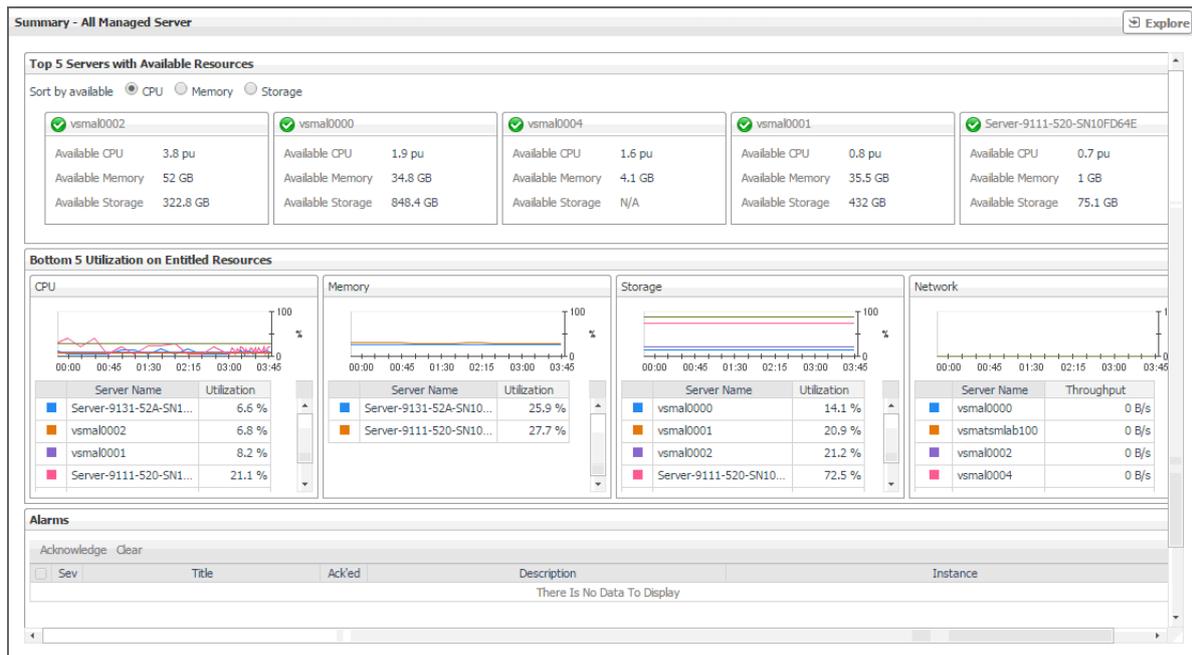


Table 3. Summary - All Managed Server view

Top 5 Servers with Available Resource	The managed servers with the highest amounts of available processor, memory, and disk resources.
Sort by available	Select CPU , Memory , or Storage , to sort the list.
Available CPU	The number of processing units that are allocated to all managed servers, and are available for use.
Available Memory	The amount of memory in GB that is allocated to all managed servers, and is currently available for use.
Available Storage	The amount of disk storage in GB that is allocated to all managed servers, and is currently available for use.
Bottom 5 Utilization on Entitled Resource	The managed servers with the lowest use of allocated processor, memory, and disk resources.
CPU	
%	The percentage of the allocated managed server processor resources that are in use over the selected time range.
Server Name	The managed server name.
Utilization	The percentage of the allocated managed server processor resources that are currently in use.
Memory	
%	The percentage of the allocated managed server memory resources that are in use over the selected time range.
Server Name	The managed server name.
Utilization	The percentage of the allocated managed server memory resources that are currently in use.
Storage	
%	The percentage of the allocated managed server disk resources that are in use over the selected time range.
Server Name	The managed server name.

Table 3. Summary - All Managed Server view

Utilization	The percentage of the allocated managed server disk resources that are currently in use.
Network	
KB/s	The rate at which the managed server transfers data to or from the network during the selected time range.
Server Name	The managed server name.
Throughput	The current rate at which the managed server transfers data to or from the network.
Alarms	For more information, see Observing alarms .

Viewing individual managed server details

The managed server **Summary** view displays the resource utilization for the selected managed server. This view shows the levels of processor, memory, disk, and network usage. It also identifies the elements of your PowerVM infrastructure associated with the selected managed server, such as the HMC, VIOS, logical partitions, and others.

This view appears in the **Selected Service PowerVM** view when you select the **Managed Server** tile, and then click a managed server in the **Managed Server** view on the left. Use it to review the trends in usage of the Managed server's system resources, and to review any generated alarms, if they exist. For example, high peaks in the processor utilization chart could result in performance degradation and should be investigated.

Figure 20. Managed Server Summary view

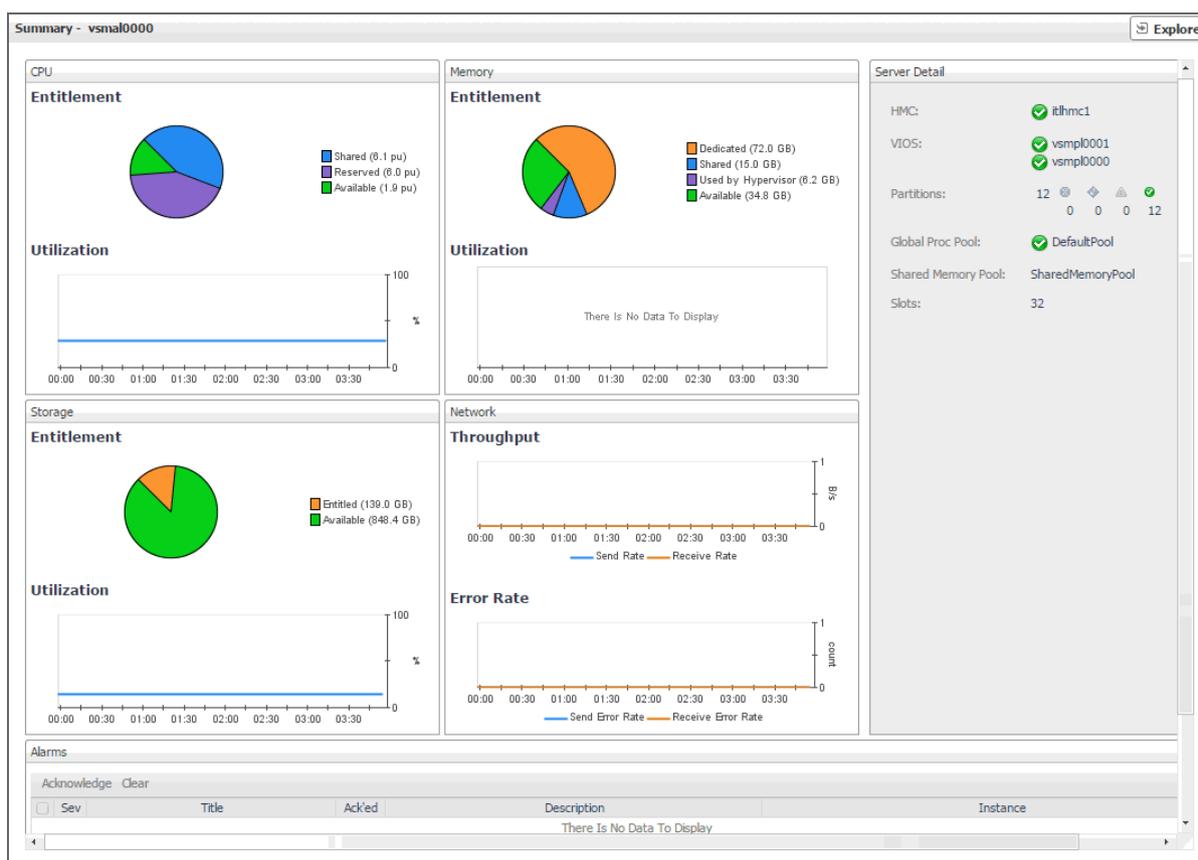


Table 4. Managed Server Summary view

CPU
Entitlement

Table 4. Managed Server Summary view

Shared	The number of processing units that are allocated to the managed server, and are currently in use in shared mode by the associated PowerVM partitions.
Dedicated	The number of processing units that are allocated to the managed server, and are currently in use in dedicated mode by the associated PowerVM partitions.
Available	The number of processing units that are allocated to the managed server, and are available for use.
Utilization	The percentage of the managed server processor resources that are in use over the selected time range.
Memory	
Entitlement	
Shared	The amount of memory that is allocated to the managed server, and is currently in use in shared mode by the associated PowerVM partitions.
Dedicated	The amount of memory that is allocated to the managed server, and is currently in use in dedicated mode by the associated PowerVM partitions.
Hypervisor Used	The amount of memory that is allocated to the managed server, and is currently in use by the PowerVM Hypervisor software layer.
Available	The amount of memory that is allocated to the managed server, and is available for use.
Utilization	The percentage of the managed server memory resources that are in use over the selected time range.
Storage	
Entitlement	
Entitled	The amount of disk space that is allocated to the managed server, and is currently in use by the associated PowerVM partitions.
Available	The amount of disk space that is allocated to the managed server, and is available for use.
Utilization	The percentage of the managed server disk space that is in use over the selected time range.
Network	
Throughput	
Send Rate	The rate at which the managed server transfers data to the network during the selected time range.
Receive Rate	The rate at which the managed server transfers data from the network during the selected time range.
Error Rate	
Send Error Rate	The rate at which the managed server encounters errors while transferring data to the network during the selected time range.
Receive Error Rate	The rate at which the managed server encounters errors while transferring data from the network during the selected time range.
Related Info	
HMC	The name of the HMC instance to which the managed server belongs, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
VIOS	The names of the virtual I/O servers associated with this managed server, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Partitions	The number of PowerVM partitions that use the managed server's resources, followed by the numbers of partitions in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Global Proc Pool	The name of the global processor pool associated with this managed server, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .

Table 4. Managed Server Summary view

Shared Memory Pool	The name of the global memory pool associated with this managed server, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Slot	The slot number of the managed server in your monitored system.
Alarms	For more information, see Observing alarms .

Identifying top consumers of PowerVM partition resources

A PowerVM[®] partition encapsulates operating system and application components. It runs inside a physical server and is dedicated a share of its physical resources. PowerVM partitions are managed by the PowerVM Hypervisor software. You can monitor the performance of PowerVM partitions when you select the **Partition** tile on the Infrastructure Environment dashboard.

The **Summary - All Partition** view identifies five partitions with the highest processor, memory, and disk entitlement, the highest processor, memory, and disk utilization, and the highest network throughput. This view appears in the **Selected Service PowerVM** view when you select **All Partition** in the **Partition** view on the left. Use it to quickly locate a PowerVM partition system with available resources that can be put to specific use in your integrated environment.

Figure 21. Summary - All Partition view



Table 5. Summary - All Partitions view

Top CPU Utilization	The PowerVM partitions consuming the most of the processor resources allocated to them.
%	The percentage of all processing units that are allocated to the PowerVM partition and are in use over the selected time range.
Partition Name	The PowerVM partition name.

Table 5. Summary - All Partitions view

Utilization	The percentage of all processing units that are allocated to the PowerVM partition and are currently in use.
Top Memory Utilization	The PowerVM partitions consuming the most of the memory resources allocated to them.
%	The percentage of memory resources that is allocated to the PowerVM partition and is in use over the selected time range.
Partition Name	The PowerVM partition name.
Utilization	The percentage of the memory resources that is allocated to the PowerVM partition and is currently in use.
Top Storage Utilization	The PowerVM partitions consuming the most of the disk space allocated to them.
%	The percentage of disk space that is allocated to the PowerVM partition and is in use over the selected time range.
Partition Name	The PowerVM partition name.
Utilization	The percentage of the memory resources that is allocated to the PowerVM partition and is currently in use.
Top Network Throughput	The PowerVM partitions with the highest network throughput.
b/s	The rate at which the PowerVM partition transfers data to or from the network during the selected time range.
Partition Name	The PowerVM partition name.
Throughput	The current rate at which the PowerVM partition transfers data to or from the network.
Top CPU Entitlement	The PowerVM partitions with the highest numbers of allocated processing units.
pu	The number of processing units allocated to the PowerVM partition over the selected time range.
Partition Name	The PowerVM partition name.
Entitlement	The number of processing units allocated to the PowerVM partition.
Top Memory Entitlement	The PowerVM partitions with the highest amounts of allocated memory resources.
GB/MB	The amount of memory allocated to the PowerVM partition over the selected time range.
Partition Name	The PowerVM partition name.
Entitlement	The amount of memory allocated to the PowerVM partition.
Top Storage Entitlement	The PowerVM partitions with the highest amounts of allocated disk space.
GB/MB	The amount of disk space allocated to the PowerVM partition over the selected time range.
Partition Name	The PowerVM partition name.
Entitlement	The amount of disk space allocated to the PowerVM partition.
Alarms	For more information, see Observing alarms .

Viewing individual PowerVM partition details

The PowerVM[®] Partition **Summary** view displays the resource utilization for the selected partition. This view shows the levels of processor, memory, disk, and network usage. It also identifies the elements of your PowerVM infrastructure associated with the selected PowerVM partition, such as the HMC, managed server, VIOS, and others. This view appears in the **Selected Service PowerVM** view when you select the **Partition** tile, and a PowerVM partition in the **Selected Service PowerVM** view, on the left.

Use this view to review the trends in usage of the partition's system resources, and to review any generated alarms, if they exist. For example, high peaks in the processor usage chart could cause performance degradation and should be investigated.

Figure 22. Partition Summary view



Table 6. Partition Summary view

CPU

Alarm state	The alarm state of the processor resources allocated to the PowerVM partition.
Alarms	The counts of alarms generated against the processor resources allocated to the PowerVM partition, given for each severity state: Warning, Critical, and Fatal.
History	A color-coded bar, representing the alarm state of the processor resources allocated to the PowerVM partition over the selected time range. The color of the bar changes depending on the alarm state. Red indicates a Fatal state, orange indicates Critical, yellow means Warning, and green is the Normal State.
CPU Utilization	The percentage of all processing units that are allocated to the PowerVM partition and are currently in use.
Baseline	An envelope indicating the expected processor utilization range based on historical data.
Run Queue Length	The number of processes that are either running or waiting in the queue. A high number of processes in the run queue means the processor is busy. A consistently high number can indicate that the PowerVM partition needs additional processor resources.
CPU Usage by Process	Click this link to display the Processes dashboard and find out more about the processor and memory usage by each process running on this PowerVM partition. For more information about this dashboard, see the <i>Foglight™ for Infrastructure User and Reference Guide</i> .

Memory

Alarm state	The alarm state of the memory resources allocated to the PowerVM partition.
Alarms	The counts of alarms generated against the memory resources allocated to the PowerVM partition, given for each severity state: Warning, Critical, and Fatal.
History	A color-coded bar, representing the alarm state of the memory resources allocated to the PowerVM partition over the selected time range. The color of the bar changes depending on the alarm state. Red indicates a Fatal state, orange indicates Critical, yellow means Warning, and green is the Normal State.

Table 6. Partition Summary view

Memory Utilization	The percentage of all memory resources that are allocated to the PowerVM partition and are currently in use.
Baseline	An envelope indicating the expected memory utilization range based on historical data.
Memory Usage by Process	Click this link to display the Processes dashboard and find out more about the processor and memory usage by each process running on this PowerVM partition. For more information about this dashboard, see the <i>Foglight for Infrastructure User and Reference Guide</i> .
Storage	
Alarm state	The alarm state of the disk resources allocated to the PowerVM partition.
Alarms	The counts of alarms generated against the disk resources allocated to the PowerVM partition, given for each severity state: Warning, Critical, and Fatal.
History	A color-coded bar, representing the alarm state of the disk resources allocated to the PowerVM partition over the selected time range. The color of the bar changes depending on the alarm state. Red indicates a Fatal state, orange indicates Critical, yellow means Warning, and green is the Normal State.
Disk Utilization	The percentage of all disk resources that are allocated to the PowerVM partition and are currently in use.
Baseline	An envelope indicating the expected disk utilization range based on historical data.
Disk Transfer Rate	The current rate at which the PowerVM partition transfers data to or from the disk.
Network	
Alarm state	The alarm state of the network interfaces allocated to the PowerVM partition.
Alarms	The counts of alarms generated against the network interfaces allocated to the PowerVM partition, given for each severity state: Warning, Critical, and Fatal.
History	A color-coded bar, representing the alarm state of the network interfaces allocated to the PowerVM partition over the selected time range. The color of the bar changes depending on the alarm state. Red indicates a Fatal state, orange indicates Critical, yellow means Warning, and green is the Normal State.
Network Utilization	The percentage of the network resources that are allocated to the PowerVM partition and are currently in use.
Baseline	An envelope indicating the expected network utilization range based on historical data.
Related Info	
HMC	The name of the HMC to which the PowerVM partition belongs, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Managed Server	The name of the managed server associated with this PowerVM partition, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
VIOS	The name of the virtual I/O server associated with this PowerVM partition, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Proc Pool	The name of the processor pool associated with this PowerVM partition, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Shared Memory Pool	The name of the global memory pool associated with this PowerVM partition, preceded by its alarm state: Normal  , Warning  , Critical  , or Fatal  .
Slot	The slot number dedicated to the PowerVM partition in your monitored system.
Other Alarms	For more information, see Observing alarms .

Identifying top consumers of PowerVM virtual I/O server resources

A PowerVM® Virtual I/O Server (VIOS) enables logical partitions to share physical resources of a managed server among multiple PowerVM partitions. It can provide virtualized storage along with network adapters, allowing you to facilitate the LAN and disk resources. You can monitor the performance of VIOS instances when you select the **VIOS** tile on the Infrastructure Environment dashboard.

The **Summary - All VIOS** view identifies five servers with the highest processor and memory utilization, and the most or least available storage space, and the most or least network throughput. This view appears in the **Selected Service PowerVM** view when you select **All VIOS** in the **VIOS** view on the left. Use it to quickly locate a VIOS with available resources that can be put to better use in your integrated environment.

Figure 23. Summary - All VIOS view

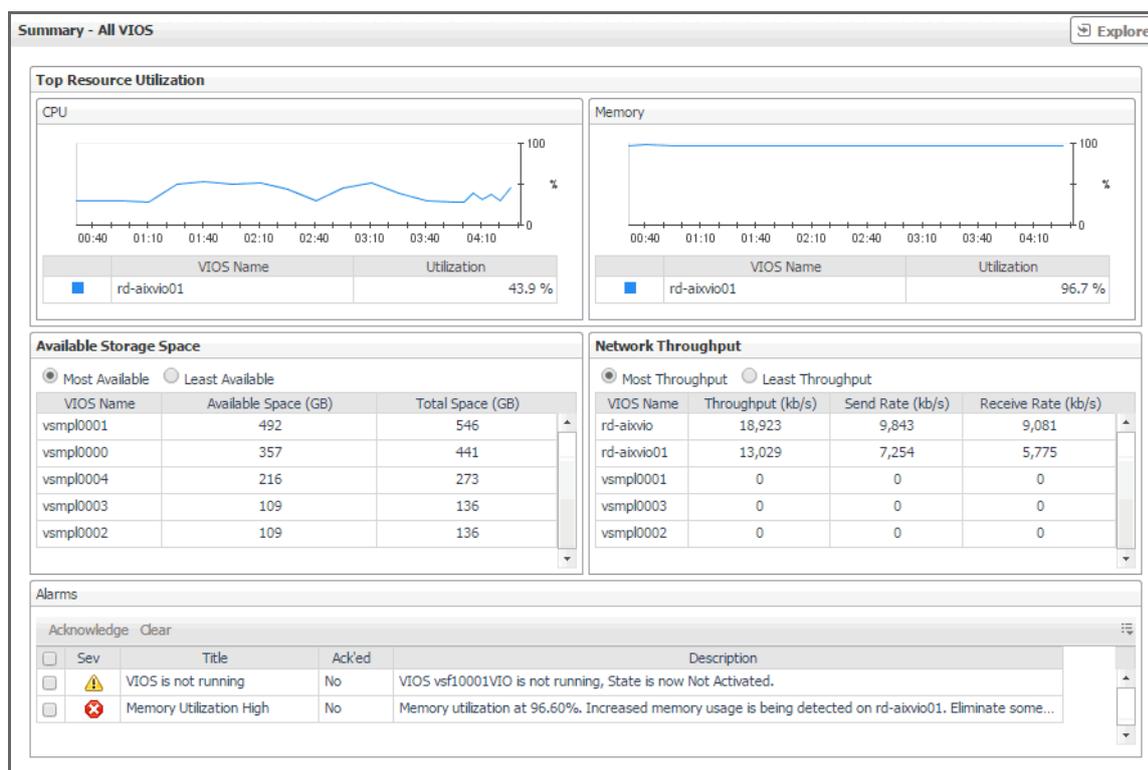


Table 7. Summary - All VIOS view

Top Resource Utilization

CPU	The VIOS instances with the highest utilization of allocated processor resources.
%	The percentage of the processor resources allocated to the VIOS that are in use over the selected time range.
VIOS Name	The VIOS name.
Utilization	The percentage of the allocated processor resources that are currently in use by the VIOS.
Memory	The VIOS instances with the highest utilization of allocated memory resources.
%	The percentage of the memory resources allocated to the VIOS that are in use over the selected time range.
VIOS Name	The VIOS name.
Utilization	The percentage of the allocated VIOS memory resources that are currently in use.

Table 7. Summary - All VIOS view

Available Storage Space

- Most Available** Select this option to sort the VIOS instances by the most available disk space.
- Least Available** Select this option to sort the VIOS instances by the least available disk space.
- VIOS Name** The VIOS name.
- Available Space (GB)** The amount of disk space allocated to the VIOS that is available for use.
- Total Space (GB)** The total amount of disk space allocated to the PowerVM VIOS.

Network Throughput

- Most Throughput** Select this option to sort the VIOS instances by the highest network throughput.
- Least Throughput** Select this option to sort the VIOS instances by the lowest network throughput.
- VIOS Name** The VIOS name.
- Throughput (kb/s)** The current rate at which the VIOS transfers data to or from the network.
- Send Rate (kb/s)** The current rate at which the VIOS transfers data to the network.
- Receive Rate (kb/s)** The current rate at which the VIOS transfers data from the network.

Alarms For more information, see [Observing alarms](#).

Viewing individual PowerVM VIOS details

The VIOS **Summary** view displays the resource utilization for the selected PowerVM® VIOS. This view identifies the resource components managed by this VIOS, the levels of available and assigned disk resources, the network throughput rate and numbers of errors and dropped packets over the selected time range, and the trends in the host processor and memory utilization.

This view appears in the **Selected Service PowerVM** view when you select the **VIOS** tile, and then click a VIOS in the **VIOS** view on the left. Use it to review the trends in usage of the resources managed by the selected PowerVM VIOS, and to review any generated alarms, if they exist. For example, high peaks in the **VIOS Managed Network Resource** chart could indicate performance degradation and should be investigated.

Figure 24. VIOS Summary view

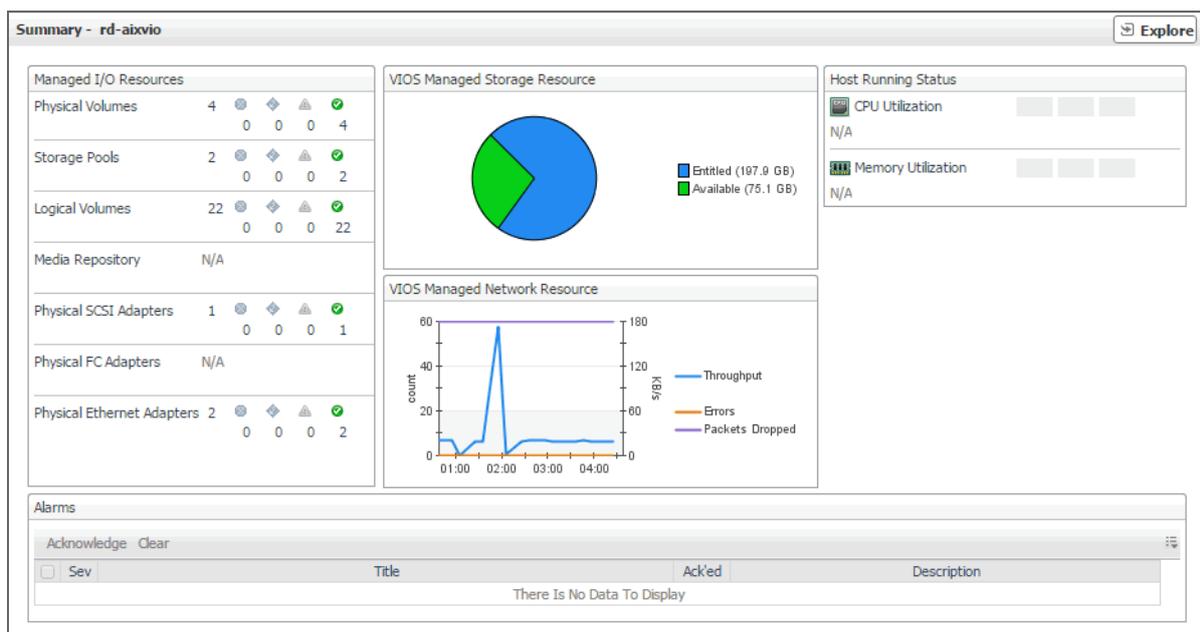


Table 8. VIOS Summary view

Managed I/O Resources	
Physical Volumes	The number of physical disks available to the PowerVM virtual I/O server, followed by the numbers of physical disks in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Storage Pools	The number of storage pools available to the PowerVM virtual I/O server, followed by the numbers of storage pools in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Logical Volumes	The number of logical disks available to the PowerVM virtual I/O server, followed by the numbers of logical disks in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Media Repository	The number of media repositories available to the PowerVM virtual I/O server, followed by the numbers of media repositories in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Physical SCSI Adapter	The number of physical SCSI (Small Computer System Interface) adapters available to the PowerVM virtual I/O server, followed by the numbers of physical SCSI adapters in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Physical FC Adapter	The number of physical FC (fiber channel) adapters available to the PowerVM virtual I/O server, followed by the numbers of adapters in each severity state: Normal  , Warning  , Critical  , and Fatal  .
Physical Ethernet Adapter	The number of physical Ethernet adapters available to the PowerVM virtual I/O server, followed by the numbers of adapters in each severity state: Normal  , Warning  , Critical  , and Fatal  .
VIOS Managed Storage Resource	
Entitled Available	The amount of disk space that is managed by the PowerVM VIOS.
	The amount of disk space that is managed by the PowerVM VIOS, and is available for entitlement to the associated PowerVM partitions, or for its own use.
VIOS Managed Network Resource	
Throughput	The rate at which the PowerVM VIOS transfers data to and from the network during the selected time range.
Errors	The rate at which the PowerVM VIOS encounters errors while transferring data to and from the network during the selected time range.
Packets Dropped	The rate at which the PowerVM VIOS drops data packets while transferring data to and from the network during the selected time range. By default, network adapters send or receive multiple data packets in a single request. When data packets are dropped before reaching their destinations, this may indicate a problem with the network connection or the network adapters.
Host Running Status	
CPU Utilization	The percentage of the allocated processor resources that the host associated with the PowerVM VIOS currently uses.
Memory Utilization	The percentage of the allocated memory resources that the host associated with the PowerVM VIOS currently uses.
Alarms	For more information, see Observing alarms .

Investigating additional managed server, partition, and VIOS details

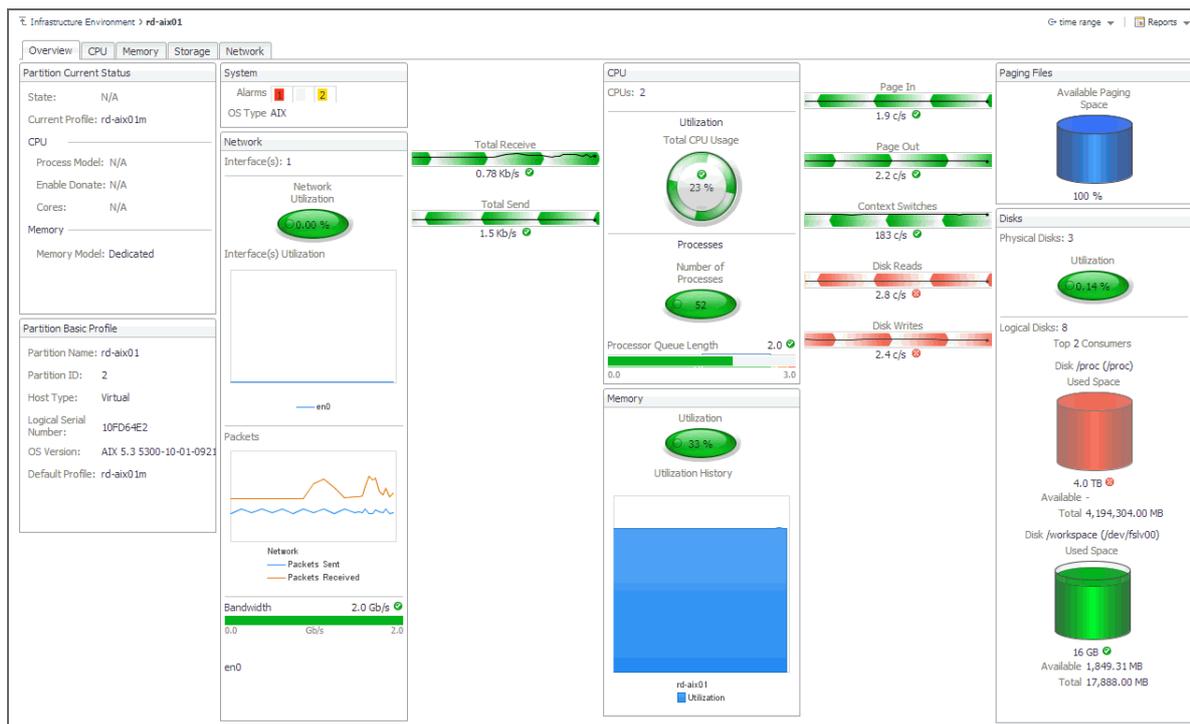
If you see any indicators that can potentially lead to performance degradation, you can explore managed servers, partitions, and virtual I/O servers in more detail to find out more information that can help you prevent service interruptions. To do that, select a desired managed server, partition, or virtual I/O server object, and click **Explore** in the top-right corner of the **Selected Service PowerVM** view.

Figure 25. Exploring monitored PowerVM objects in more detail



The type and extent of information appearing in the resulting **Details** view depends on the type of the selected object. For example, if you drill down on a PowerVM partition, the **Details** view consists of four tabs: **Overview**, **CPU**, **Memory**, **Storage**, and **Network**, each describing the relevant resource-related details.

Figure 26. PowerVM Partition Details view



Observe the information on these tabs when you want to understand the underlying cause of performance degradation. For example, unusually high values of memory usage can lead to system degradation and should be further investigated.

Exploring PowerVM partition and PowerVM VIOS overviews

The **Overview** tab appears open when you explore a PowerVM® partition or a PowerVM VIOS for the first time. It displays the state of the partition's processor, memory, disk, and network resources. These visual elements are connected with graphical flows that illustrate the flow of data in real time. For example, you can review the rates of incoming and outgoing network data.

To navigate to this tab, in the **Selected Service PowerVM**, select a PowerVM partition or a PowerVM VIOS, and click **Explore**.

Figure 27. PowerVM Partition Details view

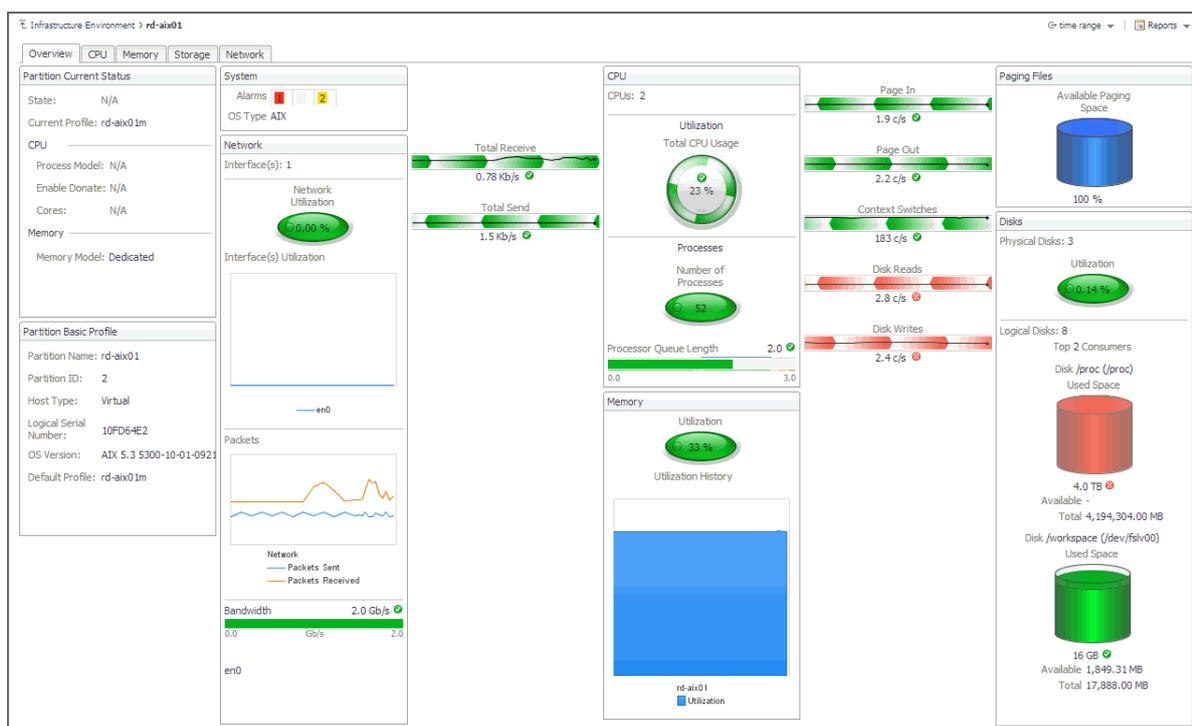


Table 9. PowerVM Partition Details and PowerVM VIOS Details views

Partition/VIOS Current Status

State	The running state of the PowerVM Partition or VIOS. Possible values include: Not Activated, Starting, Running, Shutting Down, Error, Open Firmware, Not Available.
Current Profile	The name of the currently used PowerVM Partition or VIOS profile. A profile specifies how many processors, how much memory, and which I/O devices and slots are to be allocated to the partition. A PowerVM Partition or VIOS can have multiple profiles.
CPU	
Process Model	The process model of the PowerVM Partition or VIOS. There are two types of processing modes: Shared and Dedicated. Shared mode allows assigning partial processor units from the shared processor pool. Dedicated mode indicates that all processor units can be assigned to one PowerVM partition or VIOS, for its own use.
Enable Donate	Indicates if the PowerVM Partition or VIOS can donate spare processor cycles to the shared pool.
Smt	The number of threads the PowerVM Partition or VIOS uses.
Memory	

Table 9. PowerVM Partition Details and PowerVM VIOS Details views

Memory Model	The memory model the PowerVM Partition or VIOS uses: Dedicated or Shared .
Partition/VIOS Basic Profile	
Partition/VIOS Name	The name of the PowerVM Partition or VIOS
Partition ID	The ID of the PowerVM Partition or VIOS object in your monitored system.
Host Type	The type of the system host on which the PowerVM Partition or VIOS object is running: Virtual or Physical .
Logical Serial Number	The serial number of the PowerVM Partition or VIOS.
OS Version	The OS name and version running on the PowerVM Partition or VIOS.
Default Profile	The default PowerVM Partition or VIOS profile, indicating how resources are assigned to that component (selected or all resources). If the current profile does not exist, the system loads this default profile. A profile specifies how many processors, how much memory, and which I/O devices and slots are to be allocated to the partition. A PowerVM Partition or VIOS can have multiple profiles.

NOTE: For information about the other views also appearing in **PowerVM Partition Details** view, see the *Foglight User Guide*.

Exploring managed server profiles

The **Server Profile** tab appearing on the managed server view displays system configuration information about the selected managed server. Use it to find out the overall capacity of the physical system hosting multiple logical partitions.

To navigate to this tab, in the **Selected Service PowerVM** view, select a managed server, and click **Explore**.

Figure 28. Server Profile tab

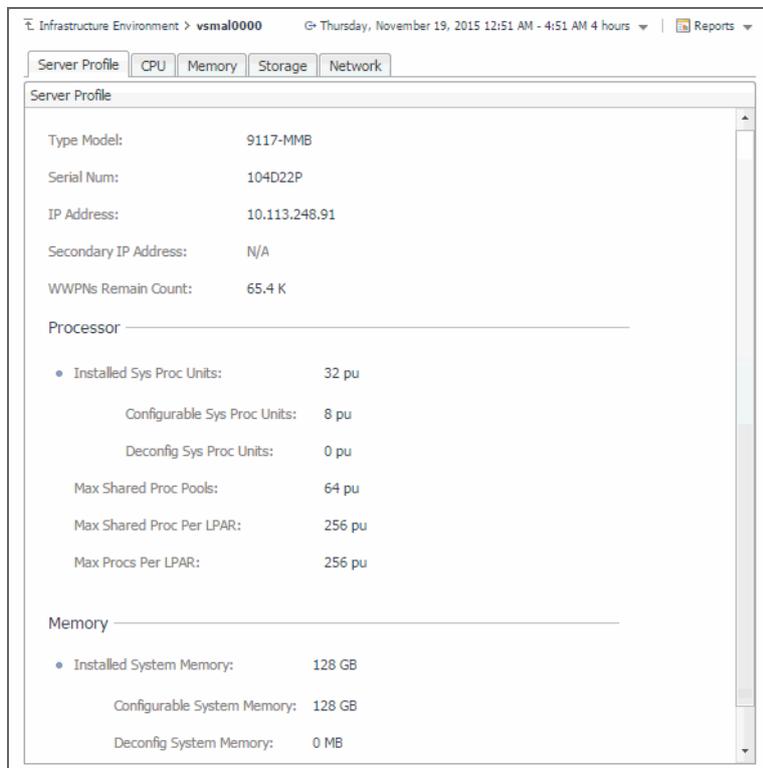


Table 10. Server Profile tab

Type Model	The managed server model number.
Serial Num	The managed server serial number.
IP Address	The primary IP address of the managed server in your monitored environment.
Secondary IP Address	The secondary IP address of the managed server.
WWPNs Remain Count	The remaining number of world wide port names (WWPNs).
Processor	
Installed Sys Proc Units	The number of processor units allocated to the managed server.
Configurable Sys Proc Units	The number of processor units that can be allocated to the shared processor pool and reallocated to PowerVM partitions associated with the managed server.
Deconfig Sys Proc Units	The number of processor units that cannot be allocated to the shared processor pool.
Max Shared Proc Pools	The maximum number of shared processor pools associated with the managed server.
Max Shared Proc Per LPAR	The maximum number of shared processor units allocated to each PowerVM partition associated with the managed server.
Max Procs Per LPAR	The maximum number of processor units allocated to each PowerVM partition associated with the managed server.
Memory	
Installed System Memory	The amount of system memory allocated to this managed server.
Configurable System Memory	The portion of system memory that can be allocated to the shared memory pool and reallocated to PowerVM partitions associated with the managed server.
Deconfig System Memory	The portion of system memory that cannot be allocated to the shared memory pool.

Investigating the use of PowerVM partition and VIOS processor resources

The **CPU** tab of the Partition detail view or VIOS detail view shows the amount of processor resources dedicated to the selected object and their entitlement mode. It also shows the top processor consumers on the host, the top processes' utilization trends, and illustrates the processor load.

To navigate to this tab, in the **Selected Service PowerVM** view, select a PowerVM® partition or a VIOS, click **Explore**, and in the view that appears, open the **CPU** tab.

Figure 29. CPU tab

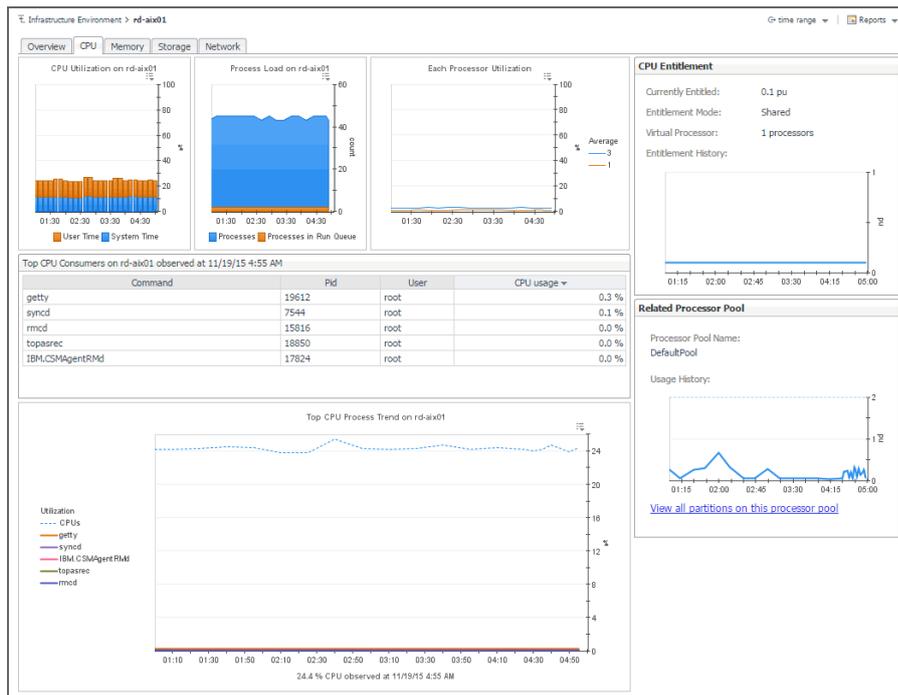


Table 11. CPU tab

CPU Entitlement

- Currently Entitled** The number of processing units allocated to this PowerVM partition or VIOS.
- Entitlement Mode** Indicates if the processor resources allocated to PowerVM partition or VIOS can be allocated to the shared processor pool and reallocated to other PowerVM partitions or virtual I/O servers.
- Virtual Process** The number of processors allocated to the virtual process associated with this PowerVM partition or VIOS.
- Entitlement History** The number of processing units allocated to this PowerVM partition or VIOS over the selected time range.

Related Processor Pool

- Processor Pool Name** The name of the processor pool associated with this PowerVM partition or VIOS.
- Usage History** The number of processing units allocated to this processor pool that are used over the selected time range.

View all partitions on this processor pool Clicking this link displays the **All Partitions in Pool** dialog box that lists all PowerVM partitions that use this processor pool, their entitlement, utilization, and usage.

Partition Name	Entitlement	Utilization	Usage
rd-aixp04	0.4 pu	2 %	0.0 pu
rd-aixivio	0.3 pu	3 %	0.0 pu
rd-aix02	0.2 pu	3 %	0.0 pu
rd-aix03	0.1 pu	5 %	0.0 pu
rd-aixp02	0.1 pu	4 %	0.0 pu
rd-aix01	0.1 pu	55 %	0.1 pu
torbld05	0.1 pu	n/a	n/a

NOTE: The other views also appearing on this tab are defined on the Management Server and as such are documented in the *Foglight™ User Guide*.

Investigating the use of PowerVM partition and VIOS memory resources

The **Memory** tab of the Partition detail view or VIOS detail view shows the amount of memory resources dedicated to the selected PowerVM® partition or PowerVM VIOS, their entitlement model, and identifies the managed server on which the partition is running. It also shows the overall memory utilization on the selected partition and the amount of memory the top processes are consuming. The **Top Memory Consumers** table lists both resident and virtual sizes for the top processes.

To navigate to this tab, in the **Selected Service PowerVM** view, select a PowerVM partition or PowerVM VIOS, click **Explore**, and in the view that appears, open **Memory** tab.

Figure 30. Memory tab

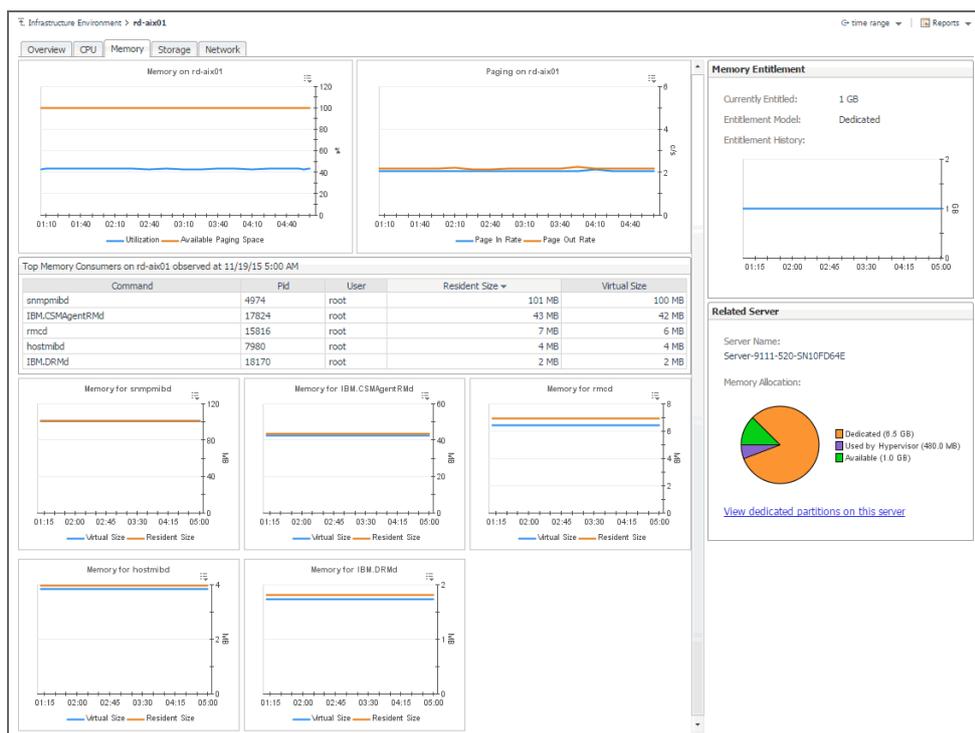


Table 12. Memory tab

Memory Entitlement

- Currently Entitled** The amount of memory allocated to this PowerVM partition or VIOS.
- Entitlement Model** Indicates if the memory resources allocated to PowerVM partition or VIOS can be allocated to the shared memory pool and reallocated to other PowerVM partitions or virtual I/O servers.
- Entitlement History** The amount of memory allocated to this PowerVM partition or VIOS over the selected time range.

Related Server

- Server Name** The name of the PowerVM server this PowerVM partition or VIOS belongs to.
- Memory Allocation**
 - Dedicated** The amount of memory that is allocated to the PowerVM server this PowerVM partition or VIOS belongs to, and is currently in use by the associated PowerVM partitions.
 - Hypervisor Used** The amount of memory that is allocated to the PowerVM server this PowerVM partition or VIOS belongs to, and is currently in use by the PowerVM Hypervisor software layer.

Table 12. Memory tab

Available

The amount of memory that is allocated to the PowerVM server this PowerVM partition or VIOS belongs to, and is available for use.

View dedicated partitions on this server

Click this link to display the **All Dedicated Partitions in Server** dialog box that lists all dedicated PowerVM partitions that belong to this PowerVM server, their entitlement, utilization, and usage.

Partition Name	Entitlement	Utilization	Usage
rd-aix02	2 GB	n/a	n/a
rd-aix03	2 GB	n/a	n/a
rd-axlp04	2 GB	38 %	583 MB
rd-aix01	1 GB	26 %	265 MB
rd-axlp02	448 MB	80 %	359 MB
torbld05	0.0 MB	n/a	n/a

NOTE: The other views also appearing on this tab are defined on the Management Server and as such are documented in the *Foglight User Guide*.

Exploring processor, memory, and storage breakdowns

The **Breakdown** views on the **CPU**, **Memory**, and **Storage** tabs of the managed server details view, and on the **Storage** tab of the VIOS details view show the amount of processor, memory, and storage resources the selected managed server or PowerVM[®] VIOS uses from the Global Default Pool over the selected time range. Use this view to find out how the selected object uses the allocated resources, and if there are any patterns indicating their usage levels that may potentially compromise the stability of your infrastructure.

To navigate to this tab, in the **Selected Service PowerVM** view, select a managed server, or a PowerVM VIOS, and click **Explore**. Open the **CPU** (PowerVM servers only), **Memory** (PowerVM servers only), or **Storage** tab, and then open the **Breakdown** tab.

Figure 31. Breakdown tab

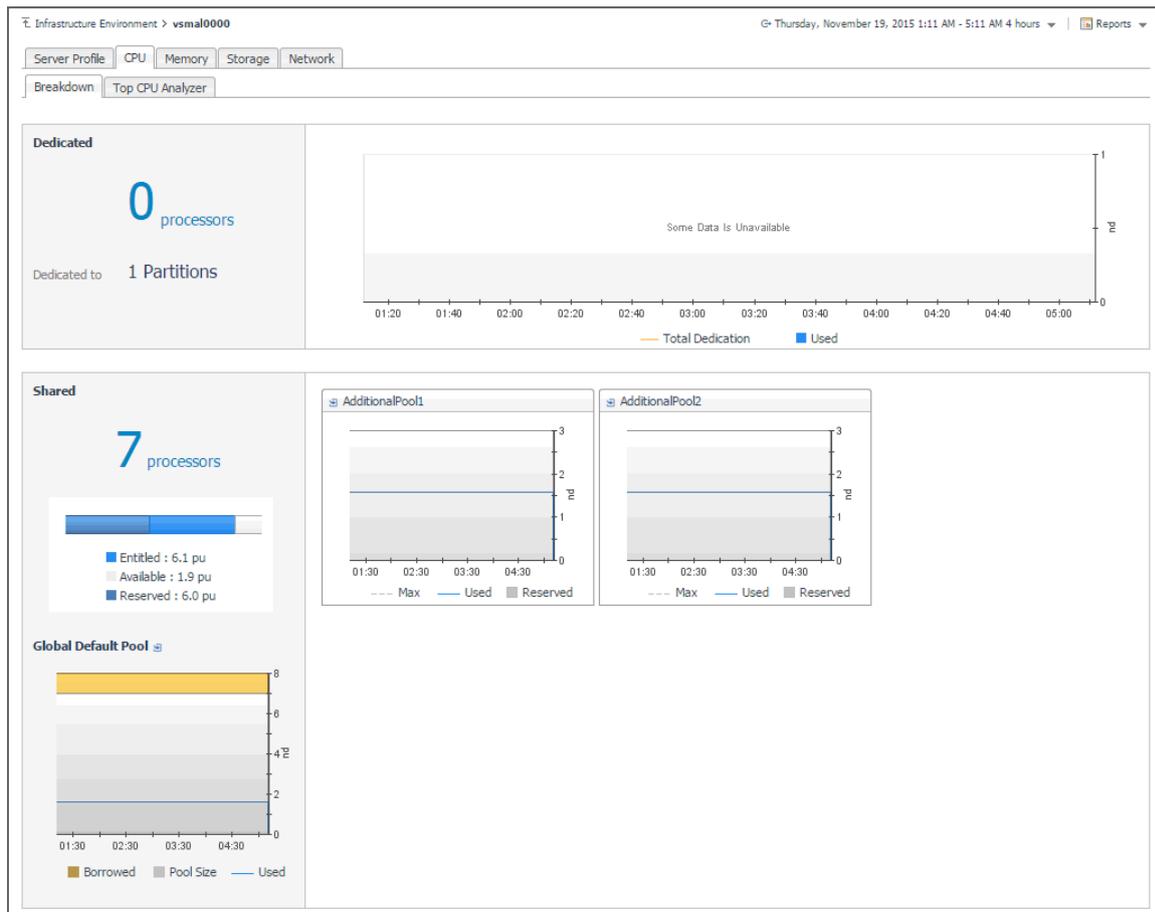


Table 13. Breakdown tab

CPU tab: Breakdown view

Dedicated	Information about the processor dedicated to the PowerVM partitions that are running on the selected managed server.
Shared	The number of processing units shared among the PowerVM partitions that are running on the selected managed server.
Entitled	The number of processing units that are allocated to the PowerVM partitions that are running on the selected managed server.
Available	The number of processing units that are available for allocation to the PowerVM partitions that are running on the managed server.
Reserved	The number of processing units that are reserved by the Shared Processor Pool associated with the selected managed server.
Global Default Pool	
Borrowed	The total number of processing units that are borrowed from donating PowerVM partitions and of all processing units.
Pool Size	The total number of processing units contributed to the global default pool, over the selected time range.
Used	The number of processing units contributed to the global default pool that are in use by the associated PowerVM partitions, over the selected time range.
AdditionalPool	
Max	The highest number of processing units allocated to the additional pool, over the selected time range.

Table 13. Breakdown tab

Used	The number of processing units allocated to the additional pool that are in use, over the selected time range.
Reserved	The number of processing units allocated to the additional pool that are reserved, over the selected time range.

Memory tab: Breakdown view

Dedicated	The amount of memory that is allocated specifically to the partitions running on the managed server, and not participating in shared memory pools.
Dedicated to	The number of partitions to which the managed server allocates memory resources that do not participate in shared memory pools.
Total Dedication	The total amount of memory that is allocated specifically to the partitions running on the managed server, and not participating in shared memory pools, over the selected time range.
Used	The amount of memory that is allocated specifically to the partitions running on the managed server, and used by those partitions, over the selected time range.
Shared	Information about the memory resources that the managed server allocates to the shared memory pool.

Storage tab: Breakdown view

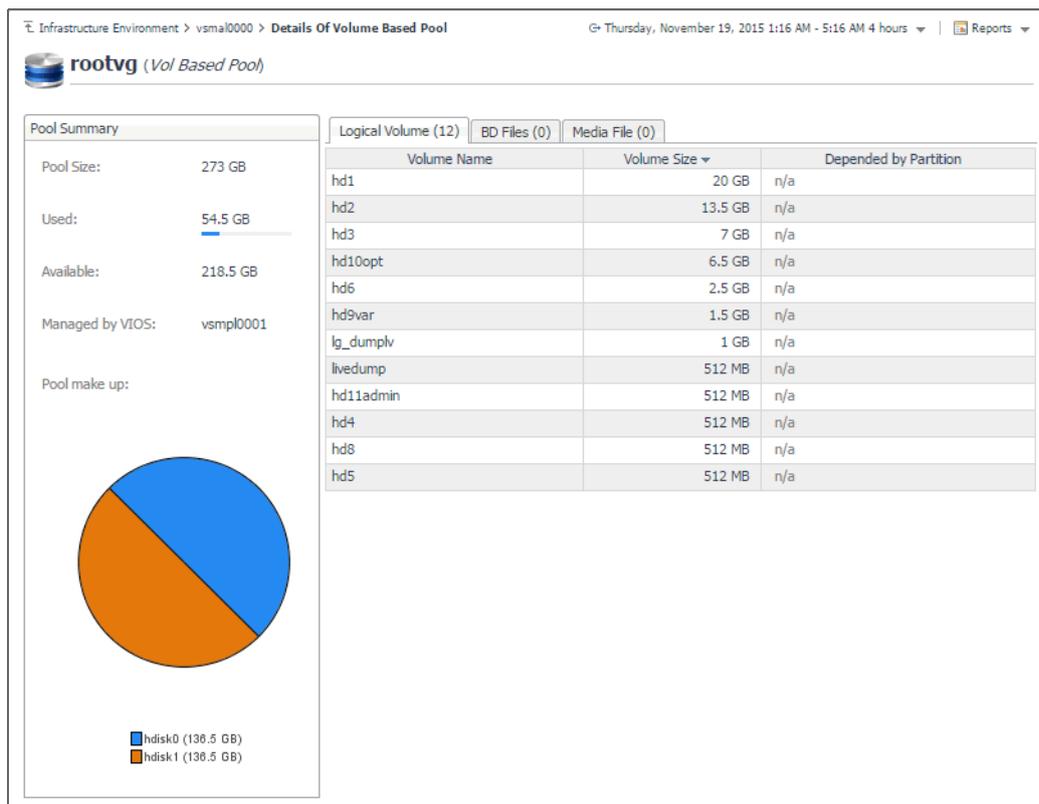
Unentitled Volume	Information about the disk volume that is not allocated to the partitions running on the managed server or PowerVM VIOS.
Dedicated Volume	Information about disk volumes that are dedicated and in use.
Volume Based Pools Entitled	
Physical Volume	The number of physical disks that are used to create volume-based pools.
Volume Based Pools	The number of volume-based pools associated with the selected managed server or PowerVM VIOS.
Used	The amount of disk space that is allocated and used by to the PowerVM partitions running on the managed server, over the selected time range.
Total Size	The total size of the volume-based pool.

Investigating volume-based storage pools

PowerVM® managed servers and PowerVM virtual I/O servers typically rely on volume-based shared storage pools to efficiently use system resources and increase their overall utilization. When you explore the managed server and PowerVM VIOS **Storage Breakdown** views, you can drill down on a volume-based storage pool to see the logical volumes that belong to it, the PowerVM partition associated with them, their size, and other details, in the **Details Of Volume Based Pool** view. This view also displays general information about the selected disk pool, such as its size, the managing VIOS, and identifies the hard disks on which the containing logical disk partitions. Use this view to find out the sizes of individual volumes and their association with PowerVM partitions, when needed.

To navigate to this view, in the **Selected Service PowerVM** view, select a managed server or a PowerVM VIOS, and click **Explore**. Open the **Storage** tab, and in the **Volume Based Pool Entitled** view, click the title bar of the volume-based pool that you want to investigate.

Figure 32. Details Of Volume Based Pool view



Use this view to find out how the selected object uses the allocated resources, and if there are any patterns indicating their usage levels that may potentially compromise the stability of your infrastructure.

Table 14. Details Of Volume Based Pool view

Pool Summary

Pool Size	The total amount of disk space participating in the volume-based pool.
Used	The total of disk space participating in the volume-based pool that is currently in use.
Available	The total of disk space participating in the volume-based pool that is currently available.
Managed by VIOS	The PowerVM VIOS that is managing the volume-based pool.
Pool make up	A pie chart indicating the amount of disk space each physical disk contributes to the volume-based pool.

Logical Volume

Volume Name	The name of the logical disk.
Volume Size	The size of the logical disk.
Depended by Partition	The name of the PowerVM partition using the logical disk.

BD File

BD File	The name of the Backing Device (BD) file, used for storage virtualization.
File Size	The size of the BD file.
Depended by Partition	The name of the PowerVM partition using the BD file.

Media Files

Media File	The name of the file created by the media repository, used for optical virtualization.
-------------------	--

Table 14. Details Of Volume Based Pool view

File Size	The size of the media file.
Depended by Partition	The name of the PowerVM partition using the media file.

Exploring top resource analyzers

The **Top Analyzer views** appearing on the individual tabs of the managed server details view and PowerVM VIOS details view identify the partitions with the highest resource entitlement, utilization, and usage, and the physical adapters with the highest network throughput. Use them to find out if any of the logical partitions running on the selected managed server or associated with the selected PowerVM VIOS show the signs of processor, memory, or disk space exhaustion or poor utilization, or unusually high network activity. This can help you maintain the stability of your system and ensure optimum resource usage.

To navigate to this tab, in the **Selected Service PowerVM** view, select a managed server or a VIOS, and click **Explore**. Open the **CPU** (managed servers only), **Memory** (managed servers only), **Storage**, or **Network** tab, and open the **Top CPU Analyzer** (managed servers only), **Top Memory Analyzer** (managed servers only), **Top Storage Analyzer**, or **Top Network Analyzer** tab.

Figure 33. Top CPU Analyzer tab

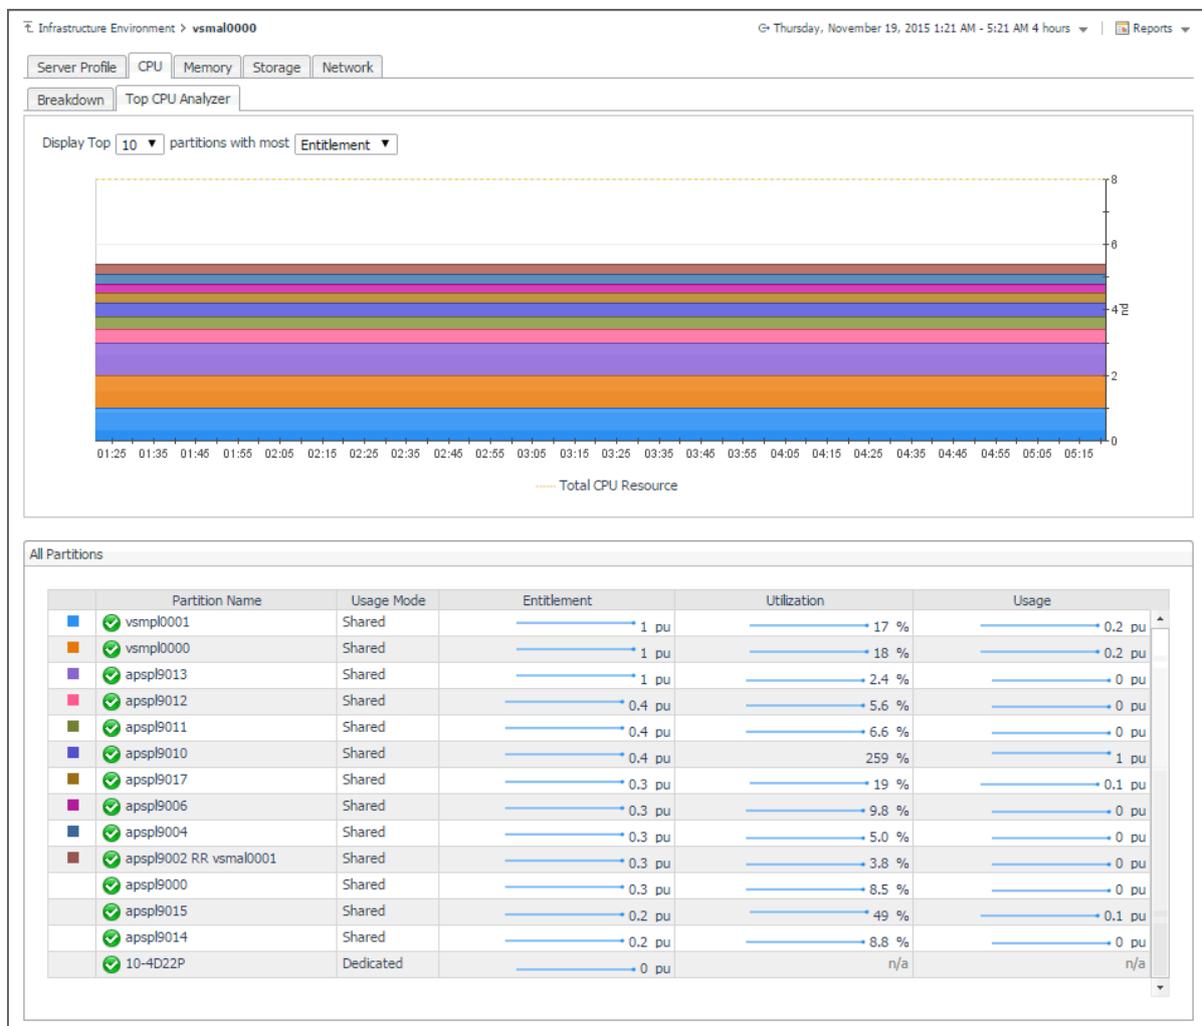


Table 15. Top CPU Analyzer, Top Memory Analyzer, and Top Storage Analyzer tabs

Top CPU Analyzer, Top Memory Analyzer, and Top Storage Analyzer

Display Top *n* partitions with most... Select the number of partitions that you want to display and the metric type that you want to use as a filter.

Top CPU Analyzer chart When **Entitlement** is selected, the **Total CPU Resource** areas in the graph represent the number of processing units allocated to the PowerVM partitions that are associated with the managed server.

When **Utilization** is selected, the areas in the graph represent the percentage of the managed server processor resources allocated to the partitions that are associated with the server and in use over the selected time range.

When **Usage** is selected, the areas in the graph represent the number of processing units allocated to the partitions associated with the managed server that are in use over the selected time range.

Total Memory Analyzer chart When **Entitlement** is selected, the **Total Memory Resource** areas in the graph represent the amount of memory resources allocated to the PowerVM partitions that are associated with the managed server.

When **Utilization** is selected, the areas in the graph represent the percentage of the managed server memory resources allocated to the partitions associated with the server that are in use over the selected time range.

When **Usage** is selected, the areas in the graph represent the amount of memory resources allocated to the partitions associated with the managed server that are in use over the selected time range.

Total Storage Analyzer chart When **Usage** is selected, the areas in the graph represent the disk space allocated to the partitions associated with the managed server or PowerVM VIOS, that are in use over the selected time range.

All Partitions

Partition Name The name of the partition.

Usage Mode **Top CPU Analyzer and Top Memory Analyzer only.** Indicates if the processor or memory resources allocated to the PowerVM partition are participated in shared processor or memory pools (**Shared**), or are dedicated specifically to that partition (**Dedicated**).

Entitlement A spark line indicating the amount of processor, memory, or storage resources that are allocated to the PowerVM partition.

Utilization A spark line indicating the percentage of processor or memory resources allocated to the partition that are in use.

Usage A spark line indicating the amount of processor, memory, or storage resources that are currently in use by the PowerVM partition.

Top Network Analyzer

Top 10 Highest Throughput Physical Adapters The areas in the chart represent the rate at which each identified physical adapter transfers data to and from the network.

All Physical Adapters

Physical Adapter The name of the physical adapter.

Throughput A sparkline indicating the rate at which the physical adapter transfers data to and from the network.

Table 15. Top CPU Analyzer, Top Memory Analyzer, and Top Storage Analyzer tabs

Packets Dropped

A sparkline indicating the rate at which the physical adapter drops data packets while transferring data to and from the network. By default, network adapters send or receive multiple data packets in a single request. When data packets are dropped before reaching their destinations, this may indicate a problem with the network connection or the network adapters.

Errors

A sparkline indicating the rate at which the physical adapter encounters errors while transferring data to and from the network.

Exploring PowerVM VIOS network summaries

The **Summary** tab on the **Network** tab of the PowerVM VIOS detail view provides information all network adapters associated with the selected PowerVM® VIOS, shared ethernet adapters, virtual ethernet adapters, and physical ethernet adapters. Use it to find out if any specific collection of adapters exhibits low throughput, high levels of dropped packets or errors, to maintain the stability of your system and ensure optimum network throughput levels.

To navigate to this tab, in the **Selected Service PowerVM** view, select a PowerVM VIOS, and click **Explore**. Open the **Network** tab, and open the **Summary** tab.

Figure 34. Summary tab

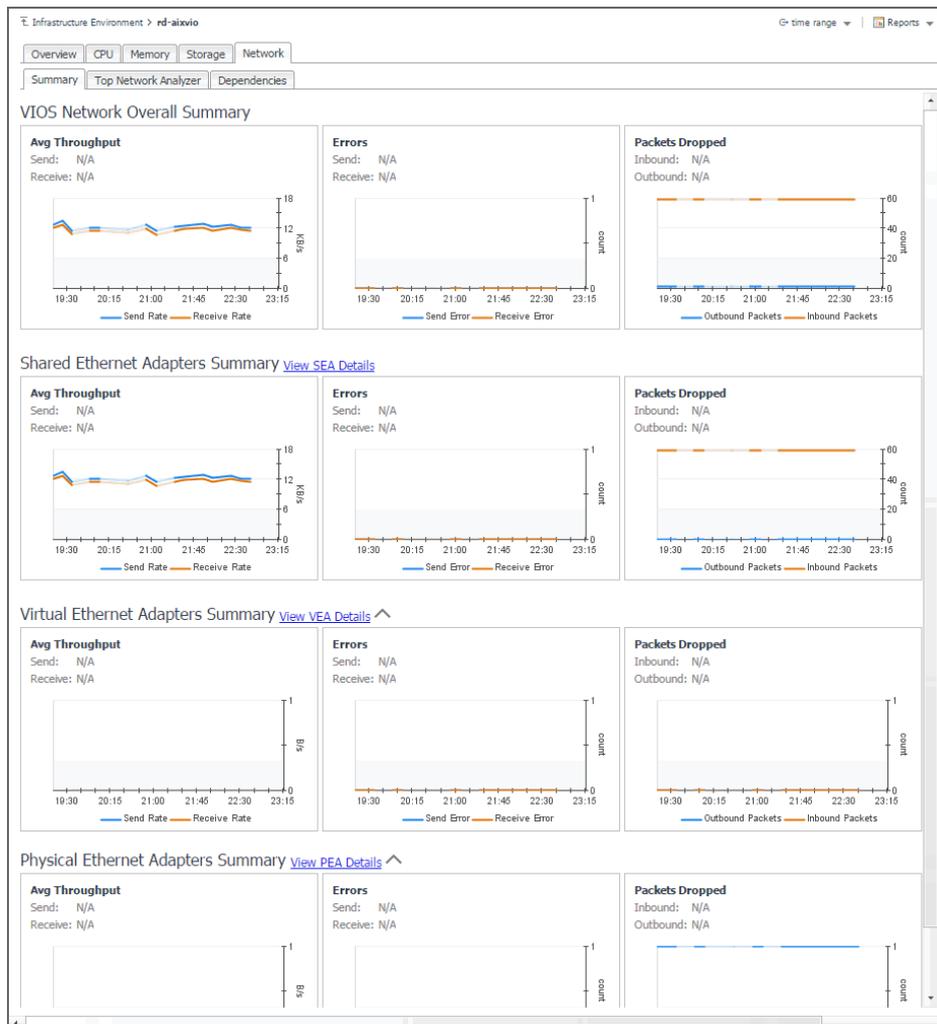


Table 16. Summary view

VIOS Network Overall Summary	Displays information about all network adapters associated with the selected PowerVM VIOS.
Shared Ethernet Adapters Summary	Displays information about shared network adapters associated with the selected PowerVM VIOS.
Virtual Ethernet Adapters Summary	Displays information about virtual network adapters associated with the selected PowerVM VIOS.
Physical Ethernet Adapters Summary	Displays information about physical network adapters associated with the selected PowerVM VIOS.
Avg Throughput	
Send	The current rate at which monitored network adapters send data to the network.
Receive	The current rate at which monitored network adapters receive data from the network.
Send Rate	The average rate at which monitored network adapters send data to the network during the selected time range.
Receive Rate	The average rate at which monitored network adapters receive data from the network during the selected time range.
Errors	
Send	The current rate at which monitored network adapters encounter errors while sending data to the network.
Receive	The current rate at which monitored network adapters encounter errors while receiving data from the network.
Send Error	The average rate at which monitored network adapters encounter errors while sending data to the network during the selected time range.
Receive Error	The average rate at which monitored network adapters encounter errors while receiving data from the network during the selected time range.
Packets Dropped	
Inbound	The current rate at which monitored network adapters drop data packets while sending data to the network. By default, network adapters send or receive multiple data packets in a single request. When data packets are dropped before reaching their destinations, this may indicate a problem with the network connection or the network adapters.
Outbound	The current rate at which monitored network adapters drop data packets while receiving data from the network.
Inbound Packets	The average rate at which monitored network adapters drop data packets while sending data to the network during the selected time range.
Outbound Packets	The average rate at which monitored network adapters drop data packets while receiving data from the network during the selected time range.

Drilling down on shared network adapters

The **Shared Ethernet Adapter Details** view provides information about all shared network adapters associated with the selected PowerVM® VIOS. A shared network adapter enables multiple PowerVM partitions to share one adapter, to optimize the overall use of physical resources. Use it to find out if the adapters exhibits low throughput, high levels of dropped packets or errors, to maintain the stability of your system and ensure optimum network throughput levels. You can also use this view to review the PowerVM partitions that use these adapters, and to review their settings.

To navigate to this view, in the **Selected Service PowerVM**, select a PowerVM VIOS, and click **Explore**. Open the **Network** tab, click **Summary**, and in the **Shared Ethernet Adapters Summary** area, click **View SEA Details**.

Figure 35. Shared Ethernet Adapter Details view

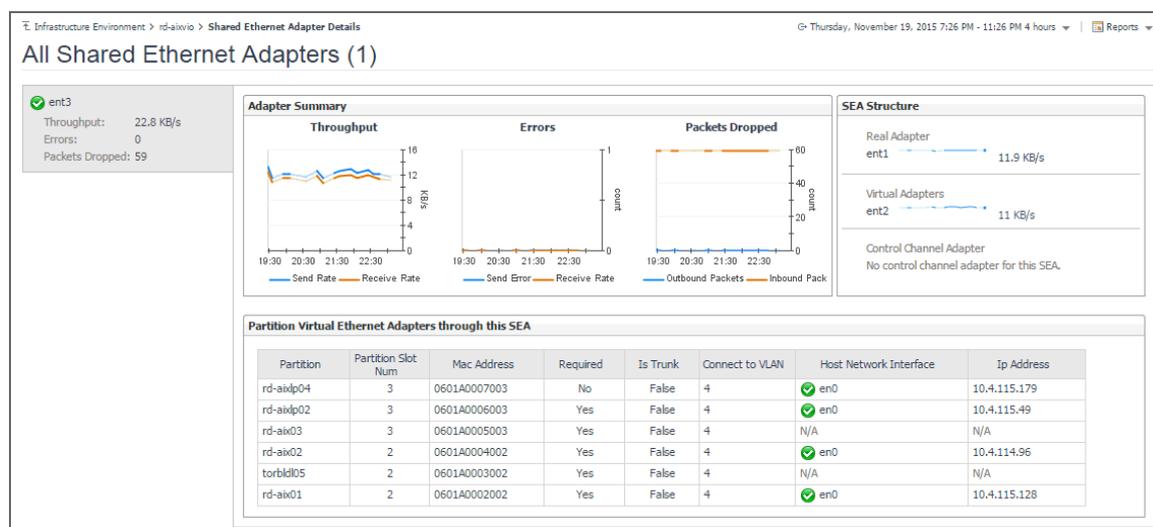


Table 17. Shared Ethernet Adapter Details view

Adapter Summary

Throughput

Send

The average rate at which shared network adapters send data to the network during the selected time range.

Receive

The average rate at which shared network adapters receive data from the network during the selected time range.

Errors

Send Error

The average rate at which shared network encounter errors while sending data to the network during the selected time range.

Receive Error

The average rate at which shared network adapters encounter errors while receiving data from the network during the selected time range.

Packets Dropped

Inbound Packets

The average rate at which shared network adapters drop data packets while sending data to the network during the selected time range. By default, network adapters send or receive multiple data packets in a single request. When data packets are dropped before reaching their destinations, this may indicate a problem with the network connection or the network adapters.

Outbound Packets

The average rate at which shared network adapters drop data packets while receiving data from the network during the selected time range.

SEA Structure

Real Adapter

A sparkline indicating the rate at which physical, virtual, Ether Channel, or control channel network adapters send and receive data from the network

Virtual Adapter

Channel Control Adapter

Partition Virtual Ethernet Adapters through this SEA

Partition

The PowerVM partition using this network adapter.

Partition Slot Num

The slot number associated with the PowerVM partition.

Mac Address

The MAC address of the PowerVM partition.

Table 17. Shared Ethernet Adapter Details view

Required	Indicates whether the I/O slot or virtual I/O adapter is required for the partition. Valid values: <ul style="list-style-type: none"> 0: Required 1: Not required
Is Trunk	Indicates whether the virtual Ethernet adapter is the trunk or uplink adapter for the virtual LAN. Valid values: <ul style="list-style-type: none"> 0: Yes 1: No
Connect to VLAN	The virtual LAN ID.
Host Network Interface	The physical network adapter associated with the PowerVM partition.
Ip Address	The IP address of the PowerVM partition.

Drilling down on physical and virtual network adapters

The **Network - Virtual Ethernet Adapters Detail** and **Network - Physical Ethernet Adapters Detail** views provide information about physical and virtual network adapters associated with the selected PowerVM® VIOS. Use it to find out if the adapters exhibits low throughput, high levels of dropped packets or errors, to maintain the stability of your system and ensure optimum network throughput levels.

To navigate to these views, in the **Selected Service PowerVM** view, select a PowerVM VIOS, and click **Explore**. Open the **Network** tab, and then open the **Summary** tab. In the **Virtual Ethernet Adapters Summary** area, click **View VEA Details**, to see the **Network - Virtual Ethernet Adapters Detail** view. To access the **Network - Physical Ethernet Adapters Detail** view, in the **Physical Ethernet Adapters Summary** area, click **View PEA Details**.

Figure 36. Network - Physical Ethernet Adapters Detail view

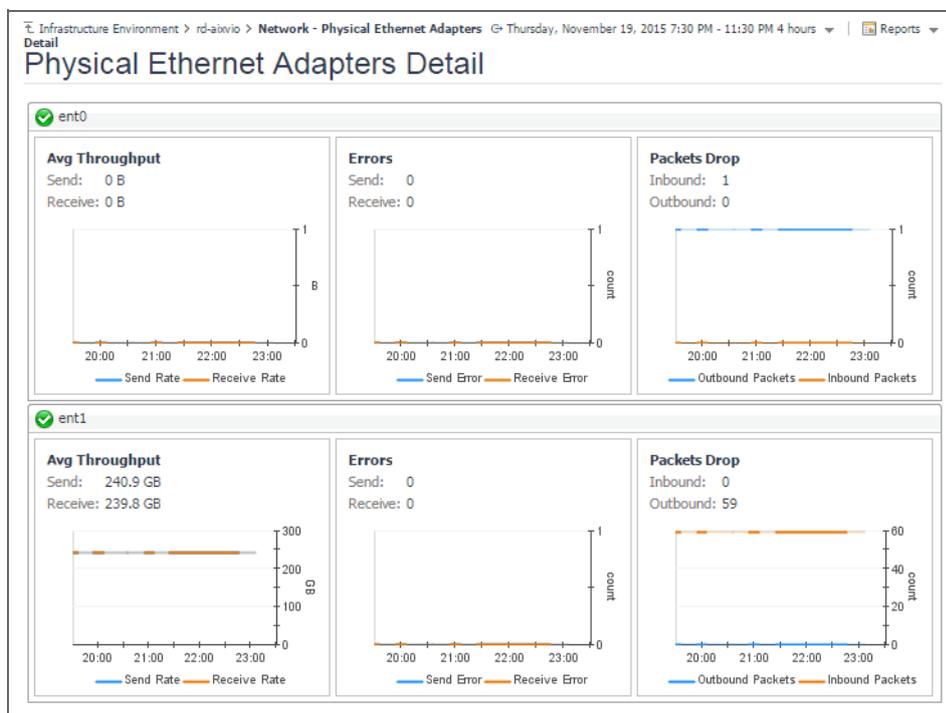


Table 18. Network - Virtual/Physical Ethernet Adapters Detail views

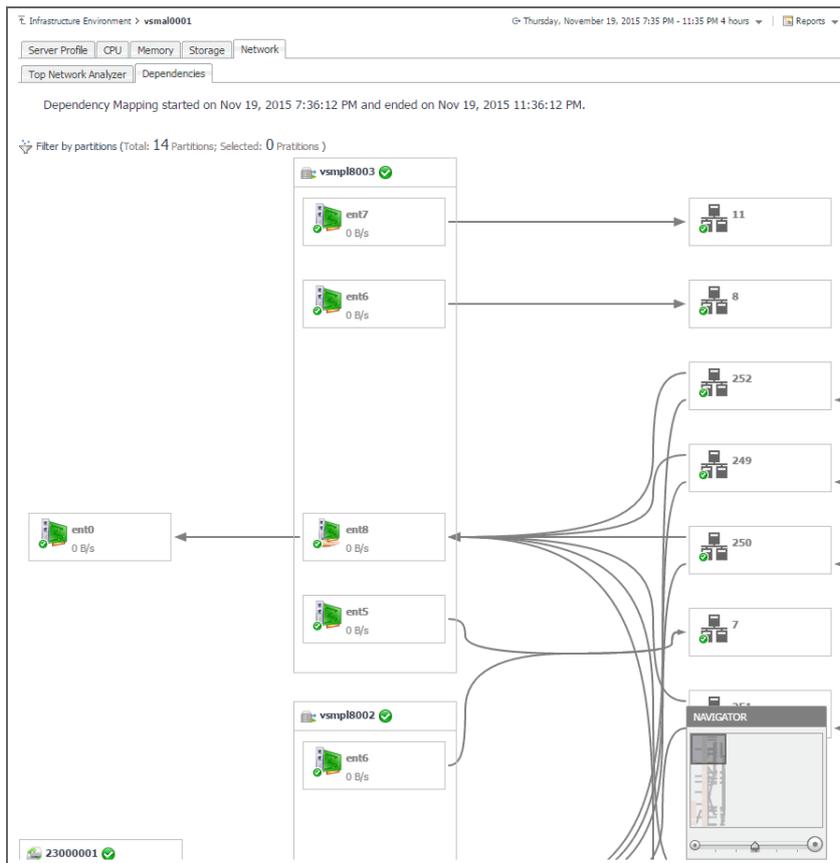
Avg Throughput	
Send	The current rate at which virtual network adapters encounter errors while sending data to the network.
Receive	The current rate at which virtual network adapters encounter errors while receiving data from the network.
Send Rate	The average rate at which virtual network adapters encounter errors while sending data to the network during the selected time range.
Receive Rate	The average rate at which virtual network adapters encounter errors while receiving data from the network during the selected time range.
Errors	
Send	The current rate at which virtual network adapters encounter errors while sending data to the network.
Receive	The current rate at which virtual network adapters encounter errors while receiving data from the network.
Send Error	The average rate at which virtual network adapters encounter errors while sending data to the network during the selected time range.
Receive Error	The average rate at which virtual network adapters encounter errors while receiving data from the network during the selected time range.
Packets Dropped	
Inbound	The current rate at which virtual network adapters drop data packets while sending data to the network. By default, network adapters send or receive multiple data packets in a single request. When data packets are dropped before reaching their destinations, this may indicate a problem with the network connection or the network adapters.
Outbound	The current rate at which virtual network adapters drop data packets while receiving data from the network.
Inbound Packets	The average rate at which virtual network adapters drop data packets while sending data to the network during the selected time range.
Outbound Packets	The average rate at which virtual network adapters drop data packets while receiving data from the network during the selected time range.

Investigating object dependencies

Dependency views visualize the relationships between the selected memory, disk, and network elements and other objects in your integrated environment, through an interactive map. Use these view to better understand resource-related dependencies, to predict the impact a potential outage may have on your environment, and to prevent such events, by reallocating resources where they are most needed.

To navigate to dependency views, in the **Selected Service PowerVM** view, select a PowerVM® partition, managed server, or PowerVM VIOS, and click **Explore**. Open the **Memory** tab (managed servers only), **Storage** tab, or **Network** tab, and then open the **Dependencies** tab.

Figure 37. Dependency tab



A typical PowerVM environment consists of many interrelated components. Understanding the dependencies between logical and virtual components in your monitored environment and the levels of resources they consume allows you to project possible performance bottlenecks that can affect the stability of your system. This can help you predict the impact a potential outage may have on your environment, and to prevent such events, by reallocating resources where they are most needed.

The Dependencies dashboard visualizes the relationships between the objects in your environment through an interactive map. The map illustrates how different components relate to each other, and the levels of the available resources available to them.

To access a Dependency map:

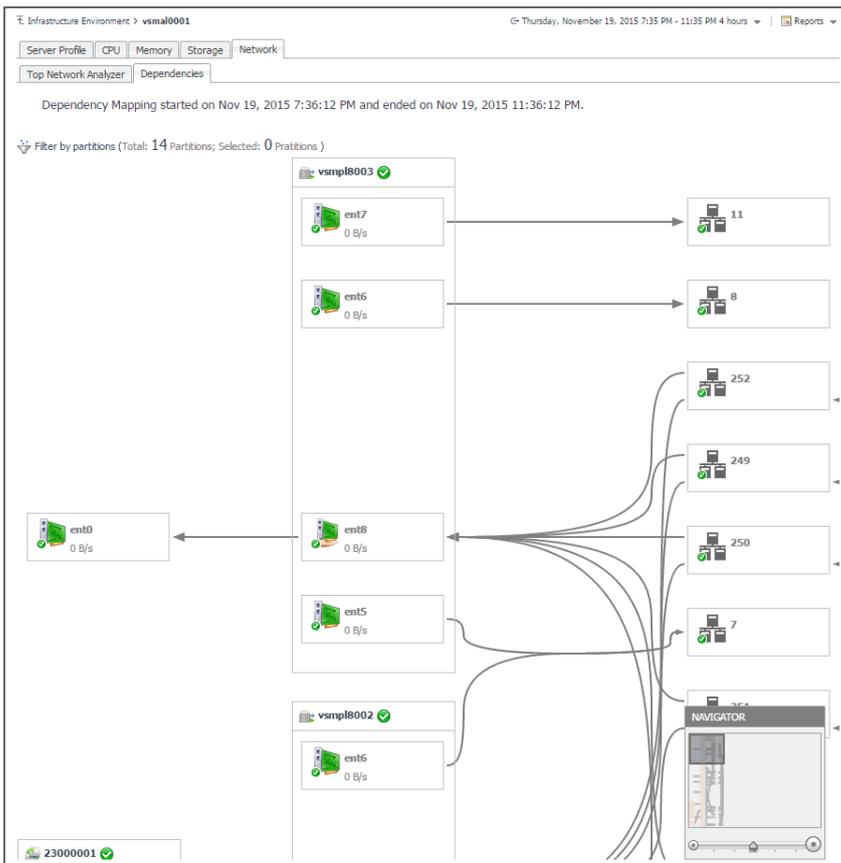
- 1 On the Infrastructure Environment dashboard, on the **Monitoring** tab, select a managed server, partition, or VIOS, and click **Explore**.
- 2 On the Detail view appearing in the display area, open the **Storage** or **Network** tab, and click **Dependency**.

Figure 38. Accessing a dependency map



A Dependency map appears in the display area.

Figure 39. Dependency map



Navigating dependency maps

The complexity of the information appearing in a dependency map depends on the selected object and the dependencies that object has with other objects within your integrated infrastructure.

In a large multi-component environment, dependency maps are likely complex and may not fit your screen. The NAVIGATOR in the top-right corner allows you to easily set the zoom level by dragging the slider into a desired position.

Figure 40. NAVIGATOR



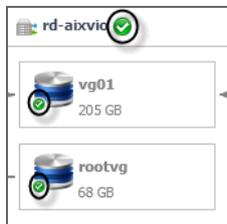
In a dependency map, some objects are represented with container tiles, others with single-object tiles or just icons. Every object is represented with an icon, representing its type. This allows you to quickly identify the elements of a dependency map.

Table 19. Object types appearing in dependency map

	Partition		VIOS
	Host		Host's disk
	Disk pool		Physical disk
	Network adapter		SCSI adapter
	VLAN		

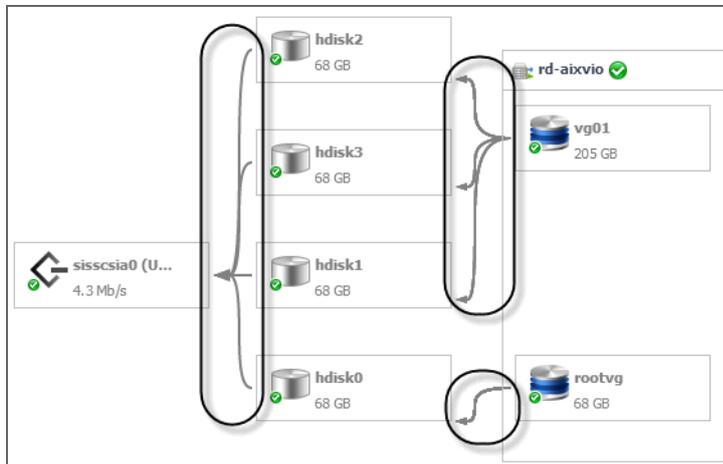
Every object appearing in the map includes an indicator of its health.

Figure 41. Health indicator



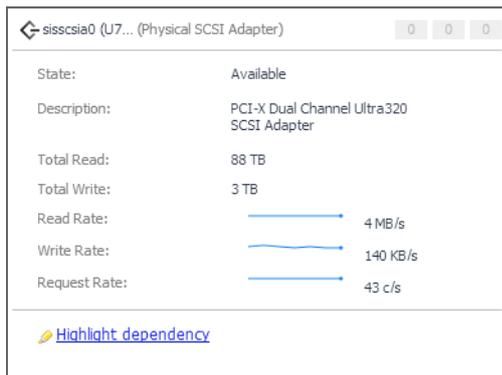
Dependencies between the objects in a map are illustrated with single-directional arrows.

Figure 42. Object dependencies



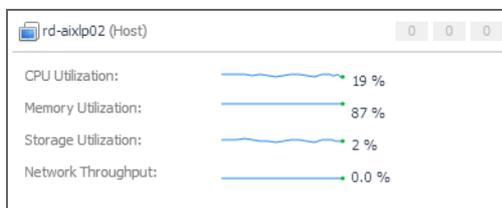
To find out more about an object appearing in the dependency map, click the object icon. A dwell appears, displaying more details about that object. The type and range of information appearing in the dialog box depends on the selected object's type. For example, drilling down on a SCSI Adapter shows the counts of alarms generated against the selected adapter in each severity state, the adapter's state and description, and displays information about the data traffic processed by the adapter.

Figure 43. SCSI Adapter drilldown



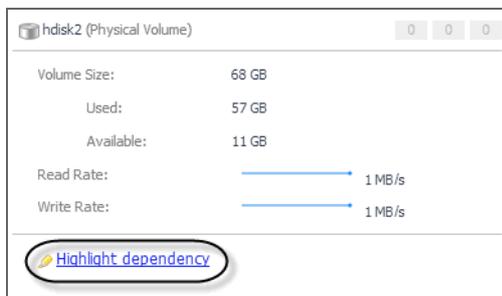
However, if you drill down on a host object the resulting dialog box displays the counts of alarms generated against the selected host in each severity state, and the levels of processor, memory, network, and datastore resources the currently host uses.

Figure 44. Host object drilldown



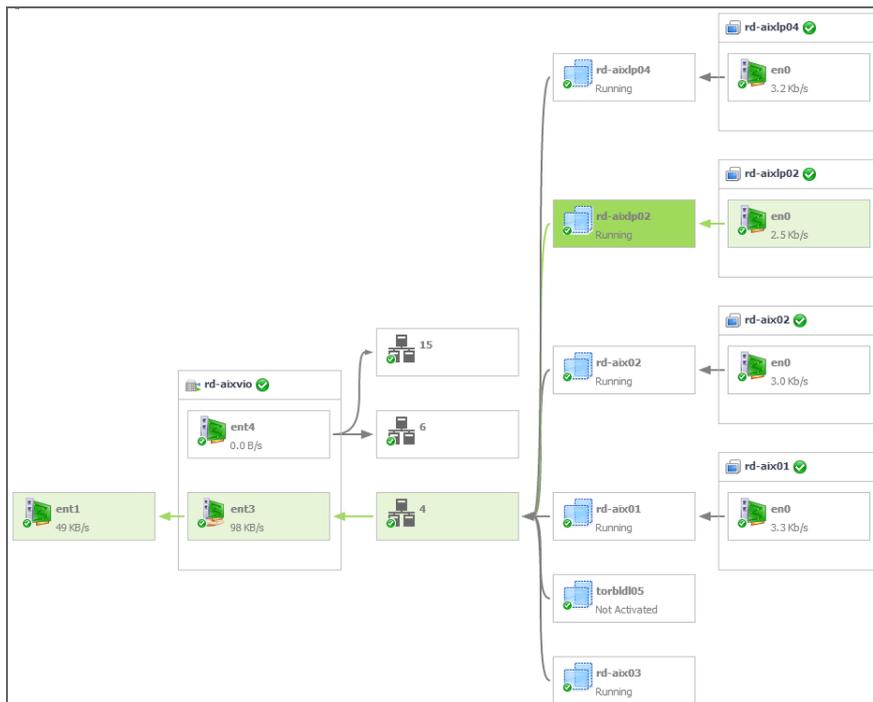
Certain object types allows you to trace down the dependencies an object has with multiple map elements. Use this feature if you want to find out which elements are associated with a particular object whose performance you want to investigate. To do that, click the object on the map, and in the dwell that appears, click **Highlight dependency**.

Figure 45. Highlighting dependencies



This causes all of the related objects and dependencies to appear highlighted in the map.

Figure 46. Highlighted dependencies



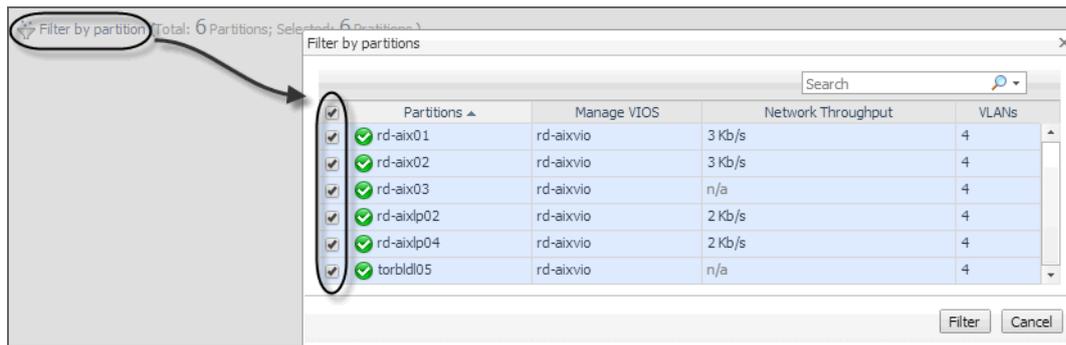
In addition to **Highlight dependency**, some object types offer other options when selected. For example, when you click a partition, you can click **Check partition details** to drill down on the partition, or **Expand partition** or **Collapse partition** for contained adapters to show or hide the network adapters the selected partition uses.

Figure 47. Check partition details and Expand partition links



When you drill down on a managed server or a VIOS, their dependency maps allow you to filter the map by a partition. Doing so allows you to display only those elements associated with one or more desired partitions. This can help you to isolate specific objects in your integrated environment, and to quickly focus your attention on possible bottlenecks.

Figure 48. Filtering objects by partition

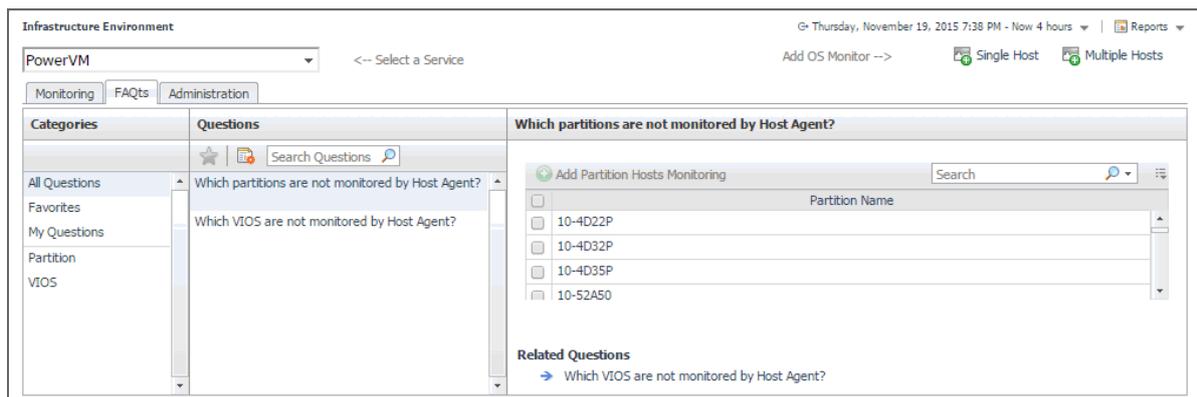


Reviewing frequently asked questions

Foglight for PowerVM™ offers a collection of frequently asked questions that provide quick insight into resource utilization levels for the applications, desktops, user sessions, and the overall infrastructure in your monitored system. The question mechanism is interactive, guiding you to choose a category and specify additional parameters.

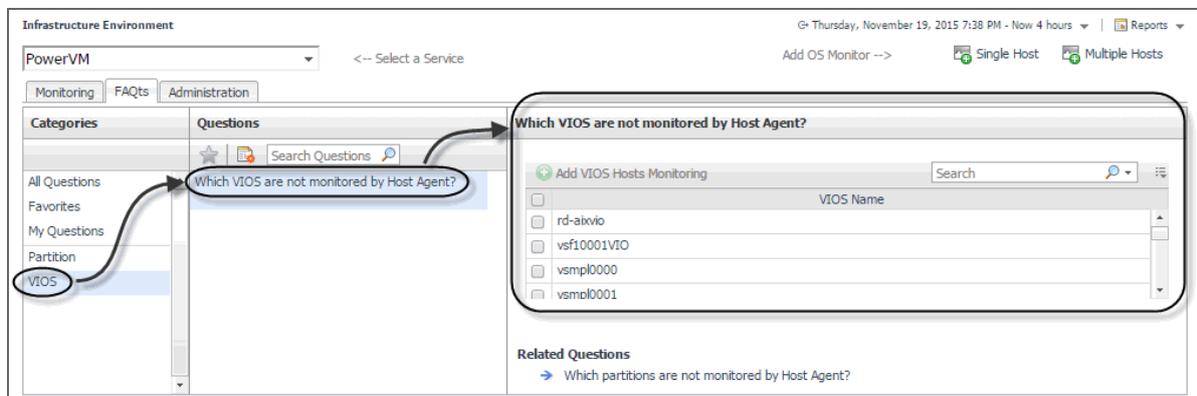
You can find the available questions on the **FAQs** tab of the Infrastructure Environment dashboard.

Figure 49. FAQs tab



On this tab, the **Categories** pane contains several question groups. Selecting a category shows the questions belonging to that category in the **Questions** pane. From there, clicking a question shows the answer on the right.

Figure 50. Finding answers



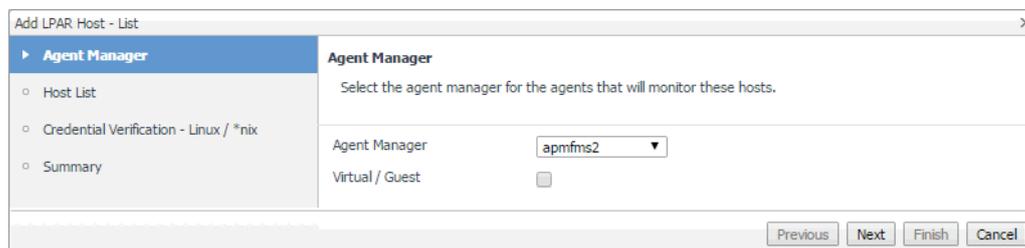
Monitoring a partition or VIOS with Foglight for Infrastructure

If you want to monitor a VIOS or a partition with Foglight for Infrastructure, you can quickly create a Host agent. Foglight for PowerVM allows you to monitor the health of the system infrastructure associated with a monitored VIOS or a partition. For complete information about host agents, see [Using Foglight for Infrastructure agents](#).

To create a Host agent to monitor a VIOS or a partition system infrastructure:

- 1 On the Infrastructure Environment dashboard, ensure that the PowerVM service is selected, and open the **FAQts** tab.
- 2 Select a VIOS or a partition that you want to monitor with Foglight for Infrastructure.
 - To select a VIOS:
 - a On the **FAQts** tab, in the **Categories** pane, select **VIOS** or **All Questions**.
 - b In the **Questions** pane, select **Which VIOS are not monitored by Host Agent**.
 - c In the **Which VIOS are not monitored by Host Agent** pane, select a VIOS, and click **Add VIOS Hosts Monitoring**.
 - To select a partition:
 - a On the **FAQts** tab, in the **Categories** pane, select **Partition** or **All Questions**.
 - b In the **Questions** pane, select **Which partitions are not monitored by Host Agent**.
 - c In the **Which partitions are not monitored by Host Agent** pane, select a partition, and click **Add Partition Hosts Monitoring**.

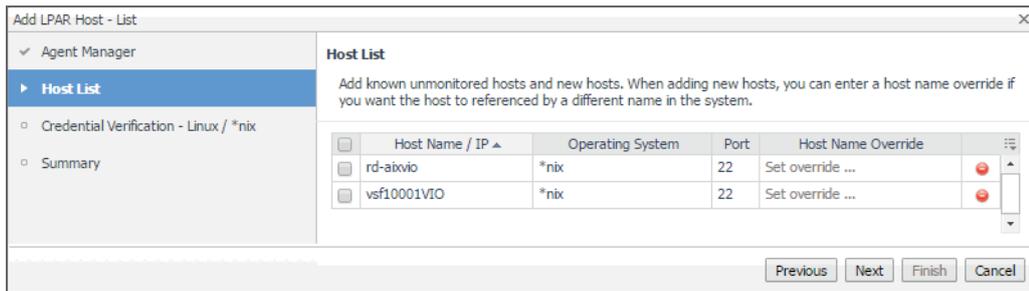
The **Add LPAR Host - List** wizard appears.



The screenshot shows a wizard window titled "Add LPAR Host - List". The left sidebar has a tree view with "Agent Manager" selected. The main area is titled "Agent Manager" and contains the instruction "Select the agent manager for the agents that will monitor these hosts." Below this, there is a dropdown menu for "Agent Manager" with "apmfms2" selected, and a "Virtual / Guest" checkbox which is currently unchecked. At the bottom right, there are four buttons: "Previous", "Next", "Finish", and "Cancel".

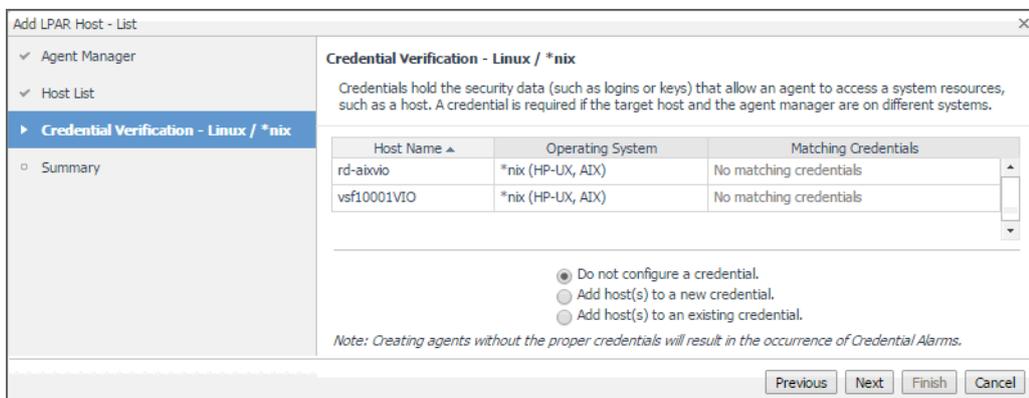
- 3 In the **Add LPAR Host - List** wizard, on the Agent Manager page, select the Agent Manager that you want to associate with the Host agent you are about to create.
 - a Click **Agent Manager** and select the Agent Manager host that you want to use to associate with the Host agent. The list includes only the Foglight Agent Managers to which the Foglight for Infrastructure agent package is already deployed.
 - b Ensure that the **Virtual / Guest** check box is cleared.
 - ⓘ **NOTE:** Ability to monitor virtual hosts is currently implemented only in Foglight for Virtualization, Enterprise Edition. For information about how to configure a virtual agent, refer to the Foglight for Virtualization, Enterprise Edition documentation.
 - c Click **Next**.

The **Add LPAR Host - List** wizard refreshes, showing the **Host List** page.

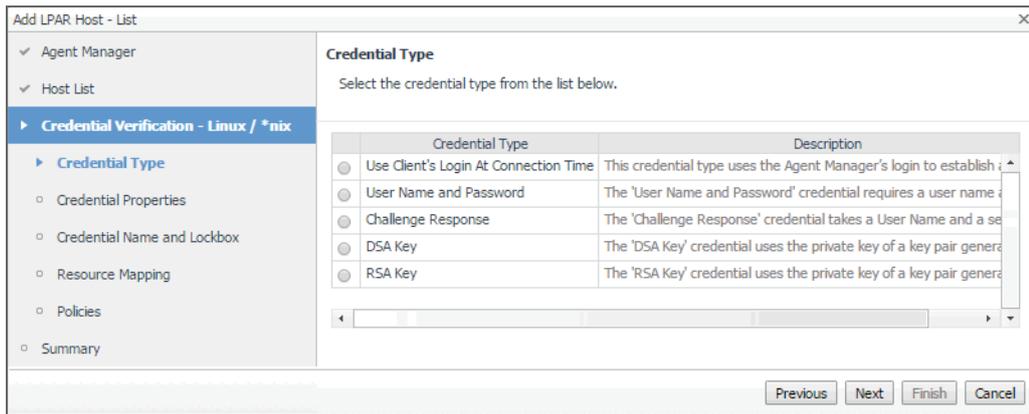


- 4 In the **Add LPAR Host - List** wizard, on the **Host List** page, review the list of VIOS or partition instances that you selected for monitoring.
 - a By default, the VIOS or partition name is used as a host name. If you want to display one or more of these components under a different name in the browser interface and reports, click the **Host Name Override** column, and type a new name when prompted.
 - b Click **Next**.

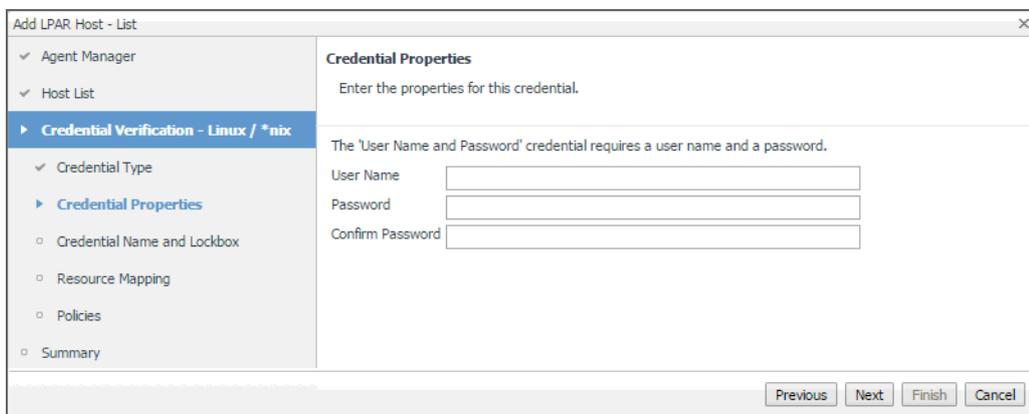
The wizard checks if the Agent Manager has any credentials configured for the selected components, and displays the results on the **Credential Verification** page.



- 5 On the **Credentials Verification** page, select one of the following options:
 - **Do not configure a credential:** Select this option if you want to configure the credential for this component at a later time. Click **Next** and continue to [Step 8](#).
 - **Add host(s) to a new credential:** Select this option if you want to create a new credential for this component. This option is suitable if **none** of the existing credentials have the connection details needed to access the new host. Click **Next** and continue to [Step 6](#).
 - **Add host(s) to an existing credential:** Select this option if you want the Host agent to use an existing credential to access this component. This option is suitable if an existing credential has the security data needed to access this component (such as user name and password), but you need to edit its resource mappings to include this component. Click **Next** and continue to [Step 7](#).
 - For complete information about Foglight credentials, see the *Administration and Configuration Help*.
- 6 Create a new credential.
 - a On the **Credential Type** page that appears, select the desired credential type.

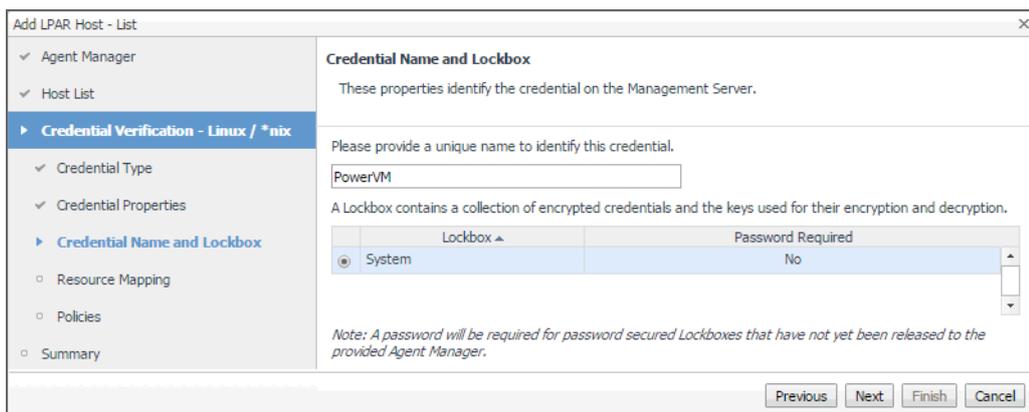


- b On the **Credential Type** page, click **Next**.
The **Credential Properties** page appears.



- c On the **Credential Properties** page, provide the required information. The level of required data depends on the selected credential type. For example, a **User Name and Password** credential needs a user name and a password, while a **Challenge Response** credential requires a user name along with a question/response pair.
- d On the **Credential Properties** page, click **Next**.

The **Credential Name and Lockbox** page appears.

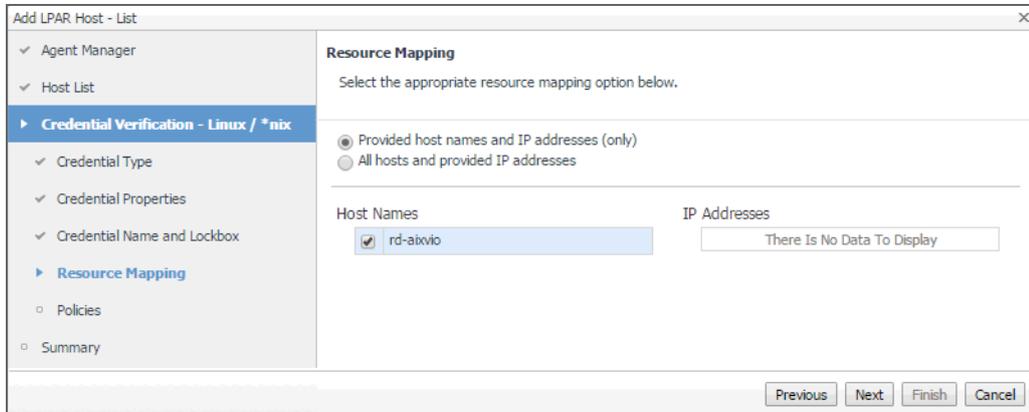


- e On the **Credential Name and Lockbox** page, provide a name to identify the credential, and select a lockbox in which you want to keep the credential. A lockbox can be used to group credentials for access and/or security. The default System lockbox is sufficient for smaller installations.

i | **NOTE:** If a lockbox is password protected and is not released to the Agent Manager, you can provide the lockbox password on the last page of the wizard.

f On the **Credential Name and Lockbox** page, click **Next**.

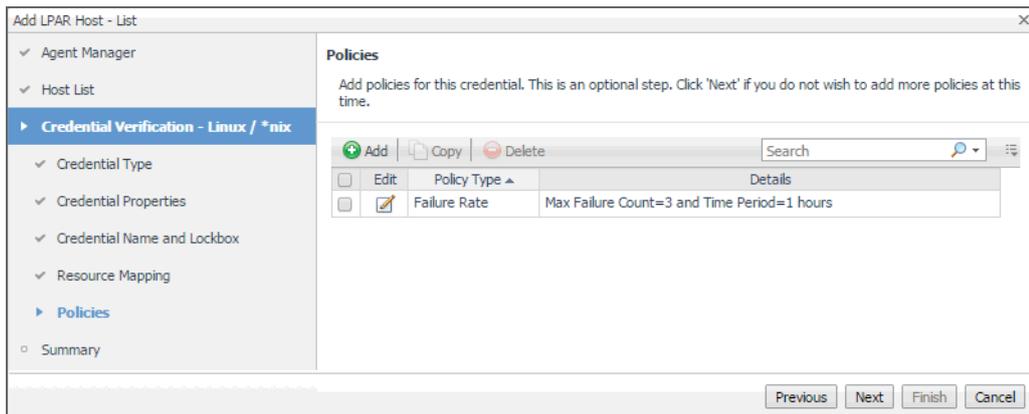
The **Resource Mapping** page appears.



g On the **Resource Mapping** page, ensure that the VIOS or partition components that you want to associate with this credential are selected.

h On the **Resource Mapping** page, click **Next**.

The **Policies** page appears.

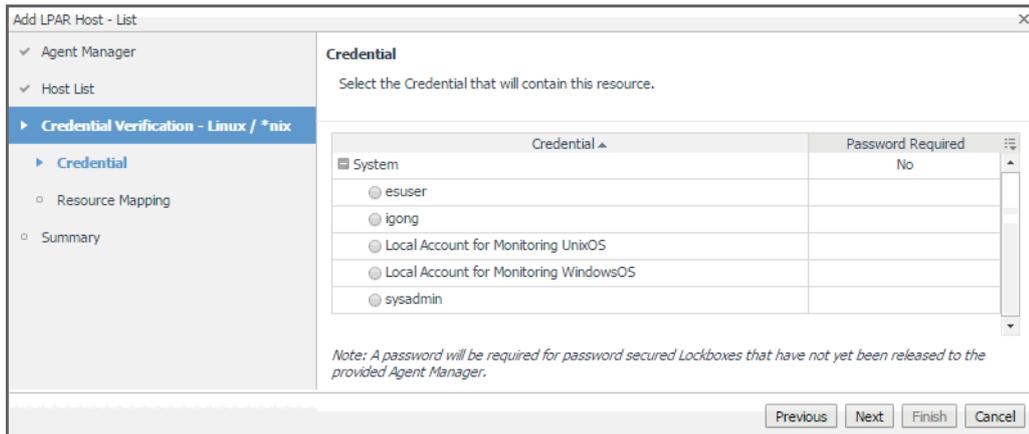


i **Optional**—On the **Policies** page, define one or more policies for this credential. A policy defines the number of times a credential can be used, the number of allowed authentication failures, the time range during which the credential is valid, or the length of time the credential data can be cached on the client. For example, you can specify the number of times the credential can be used, or the time period during which it can be used. For complete information about the available credential policies, see the *Administration and Configuration Help*.

j On the **Policies** page, click **Next**. Continue to [Step 8](#).

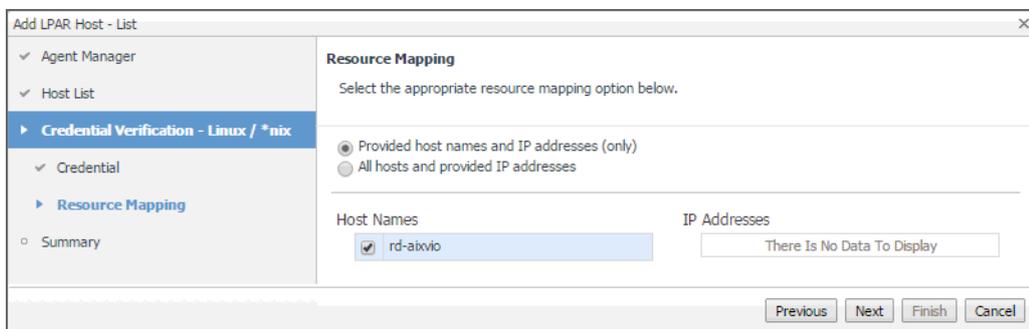
7 Select an existing credential.

a On the **Credential** page that appears, select an existing credential that you want the Host agent to use to access this component.



b On the **Credential** page, click **Next**.

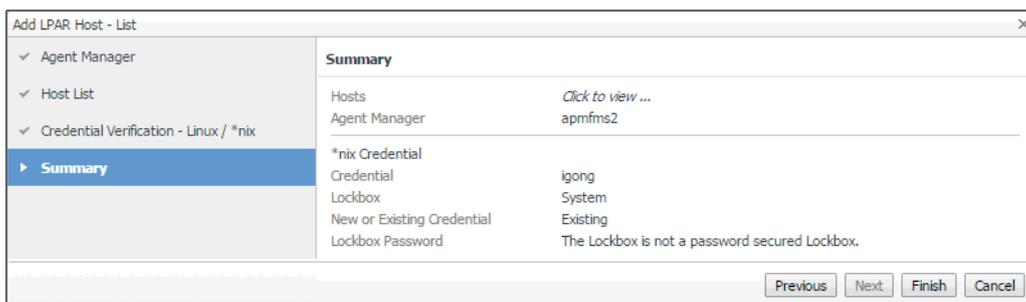
The **Resource Mapping** page appears.



c On the **Resource Mapping** page, ensure that the VIOS or partition components that you want to associate with this credential are selected.

d On the **Resource Mapping** page, click **Next**. Continue to [Step 8](#).

8 On the **Summary** page that appears, review the information about the monitored component, selected Agent Manager, and the credential.



9 Click **Finish**.

After a few moments, a new host appears on the Hosts dashboard and a new instance of the monitoring Host agent is created. The **Agent Creation Successful** dialog box appears. Review the information in this dialog box.

Advanced system configuration and troubleshooting

This chapter contains platform-specific configuration information for configuring Foglight® for Infrastructure on Windows® when using Windows Management Instrumentation (WMI) or Windows Remote Management (WinRM) for remote monitoring access. And how to configure default local user credentials for Infrastructure Agents.

i | **NOTE:** WMI and WinRM are two different mechanisms that monitoring agents can use to establish remote connections. In most scenarios, only one of these mechanisms needs to be configured. The preferred mechanism is WinRM because of WMI scalability limitations.

For more details, see the following topic:

- [Advanced system configuration for WinRM](#)
- [Configuring default local user credentials for Infrastructure Agents](#)

Advanced system configuration for WinRM

Even though the local user will now have access to Windows Remote Management (WinRM), not all performance monitoring classes allow non-administrative users to access their instances. Some performance classes will still need advanced configuration for the non-administrative user to be able to perform queries/execute methods on their object instances. To perform the advanced configuration for non-administrative users, make sure to log into the remote machine using the Administrator account.

The advanced configuration for non-administrative users includes the following steps:

- 1 [Adding a non-administrative user to user groups](#)
- 2 [Setting WinRM RootSDDL for a non-administrative user](#)
- 3 [Granting permission to the namespace](#)
- 4 [Granting permission to the service](#)
- 5 [Additional WinRM configuration in FIPS-compliant mode](#)

Adding a non-administrative user to user groups

The non-administrative users must be added into the Performance Monitor Users group that enables the access to the object instances and the Event Log Readers group that grants the permission for reading event logs.

i | **NOTE:** The non-administrative users should have been added to the Distributed COM Users group by default, which enables the remote connection using Windows Remote Management (WinRM).

To add a non-administrative user to the above user groups:

- 1 Run `lusrmgr.msc`.
 - 2 From the **Local Users and Groups (Local) > Users** list, right click the non-administrative user who you want to add, and then click **Properties** from the context menu.
 - 3 In the *<non-administrative user> Properties* dialog box, click **Member of**.
 - 4 Click **Add**.
- The *Select Groups* dialog box appears.
- 5 In the *Select Groups* dialog box, type *Performance Monitor Users;Event Log Readers* in the *Enter the object names to select* field, and then click **Check Names**.
 - 6 Click OK to return to the *<non-administrative user> Properties* dialog box.
 - 7 Click OK.

The selected user has been added to both user groups.

Setting WinRM RootSDDL for a non-administrative user

After adding a non-administrative user to the needed user groups, you will see the following result if the user permission is not allowed through SDDL.

```
winrm enum wmicimv2/Win32_ComputerSystem -u:<non-administrator user>
WSManFault
    Message = Access is denied.
Error number: -2147024891 0x80070005
Access is denied.
```

To grant the WinRM RootSDDL permission to a non-administrative user:

- 1 Execute the following command to get the RootSDDL value.

```
winrm get winrm/config
```

- 2 You will see the following default RootSDDL value.

```
...
RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
Notes:
O = Owner
G = Primary Group
D = DACL Entries
S = SACL Entries
...
```

- 3 Execute the following command, in order to modify the SDDL permission in DACL Entries and to allow the full control access to the Performance Monitor Users group (which SID is S-1-5-32-558).

i | NOTE: The SID of the Performance Monitor Users group will be varied. Change the SID to the actual SID used in your environment, as needed.

```
winrm set winrm/config/service @{RootSDDL="O:NSG:BAD:P(A;;GA;;;BA)(A;;GA;;;S-1-5-32-558)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)"}
}
```

- 4 You will see the following RootSDDL value.

```
RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)(A;;GA;;;S-1-5-32-558)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
```

Granting permission to the namespace

After setting WinRM RootSDDL for a non-administrative user, log into the remote machine using the administrative user credentials and execute the following command:

```
winrm enum wmicimv2/Win32_ComputerSystem -u:<non-administrator user>
```

You will see the following result, if the non-administrative user does not have the permission to the namespace.

```
winrm enum wmicimv2/Win32_ComputerSystem -u:<non-administrator user>
WSManFault
  Message
    ProviderFault
      WSMANFault
        Message = The WS-Management service cannot process the request.
The WMI service returned an 'access denied' error.
Error number: -2144108284 0x80338104
The WS-Management service cannot process the request. The WMI service returned an
'access denied' error.
```

To grant the permission to the namespace for the current non-administrative user:

- 1 Run `wmimgmt.msc`.
The WMI Control (Local) Properties dialog box appears.
- 2 Right click WMI Control, and click Properties from the context menu.
- 3 Click Security > Root > cimv2, and then click Security.
The Security for ROOT\cimv2 dialog box appears.
- 4 Click Add.
The Select Users, Computers, Service Accounts, or Groups dialog box appears.
- 5 Type Performance Monitor Users;Event Log Readers in the Enter the object names to select field, and then click Check Names.
- 6 Click OK to return the Security for ROOT\cimv2 dialog box.
- 7 Selected the added user group, and then grant all permissions for this user group.
- 8 Click OK to return the WMI Control (Local) Properties dialog box.
- 9 Repeat [Step 3](#) to [Step 8](#) to grant all permissions for Security > Root and Security > Root > mscluster.

Granting permission to the service

Most services data can be collected if the Performance Monitor Users group has been granted the permission to services. If you still fail to collect some service data because of permission issues, on the remote machine, execute the following command where `<service_name>` is the Service name in the Service Property dialog box:

```
sc sdshow <service_name>
```

The permissions of the Built-in Administrator (BA) will appear after executing the above command. Execute the following command to grant the additional permissions, as needed.

```
sc sdset <service_name>...
```

- i** | **NOTE:** Ensure that the Administrators have been granted with sufficient permissions; otherwise you will fail to monitor some services.

The following sample demonstrates how to grant permission to the SCMANAGER service:

- 1 Execute the following command:

```
sc sdshow SCMANAGER
```

- 2 You will see the following result:

```
D: (A;;CC;;;AU) (A;;CCLCRPRC;;;IU) (A;;CCLCRPRC;;;SU) (A;;CCLCRPWPRC;;;SY) (A;;KA;;;BA)
```

- 3 Add "(A;;CCLCRPRC;;;S-1-5-32-558)" between "D:" and "(A;;CC;;;AU)", and execute the `sc sdset` command to grant the permission to the SCMANAGER service. For example, execute the following command:

```
sc sdset SCMANAGER D: (A;;CCLCRPRC;;;S-1-5-32-558) (A;;CC;;;AU) (A;;CCLCRPRC;;;IU) (A;;CCLCRPRC;;;SU) (A;;CCLCRPWPRC;;;SY) (A;;KA;;;BA)
```

Additional WinRM configuration in FIPS-compliant mode

Foglight for Infrastructure WindowsAgent supports to establish Windows Remote Management (WinRM) connections in FIPS-compliant mode. However, when establishing the WinRM connection using Negotiate authentication scheme in FIPS-compliant mode, the password of the credential is required to be no less than 14 characters, otherwise, the connection will be rejected.

If you figured out that your agent's WinRM connection is not working, check below list to see if it is caused by the password length restriction issue:

- Access the Foglight Management Server's **Homes > Alarms** dashboard and check if there is an "Insufficient Password Length" alarm related to your monitored host. See the example as below:

The screenshot shows the Foglight Alarms dashboard. At the top, it says "Alarms" and "This dashboard shows the information of system alarms and changes, and facilitates the investigation of top issues in your environment." Below that, there are tabs for "Alarms", "Alarms Analysis", and "Blackouts". An "Alarm Filter" is set to "Unset". A summary bar indicates "26 Alarm(s)", "19 Error Instance(s)", and "10 Related Host(s)". A table lists various alarms with columns for Severity, Time, Ack'd, Cleared, and Host. One alarm is highlighted, and its details are shown in a pop-up window. The details include: Host (n/a), Instance (zhvm-fog-2332.domain.fglam (fglam) ...), Origin (By System), Agent (n/a), and Agent Type (Default Drilldown). The message states: "Unable to establish WinRM connection to http://zhvm-fog-2332.domain.fglam:5985 using Negotiate authentication due to insufficient password length. The password for the account [domain:DOMAIN.FGLAM, username:13bit] must be at least 14 characters." There are also buttons for "Acknowledge", "Acknowledge Until Normal", "Clear", "Find Historic Occurrences", and "Cancel".

- Or check the logs of Foglight Agent Manager where the WindowsAgent instance is deployed and see if there are log messages as below.

```
2019-12-23 10:49:09.713 WARN [WinRMWebServiceManagementConnection[25]-0] com.quest.glue.core.remoteconnection.negotiate.SPNEGOAuthenticationScheme - Authentication failed: unable to log in to https://hostname...:5986
```

```
(org.bouncycastle.crypto.fips.FipsUnapprovedOperationError: password must be at least 112 bits
```

If either of above cases exists, you need to change your monitored host password to no less than 14 characters in order to successfully monitor your host with WinRM connection using Negotiate authentication scheme in FIPS-compliant mode.

Configuring default local user credentials for Infrastructure Agents

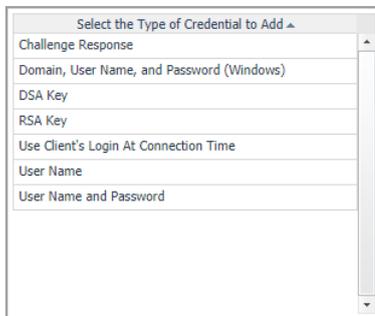
To monitor target hosts, Infrastructure Agent instances require user credentials to get access to the system resources. If the Foglight Agent Manager is physically located on the host being monitored, default local user credentials are automatically created while deploying the Infrastructure cartridge.

However, under certain circumstances, such as the default lockbox cannot be found in the server, the cartridge deploying process will not create local user credentials. Then, if you want to do local monitoring with Infrastructure Agents, you will have to manually add the credentials to the server.

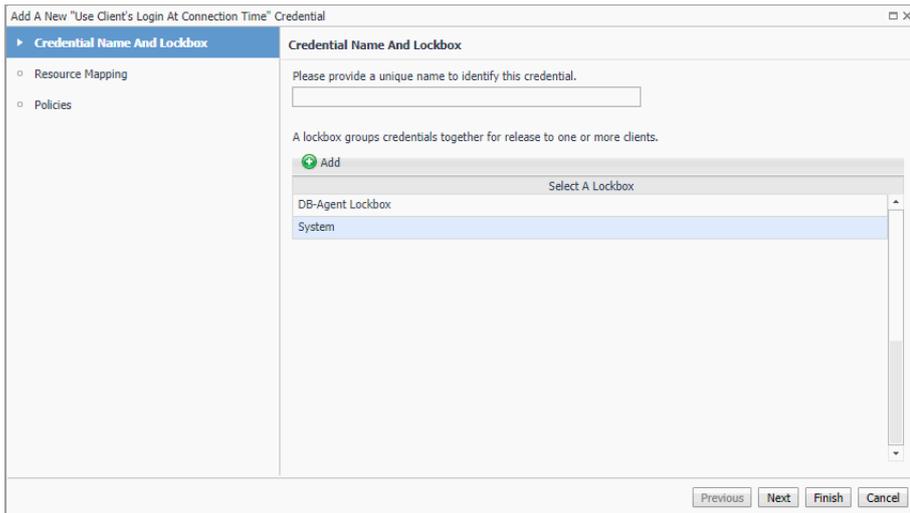
To create default local user credential for monitoring local host:

- 1 Log in to the Foglight browser interface.
- 2 On the navigation panel, under **Dashboards**, click **Administration > Credentials**.
- 3 On the *Credentials* page that appears in the display area, click **Manage Credentials**.
- 4 On the *Manage Credentials* dashboard that appears in the display area, click **Add**.

The *Select the Type of Credential to Add* list appears.



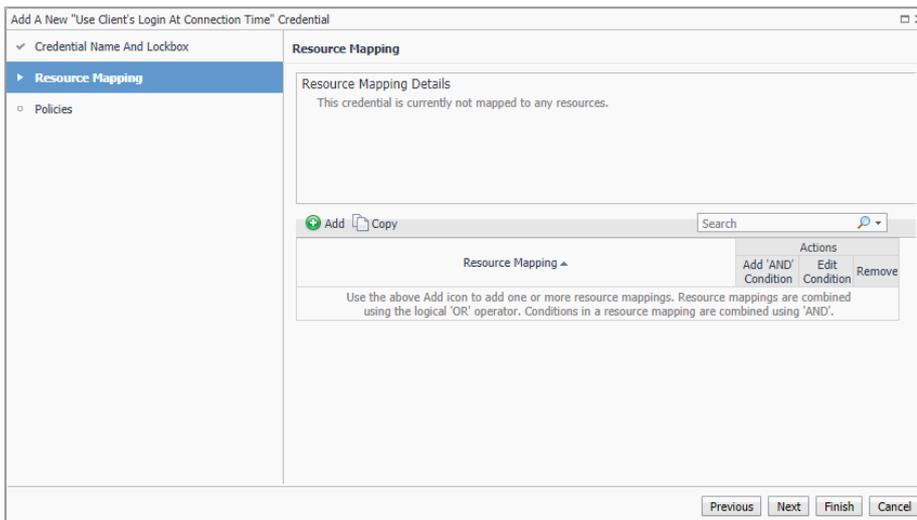
- 5 In the *Select the Type of Credential to Add* list, click *User Client's Login At Connection Time*.
The *Add A New "User Client's Login At Connection Time" Credential* wizard appears.



- 6 On the *Credential Name and Lockbox* page, select the lockbox in which you want to store the local user credential, and specify a unique credential name. Click **Next**.

i | **TIP:** If you cannot find a suitable lockbox in the list, click **Add** to create a new one.

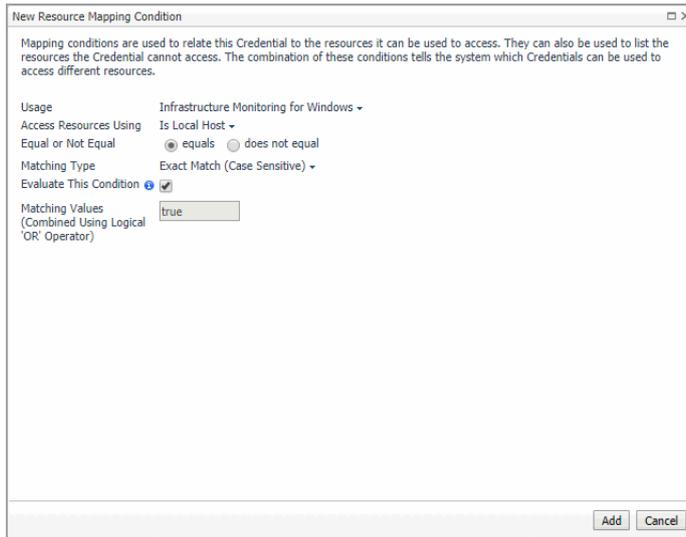
The *Resource Mapping* page appears.



- 7 Specify mapping conditions of the credential.

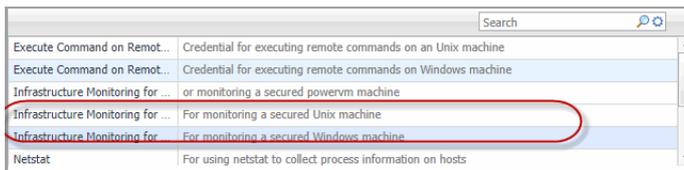
- a On the *Resource Mapping* page, click **Add**.

The *New Resource Mapping Condition* dialog box appears.

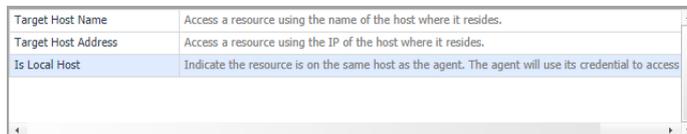


b Select the appropriate Usage according to your operating system. To open the **Usage** drop-down list, click the down-facing arrow on the right. Select either of the following:

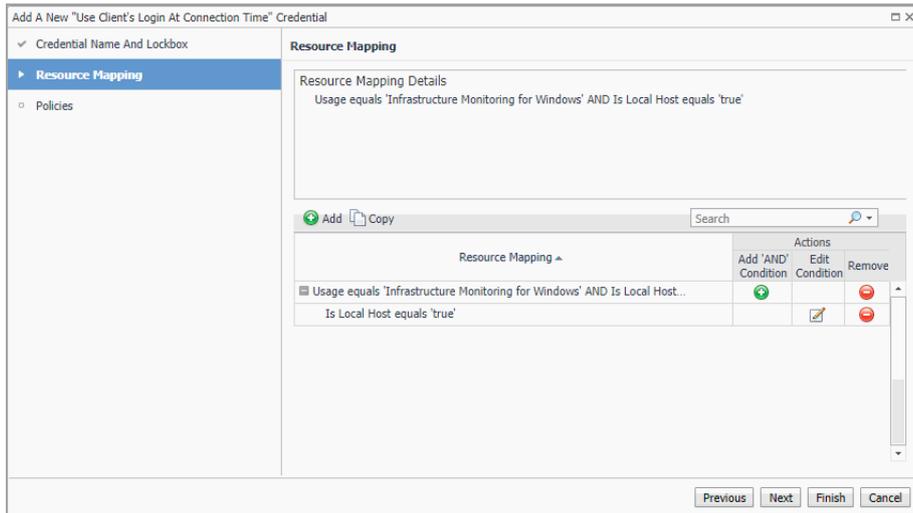
- **Local Unix machine** - Infrastructure Monitoring for Unix
- **Local Windows machine** - Infrastructure Monitoring for Windows



c Click the down-facing arrow on the right to open the **Access Resources Using** drop-down list. Select the **Is Local Host**.



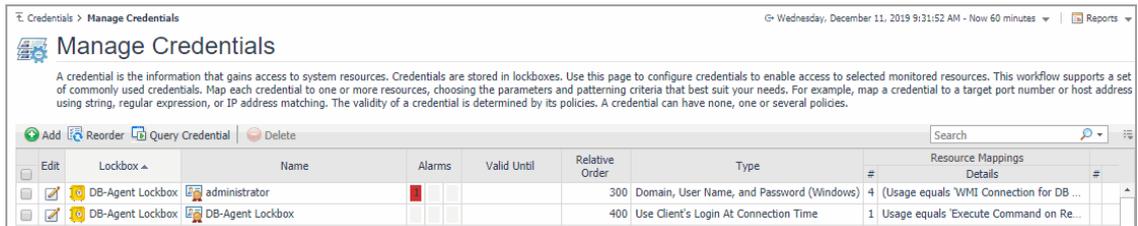
- d Ensure that **equals** and **Evaluate This Condition** are selected. Click **Add**.
- e The *New Resource Mapping Condition* dialog box closes and the *Resource Mapping* page refreshes. The newly specified resource mapping is displayed.



- f *Optional*—you can refine your credential settings. For example, to specify the time during which the credential is valid, the number of failed attempts after which the credential will be locked, the number of times the credential can be used, or the period of time during which the credential data is cached on the server. For complete information, see the *Administration and Configuration Help*.

8 Click **Finish**.

The *Add A New "User Name and Password" Credential* wizard closes and the *Manage Credentials* dashboard refreshes. The newly added local user credential is displayed in the list.



Reference

Foglight displays monitoring data in views that group, format, and display data. The main types are described below.

Dashboards are top-level views that contain lower-level views. The dashboards supplied with Foglight, as well as those created by users, are accessible from the navigation panel.

Lower-level views in Foglight can be added to dashboards or can be accessed by drilling down from a dashboard. They receive and display data directly from the Management Server or from other views. Some views filter or select data that appears in other views in the same dashboard. Some are tree views with expandable nodes for selecting servers, applications, or data.

For more details, see these topics:

- [Foglight for Infrastructure views](#)
- [Foglight Log Monitor views](#)
- [Rules](#)
- [Metrics](#)

Foglight for Infrastructure views

Foglight for Infrastructure ships with several predefined views, to help you monitor your infrastructure environment:

- [Quick View](#)
- [Summary view](#)
- [Resource Utilizations view](#)
- [Host Monitor views](#)
- [Processes views](#)

Quick View

Purpose

Display the list of objects associated with a selected domain and a summary of the domain's host status.

It includes the following embedded views:

- **Host List view:** displays the list of objects associated with the selected resource object.
- **FAQ Mini Viewer:** displays the list of questions relevant to the selected object.
- **Summary view:** displays the utilization summary for all monitored hosts in the selected object.
- **Resource Utilizations view:** displays the resource utilization values for the selected object or group of objects.

How to get here

When you click a tile in the Infrastructure Environment dashboard, the Quick view (at the bottom of the tab) is refreshed to display the information relevant to the selected tile.

Summary view

Purpose

This view displays the utilization summary for infrastructure resource objects.

How to get here

The *Summary* view is the default view displayed when you select a tile in the Monitoring tab.

- i** | **TIP:** If the *Resource Utilizations* view is displayed for a selected monitored host, you can return to displaying the *Summary* view by clicking **All <object_type>** in the Host List view.

Description of embedded views

The Summary view is view is made up of the following embedded views:

- [Top CPU Utilization](#)
- [Top Network Utilization](#)
- [Top Memory Utilization](#)
- [Top Disks Utilization](#)
- [Alarms](#)

From the Summary view, you can drill down into the All Hosts view by clicking the **Explore**  icon in the upper-right corner of the view.

Top CPU Utilization

Table 1. Top CPU Utilization view

Description	Displays the top CPU consumers among the hosts monitored for the selected domain.
Data Displayed	<ul style="list-style-type: none">• Health: host's health state severity (Normal, Warning, Critical, or Fatal).• Name: host's name.• Current: the percentage of the host's CPU utilization.
Where to Go Next	To drill down into the Metric Analyzer dashboard and view the metrics collected for this HostCPUs object, click on the chart.

Top Network Utilization

Table 2. Top Network Utilization view

Description	Displays the top consumers of network resources among the hosts monitored for the selected domain.
--------------------	--

Table 2. Top Network Utilization view

Data Displayed	<ul style="list-style-type: none"> • Health: host's health state severity (Normal, Warning, Critical, or Fatal). • Name: host's name. • Current: the percentage of the host's network utilization.
Where to Go Next	To drill down into the Metric Analyzer dashboard and view the metrics collected for this HostNetwork object, click on the chart.

Top Memory Utilization

Table 3. Top Memory Utilization view

Description	Displays the top memory consumers among the hosts monitored for the selected domain.
Data Displayed	<ul style="list-style-type: none"> • Health: host's health state severity (Normal, Warning, Critical, or Fatal). • Name: host's name. • Current: the percentage of the host's memory utilization.
Where to Go Next	To drill down into the Metric Analyzer dashboard and view the metrics collected for this Memory object, click on the chart.

Top Disks Utilization

Table 4. Top Disks Utilization view

Description	Displays the top disk space consumers among the hosts monitored for the selected domain.
Data Displayed	<ul style="list-style-type: none"> • Health: host's health state severity (Normal, Warning, Critical, or Fatal). • Name: host's name. • Current: the percentage of the host's disk space utilization.
Where to Go Next	To drill down into the Metric Analyzer dashboard and view the metrics collected for this HostStorage object, click on the chart.

Alarms

Table 5. Domain Alarms view

Description	Displays the list of alarms fired for all hosts in the selected domain.
Data Displayed	<ul style="list-style-type: none"> • Sev: alarm's state severity (Normal, Warning, Critical, or Fatal). • Host: host's name. • Title: alarm title, as generated by the Alarm Detail dialog. • Ack'ed: set to <i>true</i> or <i>false</i>. • Ack'ed Time: time when the alarm was acknowledged. • Description: alarm description, as generated by the Alarm Detail dialog. • Time: time when the alarm was fired. • Instance: instance's name. • Origin: source of the alarm. • Agent: agent's name. • Agent Type: agent's type.
Where to Go Next	To drill down into the Alarm view for a specific host, click the alarm name in the alarms table.

Resource Utilizations view

Purpose

This view displays the utilization summary for a monitored host.

How to get here

In the Host List view, click the name of the host.

Description of embedded views

This view is made up of the following embedded views:

- [CPU](#)
- [Memory](#)
- [Storage](#)
- [Network](#)
- [Other Alarms](#)

From the *Resource Utilizations* view you can drill down into the selected Host view by clicking the **Explore**  button in the upper-right corner of the view.

CPU

Table 6. CPU view

Description	Displays the CPU utilization incurred by the monitored host (as a percentage of the total CPU) and the RunQueueLength.
Data Displayed	<ul style="list-style-type: none">• Alarms: number of alarms generated for this host and their severity (Warning, Critical, or Fatal).• Chart: the percentage of the host's CPU utilization for the interval indicated in the History field.
Where to Go Next	<p>To drill down into the Alarm view for this monitored host, click a specific alarm icon.</p> <p>To drill down into the Metric Analyzer dashboard and view the metrics collected for a topology object, click the data series on the chart.</p> <p>To drill down into the Processes dashboard and view the process details, click the <i>CPU Usage by Process</i> link at the bottom of the view.</p>

Memory

Table 7. Memory view

Description	Displays the memory utilization incurred by the monitored host, as a percentage of the total memory.
--------------------	--

Table 7. Memory view

Data Displayed	<ul style="list-style-type: none"> • Alarms: number of alarms generated for this host and their severity (Warning, Critical, or Fatal). • Chart: the percentage of the host's memory utilization for the interval indicated in the History field.
Where to Go Next	<p>To drill down into the Alarm view for this monitored host, click a specific alarm icon.</p> <p>To drill down into the Metric Analyzer dashboard and view the metrics collected for a topology object, click the data series on the chart.</p> <p>To drill down into the Processes dashboard and view the process details, click the <i>Memory Usage by Process</i> link at the bottom of the view.</p>

Storage

Table 8. Storage view

Description	Displays the disk space utilization incurred by the monitored host (as a percentage of the total disk space) and the DiskTransferRate (for UnixAgents only).
Data Displayed	<ul style="list-style-type: none"> • Alarms: number of alarms generated for this host and their severity (Warning, Critical, or Fatal). • Chart: the percentage of the host's disk utilization for the interval indicated in the History field.
Where to Go Next	<p>To drill down into the Alarm view for this monitored host, click a specific alarm icon.</p> <p>To drill down into the Metric Analyzer dashboard and view the metrics collected for a topology object, click the data series on the chart.</p>

Network

Table 9. Network view

Description	Displays the network utilization incurred by the monitored host, as a percentage of the total network.
Data Displayed	<ul style="list-style-type: none"> • Alarms: number of alarms generated for this host and their severity (Warning, Critical, or Fatal). • Chart: the percentage of the host's network utilization for the interval indicated in the History field.
Where to Go Next	<p>To drill down into the Alarm view for this monitored host, click a specific alarm icon.</p> <p>To drill down into the Metric Analyzer dashboard and view the metrics collected for a topology object, click the data series on the chart.</p>

Other Alarms

Table 10. Other Alarms view

Description	Displays all alarms on the selected host, other than CPU, Memory, Network, and Storage alarms.
Data Displayed	<ul style="list-style-type: none">• Sev: alarm's state severity (Normal, Warning, Critical, or Fatal).• Title: alarm title, as generated by the Alarm Detail dialog.• Ack'ed: time when the alarm was acknowledged.• Instance: instance's name.• Origin: source of the alarm.• Description: alarm description, as generated by the Alarm Detail dialog.• Time: time when the alarm was fired.• Agent: agent's name.• Agent Type: agent's type.
Where to Go Next	To drill down into the Alarm view for this monitored host, click a specific alarm on the list.

Host Monitor views

Purpose

These views display CPU, memory, and other system metrics, along with a series of views connected with graphical flows, illustrating the selected host's performance. Explore these views when you want to investigate a selected host's statistics in real time.

i **TIP:** The value of Processor Queue Length metric is color-coded to indicate the threshold severity. For generic information about metrics' thresholds, see "Working with Metric Thresholds" in the Administration and Configuration Help.

TIP: The value of Available Paging Space metric is color-coded to indicate the threshold severity. For generic information about metrics' thresholds, see "Working with Metric Thresholds" in the Administration and Configuration Help.

How to get here

- 1 On the Infrastructure Environment dashboard, click the tile reflecting the domain of the host that you want to drill down to. For example, to drill down on a Windows® host, click the **Windows** tile.
- 2 In the Quick View, in the left pane, select the host.
- 3 In the top-right corner of the **Resource Utilizations** view, click **Explore**.
- 4 Click the **Monitor** tab on the top-left corner.

The Host Monitor views appear in the display area.

i **TIP:** This view also appears when you drill down on a monitored host from the Hosts dashboard. For more information about this dashboard, see the *Foglight User Help*.

Description of embedded views

This view is made up of the following embedded views:

- [CPU](#)
- [Data Flows](#)

- [Disks](#)
- [Memory](#)
- [Network](#)
- [Paging Files](#)
- [System](#)

CPU

Table 11. CPU view

Description	Displays the aggregate processor utilization percentage and the number of engines configured. Displays information about number of processors, processor queue length, the number of processes and top CPU consumers.
Data Displayed	<ul style="list-style-type: none"> • Number of Processes: The number of processes that are waiting in the run queue. • Processor Queue Length: The number of processes that are either running or waiting in the queue. A high number of processes in the run queue means the CPU is busy. A consistently high number can indicate that host needs more CPU power. • Total CPU Usage: The percentage of time the CPU executes system code and user programs. This time includes both system and user time.
Where to Go Next	To drill down to CPU details, click any of the available metrics. For complete information about the CPU details that appear, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i> .

Data Flows

- i** | **NOTE:** UnixAgents that monitor AIX hosts will not return data on Disk Reads or Disk Writes for logical disks. As a result, when you drill down for more information by clicking on the data flow view for Disk Reads or Disk Writes, the third column of charts, Logical_Disk Reads/Writes, will always show 'There is No Data to Display' for AIX hosts.

Table 12. Data Flows view

Description	Displays a collection of graphical flows, indicating how quickly the data is moving through the monitored host.
Data Displayed	<ul style="list-style-type: none"> • Context Switches. The rate of switches from one thread to another. Thread switches can occur inside a single process or across processes. Context switches occur because one thread requests information from another or one thread is preempted by a higher priority thread. • Disk Reads. The rate at which the system reads data to disk. The unit of this value can be either of the following: c/s (count per second), k/s (thousand per second), m/s (million per second), g/s (billion per second), or t/s (trillion per second). • Disk Writes. The rate at which the system writes data to disk. The unit of this value can be either of the following: c/s (count per second), k/s (thousand per second), m/s (million per second), g/s (billion per second), or t/s (trillion per second). • Page In. The number of pages that the system reads from the disk every second, to resolve memory references to pages that were not in memory at the time of the reference. This counter includes paging traffic on behalf of the system cache to access file data for applications. Observe this rate if you are concerned about excessive memory pressure (called thrashing) and the excessive paging that might result. • Page Out. The number of pages that the system writes to the disk every second, after they are modified in the memory. • Total Receive. The rate at which the system is receiving data from the network, measured in megabytes per second. It includes all application data as well as network protocol information, such as packet headers. • Total Send. The rate at which the system is sending data to the network, measured in megabytes per second. It includes all application data as well as network protocol information, such as packet headers.
Where to Go Next	<p>Drill down on:</p> <ul style="list-style-type: none"> • Disk Reads or Disk Writes. Links to the Disk details view. • Page In or Page Out. Links to the Memory details view. • Total Receive or Total Send. Links to the Network details view. <p>For more information about these views, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i>.</p>

Disks

- i** | **NOTE:** UnixAgents that monitor AIX hosts will not return data on Disk Reads or Disk Writes for logical disks. As a result, when you drill down for more information by clicking on the Disks view, the third column of charts, Logical_Disk Reads/Writes, will always show 'There is No Data to Display' for AIX hosts.

Table 13. Disks view

Description	Displays the utilization of physical disks, the number of logical disks, and identifies the top logical disk consumers.
--------------------	---

Table 13. Disks view

Data Displayed	<ul style="list-style-type: none"> • Available. The amount of the logical disk space that is available for use. • Logical Disks. The number of logical disks used by the selected host. • Physical Disks. The number of physical disks used by the selected host. • Total. The total amount of the logical disk space, including available and used space. • Utilization. The percentage of time spent on servicing physical disk requests. • Used Space. The amount of the logical disk space that is currently in use.
Where to Go Next	<p>Drill down on Used Space. Links to the Disk details view.</p> <p>Drill down on Utilization. Links to the Storage Details view.</p> <p>NOTE: The Foglight for Infrastructure agent only supports collecting the “Available” metric for ZFS[®] mounted as Loopback File Systems (LOFS), due to the limitation of the available commands within the zone.</p> <p>The Disk Name was updated to include “(Loopback ZFS) after the disk name, to identify the Loopback ZFS. The Logical Disks table now includes two new columns, including “Available” and “Filesystem Type”.</p> <p>For more information about these views, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i>.</p>

Memory

Table 14. Memory view

Description	Displays the total amount of memory usage for the host, the top memory consumers and the top I/O consumers.
Data Displayed	<ul style="list-style-type: none"> • Utilization. The percentage of memory that is currently in use. • Utilization History. The host’s memory utilization over the selected time range. • Available. Free memory of the host. • Total. Total memory of the host.
Where to Go Next	<p>Drill down on Utilization or Utilization History. Links to the Memory details view.</p> <p>For more information about this view, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i>.</p>

Network

Table 15. Network view

Description	Displays information about the existing network interfaces and network utilization in general.
Data Displayed	<ul style="list-style-type: none"> • Bandwidth. The data transfer rate provided to the existing network interfaces. • Interface Utilization. The bandwidth of each network interface in bytes per second during the selected time range. • Network Utilization. The network utilization incurred by the monitored device, as a percentage of the total network utilization. • Packets Received. The number of inbound data packets that the device transmits with success over the selected time range. • Packets Sent. The number of outbound data packets that the device transmits with success over the selected time range.
Where to Go Next	<p>Drill down on Network Utilization. Links to the Network details view. For more information about this view, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i>.</p>

Paging Files

Table 16. Paging Files view

Description	This view displays the percentage of paging space that is currently available.
Where to Go Next	Drill down on Available Paging Space . Links to the Memory details view. For more information about this view, see “Viewing A Host’s Performance” in the <i>Foglight User Guide</i> .

System

Table 17. System view

Description	This view shows the number of alarms associated with the monitored device.
Data Displayed	<ul style="list-style-type: none">• Alarms. The counts of Fatal, Critical, and Warning alarms. Click a number to display the Outstanding Alarm(s) dialog box that lists the alarms. For more information about alarms in Foglight, see the <i>Foglight User Help</i>.

Processes views

Purpose

The *Processes* dashboard provides a summary of CPU and memory usage by process, as well as a table of process details. If a host does not have process data (for example, because it is monitored by a virtual agent which does not collect process data), then it is possible to create an agent to monitor processes directly from this drill-down view.

How to get here

- 1 On the Infrastructure Environment dashboard, click the tile reflecting the domain of the host that you want to drill down to. For example, to drill down on a Windows® host, click the **Windows** tile.
- 2 Do one of the following:
 - In the Monitoring tab > CPU View, click **CPU Usage by Process**.
 - In the Monitoring tab > Memory View, click **Memory Usage by Process**.
 - In the Monitoring tab > Storage View, click **Disk IO by Process**.
 - Click **Explore** on the top-right corner of **Resource Utilization**, and then click **Processes** on the top-left corner of the display area.

The Processes dashboard is displayed.

Description of embedded views

This dashboard is made up of the following embedded views:

- [CPU At A Glance](#)
- [CPU Usage by Process](#)
- [Memory At A Glance](#)
- [Memory Usage by Process](#)
- [Monitored Processes](#)
- [User Defined Processes \(Process Availability Config\)](#)

CPU At A Glance

Table 18. CPU At A Glance view

	This view displays how the CPU resources are being used based on the following:
Description	<ul style="list-style-type: none">• The top five processes using the most CPU. These five processes are the “top consumers” of CPU for the host.• The top five processes that have had the largest recent change in CPU usage.
Data Displayed	<ul style="list-style-type: none">• AVG CPU. The average amount of CPU resources spent by the process.• Change/hr. The change in the usage of CPU resources.• Process Name. The name of the process.• Trend. A sparkline indicating the change in the usage of CPU resources.

CPU Usage by Process

Table 19. CPU Usage by Process view

Description	This view displays the cumulative CPU usage by processes over time. It shows the total process load on the selected system.
--------------------	---

Memory At A Glance

Table 20. Memory At A Glance view

	This view displays how the memory resources are being used based on the following:
Description	<ul style="list-style-type: none">• The top five processes using the most memory. These five processes are the “top consumers” of memory for the host.• The top five processes that have had the largest recent change in memory usage.
Data Displayed	<ul style="list-style-type: none">• AVG Memory. The average amount of memory resources spent by the process.• Change/hr. The change in the usage of memory resources.• Process Name. The name of the process.• Trend. A sparkline indicating the change in the usage of memory resources.

Memory Usage by Process

Table 21. Memory Usage by Process view

Description	This view displays the memory usage by processes over time. It shows the load on memory for processes on the selected system.
--------------------	---

Monitored Processes

Table 22. Monitored Processes view

Description	This view lists all running processes on a given host. The list is sortable, filterable, and allows you to do process searches.
Data Displayed	<ul style="list-style-type: none">• Process Name. The name of the process.• Instances. The number of running process instances.• CPU Usage, Average. The average percentage of CPU resources that the process is using.• CPU Usage, Time. The average amount of time the process is running.• Memory, Average. The average percentage of memory resources that the process is using.• Memory, Swap Size (Avg). The average amount of memory swap space this process has in one or more paging files.• Memory, Working Set Size. The amount of memory currently allocated to this process.• Memory, Virtual Size (Avg). The current size of the virtual address space of this process.

User Defined Processes (Process Availability Config)

Table 23. User Defined Processes (Process Availability Config) view

Description	This view lists all user-defined processes running on a given host. The list is sortable, filterable, and allows you to do process searches. NOTE: Make sure to set Report only aggregate process metrics to False, in order to collect all details about process instances.
Data Displayed	<ul style="list-style-type: none">• Process Name. The name of the process.• Command Line. The Command Line works with the Process Name to filter out process instances, as needed.• Expected Process Count. The number of expected process instances.• Matched Process Count. The number of running process instances that match the user-defined process name and command line.

Foglight Log Monitor views

Foglight Log Monitor allows you to monitor and configure your virtual environment. For more information about the Log Monitor dashboard, see [Investigating log records](#).

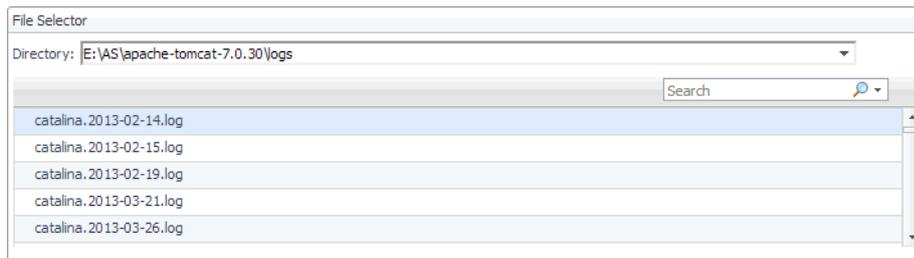
The Log Monitor dashboard contains the following views:

- [File Selector view](#)
- [Log Records view](#)
- [Record Details view](#)

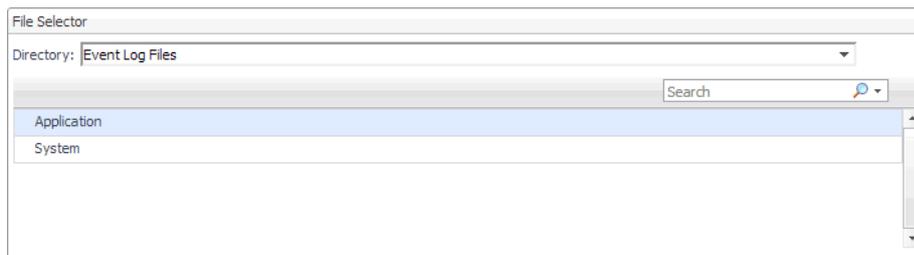
File Selector view

The **File Selector** view allows you to select a log. You have an option to select a file log monitored by the File Log Monitor Agent, or a Windows Event Log, monitored by the Windows Event Log Monitor Agent.

Selecting a file log



Selecting a Windows Event Log



How to Get Here

- This view appears in the upper part of the Log Monitor dashboard.

Description of the View

Table 24. File Selector view

- Data displayed**
- Directory.** Lists the directories containing monitored log files, and the monitored Windows Event Logs.

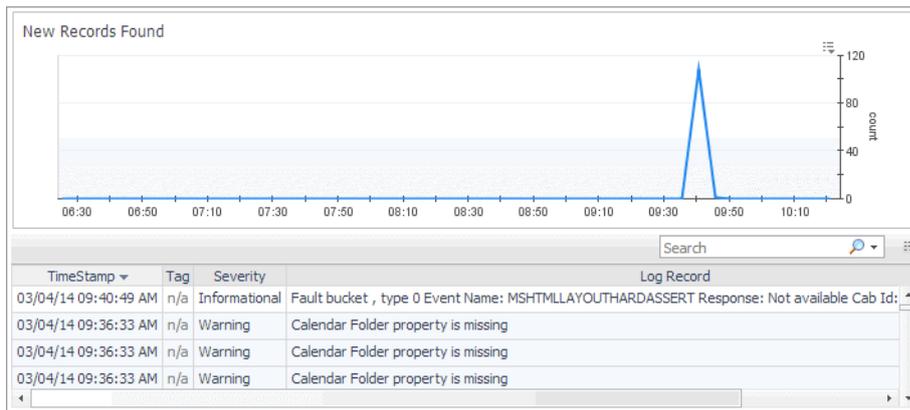
Log Records view

The **Log Records** view lists the records available in the selected log and shows the growth of log records over time. Use this view to review the contents of your logs, and look for signs that indicate potential issues. The layout of this view depends on the selected log type.

File logs



Windows Event Logs



How to Get Here

- On the Log Monitor dashboard, in the File Selector view, select a log whose records you want to display.

Description of Embedded Views

This view is made up of the following embedded views:

- [Description of File Size Increase view](#)
- [Description of Log Records table view](#)
- [Description of New Records Found \(file logs only\) view](#)
- [Description of New Records Found \(Windows Event Logs only\) view](#)

File Size Increase (file logs only)

Table 25. Description of File Size Increase view

Description Shows the increase in the size of the monitored log file over the selected time range.

Log Records table (file logs and Windows Event Logs)

Table 26. Description of Log Records table view

Description Lists the records that exist in the selected log.

Data displayed

- **Log Record.** The message contained in the log record.
- **Severity.** The severity of the log message.
- **Tag.** The tag associated with the record, as specified in the agent properties.
- **TimeStamp.** The date and time when the log record is created.

New Records Found (file logs only)

Table 27. Description of New Records Found (file logs only) view

Description Shows the increase in the number of log records in the selected log file over the selected time range.

New Records Found (Windows Event Logs only)

Table 28. Description of New Records Found (Windows Event Logs only) view

Description	Shows the increase in the number of log records in the selected Windows Event Log over the selected time range.
--------------------	---

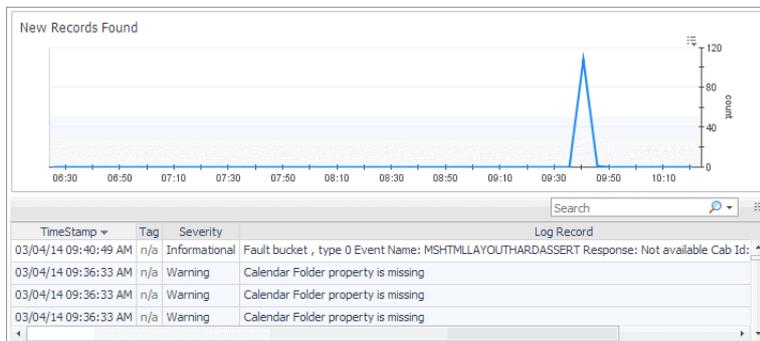
Record Details view

The **Record Details** view displays the contents of a selected record. The layout of this view depends on the selected log type. Windows Event Log records appear in a table, while file log records appear in text form.

File log record contents



Windows Event Log record contents



How to Get Here

- 1 On the Log Monitor dashboard, in the File Selector view, select a log whose records you want to display.
- 2 In the [Log Records view](#), select a log record.

The **Record Details** view refreshes, showing the contents of the selected log record.

Description of the View (log file records)

Table 29. Record Details view (log file records)

Description	Displays the contents of the selected record.
--------------------	---

Description of the View (Windows Event Log records)

Table 30. Record Details view (Windows Event Log records)

Data displayed	<ul style="list-style-type: none"> • Message. The message contained in the Windows Event Log record. • Severity. The severity of the Windows Event Log record. • TimeStamp. The date and time when the Windows Event Log record is created. • Computer. The name of the computer on which the Windows Event Log record is created. • ID. The ID of the Windows Event Log record.
-----------------------	--

Table 30. Record Details view (Windows Event Log records)

- **Source.** The application that generated the Windows Event Log record.
- **User.** The Foglight severity of the Windows Event Log record.

Rules

Foglight for Infrastructure includes some built-in rules that monitor the health of your infrastructure environment. You can and should modify rules to better suit your environment.

To modify Foglight for Infrastructure rule threshold values and alarm template settings, refer to “Viewing, Creating, and Managing Alarm Templates” in the *Foglight User Guide*.

i **NOTE:** Foglight 6.0 introduces alarm templates to gather alarm rules into a domain-specific template that is easily modified and applied to targets.

Avoid editing Foglight for Infrastructure rules in the **Administration > Rules & Notifications > Rule Management** dashboard. Default rules may be modified during regular software updates and your edits will be lost. Always use the Alarm Templates dashboard.

In the Rules Management dashboard, all Foglight rules are displayed in the list. You can restrict the list to show only the rules installed with a cartridge. You can also search the list for a single rule or a group of rules.

To review rules:

- 1 In the navigation panel, under **Homes**, click **Administration > Rules & Notifications > Rules**.
The Rules dashboard opens.
- 2 From the Cartridge list, select *Infrastructure*.
The dashboard refreshes to display only the Foglight for Infrastructure rules.
- 3 From here, you can perform the following tasks:
 - Review a short description of the rule.
 - Associate actions with rules.
 - Create user-defined rules.

For help with these tasks, open the online help from the Rules dashboard.

Metrics

Foglight for Infrastructure agents collect metrics and store them in the Foglight topology model as topology objects:

- [AIX metrics](#)
- [HP-UX metrics](#)
- [Linux metrics](#)
- [Solaris metrics](#)
- [Windows metrics](#)

AIX metrics

The following topology objects are collected by the UnixAgent when monitoring AIX® systems:

- AIXHostDetails topology object
- AIXMemoryDetails topology object
- AIXPhysicalCPUUsage topology object
- AIXUsedMemoryDetails topology object
- CPUCounts topology object
- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- HostTopProcessEntry topology object
- LPARNetworkInterfaceDetails topology object
- LPARPhysicalDiskDetails topology object
- LogicalDisk topology object
- Memory topology object
- NetworkInterface topology object
- NetworkInterfaceDetails topology object
- OperatingSystem topology object
- PhysicalDisk topology object
- Processor topology object

i | **NOTE:** Some of these metrics can only be collected if the native agent has been deployed to the monitored host.

The following topology objects may be collected by the UnixAgent when monitoring AIX® systems, when users set the property “Collect Hypervisor metrics” for VIOS host monitoring:

- PowerVMVIOS topology object
- PowerVMVIOSCPUs topology object
- PowerVMVIOSMemory topology object
- PowerVMVIOSNetwork topology object
- PowerVMVIOSStorage topology object

HP-UX metrics

The following topology objects are collected by the UnixAgent when monitoring HP-UX systems:

- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- HostTopProcessEntry topology object
- LogicalDisk topology object

- Memory topology object
- NetworkInterface topology object
- NetworkInterfaceDetails topology object
- OperatingSystem topology object
- PhysicalDisk topology object
- Processor topology object

i | **NOTE:** Many of these metrics can only be collected if the native agent has been deployed to the monitored host.

The following topology objects may be collected by the UnixAgent when monitoring HP-UX systems, when users set the property “Collect Hypervisor metrics” for VIOS host monitoring:

- PowerVMVIOS topology object
- PowerVMVIOSCPUs topology object
- PowerVMVIOSMemory topology object
- PowerVMVIOSNetwork topology object
- PowerVMVIOSStorage topology object

Linux metrics

The following topology objects are collected by the UnixAgentPlus when monitoring Linux® systems:

- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- HostTopProcessEntry topology object
- LogicalDisk topology object
- Memory topology object
- NetworkInterface topology object
- NetworkInterfaceDetails topology object
- NixHostProcessInstance topology object
- OperatingSystem topology object
- PhysicalDisk topology object
- Processor topology object

The following topology objects are collected by the UnixAgent when monitoring Linux® systems:

- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- LogicalDisk topology object

- Memory topology object
- NetworkInterface topology object
- NetworkInterfaceDetails topology object
- OperatingSystem topology object
- PhysicalDisk topology object
- Processor topology object

The following topology objects may be collected by the UnixAgentPlus and UnixAgent when monitoring Linux® systems, when users set the property “Collect Hypervisor metrics” for VIOS host monitoring:

- PowerVMVIOS topology object
- PowerVMVIOSCPUs topology object
- PowerVMVIOSMemory topology object
- PowerVMVIOSNetwork topology object
- PowerVMVIOSStorage topology object

Solaris metrics

The following topology objects are collected by the UnixAgentPlus when monitoring Sun® Solaris® systems:

- CPUCounts topology object
- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- HostTopProcessEntry topology object
- LogicalDisk topology object
- Memory topology object
- NetworkInterface topology object
- NetworkInterfaceDetails topology object
- OperatingSystem topology object
- PhysicalDisk topology object
- Processor topology object

The following topology objects are collected by the UnixAgent when monitoring Sun® Solaris® systems:

- Host topology object
- HostCPUs topology object
- HostNetwork topology object
- HostProcess topology object
- HostProcessInstance topology object
- HostTopProcessEntry topology object
- LogicalDisk topology object
- Memory topology object

- [NetworkInterface](#) topology object
- [NetworkInterfaceDetails](#) topology object
- [OperatingSystem](#) topology object
- [PhysicalDisk](#) topology object
- [Processor](#) topology object

The following topology objects may be collected by the UnixAgentPlus and UnixAgent when monitoring Sun® Solaris® systems, when users set the property “Collect Hypervisor metrics” for VIOS host monitoring:

- [PowerVMVIO](#) topology object
- [PowerVMVIOCPUs](#) topology object
- [PowerVMVIOMemory](#) topology object
- [PowerVMVIONetwork](#) topology object
- [PowerVMVIOStorage](#) topology object

Windows metrics

The following topology objects are collected by the WindowsAgent:

- [Host](#) topology object
- [HostCPUs](#) topology object
- [HostNetwork](#) topology object
- [HostProcess](#) topology object
- [HostProcessInstance](#) topology object
- [HostService](#) topology object
- [HostTopProcessEntry](#) topology object
- [LogicalDisk](#) topology object
- [MSCluster](#) topology object
- [MSClusterNode](#) topology object
- [MSClusterResource](#) topology object
- [MSClusterResourceGroup](#) topology object
- [MSHostedResourceGroup](#) topology object
- [Memory](#) topology object
- [NetworkInterface](#) topology object
- [OperatingSystem](#) topology object
- [PhysicalDisk](#) topology object
- [Processor](#) topology object
- [WinHostProcessInstance](#) topology object

AIXHostDetails topology object

The AIXHostDetails type contains additional AIX®-specific host information, relating to partitions. A single instance is created as a child of a Host’s details node.

Table 31. AIXHostDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
Capped	Indicates whether the partition is capped or uncapped.	AIX®
entitlement	Indicates the entitlement of processing capacity allocated in processor units to the partition.	AIX
isLPAR	Indicates whether the monitored host represents a logical partition.	AIX
lparID	Indicates logical partition ID if the host is a LPAR.	AIX
lparName	Indicates logical partition name if the host is a LPAR.	AIX
SerialNum	Indicates the serial number where the AIX host is running.	AIX
Shared	Indicates whether the partition is shared or dedicated.	AIX
SMT	Indicates whether the partition has simultaneous multi-threading enabled.	AIX
typeModel	Indicates the IBM® machine type where the AIX host is running.	AIX

AIXMemoryDetails topology object

The AIXMemoryDetails type contains additional details specific to AIX® hosts, relating to memory activity. A single instance is associated with the monitored host's Memory (as a child of details, called "MemoryDetails"). These metrics are only available if using the native collector.

Table 32. AIXMemoryDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
cycles	Number of page replacement cycles over the polling interval.	AIX®
numperm	Number of frames used for files.	AIX
pgbad	Number of bad pages occurring over the polling interval.	AIX
pgexct	Number of page faults over the polling interval.	AIX
pgsp_rsvd	Size of reserved paging space (in 4 KB pages).	AIX
pgsteals	Number of page steals over the polling interval.	AIX
real_process	Number of pages used by process segments. This is defined as: $real_total - real_free - numperm - real_system$ Because $real_system$ is an approximation, this number is also an approximation.	AIX
real_system	Number of pages used by system segments. This is the sum of all the used pages in segments marked for system usage. Because segment classifications are not always guaranteed to be accurate, this number is only an approximation.	AIX
real_user	Number of pages used by non-system segments. This is the sum of all pages used in segments not marked for system usage. Because segment classifications are not always guaranteed to be accurate, this number is only an approximation.	AIX
scans	Number of page scans by clock over the polling interval.	AIX

AIXPhysicalCPUUsage topology object

The AIXPhysicalCPUUsage type contains additional details regarding Physical CPU usage on AIX® systems. Summary and per-processor metrics are reported scaled to the LPAR entitlement. These are the raw values.

Table 33. AIXPhysicalCPUUsage topology object

Metric Name	Metric Description	Metric Collected for These Platforms
idle	Percentage of the collection period the physical CPUs on the system spent waiting for I/O.	AIX®
physical_processors_used	How many physical CPUs were consumed during the collection period. A value of (for example) 0.2 on a 2-processor system means that (0.2/2.0) processors were used up, or about 10% physical CPU utilization.	AIX
system	Percentage of the collection period the physical CPUs on the system spent executing in system mode.	AIX
user	Percentage of the collection period the physical CPUs on the system spent executing in user mode.	AIX
system	Percentage of the collection period the physical CPUs on the system spent idle.	AIX

AIXUsedMemoryDetails topology object

The AIXUsedMemoryDetails type contains additional details regarding used memory on AIX®. This information may require use of sudo for the agent to collect. For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).

Table 34. AIXUsedMemoryDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
client	Amount of RAM in use as client pages.	AIX®
persistent	Amount of RAM in use as persistent pages.	AIX
working	Amount of RAM in use as working pages.	AIX

CPUCounts topology object

The CPUCounts type is a topology detail stored on the *HostCPUs*, which contains the physical, logical, and virtual CPU counts for the system should they be available.

Table 35. CPUCounts topology object

Metric Name	Metric Description	Metric Collected for These Platforms
logicalCPUs	If a physical chip has multiple cores on it, each core is considered a logical CPU. This is the total number on the system.	AIX® Solaris®

Table 35. CPUCounts topology object

Metric Name	Metric Description	Metric Collected for These Platforms
physicalCPUs	The number of physical CPUs on the system. These are actual processor chips you can pick up.	AIX Solaris
virtualCPUs	If the CPU support HyperThreading (or the platform equivalent), this is the total number of such CPUs. Generally, this is the number of CPUs that appear to be available to user processes.	AIX Solaris

Host topology object

Host objects are identified in Foglight by their `Name` property. This property typically contains the host's fully qualified domain name. The fully qualified name is also decomposed into the `localName` and `domainName` properties. The IP addresses assigned to a host are captured in the `primaryIpAddress` and `ipAddresses` properties. The addresses are not unique to a particular host.

Table 36. Host topology object

Metric Name	Metric Description	Metric Collected for These Platforms
availablePagingSpace	The percentage of paging space available on the host.	AIX® HP-UX Linux® Solaris® Windows®
bootTime	The time at which the system was booted, in seconds since UTC epoch (January 1, 1970).	AIX ^a HP-UX ^b Linux Solaris Windows
contextSwitches	The rate of context switching on the host (count/second).	AIX HP-UX Linux Solaris Windows
fiveMinuteLoadAverage	The average system load over a period of five minutes.	AIX HP-UX Linux Solaris Windows
interrupts	The interrupt rate (count/second).	AIX HP-UX Linux Solaris Windows

Table 36. Host topology object

Metric Name	Metric Description	Metric Collected for These Platforms
ipAddresses	A list of all IP addresses assigned to the host.	AIX HP-UX Linux Solaris Windows
monitored	Shows whether the agent is able to connect to the host for monitoring. Monitored state is MONITORED, UNMONITORED, or UNAVAILABLE.	AIX HP-UX Linux Solaris Windows
name	The name of this host.	AIX HP-UX Linux Solaris Windows
numProcesses	The number of processes running on the host.	AIX HP-UX Linux Solaris Windows
primaryIpAddress	The primary or canonical IP address for this host.	AIX HP-UX Linux Solaris Windows
rebootCount	The host system reboot count.	AIX HP-UX Linux Solaris Windows
runQueueLength	The run queue length. NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows

Table 36. Host topology object

Metric Name	Metric Description	Metric Collected for These Platforms
systemId	The unique identifier for the host.	AIX
		HP-UX
		Linux
		Solaris
uptime	The duration since the system was last started.	Windows
		AIX
		HP-UX
		Linux
		Solaris
		Windows

a. This metric is available only on AIX systems using native collectors. It is not collected for command line.

b. This metric is available only on HP-UX systems using native collectors. It is not collected for command line.

HostCPUs topology object

There is a single HostCPUs instance attached to a Host that provides host level summary metrics for the processors on a host. The HostCPUs instance is identified by the reference to the associated Host object and has its name set with the constant string `CPUs`.

Table 37. HostCPUs topology object

Metric Name	Metric Description	Metric Collected for These Platforms
percentIOWaitTime	The average percentage of time that processors on the host were idle while there was at least one I/O in progress.	AIX®
		HP-UX
		Linux®
		Solaris®
percentIdleTime	The average percentage of time that processors on the host are idle.	AIX
		HP-UX
		Linux
		Solaris
percentSystemTime	The average percentage of time that processors on the host spend executing system code in user mode.	Windows®
		AIX
		HP-UX
		Linux
percentUserTime	The average percentage of time that processors on the host spend executing user code.	Solaris
		Windows
		AIX
		HP-UX
		Linux
		Solaris
		Windows

Table 37. HostCPUs topology object

Metric Name	Metric Description	Metric Collected for These Platforms
totalHz	The total speed (measured in Hertz) of all physical processors installed on a host, or the total number of cycles allocated to a virtual machine.	AIX HP-UX Linux Solaris Windows
usedHz	The total amount of clock cycles (measured in Hertz) being used by processors on the host.	AIX HP-UX Linux Solaris Windows
utilization	The average percentage of time that processors on the host are utilized (that is, not idle). NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX Linux ^a Windows

a. Collected for Linux platforms only by the UnixAgentPlus.

HostCPUs — Physical CPU Utilization

The CPU usage of AIX® LPAR (Logical Partition) is not easily or consistently measurable via a percentage. Depending on the LPAR configuration, the numbers can vary widely (often exceeding 100%, as in prior releases).

The current formula for calculating the Physical CPU Utilization provides a more meaningful value for this metric, which is reflective of physical CPU allocation (regardless of virtualization setup) and does not exceed 100%.

Customers who prefer using the old formula from the legacy agent (delivered with the Cartridge for Operating Systems) can revert to the old calculation by using the following formulas:

- For capped lpar:

$$\text{CPU Utilization (oldMetric)} = \text{Physical CPU Utilization (newMetric)} / \text{entitlement}$$

- For uncapped lpar:

- if virtualCPUs >= physicalCPUs

$$\text{oldMetric} = (\text{newMetric} * \text{physicalCPUs}) / \text{entitlement}$$

- if virtualCPUs < physicalCPUs

$$\text{oldMetric} = (\text{newMetric} * \text{virtualCPUs}) / \text{entitlement}$$

HostNetwork topology object

There is a single HostNetwork instance attached to a host that provides host level summary metrics for the network interfaces on the host. The HostNetwork instance is identified by the reference to the associated Host object and has its name set with the constant string `Network`.

Table 38. HostNetwork topology object

Metric Name	Metric Description	Metric Collected for These Platforms
packetsReceived	The total number of packets received on all interfaces (count/second). NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX® HP-UX Linux® Solaris® Windows®
packetsSent	The total number of packets sent on all interfaces (count/second). NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows
receiveRate	The combined receive rate on all network interfaces (bit/second). NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows
sendRate	The combined send rate on all network interfaces (bit/second). NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows
tcpConnectionsActive	The number of open and active TCP connections.	AIX HP-UX Linux Solaris Windows
tcpConnectionsEstablished	The number of TCP connections last observed to be in the ESTABLISHED or CLOSE-WAIT state.	AIX HP-UX Linux Solaris Windows
tcpConnectionsFailures	The number of TCP connections that have failed since the service was started. TCP considers a connection as failed when it goes directly from sending (SYN-SENT) or receiving (SYN-RCVD) to closed (CLOSED) state, or from receiving (SYN-RCVD) to listening (LISTEN).	AIX HP-UX Linux Solaris Windows
tcpConnectionsPassive	The number of open and passive TCP connections.	AIX HP-UX Linux Solaris Windows

Table 38. HostNetwork topology object

Metric Name	Metric Description	Metric Collected for These Platforms
tcpConnectionsReset	The number of connections that have been reset since the service was started (regardless of when the System Monitor was started). TCP considers a connection as reset when it goes directly from ESTABLISHED or CLOSE-WAIT to CLOSED state. In Linux, this number represents the count of receiving resets.	AIX HP-UX Linux Solaris Windows
transferRate	The total rate on all transfers on all network interfaces (bit/second).	AIX HP-UX Linux Solaris Windows
utilization	The average utilization of network interfaces on the host. NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux ^a Solaris Windows

a. This metric is unavailable when the network card's maximum bandwidth is not reported by the kernel.

HostProcess topology object

The HostProcess type captures aggregate metrics for all processes of a given type on the host. Processes are aggregated according to executables.

Table 39. HostProcess topology object

Metric Name	Metric Description	Metric Collected for These Platforms
dataOther	The rate (byte/second) at which instances of the process transfer bytes in I/O operations that are neither read nor write operations (for example, control operations).	AIX® HP-UX Solaris® Windows®
dataRead	The rate (byte/second) at which instances of the process read data in I/O operations to file, network, etc.	AIX HP-UX Linux® Solaris Windows
dataWritten	The rate (byte/second) at which instances of the process write data in I/O operations to file, network, etc.	AIX HP-UX Linux Solaris Windows

Table 39. HostProcess topology object

Metric Name	Metric Description	Metric Collected for These Platforms
name	The name of the HostProcess group.	AIX HP-UX Linux Solaris Windows
operationsOther	The rate (count/second) of I/O operations (other than read/write operations) performed by instances of this process.	AIX HP-UX Solaris Windows
operationsRead	The rate (count/second) of read I/O operations performed by instances of the process.	AIX HP-UX Linux Solaris Windows
operationsWritten	The rate (count/second) of write I/O operations performed by instances of the process.	AIX HP-UX Linux Solaris Windows
pageFaults	The total number of page faults for this process (count/second).	AIX HP-UX Linux Solaris Windows
percentAvailability	The percentage of processes found divided by the number of processes expected. (percent)	AIX HP-UX Linux Solaris Windows
percentMemory	The percentage of real memory that has been consumed by instances of the process.	AIX HP-UX Linux Solaris Windows
percentProcessorTime	The percentage of the total processor time that has been consumed by instances of the process.	AIX HP-UX Linux Solaris Windows
swapSize	The total size (kilobyte) of swap or paging space used by instances of the process.	AIX HP-UX Solaris Windows

Table 39. HostProcess topology object

Metric Name	Metric Description	Metric Collected for These Platforms
totalProcessorTime	The total processor time (seconds) that has been consumed by instances of the process.	AIX HP-UX Linux Solaris Windows
virtualSize	The total virtual size (kilobyte) of instances of the process.	AIX HP-UX Linux Solaris Windows
workingSetSize	The total amount of real memory (kilobyte) consumed by instances of the process.	AIX HP-UX Linux Solaris Windows

HostProcessInstance topology object

The *HostProcessInstance* complex observation captures detailed per-process statistics, but that observation may not be produced at all times.

Table 40. HostProcessInstance topology object

Metric Name	Metric Description	Metric Collected for These Platforms
dataOther	The rate at which instances of the process transfer bytes in I/O operations that are neither read nor write operations (for example, control operations).	AIX ^a Windows [®]
dataRead	The rate at which instances of the process read data in I/O operations to file, network, etc.	Linux [®] Solaris ^{®b} Windows
dataWritten	The rate at which instances of the process write data in I/O operations to file, network, etc.	Linux Solaris ^c Windows
operationsOther	The rate of I/O operations (other than read/write operations) performed by instances of this process.	Windows
operationsRead	The rate of read I/O operations performed by instances of the process. (count/second)	HP-UX ^d Linux Solaris ^e Windows

Table 40. HostProcessInstance topology object

Metric Name	Metric Description	Metric Collected for These Platforms
operationsWritten	The rate of write I/O operations performed by instances of the process. (count/second)	HP-UX ^f Linux Solaris ^g Windows
pageFaults	The total number of page faults for this process instance (count/second).	AIX HP-UX Linux Solaris Windows
parentPid	The PID of the process that launched this instance.	AIX HP-UX Linux Solaris Windows
percentMemory	The percentage of real memory that has been consumed by instances of the process.	AIX HP-UX Linux Solaris Windows
percentProcessorTime	The percentage of the total processor time that has been consumed by instances of the process.	AIX HP-UX Linux Solaris Windows
pid	The PID of the identified process. Note that PIDs can be recycled once their process exits.	AIX HP-UX Linux Solaris Windows
swapSize	The total size of swap or paging space used by instances of the process.	Windows
totalProcessorTime	The total processor time that has been consumed by instances of the process. (seconds)	AIX HP-UX Linux Solaris Windows

Table 40. HostProcessInstance topology object

Metric Name	Metric Description	Metric Collected for These Platforms
virtualSize	The total virtual size of instances of the process. (kilobytes)	AIX HP-UX Linux Solaris Windows
workingSetSize	The total amount of real memory consumed by instances of the process. (kilobytes)	AIX HP-UX Linux Solaris Windows

a. This metric is available only on AIX systems using native collectors. The value represents the sum of bytes read and written by the process.

b. Some data necessary to calculate these metrics require use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).) These values and the information derived from them, such as the Top 5 Processes by Disk I/O, cannot be considered fully accurate, because the Solaris kernel does not always update these metrics when the disk I/O is done on behalf of a process. As a result, these metrics should be treated as indications of the general amount of disk I/O that is done, rather than hard numbers.

c. Some data necessary to calculate these metrics require use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).) These values and the information derived from them, such as the Top 5 Processes by Disk I/O, cannot be considered fully accurate, because the Solaris kernel does not always update these metrics when the disk I/O is done on behalf of a process. As a result, these metrics should be treated as indications of the general amount of disk I/O that is done, rather than hard numbers.

d. This metric is not available on HP-UX systems using non-native collectors.

e. Some data necessary to calculate these metrics require use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).) These values and the information derived from them, such as the Top 5 Processes by Disk I/O, cannot be considered fully accurate, because the Solaris kernel does not always update these metrics when the disk I/O is done on behalf of a process. As a result, these metrics should be treated as indications of the general amount of disk I/O that is done, rather than hard numbers.

f. This metric is not available on HP-UX systems using non-native collectors.

g. Some data necessary to calculate these metrics require use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo](#).) These values and the information derived from them, such as the Top 5 Processes by Disk I/O, cannot be considered fully accurate, because the Solaris kernel does not always update these metrics when the disk I/O is done on behalf of a process. As a result, these metrics should be treated as indications of the general amount of disk I/O that is done, rather than hard numbers.

HostService topology object

The *HostService* complex topology object captures information about the Windows® service.

Table 41. HostService topology object

Metric Name	Metric Description	Metric Collected for These Platforms
displayName	The display name of the service.	Windows®
name	The name of the service.	Windows
pid	The process ID of the service.	Windows
process	The process of the service.	Windows
startupType	The Startup Type of the service.	Windows
state	The state of the service.	Windows

HostTopProcessEntry topology object

A process detected in one of the Top-N lists, such as the top ten CPU consumers or disk users.

Table 42. HostTopProcessEntry topology object

Metric Name	Metric Description ^a	Metric Collected for These Platforms
command	The command line used to launch this process, if available. Otherwise just the name of the process.	AIX® HP-UX Linux® Solaris® Windows®
userName	The name of the user that launched this process, if available.	AIX HP-UX Linux Solaris® Windows

a. For Linux platforms, these metrics are collected only by the UnixAgentPlus.

LPARNetworkInterfaceDetails topology object

Extra properties related to the network interface if the Host is LPAR.

Table 43. LPARNetworkInterfaceDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
parentAdapterLocationCode	Virtual Client adapter location code mapping to interface.	AIX®
parentAdapterName	Virtual Client adapter name mapping to interface.	AIX®

LPARPhysicalDiskDetails topology object

Extra properties related to the physical disk if the Host is LPAR.

Table 44. LPARPhysicalDiskDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
locationCode	Physical disk location code.	AIX®
parentAdapterLocationCode	Physical adapter location code mapping to physical disk.	AIX®
parentAdapterName	Physical adapter Name mapping to physical disk.	AIX®
parentAdapterType	Physical adapter type mapping to physical disk.	AIX®
pvid	Physical volume pvid.	AIX®

LogicalDisk topology object

A Windows® partition or UNIX® filesystem. Multiple LogicalDisk instances are automatically aggregated by the server into a HostStorage object, of which there is one per Host. A LogicalDisk instance is identified by its name (for example, *hda1*) and a reference to the Host object in which it resides.

- i** | **IMPORTANT:** When monitoring the Oracle Solaris® ZFS® file systems, it is expected that:
- The `capacityUsed`, `capacityAvailable`, `spaceUsed`, and `spaceAvailable` metrics do not match with the data shown by running the following command: `df -h`
 - The sum of `capacityUsed` and `capacityAvailable` is not equal to 100%.
- For more information, see the Knowledge Base article <https://support.quest.com/kb/206363>.

Table 45. LogicalDisk topology object

Metric Name	Metric Description	Metric Collected for These Platforms
averageReadTime	The average amount of time (millisecond) for a disk read to complete.	Linux® ^a Solaris® ^b
averageTransferTime	The average amount of time (millisecond) for a disk transfer, either read or write.	Solaris® Windows®
averageWriteTime	The average amount of time (millisecond) for a disk write to complete.	Linux ^c Solaris ^d
bytesRead	The volume of data read from the filesystem (kilobyte/second).	AIX® HP-UX ^e Linux Solaris Windows
bytesWritten	The volume of data written to the filesystem (kilobyte/second).	AIX HP-UX ^f Linux Solaris Windows

Table 45. LogicalDisk topology object

Metric Name	Metric Description	Metric Collected for These Platforms
capacityAvailable	The amount of space available in the filesystem as a percentage of the total space. (percent)	AIX HP-UX Linux Solaris Windows
capacityUsed	The amount of space used in the filesystem as a percentage of the total space. (percent)	AIX HP-UX Linux Solaris Windows
Detail/readable	Indicates whether or not data can be read from disk.	AIX HP-UX Linux Solaris
Detail/writable	Indicates whether or not data can be written to disk.	AIX HP-UX Linux Solaris
Detail/writeOnce	Indicates whether or not the disk can be written to only once.	Linux ^g
filesystemType	The type of file system used for this logical disk.	AIX HP-UX Linux ^h Windows
isRemote	Whether or not the disk is mounted from a remote source (a network file system).	AIX HP-UX Linux ⁱ Windows
name	The name of the partition, such as "C:", or "hda1", or "sd1,a".	AIX HP-UX Linux Solaris Windows
percentInodesRemaining	The percentage of inodes available on the filesystem. (percent)	AIX HP-UX Linux Solaris Windows
queueLength	The length of the queue containing I/O requests that have begun processing but have not yet finished.	Solaris Windows

Table 45. LogicalDisk topology object

Metric Name	Metric Description	Metric Collected for These Platforms
reads	The rate (count/second) of read operations.	HP-UX ^j Linux Solaris Windows
spaceAvailable	The amount of space available in the filesystem (megabyte).	AIX HP-UX Linux Solaris Windows
spaceUsed	The amount of space used in the filesystem (megabyte).	AIX HP-UX Linux Solaris Windows
totalSpace	The size of the filesystem (megabyte).	AIX HP-UX Linux Solaris Windows
utilization	The utilization in terms of available bandwidth.	Solaris
writes	The (count/second) rate of write operations.	HP-UX ^k Linux Solaris Windows

a. Collected for Linux platforms only by the UnixAgentPlus.

b. Collected for Solaris platforms only by the UnixAgentPlus.

c. Collected for Linux platforms only by the UnixAgentPlus.

d. Collected for Solaris platforms only by the UnixAgentPlus.

e. This metric is not available on HP-UX systems using non-native collectors.

f. This metric is not available on HP-UX systems using non-native collectors.

g. Collected for Linux platforms only by the UnixAgentPlus. This value is not collected for logical disks because it does not exist.

h. Collected for Linux platforms only by the UnixAgentPlus.

i. Collected for Linux platforms only by the UnixAgentPlus.

j. This metric is not available on HP-UX systems using non-native collectors.

k. This metric is not available on HP-UX systems using non-native collectors.

MSCluster topology object

An object that represents a node in a Microsoft® Windows® cluster. Typically these will share names with actively monitored hosts. The WindowsAgent's cluster data collection must be enabled for these to be created. The MScClusterNode is placed under a Host's details, and it references the MScCluster object for the cluster of which it is a part.

Table 46. MSCluster topology object

Metric Name	Metric Description	Metric Collected for These Platforms
desc	A brief description of the cluster, offered by the user.	Windows®
maxNumberOfNodes	The maximum number of nodes supported by the cluster.	Windows
name	A string offering the name of the cluster.	Windows
primaryOwnerContact	The primary owner's contact information.	Windows
primaryOwnerName	The primary owner of the cluster.	Windows
status	The status of the cluster.	Windows

MSClusterNode topology object

An object that represents a node in a Microsoft® Windows® cluster. Typically these will share names with actively monitored hosts. The WindowsAgent's cluster data collection must be enabled for these to be created. The MSClusterNode is placed under a Host's details, and it references the MSCluster object for the cluster of which it is a part.

Table 47. MSClusterNode topology object

Metric Name	Metric Description	Metric Collected for These Platforms
desc	A brief description of the node, offered by the user.	Windows®
name	A string offering the name of the node.	Windows
state	The state of the cluster.	Windows
status	The status of the cluster.	Windows

MSClusterResource topology object

An object that represents a resource in a Microsoft® Windows® cluster. The WindowsAgent's cluster data collection must be enabled for these to be created. Resources are held by resource groups, which are accessible through the MSCluster nodes on which they are active.

Table 48. MSClusterResource topology object

Metric Name	Metric Description	Metric Collected for These Platforms
desc	A brief description of the resource, offered by the user.	Windows®
name	A string offering the name of the resource.	Windows
state	The state of the resource.	Windows
type	The type of resource.	Windows

MSClusterResourceGroup topology object

An object that represents a resource group in a Microsoft® Windows® cluster. The WindowsAgent's cluster data collection must be enabled for these to be created. Resource groups hold a collection of resources; resource groups are reachable through the node on which they are actively hosted.

Table 49. MSClusterResourceGroup topology object

Metric Name	Metric Description	Metric Collected for These Platforms
desc	A brief description of the resource, offered by the user.	Windows®
name	A string offering the name of the resource group.	Windows
state	The state of the resource.	Windows
status	The status of the resource.	Windows

MShostedResourceGroup topology object

An object that represents a resource group, but treated as a Host, since it extends HostExtension. The WindowsAgent's *'treatResourceGroupAsHost'* boolean must be enabled in order for these to be created. Along with the MShostedResourceGroup, a new Host object of equivalent name is created. The two objects together can be used to represent the host on which a resource group is running. The actual host running the resource group can be determined from the object's *'activeHost'* property.

Table 50. MShostedResourceGroup topology object

Metric Name	Metric Description	Metric Collected for These Platforms
activeHost	The host currently hosting the resource group. If performance metrics (such as CPU metrics) should be associated with the resource group, the values on the activeHost may be helpful.	Windows®
clusterName	The cluster of which the resource group is a part.	Windows
host	A fabricated host, with name matching that of the MShostedResourceGroup.	Windows
name	The name of the hosted resource group, shared by the host property's name. It is a combination of the clusterName and resource group name.	Windows
resourceGroupName	The name of the resource group.	Windows

Memory topology object

There is a single Memory object attached to a host that is identified by that Host instance. The name property of a Memory object is set with the constant string `Memory`.

Table 51. Memory topology object

Metric Name	Metric Description	Metric Collected for These Platforms
capacity	The amount of memory (megabyte) installed on or allocated to the host.	AIX® HP-UX Linux® Solaris® Windows®
consumed	The amount of memory (megabyte) that is being used on the host.	AIX HP-UX Linux Solaris Windows
fileSystemCache	The amount of physical memory (Mb) used as a file system cache.	Linux ^a Solaris ^b
nonPagedKernelMemory	The amount of memory (Mb) allocated by the kernel that cannot be paged out.	none
pageInRate	The rate (count/second) at which pages are being read from disk. NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows
pageOutRate	The rate (count/second) at which pages are being written to disk. NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX HP-UX Linux Solaris Windows
pageScanRate	The number of pages scanned by the clock per second.(count/second)	AIX HP-UX Linux Solaris Windows
pagedKernelMemory	The amount of physical memory (Mb) used by the kernel for data that can be paged out.	none
utilization	The percentage of physical memory that is currently utilized. NOTE: IntelliProfile thresholds are defined for this metric. For details, review the thresholds list in the Manage Thresholds dashboard.	AIX® HP-UX Linux Solaris Windows
virtualMemorySize	The total amount of virtual memory (Mb) allocated on the host. This is physical memory size plus the size of the paging/swap space.	AIX HP-UX Solaris
hugePagesTotal	The pool size of huge pages	Linux
hugePagesFree	The available pool size of huge pages	Linux
hugePageSize	The default size of huge page	Linux

a.Collected for Linux platforms only by the UnixAgentPlus.

b.Collected for Solaris platforms only by the UnixAgentPlus.

NetworkInterface topology object

A single network interface. Multiple network interfaces can be assigned to a single host object.

Table 52. NetworkInterface topology object

Metric Name	Metric Description	Metric Collected for These Platforms
bandwidth	The bandwidth (bit/second) of the network interface (for example, 100Mb/s).	AIX ^a HP-UX ^b Linux ^c Solaris [®] Windows [®]
collisions	The number of collisions (requiring packet resend) that occurred on the network interface. (count)	AIX HP-UX ^d Linux Solaris
discoveryStatus	This value is set to Active if the interface was recently discovered and/or verified to exist. The value is set to Inactive if the interface is considered obsolete.	AIX ^e HP-UX
inboundErrors	The number (count) of inbound packets that were dropped because of errors.	AIX HP-UX Linux Solaris Windows
inboundPacketsDropped	The number of inbound packets that were dropped even though no error that would have prevented the packet from being received was detected. (count)	HP-UX ^f Linux Windows
interfaceIndex	A unique number (> 0) identifying the interface within the device.	none
interfaceType	The type of interface as assigned by the Internet Assigned Numbers Authority (IANA).	AIX HP-UX Linux
ipAddress	IP address of the interface.	AIX HP-UX Linux Solaris
isLayer2	This property is set to true if the interface is a layer 2 interface.	none
layer2Segment	The layer 2 segment to which the interface belongs.	none
macAddress	The MAC or physical address assigned to the interface.	AIX HP-UX Linux Solaris

Table 52. NetworkInterface topology object

Metric Name	Metric Description	Metric Collected for These Platforms
netmask	The subnet mask associated with the IP address of the interface.	AIX HP-UX Linux Solaris
networkSegment	The network segment to which the interface belongs.	AIX HP-UX Solaris
outboundErrors	The number (count) of outbound packets that could not be transmitted because of errors.	AIX HP-UX Linux Solaris Windows
outboundPacketsDropped	The number (count) of outbound packets that were dropped even though no error that would have prevented the transmission was detected.	HP-UX ^g Linux Windows
outputQueueLength	The length of the output packet queue (count).	HP-UX ^h Windows
packetsReceived	The number of packets received on this interface (count/second).	AIX HP-UX Linux Solaris Windows
packetsSent	The number of packets sent on this interface (count/second).	AIX HP-UX Linux Solaris Windows
portNumber	The port number associated with the interface.	none
receiveRate	The receive rate on this network interface (bit/second).	AIX ⁱ HP-UX ^j Linux Solaris Windows
sendRate	The send rate on this network interface (bit/second).	AIX ^k HP-UX ^l Linux Solaris Windows

Table 52. NetworkInterface topology object

Metric Name	Metric Description	Metric Collected for These Platforms
utilization	The utilization of the network interface's bandwidth.	AIX ^m HP-UX ⁿ Linux ^o Solaris Windows
collisionRate	The percentage of packet collisions on this network interface.	AIX HP-UX ^p Linux Solaris
errorPercentage	The percentage of error packets on this network interface.	AIX HP-UX Linux Solaris Windows

- a. This metric is not available on AIX systems using non-native collectors.
- b. This metric is not available on HP-UX systems using non-native collectors.
- c. This metric is reported by ethtool (or, if unavailable, by mii-tool). This metric is not available for loopback interfaces. It requires use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo.](#))
- d. This metric is not available on HP-UX systems using native collectors.
- e. This metric is not available on AIX systems using native collectors.
- f. This metric is not available on HP-UX systems using non-native collectors.
- g. This metric is not available on HP-UX systems using non-native collectors.
- h. This metric is not available on HP-UX systems using non-native collectors.
- i. This metric is not available on AIX systems using non-native collectors.
- j. This metric is not available on HP-UX systems using non-native collectors.
- k. This metric is not available on AIX systems using non-native collectors.
- l. This metric is not available on HP-UX systems using non-native collectors.
- m. This metric is not available on AIX systems using non-native collectors.
- n. This metric is not available on HP-UX systems using non-native collectors.
- o. This metric requires bandwidth from ethtool or mii-tool. It requires use of sudo for the agent to collect. (For more information about sudo commands that require root access, see [Configuring secure launcher permissions using sudo.](#))
- p. This metric is not available on HP-UX systems using native collectors.

NetworkInterfaceDetails topology object

This topology object houses additional metrics related to the network interface.

Table 53. NetworkInterfaceDetails topology object

Metric Name	Metric Description	Metric Collected for These Platforms
collisions	The number of collisions (requiring packet resend) that occurred on the network interface. (count)	AIX® HP-UX ^a Linux® Solaris®

a.This metric is not available on HP-UX systems using native collectors.

NixHostProcessInstance topology object

On UNIX®, instances of NixHostProcessInstance are used instead of HostProcessInstance. These topology objects extend HostProcessInstance with a few more details that can be captured on UNIX platforms.

Table 54. NixHostProcessInstance topology object

Metric Name	Metric Description ^a	Metric Collected for These Platforms
commandLine	The full command line used to launch the executable.	Linux®
creationDate	The time the process instance was created. (date)	Linux
executablePath	The full path to the executable.	Linux
sessionId	The session id of the process.	Linux

a.For Linux platforms, these metrics are collected only by the UnixAgentPlus.

OperatingSystem topology object

There is a single instance of the OperatingSystem type attached to a host. The instance is identified by the reference to that Host and provides additional details about the operating system.

Table 55. OperatingSystem topology object

Metric Name	Metric Description ^a	Metric Collected for These Platforms
architecture	The architecture (for example, "ia32").	AIX® HP-UX Linux® Solaris® Windows®
buildNumber	The OS build number. Typically found on Windows machines.	Solaris Windows
name	The operating system name (for example, "Windows XP").	AIX HP-UX Linux Solaris Windows
release	The Linux distribution name.	Linux

Table 55. OperatingSystem topology object

Metric Name	Metric Description ^a	Metric Collected for These Platforms
servicePackMajorVersion	The major version of the service pack. Typically found on Windows machines.	Windows
servicePackMinorVersion	The minor version of the service pack. Typically found on Windows machines.	Windows
type	The general type of operating system. This is typically one of "Windows", "Linux", "AIX", "HPUX", or "Sun®".	AIX HP-UX Linux Solaris Windows
version	The OS version.	AIX HP-UX Linux Solaris Windows

a.For Linux platforms, metrics are collected only by the UnixAgentPlus.

PhysicalDisk topology object

The PhysicalDisk represents a disk that is installed in the machine or configured for a virtual machine. It is identified by its name (for example, *hda*) and a reference to its parent Host.

Table 56. PhysicalDisk topology object

Metric Name	Metric Description	Metric Collected for These Platforms
averageReadTime	The average amount of time (millisecond) for a disk read to complete.	Linux ^a Solaris ^b
averageTransferTime	The average amount of time (millisecond) for a disk transfer, either read or write.	AIX ^c HP-UX ^d Linux Solaris [®] Windows [®]
averageWriteTime	The average amount of time (millisecond) for a disk write to complete.	Linux ^e Solaris ^f
bytesRead	The volume of data read from the filesystem (kilobyte/second).	AIX Linux Solaris Windows
bytesWritten	The volume of data written to the filesystem (kilobyte/second).	AIX Linux Solaris Windows

Table 56. PhysicalDisk topology object

Metric Name	Metric Description	Metric Collected for These Platforms
capacityAvailable	The amount of space available in the filesystem as a percentage of the total space (percent).	AIX HP-UX Windows ^g
capacityUsed	The amount of space used in the filesystem as a percentage of the total space (percent).	AIX HP-UX Windows ^h
Detail/readable	Indicates whether or not the disk can be read from.	none
Detail/writable	Indicates whether or not the disk can be written to.	none
Detail/writeOnce	Indicates whether or not the disk can be written to only once.	none
name	The name of the disk, such as "C:", or "hda", or "sd1".	AIX HP-UX Linux ⁱ Solaris Windows
queueLength	The length of the queue containing I/O requests that have begun processing but have not yet finished.	AIX ^j HP-UX ^k Linux Solaris Windows
reads	The rate (count/second) of read operations.	AIX ^l Linux Solaris Windows
spaceAvailable	The amount of space available in the filesystem (megabyte).	AIX ^m
spaceUsed	The amount of space used in the filesystem (megabyte).	AIX ⁿ
totalSpace	The size of the filesystem (megabyte).	AIX HP-UX
utilization	The utilization in terms of available bandwidth.	AIX ^o Linux Solaris Windows
writes	The rate (count/second) of write operations.	AIX ^p Linux Solaris Windows

a. Collected for Linux platforms only by the UnixAgentPlus.

b. Collected for Solaris platforms only by the UnixAgentPlus.

c. This metric is not available when using command line mode.

d. This metric is not available when using command line mode.

- e. Collected for Linux platforms only by the UnixAgentPlus.
- f. Collected for Solaris platforms only by the UnixAgentPlus.
- g. This metric is only available for logical filesystems.
- h. This metric is only available for logical filesystems.
- i. Collected for Linux platforms only by the UnixAgentPlus.
- j. This metric is not available when using command line mode.
- k. This metric is not available when using command line mode.
- l. This metric is not available when using command line mode.
- m. This metric is not collected for native.
- n. This metric is not collected for native.
- o. This metric is not available when using command line mode.
- p. This metric is not available when using command line mode.

PowerVMVIOS topology object

As part of PowerVM, the Virtual I/O Server is a software appliance with which you can associate physical resources and that allows you to share these resources among multiple client logical partitions. The Virtual I/O Server can use both virtualized storage and network adapters, making use of the virtual SCSI and virtual Ethernet facilities.

Table 57. PowerVMVIOS topology object

Metric Name	Metric Description	Metric Collected for These Platforms
cores	The count for cores of the VIOS.	none
	The operating system version for this VIOS.	AIX® HP-UX Linux® Solaris® Windows®
osVersion	The current processor mode of this VIOS which is “ded” or “shared”.	AIX HP-UX Linux Solaris Windows
	The processing sharing mode of this VIOS which is “cap” or “uncap”.	AIX HP-UX Linux Solaris Windows
procMode	The state of this VIOS whether is “Running” or “Not Activated”.	AIX HP-UX Linux Solaris Windows
sharingMode		
state		

Table 57. PowerVMVIOS topology object

Metric Name	Metric Description	Metric Collected for These Platforms
viosID	The unique integer identifier for the VIOS.	AIX HP-UX Linux Solaris Windows
viosName	The user-defined name of the VIOS.	AIX HP-UX Linux Solaris Windows

PowerVMVIOSCPUs topology object

There is a single *PowerVMVIOSCPUs* instance attached to a *PowerVMVIOS* that provides *PowerVMVIOS* level summary metrics for the processors interfaces on the *PowerVMVIOS*.

Table 58. PowerVMVIOSCPUs topology object

Metric Name	Metric Description	Metric Collected for These Platforms
desiredProcUnits	The desired processor units of this VIOS. (count)	AIX® HP-UX Linux® Solaris® Windows®
desiredProcs	The desired processors of this VIOS. (count)	AIX HP-UX Linux Solaris Windows
percentageUnusedQuota	The percentage of the unused capacity.	AIX HP-UX Linux Solaris Windows
uncappedWeight	The uncapped weight must be a whole number from 0 to 255. The default uncapped weight for uncapped micro-partitions is 128. (count)	AIX HP-UX Linux Solaris Windows

Table 58. PowerVMVIOSCPUs topology object

Metric Name	Metric Description	Metric Collected for These Platforms
utilization	The processor utilization of this VIOS.	AIX
		HP-UX
		Linux
		Solaris
		Windows
utilizationProcUnits	The processor units utilized of this VIOS. (count)	AIX
		HP-UX
		Linux
		Solaris
		Windows

PowerVMVIOSMemory topology object

There is a single *PowerVMVIOSMemory* instance attached to a *PowerVMVIOS* that provides *PowerVMVIOS* level summary metrics for the memory interfaces on the *PowerVMVIOS*.

Table 59. PowerVMVIOSMemory topology object

Metric Name	Metric Description	Metric Collected for These Platforms
entitledMem	The current entitled memory of VIOS. (megabyte)	AIX® HP-UX Linux® Solaris® Windows®

PowerVMVIOSNetwork topology object

There is a single *PowerVMVIOSNetwork* instance attached to a *PowerVMVIOS* that provides *PowerVMVIOS* level summary metrics for the network interfaces on the *PowerVMVIOS*.

Table 60. PowerVMVIOSNetwork topology object

Metric Name	Metric Description	Metric Collected for These Platforms
networkReceiveRate	The combined receive rate on all Ethernet adapters.(byte/second) (byte/second)	AIX®
		HP-UX
		Linux®
		Solaris®
		Windows®
networkSendRate	The combined send rate on all Ethernet adapters.(byte/second) (byte/second)	AIX
		HP-UX
		Linux
		Solaris
		Windows

PowerVMVIOSStorage topology object

There is a single *PowerVMVIOSStorage* instance attached to a *PowerVMVIOS* that provides *PowerVMVIOS* level summary metrics for the storage interfaces on the *PowerVMVIOS*.

Table 61. PowerVMVIOSStorage topology object

Metric Name	Metric Description	Metric Collected for These Platforms
availableSpace	The total available space of all physical volumes on this VIOS. (gigabyte)	AIX® HP-UX Linux® Solaris® Windows®
diskBytesReadRate	The total bytes read rate of all physical volumes on this VIOS.(kilobyte/second) (kilobyte/second)	AIX HP-UX Linux Solaris Windows
diskBytesWriteRate	The total bytes write rate of all physical volumes on this VIOS.(kilobyte/second) (kilobyte/second)	AIX HP-UX Linux Solaris Windows
diskTransmitRate	The total rate sum from diskBytesReadRate and diskBytesWriteRate.(kilobyte/second) (kilobyte/second)	AIX HP-UX Linux Solaris Windows
entitledSpace	The total used space of all physical volumes on this VIOS. (gigabyte)	AIX HP-UX Linux Solaris Windows
totalSpace	The total space of all physical volumes on this VIOS. (gigabyte)	AIX HP-UX Linux Solaris Windows

Processor topology object

Tracks the metrics for a single processor. Multiple processors can be attached to a Host, each identified with a unique name (or number).

Table 62. Processor topology object

Metric Name	Metric Description	Metric Collected for These Platforms
name	The name of the processor, typically just a unique number in the system (0, 1, 2, etc).	AIX® HP-UX Linux® Solaris® Windows®
percentIOWaitTime	The percentage of time that the processor was idle while there was at least one I/O in progress.	AIX HP-UX Linux Solaris
percentIdleTime	The percentage of time that the processor was idle.	AIX HP-UX Linux Solaris Windows
percentSystemTime	The percentage of time that the processor spent executing system code in user mode.	AIX HP-UX Linux Solaris Windows
percentUserTime	The percentage of time that the processor spent executing user code.	AIX HP-UX Linux Solaris Windows
processorType	The type of processor.	AIX HP-UX Linux Solaris
stepping	The stepping revision.	AIX ^a HP-UX ^b Linux
totalHz	The speed (measured in Hertz) of a physical processor, or the number of cycles allocated to a virtual processor.	AIX HP-UX Linux Solaris Windows

Table 62. Processor topology object

Metric Name	Metric Description	Metric Collected for These Platforms
usedHz	The amount of clock cycles (measured in Hertz) being used on a processor.	AIX HP-UX Linux Solaris Windows
utilization	The percentage of time that the processor was utilized (that is, not idle).	AIX HP-UX Linux Solaris Windows

a.This metric is not collected for native.

b.This metric is only available on Itanium systems using non-native collectors.

WinHostProcessInstance topology object

On Windows®, instances of WinHostProcessInstance are used instead of HostProcessInstance. These topology objects extend HostProcessInstance with a few more details that can be captured on Windows platforms.

Table 63. WinHostProcessInstance topology object

Metric Name	Metric Description	Metric Collected for These Platforms
commandLine	The full command line used to launch the executable.	Windows®
creationDate	The time the process instance was created.	Windows
executablePath	The full path to the executable.	Windows

Appendix: Building regular expressions in Foglight

This section describes regular expression basics and gives hands-on examples. For more details, see these topics:

- [What is a regular expression?](#)
- [Where can I find regular expressions?](#)
- [Regular expression basics](#)

What is a regular expression?

A regular expression is a sequence of characters used to describe text ranges, patterns, and various kinds of special conditions. Regular expressions are supported in many programming languages, including Java® and Groovy, the core languages that Foglight uses.

The syntax of Groovy regular expressions comes from Java, so the syntax of Java and Groovy regular expressions is the same.

Where can I find regular expressions?

Regular expressions can be found in many parts of Foglight. For example, credential mappings use the regular expression syntax to select the hosts accessible with a specific credential.

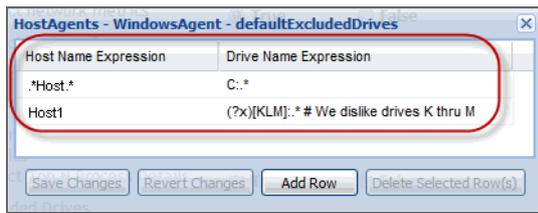
Figure 1. Edit Resource Mapping Condition

Figure 1 shows the 'Edit Resource Mapping Condition' dialog box. The dialog contains the following elements:

- Access Resources Using:** Target Host Name
- Equal or Not Equal:** equals (selected), does not equal
- Matching Type:** Regular Expression (selected)
- Evaluate This Condition:**
- Matching Values (Combined Using Logical 'OR' Operator):**
 - Add a New Regular Expression
 - Table with 1 row: (?)\.*rd-qant.*

The Excluded Drives property of the WindowsAgent is another example of a regular expression usage in Foglight.

Figure 2. Excluded Drives property

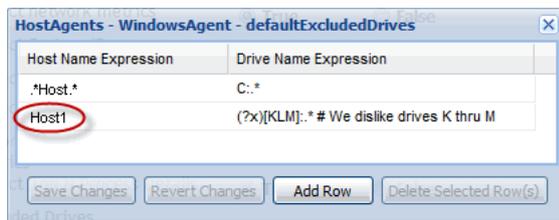


In addition, regular expressions can be found in LogFilter expressions, FxV Hit Analysis, and other parts of Foglight.

Regular expression basics

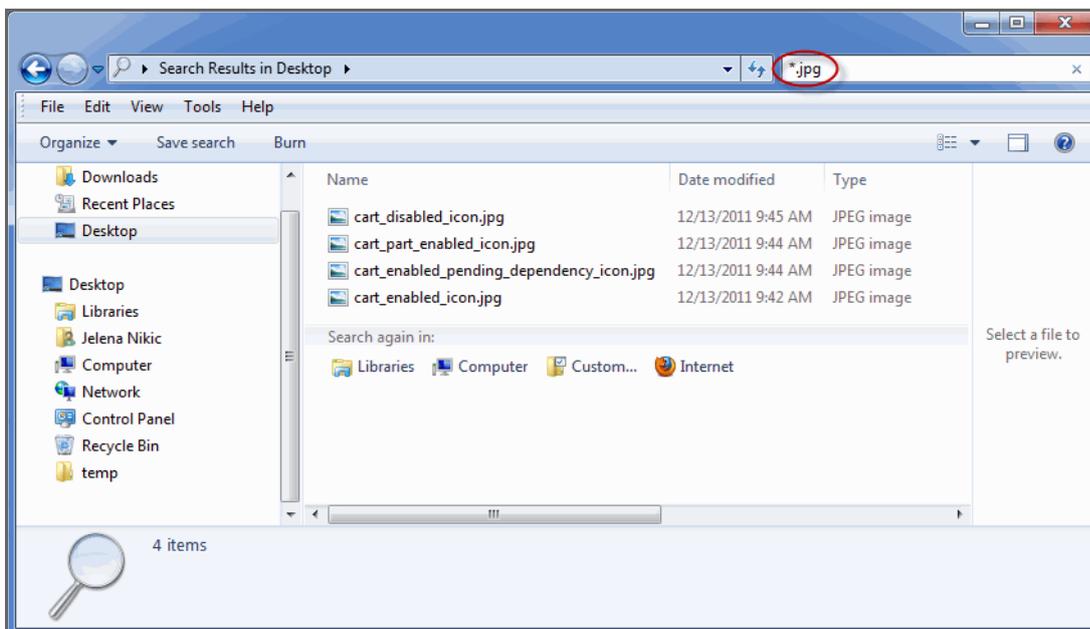
The simplest regular expression is a basic text string containing only alphanumeric characters. For example, Host1, matches the same value, Host1.

Figure 3. Example of a simple regular expression



A regular expression is fundamentally different from a pattern using file name wild cards such as an asterisk '*' or a question mark '?'. For example, issuing a search in Windows® Explorer using *.jpg as a filter matches all files with the .jpg extension.

Figure 4. Windows Explorer search using a wild card



This is **not** a regular expression, it is a simple filter that uses wild cards. An equivalent regular expression that selects all files with the .jpg extension is `.*\.\.jpg`. The details of this regular expression are covered next.

For more details, see these topics:

- [Building a simple pattern](#)
- [Building a pattern that matches a specific character](#)
- [Building a pattern that matches multiple characters](#)
- [Using advanced quantifiers](#)
- [Using special characters and regular expression flags](#)
- [Grouping elements in a pattern](#)
- [Additional information](#)

Building a simple pattern

In this section we look at a simple regular expression pattern that selects all hosts whose names start with `Host`. The completed expression is `Host.*`. This expression contains two special characters:

- A period `.`. This is a special matching symbol that matches any single character.
- An asterisk `*`. This is a quantifier that instructs the regular expression interpreter to match the previous pattern zero or more times.

The final expression, `Host.*`, results in matching any strings that start with `Host` and match any of the following host names:

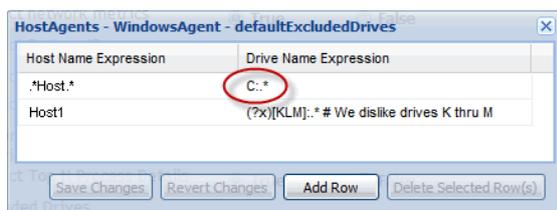
- `Host01`
- `Host-test`
- `Host`

However, the following host names do not match this expression:

- `MyTestHost1`
- `ProductionHst-9000`

We now take a look at a simple pattern that matches a group of similar text strings. A common pattern for selecting Windows® drive names is `C:.*`. A typical usage of this expression is in the Excluded Drives property of the WindowsAgent.

Figure 5. WindowsAgent's Excluded Drive property



Unlike in a simple file matching filter, `C:*`, the equivalent regular expression requires a period between the colon `:` and the asterisk `*`: `C:.*`.

! | CAUTION: Failing to include a period in this particular example can result in data loss.

Next, we build a regular expression that selects all hosts whose name include the string `Host`, not just the ones that start with `Host`. To do that, simply add a prefix to the above expression, `Host.*`, resulting in `.*Host.*`.

This expression matches any hosts that include `Host`, but not necessarily begin with it. That is because the prefix `.*` translates to any combination of zero or more characters appearing before the string `Host`. The expression matches each of the following host names:

- `HostProd01`

- `Host-test`
- `Host`
- `MyHost`
- `DBHost-9000`, but not `DBHst-9000`

Building a pattern that matches a specific character

A regular expression pattern can be used to match a string that includes a specific character. For example, you can build a pattern that matches only the drive letters C through H. To do this, you have two options:

- Specify the entire set of valid choices enclosed in square brackets: `[CDEFGH]`
- Specify a range using the first and last character only: `[C-H]`

Choosing the second option, the resulting regular expression is: `[C-H]:.*`. This means, any letter in the range and including 'C' through 'H', followed by a colon ':', and optionally by more characters.

Building a pattern that matches multiple characters

Using regular expressions you can define a pattern to match a multi-character pattern. For example, you can write a regular expression to match all hosts whose names contain `Host` and are followed by exactly two digits. That means you want to match the following strings:

- `Host01`
- `Host03`

But not:

- `Host01-bckp`

The expression that matches this pattern is: `Host[0-9][0-9]`.

i | **TIP:** Because there is no trailing `.*` in the expression, the host name with the `-bckp` suffix does not match.

Going further, you can write an expression that matches all hosts whose names contain `Host`, followed by exactly two digits, and optionally a lowercase letter. That means you want to match the following strings:

- `Host01`
- `Host03a`

But still not:

- `Host01-bckp`

The expression that matches this pattern is: `Host[0-9][0-9][a-z]?`.

i | **NOTE:** The question mark '?' in regular expressions is also a quantifier, and not a single character wild character like in some file systems.

Using advanced quantifiers

In addition to asterisk `*` and question mark `'?'`, there are additional quantifiers that are supported in regular expressions:

- The plus sign '+' means one or more times. For example, `case[0-9]+` matches `case1`, `case12345`, but not simply `case`.
- The braces "{}" express a range. For example:
 - `{3}` occur exactly three times.
 - `{2,4}` occur at least two times and as many as four times.
 - `{3,}` occur at least three times up to infinity.

You can use this notation instead of the common quantifiers:

- An asterisk '*' has the same meaning as `{0,}`.
- A plus sign '+' has the same meaning as `{1,}`.
- A question mark '?' has the same meaning as `{0,1}`.

Using special characters and regular expression flags

Special characters in regular expressions the regular expressions further extend their flexibility in advanced use cases.

- The caret '^' reverses the meaning of a regular expression element. For example, `[^KLM]` matches a single character that is not K, L, or M.

In the Windows® drive mapping example, the following expression excludes all drives except K, L, or M:
`(?x)[^KLM]:.* # Exclude all drives except K thru M`

- The backslash '\' is an escape character. It indicates that the next character is a special character and interpreted literally.
- A backslash '\' followed by a lowercase 'd', `\d`, means a digit.
- A backslash '\' followed by a lowercase 'w', `\w`, means a word character (an alphanumeric character or an underscore '_'). It has the same meaning as `[a-zA-Z_0-9]`.
- A backslash '\' followed by a lowercase 's', `\s`, means a white space character. It has the same meaning as `[\t\n\x0b\r\f]`.
- Two backslashes "\\" followed by a period '.', `\\.` , means a literal dot.
- The flag `(?i)` makes the regular expression case insensitive. In Foglight for Infrastructure, the Add OS Monitor wizard uses this flag in the Resource Mapping regular expression. For example: `(?i).*host.*`.
- The flag `(?x)` allows you to add comments to explain a complex or unusual pattern. The comment starts with a number sign '#'. For example: `(?x)[KLM]:.* # We dislike drives K through M.`

Grouping elements in a pattern

The regular expression syntax provides a way to group elements together and use groups together with other operators, as required.

- Use parentheses "()" to group elements together.
 - For example, to match a period '.' followed by one or more alphabetic characters, write `(\\.[A-Za-z]+)`.
- Use a common quantifier with a group.
 - For example, to match the above group `((\\.[A-Za-z]+)` zero or more times, write `(\\.[A-Za-z]+)*`.
- Use a pipe '|' with a set of alternatives.

- For example, to use a prefix that is either `dev` or `prod`, write `(dev|prod)`.

Additional information

For more information and more advanced examples, you can refer to the following topics:

- “Appendix: Regular Expressions” in the *Installing the Java EE Technologies Management Capabilities* guide
- “Appendix: Java Regular Expressions in FxV Hit Analysis” in the *Foglight Experience Viewer User Guide*

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit <https://www.quest.com/>.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.