

Foglight[®] 7.3.0
Administration and Configuration
Guide

© 2024 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. Tevron, the Tevron logo, and CitraTest are registered trademarks of Tevron, 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.

Contents

Administering and Configuring Foglight	6
Configuring Foglight for initial use	6
Logging in to Foglight	7
Managing Licenses	11
Understanding subscription licenses	12
Allowing URLs for subscription license management	12
Configuring email notifications	12
Managing Users and Security	13
Suspending Alarms and Data Collection	14
Managing Support Bundles	15
Viewing Client Connection Status	15
Viewing Audited Entries	15
Viewing Configuration Details	16
Enabling Usage Feedback	16
Setting up Proxy Configuration	16
Enabling Automatic Communication with Quest	17
Extending Your Monitoring Reach with Foglight Cartridges	19
Configuring Foglight Agents for Host Monitoring	19
Editing Agent Properties by Type	20
Moving FglAM agents	21
Viewing Agent Adapters	22
Controlling System Access with Credentials	22
Administering Foglight	25
Recommended Self-Monitoring Automation	26
Cyclic Maintenance	26
Using a single pane of glass for your administration needs	28
Manage Foglight Database Performance	28
Monitor Server Performance	29
Monitor Agent Manager Performance	30
Associate Service Objects with Groups and Tiers	31
Back Up and Restore Foglight	31
Configure Rules and Metric Calculations to Discover Bottlenecks	32
Working with Derived Metrics	33
Working with Metric Thresholds	34
Analyze Activity Levels with IntelliProfile	35
Working with Foglight Registry Variables	36
Associate Metric Calculations with Schedules	37
Manage Data Retention	38
Expand Your Collection of Topology Types	39
Customizing Your Foglight Environment with Tooling	40
Merging Host Objects	40
Building Script Agents	41
Using the Query Language	42

Retrieving Data with the REST API	42
Retrieving Data with Scripts and Queries	42
About Us	44
Technical support resources	44

Administering and Configuring Foglight

Foglight comes with the following additional built-in roles which control access to the dashboards in Foglight. This *Administration and Configuration Guide* provides conceptual information about Foglight® administration, and instructions for using the administration dashboards. It contains an overview of the administration features and their location in the browser interface.

This guide is intended for Foglight system administrators who need to administer and configure Foglight.

- Administrators who are new to Foglight can find information related to first-time use in the first two topics: [Configuring Foglight for initial use](#) on page 6, and [Extending Your Monitoring Reach with Foglight Cartridges](#) on page 19.
- For day-to-day tasks, see [Administering Foglight](#) on page 25 and [Configure Rules and Metric Calculations to Discover Bottlenecks](#) on page 32.
- Advanced Foglight administrators can find information on key tools used to manage Foglight in [Customizing Your Foglight Environment with Tooling](#) on page 40.

For more information about specific administration tasks, or additional technical information that further describes Foglight administration features, see the *Administration and Configuration Help*.

- [Configuring Foglight for initial use](#)
- [Logging in to Foglight](#)
- [Managing Licenses](#)
- [Configuring email notifications](#)
- [Managing Users and Security](#)
- [Suspending Alarms and Data Collection](#)
- [Managing Support Bundles](#)
- [Viewing Client Connection Status](#)
- [Viewing Audited Entries](#)
- [Viewing Configuration Details](#)
- [Enabling Usage Feedback](#)
- [Setting up Proxy Configuration](#)

Configuring Foglight for initial use

Foglight® collects data from monitored hosts and builds models with tree-like structures in real time. Foglight administration capabilities allow you to configure the hosts for monitoring, dictate how the data is collected, restrict user access, and build and edit flexible rules to implement your business logic. The type and range of administration steps depends on the complexity of your monitoring needs.

The Foglight browser interface includes a set of dashboards that have administration capabilities. To access them, your user account must belong to a group with the Administration role. Administrators can manipulate agents, rules, derived metrics, registry variables, cartridges, types, and scripts.

In most environments, one of the first things you are prompted to do after installing Foglight and logging in to the browser interface is to install a valid license. Next, you configure email actions, to ensure that Foglight can send

email to interested parties when pre-defined thresholds are reached. If you also manage user access, you can create user accounts and assign the appropriate permissions.

i | **IMPORTANT:** For more information about specific tasks, or additional technical information about the features described in this topic, see the reference topics in the Administration and Configuration Help. For example, to find out more about Foglight licensing, in the browser interface, open the Help tab in the action panel. From there, navigate to **Administration and Configuration Help > Administering and Configuring Foglight > Managing Licenses**, to find a list of reference topics.

Logging in to Foglight

The Foglight® user interface runs inside a web browser. Before you log in, ensure that your Foglight Management Server is up and running, and obtain your user name and password. The default account, *foglight/foglight*, provides full access to the browser interface.

You can access the browser interface by opening a web browser instance and navigating to the Management Server URL, which uses the following syntax: <http://<localhost>:<port>>, where *localhost* and *port* are the name of the computer and port number on which the Management Server is running. The security settings associated with your user account determine which dashboards you can access.

The first page you see also depends on the edition of Foglight that you are running:

- [Foglight users](#): on page 7
- [Foglight Evolve users](#): on page 7

Foglight users:

The first time that you log in to Foglight, the **Welcome** page appears in the display area.

The appearance of the Welcome page depends on your user permissions. If your user account belongs to a group that has the Administrator role, you see the Welcome page with a list of the common administration tasks that you typically perform upon logging in to Foglight.

To access the full Administration dashboard, in the navigation panel, under **Homes**, click **Administration**.

Foglight Evolve users:

The first time that you log in to Foglight Evolve, the **Getting Started** tab of the Environment Overview page appears in the display area.

To access the Environment Overview dashboard, in the navigation panel, under **Homes**, click **Environment Overview**.

Figure 1. Foglight Evolve Environment Overview dashboard

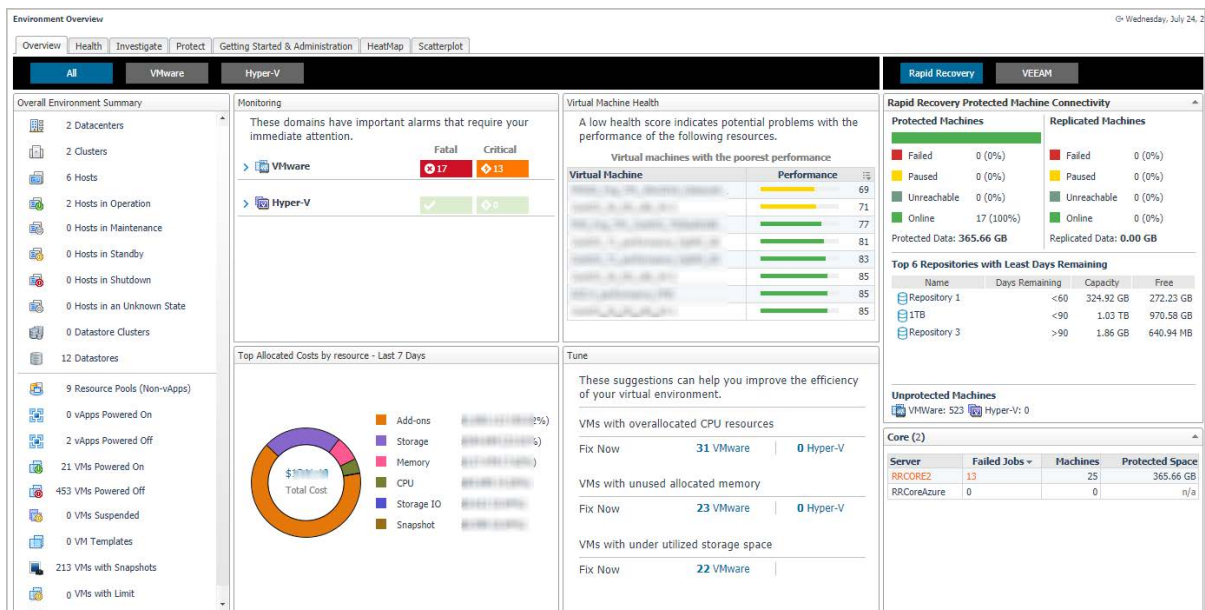


Table 1. Environment Overview dashboard

Tab	Views or tiles in the tab
Overview	<ul style="list-style-type: none"> • Overall Environment Summary. Provides a summary of your virtual environment. • Monitoring. Highlights the alarms that need immediate attention, for each monitored domain. Drill down into the available links, to review the alarms in more detail. • Top Allocated Costs by resource - Last 7 Days. Provides a list of the top allocated costs by resource within the last 7 days. • Virtual Machine Health. Provides a list of virtual machine with the lowest health score in your environment. To investigate the potential problems affecting the health of a resource, click the virtual machine name or the performance icon. • Tune. Provides a list of suggestions for making your virtual environment more efficient. Clicking the Fix Now link opens the Optimizer Main View dashboard, which allows you to optimize your environment, as necessary. • Rapid Recovery/VEEAM. Provides an overview about the protected and unprotected virtual machines in the monitored environment.
Health	<ul style="list-style-type: none"> • VMware. Highlights the alarms that need immediate attention, in your VMware® environment. Drill down into the available links, to review the alarms in more detail. • Hyper-V. Highlights the alarms that need immediate attention, in your Hyper-V® environment. Drill down into the available links, to review the alarms in more detail. • Storage. Highlights the alarms that need immediate attention, in your Storage Management environment. Drill down into the available links, to review the alarms in more detail.

Table 1. Environment Overview dashboard

Tab	Views or tiles in the tab
Investigate	<ul style="list-style-type: none">• Resource Efficiency. Provides a list of suggestions for making your virtual environment more efficient, for each type of virtual object. Clicking the Optimize or View more details links opens the Optimizer Main View dashboard, which allows you to optimize your environment, as necessary.• Highest Impact Changes (VMware only). Displays the change events of objects and the performance impact to the change objects.• VMs with CPU Problems. Provides a list of virtual machines experiencing CPU problems. Clicking a virtual object to drill down and investigate the issues in more details.• VMs with Memory Problems. Provides a list of virtual machines experiencing memory problems. Clicking a virtual object to drill down and investigate the issues in more details.• VMs with Storage Problems. Provides a list of virtual machines experiencing storage problems. Clicking a virtual object to drill down and investigate the issues in more details.• VMs with Network Problems. Provides a list of virtual machines experiencing network problems. Clicking a virtual object to drill down and investigate the issues in more details.

Table 1. Environment Overview dashboard

Tab	Views or tiles in the tab
Protect	<ul style="list-style-type: none"> <li data-bbox="448 277 671 300">• Rapid Recovery <ul style="list-style-type: none"> <li data-bbox="528 322 1378 376">▪ <i>Core</i>: This table lists the monitored Rapid Recovery Core Servers. Click this link to navigate to the Rapid Recovery Infrastructure Tab. <li data-bbox="528 398 1390 479">▪ <i>Machine Connectivity</i>: This table shows the connectivity state of machines protected and replicated on the monitored Rapid Recovery core. It also shows connectivity for data on recovery points-only machine. <li data-bbox="528 501 1374 613">▪ <i>Trouble Monitor</i>: This table shows job activity, connections with the license portal, and transfer activity to detect early on the monitored Rapid Recovery core. The time range is configurable, defaulting to last 24 hours. Click this graph to navigate to the Events > Journal tab. <li data-bbox="528 636 1353 748">▪ <i>Top 8 Repositories with Least Days Remaining</i>: This tables shows the top eight repositories with the least days remaining on the monitored Rapid Recovery core. Click this link to navigate to the Repositories Capacity Planning Tab - Rapid Recovery. <li data-bbox="528 770 1374 904">▪ <i>Transfer Job per Machine</i>: This table shows, by protected machine of which the latest transfer job is failed, the number of successful and failed transfer jobs in the specified time range. This table shows the top ten machines with the most failed transfer job. The time range is configurable, defaulting to last 24 hours. Click this graph to navigate to the Events > Tasks tab. <li data-bbox="528 927 1366 1039">▪ <i>Transfer Jobs</i>: This table shows all snapshot data transfers (including base images and incremental snapshots) that completed in the specified time range. The time range is configurable, defaulting to last 24 hours. Click this graph to navigate to Events > Tasks tab. <li data-bbox="448 1061 576 1084">• VEEAM <ul style="list-style-type: none"> <li data-bbox="528 1106 1378 1160">▪ <i>Backup Servers</i>: This table lists the monitored Veeam Backup Core Servers. Click this link to navigate to the Veeam Infrastructure Tab. <li data-bbox="528 1182 1385 1330">▪ <i>Protected VMs Overview</i>: This table presents the information about how your VMs are protected, number of protected VMs (backed up or replicated), number of restore points available, source VM size, full and incremental backup size, and successful backup sessions ratio on the monitored Veeam Backup Server. <li data-bbox="528 1352 1278 1406">▪ <i>Backup Window</i>: This table shows the total duration of Backup and Replication jobs. <li data-bbox="528 1429 1378 1541">▪ <i>Top 8 Repositories with Least Free</i>: This tables shows the top eight repositories with the least days remaining on the monitored Veeam Backup server. Click this link to navigate to the Repositories Capacity Planning Tab - Veeam. <li data-bbox="528 1563 1374 1697">▪ <i>Transfer Job per Machine</i>: This table shows, by protected machine of which the latest transfer job is failed, the number of successful and failed transfer jobs in the specified time range. This table shows the top ten machines with the most failed transfer job. The time range is configurable, defaulting to last 24 hours. Click this graph to navigate to the Jobs tab. <li data-bbox="528 1720 1366 1832">▪ <i>Transfer Jobs</i>: This table shows all snapshot data transfers (including base images and incremental snapshots) that completed in the specified time range. The time range is configurable, defaulting to last 24 hours. Click this graph to navigate to Jobs tab.

Table 1. Environment Overview dashboard

Tab	Views or tiles in the tab
Getting Started & Administration	<ul style="list-style-type: none"> • Monitor a new virtual domain. Depending on the virtual technology that you want to monitor, select one of the following options: <ul style="list-style-type: none"> ▪ Connect to VMware Virtual center. Launches a wizard to start the creation and configuration of a VMware Performance agent. ▪ Connect to Hyper-V servers. Launches a wizard to start the creation and configuration of a Hyper-V agent. • Monitor SAN Storage. Allows you to monitor the SAN Storage, provided that the Storage Management license trial has been activated. <ul style="list-style-type: none"> ▪ Start monitoring storage devices. Launches a wizard to start creating Storage Management agents. For more information, see the <i>Foglight for Storage Management User and Reference Guide</i>. • License Management. All Foglight components are license-protected. Some components are covered by the base Foglight license, while others require an additional license. Foglight restricts access to only those features enabled by active license files. In a typical installation, you need a license for the Management Server, along with a license for each license-protected component that exists in your monitoring environment. License-protected components may be installed on the server prior to installing the corresponding license, however such components will be disabled until a valid license is installed. <ul style="list-style-type: none"> ▪ Add a new license. Allows you to quickly access the Manage Licenses dashboard. For more information on licenses, see the <i>Administration and Configuration Guide</i>. ▪ FE Licenses Overview. Get an overview of Foglight Evolve licenses information from this page.
HeatMap	Shows a Heat Map of the virtual machines monitored in the VMware or Hyper-V environment.
Scatterplot	Shows a Scatterplot of the virtual machines monitored in the VMware or Hyper-V environment.

Managing Licenses

Foglight® includes a licensing capability that restricts access to only those features that are defined in the license file. A server installation requires a license file that provides access to the server-specific part of the browser interface and the features associated with it.

All Foglight cartridges are license-protected. The base Foglight license covers some cartridges, while others may require an extra license. The cartridges included with the server do not require any additional license. The Foglight Agent Manager and Infrastructure cartridges fall into this category. Some cartridges installed on the server require an extra license.

In a typical installation, you need a license for the Management Server, along with a license for each license-protected cartridge that exists in your monitoring environment. If a cartridge requires a license, install the license on the server immediately after installing or upgrading that cartridge. Foglight allows you to install a license-protected cartridge on the server before installing its license, however it disables the cartridge until a valid cartridge license is installed.

Foglight can limit the number of monitored instances of certain object types according to the restrictions in the associated licenses. When that limit is reached, the server stops creating instances of that type and triggers an alarm with a message indicating the problem. For example:

```
The licensed limit of 5 instances of type Host has been reached. New instance of type Host is rejected.
```

You can install and manage Foglight licenses using the Manage Licenses dashboard. To access this dashboard, from the **Administration** dashboard, in the **Setup** column, click **Manage Licenses**.

Understanding subscription licenses

Foglight Flex subscription licenses calculate usage differently for Foglight Evolve and Foglight for Databases. Review the content below to ensure you understand license consumption.

Foglight Evolve

Foglight Evolve users who are subscribed to use a Foglight Evolve Flex subscription will be billed at the end of the billing period following your actual usage.

Your usage will be calculated based on the number of hosts you are monitoring.

- i** | **NOTE:** If you are monitoring a Container Host which is running as a Hyper-V or Google Cloud virtual machine, the container host and Hyper-V VM or Google Cloud virtual machine will each be counted as one usage unit.
- i** | **NOTE:** If you are monitoring a Google Cloud VM without Stackdriver, and if it will also be monitored via an Infrastructure Cartridge or AD/Exchange agent, each one of them will be counted as one usage unit.

Foglight Databases

Foglight for Databases users who are subscribed to use a Foglight for Databases Flex subscription will be billed at the end of the billing period following your actual usage.

Your usage will be calculated based on the number of instances you are monitoring.

- i** | **NOTE:** When using Foglight for Azure SQL Database, each database will be counted as a monitored instance.

Allowing URLs for subscription license management

In order to ensure that subscription licenses function properly, ensure that the URLs listed below are not blocked by your firewall.

- License management: <https://msp-api.licenseportal.com:8093>
- Usage data collection: <https://msp-api.licenseportal.com:8091>

Configuring email notifications

Foglight® uses email notifications to send reports or alarm-related messages to email recipients when certain thresholds are reached. This notification can happen, for example, when a rule enters a particular state. Foglight can also send reports to email recipients.

Email settings are stored in the Foglight registry. To ensure delivery of email messages to selected recipients, configure Foglight to use your email server along with an existing email account. Use the Email Configuration dashboard to configure your email. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Setup > Email Configuration**.

Managing Users and Security

Foglight® controls user access using the concept of users, groups, and roles. Each user can belong to one or more groups. The roles assigned to those groups determine the set of actions that the user can access. For example, if your user account belongs to a group that includes the Administration role, you can access the administrative dashboards in the browser interface.

Each Foglight user has a user name and a password and can belong to one or more groups. Foglight can store user passwords on the Management Server, or in an external directory.

The Users & Security dashboard allows you to manage user access. To access this dashboard, your user account must belong to a group with the Security Administration role.

To access this dashboard:

- 1 On the navigation panel, click **Dashboards > Administration**.
- 2 In the **Administer Server** column, click **Users & Security**.
- 3 To start managing user access, click **Manage Users, Groups, Roles**.

i | **NOTE:** The Administration dashboard has been updated for this release. If you prefer to continue working with the previous version of the dashboard, click **Use 5.6 Administration view**. This link is located in the lower right corner of the dashboard.

Configuring Password settings

Password settings define the restrictions that apply to passwords for Foglight users. Foglight allows you to specify similar policies for its passwords and users that are likely in place in your corporate environment, including:

- password expiry dates and user warning
- password retries before user lockout
- password length, complexity, and reuse of old passwords
- User name length policies

To view and edit Password settings, on the main Users & Security Management dashboard, click **Password Policy Settings**.

Configuring directory services

Foglight supports the Lightweight Directory Access Protocol (LDAP version 3). This security feature allows Foglight to access user account information that is stored in an external directory. The following directory services are supported:

- Microsoft® Active Directory®
- Oracle® Directory Server Enterprise Edition
- OpenLDAP®
- Novell® eDirectory™

To configure directory services, you must be familiar with the details of your LDAP directory service. After configuring the LDAP directory service, Foglight creates a user account each time an LDAP user successfully logs in to Foglight for the first time. Any password changes in the LDAP directory service are transparent to Foglight. After a user's password changes in the directory service, that user can log in to Foglight with the new password while any attempts to use the old password fail. If a user account is removed from the directory service, any login requests with those credentials result in a failure. Similarly, if the LDAP Authentication Service is down, Foglight cannot authenticate any of the users whose accounts are defined in the LDAP directory service. Any internal Foglight users, such as the default *foglight* account, or any accounts that you create, are unaffected during LDAP authentication interruptions.

i | **IMPORTANT:** After configuring the LDAP directory service, import LDAP groups into Foglight and grant them access permissions to enable their users to access the browser interface. Failure to do so prevents them from using the browser interface. For complete information, see “Importing and configuring LDAP groups” in the Online help.

The following considerations are important when planning to integrate an external directory service with the Management Server:

- Secure LDAP is supported, but not required.
- LDAP with Transport Layer Security is not supported.
- A persistent connection to the LDAP server is not required.

You can track user login credentials using the **Users** tab, accessible from the **User Management** view. This tab lists the users who have logged in to Foglight using their external account credentials.

To view and edit external directory settings, on the main Users & Security Management dashboard, click **Directory Services Settings**.

Configuring user sessions

A user session takes place during the time a user is logged in to the Management Server. Depending on your needs, you can configure Foglight to log out any users that are inactive after a specific length of time, or have user sessions that never time out.

Assigning an ID to a Foglight Server

An administrator is now able to assign an ID to the Foglight Server. This shows up on the login page and in the masthead. This allows users with multiple Foglight Management Servers to distinguish between them.

It can be set at the bottom of the Administration screen. Click the pencil icon, enter the server identifier and save. A page reload or logout/login is needed to refresh the masthead.

Suspending Alarms and Data Collection

A blackout is a period where normal monitoring activities are suspended due to some administrative preference. Blackouts are commonly created to prevent frequent alerts during scheduled maintenance periods.

Foglight® collects data about your system and dynamically builds topology models at run-time. A topology model consists of nodes, where each node is a topology object instance. A set of blackout management dashboards allow you to disable alarms and data collection for a specific period. Suspending alarms involves assigning blackout periods to topology objects. Suspending data collection is slightly different in that it involves assigning blackout periods to specific agent instances. An agent blackout is a scheduled event during which the agent does not collect data. Unlike agent blackouts, topology object blackouts do not interrupt the data collection for the object to which the blackout is assigned. Blacking out a topology object means that no rules analyze that object during the blackout. For more information about topology models, see the *Data Model Guide*.

Blackouts can be applied to dynamic managed components or services. If an object becomes part of a blacked out dynamic managed component or service, it is included in the blackout, even if the blackout is already in effect. Similarly, if an object is removed from a blacked out service or dynamic managed component during the blackout, it ceases to be blacked out.

i **IMPORTANT:** In addition to the features available on the blackout management dashboards, topology and agent blackouts can also be configured using the command-line interface. However, the mechanism for creating blackouts from the command line is independent. It is not recommended to use both methods on the same Foglight Management Server. If you choose to use the command line for creating blackouts, delete all blackouts created with the command line before using the browser interface. If you want to switch from the command line to the blackout management dashboards, use the conversion script to convert the existing blackouts created with the command line. This way all blackouts can be managed in one location. To see a list of existing blackouts that were created using the command line, issue the `topology:blackouts` and `agent:showschedule fglcmd` commands. For more information about these commands, see the *Command-Line Reference Guide*. For more information about the conversion script, see the *Foglight Upgrade Guide*.

You can configure blackouts as recurring events, by associating them with an existing schedule, or as one-time events. Existing blackouts can be deleted or edited, as required.

To access the Blackouts dashboard, from the **Administration** dashboard, click **Blackouts**.

Managing Support Bundles

Foglight® allows you to gather diagnostic data and save it as a collection of files, called a *support bundle*. Support bundles can be forwarded to Quest Support, upon their request. There are two types of support bundles: *server support bundles* and *Foglight Agent Manager support bundles*.

Each server support bundle contains a diagnostic snapshot of the Management Server, log files, and a list of cartridges installed on the Management Server. The logs contain information specific to the operating environment like IP addresses, host names and user actions that have caused changes to the monitoring environment. The bundle is a simple zip file and can be expanded to view the contents which consist of plain text and PDF files. Foglight saves support bundles as ZIP files in the `<foglight_home>/support/<user_name>` directory on the machine hosting the Management Server.

The Support dashboard allows you to create server support bundles, and to download, or delete Foglight Agent Manager and server support bundles. You can also use this dashboard to generate and download host (Agent Manager) support bundles.

To access the Support dashboard, from the **Administration** dashboard, in the **Support** column, click **Support Bundles**.

Viewing Client Connection Status

Foglight® tracks connections to the Management Server for security purposes. The Connection Status dashboard lists browsers that are connected to the server. For each browser instance, the list shows the client's IP address, login time, request name, and request time.

To access the Connection Status dashboard, from the **Administration** dashboard, in the **Server** column, click **Connections**.

Viewing Audited Entries

Foglight® generates security and change logs that contain information about the users who are authenticated upon logging in to Foglight, and user management or configuration changes, such as changes to Foglight registry and rules. You can use the View Audit Information dashboard to look at individual log entries.

Each entry shows:

- the date and time at which the operation is performed
- the name of the user who initiated the operation
- the name of the service that performed the operation
- the operation name

To access the View Audit Information dashboard, from the **Administration** dashboard, in the **Server** column, click **View Audit Information**.

Viewing Configuration Details

Foglight® configuration settings include the basic environment parameters for host and port settings, virtual memory, server federation, and many others. Other types of settings reflect the version and patch level of various components such as the Management Server, WCF, and JVM versions; these settings cannot be changed unless you choose to upgrade to a later version of Foglight. More display-only settings indicate the OS of the computer on which the Management Server is installed, and its patch level.

The configuration values are set in the *server.config* file and the Foglight registry. For example, the database settings are typically set in *server.config*, while global email settings are specified in the Foglight registry. Editing the configuration file requires a restart of the Management Server in order for these changes to take effect. Changes to the Foglight registry do not require a system restart.

Foglight configuration settings can be viewed using the Management Server Configuration dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Setup > Management Server Configuration**.

Enabling Usage Feedback

The Foglight® Management Server can collect usage data about your environment and send it to Quest to improve support response. This data helps with identifying potential bottlenecks, improving the overall Management Server performance and server versions going forward.

The collected usage data contains information about the visited dashboards, unique ID of management server, server version, along with configuration details. It does not identify any users or provide additional information about their actions in the user interface. For additional information about the overall protection of user-related data, see the *Security and Compliance Guide*.

This feature may be enabled by default. To turn it off, click **Disable** on the *Automatic Communication with Quest* dashboard. To access this dashboard, from the **Administration** dashboard, in the **Support** column, click **Support Notifications**.

Setting up Proxy Configuration

The Foglight Management Server often reside behind a firewall in some production environments, therefore some Foglight functionalities, for example License Entitlement, do not work within these environments. Foglight provides a new Proxy Configuration feature, which enables the Foglight Management Server to connect to QorePortal and collect Foglight usage data through the preconfigured proxy server.

This feature may be disabled by default. To turn it on, click **Enable** on the *Proxy Configuration* dashboard. To access this dashboard, from the **Administration** dashboard, in the **Administer Server** column, click **Proxy**.

To add or edit proxy settings:

- 1 On the navigation panel, click **Dashboards > Administration**, and then click **Proxy** in the **Administer Server** column.
 - i** **NOTE:** The Administration dashboard has been updated for this release. If you prefer to continue working with the previous version of the dashboard, click **Use 5.6 Administration view**. This link is located in the lower right corner of the dashboard.
- 2 The *Proxy Configuration* dashboard opens with the proxy settings disabled by default.
- 3 To add or edit proxy settings, click **Edit Proxy Settings**.
- 4 In the **Proxy Settings Editor** dialog box, specify the following fields, as needed:
 - *Proxy Type*: the proxy type can be either *HTTP* or *Socks*.
 - *Proxy Server Address*: input the host name or IP address of the proxy server.
 - *Proxy Port*: input the port number of the proxy server.
 - *Credential Required*: indicates whether user authentication is required for the proxy server.
 - *Username*: input the user name. This field is available only when the *Credential Required* field is selected.
 - *Password*: input the password. This field is available only when the *Credential Required* field is selected.
 - *Confirm Password*: input the password again to confirm. This field is available only when the *Credential Required* field is selected.
- 5 (Optional) Click **Test Connection** to check if the proxy server works. For more information, see [To check if the proxy server works](#): on page 17.
- 6 Click **Save**. The Foglight Management Server refreshes and changes are saved.

To check if the proxy server works:

- 1 On the navigation panel, click **Dashboards > Administration**, and then click **Proxy** in the **Administer Server** column.
 - i** **NOTE:** The Administration dashboard has been updated for this release. If you prefer to continue working with the previous version of the dashboard, click **Use 5.6 Administration view**. This link is located in the lower right corner of the dashboard.
- 2 The *Proxy Configuration* dashboard opens with the proxy settings disabled by default.
- 3 To test the proxy connection, click **Test Proxy Connection**.
- 4 In the **Test Proxy Connection** dialog box, input an available URL, as needed.
- 5 Click **Test**. The Foglight Management Server checks the proxy connection in the background and returns the test result accordingly.

Enabling Automatic Communication with Quest

The Foglight® Management Server can communicate with Foglight to check for important updates and to send usage data to Quest. Enabling this setting also allows usage data to be sent to Quest including visited dashboards, unique ID of Management Server, Server version and configuration details. This information is then displayed on the *Automatic Communication with Quest* dashboard or can be viewed in a popup by hovering over the Update link on the Administration page.

When this setting is enabled, you will receive messages:

- If an update is available for your Foglight Management Server.
- To communicate important messages about Foglight. For example, “We are aware of the Heart Bleed bug”.
- To communicate known issues to the user about a component you already have installed.

Notifications are sent to users with the Administrator role. By default, this feature is turned on. To turn it off, from the *Automatic Communication with Quest* dashboard, click **Disable** next to the *Communication with Foglight.com* part. To access the *Automatic Communication with Quest* dashboard, from the **Administration** dashboard, in the **Support** column, click **Support Notifications**.

The Foglight Management Server also integrates with the Quest QorePortal that enables access to your Foglight installation from any Internet enabled device. Integration with the Quest QorePortal is available for users with the Administrator role and is turned off by default. To turn in on, implement either of the following:

- Click **Enable** under the QorePortal *Integration* part on the *Automatic Communication with Quest* dashboard.
- Click **Yes** in the prompted **QorePortal Integration** dialog box when starting up the Foglight Management Server after installation.

If you need more information about how to set up the connection with this portal or see the *Last Connection Status: Connection failed.* message under the *QorePortal Integration* part, refer to [Setting up Proxy Configuration](#) on page 16.

Extending Your Monitoring Reach with Foglight Cartridges

Each Foglight® cartridge contains extensions for monitoring a specific environment, such as applications, operating systems, or database management systems. Cartridges are installed on the server. A cartridge can contain one or more agents that are used to collect data from monitored environments.

i | **IMPORTANT:** For more information about specific tasks, or additional technical information about the features described in this chapter, see the reference topics accessible from the applicable section in the Administration and Configuration Help. For example, to find out more about managing Foglight cartridges, in the browser interface, open the Help tab in the action panel, and from there, navigate to **Administration and Configuration Help > Extending Your Monitoring Reach with Foglight Cartridges > Configuring Foglight Agents for Host Monitoring**. You will find a list of reference topics at the bottom of the page.

The Foglight Management Server includes a set of default cartridges that contain models, views and monitoring policies. In addition to these server-specific cartridges, there are other types of cartridges that are distributed separately from the server. They extend the server's monitoring capabilities, allowing you to monitor specific types of environments, such as operating systems, processes, databases, applications, hosts, and others. These types of cartridges typically include agent packages, agent adapters, monitoring policies, and dashboards.

A cartridge monitoring policy contains settings that help Foglight analyze the data that the agents collect, such as rules, registry variables, schedules, and derived metrics. The items included in the monitoring policy are specific to each cartridge type. The dashboards included with a cartridge allow the information collected by the agents to be displayed in a unified view.

The Management Server includes the Cartridge Inventory dashboard, which contains controls for installing, enabling, disabling, and uninstalling cartridges, and for viewing information about the installed cartridges. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Cartridges > Cartridge Inventory**.

Some cartridges include additional components, such as agent installers or additional configuration files. After cartridge installation, these components are available for download from the Foglight Management Server using the Components For Download dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration**. On the Administration dashboard, click **Component Download** in the **Cartridges** column.

Configuring Foglight Agents for Host Monitoring

Foglight® agents collect data from monitored environments and send it to the Management Server. Each agent type can monitor a specific part of your environment, such as an operating system, application, or server. Foglight cartridges that you install on the server include one or more agent types. Foglight also includes internal agents that monitor Foglight components and services.


Foglight uses the Foglight Agent Manager to communicate with monitored hosts. A server installation includes an embedded Foglight Agent Manager. The embedded Agent Manager starts up and stops with the Foglight Management Server. This embedded instance can be used to monitor the host on which the Management Server is installed. To monitor additional hosts in your environment, most cartridges require an Agent Manager installation on each host computer.

After installing and enabling a cartridge, and downloading remaining cartridge components, deploy its agent package to each host running an Agent Manager that you want to monitor. After deployment, create an agent

instance for each monitored host, edit the agent's properties, and start its data collection. See the *Installation and Setup Guide* set for information on installing the Agent Manager on the hosts that you want to monitor.


There are two dashboards that allow you to deploy agent packages and create agent instances: Agent Status and Agent Managers. To access the Agent Managers dashboard, from the navigation panel, click **Dashboards > Administration > Agents > Agent Managers**.

Figure 2. The Agent Managers dashboard in the Foglight browser interface.



To access the *Agent Status* dashboard, from the navigation panel, click **Dashboards > Administration > Agents > Status**.

Figure 3. The Agent Status dashboard.



Editing Agent Properties by Type

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. Foglight stores agent properties on the Management Server.

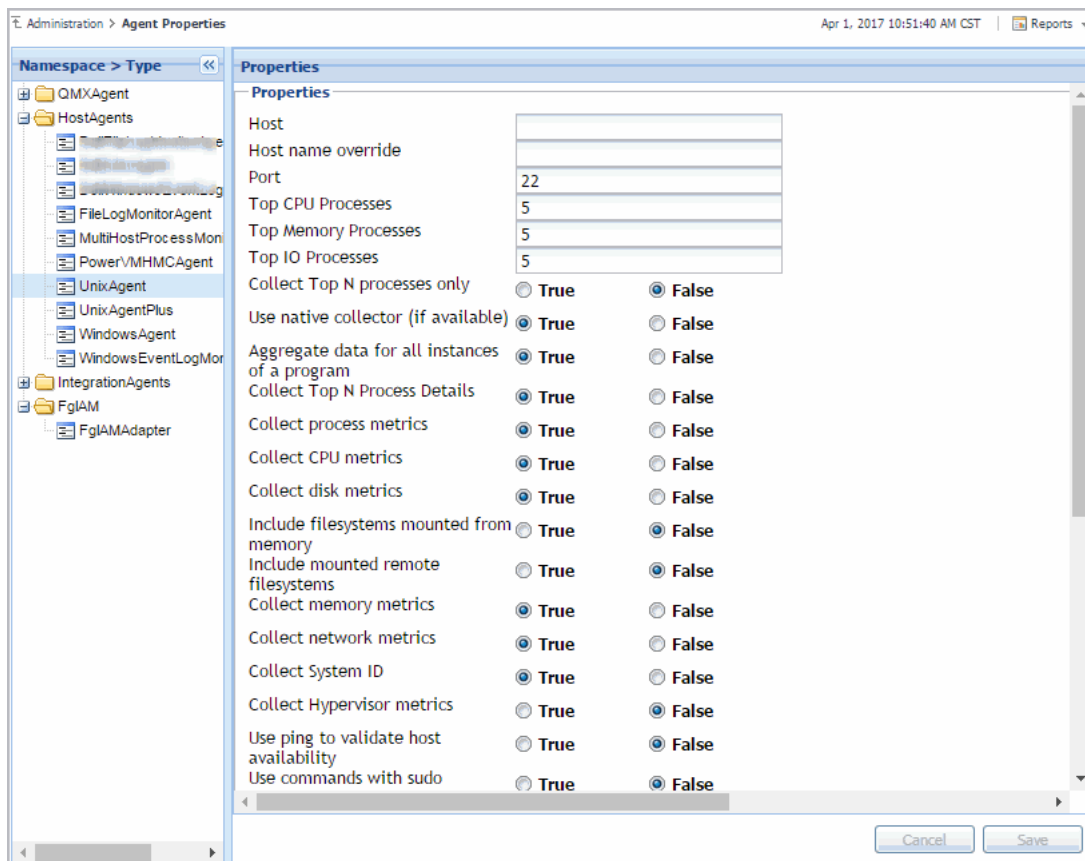
You can edit the properties of agent types (*global properties*), or the properties of individual agent instances (*private properties*). Typically, after creating an agent instance, before its activation, you edit the agent properties. This causes the agent property values to become private. Agent instances with private properties are not affected by any subsequent changes made to its global properties. Changes to global properties are only carried over to new agent instances that are created after those changes.

Another classification of agent properties refers to their format, complexity, and shareability. *Primary* properties are typically text, number, or Boolean values. These types of values are easily specified and changed. *Secondary* or *shared* properties are in list form and contain multiple values organized into rows and columns. They can be easily cloned and modified to suit the needs of individual agent instances. The difference between primary and secondary (shared) properties is that secondary properties can be shared amongst multiple agent instances, meaning that a list can be associated with multiple agent instances of the same type. However, any changes made to a shared list affect all of those agents.

Any passwords that are defined in agent properties are encrypted. This feature is useful in situations when a database password is defined in agent properties, and there are multiple databases in your environment, for example, a Foglight database and a production database. Encrypting database passwords prevents unauthorized database administrators from accessing the database.

Agent properties can be edited using the Agent Properties dashboard. You can use it to edit the properties globally, for all instances of the same type, or only for a specific agent instance. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Agents > Agent Properties**, or select an agent instance on the Agent Status dashboard and click **Edit Properties**.


Figure 4. The Agent Properties dashboard.



Moving FglIAM agents

The *Agent Status* dashboard allows you to move the FglIAM agents within various FglIAM hosts, if needed. To access the *Agent Status* dashboard, from the navigation panel, click **Dashboards > Administration > Agents > Status**.

To move an FglIAM agent:

- NOTE:** The agent movement function is only applicable for the FglIAM agent which type is one of the following: *FileLogMonitorAgent*, *UnixAgentPlus*, *WindowsEventLogMonitorAgent*, *MultiHostProcessMonitorAgent*, *PowerVMHMCAGENT*, *UnixAgent*, and *WindowsAgent*.
- 1 On the navigation panel, under *Dashboards*, click **Administration > Agents > Status**.
The *Agent Status* dashboard opens.
- 2 Select an agent that you want to move from the agents list, and click the move icon .

The **Move selected agents to a different Agent Manager** dialog box opens.
- 3 Use the group selector to select the target FglIAM host from the **Select Agent Manager** drop-down list, and click **Move**.
- NOTE:** The target FglIAM host must have the same lockbox as the source FglIAM host. For more information, see [Managing lockboxes](#) on page 23.

Several Operating System (OS) and environment-specific issues, which include the following, may arise when moving an FglIAM agent.

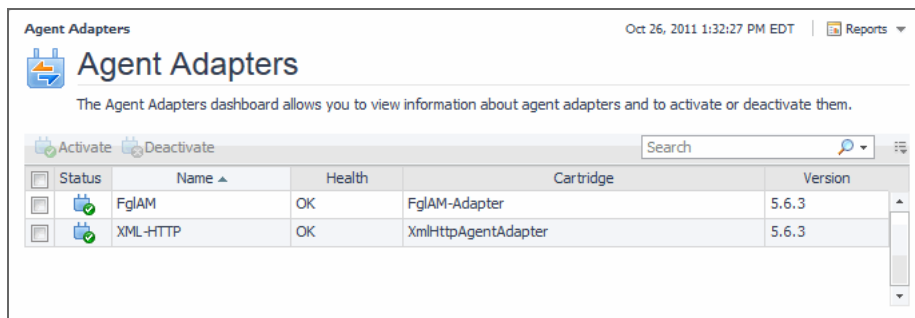
- The WindowsAgent stops collecting the data, if this WindowsAgent is moved from FglAM-A (an FglAM host installed with the Windows environment) to FglAM-B (an FglAM host installed with the Linux environment). See the “Access to DCOM objects and registry is denied” section in the *Foglight Agent Manager Guide* to resolve this issue.
- When moving an FglAM agent which ASP configurations include a value of “localhost”, it is strongly recommended that you change the corresponding ASP value to the name or IP address of the local host; otherwise the unexpected result may occur.
- The agent movement fails if an FglAM agent is moved from Solaris OS to Windows 2008 OS. This issue is known to exist. For more information, see the *Foglight 5.7.6 Release Notes* document.
- After moving an FglAM agent to an FglAM host that has the invalid lockbox, this FglAM agent cannot be moved to other Foglight hosts any more. The invalid lockbox of the target FglAM host causes that the target FglAM host cannot release any FglAM agent.
- Before moving an FglAM agent, you need be aware that all lockboxes should be released from the source FglAM host to the target FglAM host, in order to achieve a successful FglAM agent movement.

Viewing Agent Adapters

Foglight uses agent adapters to communicate with agents that collect information from monitored hosts. The majority of Foglight agents use the Foglight Agent Manager (FglAM). There are some agents that use other types of agent adapters.

The Agent Adapters dashboard allows you to view information about agent adapters and to activate or deactivate them, if instructed by Quest Support. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Agents > Adapters**.

Figure 5. The Agent Adapters dashboard.



Controlling System Access with Credentials

Foglight can monitor different resources in an organization. Some of these resources can be secured while others are not. Your organization has specific security policies for those resources that require authorized access. For example, Windows hosts are secured resources that require a Windows login. Unix hosts are also secured resources that require a valid Unix login. Foglight allows you to store system credentials for accessing secured resources in a central location.

A credential is a piece of information that an agent instance needs to gain access to the monitored system. For example, you can associate one client with a database server and another with a production server, and have an agent instance monitoring the database server connect using a user name and password. Another agent instance monitoring the production server can connect to it using Windows-based login information.

Different cartridges support different types of credentials. Some cartridges, for example, support the use of Windows and Unix credentials, while others can only authenticate local users. The credential type determines which parts of the monitored system are used to connect to a resource, such as host names or IP addresses. For complete information about cartridge-specific credential types, see your cartridge documentation.

Credentials are encrypted and stored in lockboxes. Lockboxes are released to credential clients, such as agent managers.

The Credentials dashboard provides quick access to credentials and lockboxes. This dashboard provides at-a-glance information about the current state of credentials, lockboxes, credential clients, the alarms they generate, and cartridge-specific credential views. Use it as a starting point for your credential management needs. To access this dashboard, on the navigation panel, click **Dashboards > Administration > Credentials**.

i | **TIP:** Another way to access the Credentials dashboard is from the Administration dashboard, under Credentials > Credential Management.

i | **NOTE:** In a federated environment, credentials can only be managed on Federated Child servers. Credentials cannot be administered from the Federation Master.

Managing credentials

A credential is a piece of information required to gain access to system resources. Foglight agents need access to this information when monitoring systems that require credential verification.

Foglight supports a set of commonly used credentials such as currently logged in user, password-based, user name with or without password, and Windows credentials. Each credential can have one or more authentication policies, based on the desired usage count, failure rate, the time range during which the credential can be used, and the amount of time during which the credential information is cached locally. Credentials can apply to specific parts of the monitored environment, such as hosts and ports.

You create and manage credentials, as well as edit their type, authentication policies, and target resources, using the Manage Credentials dashboard. To access this dashboard, on the main Credentials dashboard, click **Manage Credentials**.

Monitoring credential alarms

Occasionally, credential clients may encounter errors. For example, a credential client can fail to start a monitoring agent due to a credential failure. The Monitor Credential Alarms dashboard lists all alarms that are raised by credential clients and provides additional information about each alarm, such as the severity, alarm message, event or rule that generated it, and other information. To access this dashboard, on the main Credentials dashboard, click **Monitor Credential Alarms**.

Managing lockboxes

A lockbox can be password-protected, and contains a collection of credential keys used for encryption and decryption.

You can create, edit, and manage lockboxes, change their passwords, and release them to credential clients using the Manage Lockboxes dashboard. To access this dashboard, on the main Credentials dashboard, click **Manage Lockboxes**.

Viewing credential clients

The View Clients dashboard lists all credential clients that exist in Foglight, and provides additional information about each client, such as its name, type, and assigned lockboxes, along with other information. To access this dashboard, on the main Credentials dashboard, click **View Clients**.

Using domain-specific credential views

In addition to the credential dashboards included with the Management Server, some cartridges may include their own credential views. If your system includes any domain-specific credential views, the links to these views are listed at the bottom of the Credentials dashboard.

Administering Foglight

This chapter focuses on recommended maintenance tasks that ensure optimal Foglight® performance. It also describes the starting points in Foglight administration.

i **IMPORTANT:** For more information about specific tasks, or additional technical information about the features described in this chapter, see the reference topics accessible from the applicable section in the Administration and Configuration Help. For example, to find out more about the Administration home page, in the browser interface, open the Help tab in the action panel, and from there, navigate to **Administration and Configuration Help > Administering Foglight > Using a single pane of glass for your administration needs**. You will find a list of reference topics at the bottom of the page.

The following standard maintenance tasks should be performed to ensure a stable Foglight system. This simple guidance can help Foglight administrators to optimally and consistently perform their administrative tasks. Successful ongoing management of an enterprise-class Foglight implementation requires formal Foglight administrator training from Quest professional services.

As a starting point, this section assumes a stable installation which, at minimum, contains the following configuration components:

- A physical or virtual server, appropriately sized to ensure that it has appropriate resources required to support the monitoring requirements.
- An installed Foglight Management Server, properly tuned to meet the monitoring requirements.
- The installation of Foglight Agent Manager components and Foglight monitoring agents in the monitored landscape.
- Tuned collection rates and other agent properties. This is required for Java environments.
- A configured SMTP server, to ensure that the system is capable of sending email notifications.
- A defined back-up strategy that can be regularly executed.
- Tuning of the rules to adjust alarm thresholds.

Finally, this section addresses the ongoing maintenance tasks required to ensure your Foglight Management Server stays operationally healthy.

Generally speaking, most of the tasks required to ensure a stable installation can and should be automated. For example, there is not much value in asking a Foglight administrator to log in only to ensure it is running or check to see if there is enough memory when rules and automated emails can be generated to notify administrators of potentially worrisome conditions. As such, the first part of this section identifies the automatable self-monitoring options that should be configured. The rest of the section covers:

- The manual processes to check health state of the automatable functions in the event that you do not choose to configure the notifications.
- The manual processes that are required and cannot be entirely automated.

For more information, see the following topics:

- [Recommended Self-Monitoring Automation](#)
- [Using a single pane of glass for your administration needs](#)
- [Manage Foglight Database Performance](#)
- [Monitor Server Performance](#)
- [Monitor Agent Manager Performance](#)
- [Associate Service Objects with Groups and Tiers](#)
- [Back Up and Restore Foglight](#)

Recommended Self-Monitoring Automation

Foglight Management Server up/down status

A separate Remote Monitor process can be configured to watch the Foglight® High Availability (HA) process and notify an administrator in the event that the Foglight Management Server process unexpectedly shuts down. Configure the Remote Monitor process to ensure that an administrator is notified when the Foglight Management Server shuts down. For more information about the Remote Monitor process and running Foglight in HA mode, see the *Foglight High Availability Field Guide*.

Core-Monitoring Policy rules

A set of rules covering the critical health items for a Foglight Management Server is delivered with the Core-Monitoring Policy cartridge, included with the server install. This cartridge is installed and enabled during the server installation. Email notifications should be set-up for each of the rules delivered in this cartridge. To configure email notifications, use the Email Configuration dashboard. To access this dashboard, from the **Administration** dashboard, under **Support**, click **Email**.

Cyclic Maintenance

Each of the following tasks helps to ensure that the Foglight® Management Server is stable and is operating normally.

Daily Maintenance Tasks

Assuming you automate self-monitoring as described in [Recommended Self-Monitoring Automation](#), there is no need to perform the checks described below.

If you do not have automated self-monitoring in place, you must use the manual technique below, performing each item at least once daily.

Foglight Management Server

- Validate that the Management Server is running by logging into the Management Server on a daily basis.
- Validate that the Management Server is operating within basic resource and operational guidelines by checking the Alarms dashboard for any alarms raised by the following rules:
 - Catalyst Data Service Discarding Data
 - Catalyst Database Space Checking
 - Catalyst Free Database Space Checking
 - Foglight Garbage Collector Check
 - Foglight Memory Usage Check
 - Foglight Topology Size Limit Reached

For context on the key resource requirements and their effect on the Foglight Management Server, see the *Foglight Performance Tuning Field Guide*.

Foglight Agents and Agent Manager

- Validate that all Agent Manager Instances are running by checking the Alarms dashboard for any alarms raised by the following rules:
 - Remote Agent Manager State
 - Remote Agent Managers State per Host
- Validate that all agents are running and collecting data by checking the Alarms dashboard for any alarms raised by the following rules:
 - Agent Health State
 - Idle Agent

Weekly Maintenance Tasks

Perform the following tasks once every week:

- Back-up the Foglight server and repository. This can be done more or less frequently, depending on the specific backup and recovery strategy for your Foglight installation.
- Review Foglight resource utilization trends to ensure there is not an increasing resource consumption trend.
 - Generate and review graphs from the seven-day performance report.
Automated approach: Schedule this to generate weekly and email.
Manual approach: Manually create a report using the Report Manager and review it.

Monthly Maintenance Tasks

Perform the following tasks manually, once every month:

- Identify and adjust thresholds for rules that may be firing too frequently.
- Evaluate the browser interface performance trends to ensure there is no negative trend in the performance. While there is inherent variability in all Foglight installations, we expect a well-tuned Foglight installation to show less than five second average response times on the following benchmark dashboards:
 - Active Hosts Summary
 - Administration
 - Agents
 - Agent Status
 - Alarms
 - All Hosts (All Hosts)
 - Edit rule view
 - Manage Reports
 - Monitored Hosts Only (All Hosts)
 - Service Operations Console
 - Rule Management

For more information about these dashboards, see the related online help topics.

Using a single pane of glass for your administration needs

The Administration dashboard can be used as a front-end for most Foglight® administration tasks. This dashboard contains links to most of the administration dashboards and shows configuration specifics that may be critical to your day-to-day operation. If your Foglight role involves daily administration, this is probably the best place to start as it gives you a quick insight into the system complexity and its health, along with close-at-hand links to most administration dashboards.

You can access this dashboard from the navigation panel, by clicking **Homes > Administration**.

In addition to the administration dashboards accessible from the Administration home page and the Administration node on the navigation panel, Foglight includes another set of administration-level dashboards that can be used to analyze logs, monitor the server performance, and service levels. These dashboards are described in this chapter.

Manage Foglight Database Performance

In Foglight, retention policies define time periods during which the collected data is sampled, persisted into the database, aggregated, or purged from your system. The Object Cleanup dashboard is useful for inspecting, purging, and deleting data objects, and particularly for cleaning up objects that are no longer needed. For instance, if a set of agents is no longer required, the objects created by those agents are still visible in the browser interface. Removing these objects can have a positive impact on performance.

- i** | **NOTE:** In previous releases, the Data Management dashboard allowed you to manage both retention policies and object cleanup. Beginning with Foglight 7.1.0, there are now two separate dashboards:
- **Retention Policies**
 - **Object Cleanup**

Using these dashboards you can manually tune the size and performance of your environment. Tuning performance manually allows you to:

- Control the persistence policies of the Foglight Management Server by setting retention policies and purging.
- Set a global default, which auto-deletes objects after a specified period of time.
- Manually inspect and adjust the amount of data that has been saved by showing all topology objects in the server and the metrics attached to them. You can then delete unneeded objects and metrics.
- Purge objects that have been captured by the Foglight Management Server.

To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Servers**. Then select either **Object Cleanup** or **Retention Policies**.

Another way to control the retention policies on a per-object basis is using the Manage Retention Policies dashboard. For more information, see [Manage Data Retention](#) on page 38.

In addition to the Object Cleanup and Retention Policies dashboards that allow you to inspect and tune the overall performance, the Database Overview dashboard summarizes the database activities such as data row operations, database buffer pool, and any inserts, deletes, and updates. To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Servers > Database Overview**.

Monitor Server Performance

Foglight manages host data sent from agents, and evaluates rule conditions and metric calculations. It also provides browser interface access to remotely monitored servers. The browser interface includes a set of dashboards that allow you to monitor the server state and prevent potential bottlenecks.

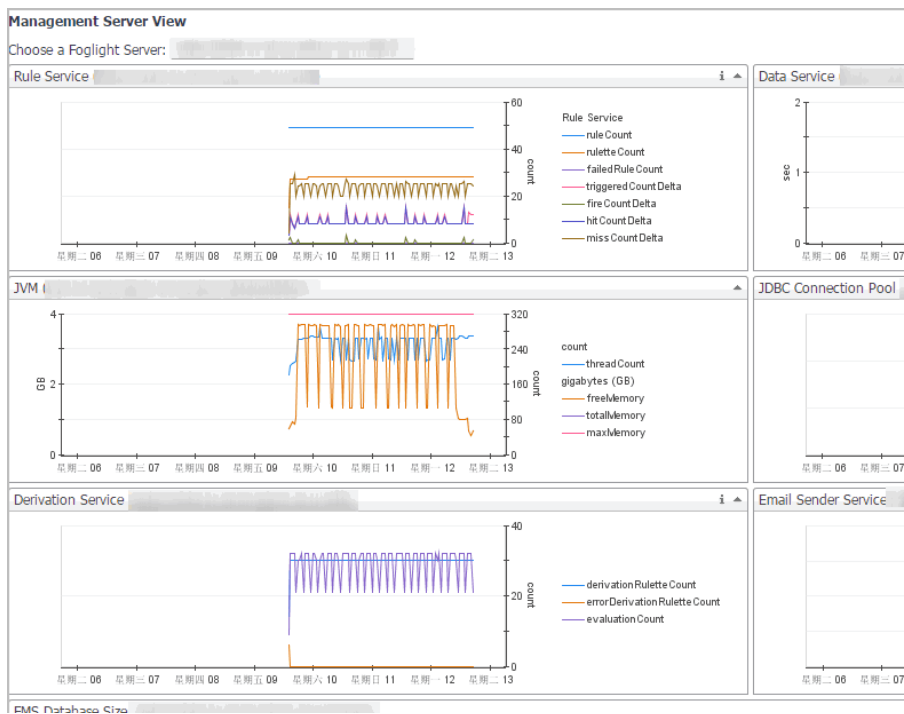
The Performance dashboard contains at-a-glance view of Foglight diagnostics. It shows the rate of database inserts, data processing activity, JVM memory usage, server load, and other combinations of views. Certain types of metric patterns displayed on this dashboard can be useful in troubleshooting specific performance problems. For example, a sudden increase in free memory utilization is a good indicator that the amount of incoming data exceeds typical thresholds. To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Diagnostic > Performance**.

Figure 6. The Performance dashboard.



The Management Server View dashboard is useful for examining server performance. Use it to look for root causes of server-related performance problems. To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Servers > Management Server View**.

Figure 7. The Management Server View.




Server log files are another source of information that can help you diagnose the root cause of performance-related bottlenecks. They contain information about known events and error conditions as well as verbose or informational messages. The Log Analyzer dashboard allows you to analyze generated log files or download a selected log file to a desired location. To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Diagnostic > Log Analyzer**.

Monitor Agent Manager Performance

Foglight uses the Agent Manager to communicate with monitored hosts. The embedded Agent Manager can be used to monitor the host on which the Management Server is installed. Your monitoring environment typically includes a number of Agent Manager instances that are installed on different hosts. You can monitor their state using the Performance Overview dashboard. For example, an unusually high number of pending messages in the queue indicates a potential performance bottleneck. To access this dashboard, from the navigation panel, click **Dashboards > Management Server > Diagnostic > Agent Manager**.

Figure 8. The Foglight Agent Manager Performance Overview dashboard.



Associate Service Objects with Groups and Tiers

Foglight monitors specific parts of your environment based on the concept of services. A service is a collection of monitored objects. An object group is a mechanism that assists in service creation and monitoring. It is a logical way of grouping objects that are of interest to an individual user (for example, an Oracle database administrator), or to multiple users of a system (for example, Oracle databases).

Object groups can be associated with another logical component, tiers. Each tier is a logical representation of a service component. A Foglight service can have one or more tiers. By default, Foglight organizes data into default tiers, including User, Web, Application, Database, Host, and Agent. Tiers allow you to structure services in a way that best represents your monitored environment. This type of logical structure helps you isolate performance problems associated with a specific service tier. For example, to investigate the state of the monitored hosts within the Host tier for a service, drill down on the Host tier and investigate the hosts that are related to it. For more information about services, see the *Foglight User Guide*.

The object groups that are needed for most service monitoring already ship with Foglight. You can create additional object groups and associate groups with tiers using the Object Groups and Tier Definitions dashboards.

To access the Object Groups dashboard, from the navigation panel, click **Dashboards > Services > Object Groups**.

To access the Tier Definition dashboard, from the navigation panel, click **Dashboards > Services > Tier Definition**.

Back Up and Restore Foglight

Backup and restore processes are important aspects of database administration. The term *backing up* refers to making copies of data that can be used to restore your system after a data loss event. The backup process includes:

- Archiving the Foglight configuration file, scripts, and installed cartridges
- Backing up the entire database (Microsoft SQL Server, MySQL, PostgreSQL or Oracle)
- Verifying the settings of environment variables (Oracle)
- Saving the archive in a safe location, outside of the original installation directory

Restoring a physical backup means reconstructing it and making it available to users. You can restore a previous Foglight installation from a backed up copy of the original environment.

For more information about backing up and restoring Foglight, see the related help topics accessible from this section in the online help.

Configure Rules and Metric Calculations to Discover Bottlenecks

The Foglight® Management Server and individual cartridges each come with a set of topology types, data retention policies, rules, and calculations, that, in most cases, work out-of-the-box. In more complex environments, your business case may require additional tuning.

i **IMPORTANT:** For more information about specific tasks, or additional technical information about the features described in this chapter, see the reference topics accessible from the applicable section in the Administration and Configuration Help. For example, to find out more about the Foglight registry, in the browser interface, open the Help tab in the action panel, and from there, navigate to **Administration and Configuration Help > Configuring Rules and Metric Calculations to Discover Bottlenecks > Working with Foglight Registry Variables**. You will find a list of reference topics at the bottom of the page.

Foglight uses flexible rules to apply your business logic to complex, interrelated data from multiple sources within your distributed system. A rule is a piece of business logic that links a condition with a result, for example if HTTP response time exceeds 500 ms, fire a warning alarm. A rule can include a scope and one or more conditional expressions, alarm messages, and actions. The scope defines the set of topology objects against which it runs. The conditional expression defines the thresholds, that, if reached, cause the rule to fire. Conditional expressions can include registry variables, raw metrics, derived metrics, and topology object properties.

There are two types of rules in Foglight: *simple rules* and *multiple-severity rules*. Simple rules do not generate alarms, they fire and invoke actions when their conditions are met. Multiple-severity rules include up to five severity levels and generate alarms when the condition associated with any one of its severity levels is met.

Each severity level can have its own set of variables that you can use in alarm messages. Unlike registry variables, which are global in nature, severity-level variables are only accessible to the severity level in which they exist. For example, a Warning-level variable that contains alarm text can only be referenced by the alarm message defined for the Warning severity. Critical- or Fatal-level alarm messages, associated with the same rule, do not have access to this variable.

Severity levels can be associated with actions, causing them to occur each time a threshold is reached. Foglight comes with different types of actions, such as email, command, script, and other types of actions.

Rules have four types of triggers: *data-*, *time-*, *schedule-*, and *event-driven triggers*.

- If a rule has a data-driven trigger, one or more of its conditions are evaluated every time the data associated with the rule is collected. This is the default trigger.
- A time-driven trigger causes one or more of a rule's conditions to be evaluated once per pre-defined interval. By default, Foglight evaluates time-driven rules only if the evaluation data is available.
- Event-driven triggers cause the rule conditions to be evaluated as a response to one of the following events: *AgentManagementSystemEvent*, *AlarmSystemEvent*, or *ReportGeneratedEvent*.
- Schedule-driven triggers cause the rule conditions to be evaluated during a selected schedule.

Many rules come included with Foglight and installed cartridges, such as *Agent Health State*, *BSM All Events*, *Catalyst Data Service Discarding Data*, and others. If the existing rules do not meet your needs, you can create a new one and add it to the rule collection.

A typical installation can include a large number of rules. You can create and manage multiple-severity rules using the Rule Management dashboard. To access this dashboard, on the navigation panel, click **Dashboards > Administration > Rules & Notifications > Rules**. The Rules dashboard provides access to modifying threshold variables appearing in rule conditions. For an extended set of rule management tasks, such as viewing and editing

rule definitions, copying, disabling, and enabling rules, suspending or resuming rule actions, use the Manage Rules dashboard. Click **Old Manage Rules** to access the Manage Rules dashboard.

Figure 9. The Rule Management dashboard.

Rule	90.0	80.0	70.0	Other	Alarms
Additional Process Pool Utilization	90.0	80.0	70.0		0
Agent Health State					1
Agent Manager Error Message Alarm					0
Available Paging Space	5	10	20		0
Blocked Transaction Alarm Generator					0

To obtain additional diagnostics about rule behavior and how they affect your monitoring environment, use the Rule Diagnostics dashboard. From here, drill down to a specific rule and explore the objects are affected by the rule, or find out how many times a rule was executed against a specific object. This can help you understand rule behavior and debug any problems associated with a particular rule. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Rules & Notifications > Rules**, then click **Diagnostics** from the rule’s context menu.

Related topics:

- [Working with Derived Metrics](#)
- [Working with Metric Thresholds](#)
- [Analyze Activity Levels with IntelliProfile](#)
- [Working with Foglight Registry Variables](#)
- [Associate Metric Calculations with Schedules](#)
- [Manage Data Retention](#)
- [Expand Your Collection of Topology Types](#)

Working with Derived Metrics

In addition to building topology models at run-time using the data collected from monitored systems, Foglight has a unique capability to apply pre-defined calculations to the collected metrics. Metric calculations are typically scoped to specific parts of the topology and their values can change over time. They can be reused in expressions to simplify their syntax and speed up system deployment.

A metric is a specific value that is measured over time. There are two types of metrics in Foglight: *raw metrics* and *derived metrics*. Raw metrics are collected directly from your monitored environment and sent to the Foglight Management Server. Derived metrics are calculated from one or more raw or derived metrics. They are scoped to a topology type and can optionally be scoped to specific objects of that type. Many derived metrics come included with Foglight and installed cartridges. If none of the existing derived metrics meet your needs, you can create a new one and add it to the derived metric collection.

There are many reasons why it can be useful to create derived metrics. For example, creating derived metrics can make managing rules simpler by reusing metric expressions.

You can create and manage derived metrics using the Manage Derived Metrics dashboard. A typical installation can include a large number of rules. The Manage Derived Metrics dashboard lists all derived metrics that exist in your environment, and allows you to drill down to derived metric definitions. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Data > Derived Metrics**.

Figure 10. The Manage Derived Metrics dashboard.

The screenshot shows the 'Derived Metrics' dashboard. At the top, it says 'Derived Metrics' and 'Nov 11, 2015 10:41:45 AM EST'. Below that is a search bar and a 'Reports' dropdown. The main content is a table with columns: Name, Scope, and Cartridge. The table lists various metrics like activeAgentCount, availability, averageAvailability, etc.

	Name	Scope	Cartridge
<input type="checkbox"/>	activeAgentCount	CatalystDiagnosticAgentType	Diagnostic
<input type="checkbox"/>	availability	FSMServiceLevelPolicy : sourceIds = \$null	Core-ServiceModel-Extensions
<input type="checkbox"/>	averageAvailability	NetMonitorDevice	NetMonitor-UI
<input type="checkbox"/>	averageExpectedResponseTime	NetMonitorDevice	NetMonitor-UI
<input type="checkbox"/>	averageHopCount	NetMonitorDevice	NetMonitor-UI
<input type="checkbox"/>	averagePacketLoss	NetMonitorDevice	NetMonitor-UI
<input type="checkbox"/>	averageResponseTime	NetMonitorDevice	NetMonitor-UI
<input type="checkbox"/>	avg_inserts_per5min	CatalystDatabaseStatus	Diagnostic
<input type="checkbox"/>	baselineAvailability	ServiceLevelPolicy : sourceIds = \$null	Core-ServiceLevelPolicy
<input type="checkbox"/>	collectionCountDelta	(CatalystServer).jvm.garbageCollectors	Diagnostic
<input type="checkbox"/>	collectionTimeDelta	(CatalystServer).jvm.garbageCollectors	Diagnostic
<input type="checkbox"/>	databaseInsertRate	CatalystPersistenceService	Core-MonitoringPolicy

To obtain additional diagnostics about derived metric behavior and how they affect your monitoring environment, use the Derived Metrics Diagnostics dashboard. From here, drill down to a specific derived metric and explore the objects that are affected by the derived metric, or find out how many times a derived metric has been calculated against a specific object. This can help you understand derived metric behavior and debug any problems associated with a particular derived metric. To access this dashboard, in the Administration dashboard, click **Metrics Diagnostics** in the **Data** column.

Working with Metric Thresholds

Threshold levels in metrics are useful in situations when you need to reference a specific metric value multiple times, for example in derived metrics or rules. Each metric can have one threshold associated with it. A threshold is always associated with a threshold level. Threshold levels refer to a particular state of monitoring entities, such as agent states, alarm severities, and others. Each threshold level includes a unique set of threshold bound levels that are specific to that level. For example, the threshold level `AgentState` comes with several bound levels that relate to agent states, such as `Running` and `Collecting Data`.

Creating a threshold involves selecting a metric and defining values for threshold bounds. A bound level value can be a constant value, a registry variable, or another metric of the same topology type. As data is sampled, Foglight evaluates the metrics for which thresholds are defined, matching their run-time values with bound-specific values in pre-defined order, and performs actions when specific bound levels are reached.

You can create and manage thresholds using the Manage Thresholds dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Data > Manage Thresholds**.

Figure 11. The Manage Thresholds dashboard.

The screenshot shows the 'Manage Thresholds' dashboard. At the top, it says 'Manage Thresholds' and has 'Add' and 'Delete' buttons. Below that is a table with columns: Metric, Topology Type, and Summary. The table lists thresholds for metrics like active_time, adh_service, agentMemoryConsumption, and agentRestarts.

	Metric	Topology Type	Summary
<input type="checkbox"/>	active_time	DBSS_Wait_Event_Category	Normal: 0.0 (c), Normal: 0.01 (c), Warning: IntelliProfi
<input type="checkbox"/>	adh_service	DBSS_SQL_Server_Services	Normal: 0.0 (c), Warning: DBSS-ServicesADHDown_Lo
<input type="checkbox"/>	agentMemoryConsumption...	FxMSystemHealth	Normal: 0.0 (c), Fatal: fxm.agentMemoryConsumption
<input type="checkbox"/>	agentRestarts	FxMSystemHealth	Normal: 0.0 (c), Fatal: fxm.agentRestarts.fatal (v)

Analyze Activity Levels with IntelliProfile

Foglight uses the IntelliProfile technology to estimate the system performance and help with system monitoring and planning. IntelliProfile uses data collected during a desired time period and generates a baseline operating range based on the collected metrics.

i | **NOTE:** You can only use this feature after installing a baseline cartridge that contains definitions for a selected metric property. Agent developers can write baseline cartridges. For more information about metric thresholds, see [Working with Metric Thresholds](#).

A baseline establishes expected data patterns for a given time period. Baselines are periodically evaluated during IntelliProfile learning cycles, to reflect changes in data patterns.

During operating cycles, IntelliProfile compares incoming data for those metrics that have IntelliProfile threshold levels configured. Metric threshold states reflect the degree of deviation from the baseline, and can indicate potential performance bottlenecks. If there are any rule conditions that evaluate threshold states for such metrics, Foglight can generate alarms when a metric enters a certain threshold state.

IntelliProfile settings can be viewed and managed using the IntelliProfile dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Data > IntelliProfile**.

Figure 12. The IntelliProfile dashboard.

The screenshot shows the IntelliProfile dashboard interface. At the top, it displays the title 'IntelliProfile' and the date 'Dec 12, 2011 9:27:26 AM EST'. Below the title, there are two main sections: 'General' and 'Thresholds'.

General Section:

- IntelliProfile is ready after first hours of work.
- Allow IntelliProfile to automatically calculate the optimal based on the instance activity of the past days.

Thresholds Section:

IntelliProfile qualifies activity levels relative to a normative baseline. Deviation messages are presented as information, warning or critical according to pre-configured activity range thresholds.

Use the table below to modify the threshold range configurations.

Severity	Above (%)	Below (%)
Information	5	95
Warning	3	97
Critical	1	99

Buttons: 'Restore Default Thresholds' and 'Save Changes'.

Working with Foglight Registry Variables

The Foglight registry is a collection of variables. Their values can be used for evaluation or comparison in monitoring entities and thus often determine evaluation outcome and triggered actions. The Foglight registry is not related to the host's OS registry (for example, the Windows registry).

Rule conditions, for example, and their expressions can reference registry variables. A registry variable can have a global value that is available to all topology types and objects. It can also have one or more values associated with specific topology types or objects, or calendar dates. For example, your organization can have different administrators looking after different hosts. To configure Foglight to send host-related emails to appropriate recipients, scope the `SYSADMIN` variable to the monitored host instances and associate an email address with each host.

Many registry variables come included with Foglight and installed cartridges, including *AvailabilityCritical*, *AvailabilityFatal*, *AvailabilityTarget*, and many others. If the existing registry variables do not meet your needs, you can create a new one and add it to the registry variable collection.

You can create and manage registry variables using the Manage Registry Variables dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Rules & Notifications > Manage Registry Variables**.

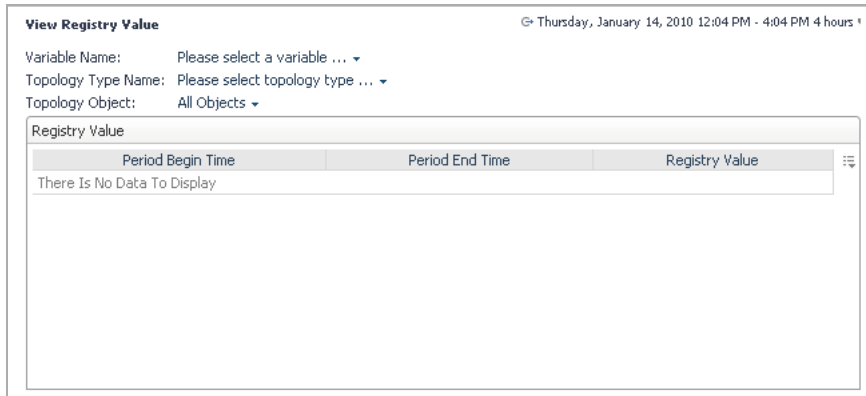
Figure 13. The Manage Registry Variables dashboard.

The screenshot shows the 'Manage Registry Variables' dashboard. At the top, it says 'Manage Registry Variables' and 'Thursday, February 3, 2011 6:38 AM - 10:38 AM 4 hours'. Below the title is a description: 'The Foglight registry is a collection of variables. Their values can be used for evaluation or comparison in monitoring entities to determine evaluation outcome and trigger selected actions. A registry variable can have a global value that is available to all topology types and objects. It can also have one or more values associated with specific topology types or objects, or calendar dates. This dashboard lists all registry variables that exist in your monitoring environment.' Below the description is a table with columns: Variable Name, Type, Cartridge Name, Global Default, Scoped Default (TopologyObject), Schedule, and Active Schedule. The table contains several rows of variables.

Variable Name	Type	Cartridge Name	Global Default	Scoped Default (TopologyObject)	Schedule	Active Schedule
AvailabilityCritical	Doubl	Core-ServiceLeve	85.0	85.0	No	
AvailabilityFatal	Doubl	Core-ServiceLeve	70.0	70.0	No	
AvailabilityTarget	Doubl	Core-MonitoringPi	95.0	95.0	No	
AvailabilityWarning	Doubl	Core-ServiceLeve	95.0	95.0	No	
BaselineFactor	Doubl	Core-MonitoringPi	2.0	2.0	No	
BSM URL	String	BSM-Integration			No	
CATALYST_URL	String	System	http://host1.exempl	http://host1.exempl	No	
DBADMIN	String	Core-MonitoringPi	unset	unset	No	

To find out the value of a registry variable for a particular topology type or object during a specific time period, use the Check Registry Value dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Rules & Notifications > Check Registry Value**.

Figure 14. The Check Registry Value dashboard.



Associate Metric Calculations with Schedules

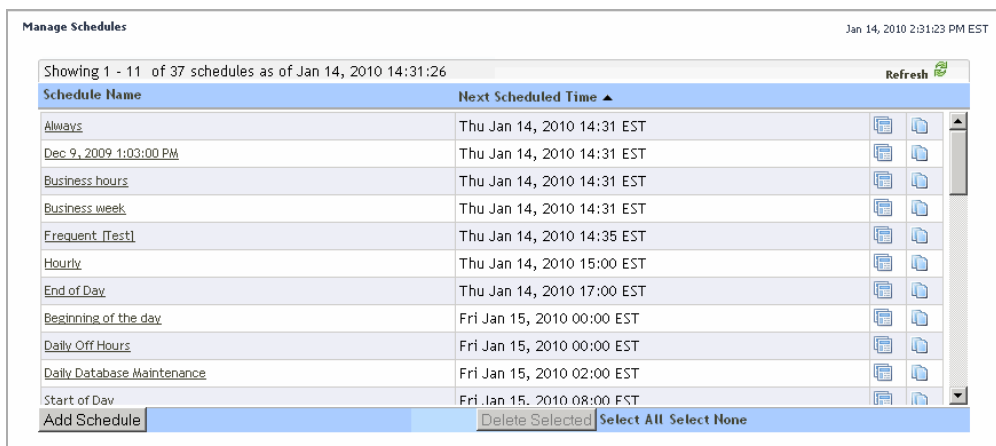
A schedule consists of one or more schedule items. Each schedule item includes a start date, an end date or a time range during which it runs, a recurrence pattern, and the range of recurrence. A default Foglight installation includes a number of schedules, including *Always*, *Business hours*, *Business week*, and many others.

Rules, registry variables, derived metrics, and other Foglight components use schedules to initiate calendar-driven actions. For example, a registry variable can have multiple values, each associated with a specific schedule. If none of the existing schedules meet your needs, add a new schedule to the existing collection and associate it with the registry value.

You can create and manage schedules using the Manage Schedules dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Schedules > Manage Schedules**.

i | **NOTE:** If you accidentally delete the Daily Database Maintenance schedule, the server will recreate it within 24 hours.

Figure 15. The Manage Schedules dashboard.



Manage Data Retention

Retention policies allow you to define how monitoring data is aggregated and for long it is kept before being purged from Foglight. All topology objects in Foglight form a hierarchy whose root is the super-type `TopologyObject`. Retention policies are inherited from the object's type. These policies may be overwritten, in which case the modification applies to all child types in the hierarchy.

In addition to retention policies, the collected data has additional life-cycle properties that are defined in *storage-config.xml*. The life cycle involves several iterations of data collection, aggregation, and storage in *database generations*. Database generations are database structures that store aggregated data for a specific period of time.

For example, the default retention policy associated with `TopologyObject` causes the collected data to be rolled up to 15-minute periods after the age of 15 minutes, and stored in Generation 1 for three days. From there, four-hour interval data is rolled up to one-hour periods, and then stored in Generation 2. After 14 days, 5-day interval data from Generation 2 is rolled up to four-hour periods and stored in Generation 3 indefinitely, or until it is purged.

If there is no existing retention policy for a topology type, that type inherits the retention policy from its parent type. If no policies exist within the entire hierarchy, the type inherits the policy from the `TopologyObject` type. Conversely, setting a retention policy for a topology type completely overrides any policy it inherits from a super-type, and is applied to all sub-types of that topology type.

You create and manage data retention policies using the Retention Policies dashboard. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Data > Manage Retention Policies**.

Figure 16. The Manage Retention Policies dashboard.

	Topology Type - Property Name ▲	Age	Roll-up Period
<input type="checkbox"/>	Agent		
<input type="checkbox"/>	AgentHealthState		
<input type="checkbox"/>	AgentManagementSystemEventType		
<input type="checkbox"/>	AgentState		
<input type="checkbox"/>	AgentTypeLicense		
<input type="checkbox"/>	AggregateModelInstance		
<input type="checkbox"/>	AggregateModelRoot		
<input type="checkbox"/>	AlarmChangeType		
<input type="checkbox"/>	AlarmRuleBasedView		
<input type="checkbox"/>	AlarmSeverity		

Another way to control the retention policies is through the Retention Policies dashboard. This dashboard allows you to control system-wide retention policies and to delete unwanted data from the database as a performance-tuning measure. For more information, see [Manage Foglight Database Performance](#) on page 28.

Expand Your Collection of Topology Types

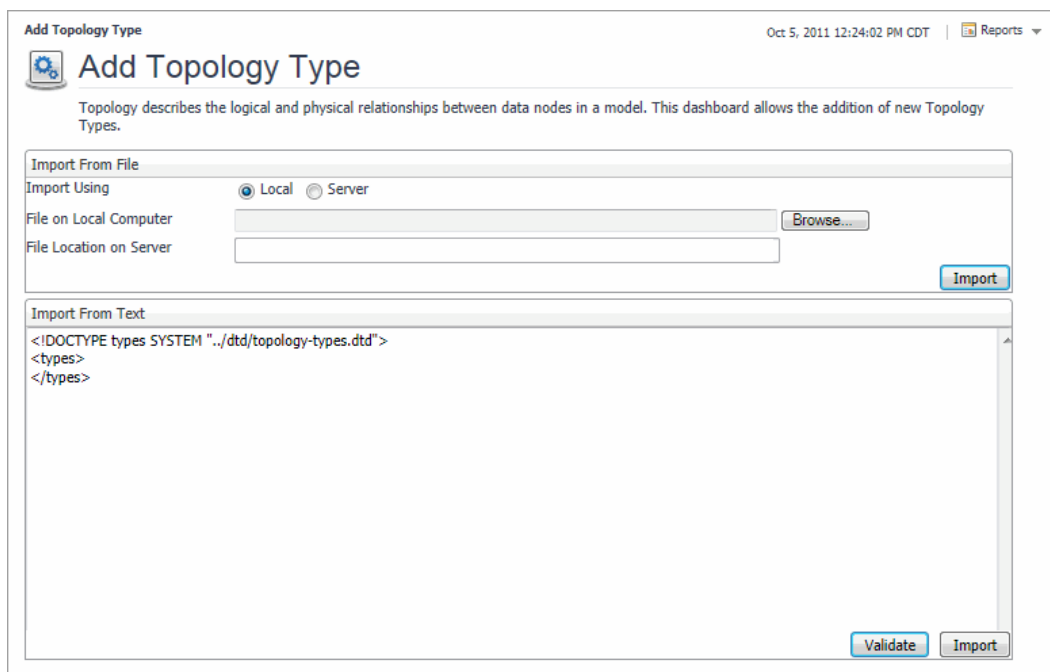
The set of topology types that exist in your environment depends on your monitoring needs, reflected in the type and nature of cartridges that you use for data collection. If you need additional topology types, you can add them to Foglight.

The following example shows the syntax for defining a topology type:

```
<type name="ApacheSvr_Transactions" extends="F4Table">
<property name="IntervalTransactions" type="Metric" is-containment="true" />
<property name="TransactionRate" type="Metric" is-containment="true" />
<property name="TransactionTag" type="String" is-identity="true" />
<property name="TransactionThroughput" type="Metric" is-containment="true" />
<property name="TransactionThroughputRate" type="Metric"
  is-containment="true" />
</type>
```

The Add Topology Types dashboard allows you to add new topology types to your topology model and to validate them. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Data > Add Topology Type**.

Figure 17. The Add Topology Types dashboard.



You can explore your database schema using the Schema Browser. You can access the Schema Browser using the Dashboard Development page. For additional details, see the *Dashboard Support Guide*.

Customizing Your Foglight Environment with Tooling

Foglight® includes a set of advanced administration features that allow you to address your monitoring needs beyond typical day-to-day use. These features are described in this chapter.

For more information, see the following topics:

- [Merging Host Objects](#)
- [Building Script Agents](#)
- [Using the Query Language](#)
- [Retrieving Data with the REST API](#)
- [Retrieving Data with Scripts and Queries](#)

i **IMPORTANT:** For more information about specific tasks, or additional technical information about the features described in this chapter, see the reference topics accessible from the applicable section in the *Administration and Configuration Help*. For example, to find out more about script agents, in the browser interface, open the **Help** tab in the action panel, and from there, navigate to **Administration and Configuration Help > Customizing Your Foglight Environment with Tooling > Building Script Agents**. You will find a list of reference topics at the bottom of the page.

Merging Host Objects

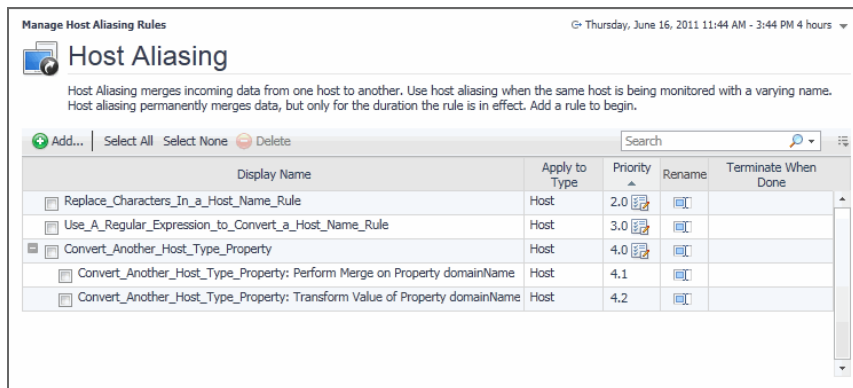
Merging two or more hosts refers to the ability to consolidate data for those host objects using host aliasing rules. A host aliasing rule includes one or more property matching filters that select the topology objects that are to be merged, along with the logical definition of the merge operation. Property matching filters can only reference a subset of the entire property set for a topology type such as String or Boolean properties.

Simple merging rules contain one stand-alone rule. They are used to merge one or more host objects, or to rename a host object in the model. *Advanced merging rules* consist of a group of individual rules that are executed in pre-defined order. They can merge one or more topology objects. For example, merging two agent instances involves a rule for transforming the instance name and another one for merging the two instances.

Merging rules are useful in situations when a host name changes and there is a need to consolidate the data under a single host object. Consider for example an Agent Manager installed on a host whose name is `Toronto123`. The host reports into Foglight as `Toronto123`, which creates a new `Host` object, `Toronto123`. The system administrator modifies the host's configuration which causes `Toronto123` to start reporting itself using its IP address, `10.1.234.56`. When the Agent Manager collects information from the newly-renamed host, a new `Host` object is created on the server, with the name `10.1.234.56`. After noticing the problem, the Foglight administrator solves it by creating an alias for the host, mapping it to its original name, `Toronto123`.

Use the Manage Host Aliasing Rules dashboard to create host aliasing rules or to manage the existing ones. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Tooling > Manage Host Aliasing Rules**.

Figure 18. The Manage Host Aliasing dashboard.



Building Script Agents

A script agent is a special type of agent that is used to execute a script on a monitored host. The script can be in any language, but it must be written in a way that provides specifically formatted output to `STDOUT`. Creating script agents is the easiest mechanism for bringing custom data into Foglight®. Once the script agent data are in the system, Foglight treats them the same way as any other data collected by any other agent type. Data collected by script agents can then be used in rules, derived metrics, dashboards, and other Foglight components.

Custom script agents interact with the Agent Manager through the Foglight collector executable. Script-based custom agents output data to standard output (`STDOUT`). The Foglight collector reads this data and retransmits it to the Agent Manager.

The output format is straightforward:

```
TABLE TableName
START_SAMPLE_PERIOD
Field = Value
END_SAMPLE_PERIOD
END_TABLE
```

There are two types of scripts: *Type 1* and *Type 2* scripts. Foglight calls *Type 1* scripts every time they need to collect data. In *Type 1* scripts, the collector executes the script, then stands by for a time period specified in the agent properties. When the standby period ends, the collector becomes active and reruns the script. *Type 1* scripts are useful for collecting data that does not require calculations from multiple collection periods. Sample *Type 1* scripts are available in the server installation directory: *Type1_NT_Script.bat* (Windows) and *Type1_Unix_Script.sh* (Unix).

Type 2 scripts control their own collection frequency cycle. In *Type 2* scripts, the Foglight collector executes the script and remains open. The script controls the standby period instead of the agent properties. *Type 2* scripts perform data calculations before the data enters the database and measure changes between collection periods. A sample Windows *Type 2* script is available in the server installation directory: *Type2_NT_Script.bat*.

Type 2 scripts are more complex because the script writer must handle looping and honor the sampling interval from the server. This might be necessary if the length of the loop is important for calculating rates. For the purposes of getting started, use *Type 1*. This will minimize the complexity. Switch to *Type 2* once you have a reason for hand-coding the loop.

You can use the Build Script Agent dashboard to upload custom agent script to the server. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Tooling > Script Agent Builder**.

Figure 19. The Script Agent Builder dashboard.

Script Agent Builder

Jun 25, 2013 3:17:59 PM EDT | Reports

Script Agent Builder

Custom script agents interact with the Foglight Agent Manager through the Foglight collector executable. User can use any scripting language to write scripts. Scriptbased custom agents output to standard output (stdout) and the Foglight collector reads the data and retransmits it to the Foglight Agent Manager. The Build Script Agent dashboard allows to upload an agent script to the Foglight Management Server.

Script File Browse...

Cartridge Version

Auto Generate Version

Build Script Agent

Using the Query Language

The Query Service gives users and Management Server services a way to make queries about topology objects and monitored metrics using the Foglight query language. Query language is used in rule conditions and derived metric expressions.

Typically, when working with rules and derived metrics, you first write a *topology query* to scope on a specific subset of the topology model, then write a script that performs a mathematical operation against that data subset. *Metric queries* retrieve metric values from one or more objects over a specified period of time. They cannot be used to set the scope of rules and derivations, but rather to query the database for the value of a particular metric over a specific period of time.

For more information about the query language, its syntax, and examples, see the related help topics accessible from this section in the help.

Retrieving Data with the REST API

The Foglight REST API is an application programming interface (API) that uses HTTP requests to GET, PUT, POST, and DELETE data. REST APIs are protected by authentications, which means you need retrieve an access token before using REST APIs. For more information about the Foglight REST API, refer to the *Foglight REST API Reference Guide*.

Retrieving Data with Scripts and Queries

In some cases, you may be required to run scripts, at the request of Quest Support, or for other maintenance functions. You can use the Script Console dashboard to test sample scripts. This dashboard is accessible to users with the Administrator and Cartridge Developer roles only. To access this dashboard, from the navigation panel, click **Dashboards > Administration > Tooling > Script Console**.

Figure 20. The Script Console.

The screenshot shows the 'Script Console' interface. At the top, it displays the title 'Script Console' and a timestamp 'Wednesday, October 12, 2011 6:37 AM - 10:37 AM 4 hours'. Below the title is a description: 'Query and scripting service tool that allows testing of queries and scripts, and inspection of objects. This dashboard replaces and extends the Script Editor (deprecated)'. The main area is divided into several sections:

- Abstract Section:** A dropdown menu is set to 'Abstract' with a 'List Instances' button next to it.
- Table Section:** A table with columns 'ID', 'Name', and 'Type'. The table is currently empty, displaying the message 'There Is No Data To Display'. A search bar is located above the table.
- Object ID Section:** A text input field for 'Object ID' and a 'Display' button.
- Scripts Section:** A tabbed interface with 'Object Panel' and 'Scripts' tabs. Below the tabs, it says 'No Scoping Object Selected' and provides a 'Launch Service Layer Documentation' button.
- Scripts List Section:** A table with columns 'Script Name', 'Associated Cartridge Name', and 'Scope Topology Type'. The table contains two entries: 'My Sample Script' and 'Alarms'. Above the table are buttons for 'Add', 'Remove Selected', 'Export Scripts', 'Import Scripts', 'Select All', and 'Select None', along with a search bar.

About Us

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.