QUEST SOFTWARE®

Foglight® 5.5.8

Administration and Configuration Guide

The information in this document is provided in connection with Quest 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 products. EXCEPT AS SET FORTH IN QUEST'S TERMS AND CONDITIONS AS SPECIFIED IN THE LICENSE AGREEMENT FOR THIS PRODUCT, QUEST 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 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 HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. Quest 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 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 World Headquarters
LEGAL Dept
5 Polaris Way
Aliso Viejo, CA 92656
www.quest.com
email: legal@quest.com

Refer to our Web site for regional and international office information.

Patents

This product includes patent pending technology.

Trademarks

Quest, Quest Software, the Quest Software logo, Foglight, IntelliProfile, PerformaSure, Spotlight, StealthCollect, TOAD, Tag and Follow, Vintela Single Sign-on for Java, and vFoglight are trademarks and registered trademarks of Quest Software, Inc in the United States of America and other countries. For a complete list of Quest Software's trademarks, please see http://www.quest.com/legal/trademark-information.aspx. Other trademarks and registered trademarks are property of their respective owners.

Third Party Contributions

Foglight contains some third party components. For a complete list, see the License Credits page in Foglight online help.

Administration and Configuration Guide
November 2010
Version 5.5.8

Table of Contents

Introduction to this Guide

This *Administration and Configuration Guide* provides conceptual information about Foglight administration, and instructions on how to use 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 chapters of this guide. The following chapters focus on day-to-day administration and fine-tuning. Advanced Foglight administrators can find information on key tools used to manage Foglight in the final chapter.

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

About Quest Software, Inc.

Quest Software simplifies and reduces the cost of managing IT for more than 100,000 customers worldwide. Our innovative solutions make solving the toughest IT management problems easier, enabling customers to save time and money across physical, virtual and cloud environments. For more information about Quest go to *www.quest.com*.

Contacting Quest Software

Email	*info@quest.com*
Mail	Quest Software, Inc. World Headquarters 5 Polaris Way Aliso Viejo, CA 92656 USA
Web site	*www.quest.com*

Refer to our Web site for regional and international office information.

Contacting Quest Support

Quest Support is available to customers who have a trial version of a Quest product or who have purchased a Quest product and have a valid maintenance contract. Quest Support provides unlimited 24x7 access to SupportLink, our self-service portal. Visit SupportLink at *http://support.quest.com*.

From SupportLink, you can do the following:

- Retrieve thousands of solutions from our online Knowledgebase
- Download the latest releases and service packs
- Create, update and review Support cases

View the *Global Support Guide* for a detailed explanation of support programs, online services, contact information, policies and procedures. The guide is available at: *http://support.quest.com*.

1

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

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 Foglight licensing, in the browser interface, open the **Help** tab in the action panel, and from there, navigate to **Using Foglight > Administration and Configuration Help > Configuring Foglight for Initial Use > Managing Users and Security**. You will find a list of reference topics at the bottom of the page.

Logging in to Foglight

Foglight user interface runs inside a Web browser. Before you log in, you need to ensure that your Foglight Management Server is up and running, and to 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 Foglight 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 Foglight Management Server is running. The security settings associated with your user account determine which dashboards you can access.

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

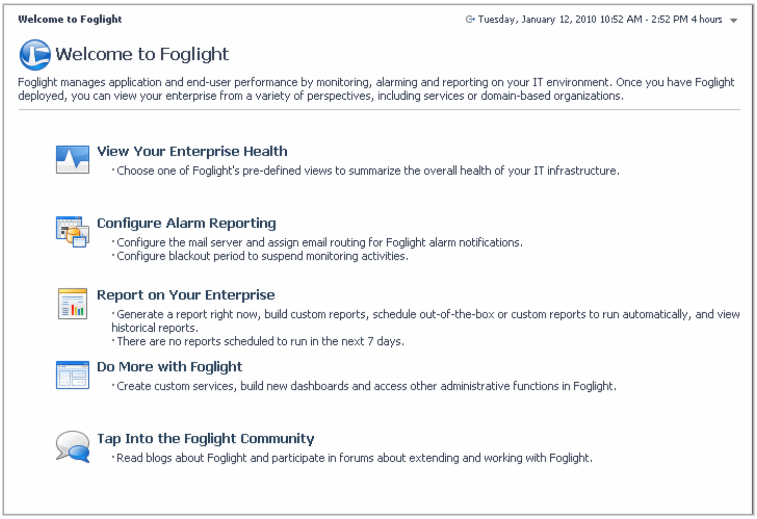

Welcome to Foglight

Tuesday, January 12, 2010 10:52 AM - 2:52 PM 4 hours

Welcome to Foglight

Foglight manages application and end-user performance by monitoring, alarming and reporting on your IT environment. Once you have Foglight deployed, you can view your enterprise from a variety of perspectives, including services or domain-based organizations.

View Your Enterprise Health
- Choose one of Foglight's pre-defined views to summarize the overall health of your IT infrastructure.

Configure Alarm Reporting
- Configure the mail server and assign email routing for Foglight alarm notifications.
- Configure blackout period to suspend monitoring activities.

Report on Your Enterprise
- Generate a report right now, build custom reports, schedule out-of-the-box or custom reports to run automatically, and view historical reports.
- There are no reports scheduled to run in the next 7 days.

Do More with Foglight
- Create custom services, build new dashboards and access other administrative functions in Foglight.

Tap Into the Foglight Community
- Read blogs about Foglight and participate in forums about extending and working with Foglight.

The appearance of the Welcome page depends on your user permissions. If your user account belongs to group that has the Administrator role, this is what you see when you log into Foglight for the first time. The Welcome page lists the common administration tasks that you typically perform upon logging in to Foglight.

Managing Foglight Licenses

Foglight includes a licensing capability that restricts access to only to 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.

Some Foglight cartridges are license-protected, while others do not require a license. The cartridges included with the server do not require any additional license. The Foglight Agent Manager and OS cartridges fall into this category. Some cartridges installed on top of the server require a license.

In a typical installation, you need a license for the Foglight Management Server, along with the license for each license-protected cartridge that exists in your monitoring environment. If a cartridge requires a license, you must 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 prior to installing its license, however it disables the cartridge until a valid cartridge license is installed.

You can install and manage Foglight licenses using the Manage Licenses dashboard. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > Manage Licenses**.

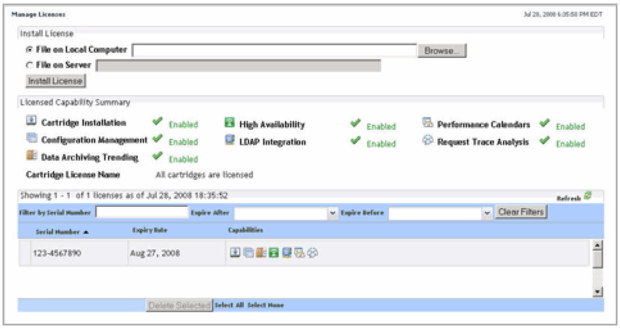

Cartridge Installation	✓ Enabled	High Availability	✓ Enabled	Performance Calendars	✓ Enabled
Configuration Management	✓ Enabled	LDAP Integration	✓ Enabled	Request Trace Analysis	✓ Enabled
Data Archiving Trending	✓ Enabled				

Cartridge License Name All cartridges are licensed

Showing 1 - 1 of 1 licenses as of Jul 28, 2008 18:35:52

Serial Number ▲	Expiry Date	Capabilities
123-4567890	Aug 27, 2008	

Configuring Email Notifications

Foglight uses email notifications to send reports or alarm-related messages to email recipients when certain thresholds are reached. This 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. This can be achieved using the Email Configuration dashboard. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > Email Configuration**.

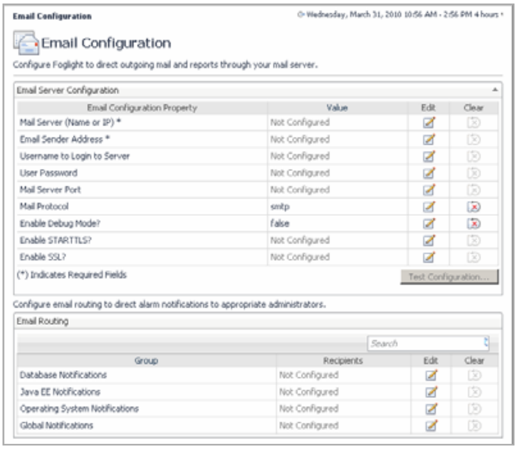

Email Configuration

Wednesday, March 31, 2010 10:56 AM - 2:56 PM 4 hours

Email Configuration

Configure Foglight to direct outgoing mail and reports through your mail server.

Email Server Configuration

Email Configuration Property	Value	Edit	Clear
Mail Server (Name or IP) *	Not Configured	✎	☒
Email Sender Address *	Not Configured	✎	☒
Username to Login to Server	Not Configured	✎	☒
User Password	Not Configured	✎	☒
Mail Server Port	Not Configured	✎	☒
Mail Protocol	smtp	✎	☒
Enable Debug Mode?	false	✎	☒
Enable STARTTLS?	Not Configured	✎	☒
Enable SSL?	Not Configured	✎	☒

(*) Indicates Required Fields

Configure email routing to direct alarm notifications to appropriate administrators.

Email Routing

Group	Recipients	Edit	Clear
Database Notifications	Not Configured	✎	☒
Java EE Notifications	Not Configured	✎	☒
Operating System Notifications	Not Configured	✎	☒
Global Notifications	Not Configured	✎	☒

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 has access to. 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 Foglight 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, on the navigation panel, choose **Dashboards > Administration > Users & Security**.

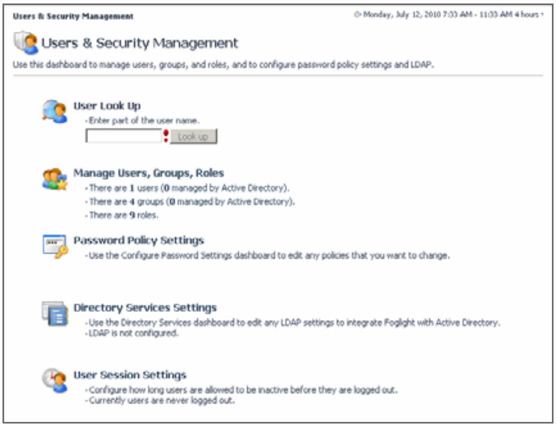

Users & Security Management

Monday, July 12, 2010 7:03 AM - 11:03 AM 4 hours

Users & Security Management

Use this dashboard to manage users, groups, and roles, and to configure password policy settings and LDAP.

User Look Up
- Enter part of the user name.

Look up

Manage Users, Groups, Roles
- There are **1** users (**0** managed by Active Directory).
- There are **4** groups (**0** managed by Active Directory).
- There are **9** roles.

Password Policy Settings
- Use the Configure Password Settings dashboard to edit any policies that you want to change.

Directory Services Settings
- Use the Directory Services dashboard to edit any LDAP settings to integrate Foglight with Active Directory.
- LDAP is not configured.

User Session Settings
- Configure how long users are allowed to be inactive before they are logged out.
- Currently users are never logged out.

From there, to start managing user access, click **Manage Users, Groups, Roles**.

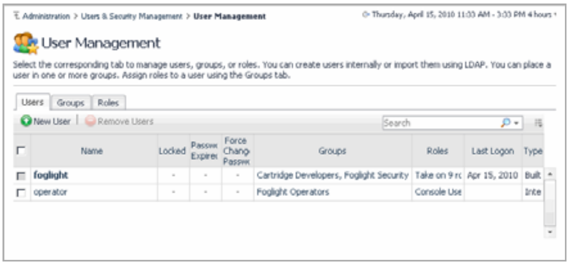

Administration > Users & Security Management > **User Management**

Thursday, April 15, 2010 11:03 AM - 3:03 PM 4 hours

User Management

Select the corresponding tab to manage users, groups, or roles. You can create users internally or import them using LDAP. You can place a user in one or more groups. Assign roles to a user using the Groups tab.

□	Name	Locked	Passw Expire	Force Chang Passw	Groups	Roles	Last Logon	Type
□	**foglight**	-	-	-	Cartridge Developers, Foglight Security	Take on 9 rc	Apr 15, 2010	Built
□	operator	-	-	-	Foglight Operators	Console Use		Inte

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**.

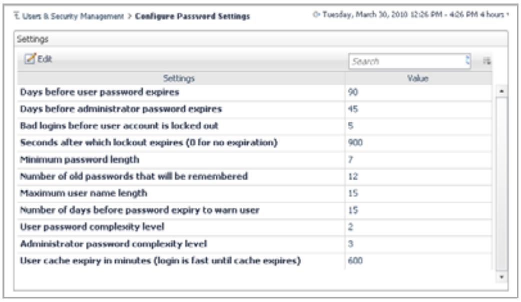

Users & Security Management > Configure Password Settings
Tuesday, March 30, 2010 12:26 PM - 4:26 PM 4 hours
Settings
Edit
Search

Settings	Value
Days before user password expires	90
Days before administrator password expires	45
Bad logins before user account is locked out	5
Seconds after which lockout expires (0 for no expiration)	900
Minimum password length	7
Number of old passwords that will be remembered	12
Maximum user name length	15
Number of days before password expiry to warn user	15
User password complexity level	2
Administrator password complexity level	3
User cache expiry in minutes (login is fast until cache expires)	600

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:

- Active Directory
- Sun Java Systems Directory Server
- OpenLDAP
- Novell eDirectory

You need to be familiar with the details of your LDAP directory service to perform this configuration. After configuring the LDAP directory service, Foglight creates a user account each time an LDAP user successfully logs into 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 into 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.

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 and Security dashboard, click **Directory Services Settings**.

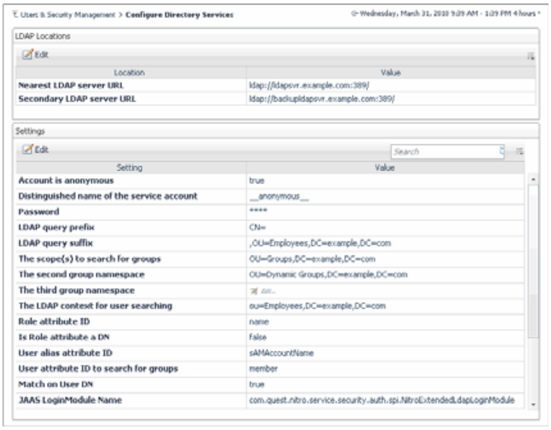

Users & Security Management > **Configure Directory Services**

Wednesday, March 31, 2010 9:09 AM - 1:09 PM 4 hours

LDAP Locations

Location	Value
Nearest LDAP server URL	ldap://ldapsvr.example.com:389/
Secondary LDAP server URL	ldap://backupldapsvr.example.com:389/

Settings

Setting	Value
Account is anonymous	true
Distinguished name of the service account	__anonymous__
Password	****
LDAP query prefix	CN=
LDAP query suffix	,OU=Employees,DC=example,DC=com
The scope(s) to search for groups	OU=Groups,DC=example,DC=com
The second group namespace	OU=Dynamic Groups,DC=example,DC=com
The third group namespace	Edit...
The LDAP context for user searching	ou=Employees,DC=example,DC=com
Role attribute ID	name
Is Role attribute a DN	false
User alias attribute ID	sAMAccountName
User attribute ID to search for groups	member
Match on User DN	true
JAAS LoginModule Name	com.quest.nitro.service.security.auth.spi.NitroExtendedLdapLoginModule

Configuring user sessions

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

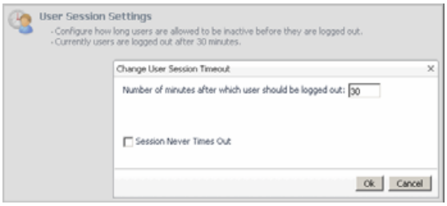

User Session Settings

- Configure how long users are allowed to be inactive before they are logged out.
- Currently users are logged out after 30 minutes.

Change User Session Timeout

Number of minutes after which user should be logged out: 30

☐ Session Never Times Out

Suspending Alarms and Data Collection

A blackout is a period of time 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. The Blackout Configuration dashboard allows you to suspend alarms and data collection for a specific period of time. 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 for the duration of the schedule. 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 for the duration of the blackout. For more information about topology models, see the *Getting Started 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.

Caution In addition to the features provided by the Blackout Configuration dashboard, topology and agent blackouts can also be configured using the command-line interface. However, the mechanism for creating blackouts using this other method 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 Configuration dashboard, 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 are 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 this dashboard, from the navigation panel, choose **Dashboards** > **Administration** > **Setup & Support** > **Blackout Configuration**.

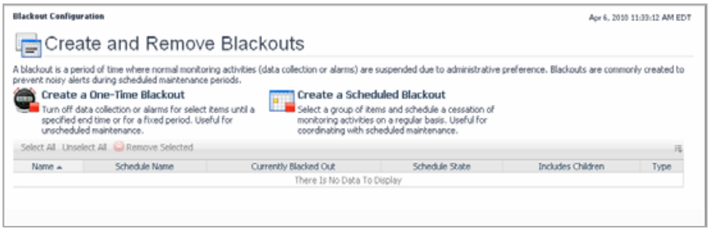

A blackout is a period of time where normal monitoring activities (data collection or alarms) are suspended due to administrative preference. Blackouts are commonly created to prevent noisy alerts during scheduled maintenance periods.

Create a One-Time Blackout
Turn off data collection or alarms for select items until a specified end time or for a fixed period. Useful for unscheduled maintenance.

Create a Scheduled Blackout
Select a group of items and schedule a cessation of monitoring activities on a regular basis. Useful for coordinating with scheduled maintenance.

Name ▲	Schedule Name	Currently Blacked Out	Schedule State	Includes Children	Type
There Is No Data To Display					

Additional Configuration Features

The administration interface comes with a set of dashboards that show run-time details, such as audit logs and diagnostic data that can be sent to Quest Support. Additionally, you can view host connection status, port numbers, database properties, and manage data collection blackouts and server support bundles.

Managing Support Bundles

Foglight allows you to gather diagnostic data and save it as a collection of files, called *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 Foglight Management Server, log files, and a list of cartridges installed on the Foglight Management Server. Foglight saves each server support bundle as a .ZIP file in the *<foglight_home>/support/<user_name>* directory on the machine hosting the Foglight Management Server. Foglight Agent Manager support bundles contain diagnostic data about the monitored host. When you create a monitored host support bundle, Foglight Agent Manager saves this data in a ZIP file in the *<foglight_agent_mgr_home>/state/default/support* directory on the computer hosting the Foglight Agent Manager.

The Manage Support Bundles dashboard allows you to create server support bundles, and to download, or delete Foglight Agent Manager and server support bundles. Use the Foglight Agent Manager Support Bundle dashboard to generate and download host support bundles. To access the Manage Support Bundles dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > Manage Support Bundles**.

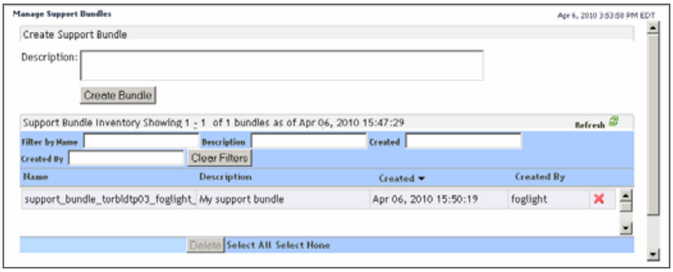

Manage Support Bundles
Apr 6, 2010 3:53:58 PM EDT

Create Support Bundle

Description:

Support Bundle Inventory Showing 1 - 1 of 1 bundles as of Apr 06, 2010 15:47:29

Name	Description	Created	Created By
support_bundle_torbldtp03_foglight_	My support bundle	Apr 06, 2010 15:50:19	foglight

To access the FglAM Support Bundle dashboard, from the navigation panel, choose **Dashboards > Administration > Agents > FglAM Support Bundle**.

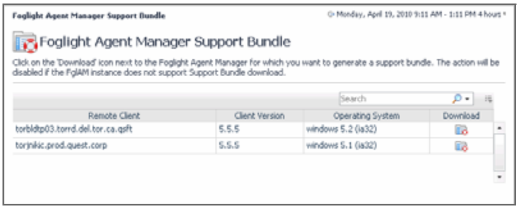

Foglight Agent Manager Support Bundle

Monday, April 19, 2010 9:11 AM - 1:11 PM 4 hours

Foglight Agent Manager Support Bundle

Click on the 'Download' icon next to the Foglight Agent Manager for which you want to generate a support bundle. The action will be disabled if the FglAM instance does not support Support Bundle download.

Remote Client	Client Version	Operating System	Download
torbldtp03.torrd.del.tor.ca.qsft	5.5.5	windows 5.2 (ia32)	
torjnikic.prod.quest.corp	5.5.5	windows 5.1 (ia32)	

Viewing Host Connection Status

Foglight uses the Foglight Agent Manager to communicate with monitored hosts. The Connection Status dashboard lists agent manager components that are connected to the server. For each agent manager instance, the list shows the host’s IP address, login time, request name, and request time.

To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > Connection Status**.

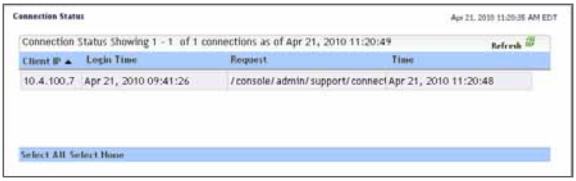

Connection Status

Apr 21, 2010 11:20:35 AM EDT

Connection Status Showing 1 - 1 of 1 connections as of Apr 21, 2010 11:20:49

Client IP ▲	Login Time	Request	Time
10.4.100.7	Apr 21, 2010 09:41:26	/console/admin/support/connect	Apr 21, 2010 11:20:48

Viewing Audited Entries

Foglight generates security and change logs that contain information about the users who are authenticated upon logging in to Foglight, 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, name of the user who initiated the operation, name of the service that performed the operation, and the operation name.

To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > View Audit Information**.

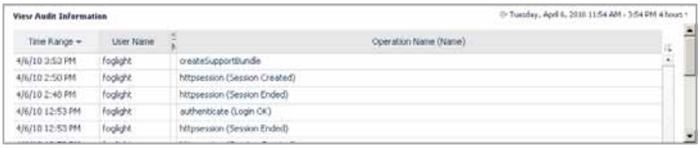

View Audit Information

Tuesday, April 6, 2010 11:54 AM - 3:54 PM 4 hours

Time Range	User Name	Operation Name (Name)
4/6/10 3:53 PM	foglight	createSupportBundle
4/6/10 2:50 PM	foglight	httpsession (Session Created)
4/6/10 2:40 PM	foglight	httpsession (Session Ended)
4/6/10 12:53 PM	foglight	authenticate (Login OK)
4/6/10 12:53 PM	foglight	httpsession (Session Ended)

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 Foglight Management Server, WCF, and JVM versions; these settings cannot be changed unless you choose to upgrade to a higher version of Foglight. Additional display-only settings indicate the OS of the computer on which the Foglight Management Server is installed, and its patch level.

The configuration values are set in the *foglight.config* file and the Foglight registry. For example, the database settings are typically set in *foglight.config*, while global mail settings are specified in the Foglight registry. Editing the configuration file requires a restart of the Foglight 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 Foglight Configuration dashboard. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Setup & Support > Management Server Configuration**.

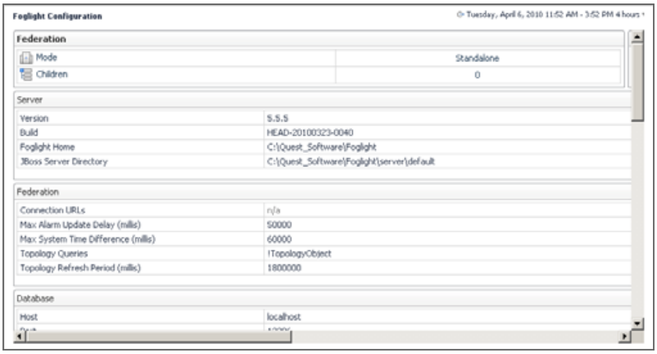

Foglight Configuration

Tuesday, April 6, 2010 11:52 AM - 3:52 PM 4 hours

Federation

Mode	Standalone
Children	0

Server

Version	5.5.5
Build	HEAD-20100323-0040
Foglight Home	C:\Quest_Software\Foglight
JBoss Server Directory	C:\Quest_Software\Foglight\server\default

Federation

Connection URLs	n/a
Max Alarm Update Delay (millis)	50000
Max System Time Difference (millis)	60000
Topology Queries	!TopologyObject
Topology Refresh Period (millis)	1800000

Database

Host	localhost

2

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.

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 **Using Foglight > Administration and Configuration Help > Extending Your Monitoring Reach with Foglight Cartridges > Managing Foglight Cartridges**. You will find a list of reference topics at the bottom of the page.

Managing Foglight Cartridges

A Foglight Management Server installation includes a set of default cartridges that contain models, views and monitoring policies. In addition to the 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 Foglight 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, choose **Dashboards > Administration > Cartridges > Cartridge Inventory**.

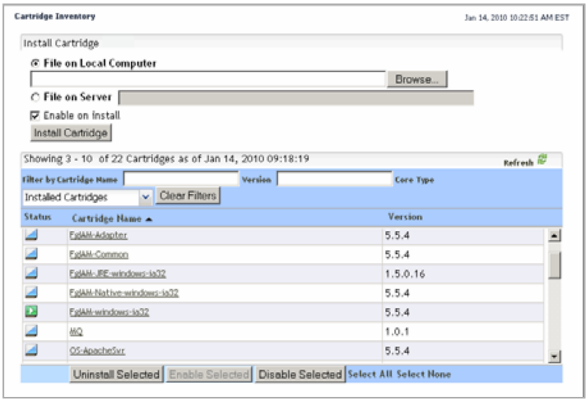

Cartridge Inventory

Jan 14, 2010 10:22:51 AM EST

Install Cartridge

● File on Local Computer

○ File on Server

☑ Enable on install

Showing 3 - 10 of 22 Cartridges as of Jan 14, 2010 09:18:19

Filter by Cartridge Name · Version · Core Type

Installed Cartridges

Status	Cartridge Name ▲	Version
	FglAM-Adapter	5.5.4
	FglAM-Common	5.5.4
	FglAM-JRE-windows-ia32	1.5.0.16
	FglAM-Native-windows-ia32	5.5.4
	FglAM-windows-ia32	5.5.4
	MQ	1.0.1
	OS-ApacheSvr	5.5.4

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

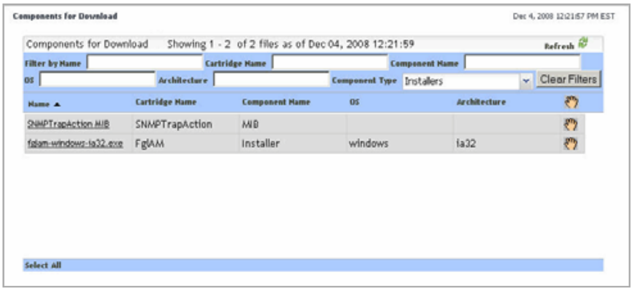

Components for Download

Dec 4, 2008 12:21:57 PM EST

Components for Download Showing 1 - 2 of 2 files as of Dec 04, 2008 12:21:59

Filter by Name · Cartridge Name · Component Name

OS · Architecture · Component Type: Installers

Name ▲	Cartridge Name	Component Name	OS	Architecture	
SNMPTrapAction.MIB	SNMPTrapAction	MIB			
fglam-windows-ia32.exe	FglAM	Installer	windows	ia32	

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 Foglight Agent Manager starts up and stops with the Foglight Management Server. This embedded instance can be used to monitor the host on which the Foglight Management Server is installed. To monitor additional hosts in your environment, most cartridges first require a Foglight 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 Foglight Agent Manager 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 Foglight Agent Manager on the hosts you want to monitor.

There are two dashboards that allow you to deploy agent packages and create agent instances: Agent Status and Agent Hosts. Use the Agent Hosts dashboard to deploy agent packages and create agent instances on multiple hosts. The Agent Status dashboard allows you to perform these operations one host at a time.

To access the Agent Host dashboard, from the navigation panel, choose **Dashboards > Administration > Agents > Agent Hosts**.

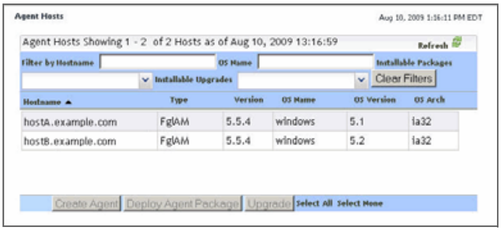

Agent Hosts

Aug 10, 2009 1:16:11 PM EDT

Agent Hosts Showing 1 - 2 of 2 Hosts as of Aug 10, 2009 13:16:59

Hostname ▲	Type	Version	OS Name	OS Version	OS Arch
hostA.example.com	FglAM	5.5.4	windows	5.1	ia32
hostB.example.com	FglAM	5.5.4	windows	5.2	ia32

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

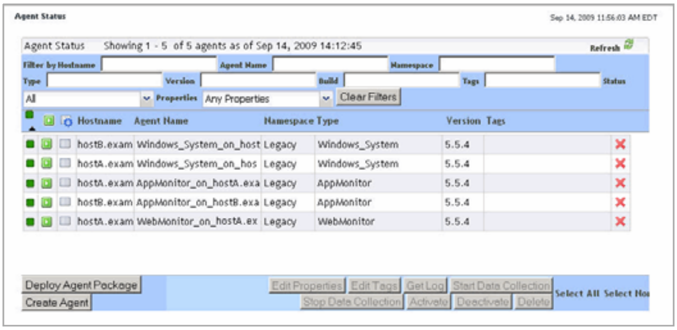

Agent Status Showing 1 - 5 of 5 agents as of Sep 14, 2009 14:12:45

Type: All
Properties: Any Properties

			Hostname	Agent Name	Namespace	Type	Version	Tags	
🟩	▶	□	hostB.exam	Windows_System_on_host	Legacy	Windows_System	5.5.4		✕
🟩	▶	□	hostA.exam	Windows_System_on_hos	Legacy	Windows_System	5.5.4		✕
🟩	▶	□	hostA.exam	AppMonitor_on_hostA.exa	Legacy	AppMonitor	5.5.4		✕
🟩	▶	□	hostB.exam	AppMonitor_on_hostB.exa	Legacy	AppMonitor	5.5.4		✕
🟩	▶	□	hostA.exam	WebMonitor_on_hostA.ex	Legacy	WebMonitor	5.5.4		✕

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 Foglight Management Server.

Default versions of agent properties are installed with the cartridge. You can edit the default properties, create sets of properties that apply only to a specific agent instance, and create edited clones of property sets that are used by a subset of the agents of a certain type.

There are two types of agent properties: *Primary properties* are included in the agent component and their settings can be specific to the agent type or the agent instance. If you do not change agent properties for an instance, Foglight uses the default properties that come with that agent type. *Secondary properties* are type-specific and are global in nature, which means that any changes to them affect all instances of that agent type. To make lists instance-specific, clone a list and set the agent property to use the cloned list.

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, choose **Dashboards** > **Administration** > **Agents** > **Agent Properties**, or select an agent instance on the Agent Status dashboard and click **Edit Properties**.

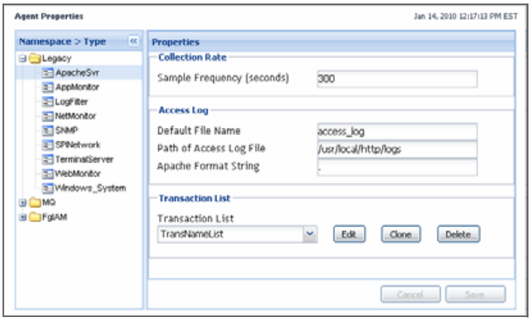

Collection Rate

Sample Frequency (seconds): 300

Access Log

Default File Name: access_log
Path of Access Log File: /usr/local/http/logs
Apache Format String: -

Transaction List

Transaction List: TransNameList

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. There are some agents that use other types of agent adapters, such as the Java EE Agent.

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, choose **Dashboards** > **Administration** > **Agents** > **Agent Adapters**.

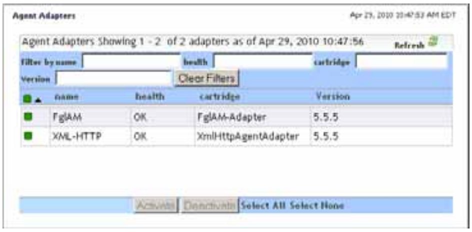

Agent Adapters Showing 1 - 2 of 2 adapters as of Apr 29, 2010 10:47:56

🟩	name	health	cartridge	Version
🟩	FglAM	OK	FglAM-Adapter	5.5.5
🟩	XML-HTTP	OK	XmlHttpAgentAdapter	5.5.5

3

Administering Foglight

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

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 **Using Foglight > 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.

Performing Standard Maintenance Tasks

This section describes the standard maintenance tasks that 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.

Recommended Self-Monitoring Automation

Foglight Management Server up/down status

A separate Remote Monitor process can be configured to watch the Foglight HA (High Availability) 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 *Installation and Setup Guide* set.

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 navigation panel, choose **Dashboards** > **Administration** > **Setup & Support** > **Email Configuration**

Daily Maintenance Tasks

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

Assuming you have automated self-monitoring as directed above, there is no need to perform the checks described below.

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

Foglight Management Server

- Validate that the Foglight Management Server is running by logging into the Foglight Management Server on a daily basis.
- Validate that the Foglight 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 *Performance Tuning Field Guide*.

Foglight Agents and Agent Manager

- Validate that all Foglight 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
  - Edit rule view
  - Services (All Alarms)
  - Hosts (all hosts)
  - Hosts (monitored hosts only)

- Manage Rules
- Reports Manager

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 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.

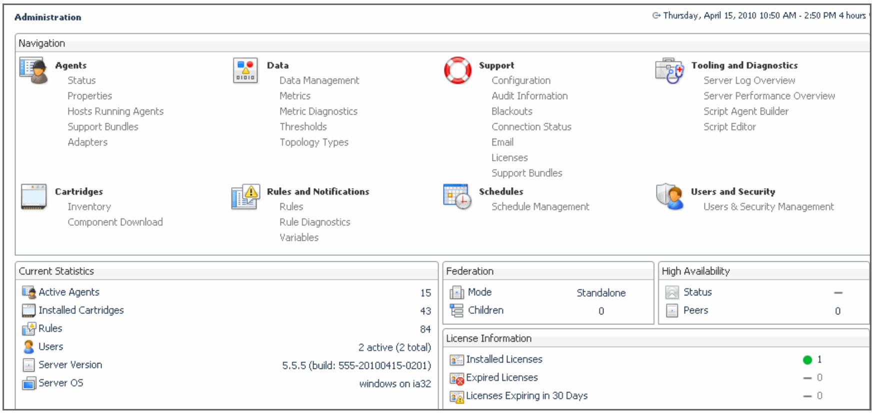

Current Statistics

Active Agents	15
Installed Cartridges	43
Rules	84
Users	2 active (2 total)
Server Version	5.5.5 (build: 555-20100415-0201)
Server OS	windows on ia32

Federation

Mode	Standalone
Children	0

High Availability

Status	—
Peers	0

License Information

Installed Licenses	● 1
Expired Licenses	— 0
Licenses Expiring in 30 Days	— 0

You can access this dashboard from the navigation panel, by choosing **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, described in this chapter.

Exploring the Data Model

Foglight collects data about your system and dynamically builds topology models at run-time. A topology model organizes the data in a way that represents the logical and physical relationship between items in your monitored environment and provides the context for the collected metrics.

Topology models consist of nodes, where each node is an instance of a topology object. The nature of the monitored environment dictates the structure and complexity of each topology model and the collection of available topology types. A basic server installation includes a set of core topology types, and each installed cartridge adds to that collection.

Use the Schema Browser dashboard to view information about the available data types, their relationships in the data model, properties, and object instances. This dashboard can help you to better understand the data model structure and learn about existing object dependencies. To access this dashboard, from the navigation panel, choose **Dashboards > Management Server > Schema > Schema Browser**.

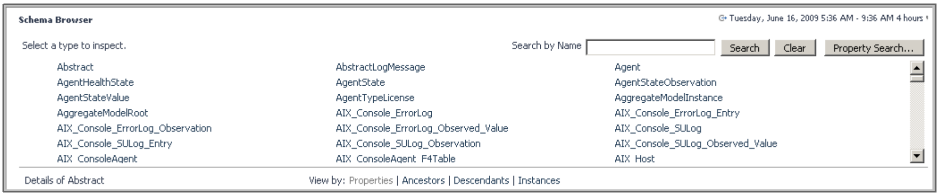

Schema Browser

Tuesday, June 16, 2009 5:36 AM - 9:36 AM 4 hours

Select a type to inspect.

Search by Name

Abstract	AbstractLogMessage	Agent
AgentHealthState	AgentState	AgentStateObservation
AgentStateValue	AgentTypeLicense	AggregateModelInstance
AggregateModelRoot	AIX_Console_ErrorLog	AIX_Console_ErrorLog_Entry
AIX_Console_ErrorLog_Observation	AIX_Console_ErrorLog_Observed_Value	AIX_Console_SULog
AIX_Console_SULog_Entry	AIX_Console_SULog_Observation	AIX_Console_SULog_Observed_Value
AIX ConsoleAgent	AIX ConsoleAgent F4Table	AIX Host

Details of Abstract

View by: Properties | Ancestors | Descendants | Instances

If required, you can add custom topology types to the schema. For details, see “Expanding Your Collection of Topology Types” on page 45.

Managing 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 Data Management 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.

Using this dashboard 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.
- 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, choose **Dashboards > Management Server > Servers > Data Management**.

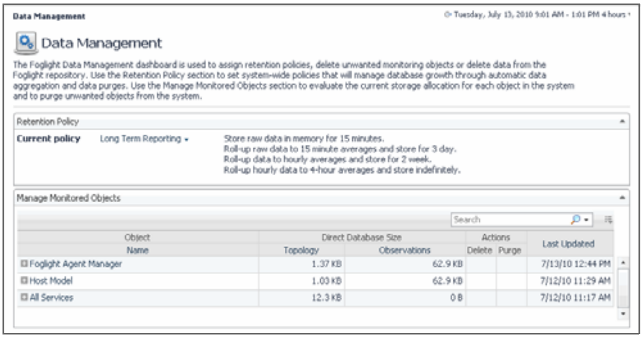

Data Management

Tuesday, July 13, 2010 9:01 AM - 1:01 PM 4 hours

Data Management

The Foglight Data Management dashboard is used to assign retention policies, delete unwanted monitoring objects or delete data from the Foglight repository. Use the Retention Policy section to set system-wide policies that will manage database growth through automatic data aggregation and data purges. Use the Manage Monitored Objects section to evaluate the current storage allocation for each object in the system and to purge unwanted objects from the system.

Retention Policy

Current policy Long Term Reporting

Store raw data in memory for 15 minutes.
Roll-up raw data to 15 minute averages and store for 3 day.
Roll-up data to hourly averages and store for 2 week.
Roll-up hourly data to 4-hour averages and store indefinitely.

Manage Monitored Objects

Search

Object	Direct Database Size		Actions		Last Updated
Name	Topology	Observations	Delete	Purge	
⊞ Foglight Agent Manager	1.37 KB	62.9 KB			7/13/10 12:44 PM
⊞ Host Model	1.03 KB	62.9 KB			7/12/10 11:29 AM
⊞ All Services	12.3 KB	0 B			7/12/10 11:17 AM

Another way to control the retention policies on a per-object basis is using the Manage Retention Policies dashboard. For more information, see “Managing Data Retention” on page 44.

In addition to the Data Management dashboard that allows 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, choose **Dashboards > Management Server > Servers > Database Overview**.

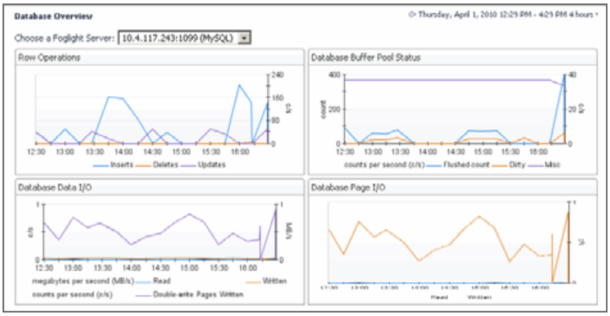

Database Overview

Thursday, April 1, 2010 12:29 PM - 4:29 PM 4 hours

Choose a Foglight Server: 10.4.117.243:1099 (MySQL)

Monitoring 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, choose **Dashboards** > **Management Server** > **Diagnostic** > **Performance**.

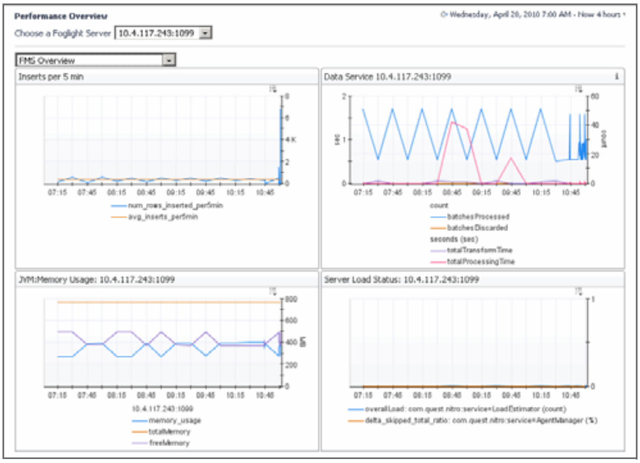

Performance Overview

Choose a Foglight Server 10.4.117.243:1099

FMS Overview

Wednesday, April 20, 2010 7:00 AM - Now 4 hours
Inserts per 5 min
8
6
4 K
2
0
07:15
07:45
08:15
08:45
09:15
09:45
10:15
10:45
num_rows_inserted_per5min
avg_inserts_per5min
Data Service 10.4.117.243:1099
2
1
0
sec
60
40
20
0
count
batchesProcessed
batchesDiscarded
seconds (sec)
totalTransformTime
totalProcessingTime
JVM:Memory Usage: 10.4.117.243:1099
800
600
400
200
0
MB
10.4.117.243:1099
memory_usage
totalMemory
freeMemory
Server Load Status: 10.4.117.243:1099
1
0
overallLoad: com.quest.nitro:service=LoadEstimator (count)
delta_skipped_total_ratio: com.quest.nitro:service=AgentManager (%)

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, choose **Dashboards** > **Management Server** > **Servers** > **Management Server View**.

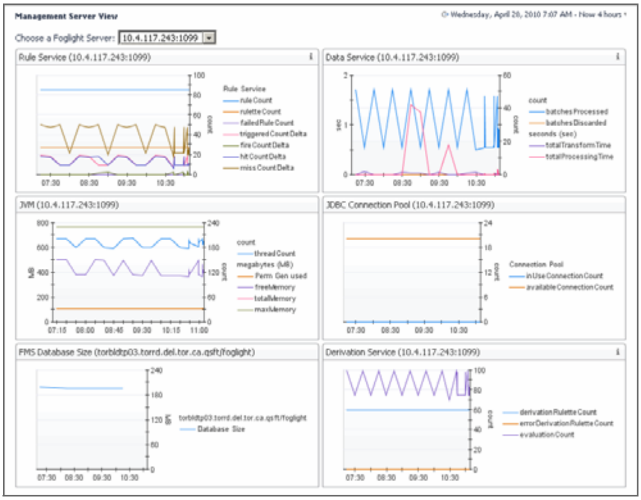

The Management Server Metrics dashboard contains a data tree and a metric browser that shows the metrics associated with the server. Use it to quickly determine if the resources available to the server are in line with the pattern of process requests for a given time period. For example, a performance slowdown can be caused by a decrease in the amount of paging space or a high number of processes in the queue. To access this dashboard, from the navigation panel, choose **Dashboards > Management Server > Servers > Management Server Metrics**.

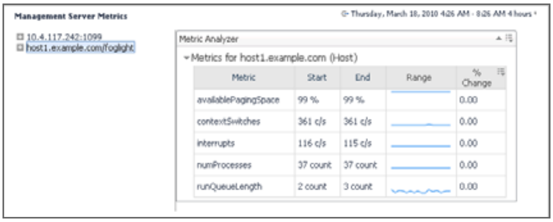
Management Server Metrics

Thursday, March 18, 2010 4:26 AM - 8:26 AM 4 hours

⊞ 10.4.117.242:1099
⊞ host1.example.com/foglight

Metric Analyzer

Metrics for host1.example.com (Host)

Metric	Start	End	Range	% Change
availablePagingSpace	99 %	99 %		0.00
contextSwitches	361 c/s	361 c/s		0.00
interrupts	116 c/s	115 c/s		0.00
numProcesses	37 count	37 count		0.00
runQueueLength	2 count	3 count		0.00

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, choose **Dashboards > Management Server > Diagnostic > LogAnalyzer**.

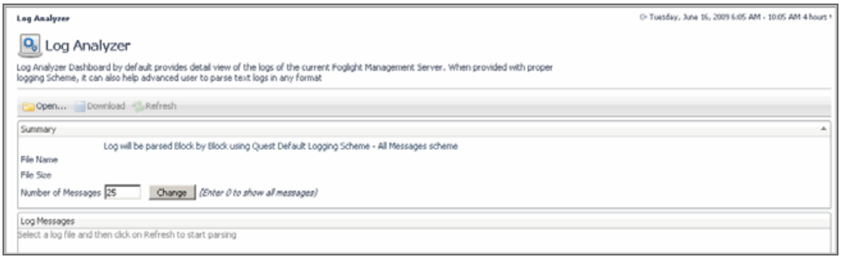

Log Analyzer

Tuesday, June 16, 2009 6:05 AM - 10:05 AM 4 hours

Log Analyzer

Log Analyzer Dashboard by default provides detail view of the logs of the current Foglight Management Server. When provided with proper logging Scheme, it can also help advanced user to parse text logs in any format

Summary

Log will be parsed Block by Block using Quest Default Logging Scheme - All Messages scheme

File Name

File Size

Number of Messages 25 Change *(Enter 0 to show all messages)*

Log Messages

Select a log file and then click on Refresh to start parsing

Monitoring Foglight Agent Manager Performance

Foglight uses Foglight Agent Manager to communicate with monitored hosts. The embedded Foglight Agent Manager can be used to monitor the host on which the Foglight Management Server is installed. Your monitoring environment typically includes a number of Foglight Agent Manager instances that are installed on different hosts. You can monitor their performance using the Foglight Agent Manager 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, choose **Dashboards** > **Management Server** > **Diagnostic** > **Foglight Agent Manager**.

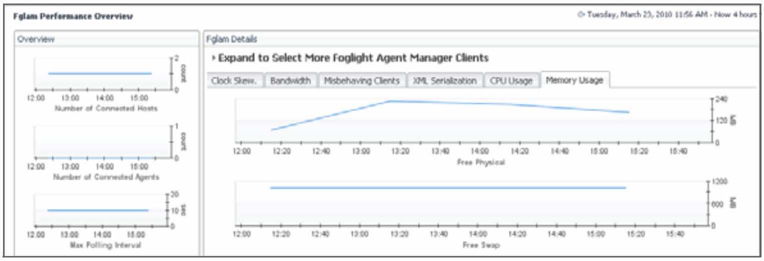

Fglam Performance Overview

Tuesday, March 23, 2010 11:56 AM - Now 4 hours

Associating 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, choose **Dashboards > Services > Object Groups**.

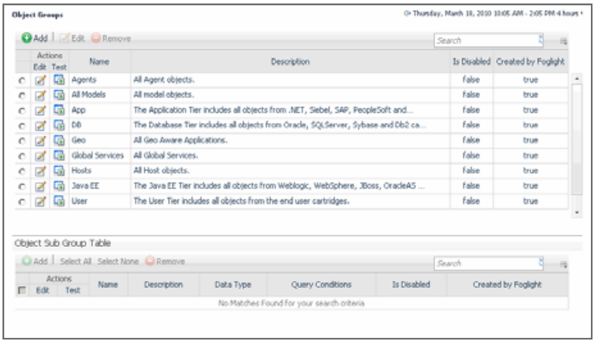

	Actions		Name	Description	Is Disabled	Created by Foglight
	Edit	Test				
○			Agents	All Agent objects.	false	true
○			All Models	All model objects.	false	true
○			App	The Application Tier includes all objects from .NET, Siebel, SAP, PeopleSoft and...	false	true
○			DB	The Database Tier includes all objects from Oracle, SQLServer, Sybase and Db2 ca...	false	true
○			Geo	All Geo Aware Applications.	false	true
○			Global Services	All Global Services.	false	true
○			Hosts	All Host objects.	false	true
○			Java EE	The Java EE Tier includes all objects from Weblogic, WebSphere, JBoss, OracleAS ...	false	true
○			User	The User Tier includes all objects from the end user cartridges.	false	true

□	Actions		Name	Description	Data Type	Query Conditions	Is Disabled	Created by Foglight
	Edit	Test						
No Matches Found for your search criteria								

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

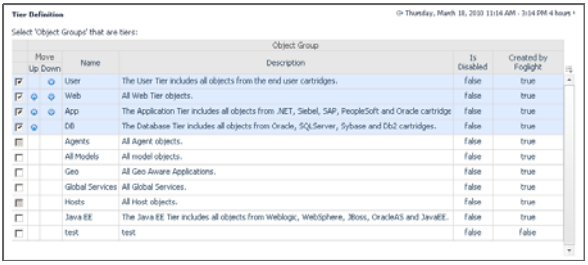

Select 'Object Groups' that are tiers:

	Object Group					
	Move		Name	Description	Is Disabled	Created by Foglight
	Up	Down				
☑		○	User	The User Tier includes all objects from the end user cartridges.	false	true
☑	○	○	Web	All Web Tier objects.	false	true
☑	○	○	App	The Application Tier includes all objects from .NET, Siebel, SAP, PeopleSoft and Oracle cartridge	false	true
☑	○		DB	The Database Tier includes all objects from Oracle, SQLServer, Sybase and Db2 cartridges.	false	true
□			Agents	All Agent objects.	false	true
□			All Models	All model objects.	false	true
□			Geo	All Geo Aware Applications.	false	true
□			Global Services	All Global Services.	false	true
□			Hosts	All Host objects.	false	true
□			Java EE	The Java EE Tier includes all objects from Weblogic, WebSphere, JBoss, OracleAS and JavaEE.	false	true
□			test	test	false	false

Backing Up and Restoring 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 (MS SQL, MySQL, 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 help.

4

Tuning Foglight for Optimal Performance

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.

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 **Using Foglight > Administration and Configuration Help > Tuning Foglight for Optimal Performance > Using Metric Calculations in Foglight > Working with Foglight Registry Variables**. You will find a list of reference topics at the bottom of the page.

Using Foglight Rules to Report on Bottlenecks

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 includes 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.

You can create and manage rules using the Manage Rules dashboard. A typical installation can include a large number of rules. The Manage Rules dashboard lists all rules that exist in your environment, and allows you to drill down to rule definitions. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Rules & Notifications > Manage Rules**.

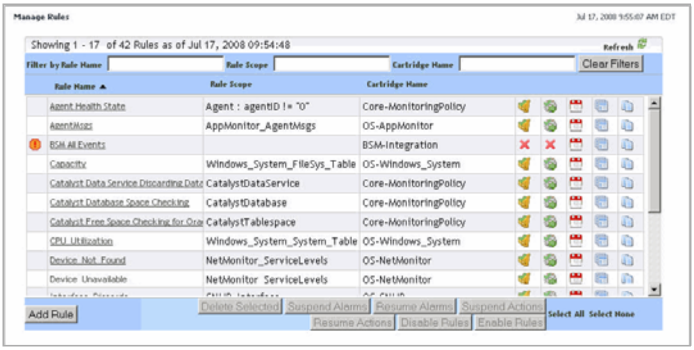

Manage Rules

Jul 17, 2008 9:55:07 AM EDT

Showing 1 - 17 of 42 Rules as of Jul 17, 2008 09:54:48

Rule Name ▲	Rule Scope	Cartridge Name
Agent Health State	Agent : agentID != "0"	Core-MonitoringPolicy
AgentMsgs	AppMonitor_AgentMsgs	OS-AppMonitor
BSM All Events		BSM-Integration
Capacity	Windows_System_FileSys_Table	OS-Windows_System
Catalyst Data Service Discarding Data	CatalystDataService	Core-MonitoringPolicy
Catalyst Database Space Checking	CatalystDatabase	Core-MonitoringPolicy
Catalyst Free Space Checking for Ora	CatalystTablespace	Core-MonitoringPolicy
CPU_Utilization	Windows_System_System_Table	OS-Windows_System
Device_Not_Found	NetMonitor_ServiceLevels	OS-NetMonitor
Device_Unavailable	NetMonitor_ServiceLevels	OS-NetMonitor

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, choose **Dashboards** > **Administration** > **Rules & Notifications** > **Rule Diagnostics**.

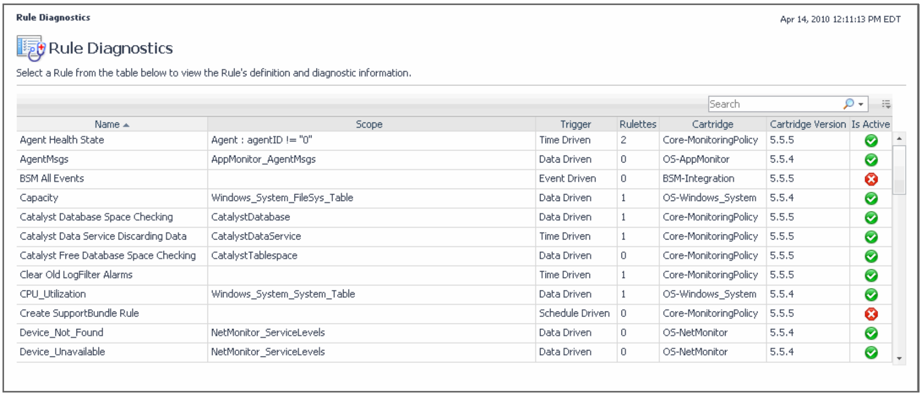

Rule Diagnostics

Apr 14, 2010 12:11:13 PM EDT

Rule Diagnostics

Select a Rule from the table below to view the Rule's definition and diagnostic information.

Name ▲	Scope	Trigger	Rulettes	Cartridge	Cartridge Version	Is Active
Agent Health State	Agent : agentID != "0"	Time Driven	2	Core-MonitoringPolicy	5.5.5	✓
AgentMsgs	AppMonitor_AgentMsgs	Data Driven	0	OS-AppMonitor	5.5.4	✓
BSM All Events		Event Driven	0	BSM-Integration	5.5.5	✗
Capacity	Windows_System_FileSys_Table	Data Driven	1	OS-Windows_System	5.5.4	✓
Catalyst Database Space Checking	CatalystDatabase	Data Driven	1	Core-MonitoringPolicy	5.5.5	✓
Catalyst Data Service Discarding Data	CatalystDataService	Time Driven	1	Core-MonitoringPolicy	5.5.5	✓
Catalyst Free Database Space Checking	CatalystTablespace	Data Driven	0	Core-MonitoringPolicy	5.5.5	✓
Clear Old LogFilter Alarms		Time Driven	1	Core-MonitoringPolicy	5.5.5	✓
CPU_Utilization	Windows_System_System_Table	Data Driven	1	OS-Windows_System	5.5.4	✓
Create SupportBundle Rule		Schedule Driven	0	Core-MonitoringPolicy	5.5.5	✗
Device_Not_Found	NetMonitor_ServiceLevels	Data Driven	0	OS-NetMonitor	5.5.4	✓
Device_Unavailable	NetMonitor_ServiceLevels	Data Driven	0	OS-NetMonitor	5.5.4	✓

Using Metric Calculations in Foglight

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.

Working with Derived Metrics

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, choose **Dashboards** > **Administration** > **Data** > **Manage Derived Metrics**.

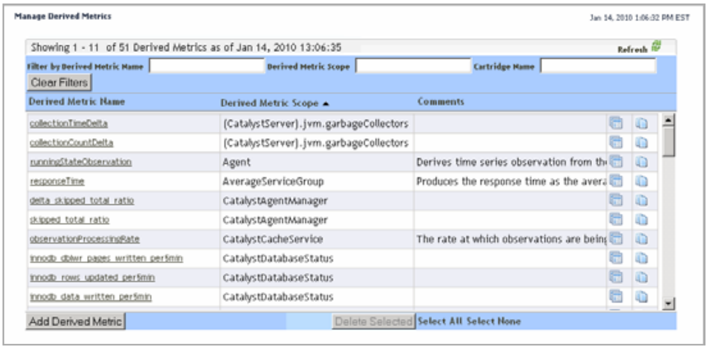

Manage Derived Metrics

Jan 14, 2010 1:06:32 PM EST

Showing 1 - 11 of 51 Derived Metrics as of Jan 14, 2010 13:06:35

Filter by Derived Metric Name: ______ Derived Metric Scope: ______ Cartridge Name: ______

Derived Metric Name	Derived Metric Scope ▲	Comments
collectionTimeDelta	(CatalystServer).jvm.garbageCollectors	
collectionCountDelta	(CatalystServer).jvm.garbageCollectors	
runningStateObservation	Agent	Derives time series observation from th
responseTime	AverageServiceGroup	Produces the response time as the avera
delta_skipped_total_ratio	CatalystAgentManager	
skipped_total_ratio	CatalystAgentManager	
observationProcessingRate	CatalystCacheService	The rate at which observations are being
innodb_dblwr_pages_written_per5min	CatalystDatabaseStatus	
innodb_rows_updated_per5min	CatalystDatabaseStatus	
innodb_data_written_per5min	CatalystDatabaseStatus	

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, from the navigation channel, choose **Dashboards > Administration > Data > Derived Metrics Diagnostics**.

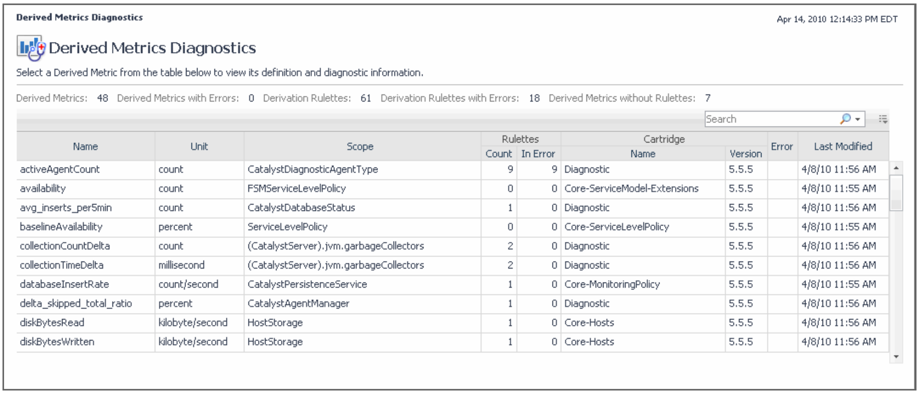

Derived Metrics Diagnostics

Apr 14, 2010 12:14:33 PM EDT

Derived Metrics Diagnostics

Select a Derived Metric from the table below to view its definition and diagnostic information.

Derived Metrics: 48 Derived Metrics with Errors: 0 Derivation Rulettes: 61 Derivation Rulettes with Errors: 18 Derived Metrics without Rulettes: 7

Name	Unit	Scope	Rulettes		Cartridge		Error	Last Modified
			Count	In Error	Name	Version		
activeAgentCount	count	CatalystDiagnosticAgentType	9	9	Diagnostic	5.5.5		4/8/10 11:56 AM
availability	count	FSMServiceLevelPolicy	0	0	Core-ServiceModel-Extensions	5.5.5		4/8/10 11:55 AM
avg_inserts_per5min	count	CatalystDatabaseStatus	1	0	Diagnostic	5.5.5		4/8/10 11:56 AM
baselineAvailability	percent	ServiceLevelPolicy	0	0	Core-ServiceLevelPolicy	5.5.5		4/8/10 11:55 AM
collectionCountDelta	count	(CatalystServer).jvm.garbageCollectors	2	0	Diagnostic	5.5.5		4/8/10 11:56 AM
collectionTimeDelta	millisecond	(CatalystServer).jvm.garbageCollectors	2	0	Diagnostic	5.5.5		4/8/10 11:56 AM
databaseInsertRate	count/second	CatalystPersistenceService	1	0	Core-MonitoringPolicy	5.5.5		4/8/10 11:55 AM
delta_skipped_total_ratio	percent	CatalystAgentManager	1	0	Diagnostic	5.5.5		4/8/10 11:56 AM
diskBytesRead	kilobyte/second	HostStorage	1	0	Core-Hosts	5.5.5		4/8/10 11:56 AM
diskBytesWritten	kilobyte/second	HostStorage	1	0	Core-Hosts	5.5.5		4/8/10 11:56 AM

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, choose **Dashboards > Administration > Data > Manage Thresholds**.

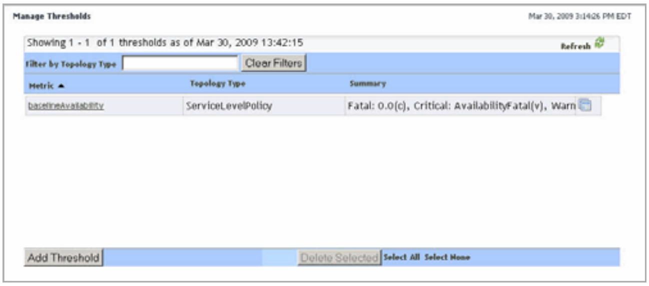

Showing 1 - 1 of 1 thresholds as of Mar 30, 2009 13:42:15

Metric ▲	Topology Type	Summary
baselineAvailability	ServiceLevelPolicy	Fatal: 0.0(c), Critical: AvailabilityFatal(v), Warn

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, choose **Dashboards > Administration > Rules & Notifications > Manage Registry Variables**.

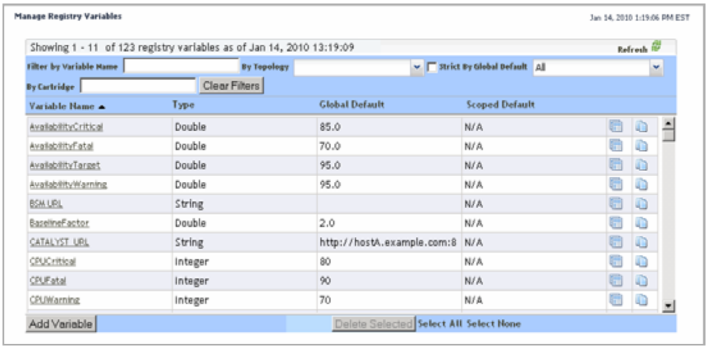

Manage Registry Variables

Jan 14, 2010 1:19:06 PM EST

Showing 1 - 11 of 123 registry variables as of Jan 14, 2010 13:19:09

Filter by Variable Name: By Topology: □ Strict By Global Default: All

By Cartridge:

Variable Name ▲	Type	Global Default	Scoped Default
AvailabilityCritical	Double	85.0	N/A
AvailabilityFatal	Double	70.0	N/A
AvailabilityTarget	Double	95.0	N/A
AvailabilityWarning	Double	95.0	N/A
BSM URL	String		N/A
BaselineFactor	Double	2.0	N/A
CATALYST_URL	String	http://hostA.example.com:8	N/A
CPUCritical	Integer	80	N/A
CPUFatal	Integer	90	N/A
CPUWarning	Integer	70	N/A

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, choose **Dashboards > Administration > Rules & Notifications > Check Registry Value**.

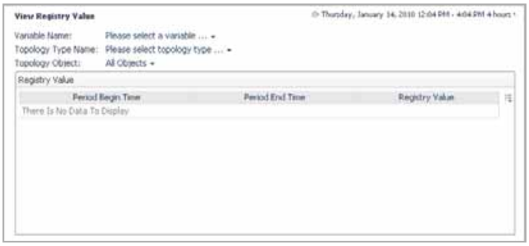

View Registry Value

Thursday, January 14, 2010 12:04 PM - 4:04 PM 4 hours

Variable Name: Please select a variable ...

Topology Type Name: Please select topology type ...

Topology Object: All Objects

Registry Value

Period Begin Time	Period End Time	Registry Value
There Is No Data To Display		

Associating 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, choose **Dashboards > Administration >Schedules > Manage Schedules**.

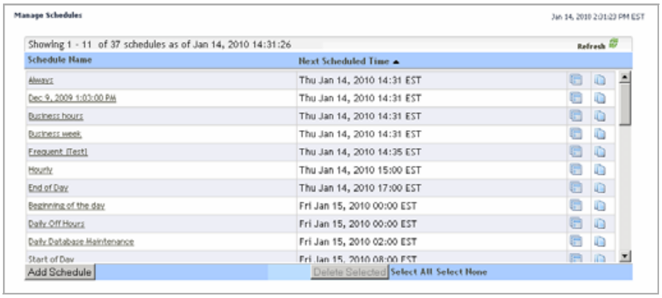

Showing 1 - 11 of 37 schedules as of Jan 14, 2010 14:31:26

Schedule Name	Next Scheduled Time ▲
Always	Thu Jan 14, 2010 14:31 EST
Dec 9, 2009 1:03:00 PM	Thu Jan 14, 2010 14:31 EST
Business hours	Thu Jan 14, 2010 14:31 EST
Business week	Thu Jan 14, 2010 14:31 EST
Frequent (Test)	Thu Jan 14, 2010 14:35 EST
Hourly	Thu Jan 14, 2010 15:00 EST
End of Day	Thu Jan 14, 2010 17:00 EST
Beginning of the day	Fri Jan 15, 2010 00:00 EST
Daily Off Hours	Fri Jan 15, 2010 00:00 EST
Daily Database Maintenance	Fri Jan 15, 2010 02:00 EST
Start of Day	Fri Jan 15, 2010 08:00 EST

Managing 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 Manage Retention Policies dashboard. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Data > Manage Retention Policies**.

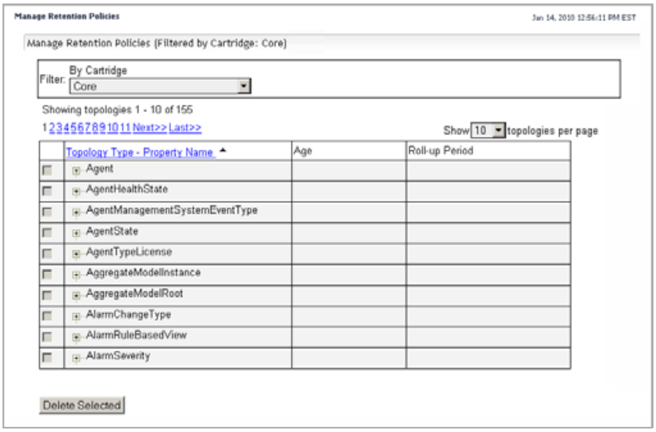

Manage Retention Policies

Jan 14, 2010 12:56:11 PM EST

Manage Retention Policies (Filtered by Cartridge: Core)

Filter: By Cartridge: Core

Showing topologies 1 - 10 of 155

1 2 3 4 5 6 7 8 9 10 11 Next>> Last>>

Show 10 topologies per page

	Topology Type - Property Name ▲	Age	Roll-up Period
□	⊞ Agent		
□	⊞ AgentHealthState		
□	⊞ AgentManagementSystemEventType		
□	⊞ AgentState		
□	⊞ AgentTypeLicense		
□	⊞ AggregateModelInstance		
□	⊞ AggregateModelRoot		
□	⊞ AlarmChangeType		
□	⊞ AlarmRuleBasedView		
□	⊞ AlarmSeverity		

Another way to control the retention policies is through the Data Management 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 “Managing Foglight Database Performance” on page 30.

Expanding 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, choose **Dashboards** > **Administration** > **Data** > **Add Topology Types**.

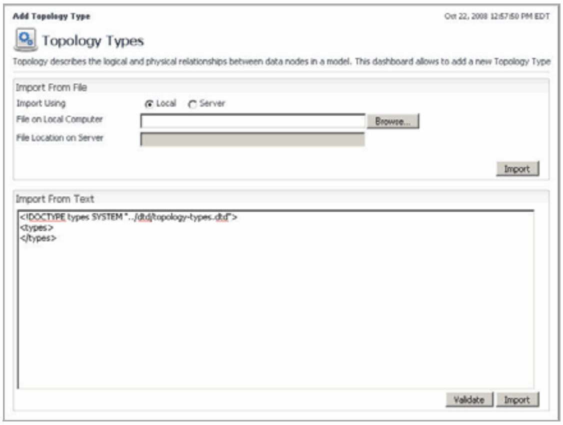

Add Topology Type Oct 22, 2008 12:57:50 PM EDT

Topology Types

Topology describes the logical and physical relationships between data nodes in a model. This dashboard allows to add a new Topology Type

Import From File

Import Using ● Local ○ Server

File on Local Computer

File Location on Server

Import From Text

```
<!DOCTYPE types SYSTEM "../dtd/topology-types.dtd">
<types>
</types>
```

You can explore your database schema using the Schema Browser. For details, see "Exploring the Data Model" on page 30.

5

Customizing Your Environment with Foglight Tooling

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

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 **Using Foglight > Administration and Configuration Help > Customizing Your Environment with Foglight 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 a Foglight 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 Foglight 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, choose **Dashboards > Administration > Tooling > Manage Host Aliasing Rules**.

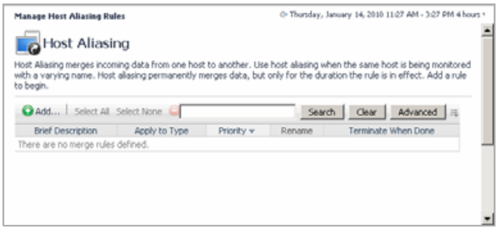

Manage Host Aliasing Rules

Thursday, January 14, 2010 11:27 AM - 3:27 PM 4 hours

Host Aliasing

Host Aliasing merges incoming data from one host to another. Use host aliasing when the same host is being monitored with a varying name. Host aliasing permanently merges data, but only for the duration the rule is in effect. Add a rule to begin.

Brief Description	Apply to Type	Priority ▾	Rename	Terminate When Done
There are no merge rules defined.				

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 such 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 Foglight 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 Foglight 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.bat* (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, choose **Dashboards > Administration > Tooling > Script Agent Builder**.

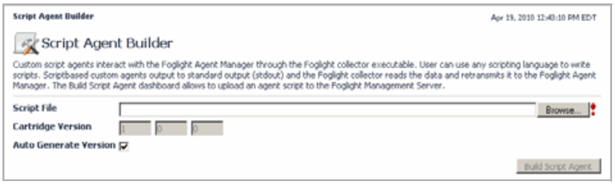

Script Agent Builder

Apr 19, 2010 12:43:10 PM EDT

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

Cartridge Version 1 0 0

Auto Generate Version ☑

Using the Query Language

The Query Service gives users and Foglight 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 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 Editor dashboard to test sample scripts. This tool has no access restrictions, but is recommended for advanced users. To access this dashboard, from the navigation panel, choose **Dashboards > Administration > Tooling > Script Editor**.

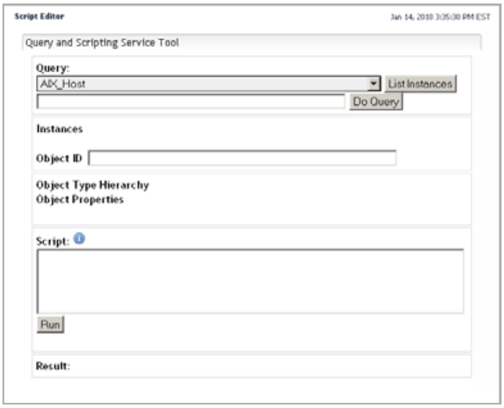

Script Editor

Jan 14, 2010 3:35:30 PM EST

Query and Scripting Service Tool

Query: AIX_Host

Index