

Foglight® 5.9.x

High Availability Field Guide



© 2018 Quest Software Inc.

ALL RIGHTS RESERVED.

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

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

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

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

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

Patents

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

Trademarks

Quest, the Quest logo, and Join the Innovation are trademarks and registered trademarks of Quest Software Inc. For a complete list of Quest marks, visit <https://www.quest.com/legal/trademark-information.aspx>. "Apache HTTP Server", Apache, "Apache Tomcat" and "Tomcat" are trademarks of the Apache Software Foundation. Google is a registered trademark of Google Inc. Android, Chrome, Google Play, and Nexus are trademarks of Google Inc. Red Hat, JBoss, the JBoss logo, and Red Hat Enterprise Linux are registered trademarks of Red Hat, Inc. in the U.S. and other countries. CentOS is a trademark of Red Hat, Inc. in the U.S. and other countries. Fedora and the Infinity design logo are trademarks of Red Hat, Inc. Microsoft, .NET, Active Directory, Internet Explorer, Hyper-V, Office 365, SharePoint, Silverlight, SQL Server, Visual Basic, Windows, Windows Vista and Windows Server are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries. AIX, IBM, PowerPC, PowerVM, and WebSphere are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Java, Oracle, Oracle Solaris, PeopleSoft, Siebel, Sun, WebLogic, and ZFS are trademarks or registered trademarks of Oracle and/or its affiliates in the United States and other countries. SPARC is a registered trademark of SPARC International, Inc. in the United States and other countries. Products bearing the SPARC trademarks are based on an architecture developed by Oracle Corporation. OpenLDAP is a registered trademark of the OpenLDAP Foundation. HP is a registered trademark that belongs to Hewlett-Packard Development Company, L.P. Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both. MySQL is a registered trademark of MySQL AB in the United States, the European Union and other countries. Novell and eDirectory are registered trademarks of Novell, Inc., in the United States and other countries. VMware, ESX, ESXi, vSphere, vCenter, vMotion, and vCloud Director are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Sybase is a registered trademark of Sybase, Inc. The X Window System and UNIX are registered trademarks of The Open Group. Mozilla and Firefox are registered trademarks of the Mozilla Foundation. "Eclipse", "Eclipse Foundation Member", "EclipseCon", "Eclipse Summit", "Built on Eclipse", "Eclipse Ready" "Eclipse Incubation", and "Eclipse Proposals" are trademarks of Eclipse Foundation, Inc. IOS is a registered trademark or trademark of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries. Apple, iPad, iPhone, Mac OS, Safari, Swift, and Xcode are trademarks of Apple Inc., registered in the U.S. and other countries. Ubuntu is a registered trademark of Canonical Ltd. Symantec and Veritas are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. OpenSUSE, SUSE, and YAST are registered trademarks of SUSE LCC in the United States and other countries. Citrix, AppFlow, NetScaler, XenApp, and XenDesktop are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. AlertSite and DéjàClick are either trademarks or registered trademarks of Boca Internet Technologies, Inc. Samsung, Galaxy S, and Galaxy Note are registered trademarks of Samsung Electronics America, Inc. and/or its related entities. MOTOROLA is a registered trademark of Motorola Trademark Holdings, LLC. The Trademark BlackBerry Bold is owned by Research In Motion Limited and is registered in the United States and may be pending or registered in other countries. Ixia and the Ixia four-petal logo are registered trademarks or trademarks of Ixia. Opera, Opera Mini, and the O logo are trademarks of Opera Software ASA. Tevron, the Tevron logo, and CitraTest are registered trademarks of Tevron, LLC. PostgreSQL is a registered trademark of the PostgreSQL Global Development Group. MariaDB is a trademark or registered trademark of MariaDB Corporation Ab in the European Union and United States of America and/or other countries. Vormetric is a registered trademark of Vormetric, Inc. Intel, Itanium, Pentium, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. Debian is a registered trademark of Software in the Public Interest, Inc. OpenStack is a trademark of the OpenStack Foundation. Amazon Web Services, the "Powered by Amazon Web Services" logo, and "Amazon RDS" are trademarks of Amazon.com, Inc. or its affiliates in the United States and/or other countries. Infobright, Infobright Community Edition and Infobright Enterprise Edition are trademarks of Infobright Inc. POLYCOM®, RealPresence® Collaboration Server, and RMX® are registered trademarks of Polycom, Inc. All other trademarks and registered trademarks are

property of their respective owners.

Legend

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

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

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

Foglight High Availability Field Guide
Updated - June 2018
Software Version - 5.9.x

High Availability mode overview	5
Capabilities	5
Use cases	6
Known issues	6
Requirements	6
Managing High Availability mode	7
Enabling High Availability mode	7
Starting and stopping the server in High Availability mode	9
Redirecting secondary Management Servers to the correct URL	10
Remote Monitor	11
Enabling email notification for High Availability	12
Starting and stopping the HA server from the command line	12
fmsha	12
remotemonitor	14
Configuring High Availability mode ports	16
Configuring the cluster multicast port used in HA mode	17
Important considerations for configuring HA over WAN	18
Security considerations for HA configuration	18
Exploring the High Availability view	18
Configuring Foglight Agent Manager for HA mode	19
Upgrading cartridges in a High Availability environment	19
Troubleshooting FAQ	19
Common configuration scenarios	21
Running on Oracle RAC	21
Running behind an HTTP proxy	22
Running in a cluster	22
Tuning connection issues in HA implementations	24
Tuning the Management Server	24
Managing hosts with multiple network interfaces	25
JDK with IPv6 on Linux	26
Foglight Management Server automatically restarted	26
Other JGroup-related issues	27
About Us	28
We are more than just a name	28
Our brand, our vision. Together.	28
Contacting Quest	28
Technical support resources	28

High Availability mode overview

Welcome to the Foglight® *High Availability Field Guide*. This Field Guide provides step-by-step instructions on how to configure your Foglight 5 environment to use High Availability (HA) Mode. Use it to enable, configure, and tune Foglight Management Server(s) in HA Mode.

This guide is intended for technical audiences and Quest SC/PSO (System Consultant/Professional Services Organization) personnel working in the field.

High Availability mode is a configuration of multiple Foglight Management Servers in a cluster where one Management Server is the primary server and the others are standby servers. If the primary server stops responding, one of the standby servers in the HA cluster takes over monitoring responsibility. HA mode uses redundancy and failover to ensure that Foglight monitoring continues uninterrupted.

High Availability mode provides different functionality than Federation. Federation is a Foglight feature that addresses the needs of customers who monitor large-scale environments that are naturally partitioned into logical units. Each partition is traditionally served by one Management Server instance and its distributed clients/agents. For more information, see the *Federation Field Guide*.

NOTE: The vFoglight Cartridge for Automation is not supported with Foglight when Foglight is configured to use the High Availability feature. This limitation will be addressed in a future release. For more information, see the *Foglight Release Notes*.

- [Capabilities](#)
- [Use cases](#)
- [Known issues](#)
- [Requirements](#)

Capabilities

Configuring multiple servers using HA mode offers the following capabilities:

- **Automatic restart:** The restart process (FoglightHA Daemon) runs on the same host that the Foglight® Management Server it is monitoring and monitors the state of the server. If the server stops responding or has a fatal error, then the restart process stops and restarts the server automatically. The process can also send out email notifications about the server condition.
 - **NOTE:** The default value for the restart process is ten minutes (600 seconds). Machines with a slower processing speed take longer to reboot. You can modify the default time by editing the `startup.grace` variable in the `<foglight_home>/config/restart_monitor.config` file.
- **Redundancy:** To reduce outage times, a standby server instance is created (for a particular Management Server). One server becomes the primary server and communicates with the agents, receives and processes data, and processes rules. All other servers are standby servers. If a primary server stops responding, one of the standby servers resumes service and becomes the primary server.

Use cases

High Availability is a Foglight® feature that addresses the needs of customers to ensure that a Management Server is always available, even in the event of a primary server failure.

HA mode has been tested in the following [Common configuration scenarios](#):

- [Running on Oracle RAC](#)
- [Running behind an HTTP proxy](#)
- [Running in a cluster](#)

Known issues

Foglight® has the following known issues with respect to High Availability mode:

- The Foglight HA implementation uses UDP for communication by default. It is recommended that you reconfigure Foglight to use TCP instead. For more information, see [Tuning connection issues in HA implementations](#) on page 24.
- There is a JDK issue with IPv6 on Linux® that affects both HA and standalone server. For more information, see [JDK with IPv6 on Linux](#) on page 26.
- Running a Federation Master in High Availability mode is not supported. Only Federated Children can be run in High Availability mode.

Requirements

If you plan to run the Foglight® Management Server in HA mode, there are certain considerations:

- A server running in HA mode can only use an external database. An external instance of MySQL with an InnoDB storage engine can be used if your database administrator has installed one. Please see the *Foglight System Requirements and Platform Support Guide* for information about the external databases supported for Foglight.

If you have been using a Management Server with the embedded database, refer to the Appendix: “Switching from an Embedded to an External Database” in the *Foglight Installation and Setup Guide*.

- The database must be running for Foglight to operate. You may have to implement a combined database-backup and High Availability solution to achieve true redundancy of all components.
- Specific ports are required for the HA setup to work, all of which must be set to the same values for all Management Servers in the HA cluster. For more information, see [Configuring High Availability mode ports](#) on page 16.
- By default, HA mode works only within the local subnet and uses UDP (User Datagram Protocol) for communication. You can reconfigure your HA cluster to use TCP (Transmission Control Protocol), which can work outside of the local subnet; this is the recommended alternative if your HA cluster consists of a small number of designated HA Management Servers.

Configuring HA outside of a subnet (across a WAN) introduces some potential performance issues. Review the information in [Important considerations for configuring HA over WAN](#) on page 18 before you configure the HA cluster.

- To use the Remote Monitor, JRE 1.5 or later must be installed on the machine hosting the Remote Monitor, and the `JAVA_HOME` environment variable must be set to point to the installation directory. For more information, see [Remote Monitor](#) on page 11.

Managing High Availability mode

High Availability (HA) mode enables you to install multiple Foglight® Management Servers, where one server is the primary server and the others are standby servers. If the primary server stops responding, then the responsibility is taken over by one of the standby servers.

This chapter describes the common management tasks for HA servers.

For information on tuning an HA cluster for better performance, see [Tuning connection issues in HA implementations](#) on page 24.

- [Enabling High Availability mode](#)
- [Starting and stopping the server in High Availability mode](#)
- [Redirecting secondary Management Servers to the correct URL](#)
- [Starting and stopping the HA server from the command line](#)
- [Configuring High Availability mode ports](#)
- [Important considerations for configuring HA over WAN](#)
- [Exploring the High Availability view](#)
- [Configuring Foglight Agent Manager for HA mode](#)
- [Upgrading cartridges in a High Availability environment](#)
- [Troubleshooting FAQ](#)

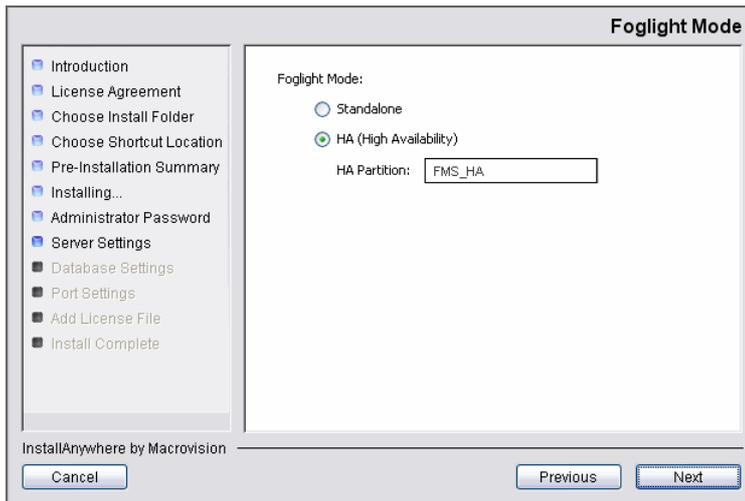
Enabling High Availability mode

You can enable High Availability mode during a Foglight® Management Server installation, or afterwards by editing the `server.config` file.

Enabling During Installation

To enable HA mode during installation of a new Management Server:

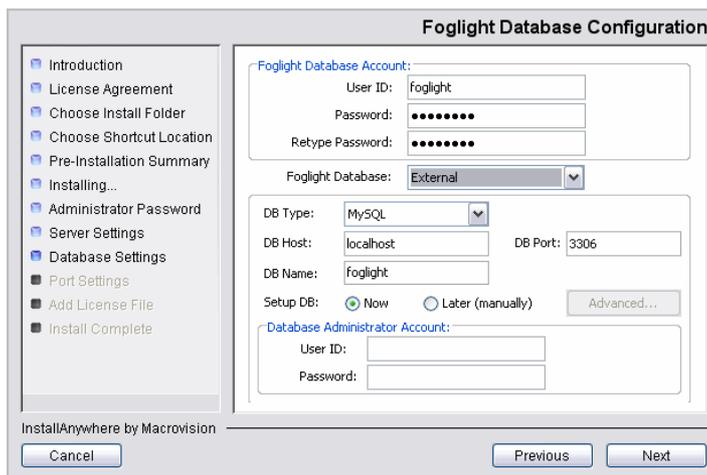
- 1 On the Foglight Mode page of the installer, select **HA** (High Availability) as the server startup mode.



- 2 In the **HA Partition** field, type the partition name.

i | IMPORTANT: All Management Servers that you want to act as either primary or secondary servers within the same HA cluster must be configured with the same partition name.

- 3 You may also need to configure additional database settings in the Foglight Database Configuration step of the Management Server installer.



Configure the server to connect to the same database instance as all other members of the cluster, by setting the **DB Type**, **DB Host**, **DB Port or Instance** and **DB Name** to the same values as those used for the other members.

i | NOTE: In some rare cases, certain settings might be different. For example, the value for DB Host might be different if there is a firewall between the database and certain members of the cluster in your environment.

You must also set the Database Administrator Account **User ID** and **Password** to the same values as those used for the other Management Servers that share the same HA partition name.

i | IMPORTANT: The Database Administrator Account User ID cannot be the same as the Foglight Database Account User ID.

Enabling HA mode after installation

To enable HA mode after installation:

- 1 Open the file `<foglight_home>/config/server.config` on each Management Server that you want to participate in the HA cluster.
- 2 Locate and uncomment the parameter `server.ha.partition.name`.
- 3 Set the argument of `server.ha.partition.name` to the partition name for your HA cluster.

i | **IMPORTANT:** All Management Servers that you want to act as either primary or secondary servers within the same HA cluster must be configured with the same partition name.

- 4 Configure the Foglight Management Server to connect to the same database instance as all other members of the HA cluster, by setting the parameters listed below to the same values as the other members:

i | **IMPORTANT:** In some rare cases, certain parameters might be different. For example, the value for `server.database.host` might be different if there is a firewall between the database and certain members of the cluster in your environment.

- a `server.database.host`: Set the argument of this parameter to the Foglight database host name (or IP address).
- b `server.database.port`: Set the argument to the Foglight database port number.
- c `server.database.name`: Set the argument to the Foglight database name.
- d `server.database.user`: Set the argument to the Foglight database user name.
- e `server.database.password`: Set the argument to the Foglight database user password.
- f `server.database.type`: Set the argument to the type of database used (`mysql`, `oracle`, or `sqlsvr`).

- 5 Specific ports are required for the HA setup to work. These are listed in [Configuring High Availability mode ports](#) on page 16.

Ensure that all Management Servers in the HA cluster have these ports set to the same values in `<foglight_home>/config/server.config`.

i | **NOTE:** When running in HA mode, the Foglight Management Server may use ports in addition to the ones that you can set using `server.config`. For more information, see [Configuring High Availability mode ports](#) on page 16.

- 6 Restart the Foglight Management Server in HA mode by navigating to `<foglight_home>/bin` and running the following command:

```
fmsha
```

i | **NOTE:** For more information about command-line options, see [Starting and stopping the HA server from the command line](#) on page 12.

Starting and stopping the server in High Availability mode

To run the Foglight® Management Server in HA mode, you must use an external database, which must be running before you start the server.

To start the server in HA mode:

Do one of the following:

- (Windows only) Choose **Start > Programs > Quest > Foglight > High Availability > Start Foglight In HA Mode**.
- Enter the following from the command-line:

```
fmsha
```

i | **NOTE:** For more information about command-line options, see [Starting and stopping the HA server from the command line](#) on page 12.

After the server starts successfully, the following message appears in the command shell or window:

```
Forge Server startup completed.
```

To stop the High Availability Server:

Do either of the following:

- Type **Ctrl-C** in the command shell in which the Foglight HA server started.
- (Windows only) Choose **Start > Programs > Quest > Foglight > High Availability > Stop Foglight In HA Mode**.
- Open a second command shell, navigate to the directory `<foglight_home>/bin`, and execute the following command:

```
fmsha --stop
```

After the server stops successfully, the command shell closes.

i | **NOTE:** For more information about command-line options, see [Starting and stopping the HA server from the command line](#) on page 12.

Redirecting secondary Management Servers to the correct URL

When Foglight® Management Servers are configured as an HA cluster, the browser interface for each of the secondary Management Servers redirects to the primary Management Server. By default, the URL for the browser interface to which the secondary Management Servers are redirected is generated automatically by Foglight based on the primary Management Server machine's host name.

You can override this URL and point Foglight to a different URL by starting the secondary Management Servers from the command-line with the following option:

```
-Dquest.host.name=<host_name>
```

where `<host_name>` is the desired host name or Fully Qualified Domain Name (FQDN).

Use this option in cases where a single machine hosts multiple applications (including Foglight) and is configured with multiple host names and IP addresses. In such cases:

- Each Web server is bound to a single IP address and responds to a unique FQDN or alias.
- The default URL used by Foglight to redirect the secondary Management Servers is based on the host name of the primary Management Server machine and may be bound to a different alias. This causes so might not present the correct URL since Foglight

You can base the URL on the alias using the `-D` option.

Example:

You have two machines, whose host names are *server1* and *server2*. You have an HA cluster running on these machines. The primary Management Server runs on *server1* and the secondary Management Server runs on *server2*.

Both machines have two IP addresses: one internal and one external. The external IP addresses are bound to the DNS (Domain Name System) names *foglight1.example.com* and *foglight2.example.com*. Foglight has been bound to the external IP addresses and is not listening on the internal IP addresses. In this case, the Foglight browser interface is available at `http://foglight1.example.com:8080`.

However, the secondary server redirects to the default URL `http://server1:8080`. In this case, the Foglight browser interface is not available at this URL. To resolve this issue, override the URL with: –
`Dquest.host.name=foglight1.example.com`.

Remote Monitor

The Remote Monitor is a Java™ program that regularly communicates with a set of High Availability (HA) servers. If there is no reply after it attempts to contact the server, then it logs the event and (optionally) sends e-mail notifications to an administrator.

The Remote Monitor takes a list of host names or host name and port pairs as parameters. When invoked without parameters, Remote Monitor takes the list of host names or host name and port pairs from the file `<foglight_home>/config/remote_monitor.config` under the entry `server.urls`.

Each parameter in the entry `server.urls` points to a specific HA server. If a port is not specified as part of this entry, the default port 51231 is used. The default port can be configured in `<foglight_home>/config/remote_monitor.config`, under the entry `health.monitor.port`. This entry should match the `health.monitor.port` entry in the file `<foglight_home>/config/restart_monitor.config` on your servers.

NOTE: To use the Remote Monitor, JRE 1.5 or later must be installed on the machine hosting the Remote Monitor and the `JAVA_HOME` environment variable must be set to point to the installation directory.

To install the Remote Monitor:

- Copy the `<foglight_home>/tools/remote_monitor.zip` file to your target machine and unzip the contents to a directory.

To start the Remote Monitor:

NOTE: The Foglight Management Server must be running before you can start the Remote Monitor.

On UNIX®, enter the command:

- `<foglight_home>/bin/remotemonitor <hostname[:port]>>`

On Windows®, do one of the following:

- Enter the command: `<foglight_home>\bin\remotemonitor.exe`
- Choose **Start > Programs > Quest > Foglight > High Availability > Start Remote Monitor**.

To stop the Remote Monitor:

Do one of the following:

- Enter **Ctrl-C** from the command shell or window where the Remote Monitor started.
- On UNIX®, enter the command: `<foglight_home>/bin/remotemonitor <hostname[:port]>> -q`

On Windows®, choose **Start > Programs > Quest > Foglight > High Availability > Stop Remote Monitor**

Or

Enter the command: <foglight_home>\bin\remotemonitor.exe -q

Enabling email notification for High Availability

Servers running in High Availability (HA) mode can send out email notifications in various situations, such as when servers stop responding or on failure to restart a server. The files <foglight_home>/config/restart_monitor.config and <foglight_home>/config/remote_monitor.config control how email notification functions, such as email protocols, email server information, and the message and recipients for each situation. Email can be sent out from the HA server or from the Remote Monitor. See the comments in these configuration files for more information about email notification.

Starting and stopping the HA server from the command line

fmsha

The `fmsha` command provides a command-line interface to the Management Server process running in High Availability (HA) mode. Running Foglight® in HA mode allows you to manage multiple instances of the Management Server in a JGroup partition that supports the HA feature.

This command offers a set of options that you can use to perform any of the following operations as required:

- Start or stop the Management Server in HA mode
- Install and start the Management Server in HA mode as a Windows® service
- Stop and remove a Management Server HA Windows service
- Configure Java™ Virtual Machine (JVM) options and add entries to the Foglight classpath
- Assign different names to different Management Server process launchers
- Display version information or a list of arguments and their descriptions

Syntax

Windows only

```
fmsha [-s|--start] [-q|--stop] [-w|--wait] [-n|--name process_name]
      [-i|--install-service] [-r|--remove-service] [-b|--start-service]
      [-j|--jvm-argument JVM_options] [-p|--classpath class_path] [-v|--version]
      [-h|--help] [-t|--thread-dump] [-m|--javavm path_to_JAVA_HOME]
      [-D|-X JVM_option] [--set-global-debug-level debug_level]
      [--add-debug-level package_1.=debug_level_1 ... package_n.=debug_level_n]
      [--remove-debug-level package_1. ... package_n.]
```

Unix only

```
fmsha [-d|--daemon] [-s|--start] [-q|--stop] [-w|--wait]
      [-n|--name process_name] [-j|--jvm-argument JVM_options]
      [-p|--classpath class_path] [-v|--version] [-h|--help] [-t|--thread-dump]
      [-m|--javavm path_to_JAVA_HOME] [-D|-X JVM_option]
```

```
[--set-global-debug-level debug_level]
[--add-debug-level package_1.=debug_level_1 ... package_n.=debug_level_n]
[--remove-debug-level package_1. ... package_n.]
```

i | **NOTE:** If you do not specify any options, `fmsha` uses the default option, `s`, and starts an instance of the Management Server in HA mode.

Table 1. Options and arguments

Options	Argument	Description
UNIX® and Windows®		
add-debug-level	<code>package.=debug_level</code>	Specifies a debug level for one or more packages. Higher debug levels result in more detailed logging. For example, you can set the <code>debug_level</code> to one of the following values: 0: No debugging 1: Minimal debugging 2: Detailed debugging The period '.' following the package name is mandatory, otherwise the package is treated as a class. The <code>package.=debug_level</code> argument can be specified multiple times.
D X	<code>JVM_option</code>	Passes an option to the JVM.
h help	None	Displays a list of arguments and their descriptions.
j jvm-argument	<code>JVM_options</code>	Specifies one or more JVM options.
m javavm	<code>path_to_JAVA_HOME</code>	Points to the JVM for the Management Server process.
n name	<code>process_name</code>	Specifies a unique process name for the current instance of the Management Server. Foglight uses process names to distinguish between different instances of the same process launcher.
p classpath	<code>class_path</code>	Adds entries to the JVM classpath.
q stop	None	Stops the running Management Server process.
remove-debug-level	<code>package.</code>	Removes a debug level associated with one or more packages. The period '.' following the package name is mandatory, otherwise the package is treated as a class. The <code>package.</code> argument can be specified multiple times.
s start	None	Starts the Management Server.
set-global-debug-level	<code>debug_level</code>	Sets the global debug level. The <code>debug_level</code> argument must be a non-negative integer. Higher debug levels result in more detailed logging. For example, you can set the <code>debug_level</code> to one of the following values: <ul style="list-style-type: none"> 0: No debugging 1: Minimal debugging 2: Detailed debugging
t thread-dump	None	Requests a thread output from the running application. This option writes the output to a separate log file in the application's installation directory.
v version	None	Displays the version number, copyright, build number, and the installation directory.
w wait	None	When sending a shutdown command to an existing Management Server process, this option instructs the command to wait indefinitely for the process to exit before shutting it down.

Unix only

Table 1. Options and arguments

Options	Argument	Description	
d	daemon	None	Starts the Management Server as a daemon process.
<p>NOTE: Unix installations also include the following scripts for starting and stopping the Management Server in HA mode:</p> <ul style="list-style-type: none"> • <code>fmsStartupHA.sh</code> • <code>fmsShutdownHA.sh</code> 			
Windows only			
b	start-service	None	Starts the Management Server Windows service.
i	install-service	None	Installs the Management Server as a Windows service.
r	remove-service	None	Stops and removes the Management Server Windows service.

Examples

Starting the server in HA mode

```
C:\Quest\Foglight\bin>fmsha
2017-07-03 17:04:33.000 INFO [native] Inter-launcher communications channel
active at:
C:\Quest\Foglight\bin\fmsha.exe\.FoglightHA-ZHUVN-FOG-2769.msg
2017-07-03 17:04:37.612 INFO Starting Forge Server with the command bin\fms
-Dfoglight.cluster.mode=true -Dquest.common.process-runner=false -Dquest.
native.launcher.io=true...
```

Installing the server in HA mode as a Windows service

```
C:\Quest\Foglight\bin>fmsha -i
2017-07-03 17:02:36.000 INFO [native] High Availability Foglight (FoglightHA)
service installed
```

Removing the server HA Windows service

```
C:\Quest\Foglight\bin>fmsha -r
2017-07-03 17:03:38.000 INFO [native] Removed the High Availability Foglight
(FoglightHA) service installed from 'C:\Quest\Foglight'
```

remotemonitor

The `remotemonitor` command provides command-line interface to the Remote Monitor application. The Remote Monitor communicates with multiple instances of the Management Server running in HA mode. If a server fails to reply, the Remote Monitor logs an event and sends e-mails to the server administrator. The Remote Monitor uses a list of host names or host names and port numbers to identify High Availability servers that it communicates with. This information is stored in the `<foglight_home>/config/remote_monitor.config` file under the `server.urls` entry.

For information on how to install and configure the Remote Monitor, see [Remote Monitor](#) on page 11; for additional information about the Remote Monitor application, see the *Foglight Installation and Setup Guide*.

The `remotemonitor` command offers a set of options that you can use to perform any of the following operations as required:

- Start or stop the Remote Monitor
- Install and start the Remote Monitor as a Windows® service
- Stop and remove the Remote Monitor Windows service

- Configure Java™ Virtual Machine (JVM) options and add entries to the Remote Monitor classpath
- Assign different names to different Remote Monitor process launchers
- Display version information or a list of arguments and their descriptions

Syntax

Windows only

```
remotemonitor [-s|--start] [-q|--stop] [-w|--wait] [-n|--name process_name]
  [-i|--install-service] [-r|--remove-service] [-b|--start-service]
  [-j|--jvm-argument JVM_options] [-p|--classpath class_path] [-v|--version]
  [-h|--help] [-t|--thread-dump]
```

Unix only

```
remotemonitor [-s|--start] [-q|--stop] [-w|--wait]
  [-n|--name process_name] [-j|--jvm-argument JVM_options]
  [-p|--classpath class_path] [-v|--version] [-h|--help]
  [-t|--thread-dump]
```

i | **NOTE:** If you do not specify any options, remotemonitor uses the default option, s, and starts an instance of the Remote Monitor utility.

Table 2. remotemonitor options and arguments

Options	Argument	Description	
UNIX® and Windows®			
h	help	None	Displays a list of arguments and their descriptions.
j	jvm-argument	<i>JVM_options</i>	Specifies one or more Java Virtual Machine (JVM) options.
n	name	<i>process_name</i>	Specifies a unique process name for the current instance of the Management Server. Foglight uses process names to distinguish between different instances of the same process launcher.
p	classpath	<i>class_path</i>	Adds entries to the JVM classpath.
q	stop	None	Stops the running Management Server process.
s	start	None	Starts the Management Server.
t	thread-dump	None	Requests a thread output from the running application. This option writes the output to a separate log file in the application's installation directory.
v	version	None	Displays the version number, copyright, build number, and installation directory.
w	wait	None	When sending a shutdown command to an existing Management Server process, this option instructs the command to wait indefinitely for the process to exit before shutting it down.
Windows only			
b	start-service	None	Starts the Management Server Windows service.
i	install-service	None	Installs the Management Server as a Windows service.
r	remove-service	None	Stops and removes the Management Server Windows service.

Examples

Installing Remote Monitor as a Windows service

```
C:\Quest\Foglight\bin>remotemonitor --install-service
2017-07-03 17:16:13.000 INFO [native] Foglight Remote Monitor (RemoteMonitor)
service installed
```

Removing the Remote Monitor Windows service

```
C:\Quest\Foglight\bin>remotemonitor --remove-service
2017-07-03 17:17:24.000 INFO [native] Removed the Foglight Remote Monitor
(RemoteMonitor) service installed from 'C:\Quest\Foglight'
```

Displaying Remote Monitor version information

```
C:\Quest\Foglight\bin>remotemonitor --version
Foglight RemoteMonitor 5.9.x
Copyright (c) 2017 Quest Software Inc.
Build Number: 5.9.1-201706300443-eebb0247-166 (64-bit)
Client Binary Directory: C:\Quest\Foglight
```

Configuring High Availability mode ports

The following table shows the default port assignments that are used in HA mode. The port numbers can be specified at installation time, or after the installation using the configuration parameters in the file `<foglight_home>/config/server.config`.

i **NOTE:** In addition to the ports listed below, in some configurations, the Foglight® Management Server may use additional ports, such as the ports used by JGroups which is the underpinning component of Foglight HA. JGroups is a toolkit for reliable multicast communication. In a default configuration, it dynamically allocates some UDP ports for the communication between cluster members. Network traffic from those ports does not go beyond the local subnet. Therefore, no special firewall configuration is needed.

Table 3. High Availability mode ports

Port	Server/Client Communication				
	Default Number	Name	Listens On	Used By	Communication Direction
45566	Cluster Mcast Port	Management Server	Management Server	Management Server	Bidirectional between the HA Management Server Primary and Secondary servers.
Description: Cluster multicast port, used when Foglight is running in HA mode.					
Configuration parameter: <code>server.cluster.mcast_port</code>					

Table 3. High Availability mode ports

Port	Server/Client Communication			
	Default Number	Name	Listens On	Used By
7800	TCP port	Management Server	Management Server	Bidirectional between the HA Management Server Primary and Secondary servers.

Description: Port for TCP communication between two Management Servers when Foglight is running in HA mode.

NOTE: The TCP port number is not configured in the `<foglight_home>/config/server.config` file, like the rest of the ports described in this table. It is specified in the following line of the `<foglight_home>/config/jgroups-config.xml` file:

```
<TCPPING initial_hosts="{jgroups.bind_addr} [7800],otherhost[7800]" ...
```

To enable TCP communication, you must first uncomment the appropriate lines of code in the `jgroups-config.xml` file, as instructed in [Tuning connection issues in HA implementations](#) on page 24.

In some cases, the system may use a different port number for TCP communication, for example, if the port 7800 is unavailable due to a conflict with another application. To find out the actual TCP port number, look through the Management server logs. For example, the following log entry indicates that the current TCP port number is 3307:

```
2014-04-07 16:53:01.125 VERBOSE [main] STDOUT -
-----
GMS: address is 127.0.0.1:3307 (cluster=my-cluster-name)
-----
```

Configuration parameter: N/A

Configuring the cluster multicast port used in HA mode

You can change the cluster multicast port that is used when Foglight® is running in High Availability (HA) mode. This port is used by members of an HA cluster to communicate with each other. The primary and secondary Management Servers within each HA partition must be configured to use the same cluster multicast port.

If you plan to have multiple HA partitions on the same sub-network, configure a different cluster multicast port for each HA partition on the sub-network.

To change the cluster multicast port:

- 1 Stop the primary and secondary Management Servers that are part of the HA cluster.
- 2 Open the file `<foglight_home>/config/server.config` on each machine hosting a Management Server that is part of the HA cluster.
- 3 Set the parameter `server.cluster.mcast_port` to the desired value in these files.
 - i** | **IMPORTANT:** Ensure that the primary server and secondary servers within each HA partition are configured to use the same cluster multicast port.
- 4 Save the `server.config` files.
- 5 Restart the Management Servers.

! CAUTION: Errors can occur if servers that are not part of the same HA partition are configured to multicast on the same port on the same sub-network as the HA partition. For more information, see [Troubleshooting FAQ](#) on page 19.

Important considerations for configuring HA over WAN

Although HA outside of a single subnet (that is, across a WAN) is a supported configuration, it is important to first understand the performance and reliability implications of using this configuration.

Before distributing a Foglight® HA installation across data centers, consider the following:

- WAN networks typically have lower bandwidth than LAN networks. Low bandwidth and network congestion can lead to false failovers.
- If the Foglight database is not distributed across data centers, one (or more) of the HA nodes must access the shared database remotely. This may cause poor performance.
- If the Foglight database is not distributed across data centers, an outage in the data center that hosts the shared database results in total failure of the cluster.

For more information, see [Troubleshooting FAQ](#) on page 19, and [Tuning connection issues in HA implementations](#) on page 24.

Security considerations for HA configuration

JGroups provides the *ENCRYPT* class with the capability of encryption and decryption communications at any layer, without the need for a coordinator. For more information about the JGroups *ENCRYPT* class, visit the [JGroups](#) website. The Foglight Management Server adopts the option of configuring a secretKey in a keystore to secure the HA mode. This security implementation is specified by the following lines in the `<foglight_home>/config/jgroups-config.xml` file.

```
<config ....>
  <UDP.../>
  <!-- encryption -->
  <ENCRYPT
    encrypt_entire_message="true"
    key_store_name="foglight.ha.keystore"
    store_password="foglightha"
    key_password="foglightha"
    alias="foglightha"
  />
</config>
```

Exploring the High Availability view

The High Availability View, located on the Management Server Administration dashboard, indicates if Foglight® is running in High Availability (HA) mode.



This view displays the following information:

- **Status** shows if Foglight is running in HA mode (●) or not (—).
- **Peers** shows the number of servers in the cluster, not including the current server.

Configuring Foglight Agent Manager for HA mode

The Foglight® Agent Manager client fully supports communicating with a Management Server HA cluster. During Agent Manager installation, you can configure multiple HA URLs for the upstream Management Server. In addition, if you connect to the active Management Server in an HA cluster, the Agent Manager installer can automatically query for and configure the other HA peer URLs. This ensures that, if the active (primary) Management Server fails over to an HA peer (secondary server), the Agent Manager clients will connect to the new active peer seamlessly.

For more information about installing and configuring the Agent Manager, see the *Foglight Agent Manager Installation Guide*.

Upgrading cartridges in a High Availability environment

To upgrade cartridges in an HA environment, you only need to install the cartridge you want to upgrade on the primary server while both the primary and the secondary Management Servers are running. The cartridges are then copied over to the secondary server automatically.

The management log reports on the copying of the cartridges.

Troubleshooting FAQ

I am running Foglight® in High Availability (HA) mode. When I start my Management Servers, “UDP mcast receiver” errors appear in the logs for the primary and secondary Management Servers. How can I resolve these errors?

If Management Servers that are not part of a particular HA partition are configured to multicast using the same port on the same subnet as those within the HA partition, errors like the following appear in the logs for the Management Servers that are configured in this way:

```
ERROR [UDP mcast receiver] org.jgroups.protocols.UDP - message does not have a UDP header
```

or

```
ERROR [UDP mcast receiver] org.jgroups.protocols.UDP - discarded message from different group
```

If you see this type of message, change the cluster multicast port of the servers that are reporting the error. In addition, if you plan to have multiple High Availability (HA) partitions on the same subnet, configure a different cluster multicast port for each HA partition. See [Configuring the cluster multicast port used in HA mode](#) on page 17 for instructions.

i | **IMPORTANT:** In both cases described above (resolving errors and planning to have multiple HA partitions on the same subnet), the primary server and secondary servers within each HA partition must be configured to use the same cluster multicast port.

Why does a Management Server running in HA mode restart when it is configured to use SSL?

Each Management Server in an HA cluster checks its own internal `health.monitor.port` to try and determine the server health. If this check fails, the server is automatically restarted. The HA Health Check Monitor does not use HTTPS, so when the Management Server is configured to use HTTPS/SSL, the port health check fails and the server is restarted.

To prevent this issue, configure the Health Check URL in the `restart_monitor.config` to use HTTP and port 8080.

By default, the URL is:

```
health.check.url="https://localhost:8443/foglight-sl/HealthCheck";
```

Specify HTTP as the transport protocol, and define the actual Management Server host name and port 8080. For example:

```
health.check.url="http://hostname:8080/foglight-sl/HealthCheck";
```

Ensure that the URL is updated on each server in the cluster.

Common configuration scenarios

This section outlines three common configuration scenarios for HA mode. It is intended only to illustrate possible configurations and processes and does not imply that Quest favors or recommends any of the platforms, technologies, or procedures described.

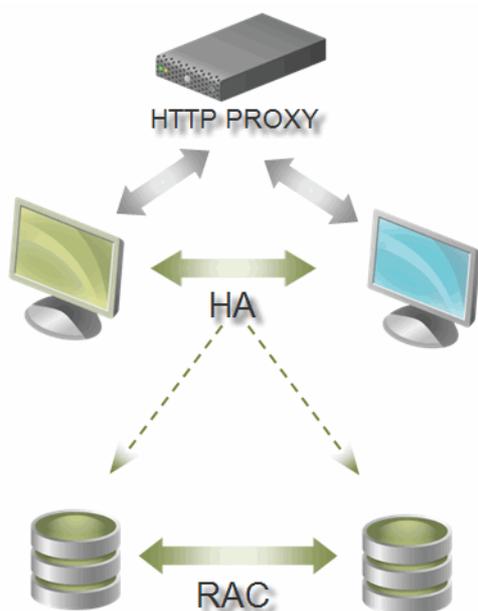
This information is intended for anyone taking the role of architect or installer for a Foglight® 5.x project.

- [Running on Oracle RAC](#)
- [Running in a cluster](#)

Running on Oracle RAC

In this configuration, the Management Server was installed on two distinct physical hosts, one acting as HA mode primary, and the other as secondary. This setup was run over an Oracle® RAC cluster to provide HA at the database layer. In addition, an HTTP proxy was configured to provide end users with a single URL for access to the Foglight® browser interface.

Figure 1. Oracle RAC configuration



i | **NOTE:** For more information about Oracle RAC, see:
<http://www.oracle.com/technetwork/database/clustering/overview/index.html>.

To configure an HA cluster on RAC:

- 1 During installation of the primary Management Server, choose Oracle as the external database.
- 2 Choose the **Configure Later** option.

- 3 Ensure that you do not start the Management Server automatically when installation is complete.
- 4 Before starting the server, open the `datasource-oracle.properites` file and edit the `url` section:


```
url=jdbc:oracle:thin:@(DESCRIPTION= (FAILOVER=ON)
(AADDRESS_LIST=(LOAD_BALANCE=ON)
(AADDRESS=(PROTOCOL=TCP) (HOST=primary.domain.com) (PORT=1521))
(AADDRESS=(PROTOCOL=TCP) (HOST=secondary.domain.com) (PORT=1521)))
(CONNECT_DATA=(SERVICE_NAME=FGLRAC.domain.com)))
```
- 5 Install the Management Server on each machine in the cluster in turn.
- 6 Replace the `datasource-oracle.properites` file on each machine with the modified one shown partially in [Step 4](#).
- 7 Start the primary and secondary Management Servers.

NOTE: The previous procedure outlines only one method of configuration. Other methods include: using the SQL scripts included with Foglight to do the first-time configuration, and further configuration of the `oracle-ds.xml` file. These methods are possible, but have not been tested.

Running behind an HTTP proxy

One of the most common configurations is running the Management Server in HA mode behind a proxy so that end users and the agent configuration only need one URL for connection to the browser interface.

In the following example, a third-party software program called *balance* is used as a proxy.

To run the Management Server behind an HTTP proxy:

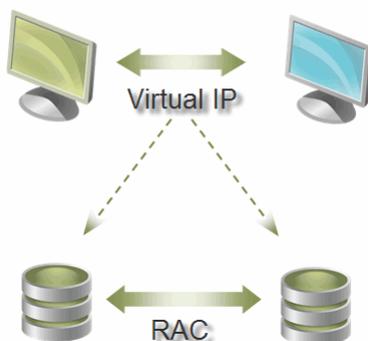
- 1 Install *balance*, available from: <http://www.inlab.de/balance.html>.
- 2 Configure *balance* to forward connections made to the server's port 8888 to the Foglight® UI port on the two machines in the HA cluster.
- 3 Configure the proxy to regularly test the two Management Servers to determine which one is active and to send all requests to the active Management Server.

NOTE: On failover to another member of the cluster, the Foglight browser interface session will restart.

Running in a cluster

This section describes one scenario for running the Foglight® Management Server in HA mode with clustering at the application layer. In particular, the focus is on operating system level clustering used in commercial systems such as Oracle®, Microsoft®, and Veritas™ clusters.

Figure 2. Cluster configuration



This scenario uses the Linux® HA Project's *heartbeat* software for SuSE ES 10.1. The *heartbeat* package includes a virtual IP resource called *IPaddr2*.

Configuration:

- 1 One virtual IP resource was set up in the cluster to migrate between two hosts.
- 2 The Management Server was installed on a shared disk so that it could run on either node in the cluster.
- 3 Agents were configured with the virtual IP.
- 4 All interaction with the Foglight browser interface was done through the virtual IP.

i | **NOTE:** The agent has no mechanism to tie it to a particular interface, and it stops sending data to the Management Server at failover until restarted.

For more information about the Linux HA Project and heartbeat, see <http://linux-ha.org>.

Tuning connection issues in HA implementations

In some cases, the communication between servers in the cluster is not reliable. Sometimes unwanted behaviors occur, such as:

- A secondary server takes over while the primary server is still running.
- Messages from an HA member are discarded as “message from non-member”.
- When starting up a secondary server, it fails to recognize the primary server.

These issues may be attributable to JGroup (the underlying communication package that the Foglight Management Server uses for its HA implementation) and to the fact that a Foglight HA implementation uses UDP for communication by default and UDP is by nature an unreliable protocol.

- [Tuning the Management Server](#)
- [Managing hosts with multiple network interfaces](#)

Tuning the Management Server

If you encounter these HA issues, consider performing the following tuning exercise.

To tune the Management Server:

- 1 Ensure the servers have synchronized system clocks. Out of sync server system clocks greatly increase the risk of communication errors.

i | IMPORTANT: If the system clock is out of sync by 5 seconds or more during HA server startup, a warning message similar to the following appears in the log:

```
Warning: system time not in sync with primary server, Mon Aug 30 11:47:57
ECT 2010 vs Mon Sug 30 11:48:52 EDT 2010.)
```

- 2 Reconfigure JGroup to use TCP instead of UDP as the communication protocol. This is suitable for clusters with a small number (less than five) of predetermined servers.

i | IMPORTANT: You must make the following changes to each server in the cluster.

- a Shut down the servers.
- b Edit the `config/jgroups-config.xml` file.
 - a Uncomment the TCP block:


```
<TCP
      . . . .
/>
```
 - b Uncomment the TCPPING block:

```
<TCPPING
    ....
/>
```

- c Comment out the UDP block:

```
<UDP
    ....
/>
```

- d Comment out the PING block:

```
<PING
    ....
/>
```

- c Make these changes on both your primary and secondary HA servers:

- a Find the line beginning with:

```
<TCPPING initial_hosts="${jgroups.bind_addr}[7800],otherhost[7800]"
    port_range="3"
```

- b Change `otherhost` to the host name of the secondary server in the cluster. For example:

```
<TCPPING initial_hosts="${jgroups.bind_addr}[7800],host2[7800]"
    port_range="3"
```

- c If there are more than two servers in the cluster, add those servers to the list as well. For example:

```
<TCPPING initial_hosts="${jgroups.bind_addr}[7800],host2[7800],
    host3[7800]" port_range="3"
```

- d Repeat the edit on your secondary HA host, to the initial HA host name. For example:

```
<TCPPING initial_hosts="${[7800],host1[7800]" port_range="3"
```

If there are more than two servers in the cluster, add those servers to the list as well. For example:

```
<TCPPING initial_hosts="$ {jgroups.bind_addr}[7800],host1[7800],
    host3[7800]"
```

i | **NOTE:** In these examples, 7800 is the port number used for TCP communication. For more information about the HA ports used by the Management Server, see [Configuring High Availability mode ports](#) on page 16.

- e Save your changes.

- f Start the primary server. Wait for the primary server to start up completely, then restart the secondary server(s).

Managing hosts with multiple network interfaces

When configuring Foglight[®] Management Servers that are installed on hosts with multiple network interfaces (that is, with multiple IP addresses and host names), you can specify the IP address that is to be used for communication with the server. To specify the IP address and host name, configure the following parameters in the `server.config` file:

```
server.bind.address = "X.X.X.X";
server.remote.address = "host_name";
```

Where:

- `X.X.X.X` is the desired IP address.
- `host_name` is the DNS (Domain Name System) host name or the IP address.

Example A

```
server.bind.address = "192.0.2.2";
server.remote.address = "host1.example.com";
```

If a DNS host name is not available for the bind address, the IP address can be used:

```
server.bind.address = "192.0.2.2";
server.remote.address = "192.0.2.2";
```

When the Management Server is configured to run in HA mode, you also need to reconfigure the restarter to perform server health check with the same host name or IP address. To do that, open the `restart_monitor.config` file, locate the `health.check.url` and `server.host` properties, and set their values to the one specified by the `server.remote.address` parameter (in `server.config`).

Example B

```
server.bind.address = "host1.example.com";
```

If a DNS name is available for the desired bind address, the name can be used instead of the raw address.

For more information about `server.bind.address` and `server.remote.address`, see the comments in `server.config`.

JDK with IPv6 on Linux

Management Servers running on a VMware® Linux® x86_64 image fail to start up and a “`java.net.BindException: Cannot assign requested address exception`” is raised. (85262)

This is a JDK issue with IPv6 on Linux, not a true High Availability issue. It also impacts standalone servers. However, the issue is documented here because it occurs when the JGroup cluster service is starting up.

Workaround: Start the Management Server with the following command:

```
-Djava.net.preferIPv4Stack=true
```

For more information, see: the JGroups FAQ (<http://community.jboss.org/wiki/JGroupsFAQ>) and Oracle® Bugs for it (http://bugs.sun.com/bugdatabase/view_bug.do?bug_id=6521974).

Foglight Management Server automatically restarted

Sometimes, after what appears to be a normal and successful startup, an HA server is automatically shut down and restarted. The most likely reason for this is a misconfiguration of health check URL in the `restart_monitor.config` file.

For example, you may have reconfigured the HTTP port of the Management Server or reconfigured the IP address that the server is bound to, but forgotten to reconfigure the health check URL of the restarter. If the restarter cannot contact the Management Server for health check, it considers the Management Server not responsive and restarts it.

You can check the `server_restarter_xxxx-xx-xx_xxxxx_xxx.log` file to determine if this scenario is what caused the restart. If so, edit the `restart_monitor.config` file by locating the line beginning with `health.check.url` and configuring the URL properly.

Other JGroup-related issues

ERROR org.jgroups.protocols.UDP max_bundle_size (64000_ is greater than the largest TP fragmentation size (8000):

<https://jira.jboss.org/browse/JGRP-798>

“Cross talk” can occur between servers on different clusters:

<https://jira.jboss.org/browse/JGRP-614>

We are more than just a name

We are on a quest to make your information technology work harder for you. That is why we build community-driven software solutions that help you spend less time on IT administration and more time on business innovation. We help you modernize your data center, get you to the cloud quicker and provide the expertise, security and accessibility you need to grow your data-driven business. Combined with Quest's invitation to the global community to be a part of its innovation, and our firm commitment to ensuring customer satisfaction, we continue to deliver solutions that have a real impact on our customers today and leave a legacy we are proud of. We are challenging the status quo by transforming into a new software company. And as your partner, we work tirelessly to make sure your information technology is designed for you and by you. This is our mission, and we are in this together. Welcome to a new Quest. You are invited to Join the Innovation™.

Our brand, our vision. Together.

Our logo reflects our story: innovation, community and support. An important part of this story begins with the letter Q. It is a perfect circle, representing our commitment to technological precision and strength. The space in the Q itself symbolizes our need to add the missing piece—you—to the community, to the new Quest.

Contacting Quest

For sales or other inquiries, visit www.quest.com/contact.

Technical support resources

Technical support is available to Quest customers with a valid maintenance contract and customers who have trial versions. You can access the Quest Support Portal at <https://support.quest.com>.

The Support Portal provides self-help tools you can use to solve problems quickly and independently, 24 hours a day, 365 days a year. The Support Portal enables you to:

- Submit and manage a Service Request.
- View Knowledge Base articles.
- Sign up for product notifications.
- Download software and technical documentation.
- View how-to-videos.
- Engage in community discussions.
- Chat with support engineers online.
- View services to assist you with your product.