

Foglight® 7.3.0
Agent Manager Guide



© 2024 Quest Software Inc.

ALL RIGHTS RESERVED.

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

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

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

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

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

Patents

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

Trademarks

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

their respective owners.

Legend

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

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

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

Contents

| | |
|--|-----------|
| Configuring the embedded Agent Manager | 7 |
| Getting started | 7 |
| Interacting with the embedded Agent Manager | 8 |
| Automatic configuration of the Management Server host name and port number | 9 |
| Installing external Agent Managers | 10 |
| Understanding how the Agent Manager communicates with the Management Server | 10 |
| Deploying the Agent Manager cartridge | 11 |
| Downloading the Agent Manager installer | 12 |
| Installing the Agent Manager | 14 |
| Installing the Agent Manager using the installer interface | 15 |
| Installing the Agent Manager from the command line | 22 |
| Using the Agent Manager silent installer | 31 |
| Installing the Agent Manager as a Windows service | 34 |
| Starting or stopping the Agent Manager process | 34 |
| Identifying the Agent Manager process | 35 |
| About platform-specific identification | 36 |
| Frequently asked questions | 36 |
| Configuring the Agent Manager | 40 |
| Operating system patches | 40 |
| Launching the Agent Manager installation interface | 40 |
| Configuring the Agent Manager to run in FIPS-compliant mode | 41 |
| Configuring the Agent Manager from the command line | 42 |
| Configuring the Agent Manager to use SSL certificates | 43 |
| Configuring an Agent Manager instance as a Concentrator | 44 |
| Configuring the concentrator | 45 |
| Configuring downstream Instances | 47 |
| Creating a secure connection with downstream instances | 47 |
| Excluding SSL ciphers from upstream or downstream Connections | 49 |
| Excluding Specific SSL Protocols from Downstream Connections | 49 |
| Configuring the Agent Manager to accept connections from the Management Server | 50 |
| Configuring the Agent Manager to execute commands on remote hosts | 53 |
| Configuring multiple Agent Manager instances | 54 |
| Example: Running multiple instances in a cluster environment | 55 |
| Controlling the polling rate | 56 |
| Configuring the Agent Manager to work in HA mode | 56 |
| Assigning Agent Managers to HA partitions | 57 |
| Adding cartridges to the HA deployment whitelist | 58 |
| About agent fail-over | 58 |
| About non-HA deployments | 59 |
| Negotiating Agent Manager resources at runtime | 59 |
| Disabling runtime resource negotiation | 59 |

| | |
|---|-----------|
| Configuring credentials | 60 |
| Configuring anti-virus exclusion settings | 62 |
| Troubleshooting | 63 |
| Errors related to Windows WMI and DCOM configuration | 63 |
| Resolving DATA_COLLECTION_FAILED errors | 63 |
| Adjusting the maximum polling interval | 63 |
| Advanced system configuration and troubleshooting | 65 |
| Configuring Windows Management Instrumentation (WMI) | 65 |
| WMI IPv6 connection support | 66 |
| Windows Firewall interference | 66 |
| Minimum requirements for Windows Management Instrumentation | 66 |
| WMI access violation and OS connectivity verification failure | 68 |
| WMI and Quota Violation error | 69 |
| Known WMI issues in Windows Server 2008 | 69 |
| Tuning WMI connections | 69 |
| Modifying registry key ownership on Windows Server 2008 R2 | 70 |
| Configuring registry settings for Windows Server 2008 R2 and Windows 7 | 70 |
| Resolving Access Denied errors when connecting to Windows XP Professional | 71 |
| OS collection fails with a <i>Local_Limit_Exceeded</i> error | 72 |
| Access to DCOM objects and registry is denied | 72 |
| Configuring registry settings for WinShell access through DCOM | 73 |
| Permissions on registry keys to configure DCOM command shell connection | 74 |
| Enabling agents to connect from UNIX machines | 75 |
| Enabling agents to connect locally on Windows | 76 |
| Releasing a locked MySQL process | 77 |
| Configuring Windows Remote Management (WinRM) | 77 |
| WinRM IPv6 connection support | 78 |
| Understanding Negotiate/Kerberos authentication | 78 |
| Understanding Basic authentication | 82 |
| About WinRM authentication and the Agent Manager | 82 |
| Configuring the target (monitored) system | 83 |
| Configuring the Agent Manager (monitoring) system | 85 |
| Generating a configuration file required for WinRM Negotiate authentication | 85 |
| Configuring command-shell connection settings | 87 |
| About WinRM connection ports | 88 |
| Troubleshooting | 89 |
| UNIX- and Linux-specific configuration | 91 |
| Agent Manager service can't start automatically when the operating system restarts .. | 91 |
| SSH IPv6 connection support | 93 |
| About supported remote monitoring protocols | 93 |
| Configuring the Agent Manager to run as a daemon | 93 |
| Configuring Agent Manager agent privileges | 94 |
| Preventing Agent Manager core dumps on Linux | 97 |
| Monitoring the Agent Manager performance | 99 |
| Investigating Agent Manager diagnostics | 99 |
| Exploring the Adapter Details tab | 100 |

| | |
|---|------------|
| Exploring the Client Details tab | 101 |
| Deploying the Agent Manager to large-scale environments | 107 |
| Using Agent Manager silent installer Parameters | 107 |
| Example: Deploying the Agent Manager to multiple UNIX hosts | 107 |
| Working with this example | 108 |
| Example: Creating and running an Agent Manager deployment script | 108 |
| Example: Deploying the Agent Manager to multiple Windows hosts | 112 |
| Example: Using a software deployment tool and silent installer parameters | 112 |
| Next steps | 113 |
| About Us | 114 |
| Technical support resources | 114 |

Configuring the embedded Agent Manager

This guide provides instructions for installing, configuring, and starting the Foglight® Agent Manager. Before you begin, refer to the *Foglight System Requirements and Platform Support Guide*, and the *Foglight Installation and Setup Guide* set.

This guide provides installation instructions for both Windows® and UNIX® operating systems. Commands are listed separately for each OS. Generic examples follow the UNIX path structure that uses forward slashes '/' to separate directories; substitute the back slash where necessary for your operating system.

- [Getting started](#)
- [Interacting with the embedded Agent Manager](#)

Getting started

An instance of the Foglight® Agent Manager is automatically installed with new installations of the Foglight Management Server. This embedded Agent Manager instance runs on the Management Server machine. You can deploy agents to the embedded Agent Manager if you want to monitor the machine on which the Management Server runs.

In certain environments, Foglight starts and stops the embedded Agent Manager along with the Management Server by default. You can configure whether or not you want the embedded Agent Manager to run in tandem with the Management Server.

i **NOTE:** Although the Agent Manager Adapter cartridge is installed by default with the embedded Foglight, you must still install a Foglight cartridge that contains installers for your supported platforms in order to deploy the Agent Manager to remote hosts. See “Installing External Agent Managers” on page 11 for information about selecting a Agent Manager cartridge, downloading installers, installing, configuring, and running the Agent Manager.

By default, Foglight starts and stops the embedded Agent Manager along with the Management Server when the Management Server is installed in standalone mode.

The embedded Agent Manager does not start automatically when:

- Foglight is installed in High Availability (HA) mode. If you want to run the Agent Manager on the Management Server machines in your HA cluster, you must install, configure, and run separate (non-embedded) Agent Manager instances on these machines. See [Installing external Agent Managers](#) on page 10 for instructions. For more information about High Availability mode, see the *High Availability Field Guide*.
- The Management Server is a Federation Master. The embedded Agent Manager does not start if you install Foglight as a Federation Master. See the *Federation Field Guide* for more information about federation.

i **CAUTION:** The `fglam --create-state` and `fglam --location` command-line options are not recommended for use with the embedded Agent Manager. For more information, see [Configuring multiple Agent Manager instances](#) on page 54.

If you are running the Management Server as a Windows[®] service, the embedded Agent Manager continues to run when you log out of the machine on which the embedded Agent Manager and the Management Server are running.

An embedded instance of the Agent Manager automatically installs with the Management Server. You can deploy agents to the embedded Agent Manager. You can specify whether or not the embedded Agent Manager runs in tandem with the Management Server by changing the value of the `server.fglam.embedded` parameter in the `server.config` file.

To configure whether or not the embedded Agent Manager runs automatically with the Management Server:

- 1 Stop the Management Server.
- 2 On the Management Server machine, open `<foglight_home>/config/server.config` for editing.
- 3 Set the parameter `server.fglam.embedded` to the desired value:
 - To have the embedded Agent Manager stop and start automatically with the Management Server, set `server.fglam.embedded` to `true`.
 - To disable the embedded Agent Manager, set `server.fglam.embedded` to `false`.
- 4 Save your changes to the `server.config` file.
- 5 Restart the Management Server.

Interacting with the embedded Agent Manager

You can run the configuration interface from the embedded Agent Manager installation directory, `<foglight_home>/fglam/bin`.

To interact with the embedded Agent Manager:

- 1 Launch a command shell on the machine hosting the embedded Agent Manager machine and navigate to the `bin` directory.
- 2 Stop the embedded Agent Manager by issuing the following command:

```
fglam --stop
```
- 3 Run the embedded Agent Manager from the command line with the appropriate set of options, depending on how you want to run the configuration interface.
 - To launch the Agent Manager configuration GUI, issue the following command:

```
fglam --configure
```
 - To launch the Agent Manager configuration command-line interface, issue the following command:

```
fglam --headless --configure
```
- 4 Follow the appropriate set of instructions provided in [To launch the Agent Manager Installation and Configuration interface](#): on page 41 or [To launch the Agent Manager configuration command-line interface](#): on page 42 for proceeding through the configuration interface steps.

i | IMPORTANT: On UNIX[®] platforms, if you attempt to run the installation interface, and the `DISPLAY` environment variable is not set on the machine, the Agent Manager defaults to the command-line interface.

Aside from setting the host name and port for the Management Server with which it communicates (described in [Automatic configuration of the Management Server host name and port number](#) below), you can configure the embedded Agent Manager as you would a standalone Agent Manager.

- 5 Restart the embedded Agent Manager by issuing the following command:

```
fglam --start
```

Automatic configuration of the Management Server host name and port number

The embedded Agent Manager only communicates with the Management Server that starts it (the Management Server with which it was installed).

When configuring the embedded Agent Manager, you do not need to set the host name and port number of the Management Server with which it communicates. The embedded Agent Manager automatically sets the host name and port number of the Management Server that starts it in its configuration file, *fglam.config.xml*, which is located in the `<foglight_home>/fglam/state/default/config/` directory.

If you configure a host name and port for a different Management Server, these settings are overwritten when the embedded Agent Manager starts.

Installing external Agent Managers

A Foglight® agent is a component hosted by the Agent Manager that:

- collects data from monitored resources
- if necessary, manipulates data into the format required by the Management Server
- submits the data to the Management Server using the Agent Manager

The Agent Manager is a middleware application that:

- hosts agents
- manages agent lifecycle (such as start, stop, and restart)
- manages the connection to the Management Server
- provides centralized services to agents

The Agent Manager supplies a centralized communications link between the Management Server and monitoring agents. It also provides a number of support services such as installation, upgrade, and the ability to configure agents.

- [Understanding how the Agent Manager communicates with the Management Server](#)
- [Deploying the Agent Manager cartridge](#)
- [Downloading the Agent Manager installer](#)
- [Installing the Agent Manager](#)
- [Starting or stopping the Agent Manager process](#)
- [Frequently asked questions](#)

Understanding how the Agent Manager communicates with the Management Server

By default, the Agent Manager initiates communication with the Management Server. However, this communication is not always uni-directional. In some cases, the Management Server can initiate a reverse connection to the Agent Manager.

If the Agent Manager is configured with an open concentrator port, and the Management Server can resolve and connect to the Agent Manager host address, then the Management Server attempts to open a reverse connection to the Agent Manager as an optimization. This enables the Management Server to send agent control messages (for example, for starting or stopping the agent, configuration changes, or callbacks) without waiting for the Agent Manager to initiate a connection.

Configuring an open concentrator port is an optional step. If there is no concentrator, or if the Agent Manager cannot be reached by the Management Server due to NAT or a firewall, the Management Server buffers the messages until the Agent Manager polls. Therefore, Agent Manager connections through NAT or a firewall are fully supported.

Embedded Agent Managers always open a private concentrator port that is restricted to connections from *localhost* only. This automatically enables reverse connections without requiring any configuration.

For more information, see [Configuring an Agent Manager instance as a Concentrator](#) on page 44.

By default, the Agent Manager installer uses the system-defined temporary directory (for example, `/tmp` or `C:\TEMP`) to unpack and execute the files required during installation. In some cases, such as when `/tmp` is mounted `noexec`, you may need to specify an alternate temporary location. To override the default temporary installation location, simply set the `TMP`, `TEMP`, or `TMPDIR` environment variable to point to a different location.

i | **NOTE:** These variables are checked in the order listed. If more than one is set, the first valid setting is used.

Syntax conventions

- `[option1|option2]` is used in file names at points where you must type, select, or otherwise specify one of multiple options. For example, `FglAM-[all|<platform>]-<version>.car` might be used in a step where you need to choose between the file `FglAM-all-<version>.car` or the file `FglAM-<platform>-<version>.car`.
- `<platform>` represents the appropriate name for your platform. For example, `FglAM-linux-x86_64-<version>.car`, for a 64-bit Linux® platform. For more information, see [Downloading the Agent Manager installer](#) on page 12.
- `<version>` represents the version number of the Agent Manager that you are installing. For example, `FglAM-<platform>-5_8_5_5.car`.

Deploying the Agent Manager cartridge

Several types of Agent Manager cartridges are available in the `FoglightClient-FglAM-<version>` folder in the `Server` directory on `Disk1` of your Foglight® installation media or for download from our Support Portal.

The type of Agent Manager cartridge you deploy depends on your needs and the monitored environment:

- The cartridge file `FglAM-all-<version>.car` contains Agent Manager installers for all supported platforms. Deploy this cartridge if you plan to install the Agent Manager on multiple platforms in your environment.
- The cartridge files `FglAM-<platform>-<version>.car` contain platform-specific Agent Manager installers. Deploy a platform-specific cartridge if you do not plan to install the Agent Manager on more than one platform in your environment. Consult the table below for the appropriate installer name for your platform.
- The cartridge file `FglAM-patch-<version>.car` can only be used to upgrade certain versions of the Agent Manager to the current version; it cannot be used for a new install. See the *Foglight Upgrade Guide* for more information.

The following table outlines which cartridges contain the appropriate Agent Manager installers for your platforms.

i | **NOTE:** The cartridge `FglAM-patch-<version>.car` is not listed below. See the *Foglight Upgrade Guide* for more information about using this cartridge during the upgrade process.

i | **NOTE:** Support of Foglight Agent Manager (FglAM) on AIX, Sun Solaris, Oracle Solaris, and HP-UX platforms was discontinued as of version 7.1.0. FglAM installers are only provided for Windows and Linux platforms.

Table 1. FglAM Cartridges available per OS

| | | |
|---------------------|--------|--|
| Linux® | x86-64 | <code>FglAM-linux-x86_64-<version>.car</code> |
| | ia32 | <code>FglAM-windows-ia32-<version>.car</code> |
| Microsoft® Windows® | ia64 | <code>FglAM-windows-ia32-<version>.car</code> (emulated) |
| | x86-64 | <code>FglAM-windows-x86_64-<version>.car</code> |

To deploy the Agent Manager cartridge:

- 1 Log in to Foglight.
- 2 Navigate to the Cartridge Inventory dashboard (**Administration > Cartridges > Cartridge Inventory**).
- 3 On the **Installed Cartridges** tab, click **Install Cartridge**.
- 4 In the **Install Cartridge** dialog box, click **Browse** to navigate to and select the appropriate *FglAM-[all]<platform>]-<version>.car* file.
- 5 Ensure that the **Enable on Install** check box is selected.

i | **IMPORTANT:** If the Enable on Install check box is not selected, you must manually enable all cartridges that are contained within the *FglAM-[all]<platform>]-<version>.car* file you are deploying.

- 6 Click **Install Cartridge**.

After you deploy the applicable cartridge for the Agent Manager, you can download the installer for your platform and install the Agent Manager. See [Downloading the Agent Manager installer](#) on page 12, and [Installing the Agent Manager](#) on page 14.

Downloading the Agent Manager installer

The sections below describe the different methods to download and run the Agent Manager installer on a target system:

- Using the Components for Download dashboard
- Using the Components for Download page
- Using the Agent Managers or Agent Status dashboard
- Using a file retrieval tool

Before you begin, review the matrix below to determine the appropriate installer for your platform. The following table lists the installers available for the platforms on which the Agent Manager is supported.

Table 2. Agent Manager Installers


| Operating System | Architecture | FglAM Installer |
|---------------------|--------------|--|
| Linux® | ia32 | <i>FglAM-<version>-linux-ia32.bin</i> |
| | x86-64 | <i>FglAM-<version>-linux-x86_64.bin</i> |
| Microsoft® Windows® | ia32 | <i>FglAM-<version>-windows-ia32.exe</i> |
| | ia64 | <i>FglAM-<version>-windows-ia32.exe (emulated)</i> |

After deploying the *FglAM-[all]<platform>]-<version>.car* file on the Management Server, you can download the appropriate platform-specific installer to the target system from the Components for Download dashboard.

In situations where unauthenticated or headless access to the installers is required, you can download the appropriate platform-specific Agent Manager installer from a separate Components for Download page (after deploying the *FglAM-all-<version>.car* or *FglAM-<platform>-<version>.car* file on the Management Server).

i | **NOTE:** Components for Download is a stand-alone page. It does not provide access to any Foglight dashboards. If you click your Web browser's Back and Forward buttons while viewing this page, the pages you visited in your current browser session open; you do not navigate to Foglight dashboards.

To download the Agent Manager installer using the Components for Download dashboard:

- 1 Log in to Foglight®.
- 2 Navigate to the Components for Download dashboard (**Dashboards > Administration > Cartridges > Components for Download**).
- 3 On the Components for Download dashboard, in the row containing the Agent Manager installer for your platform, click .

i | TIP: For the appropriate installer name for your platform, consult the above table.

- 4 When prompted, save the installer.
- 5 Run the installer. See [Installing the Agent Manager](#) on page 14 for details.

To download the Agent Manager installer using the Components for Download page:

- 1 Open a Web browser and navigate to the following page:

```
[http|https]://<hostname>:<port>/console/installer/list
```

Where:

- <hostname> is the name of the machine where the Management Server is installed.
- <port> is the HTTP port specified during installation (the default is 8080).

If the Management Server has been configured to use HTTPS, specify `https` as the protocol.

i | IMPORTANT: No login is required to access this page.

- 2 Click the appropriate Agent Manager installer for your platform. For the appropriate installer name for your platform, consult the above table.
- 3 When prompted, save the installer.
- 4 Run the installer. See [Installing the Agent Manager](#) on page 14 for details.

To download and run the Agent Manager installer using the Agent Managers or Agent Status dashboard:

- 1 Log in to Foglight.
- 2 Navigate to the Agent Managers or Agent Status dashboard (**Dashboards > Administration > Agents > Agent Managers or Agent Status**).
- 3 In the top right corner, click Download Agent Manager Software.
- 4 In the dialog box that appears, click the appropriate Agent Manager installer for your platform. For the appropriate installer name for your platform, consult the above table.
- 5 Click Download to start the process.
- 6 Run the installer. See [Installing the Agent Manager](#) on page 14 for details.

To download and run the Agent Manager installer using a file retrieval tool:

- 1 Use a file retrieval tool to obtain the appropriate Agent Manager installer for your platform by specifying a path formatted as:

```
[http|https]://<hostname>:<port>/console/installer/download?downloadId=fglam-<version>-<platform>
```

Where:

- <hostname> is the name of the machine where the Management Server is installed
- <port> is the HTTP port specified during installation (the default port is 8080)

- `<version>` is the version of the Agent Manager
- `<platform>` is the appropriate Agent Manager installer for your operating system

For the appropriate installer name for your platform, consult the above table.

If the Management Server is configured to use HTTPS, specify `https` as the protocol.

For example, to retrieve the file `fglam-5_8_5_5-linux-ia32.bin` from a machine called `server1`, that uses the default port 8080, using Wget, issue the following command:

```
wget -O fglam-5_8_5_5-linux-ia32.bin
"http://server1:8080/console/installer/download-
installer.action?downloadId=fglam-5_8_5_5-linux-ia32.bin"
```

- 2 **UNIX[®] only.** Grant execute access to the installer by issuing the following command:

```
chmod +x FglAM-<version>-<platform>.bin
```

- 3 Run the installer. See [Installing the Agent Manager](#) on page 14 for details.

Installing the Agent Manager

After downloading the installer to the machine on which you want to install the Agent Manager, you can use the installer interface, run the installer from the command-line, or use the silent (non-interactive) installer.

Install the Agent Manager on each host that you want to monitor locally. Many agents allow hosts to be monitored remotely, from an Agent Manager installed on another machine. A local Agent Manager is not required on remotely monitored machines.

The Agent Manager should be installed in a directory that is local to the system. It should also run using a local account, not a network or domain account. This should also include a local user home directory. Because the Agent Manager monitors and detects problems such as network and disk failures, having the Agent Manager installed in a local directory and running it as a local user makes the Agent Manager more resistant to failures in those services and better able to detect and report those failures. Otherwise, having the Agent Manager installed on a network drive, could cause the Agent Manager to lock itself when the network drive fails, preventing this failure from being reported.

In some environments, the need for more efficient credentials management may require the use of domain accounts. In these environments, the use of gMSA (Group Managed Service Accounts) may provide additional security and efficiency. Refer to **Planning for authentication with gMSA** in the Foglight Installation and Setup Guide: Installing on Windows with an External Microsoft SQL Server Database.

Installing the Agent Manager using the installer interface

To install the Foglight[®] Agent Manager using the Agent Manager Installation and Configuration interface:

i | **NOTE:** On UNIX[®] platforms, if necessary, change the permissions for the installer file so that it is executable (as described in [To download and run the Agent Manager installer using the Agent Managers or Agent Status dashboard](#): on page 13).

- Run the installer executable in GUI mode on the target machine. To start the installer interface from the command line, run the installer executable with no options.

i | **IMPORTANT:** If you want to configure the Agent Manager to accept upstream HTTP connections during the installation, start the installer on the command line with the `--allow-unsecured` option:

```
FglAM-<version>-<platform>.exe --allow-unsecured
```

i | **NOTE:** If you want to install the Agent Manager as a Windows service, or as a Unix daemon, the installation interface appears with pre-selected options, indicating that the Agent Manager will start immediately after the installation. You can force these options to appear disabled by default (and enable them, if required, during the installation), if you start the installer executable with the `--no-start-on-exit` option:

```
FglAM-<version>-<platform>.exe --no-start-on-exit
```

For more information about these options in the installation interface, see [Step 7: Install init.d Script](#) on page 18 (Unix) and [Step 9: Windows Service](#) on page 20.

The installer program prompts you for information, and informs you of the progress of your installation.

On some platforms, a command shell may appear while the installer loads and extracts files to a temporary location on your machine.

i | **IMPORTANT:** On Windows® 7 and Vista, you must run the installer GUI as an administrator if you want it to automatically install the Agent Manager as a service. To do so, right-click the installer executable and select Run as Administrator.

The Agent Manager installation program consists of several steps. Each installation screen includes a **Previous** button, allowing you to go back and adjust the information you specified.

Step 1: Introduction

The **Agent Manager Installation and Configuration** window opens, showing the **Introduction** step.

Read the information in the **Introduction** step and click **Next**.

Step 2: License Agreement

i | **NOTE:** You must accept the license agreement before you can install the product.

- Read the information in the **License Agreement** step, enable the check box to accept the terms of the license agreement, then click Next.

Step 3: Choose Install Location

- In the **Choose Install Location** step, choose the directory where you want to install the Agent Manager and click Next.

If the selected directory does not exist, the installer informs you of this and prompts you regarding whether or not you want the directory to be created. To create the directory, click **Yes**. To return to the **Installation Directory** step and specify a different directory, click **No**.

If the Agent Manager is already installed in the directory, you must specify a different directory. The installer informs you if the Agent Manager is already installed in the directory, provides information about performing upgrades, and prompts you to select a different directory.

i | **NOTE:** If you are using gMSA authentication, the gMSA account must have read & write permission on the Agent Manager installation directory.

Step 4: Host Display Name

The **Host Display Name** step allows you to configure the host name that the Agent Manager uses to identify itself. This is also the name under which the Agent Manager submits metrics to the Management Server.

By default, the Agent Manager uses the host name that is automatically detected for the machine on which it is being installed. This host name initially appears in the **Host Display Name** box.

There are certain cases in which you should explicitly set the host name in this box: for example, if the host name is already in use by another machine. If necessary, you can replace the host name with a different (non-host

name) value that suits the needs of your environment; for example, `WebServer (Unix Cluster 1)` or `12345.example.com (Databases)`.

By default, the Agent Manager logs a warning when it starts if the host name you set differs from the automatically detected host name. This warning message appears in the console when you start the Agent Manager and in the Agent Manager log file. If you do not want the warning logged, clear the check box in this step.

i | **TIP:** If the machine on which you are installing the Agent Manager is configured with multiple IP addresses or host names, clear the check box to suppress the warning messages.

- Configure the host name settings, as required, and click **Next**.

Step 5: Server URLs

The **Server URLs** step provides multiple ways to configure the connection between the Agent Manager and the Management Server. For example, you can specify the URL of a single Management Server to which you want the Agent Manager to connect, or configure multiple Management Server URLs for failover purposes. You can also specify the URL of an Agent Manager concentrator to which you want the Agent Manager to connect.

i | **NOTE:** You can also configure Management Server URLs at a later time using the Agent Manager configuration interface. See [Configuring the Agent Manager](#) on page 40 for information about launching this interface post-installation.

From this step, you can also add SSL certificates to the Agent Manager's certificate store, or configure reverse polling.

- 1 By default, the Agent Manager uses secure connections (HTTPS) with the Management Server. If you need to use an unsecured connection, click **Allow Unsecured Connections** to enable this option.

i | **TIP:** To complete this configuration, you must clear the **Connect Using HTTPS** check box in [Step 2](#).

- 2 In the **Server URLs** step, click **Add**.

- a In the **Edit Server URL** dialog box that appears, specify the host name and port number that you want the Agent Manager to use when connecting to the Management Server.

i | **TIP:** You can also configure other Agent Manager connection options in this dialog. See [Configuring Management Server URLs using the installer interface](#) on page 21 for information about these options.




- b To use an unsecured connection with the Management Server, clear the **Connect Using HTTPS** check box.

i | **TIP:** This option is only available if you selected **Allow Unsecured Connections** in [Step 1](#).

- c Once you have specified the desired connection options, click **OK**.

Repeat these steps for each Management Server URL that you want to add. As you add the URLs, they appear listed in the **Server URLs** step. If you want to remove a URL, select the URL from the list and click **Delete**.

- 3 Test the connection between the Agent Manager and the Management Server. Select a URL in the list and click **Test**. An icon on the left of each URL indicates if the URL is tested and the outcome of the connection test:

- : The URL passed the connection test.
- : The URL is not tested.
- : The URL failed the connection test.

See [Frequently asked questions](#) on page 36 for information about why the URL may fail the connectivity test.

- 4 After you specify at least one Management Server for the Agent Manager to connect to, you can search for other Management Servers that have been configured to be part of the same HA (High Availability) partition by clicking **Find HA Servers**.

i | **IMPORTANT:** You can specify URLs for both Management Servers and Agent Manager concentrators. However, only Management Servers can have HA (High Availability) peers, not concentrators. Clicking Find HA Servers does not cause any concentrators to be added to the list of URLs.

If you specify a localhost address in the Edit Server URL dialog box and then search for HA servers, a URL that shows the real machine name is displayed as well as the URL for *localhost*. See [Frequently asked questions](#) on page 36 for more information.

- 5 To add SSL certificates, click **SSL Certificates**.

The **Manage SSL Certificates** dialog box appears. Use this dialog box, to add and remove SSL certificates from trusted certificate authorities. Only certificates that you manually add appear in this list. The default set of trusted CA certificates is not included. For more information, see [Configuring the Agent Manager to use SSL certificates](#) on page 43.

- a Click **Add**.
- b Use the file chooser to select an SSL certificate.
- c In the **SSL Certificate Alias** dialog box, type a name (alias) to identify the certificate you are adding, and click **OK**.

i | **NOTE:** The alias must be unique.

A summary of the new certificate appears in the **Manage SSL Certificates** dialog box.

- d To add another certificate, repeat [Step a](#) through [Step c](#).
 - e When done, click **OK**.
- 6 Click **Next**.

If you choose not to configure any Management Server URLs, add any URLs without testing them, or if there are URLs listed that failed the connectivity test, a **Warning** message box appears, asking you to confirm that you want to continue.
 - 7 If you are installing the Agent Manager on Windows®, proceed to [Step 9: Windows Service](#). Otherwise, continue with [Step 6: Configure Secure Launcher](#) on page 18 and then [Step 7: Install init.d Script](#) on page 18.

Step 6: Configure Secure Launcher

UNIX® platforms only.

The **Configure Secure Launcher** step allows you to configure the external loader used by the Agent Manager to provide certain Foglight agents with permissions required to gather system-level metrics. See [Configuring Agent Manager agent privileges](#) on page 94 for more information.

- 1 In the **Configure Secure Launcher** step, complete one of the following steps:
 - Accept the default setting.
 - Edit the path to point to a different *sudo* executable.
 - Edit the path to point to the executable for a *sudo*-like application.

i | **IMPORTANT:** After the installation is complete, you must edit the sudoers file for your system. If you are using a privilege-escalation tool other than sudo (for example, *setuid_launcher*), you must make changes related to that application. Follow the instructions in [Configuring Agent Manager agent privileges](#) on page 94.

- 2 When you have finished making changes to the Configure Secure Launcher screen, click **Next**.

Step 7: Install init.d Script

UNIX platforms only.

The **UNIX init.d Script** step allows you to configure the Agent Manager to run as a daemon. You do that by instructing the installer to installing an *init.d*-style script called *quest-fglam* in the *init.d* directory on your system. See [Locating the init.d script](#) on page 94 for the location of this directory on your operating system.

The system calls the *quest-fglam* script when the host on which the Agent Manager is installed starts up or shuts down.

- 1 In the **UNIX init.d Script** step, complete one of the following steps:
 - If you want to use the default configuration options, ensure the **Would you like to customize the start-up script** check box is cleared, and click **Next**. Then, continue to [Step 8: Downstream Connection Configuration](#).
 - To apply customized configuration options, click the **Would you like to customize the start-up script** check box, and proceed to [Step 2](#).
- 2 Complete one of the following steps:
 - To install the script, select the **Yes, install init.d scripts** check box.
 - i** | **NOTE:** This check box appears disabled if you are not running the Agent Manager installer as root.
 - If you do not want to install the scripts, ensure that the **Yes, install init.d scripts** check box is clear (the default setting).
 - i** | **NOTE:** Even if you choose not to install the *init.d* script, or if you are not performing the installation as the root user, two scripts to perform the necessary setup are generated for later use. See [Configuring the Agent Manager to run as a daemon](#) on page 93 for more information.
If you are not installing the *init.d* script at this time, it is recommended to configure as many options as possible in this step. They are referenced when generating these scripts for a later use.
- 3 **Linux only.** If the correct operating system is not auto-detected by the installer, from the **Select OS Type** menu, select your OS.
- 4 In the **Run as user** box, type the name of the user account used to run the Agent Manager.

If you choose to have the Agent Manager run as a different user than the one who performed the installation, then that user account must already exist in the system. This user becomes the owner of the `<fglam_home>` directory and all files within it, including all *state* directories.
- 5 Select the one or more of desired numeric run levels for the Agent Manager.
 - i** | **IMPORTANT:** These numbers signify different run levels on different UNIX platforms. Consult your UNIX system administrator for more information.
- 6 If you want the Agent Manager to start immediately after the installation, ensure that **Start Foglight Agent Manager at the end of the installation** is selected.
 - i** | **NOTE:** If the installer is started on the command line with the `--no-start-on-exit` option, this check box does not appear selected in the **Install init.d Script** step. Starting the installer without the `--no-start-on-exit` option causes the **Start Foglight Agent Manager at the end of the installation** check box to be selected by default. For more information, see [Installing the Agent Manager using the installer interface](#) on page 15.
- 7 Click **Next**.

If the user specified in the **Run as user** box is not a local user, the installer displays a Warning and prompts you for further actions.

 - To continue with the specified user, click **Yes**.

- To specify a different user, click **No**.

8 Continue to [Step 8: Downstream Connection Configuration](#).

Step 8: Downstream Connection Configuration

Foglight Agent Manager can accept incoming connections and be configured as a concentrator that acts as an intermediary connection to aggregates one or more downstream Agent Manager clients. A concentrator configuration provides a single connection through either a firewall or proxy for all downstream clients, or as an aggregated connection directly to the server.

Configuring the Agent Manager to act as a concentrator involves configuring queue and heap sizes to adequately support accepting and transferring data from one or more downstream connections. You can configure downstream connections when the Agent Manager needs to accept connections from the Management Server and enable reverse data polling. This is useful when the Agent Manager cannot connect to the Management Server due to its location.

Using this installer step, you can configure:

- Downstream SSL connections, when a certificate host name and a password are provided.
- Downstream non-SSL connections, but only if you started the installer on the command line with the `--allow-unsecured` option (as instructed in [Installing the Agent Manager using the installer interface on page 15](#)).

User-provided certificates or keystores are supported, but can be configured after the installation.

If you are configuring the Agent Manager as a concentrator in order to enable connections from the Management Server, additional setup is required. For more information about this procedure, or to find out to configure non-SSL connections and user-provided certificates, see [Configuring the Agent Manager to accept connections from the Management Server on page 50](#).

- 1 If you want to enable downstream connections, in the **Downstream Connection Configuration** step, select the **Enable concentrator/downstream connection support** check box.
- 2 Drag the **Pre-Configured Size** slider to set the desired queue and memory sizes.
- 3 In the **Port** box, type the port number that you want the Agent Manager to use to listen for downstream connections.
- 4 In the **Certificate Host Name** box, type the name of the host on which you are installing the Agent Manager. The host name you specify here is added to the SSL certificate that is to be generated.
- 5 In the **New Certificate Password** and **Re-enter Password** boxes, type the password of the SSL certificate keystore.
- 6 Click **Add**.

The **Allowed Downstream Connections** box refreshes, showing the newly created downstream connection URL.

- 7 If needed, create additional downstream connections.
To delete any downstream connections, select them in the **Allowed Downstream Connections** box, and click **Delete**.
- 8 Click **Next**.
- 9 If you are installing the Agent Manager on Windows, proceed to [Step 9: Windows Service on page 28](#). Otherwise, continue with [Step 7: Secure Launcher](#) and [Step 8: Install init.d Script](#).

Step 9: Windows Service

Windows® platforms only.

The **Windows Service** step allows you to specify if you want to install the Agent Manager as a Windows service. A Windows service operates in the background while the system on which it is installed is running. Installing the Agent Manager as a Windows Service causes the Agent Manager to start automatically on your system startup.

- 1 In the **Windows Service** step, complete one of the following steps:

- To install the Agent Manager as a Windows service, leave the check box selected.
 - If you do not want to install the Agent Manager as a Windows service, clear the check box.
- 2 If you want install the Agent Manager as a Windows service, and you want that service to start immediately after the installation, ensure that **Start Foglight Agent Manager Windows Service at the end of the installation** is selected.
 - i** | **NOTE:** If the installer is started on the command line with the `--no-start-on-exit` option, this check box does not appear selected in the **Windows Service** step. Starting the installer without the `--no-start-on-exit` option causes the **Start Foglight Agent Manager Windows Service at the end of the installation** check box to be selected by default. For more information, see [Installing the Agent Manager using the installer interface](#) on page 15.
 - 3 Click **Next**.

Step 10: Summary

The **Summary** step informs you that you can complete installation.

- 1 In the **Summary** step, click Finish.

The Agent Manager installer copies the Agent manager files to the machine on which it is being installed.

 - i** | **IMPORTANT:** On UNIX platforms, if the Agent Manager installation is performed by a user without root privileges, the `init.d`-style script is not installed. One or more messages appear, informing you of this, and also of the location of the script installer and a copy of the script itself.

The message *The Agent Manager has been installed* appears.

- 2 Click OK to close the message box.
 - i** | **IMPORTANT:** On some platforms, the Agent Manager Installer dialog might close automatically when the installation is complete or you might need to close it manually. If a command shell appeared while the installer was loading, this shell might also need to be closed manually.

On Windows operating systems, if you chose to install the Agent Manager as a Windows service, it starts automatically.

- i** | **IMPORTANT:** The Agent Manager starts automatically when the installation is complete only when it is installed as a service on Windows. It does not start automatically at the end of an installation on a UNIX platform.

Configuring Management Server URLs using the installer interface

As described in [Step 5: Server URLs](#) on page 16 of [Configuring Management Server URLs using the installer interface](#), you can configure the Agent Manager connection parameters using the installer interface.

Connecting to the Management Server using HTTPS

If the Management Server is configured to use HTTPS, then HTTPS can be used by the Agent Manager to connect to the Management Server.

- 1 In the **Server URLs** step, (see page 16), complete one of the following steps.
 - Double-click a Management Server URL.
 - Click **Add** to create a new Management Server URL.
- 2 In the **Edit Server URL** dialog box that appears, select the **Connect using HTTPS** check box
- 3 Select any of the following options, as required:

- **Allow self-signed certificates:** Select to enable the Agent Manager to accept self-signed certificates from the Management Server.
 - **NOTE:** It is not recommended to enable this configuration in FIPS-compliant mode for security consideration.
- **Allow a certificate with an unexpected common name:** Select to enable the Agent Manager to accept a certificate with a common name (host name) different than the one reported by the Management Server. Specify the name in the **Certificate Common Name** box.

Connecting to the Management Server using a proxy

You can specify whether the Agent Manager should connect to the Management Server using a proxy.

- 1 In **Server URLs** step, (see page 16), double-click a Management Server URL, or click **Add** to create a new one, as required.
- 2 In the **Edit Server URL** dialog box that appears, select the **Connect using a proxy** check box
- 3 Configure the following settings:
 - **Proxy URL:** Type the URL of the proxy used to connect to the Management Server.
 - **Username:** Type the user name needed to access the proxy.
 - **Password:** Type the password associated with the user name. The password is saved encrypted in the Agent Manager configuration file (`<fglam_home>/state/<state name>/config/fglam.config.xml`) the next time you start or restart the Agent Manager.
 - **NTLM Domain:** If the proxy uses Windows authentication, specify the Windows domain.

Binding to a local address

You can specify a local network address from which you want the Agent Manager to connect to the Management Server.

- 1 In **Server URLs** step, (see page 16), double-click a Management Server URL, or click **Add** to create a new one, as required.
- 2 In the **Edit Server URL** dialog box that appears, select the **Bind to a local address** check box
- 3 In the **Local Address** box, type the IP address of a NIC (network interface card) on the machine hosting the Agent Manager that you want to use to establish outbound connections to the Management Server.

Using GZIP compression

You can configure the Agent Manager to establish HTTP-compressed communication with the Management Server, if required.

- 1 In **Server URLs** step, (see page 16), double-click a Management Server URL, or click **Add** to create a new one, as required.
- 2 In the **Edit Server URL** dialog box that appears, the **Use GZIP compression** check box is selected by default. This causes HTTP-compressed communication between the Agent Manager and the Management Server, for both requests and responses.

Selecting this option establishes HTTP-compressed communication regardless of whether the Agent Manager connects to an Agent Manager concentrator or directly to the Management Server.

Installing the Agent Manager from the command line

The command-line installer prompts you for information and informs you of the progress of your installation.

To install the Agent Manager from the command line:

- 1 **UNIX® platforms only.** If required, change the permissions for the installer file to make it executable (as described in [To download and run the Agent Manager installer using the Agent Managers or Agent Status dashboard](#): on page 13).
- 2 Launch a command shell on the target machine and navigate to the directory to which you downloaded the installer.

i | **IMPORTANT:** On Windows® 7 and Vista, you must run the command-line installer from an administrator version of *cmd.exe* or PowerShell (not just logged in as administrator) if you want the installer to install the Agent Manager as a service.

- 3 Run the installer executable by issuing the following command:

```
FglAM-<version>-<platform> --headless
```

Where *<platform>* and *<version>* reflect the platform of the machine on which the Agent Manager is about to be installed and its version number. Consult the matrix in [Downloading the Agent Manager installer](#) on page 12 for more information.

i | **IMPORTANT:** If you want to configure the Agent Manager to accept upstream HTTP connections during the installation, you must start the installer on the command line with the `--allow-unsecured` option:

```
FglAM-<version>-<platform>.exe --headless --allow-unsecured
```

i | **NOTE:** If you want to install the Agent Manager as a Windows service, or as a Unix daemon, the default values in the related installation steps, if selected, cause the Agent Manager Windows service or Unix daemon to start immediately after the installation. You can override these defaults during the installation, or by starting the installer executable with the `--no-start-on-exit` option:

```
FglAM-<version>-<platform>.exe --headless --no-start-on-exit
```

For more information about these options in the installation interface, see [Step 8: Install init.d Script](#) on page 27 (Unix) and [Step 9: Windows Service](#) on page 28.

The command shell displays messages indicating that the installer files are being extracted to a temporary directory, and that the installer is starting up.

i | **TIP:** To cancel the installation at any time, press Ctrl-C.

Step 1: Introduction

When the installer finishes loading, the **Introduction** step appears.

- Review the information in the **Introduction** step and press Enter.

Step 2: License Agreement

i | **NOTE:** You must accept the license agreement before you can install the product.

- 1 When the first part of the **License Agreement** step appears, press Enter to page through the license agreement.
- 2 At the prompt *Do you accept the terms of the license agreement? [Y/N]*, type Y and press Enter to accept the terms, and to continue the installation.

Step 3: Installation Directory

The **Installation Directory** step allows you to specify the directory where you want to install the Agent Manager.

- Complete one of the following steps:
 - To accept the default installation directory, press Enter.

- To specify a different installation directory, type it at the command prompt, and press Enter.

If the specified directory does not exist, the installer prompts you regarding whether or not you want it to be created. To create the directory, press Enter. To return to the Install directory prompt and specify a different directory, type **N** and then press Enter.

If the Agent Manager is already installed in the directory, you must specify a different directory. The installer informs you if the Agent Manager is already installed in the directory, provides information about performing upgrades, and prompts you to select a different directory.

i | **NOTE:** If you are using gMSA authentication, the gMSA account must have read & write permission on the Agent Manager installation directory.

Step 4: Host Display Name

The **Host Display Name** step allows you to configure the host name that the Agent Manager uses to identify itself. This is also the name under which the Agent Manager submits metrics to the Management Server.

By default, the Agent Manager uses the host name that is automatically detected for the machine on which it is being installed.

There are certain cases in which you should explicitly set the host name in this box: for example, if the host name is already in use by another machine. If necessary, you can replace the host name with a different (non-host name) value that suits the needs of your environment; for example, `WebServer (Unix Cluster 1)` or `12345.example.com (Databases)`.

- 1 Complete one of the following steps:
 - To accept the detected host name, press Enter.
 - To use a different host name, type the host name and press Enter.

The *Log a warning if the detected host name changes* prompt appears.

This allows you to configure the Agent Manager to log a warning message when it starts if the specified host name differs from the automatically detected host name. If logged, this message appears in the console when you start the Agent Manager and in the Agent Manager log file.

- 2 If you want the Agent Manager to log a warning, press Enter.

If you want to suppress warning messages, or the machine on which you are installing the Agent Manager is configured with multiple IP addresses or host names, type **N** and press Enter.

Step 5: Server URLs

The **Server URLs** step provides multiple ways to configure the connection between the Agent Manager and the Management Server. For example, you can specify the URL of a single Management Server to which you want the Agent Manager to connect, or configure multiple Management Server URLs for failover purposes. You can also specify the URL of an Agent Manager concentrator to which you want the Agent Manager to connect.

i | **NOTE:** You can configure Management Server URLs at a later time using the Agent Manager configuration interface. See [Configuring the Agent Manager](#) on page 40 for information about launching this interface post-installation.

- 1 Add a new Management Server URL.
 - a Type 1 and press Enter.

Review the information on the screen about the available parameters. See also [Step 10: Change service credentials \[Optional\]](#) on page 28 for information about these parameters.
 - b When prompted, type the URL of a Management Server or Agent Manager concentrator to which you want the Agent Manager to connect, followed by the applicable parameters. For example:


```
url=http://localhost:8080,address=127.0.0.1,proxy=http://proxy.server
```
 - c Press Enter.

A message appears, indicating that the Management Server URL is added.

- d If you want this Agent Manager to connect to additional Management Server URLs, for example, for failover purpose repeat these steps.
- 2 To review a list of newly added URLs, type 4 and press Enter.
A numbered list of configured Management Server URLs appears on the screen.
- 3 To remove a URL, complete the following steps.
 - a Type 7 and press Enter.
A numbered list of configured Management Server URLs appears on the screen.
 - b Type the number identifying the URL that you want to delete and press Enter.
A message appears, indicating that the Management Server URL is deleted.
- 4 Test the connections between the Agent Manager and the Management Servers.
 - a Type 3 and press Enter.
A message appears, indicating the test progress. When the test is complete, the menu options appear on the screen.
 - b To see if the configured URLs passed the connectivity test, type 4 and press Enter.
A numbered list of configured Management Server URLs appears on the screen. Any URLs that fail the connectivity test are marked with an 'x' on the left. To find why a URL may fail the connectivity test, see [Frequently asked questions](#) on page 36.
The menu options appear on the screen.
- 5 If required, search for other Management Servers that are configured to be part of the same HA (High Availability) partition.

i **IMPORTANT:** You can specify URLs for both Management Servers and Agent Manager concentrators. However, only Management Servers can have HA (High Availability) peers, not concentrators. Searching for HA servers does not cause any concentrators to be added to the list of URLs.

- a To search for additional HA servers, type 3 and press Enter.
The installer searches for HA peers and tests the connections. A message appears informing you of the progress.
If you specify a *localhost* address at the URL prompt (described in [Step 1](#)), and then search for HA peers, a URL that shows the real machine name appears along with the *localhost* URL. See [Frequently asked questions](#) on page 36 for more information.
When the search is complete, the menu options appear on the screen.
 - b To verify if the installer found HA peers and added them to the list, type 4 and press Enter.
A numbered list of configured Management Server URLs appears on the screen. Any URLs that fail the connectivity test have an 'x' on their left. To find why a URL may fail the connectivity test, see [Frequently asked questions](#) on page 36.
The menu options appear on the screen.
- 6 You can also manage SSL Certificate CAs from the Server URLs menu. For more information, see [Configuring the Agent Manager to use SSL certificates](#) on page 43.

- a To view a list of all certificates currently in the certificate store, type **6** and press **Enter**.
- b To add a new SSL Certificate Authority (CA), type **2** and press **Enter**.

Type the alias and file names of the SSL CA certificate using the following syntax:

```
alias=file_name
```

Where

`alias` is the name that you want to associate with the certificate.

i | **NOTE:** The alias must be unique.

`file_name` is the full path to the certificate file.

For example, on a Windows machine, you can type the following:

```
NewCert1=C:\certificates\example_ca_certificate.crt
```

- c Press **Enter**.
 - d To remove a certificate from the certificate store, type **8** and press **Enter**.
- 7 If you want to enable the Agent Manager to accept connections from the Management Server, or downstream Agent Managers that use it as a concentrator, configure the ports that you want to use to listen for these types of connections. For more information, see [Configuring the Agent Manager to accept connections from the Management Server](#) on page 50.
- a Type **9** and press **Enter**.
 - b Specify the protocol type (HTTP or HTTPS) and the port number that you want the Agent Manager to listen on. Type them, as prompted, using the following syntax:

```
port=<port>, type=<http|https>
```
 - c Press **Enter**.
 - d To specify additional ports, repeat [Step a](#) through [Step c](#).
- 8 When you are finished adding Management Server URLs and managing SSL certificate CAs, type 0 and press Enter.

i | **IMPORTANT:** If you did not configure any Management Server URLs, if there are URLs listed that have not been tested, or if there are URLs listed that failed the connectivity test, the installer prompts you to confirm that you want to continue.

Continue with [Step 6: Downstream Connection Configuration](#).

Step 6: Downstream Connection Configuration

Foglight Agent Manager can accept incoming connections and be configured as a concentrator that acts as an intermediary connection to aggregates one or more downstream Agent Manager clients. A concentrator configuration provides a single connection through either a firewall or proxy for all downstream clients, or as an aggregated connection directly to the server.

Configuring the Agent Manager to act as a concentrator involves configuring queue and heap sizes to adequately support accepting and transferring data from one or more downstream connections. You can configure downstream connections when the Agent Manager needs to accept connections from the Management Server and enable reverse data polling. This is useful when the Agent Manager cannot connect to the Management Server due to its location.

Using this installer step, you can configure:

- Downstream SSL connections, when a certificate host name and a password are provided.
- Downstream non-SSL connections, but only if you started the installer on the command line with the `--allow-unsecured` option (as instructed in [Installing the Agent Manager using the installer interface](#) on page 15).

User-provided certificates or keystores are supported, but can be configured after the installation.

If you are configuring the Agent Manager as a concentrator in order to enable connections from the Management Server, additional setup is required. For more information about this procedure, or to find out to configure non-SSL connections and user-provided certificates, see [Configuring the Agent Manager to accept connections from the Management Server](#) on page 50.

- 1 To enable downstream connections, type **Y**, and press **Enter**.
- 2 To set the queue and memory sizes, type the number associated with the desired size of the queue and heap, and press **Enter**.

- 3 Type the port number that you want the Agent Manager to use to listen for downstream connections, and press **Enter**.
- 4 Type the SSL certificate password, and press **Enter**. Retype the password, and press **Enter**.

i | **NOTE:** The passwords are not visible on the screen as you type them.

- 5 Type the name of the host on which you are installing the Agent Manager. The host name you specify here is added to the SSL certificate that is to be generated.

If you are installing the Agent Manager on Windows, proceed to [Step 9: Windows Service](#) on page 28. Otherwise, continue with [Step 7: Secure Launcher](#) and [Step 8: Install init.d Script](#).

Step 7: Secure Launcher

UNIX platforms only.

The **Secure Launcher** step defines the external launcher used by the Agent Manager to provide certain Foglight agents with the required permissions to gather system metrics. See [Configuring Agent Manager agent privileges](#) on page 94 for more information.

- Complete one of the following steps:
 - To accept the default settings, press **Enter**.
 - To edit the path to point to a different *sudo* executable, type that path and press **Enter**.
 - To edit the path to point to the executable for a *sudo*-like application, type that path and press **Enter**.

i | **IMPORTANT:** After the installation is complete, you must edit the sudoers file for your system. If you are using a privilege-escalation tool other than *sudo* (for example, *setuid_launcher*), you must make changes related to that application. Follow the appropriate set of instructions in “Using Sudo to Configure Secure Launcher Permissions” on page 88.

Step 8: Install init.d Script

UNIX platforms only.

The **Install init.d Script** step allows you to configure the Agent Manager to run as a daemon. You do that by instructing the installer to installing an *init.d*-style script called *quest-fglam* in the *init.d* directory on your system. See [Locating the init.d script](#) on page 94 for the location of this directory on your operating system.

The system calls the *quest-fglam* script when the host on which the Agent Manager is installed starts up or shuts down. See [Configuring the Agent Manager to run as a daemon](#) on page 93 for more information.

Installation as the root user only: The installer prompts you to specify whether you want to install the *init.d* script.

Installing as a non-root user only: The installer does not prompt you to specify whether you want to install the *init.d* script. However, the installer generates two scripts that perform the necessary setup for later use. See [Configuring the Agent Manager to run as a daemon](#) on page 93 for more information.

- 1 In the *UNIX init.d Script* step, complete one of the following steps:
 - If you want to use the default configuration options, type **Y**, and press **Enter**. Then, continue to [Step 10: Summary](#).
 - To apply customized configuration options, type **N**, and press **Enter**. Then, proceed to [Step 2](#).
- 2 Using the *Select the Unix/Linux distribution* menu, specify the number in the identifies your operating system.

i | **IMPORTANT:** Even if you choose not to install the *init.d* script, or if you are not performing the installation as the root user, it is recommended that you configure as many options as possible in this step, and in the subsequent Install UNIX/Linux *init.d* Scripts steps. These options are used when generating the two scripts for later use.

- To accept the detected value, press **Enter**.

- To specify a different operating system, type the number associated with it, and press **Enter**.
- 3 At the *Enter the user the service will run as* prompt, specify the user account that runs the Agent Manager daemon.
 - To accept the default (current) user, press **Enter**.
 - To specify a different user account, type its name and press **Enter**.

If you choose to have the Agent Manager run as a different user than the one who performed the installation, then that user account must already exist in the system. This user becomes the owner of the `<fglam_home>` directory and all files within it, including all *state* directories.

If the user specified at this prompt is not a local user, the installer displays a warning and prompts you to specify whether you want to continue.

To continue with the specified user, press **Enter**.

To specify a different user, type **N** and press **Enter**. At the prompt, specify the user name.
 - 4 Type the number that corresponds with a run level to select a run level and press **Enter**.
 - i** | **IMPORTANT:** These numbers signify different run levels on different UNIX platforms. Consult your UNIX system administrator for more information.
 - 5 Type **0** and press **Enter**.
 - 6 If you want the Agent Manager process to start immediately after the installation, press **Y**, and then Enter.
 - i** | **NOTE:** The default answer to this question is **Y** (Yes). However, if you started the installer with the `--no-start-on-exit` option, the default answer to this question is set to **N** (No). For more information, see [Installing the Agent Manager from the command line](#) on page 22.
 - 7 Continue with [Step 10: Summary](#).

Step 9: Windows Service

Windows platforms only.

The **Windows Service** step allows you to specify if you want to install the Agent Manager as a Windows service. A Windows service operates in the background while the system on which it is installed is running. Installing the Agent Manager as a Windows Service causes the Agent Manager to start automatically on your system startup.

i | **NOTE:** If you intend to use gMSA authentication you must install the Agent Manager as a service.

- 1 Complete one of the following steps:
 - To install the Agent Manager as a Windows service, press Enter.
 - If you do not want to install the Agent Manager as a Windows service, type N and press Enter.
- 2 If you want the Agent Manager process to start immediately after the installation, press **Y**, and then Enter.

i | **NOTE:** The default answer to this question is **Y** (Yes). However, if you started the installer with the `--no-start-on-exit` option, the default answer to this question is set to **N** (No). For more information, see [Installing the Agent Manager from the command line](#) on page 22.

Step 10: Summary

The **Summary** step informs you that the installer has sufficient information to complete the Agent Manager installation.

- Press Enter to complete the installation and exit the command-line installer.

The Agent Manager installer calculates the amount of disk space required to complete the installation, and copies the required to the machine on which it is being installed.

i | **NOTE:** Unix platforms only. If the user account that you used to install the Agent Manager has no root privileges, the init.d-style script is not installed. Messages appear that inform you that the script is not installed due to insufficient user permissions and provide the path to the script installer and a copy of the script file.

If you choose to install it as a Windows service, the Agent Manager starts automatically.

i | **NOTE:** The Agent Manager starts automatically when the installation is complete only when it is installed as a service on Windows. It does not start automatically at the end of an installation on a UNIX platform.

Step 10: Change service credentials [Optional]

When using gMSA authentication, the Log On user for the Agent Manager service must be changed after installation.

- 1 Open Services (services.msc)
- 2 Find the Agent Manager service. \
- 3 Right-click on the service and choose Properties

Update the Log On User with the gMSA account. Leave the password field empty.

Configuring Management Server URLs from the command line

As described in [Step 5: Server URLs](#) on page 24 of [Installing the Agent Manager from the command line](#) and [Using the Agent Manager silent installer](#) on page 31, you can configure the Foglight® Agent Manager connection parameters using the `fglam` command and the arguments passed to its `fms` command option. For example:

```
fglam --headless --configure --fms url=http://localhost:8080, address=127.0.0.1,
proxy=http://proxy.server
```

i | **NOTE:** For complete information about the `fglam` command, see “Managing the Foglight Agent Manager” in the *Command-Line Reference Guide*.

Syntax

```
fglam --headless --configure [--fms url={http|https}://host:port
[proxy={http|https}://host:port [proxy-user=user_name] [proxy-pass=password]
[proxy-ntlm-domain=domain]] [address=IP_address]
[ssl-allow-self-signed={true|false}] [ssl-cert-common-name=name]
[compressed={true|false}]] [--downstream "<port=<port>,key-password=<password>>
[,<host=<host>,type=<https|http>,size=Small|Medium|Large|Huge|Maximum>"] |
[--deletedownstream <port>]| [--deletealldownstream]]
```

Table 3. Options and arguments

| Option | Argument | | Mandatory or Optional? | Description |
|------------------------|----------|-------|------------------------|---|
| | Name | Value | | |
| <code>headless</code> | | | Mandatory | Launches the Agent Manager and configuration interface on the command line, when used with the <code>configure</code> option. |
| <code>configure</code> | | | Mandatory | Launches the Agent Manager and configuration interface. |
| <code>fms</code> | | | Mandatory | Specifies the URL to the Management Server that you want to configure. |

Table 3. Options and arguments

| Option | Argument Name | Value | Mandatory or Optional? | Description |
|--------|------------------------------|--|------------------------|---|
| | url | http https: <i>//host:port</i> | Mandatory | The URL to the Management Server, where <i>host</i> and <i>port</i> specify the host name of the machine on which the Foglight Management Server is installed, and the port number the server uses to communicate with the Foglight Agent Manager. |
| | proxy | http https: <i>//host:port</i> | Optional | The URL to the proxy server needed to connect to the Management Server. |
| | proxy-user | <i>user_name</i> | Optional | The user name needed to connect to the proxy server. |
| | proxy-pass | <i>password</i> | Optional | The password associated with the user name needed to connect to the proxy server. The password is saved encrypted in the Agent Manager configuration file (<i><fglam_home>/state/<state name>/config/fglam.config.xml</i>) the next time you start or restart the Agent Manager |
| | proxy-ntlm-domain | <i>domain</i> | Optional | The Windows domain to which the proxy server belongs. |
| | address | <i>ip_address</i> | Optional | The IP address of the Foglight Agent Manager that is used to connect with the Foglight Management Server. |
| | ssl-allow-self-signed | true false | Optional | Indicates if self-signed certificates are accepted (true) or not (false). NOTE: It is not recommended to enable this configuration in FIPS-compliant mode for security consideration. |
| | ssl-cert-common-name | <i>name</i> | Optional | The common name of the expected certificate. Specifying this argument causes a certificate with a common (host) name, that is different than the one reported by the Management Server, to be accepted. |
| | compress | true false | Optional | Indicates if HTTP compression is enabled (true) or disabled (false). HTTP compression is enabled by default. When enabled, it is applied to both request and responses, and used when connecting to either Agent Management concentrators or Management Servers. |

Table 3. Options and arguments

| Option | Argument | | Mandatory or Optional? | Description |
|----------------------------|----------------|-----------------|--|--|
| | Name | Value | | |
| downstream | port | <i>port</i> | Mandatory | When used with --configure , this option creates a downstream connection. The port argument specifies the number of the port the Agent Manager uses to listen for downstream connections. |
| | key-password | <i>password</i> | Mandatory | This argument specifies the password needed to access the private key contained in the keystore. |
| | host | host | Optional | Specifies the host name to be set in the certificate. |
| | type | http https | Optional | Specifies the type of the supported protocol. |
| | size | Small | Optional | Specifies the amounts of the disk and memory resources. Small allocates 10 MB for the disk queue and 512 MB of memory. |
| | | Medium | | Medium allocates 100 MB for the disk queue and 768 MB of memory. |
| | Large | | Large allocates 500 MB for the disk queue and 1 GB of memory. | |
| | Huge | | Huge allows unlimited amounts of disk space for the queue, and up to 4 GB or 85% of system memory, whichever is less. | |
| | Maximum | | Maximum allows unlimited amounts of disk space for the queue, and 85% of available system memory. | |
| deletedownstream | port | <i>port</i> | Mandatory | When used with --configure , this option deletes a downstream connection given its port number. |
| deletealldownstream | | | | When used with --configure , this option deletes all downstream connections. |

Removing Management Server URLs

In addition to managing connections to the Management Server from the command line, the `fglam` command also provides arguments for removing Management Server URLs, when required. You can delete one URL at a time, or all of them, as required. Deleting all Management Server URLs can, for example, be useful in situations when you need to migrate an Agent Manager from one Management Server to another.

To delete a single Management Server URL:

- Issue the following command:

```
fglam --headless --configure --deletefms url=<http|https>://
<hostname>:<port>
```

To delete all configured Management Server URLs:

- Issue the following command:

```
fglam --headless --configure --deleteallfms
```

In some configuration scenarios, after dissociating your Agent Manager from one or all of your Management Servers, you typically need to connect it to another Management Server. The `fglam` command is flexible enough to allow for multiple operations on a single command line, as long as you first specify the arguments to first delete the existing URLs, and then add a new one. For example:

```
fglam --headless --configure --deleteallfms --fms url=http://Host1:8080
```

Detecting HA Servers

The `fglam` command also allows you to detect High Availability (HA) failover servers (peers) using its `detecttha` argument. When used with the `configure` option, this option instructs the installer to detect and test any available Management Server HA peers.

```
fglam --headless --configure --detecttha
```

Using the Agent Manager silent installer

The Agent Manager silent installer allows you to install the Agent Manager non-interactively—for example, to install the Agent Manager from the command line onto a remote, headless machine.

CAUTION: The silent installer is an advanced tool. It should be used cautiously and only by advanced Foglight administrators. Do not run the installer in silent mode unless you are an advanced Foglight administrator who is familiar with configuring the Agent Manager and you are certain what setup is required for your environment. If you are not certain which installation options you need, use the installer interface or command-line installer instead.

If you are running Foglight® in High Availability (HA) mode, you can configure the Agent Manager to work with a set of primary and secondary Management Servers in an HA cluster.

To install the Agent Manager using the silent installer:

- 1 If you are installing the Agent Manager onto a remote machine, log in to the target machine (for example, using SSH).
- 2 Launch a command shell on the target machine and navigate to the directory to which you downloaded the installer.
- 3 On UNIX® platforms, if necessary, change the permissions for the installer file so that it is executable (as described in [To download and run the Agent Manager installer using the Agent Managers or Agent Status dashboard](#): on page 13).
- 4 Run the installer executable with the `--silent` option.

IMPORTANT: On Windows® 7 and Windows Vista®, you must run the silent installer from an administrator version of `cmd.exe` or PowerShell (not just logged in as administrator) if you want the installer to install the Agent Manager as a service.

IMPORTANT: The `--silent` option is an advanced option. It should be used cautiously and only by advanced Foglight administrators.

All desired installation parameters must be included in the command with the `--silent` option, using the following syntax:

```
fglam-<version>-<platform> --silent --installdir <install_directory_path>
  [--fms <url_and_other_parameters>] [--noservice] [--host-display-name
  <display_name>] [--spid <path_to_SPID>] [--certificate <alias=path>]
  [--downstream "<port=<port>,key-password=<password>>[,type=<https|http>,
  host=<host>,size=<Small|Medium|Large|Huge|Maximum>"] ] [--allow-unsecured]
  [--no-start-on-exit] [-h|--help] [-v|--version]
  [<-m|--javavm> <path>] [--installer-properties <file_path>]
  [--headless] [auth-token <token>] [--noservice] [--host-display-name <name>]
  [--detecttha]
```

Where:

- `silent` prevents the installer from prompting for configuration options. It uses default values unless they are specified on the command line.
- `<version>` is the version number of the Agent Manager.

- `<platform>` is the appropriate installer name for your operating system. Consult the matrix in [Downloading the Agent Manager installer](#) on page 12 for the installer name.
- `<install_directory_path>` is the full path to the directory where you want to install the Agent Manager. The `installdir` option is mandatory.

i **IMPORTANT:** Ensure that there is no existing installation of the Agent Manager in the directory that you specify with the `--installdir <install_directory_path>` parameter. If there is an Agent Manager installation in the directory, you must specify a different directory or the installer fails. The installer aborts, to protect you from overwriting your existing installation. Consult the *Foglight Upgrade Guide* for upgrade instructions.

You can also include the following optional command-line options:

- `--fms <url_and_other_parameters>`, where `<url_and_other_parameters>` is a comma-separated list of parameters for configuring the connection to a Management Server. See [Step 10: Change service credentials \[Optional\]](#) on page 28 for a description of the parameters for the `--fms` option.

If you are running Foglight in High Availability (HA) mode, add the `--fms` option once for each member of the HA cluster to which you want the Agent Manager to connect. You must do so to configure the connection to more than one Management Server (or Agent Manager concentrator) within the cluster.

- `--noservice` (on UNIX) prevents an *init.d*-style script from being installed (that would automatically start the Agent Manager). On Windows®, this option causes the Agent Manager not to be installed as a service.
- `--host-display-name <display_name>`, where `<display_name>` is the host display name that you want to set manually for the Agent Manager. This is the host name that the Agent Manager uses to identify itself and the name under which it submits metrics to the Management Server. By default, the Agent Manager uses the host name that it automatically detects for the machine on which it is being installed. There are certain cases in which you should explicitly set the host display name: for example, if the host name is already in use by another machine.
- `--spid <path_to_SPID>`, where `<path_to_SPID>` is the path to an existing SPID installation. This allows you to (optionally) migrate agents from an existing SPID installation to the new Agent Manager installation. The Agent Manager automatically detects any agents managed by SPID from the old installation and copies the agent instances to the new Agent Manager installation.
- `--certificate alias=<path>` allows you to add an SSL certificate, where `<path>` is the directory path to the SSL certificate file.
- `--downstream` allows you to create a downstream connection.
 - `port=<port>` specifies the number of the port the Agent Manager uses to listen for downstream connections. This argument is mandatory.
 - `key-password=<password>` specifies the password needed to access the private key contained in the keystore.
 - `type=<https|http>` specifies the type of the supported protocol.
 - `host=<host>` specifies the host name to be set in the certificate.
 - `size=<Small|Medium|Large|Huge|Maximum>` specifies the amounts of disk and memory resources that are allocated to downstream connections.
- `--allow-unsecured` enables the configuration of HTTP downstream connections.
- `--no-start-on-exit` prevents the Agent Manager Windows service or Unix daemon from starting immediately after the installation.
- `-h` or `--help` lists the arguments available with the installer executable and exits.
- `-v` or `--version` display the Agent Manager version number and exits.

- `-m` or `--javavm` sets the location of the Java virtual machine, for example, the directory that `JAVA_HOME` points to.
 - **NOTE:** Java 7 or lower is not supported.
- `--installer-properties` sets the path to the installer properties file. This file contains the default installer values used during installation. Its contents must be in a Java Properties format and can include a mix of installer arguments and arbitrary properties for use by the installer runtime. All installer options can be defined here but must be prefixed with “`installer.`” and have the “`--`” removed from the argument name. All other argument value setting constraints still apply. For installer arguments that can be declared multiple times, a numeric value must be appended to the end of the property name in order to make the key unique. For example:
 - `installer.fms.0=url=https://localhost:8443,ssl-allow-self-signed=true`
 - `installer.fms.1=url=https://127.0.0.1:8443,ssl-allow-self-signed=true`
 - `installer.headless`
 - `installer.downstream=type=http,port=15872,size=Large`
 - `custom.property.one=true`

If the `--installer-properties` argument is not declared, the installer runtime searches for a file named *installer.properties* in the root directory of the installer binary, or in the extraction root directory of the installer payload. When located, the properties file is automatically loaded during startup.

- `--headless` launches the Agent Manager and configuration interface on the command line. If not specified, the graphical interface is displayed.
- `--auth-token` registers an authentication token during the installation. The token is generated from the Management Server and provides authorization for this Agent Manager to connect.
- `--noservice` prevents the Agent Manager service from being installed on Windows. On UNIX, it prevents the installer from installing an `init.d` script that automatically starts the Agent Manager.
- `--host-display-name` specifies the display name used to identify this Agent Manager instance.
- `--detectha` locates and configures additionally available high availability (HA) servers.

For example, in stand-alone mode:

```
FglAM-<version>-windows-x86_64.exe --silent --fms url=http://serverA:8080
--installdir C:\Quest\Foglight_Agent_Manager --noservice

./FglAM-<version>-linux-x86_64.bin --silent --fms url=http://serverA:8080
--installdir /opt/Quest/Foglight_Agent_Manager --noservice
```

In HA mode:

```
FglAM-<version>-windows-x86_64.exe --silent --fms url=http://serverA:8080/
--fms url=http://serverB:8080/ --installdir C:\Quest\Foglight_Agent_Manager

./FglAM-<version>-linux-x86_64.bin --silent --fms url=http://serverA:8080/
--fms url=http://serverB:8080/ --installdir /opt/Quest/Foglight_Agent_Manager
```

Messages appear in the command-line window while the installer starts. Installer files are extracted to the location you specified in the `--installdir` parameter and the installer runs.

- **NOTE:** Install the Agent Manager on each host that you want to monitor with local agents.

Installing the Agent Manager as a Windows service

If you did not install the Agent Manager as a Windows[®] service using the installer, you can do so from the command line after installing the Agent Manager.

To install the Foglight[®] Agent Manager Windows service:

- 1 Launch a Command Prompt window on the Agent Manager host machine and navigate to the `<fglam_home>\bin` directory.

i | **IMPORTANT:** On Windows 7 and Vista, you must issue the command to manually install the Agent Manager as a Windows service using an administrator version of `cmd.exe` or PowerShell (not just logged in as administrator).

- 2 Issue the following command to install the Agent Manager as a Windows service:

```
fglam --install-service
```

Alternatively, to create multiple Windows services on the same machine, for each Agent Manager service that you want to install, issue a command using the following syntax:

```
fglam --location <STATE> --install-service
```

Where `STATE` is the name of the Agent Manager instance that you want to install as a Windows service. For more information, see [Configuring multiple Agent Manager instances](#) on page 54.

- 3 To start or stop the Foglight Agent Manager service manually, follow the instructions in [To run the Agent Manager as a Windows service](#): on page 35.

To remove the Foglight Agent Manager Windows service, follow the instructions in [To remove the Foglight Agent Manager Windows service](#): on page 36.

Starting or stopping the Agent Manager process

The section below describes basic options for running the Agent Manager. See the *Command-Line Reference Guide* for additional options that you can use with the `fglam` command.

The Agent Manager should be installed in a directory that is local to the system. It should also run using a local account, not a network or domain account. This should also include a local user home directory. Because the Agent Manager monitors and detects problems such as network and disk failures, having the Agent Manager installed in a local directory and running it as a local user makes the Agent Manager more resistant to failures in those services and better able to detect and report those failures. Otherwise, having the Agent Manager installed on a network drive, could cause the Agent Manager to lock itself when the network drive fails, preventing this failure from being reported.

In a default installation, the Agent Manager is installed as a Windows service or a Unix daemon, this process starts immediately after the installation. You can override this default behavior by using the `--no-start-on-exit` option. For more information, see [Installing the Agent Manager using the installer interface](#) on page 15.

To start the Agent Manager:

- Navigate to the `bin` directory of an Agent Manager installation and run the `fglam` executable.

For complete information about the `fglam` command and the available command-line options, see the *Command-Line Reference Guide*.

To restart the Agent Manager:

- In the Foglight browser interface, navigate to the Agent Managers dashboard, select the Agent Manager host, and click **Restart**.

To stop the Agent Manager:

- Navigate to the *bin* directory of an Agent Manager installation and issue the `fglam` executable with the `--stop` option:

```
fglam --stop
```

To run the Agent Manager as a daemon on UNIX®:

- 1 Follow the instructions in [To install the init.d script](#): on page 93. See [Configuring the Agent Manager to run as a daemon](#) on page 93 for more information about this script.
- 2 Navigate to the location in which the *init.d* script `quest-fglam` was installed. See [Locating the init.d script](#) on page 94 for more information.
- 3 Run the script with a required option.
 - To start the Agent Manager daemon, run the `quest-fglam` script with the `start` option. For example (on Linux):

```
/etc/init.d/quest-fglam start
```
 - To stop the Agent Manager daemon, run the `quest-fglam` script with the `stop` option. For example (on Linux):

```
/etc/init.d/quest-fglam stop
```

To run the Agent Manager as a Windows service:

- 1 Follow the instructions in [Installing the Agent Manager as a Windows service](#) on page 34.
- 2 Launch a command-line window on the Agent Manager machine and navigate to the `<fglam_home>/bin` directory.
 - i** | **IMPORTANT:** On Windows 7 and Vista, you must issue the commands to manually start and stop the Agent Manager as a Windows service using an administrator version of `cmd.exe` or PowerShell (not just logged in as administrator).
- 3 Start or stop the Foglight Agent Manager service from the command-line using a required option.
 - To start the Foglight Agent Manager service, run `fglam.exe` with the `--start-service` option:

```
fglam --start-service
```
 - To stop the Foglight Agent Manager service, run `fglam.exe` with the `--stop` option:

```
fglam --stop
```

Identifying the Agent Manager process

The Agent Manager has different process names on different operating systems.

- On Windows® operating systems, the process name is `fglam.exe`.
- On Linux® operating systems, the process name is `Foglight <version>: FoglightAgentManager [Daemon] on <machine_name>`.

About platform-specific identification

The Agent Manager determines a unique ID for each system on which it runs, and includes that ID with the data submission from each agent. On some Linux[®] systems, however, the Agent Manager may be unable to determine a unique system ID. In such cases, the Agent Manager does not return any system ID in the data submission.

Frequently asked questions

How do I upgrade the Agent Manager?

Consult the *Foglight Upgrade Guide* for detailed upgrade instructions.

How do I uninstall the Agent Manager?

This section topic describes how to completely uninstall the Agent Manager and remove the `init.d` script used to run the Agent Manager as a daemon or the Foglight[®] Agent Manager Windows[®] service.

To remove the `init.d` script used to run the Agent Manager as a daemon on UNIX[®]:

i | **IMPORTANT:** Stop the Agent Manager using the `init.d` script, then remove the `init.d` script before uninstalling the Agent Manager.

- 1 Launch a command shell on the Agent Manager machine and navigate to `<fglam_home>/state/default/`.
- 2 Switch to the root user and run the script `fglam-init-script-installer.sh` with the `remove` option:

```
./fglam-init-script-installer.sh remove
```

i | **IMPORTANT:** This script must be run as root.

The setup script `fglam-init-script-installer.sh` removes the `init.d` script `quest-fglam` and all known symlinks to the `quest-fglam` script. See [Installing the Agent Manager using the installer interface](#) on page 15 for the location from which it is removed.

To remove the Foglight Agent Manager Windows service:

i | **IMPORTANT:** Remove the Agent Manager Windows service before uninstalling the Agent Manager.

- 1 Launch a Command Prompt window on the Agent Manager machine and navigate to the `<fglam_home>\bin` directory.

i | **IMPORTANT:** On Windows 7 and Vista, you must issue the command to manually remove the Agent Manager as a Windows service using an administrator version of `cmd.exe` or PowerShell (not just logged in as administrator).

- 2 Run the Agent Manager from the command-line with the `--remove-service` option:

```
fglam.exe --remove-service
```

- 3 Close the Command Prompt window.

To uninstall the Agent Manager:

- 1 Stop the Agent Manager.
- 2 Delete the installation directory (referred to as `<fglam_home>` in this chapter section) and any state directories related to this installation.

Where can I find the Agent Manager installation log files?

The Agent Manager saves its log files in the `<fglam_home>/state/<state_name>/logs` directory.

i | **NOTE:** If you did not specify the `<state_name>` using the `--location` command, the default name is `default`.

In addition to the installation log file, `Install*.log`, that contains messages logged during the installation, the Agent Manager also provides the following log files:

- `FglAM*.log` contains messages logged during run-time.
- `quest-runner*.log` contains messages logged by the external process runner.
- `quest-watchdog*.log` contains information from the Agent Manager self-monitoring watchdog process.
- `auditor/SecurityAudit*.log` contains information about agent activities needed for agent developers and Quest Support.

Are the passwords that are stored in the configuration file, `fglam.config.xml`, encrypted?

Certain passwords that are specified (or stored) in the Agent Manager configuration file `<fglam_home>/state/<state name>/config/fglam.config.xml` are automatically encrypted when the Agent Manager restarts.

The passwords that are encrypted are the ones set as arguments for the `proxy-pass`, `key-password`, `keystore-password`, and `truststore-password` attributes in `fglam.config.xml`. These passwords are encrypted after you manually edit them in `fglam.config.xml` and then restart the Agent Manager.

The passwords stored as arguments for the `key-password`, `keystore-password`, and `truststore-password` attributes are also encrypted when you start an Agent Manager concentrator for the first time after configuring it to communicate with downstream instances using HTTPS. See [Creating a secure connection with downstream instances](#) on page 47 for more information about this type of configuration.

In addition, you can specify the password that the Agent Manager uses when connecting to the Management Server using a proxy during or after the installation (using the Agent Manager configuration interface). This password is stored in an encrypted form in `fglam.config.xml` (as the argument for the `proxy-pass` attribute) when you start the Agent Manager after installation or restart after using the Agent Manager configuration interface. See [Configuring Management Server URLs using the installer interface](#) on page 21 and [Step 10: Change service credentials \[Optional\]](#) on page 28 for more information about setting the proxy password through the Agent Manager installer or configuration interface.

How can I see what parameters are available for the silent installers?

You can see the list of parameters available for the silent Agent Manager installer by running the `FglAM-<version>-<platform>` command with the `--help` option.

Why are two URLs displayed for localhost when I search for HA peers?

If you specify a `localhost` address as the Management Server URL and then search for HA peers (while installing or configuring the Agent Manager), two URLs appear: one that shows the real machine name and one for `localhost`.

For example, you are installing the Agent Manager on *server1*, the same machine on which the Management Server is running. You type `localhost` and `8080` as the host name and port used by the Agent Manager to connect to the Management Server. After you search for HA peers, two URLs are listed: <http://localhost:8080> and <http://server1.example.com:8080>.

I tested the connection between the Agent Manager and a Management Server, but the Management Server URL failed the connectivity test. Why did this happen?

There are several reasons why a Management Server (or Agent Manager concentrator) URL may fail the connectivity test, including the following:

- The Management Server is not running. Verify that the Management Server is running by navigating to the Management Server URL in a browser:

`<http|https>://<management_server_host_name>:<management_server_port>`

If the Management Server is running, the Foglight Login Page appears. If the Login Page does not appear, follow the applicable instructions in “Starting and Stopping the Foglight Management Server” in the *Installation and Setup Guide* for the platform and database you are running to start the Management Server.

- The Agent Manager adapter is either deactivated or uninstalled. Verify that the Agent Manager adapter is running by logging in to Foglight and navigating to **Dashboards > Administration > Agents > Agent Adapters** using the navigation panel.

If the Agent Manager adapter is running, the Agent Manager adapter (named FglAM) is listed in the Agent Adapters table and an Active icon (■) appears in the row for the adapter.

If the Agent Manager adapter is listed but no Active icon (■) appears in the row for the adapter, select the adapter and click **Activate** at the bottom of the dashboard.

If the Agent Manager adapter has been uninstalled (it is not listed in the Agent Adapters table), re-install the Agent Manager cartridge. See [Deploying the Agent Manager cartridge](#) on page 11 for instructions.

- You are attempting to connect the Agent Manager to a concentrator or Management Server over HTTPS, but your certificate has expired. Renew your certificate and test the connection again.
- You are attempting to connect the Agent Manager to a concentrator or Management Server through a proxy configuration but your proxy is rejecting connections. Ensure that the proxy you specify is accepting connections and test the connection again.

Configuring the Agent Manager

This section contains information about configuring the Agent Manager after the installation.

- [Operating system patches](#)
- [Launching the Agent Manager installation interface](#)
- [Configuring the Agent Manager to run in FIPS-compliant mode](#)
- [Configuring the Agent Manager from the command line](#)
- [Configuring the Agent Manager to use SSL certificates](#)
- [Configuring an Agent Manager instance as a Concentrator](#)
- [Configuring the Agent Manager to accept connections from the Management Server](#)
- [Configuring the Agent Manager to execute commands on remote hosts](#)
- [Configuring multiple Agent Manager instances](#)
- [Controlling the polling rate](#)
- [Configuring the Agent Manager to work in HA mode](#)
- [Negotiating Agent Manager resources at runtime](#)
- [Configuring credentials](#)
- [Troubleshooting](#)

Operating system patches

As noted in the *System Requirements and Platform Support Guide*, Foglight® requires that the operating systems on which it runs have all vendor-recommended patches applied for running the Oracle® Java™ Virtual Machine.

Launching the Agent Manager installation interface

You can change many of the settings available in the Agent Manager installer at a later time using the Agent Manager Installation and Configuration interface.

i | **NOTE:** All of the command-line options that you can specify when starting the Agent Manager installer also work with the `--configure` option that is used to launch the Agent Manager Installation and Configuration interface.

To launch the Agent Manager Installation and Configuration interface:

- 1 If you are running multiple instances of the Agent Manager from one installation directory, ensure that no other instances are running from the same directory as the instance you want to configure.

- 2 Stop the Agent Manager instance you want to configure.
- 3 Run the Agent Manager from the command-line with the `--configure` option:

```
<fglam_home>/bin/fglam --configure
```

i | **IMPORTANT:** On Windows 7 and Vista, it is recommended to run the command `fglam --configure` from an administrator version of `cmd.exe` or PowerShell (not just logged in as administrator) if you installed the Agent Manager as administrator.

The Agent Manager Installation and Configuration dialog box appears.

- 4 Review the information in the **Introduction** step, and click **Next**.

The **Host Display Name** step appears. For more information, see [Installing the Agent Manager using the installer interface](#) on page 15.

- 5 In the **Host Display Name** step, set the desired options, if applicable, and click **Next**.

The **Update Client ID** step appears. This step allows you to reset the unique identifier assigned to this Agent Manager. Change the Agent Manager ID if you discover that its ID is the same as another Agent Manager's ID.

i | **IMPORTANT:** Do not reset the Agent manager ID unless Quest Support instructs you to do that.

- 6 If Quest Support instructs you to reset the unique ID assigned to this Agent Manager instance, select the **Yes, reset this client's ID** check box. Otherwise, ensure that the check box is clear.
- 7 Click **Next**.

The **Server URLs** step appears.

- 8 Continue to navigate through the remaining steps and set the desired options, as prompted. Use the **Next** and **Previous** buttons to navigate through the steps. The remaining steps are the same as in the Agent Manager installer interface. For more information, see [Installing the Agent Manager using the installer interface](#) on page 15.
- 9 When you navigate to the **Summary** step, click **Finish**.
- 10 Restart the Agent Manager instance in the desired mode. For information about the available modes, see [Starting or stopping the Agent Manager process](#) on page 34.

Configuring the Agent Manager to run in FIPS-compliant mode

Whether the Agent Manager is FIPS-compliant is determined by the Foglight Management Server from which the Agent Manager installer is downloaded. That is to say if the Agent Manager installer is downloaded from an FIPS-compliant Foglight Management Server, the Agent Manager will be configured to FIPS-compliant automatically, and vice versa.

You can check the value of the property `fips.approved.mode.enabled` in `<fglam_home>/state/default/config/client.config` file to see in which mode this Agent Manager is running. If the property is `True`, it means this Agent Manager is FIPS-compliant, and vice versa. In case the property is not found, it means this Agent Manager is not FIPS-compliant as well.

i | **NOTE:** Do NOT change the value of `fips.approved.mode.enabled` property, otherwise the Agent Manager won't work with the Foglight Management Server if their FIPS-compliant modes are inconsistent.

Configuring the Agent Manager from the command line

If you are running multiple instances of the Agent Manager from one installation directory, before you begin, ensure that no other instances are running from the same directory as the instance that you want to configure.

To launch the Agent Manager configuration command-line interface:

- 1 Stop the Agent Manager instance that you want to configure.
- 2 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/bin` directory.
- 3 Issue the following command:

```
fglam --configure --headless
```

i | **NOTE:** If you want to install Agent Manager as a Windows service, or as a Unix daemon, the configuration interface appears with pre-selected options, indicating that the Agent Manager will start immediately after the installation. You can force these options to appear disabled by default (and enable them, if required, during the installation), if you start the `fglam` command with the `--no-start-on-exit` option:

```
fglam --configure --headless --no-start-on-exit
```

For more information about these options in the installation interface, see [Step 7: Install init.d Script](#) on page 18 (Unix) and [Step 9: Windows Service](#) on page 20.

The **AuthToken** step appears.

- 4 Review the information provided in this step. If you do not want to make any changes to the `auth-token`, press **Enter**.

A message appears while the configuration interface starts. When the configuration interface finishes loading, the **Introduction** step appears.

- 5 Review the information in the **Introduction** step, and press Enter.

The **Host Display Name** step appears. For more information, see [Installing the Agent Manager from the command line](#) on page 22.

- 6 In the **Host Display Name** step, set the desired options, if applicable, and press Enter.

The **Update Client ID** step appears. This step allows you to reset the unique identifier assigned to this Agent Manager if you discover that this Agent Manager is using the same identifier as another Agent Manager.

i | **IMPORTANT:** Do not reset the Agent manager ID unless Quest Support instructs you to do that.

- 7 If Quest Support instructs you to reset the unique ID assigned to this Agent Manager instance, type `Y` at the prompt and press Enter.

Otherwise, accept the default option (`N`), and press Enter.

The **Server URLs** step appears.

- 8 Continue to navigate through the remaining steps and set the desired options, as prompted. The steps are the same as in the Agent Manager installer interface. For more information, see [Installing the Agent Manager from the command line](#) on page 22.

- 9 When you navigate to the **Summary** step, press Enter.

- 10 Restart the Agent Manager instance in the desired mode. See [Starting or stopping the Agent Manager process](#) on page 34 for information about the different modes in which you can run the Agent Manager.

Configuring the Agent Manager to use SSL certificates

You can configure the Agent Manager to communicate with the Management Server using an HTTPS connection.

You can set this option either while installing the Agent Manager, or after installation. See [Installing the Agent Manager](#) on page 14, or [Configuring the Agent Manager](#) on page 40, for more information about configuring the Agent Manager to connect to the Management Server using HTTPS.

By default, Foglight® ships with a self-signed SSL certificate. If you configure the Management Server to use an SSL certificate signed by a third-party Certificate Authority (CA), whose root certificate is already included in the JRE used by the Agent Manager, you do not need to add a new CA to the Agent Manager keystore. Instead, ensure that the Agent Manager connects to the Management Server using HTTPS.

i | **NOTE:** If the root certificate for the third-party CA is not included in the JRE, follow the instructions in [To add a new CA to the Agent Manager certificate store:](#) on page 43.

You must add a new CA to the JRE used by the Agent Manager if:

- You want the Agent Manager to communicate with the Management Server using an HTTPS connection.
and
- The Management Server uses an SSL certificate signed by a private CA. This certificate must be checked for a valid signer.

The Agent Manager includes command-line options for managing certificates in the Agent Manager keystore.

You add a new CA by importing a new root certificate for the CA into the certificate store used by the Agent Manager, as described below.

To check if a third-party CA is included in the JRE:

- 1 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/jre/<jre_version>/jre/bin` directory.
- 2 Issue the following command:

```
keytool --keystore ..\path\cacerts -storepass changeit -list | findstr  
  <3rd_party_signer>
```

For example:

```
keytool --keystore ..\lib\security\cacerts -storepass changeit -list | findstr  
  godaddy  
godaddyclass2ca, Jan 20, 2005, trustedCertEntry,  
godaddyrootg2ca, Jul 18, 2014, trustedCertEntry,
```

To add a new CA to the Agent Manager certificate store:

- 1 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/bin` directory.
- 2 Import a new root certificate for a CA that you want to add. To do that, issue the following command:

```
fglam --add-certificate <alias=/path/to/certificate>
```

i | **IMPORTANT:** The certificate file that you import should be the public certificate for the CA that signed the SSL certificate, not the SSL certificate itself.

- 3 Stop the Agent Manager.
- 4 Configure the Agent Manager to connect to the Management Server using HTTPS without allowing self-signed certificates (since the certificate from the server is no longer considered signed by a private CA).

- a Run the Agent Manager configuration interface. See [Configuring the Agent Manager](#) on page 40, [Configuring Management Server URLs using the installer interface](#) on page 21, or [Step 10: Change service credentials \[Optional\]](#) on page 28.
 - b Configure the Agent Manager to connect to the Management Server using HTTPS.

If you are using the Agent Manager Installation and Configuration interface, select the **Connect using HTTPS** check box.

If you are using the command-line version of the configuration interface, set the start of the `url` parameter to `https` (for example, `url=https://server1.example.com:8443`).
 - c Change the setting for self-signed certificates so that they are not allowed.

If you are using the Agent Manager Installation and Configuration interface, clear the **Allow self-signed certificates** check box.

If you are using the command-line version of the configuration interface, set `ssl-allow-self-signed` to `false`.
- 5 Restart the Agent Manager.
 - 6 If this Agent Manager is a concentrator and you want the Agent Manager instances that connect to it to use HTTPS, follow the instructions in [To create a secure connection between the concentrator and downstream instances](#): on page 47.

Configuring an Agent Manager instance as a Concentrator

A concentrator is an Agent Manager instance that functions similarly to an HTTP proxy. Configure it to accept connections from other Agent Manager instances (called downstream instances) and forward these connections to an upstream target, either the Management Server or another Agent Manager concentrator.

i | **NOTE:** Concentrator ports can also be configured to optimize transmission of control messages from the Management Server to the Agent Manager. For more information, see [Understanding how the Agent Manager communicates with the Management Server](#) on page 10..

Figure 1. A simple Agent Manager concentrator configuration



You can configure one or more Agent Manager instances to act as a concentrator in situations where:

- You do not want the Agent Manager instances on your monitored hosts to connect directly to the Management Server—for example, if you have a large number of Agent Manager instances running in your monitored environment and you want to reduce the number of connections to the Management Server.

i | **IMPORTANT:** Using a concentrator reduces the number of connections to the Management Server, not the volume of data sent to the Management Server.
- The Management Server is not co-located with the monitored hosts and you want to make only a single connection from the remote hosts to the Management Server—for example, if your Management Server is in Sydney and your monitored hosts are in Vancouver, and you only want to make one transpacific connection.
- The Agent Manager instances cannot connect directly to the Management Server (as in the example below).

Example:

Your firewall configuration does not allow the Agent Manager instances on your monitored hosts to connect directly to the Management Server (running on *ManagementServerHost*). However, there is an intermediate host in your environment (*IntermediateHost*) that can accept connections from your monitored hosts and also communicate with the Management Server.

To allow connections from your monitored hosts to be forwarded to the Management Server, you install an Agent Manager instance on *IntermediateHost* and configure it as a concentrator:

- 1 While installing the instance on *IntermediateHost* (using the GUI installer), you specify the host name and port (*ManagementServerHost* and 8080) of the Management Server to which you want this concentrator to connect in the Configure Server URLs step.
- 2 When the installation is complete, you ensure that the instance is shut down and configure it as a concentrator by editing its *fglam.config.xml* file so that it listens for connections from downstream instances on a specified port (8081).
- 3 You restart the Agent Manager instance on *IntermediateHost*. This instance is now configured as a concentrator: it listens for connections from downstream instances on port 8081 and forwards data to the Management Server on port 8080.

Now that the concentrator is set up on *IntermediateHost*, you configure the Agent Manager instances on the monitored hosts to connect to the concentrator:

- 1 After stopping these Agent Manager instances, you run the Agent Manager configuration interface for each instance.
- 2 In the configuration interface, you specify the concentrator's host name and the port on which it is listening (*IntermediateHost* and 8081) when setting the URL to which the instances connect.
- 3 Once the instances are re-configured to connect to the concentrator, you restart them.

The Agent Manager instances on the monitored hosts can now connect to the Management Server through the concentrator. You can also perform agent management tasks from the Management Server, such as deploying agent packages to the monitored hosts and creating new agent instances on them. There is no indication that the downstream instances are not connected directly to the Management Server.

Configuring the concentrator

This section describes how to configure the concentrator to connect to the upstream target (either the Management Server or another Agent Manager concentrator) and to listen for connections from downstream Agent Manager instances.

A concentrator's upstream connection is independent of the downstream connections. For example, several Agent Manager instances on a local subnet can communicate to a concentrator using HTTP while the concentrator forwards requests over an non-secure network to the Management Server using HTTPS (or the other way around).

CAUTION: Do not set up a concentrator to forward messages to itself or create any kind of loop or cycle. This causes the Agent Manager to indefinitely feeds messages through the loop.

To configure a concentrator to connect to the upstream target:

You can configure the concentrator to connect to the upstream target in different ways:

- **Using HTTP:** Set the upstream target of the concentrator in the same way you typically set the Management Server URL:
 - During installation, use the Agent Manager Installation and Configuration interface, command-line interface, or silent installer. See [Installing the Agent Manager](#) on page 14 for more information.
 - After installation, use the Agent Manager configuration interface. See [Configuring the Agent Manager](#) on page 40 for more information.

- **Using HTTPS:** To configure a concentrator connection to the Management Server using a secure connection, follow the instructions in [Configuring the Agent Manager to use SSL certificates](#) on page 43.

Between connections, the Agent Manager collects all upstream and downstream messages in queues. Queuing messages prevents them from getting lost in the event of a disconnection.

When running the Agent Manager as a concentrator, you must increase the default disk cache sizes.

CAUTION: Increasing disk cache sizes is an advanced activity that involves custom tuning. The size of cache that you need depends on the number of hosts that connect to the concentrator and the type of load they put on the concentrator. The 1 GB value shown below is an example only.

To configure the size of the disk cache used to store messages:

- 1 Open the `<fglam_home>/state/<state name>/config/fglam.config.xml` file for editing.
- 2 Locate the `<queue-sizes>` XML element.
- 3 Edit the `<upstream/>` and `<downstream/>` blocks that appear after the `<documentation>` block:

- Change the argument for the `max-disk-space` attribute in both the `<upstream/>` and `<downstream/>` blocks to a value larger than the default setting (1024 KB). For example, to change the default disk cache size to 1 GB, set the `max-disk-space` attribute in both the `<upstream/>` and `<downstream/>` blocks as follows:

```
max-disk-space="1048576"
```

The `max-disk-space` argument sets the amount of disk space (in KB) that can be used to store messages.

- 4 Save your changes to the `fglam.config.xml` file.
- 5 Restart the concentrator.

To configure a concentrator to listen for connections from downstream instances:

NOTE: If you want to create a secure connection between the concentrator and downstream Agent Manager instances using HTTPS, follow the instructions in [Creating a secure connection with downstream instances](#) on page 47.

- 1 Open the `<fglam_home>/state/<state name>/config/fglam.config.xml` file for editing.
- 2 Locate the `<http-downstreams>` XML element.
- 3 After the `<documentation>` block, add an `<http-downstream/>` child element:

```
<config:http-downstream port="port_number" [address="network_address"]/>
```

- a Replace `port_number` with an available port number. This is the port on which the concentrator listens for connections from downstream Agent Manager instances.
- b **Optional.** If required, you can also bind the concentrator to single network address. To do so, include the attribute `address="network_address"` in the `http-downstream` child element (shown as an optional attribute in Step 3), replacing `network_address` with the network address where you want the concentrator to receive connections from the downstream instances.

The optional `address` attribute is useful when a machine has two or more network addresses, and you want the connections to the Management Server to go out from one, and the connections from the downstream instances to come in to another.

- 4 If required, configure the concentrator to listen for connections on multiple different ports by adding additional `<http-downstream/>` elements and setting the port number (and, optionally, the network address), as described above.

This is useful in situations where the concentrator machine has multiple network connections, and you want the concentrator to listen on different network connections with different ports. For example, the concentrator listens on network connection 1 with port 8081 and listens on network connection 2 with port 8082.

- 5 Restart the concentrator.

Configuring downstream Instances

This section describes how to configure the downstream Agent Manager instances to connect to the concentrator.

To configure the downstream instances to connect to the concentrator:

- 1 Configure downstream instances in the same way you configure any Agent Manager:
 - During installation, use the Agent Manager GUI, command-line, or silent installer. See [Installing the Agent Manager](#) on page 14 for more information.
 - After installation, use the Agent Manager configuration interface. See [Configuring the Agent Manager](#) on page 40 for more information.
- 2 At the step where you set the URL for Management Server, specify the host name and port for the concentrator instead of the Management Server.
- 3 If you want the downstream instance to connect to the concentrator using a secure connection, follow the instructions in [Creating a secure connection with downstream instances](#) on page 47.

Creating a secure connection with downstream instances

The following procedure can be used to create a secure connection between the concentrator and downstream Agent Manager instances using HTTPS.

To create a secure connection between the concentrator and downstream instances:

- 1 Stop the Agent Manager concentrator.
- 1 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/jre/<jre_version>/<jre>/bin/` directory.
- 2 If you do not already have an SSL certificate for the concentrator host, you can create a self-signed certificate by executing the following command on the concentrator machine, where `<password>` is replaced with your desired password:

```
keytool -genkeypair -alias fglam -keystore </path/to/fglam.keystore> -storepass <password>
```
- 3 Respond to the prompts from `keytool`. Only the “first and last name” are required, all other fields can be left blank. The “first and last name” form the common name (CN) for this key pair and this needs to be provided to the Management Server (for reverse polling) or downstream Agent Managers (as the `ssl-cert-common-name`). You can type anything you want into this field, but the host name is the most common choice. The default value, if left blank, is `Unknown`.
- 4 Press Enter to set the key password to the same value as the keystore.
- 5 Open the file `<fglam_home>/state/<state name>/config/fglam.config.xml` for editing.
- 6 Between the existing `<http-downstreams>` and `</http-downstreams>` tags, add an `<https-downstream/>` child element:

```
<https-downstream key-password=<password> keystore=<path_to_keystore>
port=<port_number> [address=<network_address>] />
```

Where:

- `<password>` is the same password you specified in [Step 2](#) for `-storepass`. The password is saved in an encrypted form in `fglam.config.xml` when you restart the Agent Manager.

- `<path_to_keystore>` is the path to the Agent Manager keystore.
- `<port_number>` is the port number on which you want the concentrator to listen for connections from downstream Agent Manager instances.
- `<network_address>` is the network address, to which the concentrator is bound when receiving connections from the downstream instances. This argument is optional. It is useful when a machine has two or more network addresses and you want the connections to the Management Server to go out from one, and the connections from the downstream instances to come in to another.

i | **IMPORTANT:** Other optional attributes are available for the `<https-downstreams>` element. See the file `fglam.config.xml` for details.

- 7 If required, configure the concentrator to listen for connections on multiple different ports by adding additional `<https-downstream/>` elements and setting the arguments as described above.

This is useful in situations where the machine on which the concentrator runs has multiple network connections and you want the concentrator to listen on different network connections with different ports. For example, the concentrator listens on network connection 1 with port 8081 and listens on network connection 2 with port 8082.

- 8 Restart the Agent Manager concentrator.
- 9 Configure each downstream Agent Manager instance that connects to the Agent Manager concentrator using this secure connection so that:
 - It connects using HTTPS.
 - It accepts self-signed certificates.
 - It accepts a certificate with an unexpected common name (host name), and the common name for the certificate is set to `Unknown`.

See [Configuring Management Server URLs using the installer interface](#) on page 21 or [Step 10: Change service credentials \[Optional\]](#) on page 28 for information about these parameters, which you can set through the Agent Manager installer or configuration interface.

i | **NOTE:** It is not recommended to enable the `ssl-allow-self-signed` configuration when the downstream Agent Manager is running in FIPS-compliant mode. If this configuration is disabled, you have to add the concentrator's certificate to the downstream Agent Manager's keystore in order to connect to the concentrator using HTTPS.

To export certificate from concentrator:

1. Locate the element `<config:http-downstream>` in `<fglam_home>/state/default/config/fglam.config.xml` file on concentrator Agent Manager, and get the path of the keystore corresponding to the downstream Agent Manager. If it is a relative path, it is relative to the path of "`<fglam_home>/state/default`".
2. Launch a command shell and navigate to the `<fglam_home>/jre/<jre_version>/jre/bin` directory.
3. Issue the following command to export concentrator's certificate:

```
keytool -exportcert -noprompt -rfc -alias fglam-cert -file <exported-cert-
filename> -keystore </path/to/keystore> -storepass <key-password> -
storetype BCFKS -providerpath "<fglam_home>\client\<build-version>\lib\bc-
fips.jar" -providername BCFIPS -providerclass
org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider
```

To import the exported certificate to downstream Agent Manager:

1. Launch a command shell and navigate to the `<fglam_home>/bin` on the downstream Agent Manager.

2. Issue the following command to import certificate:

```
fglam --add-certificate <alias=/path/to/exported-cert-filename>
```

Excluding SSL ciphers from upstream or downstream Connections

You can exclude SSL cipher suites from both upstream Agent Manager connections (to the Management Server or an Agent Manager concentrator), or downstream connections (as a concentrator).

If you need to exclude one or more ciphers from the SSL encryption used for SSL connections, you can do so using one or more `excluded-ssl-cipher` elements in the `fglam.config.xml` file. For example, you may want to exclude lower encryption strength ciphers, or ciphers with security vulnerabilities.

To exclude specific SSL ciphers from an upstream connection:

- 1 Open the `<fglam_home>/state/<state name>/config/fglam.config.xml` file for editing.
- 2 Between the existing `<config:http-upstreams>` and `</config:http-upstreams>` tags, add an `<config:http-upstream/>` child element:

```
<config:http-upstream url="https://secure_server_URL:port_number">
  <config:excluded-ssl-cipher name="SSL_RSA_WITH_RC4_128_MD5"/>
  <config:excluded-ssl-cipher name="SSL_RSA_EXPORT_WITH_DES40_CBC_SHA"/>
  <config:excluded-ssl-cipher name="SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA"/>
</config:http-upstream>
```
- 3 Save your changes.
- 4 Restart the Agent Manager.

To exclude specific SSL ciphers from a downstream connection:

- 1 Open the `<fglam_home>/state/<state name>/config/fglam.config.xml` file for editing.
- 2 Between the existing `<config:http-downstreams>` and `</config:http-downstreams>` tags, add an `<config:https-downstream/>` child element:

```
<config:https-downstream port="port_number">
  <config:excluded-ssl-cipher name="SSL_RSA_WITH_RC4_128_MD5"/>
  <config:excluded-ssl-cipher name="SSL_RSA_EXPORT_WITH_DES40_CBC_SHA"/>
  <config:excluded-ssl-cipher name="SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA"/>
</config:https-downstream>
```
- 3 Save your changes.
- 4 Restart the Agent Manager.

i | **NOTE:** For a complete list of SSL cipher suite names, consult your vendor-specific JRE documentation.

Excluding Specific SSL Protocols from Downstream Connections

If you need to exclude one or more protocols from the SSL protocol negotiation, you can do so using one or more `excluded-ssl-protocol` elements. Some common values are `SSLv2Hello`, `SSLv3`, `TLSv1`, `TLSv1.1`, `TLSv1.2`.

If none are specified, then `SSLv2Hello` and `SSLv3` are disabled by default. Otherwise only those protocols listed will be excluded.

To exclude specific SSL protocols from a downstream connection:

- 1 Open the `<fglam_home>/state/<state name>/config/fglam.config.xml` file for editing.
- 2 Between the existing `<config:http-downstreams>` and `</config:http-downstreams>` tags, add an `<config:https-downstream/>` child element:


```
<config:https-downstream port="8443">
  <config:excluded-ssl-protocol name="SSLv2Hello"/>
  <config:excluded-ssl-protocol name="SSLv3"/>
  <config:excluded-ssl-protocol name="TLSv1"/>
</config:https-downstream>
```

- 3 Save your changes.
- 4 Restart the Agent Manager.

Configuring the Agent Manager to accept connections from the Management Server

You can configure the Foglight® Agent Manager to accept connections from the Management Server and enable reverse data polling. This is useful in situations when the Agent Manager cannot connect to the Management Server due to its location. For example, when the Agent Manager is located in the cloud and the Management Server runs on premises, the Agent Manager has no means to connect to the Management Server and pass on the collected data. Another example is when the Agent Manager resides in a demilitarized zone (DMZ), exposed to untrusted networks, and the Management Server is behind a firewall.

To enable this feature, you must instruct the Agent Manager to accept connections from the Management Server in order to facilitate normal message passing and data polling.

i | **NOTE:** Only the Management Server can reverse poll an Agent Manager instance. An Agent Manager instance cannot reverse poll another Agent Manager instance.


You do this by performing the following steps:

- Configure downstream SSL connections using the configuration interface. You can do that either using the installer interface (see [Step 8: Downstream Connection Configuration](#) on page 19), or the command-line interface (see [Step 6: Downstream Connection Configuration](#) on page 26).

Or:

Configure downstream non-SSL connections, or connections requiring user-provided certificates or keystores. For instructions, see [To configure non-SSL connections or connections using user-provided certificates or keystores](#).

- Using the *fglam.config.xml* file, disable upstream connections to the Management Server. For instructions, see [To prevent the Agent Manager from connecting to the Management Server](#).
- On the Management Server, configure the FglAM Adapter, to instruct the Management Server to connect to this Agent Manager. For instructions, see [To configure the Management Server to connect to the Agent Manager](#).



To prevent the Agent Manager from connecting to the Management Server:

- 1 Open the *fglam.config.xml* file for editing. This file is located in the `<fglam_home>/state/default/config` directory.
- 2 In the *fglam.config.xml* file, locate the `<config:http-upstreams>` XML element, and within that element, declare a new `<config:http-upstream>` element using the following lines of code:

```
<config:http-upstream>
  <no-connection/>
</config:http-upstream>
```

The `no-connection` element prevents the Agent Manager from connecting to the upstream Management Server.

- 3 Save your changes and restart the Agent Manager.

To configure non-SSL connections or connections using user-provided certificates or keystores:

- 1 Open the *fglam.config.xml* file for editing. This file is located in the `<fglam_home>/state/default/config` directory.
- 2 In the *fglam.config.xml* file, locate the `<config:http-downstreams>` XML element, and within that element, declare a new `<config:http-downstream>` sub-element for a non-SSL connection or `<config:https-downstream>` for an SSL connection.
- 3 **Non-SSL connections only.** Within the newly created `<config:http-downstream>` element, provide a port number that the Agent Manager will use to listen for incoming connections, and optionally the IP address of the network interface. For example:

```
<http-downstream port="9090" address="127.0.0.1"/>
```

- 4 **User-provided certificates or keystores only.** Within the newly created `<config:https-downstream>` element, provide the information about the certificate and keystore you want to use. There is a wide range of attributes that you can use. For complete instructions, review the `<config:documentation>` element under `<config:http-downstreams>`.

NOTE: When Agent Manager runs in FIPS-compliant mode, only the BCFKS keystore type can be used to store the key pair. Follow below steps to generate a key pair and BCFKS keystore on Agent Manager, and then import its certificate to Management Server for SSL connections:

- 1 Launch a command shell and navigate to the `<fglam_home>/jre/<jre_version>/jre/bin` directory on Agent Manager. And then issue the following command to generate the keypair and BCFKS keystore.

```
keytool -genkeypair -noprompt -keyalg RSA -keysize 2048 -sigalg
SHA256withRSA -dname "CN=<fglam_host_name>" -validity 365 -alias
<keypair_alias_name> -keystore </path/to/fglam.kestyore> -storepass
<password> -storetype BCFKS -providerpath <fglam_home>\client\<build-
version>\lib\bc-fips.jar -providername BCFIPS -providerclass
org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider
```

- 2 Export the certificate from BCFKS keystore:

```
keytool -exportcert -noprompt -rfc -alias <keypair_alias_name> -file
</path/to/exported-cert-filename> -keystore </path/to/fglam.kestyore> -
storepass <password> -storetype BCFKS -providerpath
<fglam_home>\client\<build-version>\lib\bc-fips.jar -providername BCFIPS
-providerclass org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider
```

- 3 Import the certificate to Management Server:

- a Locate the property 'Trust Store' in **Administration > Setup > Management Server Configuration** dashboard, and get the path of current trust store used by Management Server.
- b Issue the following command to import the certificate to Management Server:

- The JRE cacerts is the default trust store if Management Server runs in non-FIPS mode. Issue the following command to import the certificate to Management Server:

```
keytool -import -alias <alias_name> -file </path/to/exported-cert-filename> -keystore <fms_home>/jre/lib/security/cacerts -storepass changeit
```
- The trust.fips.keystore is the default trust store if Management Server runs in FIPS-compliant mode. Issue the following command to import the certificate to Management Server:

```
keytool -import -alias <alias_name> -file </path/to/exported-cert-filename> -keystore <fms_home>/config/security/trust.fips.keystore -deststoretype BCFKS -provider org.bouncycastle.jcajce.provider.BouncyCastleFipsProvider -providerpath <fms_home>/server/core/bc-fips.jar -storepass nitrogen
```

To configure the Management Server to connect to the Agent Manager:

- 1 Log in to the Foglight browser interface and navigate to the Agent Properties dashboard.
- 2 On the Agent Properties dashboard, under **Agent Type**, select **FglAM Adapter**, and in the pane on the right, click **Edit**.
- 3 In the **Agent Type Properties** dialog box that appears, under **Hosts to Pull Data From**, click **Edit**.
- 4 In the **Edit List of Hosts to Pull Data From** dialog box that appears, click **Add**.
The table refreshes, showing a newly added row.
- 5 Specify the connection information to the Agent Manager that you want the Management Server to connect to by populating any of the following columns, as required.
 - **Enabled:** Select this check box if you want the Management Server to connect to this Agent Manager.
 - **URL:** Type the URL the Agent Manager uses to communicate with the Management Server.
 - **Local Address:** To specify a local network address for the Management Server connection to the Agent Manager, type the IP address of a NIC (network interface card) on the machine hosting the Agent Manager required to establish connections to the Management Server.
 - **Proxy URL:** If you want the Management Server to connect to the Agent Manager using a proxy, type the URL of the proxy server.
 - **Proxy NTLM Domain:** If you are using a proxy server for communication, and the proxy uses Windows authentication, type the Windows domain.
 - **Proxy User Name:** If you are using a proxy server for communication, type the user name needed to access the proxy server.
 - **Proxy Password:** If you are using a proxy server for communication, type the password associated with the user name.
 - **Allow Self Signed SSL Certificates:** Select this check box if you want to enable the Management server to accept self-signed-certificates from the Agent Manager. It is not recommended to enable this configuration in FIPS-compliant mode for security consideration. When Management Server is running in FIPS-compliant mode, you need to add the Agent Manager's public certificate to Management Server's jre keystore. For more information, see [To configure non-SSL connections or connections using user-provided certificates or keystores:](#) on page 51.
 - **SSL Certificate Common Name:** If you want to enable the Management Server to accept self-signed certificates from the Agent Manager, and the certificate has a different common (host) name than the one reported by the Agent Manager, type the certificate common name.
 - **Compressed Connection:** Select this check box if you want the Management Server to establish HTTP-compressed communication with the Agent Manager.

- **Chunked HTTP Connection:** Select this check box if you want to use an HTTP connection with chunked transfer encoding enabled.

Configuring the Agent Manager to execute commands on remote hosts

ExecuteCommandOnRemoteHostsAction is an action that is used for the Agent Manager to execute a command on remote hosts.

To configure this action for the Agent Manager:

NOTE: Both Adapter and the Agent Manager must be upgraded to version 5.9.2 or later before the following procedure.

- 1 Add the *ExecuteCommandOnRemoteHostsAction* action to your rule list. For more information about how to add or edit an action in a rule, refer to the *Getting Started: create a new rule* or *Getting Started: view and edit rule definitions* section in the *Foglight Administration and Configuration Help*.
- 2 Specify the following parameters for the action as needed, then save your changes:

The screenshot shows the 'Condition & Actions' configuration page for a rule named 'Email rule'. The 'Action' tab is selected, showing the 'ExecuteCommandOnRemoteHostsAction' configuration. Below the description, there is a table of action parameters:

| Name | Default Value | Required | Type | Value | Parameter Description |
|---------------------------|---------------|-----------|---------|---------|--------------------------------|
| AgentManagerName | Undefined | mandatory | String | Default | The Agent Manager uses t... |
| CommandLine | Undefined | mandatory | String | Default | Command to be executed ... |
| RemoteHosts | Undefined | mandatory | String | Default | Target host address and p... |
| AgentManagerNameUseRegExp | false | optional | String | Default | A flag indicating whether t... |
| MatchAll | false | optional | String | Default | A flag indicating whether t... |
| Port | 22 | optional | Integer | Default | SSH connection port, defa... |
| Timeout | 0 | optional | Long | Default | Command executing timeo... |

- **AgentManagerName (Mandatory):** The Agent Manager name that is delegated to invoke the command on remote hosts.
- **CommandLine (Mandatory):** The command to be executed remotely.
- **RemoteHosts (Mandatory):** Target host addresses and platforms map, for example, *hostName=Windows!hostIP=Linux*. The platform value can be either of the following: *Windows* or *Linux*.
- **AgentManagerNameUseRegExp (Optional):** The flag indicating whether the *AgentManagerName* parameter should be treated as regular expressions. This value is either true or false.
- **MatchAll (Optional):** The flag indicating whether to run the command on all Agent Managers that matches the selection criteria. If set to false, the command will be executed only on the first matching Agent Manager.
- **Port (Optional):** SSH connection port, default is 22. Windows platform does not need to configure this value.

- *Timeout* (Optional): Command executing timeout value. Default is 0, which will use the Agent Manager default timeout value.
- 3 Go to the *Dashboards > Administration > Credentials > Manage Credentials* dashboard, add credentials for those remote hosts. The credentials usage must set to either “*Execute Command On Remote Hosts For Windows*” or “*Execute Command On Remote Hosts For Unix*”.

After saving your configurations and the rule is triggered, you will see the executed results on the server console or logs.

Configuring multiple Agent Manager instances

The files related to the Agent Manager’s run-time state (for example, configuration and log files), are saved in the `<fglam_home>/state/` directory tree. Under the *state* directory, there is a sub-directory for each available Agent Manager instance.

You can run multiple instances of the Agent Manager using a single Agent Manager *bin* directory.

CAUTION: Running multiple instances of the Agent Manager on the same host is intended for cluster and failover support only and is not intended to be used for any other purposes. Installing separate instances of the Agent Manager on the same system is therefore neither recommended nor supported.

In this type of configuration, you create multiple instances and each instance uses a different *state* directory but runs from a single Agent Manager *bin* directory. One example of this type of configuration is to test new agent settings without making changes to the agents you are currently using to monitor your production environment.

CAUTION: In a configuration where multiple Agent Manager instances share one *bin* directory, each time you deploy agents, upgrade agents, or deploy Agent Manager upgrades to one instance, all other instances are also updated.

You can also configure multiple physical installations of the Agent Manager to use a corresponding state directory that exists on a single shared drive. One example use of this functionality is running the Agent Manager in cluster environments. See [Example: Running multiple instances in a cluster environment](#) on page 55 for more information.

As described below, you create a new instance (and its associated state sub-directory) by including the `--create-state` and `--location "<state_name>"` (or `-l "<state_name>"`) options with the `fglam` command; you then use the `fglam --location "<state_name>"` command to run that new instance.

CAUTION: The `--create-state` and `--location` command-line options are not recommended for use with the embedded Agent Manager, because the embedded Agent Manager does not include its own JRE. Instead, it is configured by the Management Server to use the same JRE that the Management Server uses. Since the Management Server is not aware of additional state directories that you may have created, it cannot properly configure them. Instead, you need to manually configure and maintain the JVM location in the Management Server installation.

If you do not create multiple instances (by following the instructions below), the Agent Manager creates an instance called *default* and stores the state files in the `state/default` directory.

IMPORTANT: On UNIX® platforms, the entire Agent Manager installation — including all state directories — must be owned by the same system user.

To create a new Agent Manager instance:

- 1 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/bin` directory.
- 2 Issue the following command:

```
fglam --create-state --location "<state_name>"
```

Where `<state_name>` is the name of the new instance.

A new state directory is created in `<fglam_home>/state/<state_name>`.

To run an Agent Manager instance:

- 1 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/bin` directory.
- 2 Issue the following command:

```
fglam --location "<state_name>"
```

Where `<state_name>` is the same instance name you specified above.

When you deploy agents to the instance, files related to the run-time state for these agents (including log files) are stored under the `<fglam_home>/state/<state_name>/agents` directory for that instance.

Example: Running multiple instances in a cluster environment

i | **IMPORTANT:** This section is provided as an example only. The commands shown are fully supported, but the example itself does not describe a configuration that is supported as part of the product warranty.

One example use of this functionality is running the Agent Manager in cluster environments, since it allows the next assigned host in the cluster to relaunch an Agent Manager instance—and the specific agents it manages—when cluster failover occurs.

In this type of installation, there are multiple physical installations of the Agent Manager on different failover nodes. When one node fails and shuts down, the next one starts and its Agent Manager instance accesses the latest changes stored in its `state` directory on a shared drive.

The process of running multiple Agent Manager instances in a cluster environment follows the outline presented below.

Part 1: Install the Agent Manager on each node in the cluster

Begin by installing the Agent Manager on each node in your cluster. See [Installing the Agent Manager](#) on page 14 for installation instructions.

Part 2: Initialize a state directory on the shared drive

When the Agent Manager installation is available to the nodes in the cluster, the next step is to initialize a state directory for an instance on the shared drive that is used by the cluster. When setting the state location locally from one of the nodes, you must define the full path to the remotely-mounted state directory.

In the following example, `<state_dir>` is a path to a state directory on a shared network server that is accessible locally from all machines. For example: on Windows clients, the `<state_dir>` can be `f:\cluster_shared_dir\fglam_states\STATENAME_A`, while on UNIX® clients, it is `/mnt/cluster_shared_dir/fglam_states/STATENAME_A`.

To set the state location, use the following command:

```
fglam --create-state --location <state_dir>
```

i | **NOTE:** When you provide a full path with the `--location` parameter, the state directory can be located anywhere on the shared drive, it does not need to be in a subdirectory of an Agent Manager installation.

Part 3: Launch the Agent Manager from the active node

Run the Agent Manager from the active node and provide the full path to this instance's state directory on the shared drive. For example:

```
fglam --location <state_dir>
```

The files related to the Agent Manager instance's run-time state—for example, configuration and log files—are stored under its remote *state* directory on the shared drive.

Part 4: Deploy agents and create agent instances

When the Agent Manager is running, you can deploy agents to it and create agent instances. Files related to the run-time state for these agents (including log files) are stored under the remote *state* directory for this Agent Manager instance. Using the example above, they are stored in <state_dir>.

Part 5: Ensure exclusive access to the shared state directory

Ensure that only one instance of the Agent Manager that uses a particular *state* directory is running at a time. Do not run two instances of the Agent Manager on separate machines (or separate active nodes in the cluster) and cause these instances to use the same shared *state* directory simultaneously.

Controlling the polling rate

The FglAMAdapter, a component included with the Management Server, controls how often the connected downstream Agent Managers and agent instances connect and poll for messages. In general, the more hosts that are connected to the server, the less often they should be instructed to poll. The properties included with the FglAMAdapter control the polling behavior. They can be found on the Agent Properties dashboard. In most cases, changes to these properties are not required. Doing so is only recommended when instructed by Quest Support.

The polling rate is controlled by the following properties:

- **Minimum Polling Interval (seconds):** The minimum polling interval, in seconds.
- **Maximum Polling Interval (seconds):** The maximum polling interval, in seconds.
- **Polling Timeout (seconds):** A time-out/grace period (in seconds) that the FglAMAdapter waits for a host to respond, before considering it as disconnected. This is used to account for clock skews and changes in timing typically seen on heavily loaded VMware images.

For more information about the Agent Properties dashboard, see the *Administration and Configuration Help*.

Configuring the Agent Manager to work in HA mode

High Availability (HA) mode is a configuration in which multiple Agent Managers work together in an HA Partition, where one Agent Manager is a primary host (HA Primary), and others are standby hosts (HA Peers). When configured, agent instances whose types are configured as HA Aware and belong to the same HA Partition are managed by the HA Primary host. If that Agent Manager stops responding or goes offline, the agent instances fail over to another Agent Manager.

Under a configured HA Partition, a common deployment set of agent types is kept in sync across all HA Peers. Agent packages deployed to that HA Partition are checked for any HA Aware agent types. Any detected HA Aware types are automatically deployed to all other HA Peers.

The FglAM Adapter monitors the deployments of the each Agent Manager host within the named HA Partition. The HA Primary is considered the master in terms of the deployment set and automatically deploys (or undeploys) HA

Aware cartridges to each HA Peer. This also happens during cartridge upgrades, when the Adapter automatically pushes out the updates to all of the HA Peers in that HA Partition.

CAUTION: Do not create Agent Manager HA clusters on Agent Manager installations that already contain agent packages. Doing so causes any HA-aware packages, that only exist on the secondary node, to become automatically undeployed. This is because the primary node controls the deployment all HA-aware packages.

Assigning Agent Managers to HA partitions

HA mode is configured through the FglAM Adapter agent properties. You can use these properties to assign an Agent Manager to HA partitions, and define the priorities for promoting HA Peers to HA Primary hosts.

Start by navigating to the Agent Properties dashboard, the **FglAM** namespace, and the **FglAMAdapter** properties. From there, you can edit the **High Availability Host Config** list-based property to assign an Agent Manager to HA Partitions and define their eligibility for becoming HA Primary hosts.

To assign an Agent Manager to an HA Partition:

- 1 Ensure that the Agent Manager is connected to the Management Server.
- 2 Log in to the Foglight browser interface.
- 3 On the navigation panel, under Dashboards, navigate to **Administration > Agents > Agent Properties**.
- 4 On the Agent Properties dashboard that appears in the display area, in the **Namespace/Type** view, expand the **FglAM** node, and click the **FglAMAdapter** node.
- 5 Assign the Agent Manager to a desired HA Partition by editing its entry in the High Availability Host Config list. This list contains all Agent Managers that are currently connected to the Adapter, and is accessible through the **High Availability Host Config** property. The list also identifies the names of their respective HA Partitions, and the priorities for considering Agent Managers as potential HA Primary hosts.

- a Get started with editing the **High Availability Host Config** list-based property.

In the **Properties** view, under **High Availability**, on the right of the **High Availability Host Config** property, click **Edit**.

- b Assign the Agent Manager to an HA Partition.

In the dialog box that appears, locate the Agent Manager entry that you want to add to an HA Partition, and in its **HA Partition Name** column, type the name of that HA Partition. Do not add an Agent Manager to this list if the client you want to assign the HA Partition is not listed here. Agent Managers that connected to this Management Server will be automatically added to this list. Manually adding Agent Managers is not supported.

To assign more Agent Managers to this or other HA Partitions, repeat this step.

- c Define the priority used to consider this Agent Manager as an HA Primary host.

When promotions to the HA Primary role are evaluated, the FglAM Adapter gives preference to the HA Peer with the highest priority. Acceptable values are in the minus ten to ten range. Use them to decrease or increase the likelihood of an HA Primary role assignment. Type the desired value in the **Priority** column, or leave the default value of zero.

To define the priorities of other Agent Managers, repeat this step.

- d Click **Save Changes**.

TIP: To get an overview of the HA Configuration in real time (for example, to see the current HA Primary hosts, and the HA State of each HA Peer), you can view the `HAManagerMBean`.

Using the JMX-Console, click `FglAM:name=HAManager` and invoke the `diagnosticSnapshotAsString()` method. The resulting output lists each of the known Agent Managers, which (if any) HA Partition they are assigned to, what deployment set they have, who is the HA Primary and what HA State they are in.

Adding cartridges to the HA deployment whitelist

In addition to assigning an Agent Manager to an HA Partition, real HA activities do not take place until HA Aware agent types are deployed.

To provide eligibility for HA agent fail-over to existing agents that primarily function as remote monitors, the **HA Deployment Whitelist** property is included. This list-based property identifies prior cartridge deployments that were not originally marked as eligible for HA agent fail-over. In this list, each row entry contains a regular expression pattern that matches the agent package, version and agent types that are permitted to be managed in an HA Partition.

Understanding the criteria for HA eligibility

The criteria for inclusion into this list is that the agent type must be capable of monitoring remote resources. It does not make sense to enable HA fail-over for an agent that is collecting data that can only be retrieved from the Agent Manager host on which it is running. If such an agent is added to the Whitelist, when the host fails over, the agent can potentially start monitoring the resources of the newly promoted HA Primary, which can result in unexpected outcome when viewing the dashboards that display the collected metrics.

i | **IMPORTANT:** To effectively be eligible for fail-over functionality, agents must be using one of the `RemoteConnection` Services available in the Agent Manager Development Kit.

To add an entry to the Whitelist:

- 1 Ensure that the Agent Manager is connected to the Management Server.
- 2 Log in to the Foglight browser interface.
- 3 On the navigation panel, under Dashboards, navigate to **Administration > Agents > Agent Properties**.
- 4 On the Agent Properties dashboard that appears in the display area, in the **Namespace/Type** view, expand the **FglIAM** node, and click the **FglIAMAdapter** node.
- 5 Add a new regular expression to select a desired agent type and add it to the Whitelist.
 - a Get started with editing the **HA Deployment Whitelist** property.

In the **Properties** view, under **High Availability**, on the right of the **HA Deployment Whitelist** property, click **Edit**.
 - b Add a new regular expression to the table.

In the dialog box that appears, click **Add Row**, and in the new row that appears, type a regular expression using the following syntax:

```
[Agent_package]/[Agent_version]/[Agent_type]
```

i | **IMPORTANT:** The elements in the expression must be separated with slashes '/
 - c Click **Save Changes**.
- 6 If the cartridge containing the specified agent type is not already installed, install it using the Cartridge Inventory dashboard, and deploy its agent package to an Agent Manager that is assigned to an HA partition.

About agent fail-over

When the HA Primary host goes off-line, a schedule is registered that delays the transfer of the agents (and the promotion of another HA Primary) for ten minutes. This delay is in place to allow the HA Primary host to be restarted and retain its HA status and assigned agents. If you find that the duration is too long, the following startup parameter can be set when launching the Management Server:

```
-Dfglam.ha.agent.transfer.delay.milliseconds=XXXX
```

About non-HA deployments

Agent Managers that are assigned to an HA Partition can still manage deployments and agents that are not HA Aware. When these cartridges are deployed to an Agent Manager, they by-pass the HA Deployment Set checks and are only deployed to this host. If this host goes offline, any non-HA agents also become off-line.

Negotiating Agent Manager resources at runtime

Monitoring agents often have specific Agent Manager resource requirements. Some agents need more memory, others require larger queue sizes, or more lenient time-outs. The Agent Negotiation feature allows monitoring agents to request resource changes during deployment, activation, and runtime. During deployment, an agent package can request specific resource settings to be applied as part of the deployment process. Depending on the current Agent Manager configuration, a restart may be required in order to apply the requested changes. This restart of the Agent Manager is automatic and is part of the deployment cycle.

Disabling runtime resource negotiation

Disabling runtime resource negotiation causes the Agent Manager to revert to running under its default startup configuration.

To disable runtime resource negotiation on startup:

- Set up the following runtime switch on the command line:
`./fglam --disable-overrides`

To permanently disable runtime resource configuration:

- 1 If the Agent Manager process is running, shut it down.
- 2 Edit the `<FGLAM_STATE>/config/client.config` file to include the following block of code:

```
# *** Disable Deployment Overrides ***  
#  
# Set this value to true in order to disable the resource settings  
# negotiation from occurring when an AgentPackage is deployed.  
disable.overrides = true;
```
- 3 Save the changes.
- 4 Start the Agent Manager.

Disabling agent-specific changes to the upstream queue

The `config/fglam-config.xml` file contains an attribute that allows you to disable any agent-specific changes to the upstream queue settings at runtime. By default this attribute is set to `true`. You can use this feature, for example, to prevent individual agents to negotiate allocating large amounts of memory.

```
<config:upstream max-queue-size="-1" max-disk-space="1024" max-batch-size="500"  
  allow-runtime-change="false"/>>
```

Configuring credentials

The Management Server includes a credential management system that enables you to create, store, and manage credentials through the Foglight® browser interface.

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

Credentials are encrypted and stored in lockboxes. Lockboxes are released to credential clients, such as the Agent Manager. Agents, in turn, request credentials from the Agent Manager.

For detailed information about managing credentials in Foglight, see “Controlling System Access with Credentials” in the *Administration and Configuration Guide*.

Foglight agents need access to credentials when monitoring systems that require credential verification. Credential information consists of a name, type, policies, and resource mappings. You can create and manage credentials through the Management Server browser interface.

Foglight supports the following commonly used credential types:

- **Challenge Response:** Uses one or more challenge and response pairs to grant access without requiring any interaction in the browser interface. The answers are sent by the agent as part of the agent configuration.
- **Domain, User Name, and Password (Windows):** Requires a user name and password to access a monitored resource. The domain name is optional.

i **IMPORTANT:** When specifying a domain name in this credential type, a fully qualified domain name is required. Failing to use a fully qualified domain name may prevent the Agent Manager from establishing a connection to a remote monitored resource. For example, if the full domain name is `prod.example.com`, use `prod.example.com` as the domain name instead of just `prod`, when configuring the credential.

- **DSA Key:** Uses the Digital Signature Algorithm (DSA) Key for authentication.
- **RSA Key:** Uses the RSA (Rivest, Shamir, and Adleman) Key for authentication.
- **Use Client’s Login At Connection Time:** Uses the currently logged in user’s account to access secured resources. This is not the user currently logged into the Management Server, but the user under which the credential client is running. For example, a credential provided to an Agent Manager instance launched by a user on a remote machine, causes the connection to the secured resource to be made using this user’s identity.
- **User Name:** Requires a user name to access a monitored resource.
- **User Name and Password:** Requires a user name and password to access a monitored resource.

Each credential can have one or more authentication policies, based on the desired usage count, failure rate, the time range during which the credential can be used, and the amount of time during which the credential information is cached locally. Credentials can apply to specific parts of the monitored environment, such as hosts and ports. Resource mappings identify secured resources. The mappings typically contain a combination of literal expressions, regular expressions, or an IP address range.

For more information about creating and managing credentials, including detailed examples of configuring a credential, see “Exploring the Manage Credentials Dashboard” in the *Foglight Administration and Configuration Guide*.

Managing lockboxes

A lockbox can be password-protected, and contains a collection of credential keys used for encryption and decryption. A lockbox can encrypt one or more credentials. All lockboxes, except the System lockbox, are password-protected.

You can create, edit, and manage lockboxes, change their passwords, and release them to credential clients (such as the Agent Manager) using the Manage Lockboxes dashboard in the Management Server browser interface.

Releasing lockboxes to the Agent Manager

Each lockbox in the Management Server contains a set of credentials and keys for their encryption and decryption. Credentials are released to the Agent Manager unencrypted. When a lockbox is released to the Agent Manager, the Agent Manager passes the credential information to its agents. The agents use this information to establish connection with target resources.

When you start the Agent Manager without having first released a lockbox to it from the Management Server, the following message appears in the startup log:

```
INFO The Credential Manager has not been assigned any lockboxes. Lockboxes are used to decrypt credentials received as a result of an Agent Credential Query. Without any lockbox assignments, credentials received within a credential query result-set will be discarded. You can grant lockboxes to this Agent Manager through the Credential Administrator on the Server.
```

The lockbox you release to the Agent Manager must contain the credentials necessary for the agents to access the monitored resources.

CAUTION: Any agents that have access to an Agent Manager with a released lockbox can potentially query and obtain credential information stored within that lockbox.

To release a lockbox to the Agent Manager:


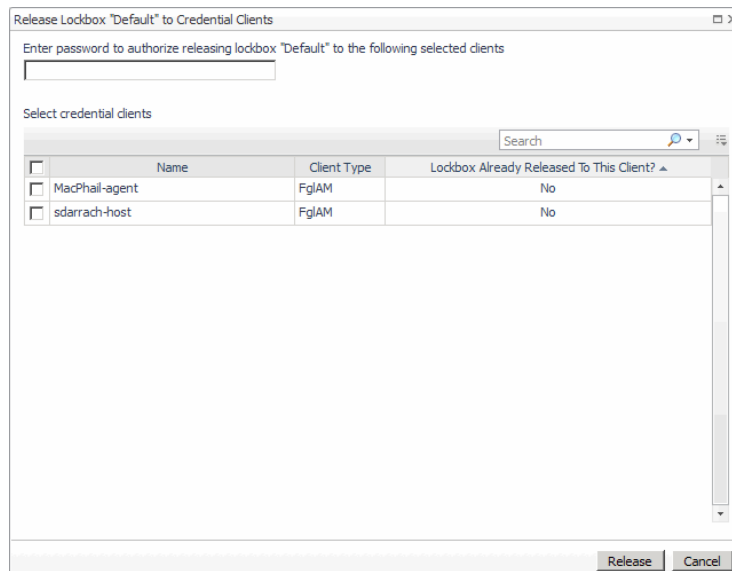
- 1 Log in to the Foglight browser interface.
- 2 On the navigation panel, click **Dashboards > Administration > Credentials > Manage Lockboxes**.
- 3 On the Manage Lockboxes dashboard, in the row containing the lockbox that you want to release, click  the **Release to Credential Clients** icon.
- 4 In the **Release Lockbox to Credential Clients** dialog box, type the lockbox password (if one exists) and select one or more credential clients (that is, Agent Managers) for lockbox release.

Figure 2. Release Lockbox to Credential Clients dialog box



IMPORTANT: The System lockbox that is included by default with the Management Server is not password-protected. Its contents are accessible to all clients in your system.

- 5 Click **Release**.

The **Release Lockbox to Credential Clients** dialog box closes, indicating success.

- 6 **Optional**—ensure the Credential Clients column is populated.
 - a Using the breadcrumb trail, return to the main Credentials dashboard, and navigate to the **View Clients** dashboard.
 - b On the **View Clients** dashboard, ensure that the **Show lockboxes currently assigned to each client** check box is selected.

NOTE: This functionality consumes server resources. It can be significant depending on the size of your client list.

The view refreshes, with the **Assigned Lockboxes** column populated.

- c Return to the main **Credentials** dashboard.
- d Navigate to the **Manage Lockboxes** dashboard.
- e On the **Manage Lockboxes** dashboard, observe the **Credential Clients** column of the newly released lockbox entry. The column lists the credential clients to which the lockbox is assigned.

When the lockbox is released to the Agent Manager, any credentials that are later added to the same lockbox are also accessible to the Agent Manager and its monitored agents.

Configuring anti-virus exclusion settings

Anti-virus software may negatively impact the CPU and system performance of machines running Foglight. To reduce resource consumption, it is highly recommended to exclude the relevant directory, processes, and executables from being scanned by the anti-virus software.

- The common installation directory is as follows:

```
%fglam_home%
```

- FglAM related processes and executables are as follows:

- For Windows:

- <FglAM Base Folder>\bin\fglam.exe
- <FglAM Base Folder>\client\<version>\bin\bindDA.exe
- <FglAM Base Folder>\client\<version>\bin\dcmlist.exe
- <FglAM Base Folder>\client\<version>\bin\fog4_launcher.exe
- <FglAM Base Folder>\client\<version>\bin\installer.exe
- <FglAM Base Folder>\client\<version>\bin\qcn_relauncher.exe
- <FglAM Base Folder>\client\<version>\bin\qcn_runner.exe
- <FglAM Base Folder>\client\<version>\bin\qcn_watchdog.exe
- <FglAM Base Folder>\client\<version>\bin\setDA.exe
- <FglAM Base Folder>\client\<version>\bin\udp2icmp.exe
- Non-embedded FglAM: <FglAM Base Folder>\jre\<jre_version>\bin\java.exe

- For other operating systems:

- <FglAM Base Folder>/bin/fglam
- <FglAM Base Folder>/bin/setuid_launcher

- <FglAM Base Folder>/client/<version>/bin/bindDA
- <FglAM Base Folder>/client/<version>/bin/dcmlist
- <FglAM Base Folder>/client/<version>/bin/fog4_launcher
- <FglAM Base Folder>/client/<version>/bin/installer
- <FglAM Base Folder>/client/<version>/bin/qcn_relauncher
- <FglAM Base Folder>/client/<version>/bin/qcn_runner
- <FglAM Base Folder>/client/<version>/bin/qcn_watchdog
- <FglAM Base Folder>/client/<version>/bin/setDA
- <FglAM Base Folder>/client/<version>/bin/udp2icmp
- Non-embedded FglAM: <FglAM Base Folder>/jre/<jre_version>/bin/java

Troubleshooting

This section provides information about problems that you may encounter while running the Agent Manager and describes the solutions available for these problems.

Errors related to Windows WMI and DCOM configuration

If you encounter errors with WMI or DCOM configuration for remote agents, see [Configuring Windows Management Instrumentation \(WMI\)](#) on page 65 and [Configuring Windows Remote Management \(WinRM\)](#) on page 77.

Resolving DATA_COLLECTION_FAILED errors

Credentials that are used by host agents to monitor remote hosts can have a `Use Count` policy configured through the Management Server browser interface. In such cases, when the number of collections exceeds the `Use Count` number, the host agent fails and enters the `DATA_COLLECTION_FAILED` mode. A `Use Count` policy violation event is also generated on the Management Server and can be viewed in the **Monitor Credential Alarms** dashboard.

To restart data collection:

- Delete the `Use Count` policy associated with the credential, through the Management Server browser interface (**Dashboards > Credentials > Manage Credentials > Edit > Policies**). For more information about credential management, see the *Foglight Administration and Configuration Guide*.

Adjusting the maximum polling interval

The **Maximum Number of Polls (per minute)** property sets the maximum number of polls that Foglight accepts from Agent Manager instances per minute.

In environments with a large number of Agent Manager instances, running Foglight® with the default value of 500 for **Maximum Number of Polls (per minute)** can increase Agent Manager communication latency to unacceptable levels.

If you are running a large number of Agent Manager instances in your environment, Quest recommends that you change the value of this property so that it is higher than 500 (the default value). Doing so increases the number of polls allowed per minute and decreases the maximum polling interval.

To adjust the **Maximum Number of Polls (per minute)** setting, in the Foglight browser interface, navigate to **Administration > Agents > Agent Properties > FglAM > FglAMAdapter**.

! **CAUTION:** Allowing a higher number of polls per minute increases the polling load on the Management Server. The value of Maximum Number of Polls (per minute) should only be increased on machines and in networks that can handle the increased load.

Advanced system configuration and troubleshooting

This chapter contains platform-specific configuration information for configuring Foglight® Agent Manager on Windows® when using Windows Management Instrumentation (WMI) or Windows Remote Management (WinRM) for remote monitoring access.

i | **NOTE:** WMI and WinRM are two different mechanisms that monitoring agents can use to establish remote connections. In most scenarios, only one of these mechanisms needs to be configured. The preferred mechanism is WinRM because of WMI scalability limitations.

This chapter also describes platform-specific instructions for configuring the Foglight Agent Manager on UNIX®.

Throughout this chapter, “the agent” is used as a placeholder for any Foglight agent that encounters these issues.

- [Configuring Windows Management Instrumentation \(WMI\)](#)
- [Configuring Windows Remote Management \(WinRM\)](#)
- [UNIX- and Linux-specific configuration](#)

Configuring Windows Management Instrumentation (WMI)

To remotely monitor a Windows® machine, some Foglight® agents require access to the WMI services, and specifically to the `WMIConnectionService`. This section outlines potential issues that may be encountered by agents attempting to access WMI, and provides solutions or workarounds.

i | **NOTE:** WMI is not allowed for remote connection in FIPS-compliant mode since the NTLM authentication it uses is non-FIPS compliant.

In order to access WMI, the user that the monitoring agent connects as must have sufficient credentials. Any user included in the Administrator group on the monitored machine already has the required access levels. For more information, see [Minimum requirements for Windows Management Instrumentation](#) on page 66.

There are several OS and environment-specific issues that may arise when using a Foglight agent to monitor a Windows machine remotely. This section provides solutions for the following issues:

- [WMI IPv6 connection support](#)
- [Windows Firewall interference](#)
- [Minimum requirements for Windows Management Instrumentation](#)
- [WMI access violation and OS connectivity verification failure](#)
- [WMI and Quota Violation error](#)
- [Known WMI issues in Windows Server 2008](#)
- [Tuning WMI connections](#)

- [Modifying registry key ownership on Windows Server 2008 R2](#)
- [Configuring registry settings for Windows Server 2008 R2 and Windows 7](#)
- [Resolving Access Denied errors when connecting to Windows XP Professional](#)
- [OS collection fails with a Local_Limit_Exceeded error](#)
- [Access to DCOM objects and registry is denied](#)
- [Configuring registry settings for WinShell access through DCOM](#)
- [Enabling agents to connect from UNIX machines](#)
- [Enabling agents to connect locally on Windows](#)
- [Releasing a locked MySQL process](#)

WMI IPv6 connection support

Starting with the Foglight Agent Manager version 5.9.1, the WMI connection with unique local IPv6 Address and link-local IPv6 Address is supported on the Agent Manager running on Windows and Linux.

Windows Firewall interference

Since the agent connects remotely (that is, from an external source) the Windows[®] Firewall can interfere with operations. In such cases, it is recommended that you initially try disabling the firewall to determine if that allows the agent to connect. When the agent can connect with the firewall disabled, re-enable it and open the following ports:

- TCP Port 135 (DCE/RPC Locator service, WindowsShellService, WMIConnectionService)
- TCP Port 139 (NetBIOS Session Service)
- TCP Port 445 (Windows shares)
- "Dynamic RPC" local ports

Minimum requirements for Windows Management Instrumentation

In order for the agent to have access to query WMI to collect OS and database metrics, the agent must have permission to access both DCOM and WMI. By default, any user in the Local Administrators group on the monitored host has the required permissions. Therefore, the best practice is to use a Local Administrator account on the monitored host as the agent OS user.

Promoting remote users to administrators on local machines through the Domain Controller

The recommended way of making users the administrators of their local machines is through Active Directory on the domain controller. Using the Domain Controller, you can:

- Set up local administrators for specific machines in the domain
- Promote local users to administrators for specific machines in the domain
- Make domain users administrators of all machines in the domain by adding them to the *Domain Admins* group

- Make domain users administrators of specific machines in the domain by adding them to the *Domain Admins* group

To promote a user to an administrator on a local machine using the Domain Controller:

- 1 Choose **Control Panel > Administrative Tools > Active Directory Users and Computers**.
- 2 In the **Active Directory Users and Computers** window that appears, in the left pane, under the domain node, click **Computers**.
- 3 In the right pane, right-click the machine whose local user you want to promote to an administrator, and choose **Manage**.
- 4 In the Computer Management window that appears, choose **System Tools > Local Users and Groups**.
- 5 You can now do any of the following:
 - Using the **Users** node in the right pane, make an existing or a new user an Administrator.
 - Using the **Groups** node, add an existing user to the Administrators group.

Granting required permissions to individual remote users

When making users the administrators of their local machines is not possible, you can grant required permissions to individual remote users using the following procedures.

To grant DCOM permissions to a user:

- 1 Add the local user to the "Distributed COM Users" group and the "Performance Monitor Users" group.
- 2 **If the Agent Manager is installed on a UNIX[®] machine:**
 - a On the monitored host machine, at the Windows[®] Run prompt, type **DCOMCNFG** and press **Enter**.
 - b In the **Component Services** window that appears, navigate to **Component Services > Computers > My Computer**.
 - c Right-click **My Computer** and click **Properties**.
 - d In the **My Computer Properties** dialog box that appears, open the **COM Security** tab.
 - e In the **Access Permissions** area, click **Edit Defaults**.
 - f In the **Access Permission** dialog box that appears, add the **Distributed COM Users** group to the list and grant it all permissions.
 - g Click **OK** to save your changes and close the **Access Permission** dialog box.
 - h In the **Launch and Activation Permissions** area, click **Edit Defaults**.
 - i In the **Launch and Activation Permissions** dialog box that appears, add the **Distributed COM Users** group to the list and grant it all permissions.
 - j Click **OK** to save your changes and close the **Launch and Activation Permissions** dialog box.
 - k In the **My Computer Properties** dialog box, click **OK** to close it.
 - l Close the **Component Services** window.

To grant minimum WMI permissions to a remote user:

- 1 On the monitored host machine, right-click **My Computer**, and navigate to **Manage > Services and Applications > WMI Control**.
- 2 Right-click **WMI Control** and click **Properties**.
- 3 In the **WMI Control Properties** dialog box, open the **Security** tab.
- 4 Expand the **Root** node and select **CIMV2**, then click **Security**.
- 5 In the **Security for ROOT\CIMV2** dialog box, add the **Distributed COM Users** group

- 6 Grant the required permissions to the remote user by enabling the following check boxes in the **Allow** column:
 - **Execute Methods**
 - **Enable Account**
 - **Remote Enable**
 - **Read Security**
- 7 Click **Apply** and then click **OK**.

To add subsequent users, they only need to be added to the two groups, **Distributed COM Users** and **Performance Monitor Users**, since these groups are already granted the required permissions.

Even though the local user is now granted access to WMI with the above configuration, not all performance monitoring classes allow non-administrative users to access their instances. Some performance classes need special permission to enable non-administrative users to perform queries or execute methods on their object instances. Some of these queries can fail clearly with an error code (for example, by the Agent Manager service throwing a Java exception), but some of them can fail without returning any data or error codes. Therefore, this setup must be used carefully, as query results can be unpredictable. From the system security perspective, there is still only so much a non-administrative user can do.

WMI access violation and OS connectivity verification failure

Access to WMI is required for both OS verification while creating an agent, and for collecting OS metrics from the monitored host. In some cases, an access violation on the WMI namespace (error 0x00000005) may occur when the agent connectivity verification fails. The access violation error can occur when the permissions or credentials used to access WMI, and specifically the *root\cimv2* namespace and *Root\MSCluster* (if the machine is included in a cluster), are changed or invalid.

In order to view and change namespace security, the user must have Read Security and Edit Security permissions. Administrator accounts have these permissions by default, and can assign them to other users if necessary.

To grant namespace permissions to a user:

- 1 Log in to the monitored host machine.
- 2 Right-click **My Computer** and click **Manage**.
- 3 In the **Computer Management** dialog box that appears, expand the **Services and Applications** node.
- 4 Right-click **WMI Control**, and click **Properties**.
- 5 In the **WMI Control Properties** dialog box that appears, open the **Security** tab.
- 6 Expand the **Root** node and select **CIMV2**, then click **Security**.
- 7 In the **Security for ROOT\CIMV2** dialog box that appears, modify or assign permissions as necessary.

i | **IMPORTANT:** If the user accesses the namespace remotely, you must select the Remote Enable permissions check box.

- 8 Click **OK** on each of the dialog boxes to close them.
- 9 If the machine is part of a cluster, repeat the procedure for the *ROOT\MSCluster* namespace.

i | **NOTE:** User permissions set on a namespace apply only to the namespace and not to any sub-namespaces. If the user also needs to access sub-namespaces, you can enable access from the Advanced Security Settings dialog box (accessible by clicking **Advanced** on the Security for [namespace] dialog box).

WMI and Quota Violation error

In some cases, when the agent queries WMI to collect metrics on the monitored host, the host's performance may be affected. When this occurs, the agent log includes the message: `Quota violation [SWbemServicesEx]`. This error stems from WMI rather than the agent, and has been identified by Microsoft® as a known issue in Windows® Server 2003.

To correct this error:

- 1 Ensure that the following hotfix, available from Microsoft in KB 828653, is applied:
<http://support.microsoft.com/kb/828653>
- 2 Restart the WMI service on the monitored host.

Known WMI issues in Windows Server 2008

Microsoft® has identified a number of known issues with WMI services in Windows® Server 2008. The following issues may occur:

- The database agent OS validation or collection has errors, even when the OS credentials are correct.
- The WMI reports incorrect data from the performance classes.
- Remotely accessing WMI causes memory leaks in Windows Server 2008.
- The WMI reports are unstable.
- In SQL Server 2008, the *perfmon* counters do not appear under WMI because they do not exist.

To resolve these issues:

- 1 Using the same OS credentials as the agent uses, run `perfmon /wmi` on the monitored host.
- 2 Install the hotfixes available in the following Microsoft support articles on the Windows 2008 machine:
 - <http://support.microsoft.com/?id=961435>
 - <http://support.microsoft.com/?id=977357>
 - <http://support.microsoft.com/?id=970520>
- 3 In SQL Server 2008 on Windows Server 2008, if the *perfmon* counters are missing from WMI, install the following hotfixes:
 - To install SQL Server 2008 Service Pack 1, go to <http://support.microsoft.com/kb/968382>.
 - To install cumulative update package 3 for SQL Server 2008 Service Pack 1, go to <http://support.microsoft.com/kb/971491>.

Tuning WMI connections

The following parameters can be used to fine tune the Windows Management Instrumentation (WMI) connections made by Foglight® agents. You can set these parameters for Foglight Agent Manager either in the `baseline.jvmargs.config` file as `vmparameter` options, or specify them directly as `jvm` options when starting up Foglight Agent Manager via command line. The values used for these parameters vary from environment to environment and so should be used as reference only.

- **Max Active Connections:** Defines the number of simultaneous outbound connections that can be made by all agents running on a particular instance of Foglight Agent Manager. The default value is 100.

Sample: `-Dcom.quest.connection.regulator.maxActiveConnectionsCap=500`

- **Expire connection based on no. of executions:** Defines the number of executions that can be made by a connection before forcibly closing it. The default value is 50.

Sample: `-Dquest.debug.poolable.wmi.time.to.live.executions=200`

- **Timeout by elapsed time:** Defines the time (in milliseconds) after which a connection is expired. The default value is 600000 (10 mins).

Sample: `-Dquest.debug.poolable.wmi.time.to.live.timeout.millis=900000`

Modifying registry key ownership on Windows Server 2008 R2

Configuring the agent to connect to the remote Windows® Management Instrumentation (WMI) and WinShell components on Windows Server 2008 R2 systems requires certain modifications to the registry keys. In some situations, these registry keys may be owned by `TrustedInstaller` instead of the Administrators group, which prevents modifications from being made even by Administrator accounts.

i | **TIP:** It is recommended that you create a backup copy of the Windows Registry that you can revert to prior to making any changes.

To modify the registry key ownership:

- 1 Using the registry editor, locate the registry key that you need to modify.
- 2 Right-click the registry key and click **Permissions**.
- 3 On the **Permissions** dialog box, click **Advanced**.
- 4 On the **Advanced Security Settings** dialog box, open the **Owner** tab.
- 5 In the **Change owner to list**, select the **Administrators** group.

i | **IMPORTANT:** An error indicating insufficient permissions may appear. To resolve, a user with the administrative account intended to use for profiling can change the ownership of the keys to that administrative account, and then set the permissions to Full on the appropriate registry keys.

- 6 Click **OK**.

Ownership of the registry key is assigned to the Administrators group, allowing any member of that group to modify the key's permissions. For more information about required registry changes, see [Configuring registry settings for Windows Server 2008 R2 and Windows 7](#).

Configuring registry settings for Windows Server 2008 R2 and Windows 7

In order to allow the agent to connect to the remote Windows® Management Instrumentation (WMI) components on Windows Server 2008 R2 or Windows 7 systems, special registry settings are required.

To configure Windows Server 2008 R2 registry settings:

- 1 If necessary, add the required domain\user to the Administrators group.
 - a On the desktop, right-click **My Computer** and navigate to **Groups (My Computer > Manage > Configuration > Groups)**.
 - b Double-click the **Administrators** group.
 - c Add the required `domain\user` to the Administrators group.
 - d Click **Apply**, and then click **OK** to close the dialog box.

2 Edit the `Wbem Scripting Locator` registry key.

i | **IMPORTANT:** There may be multiple instances of each key, especially for 64-bit versions of Windows Server 2008 R2 and Windows 7. Each instance must be updated.

- a Run `regedit`.
- b Search for all instances of the keys named:
 - {72C24DD5-D70A-438B-8A42-98424B88AFB8} (Windows Script Host Shell Object)
 - {76A64158-CB41-11d1-8B02-00600806D9B6} (WBEM Scripting Locator)
 - {0D43FE01-F093-11CF-8940-00A0C9054228} (Windows Script FileSystem Object)

There may be multiple instances of each key, especially for 64-bit versions of Windows Server 2008 R2 and Windows 7. Each instance must be updated.

- c For each key, right-click the key and navigate to the owner permissions (**Permissions > Advanced > Owner**).
- d Click **Other Users and Groups**, and add the Local Administrators group.
- e Select the **Replace Owner on subcontainers and objects** check box.
- f Click **Apply**, and then click **OK**.

The security screen appears.

- g Select the Local Administrators group and grant full permissions.

i | **NOTE:** An error indicating insufficient permissions may appear. To resolve, a user with the administrative account intended to use for profiling can change the ownership of the keys to that administrative account, and then set the permissions to Full on the appropriate registry keys.

3 Close `regedit`.

If configuring the registry does not solve the connection issue, you may need to disable Windows User Account Control (UAC).

To disable Windows UAC:

- Follow the instructions provided on the following website:

http://technet.microsoft.com/en-us/library/cc709691%28WS.10%29.aspx#BKMK_S3

Resolving Access Denied errors when connecting to Windows XP Professional

When attempting to monitor a Windows® XP Professional machine, the remote agent sometimes fails to connect with an `Access Denied` error message. This error can be caused by the Windows `ForceGuest` setting, which is enabled by default on machines that are not part of a domain. To resolve the issue, ensure that remote connections are not being coerced to log on as the guest account by following the procedure below.

To disable ForceGuest for remote connections:

- 1 In the **Start** menu, type `run`.
- 2 In the **Run** dialog box that appears, type `secpol.msc` and click **OK**.
The **Local Security Policy** editor appears.
- 3 Choose **Local Policies > Security Options**.
- 4 Select the entry: **Network access: Sharing and security model for local accounts**.

- 5 If this entry's security setting is **Guest only**, right-click the entry and click **Properties**.
- 6 In the dialog box that appears, on the **Local Security Setting** tab, select **Classic - local users authenticate as themselves** from the list.
- 7 Click **OK**.
- 8 Restart the computer.

For more information, see step 5 in the following VMWare KB article:

http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=2013

OS collection fails with a *Local_Limit_Exceeded* error

The agent uses Windows® authentication to Negotiate the monitored instance. In some cases, the negotiation can fail if there is a mismatch between the authentication types used by the client and the server.

The following symptoms indicate an authentication issue:

- NTLMv2 security is used for Windows authentication
- Operating system collection fails
- The agent log includes the following message:

```
WMIconnection - WMIconnection.An internal error occurred.  
...  
org.jinterop.dcom.common.JIException: An internal error occurred.  
    [0x8001FFFF]  
...  
Caused by: rpc.BindException: Unable to bind. (LOCAL_LIMIT_EXCEEDED)
```

To resolve these issues, you may need to disable the NTLMv2 authentication.

To turn off NTLMv2 authentication:

- 1 Run *regedit* to edit the registry.
- 2 Locate the following registry key:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa
- 3 Locate the value named *LMCompatibilityLevel*, and change the `DWORD` value to 2 (send NTLM authentication only).
- 4 Close *regedit* and restart the machine.

Access to DCOM objects and registry is denied

When the agent attempts to access the Windows® registry and enter the Windows Management Instrumentation (WMI) component, the attempt can fail if the agent user does not have the required permissions.

The following symptoms indicate a permissions issue:

- The database agent fails even when the user credentials are valid.
- The agent log includes the following message:
The user account <username> has invalid login credentials.
- The agent log includes the error `WMIconnection.Access is denied` and error code `0x00000005`.
- The agent OS verification fails on Windows Vista or Windows 7 installed with a non-English character set.

To enable agent verification:

- 1 Ensure that the following two registry keys exist on the monitored host, in `HKEY_CLASSES_ROOT\CLSID\`:

i | **IMPORTANT:** The difference between these two keys is the bold character.

- 76A6415**B**-CB41-11d1-8B02-00600806D9B6
- 76a6415**8**-cb41-11d1-8b02-00600806d9b6

i | **IMPORTANT:** There may be multiple instances of each key, especially for 64-bit versions of Windows Server 2008 R2 and Windows 7. Each instance must be updated.

If one or both of these keys does not exist on the monitored host, export them from a machine (that is, running the same Windows OS) and import them to the monitored host.

- 2 Set permissions for the registry keys from [Step 1](#).

i | **IMPORTANT:** An error indicating insufficient permissions may appear. To resolve, a user with the administrative account intended to use for profiling can change the ownership of the keys to that administrative account, and then set the permissions to Full on the appropriate registry keys.

- a Log in to the target host with an account that has Administrator privileges.
- b Start *regedit*, and from the **Edit** menu, use **Find** to search for the following key:
76A6415B-CB41-11d1-8B02-00600806D9B6.
- c Right-click the **Class ID**, and click **Permissions**.
- d In the **Permissions** dialog box, click **Advanced**.
- e In the **Advanced Security Settings** dialog box, open the **Owner** tab.
- f On the **Owner** tab, in the **Change owner to area**, select the account with which you are currently logged in.
- g Click **OK**.
The **Advanced Security Settings** dialog box closes.
- h In the **Permissions** dialog box, select the **Administrators** group.
- i In the **Permissions for Administrators** area, in the **Allow** column, select the **Full Control** check box.
- j Click **OK**. The **Permissions** dialog box closes.
- k Repeat [Step a](#) through [Step j](#) for the second key: 76a64158-cb41-11d1-8b02-00600806d9b6

Configuring registry settings for WinShell access through DCOM

Any *WindowsShell* connection made to a non-local host requires DCOM access to that machine, regardless of whether the user establishing the connection is a local or third-party user.

Therefore, agents that connect to Windows® machines using the Agent Manager's *WindowsShellService* need to make the following specific registry changes to allow the connection.

To enable *WindowsShellService* access:

- 1 Ensure that the following two registry keys exist on the monitored host, in `HKEY_CLASSES_ROOT\CLSID\`:

- 72C24DD5-D70A-438B-8A42-98424B88AFB8
- 0D43FE01-F093-11CF-8940-00A0C9054228

i | **IMPORTANT:** There may be multiple instances of each key, especially for 64-bit versions of Windows Server 2008 R2 and Windows 7. Each instance must be updated.

If one or both of these keys does not exist on the monitored host, export them from an identical machine (that is, running the same Windows OS) and import them to the monitored host.

2 Set permissions for the registry keys from [Step 1](#).

i | **IMPORTANT:** An error indicating insufficient permissions may appear. To resolve, a user with the administrative account intended to use for profiling can change the ownership of the keys to that administrative account, and then set the permissions to Full on the appropriate registry keys.

- a Log in to the target host with an account that has Administrator privileges.
- b Start *regedit*, and from the **Edit** menu, use **Find** to search for the following key: 72C24DD5-D70A-438B-8A42-98424B88AFB8.
- c Right-click **Class ID**, and click **Permissions**.
- d In the **Permissions** dialog box that appears, click **Advanced**.
- e In the **Advanced Security Settings** dialog box, open the **Owner** tab.
- f On the **Owner** tab, in the **Change owner to area**, select the account with which you are currently logged in.
- g Click **OK**.
The **Advanced Security Settings** dialog box closes.
- h In the **Permissions** dialog box, select the **Administrators** group.
- i In the **Permissions for Administrators** area, in the **Allow** column, enable the **Full Control** check box.
- j Click **OK**.
The **Permissions** dialog box closes.
- k Repeat [Step a](#) through [Step j](#) for the second key:
0D43FE01-F093-11CF-8940-00A0C9054228.

Permissions on registry keys to configure DCOM command shell connection

A Windows® operating system user needs full control permissions on the following registry keys to monitor the operating system:

- 76A64158-CB41-11D1-8B02-00600806D9B6 (WBEM Scripting Locator)
- 72C24DD5-D70A-438B-8A42-98424B88AFB8 (Windows Script Host Shell Object)
- 0D43FE01-F093-11CF-8940-00A0C9054228 (FileSystem Object)
- 76A6415B-CB41-11D1-8B02-00600806D9B6

According to the COM specification, the full control permission to the registry keys are required to write values to the registry keys. The values written to the registry key are as follows:

- HKEY_CLASSES_ROOT\AppID\{key}: Need to write the string value name to DIISurrogate and leave the value to blank.
- HKEY_CLASSES_ROOT\CLSID\{key}: Need to write the string value name to AppID and set the value to {key}.

For 64-bit Windows operating system, there might be two directories of AppID and CLSID, then the written values are:

- `HKEY_CLASSES_ROOT\AppID\{key}`: Need to write the string value name to `DIISurrogate` and leave the value to blank.

- `HKEY_CLASSES_ROOT\Wow6432Node\AppID\{key}`: Need to write the string value name to `DIISurrogate` and leave the value to blank.
- `HKEY_CLASSES_ROOT\CLSID\{key}`: Need to write the string value name to `AppID` and set the value to `{key}`.
- `HKEY_CLASSES_ROOT\Wow6432Node\CLSID\{key}`: Need to write the string value name to `AppID` and set the value to `{key}`.

i NOTE:

1. If the keys under `HKEY_CLASSES_ROOT\AppID` do not exist, manually add the keys to the written value by default permission.
2. If the keys under `HKEY_CLASSES_ROOT\CLSID` and `HKEY_CLASSES_ROOT\Wow6432Node\CLSID` do not exist, and you do not have permission to add a new `String Value` or edit the `Value data`, change the `Owner` from `TrustedInstaller` to `Administrators`, then grant the `Set Value` permission first.

Providing the full control permissions to a Foglight Agent Manager (FglAM) user is the most convenient way to write these values, which will be generated automatically. If you don't want to provide the full control permissions to the FglAM user, do either of the following:

- Manually write the values to those keys, and then remove the full control permission. If the full control permissions cannot be deselected, select `Deny Permission` entry to remove all the permissions, and keep permissions for the entries `Query Value`, `Enumerate Subkeys`, `Notify`, and `Read control` to `Read only`. To set deny permission, right click on the registry key and select **Permissions**. Click **Advanced** on the popup dialogue box, then double click on the FglAM user, and check `Deny Permission` entry.
- Give permission to `Set Value` only. After writing the value name and value data, revoke the `Set Value` permission.

For `FileLogMonitorAgent` and `WindowsEventLogMonitorAgent`:

- If the Agent Manager is running in Windows, write values for the following two keys:
 - `72C24DD5-D70A-438B-8A42-98424B88AFB8`
 - `0D43FE01-F093-11CF-8940-00A0C9054228`
- If the Agent Manager is running in Linux, write values for following three keys:
 - `76A64158-CB41-11D1-8B02-00600806D9B6` (For `j-interop WMIJavaConnection`)
 - `72C24DD5-D70A-438B-8A42-98424B88AFB8`
 - `0D43FE01-F093-11CF-8940-00A0C9054228`

The key `76A64158-CB41-11D1-8B02-00600806D9B6` is used for the Agent Managers installed on Unix or Linux machine to establish the `WMIJavaconnection`, which requires the administrator privilege to monitor.

Enabling agents to connect from UNIX machines

When an agent connects to a monitored Windows® host from a UNIX® machine, you must make certain registry changes in order to allow the required COM services to run.

To add the following registry key:

- 1 Click **Start > Run**.
- 2 Input `regedit` in the dialog box and click **OK**.

- 3 Add the following registry key to Windows if it does not exist: `HKEY_CLASSES_ROOT\AppID{76A64158-CB41-11D1-8B02-00600806D9B6}`. Create a new string value named `DllSurrogate` under that key and leave it blank.
- 4 Add the following registry key to Windows if it does not exist: `HKEY_CLASSES_ROOT\CLSID{76A64158-CB41-11D1-8B02-00600806D9B6}`. Create a new string value named `AppID` under that key and modify the data to: `{76A64158-CB41-11D1-8B02-00600806D9B6}`

To allow the agent to connect from a UNIX machine to a monitored Windows host:

- 1 Enable the Remote Registry Service.

Once the agent has successfully connected from a UNIX machine and the Agent Manager connection services have made the required changes, the Remote Registry Service can be disabled.

- 2 Ensure that the Server service is running.

i | IMPORTANT: Normally the Server service starts automatically. If it is stopped, or fails to start, it must be manually restarted before the Agent Manager can connect from a UNIX machine.

Disabling UAC

When an agent connects to a monitored Windows host from a UNIX machine, user access control (UAC) must also be disabled in order for WMI connections to succeed.

This requirement affects: Windows Vista, Windows Server 2008, and Windows 7.

To turn off UAC on Windows 7:

- Navigate to **Control Panel > User Accounts and Family Safety > User Accounts > Change User Account Control** Settings, and change the setting to **Never Notify**.

Granting access to *dllhost.exe* when Windows Firewall is enabled

When an agent connects to a monitored Windows host from a UNIX machine, and the Windows firewall is enabled, access to *dllhost.exe* must be allowed through the firewall.

To grant access to *dllhost.exe*:

- 1 Issue the following command on the command-line of the monitored Windows host:

```
netsh firewall add allowedprogram program=%windir%\system32\dllhost.exe
name=Dllhost
```

- 2 Ensure that Windows UAC is disabled. See [Disabling UAC](#) on page 76 for details.
- 3 Restart the monitored host.

Enabling agents to connect locally on Windows

When a WMI agent connects to the same machine it is running on (that is, *localhost*) using credentials that explicitly specify a user other than the currently logged on user, you must make certain registry changes to allow the required COM services to run.

To allow the agent to connect locally on Windows®:

- 1 Enable the Remote Registry Service.

Once the agent has successfully connected and the Agent Manager connection services have made the required changes, the Remote Registry Service can be disabled.

- 2 Ensure that the Server service is running.

i | **IMPORTANT:** Normally, the Server service starts automatically. If it is stopped, or fails to start, it must be manually restarted before the Agent Manager can connect locally using credentials for a user other than the currently logged on user.

Releasing a locked MySQL process

The Agent Manager uses the `wmiprvse.exe` process to make use of WMI for remote Windows® monitoring. However, in some situations, this process can lock the MySQL process, `mysqld.exe`, preventing it from being uninstalled, deleted, moved, or updated.

To release a locked MySQL process:

- Stop the `wmiprvse.exe` process.

i | **IMPORTANT:** Stopping the `wmiprvse.exe` process does not affect the running state of the WMI service. The process starts again automatically when an application requires WMI.

Configuring Windows Remote Management (WinRM)

Windows Remote Management is the Microsoft® implementation of the Web Services Management Protocol (WSMAN) which is a Simple Object Access Protocol (SOAP) based protocol over HTTP/HTTPS and is used for system management. For more information, visit <https://msdn.microsoft.com/en-us/library/aa384470%28v=vs.85%29.aspx>.

WinRM has two authentication mechanisms that are used by the Agent Manager to establish connections:

- *Negotiate authentication* is based on Kerberos authentication, involving tickets/keys obtained from a Key Distribution Center (KDC).
- *Basic authentication* uses standard HTTP headers to communicate directly with the remote machine.

This section provides solutions for the following issues:

- [WinRM IPv6 connection support](#)
- [About WinRM authentication and the Agent Manager](#)
- [Configuring the target \(monitored\) system](#)
- [Configuring the Agent Manager \(monitoring\) system](#)
- [Generating a configuration file required for WinRM Negotiate authentication](#)
- [Configuring command-shell connection settings](#)
- [About WinRM connection ports](#)
- [Troubleshooting](#)

WinRM IPv6 connection support

Starting with the Foglight Agent Manager version 5.9.1, the WinRM connection with unique local IPv6 Address and link-local IPv6 Address is supported on the Agent Manager running on Windows and Linux.

Understanding Negotiate/Kerberos authentication

Kerberos is a network security protocol that involves three elements: the KDC, the client user, and the server with the desired service to access. The KDC is installed as part of the Domain Controller and acts as the authentication service and the ticket-granting service.

Each administrative domain has its own KDC, which contains information about the users and services for that particular domain. This administrative domain is a Kerberos realm.

The steps involved in Kerberos authentication involve the following actions:

- 1 Initial Authentication for a Kerberos Session (*Kerberos step 1*)
 - a A client begins a Kerberos session by requesting a ticket-granting-ticket (TGT) from the KDC based on the domain user name. This is done automatically on Windows login, but can be repeated when a different KDC is needed for a service on a different realm.
 - b The KDC creates a TGT and sends it back to the client, in encrypted form, based on the client's password. The client decrypts the TGT using the client's password.
 - c In possession of a valid TGT, the client can request tickets for various services.
- 2 Subsequent Kerberos Authentications (*Kerberos step 2*)
 - a When the client requires access to a service using Kerberos authentication, the client requests a ticket for the particular service, from the KDC by sending the TGT as proof of identity.
 - b The KDC sends the ticket for the specific service to the client.
 - c The client can then send this ticket to the remote server to establish a session with the server's service.

Finding the Kerberos authentication file

Foglight Agent Manager always generates an *auth.login.config* file that is used for Kerberos. It is generated by the Agent Manager on both UNIX® and Windows®, and is used to configure the Kerberos module that the Agent Manager uses for authentication. This file is located in the *<fglam_dir>/state/default/config* directory, and must never be modified.

Observing the contents of the Kerberos configuration file

The Kerberos configuration file specifies the KDC from which tickets are obtained. Operating systems sometimes have their own Kerberos configuration files. If present, the Agent Manager uses them by default. They can be found in the following locations:

- **Windows:** *%WINDIR%\krb5.ini* which typically translates to *C:\Windows\krb5.ini*
- **UNIX:**
 - */etc/krb5.conf*
 - Or:
 - */etc/krb5/krb5.conf*

If none of these files are found, the Agent Manager attempts to create its own kerberos configuration file, based on the detected settings. The detection can only be done on Windows, and on Unix, the empty file is generated.

To detect the settings on Windows, the following environment variables are checked:

- **%LOGONSERVER%:** Provides the name of the domain controller that authenticated the client's logon to the machine. This value is just the simple name of the KDC, but the fully qualified name must be used in the configuration file.

- `%USERDNSDOMAIN%`: Provides the fully qualified DNS domain that the currently logged on user's account belongs to.

These two environment variables are only present when the user credential used to login to the machine, is a domain credential, for example, when the user account belongs to the domain. If a local user account is used, these environment variables are not present, and an attempt to generate the Kerberos configuration file is made using the network information. If not found, the empty file is generated.

If you are running the Agent Manager on a UNIX machine, to determine the values to use in the Kerberos configuration file for the remote machine, simply log in to the remote machine using a domain user credential and check the values. The format of the file is as follows:

```
[libdefaults]
    default_realm = <REALM_NAME_IN_CAPS>

[realms]
    <REALM_NAME_IN_CAPS> = {
        kdc = <fully_qualified_kdc_name>
    }

[domain_realm]
    .<domain_in_lower_case> = <REALM_NAME_IN_CAPS>
```

If the Kerberos configuration file is generated by the Agent Manager, it is placed in the `<fglam_dir>/state/default/config/krb5.config` file, and an entry is added to the `<fglam_dir>/state/default/config/fglam.config.xml` file so that the Agent Manager is aware of the file location. An example of this entry on Windows is as follows:

```
<config:krb5-config-
file>C:\Quest_Software\fms_5_7_5_5\state\fglam\state\default\.\config\krb5.config</
config:krb5-config-file>
```

If the file is not generated, you can generate your own file, add a value for the `krb5-config-file` entry in the `fglam.config.xml` file, and restart the Agent Manager.

About the Kerberos Configuration File Format

- **NOTE:** If current OS user have permission to modify `krb5.config` file, FglAM will automatically specify domain name as the `kdc` in the file if the `kdc` entry is not found when making Kerberos authentication.

The Kerberos configuration file typically looks like:

```
# Copyright 2017 Quest Software Inc.
# ALL RIGHTS RESERVED.
[libdefaults]
    default_realm = EXAMPLE.COM

[realms]
    EXAMPLE.COM = {
        kdc = HOST1.EXAMPLE.COM
    }

[domain_realm]
    .example.com = EXAMPLE.COM
```

- Any line that begins with a number sign '#' is considered a comment, and is ignored.
- The `default_realm` value is used to determine what KDC should be used, if the realm cannot be determined from the domain.
- The `[realms]` section is used to provide the KDC for the specified realm.
- The `[domain_realm]` section is used to map the domain to the realm to use.

So for example, if connecting to a host **A** with user credential `example.com\UserX`, the kerberos file is used as follows:

- 1 The domain is determined from the user credential, so in this case the domain name is `example.com`.
- 2 The domain `example.com` maps to the realm `EXAMPLE.COM` in the `domain_realm` section.
- 3 The realm `EXAMPLE.COM` is then found, and its KDC value is used to determine the KDC to use for authentication.
- 4 The KDC, `HOST1.EXAMPLE.COM`, is then communicated with for authentication.

In another example, if connecting to a host **B** with user credential `other.domain\UserY`, the same Kerberos file is used as follows:

- 1 The domain is determined from the user credential, so in this case the domain name is `other.domain`.
- 2 The domain `other.domain` does not map to any realm in the `domain_realm` section, so the KDC is attempted to be resolved from the DNS.
- 3 Typically, the DNS does not find the KDC for the different domain, and so the `default_realm` value, `EXAMPLE.COM` is used instead.
- 4 The realm `EXAMPLE.COM` is then found, and its KDC value is used to determine the KDC to use for authentication.
- 5 The KDC, `HOST1.EXAMPLE.COM`, is then communicated with for authentication of the user credential.

To specify a non-default realm to use for the `other.domain` value in the second example, the Kerberos configuration file can be modified as follows

```
# Copyright 2017 Quest Software Inc.
# ALL RIGHTS RESERVED.
[libdefaults]
default_realm = EXAMPLE.COM

[realms]
EXAMPLE.COM = {
    kdc = HOST1.EXAMPLE.COM
}
OTHER.DOMAIN = {
    kdc = MYKDC.OTHER.DOMAIN
}

[domain_realm]
.example.com = EXAMPLE.COM
.other.domain = OTHER.DOMAIN
```

Now, if a domain of `other.domain` is encountered, the realm used will be `OTHER.DOMAIN` instead of the `default_realm` value since there is a domain mapping entry. This can be repeated for other domains and realms.

This is the reason that fully qualified domain names be used for the domain value of the credential. Also, to prevent any possible DNS issues with the KDC, the fully qualified name of the host should be used, such as `A.example.com` or `B.other.domain`.

There are other sections and properties that can be used in the Kerberos configuration file, but for the Agent Manager purposes, the ones described above are sufficient.

Configuring multiple KDC entries in the same realm

The Kerberos configuration file supports specifying multiple KDC entries in one realm, which enables the WinRM connections to try to obtain the tickets from one of these KDCs, in order to prevent the single point of failure. The multiple KDC entries configuration typically looks like:

```
[libdefaults]
default_realm = DOMAIN.FGLAM
kdc_timeout = 30000
```

```

max_retries = 3
[realms]
  WINRM.FGLAM = {
    kdc = FGLAMDCS01.WINRM.FGLAM
    kdc = FGLAMDCS02.WINRM.FGLAM
  }
[domain_realm]
  .domain.fglam = DOMAIN.FGLAM
  .winrm.fglam = WINRM.FGLAM

```

- `kdc_timeout`: This parameter sets the maximum number of milliseconds to wait for a reply from a KDC, the default value is 30000 milliseconds.
- `max_retries`: This parameter sets the maximum number of times each KDC will be tried, the default value is 3.

Each KDC listed in the same realm will be tried up to `<max_retries>` times and each time will wait for up to `<kdc_timeout>` milliseconds until one succeeds.

NOTE: Multiple KDC entries are not supported with the default settings because the total timeout number (`<KDC_number> * <max_retries> * <kdc_timeout>`) for KDC listed in the same realm is larger than the default timeout limit (60000 milliseconds) for FglAM WinRM connection.

To fix this issue, follow the instructions to achieve that total timeout number for KDC listed in the same realm does not exceed the timeout limit set by WinRM connection:

- Decrease the parameter values of `max_retries` and `kdc_timeout` in `krb5` file. For example, `max_retries=1` and `kdc_timeout=10000`.
- Run `fglam` with following command-line parameter to increase timeout number for FglAM WinRM connection.

```
-Dwinrm.connection.timeout.milliseconds=<time>
```

About the cross-realm negotiate authentication behavior

Cross-realm or cross-domain is the mechanism of using WinRM Negotiate authentication to establish a connection to a machine in a different domain. There are also some differences here between behavior when the Agent Manager uses Java 6 and Java 7.

Enable DNS reverse lookup

DNS reverse lookup is the querying technique of the Domain Name System (DNS) to determine the domain name associated with an IP address – the reverse of the usual “forward” DNS lookup of an IP address from a domain name. The process of reverse resolving of an IP address uses PTR records.

The *Enable DNS Reverse Lookup* switch is under **Dashboards > Administration > Agents > Agent Properties > FglAM > FglAMadapter**.

- True: Reverse lookup will be used along with forward lookup for canonicalizing host names used in SPN (service principal names).
- False: Reverse lookup is disabled, only forward lookup will be used.

Understanding Basic authentication

The Basic authentication mechanism simply provides HTTP headers to the remote machine, to provide the credentials that should be used for authentication. There is no use of a configuration file or a KDC. The

mechanism provides no protection for the transmitted credentials, since they are simply encoded in *base64* and not encrypted or hashed. To address any security concerns, it is recommended that Basic authentication attempts are made over an HTTPS and not HTTP connection. Since the KDC is not involved in the authentication process, the credentials used for Basic authentication must be local user credentials, such as local user credentials for the remote machine, and not domain user credentials.

When the remote machine is contacted, the remote machine responds indicating that Basic authentication needs to be used by the client. The client side then base64-encodes the user name and password and sends an HTTP request with an HTTP header containing the base64-encoded result to the remote machine, which then validates the provided credential.

Since local user credentials are used, and a KDC is not needed, the domain the host is on is irrelevant for Basic authentication.

About WinRM authentication and the Agent Manager

The Agent Manager supports *Basic* and *Negotiate* WinRM authentication schemes with Windows credentials. The Negotiate authentication scheme is enabled by default in WinRM and is the recommended way to authenticate in most environments.

In order to establish connections over Windows® Remote Management (WinRM), you must provide a Windows credential. For more information, see [Configuring credentials](#) on page 60.

The Negotiate authentication scheme requires fully-qualified domain accounts. For example, instead of using the NetBIOS domain, the Windows credential should be configured with the fully qualified domain.

i | **NOTE:** When running in FIPS-compliant mode, the password for the account must be at least 14 characters (112 bits).

The Basic authentication scheme requires local Administrator accounts; you cannot use domain accounts. For more information, see [Promoting remote users to administrators on local machines through the Domain Controller](#) on page 66. Basic authentication is insecure because it transmits user names and passwords in an easily decoded string, and therefore it should not be used on an untrusted network. If Basic authentication is required, and security is a concern, configure the target system to accept only HTTPS traffic. For more information, see [Manually configuring WinRM HTTPS access](#). If Basic authentication is not acceptable in your environment because of some specific security concerns, it can always be disabled.

Basic and Negotiate authentication schemes are not mutually exclusive when it comes to their configuration. The remote system can be enabled to use both authentication schemes, given a specific order of preference. For example, a remote machine can be configured to first attempt the Negotiate authentication, and it fails, to subsequently try the Basic authentication. Therefore, if you enable Basic authentication, there is no need to disable Negotiate authentication requests, and the other way around. However, you can disable either scheme, as required.

In general, if Basic authentication is not suitable, it should be disabled. If both Negotiate and Basic authentication are enabled, and Negotiate fails, Basic authentication is still attempted, and the credentials are transmitted in an easily decoded string. Because of different credential requirements for Negotiate (a domain user credential) and Basic (a local administrator credential), a credential only applies to one authentication type. Therefore a general recommendation is that if Basic authentication is required, Negotiate authentication does not have to be disabled. But if Negotiate authentication is needed, best practice is to disable Basic authentication requests.

You can also use Windows Group Policy Objects to automatically configure HTTP or HTTPS listeners in WinRM. For more information, see [Using Group Policy Objects to configure WinRM](#) on page 84.

You can also use the **Enable WinRM authentication based on only Basic or Negotiate type** switch in the *FglAMAdapter Properties* view to decide which WinRM authentication scheme will be used.

- *False*: Negotiate authentication scheme will be attempted at first. If Negotiate fails, then Basic authentication will be attempted.
- *True*: Either Negotiate or Basic scheme will be attempted, depending on the type of user credentials:

- WinRM connection will only attempt Basic scheme if it is a local user account.
- WinRM connection will only attempt Negotiate scheme if it is a domain user account.

i | **IMPORTANT:** For WinRM connections on Windows Vista, Service Pack 1 (SP1) or later must be installed.

Configuring the target (monitored) system

Recent versions of Windows® OS include WinRM, but it is disabled by default. There are two ways to configure HTTP or HTTPS: manually or using Group Policy Objects.

Manually configuring WinRM HTTP access

To manually configure the target machine for WinRM authentication:

- 1 Open a command prompt window on the target machine.
- 2 Type the following:

```
winrm quickconfig
```

This enables the default WinRM configuration and disables payload-level encryption.

i | **IMPORTANT:** Payload-level encryption encrypts only the contents of the messages that are exchanged between the Agent Manager and the monitored system. It does not affect the security of any credentials passed between the target system and the Agent Manager. If you require full encryption, configure the WinRM server to use HTTPS. HTTPS safely and securely encrypts all data transmitted between the Agent Manager and the monitored system.

The default WinRM settings allow only Negotiate authentication.

i | **NOTE:** You can enable Basic authentication without disabling Negotiate authentication. Some systems can have both Negotiate and Basic authentication enabled, in that order of preference. If Negotiate authentication fails, a Basic authentication request is attempted.

To enable Basic authentication:

- 1 Open a command prompt window on the target machine.
- 2 Type the following:

```
winrm set winrm/config/service/auth @{Basic="true"}
winrm set winrm/config/service @{AllowUnencrypted="true"}
```

- 3 **Optional.** If Negotiate authentication is enabled, and you want to disable it, type the following:

```
winrm set winrm/config/service/auth @{Negotiate="false"}
```

i | **TIP:** To enable or disable either Basic or Negotiate authentication, use the following command syntax:

```
winrm set winrm/config/service/auth @{<Basic|Negotiate>=<"true|false">}
```

Manually configuring WinRM HTTPS access

If additional security is required, for example if data is transferred across untrusted networks, you can configure WinRM to use HTTPS. A valid server authentication certificate must be installed on the target machine in order to enable HTTPS.

i | **NOTE:** You cannot use self-signed certificates.

NOTE: The target machine is required to enable TLSv1.2 when it is monitored via WinRM HTTPS in FIPS-compliant mode.

The certificate must be granted by a recognized certificate-granting authority (CA) in order for the Agent Manager to authenticate it. Otherwise you must install the root CA certificate in the Agent Manager's trusted keystore, as described in [Installing HTTPS certificates](#).

i | TIP: Install the certificates in the following location: Certificates (Local computer)/Personal/Certificates

To enable HTTPS access:

- 1 Open a command prompt window on the target machine.
- 2 Issue the following command:

```
winrm quickconfig -transport:https
```

The above command enabled HTTPS access using the certificate installed on the host.

If you want to use a different certificate, you can create a new HTTPS listener and specify the certificate:

```
winrm create winrm/config/listener?Address=*&Transport=HTTPS  
@{Hostname="<host>";CertificateThumbprint="<thumbprint>"}
```

Where:

- *host* is a fully qualified host name, as it appears in the certificate.
- *thumbprint* is the certificate thumbprint, with spaces removed.

Using Group Policy Objects to configure WinRM

You can use Windows Group Policy Objects to automatically configure HTTP or HTTPS listeners in WinRM. Enable or disable the appropriate default group policies that are pre-installed with Windows, or use the defaults as a template to develop new policies.

To view the default policies:

- 1 On the target machine, click **Start**.
- 2 Type `run` and press **Enter**.
- 3 In the Run dialog box that appears, type `mmc` and click **OK**.
- 4 The **Console Root** window appears.
- 5 In the **Console Root** window, choose **File > Add/Remove Snap-In**.
- 6 In the **Add or Remove Snap-ins** dialog box that appears, in the **Available** snap-ins area, select **Group Policy Object**, and click **Add**.
- 7 In the **Select Group Policy Object** dialog box that appears, click **Finish** to close it.
- 8 Click **OK** to close the **Add or Remove Snap-ins** dialog box.
- 9 In the **Console Root** window, in the navigation tree on the left, choose **Local Computer Policy > Computer Configuration > Administrative Templates > Windows Components > Windows Remote Management (WinRM) > WinRM Service**.
- 10 Review the settings on the right. Double-click a setting to edit its state.
- 11 After you have edited the settings as necessary for your environment, close the **Console Root** window.

Installing HTTPS certificates

In environments where an in-house certificate granting authority (CA) is in use, the CA's certificate must be added to the Agent Manager's trust keystore. The Agent Manager then assumes that the authority, and any signed certificates it issues, are trusted.

To add a certificate to the keystore:

- 1 Ensure that the Agent Manager is running.

- 2 Launch a command shell on the Agent Manager machine, and navigate to the `<fglam_home>/bin` directory.
- 3 Issue the following command:

```
fglam --add-certificate alias=/path/to/saved.ca.certificate
```

If the import succeeds, the Agent Manager automatically recognizes and uses the certificate.

Configuring the Agent Manager (monitoring) system

The Agent Manager automatically attempts to perform the necessary configuration on the monitoring host machine. If for any reason the changes cannot be made, you have to manually update the following settings.

Configuring Kerberos

The WinRM Negotiate authentication scheme uses Kerberos. By default, the Agent Manager searches the following files for information about the location of the Key Distribution Center (KDC):

- `%WINDIR%/krb5.ini`
- `/etc/krb5.conf`
- `/etc/krb5/krb5.conf`

If the Agent Manager fails to find a configuration file, it attempts to automatically detect the required settings and writes them to `$FGLAM/state/default/config/krb5.config`.

i | **NOTE:** The `fglam.config.xml` file specifies the location of the `krb5` file in the `<kerberos-config-file>` element. If the element is empty or omitted, no configuration file is used.

You can manually override the location of `krb5` file with the following command-line parameter:

```
-Djava.security.krb5.conf=</path/to/file>
```

Generating a configuration file required for WinRM Negotiate authentication

WinRM connections using the default Negotiate authentication require a copy of the `krb5.config` file. The Agent Manager attempts to auto-generate this file and places it under `<fglam_home>/state/default/config/krb5.config`.

If the file needs to be created, the format of the `krb5.config` file for the WinRM Negotiate authentication is as follows:

```
[libdefaults]
default_realm = <dns_suffix_upper_case>
[realms]
<dns_suffix_upper_case> = {
    kdc = <DNS_Server_for_dns_suffix_upper_case>
}
[domain_realm]
.<dns_suffix_lower_case> = <dns_suffix_upper_case>
```

Where:

The values `dns_suffix_upper_case`, `dns_suffix_lower_case`, and `DNS_Server_for_dns_suffix_upper_case` must be replaced with their actual values.

The `[domain_realm]` section in the file maps the domain of the host being connected to, to a *realm*.

The `[realm]` section provides the relevant `kdc` (key distribution center) server with a specific realm to use for kerberos authentication. This is generally the DNS server for the relevant domain.

The `default_realm` value in the `libdefaults` section is the *realm* mapping to use when the domain of the host cannot be matched to a *realm*.

For example, for connecting to hosts on the `sample.domain.com` domain with the `dnsserver.sample.domain.com` DNS Server, the contents of the `krb5.config` file should be as follows:

```
[libdefaults]
default_realm = SAMPLE.DOMAIN.COM
[realms]
SAMPLE.DOMAIN.COM = {
    kdc = DNSSERVER.SAMPLE.DOMAIN.COM
}
[domain_realm]
.sample.domain.com = SAMPLE.DOMAIN.COM
```

When connecting to a `host1.sample.domain.com`, the `host1`'s domain is mapped to the `SAMPLE.DOMAIN.COM` realm, which maps to the `DNSSERVER.SAMPLE.DOMAIN.COM` `kdc` to use for kerberos authentication.

After the `krb5.config` file is created the absolute path to the generated `krb5.config` file should be provided in the `<config:krb5-config-file>` tag value of the `<fglam_home>/state/default/config/fglam-config.xml` file, so that it can be accessed by the Agent Manager. Any changes to the `fglam-config.xml` file require the Agent Manager to be restarted in order for those changes to take effect. Therefore, if the Agent Manager is running while you are making these changes, you must restart it.

Configuring Kerberos on a Windows host

Additional registry keys may be required when you deploy an Agent Manager on a Windows host. The Agent Manager installer attempts to make the changes automatically. If the Agent Manager is unable to establish a WinRM connection, check that the following changes were made correctly.

- In Windows Vista and later versions, User Account Control (UAC) affects access to the WinRM service. When the Negotiate authentication scheme is used in a workgroup only, the built-in Administrator account can access the service.

To allow all accounts in the Administrators group to access the service, set the value of the following registry key to one '1':

```
HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Policies\system\LocalAccountTokenFilterPolicy
```

- By default, Windows does not allow Java™ to access certain required session keys when Java™ attempts to authenticate with Kerberos. The following registry keys should be added to ensure that the required sessions keys are available. The Agent Manager attempts to detect and update these registry keys automatically the first time a WinRM connection is attempted.

Windows 2003, Windows Vista, and later:

- **Key:**
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Kerberos\Parameters
- **Value Name:** allowtgtsessionkey
- **Value Type:** REG_DWORD
- **Value:** 0x01

Windows XP and Windows 2000:

- **Key:** HKEY_LOCAL_MACHINE\System\CurrentControlSet\Control\Lsa\Kerberos
- **Value Name:** allowtgtsessionkey
- **Value Type:** REG_DWORD
- **Value:** 0x01

Configuring Kerberos during runtime

The `KerberosConfigurationService` API provides the ability for agents to modify or create a Kerberos configuration file during runtime.

The Agent Manager uses the Kerberos configuration file to establish WinRM Negotiate connections to hosts. In most cases, the Agent Manager can create the Kerberos configuration for the current domain to which the machine running the Agent Manager belongs. However, the Kerberos configuration typically needs to be modified when cross-domain WinRM connections are required. This can be done by modifying the Kerberos configuration file manually, to add the new domain properties, and restarting the Agent Manager. If no instance of the previous Kerberos configuration file is found, the `fglam.config.xml` file needs to be updated to instruct the Agent Manager which Kerberos configuration file to use for WinRM connections.

All of these actions can also be performed during runtime, without requiring any manual changes, or an Agent Manager restart. The `KerberosConfigurationService` allows agents to make these changes during runtime and have the changes take effect immediately. If a new configuration file is created, `fglam.config.xml` file is updated automatically.

For complete information about this service, see the Foglight Agent Manager Devkit and Javadoc documentation.

Configuring command-shell connection settings

WinRM relies on a set of configuration parameters that establish the level of system resources the WinRM service needs to address incoming requests. In certain cases, some parameter values do not provide sufficient configuration levels which can lead to run-time errors.

Depending on how WinRM is used, some parameter values may not provide sufficient configuration levels which can lead to connection issues. The Agent Manager makes an attempt to diagnose some of these situations and communicate appropriate recommendations using Warning messages. The configuration levels that the Agent Manager attempts to diagnose are:

- **MaxConcurrentOperationsPerUser:** This parameter specifies the maximum number of concurrent Enumeration operations allowed by an individual user. The value must be in the range of 1 to 4294967295.


This parameter is only available for WinRM version 2.0 and later. In version 1.1, the `MaxConcurrentOperations` parameter is used instead.

To increase the value assigned to this parameter, issue the following command:

```
winrm set winrm/config/service @{MaxConcurrentOperationsPerUser="<number>"}
```

i **TIP:** WinRM parameters can also be edited using the Group Policy Object Editor. To start the editor, type `gpedit.msc` at the command line, and then navigate to **Local Computer Policy > Computer Configuration > Administrative templates > Windows Components > Windows Remote Management (WinRM) and Windows Remote Shell**.

Figure 1. Windows Components in the Local Group Policy Editor



- **MaxConcurrentOperations:** This parameter specifies the maximum number of concurrent Enumeration operations allowed by an individual user. Any number from 1 to 4294967295 can be used. For more information about this parameter, you can visit the following Web page: <http://msdn.microsoft.com/en-us/library/cc251426.aspx>.

This parameter is only available for WinRM version 1.1. It is deprecated for version 2.0 and later, and `MaxConcurrentOperationsPerUser` is used instead.

To increase the value assigned to this parameter, issue the following command:

```
winrm set winrm/config/service @{MaxConcurrentOperations="<number>"}
```

- **MaxShellsPerUser:** This parameter specifies the maximum number of concurrent shells any user can remotely open on the same system. Any number from 0 to 2147483647 can be used, where 0 means unlimited number of shells. If this policy setting is enabled, the user cannot to open new remote shells if the count exceeds the specified limit.

To increase the value assigned to this parameter, issue the following command:

```
winrm set winrm/config/winrs @{MaxShellsPerUser="<number>"}
```

- **AllowRemoteShellAccess:** This parameter controls access to the remote shell. It must be set to `true`.

To set this parameter to `true`, issue the following command:

```
winrm set winrm/config/winrs @{AllowRemoteShellAccess="true"}
```

For additional information, visit the following Web page:

<http://msdn.microsoft.com/en-us/library/windows/desktop/aa384372%28v=vs.85%29.aspx>

About WinRM connection ports

WinRM uses a set of default ports for communication. Depending on the WinRM version, the following port numbers are used:

- **WinRM 1.1 and earlier:** The default HTTP port is 80, and the default HTTPS port is 443.
- **WinRM 2.0 and later:** The default HTTP port is 5985, and the default HTTPS port is 5986.

After issuing the `winrm quickconfig` command, the listener port number can be determined using the `winrm enum winrm/config/listener` command. For example:

```
> winrm enum winrm/config/listener
Listener
Address = *
```

```
Transport = <HTTP|HTTPS>
Port = <port>
Hostname
Enabled = true
URLPrefix = wsman
CertificateThumbprint
ListeningOn = <ip_addresses>
```

Troubleshooting

If you have verified all of the WinRM configuration information and are still experiencing connection issues, the following techniques may be helpful for diagnosing the problem.

Enabling specific TLS protocols

If an agent could not establish secure connection to a target server, verify if the target server supports TLS protocol version negotiation. To enable specific TLS protocols on the agent manager, run the Agent Manager with the following switch:

```
-Djdk.tls.client.protocols=TLSv1, TLSv1.1, TLSv1.2
```

i | **NOTE:** To enable specific TLS protocols, specify them in a comma-separated list within quotation marks and all other supported protocols will be disabled. For example, if the value of this property is “TLSv1,TLSv1.1”, then the default protocol settings are only for TLSv1 and TLSv1.1 while other protocols are unavailable.

Verifying setup

To check whether a listener is configured for WinRM, you can issue the following command and observe its output:

```
$ winrm enum winrm/config/listener
Listener
  Address = *
  Transport = HTTP
  Port = 5985
  Hostname
  Enabled = true
  URLPrefix = wsman
  CertificateThumbprint
  ListeningOn = 10.4.114.29, 127.0.0.1, ::1, fe80::100:7f:fffe%12, fe80::5efe:
  10.4.114.29%14, fe80::38a7:8fc9:3d7d:f4d7%13
```

The `Port` and `Transport` elements contain important information. The above command output identifies an HTTP listener on port 5985.

To see a full list of WinRM configuration values for the WinRM service that the Agent Manager is to use on the remote machine, you can issue the following command and observe its output:

```
$ winrm get winrm/config/service
Service
  RootSDDL = O:NSG:BAD:P(A;;GA;;;BA)S:P(AU;FA;GA;;;WD)(AU;SA;GWGX;;;WD)
  MaxConcurrentOperations = 4294967295
  MaxConcurrentOperationsPerUser = 15
  EnumerationTimeoutms = 60000
  MaxConnections = 25
  MaxPacketRetrievalTimeSeconds = 120
  AllowUnencrypted = true
  Auth
    Basic = true
    Kerberos = true
    Negotiate = true
```



```
Certificate = true
CredSSP = false
CbtHardeningLevel = Relaxed
DefaultPorts
  HTTP = 5985
  HTTPS = 5986
IPv4Filter = *
IPv6Filter = *
EnableCompatibilityHttpListener = false
EnableCompatibilityHttpsListener = false
CertificateThumbprint = ce 20 a6 47 d8 0d 71 b3 e2 7c dd f7 70 d9 57 d1 3f a5 65
df
```

The important properties are `AllowUnencrypted` (it indicates whether HTTP is allowed or not), and the `Auth` values that are set to true (enabled), namely `Basic` and `Negotiate`. In this example, both authentication types are enabled.

For more information on setting up for HTTPS, see [About WinRM authentication and the Agent Manager](#) on page 82.

Identifying common causes of WinRM failures

The following conditions can result in a WinRM failure:

- The WinRM listener is not configured on the remote machine.
- The `AllowUnencrypted` WinRM service configuration setting is `false`, and HTTP is used on the remote machine.
- The expected authentication type is not enabled for the WinRM service on remote machine. For example, Basic authentication is expected to be used, but it is not enabled.
- A fully qualified domain name is not specified in the user credential created on the Management Server.
- An incorrect Kerberos configuration file is used. For example, the file does not exist on a UNIX system, or the realm is not set up for cross-realm authentication.

Reviewing application event logs

WinRM logs activity to an event log on the target machine. This includes both success and failure messages for authentication.

To view application event logs:

- 1 On the target machine, right-click **My Computer** and select **Manage**.
- 2 In the navigation tree on the left, choose **System Tools > Event Viewer > Applications and Services Logs > Microsoft > Windows > Windows Remote Management > Operational**.

The default *Operational* log contains the most common events.

To enable additional debug logging information:

- 1 Click **View**.
- 2 Click **Show Analytic and Debug Logs**.
- 3 Right-click the log file you want to view.
- 4 Select **Enable Log**.

Enabling connection type debugging

If the only information you are interested in is the types of connections that are being established, there is a command-line setting that enables logging the connection types.

Run the Agent Manager with the following switch:

```
-Dquest.debug.windowsinfo.types
```

i | **NOTE:** This logging occurs every time a connection is established and can be very verbose. It is recommended for debugging purposes only.

UNIX- and Linux-specific configuration

This section contains platform-specific configuration information for configuring the Foglight Agent Manager on UNIX® or Linux®.

This section provides solutions for the following issues:

- [Agent Manager service can't start automatically when the operating system restarts](#)
- [SSH IPv6 connection support](#)
- [About supported remote monitoring protocols](#)
- [Configuring the Agent Manager to run as a daemon](#)
- [Configuring Agent Manager agent privileges](#)
- [Preventing Agent Manager core dumps on Linux](#)

Agent Manager service can't start automatically when the operating system restarts

When the Agent Manager service is running in the following platforms, it might not be able to start automatically when the operating system restarts.

| Operating system platform | Operating system version |
|---------------------------|--------------------------|
| CentOS Linux | 8.0 |
| | 8.1 |
| | 8.2 |
| Red Hat Linux | 8.1 |
| | 8.2 |
| Oracle Linux | 8.0 |
| | 8.1 |
| | 8.2 |
| SLES Linux | 15 |
| | 15 SP1 |
| | 15 SP2 |

To fix this issue, follow the instructions provided below:

Use the `ausearch` utility to check the Access Vector Cache (AVC) messages and see if SELinux denies any of the FglAM actions:

```
# ausearch -m AVC,USER_AVC -ts today
time->Wed Nov 4 11:18:11 2020 type=AVC msg=audit(1604459891.164:117): avc: denied {
open } for pid=1311 comm="fglam" path="/root/
5981fips/jre/1.8.0.265/jre/lib/jce.jar" dev="dm-0" ino=11429653
scontext=system_u:system_r:init_t:s0 tcontext=unconfined_u:object_r:admin_home_t:s0
tclass=file permissive=1
```

The `-m` option specifies what kind of information `ausearch` returns. The `-ts` option specifies the time stamp. For example, `-ts today` returns messages from the whole day.

- If any FglAM action has been denied, there are two options to fix the issue:
 - Disable SELinux service.
 - If you don't want to disable SELinux, do the following:
 - a Open the `/etc/selinux/config` file and change SELinux mode to permissive. Using permissive mode will force SELinux to accept all FglAM actions. SELinux will log all the denials regarding to FglAM actions that would have been denied in enforcing mode, by identifying them one at a time as the FglAM gets permissions granted individually.
 - b Restart FglAM machine.
 - c Ensure FglAM service starts automatically. Then try all of the functions that FglAM and agents would perform, such as deploying agent gars, creating agent instances, releasing lockbox to FglAM, and so on. Therefore, it will reveal all the FglAM actions that would have been denied by SELinux if running in enforcing mode.
 - d Use the `'journalctl -t setroubleshoot --since= [time]'` utility to view more information about the AVC message:

```
# journalctl -t setroubleshoot --since=11:18
- Logs begin at Tue 2020-11-03 10:37:14 CST, end at Wed 2020-11-04
11:19:27 CST. - Nov 04 11:18:30 centos82-s1 setroubleshoot[1416]:
SELinux is preventing quest-fglam from execute access on the file
fglam. For complete SELinux messages run: sealert -l 06149362-e530-
4f52-a081-53751a98eab7
Replace [time] with the machine restart time.
```
 - e Use the `'sealert -l [AVC message ID]'` utility to further inspect the AVC message:

```
# sealert -l 06149362-e530-4f52-a081-53751a98eab7
SELinux is preventing quest-fglam from execute access on the file
fglam.
***** Plugin catchall (100. confidence) suggests*****
If you believe that quest-fglam should be allowed execute access on
the fglam file by default. Then you should report this as a bug. You
can generate a local policy module to allow this access. Do allow
this access for now by executing:
# ausearch -c 'quest-fglam' --raw | audit2allow -M my-questfglam
# semodule -X 300 -i my-questfglam.pp
[trimmed for clarity]
```
 - f Perform actions according to suggestions provided in [Step e](#).
 - g Repeat [Step e](#) to [Step f](#) for all FglAM action denials AVC messages found in [Step d](#).
 - h Restore SELinux to enforcing mode and restart FglAM machine.
 - i Check if there are still denials about FglAM actions. If yes, repeat [Step a](#) to [Step i](#) until no denials to FglAM actions are found.
- If it is not caused by SELinux, perform below command to check if it works.

```
systemctl enable quest-fglam.service
```

If it doesn't work, fix the issue according to the error. For example, it may report below error, which instructs to install the tool `insserv` first and then run above command again to fix this issue.

```
Executing: /usr/lib/systemd/systemd-sysv-install enable quest-fglam
```

```
/sbin/insserv: No such file or directory
```

SSH IPv6 connection support

Starting with the Foglight Agent Manager version 5.9.1, the SSH connection with unique local IPv6 Address and link-local IPv6 Address is supported on the Agent Manager running on Windows and Linux.

About supported remote monitoring protocols

The Agent Manager supports the SSH (secure shell) protocol for remote monitoring of hosts running Linux[®] and UNIX[®] operating systems. SSH is a protocol which encrypts all traffic between the client and the server, and supports a wide variety of secure authentication mechanisms. SSH is available for installation on all platforms supported for remote monitoring by Foglight.

The Agent Manager does not support the older Telnet protocol for remote monitoring. Telnet is an insecure protocol which does not encrypt traffic, and requires that the passwords used to authenticate collections are sent in the clear. For that reason, supporting Telnet as a remote monitoring protocol potentially exposes monitored systems to trivial network eavesdropping attacks, disclosing passwords to attackers.

While it is possible to create a closed network with strong security boundaries within which it is safe to run the Telnet protocol, this use case is not common, and it is impossible for the Agent Manager to determine if Telnet is safe to use before offering it as an option for connectivity. Further, the effort required on the part of the user to maintain such a secure environment is far less than that required to simply enable SSH connections on a host. For these reasons, the Agent Manager does not support Telnet as a remote monitoring protocol.

Configuring the Agent Manager to run as a daemon

As described in [Installing the Agent Manager using the installer interface](#) on page 15 and [Installing the Agent Manager from the command line](#) on page 22, you can install an *init.d*-style script called *quest-fglam* in the *init.d* directory on your system. This script is called when the host on which the Agent Manager is installed starts or shuts down, allowing it to run as a daemon.

Even if you choose not to install the *init.d* script during the installation, or if you do not perform the installation as the root user, the installer generates scripts that can perform the necessary setup.

These scripts are *fglam-init-script-installer.sh* and *fglam-init-script.sh*, and they are located in the `<fglam_home>/state/default/` directory.

The script *fglam-init-script-installer.sh* installs the script *fglam-init-script.sh* into your system's *init.d* directory as *quest-fglam*. Your system's *init.d* process then uses *quest-fglam* to run the Agent Manager as a daemon.

To install the *init.d* script:

- 1 Launch a command shell on the Agent Manager machine and navigate to the `<fglam_home>/state/default/` directory.
- 2 **Optional.** If you want to make any edits to *fglam-init-script.sh* to customize it for your system, do so prior to running *fglam-init-script-installer.sh*.

i | **IMPORTANT:** Any customizations that you make to the script *fglam-init-script.sh* are not supported by Quest Software Inc..

- 3 Switch to the root user.
- 4 From the command shell, run the script *fglam-init-script-installer.sh* with the `install` option:

```
./fglam-init-script-installer.sh install
```

i | **IMPORTANT:** This script must be run as root.

The setup script `fglam-init-script-installer.sh` installs the `init.d` script `quest-fglam`. See [Locating the `init.d` script](#) on page 94 for the location in which it is installed.

- 5 To start or stop the Agent Manager daemon manually, follow the instructions in [To run the Agent Manager as a daemon on UNIX®](#): on page 35.

To remove the `init.d` script, follow the instructions in [To remove the `init.d` script used to run the Agent Manager as a daemon on UNIX®](#): on page 36.

Locating the `init.d` script

Depending on the operating system you are running, the `init.d`-style script `quest-fglam` is installed to a different location either by the Agent Manager installer or after you run the script `fglam-init-script-installer.sh`.

The location of the installed `init.d` script (listed by operating system) is:

- **All Linux® operating systems:** `/etc/init.d`

i | **NOTE:** The location of the `init.d` script depends only on the type of operating system, not on the specific architecture.

Obtaining the Agent Manager daemon status

In addition to starting or stopping the Agent Manager process, the `init.d` script allows you to obtain the status of the daemon process when you run the script with the `status` option. When the `status` option is specified with the `init.d` script, the script returns one of the following status codes:

- **0:** The Agent Manager daemon process is running.
- **1:** The Agent Manager daemon process is dead and a `pid` file is generated.
- **3:** The Agent Manager daemon process is not running.
- **4:** The Agent Manager daemon process status is unknown.

Configuring Agent Manager agent privileges

On UNIX® systems, certain Foglight® agents require elevated privileges in order to gather the required system metrics. This is achieved by having the Agent Manager launch these agents with root privileges.

To give these agents the required access, the Agent Manager is configured to launch these agents using an external application like `sudo`, `setuid_launcher`, or any other tool that allows privilege escalation (without a password) and supports the same command-line semantics as `sudo`.

i | **NOTE:** The tool `setuid_launcher` is included with the Agent Manager, in the `<fglam_home>/bin/setuid_launcher` directory.

Instructions for using `sudo` and `setuid_launcher` to give these agents the necessary privileges are provided below.

i | **NOTE:** Certain agents that require root privileges to gather a more complete set of system metrics are able to function without these privileges. See the *Managing Operating Systems User Guide* for more information. If an agent is configured to be launched by an external application and fails to start, the Agent Manager logs a warning and then tries starting the agent without the launcher and without root privileges. The agent does not collect as much data as when it is run with root privileges.

Using *sudo* to configure secure launcher permissions

This section contains instructions for using *sudo* to give agents elevated permissions.

To set up secure launcher permissions using the configuration interface and *sudo*:

- 1 Follow the instructions in [To launch the Agent Manager Installation and Configuration interface](#): on page 40 or [To launch the Agent Manager configuration command-line interface](#): on page 42.
- 2 Navigate to the **Configure Secure Launcher** or **Secure Launcher** step.
- 3 Set the path to point to the *sudo* executable. This executable is typically located in */usr/bin/sudo* (the default path provided by the Agent Manager installer).
- 4 Exit from the configuration interface as described in [To launch the Agent Manager Installation and Configuration interface](#): on page 40 or [To launch the Agent Manager configuration command-line interface](#): on page 42.

- 5 Edit the *sudoers* file for your system to allow `<fglam_home>/client/<fglam_version>/bin/fog4_launcher` to be run as root by a specific user, without requiring a password, and only for the agents that require root privileges.

For example, to allow the user *foglight* to execute *fog4_launcher* for two specific agents without being prompted for a password:

```
foglight    ALL = NOPASSWD: \
/<fglam_home>/client/*/bin/fog4_launcher
/<fglam_home>/state/default/<cartridge>*/bin/<agent> ?*@?* bin/<agent>, \
/<fglam_home>/client/*/bin/fog4_launcher
/<fglam_home>/state/default/<cartridge>*/bin/<agent2> ?*@?* bin/<agent2>
```

The example above also limits the acceptable arguments to match the expected pattern when the Agent Manager runs the agents.

- 6 Ensure that the `requiretty` option is disabled in the *sudoers* file. For example, to disable this option for the *foglight* user, add the following entry to the file:

```
Defaults:foglight !requiretty
```

- 7 If the agent uses an ICMP ping service, edit the *sudoers* file for your system to allow `<fglam_home>/client*/bin/udp2icmp` to be run as root by a specific user, without requiring a password.

For detailed examples of how to edit the *sudoers* file to restrict the granted permissions to a specific set of agents, see the *Foglight for Infrastructure User and Reference Guide*.

i **TIP:** For *sudo* configuration, it is a best practice to use a wildcard for the version-specific Agent Manager and cartridge directories, as shown in the example above. Using a wildcard in a path is described in the *Sudoers Manual* located at:

<http://www.gratisoft.us/sudo/man/sudoers.html#wildcards>

Using a wildcard for the version-specific directories allows you to avoid updating each *sudoers* file that references these directories when you upgrade the Agent Manager or the agents.

If these permissions are no longer needed, remove the lines that you added to run *fog4_launcher* or *udp2icmp* with root permissions.

To set up secure launcher permissions using *fglam.config.xml* and *sudo*:

- 1 Navigate to `<fglam_home>/state/default/config`.
- 2 Open the *fglam.config.xml* file for editing.
- 3 Edit the `<config:path>` element under `<config:secure-launcher>` to point to the *sudo* executable. This executable is typically located in */usr/bin/sudo* (the default path provided by the Agent Manager installer).

- 4 Edit the *sudoers* file for your system to allow `<fglam_home>/client/<fglam_version>/bin/fog4_launcher` to run as root by a specific user, without requiring a password, and only for the agents that require root privileges.

For example, to allow the user *foglight* to execute *fog4_launcher* for two specific agents without being prompted for a password:

```
foglight    ALL = NOPASSWD: \  
/<fglam_home>/client/*/bin/fog4_launcher  
/<fglam_home>/state/default/<cartridge>*/bin/<agent> ?*@?* bin/<agent>, \  
/<fglam_home>/client/*/bin/fog4_launcher  
/<fglam_home>/state/default/<cartridge>*/bin/<agent2> ?*@?* bin/<agent2>
```

The example above also limits the acceptable arguments to match the expected pattern when the Agent Manager runs the agents.

- 5 If the agent uses an ICMP ping service, edit the *sudoers* file for your system to allow `<fglam_home>/client/*/bin/udp2icmp` to be run as root by a specific user, without requiring a password.

See the *Managing Operating Systems User Guide* for detailed examples of how to edit the *sudoers* file to restrict the granted permissions to a specific set of agents.

i **TIP:** For *sudo* configuration, it is a best practice to use a wildcard for the version-specific Agent Manager and cartridge directories, as shown in the example above. Using a wildcard in a path is described in the *Sudoers Manual* located at:

<http://www.gratisoft.us/sudo/man/sudoers.html#wildcards>

Using a wildcard for the version-specific directories allows you to avoid updating each *sudoers* file that references these directories when you upgrade the Agent Manager or the agents.

Using *setuid_launcher* to configure secure launcher permissions

This section contains instructions for using *setuid_launcher* to give agents elevated permissions.

To set up secure launcher permissions using the configuration interface and *setuid_launcher*:

- 1 Follow the instructions in [To launch the Agent Manager Installation and Configuration interface](#): on page 40 or [To launch the Agent Manager configuration command-line interface](#): on page 42.
- 2 Navigate to the Configure Secure Launcher screen or the Secure Launcher step.
- 3 Set the path to point to the *setuid_launcher* executable. This executable is located in `<fglam_home>/bin/setuid_launcher`.
- 4 Exit from the configuration interface as described in [To launch the Agent Manager Installation and Configuration interface](#): on page 40 or [To launch the Agent Manager configuration command-line interface](#): on page 42.
- 5 Use the command `chmod u+s` to set the sticky bit on `<fglam_home>/bin/setuid_launcher`.
- 6 Change the owner of `<fglam_home>/bin/setuid_launcher` to *root*. This permits the agents that need root privileges to be run as the *root* user without requiring a password.

If these permissions are no longer needed, issue the following command:

```
chmod u-s <fglam_home>/bin/setuid_launcher
```

To set up secure launcher permissions using *fglam.config.xml* and *setuid_launcher*:

- 1 Navigate to `<fglam_home>/state/default/config`.
- 2 Open the *fglam.config.xml* file for editing.
- 3 Edit the `<config:path>` element under `<config:secure-launcher>` to point to your local *setuid_launcher* executable. This executable is located in `<fglam_home>/bin/setuid_launcher`.

- 4 Issue the command `chmod u+s` to set the sticky bit on `<fglam_home>/bin/setuid_launcher`.
- 5 Change the owner of `<fglam_home>/bin/setuid_launcher` to `root`. This permits the agents that need root privileges to be run as the `root` user without requiring a password.
- 6 If these permissions are no longer needed, issue the command:

```
chmod u-s <fglam_home>/bin/setuid_launcher
```

Preventing Agent Manager core dumps on Linux

Installing and running the Agent Manager on a Linux[®] machine with several interfaces can result in a core dump with the following console output:

```
*** glibc detected ***  
...  
malloc(): memory corruption:  
...  
Java_java_net_NetworkInterface_getAll+0x8c
```

This error is related to a known Java[™] 6 issue, JDK-7078386. This issue is resolved in Java 7. For more information about JDK-7078386, you can visit http://bugs.java.com/bugdatabase/view_bug.do?bug_id=7078386.

To prevent an Agent Manager core dump on a Linux machine:

- 1 Before installing the Agent Manager, install Java 7 update 45 on the machine where you plan to install the Agent Manager.
- 2 On this machine, install the Agent Manager, and navigate to the following directory:
`<fglam_home>/state/default/config`.
- 3 In this directory, locate and open the `vm.config` file for editing.
- 4 In the `vm.config` file, search for the `java.vm` property, and edit its value to point to the new JDK 7 installation.
- 5 Save your changes to the `vm.config` file, and close it.
- 6 Start up the Agent Manager.

Monitoring the Agent Manager performance


Foglight® uses the Agent Manager to communicate with monitored hosts. The embedded Agent Manager can be used to monitor the host on which the Management Server is installed. Your monitoring environment typically includes a number of Agent Manager instances installed on different hosts.

- [Investigating Agent Manager diagnostics](#)

Investigating Agent Manager diagnostics

You monitor the state of your Agent Manager instances and the related Management Server adapters using the FglAM Performance Overview dashboard. Use this dashboard to better find out how these components perform over time and to look for any indicators that may predict potential bottlenecks. 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 > Agent Manager**.

Figure 1. FglAM Performance Overview dashboard



The information appearing on this dashboard appears on two major tabs, **Adapter Details** and **Client Details**, each consisting of several sub-tabs. For more information about the data appearing on this tab, see the following topics:

- [Exploring the Adapter Details tab on page 100](#)
- [Exploring the Client Details tab on page 101](#)

Exploring the Adapter Details tab

The **Adapter Details** tab allows you to investigate the performance of the Management Server adapters that communicate with the connected hosts. Use it to see the overall processing rates and to estimate your system's load. Higher processing rates may lead to decreased performance, requiring further investigation.

For more information about the data appearing on this tab, see the following topics:


- [Message Processing tab on page 100](#)
- [XML Serialization tab on page 100](#)
- [Servlets tab on page 101](#)

Message Processing tab

The **Message Processing** tab contains graphs that tell you how well the adapter is handling incoming messages. It displays the numbers of connected hosts and agents, the number of incoming queue threads, and the numbers of pushed and pending messages over time. It also shows the minimum and maximum polling intervals, the rate of incoming messages, and the percentage of time the adapter spends on message processing.

High peaks in the graphs likely indicate a potential performance bottleneck. For example, a buildup of pending messages can cause delays in your monitoring environment and should be investigated.


Figure 2. Message Processing tab



XML Serialization tab

The **XML Serialization** tab shows graphs indicating the rates and times of message serialization to and from XML used by the adapter over the selected time range. High peaks in the graphs can indicate signs of performance decrease and may need to be investigated.


Figure 3. XML Serialization tab



Servlets tab

The **Servlets** tab displays graphs that indicate the rates of read, write and request rates for the individual Adapter servlets: transfer, installer, and servlet, over the selected time range. It also shows the percentage of time the adapter spends on the requests coming from these components.

Figure 4. Servlets tab



Exploring the Client Details tab

The **Client Details** tab allows you to investigate the performance of the Agent manager clients adapters that are connected to the FglAM Adapter. Use it to see their overall performance rates and to investigate any potential bottlenecks.

For more information about the data appearing on this tab, see the following topics:

- [Agent Manager Clients view](#) on page 102

- [Message Processing tab](#) on page 102
- [Clock Skew tab](#) on page 103
- [Bandwidth tab](#) on page 103
- [Misbehaving Clients tab](#) on page 103
- [XML Serialization tab](#) on page 104
- [CPU Usage tab](#) on page 104
- [Memory Usage tab](#) on page 105
- [CDT Submission tab](#) on page 105
- [Queue Utilization tab](#) on page 106

Agent Manager Clients view

This view, accessible by clicking **Expand to Select More Foglight Agent Manager Clients** at the top of the **Client Details** tab, shows a list of the connected Agent Manager instances. For each Agent Manager instance, it shows the instance name, the numbers of pending messages and incoming queue threads, the upstream time difference, the number of disconnections, and the amounts of used physical and swap memory.

Figure 5. Agent Manager Clients view

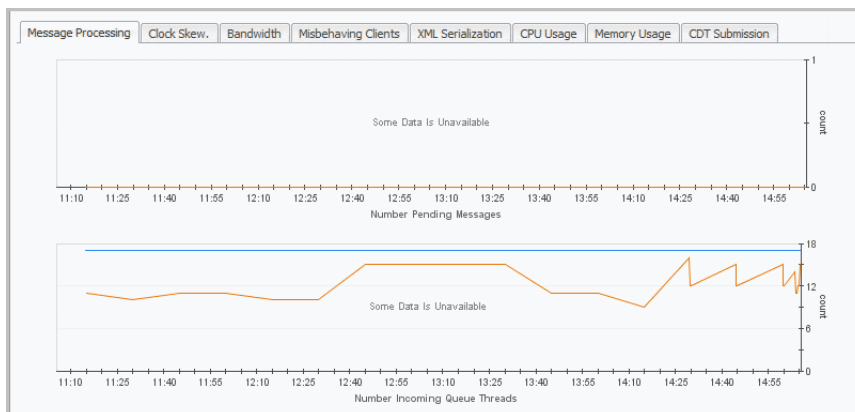
| <input checked="" type="checkbox"/> | Name | Pending Messages | Incoming Queue Threads | Upstream Diff | Disconnections | Physical Memory | Swap Memory |
|-------------------------------------|-------------------------------|------------------|------------------------|---------------|----------------|-----------------|-------------|
| <input checked="" type="checkbox"/> | foglight-nightcrawler (fglam) | 0 | 10 | -2 ms | 0 | 2.475 GB | 0.719 GB |

To review details about one or more Agent Managers, select it in the list, and review the data displayed on the tabs below.

Message Processing tab

The **Message Processing** tab shows tables for the number of pending messages and the number of incoming queue threads for a selected Agent Manager instance.


Figure 6. Message Processing tab



Clock Skew tab

The **Clock Skew** tab contains a graph that displays the upstream difference per second for a selected Agent Manager instance.

Figure 7. Clock Skew tab

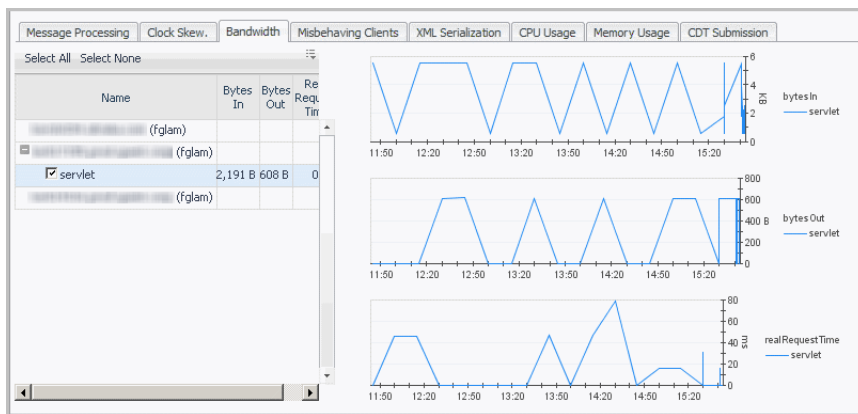


Bandwidth tab

The **Bandwidth** tab displays metrics for bandwidth usage samples taken at prescribed intervals per second for a selected Agent Manager instance, including the bytes in, bytes out, and real request time for a selected servlet.

The bandwidth statistics reflect the bandwidth (bytes per second) used by the Agent Manager instance data over a specific time interval.


Figure 8. Bandwidth tab



Misbehaving Clients tab

The **Misbehaving Clients** tab displays the disconnections per count for a selected Agent Manager instance.

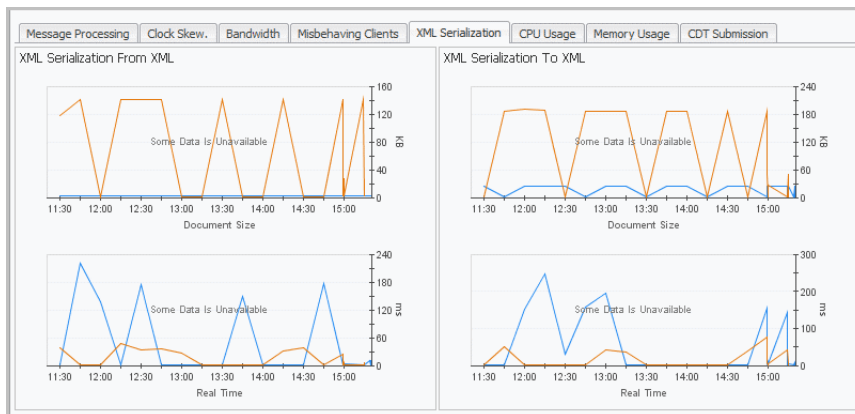
Figure 9. Misbehaving Clients tab



XML Serialization tab

The **XML Serialization** tab shows graphs indicating the rates and times of message serialization to and from XML for a selected Agent Manager instance.


Figure 10. XML Serialization tab



CPU Usage tab

The **CPU Usage** tab displays the amount of time the CPU associated with a selected Agent Manager instance spends executing active processes and their number. For example, a sudden increase in CPU time may indicate that the user code is running inefficiently or a possible runaway process. Also, high CPU loads sometimes suggest that the host needs more CPU power to run efficiently.


Figure 11. CPU Usage tab



Memory Usage tab

The **Memory Usage** tab displays the amount of available memory and swap space for a selected Agent Manager instance. For example, a shortage of swap space often suggests a memory shortage, while a decline in the available memory may indicate a memory leaking process.

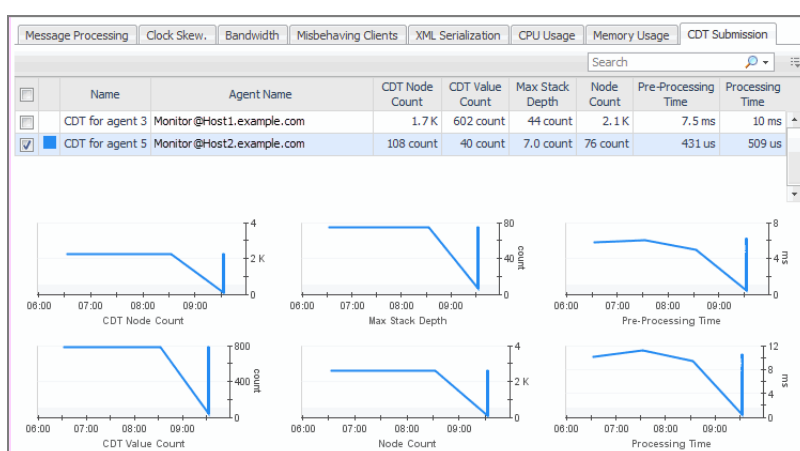
Figure 12. Memory Usage tab



CDT Submission tab

The **CDT Submission** tab displays information about the complexity of data that agents collect and submit for Canonical Data Transform (CDT) processing. Excessively large values or sudden increases in this data can result in performance problems.

Figure 13. CDT Submission tab



The following metrics are displayed for each agent:

- **CDT Node Count:** The number of internal data nodes produced by the CDT. Typically, there is one node for each topology object and its metrics, and one global (root) node in the tree.

This value takes into consideration the count of the topology object submission events that have associated observations or metrics. Topology objects that do not have any metrics associated with them may go through the CDT without affecting this count. Additionally, one topology object may be submitted several times with different metric values, then counted several times.

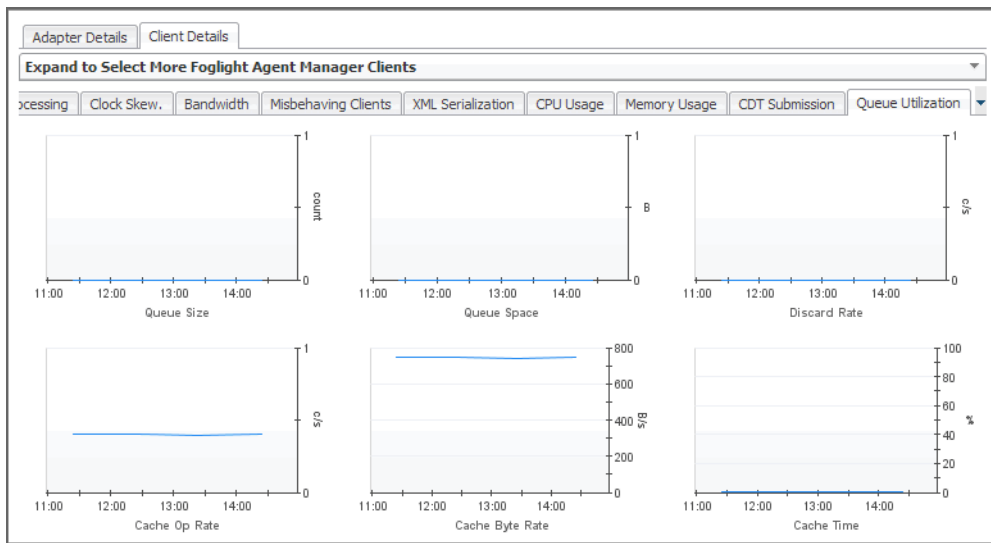
- **CDT Value Count:** The number of metrics and/or observations produced by the CDT.
- **Max Stack Depth:** The depth of the data tree structure submitted by each agent instance. When a topology object in the data submission refers to another topology object, that increases the depth by one. When the second object refers to a third, the depth is increased again.

- **Node Count:** The number of nodes in the submission. It is a measure of how large the data submission is.
- **Pre-processing time:** The Agent Manager makes two passes over the data submitted by an agent. This is the time spent on the first pass, when the Agent Manager fills in references, validates the submission, calculates missing timestamps, and performs other pre-processing tasks.
- **Processing time:** The length of time it takes the Agent Manager to transfer the data submission to the CDT processing engine of the Management Server.

Queue Utilization tab

The **Queue Utilization** tab contains graphs that tell you how well the client queue and cache are performing. It displays the count of messages in the queue over time, its size, and the rate at which it discards messages. It also shows the rates of cache-related operations and memory consumption, and the percentage of time the client spends on cache-related operations.

Figure 14. Queue Utilization tab



Deploying the Agent Manager to large-scale environments

! **IMPORTANT:** The Agent Manager deployment script discussed in this section is provided as an example only, and is not supported as part of the product warranty.

This chapter contains tips for deploying the Foglight® Agent Manager to a large number of UNIX® or Windows® hosts in your environment. It also discusses some of the deployment options you might want to consider.

Deploying the Foglight Agent Manager is a multi-step process that includes selecting and deploying an Agent Manager cartridge, downloading the Agent Manager installer, and installing the Agent Manager. The steps in this process are described in [Installing external Agent Managers](#) on page 10.

If you are planning to deploy the Agent Manager to multiple UNIX hosts, you can create a deployment script to simplify this process. An example of this type of script is outlined below, along with suggestions for how you could customize the script to suit your environment.

If you are planning to deploy the Agent Manager to multiple Windows hosts, it is likely that you need a software deployment tool. Suggestions for using this type of tool are provided below.

Using Agent Manager silent installer Parameters

The instructions in this chapter include references to the Agent Manager silent installer parameters.

These parameters are used with the `--silent` command option when you run the Agent Manager installer non-interactively—for example, to install the Agent Manager from the command-line onto remote machines, as described in the examples in this chapter.

See [Using the Agent Manager silent installer](#) on page 31 for more information about these parameters.

! **CAUTION:** As with deploying the Agent Manager to multiple hosts, using the silent installer is an advanced activity.

Example: Deploying the Agent Manager to multiple UNIX hosts

If you are planning to deploy the Agent Manager to multiple UNIX® hosts, you can simplify the process by creating a deployment script that downloads installers, installs, and starts the Agent Manager on these hosts.

This section provides an example of this type of script and suggestions for customizing it to suit your environment.

Working with this example

This example is based on using the *Wget* network utility to download Management Server installers, providing a specific set of configuration options with the Agent Manager silent installer command, and using passwords to log in to remote hosts over SSH.

However, you could instead create a script that works with any file-retrieval tool that downloads content using HTTP, such as *cURL* or a Perl script that uses the `LWP::Simple` interface.

Additionally, if you require non-default silent installer parameters, such as those used to connect to the Management Server through a proxy, you could edit the parameters used when invoking the Agent Manager silent installer to suit your environment. See [Using Agent Manager silent installer Parameters](#) on page 107 for more information.

The end of the example, [Part 5: Run the script](#) on page 112, describes using one type of login method (passwords) for connecting to remote machines using SSH. However, if you have private keys configured in your environment, you might encounter different scenarios when establishing an SSH connection, which are described in [Part 5: Run the script](#) on page 112.

! **CAUTION:** Using a deployment script for the Agent Manager is an advanced activity. You should create and use this type of script only if both of the following conditions apply:

- You are a Foglight® administrator who is familiar with configuring the Agent Manager
- You are certain of the setup that is required for your environment.

i **NOTE:** If you are not certain which installation options you need, use the Agent Manager Installation and Configuration or command-line installer on each machine, instead.

Before you begin

This example deployment script is designed to use a specific set of tools and protocols and depends on certain settings being configured in your environment.

If you decide to follow this example to create your own deployment script, you need to perform certain steps before you begin:

- Enable SSH on the hosts to which you want to deploy the Agent Manager. The example describes a script that uses the `ssh` command.
- Install the file retrieval tool *Wget* on all machines to which you want to deploy the Agent Manager. The example describes a script that uses the `wget` command to retrieve the Agent Manager installers from the Management Server.
See <http://www.gnu.org/software/wget/> for more information about *Wget*.
- Ensure that the remote hosts to which you are deploying the Agent Manager have enough space in the `/tmp` directory for the installer to be downloaded to that location.

Remember that you do not need to follow all parts of the example as described. For example, if you create a script, you might choose to use *cURL* instead of *Wget* to retrieve the Agent Manager installers.

Example: Creating and running an Agent Manager deployment script

This section outlines the creation and use of an example deployment script and contains suggestions for customizing this type of script to suit your environment.

If you create your own script by following this example, you can use it to install and run the Agent Manager on the hosts you specify.

CAUTION: Line breaks are added to the commands in the example script below. These are highlighted by backslashes (\) at the end of each broken line. If you copy and paste from this example to create your own deployment script, ensure that you remove these backslashes and the subsequent line break.

Part 1: Create the working directory and script file

On the machine from which you deploy the Agent Manager to the remote hosts, create a working directory and create the script file within it using a text editor such as *vi*.

In this example, the working directory is *fglaminstall* and the script file is *remote_install_fglam_example.sh*.

Part 2: Create a list of hosts

The *remote_install_fglam_example.sh* script that you are creating depends on a text file that lists the names of the hosts to which you are deploying the Agent Manager.

In this example, the text file is called *hostnames*, is located in the *fglaminstall* directory, and includes a list of comma-separated host name and operating system pairs. List all of the UNIX® hosts to which you want to deploy the Agent Manager. For example:

```
hostname2,linux64
hostname3,linux32
```

The same format that is used in the example above to specify the operating systems is used in [Part 4: Specify the script commands](#) on page 110.

Part 3: Add the script parameters

After performing the preliminary steps, you are ready to start configuring the Agent Manager deployment script.

In this part of the example, you will add the list of parameters for the script.

NOTE: The default values shown for these parameters are examples only. The steps below describe how to configure these example settings for your environment.

- 1 Navigate to the *fglaminstall* working directory.
- 2 Open *remote_install_fglam_example.sh*.
- 3 Add the following lines, which describe the parameters you are about to set for the script:

```
#!/bin/sh
# Script to download, install, and run the Agent
# Manager
#
# THIS CODE AND INFORMATION ARE PROVIDED "AS IS" WITHOUT
# WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED,
# INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF
# MERCHANTABILITY AND/OR FITNESS FOR A PARTICULAR PURPOSE.
#
#####
# Location of the hostnames file for mass deployment.
HOSTFILE=/fglaminstall/hostnames

# Location of the Management Server.
FMS=[http|https]://<hostname>:<port>

# Operating system user name to run the installation on the
# remote hosts.
```

```

FGLAM_USER=foglight

# Installation directory, which should be writable by the
# Agent Manager user.
INSTALL_DIR=/opt/Quest_Software/Foglight_Agent_Manager

# Management Server version (specified without punctuation, for example:
# 555 for version 5.5.5).
FMS_VER=558
#####
# The following lines do not require configuration.
# URL for downloading installers (requires version 5.2.4+ of
# the Management Server).
FMS_INSTALLERS=\
/console/installer/download-installer.action?downloadId=
if [ $FMS_VER -lt 550 ]; then
FMS_INSTALLERS=/catalyst-glue-service/installers/
fi

```

- 4 Change the value of `HOSTFILE` to the location of the text file *hostnames* that you created in [Part 2: Create a list of hosts](#) on page 109.
- 5 Change the value of `FMS` to the URL of the Management Server (version 5.2.4 or higher) from which you want to retrieve the Agent Manager installers and to which the Agent Manager instances connect. Use the format

```
<http|https>://<hostname>:<port>
```

where `<hostname>` is the resolvable DNS name or IP address and `<port>` is the HTTP port specified during installation (the default port is 8080).

If the Management Server has been configured to use HTTPS, then you can specify `https` as the protocol used by the Foglight Agent Manager to connect to the Management Server.
- 6 Change the value of `FGLAM_USER` to the name of the user that your UNIX administrator created for you to use when running the Agent Manager.
- 7 Change the value of `INSTALL_DIR` to the directory on each machine where the Agent Manager is being installed.
- 8 Change the value of `FMS_VER` to the version number of the Management Server (version 5.2.4 or higher) from which you want to retrieve the Agent Manager installers and to which the Agent Manager instances will connect. Specify the version number without punctuation, for example: 567, for Management Server version 5.6.7.
- 9 Save *remote_install_fglam_example.sh*.

Part 4: Specify the script commands

The steps below describe the process of specifying the commands to download and run the installers on the remote hosts and to run the installed Agent Manager instances.

- 1 Add the lines below after those that you added in [Part 3: Add the script parameters](#) on page 109. These lines verify the existence of the *hostnames* file and then process it.

```

#####
# Check that the HOSTFILE exists.
if [ -r $HOSTFILE ]; then

# Start the For loop to process each host in the HOSTFILE.
for i in $(cat $HOSTFILE);
do
# Parse the HOSTFILE to get the host name and then the
# installer type.
HOSTNAME=`echo $i | awk -F , '{print $1}'`

```

```

OSTYPE=`echo $i | awk -F , '{print $2}'`

if [ $FMS_VER -ge 550 ]; then
FMS_VER_MAJOR=`echo $FMS_VER | cut -c1`
FMS_VER_MINOR=`echo $FMS_VER | cut -c2`
FMS_VER_MICRO=`echo $FMS_VER | cut -c3`
INSTALLER_PREFIX=fglam-
${FMS_VER_MAJOR}_${FMS_VER_MINOR}_${FMS_VER_MICRO}-
else
INSTALLER_PREFIX=fglam-
fi

# Match the operating system to the installer.
case $OSTYPE in

linux64)
INSTALLER=${INSTALLER_PREFIX}linux-x86_64.bin
;;
linux32)
INSTALLER=${INSTALLER_PREFIX}linux-ia32.bin
;;

*)
echo "Cannot match the operating system type ($OSTYPE) for \
$HOSTNAME, skipping the deployment of the Agent \
Manager to that host."
continue;
;;
esac

```

2 Add the lines below to the end of the file. These lines run the `ssh` command that:

- Retrieves the installers using `Wget`.
- Uses the silent installer to install the Agent Manager on each remote UNIX host specified in the *hostnames* file. The Agent Manager silent installer is invoked in this example with the command `/tmp/${INSTALLER} --silent --fms url=${FMS} --installdir ${INSTALL_DIR}`.
- Runs the Agent Manager on the specified hosts.

```

#####
# Run the SSH command to install the Agent Manager.
ssh $FGLAM_USER@$HOSTNAME "cd /tmp; wget -q \
--tries=10 \"${FMS}${FMS_INSTALLERS}${INSTALLER}\" \
-O ${INSTALLER}; chmod a+x /tmp/${INSTALLER}; \
/tmp/${INSTALLER} --silent --fms url=${FMS} --installdir \
${INSTALL_DIR}; sleep 3; ${INSTALL_DIR}/bin/fglam \
--daemon; sleep 3; rm /tmp/${INSTALLER};"

echo "Installation complete on $HOSTNAME."
echo "Connect to $HOSTNAME and use tail -f \
${INSTALL_DIR}/state/default/logs/FglAM*.log to validate \
the installation."
done
else
echo "$HOSTFILE does not exist, exiting..."
exit 1
fi

exit 0

```

3 Save *remote_install_fglam_example.sh*.

Part 5: Run the script

Once you have finished editing the deployment script, it is ready to run to deploy the Agent Manager to the remote UNIX hosts.

i **NOTE:** This part of the example describes using one type of login method (passwords) for connecting to remote machines using SSH. However, if you have private keys configured in your environment, then you might encounter one of the following scenarios when establishing an SSH connection:

- You do not need to provide a password because the host to which you are connecting is a trusted host.
- You are prompted for a passphrase instead of a password.

1 Ensure that *remote_install_fglam_example.sh* is executable.

2 Execute the *remote_install_fglam_example.sh* script.

If Wget is not found on a host or if the operating system listed for a host in the *hostnames* file does not match the operating system types listed in the script, the script continues deploying the Agent Manager to the rest of the hosts that you specified in the *hostnames* file.

3 Provide a password for each host when you are prompted to do so.

The Agent Manager is installed and running on the specified hosts.

Example: Deploying the Agent Manager to multiple Windows hosts

The recommended method of deploying the Agent Manager to a large number of Windows hosts is to use a software deployment tool such as Microsoft® Systems Management Center.

Example: Using a software deployment tool and silent installer parameters

The example in this section outlines a scenario in which you use a software deployment tool and provide it with Agent Manager silent installer parameters to install the Agent Manager on multiple Windows hosts.

! **CAUTION:** Using a deployment tool and Agent Manager silent installer parameters is an advanced activity. You should do so only if both of the following conditions apply:

- You are a Foglight administrator who is familiar with configuring the Agent Manager.
- You are certain of the setup that is required for your environment.

i **NOTE:** If you are not certain which installation options you need, run the interactive GUI or command-line installer on each machine instead.

Part 1: Provide the silent installer parameters

Provide the software deployment tool with the Agent Manager silent installer parameters that are required for your environment. See [Using Agent Manager silent installer Parameters](#) on page 107 for more information.

Part 2: Make the Agent Manager installers accessible

Place the Agent Manager installers on a shared drive that is accessible to all of the Windows hosts to which you are deploying the Agent Manager.

Part 3: Run the software deployment tool

Follow the standard software-deployment steps for the tool you are using.

Next steps

To configure, start, and stop the Agent Manager, refer to [Configuring the Agent Manager](#) on page 40 and [Starting or stopping the Agent Manager process](#) on page 34.

Quest creates software solutions that make the benefits of new technology real in an increasingly complex IT landscape. From database and systems management, to Active Directory and Office 365 management, and cyber security resilience, Quest helps customers solve their next IT challenge now. Around the globe, more than 130,000 companies and 95% of the Fortune 500 count on Quest to deliver proactive management and monitoring for the next enterprise initiative, find the next solution for complex Microsoft challenges and stay ahead of the next threat. Quest Software. Where next meets now. For more information, visit <https://www.quest.com/>.

Technical support resources

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

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

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