



GlassFish Server Open Source Edition 3.1 Deployment Planning Guide

Oracle Corporation
500 Oracle Parkway
Redwood City, CA 94065
U.S.A.

Part No: 821-2448-12
July 2011

Copyright © 2010, 2011, Oracle and/or its affiliates. All rights reserved.

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

Contents

Preface	9
1 Product Concepts	15
Java EE Platform Overview	15
Java EE Applications	15
Containers	15
Java EE Services	16
Web Services	16
Client Access	16
External Systems and Resources	17
GlassFish Server Components	18
Server Instances	18
Administrative Domains	18
Clusters	19
Named Configurations	20
HTTP Load Balancer Plug-in	20
Session Persistence	21
IIOP Load Balancing in a Cluster	22
Message Queue and JMS Resources	23
Configuration Roadmap for High Availability of GlassFish Server	24
▼ To Configure GlassFish Server for High Availability	24
2 Planning your Deployment	27
Establishing Performance Goals	27
Estimating Throughput	28
Estimating Load on GlassFish Server Instances	28
Planning the Network Configuration	31

Setting Up Traffic Separation	31
Estimating Bandwidth Requirements	32
Calculating Bandwidth Required	32
Estimating Peak Load	33
Choosing Network Cards	33
Planning for Availability	34
Rightsizing Availability	34
Using Clusters to Improve Availability	34
Adding Redundancy to the System	35
Design Decisions	36
Designing for Peak or Steady State Load	36
System Sizing	36
Sizing the Administration Thread Pool	37
Planning Message Queue Broker Deployment	37
Multi-Broker Clusters	37
Configuring GlassFish Server to Use Message Queue Brokers	39
Example Deployment Scenarios	41
3 Checklist for Deployment	43
Checklist	43
Index	47

Tables

TABLE 3-1	Checklist	43
-----------	-----------------	----

Examples

EXAMPLE 2-1	Calculation of Response Time	30
EXAMPLE 2-2	Calculation of Requests Per Second	31
EXAMPLE 2-3	Calculation of Bandwidth Required	32
EXAMPLE 2-4	Calculation of Peak Load	33

Preface

The *Deployment Planning Guide* explains how to build a production deployment of GlassFish Server Open Source Edition.

This preface contains information about and conventions for the entire GlassFish Server Open Source Edition (GlassFish Server) documentation set.

GlassFish Server 3.1 is developed through the GlassFish project open-source community at <http://glassfish.java.net/>. The GlassFish project provides a structured process for developing the GlassFish Server platform that makes the new features of the Java EE platform available faster, while maintaining the most important feature of Java EE: compatibility. It enables Java developers to access the GlassFish Server source code and to contribute to the development of the GlassFish Server. The GlassFish project is designed to encourage communication between Oracle engineers and the community.

GlassFish Server Documentation Set

The GlassFish Server documentation set describes deployment planning and system installation. For an introduction to GlassFish Server, refer to the books in the order in which they are listed in the following table.

TABLE P-1 Books in the GlassFish Server Documentation Set

Book Title	Description
<i>Release Notes</i>	Provides late-breaking information about the software and the documentation and includes a comprehensive, table-based summary of the supported hardware, operating system, Java Development Kit (JDK), and database drivers.
<i>Quick Start Guide</i>	Explains how to get started with the GlassFish Server product.
<i>Installation Guide</i>	Explains how to install the software and its components.
<i>Upgrade Guide</i>	Explains how to upgrade to the latest version of GlassFish Server. This guide also describes differences between adjacent product releases and configuration options that can result in incompatibility with the product specifications.

TABLE P-1 Books in the GlassFish Server Documentation Set (Continued)

Book Title	Description
<i>Deployment Planning Guide</i>	Explains how to build a production deployment of GlassFish Server that meets the requirements of your system and enterprise.
<i>Administration Guide</i>	Explains how to configure, monitor, and manage GlassFish Server subsystems and components from the command line by using the <code>asadmin(1M)</code> utility. Instructions for performing these tasks from the Administration Console are provided in the Administration Console online help.
<i>Security Guide</i>	Provides instructions for configuring and administering GlassFish Server security.
<i>Application Deployment Guide</i>	Explains how to assemble and deploy applications to the GlassFish Server and provides information about deployment descriptors.
<i>Application Development Guide</i>	Explains how to create and implement Java Platform, Enterprise Edition (Java EE platform) applications that are intended to run on the GlassFish Server. These applications follow the open Java standards model for Java EE components and application programmer interfaces (APIs). This guide provides information about developer tools, security, and debugging.
<i>Add-On Component Development Guide</i>	Explains how to use published interfaces of GlassFish Server to develop add-on components for GlassFish Server. This document explains how to perform <i>only</i> those tasks that ensure that the add-on component is suitable for GlassFish Server.
<i>Embedded Server Guide</i>	Explains how to run applications in embedded GlassFish Server and to develop applications in which GlassFish Server is embedded.
<i>High Availability Administration Guide</i>	Explains how to configure GlassFish Server to provide higher availability and scalability through failover and load balancing.
<i>Performance Tuning Guide</i>	Explains how to optimize the performance of GlassFish Server.
<i>Troubleshooting Guide</i>	Describes common problems that you might encounter when using GlassFish Server and explains how to solve them.
<i>Error Message Reference</i>	Describes error messages that you might encounter when using GlassFish Server.
<i>Reference Manual</i>	Provides reference information in man page format for GlassFish Server administration commands, utility commands, and related concepts.
<i>Message Queue Release Notes</i>	Describes new features, compatibility issues, and existing bugs for Open Message Queue.
<i>Message Queue Technical Overview</i>	Provides an introduction to the technology, concepts, architecture, capabilities, and features of the Message Queue messaging service.
<i>Message Queue Administration Guide</i>	Explains how to set up and manage a Message Queue messaging system.

TABLE P-1 Books in the GlassFish Server Documentation Set (Continued)

Book Title	Description
<i>Message Queue Developer's Guide for JMX Clients</i>	Describes the application programming interface in Message Queue for programmatically configuring and monitoring Message Queue resources in conformance with the Java Management Extensions (JMX).
<i>Message Queue Developer's Guide for Java Clients</i>	Provides information about concepts and procedures for developing Java messaging applications (Java clients) that work with GlassFish Server.
<i>Message Queue Developer's Guide for C Clients</i>	Provides programming and reference information for developers working with Message Queue who want to use the C language binding to the Message Queue messaging service to send, receive, and process Message Queue messages.

Related Documentation

The following tutorials explain how to develop Java EE applications:

- *Your First Cup: An Introduction to the Java EE Platform* (<http://download.oracle.com/javaee/6/firstcup/doc/>). For beginning Java EE programmers, this short tutorial explains the entire process for developing a simple enterprise application. The sample application is a web application that consists of a component that is based on the Enterprise JavaBeans specification, a JAX-RS web service, and a JavaServer Faces component for the web front end.
- *The Java EE 6 Tutorial* (<http://download.oracle.com/javaee/6/tutorial/doc/>). This comprehensive tutorial explains how to use Java EE 6 platform technologies and APIs to develop Java EE applications.

Javadoc tool reference documentation for packages that are provided with GlassFish Server is available as follows.

- The API specification for version 6 of Java EE is located at <http://download.oracle.com/javaee/6/api/>.
- The API specification for GlassFish Server 3.1, including Java EE 6 platform packages and nonplatform packages that are specific to the GlassFish Server product, is located at <http://glassfish.java.net/nonav/docs/v3/api/>.

Additionally, the *Java EE Specifications* (<http://www.oracle.com/technetwork/java/javaee/tech/index.html>) might be useful.

For information about creating enterprise applications in the NetBeans Integrated Development Environment (IDE), see the *NetBeans Documentation, Training & Support page* (<http://www.netbeans.org/kb/>).

For information about the Java DB database for use with the GlassFish Server, see the *Java DB product page* (<http://www.oracle.com/technetwork/java/javadb/overview/index.html>).

The Java EE Samples project is a collection of sample applications that demonstrate a broad range of Java EE technologies. The Java EE Samples are bundled with the Java EE Software Development Kit (SDK) and are also available from the [Java EE Samples project page](http://java.net/projects/glassfish-samples) (<http://java.net/projects/glassfish-samples>).

Default Paths and File Names

The following table describes the default paths and file names that are used in this book.

TABLE P-2 Default Paths and File Names

Placeholder	Description	Default Value
<i>as-install</i>	Represents the base installation directory for GlassFish Server. In configuration files, <i>as-install</i> is represented as follows: <code>\${com.sun.aas.installRoot}</code>	Installations on the Oracle Solaris operating system, Linux operating system, and Mac OS operating system: <i>user's-home-directory/glassfish3/glassfish</i> Windows, all installations: <i>SystemDrive:\glassfish3\glassfish</i>
<i>as-install-parent</i>	Represents the parent of the base installation directory for GlassFish Server.	Installations on the Oracle Solaris operating system, Linux operating system, and Mac operating system: <i>user's-home-directory/glassfish3</i> Windows, all installations: <i>SystemDrive:\glassfish3</i>
<i>domain-root-dir</i>	Represents the directory in which a domain is created by default.	<i>as-install/domains/</i>
<i>domain-dir</i>	Represents the directory in which a domain's configuration is stored. In configuration files, <i>domain-dir</i> is represented as follows: <code>\${com.sun.aas.instanceRoot}</code>	<i>domain-root-dir/domain-name</i>

Typographic Conventions

The following table describes the typographic changes that are used in this book.

TABLE P-3 Typographic Conventions

Typeface	Meaning	Example
AaBbCc123	The names of commands, files, and directories, and onscreen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with onscreen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	A placeholder to be replaced with a real name or value	The command to remove a file is <i>rm filename</i> .
<i>AaBbCc123</i>	Book titles, new terms, and terms to be emphasized (note that some emphasized items appear bold online)	Read Chapter 6 in the <i>User's Guide</i> . A <i>cache</i> is a copy that is stored locally. Do <i>not</i> save the file.

Symbol Conventions

The following table explains symbols that might be used in this book.

TABLE P-4 Symbol Conventions

Symbol	Description	Example	Meaning
[]	Contains optional arguments and command options.	<code>ls [-l]</code>	The <code>-l</code> option is not required.
{ }	Contains a set of choices for a required command option.	<code>-d {y n}</code>	The <code>-d</code> option requires that you use either the <code>y</code> argument or the <code>n</code> argument.
<code>\${ }</code>	Indicates a variable reference.	<code>\${com.sun.javaRoot}</code>	References the value of the <code>com.sun.javaRoot</code> variable.
-	Joins simultaneous multiple keystrokes.	Control-A	Press the Control key while you press the A key.
+	Joins consecutive multiple keystrokes.	Ctrl+A+N	Press the Control key, release it, and then press the subsequent keys.
→	Indicates menu item selection in a graphical user interface.	File → New → Templates	From the File menu, choose New. From the New submenu, choose Templates.

Documentation, Support, and Training

The Oracle web site provides information about the following additional resources:

- Documentation (<http://www.oracle.com/technetwork/indexes/documentation/index.html>)
- Support (<http://www.oracle.com/us/support/index.html>)
- Training (<http://education.oracle.com/>)

Product Concepts

GlassFish Server Open Source Edition provides a robust platform for the development, deployment, and management of Java EE applications. Key features include scalable transaction management, web services performance, clustering, security, and integration capabilities.

The following topics are addressed here:

- [“Java EE Platform Overview” on page 15](#)
- [“GlassFish Server Components” on page 18](#)
- [“Configuration Roadmap for High Availability of GlassFish Server” on page 24](#)

Java EE Platform Overview

GlassFish Server implements Java platform, Enterprise Edition (Java EE) 6 technology. The Java EE platform is a set of standard specifications that describe application components, APIs, and the runtime containers and services of an application server.

Java EE Applications

Java EE applications are made up of components such as JavaServer Pages (JSP), Java servlets, and Enterprise JavaBeans (EJB) modules. These components enable software developers to build large-scale, distributed applications. Developers package Java EE applications in Java Archive (JAR) files (similar to zip files), which can be distributed to production sites. Administrators install Java EE applications onto GlassFish Server by deploying Java EE JAR files onto one or more server instances (or clusters of instances).

Containers

Each server instance includes two containers: web and EJB. A container is a runtime environment that provides services such as security and transaction management to Java EE

components. Web components, such as Java Server Pages and servlets, run within the web container. Enterprise JavaBeans run within the EJB container.

Java EE Services

The Java EE platform provides services for applications, including:

- *Naming* - A naming and directory service binds objects to names. A Java EE application can locate an object by looking up its Java Naming and Directory Interface (JNDI) name.
- *Security* - The Java Authorization Contract for Containers (JACC) is a set of security contracts defined for the Java EE containers. Based on the client's identity, containers can restrict access to the container's resources and services.
- *Transaction management* - A transaction is an indivisible unit of work. For example, transferring funds between bank accounts is a transaction. A transaction management service ensures that a transaction is either completed, or is rolled back.
- *Message Service* - Applications hosted on separate systems can communicate with each other by exchanging messages using the Java™ Message Service (JMS). JMS is an integral part of the Java EE platform and simplifies the task of integrating heterogeneous enterprise applications.

Web Services

Clients can access a Java EE application as a remote web service in addition to accessing it through HTTP, RMI/IIOP, and JMS. Web services are implemented using the Java API for XML-based web services (JAX-WS). A Java EE application can also act as a client to web services, which would be typical in network applications.

Web Services Description Language (WSDL) is an XML format that describes web service interfaces. Web service consumers can dynamically parse a WSDL document to determine the operations a web service provides and how to execute them. GlassFish Server distributes web services interface descriptions using a registry that other applications can access through the Java API for XML Registries (JAXR).

Client Access

Clients can access Java EE applications in several ways. Browser clients access web applications using hypertext transfer protocol (HTTP). For secure communication, browsers use the HTTP secure (HTTPS) protocol that uses secure sockets layer (SSL).

Rich client applications running in the Application Client Container can directly lookup and access Enterprise JavaBeans using an Object Request Broker (ORB), Remote Method

Invocation (RMI) and the internet inter-ORB protocol (IIOP), or IIOP/SSL (secure IIOP). They can access applications and web services using HTTP/HTTPS, JMS, and JAX-WS. They can use JMS to send messages to and receive messages from applications and message-driven beans.

Clients that conform to the Web Services-Interoperability (WS-I) Basic Profile can access Java EE web services. WS-I is an integral part of the Java EE standard and defines interoperable web services. It enables clients written in any supporting language to access web services deployed to GlassFish Server.

The best access mechanism depends on the specific application and the anticipated volume of traffic. GlassFish Server supports separately configurable listeners for HTTP, HTTPS, JMS, IIOP, and IIOP/SSL. You can set up multiple listeners for each protocol for increased scalability and reliability.

Java EE applications can also act as clients of Java EE components such as Enterprise JavaBeans modules deployed on other servers, and can use any of these access mechanisms.

External Systems and Resources

On the Java EE platform, an external system is called a *resource*. For example, a database management system is a JDBC resource. Each resource is uniquely identified and by its Java Naming and Directory Interface (JNDI) name. Applications access external systems through the following APIs and components:

- *Java Database Connectivity (JDBC)* - A database management system (DBMS) provides facilities for storing, organizing, and retrieving data. Most business applications store data in relational databases, which applications access via JDBC. GlassFish Server includes the Java DB database for use sample applications and application development and prototyping, though it is not suitable for deployment. GlassFish Server provides certified JDBC drivers for connecting to major relational databases. These drivers are suitable for deployment.
- *Java Message Service* - Messaging is a method of communication between software components or applications. A messaging client sends messages to, and receives messages from, any other client via a messaging provider that implements the Java Messaging Service (JMS) API. GlassFish Server includes a high-performance JMS broker, Open Message Queue.
- *Java EE Connectors* - The Java EE Connector architecture enables integrating Java EE applications and existing Enterprise Information Systems (EIS). An application accesses an EIS through a portable Java EE component called a *connector* or *resource adapter*, analogous to using JDBC driver to access an RDBMS. Resource adapters are distributed as standalone Resource Adapter Archive (RAR) modules or included in Java EE application archives. As RARs, they are deployed like other Java EE components. GlassFish Server includes evaluation resource adapters that integrate with popular EIS.
- *JavaMail* - Through the JavaMail API, applications can connect to a Simple Mail Transport Protocol (SMTP) server to send email and to an IMAP or POP3 server to receive email.

GlassFish Server Components

This section describes the components in GlassFish Server.

The following topics are addressed here:

- [“Server Instances” on page 18](#)
- [“Administrative Domains” on page 18](#)
- [“Clusters” on page 19](#)
- [“Named Configurations” on page 20](#)
- [“HTTP Load Balancer Plug-in” on page 20](#)
- [“IIOP Load Balancing in a Cluster” on page 22](#)
- [“Message Queue and JMS Resources” on page 23](#)

The administration tools, such as the browser-based Administration Console, communicate with the domain administration server (DAS), which in turn communicates with the server instances.

Server Instances

A server instance is a GlassFish Server running in a single Java Virtual Machine (JVM) process. GlassFish Server is certified with Java platform, Standard Edition (Java SE) 6.

It is usually sufficient to create a single server instance on a machine, since GlassFish Server and accompanying JVM are both designed to scale to multiple processors. However, it can be beneficial to create multiple instances on one machine for application isolation and rolling upgrades. In some cases, a large server with multiple instances can be used in more than one administrative domain. The administration tools makes it easy to create, delete, and manage server instances across multiple machines.

Administrative Domains

An *administrative domain* (or simply *domain*) is a group of server instances that are administered together. A server instance belongs to a single administrative domain. The instances in a domain can run on different physical hosts.

You can create multiple domains from one installation of GlassFish Server. By grouping server instances into domains, different organizations and administrators can share a single GlassFish Server installation. Each domain has its own configuration, log files, and application deployment areas that are independent of other domains. Changing the configuration of one domain does not affect the configurations of other domains. Likewise, deploying an application on one domain does not deploy it or make it visible to any other domain.

Note – All hosts in a domain on which the DAS and GlassFish Server instances are running must have the same operating system.

Domain Administration Server (DAS)

A domain has one Domain Administration Server (DAS), a specially designated GlassFish Server instance that hosts the administrative applications. The DAS authenticates the administrator, accepts requests from administration tools, and communicates with server instances in the domain to carry out the requests.

The administration tools are the `asadmin` command-line tool and the browser-based Administration Console. GlassFish Server also provides a RESTful API for server administration. The administrator can view and manage a single domain at a time, thus enforcing secure separation.

The DAS is also sometimes referred to as the *admin server* or *default server*. It is referred to as the default server because it is the default target for some administrative operations.

Since the DAS is a GlassFish Server instance, it can also host Java EE applications for testing purposes. However, do not use it to host production applications. You might want to deploy applications to the DAS, for example, if the clusters and instances that will host the production application have not yet been created.

The DAS keeps a repository containing the configuration of its domain and all the deployed applications. If the DAS is inactive or down, there is no impact on the performance or availability of active server instances, however administrative changes cannot be made. In certain cases, for security purposes, it may be useful to intentionally stop the DAS process, for example to reboot the host operating system to install a kernel patch or a hardware upgrade.

Administrative commands are provided to backup and restore the domain configuration and applications. With the standard backup and restore procedures, you can quickly restore working configurations. If the DAS host fails, you must create a new DAS installation to restore the previous domain configuration. For instructions, see Chapter 3, “Administering Domains,” in *GlassFish Server Open Source Edition 3.1 Administration Guide*.

Clusters

A cluster is a named collection of server instances that share the same applications, resources, and configuration information. You can group server instances on different machines into one logical cluster and administer them as one unit. You can easily control the lifecycle of a multi-machine cluster with the DAS.

Clusters enable horizontal scalability, load balancing, and failover protection. By definition, all the instances in a cluster have the same resource and application configuration. When a server

instance or a machine in a cluster fails, the load balancer detects the failure, redirects traffic from the failed instance to other instances in the cluster, and recovers the user session state. Since the same applications and resources are on all instances in the cluster, an instance can failover to any other instance in the cluster.

Note – All hosts in a cluster on which the DAS and GlassFish Server instances are running must have the same operating system.

Clusters, domains, and instances are related as follows:

- An administrative domain can have zero or more clusters.
- A cluster has one or more server instances.
- A cluster belongs to a single domain.

Named Configurations

A *named configuration* is an abstraction that encapsulates GlassFish Server property settings. Clusters and stand-alone server instances reference a named configuration to get their property settings. With named configurations, Java EE containers' configurations are independent of the physical machine on which they reside, except for particulars such as IP address, port number, and amount of heap memory. Using named configurations provides power and flexibility to GlassFish Server administration.

To apply configuration changes, you simply change the property settings of the named configuration, and all the clusters and stand-alone instances that reference it pick up the changes. You can only delete a named configuration when all references to it have been removed. A domain can contain multiple named configurations.

GlassFish Server comes with a default configuration, called `default-config`. The default configuration is optimized for developer productivity.

You can create your own named configuration based on the default configuration that you can customize for your own purposes. Use the Administration Console and `asadmin` command line utility to create and manage named configurations.

HTTP Load Balancer Plug-in

The load balancer distributes the workload among multiple physical machines, thereby increasing the overall throughput of the system. The GlassFish Server includes the load balancer plug-ins for Oracle iPlanet Web Server, Oracle HTTP Server, Apache Web Server, and Microsoft Internet Information Server.

The load balancer plug-in accepts HTTP and HTTPS requests and forwards them to one of the GlassFish Server instances in the cluster. Should an instance fail, become unavailable (due to network faults), or become unresponsive, requests are redirected to existing, available machines. The load balancer can also recognize when a failed instance has recovered and redistribute the load accordingly.

For simple stateless applications, a load-balanced cluster may be sufficient. However, for mission-critical applications with session state, use load balanced clusters with replicated session persistence.

To setup a system with load balancing, in addition to GlassFish Server, you must install a web server and the load-balancer plug-in. Then you must:

- Create GlassFish Server clusters that you want to participate in load balancing.
- Deploy applications to these load-balanced clusters.

Server instances and clusters participating in load balancing have a homogenous environment. Usually that means that the server instances reference the same server configuration, can access the same physical resources, and have the same applications deployed to them. Homogeneity enables configuration consistency, and improves the ability to support a production deployment.

Use the `asadmin` command-line tool to create a load balancer configuration, add references to clusters and server instances to it, enable the clusters for reference by the load balancer, enable applications for load balancing, optionally create a health checker, generate the load balancer configuration file, and finally copy the load balancer configuration file to your web server `config` directory. An administrator can create a script to automate this entire process.

For more details and complete configuration instructions, see Chapter 7, “Configuring HTTP Load Balancing,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

Session Persistence

Java EE applications typically have significant amounts of session state data. A web shopping cart is the classic example of a session state. Also, an application can cache frequently-needed data in the session object. In fact, almost all applications with significant user interactions need to maintain a session state. Both HTTP sessions and stateful session beans (SFSBs) have session state data.

While the session state is not as important as the transactional state stored in a database, preserving the session state across server failures can be important to end users. GlassFish Server provides the capability to save, or persist, this session state in a repository. If the GlassFish Server instance that is hosting the user session experiences a failure, the session state can be recovered. The session can continue without loss of information.

GlassFish Server supports the following session persistence types:

- Memory
- Replicated
- File
- Coherence*Web

With memory persistence, the state is always kept in memory and does not survive failure. With replicated persistence, GlassFish Server uses other server instances in the cluster as the persistence store for both HTTP and SFSB sessions. With file persistence, GlassFish Server serializes session objects and stores them to the file system location specified by session manager properties. For SFSBs, if replicated persistence is not specified, GlassFish Server stores state information in the session-store subdirectory of this location. For more information about Coherence*Web, see [Using Coherence*Web with GlassFish Server \(http://download.oracle.com/docs/cd/E18686_01/coh.37/e18690/glassfish.htm\)](http://download.oracle.com/docs/cd/E18686_01/coh.37/e18690/glassfish.htm).

Checking an SFSB's state for changes that need to be saved is called *checkpointing*. When enabled, checkpointing generally occurs after any transaction involving the SFSB is completed, even if the transaction rolls back. For more information on developing stateful session beans, see “Using Session Beans” in *GlassFish Server Open Source Edition 3.1 Application Development Guide*. For more information on enabling SFSB failover, see “Stateful Session Bean Failover” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

Apart from the number of requests being served by GlassFish Server, the session persistence configuration settings also affect the session information in each request.

For more information on configuring session persistence, see Chapter 9, “Configuring High Availability Session Persistence and Failover,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

IIOP Load Balancing in a Cluster

With IIOP load balancing, IIOP client requests are distributed to different server instances or name servers. The goal is to spread the load evenly across the cluster, thus providing scalability. IIOP load balancing combined with EJB clustering and availability features in GlassFish Server provides not only load balancing but also EJB failover.

There are two steps to IIOP failover and load balancing. The first step, bootstrapping, is the process by which the client sets up the initial naming context with one ORB in the cluster. The client attempts to connect to one of the IIOP endpoints. When launching an application client using the `appClient` script, you specify these endpoints using the `-targetserver` option on the command line or `target-server` elements in the `sun-acc.xml` configuration file. The client randomly chooses one of these endpoints and tries to connect to it, trying other endpoints if needed until one works.

The second step concerns sending messages to a specific EJB. By default, all naming look-ups, and therefore all EJB accesses, use the cluster instance chosen during bootstrapping. The client exchanges messages with an EJB through the client ORB and server ORB. As this happens, the server ORB updates the client ORB as servers enter and leave the cluster. Later, if the client loses its connection to the server from step 1, the client fails over to some other server using its list of currently active members. In particular, this cluster member might have joined the cluster after the client made the initial connection.

When a client performs a JNDI lookup for an object, the Naming Service creates an `InitialContext` (IC) object associated with a particular server instance. From then on, all lookup requests made using that IC object are sent to the same server instance. All `EJBHome` objects looked up with that `InitialContext` are hosted on the same target server. Any bean references obtained henceforth are also created on the same target host. This effectively provides load balancing, since all clients randomize the list of live target servers when creating `InitialContext` objects. If the target server instance goes down, the lookup or EJB method invocation will failover to another server instance.

Adding or deleting new instances to the cluster does not update the existing client's view of the cluster. You must manually update the endpoints list on the client side.

Message Queue and JMS Resources

The Open Message Queue (Message Queue) provides reliable, asynchronous messaging for distributed applications. Message Queue is an enterprise messaging system that implements the Java Message Service (JMS) standard. Message Queue provides messaging for Java EE application components such as message-driven beans (MDBs).

GlassFish Server implements the Java Message Service (JMS) API by integrating Message Queue into GlassFish Server. GlassFish Server includes the Enterprise version of Message Queue which has failover, clustering and load balancing features.

For basic JMS administration tasks, use the GlassFish Server Administration Console and `asadmin` command-line utility.

For advanced tasks, including administering a Message Queue cluster, use the tools provided in the `as-install/mq/bin` directory. For details about administering Message Queue, see the *Open Message Queue 4.5 Administration Guide*.

For information on deploying JMS applications and Message Queue clustering for message failover, see [“Planning Message Queue Broker Deployment” on page 37](#).

Configuration Roadmap for High Availability of GlassFish Server

The following procedure lists the major tasks for configuring GlassFish Server for high availability. The procedure also provides cross-references to detailed instructions for performing each task.

▼ To Configure GlassFish Server for High Availability

1 Determine your requirements and goals for performance and QoS.

For more information, see the following documentation:

- [“Establishing Performance Goals” on page 27](#)
- [“Planning the Network Configuration” on page 31](#)
- [“Planning for Availability” on page 34](#)

2 Size your system.

For more information, see [“Design Decisions” on page 36](#).

3 Install GlassFish Server and related subcomponents such as a web server.

For more information, see the following documentation:

- *GlassFish Server Open Source Edition 3.1 Installation Guide*
- Installation guides for related subcomponents, for example, [Oracle iPlanet Web Server 7.0.9 Installation and Migration Guide](#)

4 If you plan to administer your clusters centrally, set up secure shell (SSH) for centralized administration.

For more information, see Chapter 2, “Setting Up SSH for Centralized Administration,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

5 Configure domains, nodes, clusters, GlassFish Server instances, and virtual servers as required.

For more information, see the following documentation:

- Chapter 3, “Administering Domains,” in *GlassFish Server Open Source Edition 3.1 Administration Guide*
- Chapter 3, “Administering GlassFish Server Nodes,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*
- Chapter 4, “Administering GlassFish Server Clusters,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*

- Chapter 5, “Administering GlassFish Server Instances,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*
- “Administering Virtual Servers” in *GlassFish Server Open Source Edition 3.1 Administration Guide*

6 Configure your load balancer.

For more information, see “Administering mod_jk” in *GlassFish Server Open Source Edition 3.1 Administration Guide*.

7 Configure the web container and EJB container for replicated session persistence.

For more information, see Chapter 9, “Configuring High Availability Session Persistence and Failover,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

8 If you are using messaging extensively, configure Java Message Service (JMS) clusters for failover .

For more information, see the following documentation:

- [“Planning Message Queue Broker Deployment” on page 37](#)
- Chapter 10, “Configuring Java Message Service High Availability,” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*
- *Open Message Queue 4.5 Administration Guide*

9 Deploy applications and configure them for high availability and session failover.

For more information, see *GlassFish Server Open Source Edition 3.1 Application Deployment Guide*.

Planning your Deployment

Before deploying GlassFish Server, first determine the performance and availability goals, and then make decisions about the hardware, network, and storage requirements accordingly.

The following topics are addressed here:

- [“Establishing Performance Goals” on page 27](#)
- [“Planning the Network Configuration” on page 31](#)
- [“Planning for Availability” on page 34](#)
- [“Design Decisions” on page 36](#)
- [“Planning Message Queue Broker Deployment” on page 37](#)

Establishing Performance Goals

At its simplest, high performance means maximizing throughput and reducing response time. Beyond these basic goals, you can establish specific goals by determining the following:

- What types of applications and services are deployed, and how do clients access them?
- Which applications and services need to be highly available?
- Do the applications have session state or are they stateless?
- What request capacity or throughput must the system support?
- How many concurrent users must the system support?
- What is an acceptable average response time for user requests?
- What is the average think time between requests?

You can calculate some of these metrics using a remote browser emulator (RBE) tool, or web site performance and benchmarking software that simulates expected application activity. Typically, RBE and benchmarking products generate concurrent HTTP requests and then report the response time for a given number of requests per minute. You can then use these figures to calculate server activity.

The results of the calculations described in this chapter are not absolute. Treat them as reference points to work against, as you fine-tune the performance of GlassFish Server and your applications.

The following topics are addressed here:

- [“Estimating Throughput” on page 28](#)
- [“Estimating Load on GlassFish Server Instances” on page 28](#)
- [“Estimating Bandwidth Requirements” on page 32](#)
- [“Estimating Peak Load” on page 33](#)

Estimating Throughput

In broad terms, *throughput* measures the amount of work performed by GlassFish Server. For GlassFish Server, throughput can be defined as the number of requests processed per minute per server instance.

As described in the next section, GlassFish Server throughput is a function of many factors, including the nature and size of user requests, number of users, and performance of GlassFish Server instances and back-end databases. You can estimate throughput on a single machine by benchmarking with simulated workloads.

High availability applications incur additional overhead because they periodically save session data. The amount of overhead depends on the amount of data, how frequently it changes, and how often it is saved. The first two factors depend on the application in question; the latter is also affected by server settings.

Estimating Load on GlassFish Server Instances

Consider the following factors to estimate the load on GlassFish Server instances.

The following topics are addressed here:

- [“Maximum Number of Concurrent Users” on page 28](#)
- [“Think Time” on page 29](#)
- [“Average Response Time” on page 29](#)
- [“Requests Per Minute” on page 30](#)

Maximum Number of Concurrent Users

Users interact with an application through a client, such as a web browser or Java program. Based on the user’s actions, the client periodically sends requests to the GlassFish Server. A user is considered *active* as long as the user’s session has neither expired nor been terminated. When estimating the number of concurrent users, include all active users.

Initially, as the number of users increases, throughput increases correspondingly. However, as the number of concurrent requests increases, server performance begins to saturate, and throughput begins to decline.

Identify the point at which adding concurrent users reduces the number of requests that can be processed per minute. This point indicates when optimal performance is reached and beyond which throughput starts to degrade. Generally, strive to operate the system at optimal throughput as much as possible. You might need to add processing power to handle additional load and increase throughput.

Think Time

A user does not submit requests continuously. A user submits a request, the server receives and processes the request, and then returns a result, at which point the user spends some time before submitting a new request. The time between one request and the next is called *think time*.

Think times are dependent on the type of users. For example, machine-to-machine interaction such as for a web service typically has a lower think time than that of a human user. You may have to consider a mix of machine and human interactions to estimate think time.

Determining the average think time is important. You can use this duration to calculate the number of requests that need to be completed per minute, as well as the number of concurrent users the system can support.

Average Response Time

Response time refers to the amount of time GlassFish Server takes to return the results of a request to the user. The response time is affected by factors such as network bandwidth, number of users, number and type of requests submitted, and average think time.

In this section, response time refers to the mean, or average, response time. Each type of request has its own minimal response time. However, when evaluating system performance, base the analysis on the average response time of all requests.

The faster the response time, the more requests per minute are being processed. However, as the number of users on the system increases, the response time starts to increase as well, even though the number of requests per minute declines.

A system performance graph indicates that after a certain point, requests per minute are inversely proportional to response time. The sharper the decline in requests per minute, the steeper the increase in response time.

The point of the peak load is the point at which requests per minute start to decline. Prior to this point, response time calculations are not necessarily accurate because they do not use peak numbers in the formula. After this point, (because of the inversely proportional relationship between requests per minute and response time), the administrator can more accurately calculate response time using maximum number of users and requests per minute.

Use the following formula to determine T_{response} , the response time (in seconds) at peak load:

$$T_{\text{response}} = n/r - T_{\text{think}}$$

where

- n is the number of concurrent users
- r is the number requests per second the server receives
- T_{think} is the average think time (in seconds)

To obtain an accurate response time result, always include think time in the equation.

EXAMPLE 2-1 Calculation of Response Time

If the following conditions exist:

- Maximum number of concurrent users, n , that the system can support at peak load is 5,000.
- Maximum number of requests, r , the system can process at peak load is 1,000 per second.

Average think time, T_{think} , is three seconds per request.

Thus, the calculation of response time is:

$$T_{\text{response}} = n/r - T_{\text{think}} = (5000/1000) - 3 \text{ sec.} = 5 - 3 \text{ sec.}$$

Therefore, the response time is two seconds.

After the system's response time has been calculated, particularly at peak load, compare it to the acceptable response time for the application. Response time, along with throughput, is one of the main factors critical to GlassFish Server performance.

Requests Per Minute

If you know the number of concurrent users at any given time, the response time of their requests, and the average user think time, then you can calculate the number of requests per minute. Typically, start by estimating the number of concurrent users that are on the system.

For example, after running web site performance software, the administrator concludes that the average number of concurrent users submitting requests on an online banking web site is 3,000. This number depends on the number of users who have signed up to be members of the online bank, their banking transaction behavior, the time of the day or week they choose to submit requests, and so on.

Therefore, knowing this information enables you to use the requests per minute formula described in this section to calculate how many requests per minute your system can handle for this user base. Since requests per minute and response time become inversely proportional at peak load, decide if fewer requests per minute is acceptable as a trade-off for better response time, or alternatively, if a slower response time is acceptable as a trade-off for more requests per minute.

Experiment with the requests per minute and response time thresholds that are acceptable as a starting point for fine-tuning system performance. Thereafter, decide which areas of the system require adjustment.

Solving for r in the equation in the previous section gives:

$$r = n / (T_{\text{response}} + T_{\text{think}})$$

EXAMPLE 2-2 Calculation of Requests Per Second

For the values:

- $n = 2,800$ concurrent users
- $T_{\text{response}} = 1$ (one second per request average response time)
- $T_{\text{think}} = 3$, (three seconds average think time)

The calculation for the number of requests per second is:

$$r = 2800 / (1+3) = 700$$

Therefore, the number of requests per second is 700 and the number of requests per minute is 42000.

Planning the Network Configuration

When planning how to integrate the GlassFish Server into the network, estimate the bandwidth requirements and plan the network in such a way that it can meet users' performance requirements.

The following topics are addressed here:

- [“Setting Up Traffic Separation” on page 31](#)
- [“Estimating Bandwidth Requirements” on page 32](#)
- [“Calculating Bandwidth Required” on page 32](#)
- [“Estimating Peak Load” on page 33](#)
- [“Choosing Network Cards” on page 33](#)
- [“Identifying Failure Classes” on page 35](#)

Setting Up Traffic Separation

You can separate external traffic, such as client requests, from the internal traffic, such as session state failover, database transactions, and messaging. Traffic separation enables you to plan a network better and augment certain parts of the network, as required.

To separate the traffic, run each server instance on a multi-homed machine. A multi-homed machine has two IP addresses belonging to different networks, an external IP and an internal IP.

The objective is to expose only the external IP to user requests. The internal IP is used only by the cluster instances for internal communication. For details, see “Using the Multi-Homing Feature With GMS” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

To plan for traffic on both networks, see [“Estimating Bandwidth Requirements” on page 32](#). For external networks, follow the guidelines in [“Calculating Bandwidth Required” on page 32](#) and [“Estimating Peak Load” on page 33](#). To size the interfaces for internal networks, see [“Choosing Network Cards” on page 33](#).

Estimating Bandwidth Requirements

To decide on the desired size and bandwidth of the network, first determine the network traffic and identify its peak. Check if there is a particular hour, day of the week, or day of the month when overall volume peaks, and then determine the duration of that peak.

During peak load times, the number of packets in the network is at its highest level. In general, if you design for peak load, scale your system with the goal of handling 100 percent of peak volume. Bear in mind, however, that any network behaves unpredictably and that despite your scaling efforts, it might not always be able to handle 100 percent of peak volume.

For example, assume that at peak load, five percent of users occasionally do not have immediate network access when accessing applications deployed on GlassFish Server. Of that five percent, estimate how many users retry access after the first attempt. Again, not all of those users might get through, and of that unsuccessful portion, another percentage will retry. As a result, the peak appears longer because peak use is spread out over time as users continue to attempt access.

Calculating Bandwidth Required

Based on the calculations made in [“Establishing Performance Goals” on page 27](#), determine the additional bandwidth required for deploying GlassFish Server at your site.

Depending on the method of access (T-1 lines, ADSL, cable modem, and so on), calculate the amount of increased bandwidth required to handle your estimated load. For example, suppose your site uses T-1 or higher-speed T-3 lines. Given their bandwidth, estimate how many lines are needed on the network, based on the average number of requests generated per second at your site and the maximum peak load. Calculate these figures using a web site analysis and monitoring tool.

EXAMPLE 2-3 Calculation of Bandwidth Required

A single T-1 line can handle 1.544 Mbps. Therefore, a network of four T-1 lines can handle approximately 6 Mbps of data. Assuming that the average HTML page sent back to a client is 30 kilobytes (KB), this network of four T-1 lines can handle the following traffic per second:

EXAMPLE 2-3 Calculation of Bandwidth Required (Continued)

$6,176,000 \text{ bits} / 10 \text{ bits} = 772,000 \text{ bytes per second}$

$772,000 \text{ bytes per second} / 30 \text{ KB} = \text{approximately } 25 \text{ concurrent response pages per second.}$

With traffic of 25 pages per second, this system can handle 90,000 pages per hour (25 x 60 seconds x 60 minutes), and therefore 2,160,000 pages per day maximum, assuming an even load throughout the day. If the maximum peak load is greater than this, increase the bandwidth accordingly.

Estimating Peak Load

Having an even load throughout the day is probably not realistic. You need to determine when the peak load occurs, how long it lasts, and what percentage of the total load is the peak load.

EXAMPLE 2-4 Calculation of Peak Load

If the peak load lasts for two hours and takes up 30 percent of the total load of 2,160,000 pages, this implies that 648,000 pages must be carried over the T-1 lines during two hours of the day.

Therefore, to accommodate peak load during those two hours, increase the number of T-1 lines according to the following calculations:

$648,000 \text{ pages} / 120 \text{ minutes} = 5,400 \text{ pages per minute}$

$5,400 \text{ pages per minute} / 60 \text{ seconds} = 90 \text{ pages per second}$

If four lines can handle 25 pages per second, then approximately four times that many pages requires four times that many lines, in this case 16 lines. The 16 lines are meant for handling the realistic maximum of a 30 percent peak load. Obviously, the other 70 percent of the load can be handled throughout the rest of the day by these many lines.

Choosing Network Cards

For greater bandwidth and optimal network performance, use at least 100 Mbps Ethernet cards or, preferably, 1 Gbps Ethernet cards between servers hosting GlassFish Server.

Planning for Availability

The following topics are addressed here:

- “Rightsizing Availability” on page 34
- “Using Clusters to Improve Availability” on page 34
- “Adding Redundancy to the System” on page 35

Rightsizing Availability

To plan availability of systems and applications, assess the availability needs of the user groups that access different applications. For example, external fee-paying users and business partners often have higher quality of service (QoS) expectations than internal users. Thus, it may be more acceptable to internal users for an application feature, application, or server to be unavailable than it would be for paying external customers.

There is an increasing cost and complexity to mitigating against decreasingly probable events. At one end of the continuum, a simple load-balanced cluster can tolerate localized application, middleware, and hardware failures. At the other end of the scale, geographically distinct clusters can mitigate against major catastrophes affecting the entire data center.

To realize a good return on investment, it often makes sense to identify availability requirements of features within an application. For example, it may not be acceptable for an insurance quotation system to be unavailable (potentially turning away new business), but brief unavailability of the account management function (where existing customers can view their current coverage) is unlikely to turn away existing customers.

Using Clusters to Improve Availability

At the most basic level, a cluster is a group of GlassFish Server instances—often hosted on multiple physical servers—that appear to clients as a single instance. This provides horizontal scalability as well as higher availability than a single instance on a single machine. This basic level of clustering works in conjunction with the HTTP load balancer plug-in, which accepts HTTP and HTTPS requests and forwards them to one of the instances in the cluster. The ORB and integrated JMS brokers also perform load balancing to GlassFish Server clusters. If an instance fails, becomes unavailable (due to network faults), or becomes unresponsive, requests are redirected only to existing, available machines. The load balancer can also recognize when a failed instance has recovered and redistribute load accordingly.

Adding Redundancy to the System

One way to achieve high availability is to add hardware and software redundancy to the system. When one unit fails, the redundant unit takes over. This is also referred to as fault tolerance. In general, to maximize high availability, determine and remove every possible point of failure in the system.

Identifying Failure Classes

The level of redundancy is determined by the failure classes (types of failure) that the system needs to tolerate. Some examples of failure classes are:

- System process
- Machine
- Power supply
- Disk
- Network failures
- Building fires or other preventable disasters
- Unpredictable natural catastrophes

Duplicated system processes tolerate single system process failures, as well as single machine failures. Attaching the duplicated mirrored (paired) machines to different power supplies tolerates single power failures. By keeping the mirrored machines in separate buildings, a single-building fire can be tolerated. By keeping them in separate geographical locations, natural catastrophes like earthquakes can be tolerated.

Planning Failover Capacity

Failover capacity planning implies deciding how many additional servers and processes you need to add to the GlassFish Server deployment so that in the event of a server or process failure, the system can seamlessly recover data and continue processing. If your system gets overloaded, a process or server failure might result, causing response time degradation or even total loss of service. Preparing for such an occurrence is critical to successful deployment.

To maintain capacity, especially at peak loads, add spare machines running GlassFish Server instances to the existing deployment.

For example, consider a system with two machines running one GlassFish Server instance each. Together, these machines handle a peak load of 300 requests per second. If one of these machines becomes unavailable, the system will be able to handle only 150 requests, assuming an even load distribution between the machines. Therefore, half the requests during peak load will not be served.

Design Decisions

Design decisions include whether you are designing the system for peak or steady-state load, the number of machines in various roles and their sizes, and the size of the administration thread pool.

The following topics are addressed here:

- [“Designing for Peak or Steady State Load” on page 36](#)
- [“System Sizing” on page 36](#)
- [“Sizing the Administration Thread Pool” on page 37](#)

Designing for Peak or Steady State Load

In a typical deployment, there is a difference between steady state and peak workloads:

- If the system is designed to handle peak load, it can sustain the expected maximum load of users and requests without degrading response time. This implies that the system can handle extreme cases of expected system load. If the difference between peak load and steady state load is substantial, designing for peak loads can mean spending money on resources that are often idle.
- If the system is designed to handle steady state load, it does not have all the resources required to handle the expected peak load. Thus, the system has a slower response time when peak load occurs.

How often the system is expected to handle peak load will determine whether you want to design for peak load or for steady state.

If peak load occurs often—say, several times per day—it may be worthwhile to expand capacity to handle it. If the system operates at steady state 90 percent of the time, and at peak only 10 percent of the time, then it may be preferable to deploy a system designed around steady state load. This implies that the system’s response time will be slower only 10 percent of the time. Decide if the frequency or duration of time that the system operates at peak justifies the need to add resources to the system.

System Sizing

Based on the load on the GlassFish Server instances and failover requirements, you can determine the number of applications server instances (hosts) needed. Evaluate your environment on the basis of the factors explained in [“Estimating Load on GlassFish Server Instances” on page 28](#) to each GlassFish Server instance, although each instance can use more than one Central Processing Unit (CPU).

Sizing the Administration Thread Pool

The default `admin-thread-pool` size of 50 should be adequate for most cluster deployments. If you have unusually large clusters, you may need to increase this thread pool size. In this case, set the `max-thread-pool-size` attribute to the number of instances in your largest cluster, but not larger than the number of incoming synchronization requests that the DAS can handle.

Planning Message Queue Broker Deployment

The Java Message Service (JMS) API is a messaging standard that allows Java EE applications and components to create, send, receive, and read messages. It enables distributed communication that is loosely coupled, reliable, and asynchronous. Message Queue, which implements JMS, is integrated with GlassFish Server, enabling you to create components that send and receive JMS messages, including message-driven beans (MDBs).

Message Queue is integrated with GlassFish Server using a *resource adapter* also known as a connector module. A resource adapter is a Java EE component defined according to the Java EE Connector Architecture (JCA) Specification. This specification defines a standardized way in which application servers such as GlassFish Server can integrate with enterprise information systems such as JMS providers. GlassFish Server includes a resource adapter that integrates with its own JMS provider, Message Queue. To use a different JMS provider, you must obtain and deploy a suitable resource adapter that is designed to integrate with it.

Creating a JMS resource in GlassFish Server using the Administration Console creates a preconfigured connector resource that uses the Message Queue resource adapter. To create JMS Resources that use any other resource adapter (including `GenericJMSRA`), you must create them under the Connectors node in the Administration Console.

In addition to using resource adapter APIs, GlassFish Server uses additional Message Queue APIs to provide better integration with Message Queue. This tight integration enables features such as connector failover, load balancing of outbound connections, and load balancing of inbound messages to MDBs. These features enable you to make messaging traffic fault-tolerant and highly available.

The following topics are addressed here:

- [“Multi-Broker Clusters” on page 37](#)
- [“Configuring GlassFish Server to Use Message Queue Brokers” on page 39](#)
- [“Example Deployment Scenarios” on page 41](#)

Multi-Broker Clusters

Message Queue supports using multiple interconnected broker instances known as a *broker cluster*. With broker clusters, client connections are distributed across all the brokers in the cluster. Clustering provides horizontal scalability and improves availability.

A single message broker scales to about eight CPUs and provides sufficient throughput for typical applications. If a broker process fails, it is automatically restarted. However, as the number of clients connected to a broker increases, and as the number of messages being delivered increases, a broker will eventually exceed limitations such as number of file descriptors and memory.

Having multiple brokers in a cluster rather than a single broker enables you to:

- Provide messaging services despite hardware failures on a single machine.
- Minimize downtime while performing system maintenance.
- Accommodate workgroups having different user repositories.
- Deal with firewall restrictions.

Message Queue allows you to create conventional or enhanced broker clusters. Conventional broker clusters offer service availability. Enhanced broker clusters offer both service and data availability. For more information, see Chapter 10, “Configuring and Managing Broker Clusters,” in *Open Message Queue 4.5 Administration Guide*.

In a conventional cluster, having multiple brokers does not ensure that transactions in progress at the time of a broker failure will continue on the alternate broker. Although Message Queue reestablishes a failed connection with a different broker in a cluster, transactions owned by the failed broker are not available until it restarts. Except for failed in-progress transactions, user applications can continue on the failed-over connection. Service failover is thus ensured.

In an enhanced cluster, transactions and persistent messages owned by the failed broker are taken over by another running broker in the cluster and non-prepared transactions are rolled back. Data failover is ensured for prepared transactions and persisted messages.

Master Broker and Client Synchronization for Conventional Clusters

In a configuration for a conventional broker cluster, each destination is replicated on all of the brokers in a cluster. Each broker knows about message consumers that are registered for destinations on all other brokers. Each broker can therefore route messages from its own directly-connected message producers to remote message consumers, and deliver messages from remote producers to its own directly-connected consumers.

In a cluster configuration, the broker to which each message producer is directly connected performs the routing for messages sent to it by that producer. Hence, a persistent message is both stored and routed by the message’s *home broker*.

Whenever an administrator creates or destroys a destination on a broker, this information is automatically propagated to all other brokers in a cluster. Similarly, whenever a message consumer is registered with its home broker, or whenever a consumer is disconnected from its home broker—either explicitly or because of a client or network failure, or because its home broker goes down—the relevant information about the consumer is propagated throughout the cluster. In a similar fashion, information about durable subscriptions is also propagated to all brokers in a cluster.

A shared database of cluster change records can be configured as an alternative to using a master broker. For more information, see Chapter 10, “Configuring and Managing Broker Clusters,” in *Open Message Queue 4.5 Administration Guide* and “Using Message Queue Broker Clusters With GlassFish Server” in *GlassFish Server Open Source Edition 3.1 High Availability Administration Guide*.

Configuring GlassFish Server to Use Message Queue Brokers

By default, Message Queue brokers (JMS hosts) run in the same JVM as the GlassFish Server process. However, Message Queue brokers (JMS hosts) can be configured to run in a separate JVM from the GlassFish Server process. This allows multiple GlassFish Server instances or clusters to share the same set of Message Queue brokers.

The GlassFish Server’s Java Message Service represents the connector module (resource adapter) for Message Queue. You can manage the Java Message Service through the Administration Console or the `asadmin` command-line utility.

In GlassFish Server, a *JMS host* refers to a Message Queue broker. The GlassFish Server’s Java Message Service configuration contains a JMS Host List (also called `AddressList`) that contains all the JMS hosts that will be used.

Java Message Service Type

There are three types of integration between GlassFish Server and Message Queue brokers: embedded, local, and remote. You can set this type attribute on the Administration Console’s Java Message Service page.

Embedded Java Message Service

If the `Type` attribute is `EMBEDDED`, GlassFish Server and the JMS broker are colocated in the same virtual machine. The JMS Service is started in-process and managed by GlassFish Server. In `EMBEDDED` mode, JMS operations on stand-alone server instances bypass the networking stack, which leads to performance optimization. The `EMBEDDED` type is most suitable for stand-alone GlassFish Server instances. `EMBEDDED` mode is not supported for enhanced broker clusters.

With the `EMBEDDED` type, use the `Start Arguments` attribute to specify Message Queue broker startup parameters.

With the `EMBEDDED` type, make sure the Java heap size is large enough to allow GlassFish Server and Message Queue to run in the same virtual machine.

Local Java Message Service

If the Type attribute is LOCAL, GlassFish Server starts and stops the Message Queue broker. When GlassFish Server starts up, it starts the Message Queue broker specified as the Default JMS host. Likewise, when the GlassFish Server instance shuts down, it shuts down the Message Queue broker. The LOCAL type is most suitable for use with enhanced broker clusters, and for other cases where the administrator prefers the use of separate JVMs.

With the LOCAL type, use the Start Arguments attribute to specify Message Queue broker startup parameters.

Remote Java Message Service

If the Type attribute is REMOTE, GlassFish Server uses an externally configured broker or broker cluster. In this case, you must start and stop Message Queue brokers separately from GlassFish Server, and use Message Queue tools to configure and tune the broker or broker cluster. The REMOTE type is most suitable for brokers running on different machines from the server instances (to share the load among more machines or for higher availability), or for using a different number of brokers and server instances.

With the REMOTE type, you must specify Message Queue broker startup parameters using Message Queue tools. The Start Arguments attribute is ignored.

Managing JMS with the Administration Console

In the Administration Console, you can set JMS properties using the Java Message Service node for a particular configuration. You can set properties such as Reconnect Interval and Reconnect Attempts. For more information, see Chapter 16, “Administering the Java Message Service (JMS),” in *GlassFish Server Open Source Edition 3.1 Administration Guide*.

The JMS Hosts node under the Java Message Service node contains a list of JMS hosts. You can add and remove hosts from the list. For each host, you can set the host name, port number, and the administration user name and password. By default, the JMS Hosts list contains one Message Queue broker, called “default_JMS_host,” that represents the local Message Queue broker integrated with GlassFish Server.

In REMOTE mode, configure the JMS Hosts list to contain all the Message Queue brokers in the cluster. For example, to set up a cluster containing three Message Queue brokers, add a JMS host within the Java Message Service for each one. Message Queue clients use the configuration information in the Java Message Service to communicate with Message Queue broker.

Managing JMS with asadmin

In addition to the Administration Console, you can use the `asadmin` command-line utility to manage the Java Message Service and JMS hosts. Use the following `asadmin` commands:

- Configuring Java Message Service attributes: `asadmin set`

- Managing JMS hosts:
 - `asadmin create-jms-host`
 - `asadmin delete-jms-host`
 - `asadmin list-jms-hosts`
- Managing JMS resources:
 - `asadmin create-jms-resource`
 - `asadmin delete-jms-resource`
 - `asadmin list-jms-resources`

For more information on these commands, see *GlassFish Server Open Source Edition 3.1-3.1.1 Reference Manual* or the corresponding man pages.

Default JMS Host

You can specify the default JMS Host in the Administration Console Java Message Service page. If the Java Message Service type is LOCAL, GlassFish Server starts the default JMS host when the GlassFish Server instance starts. If the Java Message Service type is EMBEDDED, the default JMS host is started lazily when needed.

In REMOTE mode, to use a Message Queue broker cluster, delete the default JMS host, then add all the Message Queue brokers in the cluster as JMS hosts. In this case, the default JMS host becomes the first JMS host in the JMS host list.

You can also explicitly set the default JMS host to one of the JMS hosts. When the GlassFish Server uses a Message Queue cluster, the default JMS host executes Message Queue-specific commands. For example, when a physical destination is created for a Message Queue broker cluster, the default JMS host executes the command to create the physical destinations, but all brokers in the cluster use the physical destination.

Example Deployment Scenarios

To accommodate your messaging needs, modify the Java Message Service and JMS host list to suit your deployment, performance, and availability needs. The following sections describe some typical scenarios.

For best availability, deploy Message Queue brokers and GlassFish Servers on different machines, if messaging needs are not just with GlassFish Server. Another option is to run a GlassFish Server instance and a Message Queue broker instance on each machine until there is sufficient messaging capacity.

Default Deployment

Installing the GlassFish Server automatically creates a domain administration server (DAS). By default, the Java Message Service type for the DAS is EMBEDDED. So, starting the DAS also starts its default Message Queue broker.

Creating a new domain also creates a new broker. By default, when you add a stand-alone server instance or a cluster to the domain, its Java Message Service is configured as EMBEDDED and its default JMS host is the broker started by the DAS.

Using a Message Queue Broker Cluster with a GlassFish Server Cluster

In EMBEDDED or LOCAL mode, when a GlassFish Server is configured, a Message Queue broker cluster is auto-configured with each GlassFish Server instance associated with a Message Queue broker instance.

In REMOTE mode, to configure a GlassFish Server cluster to use a Message Queue broker cluster, add all the Message Queue brokers as JMS hosts in the GlassFish Server's Java Message Service. Any JMS connection factories created and MDBs deployed then uses the JMS configuration specified.

Specifying an Application-Specific Message Queue Broker Cluster

In some cases, an application may need to use a different Message Queue broker cluster than the one used by the GlassFish Server cluster. To do so, use the `AddressList` property of a JMS connection factory or the `activation-config` element in an MDB deployment descriptor to specify the Message Queue broker cluster.

For more information about configuring connection factories, see “Administering JMS Connection Factories and Destinations” in *GlassFish Server Open Source Edition 3.1 Administration Guide*. For more information about MDBs, see “Using Message-Driven Beans” in *GlassFish Server Open Source Edition 3.1 Application Development Guide*.

Application Clients

When an application client or standalone application accesses a JMS administered object for the first time, the client JVM retrieves the Java Message Service configuration from the server. Further changes to the JMS service will not be available to the client JVM until it is restarted.

Checklist for Deployment

This appendix provides a checklist to get started on evaluation and production with the GlassFish Server.

Checklist

TABLE 3-1 Checklist

Component/Feature	Description
Application	<p>Determine the following requirements for the application to be deployed.</p> <ul style="list-style-type: none"> Required/acceptable response time. Peak load characteristics. Necessary persistence scope and frequency. Session timeout in <code>web.xml</code>. Failover and availability requirements. <p>For more information see <i>GlassFish Server Open Source Edition 3.1 Performance Tuning Guide</i>.</p>
Hardware	<ul style="list-style-type: none"> Have necessary amounts of hard disk space and memory installed. Use the sizing exercise to identify the requirements for deployment. <p>For more information see <i>GlassFish Server Open Source Edition 3.1-3.1.1 Release Notes</i></p>
Operating System	<ul style="list-style-type: none"> Ensure that the product is installed on a supported platform. Ensure that the patch levels are up-to-date and accurate. <p>For more information see <i>GlassFish Server Open Source Edition 3.1-3.1.1 Release Notes</i></p>

TABLE 3-1 Checklist (Continued)

Component/Feature	Description
Network Infrastructure	<ul style="list-style-type: none"> Identify single points of failures and address them. Make sure that the NICs and other network components are correctly configured. Run <code>ttcp</code> benchmark test to determine if the throughput meets the requirements/expected result. Setup <code>ssh</code> based your preference. <p>For more information see <i>GlassFish Server Open Source Edition 3.1 Installation Guide</i>.</p>
Back-ends and other external data sources	Check with the domain expert or vendor to ensure that these data sources are configured appropriately.
System Changes/Configuration	<ul style="list-style-type: none"> Make sure that changes to <code>/etc/system</code> and its equivalent on Linux are completed before running any performance/stress tests. Make sure the changes to the TCP/IP settings are complete. By default, the system comes with lots of services pre-configured. Not all of them are required to be running. Turn off services that are not needed to conserve system resources. On Solaris, use <code>Settoolkit</code> to determine the behavior of the system. Resolve any flags that show up. <p>For more information see <i>GlassFish Server Open Source Edition 3.1 Performance Tuning Guide</i>.</p>
Installation	<ul style="list-style-type: none"> Ensure that these servers are not installed on NFS mounted volumes. Check for enough disk space and RAM when installing GlassFish Server.
GlassFish Server Configuration	<ul style="list-style-type: none"> Logging: Enable access log rotation. Choose the right logging level. WARNING is usually appropriate. Configure Java EE containers using the Administration Console. Configure HTTP listeners using the Administration Console. Configure ORB threadpool using the Administration Console. If using Type2 drivers or calls involving native code, ensure that <code>mtmalloc.so</code> is specified in the <code>LD_LIBRARY_PATH</code>. Ensure that the appropriate persistence scope and frequency are used and they are not overridden underneath in the individual Web/EJB modules. Ensure that only critical methods in the SFSB are checkpointed. <p>For more information on tuning, see <i>GlassFish Server Open Source Edition 3.1 Performance Tuning Guide</i>.</p> <p>For more information on configuration, see <i>GlassFish Server Open Source Edition 3.1 Administration Guide</i>.</p>

TABLE 3-1 Checklist (Continued)

Component/Feature	Description
Load balancer Configuration	<ul style="list-style-type: none"> ■ Make sure the Web Server is installed. ■ Make sure the load balancer plug-in into the Web Server is installed. ■ Make sure patch checks is disabled. ■ Lower the value of the <code>KeepAliveQuery</code> parameter. The lower the value, the lower the latency is on lightly loaded systems. The higher the value, the higher the throughput is on highly loaded systems.
Java Virtual Machine Configuration	<ul style="list-style-type: none"> ■ Initially set the minimum and maximum heap sizes to be the same, and at least one GB for each instance. ■ See Java Hotspot VM Options for more information. ■ When running multiple instances of GlassFish Server, consider creating a processor set and bind GlassFish Server to it. This helps in cases where the CMS collector is used to sweep the old generation.
Configuring time-outs in the load balancer	<ul style="list-style-type: none"> ■ <code>Response-time-out-in-seconds</code> - How long the load balancer waits before declaring a GlassFish Server instance unhealthy. Set this value based on the response time of the application. If set too high, the Web Server and load balancer plug-in wait a long time before marking a GlassFish Server instance as unhealthy. If set too low and GlassFish Server's response time crosses this threshold, the instance will be incorrectly marked as unhealthy. ■ <code>Interval-in-seconds</code> - Time in seconds after which unhealthy instances are checked to find out if they have returned to a healthy state. Too low a value generates extra traffic from the load balancer plug-in to GlassFish Server instances and too high a value delays the routing of requests to the instance that has turned healthy. ■ <code>Timeout-in-seconds</code> - Duration for a response to be obtained for a health check request. Adjust this value based on the traffic among the systems in the cluster to ensure that the health check succeeds.
Configuring time-outs in GlassFish Server	<ul style="list-style-type: none"> ■ <code>Max-wait-time-millis</code> - Wait time to get a connection from the pool before throwing an exception. Default is 6 s. Consider changing this value for highly loaded systems where the size of the data being persisted is greater than 50 KB. ■ <code>Cache-idle-timeout-in-seconds</code> - Time an EJB is allowed to be idle in the cache before it gets passivated. Applies only to entity beans and stateful session beans. ■ <code>Removal-timeout-in-seconds</code> - Time that an EJB remains passivated (idle in the backup store). Default value is 60 minutes. Adjust this value based on the need for SFSB failover.

TABLE 3-1 Checklist (Continued)

Component/Feature	Description
Tune VM Garbage Collection (GC)	<p>Garbage collection pauses of four seconds or more can cause intermittent problems in persisting session state. To avoid this problem, tune the VM heap. In cases where even a single failure to persist data is unacceptable or when the system is not fully loaded, use the CMS collector or the throughput collector.</p> <p>These can be enabled by adding:</p> <pre><jvm-options>-XX:+UseConcMarkSweepGC</jvm-options></pre> <p>This option may decrease throughput.</p>

Index

A

- active users, 28
- Administration Console, 19
- administrative domain, 18
- Apache Web Server, 20
- applications, 15
- asadmin command, 19
- availability, 34–35
 - and clusters, 34
 - and redundancy, 35

B

- bandwidth, 32
- broker cluster, 37, 42

C

- capacity, using spare machines to maintain, 35
- checklist, 43
- clients, 16
 - and JMS, 42
- clusters, 19
 - and availability, 34
 - Message Queue, 37, 42
- Coherence*Web, 22
- components, 18
- concurrent users, 28
- configurations, 20
 - default, 20

- connectors, 17
- containers, 15

D

- DAS, 19
- default configuration, 20
- default deployment, 41
- default JMS Host, 41
- default server, 19
- deployment planning, 27
 - checklist, 43
 - example scenarios, 41
- design decisions, 36
- domain, 18
- Domain Administration Server (DAS), 19

E

- EJB container, 15
- ethernet cards, 33

F

- failover capacity, planning, 35
- failure
 - classes, 35
 - types, 35
- fault tolerance, 35

H

HTTP sessions, 21

I

InitialContext, 23

instances of the server, 18, 36

J

Java API for XML-based RPC (JAX-RPC), 16

Java API for XML Registries (JAXR), 16

Java Authorization Contract for Containers
(JACC), 16

Java Database Connectivity (JDBC), 17

Java EE Connector architecture, 17

Java EE services, 16

Java Message Service (JMS), 16, 17, 23, 37

Java Naming and Directory Interface (JNDI), 16

Java platform, Enterprise Edition (Java EE), 15

JavaMail API, 17

L

load

server, 28, 33

load balancing

HTTP, 20

IIOP, 22

M

machines, maintaining capacity with spare
machines, 35

message broker, 38

message-driven beans, 23

Message Queue, 23, 37

Microsoft Internet Information Server, 20

mirror machines, 35

N

named configuration, 20

naming, 16

network cards, 33

network configuration, 34–35
server, 31

O

Object Request Broker (ORB), 16

P

peak load, 33

performance, 27

persistence, session, 21

R

redundancy, 35

remote browser emulator, 27

requests per minute, 30

resource adapters, 17

resources, 17

response time, 29

S

security, 16

server

clusters, 19

components, 18

containers, 15

domain administration, 19

instances, 18, 36

load, 28, 33

network configuration, 31

performance, 27

services, 16

sessions

HTTP, 21

sessions (*Continued*)

- persistence, 21

Simple Mail Transport Protocol (SMTP), 17

sizing

- administration thread pool, 37

- system, 36

spare machines, maintaining capacity with, 35

stateful session beans, 21

T

think time, 29

thread pool, sizing for administration, 37

throughput, 28

transactions, 16

types, of failure, 35

U

users, concurrent, 28

W

web container, 15

Web Server, 20

web servers, 20

web services, 16

Web Services Description Language (WSDL), 16

Web Services-Interoperability (WS-I), 17

