GemStone®

GemBuilder for Java™
Programming Guide

September 2002

Version 2.0 for GemStone/S 6.0.1

Send your comments about this manual to docs@gemstone.com

IMPORTANT NOTICE

This documentation is furnished for informational use only and is subject to change without notice. GemStone Systems, Inc. assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation. The documentation, or any part of it, may not be reproduced, displayed, photocopied, transmitted or otherwise copied in any form or by any means now known or later developed, such as electronic, optical or mechanical means, without written authorization from GemStone Systems, Inc. Any unauthorized copying may be a violation of law.

The software installed in accordance with this documentation is copyrighted and licensed by GemStone Systems, Inc. under separate license agreement. This software may only be used pursuant to the terms and conditions of such license agreement. Any other use may be a violation of law.

Copyright © GemStone Systems, Inc. 1994-2002. All Rights Reserved.

Use, duplication, or disclosure by the Government is subject to restrictions set forth in the Commercial Software - Restricted Rights clause at 52.227-19 of the Federal Acquisitions Regulations (48 CFR 52.227-19) except that the government agency shall not have the right to disclose this software to support service contractors or their subcontractors without the prior written consent of GemStone Systems, Inc.

Trademarks

GemStone, **GemBuilder**, **GemConnect**, and the GemStone logo are trademarks or registered trademarks of GemStone Systems, Inc. in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries.

UNIX is a registered trademark in the United States and other countries, licensed exclusively through X/Open Company Limited.

Sun, **Sun Microsystems** and **Solaris** are trademarks or registered trademarks of Sun Microsystems, Inc. All **SPARC** trademarks, including **SPARCstation**, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation is licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

Microsoft, **MS**, **Windows** and **Windows NT** are registered trademarks of Microsoft Corporation in the U.S.A. and other countries.

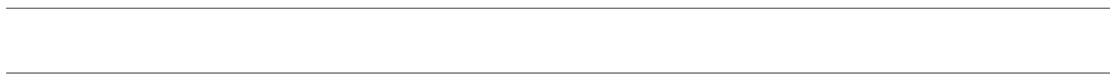

Netscape is a registered trademark of Netscape Communications Corporation in the United States and other countries.

Other company or product names mentioned herein may be trademarks or registered trademarks of their respective owners. Specifications are subject to change without notice. All terms mentioned in this documentation that are known to be trademarks or service marks have been appropriately capitalized. GemStone cannot attest to the accuracy of this information. Use of a term in this documentation should not be regarded as affecting the validity of any trademark or service mark.

Preface

About This Documentation

This documentation describes GemBuilder® for Java™ 2.0, an application programming interface (API) for developing distributed Gemstone® applications with a Java client and a GemStone/S server.

GemBuilder for Java pairs the portability and flexibility of Java with the scalability of the GemStone server. You can develop business applications around Java clients that take advantage of the server's multi-user object execution engine, transaction management facilities, and fault tolerant environment. The Java clients can be distributed as standalone applications or as applets that provide a universal client from a single, easily maintained source.

GemBuilder for Java consists of two parts: the programming interface between your Java client and the GemStone server, and Java-based tools for developing GemStone Smalltalk code and inspecting objects in the server.

The programming interface in this release lets your client do the following:

- locate GemStone server objects and obtain stub references to them;
- send messages to GemStone server objects through stub references;

- send messages to Java objects from GemStone server objects that implement an adapter interface;
- execute ad-hoc Smalltalk code on the server; and
- handle exceptions signaled by server objects in a natural fashion.

The GemBuilder for Java Tools let you work with GemStone Smalltalk in the server. These tools, which can be incorporated into a Java development environment or used independently, are described in separate documentation, *GemBuilder for Java Tools*. The toolset includes:

- A *GemStone Browser* for examining, creating, and modifying GemStone classes and methods.
- A *Workspace* for easy access to commonly used code operations.
- *Inspectors* for examining and modifying the state of GemStone objects.
- A *Debugger* for examining an execution stack left by a GemStone application.
- A *Class Version Browser* for examining a class history, inspecting instances, migrating instances, deleting versions, and moving versions to another class history.

Assumptions

To make use of the information in this documentation, you need to be familiar with the GemStone server and with GemStone's Smalltalk programming language as described in the *GemStone/S Programming Guide*. That book explains the basic concepts behind the language and describes the most important GemStone kernel classes. The documentation for this release assumes you have an existing GemStone application for which you want to create Java clients.

In addition, you should be familiar with the Java language, the Java API, and the development environment as described in your vendor's documentation.

Finally, this documentation assumes that the GemStone system has been correctly installed on your host computer as described in the *GemStone System Administration Guide* and that your system meets the requirements listed in your *GemBuilder Installation Guide*.

How This Manual is Organized

Basic Concepts describes the overall design of a GemBuilder application and presents the fundamental concepts required to understand the interface between a Java client and the GemStone server.

Communicating with the Server explains how to communicate with the GemStone server by initiating and managing GemStone sessions, and how to set up and maintain the Session Broker service through which Java clients to log in.

Interacting with Server Objects describes how to locate objects in the server and obtain stubs referencing them, how to send messages to the objects or execute ad-hoc Smalltalk code, how to handle exceptions raised on the server, and how objects are marshaled between the Java client and the server.

Forwarding Server Messages to Client Objects explains how GemStone objects can send messages to Java objects by using an adapter to compile the message into Java code.

Managing Server Transactions discusses the process of committing a transaction, the kinds of conflicts that can prevent a successful commit, and how to avoid or resolve such conflicts.

Observing Session and Server Events explains how your application can monitor events in client sessions and certain events in the server.

Deploying Your Application explains the steps you need to take to deploy your Java clients for use with the GemStone server.

Other Useful Documents

- *GemBuilder for Java Tools* online documentation describes the independent set of tools that let you explore and modify Smalltalk code in the server. These files and a file containing a printable version accompany the GemBuilder for Java and GemBuilder for the Web products.
- The *GemStone/S Programming Guide* describes the GemStone System and the GemStone Smalltalk language.
- If you will be acting as a system administrator, or developing software for someone else who must play this role, you should read the *GemStone System Administration Guide.*

Technical Support

GemStone provides several sources for product information and support. GemBuilder product manuals provide extensive documentation, and should always be your first source of information. GemStone Technical Support engineers will refer you to these documents when applicable. However, you may need to contact Technical Support for the following reasons:

- Your technical question is not answered in the documentation.
- You receive an error message that directs you to contact GemStone Technical Support.
- You want to report a bug.
- You want to submit a feature request.

Questions concerning product availability, pricing, keyfiles, or future features should be directed to your GemStone account manager.

When contacting GemStone Technical Support, please be prepared to provide the following information:

- Your name, company name, and GemStone/S license number
- The GemBuilder product and version you are using
- The hardware platform and operating system you are using
- A description of the problem or request
- Exact error message(s) received, if any

Your GemBuilder support agreement may identify specific individuals who are responsible for submitting all support requests to GemStone. If so, please submit your information through those individuals. All responses will be sent to authorized contacts only.

For non-emergency requests, you should contact Technical Support by web form, email, or facsimile. You will receive confirmation of your request, and a request assignment number for tracking. Replies will be sent by email whenever possible, regardless of how they were received.

World Wide Web: http://support.gemstone.com

This is the preferred method of contact. The Help Request link is at the top right corner of the home page—please use this to submit help requests. This form requires an account, but registration is free of charge. To get an account, just complete the Registration Form, found in the same location. You'll be able to access the site as soon as you submit the web form.

Email: support@gemstone.com

Please do not send files larger than 100K (for example, core dumps) to this address. A special address for large files will be provided as appropriate.

Facsimile: (503) 629-8556

When you send a fax to Technical Support, you should also leave a voicemail message to make sure your fax will be picked up as soon as possible.

Telephone: (800) 243-4772 or (503) 533-3503

We recommend you use telephone contact only for more serious requests that require immediate evaluation, such as a production system that is non-operational.

Emergency requests are handled by the first available engineer. If you are reporting an emergency and you receive a recorded message, do not use the voicemail option. Transfer your call to the operator, who will take a message and immediately contact an engineer.

Non-emergency requests received by telephone are placed in the normal support queue for evaluation and response.

24x7 Emergency Technical Support

GemStone offers, at an additional charge, 24x7 emergency technical support. This support entitles customers to contact us 24 hours a day, 7 days a week, 365 days a year, if they encounter problems that cause their production application to go down, or that have the potential to bring their production application down. Contact your GemStone account manager for more details.

GemStone Web Site

The GemStone Web site, at `http://support.gemstone.com`, provides a variety of information to help you use GemBuilder products. Use of this site requires an account, but registration is free of charge. To get an account, just complete the Registration Form, found in the same location. You'll be able to access the site as soon as you submit the web form.

The following types of information are provided at this web site:

Help Request is an online form that allows designated technical support contacts to submit requests for information or assistance via email to GemStone Technical Support.

Technotes provide answers to questions of general interest submitted by GemBuilder customers. They may contain coding examples, links to other sources of information, or downloadable code.

Bugnotes identify performance issues or error conditions that you may encounter when using a GemBuilder product. A bugnote describes the cause of the condition, and, when possible, provides an alternative means of accomplishing the task. In addition, bugnotes identify whether or not a fix is available, either by upgrading to another version of the product, or by applying a patch. Bugnotes are updated regularly.

Patches provide code fixes and enhancements that have been developed after product release. A patch generally addresses a specific group of behavior or performance issues. Most patches listed on the GemStone Web site are available for direct downloading.

Tips and Examples provide information and instructions for topics that usually relate to more effective or efficient use of GemStone products. Some Tips may contain code that can be downloaded for use at your site.

Release Notes and Install Guides for your product software are provided in PDF format.

Documentation for GemBuilder is provided in PDF format. This is the same documentation that is included with your GemBuilder product, with the exception of the API reference files (javadocs).

Community Links provide customer forums for discussion of GemBuilder product issues.

Technical information on the GemStone Web site is reviewed and updated regularly. We recommend that you check this site on a regular basis to obtain the latest technical information for GemStone products. We also welcome suggestions and ideas for improving and expanding our site to better serve you.

Training and Consulting

Consulting and training for all GemStone products are available through GemStone's Professional Services organization.

- Training courses are offered periodically at GemStone's offices in Beaverton, Oregon, or you can arrange for onsite training at your desired location.
- Customized consulting services can help you make the best use of GemStone products in your business environment.

Contact your GemStone account representative for more details or to obtain consulting services.

Contents

Chapter 1. Basic Concepts

The GemStone Solution 1-1
About the GemStone/S Server 1-1
About GemBuilder for Java 1-2
Integrating Information across the Enterprise 1-3
GemStone/S Sessions 1-3
The Session Broker 1-3
Development Strategy 1-4
Using GemBuilder for Java with Your Development Environment 1-5
Development Steps 1-5
Partitioning Your Application 1-6

Chapter 2. Communicating With the Server

Overview. 2-1
Opening a Session. 2-2
Creating the Session Parameters 2-2
Creating the Session and Connecting to GemStone/S 2-4

Launching Tools From Your Application 2-5
Logging of Debugging Information. 2-5
Closing a Session 2-6
Administering the Server Component 2-6
Running the Session Broker 2-7
Effect of NetLDI Mode 2-7
Configuration Files 2-8
To Start the Session Broker. 2-11
To Halt the Session Broker 2-12
Connecting to the Session Broker Gem 2-12
To Run Multiple Session Brokers 2-13
Maintaining the Log Directory 2-14
Troubleshooting 2-14
To Determine if a Session Broker Is Running 2-15
To Restart a Session Broker 2-15
To Locate Log Files 2-15
To Enable Verbose Logging 2-15

Chapter 3. Interacting with Server Objects

Overview 3-1
The Message-forwarding Interface 3-2
Using Stub Protocol to Send Messages 3-3
Sending Dynamic Messages 3-4
Handling Server Exceptions 3-5
To Invoke a Debugger on an Exception 3-8
How Objects are Marshaled 3-8
Using GemStone/S's DoubleByteString 3-9
Controlling How Objects Are Marshaled 3-9
Writing the State of an Object to Send to the Server 3-9
Reading the State of an Object Sent from the Server 3-10
Writing the State of an Object to Send to the Client. 3-10
Reading the State of an Object Sent from the Client. 3-11
Representing Server Objects in the Client 3-11
Deciding Which Objects to Represent 3-11
Obtaining GbjObject Stubs 3-12
Looking Up a Named Object in the Server. 3-12
Saving a Returned Stub 3-13

Registering a Custom Stub . . . 3-13
Accessing a Stub's Cached Value . . . 3-14
Effect of Multiple Class Versions . . . 3-15
Executing Ad-hoc Smalltalk Code . . . 3-15
Accessing Complex Objects Efficiently . . . 3-16
Getting All Named Instance Variables . . . 3-16
Flattening Objects in the Server . . . 3-17
Replicating Objects Using Holders . . . 3-17
Working with Collections . . . 3-18
The GbjCollection Protocol . . . 3-19
Protocol Examples . . . 3-20
Serializing the Collection in the Server . . . 3-20
Unpacking the Collection in the Client . . . 3-22
To Enumerate the Collection . . . 3-23
To Unpack the Collection from an Array . . . 3-23
Obtaining Application-specific Stubs . . . 3-24
Registering Stubs at Static Initialization . . . 3-24
Registering Stubs at Runtime . . . 3-25
Putting Client Data into the Server . . . 3-26

Chapter 4. Forwarding Server Messages to Client Objects

Overview . . . 4-1
Using Reflection . . . 4-2
Implementing the Adapter Interface . . . 4-3
Registering a Client Adapter . . . 4-3
Dealing with Multithreading . . . 4-3
Message-sends in the Server . . . 4-4
Exceptions Raised in the Client . . . 4-6

Chapter 5. Managing Server Transactions

Overview . . . 5-1
Operating Inside a Transaction . . . 5-2
Committing a Transaction . . . 5-4
Aborting a Transaction . . . 5-4
Handling Commit Failures . . . 5-5

Operating Outside a Transaction 5-5
Being Signaled to Abort. 5-7
Transaction Modes 5-8
Manual Transaction Mode 5-9
Automatic Transaction Mode 5-9
Transactionless Mode 5-10
Choosing Which Mode to Use 5-10
Switching Between Modes 5-10
Managing Concurrent Transactions 5-11
Read and Write Operations. 5-11
Optimistic and Pessimistic Concurrency Control 5-12
Setting the Concurrency Mode. 5-13
Setting Locks 5-13
Reduced-Conflict Classes 5-15

Chapter 6. Observing Session and Server Events

Overview 6-1
Observing Session Events 6-1
To Monitor Session Events. 6-2
Observing Server Events. 6-3

Chapter 7. Deploying Your Application

Overview 7-1
Deployment Steps 7-2
To Deploy an Applet 7-2
To Deploy a Standalone Application 7-3

Glossary

Chapter 1

Basic Concepts

The GemStone Solution

This overview describes GemStone's solution for Internet and corporate intranet applications: the GemStone/S Server and GemBuilder for Java (GBJ).

About the GemStone/S Server

The GemStone/S server provides a wide range of services to help you build object-based information systems. GemStone/S:

- supports transaction-intensive, business-critical applications involving more than 1000 concurrent users and persistent object spaces in the tens of gigabytes
- provides a distributed server architecture that allows the server and processes to be spread over multiple hardware platforms and operating systems in a heterogeneous computing environment
- provides object persistence on an instance basis (determined by reachability from other persistent objects) for greater flexibility compared to object server systems that use class-based persistence or object-relational mapping
- provides transactions having ACID properties (atomicity, consistency, isolation, durability)

- provides configurable concurrency modes and protocols for requesting locks on individual client Smalltalk objects and collections of objects, allowing developers to exercise fine-grained control over concurrent access to objects
- provides automatic referential integrity and an on-line garbage collection process that runs in the background
- supports embedded queries, path queries, collection queries, non-locking queries, multi-user indexes and extensible indexes
- can monitor events or changes in state of objects and send signals to other applications or to users, eliminating the need for inefficient polling

About GemBuilder for Java

The GemBuilder for Java API is a Java runtime package that provides a message forwarding interface between a Java client and a GemStone/S server:

- Java clients can locate GemStone/S server objects by name and obtain stub references to them.
- Java clients can send messages to GemStone/S server objects through stub references.
- GemStone/S server objects can send messages to Java client objects that implement an adapter interface.

The API does not include user interface classes or application frameworks. The focus is on transparent messaging and simple replication.

The GemBuilder for Java Tools are implemented in Java, so you can create Java clients using your preferred Java development environment, then create server-side applications in GemStone Smalltalk without leaving the Java environment. The set includes the following tools:

- A GemStone Browser lets you view the available GemStone Smalltalk classes and methods, create new classes, and add or modify methods.
- An Inspector lets you view objects residing in the GemStone/S server. You can examine the state of individual objects, modify them if desired, browse collections, or look at the contents of dictionaries.
- A Workspace lets you execute arbitrary strings of GemStone Smalltalk code. Use this tool to locate server objects to be examined in an Inspector, create sample data for testing applications, assign access control to objects or collections, and perform administration tasks in the GemStone/S server.

- A Debugger is available whenever execution of a GemStone/S method results in a run-time error.

Integrating Information across the Enterprise

Through GemStone's distributed architecture, Java applications have the use of objects residing anywhere in the enterprise. And through GemStone's connectivity tools, Java applications can also have access to business data from relational and legacy databases:

- By using GemConnect, GemBuilder for Java clients have access to Oracle and Sybase RDBMSs.
- GemStone Object Transaction Services provide complete control over transactions that include data from heterogeneous sources.

GemStone/S Sessions

All interaction with the GemStone/S server takes place in the context of a *session*, which is a login to the GemStone/S server under a particular GemStone/S User ID. Because this context remains in effect until the session is closed, it does not have to be reestablished for each individual service request.

Each session is associated with its own *Gem* process, which acts as the GemStone/S server for that session. It is the Gem process that accesses the shared objects and executes GemStone Smalltalk code.

The Session Broker

An intermediary called the *Session Broker* is responsible for assigning a Gem process to serve a Java client, and then placing the Gem and client in communication with each other. A GemStone/S administrator starts a Session Broker (an instance of GbjBroker) from a Topaz session, after choosing a TCP/IP port to serve as a well-known port. By default, the port number is 9090.

To connect to GemStone/S, a Java client must provide that port number and the name of the GemStone/S server as two of the session parameters. GemBuilder for Java then asks the Session Broker to provide a Gem session process; the Session Broker, in turn, calls on a NetLDI network server to spawn the Gem.

Once the connection has been established, further communication takes place solely between the Java client and the Gem.

Figure 1.1 Session Broker Connecting a Client to the Server

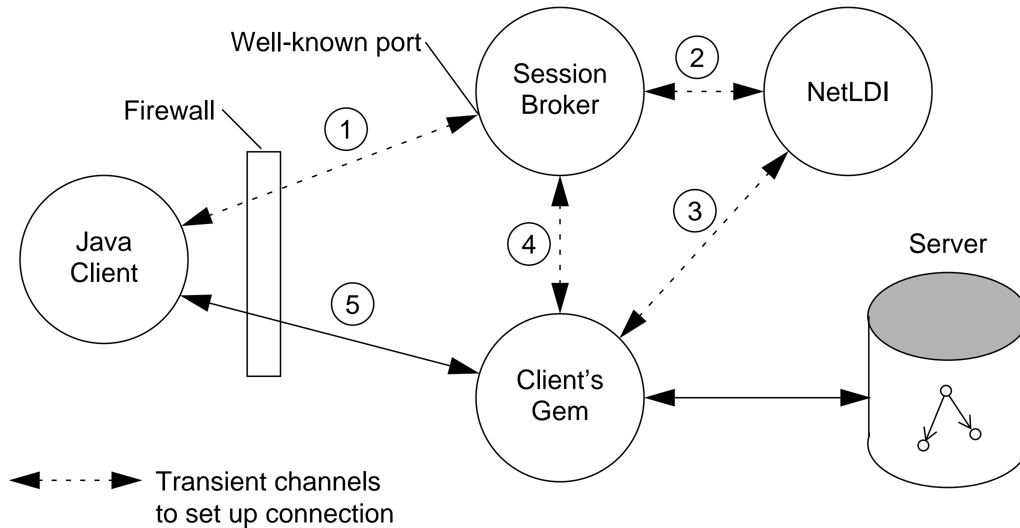

Development Strategy

This release of GemBuilder for Java supports distributed application development with the GemStone/S server. GemStone/S is used as an object management system in which Java clients can take advantage of GemStone Smalltalk method execution in the server. GBJ provides limited replication of simple data types to the client.

Application partitioning must be considered in the development of GBJ applications from the start because the client component is written in a different language than the server component.

Using GemBuilder for Java with Your Development Environment

GemBuilder for Java works in conjunction with the Java development environment (JDE) of your choice. After installing GBJ on your development platform, do the following:

- Add `gbjtools20.jar` to your CLASSPATH (deployed applications will use `gbj20.jar`, which does not include the GBJ Tools).
- Add an item to your JDE tools menu to launch the GBJ Tools. For instance:

```
java com.gemstone.tools.GbjLauncher
```

- Add "com.`gemstone.gbj.*`" to the import list for each Java class where it is appropriate.
- Optionally, add an item to the JDE tools menu to open a browser on the online documentation home page, `GbjIndex.html`. Or, add a bookmark in your Web browser. For instance:

```
netscape c:\GBJ10\doc\GbjIndex.html
```

Development Steps

The typical steps in developing an application using GemBuilder for Java are:

1. Model the system using accepted OO modeling techniques.
2. Identify objects that will provide server-based services.
3. Use the GBJ Tools to create server classes and methods.
4. Write the Java client classes using tools of your choice.
5. Test the application.

A Rapid Application Development approach could also be taken where classes and methods are built in both the client and the server during the building and testing of the application. In this approach, you would have the GemBuilder for Java Tools running along side the Java development environment. You would add or modify methods in GemStone/S, commit them to the server, and make the requisite iterative runs of the Java application within a Java development environment. Changes you make within the Tools session must be committed to the GemStone/S server before you test them in the Java application because the application being tested typically will run in a separate GemStone/S session.

Partitioning Your Application

The recommended approach to application partitioning is to keep user-interface work in the Java client, and to keep business objects — shared procedures, rules, and data — in the GemStone/S server. In general, the recommended approach is:

- Have the server handle object processing and queries. Avoid having the server side know about details of the user interface where that is not necessary. Instead, let the client request the objects the interface needs.
- Have the client make the minimum number of requests to the server by using the techniques described under “Accessing Complex Objects Efficiently” on page 3-16. Making requests efficiently is especially important for Internet applications where the network is quite slow.

The GemBuilder for Java Tools can also be used to test and debug GemStone/S server-side classes and methods independently of a Java development environment. The Workspace, Inspector, and Debugger tools are useful in performing unit tests on server classes and in examining the state of the GemStone/S server to find and correct bugs.

GbjObject, which extends java.lang.Object, is GemBuilder for Java's principal class. Instances are *stubs* that represent objects in the GemStone/S server.

Although the stub and the holder classes present one recommended approach to implementing the client, neither is inherently necessary; satisfactory results could have been achieved in other ways.

Chapter 2

Communicating With the Server

Overview

All interaction with the GemStone/S server takes place through a session that is dedicated to serve your Java client. Each session is an instance of GbjSession, which provides the environment and a number of important capabilities for interaction with the GemStone/S server. Through GbjSession, your Java client can:

- connect to and disconnect from the server
- locate server objects by name and obtain stub instances representing them
- install custom stub objects for specific kinds of server objects
- execute ad-hoc GemStone Smalltalk code in the server
- control transactions in the server
- install adapters that dispatch messages from server objects to client objects
- observe events in the server, such as a change to a specific object or a Gem-to-Gem (session to session) signal, and certain events that happen in another session object, such as the committing of a transaction.

Each session has a public variable, userData, in which you store session-specific information, such as context information that you want to be available elsewhere in the client. Because this variable is of type Object, it must be cast to the appropriate class when retrieving objects stored in it.

The session communicates with the GemStone/S server through a Gem process, which is started by a Session Broker (an instance of GbjBroker) during login, as described under "The Session Broker" on page 1-3. The Session Broker must be started before Java clients can use the server.

Opening a Session

One of the first tasks your client needs to perform is to open a session with (log in to) the appropriate GemStone/S server. You do that in two steps:

1. Create and fill in an instance of GbjParameters.
2. Use the parameters to create an instance of GbjSession, then connect the session to the GemStone/S server.

The GbjSession not only provides a connection with the server, it also provides a context, such as a name space, object access authorizations, and privileges. Once the connection is established, your Java application can access GemStone/S objects in any way permitted by the security privileges associated with that GemStone/S userId.

Creating the Session Parameters

An instance of GbjParameters defines the GemStone/S server, userId, and other information needed to log in to GemStone/S. The following code sets up parameters for a session when the user clicks the application's Connect button.

Example 2.1 Setting the Session Parameters

```
import com.gemstone.gbj.*;
GbjParameters params;
params = new GbjParameters();

// required parameters without defaults
params.userName = "DataCurator";
params.password = "swordfish";
// parameters that have defaults as shown
params.brokerMachine = "localhost";
params.brokerPort = 9090;
params.serverName = "gemserver60";
params.gemnetName = "gemnetobject";
params.timeout = 900;
params.connectTimeout = 30;
params.transactionMode = GbjParameters.ManualTransactions;
```

The **userName** and **password** fields must match an existing GemStone/S userId in the server. There are no defaults for these fields.

In the example, the GemStone/S server and Session Broker are located on this machine, so the **brokerMachine** field can be set to the machine name or to *localhost*. In other configurations, the name of a remote machine can be specified. The default is "localhost". The default for the **serverName** field is "gemserver60".

You must know, or prompt for, the TCP/IP port (**brokerPort**) at which the Session Broker is listening for connection requests. Port 9090 is commonly used, but the actual port number is set by the configuration file `gbj.ini`. The default is 9090.

The **gemnetName** field determines which Gem service start up script is read when the session starts. For most GemStone/S installations, the name is `gemnetobject` (the default, for Bourne or Korn login shells) or `gemnetobjcsh` (for users of the C shell).

The **timeout** field sets the length of time (in seconds) that the server will wait before it drops an idle session. The default is 900 seconds (15 minutes).

The **connectTimeout** field sets the length of time (in seconds) that the Java client will wait for a response from the Session Broker. The default is 30 seconds.

The **transactionMode** statement sets the way in which transactions are started in a GemStone/S session. The default mode under GemBuilder for Java is ManualTransactions, in which transactions must be started explicitly. AutomaticTransactions specifies a chained mode in which a transaction is started when the session connects to the server and new transactions begin after a commitTransaction or abortTransaction is processed. For further information, see “Transaction Modes” on page 5-8.

An RPC login may also require a host operating system user name and password as parameters (**serverOSUserName** and **serverOSPassword**), depending on the mode in which the NetLDI is running. For further information, see “Effect of NetLDI Mode” on page 2-7.

NOTE:
Ordinarily, a GemStone/S NetLDI network server must be running on the Session Broker's machine so the broker can use it to spawn Gems.

For further information about the NetLDI, refer to the *GemStone System Administration Guide*.

Creating the Session and Connecting to GemStone/S

After you have created and filled in an instance of GbjParameters, the next step is to create the session instance and connect it to the server. This example is a continuation of the previous one.

Example 2.2 Connecting to GemStone/S

```
GbjSession mySession;
mySession = new GbjSession(params);
mySession.connect();
```

Launching Tools From Your Application

You can start the GemBuilder for Java Tools from your application to facilitate debugging from the GemStone/S session in which the application is running. (The class definition must include `com.gemstone.tools.*` in place of, or in addition to `com.gemstone.gbj.*`.)

To open a Tools launcher for a new session, embed this code in your application:

```
new GbjLauncher();
```

Alternatively, provide an active, logged in session as an argument to the constructor:

```
new GbjLauncher(aGbjSession);
```

Your application can use a session's launcher text pane as a transcript by sending the message transcript() to a launcher instance. Because the transcript is a Java TextArea, you can use appropriate protocol, such as appendText(). For instance,

```
this.launcher.transcript().appendText("Done");
```

Logging of Debugging Information

When you are diagnosing problems involving client-server interaction, it may help to enable verbose logging on one or both sides. An example would be a situation in which the server component mysteriously "hangs" with no message to the client and nothing in the ordinary log.

The GbjSession method setDebuggingOptions() takes two boolean arguments, *verboseClient* and *verboseServer*. These parameters can be set programmatically or by using **File > Settings** menu picks in the GemStone Launcher.

When *verboseClient* is true, API debugging messages are displayed in System.out.

When *verboseServer* is true, server debugging messages are written to the session log file, `gbj`*nn*`.log` (for the location of the log files, see "To Locate Log Files" on page 2-15). If this argument is set prior to logging in, connection messages also appear in the broker log.

The following example logs the use of a Workspace to evaluate a GemStone Smalltalk expression:

Example 2.3

```
31/01/1997 16:58:49: reading next buffer for message
31/01/1997 16:58:49: reading packet length
31/01/1997 16:58:49: reading 2 bytes:
31/01/1997 16:58:49: aByteArray( 0, 25)
…
31/01/1997 17:11:26: Reading next key
31/01/1997 17:11:26: Reading field 24
31/01/1997 17:11:26:  value = 'evaluate:inContext:names:values:printing:'
31/01/1997 17:11:26: Reading next key
31/01/1997 17:11:26: Reading field 17
31/01/1997 17:11:26:  value = 5
31/01/1997 17:11:26: Reading next key
31/01/1997 17:11:26: Reading field 100
31/01/1997 17:11:26:  value = 'System currentSessionNames'
```

Closing a Session

When the time comes to terminate the session, send the message close() to the session object, and perform any other housekeeping necessary.

WARNING:
Closing the session logs out the user's GemStone/S session. Any uncommitted changes are lost.

Administering the Server Component

This topic provides information to help you configure the Session Broker and troubleshoot problems that may arise.

Running the Session Broker

A Session Broker can be run on the GemStone/S server's machine or on any other supported platform that has a network connection to the server. The primary consideration in locating a Session Broker is that the Gems it spawns will be located on the same machine. There are three parts to consider: the GemStone/S server, the Session Broker and Gems, and the Java clients. Each part can be on a separate machine, so several configurations are possible:

- The Session Broker can run on the Server's machine, and all Gems it spawns will be on that machine. This configuration is the simplest, and is suitable in most cases.
- The Session Broker can run on a client's machine, and the Gems will be spawned on the client's machine. Where necessary, each of several machines can have its own Session Broker.
- The Session Broker and Gems can run on an separate machine, offloading that activity from the server and client machines.

NOTE:
Ordinarily, a GemStone/S NetLDI network server must be running on the Session Broker's machine so the broker can use it to spawn Gems.

Effect of NetLDI Mode

The mode in which the NetLDI is running affects the ownership of the Gems that are spawned and the network authentication requirements. In general, the effect is the same as for other parts of a GemStone/S installation:

- The default mode requires authentication by way of a host user name and password on the Session Broker's machine. This information can be provided as session parameters or by a `.netrc` file. The Gem is owned by the specified host user.
- Guest mode makes it unnecessary to provide authentication. Under UNIX, this mode typically is combined with captive account mode, which causes the captive account to own the Gem processes. Under Windows NT, guest mode is the recommended mode; it reduces security but provides increased convenience.

Configuration Files

You can use a configuration file to provide settings for the Session Broker to use at startup. The configuration file is optional; if you specify none, the Session Broker looks for a file in the current directory with the default name `gbj.ini`. If it finds no such file, it uses default values.

You can use environment variables in a configuration file if you delimit them with the character %. For example:

```
fileDirectory=%GEMSTONE%/data
```

Lines preceded by a semi-colon are comments. The following example configuration file explains each parameter:

Example 2.4 Configuration File

```
; GBJ Broker Configuration Parameters
; start section --
[gbjbroker]
;
; serverPort is the port number to use for the broker's
; server socket.
; Permissible values: integer > 0.
; serverPort=9090
serverPort=9090
;
; ports are numbers to use for temporary server sockets
; in spawned gems.
; Permissible values: a sequence of integers or
; range of integers, separated by spaces or commas.
; For example: ports= 9091 9093-9094,9096
; ports=9091-9099
ports=9091-9099
;
; fileDirectory is the path name of the log file
; directory -- the location for log files in spawned gems
; and default location for the broker's log file.
; Permissible values: a valid directory path for
; the host OS.
; Delimit environment variables with %.
; fileDirectory=.
fileDirectory=.
;
; logFile is the name of the broker's log file.
; If no directory is specified, the file is placed in
```

```
; #fileDirectory.
; Delimit environment variables with %.
; The default log file name is 'gbjbroker.log'.
; Permissible values: a valid file name for the host OS.
; logFile=gbjbroker.log
logFile=gbjbroker.log
;
; If verbose is true, the broker writes detailed messages
; to the log file. If verbose is false, the broker writes
; error messages only.
; Permissible values: true or false.
; verbose=false
verbose=false
;
; gbjAdministrator and gbjAdministratorPassword are the
; GemStone username and password used by the startgbj and
; stopgbj scripts for starting and stopping a broker.
; Permissible values: valid GemStone username and
; password.
gbjAdministrator=DataCurator
gbjAdministratorPassword=swordfish
;
; gbjBrokerConnectToken is the GemStone username for
; connecting to a running broker.
; This need not be a valid GemStone username.
; If the GemStone username in a connection request
; matches #gbjBrokerConnectToken, the broker establishes
; a connection in the broker's gem.
; gbjBrokerConnectToken=GbjAdministrator
gbjBrokerConnectToken=GbjAdministrator
;
; Next are GemStone parameters for spawning gems.
; Permissible values: valid GemStone login parameters.
gemStoneName=gemserver60
username=DataCurator
password=swordfish
hostUsername=
hostPassword=
gemService=gemnetobject
;
; initialGemCount is the number of gems to spawn on
; startup. Gems are spawned as needed.  If you know how
; many you will need, starting all at once may save time.
; Permissible values: an integer >= 0 and >=
; #initialGemMinimum.
; initialGemCount=1
```

```
initialGemCount=1
;
; initialGemMinimum is the minimum number of gems to
; spawn on startup. If this many gems are not created
; before the number of seconds specified by
; #initialDelay, the broker quits.
; Permissible values: an integer >= 0 and <=
; #initialGemCount.
; initialGemMinimum=1
initialGemMinimum=1
;
; initialDelay is the number of seconds to wait for
; the #initialGemMinimum gems on startup.
; The broker quits if this many seconds pass before
; #initialGemMinimum gems have been created.
; 0 means no time limit.
; Permissible values: an integer >= 0
; initialDelay=180
initialDelay=180
;
; maximumGemCount is the maximum number of gems possible.
; Permissible values: an integer >= 0 and >=
; #maximumIdleGems and >= #initialGemCount.
; Default is the value of the #STN_MAX_SESSIONS GemStone
; configuration option minus one for the broker session.
; 0 means no limit.
; Note: The default value of #STN_MAX_SESSIONS is 40.
; (See GemStone Administration Guide).
; maximumGemCount=0
;
; maximumIdleGems is the maximum number of idle gems.
; If a gem goes idle, causing the number of idle gems to
; exceed this number, the gem is terminated.
; Permissible values: an integer >= 0 and <=
; #maximumGemCount.
; Default is the value of the #maximumGemCount option.
; 0 means no limit.
; maximumIdleGems=0
; End of configuration file.
```

To Start the Session Broker

You can start the Session Broker using either the startup script provided, or Topaz (GemStone's command line programming environment).

A configuration file is required if you wish to run the startup script. To do so, enter the following at the operating system command prompt:

```
C:\> startgbj gbj.ini
```

You can substitute another file name for `gbj.ini` if you wish to use a different configuration file. If you supply no file name, the GemBuilder for Java looks for a file named `gbj.ini` in the current directory. If it finds no such file, the script fails.

To use the command line:

1. Log in to the GemStone/S server as any user, using Topaz.
2. To use the default settings, enter at the command line:

   ```
   topaz 1> doit
   GbjBroker new startup
   %
   ```

3. Note the port number and the log directory that are displayed. You can minimize the window, but don't close it — that would halt the Session Broker.

To start a Session Broker from the command line using a different configuration file, enter:

```
topaz 1> doit
(GbjBroker newOnFile: 'c:\mygbj.ini')
startup
%
```

The GbjBroker class also provides methods to override specific configuration settings, either from a configuration file or the defaults. To override the server port setting, for example, enter:

```
topaz 1> doit
(GbjBroker newOnFile: 'gbj.ini')
serverPort: 9080;
startup
%
```

For related information, see "Configuration Files" on page 2-8.

To Halt the Session Broker

You can halt the Session Broker using either the shutdown script provided, or by connecting to the Session Broker as described in the next topic.

A configuration file is also required if you wish to run the shutdown script. To do so, enter the following at the operating system command prompt:

```
C:\> stopgbj gbj.ini
```

This shuts down only if all Gems are idle (not connected). If some Gems are active, use either the -**i** or -**t** option.

The -**i** option shuts down the Session Broker immediately, even if Gems are still active.

Follow -**t** with an integer representing a number of seconds. This option disables logins and then shuts down Gems as soon as they are idle, until the specified number of seconds has passed. If some Gems are still active, the Session Broker is not shut down and logins are reenabled.

An additional optional argument, -**h**, prints usage information and exits.

You can substitute another file name for `gbj.ini` if you're using a different configuration file. If you supply no file name, the GemBuilder for Java looks for a file named `gbj.ini` in the current directory. If it finds no such file, the script fails.

NOTE:
In this release, it is not possible to resume execution of the Session Broker after pressing Control-C.

Connecting to the Session Broker Gem

You can connect to the Session Broker's gem in order to shut it down, or to perform other administrative tasks. Doing so requires that you use a configuration file.

Logging In

Log in through GemBuilder for Java with a GemStone/S username that matches the gbjBrokerConnectToken parameter in the configuration file. The default is GbjAdministrator. No password is needed.

Shutting Down the Session Broker

Once connected to the Session Broker's Gem, you can shut it down by sending any of several messages either to the GbjBroker class, or to the current instance of GbjBroker.

For example:

```
GbjBroker shutdown
```

or

```
GbjBroker current shutdown
```

The shutdown messages available are:

`shutdown`	Stops the Session Broker if all Gems are unconnected.
`shutdownImmediate`	Stops the Session Broker immediately, whether Gems are unconnected or busy, terminating all Gems.
`shutdownWait`	Logins are disabled. The Session Broker waits for all Gems to disconnect before shutting down.
shutdownWait:	Takes an integer argument specifying the number of seconds to wait for all Gems to become idle before stopping the Session Broker. Logins are disabled immediately, except for other logins to the Session Broker Gem. If the number of seconds specified elapses and some Gems are still busy, logins are reenabled and the Session Broker is not shut down. To specify that the Session Broker must wait indefinitely, supply a negative number as an argument.
`shutdownCancel`	Stops the shutdown process started by `shutdownWait:`.

To Run Multiple Session Brokers

You can run more than one Session Broker on a single machine, although there is no need to do so (you can connect to any local GemStone/S server through a single Session Broker). You must be careful to use separate ports and log directories for each broker on the same machine.

1. Choose a set of unused TCP/IP ports and a log directory for the new broker.

2. Make a copy of the `gbj.ini` configuration file (in the GemBuilder for Java server directory) under another name.
3. Edit the copy to substitute the chosen ports and log directory. The lines you need to change look like these:

```
serverPort=9090
ports=9091-9099
fileDirectory=.
```

4. Follow the instructions for starting a Session Broker, logging in to a different server and using the modified configuration file.

Maintaining the Log Directory

You should plan to remove old log files periodically from the file directory specified in the configuration file, `gbj.ini`. The files that the Session Broker and each Gem create in this directory are not deleted automatically.

1. The Session Broker log, `gbjbroker.log`, continues to grow until the Broker is restarted, at which time the previous log file is replaced by a new one.
2. Each Gem creates a log file having the name `gbj`*nn*`.log`, where *nn* is a unique, sequential number that begins at 1 and increments until the Session Broker is restarted, at which time the number is reset to a value of 1.

CAUTION:
Depending on the frequency at which clients are started, the length of time the Session Broker has been running, and the chosen degree of logging detail (verbosity), it is possible for the log files to fill the disk.

Troubleshooting

If you have trouble connecting to the server, first make sure the Session Broker is running and that you have the correct machine name and port number. These can be verified by reviewing information displayed at the time the broker was started or by examining the broker's log file.

You might find it helpful to examine the broker's log file for clues. If a problem persists, try restarting the broker under verbose logging so more information is available.

To Determine if a Session Broker Is Running

1. Go to the window in which the Session Broker was started.
2. Check the status:
   - If the last status line reads “Ready” and the prompt has not returned, the Session Broker is running.
   - If the prompt has returned, the Session Broker has exited.

TIP:
Note the GbjBroker port number and the location of the broker log file. These items are part of the status information displayed by the startup script.

To Restart a Session Broker

1. Make sure the previous Session Broker was terminated cleanly. If in doubt, stop the broker as described under “To Halt the Session Broker” on page 2-12.
2. Start a new Session Broker. You can use the same port numbers.

For related information, see “To Start the Session Broker” on page 2-11.

To Locate Log Files

The log files for the Session Broker and the Gems are located in the directory specified by the `fileDirectory` instance variable of the GemStone GbjBroker. This variable is set by the broker configuration file, `gbj.ini`.

- If possible, check the window in which the Session Broker is running. The current value of `fileDirectory` is displayed at the time you start the Session Broker.
- If necessary, check the configuration file to determine the setting that was used. Its initial location is the GemBuilder for Java `server` directory.
- If the NetLDI is running under a captive account, the session log files (`gbj`*nn*`.log`) are created in the home directory of the captive account or as otherwise specified when the NetLDI was started.

To Enable Verbose Logging

1. Open the file `gbj.ini` (in the `server` directory) for editing.
2. Change the line that controls verbose logging so it looks like this:

   ```
   verbose=true
   ```

3. If the Session Broker is running, halt it.

4. Restart the broker.

For related information, see “To Halt the Session Broker” on page 2-12 and “To Start the Session Broker” on page 2-11.

You can enable verbose logging for a particular session by choosing **File** > **Settings** >**Verbose Client** or **Verbose Server**.

Chapter 3

Interacting with Server Objects

Overview

GemBuilder for Java (GBJ) provides a *message-forwarding interface* through which your Java client can interact with objects in the GemStone/S server. Your client code explicitly obtains *stubs,* which are client Java objects (instances of GbjObject or its subclasses) that represent objects in GemStone/S. A stub knows which GemStone/S object it represents, and it responds to all messages in its protocol by passing them to the appropriate GemStone/S object.

TERMINOLOGY NOTE:

Users of GemBuilder for Java for Smalltalk will notice the term stub ***as used here differs from that to which they are accustomed and corresponds to what they would call a*** forwarder. ***This difference in terminology is unavoidable because it is anchored in the Java literature.***

Because almost any message sent to a GemStone/S object is capable of raising an exception in the server, a portion of your Java code will be devoted to catching and handling these exceptions when they are returned to the client. Ordinary Java techniques are appropriate, and there is a GBJ-specific subclass, GbjException.

The Message-forwarding Interface

Once a connection to the server has been established, your client can use the session to interact with persistent objects. The process typically involves two fundamental steps:

1. locating the object by name, which returns a *stub* representing the server object, and
2. sending a message through the stub and receiving a reply in the form of another stub representing the result returned in the server. In the case of simple data types, the stub acts as a Java wrapper for a cached value (see “Accessing a Stub’s Cached Value” on page 3-14).

Figure 3.1 Stubs Representing GemStone/S Objects

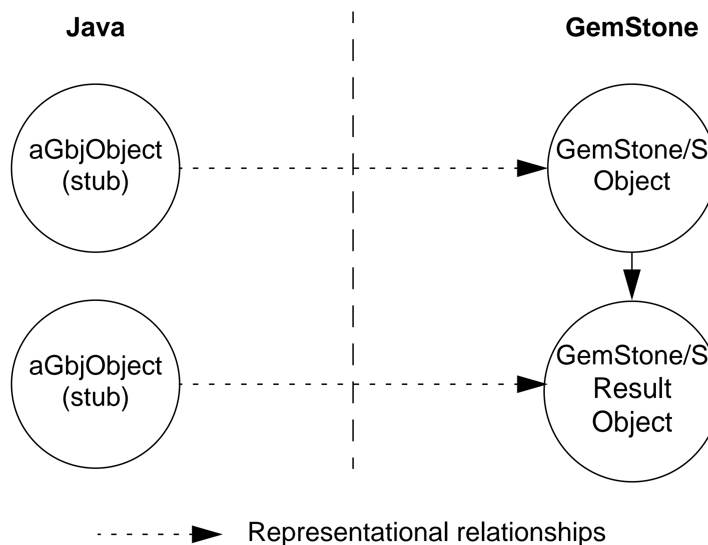

Each stub is an instance of GbjObject or one of its subclasses. GbjObject has three members, all public:

Member	Description
oop	the object identifier (OID) of the server object it represents
session	the GbjSession instance this stub uses to communicate with the server
cachedValue	if the stub represents a server object that corresponds to a fundamental Java data type, a DateTime, or a serialized collection of proxies, the cachedValue member holds that object; if not, cachedValue is null

Using Stub Protocol to Send Messages

GbjObject implements methods corresponding to the fundamental ones in the server's class Object. As a result, these are inherited by your stubs without additional effort. Here are a few examples of commonly used methods (for a complete list, see the class description):

at()	remoteEqualsIdentical()
atPut()	remoteSize()
equals()	segment()
in()	sendMsg()
isKindOf()	toString()
notNil()	

The class GbjCollection, a subclass of GbjObject, defines additional fundamental methods common to server Collection classes. For information, see “Working with Collections” on page 3-18.

NOTE:
A Java stub must explicitly implement all of the methods it wishes to forward. Java does not have a mechanism similar to Smalltalk’s doesNotUnderstand:.

Sending Dynamic Messages

The method `sendMsg()` in GbjObject lets you send GemStone Smalltalk messages to server objects without implementing corresponding protocol in Java.

The following example uses the server message firstName in a selection block to retrieve the first customer found with that name, then uses `sendMsg()` to send the message lastName to that instance. The method `detect()` is defined in GbjCollection. The object returned by `sendMsg()` is a GbjObject, and since the name is a simple data type, its value can be obtained from the stub's cachedValue member by using `stringValue()`.

Example 3.1 Sending a Message to a Server Object

```
GbjCollection myCollStub;
GbjObject aCust;

myCollStub = (GbjCollection) mySession.doit("Customer allCustomers");
aCust = myCollStub.detect("[:cust | cust firstName = 'Barney']");
System.out.println("Last Name is " +
    aCust.sendMsg("lastName").stringValue());
```

Class GbjObject implements sendMsg() several ways with different method signatures. Here is a partial list:

Method Signature	Use
sendMsg(String)	unary message
sendMsg(String, Object)	one-keyword message or binary message
sendMsg(String, Object, String, Object) ...	two-keyword message

and so forth up to a message with five keywords. For instance:

```
GbjObject obj = mySession.doit("Array new: 10");
obj.sendMsg("size");
obj.sendMsg("at:", new Integer(1));
obj.sendMsg("at:", new Integer(1), "put:", Boolean.TRUE);
```

Because the last two examples take Java String/Object pairs, the int argument must be wrapped in an Integer object. GemBuilder for Java automatically converts this Java Integer to a GemStone Integer (SmallInteger or LargeInteger) as part of the marshaling process.

The GbjObject method perform() occasionally is a useful alternative to `sendMsg()` because it takes a user-specified number of arguments and the arguments can be treated as a unit.

This example builds an Array of arguments needed to create a new engineer instance, then uses the Array as an argument to `perform()`, which creates the new instance in the server. Finally, `sendMsg()` adds the instance to the existing Collection. The client class Engineer was previously registered as a stub for the corresponding server class.

Example 3.2 Using GbjObject.perform With an Argument List

```
GbjObject engClass = null;
        Object args[] = {eml, firstNm, lastNm, ph};

        // get stub for the class
        engClass = sess.objectNamed("Engineer");
        Engineer eng = null;
        try {
            eng = (Engineer) engClass.perform(
                "email:firstName:lastName:phone:", args, 4);
        } catch (GbjEventException e) {
            new gemstone.tools.GbjDebugger(e);
        }
        engClass.sendMsg("addEngineer:", eng);
```

Handling Server Exceptions

When interaction with GemStone/S causes an error to be raised in the server, notification is returned to the client in the form of an instance of GbjException, which extends java.lang.RuntimeException. With this change from the GemBuilder for Java 1.0 superclass, the compiler ignores GbjException when enforcing catch-throw semantics. Therefore, a method that calls another method that throws a GbjException is no longer forced either to catch the exception, or to declare in its header that it throws the exception.

Here is a typical scenario: (1) the client object sends a message to a stub, which (2) the stub forwards to the server. If this message results in an error in the server, (3) the GemStone/S error information is sent back to the GemBuilder for Java, which (4) throws it as a GbjException.

Figure 3.2 Throwing a GbjException

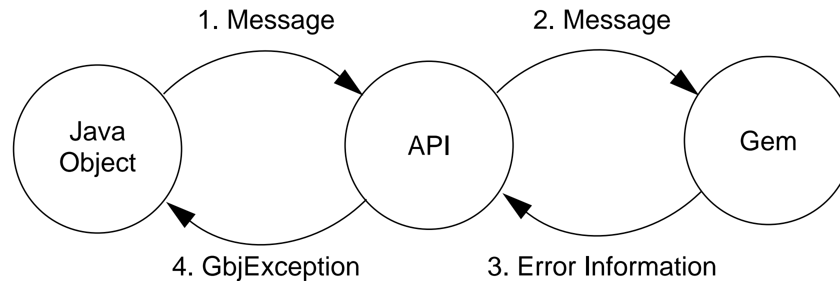

Typical code to handle the GbjException is similar to that for any other Java exception:

Example 3.3 Handling a GbjException

```
public String firstName() {
    String nm = null;
    try {
        nm = this.sendMsg("firstName").stringValue();
    } catch (GbjEventException e) {
        System.err.println("GemStone event exception occurred"
            + e.getMessage());
    } catch (GbjException e) {
        System.err.println("Engineer>firstName() GemStone exception: "
          + e.getMessage());
    }
    return nm;
    }
```

NOTE:
Most methods that access the server can cause an error to be raised in the server. Because try{}catch{} statements should be part of your GemBuilder for Java coding practice, they appear frequently in the examples in this documentation.

You can determine the nature of the exception by examining its category and number members. The class GbjGemStoneErrors provides variables for all of the GemStone/S kernel class error numbers (that is, for those associated with GbjExceptions that have category GemStoneError). The instance variable GbjException.kernel also holds this list. For instance, to find out whether a lock you obtained on an object in GemStone is dirty (that is, whether the object has changed since you started the current transaction), you could use this test:

Example 3.4 Catching a Specific GemStone/S Error

```
try {
    // obtain a lock in the server
}
catch (GbjException e) {
    if (e.number ==
        GbjException.kernel.LOCK_ERR_OBJ_HAS_CHANGED) {
        // handle dirty lock
    }
    // else throw e
}
```

However, developers can catch more specific error conditions in GemStone/S using any of the following subclasses of GbjException. These subclasses correspond to the exception categories in GbjGemStoneErrors.java:

- GbjCompilerException
- GbjRuntimeException
- GbjAbortingException
- GbjFatalException
- GbjInternalException
- GbjEventException
- GbjNonKernelException

To Invoke a Debugger on an Exception

1. Have the class import `com.gemstone.tools.*`.
2. To open the debugger on a particular exception, use the GbjDebugger constructor with the exception as an argument. For instance:

```
catch (GbjException e) {
new com.gemstone.tools.GbjDebugger(e);
}
```

How Objects are Marshaled

When a message is sent to a GemStone/S server object, the arguments to the message can be any kind of Object. GemBuilder for Java marshals these objects (serializes them into a stream) and transmits them with the message. In this release, marshaling performs these conversions between Java and GemStone/S in either direction:

Java Object	GemStone Smalltalk Equivalent
Byte	SmallInteger
Integer	Integer (may be Small or Large)
Long	Integer (may be Small or Large)
Float	SmallFloat
Double	Float
Character	Character
Boolean	Boolean
null	nil
Date	DateTime
String (one-byte strings)	String (one-byte strings)
DoubleByteString (two-byte strings)	DoubleByteString

True and false are objects in GemStone/S (instances of Boolean) but not in Java. In Java, you must wrap true or false in a Boolean object, which is then marshaled as shown above.

Instances of GemStone DecimalFloat, ScaledDecimal, and Fraction don't have direct counterparts in Java. Any of these objects are marshaled as a GbjObject reference. You can send these objects a conversion message in the server to obtain a binary Float, which will be marshaled according to the preceding table.

Using GemStone/S's DoubleByteString

GemBuilder for Java assumes Java String and StringBuffer objects are single-byte strings, which matches GemStone/S's implementation of String. If one of these client objects actually holds a two-byte string, you must specify that by wrapping the object in an instance of com.gemstone.gbj.DoubleByteString. This class is comparable to wrappers like Integer that are provided in java.lang for wrapping Java primitive data types. For instance:

```
new DoubleByteString(aJavaString)
```

Unless a Java String containing two-byte characters is wrapped in a DoubleByteString, only the low byte of each character will be transferred to the server. The server will see a GemStone/S String, which inherently is a single-byte string.

Both GemStone/S's String and DoubleByteString classes are ultimately mapped to the Java String class, since it can be used for either. The cachedValue of a stub for a GemStone/S DoubleByteString will hold a Java DoubleByteString wrapper. Clients must use dbStringValue() to access the String object held in a stub for a GemStone/S DoubleByteString.

Controlling How Objects Are Marshaled

Argument marshaling in both the client and server allows objects to control how they are marshaled. Client and server both use the same concepts, but require slightly different implementations due to the nature of the two languages. On both client and server, the object must implement an externalization interface whose methods read and write state to and from an object stream.

In Java a class must implement the GbjExternalizable interface. This interface has three methods:

- gbjServerObject()
- gbjWriteExternal()
- gbjReadExternal()

Writing the State of an Object to Send to the Server

The method gbjServerObject() returns the name of a class or object on the server that implements the method #gbjReadExternal: to read the externalized state of a Java object. gbjServerObject() can return a GbjObject stub for such a server object, instead of a name.

If the method returns an object name, it is resolved in the server's symbol list. If a class name is returned, the class is sent the message #new in the server. The #new method must return an object to perform the unmarshaling. To perform the unmarshaling, GemBuilder for Java sends the unmarshaling object the message #gbjReadExternal: with a Smalltalk GbjObjectInputStream as an argument. Methods in GbjObjectInputStream are used to read the state of the object.

The method gbjWriteExternal() is used to write the state of the object to the marshaling stream when sending a message to the server. The GBJ API uses the class GbjObjectOutputStream for marshaling. An instance of this class is sent as an argument to gbjWriteExternal(). Public methods in GbjObjectOutputStream are used to write the state of the object. The protocol is similar to the Java Object Serialization ObjectOutputStream class.

Reading the State of an Object Sent from the Server

The method gbjReadExternal() is used for reading the state of the object. An instance of GbjObjectInputStream is sent as an argument. Public methods in this class are used to read the exact state, no more and no less, that was written for the object in the server.

NOTE:
Because gbjServerObject(), gbjWriteExternal(), and gbjReadExternal() methods are sent during message marshaling and unmarshaling, the communications channel to the server is busy while these methods are executing. Therefore, avoid sending messages to GemStone/S, fetching objects from GemStone/S, or any other use of the server in the implementation of these methods.

Writing the State of an Object to Send to the Client

In the server, an object that reimplements the method #gbjExternalizable to return true must write its state to the client. The object must also implement the methods:

- #gbjClientObject
- #gbjReadExternal:
- #gbjWriteExternal:

The method #gbjClientObject returns a reference to a client object to be used for unmarshaling. This can be either a fully qualified name of a class that implements GbjExternalizable, or a client forwarder to an instance of such a class. If the name of a class is returned, an instance of this class is created in the client to perform the unmarshaling.

The API uses the method #gbjWriteExternal: to write the state of the Smalltalk object. It is passed an instance of the Smalltalk class GbjObjectOutputStream. Public methods in this class are used to write the state of the object.

Reading the State of an Object Sent from the Client

The method #gbjReadExternal: is used to read the object state written by a Java object in its gbjWriteExternal() method as described above. An instance of the Smalltalk class GbjObjectInputStream is used to retrieve the object state.

NOTE:
Make sure you marshal and unmarshal the same variables in the same order, so that the server reads exactly what the client has written to the stream, no more and no less.

The marshaling mechanism described here is based on concepts from the *Java Object Serialization Specification*, Revision 1.3, JDK 1.1, February, 1997.

You can find examples of marshaling and unmarshaling in the Examples installation directory. See the files `EFlattener.gs`, `EFlattener.java`, and `EFlattenerTest.java`, as well as `EVector.gs`, `EVector.java`, and `EVectorTest.java`.

CAUTION:
The file EFlattener.gs ***defines the GemStone Smalltalk classes EFlattener and Address in the Published dictionary. Filing in this example will replace an existing Address class is that dictionary.***

Representing Server Objects in the Client

Deciding Which Objects to Represent

GemBuilder for Java uses stubs (instances of GbjObject and its subclasses) to represent objects stored in GemStone/S to your Java client.

Because of the hierarchical structure of complex objects, you should begin by identifying the subsystems in your application that define persistent objects, and then identify a *root* object in each subsystem. The root objects of an application are the persistent objects from which other persistent objects can be reached by transitive closure; that is, either by direct reference or indirectly through any number of layers of references.

Each root object in the GemStone/S server in effect represents all other objects to which that object refers, such as its instance variables. And because those instance variables are represented, their instance variables are also represented, and so on, until you reach atomic objects that refer to no others, such as characters, integers, strings, booleans, or nil. The entire network of related objects forms a tree structure whose leaves are the final objects reached.

Obtaining a stub for the root object makes the entire subsystem in the server accessible to the client, since you can easily get stubs for other elements once you have a stub for the root. The most common kinds of root objects in the server are:

- global variables
- class variables
- class instance variables.

Obtaining GbjObject Stubs

There are three basic ways to obtain stubs:

- by looking up a named object in the server
- by sending a message to a server object through its stub and saving the stub that represents the object returned in the server
- by registering a Java class as a stub class for a corresponding server class

For related information, see “Accessing a Stub’s Cached Value” on page 3-14 and “Effect of Multiple Class Versions” on page 3-15.

Looking Up a Named Object in the Server

Performing an explicit name lookup in the server returns an instance of GbjObject representing the server object. This lookup is performed in the context of the current session logged in to GemStone/S; that is, it uses that session’s symbol list for name resolution. For instance, where AllEngineers has been defined previously as a global variable in the server:

Example 3.5 Resolving a Named Object

```
GbjObject myStub;
myStub = mySession.objectNamed("AllEngineers");
// use the stub as desired
System.out.println("Size: " +myStub.sendMsg("size").intValue());
```

Saving a Returned Stub

Sending a message to the server returns an instance of GbjObject (that is, a stub) to which you can send messages if appropriate. For instance, this code obtains a stub by using server class protocol to access AllEngineers; the stub is cast to type GbjCollection to make available the additional protocol defined for that class:

Example 3.6 Using the Result of GbjSession.doit

```
GbjCollection myCollStub;
myCollStub = (GbjCollection) mySession.doit( "Engineer allEngineers");
// use the stub as desired
System.out.println("Size: " +myCollStub.sendMsg("size").intValue());
```

The result that is represented need not be a persistent object (one which is part of the committed repository).

Registering a Custom Stub

Registering a custom stub creates a correspondence between specific Java and GemStone/S *classes*, which can provide for more natural coding within your application. The stub class must be a subclass of GbjObject or GbjCollection. Message sends to the server that correspond to the designated GemStone/S class return a stub that is instantiated in Java as an instance of the stub class.

Registration can be performed during static initialization or on an session-specific basis at runtime. For instance, to register the Java class Engineer (a subclass of GbjObject) as a stub for the server class of the same name:

```
mySession.registerStaticStub(new Engineer(), "Engineer");
```

At runtime, you could use inherited GbjObject protocol to send messages to an instance of the class. For instance, this expression in class Engineer makes the engineer available to accept an assignment:

```
this.execute("self available: true");
```

The method `execute()` is inherited from GbjObject, and available: is an instance method defined by the server class.

By implementing the method `available()` in the Java class itself, the expression could become more natural:

```
this.available(true);
```

Implement `available()` as shown below:

Example 3.7 Adding Server Protocol to the Java Side

```
public void available(boolean avail) {
        if (avail)
            this.execute("self available: true");
        else
            this.execute("self available: false");
    }
```

For more information about registering custom stubs, see “Obtaining Application-specific Stubs” on page 3-24.

Accessing a Stub’s Cached Value

When a result of a request is retrieved from GemStone/S, and that GemStone/S object is a kind of one of the classes listed under “How Objects are Marshaled” on page 3-8, the resulting Java object holds the pre-fetched value of the GemStone/S object it represents. This action minimizes the number of round trips to the server needed to get a usable form of the result.

GbjObject provides methods for getting at this cached value. The value itself is stored in the GbjObject variable *cachedValue*, which is public and may be accessed directly. GbjObject also supplies the following methods to retrieve a cachedValue in an appropriate way:

`booleanValue()`	`floatValue()`
`charValue()`	`intValue()`
`dateValue()`	`longValue()`
`dbStringValue()`	`stringValue()`
`doubleValue()`	`stringBufferValue()` (cachedValue holds a String that is converted to a StringBuffer)

Note that if the cachedValue member is directly accessed and holds an integer, an Integer wrapper is what will actually be found. This is also true of the other non-object values that are held in cachedValue. The cachedValue of a stub for a GemStone/S DoubleByteString holds a

`com.gemstone.gbj.DoubleByteString` wrapper. Clients must use `dbStringValue()` to access the String object held in a stub for a GemStone DoubleByteString.

Effect of Multiple Class Versions

GemBuilder for Java considers both the server class's classHistory and its class hierarchy in deciding which server objects a stub represents. For instance, a stub obtained for Engineer also represents any version of class Engineer as well as any subclasses of Engineer and its various versions.

Executing Ad-hoc Smalltalk Code

GemBuilder for Java provides two ways to execute *ad-hoc* code in the server, code that is not an existing compiled method.

The most general approach is to send doit() to a session object, which causes the code to be executed in that session's environment on the server. This example creates a key and value in the SymbolDictionary UserGlobals. The expression is sent to the session object for evaluation. The expression may contain temporaries and ^ (return) statements, but it may not contain refer to *self* or *super*.

Example 3.8 Executing Ad-hoc Code in the Server

```
mySession.doit(
   "UserGlobals at: #AllEngineers put: (Engineer allEngineers)"
   );
```

Alternatively, the method `GbjObject.execute()` permits you to send an ad hoc message to an object. The receiver, the server object represented by the instance of GbjObject to which the message is addressed, is the execution context. The message may use *self* to refer to the receiver and may directly reference instance variables of the receiver.

Accessing Complex Objects Efficiently

While a stub provides access to all server objects reachable by a transitive closure on that object, it can be inefficient to access a number of instance variables individually. You should consider these approaches, either individually or in combination:

- Transfer all named instance variables in one request.
- Flatten the instance variables into an Array of fundamental Java data types, which are automatically stored in the stub's cachedValue elements.
- Create a *holder* object as a proxy for the server object. Have the holder store explicit replicates of a few frequently needed instance variables as simple data types, and include the stub itself so you can access instance variables when necessary.

The above techniques also provide the building blocks for handling collections efficiently. For further information, see “Working with Collections” on page 3-18.

Getting All Named Instance Variables

The method namedInstanceVariables() returns all of an object’s named instance variables in an Array of GbjObjects (stubs), including those instance variables inherited from superclasses. Instance variables that are simple data types in Java are stored in the stub's cachedValue element.

The instance variables are in the order determined by the GemStone/S message `Behavior | allInstVarNames.`

This example first obtains a stub by sending the message myUserProfile to server class System, which is represented by a variable in GbjKernelObjects. Next, stubs for all of the UserProfile's named instance variables are obtained in a single request.

Example 3.9 Getting All Named Instance Variables

```
GbjObject uprofStub, uprofVars[];

uprofStub = mySession.kernel.System.sendMsg("myUserProfile");
uprofVars = uprofStub.namedInstanceVariables();
System.out.println("UserID: " +uprofVars[1].stringValue());
```

Flattening Objects in the Server

Many applications will want to pull information from GemStone/S to display in windows or otherwise use for user interfaces.

The recommended approach to accomplishing this task is to serialize simple data objects on the server side into sequenceable collections (such as Arrays or OrderedCollections) that can then be efficiently pulled into the client for processing.

This example gets information about one of the Customer objects, where *thisCustomer* is a previously obtained stub representing a particular customer. The method `allElements()` returns an array of GbjObjects holding the contents of the array constructed in the server.

Example 3.10 Flattening a Server Object

```
GbjObject fields[] = null;
fields = ((GbjCollection) thisCustomer.execute(
    "#[firstName, lastName, emailAddress]")).allElements();
String firstNameField = fields[0].stringValue();
String lastNameField =  fields[1].stringValue();
String emailField =  fields[2].stringValue();
System.out.println("Name: " +lastNameField +", " +firstNameField);
```

Replicating Objects Using Holders

An efficient way to handle complex server objects is to create a *holder* object in Java. Populate the holder with Java objects that replicate selected server instance variables in the form in which you need them. The holder can also store the stub itself for ease in sending additional messages to the server object. The concept explained here is particularly useful when it is extended to collections (see “Working with Collections” on page 3-18).

This example modifies the previous one on flattening objects by drawing on an additional class. MyHolder is a Java class with one instance variable that combines the separate first and last names in the server, and a *stub* instance variable. Where the previous example simply displayed the instance variables in a dialog, this example stores them in a holder. The stub can be used later for such tasks as retrieving the emailAddress instance variable by sending it the appropriate message.

Example 3.11 MyHolder.java

```
import gemstone.gbj.*;

public class MyHolder {
    public String name;
    public GbjObject stub;

   // ...
}
```

Example 3.12 Replicating an Object in MyHolder

```
GbjObject fields[] = null;
MyHolder thisCustomerHolder = new MyHolder();
fields = ((GbjCollection) thisCustomerStub.execute(
    "#[firstName, lastName, emailAddress]")).allElements();
thisCustomerHolder.name = fields[0].stringValue() + ' ' +
   fields[1].stringValue();
thisCustomerHolder.stub = thisCustomerStub;
```

Working with Collections

Collections typically are the core of an application using the GemStone/S server, so it is important that Java clients deal with them efficiently. The class GbjCollection provides additional protocol for that purpose.

In general, most interaction with large collections should be by way of remote message sends, GemStone Smalltalk code that is executed in the server. Smaller collections, or selected portions of larger ones, may be explicitly replicated in the client using this two-phase approach:

1. Serialize the desired objects in the server so elements can be brought to the client efficiently. Typically, this step involves selecting and flattening the desired instances, then bringing them together in a new collection.
2. In the client, iterate over the resulting collection, placing a replicate of each server object into a Java holder.

The GbjCollection Protocol

The GbjCollection class implements most of the protocol that is common to all Collection classes in GemStone/S. Your client can invoke these methods directly from Java, letting GemBuilder for Java forward the appropriate Smalltalk message. Here are some representative methods; for a complete list, see the description of GbjCollection.

Protocol Category	Representative Methods
Adding	add() addAll()
Converting	asArray() asBag() asIdentitySet() asSortedCollection()
Enumerating	collect() detect() reject() select()
Removing	remove() removeAll() removeIdentical()
Searching	includes() occurrencesOf()
Sorting	sortAscending() sortDescending() sortWith()

Protocol Examples

The example below (shown in the example “Using GbjObject.perform With an Argument List” on page 3-5) uses GbjCollection protocol to add a new instance of Engineer.

Example 3.13 Using GbjCollection.add()

```
/* Create the object in the server. */
GbjObject engClass = mySession.objectNamed("Engineer");
Object args[] = {"johnd", "John", "Doe", "(555) 123-4567"};
GbjObject newEngStub =
   engClass.perform("email:firstName:lastName:phone:", args, 4);

/* Get stub for server collection and add object. */
GbjCollection allEngineers = (GbjCollection)
   engClass.sendMsg("allEngineers");
allEngineers.add(newEngStub);
```

Some methods in GbjCollection require arguments that implicitly involve messages to server objects. Selection blocks are one example; for instance, class Engineer on the server implements the messages firstName and lastName. In the following code, the block argument to detect: sends the message "firstName" to each element of the collection until it obtains a match, then removes that element.

Example 3.14 Sending Method Arguments to the Server

```
GbjCollection allEng = (GbjCollection)
   mySession.doit("Engineer allEngineers");
GbjObject anEngr = allEng.detect("[:eng | eng firstName = 'John']");
allEng.remove(anEngr);
System.out.println("Removed John");
```

Serializing the Collection in the Server

The recommended approach to copying small collections to the client is to traverse them in the server, serializing the objects into a stream that GemBuilder for Java can transmit to the client in the minimum number of request round trips. For instance, this part could be encapsulated in a Smalltalk method in the server:

Example 3.15 Serializing a Collection

```
serializeCustomers
^allCustomers inject: Array new into: [:array :each |
    array add: each;
          add: (each firstName); add: (each lastName);...;
    yourself ]
```

The resulting Array contains the series *customer1*, *firstName1*, *lastName1*, ..., *customer2*, *firstName2*, *lastName2*, and so forth. Notice that the first array element for each customer is the customer object itself; in the client, this element will become a stub representing that customer. Apart from the element stubs, each field should be one of the simple data types that can be stored in a stub's cachedValue field.

Another approach (used in Customer) is to formulate a String to be sent as an ad-hoc message to the current session. In this example, the message causes the server to sort the customers and return two elements for each, the customer's name (by concatenating firstName and lastName) and the instance itself. Because most of the work is performed in the server, the method that sends the String is a doit().

Example 3.16 Serializing the AllCustomers Collection

```
static String execString =
    "  | custs custsAndNames | " +
    "  custs := Customer allCustomers asSortedCollection: [:a :b | " +
    "     (a lastName < b lastName) or: " +
    "     [(a lastName = b lastName) and: [a firstName < b firstName]]] ." +
    "  custsAndNames := OrderedCollection new." +
    "  custs do: [:cust | " +
    "      custsAndNames add: " +
    "          (cust firstName + Character space + cust lastName)." +
    "      custsAndNames add: cust. ]." +
    "  custsAndNames ";

   /**
    * Return a stub representing the result of performing the
    * above String in the server.
    */
    public static GbjCollection allCustomers(GbjSession sess) {
        return (GbjCollection) sess.doit(execString);
    }
```

Unpacking the Collection in the Client

Once the fields have been serialized in the server, GbjCollection provides two ways to access them from the client:

- The method `elements()` returns a Java Enumeration of GbjObjects (java.util.Enumeration).
- The method `allElements()` returns an array of GbjObjects.

These approaches differ primarily in the interface they provide; the underlying transport mechanism is much the same except that `elements()` caches a buffer of elements at a time, while `allElements()` retrieves all remaining elements in one trip.

To Enumerate the Collection

The method `elements()` in GbjCollection creates an Enumeration that performs efficient caching of the server object's contents for minimal client-server traffic during the enumeration operation.

In our example, the method `allCustomers()` in class CustomerHolder creates a holder object for each customer, unpacking into it the customer's name and the stub for the server element. Each holder becomes an element of a Vector that is returned. The call to Customer.allCustomers() performs the serialization on the server side, as shown previously. Customer has been registered as a stub class.

Example 3.17 Enumerating the Serialized AllCustomers

```
    public static Vector allCustomers(GbjSession sess) {
       GbjCollection allCustsStub;
       Customer custStub;
       Vector allCustomers = new Vector();
       String name;

       allCustsStub = Customer.allCustomers(sess);
       for (Enumeration e = allCustsStub.elements(); e.hasMoreElements();)
{
           name = ((GbjObject) e.nextElement()).stringValue();
           custStub = (Customer) e.nextElement();
           allCustomers.addElement(new CustomerHolder(name, custStub));
       }
       return allCustomers;
    }
```

To Unpack the Collection from an Array

The method `allElements()` in GbjCollection returns an array of GbjObject stubs. The following would iterate over such an array of customer fields, creating a CustomerHolder for each pair of elements and placing the holders in a Vector. The result would be much the same as using the Enumeration interface. Again, the call to `Customer.allCustomers()` performs the serialization on the server side, as shown previously.

Example 3.18 Unpacking AllCustomers from an Array

```
public static Vector allCustomers(GbjSession sess) {
       GbjCollection allCustsStub;
       Customer custStub;
       GbjObject fields[];
       Vector allCustomers = new Vector();
       allCustsStub = Customer.allCustomers(sess);
       fields = allCustsStub.allElements();
       for (int i=0; i<fields.length; i+=2) {
           allCustomers.addElement( new CustomerHolder(
               fields[i].stringValue(), (Customer) fields[i+1]));
       }
       return allCustomers;
   }
```

Obtaining Application-specific Stubs

You can supplement the GbjObject and GbjCollection classes provided in this release by creating your own Java classes and registering them as stubs representing specific GemStone/S classes. Two mechanisms are provided so you can register the stubs at initialization time (called *static stubs*) or at runtime (called *session-specific stubs*). Because these mechanisms build a correspondence between a Java class and a GemStone/S class, they permit you to code your Java client in a more natural way.

Registering Stubs at Static Initialization

During static initialization, you can register a class correspondence by providing an instance of the client (stub) class and the name of the server class. All sessions will have access to these stubs. In this release, the client class must be a subclass of GbjObject as shown below.

Example 3.19 Registering a Static Stub

```
public class Engineer extends GbjObject {
    // other useful code
    static {
        //during static initialization, register class as a stub
        GbjSession.registerStaticStub(new Engineer(),
            "Engineer");
    }
}
```

When a session connects to the server, GemBuilder for Java creates a registry specific to that session to map instances of Engineer and its subclasses in the server to the client stub class, using a hierarchy returned by the GemBuilder for Java server component. The server class must be in the session's symbol list.

GemBuilder for Java considers both the server class's classHistory and its class hierarchy in deciding which stub class to use in representing the server object. For instance, setting the stub class for Engineer also sets the stub class for any version of class Engineer as well as for any subclasses of Engineer and their various versions.

Because the stub registry is based on the inheritance hierarchy returned at connect time, it is reinitialized if the session is disconnected and then reconnected. Clients that modify the class hierarchy at runtime do not see the stub mapping change until they disconnect and reconnect the session.

Registering Stubs at Runtime

You can register a stub at runtime, but the stub will exist only in sessions where it is registered explicitly. To ensure that the mappings are handled correctly, they should be registered just after connection to the server. The registry cannot be created or modified while the session is in an unconnected state.

Two method signatures are available:

- registerStub(GbjObject, String) is similar to registerStaticStub() in that it takes the name of the server class as String.
- registerStub(GbjObject, GbjObject) takes a stub reference to the server class, which allows the mapping to a server class that is not in the session's symbol list. For instance, the following obtains the class stub by sending the message remoteClass() to an instance of that class.

```
GbjObject objStub, classStub;
objStub = mySession.objectNamed("someObject");
classStub = objStub.remoteClass();
mySession.registerStub((new SomeClass(), classStub);
```

Putting Client Data into the Server

There are several ways by which your client can place data in the GemStone/S server, all involving explicit action:

- Sending a message to a stub is the typical way to add an object to an existing Collection or update an anonymous instance variable. GbjObject and GbjCollection provide a number of methods for this purpose, such as atPut() and add(). If the existing object is persistent (that is, part of the committed repository), the addition also becomes persistent when the transaction is committed.
- The method putInServer() in class GbjSession places its argument in the server and returns a GbjObject representing it. This action by itself does not make the object persistent. To make it persistent, you must name it or make it reachable from a named object.
- Ad-hoc GemStone Smalltalk code, such as that executed by GbjSession.doit(), returns an object in the server, which GemBuilder for Java represents by a GbjObject. Again, this action by itself does not make the object persistent. To make it persistent, you must name it or make it reachable from a named object.

GemBuilder for Java maintains a Saved Objects set in which it tracks non-persistent objects that have been exported by the server to the client. The purpose of this set is to keep the objects from being garbage collected during the lifetime of the session. (It is analogous to the Export Set of certain other GemBuilder products.). GbjSession provides methods for examining and modifying this set, and GbjObject provides convenience methods for adding or removing a particular object.

Chapter

4

Forwarding Server Messages to Client Objects

Overview

A typical server message scenario begins with a message from a client object to a server object, which GemBuilder for Java translates into a Smalltalk message. If the client message results in a server message back to the client, GemBuilder for Java (GBJ) must translate that Smalltalk message into an appropriate Java message and see that the message reaches the proper receiver.

Figure 4.1 Sending a Server Message to a Client Object

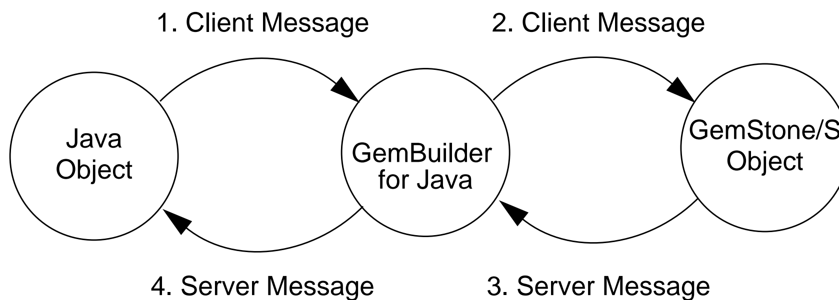

To do so, GemBuilder for Java uses the Java 1.3 Reflection API to find a matching Java method.

Alternatively, messages from server objects to client objects can be processed by a Java object that implements the GbjClientAdapter interface. The single method constituting this interface, dispatchGemStoneMessage(), translates the GemStone/S message into a Java message and sends it.

Whichever mechanism you choose, GemBuilder for Java provides default mappings for certain basic methods:

GemStone/S Selector	Java Method
=	equals()
==	==
hashCode	hashCode()
getClass	getClass()
toString	toString()

If GBJ cannot find a method, it first checks to see if the selector matches one of these. It throws an exception only if the selector does not.

Using Reflection

Reflection is used to make Java method names from Smalltalk message selectors as follows:

- Smalltalk unary and binary message selectors are used without alteration as the Java method name.
- Smalltalk keyword selectors drop all but the first keyword; this, without the colon, is used as the Java method name.

With the resulting method name, GemBuilder for Java makes two tries at method lookup. The first uses object wrapper classes for unmarshaled arguments, such as Java class Double for an instance of GemStone Float. If this fails, the second try uses the Java primitive data type, such as double for an instance of GemStone Float. Arguments that cannot be converted to Java types are passed as instances of GbjObject.

To preserve the integrity of Java objects and to avoid illegal execution, GemBuilder for Java uses the public reflection interface. Therefore, method invocation is restricted to those methods published in the public Java reflection mechanism.

Implementing the Adapter Interface

If reflection is not appropriate for your application, two other mechanisms provide for messages coming from the server:

- The receiving object in the client can implement the GbjClientAdapter interface itself, thereby serving as its own adapter.
- The client can register another class as an adapter, in which case the adapter class must implement GbjClientAdapter and must know to how to forward an appropriate message to the receiving object.

The GbjClientAdapter interface requires one method to be implemented, dispatchGemStoneMessage().

The client adapter receives the following Java objects from GemBuilder for Java:

- the session (a GbjSession)
- the receiver (an Object)
- the method selector (a String)
- the arguments (a Vector)

Registering a Client Adapter

Your client can register an adapter during static initialization or at runtime:

- The static form establishes the adapter for all sessions.

```
Adapter engAdapter = new Adapter();
GbjmySession.registerStaticAdapter("Engineer",
  engAdapter);
```

- The runtime form is effective only for the particular session in which it is invoked.

```
Adapter engAdapter = new Adapter();
mySession.registerAdapter("Engineer", engAdapter);
```

Dealing with Multithreading

Each message received from the server for local processing causes GemBuilder for Java to create a separate thread in which to handle it. This action is necessary because the client thread that sent a message to GemStone/S is blocked waiting for a response.

CAUTION:
The creation of a separate thread to handle each server message makes applications that use client adapters inherently multithreaded. You must exercise appropriate caution to process the message in the correct context.

Because GemStone/S 6.0 is not multithreaded itself, GemBuilder for Java limits access to the server from application threads so only one application thread is allowed into the communications layer at a time. Adapter threads are allowed to recursively send messages back to the server and are not blocked. Application threads are put to sleep while a message to the server from another application thread is being processed.

Message-sends in the Server

All server messages intended for a client object must be directed to a particular *client forwarder*, an instance of the private GemStone/S class GbjForwarder. In a typical scenario, the client forwarder is created from the client message that initiates the sequence.

For example, suppose much of the user interface knowledge resides in the server instead of being confined to the client. The client dialog class might implement something like the following to add an element to AllEngineers:

Example 4.1 Notifying Client from Server

```
String firstName, lastName, email, phone;
GbjCollection EngineersStub;
GbjObject result;
// code here to get data from text fields.

// now, add entry in server, close dialog if true returned:
try {
    result = EngineersStub.sendMsg("addEngrNotifying:", this,
        "firstName:", firstName, "lastName:", lastName, "email:",
        email);
    if (result.booleanValue()) {
        this.close();
    }
catch (GbjException e) {
    // handle exception
}
```

When the referent of “this” (the object to be notified) is unmarshaled in the server, the object will be represented by an instance of GbjForwarder because it is neither a stub object nor one of the simple data types for which value is cached. If the server code needs to return an error message to the client, it directs the message to this GbjForwarder. The server method might be implemented like this:

Example 4.2 Sending Notification from the Server

```
addEngrNotifying: dialog firstName: first lastName: last email: email
    | msg |
    msg := self validateFields: #[first, last, email].
    msg notNil ifTrue: [
        dialog displayError: msg.  "requires a client adapter"
        ^false ].
    [ System beginTransaction .
        "add engineer to AllEngineers class variable"
      System commitTransaction
     ] whileFalse .
    ^true
```

If the server detects invalid input, it forwards a message to the client.

If you're using reflection, you must write a method named displayError() that takes an argument of type String. GemBuilder for Java will find this method using reflection and invoke it. Here's how you can code displayError():

Example 4.3 Handling the Message Using Reflection

```
public void displayError(String msg) {
        System.err.println(msg);
}
```

If you have registered an adapter, the client object (or an instance of the registered adapter class) will receive dispatchGemStoneMessage() with the final arguments being the String "displayError:" and an array containing the contents of the server temporary variable "msg". Here's how you can code dispatchGemStoneMessage() to receive the message from the server:

Example 4.4 Handling the Message Using a Registered Adapter

```
public Object dispatchGemStoneMessage(GbjSession aSession, Object receiver,
String selector, Object args[]) {
    if (selector.equals("displayError:")) {
         this.displayError(args[0]);
    }
}
```

The GbjSession instance keeps an export set of objects represented by GbjForwarders to ensure they are not garbage-collected. GbjSession includes methods to examine and modify the export set; see, for instance, exportedReferences().

Exceptions Raised in the Client

Sending a message from the server to the client leads to the possibility of an exception being raised in the client. Such an exception will be returned to the server, where it may be transformed yet again into a Java GbjException that GemBuilder for Java throws to the client.

The following figure extends one used previously (in "Overview" on page 4-1). The sequence continues in the lower grouping, beginning with step 5 in which the Java object throws an exception in response to the server message. GemBuilder for Java catches this exception and (6) transmits it to the server. The server transforms

the exception information and (7) signals an error, #RT_ERR_CLIENT_FWD. The arguments are the detail text of the Java exception, and the Java stack. Unless the signal is handled in the server, the error information is sent back to the client (8), where GemBuilder for Java transforms it into a GbjException and (9) throws it to the application thread that started the sequence.

Figure 4.2 How a Client Exception Propagates to Calling Thread

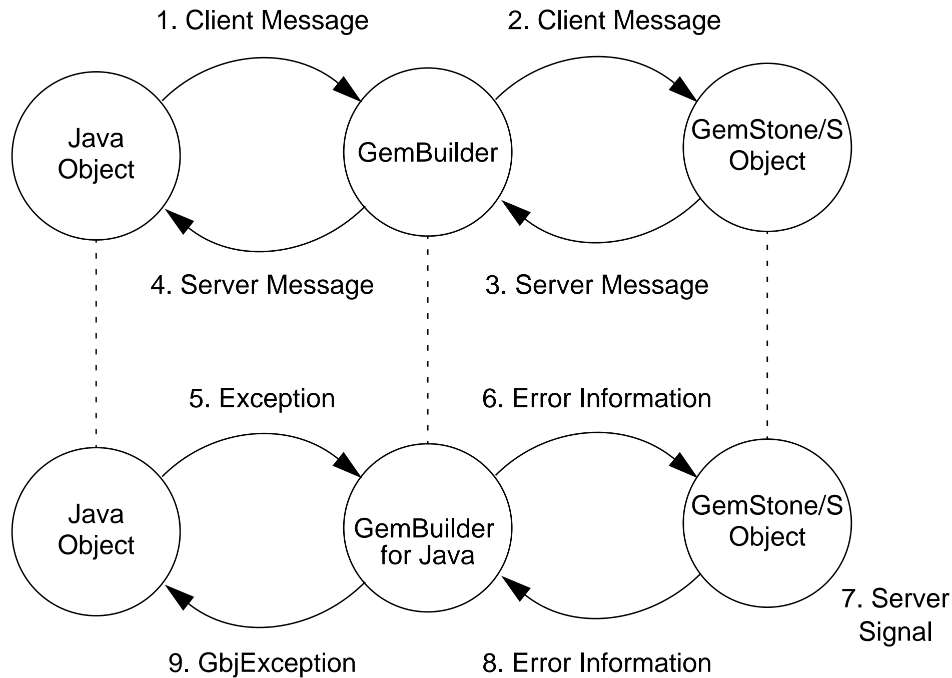

You can install an exception handler in the server with code like this:

Example 4.5 Installing an Exception Handler

```
Exception
    category: GbjSignals
    number: clientForwarderError
    do: [ :ex :cat :num :args |
        ex remove.
        ^0 ].
```

Chapter 5

Managing Server Transactions

Overview

The GemStone server provides an environment in which many users can share the same persistent objects. The server maintains a central repository of shared objects. When a GemBuilder for Java (GBJ) application needs to view or modify shared objects, it logs in to the GemStone server, starting a session as described in "Opening a Session" on page 2-2.

A GemBuilder for Java session creates a private view of the GemStone/S repository containing views of shared objects for the application's use. The application can perform computations, retrieve objects, and modify objects, as though it were working with private objects. When appropriate, the application propagates its changes to the shared repository so those changes become visible to other users.

In order to maintain consistency in the repository, GemBuilder for Java encapsulates a session's operations (computations, fetches, and modifications) in units called *transactions*. Any work done while operating in a transaction can be submitted to the server for incorporation into the shared object repository. This is called committing the transaction.

During the course of a logged-in session, an application can submit many transactions to the GemStone/S server. In a multiuser environment, concurrency conflicts will arise that can cause some commit attempts to fail. *Aborting* (rolling back) the transaction refreshes the session's view of the repository in preparation for further work.

In order to reduce operating overhead, a session by default runs outside a transaction, thereby temporarily relinquishing its ability to commit. A session running outside a transaction, called *manual transaction mode*, must explicitly begin a transaction before making changes that it will commit. Manual transaction mode is the default for GemBuilder for Java sessions, although the Development Tools override this default during login by placing the session in *automatic transaction mode*.

GemBuilder for Java provides ways of avoiding the concurrency conflicts that can cause a commit to fail. *Optimistic concurrency control* risks higher rates of commit failure in exchange for reduced transaction overhead and higher concurrency, while *pessimistic concurrency control* uses locks of various kinds to improve a transaction's chances of successfully committing. GBJ also offers *reduced-conflict classes* that are similar to familiar Smalltalk collections, but are especially designed for the demands of multiuser applications.

This chapter explains each of the topics mentioned here: transactions, committing and aborting, running outside a transaction, automatic and manual transaction modes, optimistic and pessimistic concurrency control, and reduced-conflict classes. Be sure to refer to the related topics in the *GemStone/S Programming Guide* for a full understanding of these transaction management concepts.

Separate topics explain each of the concepts mentioned here: transactions, committing and aborting, running outside a transaction, automatic and manual transaction modes, optimistic and pessimistic concurrency control, and reduced-conflict classes. Be sure to refer to the related topics in the *GemStone/S Programming Guide* for a full understanding of these transaction management concepts.

Operating Inside a Transaction

While a session is logged in to the GemStone/S server, GemBuilder for Java maintains a private view of the shared object repository for that session in the Gem. To prevent conflicts that can arise from operations occurring simultaneously in different sessions in the multi-user environment, GBJ encapsulates each session's operations in a transaction. Only when the session initiates a commit of its transaction does GemStone/S try to merge the modified objects in that session's view with the main, shared repository.

This figure shows a client and its repository, along with a common sequence of operations: (1) accessing an object from the shared repository, (2) creating an explicit replicate in the client, (3) modifying an object in the private view maintained in the Gem, and (4) committing the object's changes to the shared repository.

Figure 5.1 GemStone/S Application Workspace

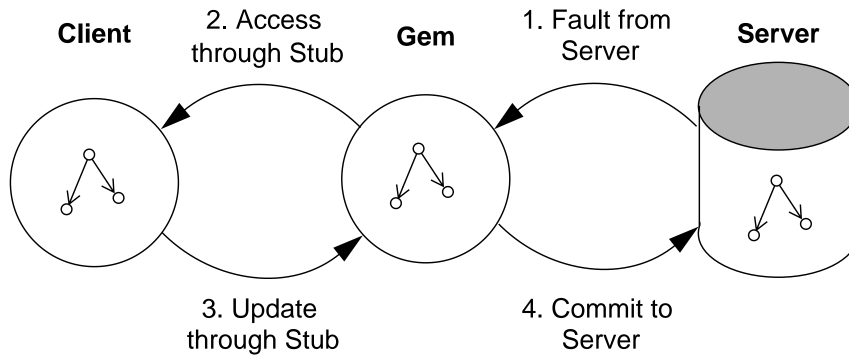

The private GemStone/S view starts each transaction as a snapshot of the current state of the repository. As the client creates and modifies shared objects, GemBuilder for Java updates the private GemStone/S view to reflect the client's changes. When your client commits a transaction, the repository is updated with the changes held in your client's private GemStone/S view.

For efficiency, GBJ does not duplicate the entire contents of the server into the Gem. GemBuilder for Java contains only those objects that have been accessed from the server or created by your client for sharing with the server. This action minimizes the amount of data that moves across the boundary from the server to the Gem.

CAUTION:

Because GemBuilder for Java does not update objects in the Java client at transaction boundaries, information copied into the client may no longer reflect the current state of the repository. Whenever you commit or abort a transaction, you should reinitialize copies of server objects to their current state in the Gem and shared repository. Alternatively, object change notification can be used for this purpose.

Committing a Transaction

When a client submits a transaction to the GemStone/S server for inclusion in the shared repository, it is said to commit the transaction. To commit a transaction, send the message:

```
aGbjSession.commitTransaction()
```

or, in the GemStone Browser, choose **Session > Commit**.

When the commit succeeds, the method returns true. Successfully committing a transaction has two effects:

- It copies the client's new and changed objects to the shared object repository, where they are visible to other users.
- It refreshes the client's private GemStone/S view by making visible any new or modified objects that have been committed by other users. *You should reinitialize objects you have copied to the client.*

A commit request fails if the server detects a concurrency conflict with the work of other users. When the commit fails, the commitTransaction() method returns false.

In order to commit, the session must be operating within a transaction. An attempt to commit while outside a transaction raises an exception.

Aborting a Transaction

A session refreshes its view of the shared object repository by aborting its transaction. Despite the terminology, a session need not be operating inside a transaction in order to abort. To abort, send the message:

```
aGbjSession.abortTransaction()
```

or, in the GemStone Browser, choose **Session > Abort**.

Aborting has these effects:

- The transaction (if any) ends. If the session's transaction mode is automatic, GemBuilder for Java starts a new transaction. If the session's transaction mode is manual (the default), the session is left outside of a transaction.
- Temporary Smalltalk objects remain unchanged.
- The session's private view of the GemStone/S shared object repository is updated to match the current state of the repository. *You should reinitialize objects you have copied to the client.*

Handling Commit Failures

If an attempt to commit fails because of a concurrency conflict, the commitTransaction() method returns false.

Following a commit failure, your session's private view of persistent objects may differ from its pre-commit state:

- The current transaction is still in effect. Nevertheless, you must end the transaction and start a new one before performing computations and before you can successfully commit.
- Temporary Smalltalk objects remain unchanged.
- Modified GemStone/S objects remain unchanged.
- Unmodified GemStone/S objects are updated with new values from the shared repository. *You should reinitialize objects you have copied to the client.*

Following a commit failure, your session **must** refresh its private view of the repository by aborting the current transaction. The uncommitted transaction remains in effect so you can save some of its contents, if necessary, before aborting.

A common strategy for handling such a failure is to abort, then reinvoke the method in which the commit occurred. Depending on your application, you may simply choose to discard the transaction and move on, or you may choose to remedy the specific transaction conflict that caused the failure, then initiate a new transaction and commit.

If you want to know why a transaction failed to commit, you can send the message:

```
aGbjSession.doit("System transactionConflicts")
```

This expression returns a stub representing a symbol dictionary whose keys indicate the kind of conflict detected and whose values identify the objects that incurred each kind of conflict. (See “Managing Concurrent Transactions” on page 5-11 for more discussion of the kinds of conflicts that can arise.)

Operating Outside a Transaction

Operating *inside* of a transaction involves a certain amount of overhead because GemBuilder for Java monitors the operations that occur, gathering all the necessary information required to prepare the transaction to be committed.

Operating *outside* of a transaction, however, saves some of the overhead of tracking changes, which may be significant in some applications. The session can view the repository, browse the objects it contains, and even make computations based upon their values, but it cannot commit any new or changed GemStone/S objects. A session operating outside a transaction can, at any time, begin a transaction.

No session is overhead-free: even a session operating outside a transaction uses GemStone/S resources to manage its objects and its view of the repository. For best system performance, all sessions, even those running outside a transaction, must periodically refresh their views of the repository by committing or aborting.

These methods in GbjSession support running outside of a transaction:

Method	Description
beginTransaction()	Aborts and begins a transaction
isAutomaticTransactionModeSet()	Returns true if the server is in automatic transaction mode and false if not
setTransactionMode("autoBegin")	Sets the transaction mode to automatic (chained transactions, with an implicit begin after a commit or abort)
setTransactionMode("manualBegin")	Sets the transaction mode to manual (explicit *beginTransaction* required)
setTransactionMode("transactionless")	Sets the transaction mode to transactionless (aborting the current transaction, if any)

To begin a transaction, send the message:

```
aGbjSession.beginTransaction()
```

or, in the GemStone Browser, choose **Session > Begin**.

This message gives you a fresh view of the repository and starts a transaction. When you abort or successfully commit this new transaction, you will again be outside of a transaction until you either explicitly begin a new one or change transaction modes.

If you are not currently in a transaction, but still want a fresh view of the repository, you can send the message aGbjSession.abortTransaction(). This message aborts your current view of the repository and gives you a fresh view, but does not start a new transaction.

Being Signaled to Abort

When you are in a transaction, GemStone/S waits until you commit or abort to reclaim storage for objects that have been made obsolete by your changes. When you are running outside of a transaction, however, you are implicitly giving GemStone/S permission to send your Gem session a signal requesting that you abort your current view so that GemStone/S can reclaim storage when necessary. When this happens, you must respond within the time period specified in GemStone/S's STN_GEM_ABORT_TIMEOUT configuration parameter. If you do not, GemStone/S forces an abort and sends your session an abortErrLostOtRoot signal (GbjGemStoneErrors.ABORT_ERR_LOST_OT_ROOT), which means that your view of the repository was lost, and any objects your client had copied may no longer be valid. When your client receives abortErrLostOtRoot, the session has been aborted; your client must reinitialize all of its data from the Gem to reflect the current state of the GemStone/S repository.

You can detect an abortErrLostOtRoot signal and control what happens when you receive a signal to abort by implementing GbjObserver.update().

For example

Example 5.1

```
public void update(GbjObservable aSession, String aString, Object arg)
{
    if (aString.equals("event")) {
        GbjException e = (GbjException) arg;
        if (e.number == GbjException.kernel.RT_ERR_SIGNAL_ABORT) {
            aSession.abortTransaction();
            // refetch objects from Gem
            try {
                //re-enable generation of the error
                aSession.doit("System enableSignaledAbortError");
            }
            catch (GbjException ex) {
                System.out.println(ex.getMessage());
            }
        }
        if (e.number == GbjException.kernel.ABORT_ERR_LOST_OT_ROOT) {
            // refetch objects from Gem
        }
    } else {
        System.out.println(aString + " " + arg);
    }
}
```

This causes your GemBuilder for Java session to abort when it receives a signal to abort, and then to reinitialize copies of server objects in the client. If an abortErrLostOtRoot signal is received, the client detects it and reinitializes its copies.

Transaction Modes

A GemBuilder for Java session, by default, always is outside a transaction when it logs in. After logging in, the session can operate in either of three transaction modes: manual, automatic, or transactionless.

Manual Transaction Mode

In manual transaction mode, the session remains outside a transaction until you begin a transaction. This is the default mode in GemBuilder for Java. In manual transaction mode, a transaction begins only as a result of an explicit request. When you abort or commit successfully, the session remains outside a transaction until a new transaction is initiated.

A new transaction always begins with an abort to refresh the session's private view of the repository. Local objects that customarily survive an abort operation, such as temporary results you have computed while outside a transaction, can be carried into the new transaction where they can be committed, subject to the usual constraints of conflict-checking. If you begin a new transaction while already inside a transaction, the effect is the same as an abort.

In manual transaction mode, as in automatic mode, an unsuccessful commit leaves the session in the current transaction until you take steps to end the transaction by aborting.

Automatic Transaction Mode

In automatic transaction mode, committing or aborting a transaction automatically starts a new transaction. In this mode, the session operates within a transaction the entire time it is logged into GemStone/S. To run this way, a session must switch to automatic transaction mode or specify that mode in the login parameters.

Being in a transaction incurs certain costs related to maintaining a consistent view of the repository at all times for all sessions. Objects that the repository contained when you started the transaction are preserved in your view, even if you are not using them and other users' actions have rendered them meaningless or obsolete. For this reason, lengthy transactions can impede garbage collection of objects in the repository that are otherwise unneeded.

Depending upon the characteristics of your particular installation (such as the number of users, the frequency of transactions, and the extent of object sharing), this burden can be trivial or significant. If it is significant at your site, you may want to reduce overhead by using sessions that run outside transactions, which is the default mode in GemBuilder for Java.

Transactionless Mode

In transactionless mode, the session remains outside a transaction. If all you need to do is browse the repository or make changes to objects in the client only, transactionless mode can be a more efficient use of system resources, because GemBuilder for Java does not need a commit page, nor will it elicit garbage collection.

Starting transactionless mode always causes an abort to refresh the session's private view of the repository.

Choosing Which Mode to Use

Use manual transaction mode if the work you are doing requires looking at objects in the repository, but only seldom requires committing changes to the repository. You will have to start a transaction manually before you can commit your changes to the repository, but the system will be able to run with less overhead.

Use automatic transaction mode if the work you are doing requires committing to the repository frequently, because you can make permanent changes to the repository only when you are in a transaction.

Use transactionless mode if the work you are doing requires looking at objects in the repository only.

Switching Between Modes

To find out if you are currently in a transaction, execute

```
aGbjSession.inTransaction()
```

This returns true if you are in a transaction and false if you are not.

To change to automatic transaction mode, execute the expression:

```
aGbjSession.setTransactionMode("autoBegin")
```

This message automatically aborts the transaction, if any, changes the transaction mode, and starts a new transaction.

To change to manual transaction mode, execute the expression:

```
aGbjSession.setTransactionMode("manualBegin")
```

This message automatically aborts the current transaction and changes the transaction mode to manual. It does not start a new transaction, but it does provide a fresh view of the repository.

To change to transactionless mode, execute the expression:

```
aGbjSession.setTransactionMode("transactionless")
```

This message automatically aborts the current transaction, if any, changes the mode to transactionless, and provides a fresh view of the repository.

Managing Concurrent Transactions

When you tell GemStone/S to commit your transaction, it checks to see if doing so presents a conflict with the activities of any other users.

- It checks to see whether other concurrent sessions have committed transactions of their own, modifying an object that you have also modified during your transaction. If they have, then the resulting modified objects can be inconsistent with each other.
- It checks to see whether other concurrent sessions have committed transactions of their own, modifying an object that you have read during your transaction, while at the same time you have modified an object that the other session has read.
- It checks for locks set by other sessions that indicate the intention to modify objects that you have read or to read objects you have modified in your view.

If it finds no such conflicts, GemStone/S commits the transaction, and your work becomes part of the permanent, shared repository. Your view of the repository is refreshed and any new or modified objects that other users have recently committed become visible in any dictionaries that you share with them.

Read and Write Operations

It is customary to consider the operations that take place within a transaction as *reading* or *writing* objects. Any operation that accesses any instance variable of an object *reads* that object, as do operations that fetch an object's size, class, or other descriptive information about that object. An object also is read in the process of being stored into another object.

An operation that stores a value in one of an object's instance variables *writes* the object. While you can read without writing, writing an object always implies reading it, because GemStone/S must read the internal state of an object in order to store a value in it.

In order to detect conflict among concurrent users, GemStone/S maintains two logical sets for each session: a set containing objects read during a transaction and a set containing objects written. These sets are called the *read set* and the *write set*. Because writing implies reading, the read set is always a superset of the write set.

The following conditions signal a possible concurrency conflict:

- An object in your write set is also in another transaction's write set (a write/write conflict).
- An object in your write set is in another transaction's read set *and* an object in your read set is in that transaction's write set (a read/write conflict).

Optimistic and Pessimistic Concurrency Control

GemStone/S provides two approaches to managing concurrent transactions: optimistic and pessimistic. An application can use either or both approaches, as needed.

Optimistic concurrency control means that you simply read and write objects as if you were the only session, letting GemStone/S detect conflicts with other sessions only when you try to commit a transaction.

Pessimistic concurrency control means that you act as early as possible to prevent conflicts by explicitly requesting locks on objects before you modify them. When an object is locked, other users may be unable to lock the object or commit changes to it.

Optimistic concurrency control is easy to implement in an application, but you run the risk of having to re-do the work you've done if conflicts are detected and you're unable to commit. When GemStone/S looks for conflicts only at commit time, your chances of being in conflict with other users increase with the time between commits and the size of your read and write sets. Under optimistic concurrency control, GemStone/S detects conflict by comparing your read and write sets with those of all other transactions committed since your transaction began.

Running under optimistic concurrency control is the most convenient and efficient mode of operation when:

- you are not sharing data with other sessions, or
- you are reading data but not writing, or
- you are writing a limited amount of shared data and you can tolerate not being able to commit your work sometimes.

If you take a pessimistic approach, you act as early as possible to prevent conflicts by explicitly requesting locks on objects before you modify them. When an object is locked, other people are unable to lock the object, and they cannot optimistically commit changes to the object. Also, when you encounter an object that someone else has locked, you can abort the transaction immediately instead of wasting time on work that can't be committed.

Locking improves one user's chances of committing, but at the expense of other users, so you should use locks sparingly to prevent an overall degradation of system performance. Still, if there is a lot of competition for shared objects in your application, or if you can't tolerate even an occasional inability to commit, then using locks might be your best choice.

Locks do not prevent read-only access to objects, so read-only query transactions are not affected by modification transactions.

Setting the Concurrency Mode

Any shared object that is not explicitly locked is treated optimistically. For objects under optimistic concurrency control, GemStone/S's level of checking for concurrency conflicts is configurable. You can set the level of checking for concurrency conflicts by specifying one of the following values for the CONCURRENCY_MODE configuration parameter in your application's configuration file. There are two levels:

- FULL_CHECKS (the default mode), which checks for both *write/write* and *read/write* conflicts. If either type of conflict is detected your transaction cannot commit.
- NO_RW_CHECKS, which performs *write/write* checking only.

Locking methods override the configured optimistic CONCURRENCY_MODE by stating explicitly the kind of pessimistic control they implement.

Setting Locks

GemBuilder for Java provides locking protocol that allows application developers to write client Java code to lock objects.

A GbjObject or one of its subclasses is the receiver of all lock requests. Locks can be requested on a single object or on a collection of objects.

Single lock requests are made with the following statements:

```
aGbjObject.readLock()
aGbjObject.writeLock()
aGbjObject.exclusiveLock()
```

The above messages request a particular type of lock on aGbjObject. If the lock is granted, the method returns the receiver. (Lock types are described in the *GemStone/S Programming Guide*.) If you don't have the proper authorization, or if another session already has a conflicting lock, an exception will be thrown.

When you request an exclusive lock, an exception will be thrown if another session has committed a change to aGbjObject since the beginning of the current transaction. In this case, the lock is granted despite the exception, but it is seen as "dirty." A session holding a dirty lock cannot commit its transaction, but must abort to obtain an up-to-date value for aGbjObject, then refetch its value through the stub. The lock will remain, however, after the transaction is aborted.

NOTE:
GemStone Smalltalk provides a number of locking methods in the server for which there is no corresponding implementation in GemBuilder for Java. For information, refer to the GemStone Kernel Reference and the GemStone/S Programming Guide.

Once you lock an object, it normally remains locked until you either log out or explicitly remove the lock; unless you specify otherwise, locks persist through aborts and commits. In general, you should remove a lock on an object when you have used the object, committed the resulting values to the repository, and no longer anticipate a need to maintain control of the object.

The following method removes a specific lock:

```
aGbjObject.removeLock()
```

To clear all locks for the session if the transaction is successfully committed, send this message:

```
aGbjSession.commitAndReleaseLocks()
```

Reduced-Conflict Classes

At times GemStone/S will perceive a conflict when two users are accessing the same object, when what the users are doing actually presents no problem. For example, GemStone/S may perceive a write/write conflict when two users are simultaneously trying to add an object to a Bag that they both have access to because this is seen as modifying the Bag.

GemStone/S provides some reduced-conflict classes that can be used instead of their regular counterparts in applications that might otherwise experience too many unnecessary conflicts. These classes include RcCounter, RcIdentityBag, RcKeyValueDictionary, and RcQueue.

Use of these classes can improve performance by allowing a greater number of transactions to commit successfully without locks, but they do carry some overhead.

For one thing, they use more storage than their ordinary counterparts. Also, you may find that your application takes longer to commit transactions when you use instances of these classes. Finally, you should be aware that under certain circumstances, instances of these classes can hide conflicts from you that you indeed need to know about.

Here are brief descriptions of the reduced-conflict classes. For details about these classes and their usage, see the *GemStone/S Programming Guide* and the *GemStone Kernel Reference*.

- RcCounter maintains an integral value that can be incremented or decremented. A single instance of RcCounter can be shared among multiple concurrent sessions without conflict.
- RcIdentityBag provides the same functionality as IdentityBag, except that no conflict occurs on instances of RcIdentityBag when a number of users read objects in the bag or add objects to the bag at the same time. Nor is there a conflict when one user removes an object from the bag while other users are adding objects, or when a number of users remove objects from the bag at the same time, so long as no more than one of them tries to remove the last occurrence of an object.
- RcKeyValueDictionary provides the same functionality as KeyValueDictionary except that no conflict occurs on instances of RcKeyValueDictionary when users read values in the dictionary or add keys and values to it (unless one tries to add a key that already exists) or when users remove keys from the dictionary at the same time (unless more than one user tries to remove the same key at the same time).

- Conflict occurs only when more than one user tries to modify or remove the same key from the dictionary at the same time.
- RcQueue represents a first-in-first-out (FIFO) queue. No conflict occurs on instances of RcQueue when multiple users read objects in or add objects to the queue at the same time, or when one user removes an object from the queue while other users are adding objects. However, if more than one user removes objects from the queue, they are likely to experience a write/write conflict.

Chapter **6**

Observing Session and Server Events

Overview

GemBuilder for Java (GBJ) provides two interfaces to support monitoring of events: GbjObserver and GbjObservable.

The GbjObserver interface consists of a single method, update(), through which objects are notified of events transpiring in other objects of interest. Each class that will be an observer must implement this interface. The action to be taken depends largely on whether the object being observed is a GemBuilder for Java session or an entity in the server.

The GbjObservable interface consists of methods that manipulate the list of observers to be notified and initiate the notification process. The class GbjSession implements this interface, and ordinarily it is the only class that needs to do so.

Observing Session Events

Your client can receive notification of significant session events by registering its interest with the object in which the action takes place. For instance, the Help Request Browser needs to be notified of a change committed through an editing dialog open on a particular help request so the browser can update its display.

Sessions report client events using the update() message, which has three parameters: *obj* (a GbjObservable), *notificationMessage* (a String), and *argument* (an Object). The second and third parameters map to events in the following way:

Event	notificationMessage	argument
Transaction commit	"commit"	null
Transaction abort	"abort"	null
Transaction begun, in manual transaction mode	"begin"	null
Informational message	"message"	message string
Close of session	"close"	null

Observers of session events are notified in the same thread that caused the event to take place. Notification is received only for events in the client session, not for similar actions taking place in Smalltalk code being executed in the server. That is, invoking `mySession.commitTransaction()` in the client results in notification to observers, but invoking `System | commitTransaction` in the server does not.

The informational message event is used internally for passing verbose (debugging) information about client and server activity (see “Logging of Debugging Information” on page 2-5). Client objects can explicitly trigger informational message events by invoking notifyObservers(), thereby passing information about events that otherwise would not be subject to notification. For example, an informational message event could be used programmatically to provide notification of a commit initiated in the server by means of doit().

To Monitor Session Events

1. Signify interest by sending addObserver() to the current session. For example:

   ```
   mySession.addObserver(this);
   ```

2. Implement GbjObserver to handle the notification. For example, when the session commits or aborts the current transaction, or begins a transaction in manual transaction mode, the Help Request Browser calls its method to reinitialize the display. If the session is closed, the browser closes itself.

Example 6.1 Handling Session Events

```
public void update(GbjObservable aSession, String message, Object argument)
{
    if ((message.equals("abort")) ||
        (message.equals("commit")) ||
        (message.equals("begin"))) {
             this.reinitialize();
    }
    if ((message.equals("close")) {
        this.sessionClosed();
    }
}
```

Observing Server Events

Clients can observe events in the server that are of general interest, such as object change notification and Gem to Gem signals. The process is much the same as for monitoring session events, except that the notification always is in a different thread from the one the application is using. (For information about monitoring session events, see “Observing Session Events” on page 6-1.)

For server events, the update() parameter *notificationMessage* always has the value "event", and *argument* is an instance of GbjException.

NOTE:
Because notification of server events arrives in a different thread, programmers must ensure actions the observer takes are thread-safe with respect to the application. FURTHER NOTE: In this release, notification of the event does not suspend execution in GemStone/S. As a result, the observer may have to wait until the currently executing requests finish.

Although the server events of interest are represented in GemStone/S as exceptions (as are errors) GemBuilder for Java distinguishes them, notifying observers of the event rather than throwing them as GbjExceptions.

Clients can identify event circumstances by comparing the number member of the exception instance with those defined in class GbjGemStoneErrors. This example obtains the GbjException, then checks whether it represents a Gem-to-Gem signal:

Example 6.2 Handling a Server Events

```
public void update(GbjObservable session, String message, Object arg) {
    if (message.equals("event")) {
        GbjException event = (GbjException)arg;
        if (event.number ==
            GbjException.kernel.RT_ERR_SIGNAL_GEMSTONE_SESSION) {
            // handle signal, then re-enable signal reception
            try {
                event.session.doit("System " +
                  "enableSignaledGemStoneSessionError");
            } catch (GbjException e) {
                System.out.println(e.getMessage());
            }
        }
    }
}
```

Chapter

7

Deploying Your Application

Overview

Clients using GemBuilder for Java (GBJ) can be deployed in two ways, as shown in the accompanying illustration:

- as applets that run in the context of a Java-enabled web browser
- as standalone Java applications

Figure 7.1 GemBuilder for Java Deployment

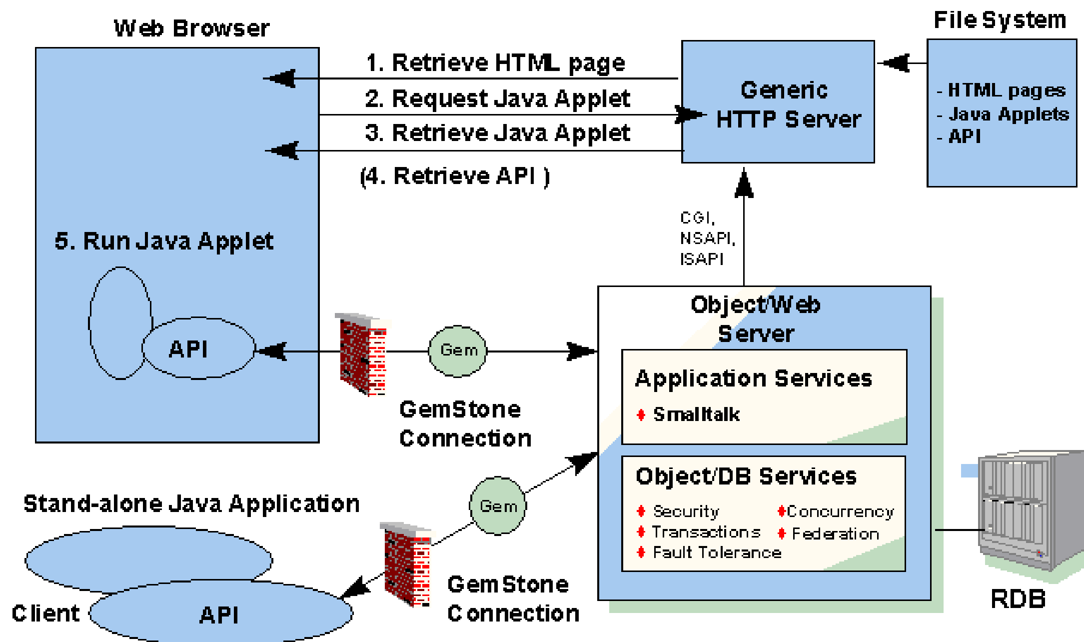

Deployment Steps

This topic explains the steps you should take in deploying your GemBuilder for Java client as a applet (to be run in a Web browser) and as a standalone application.

To Deploy an Applet

1. Install the HTML page containing the `<APPLET>` tag on the HTTP server, and install the applet class file in the location specified in the tag's `CODE` and `CODEBASE` attributes.
2. Install the GemBuilder for Java class library, `gbj20.jar`, either local to the browser or in the same location as the applet class file.
3. Start a Session Broker for the GemStone/S server. The broker can run on the server's machine or another server platform, but the broker and server machines must be on the same side of a firewall.

4. Unless they are coded in the applet, provide the server's name (such as gemserver60) and the broker's machine name and well-known port number (such as 9090) to the user. Also provide a GemStone/S userId (account name) and password.
5. Make sure a GemStone/S NetLDI is running on the broker's machine.

To Deploy a Standalone Application

1. Start a Session Broker for the GemStone/S server. The broker can run on the server's machine or another server platform, but the broker and server machines must be on the same side of a firewall.
2. Unless they are coded in the application, provide the server's name (such as gemserver60) and the broker's machine name and well-known port number (such as 9090) to the user. Also provide a GemStone/S userId (account name) and password.
3. Make sure a GemStone/S NetLDI is running on the Session Broker's machine.
4. Make sure the Java runtime system (typically `...\bin\java`) is in the user's path.
5. Install the GemBuilder for Java class library, `gbj20.jar`, where it will be accessible to the client, and add the library's location to the CLASSPATH in the manner required by your runtime Java system.

Glossary

Adapter

An object that adapts messages for another object, converting them to a form the receiving object can understand.

Applet

A Java program that is meant to be run in the context of a Java-compatible browser, rather than as a standalone program. Another kind of Java program, called an *application*, can run independently of a browser.

Enumeration

An object in Java that is used to iterate over the contents of a collection. It is similar to a Stream object in Smalltalk.

Firewall

A gatekeeper computer that protects a local network by filtering traffic to and from an external network, such as the Internet.

Gem

A GemStone server process that provides server access to clients. Currently, each session connects to the server through a dedicated Gem process.

Marshal

The process of serializing an object into a stream before sending it to another process.

NetLDI

A GemStone network server process, which provides information services and spawns other processes for GemStone clients.

Observer

A client object that receives notification of noteworthy events in a session or in the server.

Persistent Object

An object that is stored in the GemStone server and accessible through a root object.

Session Broker

A GemStone server object that supervises the connection of a Java client to the server by connecting the client to a Gem process. A broker is an instance of the GemStone class GbjBroker.

Simple Data Type

One of the following types, which can be stored in the cachedValue field of a GbjObject: Integer, Long, Float, Double, Character, Boolean, null, Date, String, gemstone.gbj.DoubleByteString, or a serialized collection of proxies.

Stone

A GemStone process that oversees other server processes; it is also known as the repository monitor.

Stub

An object that forwards messages to an object in another program, possibly on a different computer. In this release, GemBuilder for Java uses stubs to forward messages from client objects to objects in the GemStone server.

Index

A

abortErrLostOtRoot 5-7
abortTransaction 2-4, 5-4, 5-7

C

commitTransaction 2-4, 5-4, 6-2
concurrency mode 5-2, 5-4, 5-12
configuration file 2-8

E

Errors
  abortErrLostOtRoot 5-7

G

GbjBroker 2-2
GbjClientAdapter 4-2
GbjCollection 3-26
  example 3-13
GbjException 3-1, 3-7, 4-5, 5-8, 6-3
GbjExternalizable interface 3-9
GbjForwarder 4-4
GbjGemStoneErrors 3-7, 6-3
GbjGemStoneErrors.java 3-7
GbjObject 3-1
GbjObserver 5-7, 6-1
GbjParameters 2-2
GbjSession 2-1, 3-3, 6-1

H

holder 3-16, 3-23

J

java.lang.RuntimeException 3-5

L

log files
    locating 2-15
logging 2-5, 5-8

M

marshaling objects 3-8
    controlling details of 3-9

N

NetLDI 2-15

R

reduced-conflict classes 5-2, 5-15
reflection 4-2
    default mappings for basic methods 4-2
    displayError() 4-6
register 3-24, 4-3
Replicate 3-17, 5-3
RT_ERR_CLIENT_FWD 4-7

S

Session Broker 2-2, 2-6
    starting 2-11
    startup script 2-11

T

transaction
    running outside 5-5
transaction conflict 5-5
transactionless mode 5-5