

---

*GemStone*®

# **GemStone/S 64 Bit™ Programming Guide**

**Version 3.7**

September 2023



## INTELLECTUAL PROPERTY OWNERSHIP

This documentation is furnished for informational use only and is subject to change without notice. GemTalk Systems LLC assumes no responsibility or liability for any errors or inaccuracies that may appear in this documentation.

Warning: This computer program and its documentation are protected by copyright law and international treaties. Any unauthorized copying or distribution of this program, its documentation, or any portion of it, may result in severe civil and criminal penalties, and will be prosecuted under the maximum extent possible under the law.

The software installed in accordance with this documentation is copyrighted and licensed by GemTalk Systems 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.

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 GemTalk Systems.

This software is provided by GemTalk Systems LLC and contributors "as is" and any expressed or implied warranties, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose are disclaimed. In no event shall GemTalk Systems LLC or any contributors be liable for any direct, indirect, incidental, special, exemplary, or consequential damages (including, but not limited to, procurement of substitute goods or services; loss of use, data, or profits; or business interruption) however caused and on any theory of liability, whether in contract, strict liability, or tort (including negligence or otherwise) arising in any way out of the use of this software, even if advised of the possibility of such damage.

## COPYRIGHTS

This software product, its documentation, and its user interface © 1986-2023 GemTalk Systems LLC. All rights reserved by GemTalk Systems.

## PATENTS

GemStone software has been covered by U.S. Patent Number 6,256,637 "Transactional virtual machine architecture", Patent Number 6,360,219 "Object queues with concurrent updating", Patent Number 6,567,905 "Generational garbage collector with persistent object cache", and Patent Number 6,681,226 "Selective pessimistic locking for a concurrently updateable database".

## TRADEMARKS

**GemTalk**, **GemStone**, **GemBuilder**, **GemConnect**, and the GemStone and GemTalk logos are trademarks or registered trademarks of GemTalk Systems LLC, or of VMware, Inc., previously of GemStone Systems, Inc.

**UNIX** is a registered trademark of The Open Group.

**Intel** is a registered trademarks of Intel Corporation.

**Microsoft**, **Windows**, **Windows Server**, and **Azure** are registered trademarks of Microsoft Corporation.

**Linux** is a registered trademark of Linus Torvalds and others.

**Red Hat**, **Red Hat Enterprise Linux**, **RHEL**, and **CentOS** are trademarks or registered trademarks of Red Hat, Inc.

**AlmaLinux** is a trademark or registered trademark of AlmaLinux OS Foundation.

**Rocky Linux** is a trademark or registered trademark of Rocky Enterprise Software Foundation.

**Ubuntu** is a registered trademark of Canonical Ltd., Inc.

**AIX**, **Power**, **POWER**, **Power8**, **Power9**, and **VisualAge** are trademarks or registered trademarks of International Business Machines Corporation.

**Apple**, **Mac**, **macOS**, and **Macintosh** are trademarks of Apple Inc.

**Instantiations** is a registered trademarks of Instantiations, Inc.

**CINCOM**, **Cincom Smalltalk**, and **VisualWorks** are trademarks or registered trademarks of Cincom Systems, Inc.

**Raspberry Pi** is a trademark of the Raspberry Pi Foundation.

**RabbitMQ** is a trademark of VMware, Inc.

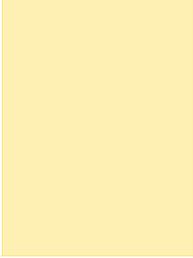
**Prometheus** is a registered trademark of The Linux Foundation.

**Grafana** is a registered trademark of Raintank, Inc. dba Grafana Labs.

Other company or product names mentioned herein may be trademarks or registered trademarks of their respective owners. Trademark specifications are subject to change without notice. GemTalk Systems cannot attest to the accuracy of all trademark information. Use of a term in this documentation should not be regarded as affecting the validity of any trademark or service mark.

**GemTalk Systems**  
15220 NW Greenbrier Parkway  
Suite 240  
Beaverton, OR 97006

---



# Preface

---

## About This Documentation

This manual describes the GemStone Smalltalk language and programming environment provided by the GemStone/S 64 Bit™ product, and how to use many of the features available in GemStone Smalltalk.

This manual is intended for users that are at least somewhat familiar with the Smalltalk programming language and with its programming environment. Appendix A includes an overview of the Smalltalk language syntax.

## Terminology Conventions

The term “GemStone” is used to refer to the server products GemStone/S 64 Bit and GemStone/S, and the GemStone family of products; the GemStone Smalltalk programming language; and may also be used to refer to the company, now GemTalk Systems, previously GemStone Systems, Inc. and a division of VMware, Inc.

## Typographical Conventions

This document uses the following typographical conventions:

- ▶ Smalltalk methods, GemStone environment variables, operating system file names and paths, listings, and prompts are shown in `monospace` typeface.
- ▶ Responses from GemStone commands are shown in an underlined typeface.
- ▶ Place holders that are meant to be replaced with real values are shown in *italic* typeface.
- ▶ Optional arguments and terms are enclosed in [square brackets].
- ▶ Alternative arguments and terms are separated by a vertical bar ( | ).

## Executing the Examples

This manual includes many examples. These can be executed using either the Topaz command-line interface, or using tools such as GemBuilder for Smalltalk (GBS) or another interface to the GemStone/S server. GBS or other IDE tools provide browsers and related tools that make it easier to define classes and methods.

The text of the GemStone Smalltalk code examples themselves (excluding the Topaz commands) is the same whichever way you enter it. The example results commonly but not always use topaz display conventions, and may be simplified for clarity.

When using Topaz, you must include extra commands to begin and end an example. An “%” is used to indicate the command is completed and executed, including in examples that omit a topaz start execution command. Refer to the *Topaz Users Guide* for more information on executing code in topaz.

## Technical Support

### Support Website

#### **[gemtalksystems.com](http://gemtalksystems.com)**

GemTalk’s website provides a variety of resources to help you use GemTalk products:

- ▶ **Documentation** for the current and for previous released versions of all GemTalk products, in PDF form.
- ▶ **Product download** for the current and selected recent versions of GemTalk software.
- ▶ **Bugnotes**, identifying performance issues or error conditions that you may encounter when using a GemTalk product.
- ▶ **Supplemental Documentation** and **TechTips**, providing information and instructions that are not in the regular documentation.
- ▶ **Compatibility matrices**, listing supported platforms for GemTalk product versions.

We recommend checking this site on a regular basis for the latest updates.

### Help Requests

GemTalk Technical Support is limited to customers with current support contracts. Requests for technical assistance may be submitted online (including by email), or by telephone. We recommend you use telephone contact only for urgent requests that require immediate evaluation, such as a production system down. The support website is the preferred way to contact Technical Support.

**Website:** **[techsupport.gemtalksystems.com](http://techsupport.gemtalksystems.com)**

**Email:** **[techsupport@gemtalksystems.com](mailto:techsupport@gemtalksystems.com)**

**Telephone:** **(800) 243-4772 or (503) 766-4702**

Please include the following, in addition to a description of the issue:

- ▶ The versions of GemStone/S 64 Bit and of all related GemTalk products, and of any other related products, such as client Smalltalk products, and the operating system and version you are using.
- ▶ Exact error message received, if any, including log files and statmonitor data if appropriate.

Technical Support is available from 8am to 5pm Pacific Time, Monday through Friday, excluding GemTalk holidays.

## **24x7 Emergency Technical Support**

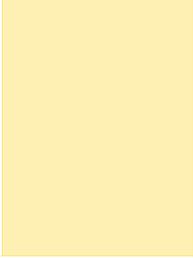
GemTalk offers, at an additional charge, 24x7 emergency technical support. This support entitles customers to contact us 24 hours a day, 7 days a week, 365 days a year, for issues impacting a production system. For more details, contact GemTalk Support Renewals.

## **Training and Consulting**

GemTalk Professional Services provide consulting to help you succeed with GemStone products. Training for GemStone/S is available at your location, and training courses are offered periodically at our offices in Beaverton, Oregon. Contact GemTalk Professional Services for more details or to obtain consulting services.



---



# Table of Contents

---

<b><i>Chapter 1. Introduction to GemStone</i></b>	<b>25</b>
1.1 GemStone Overview . . . . .	25
Multi-User . . . . .	25
Programmable . . . . .	25
Scalable . . . . .	26
Object Database . . . . .	26
Partition Between Client and Server . . . . .	26
Connect to Outside Data Sources . . . . .	27
1.2 GemStone Services . . . . .	28
Transactions and Concurrency Control . . . . .	28
Login Security and Account Management . . . . .	28
Services To Manage the GemStone Repository . . . . .	29
1.3 GemStone Smalltalk . . . . .	29
GemStone Sessions . . . . .	30
Monitoring your application . . . . .	30
Interapplication Communications . . . . .	31
1.4 Process Architecture . . . . .	31
Gem Process . . . . .	31
Stone Process . . . . .	31
NetLDI . . . . .	32
Shared Page Cache . . . . .	32
Extents and Repositories . . . . .	32
Transaction Log . . . . .	32
<b><i>Chapter 2. Class Creation</i></b>	<b>33</b>
2.1 Subclass Creation . . . . .	33
Implementation Formats . . . . .	34

Non-Indexable objects . . . . .	34
Indexable Objects . . . . .	34
NonSequencableCollection (NSC) . . . . .	35
Special . . . . .	35
Class Variables and Other Types of Variables . . . . .	35
Dynamic Instance Variables . . . . .	36
Additional Class Creation Protocol . . . . .	37
2.2 Creating Classes With Invariant Instances . . . . .	39
Per-Object Invariance . . . . .	39
Invariance for All Instances of a Class . . . . .	39
2.3 Creating Classes with Special Cases of Persistence . . . . .	40
Non-Persistent Classes . . . . .	40
DbTransient . . . . .	40
2.4 Customer-defined Special classes . . . . .	41
Money example . . . . .	42

## ***Chapter 3. Resolving Names and Sharing Objects*** **43**

3.1 Sharing Objects . . . . .	43
3.2 The UserProfile's Symbol List . . . . .	44
What's In Your Symbol List? . . . . .	44
Examining Your Symbol List . . . . .	44
Inserting and Removing Dictionaries from Your Symbol List . . . . .	46
Finding Out Which Dictionary Names an Object . . . . .	48
The Transient Symbol List . . . . .	49
Updating Symbol Lists . . . . .	50
3.3 Using Your Symbol Dictionaries . . . . .	50
Publishers, Subscribers and the Published Dictionary . . . . .	51

## ***Chapter 4. Collection and Stream Classes*** **53**

4.1 Introduction to Collections. . . . .	53
Protocol Common to All Collections . . . . .	55
Creating Instances . . . . .	55
Enumerating . . . . .	55
Collections in multi session environment . . . . .	56
Conflicting updates . . . . .	56
Visibility and ordering . . . . .	56
Collection classes . . . . .	56
Dictionary classes . . . . .	57
Internal Dictionary Structure. . . . .	57
Dictionary and KeyValueDictionary . . . . .	57
KeySoftValueDictionary . . . . .	57
SequenceableCollection classes . . . . .	58
Copying. . . . .	58

Array . . . . .	59
SortedCollection. . . . .	59
Stream Classes . . . . .	60
PositionableStream and Position . . . . .	60
AppendStream . . . . .	61
ReadByteStream. . . . .	61
UnorderedCollection classes . . . . .	61
Union, Intersection, and Difference. . . . .	62
4.2 Reduced-Conflict Collection Classes. . . . .	63
RcArray . . . . .	63
NSC/UnorderedCollection classes . . . . .	63
RcIdentityBag . . . . .	63
RcLowMaintenanceIdentityBag. . . . .	64
RcIdentitySet . . . . .	64
RcKeyValueDictionary. . . . .	64
Queue classes . . . . .	64
GsPipe . . . . .	65
RcPipe . . . . .	65
RcQueue . . . . .	65
4.3 GsBitmap . . . . .	66
GsBitmaps and C Heap memory . . . . .	67
GsBitmaps and their objects. . . . .	67
GsBitmaps methods for repository analysis. . . . .	67
Bitmap files . . . . .	68
Page order Bitmap files . . . . .	68
4.4 Sorting the objects in a collection. . . . .	69
Default Sort . . . . .	69
Sorting Application objects . . . . .	70
Sorting in multiple orders . . . . .	71
SortBlocks . . . . .	71
Sorting Large Collections . . . . .	72

## ***Chapter 5. String Classes and Collation***

**73**

5.1 Characters and Unicode . . . . .	73
Empty string canonicalization. . . . .	74
Unicode and the Unicode Database . . . . .	74
5.2 String classes . . . . .	75
Traditional Strings . . . . .	75
Unicode Strings . . . . .	76
String equality, ordering, and interoperation . . . . .	76
Other String-like classes . . . . .	76
Symbol . . . . .	76
ByteArray . . . . .	77
Utf8 . . . . .	77
Utf16 . . . . .	77

String protocol . . . . .	78
Creating Strings . . . . .	78
Concatenating Strings . . . . .	78
Converting between String classes and encodings . . . . .	78
String Transformations . . . . .	79
Equality and Identity . . . . .	79
Searching and Pattern matching . . . . .	80
5.3 String Sorting and Collation . . . . .	81
Comparison Mode . . . . .	82
StringConfiguration . . . . .	82
Auto-conversion . . . . .	82
Legacy String Comparison Mode for Traditional Strings . . . . .	83
Unicode Comparison Mode and ICU Collation . . . . .	83
IcuLocale . . . . .	84
IcuCollator . . . . .	84
Customizing Sort . . . . .	85
IcuSortedCollection . . . . .	87
ICU libraries and versioning . . . . .	88
ICU and Unicode versioning . . . . .	88
IcuLibraryVersion . . . . .	88
5.4 Encrypting Strings . . . . .	89

## Chapter 6. Numeric Classes

91

6.1 Integers . . . . .	91
SmallInteger . . . . .	92
LargeInteger . . . . .	92
Printing Integers . . . . .	92
6.2 Binary Floating Point . . . . .	92
SmallDouble . . . . .	93
Float . . . . .	93
Signalling Exception rather than returning Exceptional Float . . . . .	94
Literal Floats . . . . .	95
Printing Binary Floating Points . . . . .	95
6.3 Other Rational Numbers . . . . .	96
Fractions . . . . .	96
SmallFraction . . . . .	97
Fraction . . . . .	97
ScaledDecimals . . . . .	97
SmallScaledDecimal . . . . .	97
ScaledDecimal . . . . .	98
FixedPoints . . . . .	98
DecimalFloat . . . . .	98
Summary of literal syntax . . . . .	99
Custom numeric literals . . . . .	99
6.4 Result classes, Conversion, and Rounding . . . . .	100

Conversion . . . . .	.100
Truncation and Rounding . . . . .	.101
6.5 Dates and Times . . . . .	.101
Date . . . . .	.101
SmallDate . . . . .	.101
Instance Creation . . . . .	.102
Formatting for Instance Creation and Printing . . . . .	.102
Time . . . . .	.102
SmallTime . . . . .	.102
Instance Creation . . . . .	.103
Formatting for Instance Creation and Printing . . . . .	.103
Time offset . . . . .	.103
DateTime . . . . .	.103
Instance Creation . . . . .	.104
Formatting for Instance Creation and Printing . . . . .	.104
DateAndTime . . . . .	.105
SmallDateAndTime . . . . .	.105
Instance Creation . . . . .	.105
Formatting for Instance Creation and Printing . . . . .	.105
TimeZone . . . . .	.106
6.6 Internationalizing . . . . .	.107
Dates in GemStone log files . . . . .	.107
Internationalizing Decimal Points using Locale . . . . .	.107
Smalltalk code requires the period separator . . . . .	.108
6.7 Random Numbers . . . . .	.109
Universally unique identifiers (UUIDs) . . . . .	.109
Random Number Generator . . . . .	.109

## *Chapter 7. Files and Directories*

**113**

7.1 Accessing Files using GsFile . . . . .	.113
Client vs. Server files . . . . .	.113
Specifying Files . . . . .	.114
Creating a File. . . . .	.114
Opening a File . . . . .	.115
Closing a File or Files . . . . .	.116
Writing to a File. . . . .	.116
Reading from a File. . . . .	.117
Positioning. . . . .	.118
Testing Files . . . . .	.118
Renaming Files . . . . .	.118
Removing Files . . . . .	.119
Examining a Directory . . . . .	.119
GsFile Errors . . . . .	.120
Logging to stdout. . . . .	.120
On the server (Gem) . . . . .	.121

On the client (GCI client) . . . . .	121
7.2 FileSystem . . . . .	121
FileReference. . . . .	121
FileLocator . . . . .	122
FileSystem . . . . .	122
Specifying a FileReference . . . . .	123
Working Directory and other environment-independent paths . . . . .	123
Reading from and Writing to a File . . . . .	123
Closing file streams . . . . .	124
File encodings . . . . .	124
Binary Streams . . . . .	124
Operations on Files and Directories. . . . .	124
Create . . . . .	124
Delete . . . . .	125
Move . . . . .	125
Copy. . . . .	125
Rename . . . . .	125
Information queries . . . . .	125
Status of a file or directory . . . . .	125
Information about a file . . . . .	126
File permissions . . . . .	126
Decomposition of filename and path . . . . .	127
Listing Directories and Files . . . . .	127
Supporting Classes . . . . .	127
FsFileDescriptor . . . . .	127
Zinc Stream Classes . . . . .	128
Error Classes . . . . .	128
Path classes . . . . .	128
Opening options . . . . .	128
Store classes . . . . .	128
Resolver classes . . . . .	128
FFI support classes. . . . .	128

## ***Chapter 8. Indexes and Querying*** **129**

8.1 Overview . . . . .	130
GemStone Indexes and Queries . . . . .	131
Indexes . . . . .	131
GsQueries. . . . .	131
Deciding what to optimize . . . . .	132
Overview of the steps in creating and using indexed queries . . . . .	132
Managing Indexes . . . . .	133
Special Syntax for Indexing . . . . .	133
Historic indexing syntax . . . . .	134
Last Element Class . . . . .	134
Optimized classes . . . . .	134

Using other classes . . . . .	.135
Comparing data types . . . . .	.135
Strings in indexes . . . . .	.135
Redefining Comparison Messages . . . . .	.136
8.2 Defining Queries . . . . .	.137
Query Predicate Syntax . . . . .	.137
Predicate Terms . . . . .	.137
Combining Predicates using Boolean Logic . . . . .	.138
Combining Range Predicates . . . . .	.138
Creating a GsQuery . . . . .	.138
Query Variables . . . . .	.138
8.3 Creating Indexes . . . . .	.139
Equality and Identity Indexes . . . . .	.139
Btree and Legacy Indexes . . . . .	.140
Creating the Index . . . . .	.140
Equality Indexes on strings . . . . .	.141
Repositories in Legacy String Comparison mode . . . . .	.142
Repositories in Unicode Comparison Mode . . . . .	.142
Implicit Indexes . . . . .	.143
GsIndexOptions . . . . .	.143
Combining options . . . . .	.144
Default options . . . . .	.144
The Options in GsIndexOptions . . . . .	.144
Reduced-Conflict . . . . .	.145
Optional pathTerms . . . . .	.145
8.4 Results of Executing a GsQuery . . . . .	.146
GsQuery's Collection protocol . . . . .	.146
GsQuery enumeration methods accepting blocks . . . . .	.147
Query results as Streams . . . . .	.149
Limitations on streamable queries . . . . .	.149
8.5 Enumerated and Set-valued Indexes . . . . .	.151
Enumerated path terms in indexes and queries . . . . .	.151
Restrictions on predicates with enumerated pathTerms . . . . .	.151
Indexes and Queries with collections on the path . . . . .	.151
Set-valued query results . . . . .	.152
Restrictions on predicates in set-valued queries . . . . .	.152
8.6 Managing Indexes . . . . .	.152
While Indexes are Being Created . . . . .	.152
Queries during index creation . . . . .	.153
Auto-commit . . . . .	.153
Indexes on temporary collections . . . . .	.154
Inquiring About Indexes . . . . .	.154
Removing Indexes . . . . .	.154
To remove indexes based on a GsIndexSpec . . . . .	.154
To remove indexes using IndexManager . . . . .	.155
Rebuilding Indexes . . . . .	.155

Indexing Errors . . . . .	156
Auditing Indexes . . . . .	156
8.7 Indexing and Performance. . . . .	157
Type of index . . . . .	157
Data updates. . . . .	157
Formulating queries and performance . . . . .	158
Auto-optimize . . . . .	158
8.8 Historic Indexing API differences. . . . .	158
Index creation using UnorderedCollection protocol . . . . .	158
Internal legacy vs. btreePlus indexing structures . . . . .	159
String and Unicode Equality Indexes . . . . .	159
Reduced-conflict Equality Indexes. . . . .	159
Queries using Selection Blocks. . . . .	159
Executing Selection Block Queries. . . . .	160
Managing indexes . . . . .	161
Information about indexes . . . . .	161
Removing Indexes . . . . .	161

## ***Chapter 9. Transactions and Concurrency Control*** **163**

9.1 GemStone's Conflict Management . . . . .	163
Note on Terminology . . . . .	163
Snapshot Views and Transactions . . . . .	164
Transaction State and Transaction Modes . . . . .	166
Reading and Writing in Transactions . . . . .	167
Reading and Writing Outside of Transactions . . . . .	168
When Should You Commit a Transaction?. . . . .	168
Nested In-memory Transactions. . . . .	168
9.2 How GemStone Detects and Manages Conflict . . . . .	169
Concurrency Management . . . . .	169
Committing Transactions. . . . .	170
Handling Commit Failure in a Transaction . . . . .	170
Transaction Conflicts . . . . .	170
More details about transaction conflicts . . . . .	172
Indexes and Concurrency Control. . . . .	172
Aborting Transactions . . . . .	173
Updating the View Without Committing or Aborting . . . . .	173
Being Signaled To Abort . . . . .	174
Being Signaled to continueTransaction . . . . .	174
Handlers for abort or continueTransaction notifications . . . . .	175
9.3 Controlling Concurrent Access with Locks . . . . .	175
Lock Types . . . . .	175
Read Locks . . . . .	176
Write Locks. . . . .	176
Acquiring Locks . . . . .	177
Lock Denial. . . . .	177

Deadlocks . . . . .	178
Dirty Locks . . . . .	178
Locking Collections of Objects Efficiently . . . . .	179
Upgrading Locks . . . . .	181
Locking and Indexed Collections . . . . .	181
Removing or Releasing Locks . . . . .	182
Releasing Locks Upon Aborting or Committing . . . . .	182
Inquiring About Locks . . . . .	183
Application Write Locks . . . . .	184
9.4 Classes That Reduce the Chance of Conflict . . . . .	185
RcCounter . . . . .	185
Reduced-Conflict Collection Classes . . . . .	186
RcArray . . . . .	187
RcIdentityBag . . . . .	187
RcLowMaintenanceIdentityBag and RcIdentitySet . . . . .	188
RcKeyValueDictionary . . . . .	188
GsPipe . . . . .	188
RcPipe . . . . .	188
RcQueue . . . . .	189

## ***Chapter 10. Object Security and Authorization*** **191**

10.1 How GemStone Security Works . . . . .	191
Login Authorization . . . . .	192
The UserProfile . . . . .	192
System Privileges . . . . .	192
Object-level Security . . . . .	192
GsObjectSecurityPolicy . . . . .	193
10.2 Assigning Objects to Security Policies . . . . .	195
Default Security Policy and Current Security Policy . . . . .	195
Objects and Security Policies . . . . .	196
Configuring Authorization for an Object Security Policy . . . . .	196
How GemStone Responds to Unauthorized Access . . . . .	197
Owner, Group, and World Authorization . . . . .	197
Predefined GsObjectSecurityPolicies . . . . .	198
GsObjectSecurityPolicy names . . . . .	199
Changing the Security Policy for an Object . . . . .	199
Revoking Your Own Authorization: a Side Effect . . . . .	201
Finding Out Which Objects Are Protected by a Security Policy . . . . .	202
10.3 Application Example . . . . .	202
10.4 Development Example . . . . .	206
10.5 Planning Security Policies for User Access . . . . .	206
Protecting the Application Classes . . . . .	206
CodeModification privilege . . . . .	206
Planning Authorization for Data Objects . . . . .	207
Planning Groups . . . . .	208

Planning Security Policies . . . . .	210
Developing the Application . . . . .	210
Setting Up Security Policies for Joint Development . . . . .	210
Making the Application Accessible for Testing . . . . .	212
Moving the Application into a Production Environment . . . . .	213
Security Policy Assignment for User-created Objects . . . . .	213
10.6 Privileged Protocol for Class GsObjectSecurityPolicy . . . . .	214

## ***Chapter 11. Class versions and Instance Migration*** **217**

11.1 Versions of Classes . . . . .	217
Defining a New Version . . . . .	218
New Versions and Subclasses . . . . .	218
New Versions and References in Methods . . . . .	218
Class Variables and Class Instance Variables . . . . .	219
Class versioning and Class options . . . . .	219
11.2 ClassHistory . . . . .	219
Defining a Class as a new version of an existing Class . . . . .	220
Accessing a Class History . . . . .	221
Assigning to a Class History . . . . .	221
11.3 Migrating Objects . . . . .	221
Migration Destinations . . . . .	222
Bypassing the Migration Destination . . . . .	222
Migrating Instances that Participate in an Index. . . . .	223
Default Instance Variable Mappings . . . . .	224
Customized Instance Variable Mappings. . . . .	225
Transforming Variable Values . . . . .	226
Finding Instances . . . . .	227
Tuning migration and managing memory . . . . .	228
Using GsBitmaps to manage memory for large result sets . . . . .	229
Tuning system resource use when finding instances . . . . .	229
Committing the migration in chunks . . . . .	229
Migrating instances in Page Order. . . . .	230
11.4 Multi-threaded instance migration . . . . .	232
InstVarMappingArray . . . . .	232
Example. . . . .	233

## ***Chapter 12. Encryption and Validation*** **235**

12.1 Overview for SSL keys and certificates . . . . .	236
GsTlsCredential . . . . .	236
Creating a GsTlsCredential. . . . .	236
Verifying public/private key pairs . . . . .	237
Encryption and signing algorithms . . . . .	237
12.2 Checksums and HMAC. . . . .	238

Checksums . . . . .	.238
HMAC (Hash-based message authentication codes) . . . . .	.238
12.3 Symmetric-Key Encryption . . . . .	.239
Encryption . . . . .	.239
Example . . . . .	.240
12.4 Digital Signatures . . . . .	.240
12.5 Digital Envelopes . . . . .	.241
Creating the GsDigitalEnvelope . . . . .	.242
Using the GsDigitalEnvelope . . . . .	.243

## **Chapter 13. Operating System Access 245**

13.1 Executing Operating System Commands . . . . .	.246
Privileges and limiting OS access . . . . .	.246
performOnServer: for simple commands . . . . .	.246
Using other shells . . . . .	.247
GsHostProcess for more complex interactions . . . . .	.247
Using execute:input: to pass data to stdin . . . . .	.248
Using fork: for interactive commands . . . . .	.248
Connecting stdin, stdout, and stderr . . . . .	.248
13.2 Setting environment variables . . . . .	.249
13.3 Creating and Using Sockets . . . . .	.249
GsSocket and GsSignallingSocket . . . . .	.250
Establishing the connection . . . . .	.250
Communication on the socket . . . . .	.250
Closing the socket . . . . .	.250
Socket Configuration . . . . .	.251
GsSecureSocket . . . . .	.251
Certificates, keys, and passphrases . . . . .	.251
Enable or disable verifying CA Certificate . . . . .	.252
Set certificate, private key, and passphrase . . . . .	.253
Setup the Cipher list . . . . .	.254
Establishing the connection . . . . .	.255
Communication on the socket . . . . .	.255
Closing the socket . . . . .	.255
HTTPS connection . . . . .	.256
Error handling . . . . .	.256
GsSocket . . . . .	.256
GsSecureSocket . . . . .	.256
GsSecureSocket pre-shared keys . . . . .	.257
TLS version handling . . . . .	.257
Example in image . . . . .	.257
13.4 ssh and sftp using OpenSSL . . . . .	.258
SSH with GsSshSocket . . . . .	.258
Creating the SSH connection . . . . .	.258
SSH operations . . . . .	.258

SFTP with GsSftpSocket . . . . .	259
Creating Sftp connection . . . . .	259
Sftp operations . . . . .	259
GsSftpRemoteFile for operations on remote files . . . . .	259
13.5 Serializing Data . . . . .	261
PassiveObject . . . . .	261
JSON . . . . .	262

## ***Chapter 14. Signals and Notifiers*** **265**

14.1 Communicating Between Sessions . . . . .	265
14.2 Object Change Notification. . . . .	266
Setting Up a Notify Set . . . . .	266
Adding an Object to a Notify Set. . . . .	266
Adding a Collection to a Notify Set . . . . .	268
Listing Your Notify Set . . . . .	269
Removing Objects From Your Notify Set . . . . .	269
Notification of New Objects . . . . .	269
Receiving Object Change Notification. . . . .	270
Reading the Set of Signaled Objects . . . . .	271
Polling for Changes to Objects. . . . .	271
Troubleshooting. . . . .	272
Frequently Changing Objects. . . . .	272
Special Classes . . . . .	272
Methods for Object Notification . . . . .	274
14.3 Gem-to-Gem Signaling . . . . .	274
Sending a Signal. . . . .	275
Legacy protocol to send signals . . . . .	276
Receiving a Signal. . . . .	276
Polling. . . . .	277
Receiving a Notification. . . . .	278
14.4 Other Signal-Related Issues . . . . .	279
Inactive Gem . . . . .	279
Dealing With Signal Overflow . . . . .	279
Sending Large Amounts of Data. . . . .	280
Maintaining Signals and Notification When Users Log Out . . . . .	280

## ***Chapter 15. Handling Exceptions*** **281**

15.1 The Exception Class Hierarchy. . . . .	281
15.2 Signaling Exceptions . . . . .	283
15.3 Handling Exceptions . . . . .	284
Dynamic (Stack-Based) Handlers . . . . .	284
Selecting a Handler . . . . .	285
Flow of Control . . . . .	287

Default Handlers . . . . .	.288
Default Actions . . . . .	.289
15.4 The Legacy Exception Handling Framework . . . . .	.290
Dynamic (Stack-Based) Exception Handler . . . . .	.290
Installing a Dynamic (Stack-Based) Exception Handler . . . . .	.290
Default (Static) Exception Handlers . . . . .	.291
Installing a Default (Static) Exception Handler . . . . .	.291
GemStone Event Exceptions . . . . .	.292
Flow of Control . . . . .	.293
Signaling Other Exception Handlers . . . . .	.295
Removing Exception Handlers . . . . .	.295
Recursive Errors. . . . .	.296
Raising Exceptions . . . . .	.296
ANSI Integration . . . . .	.297

## **Chapter 16. Performance and Optimization** **299**

16.1 Profiling Smalltalk Execution . . . . .	.300
Time to execute a block . . . . .	.300
CPU Time . . . . .	.300
Elapsed Time . . . . .	.300
ProfMonitor . . . . .	.301
Sample intervals. . . . .	.301
Reporting limits. . . . .	.301
Reports. . . . .	.302
Temporary results file . . . . .	.302
Real vs. CPU time. . . . .	.302
Profiling Code. . . . .	.303
Convenience Profiling of a Block of Code . . . . .	.303
Background Profiling. . . . .	.304
Manual Profiling . . . . .	.304
Saving a ProfMonitor for later analysis. . . . .	.305
The Profile Report . . . . .	.306
Profiling Beyond Performance . . . . .	.309
Object Creation Tracking. . . . .	.310
Memory Use Profiling . . . . .	.310
16.2 Clustering Objects for Faster Retrieval . . . . .	.311
Will Clustering Solve the Problem? . . . . .	.311
Cluster Buckets . . . . .	.312
Using Existing Cluster Buckets . . . . .	.312
Creating New Cluster Buckets . . . . .	.313
Cluster Buckets and Concurrency. . . . .	.313
Cluster Buckets and Indexing . . . . .	.314
Clustering Objects . . . . .	.314
The Basic Clustering Message . . . . .	.314
Depth-First Clustering . . . . .	.316

Assigning Cluster Buckets . . . . .	316
Clustering and Memory Use . . . . .	317
Using Several Cluster Buckets . . . . .	317
Clustering Class Objects . . . . .	317
Maintaining Clusters . . . . .	318
Determining an Object's Location . . . . .	318
Why Do Objects Move? . . . . .	319
16.3 Modifying Cache Sizes for Better Performance . . . . .	320
GemStone Caches . . . . .	320
Temporary Object Space . . . . .	320
Shared Page Cache . . . . .	320
Getting Rid of Non-Persistent Objects . . . . .	321
16.4 Managing VM Memory . . . . .	322
Large Working Set. . . . .	322
Class Hierarchy . . . . .	322
UserAction Considerations . . . . .	322
Exported Set . . . . .	323
Debugging out of memory errors . . . . .	323
Signal on low memory condition . . . . .	323
Methods for Computing Temporary Object Space . . . . .	324
Statistics for monitoring memory use . . . . .	325
16.5 NotTranloggedGlobals . . . . .	327
16.6 Other Optimization Hints . . . . .	328

## ***Chapter 17. Working with Classes and Methods*** **331**

17.1 Creating and Removing Methods . . . . .	331
Defining Simple Accessing and Updating Methods. . . . .	331
Compiling Methods. . . . .	332
Removing Methods . . . . .	333
Pragmas . . . . .	333
Pragma class . . . . .	334
17.2 Information about Class and Methods . . . . .	335
Information about the Class . . . . .	335
Information about Instance, Class, and Shared Pool variables. . . . .	335
Information about Method Selectors . . . . .	335
Accessing and Managing Method Categories . . . . .	336
Specific Methods . . . . .	336
17.3 Transient Methods. . . . .	336
17.4 ClassOrganizer. . . . .	337
Examples . . . . .	338
17.5 Handling Deprecated Methods . . . . .	339
Deprecated handling . . . . .	339
Deprecation log . . . . .	340
Listing deprecated methods . . . . .	340
Determining senders of deprecated methods . . . . .	340

17.6 File In and File Out of GemStone code . . . . .	341
Fileout . . . . .	341
Filein . . . . .	341
GsFileIn . . . . .	341

## ***Chapter 18. GemStone System Features*** **343**

18.1 Hidden Sets . . . . .	343
Sets still accessed via System methods. . . . .	344
NotifySet. . . . .	344
ExportedDirtyObjs and TrackedDirtyObjs. . . . .	344
PureExportSet and GciTrackedObjs . . . . .	344
18.2 SessionTemps and access to Session State . . . . .	345
SessionState . . . . .	345
18.3 Shared Counters . . . . .	345
AppStat Shared Counters . . . . .	346
Persistent Shared Counters . . . . .	346
18.4 GsEventLog . . . . .	347
Adding events. . . . .	347
Querying and reporting . . . . .	347
Deleting events . . . . .	348

## ***Chapter 19. The Foreign Function Interface*** **349**

19.1 Overview of the Foreign Function Interface . . . . .	349
19.2 Using the FFI . . . . .	350
CLibrary: defining the compiled C or C++ library . . . . .	350
CCallout: function definitions in GemStone. . . . .	350
CHeader: Parsing the C header file. . . . .	351
CDeclaration: Finding details on a C function . . . . .	351
Creating a Class and Methods from the CHeader . . . . .	352
CByteArray: Allocating memory for data . . . . .	352
Complex data types. . . . .	353
Making FFI calls . . . . .	353
Executing a CCallout created by hand . . . . .	354
Execute CHeader-generated definition. . . . .	354
Arguments and Return Values . . . . .	354
String type arguments . . . . .	355
Variable Arguments . . . . .	355
CCalloutStructs: Using C Structures passed by value . . . . .	355
CPointer . . . . .	356
Errno Handling . . . . .	356
CCallin: Creating and invoking callbacks . . . . .	356
C type symbols . . . . .	357
Limitations with native code disabled . . . . .	359

19.3 Example using Zlib . . . . .	359
Parsing the header file . . . . .	359
Using CHeader wrapper methods to create a class . . . . .	364
19.4 Example using CCalloutStructs and variable arguments with RabbitMQ® . . . . .	366
Hand creation of CCalloutStructs . . . . .	367
Using CHeader wrapper methods to create CCalloutStructs . . . . .	368
Further refinement: creating a convenience method. . . . .	369
<b>Chapter 20. External Sessions</b>	<b>371</b>
20.1 Overview . . . . .	371
20.2 External Sessions . . . . .	372
Setup the External Session . . . . .	372
Creating the External Session . . . . .	372
Avoiding passwords in code . . . . .	373
Log in the External Session. . . . .	374
Executing Code . . . . .	374
Important caution on Export Set of remote session . . . . .	376
Exceptions . . . . .	376
20.3 NRS and Login Parameter Support . . . . .	377
For a Stone . . . . .	378
For a Gem. . . . .	378
Convenience methods for common arguments . . . . .	378
20.4 Special Cases of External Sessions . . . . .	379
Solo external sessions . . . . .	379
Primitive-based external sessions on AIX . . . . .	380
X509 external sessions . . . . .	380
<b>Chapter 21. Testing and SUnit</b>	<b>381</b>
21.1 Application Testing . . . . .	381
The value of testing . . . . .	381
Writing good tests. . . . .	382
Why SUnit? . . . . .	383
Portable SUnit classes and GemStone's SUnit classes. . . . .	383
Terminology: test errors vs. failures. . . . .	383
21.2 SUnit example: ExampleSetTest . . . . .	384
Define the new testing Class . . . . .	384
Define setUp and tearDown methods. . . . .	384
Define the actual testing methods . . . . .	385
Execute the tests . . . . .	385
Assertions and other claims in test methods . . . . .	386
Checking result value with assert: and deny:. . . . .	386
Checking for exceptions with should:raise:. . . . .	386
21.3 The SUnit Framework and Implementation . . . . .	387

SUnit Core Classes . . . . .	.387
Running a Single Test . . . . .	.388
Running a TestSuite . . . . .	.389

## ***Appendix A. GemStone Smalltalk Syntax*** **391**

A.1 GemStone and ANSI Smalltalk . . . . .	.391
GemStone and ANSI limits . . . . .	.392
A.2 GemStone Smalltalk. . . . .	.392
How to Create a New Class . . . . .	.393
Statements. . . . .	.393
Comments. . . . .	.393
Expressions . . . . .	.394
Literals . . . . .	.394
Numeric Literals . . . . .	.395
Character Literals . . . . .	.395
String Literals . . . . .	.396
Symbol Literals . . . . .	.396
Array Literals . . . . .	.396
Variables and Variable Names . . . . .	.397
Assignment . . . . .	.398
Message Expressions. . . . .	.398
Reserved and Optimized Selectors . . . . .	.399
Conditionally Optimized Selectors . . . . .	.400
Messages as Expressions. . . . .	.400
Combining Message Expressions . . . . .	.401
Cascaded Messages . . . . .	.402
Array Constructors. . . . .	.403
Path Expressions . . . . .	.404
Returning Values. . . . .	.404
A.3 Blocks. . . . .	.405
Blocks with Arguments . . . . .	.406
Blocks and Conditional Execution . . . . .	.406
Formatting Code . . . . .	.408
A.4 GemStone Smalltalk BNF. . . . .	.410



# Introduction to GemStone

---

This chapter introduces you to the GemStone/S 64 Bit™ (GemStone) system. GemStone provides a distributed, server-based, multi-user, transactional Smalltalk runtime system, with the ability to partition the application between client and server.

GemStone provides enterprise-quality security, scalability, availability, and services for managing and monitoring the repository.

## 1.1 GemStone Overview

### Multi-User

GemStone can support thousands of concurrent users, object repositories of hundreds of gigabytes, and sustained object transaction rates of hundreds of transactions per second. Server processes manage the system, while user sessions support individual user activities. Repository and server processes can be distributed among multiple machines, leveraging shared memory and SMP.

Multiple user sessions can be active at the same time, and each user may have multiple sessions open. A flexible naming scheme allows separate or shared namespaces for individual users. Changes that users make to objects are committed in transactions, with concurrency controls and locks ensuring that multi-user changes to objects are coordinated. Security is provided at several levels, from login authorization to method execution privileges and object access privileges.

### Programmable

GemStone provides data definition, data manipulation, and query facilities in a single, computationally complete language – GemStone Smalltalk. The GemStone Smalltalk language offers built-in data types (classes), operators, and control structures comparable in scope and power to those provided by languages such as C or Java, in addition to multi-user concurrency and repository management services. All system-level facilities, such as transaction control, user authorization, and so on, are accessible from GemStone Smalltalk.

## Scalable

Object programming languages such as Smalltalk are highly efficient development tools. Smalltalk exploits inheritance and code reuse and provides the flexibility of modeling real world objects with self-contained software modules. Most Smalltalk implementations, however, are memory based, and objects exist only in a single user's image.

Like a single-user Smalltalk image, GemStone consists of classes, methods, instances and meta objects. Persistence is established by attaching new objects to other persistent objects. All objects are derived from a named root (AllUsers). Objects that have been attached and committed to the repository are visible to all other authorized users.

However, since the GemStone repository is accessed through disk caches, it is not limited in size by available memory. A GemStone repository can contain billions of objects, each with a unique object identifier (known as an OOP – object-oriented pointer).

## Object Database

GemStone lets you model information in structures as simple or complex as application data requires. You can represent data objects in tables, hierarchies, networks, queues, or any other structure or nested combination of structures that is appropriate.

Because you can represent information in forms that mirror the information's natural structure, the translation of user requests into executable queries can be much easier in GemStone. You do not need to translate users' keystrokes or menu selections into relational algebra formulas, calculus expressions and procedural statements before the query can be executed. See Chapter 8, "Indexes and Querying".

## Partition Between Client and Server

GemStone applications can access objects and run their methods from a number of languages, including Smalltalk, C, Java, or any language that makes C calls. Objects created from any of these languages are interoperable with objects created from the other languages, and can run their methods within GemStone.

To provide this functionality, GemStone provides interface libraries of Smalltalk classes, Java classes, and C functions. These GemBuilder™ language interfaces allow you to move objects between an application program and the GemStone repository, and to connect client objects to GemStone objects. GemBuilder also provides remote messaging capabilities, client replicates, and synchronization of changes.

GemStone's interfaces include:

### **GemBuilder for Smalltalk**

GemBuilder for Smalltalk consists of two parts: a set of GemStone programming tools, and a programming interface between the client application code and GemStone.

GemBuilder for Smalltalk contains a set of classes installed in a client Smalltalk image that provides access to objects in a GemStone repository. Many of the client Smalltalk kernel classes are mapped to equivalent GemStone classes, and additional class mappings can be created by the application developer.

GemBuilder for Smalltalk is a separate product, and includes documentation describing installation and use.

### GemBuilder for Java

GemBuilder for Java also has two parts: a set of GemStone programming tools, and a programming interface between the client application code and GemStone.

GemBuilder for Java is a Java runtime package that provides a message-forwarding interface between a Java client and a GemStone server, allowing access to objects in a GemStone repository.

GemBuilder for Java is distributed as a separate product, and includes documentation describing installation and use.

### GemBuilder for C

GemBuilder for C is a library of C functions that provide a bridge between an application's C code and the GemStone repository. This interface allows programmers to work with GemStone objects by importing them into the C program using structural access, or by sending messages to objects in the repository through GemStone Smalltalk.

GemBuilder for C is distributed with the server product. For more information on GemBuilder for C, see the *GemBuilder for C* manual.

### GsDevKit

GsDevKit, the open-source development kit for GemStone/S 64 Bit (formerly referred to as GLASS or Seaside), provides a Pharo-compatible GemStone Smalltalk environment. With the optional Seaside framework, you can create and deploy desktop-like web applications.

GsDevKit and Seaside for GemStone are distributed as open-source products via GitHub. For more information, see [gemtalksystems.com/small-business/gsdevkit/](http://gemtalksystems.com/small-business/gsdevkit/).

### Rowan and Jadeite

Rowan is an open-source code manager in GemStone, using git as a back end; Jadeite is a Dolphin-based graphical user interface designed for Rowan. These are ongoing development projects in preview state.

In addition to these interfaces, GemStone provides a command-line tool that allows you to interact with server objects, execute code, and perform limited scripting:

### Topaz

Topaz is a GemStone programming environment that provides a scriptable command-line interface to GemStone Smalltalk. Topaz is most commonly used for performing repository maintenance operations. Topaz offers access to GemStone without requiring a window manager or additional language interfaces. You can use Topaz in conjunction with other GemStone development tools such as GemBuilder for C to build comprehensive applications.

Topaz is part of the server distribution. For more information on Topaz, see the *Topaz User's Guide*.

## Connect to Outside Data Sources

The productivity value of GemStone comes from coding in Smalltalk, but you may need or want to call out to logic written elsewhere, as for instance specialized C libraries. GemStone provides several ways to access external code from a GemStone session.

### Foreign Function Interface (FFI)

FFI classes with GemStone allow you to invoke functions in existing C libraries. The

argument and return data types are defined within GemStone Smalltalk to conform to the C function definition. The FFI interface is part of the GemStone kernel, and is documented in Chapter 19, “The Foreign Function Interface”.

#### **UserActions (C callouts from GemStone Smalltalk)**

UserActions are similar to user-defined primitives in other Smalltalks. You can use GemBuilder for C to write these user actions, and invoke these user actions from GemStone Smalltalk. The tools supporting user actions are part of the GemStone kernel, and are documented in the *GemBuilder for C* manual.

#### **GemConnect (Access to Oracle database)**

GemStone uses the User Action mechanism to build the GemConnect™ product, which provides access to relational database information from GemStone objects. GemConnect is fully encapsulated and maintained in the GemStone object server. GemConnect is distributed as a separate product, and includes documentation describing installation and use.

## 1.2 GemStone Services

### Transactions and Concurrency Control

Each GemStone session defines and maintains a consistent working environment for its application program, presenting the user with a consistent snapshot view of the object repository. The user works in an environment in which only his or her changes to objects are visible. These changes are private to the user until the transaction is committed. The effects of updates to the object repository by other users are minimized or invisible during the transaction. GemStone then checks for consistency with other users' changes before committing the transaction, and refuses to commit conflicting changes.

GemStone provides both optimistic and pessimistic approaches to managing concurrent transactions, and supports explicit object locking for read or write. To allow users to modify the same object in ways that do not actually conflict, such as two users adding to a collection, GemStone extends the Collection class hierarchy by providing reduced-conflict (Rc) classes that can be used in place of standard collection classes.

For more on transactions and reduced-conflict classes, See Chapter 9, “Transactions and Concurrency Control”.

### Login Security and Account Management

Compared to a single-user Smalltalk system, GemStone requires substantially more security mechanisms and controls. As a tool for server implementation, multi-user Smalltalk must handle requests from many users running a variety of applications, each of which can require different accessibility of objects. Authentication and authorization are the cornerstones of GemStone Smalltalk security.

#### **Login Authentication**

Before users can access system resources, they must be authenticated. Logins can be done from any of the interfaces; in each case, GemStone requires a *user ID* and a password, and a corresponding *UserProfile* must exist in GemStone. Authentication of the user ID and password can be done using GemStone's encryption, using UNIX, by Lightweight Directory Access Protocol (LDAP), or using Kerberos. GemStone uses

SRP and SSL/TLS to establish secure logins and certain types of interprocess connections. Authentication and login security features are described in the *System Administration Guide*.

#### **Object-level Authorization**

To control access to individual objects, GemStone provides object-level authorization. Authorization enforcement is implemented at the lowest level of basic object access to prevent users from circumventing the authorization checking. Read and write authorization can be granted to single objects or groups of objects, for single users or groups of users. See Chapter 10, "Object Security and Authorization".

#### **User Privileges**

GemStone defines a set of privileges for controlling the use of certain system services. Privileges determine whether the specific user is allowed to execute certain system functions, usually ones only performed by the system administrator. Privileges are described in the *System Administration Guide*.

## **Services To Manage the GemStone Repository**

GemStone is capable of managing objects shared by thousands of users, running methods that access billions of objects, and handling queries over large collections of objects by using indexes. It can support large-scale deployments on multiple machines in a variety of network configurations. All of this functionality requires a wide array of services for management of the repository, the system processes, and user sessions. These services are described in the *System Administration Guide*.

## **1.3 GemStone Smalltalk**

GemStone Smalltalk is tailored to operate in a multi-user environment, with transaction throughput and client communication as chief considerations. GemStone's class library is designed for multi-user access to objects. At the same time, its common characteristics with other Smalltalks allow you to implement shared business objects with the same language you use to build client applications. Since the same code can execute either on the client or on the object server, you can easily move behavior from the client to the server for application partitioning.

With a limited number of exceptions, GemStone Smalltalk supports the ANSI Smalltalk standard.

### **No User Interface**

Because GemStone is an object server, GemStone Smalltalk does not provide any classes for screen presentation or user interface development. Graphical user interfaces, including those for developing classes and methods as well as runtime user interfaces, are provided by the client application. The client application uses a GemBuilder interface or a web interface such as Seaside to communicate and interact with the GemStone server.

A significant part of programming with GemStone is designing the interactions between various client runtime systems and the GemStone classes, methods, and objects on the server.

## GemStone Sessions

The GemStone interfaces provide access to GemStone objects and mechanisms for running GemStone methods in the server. This access is accomplished by establishing a session with the GemStone object server. The process for establishing a session is tailored to the language or user of each interface. In all cases, however, this process requires identification of the GemStone object server to be used, the user ID for the login, and other information required for authenticating the login request.

Once a session is established, all GemStone activity is carried out in the context of that session, be it low-level object access and creation, or invocation of GemStone Smalltalk methods.

Sessions allow multiple users to share objects. In fact, different sessions can access the same repository in different ways, depending on the needs of the applications or users they are supporting. For example, an employee may only be able to access employee names, telephone extensions and department names through the human resources application, while a manager may be able to access and change salary information as well.

Sessions also control transactions, which are the only way changes to the repository can be committed. However, a *passive* session can run outside a transaction for better performance and lower overhead. For example, a stock portfolio application that reports the current value of a collection of stocks may run in a session outside a transaction until notified that a price has changed in a stock object. The application would then start a transaction, commit the change, and recalculate the portfolio value. It would then return to a passive session state until the next change notification.

A session can be integrated with the application into a single process, called a *linked* application. Each application can have only one linked session.

Alternatively, the session can run as a separate process and respond to remote procedure calls (RPC calls) from the application. These sessions are called *RPC* applications. An application may have multiple RPC sessions running simultaneously with each other and a linked session.

## System Management Classes

GemStone Smalltalk provides a number of classes that offer system management functionality.

- ▶ The class `System`, which has no instances, provides class protocol to manage the repository and individual session.
- ▶ The class `Repository`, which has a single instance named `SystemRepository`, provides protocol for data management functions, such as extent creation and access, backup and restore, and garbage collection.
- ▶ The class `UserProfileSet`, which has a single instance named `AllUsers`, provides protocol to create and manage users.

## Monitoring your application

GemStone includes `statmonitor` and `Visual Stat Display (VSD)` utilities, which allow you to monitor and record, and view statistics about your application performance. This allows precise tuning as well as detecting potential problems before they occur. GemStone also

includes profiling classes that allow you to optimize and tune your Smalltalk code for maximum performance.

## File In and File Out

GemStone Smalltalk allows you to file out source code for classes and methods, save the resulting text file, and file it in to another repository. The GemStone class `PassiveObject` also allows you to create a text representations of the binary objects, which can be written to a file and read into another repository.

## Interapplication Communications

GemStone Smalltalk provides several ways to send information from one currently logged-in session to another:

- ▶ GemStone can tell an application when an object has changed by sending the application a **notifier** at the time of commit. Notifiers eliminate the need for the application to repeatedly query the Gem for this information. Notification is optional, and can be enabled for only those objects in which you are interested.
- ▶ Applications can send messages directly to one another by using Gem-to-Gem **signals**. Sending a signal requires a specific action by the receiving Gem.

## 1.4 Process Architecture

GemStone provides the technology to build and execute applications that are designed to be partitioned for execution over a distributed network. GemStone's architecture provides both scalability and maintainability. The following sections describe the main aspects of GemStone architecture.

### Gem Process

For each login, a GemStone session is established with a Gem process. The Gem runs GemStone Smalltalk and processes messages from the client session. It provides the user with a consistent snapshot view of the repository, and it manages the user's session, keeping track of the objects the users has accessed, paging objects in and out of memory as needed, and performing dynamic garbage collection of temporary objects. A user application is always connected to at least one Gem, and may have connections to many Gems. Gems can be distributed on multiple, heterogeneous servers.

In addition to Gem Processes for user sessions, a running GemStone system includes a number of maintenance Gem processes. These system Gems include the GcGems, which handle the tasks of collecting objects that are no longer referenced and the SymbolGem, which centralizes the creation of unique, canonical symbols.

### Stone Process

The Stone process is the resource coordinator. One Stone process manages one repository. The Stone synchronizes activities and ensures consistency as it processes requests to commit transactions. Individual Gem processes communicate with the Stone through interprocess channels.

## NetLDI

Most GemStone configurations will include a network server process, known as a NetLDI (Network Long Distance Information). The NetLDI is responsible for starting up GemStone processes such as Gems, and coordinates startup when GemStone processes are needed on a node other than the one the Stone is running on.

## Shared Page Cache

The shared page cache (SPC) provides efficient retrieval of objects from disk, and the ability for multiple Gems to access the same object. The SPC is a large, contiguous area of shared memory that is shared by the Stone and each Gem process on that host. Memory is managed and allocated on pages within this shared memory. A cache is started on each machine that runs a Stone monitor, Gem session process, or linked application.

The SPC also contains buffers for communications between Gems and the Stone. The Shared Cache Monitor process initializes the shared memory cache, manages allocation to the sessions, and dynamically adjusts this allocation to fit the workload. It also makes sure that frequently accessed objects remain in memory, and that large objects queries do not flush data from the cache. These controls allow complex applications to be run on the same repository by multiple users without performance degradation.

## Extents and Repositories

Extents are composed of multiple disk files or raw partitions. A repository, which is the logical storage unit in which GemStone stores objects, is actually an ordered collection of one or more extents.

## Transaction Log

GemStone's transaction log provides complete point-in-time roll-forward recovery. The transaction log contents are composed by the Gem, and the Stone writes the tranlog using asynchronous I/O. Commit performance is improved through I/O reduction, because only log records need to be written, not many object pages. In addition, the object pages stay in memory to be reused. Transaction logs may be on file systems or on raw devices.

# Class Creation

---

The first thing you will want to do is create the classes that will implement your application. This chapter describes class creation protocol, including some special features that can apply to all instances of a class.

**Subclass Creation** (page 33)

explains how to define new GemStone classes, class implementation formats and other ways classes can store data.

**Creating Classes With Invariant Instances** (page 39)

describes how to make objects invariant.

**Creating Classes with Special Cases of Persistence** (page 40)

explains how classes can be defined so that their instances or instance variables are not stored in the repository.

**Customer-defined Special classes** (page 41)

describes how to customize the GemStone-supplied special classes to define your own special classes.

## 2.1 Subclass Creation

Almost every class in the GemStone system understands a message that causes it to create a subclass of itself.

**Example 2.1**

---

```
Object subclass: 'Animal'  
  instVarNames: #('habitat' 'name' 'predator')  
  classVars: #('AllAnimals')  
  classInstVars: #('AllOfSpecies')  
  poolDictionaries: #()  
  inDictionary: UserGlobals
```

---

This subclass creation message establishes a name ('Animal') for the new class and provides for three named instance variables ('habitat', 'name', and 'predator'), a class variable ('AllAnimals'), and a class instance variable ('AllOfSpecies'). The new class is installed in the symbolDictionary UserGlobals of the user who executes this code. You may also include reference to poolDictionaries, if this is useful for your application. Pool dictionaries are included by value, not by name; in other words, you use the reference to the pool dictionary, not a String.

The String used for the new class's name must follow the general rule for variable names – that is, it must begin with an alphabetic character and its length must not exceed 1024 characters.

There are a number of subclass creation methods. The first keyword (in the example above, `subclass:`) defines the implementation format – more on this in the next section. Subclass creation methods with additional keywords are provided to provide other information to use when creating the class.

Some GemStone server classes cannot be subclassed. This is an attribute of the class. Execute `class subclassesDisallowed` to determine if a specific class can be subclassed.

## Implementation Formats

Objects typically encapsulate data and behavior. The behavior is defined as methods on a class and the data is stored in the object. The data may be stored in named instance variables, indexed instance variables (Collection elements), or by value in specialized internal structures.

The implementation format refers to how the basic structure of the objects are defined by the class, which is done when the class is created. Implementation may be inherited from the superclass, or by using specific subclass creation methods you can specify the implementation format of the class.

### Non-Indexable objects

Many types of objects have named instance variables, but no indexable variables. Objects may have up to 2030 named instance variables, which are referred to by name in the code for that class. This limit includes all inherited instance variables as well as instance variables defined by the class.

This is the default format; subclass creation methods that begin with the `subclass:` keyword will create classes of this format, if another format is not inherited.

### Indexable Objects

Indexable objects have a variable number of elements, essentially instance variables that are referenced by an Integer index; these are may be referred to as indexed instance variables, varying instance variables, or unnamed instance variables. The number of an object's indexed instance variables can increase dynamically at run time, up to  $2^{40}-1$  (about a trillion). There are two general cases of indexable objects:

#### Pointer-format

Pointer-format indexable objects allow the instance variables to refer to any other object. Pointer-format objects may also have named instance variables.

Subclass creation methods that create indexed classes with pointer objects begin with the keyword `indexableSubclass:`.

### Byte-format

This format is used for objects with indexed instance variables that are specialized for storing byte values, `SmallIntegers` in the range 0...255. Byte-format objects may not have named instance variables.

Subclass creation methods that create byte indexable classes begin with `byteSubclass:.`

You may not create byte-indexable subclasses of pointer-indexable classes, nor vice-versa, nor can you create indexable subclasses of NSCs.

### NonSequencableCollection (NSC)

These classes store data with neither names nor indexes. They are suited to applications in which access is by value, rather than by name or position. Classes with this format are subclasses of `UnorderedCollection`, and are the classes for which `Indexes` are implemented.

You cannot directly define classes with this format, although you can subclass from existing kernel classes. Subclasses of NSC classes may have named instance variables, but not indexed instance variables.

### Special

Instances of a number of kernel classes are encoded entirely in the object identifier. Special objects do not use up an object ID (i.e., are not in the object table), do not take up separate space in the repository (beyond the original reference itself), and equal values always compare as identical.

Specials include:

- Character, Boolean, UndefinedObject,
- SmallInteger, SmallDouble, SmallFraction,
- SmallScaledDecimal, SmallDate, SmallTime, SmallDateAndTime

Many of these have a limited range in which instances are special, and a corresponding non-special class for out of range values.

You may create a limited number of your own special Classes, by modifying template classes to encode your application data with 56 bits of the OOP value. This is described under "Customer-defined Special classes" on page 41.

## Class Variables and Other Types of Variables

The implementation formats defined in the last section define several types of instance variables. Class definitions also include the following variable types:

### Class variables

A class variable is a variable whose name and value are shared by a class, all of its instances, its subclasses, and all of their instances. Both class and instance methods of the class and its subclasses can refer to the variable. You can think of these variables as falling somewhere between local and global in their scope.

### Class instance variables

A class instance variable is a variable whose name and value are shared by a class, but not by its instances. Subclasses inherit the variable's name but not its value. Only class methods of a class and its subclasses can refer to class instance variables. Class instance

variables are useful when a class and its subclasses need to share the same structure, but not the same value, for a variable.

### Pool variables

The pool variables are an Array of SymbolDictionary instances that are searched when attempting to bind a variable name during instance method compilation. Pool variables come after class variables and before globals in precedence. They are typically used when methods in a number of classes share values.

For example, one could define a SymbolDictionary with a key of #'CR' and a value of (Character codePoint: 13). If this SymbolDictionary were included in the class definition as a pool dictionary, then instance methods in the class could use CR as a way to reference the value and make the code more readable.

### Global variables

Global variables are not tied to a class. They may be entries in a SymbolDictionary referenced in the UserProfile's SymbolList.

## Dynamic Instance Variables

In addition to the fixed instance variables, which are the same for every instance of that class, you may also add dynamic instance variables to most instances.

Dynamic instance variables are key/value pairs that are stored with the instance like other instance variables, but may be added to specific instances of a class and not to other instances, without changing the class definition.

You cannot add dynamic instance variables to invariant objects, nor to Specials, nor to classes or metaclasses.

The maximum number of dynamic instance variables that can be added to an object is 255. However, the maximum may be lower for classes with many instance variables, since an object cannot be changed to a large object by adding dynamic instance variables. So, more exactly, the actual limit for the number of dynamic instance variables is calculated:

```
(255 min: ((2034 - self class instSize) / 2)
```

To add a dynamic instance variable, set the value using:

```
anObject dynamicInstVarAt: nameSymbol put: value
```

For example, say you have an instance of Animal representing the Bald Eagle. Bald Eagles are an endangered species, so you might want to add the legal and conservation information to this instance, but not to other instances of Animals.

```
theBaldEagle dynamicInstVarAt: #legalStatus
    put: 'Bald and Golden Eagle Protection Act'.
```

You can check what dynamic instance variables have been defined for an object:

```
topaz 1> printit
theBaldEagle dynamicInstanceVariables
%
an Array
  #1 legalStatus
```

and retrieve the stored value for a dynamic instance variable:

```
topaz 1> printit
theBaldEagle dynamicInstVarAt: #legalStatus
%
Bald and Golden Eagle Protection Act
```

If the Bald Eagle was no longer protected and this information was no longer needed, you could remove the dynamic instance variable

```
theBaldEagle removeDynamicInstVar: #legalStatus
```

The name and data for dynamic instance variables are persisted in the repository like any other instance variable data. Dynamic instance variables allow you to add instance variables to instance of a class, without the need to migrate. However, dynamic instance variables are less efficient than named instance variables, and make for code that is more difficult to maintain.

Note that if you add a dynamic instance variable to an object, it does not impact any existing equality semantics. If the dynamic instance variable should be considered when determining if two objects are equal, you must add or update the implementation of = that applies for that object (which may also involve updating hash); this is not recommended for GemStone kernel classes, and should be done with caution if there may be existing instance, since it can cause collection lookup problems.

## Additional Class Creation Protocol

In addition to implementation format and variables, there are other features of classes that can be, or must be, defined when the class is created. These are provided via subclass creation methods with additional keywords.

The subclass creation methods follow the form in Example 2.2.

### Example 2.2

```
Object subclass: 'Animal'
  instVarNames: #('habitat' 'name' 'predator')
  classVars: #('AllOfSpecies')
  classInstVars: #('AllAnimals')
  poolDictionaries: #()
  inDictionary: UserGlobals
  newVersionOf: Animal
  description: 'Class describing Animals'
  options: #()
```

The `newVersionOf:` allows you to create a new class that has the same classHistory as an existing class; this is covered in detail in Chapter 10. See “Versions of Classes” on page 217.

The `description:` keyword allows you to provide documentation as part of the class definition. You can also explicitly set the comment after the class has been created by using the `comment:` method. For example:

```
Animal comment: 'Class describing Animal, created for the
Programmers Guide'.
```

The `options:` keyword allows you to specify a collection of symbols to defined specific features of the new subclass. The options can include any of these:

<code>#dbTransient</code>	See “DbTransient” on page 40 for details. This option cannot be used in combination with <code>#instancesNonPersistent</code> or <code>#instancesInvariant</code>
<code>#disallowGciStore</code>	For internal use
<code>#instancesInvariant</code>	All instances of this class will be made invariant as soon as they are committed. If any class is defined with <code>instancesInvariant</code> , all its subclasses must also have <code>instancesInvariant</code> .  Cannot be used in combination with <code>#instancesNonPersistent</code> or <code>#dbTransient</code>
<code>#instancesNonPersistent</code>	See “Non-Persistent Classes” on page 40 for details. This option cannot be used in combination with <code>#dbTransient</code> or <code>#instancesInvariant</code>
<code>#logCreation</code>	Log the creation of this class, including if a new class instance/version is not created, to the gem log or linked topaz output using <code>GsFile class&gt;&gt;gciLogServer:.</code> This legacy option is not persistent and not printed; to set class creation logging for all classes, use <code>SessionTemps current at: #GsClass_logCreation put: true.</code>
<code>#modifiable</code>	If this symbol is included, the class remains modifiable after creation. No instances can be created until you make the class unmodifiable by sending it the message <code>immediateInvariant</code> .
<code>#noInheritOptions</code>	If this symbol is included, it must be first, and in this case options are not inherited from the superclass nor from an existing version of the class. This applies to the options <code>#subclassesDisallowed</code> , <code>#disallowGciStore</code> , <code>#traverseByCallback</code> , <code>#dbTransient</code> , <code>#instancesNonPersistent</code> , and <code>#instancesInvariant</code>
<code>#selfCanBeSpecial</code>	This is needed only when modifying superclass hierarchy above classes with special format. It is never inherited.
<code>#subclassesDisallowed</code>	No subclasses of the newly created class are permitted.
<code>#traverseByCallback</code>	For internal use.

Note that some of these options are handled in specific ways when new versions of classes are created. Class versioning and history are described in Chapter 11.

For more details on class creation protocol, refer to methods in the image.

In addition to the subclass creation methods described here, there are many other subclass creation methods in the image, including methods with the keywords `inClassHistory:`, `isInvariant:`, `constraints:`, `isModifiable:`, and `instancesInvariant:`. These methods are deprecated, and should not be used, although they remain to avoid problems with filing older code into the current image.

## 2.2 Creating Classes With Invariant Instances

For data that must not ever be changed, GemStone provides two ways to make objects invariant or unchangeable. These are object-level invariance, and class-level invariance.

### Per-Object Invariance

Any object can be made invariant by sending it the message `immediateInvariant` (a method defined by class `Object`). This mechanism provides a form of write-protecting objects that is useful for maintaining the integrity of your database. Once `immediateInvariant` is sent to an object, no modifications can be made to any of the object's instance variables, nor can the size or class of the object be changed. The `immediateInvariant` message takes effect immediately, but can be reversed by aborting the transaction in which it was sent. Once the transaction has been committed, you cannot reverse the effect of this message. The message `isInvariant` returns `true` if the receiver is invariant; `false` otherwise.

### Invariance for All Instances of a Class

In class-level invariance, the definition of the class specifies that all instances of the class are invariant. Such an instance can be modified only during the transaction in which it is created. When the transaction is committed, the instance becomes invariant and no further modifications can be made to any of its instance variables, nor can the size or class of the object be changed. This mechanism is useful for supporting literals in methods and in other limited situations, but is generally more cumbersome than object-level invariance.

Class-level invariance can be specified during class creation by including the `#instancesInvariant` symbol in the `options:` keyword argument. You cannot also define the class with non-persistent instances (`#instancesNonPersistent`), nor with non-persistent instances variable data (`#dbTransient`).

The following example creates a subclass of `Animal` whose instances are invariant:

#### Example 2.3

```
Animal subclass: 'InvariantAnimal'  
  instVarNames: #()  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: #()  
  inDictionary: UserGlobals  
  options: (#instancesInvariant)
```

## 2.3 Creating Classes with Special Cases of Persistence

In some cases, you may want either objects or the instance variables of objects to not be persistent, that is, not be written to disk. For example, you may want to include session-dependent information that shouldn't be read by another session, or data that is bulky and can be recreated easily. There are several ways to handle this.

### Non-Persistent Classes

You can define a class as having only non-persistent instances. This means that instances of this class cannot be committed, so you cannot include references to instances of non-persistent classes within a persistent data structure.

To create a class with non-persistent instances, in the `options:` keyword argument, include the symbol `#instancesNonPersistent`. You cannot also define the class with non-persistent instances variables (`#dbTransient`), nor with invariant instances (`#instancesInvariant`).

As discussed under “KeySoftValueDictionary” on page 57, GemStone provides a class called `KeySoftValueDictionary`, which allows you to manage non-persistent objects that are large and take time to create, but can be recreated whenever needed from small, readily available objects (tokens).

You cannot commit instances of a non-persistent class. If you attempt to do so, GemStone issues an error that indicates whether the object's class or a superclass is non-persistent. (The non-persistent status of a class is inherited by all of its subclasses.)

To determine whether a class's instances are non-persistent, you can send the following message:

```
theClass instancesNonPersistent
```

This message returns true if the class is non-persistent, false otherwise.

To make all instances of a class non-persistent, send the message:

```
theClass makeInstancesNonPersistent
```

Similarly, send this message to make all instances of a class persistent:

```
theClass makeInstancesPersistent
```

To make all instances of a class (and all of its subclasses) non-persistent, even if the class is non-modifiable:

```
ClassOrganizer makeInstancesNonPersistent: theClass
```

Similarly, you can send this message to make all instances of a class persistent, even if the class is non-modifiable:

```
ClassOrganizer makeInstancesPersistent: theClass
```

### DbTransient

Classes can also be defined as `DbTransient`. Instances of classes that are `DbTransient` can be committed — that is, there is no error if they are committed — but their instance variables are not written to disk. This is useful if you need to encapsulate objects that should not be persistent, such as semaphores, within object structures that do need to be persistent and shared.

To create a class with DbTransient instances, in the `options:` keyword argument, include the symbol `#dbTransient`. You cannot also define the class with non-persistent instances (`#instancesNonPersistent`), nor with invariant instances (`#instancesInvariant`). When a data structure containing an instance of a DbTransient class is committed, the instance variables of the DbTransient object are written to the repository as nil. Whenever a DbTransient object is read into a session, all of its instance variables are nil.

Since DbTransient instances are stored only in memory, they are affected by the in-memory GC operations. (See “Managing VM Memory” on page 322. Also see Chapter 11 of the *System Administration Guide*.)

If memory becomes low, the transient objects may be stubbed out of memory. When needed, it is re-read from the repository. However, all the instance variables will be nil after a re-read. To prevent losing non-nil instance variable values, you should keep a reference to DbTransient instances in session state.

Since the DbTransient object will remain in memory while referenced from session state, the reference from session state should be removed when the DbTransient object is no longer needed, to avoid filling up memory and causing an out of memory error.

Note that while DbTransient objects are only committed once (on creation), and so do not normally cause concurrency conflicts, if they are clustered the object will be written (still with all instance variables nil), and could potentially cause a concurrency conflict.

To set a class so all instances are DbTransient, send:

```
aClass makeInstancesDbTransient
```

*aClass* must be a non-indexable pointer class. This will cause any instance of *aClass* to be DbTransient. The change takes place immediately.

The following message:

```
aClass makeInstancesNotDbTransient
```

will cause instances to be non-DbTransient, that is, allow instance variables to be written to disk.

## 2.4 Customer-defined Special classes

Special classes are classes in which all the object data is encoded in the OOP value itself. GemStone special classes include Character, SmallInteger, SmallDouble, SmallDateAndTime, and others.

GemStone also provides 16 predefined, customizable Special classes. Each can encode 56 bits of data. You will need to be convert the actual data you are storing into a 56 bit integer, and decode back into the data to support your setter and getter instance variables.

The available classes are named Special56bit0-Special56bit15.

To create an application-specific Special class, you will modify one of these class. It is legal to create an association to this class from the same or another SymbolDictionary, and to modify the GsObjectSecurityPolicy of the class. The new class will have the name you define, but you should leave the association with the original name in the Globals dictionary, for upgrade.

The Special classes include a number of primitives, which are defined on the class, so you are not limited to this location in the hierarchy. These methods are in a category 'Base

Methods' and should not be modified. You may add all necessary methods to other categories within the class to support the specific functions you need.

These are the steps to customize a Special class for your application:

1. Create an association in the target SymbolDictionary from the target name *classNameString* to the specific *Special56bitN* class.
2. As *SystemUser*, implement *Special56bitN class >> name*, and execute *specialClass >> changeNameTo: classNameString*. If you want to edit this class as a non-*SystemUser* user, you may change it's securityPolicy using:  

```
specialClass objectSecurityPolicy: anObjectSecurityPolicy
```
3. Bit-encode the specific data for your special into the 56 available bits, invoking the Class method *value:* to create an instance, and the instance method *value* to read and decode into your specific data type.

Note that care must be taken for fileout and filein, since the base class methods are a mixture of GemStone methods and application-specific methods. There are a number of ways to manage this, depending on your source management tools. The Money example shows a workaround using customizations to base image fileout.

It is allowed to change the superclass of a Special class; the required primitive methods are defined in each class, rather than inherited. The new superclass (and each superclass up through Object) must be created with the *#selfCanBeSpecial* class creation option.

## Money example

An example of using this feature is the example at

```
$GEMSTONE/examples/smalltalk/Money.gs
```

Which demonstrates using *Special56bit0* to implement a Special Money class encoding an amount and an integer currency ID to support US and Canadian Dollars, Euro and Yen. Note that this is an example class, and does not provide the necessary currency support for a real-world application.

This example must be filed in as *SystemUser*. It creates the Money class in the Published dictionary. Methods on the resulting Money class can be edited by a user with access to the *DataCuratorObjectSecurityPolicy*.

# Resolving Names and Sharing Objects

---

This chapter describes how GemStone Smalltalk finds the objects to which your programs refer and explains how you can arrange to share (or not to share) objects with other GemStone users.

**Sharing Objects** (page 43)

explains how GemStone Smalltalk allows users to share objects of any kind.

**The UserProfile's Symbol List** (page 44)

describes the mechanism that the GemStone Smalltalk compiler uses to find objects referred to in your programs.

**Using Your Symbol Dictionaries** (page 50)

discusses how you can enable other users of your application to share information.

## 3.1 Sharing Objects

GemStone Smalltalk permits concurrent access by many users to the same data objects. For example, all GemStone Smalltalk programmers can make references to the kernel class `Object`. These references point directly to the single class `Object` – not to copies of `Object`.

GemStone allows shared access to objects without regard for whether those objects are files, scalar variables, or collections representing entire databases. This ability to share data facilitates the development of multi-user applications.

To find the object referred to by a variable, GemStone follows a well-defined search path:

1. The local variable definitions: temporary variables and arguments.
2. Those variables defined by the class of the current method definition: instance, class, class instance, or pool variables.
3. The symbol list assigned to your current session.

If GemStone cannot find a match for a name in one of these areas, you are given an error message.

## 3.2 The UserProfile's Symbol List

Each GemStone user is associated with an instance of the class UserProfile. This UserProfile stores such information as the GemStone user name, the encrypted password, and access privileges. Your UserProfile also contains the instance variable *symbolList*.

This UserProfile can be used to login multiple sessions at the same time, including logins by different people in different locations. The way the symbolLists are handled includes options to accommodate this situation.

### What's In Your Symbol List?

When creating the UserProfile's symbol list, the administrator adds specific SymbolDictionaries to the SymbolList. The SymbolDictionaries contain associations that define the names of all objects that are globally resolvable by your UserProfile. While this will vary by application, your symbol list contains at least two dictionaries:

- ▶ A "system globals" dictionary called *Globals*. This dictionary contains some or all of the GemStone Smalltalk kernel classes (Object, Class, Collection, etc.) and any other objects to which all of your GemStone users need to refer. Although you can read the objects in *Globals*, normally only SystemUser is permitted to modify them.
- ▶ A dictionary that is specific to your UserProfile, called *UserGlobals*. This dictionary can be used to store objects for your own use and new classes you do not need to share with other GemStone users. This dictionary also contains GemStone infrastructure that is UserProfile specific.

Your symbol list will usually include other application-specific dictionaries to hold the code you are developing. These may be shared with other users, so that you can all read and modify the objects they contain. An administrator can arrange for a dictionary to be shared by inserting a reference to that dictionary in each user's UserProfile symbol list.

While every user will have the shared Globals and a private UserGlobals dictionary, and by default most users will have the Published dictionary, the list of SymbolDictionaries in each user's SymbolList may otherwise be completely different.

### Examining Your Symbol List

To get a list of the dictionaries in your symbol list, send your UserProfile the message `dictionaryNames`. For example:

#### Example 3.1

```
topaz 1> printit
System myUserProfile dictionaryNames
%
 1 UserGlobals
 2 ClassesForTesting
 3 Globals
 4 Published
 5 UserClasses
```

The SymbolDictionaries listed in the example have the following function:

- ▶ **UserGlobals**  
Contains per-user application and application service objects.
- ▶ **ClassesForTesting**  
A user-defined dictionary.
- ▶ **Globals**  
Provides access for the GemStone kernel classes.
- ▶ **Published**  
Provides space for globally visible shared objects created by a user.
- ▶ **UserClasses**  
Usually only present if you are using GemBuilder for Smalltalk (GBS) to replicate classes to the server. Putting this dictionary before the Globals dictionary allows an application or user to override kernel classes without changing them. Keeping it separate from UserGlobals allows a distinction between classes and application objects.

To list the contents of a symbol dictionary:

- ▶ If you are using Topaz, execute an expression that returns the dictionary. Example 3.2 lists the dictionary keys. Alternatively, you could just execute `UserGlobals` to examine all keys and values.
- ▶ If you are running GBS, select the expression `UserGlobals` in a GemStone workspace and execute `GS-Inspect it`.

### Example 3.2

---

```
topaz 1> printit
UserGlobals keys
%
a SymbolSet
. . .
#1 GcUser
#2 UserGlobals
#3 GsPackagePolicy_Current
#4 PackageLibrary
. . .
```

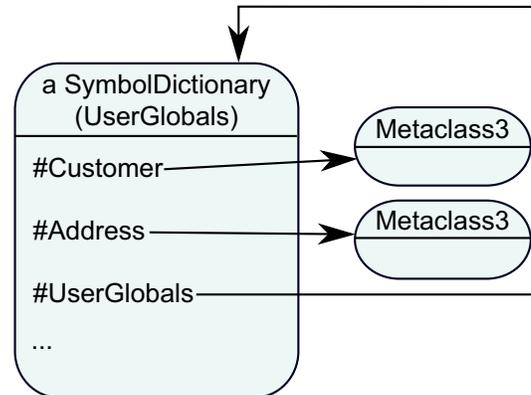
---

If you examine all of your symbol list dictionaries, you'll see that most of the kernel classes are listed. In addition, there are global variables, both public and for internal use. For a description of GemStone kernel objects, see the appropriate appendix of the *System Administration Guide*.

You'll discover that most of the dictionaries refer to themselves. Since the symbol list must contain all source code symbols that are not defined locally nor by the class of a method, the symbol list dictionaries need to define names for themselves so that you can refer to them in your code. Figure 3.1 illustrates that the dictionary named `UserGlobals` contains an association for which the key is `UserGlobals` and the value is the dictionary itself.

The object server searches symbol lists sequentially, taking the first definition of a symbol it encounters. Therefore, if a name, say “#BillOfMaterials,” is defined in the first dictionary and in the last, GemStone Smalltalk finds only the first definition.

**Figure 3.1 Self-Referencing Symbol Dictionary**



You can get the name of a SymbolDictionary by sending name. For example,

```

topaz 1> printit
System myUserProfile symbolList collect: [:ea | ea name]
%
a Array
  #1 UserGlobals
  #2 Globals
  #3 Published
  
```

## Inserting and Removing Dictionaries from Your Symbol List

Note that, to insert or remove a SymbolDictionary to/from your symbol list, you must have the necessary system privilege. For details, see "User Accounts and Security" in the *System Administration Guide*.

Creating a dictionary is like creating any other object, as the following example shows. Once you've created the new dictionary, you can add it to your symbol list by sending your UserProfile the message `insertDictionary: aSymbolDict at: anInt`.

For example:

```

System myUserProfile symbolList
  createDictionaryNamed: #NewDict at: 1.
  
```

When you specify an index for a new SymbolDictionary, the existing symbol list dictionaries are shifted as needed to accommodate the new dictionary.

Because the GemStone Smalltalk compiler searches symbol lists sequentially, taking the first definition of a symbol it encounters, your choice of the index at which to insert a new dictionary is significant.

The following example places the object `MyCollection` (a class) in the user's private dictionary named `MyClassDict`. Then it inserts `MyClassDict` in the first position of the current Session's `symbolList`, which causes the object server to search `MyClassDict` prior

to UserGlobals. This means that the GemStone object server will always find MyCollection in MyClassDict, not in UserGlobals.

### Example 3.3

```
Object subclass: 'MyCollection'
  instVarNames: #('snakes' 'snails' 'tails')
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals

! Resolves to UserGlobals
MyCollection instVarNames printString

| myClassDict |
(System myUserProfile resolveSymbol: #MyClassDict)
  ifNil:[myClassDict := (System myUserProfile createDictionary:
    #MyClassDict)]
  ifNotNil:[myClassDict := (System myUserProfile resolveSymbol:
    #MyClassDict) value].
GsSession currentSession userProfile
  insertDictionary: myClassDict at: 1.

Object subclass: 'MyCollection'
  instVarNames: #('this' 'that' 'theOther')
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: MyClassDict.

! Now resolves to different class in MyClassDict
MyCollection instVarNames printString
```

Recall that the object server returns only the *first* occurrence found when searching the dictionaries listed by the current session's symbol list. When you subsequently refer to MyCollection, the object server returns only the version in MyClassDict (which you inserted in the first position of the symbol list) and ignores the version in UserGlobals. If you had inserted MyClassDict *after* UserGlobals, the object server would only find the version of MyCollection in UserGlobals.

You may redefine any object by creating a new object of the same name and placing it in a dictionary that is searched before the dictionary in which the matching object resides. Therefore, inserting, reordering, or deleting a dictionary from the symbol list may cause the GemStone object server to return a different object than you may expect.

This situation also happens when you create a class with a name identical to one of the kernel class names.

Avoid redefining any kernel classes in other SymbolDictionaries. Their implementation may change from one version of GemStone to the next. Creating a subclass of a kernel class to redefine or extend that functionality is usually more appropriate.

To remove a symbol dictionary, send your UserProfile the message `removeDictionaryAt: anInteger`, passing in the index of the dictionary you want to remove.

## Finding Out Which Dictionary Names an Object

To find out which dictionary defines a particular object name, send your UserProfile the message `symbolResolutionOf: aSymbol`. If `aSymbol` is in your symbol list, the result is a string giving the symbol list position of the dictionary defining `aSymbol`, the name of that dictionary, and a description of the association for which `aSymbol` is a key. For example:

### Example 3.4

---

```
topaz 1> printit
System myUserProfile symbolResolutionOf: #Bag
%
2 Globals
  Bag Bag
```

---

If `aSymbol` is defined in more than one dictionary, `symbolResolutionOf:` finds only the first reference.

To find out which dictionaries stores a name for an object and what that name is, send your UserProfile the message `dictionariesAndSymbolsOf: anObject`. This message returns an array of arrays containing the dictionaries in which `anObject` is stored, and the symbols which name that object in that dictionary.

Example 3.5 uses `dictionariesAndSymbolsOf:` to find out which dictionaries in the symbol list stores a reference to class `DateTime`.

### Example 3.5

---

```
topaz 1 > printit
| anArray myUserPro |
myUserPro := System myUserProfile.

"Find the first SymbolDictionary containing DateTime."
anArray := (myUserPro dictionariesAndSymbolsOf: DateTime) first.

"Get the name of the SymbolDictionary, which is a key within
itself"
(anArray at: 1) keyAtValue: (anArray at: 1)
%
```

---

 Globals
 

---

Note that `dictionariesAndSymbolsOf:` may return zero, one, or multiple dictionaries.

## The Transient Symbol List

Since the persistent `symbolList` is shared between all sessions that use that `UserProfile` to login, you must use caution in making updates that you don't want to be picked up by other logins using the same `UserProfile`.

This is possible by using the transient `symbolList`. The transient `symbolList` is kept in the singleton instances `GsCurrentSession`, which you can access using `GsSession currentSession` or `GsCurrentSession >> currentSession`. The instance of `GsCurrentSession` is not copied into any client interface nor committed as a persistent object. Since the transient copy of the `symbolList` is transient, changes to it cannot incur concurrency conflicts, nor are they subject to rollback after an abort.

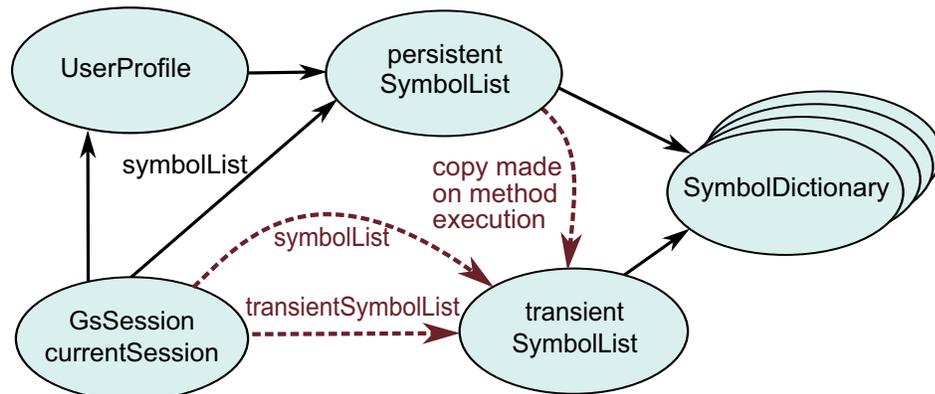
Executing any of the following methods will create a copy of the transient `symbolList` in session state:

```
System class >> refreshTransientSymbolList
GsCurrentSession >> transientSymbolList: aSymbolList
GsCurrentSession >> transientSymbolList
```

Before executing any of these methods, `GsSession currentSession symbolList` refers to the persistent `SymbolList` (`System myUserProfile symbolList`).

After executing any of these methods, `GsSession currentSession symbolList` refers to a transient copy of the persistent `SymbolList`.

**Figure 3.2** The `GsSession symbolList` – a copy of the `UserProfile symbolList`



Changes to the current session's `symbolList` (the transient `symbolList`) do not affect the `UserProfile symbolList` (the persistent `symbolList`). Thus, the `UserProfile symbolList` can continue to serve as a default list for other logins.

## Updating Symbol Lists

`GsCurrentSession >> symbolList` is used to resolve global names in the image. If your application may use transient symbolLists, this is the most reliable route to determine how symbols will resolve in the current session, since it will provide the persistent symbolList if there is no transient symbolList, otherwise the current transient symbolList.

When you have a transient symbolList (that is, if you have executed any of the methods listed on page 49), then you will need to be conscious of how you are updating the symbolList, to make sure you get the expected behavior.

- ▶ If you add a new symbolDictionary to the persistent symbolDictionary but not to the transient symbolDictionary, and these are different instances, then the new symbolDictionary is not resolvable using `GsCurrentSession >> symbolList`, and so not readily usable by name. This is the case, for example, if you make changes to the result of `System myUserProfile symbolList`. after executing `GsCurrentSession transientSymbolList`.
- ▶ If you add a new symbolDictionary to the transient symbolDictionary but not to the persistent symbolDictionary, the new symbolDictionary will be available for the session, but will disappear when the session logs out. This expression will insert a dictionary into the transient symbolList, and create the transient symbolList if it was not previously different than the persistent symbolList:

```
GsSession currentSession transientSymbolList
createDictionaryNamed:at:
```

- ▶ If you add a new symbolDictionary to both the transient symbolDictionary and to the persistent symbolDictionary, and you abort your transaction, the new symbolDictionary will disappear from the persistent symbolDictionary, but remain in the transient symbolDictionary until the session logs out.

If you want to make changes to both the persistent symbolList, and the transient symbolList, methods such as this update both the persistent and transient symbolLists.:

```
System myUserProfile insertDictionary:at:
```

## 3.3 Using Your Symbol Dictionaries

As you know, all GemStone users have access to such objects as the kernel classes Integer and Collection because those objects are referred to by the Globals dictionary that is present in every user's symbol list.

If you want GemStone users to share other objects as well, you need to arrange for references to those objects to be added to the users' symbol lists.

### NOTE

*To insert or remove a SymbolDictionary to/from your symbol list, or to make any changes to a UserProfile that is not your own, you must have the necessary system privilege. For details, see "User Accounts and Security" in the System Administration Guide.*

## Publishers, Subscribers and the Published Dictionary

The Published Dictionary, PublishedObjectSecurityPolicy, and the groups Subscribers and Publishers together provide an example of how to set up a system for sharing objects.

The Published Dictionary is an initially empty dictionary referred to by your UserProfile. You can use the Published dictionary to "publish" application objects to all users – for example, symbols that most users might need to access. The Published Dictionary is not used by GemStone classes; rather, it is available for application use.

The PublishedObjectSecurityPolicy is owned by the Data Curator and has World access set to none. Two groups have access to the PublishedObjectSecurityPolicy:

- ▶ Subscribers have read-only access.
- ▶ Publishers have read-write access.

Publishers can create objects in the PublishedObjectSecurityPolicy and enter them in the Published Dictionary. Then members of the Subscribers group can access the objects.

For example, your system administrator might add each member of a programming team to the group Publishers. After completing the definition of a new class, a programmer could make the class available to colleagues by adding it to the Published dictionary. Because this dictionary is already in each user's symbol list, whatever you add becomes visible to users the next time they obtain a fresh snapshot view of the repository. Using the Published dictionary lets you share these objects without having to put them in Globals, which contains the GemStone kernel classes, and without the necessity of adding a special dictionary to each user's symbol list.



# Collection and Stream Classes

---

Collections of objects are key features in an application. GemStone provides a variety of Collection classes, including both subclasses of Collection, and other implementations of structures with collection semantics. This chapter describes the main types of collections that are available.

Strings and ByteArrays are kinds of collections that are specialized to hold characters or bytes; these are described separately in Chapter 5.

**Introduction to Collections** (page 53)

introduces the GemStone Smalltalk objects that store groups of other objects, and describes the different kinds of collections that are available.

**Reduced-Conflict Collection Classes** (page 63)

describes specialized kinds of classes that avoid conflicts in a multi-user system.

**GsBitmap** (page 66)

describes GsBitmap, a specialized kind of collection.

**Sorting the objects in a collection** (page 69)

describes the ways to sort elements in collections.

## 4.1 Introduction to Collections

Instances of the Collection subclasses are specialized to manage an indeterminate number of objects as a group using unnamed instance variables.

Collections can be classified by whether or not they maintain a specified order for their elements, whether or not key-based lookup is supported, and the kinds of objects they can reference.

Collections can be broadly classified into basic categories:

▶ **Access by Key** – the Dictionary Classes

Instances of AbstractDictionary subclasses do not support a specific order for their elements; elements are stored and retrieved via the `at : put :` and `at : messages`, using

arbitrary objects for an element's key. Subclasses of `AbstractDictionary` are specialized based on whether key-based lookup uses equality comparison or identity comparison, the type of key, and the type of value.

Dictionaries can also have named instance variables, if you choose to define them.

▶ **Access by Position** – the `SequenceableCollection` Classes

Instances of `SequenceableCollection` classes maintain a specific order for their elements and support storage and retrieval via the `at : put :` and `at :` messages using an integer key (the one-based offset into the elements).

Byte-format classes such as `ByteArray` and `String` cannot have named instance variables. You may define named instance variables for pointer-format subclasses, such as `Array` and `OrderedCollection`.

▶ **Access by Value** – the `UnorderedCollection` Classes

Instances of `UnorderedCollection` classes – also referred to as `Non-Sequenceable Collections` or `NSCs` – do not have a specific order for their elements, and do not support storage or retrieval via the `at : put :` and `at :` messages. Objects in these collections are accessed by iterating the collection. `UnorderedCollections` support indexes, which allow ordered iteration and fast key-based lookup.

You may define named instance variables for subclasses of one of the `UnorderedCollections` subclasses.

▶ **Other kinds of Collection**

GemStone includes some collection-like classes that do not inherit from `Collection`, such as `GsBitmap`. The generalizations about collections made in this chapter may or not apply to such classes.

## Efficient Implementations of Large Collections

When you create a collection of more than about 2K objects, or a byte collection larger than 16K, GemStone internally uses a sparse tree implementation to make more efficient use of resources. These are referred to as "Large Objects", and use internal classes such as `LargeObjectNode`. This behavior occurs in a manner that is transparent to you, and you interact with Large Collections the same as smaller collections; however, these internal objects may be visible when performing object-based audit and analysis.

## Modifying objects in collections

Different kinds of collections use different criteria in which to store and locate objects. Most dictionaries use a hashed value, the result of sending the `#hash` message to each key object. If the result of sending `#hash` changes, the key may not be found in the collection using methods such as `at :`.

The ordering of a `SortedCollection` depends on the results of sending comparison methods to the objects in the collection.

If a hash method or comparison method is defined that depends on values in the instance variables of the objects, and these instance variables are modified, the object may not be found using lookup methods in the dictionary or sorted collection, though iteration will still find them. If you change the value of an instance variable of an object in such a collections, you should remove and re-insert the object in the collection so the lookup methods will work.

## Protocol Common to All Collections

Collection classes understand common protocol, inherited from the abstract superclass `Collection`. `Collection` defines methods that enable you to perform the general collection operations described in the following sections.

### Creating Instances

Collections can be created using `new`, `new:`, `with:`, and similar protocol. The most basic way to create a new collection is using the message `new`. When sent to a `Collection` class, this message causes a new instance of the class with no elements (size zero) to be created. Most kinds of collections can expand as you add additional objects.

`new: anInteger`, causes many `Collection` subclasses to create an instance that is pre-sized to hold `anInteger` elements. This avoids the need to expand the collection when elements are added. Pre-defining the size during creation is particularly important when creating a hashed collection that will hold a large number of objects. Hashed collections store elements in buckets, and the number of buckets must be increased when the number of objects in the collection reaches a threshold for the number of buckets. These expansions are expensive, since it requires that each element be re-added to the expanded collection at the recomputed hashed location.

Several kinds of `Collections` can be created as literals, using Smalltalk syntax. `Arrays`, `ByteArrays`, `Strings` and `Symbols` have literal syntax, and `Arrays` can also be created at runtime using `Array` constructors. `ByteArrays`, `Strings` and `Symbols` are discussed in Chapter 5.

### Enumerating

`Collection` defines several methods that enable you to loop through a collection's elements, evaluating a block for each element in the collection.

- ▶ The message `do: aBlock` is the most general; it evaluates `aBlock` for each element.
- ▶ Methods that iterate through the elements and return collections are `collect:`, `select:`, and `reject:`.

The class of the result collection is often, but not always, the same kind of collection as the receiver. The class methods `species`, `speciesForSelect`, and `speciesForCollect` determine the class of the result.

When sent to `SequenceableCollections`, these messages preserve the ordering of the receiver in the result. That is, if element *a* comes before element *b* in the receiver, then element *a* will come before *b* in the result.

- ▶ The messages `detect:`, `detect:ifNone:`, and `any` iterate to return a single element, based on the order the collection is enumerated. Enumeration stops after the an element is found.
- ▶ The message `anySatisfy:` enumerates each element, stopping if any is found that meet the block criteria; `allSatisfy:` enumerates, stopping if any is found that does not meet the block criteria.

To avoid unpredictable consequences, do not add elements to or remove them from a collection while you are enumerating it.

## Collections in multi session environment

In many cases, you will have collections that need to be accessed by multiple sessions in an application. Different sessions may need to read the contents or to add, remove, or modify elements.

### Conflicting updates

Due to the transactional nature of GemStone (see Chapter 9, “Transactions and Concurrency Control”), overlapping updates by two sessions may conflict, in which case the second update has failed and the work needs to be repeated.

There are some kinds of conflicts that, while they modify the same object, are not really conflicts in a logical sense. GemStone provides a number of different collection classes that avoid specific kinds of conflicts; these are described under “Reduced-Conflict Collection Classes” on page 63.

### Visibility and ordering

When you add an element to a collection, this change does not become visible to other users until you have successfully committed the transaction. While it is important in a multiuser system to avoid long periods in a transaction prior to commit, the requirements are application specific and there may be minutes or hours between the time an object is created and when it is finally committed and visible to other users. Other users, in turn, must abort or commit before they see the changes.

For ordinary (non-reduced conflict) collections, this means that object changes may become visible to other users some time after the change is actually made.

For reduced-conflict sequenceable and queue classes, the order of objects in the collection may depend on the order of the commit, rather than the order of the time in which the object was created.

## Collection classes

Collection classes can be grouped by the kinds of access methods they provide and the kinds of objects their instances can store.

- ▶ Dictionary classes, including Dictionary, KeyValueDictionaries, and KeySoftValueDictionary
- ▶ SequenceableCollection classes, including Array, OrderedCollection, and SortedCollection,
- ▶ UnorderedCollection classes, including Bag, IdentityBag, Set, and IdentitySet
- ▶ Stream Classes, including ReadStream, WriteStream, and ReadWriteStream.
- ▶ Reduced-Conflict Collection Classes, including RcArray, RcIdentityBag, RcIdentitySet, RcKeyValueDictionary, RcPipe, and RcQueue.
- ▶ Classes that aren't subclasses of Collection, including GsBitmap.

String classes, including Traditional and Unicode string classes and Symbols, are also kinds of Collections and understand many kinds of Collection messages. Concerns that are specific to Strings, including String collation, are described in Chapter 5.

This chapter does not attempt to describe all collection classes or all methods that are available; it highlights the most commonly used protocol and describes special features. Review the methods in the image for more details.

## Dictionary classes

Dictionaries provide their special facilities by storing key-value pairs instead of simple, linear lists of objects. The elements in a Dictionary collection are stored and accessed via a key; each key must be unique within that Dictionary.

While some types of dictionaries are implemented as “a collection of Associations”, the interface methods return results based on the logical contents, which are the values. Other, specialized protocol allows you to refer to the key or the value portions of the logical associations.

### Internal Dictionary Structure

For performance reasons, the internal implementation of Dictionary classes varies. Instances of Dictionary itself consist of a collection of Association objects.

KeyValueDictionary subclasses are implemented differently, as a sequence of keys and values, which may use CollisionBuckets to hold the actual values. IdentityDictionary is a sequence of keys and Associations. All these dictionaries understand common protocol, regardless of implementation.

### Dictionary and KeyValueDictionary

Dictionary class uses Associations to store the key/value pair, while subclasses of KeyValueDictionary are slot-based. KeyValueDictionary has several subclasses, divided according to the type of key used to access the information:

- ▶ IdentityKeyValueDictionary
- ▶ IntegerKeyValueDictionary
- ▶ StringKeyValueDictionary
- ▶ SymbolKeyValueDictionary
- ▶ IdentityDictionary
- ▶ SymbolDictionary

### KeySoftValueDictionary

A KeySoftValueDictionary is a subclass of KeyValueDictionary that allows the virtual machine to remove entries as needed to free up memory.

Typically, you might use a KeySoftValueDictionary to manage non-persistent objects that are large and take time to create, but that can be recreated whenever needed from small, readily available objects (tokens). For example, you might create a KeySoftValueDictionary to serve as a cache to hold large, expensive objects that are needed repeatedly. Within that dictionary, the values would be the large calculated objects, and the keys would be the corresponding tokens. If your application needs a large, expensive object but does not find it in the KeySoftValueDictionary, you can create the object and add it to the cache so that it might be available the next time it is needed.

As memory fills up, the virtual machine might remove some objects from the cache. (Remember, the contents of the cache are non-persistent and can be recreated.) The virtual machine may remove keys and values from the `KeySoftValueDictionary` until adequate memory is available. For details about how to manage the number of `KeySoftValueDictionary` entries, see “Getting Rid of Non-Persistent Objects” on page 321. Keep in mind the following:

- ▶ Entries are removed from a `KeySoftValueDictionary` only if there are no strong references to the entry’s value.
- ▶ If an entry in a `KeySoftValueDictionary` is cleared, all other entries that reference this value directly or indirectly will also have been cleared.
- ▶ Before generating an `OutOfMemory` error, the virtual machine removes all `KeySoftValueDictionary` entries that are eligible for removal.
- ▶ `KeySoftValueDictionary` entries are cleared during a mark/sweep operation, but are not cleared during a scavenge. For more about mark/sweep and scavenge operations, see the “Managing Growth” chapter of the *System Administration Guide*.
- ▶ A corresponding subclass, `IdentityKeySoftValueDictionary`, uses identity (rather than equality) comparison on keys. For details, see the image.
- ▶ A `KeySoftValueDictionary` frequently contains instances of `SoftReference`. Do not be tempted to confuse this with the notion of `WeakReference` found in many Smalltalk dialects; the two mechanisms are quite different.

## SequenceableCollection classes

`SequenceableCollections`, such as `Array` and `OrderedCollection`, let you refer to their elements with integer indexes, and they understand messages such as `first` and `last` that refer to the order of those indexed elements. Adding by default adds to the end of the collection.

### Copying

When copying a very large instance of a subclass of `SequenceableCollection`, it can be more efficient to use the method `replaceFrom:to:with:startingAt:`, which does not fault the contents into memory. This can improve performance significantly for very large collections.

This example copies two elements of an array into a different array, overwriting the target array’s original contents:

```
| numericArray |
numericArray := Array with: 55 with: 66 with: 77 with: 88.
numericArray replaceFrom: 2 to: 3
               with: #( 1 2 3 4 5 ) startingAt: 4.
numericArray
%
an Array
#1 55
#2 4
#3 5
#4 88
```

Note that, while the `replace` method does not itself fault the contents into memory, displaying the results as in the example also faults the objects into memory.

## Array

One of the most important differences between client Smalltalk arrays and a GemStone Smalltalk array is that GemStone arrays are extensible; you can increase the size of an array at any time. Sending `at:put:` will increase the size of the array, as long as the index is only one greater than the current array size. Other protocol such as `addAll:` also increase the size while adding elements.

It's also possible to change the size without explicitly storing or removing elements, using the message `size:` inherited from class `Object`. When you lengthen an array with `size:`, the new elements are set to `nil`.

## Literal Array and Array Constructors

Arrays can also be created in code without sending instance creation messages, by using literal array or array constructor syntax.

Since Array constructors perform code at runtime, it is more efficient to use Array literals if the contents are literals.

**Array Literals** are created at compile time, and hold other literal objects. These start with the pound sign, are enclosed in parenthesis and separated by white space. Array literals are defined by ANSI; syntax is described on page 396. They are invariant.

```
#( 'carrot' 'tomato' 'celery' )
```

**Array constructors** are created at runtime. These are enclosed in curly braces and separated by a period. Array constructors are GemStone-specific, not defined by ANSI; syntax is described on page 403.

```
{ Date today . Time now }
```

## SortedCollection

`SortedCollection` is a type of `SequenceableCollection` in which the elements are ordered by a specific sort order, not by the order in which they were added or by the method used to add the element. You may not send `at:put:`, `addLast:`, or similar methods to a `SortedCollection`.

Each instance of `SortedCollection` is associated with a `sortBlock`. The default block will sort elements that can be compared using `<=`, which includes strings and numbers. You can also define your own `sortBlock`, if you want elements ordered by some other criteria, such as the value of an instance variable.

For more on comparison, sorting, and sort blocks, see "Sorting the objects in a collection" on page 69.

**Example 4.1**


---

```

| scrabbleWords |
scrabbleWords := SortedCollection sortBlock:
    [:a :b | a size < b size].
scrabbleWords add: 'able'; add: 'zebra'; add: 'jumper';
    add: 'yet'.
scrabbleWords
%
aSortedCollection( 'yet', 'able', 'zebra', 'jumper')

```

---

There is overhead in always keeping the collection sorted, so it usually more efficient to sort the elements only when you need them to be sorted for presentation. Especially for large collections or collections in which objects are frequently added and removed, consider using another kind of class to store the elements, then using methods such as `sortWithBlock:` to create a new Array with the elements in sorted order.

SortedCollection sortBlocks are compiled code, and as such, may need to be recompiled on GemStone upgrade. Provided the sortBlock is simple – that is, it does not contain references to variables outside the scope of the block, nor iterative methods – the recompile can be done automatically. Since the sortBlock executes for many element pairs during sort, keeping the sortBlock simple and fast is important for performance in any case.

**Stream Classes**

A Stream acts like a SequenceableCollection that keeps track of the index most recently accessed. Streams are often used for reading characters from strings or files, but any kind of collection can be used with a Stream, and any type of object can be in that collection.

Commonly used Stream classes are ReadStream, WriteStream, and ReadWriteStream, which come in two variants; the traditional Smalltalk 1-based positional offset, and the ANSI-compliant portable streams with an 0-based offset.

**PositionableStream and Position**

PositionableStream, with its subclasses ReadStream and WriteStream, was traditionally implemented in GemStone with the position indicating an offset from 1; that is, the first position in the stream was 1.

ANSI specifies, and other Smalltalk dialects use, an offset of 0, so the first position in the stream is 0.

To allow both sets of classes to be available for use, while either one or the other uses the actual class name, GemStone includes the multiple sets of classes, implementing both interfaces. There are four sets of classes, which all exist in the image (and therefore, may have instances), with only three sets being visible at any time. The following two sets are always visible:

- ▶ Legacy-style PositionableStream classes, compatible with previous GemStone version's PositionableStream classes:

```

PositionableStreamLegacy
    ReadStreamLegacy
    WriteStreamLegacy

```

- ▶ ANSI-compliant and portable PositionableStream classes:

```
PositionableStreamPortable
  ReadStreamPortable
  WriteStreamPortable
  ReadWriteStreamPortable
```

In addition, only one of the following sets is visible, depending on how your system is configured. These are two distinct sets of instances of Class, with the same name, but different implementations.

```
PositionableStream (with legacy definition and methods)
  ReadStream
  WriteStream
PositionableStream (with portable definition and methods)
  ReadStream
  WriteStream
```

Note that there are other classes within the PositionableStream hierarchies; these examples are simplified for clarity.

The legacy versions are stored in Globals at: #GemStone\_Legacy\_Streams. The portable, ANSI-compatible versions are stored in Globals at: #GemStone\_Portable\_Streams.

To check what is currently installed, use the following methods:

```
PositionableStream class >> isLegacyStreamImplementation
PositionableStream class >> isPortableStreamImplementation
```

To install the portable version, use the method:

```
Stream class >> installPortableStreamImplementation
```

To install the legacy version, use the method:

```
Stream class >> installLegacyStreamImplementation
```

## AppendStream

AppendStream is a kind of Stream that does not maintain a position. It is designed to optimize a common use-case for streams: composing long, complex blocks of text and returning the resulting string.

Like WriteStream, you can add strings and characters to an AppendStream, and like any stream, you can get the entire contents. Many other methods commonly associated with Stream classes are not available, however.

## ReadByteStream

ReadByteStream is a kind of Stream that is optimized for reading strings and ByteArrays.

## UnorderedCollection classes

Instances of UnorderedCollection store their elements as an internal, tree-based structure referred to as a Non-Sequenceable Collection (NSC). The elements have no defined order within the collection, so methods such as `at :` and `at :put :` are disallowed.

UnorderedCollection implements protocol for indexing, which allows for large collections to be queried and sorted efficiently. Chapter 8, “Indexes and Querying”, describes the querying/sorting functions in detail. The most efficient way to handle very large collections is using UnorderedCollection, using GemStone indexes to access the contents.

UnorderedCollections cannot contain nil as an element; adding nil has no effect.

Commonly used UnorderedCollection concrete classes are Bag, Set, IdentityBag and IdentitySet. Since Bag and Set use equality for comparisons, for large collections it is much more efficient to use IdentityBag or IdentitySet, which perform comparisons based on identity (OOP).

## Union, Intersection, and Difference

Subclasses of UnorderedCollection provide messages that perform set arithmetic: union, set intersection, and set difference.

- + union, returning elements that are in either one, the other, or both.
- difference, returning elements that are in the receiver but not the argument.
- \* intersection, returning elements that are in both

### Example 4.2

---

```
| pets rodents |
pets := IdentityBag with: 'dog' with: 'cat' with: 'gerbil'.
rodents := IdentityBag with: 'rat' with: 'gerbil' with: 'beaver'.
pets * rodents
%
  anIdentityBag( 'gerbil')

pets + rodents
%
  anIdentityBag( 'beaver', 'rat', 'gerbil', 'gerbil', 'cat', 'dog')

pets - rodents
%
  anIdentityBag( 'cat', 'dog')
```

---

## Avoiding faulting contents into memory

If the argument to `addAll:` is an Array or OrderedCollection, the elements in the collection are not faulted into memory. For very large collections, or if the objects in the collection are not in the shared page cache and must be read from disk, this can be a significant advantage. Using an IdentityBag as an argument to `replaceFrom:to:with:startingAt:` allows you to get a copy of the elements without faulting the contents into memory.

### Example 4.3

---

```
| bagOfRodents |
bagOfRodents := IdentityBag withAll: #('beaver' 'rat' 'agouti'
  'chipmunk' 'guinea pig').
(Array new: 5) replaceFrom: 3 to: 5
  with: bagOfRodents startingAt: 1.
  anArray( nil, nil, 'guinea pig', 'chipmunk', 'agouti')
```

---

## 4.2 Reduced-Conflict Collection Classes

GemStone provides a variety of reduced-conflict collection classes. These classes are similar to the standard collection classes already described, but include additional processing to avoid transaction conflicts in a multi-user environment.

Each reduced-conflict class has specific types of conflicts they are designed to avoid, and the amount of internal infrastructure or the cost of resolving a conflict varies. Selection of an RC class should consider the demands of the application, and also the costs of the automatic conflict resolution.

For more on transactions and transaction conflicts, see Chapter 9. Further information on the transactional behavior of these RC classes is under the section “Classes That Reduce the Chance of Conflict” on page 185.

### RcArray

The class RcArray is similar to Array, but no conflict occurs when multiple users add objects to an RcArray. If a conflict with another update operation on the RcArray occurs, the add is replayed so that the commit can succeed.

Only the following methods support concurrent updates:

- add:
- addAll:
- at:put: (where no other session affects the element at the at: index)
- size: (when size is increased)

### NSC/UnorderedCollection classes

#### RcIdentityBag

The class RcIdentityBag provides much of the same functionality as IdentityBag, but with no conflict for multiple sessions that add objects to the bag, and a single session that removes objects.

#### Internal implementation

RcIdentityBag is internally implemented using an Array of IdentityBags. Each session number corresponds to two IdentityBags, one for additions to the RcIdentityBag, and one for removed elements. Each logged-in session only modifies the IdentityBags corresponding to its own session number. Computing the current contents of an RcIdentityBag means combining the add bags, and removing all the remove bags.

#### Maintenance

The implementation of RcIdentityBag means that reclaiming the storage of objects that have been removed from the bag actually occurs when a session performs later adds or removes, or after that session logs out, another session logs in as that session number and performs adds or removes.

If a session adds a great many objects to the RcIdentityBag, and then does not do any further adds or removes; or if it logs out and the following sessions to use that session number do not perform adds or removes on this bag, then performance can become degraded and otherwise dereferenced objects in the RcIdentityBag cannot be garbage collected.

The message `cleanupBag` may be sent to the `RcIdentityBag` to process removals for inactive sessions. This may cause conflicts if a session logs in and adds or removes an object.

### RcLowMaintenanceIdentityBag

`RcLowMaintenanceIdentityBag` is similar to `RcIdentityBag` in behavior, but does not require regular cleanup. Rather than using a per-session subcollection of add and remove elements, `RcLowMaintenanceIdentityBag` relies on replay to resolve conflicts. Like `RcIdentityBag`, it has no conflict for multiple sessions that add objects to the bag, and a single session that removes objects.

The cumbersome name is intended to be temporary, with this implementation replacing `RcIdentityBag`'s subcollection-based implementation in some future release.

### RcIdentitySet

The class `RcIdentitySet` is similar to `IdentitySet`, but no conflict occurs when multiple users add objects to an `RcIdentitySet`. If a conflict with other update operations on the `RcIdentitySet` occur, the add is replayed so that the commit can succeed.

### RcKeyValueDictionary

The class `RcKeyValueDictionary` provides the same functionality as `KeyValueDictionary`, but with no conflict for operations that involve different keys in the dictionary. As long as the keys are different, multiple sessions can add new keys to the dictionary, remove keys, or update values.

`RcKeyValueDictionary` avoids conflict by performing a selective abort and replay of the modifications to the dictionary.

### Queue classes

The Queue classes implement a first-in-first-out (FIFO) queue. These are a kind of collection that is ordered by the sequence in which objects are added to the collection. The `add:` message puts an element at the logical end of the queue, and the `remove` method returns the element at the logical head of the queue.

The following example has the same semantics for `GsPipe`, `RcPipe`, and `RcQueue`; the choice of classes depends on the transactional requirements of your application.

#### Example 4.4 FIFO Queue

```
| pipe |
pipe := RcPipe new.
pipe add: 'orange'.
pipe add: 'apple'.
pipe add: 'banana'.
pipe remove.
pipe.
%
  aRcPipe( 'apple', 'banana')
```

## GsPipe

The class `GsPipe` implements a first-in-first-out queue, with no conflict when a single session adds objects to the `RcPipe`, and only one session removes objects.

Internally, the `GsPipe` is implemented as a linked list of `GsPipeElements`. Since adds and removes only affect the respective ends of the linked list, there is no conflict between add and remove.

## RcPipe

The class `RcPipe` implements a first-in-first-out queue, with no conflict when multiple sessions add objects to the `RcPipe`, and only one session removes objects.

Internally, the `RcPipe` is implemented as a linked list of `GsPipeElements`. Unlike with `GsPipe`, if a conflict with an add by another session occurs, the add operation is replayed so that the commit can succeed. Only `add :` and operations that invoke `add :` are reduced conflict.

## RcQueue

The class `RcQueue` implements a first-in-first-out queue, with no conflict when multiple sessions add objects to the `RcQueue`, and only one session removes objects.

`RcQueue` has a more complex internal implementation, which allows it to handle high rates of concurrent updates without affecting performance. However, some usage conditions make it necessary to perform manual cleanup

### Internal implementation

Internally, `RcQueues` are implemented using an Array of `RcQueueSessionComponents`, each corresponding to a session number. The `RcQueueSessionComponents` contain `RcQueueEntry` instances, one for each object that the session with the corresponding session number has added to the queue. The `RcQueueEntry` includes timestamp and sequence number; the timestamp is used to determine the next object within the entire queue is next to be returned, and the sequence number is used to track the next element within the queue for a specific session.

When a next message causes an object to be removed, the removing session updates the `RcQueue`'s removal sequence number array corresponding to the `RcQueueSessionComponents` in which the removed object was found.

### Maintenance

Reclaiming the storage of objects that have been removed from the queue is deferred until new objects are added by a session with the same session number; this is the way the risk of conflict is avoided.

If a session adds a great many objects to the queue all at once and then does not add any more, while another session consumes the objects, performance can become degraded, particularly from the consumer's point of view. In order to avoid this, the producer can send the message `cleanupMySession` occasionally to the instance of the queue from which the objects are being removed. This causes storage to be reclaimed from obsolete objects.

To remove obsolete entries belonging to all inactive sessions, the producer can send the message `cleanupQueue`.

## 4.3 GsBitmap

A GsBitmap is quite different than the other collections that have been described. Instances of GsBitmap are objects that encapsulate an in-memory bitmap, with the presence of an object in the collection only indicated by the way a bit is set at the index for the oopNumber of the object.

GsBitmaps cannot be committed, and are designed to optimize performing tasks on very large numbers of persistent objects. In particular, repository analysis using `allInstances` and similar methods can be more easily done using GsBitmaps. The objects in a GsBitmap are not in temporary object memory, allowing arbitrary large collections. A number of repository analysis methods return GsBitmap instances, and instances of GsBitmap can be created from hidden sets (see section 18.1 on page 343).

While GsBitmap can be considered as a collection and implements some Collection protocol, it does not inherit from Collection. Methods such as `add:`, `remove:`, `includes:` and `do:` are implemented specifically for GsBitmap; see the image for specific methods. You may send `asGsBitmap` to create a GsBitmap from a collection, provided the collection only contains objects that are allowed in a GsBitmap; use `asArray` to collect the objects corresponding to the OOPs in the GsBitmap.

Since GsBitmap is intended to work with very large collections of objects, it implements set arithmetic methods, `+/union:`, `-/difference:` and `*/intersect:`.

While there are restrictions and caveats to using GsBitmap, there are significant benefits in memory use. Instances of GsBitmap use C Heap memory, not temporary object memory, to store the bit array.

The following restrictions apply to GsBitmap:

- ▶ Only committed objects can be added to a GsBitmap.
- ▶ Specials, such as Characters, Integers, and SmallDoubles, cannot be added, since they do not have POM OOPs.
- ▶ Objects can appear only once in the bitmap; duplicates are ignored.
- ▶ GsBitmaps are ordered in OOP order, regardless of the order they are added.
- ▶ GsBitmaps cannot be committed, since the underlying structure is not an object.
- ▶ Being in a GsBitmap does not count as a reference to an object, so there is a risk that objects in a GsBitmap could be garbage collected.

The following example finds all instances of Customer that are not in the AllCustomers collection:

### Example 4.5 GsBitmap

```
| bmAllInstances bmCustomerColl |
bmAllInstances := SystemRepository allInstances: Customer.
bmCustomerColl := AllCustomers asGsBitmap.
(bmAllInstances difference: bmCustomerColl) asArray.
```

## GsBitmaps and C Heap memory

While GsBitmaps do not use temporary object memory, they do still use some memory, and it is possible to run out of C Heap memory if there is extensive use of large GsBitmaps.

An instance of GsBitmap requires a minimum of 16KB (one page) of C Heap memory, which can hold up to 2K objects. A GsBitmap's memory use always grows in 16KB increments. For GsBitmap instances that contain more than 2K elements, the amount of memory used will vary, depending on how dense the OOP values are within the leaves of the internal tree structure. The best case, for very large, dense bitmaps, is about 1 bit per object. A GsBitmap that contains all the OOP in the repository (`GsBitmap allValidOops`) will take about  $(\text{System\_oopHighWaterMark} // 8)$  bytes of C heap memory.

## GsBitmaps and their objects

There is an important point to note about GsBitmaps; an object in a GsBitmap is not "referenced" by the GsBitmap in the usual way.

An object in GemStone that is not referenced by other persistent objects or by references from a session, is subject to garbage collection. In busy systems, the OOP of that object may be recycled and no longer be in use; and the OOP may be reused by this or another session for an entirely new object of any class. GemStone collections (other than GsBitmap) have references to the objects contained within them, which keeps the objects in temporary collections safe for the life of a session.

Since the references in a GsBitmap are just to the OOPS, not to the objects, objects in a GsBitmap are not safe; the reference from the bitmap is not sufficient to preserve content objects from garbage collection, if they are not referenced somewhere else in the application or session.

If your session performs commits or aborts (including automatic commits or aborts), and the objects that you are working with may become dereferenced (for example, removed from the root collection by another session), then your code should be prepared for objects to no longer exist, or to be a different object than expected.

If an object was garbage collected, and the OOP reused, it may have been used for a critical internal object, or an important object in your application. Use caution when modifying the objects returned from a GsBitmap.

## GsBitmaps methods for repository analysis

GsBitmap includes methods that enable repository-wide analysis of objects in the repository. The following methods are available; see the image for other methods.

```
GsBitmap >> referencedObjects
```

Returns a new GsBitmap containing the objects directly referenced by the objects in the receiver.

```
GsBitmap class >> allValidOops
```

Returns a GsBitmap containing the oops of all valid committed objects in the repository.

```
GsBitmap class >> transitiveReferences: aCollection
```

Returns a GsBitmap that contains all the objects which are transitively referenced from the objects in *aCollection*.

```
GsBitmap class >> allObjectsExcept: aCollection
```

Returns a GsBitmap that contains all objects that exist in the repository, which are not contained in the objects transitively referenced from *aCollection*.

## Bitmap files

In addition to standard collection protocol, GsBitmaps can be written to and read from disk, using the following methods:

```
GsBitmap >> writeToFile:
```

```
GsBitmap >> readFromFile:
```

```
GsBitmap >> readFromFile:withLimit:startingAt:
```

You may also query for information on a given bitmap file, using

```
GsBitmap >> fileInfo:.
```

 This method returns an array containing:

- ▶ number of oops in the file
- ▶ whether the file was written in page order
- ▶ number of valid oops
- ▶ number of oops that are not allocated, or in the process of being garbage collected

A GsBitmap file contains references to numeric OOPs. The caution about the risk of unreferenced OOPs being garbage collected and possibly reused, applies even more strongly when using GsBitmap files. And of course, the likelihood of incorrect results relates to the amount of garbage collection that has been done during the period between the time the file was written and when it is read.

## Page order Bitmap files

A GsBitmap is inherently in OOP order, so to use GsBitmaps with page order files, different protocol must be used.

You may write a GsBitmap file in page order using

```
GsBitmap >> writeToFileInPageOrder:
```

The bitmap file written with this method is formatted the same as that written by methods such as `Repository >> listInstancesInPageOrder:toFile:.` This format is different than the format written by `GsBitmap >> writeToFile:` and read by `GsBitmap >> readFromFile:*`.

To read a page-order GsBitmap file, and preserve the page order by returning values in an Array, use:

```
GsBitmap class >> readObjectsFromPageOrderFile: fileName  
startingAt: startIndex upTo: endIndex
```

Reads, validates and returns an array containing the valid oops in page order from a page-ordered bitmap file. *startIndex* is the index of the first object (1-based), *endIndex* is the index of the last object to return from the file.

## 4.4 Sorting the objects in a collection

You are likely at some point to want to present the contents of your Collection in a sorted order. You will have to determine how the objects in your collection should be compared to each other for the ordering you need.

### Default Sort

Many objects, such as strings, numbers, and dates, have an inherent sort ordering; they respond to `<=` in a common way, although they cannot always be compared with each other. If your collection contains only homogenous objects that share an understanding of `<=`, you can use messages such as `sortAscending`, `sortDescending`, and `asSortedCollection` to the collection.

#### Example 4.6 SmallInteger and String sorting

---

```
(Array with: 123 with: 3 with: 99 with: 10) sortDescending
%
  anArray( 123, 99, 10, 3)
```

```
(Array with: '123' with: '3' with: '99' with: '10') sortAscending
%
  anArray( '10', '123', '3', '99')
```

---

The default sort of Strings is case-insensitive, unless the only difference is case in which uppercase is first. However, in many cases you may need a different ordering, particularly when languages other than English and character outside the ASCII range are involved. GemStone provides specialized tools for this, which are described in Chapter 5.

The options depend on the type of data in your collection.

- ▶ **Sort based on predefined order of the objects.** Some objects, such as Strings, Integers, and DateTimes, have an inherent sort ordering, and GemStone provides default sorts for Collections that contain only objects that can be compared using `<=`.  
While strings have intuitive sort order, string sorting can be complex. Traditional and Unicode strings handle some cases differently. String sorting is described in section 5.3 on page 81.
- ▶ **Sort based on one or more of the predefined order of objects's instance variable values.** The sort you intend is based on the values in application objects instance variables, and these values have inherent sort order, such as sorting customers by zip code.
- ▶ **Arbitrary Sort.** *sortBlocks* allow you to specify expressions that can order any type of object according to your specific requirements.

These issues are the same when using a SortedCollection, which always maintains sort order as elements are added and removed, or when sorting another kind of collection for presentation.

## Sorting Application objects

Most likely, you will need to sort complex objects in your collection, such as Customers by name or Addresses by zip code. If the instance variables in your complex objects are objects that have a defined sort order, you can take advantage of `sortAscending:`, `sortDescending:`, and `sortWith:`, to provide a specification for the desired sort order.

You may wish to implement `<=` on your application objects, in which case you can just use `sortAscending`, `sortDescending`, and `asSortedCollection`. However, this provides a single definition of the sort order of your objects that will always be applied.

For example, say we have a class for Employee, and a Globals AllEmployees is a collection that contains instances of Employee:

### Example 4.7 Employee class and AllEmployees

```
Object subclass: 'Employee'
  instVarNames: #( 'firstName' 'lastName' 'job' 'age')
  classVars: #()
  classInstVars: #()
  poolDictionaries: #()
  inDictionary: UserGlobals

Employee compileMissingAccessingMethods

UserGlobals at: #AllEmployees put: (IdentityBag
  with: (Employee new firstName: 'Lee'; lastName: 'Smith';
    job: #librarian; age: 40)
  with: (Employee new firstName: 'Kay'; lastName: 'Adams';
    job: #clerk; age: 24)
  with: (Employee new firstName: 'Al'; lastName: 'Jones';
    job: #busdriver; age: 40))
```

To sort Employees by age and lastName, we can use the `sortAscending:` method, passing in the instance variables against which the ascending sort should be done:

### Example 4.8

```
| sorted str |
str:= String new.
sorted := AllEmployees sortAscending: #('age' 'lastName').
sorted do: [:anEmp |
  str add: (anEmp age asString); space; add: anEmp lastName; lf].
str
%
```

```
24 Adams
40 Jones
40 Smith
```

## Sorting in multiple orders

For finer control, you can use the `sortWith:` method, which allows you to define direction for each instance variable.

### Example 4.9

---

```
| sorted str |
str := String new.
sorted := AllEmployees sortWith: #('age' 'Ascending'
                                  'lastName' 'Descending').
sorted do: [:anEmp | str add: (i age asString);
           add: ' '; add: anEmp lastName; lf].

str
%
24 Adams
40 Smith
40 Jones
```

---

## SortBlocks

You can also specify sort ordering by defining a `sortBlock`. A `sortBlock` is a two-argument block that should return true if the first argument should precede the second argument, and false if not. The expressions within the block are expected to be symmetrical - i.e., for two specific arguments for which the block returns true, then the block should return false when the arguments are reversed. If values compare equal, and the block returns the same results for both argument orders, then the final ordering of the equal elements is arbitrary.

`SortedCollection` is a type of `Collection` that includes a `sortBlock`; `SortedCollection` class is discussed under “`SortedCollection`” on page 59.

You can sort the elements of a collection by creating a `SortedCollection` using `asSortedCollection: aBlock`, or by using methods such as `sortWithBlock:.` which return an `Array` with the sorted contents.

For example, to sort customers by last name:

```
AllEmployees sortWithBlock: [:a :b |
                             a lastName <= b lastName]
```

You can create sort blocks that are as elaborate as you need; however, you should observe the symmetry of the expression.

For example, this block sorts by `lastName`, with further sorting by `firstName` if the `lastNames` are the same:

```
AllEmployees sortWithBlock: [:a :b |
                             a lastName = b lastName
                             ifTrue: [a firstName <= b firstName]
                             ifFalse: [a lastName <= b lastName]
                             ].
```

## Sorting Large Collections

When sorting using the above methods, the entire collection must fit into memory. This may not be practical for very large collections.

To avoid out of memory errors when sorting large collections, you can allow the sort to issue periodic commits, which will make the sort results persistent. Persistent objects don't need to stay in memory the way temporary objects do, which reduces the demand on memory.

These intermediate commits are enabled by specifying a `persistentRoot` for the sort, and by taking advantage of the `IndexManager`'s ability to set up `autoCommit`. `IndexManager` is a class that manages `Indexes`, which you'll read more about in Chapter 8. You do not need to have an index on the collection in order to use this feature. However, you do need to set `IndexManager`'s `autoCommit` setting to true. For more information on `autoCommit`, see "Auto-commit" on page 153.

For example, the following code sorts `AllEmployees` collection using `sortWithBlock:persistentRoot:`

### Example 4.10 Sorting large collections, committed incremental results

---

```
UserGlobals at: #SortedEmployees put: Array new.  
System commitTransaction.  
AllEmployees  
    sortWithBlock: [:a :b | a lastName <= b lastName]  
    persistentRoot: SortedEmployees
```

---

# String Classes and Collation

---

String handling is an important part of most applications. While Strings are a type of Collection, they have a number of unique features and behavior.

**Characters and Unicode** (page 73)

Describes Characters.

**String classes** (page 75)

Introduces the GemStone Smalltalk objects that store collections of Characters.

**String Sorting and Collation** (page 81)

Describes collation, including traditional string collation and collation using the ICU libraries and Unicode strings.

**Encrypting Strings** (page 89)

Explains how to encrypt strings.

## 5.1 Characters and Unicode

A Character is a special object: an object whose value is encoded in the OOP. Literal Characters are formed with a leading \$.

### Code point

Each Character has a code or codePoint, which, for lower order Characters, is the ASCII value. Either of these terms may be used, though ASCII is an incorrect term for the higher code points. GemStone supports Characters with values from 0 to 16r10FFFF, the full Unicode range, except for the Unicode reserved range.

The Unicode range of codePoints from 16rD800-16rDFFF is reserved for encoding leading/trailing surrogate pairs for UTF-16 encoding. These can never be legal Unicode characters, and as such can never be present in Unicode strings.

To get the Character for a given codePoint, use the Character class methods `withValue:` or `codePoint:.`

## Attributes

Characters have “type”, and know if they are a digit, letter, separator, or other similar kind. This information is defined in the Unicode database as the Unicode general category, and a variety of testing methods are available. The Unicode database also defines the upper and lower case equivalents, and case conversion methods are available. See the image for a full list of available protocol.

For example,

```
$Z isUppercase
true

$u isDigit
false
```

## Collation

Characters are ordered (collated) using internal character tables, which provide a Unicode-like collation order for Characters up to code point 255. Characters above that are collated by code point. Character collation can be modified by installing character data tables, although this use is deprecated.

Character collation is used in collating instances of Traditional string classes, in Legacy String Comparison Mode. This character-based string collation has limitations outside the ASCII range; the ICU-library based string collation should be used if the default collation is not sufficient. For more on collation, see “String Sorting and Collation” on page 81.

## Empty string canonicalization

An empty String literal, that is, a String or Unicode7 created by evaluating "", is canonicalized to a predefined kernel OOP; these do not use additional OOPs in the repository, do not require additional space, nor affect garbage collection.

If you will be appending to a new empty string, you should start with String new, if the empty string will remain empty, using a literal String is more efficient.

## Unicode and the Unicode Database

The Unicode Consortium is an international standards organization that produces the Unicode Database. Unicode is a commonly used standard which provides unique codes for all Characters in all Character sets, in the range 0 to 0x10FFFF. It also describes the category of each Character and relationship between it and other Characters, and provides a default collation order with the Default Unicode Collation Element Table (DUCET).

For more information on this database, see

<http://www.unicode.org/Public/UNIDATA/UCD.html>

The Unicode Consortium provides code charts by script as well as a single master list of all characters, presented in an ASCII-only, comma-delimited version. The current version of this database can be found at

<http://www.unicode.org/Public/UNIDATA/UnicodeData.txt>.

GemStone’s Unicode strings and collation in Unicode Comparison Mode uses external libraries that support Unicode. Traditional Strings in Legacy Comparison Mode use a historic native GemStone implementation that broadly conforms to Unicode, but has not been updated as the standard has evolved.

## 5.2 String classes

A string is a sequence of Characters, implemented as a subclass of CharacterCollection.

Each element in a CharacterCollection is a Character. Since characters may require more than one byte of storage, the class of string may be transparently converted to an instance of the class with the appropriate capacity for that Character. The semantics of the CharacterCollection remain the same; access by index will return the Character at the given index, regardless of how many bytes the Character actually requires.

A fundamental quality of strings is collation. Since the scope of collation includes equality, the collation of strings affects a repository in many ways, such as dictionary lookups. Collation in GemStone has historically been handled using character-based tables. Unicode string-based collation using ICU open source libraries is included in recent releases and provides a much richer set of collation features. To ensure that legacy applications function correctly, GemStone supports both of these encoding/collation schemes.

- ▶ **Traditional strings and Legacy String Comparison Mode.** Traditional strings are instances of String, DoubleByteString, and QuadByteString. In Legacy String Comparison Mode, they collate using GemStone character-based tables, as in older GemStone releases.
- ▶ **Unicode strings and Unicode Comparison Mode.** Unicode strings are instance of Unicode7, Unicode16, and Unicode32. These strings use ICU string-based collation. In Unicode Comparison Mode, they can safely mix with symbols and the Traditional strings existing in the base image.

### Traditional Strings

In Legacy String Comparison Mode, Traditional strings collate using internal character-based collation tables. When the repository is in Unicode Comparison mode, however, Traditional strings use ICU-based Unicode collation.

Traditional strings are implemented in three classes:

#### String

Strings hold Characters with codepoints in the range 0..255 (8 bits).

#### DoubleByteString

DoubleByteStrings are required when one or more Characters in a string needs more than one byte of storage. DoubleByteStrings hold Characters with codepoints in the range 0...16rFFFF (64K).

#### QuadByteString

QuadByteStrings are required when one or more Characters in a string needs more than two bytes of storage. QuadByteStrings hold Characters with codepoints in the range 0...16r10FFFF.

While Traditional strings normally hold human-readable text characters, this is not a requirement. Generally, raw byte data would be held in an instance of ByteArray, but it may be more convenient to use a String. In particular, there are cases when an instance of String will be used to hold raw UTF-8 encoded bytes.

## Unicode Strings

Unicode strings always use ICU string-based collation. Like Traditional strings, there are three classes based on range, but note that the codePoint range is different than Traditional strings.

### Unicode7

A subclass of String, limited to holding Characters with codepoints in the range 0..127 that are represented in 7 bits.

### Unicode16

A subclass of DoubleByteString, holding Characters with codepoints in the range 0..16rFFFF (64K), excluding the range 16rD800-16rDFFF. This range is reserved for surrogates that allow encoding into UTF-16.

### Unicode32

A subclass of QuadByteString, holding Characters with codepoints in the range 0..16r10FFFF. Again, this excludes the range 16rD800-16rDFFF.

Unicode strings should not hold raw byte data.

## String equality, ordering, and interoperation

In Legacy String Comparison Mode, Traditional strings and symbols are compared for equality and ordered using character-based comparison, and equality includes non-printing characters as well as printing characters.

Unicode strings use the ICU string-based string collation, in which equality does not consider non-printing characters.

Since Traditional and Unicode string equality rules are different, Traditional strings and symbols (when the repository is in Legacy String Comparison Mode) may produce inconsistent results. In this mode it is an error to mix Unicode strings with Traditional strings or symbols, either for comparison or equality.

Also note that Unicode comparison of Symbols using = uses identity, while >= and <= compare according to the ICU rules for UnicodeString. Symbol comparison using = may produce inconsistent results for Symbols containing special characters, such as nuls, that are not counted by ICU comparison.

## Other String-like classes

### Symbol

A symbol is similar to a string, but each symbol with a unique set of Characters is guaranteed to have only one canonical instance in GemStone. Symbols are created by a special process, the SymbolGem, to ensure this uniqueness. Creating a new symbol will return an existing symbol, if one exists; a new symbol is only created if it has not been previously defined. Existing symbols cannot be modified.

Like strings, symbols may also contain Characters with values that require more than a byte of storage, and will convert from class Symbol into DoubleByteSymbols or QuadByteSymbols as needed. Since symbols are canonical, the class of a symbol always depends on the contents. While you can create a DoubleByteString with only characters in the range of String, you cannot create a DoubleByteSymbol that does not contain at least one character in the DoubleByte range, and the same is true for QuadByteString.

All symbols may be viewed by all users. Private information should be maintained in strings, not in symbols.

Symbols, DoubleByteSymbols, and QuadByteSymbols are restricted to 1024 or fewer characters.

Symbols that have no references from anywhere in the system may eventually be garbage collected, if the system is configured to do so. See the *System Administration Guide* for more information on symbol garbage collection.

Symbols, like strings, collate using character-based tables in Legacy String Comparison Mode and using ICU string-based collation in Unicode Comparison Mode. As a result, they cannot be compared to Unicode strings in Legacy String Comparison Mode.

Symbol equality comparisons with = compare identity, while other comparisons (> >= < <=) use the ICU string-based collation. GemStone symbols that contain the same printing characters but different non-printing characters will return false for = but compare true for >= and <=.

The literal form of a Symbol is specified using a leading #. The body of the symbol may additionally include single quotes. This is optional for symbols that are legal identifiers and keywords, but required for symbols that start with a number, include punctuation/spaces, etc. For example:

```
#'22 skidoo'  
#fooBar
```

## ByteArray

ByteArray is a specialized collection that is restricted to holding Integers between 0 and 255 (inclusive). While ByteArray is not a kind of String, the contents may be interpreted as a String.

Instances of ByteArray can be creating using literal syntax #[]. For example:

```
#[ 1 2 3 4 ]
```

## Utf8

Utf8 is a subclass of ByteArray. It is not a kind of String, but may easily be converted back and forth from a traditional or Unicode string. A Utf8 holds the UTF-8 encoded bytes created by sending encodeAsUTF8 to a string, or by reading encoded data from a GsFile using contentsAsUTF8. Utf8 instances should not be directly created or edited.

```
' amas' encodeAsUTF8  
anUtf8( 197, 161, 97, 109, 97, 115)
```

Instances of Utf8 can be read from and written to instance of GsFile, which cannot directly handle characters with codePoints over 256. FileSystem by default reads and writes in UTF-8 encoding, transparently encoding and decoding to string instances.

## Utf16

Utf16, like Utf8 is a subclass of ByteArray that is not a String, but may be converted back and forth from a traditional or Unicode string. A Utf8 holds UTF-16 encoded bytes corresponding to a traditional or Unicode String.

## String protocol

### Creating Strings

Strings created as literals, that is, in text encased in single quotes, are invariant; they cannot be modified after they are created.

In addition to creating strings as literals, you can use the inherited instance creation methods, such as `new:` and `withAll:`. For example:

```
String withAll: #($a $z $u $r $e).  
azure
```

### Concatenating Strings

A string responds to the comma operator by returning a new string in which the argument to the comma has been appended to the string's original contents. For example:

```
'String ' , 'con' , 'catenation'  
String concatenation
```

Although this technique is handy, it's not very efficient; each `#,` message send creates a new instance of `String`, so this example creates three `Strings`, returning the final one.

To build a string efficiently, by appending onto the original object, you can use `add:`, which modifies the original string. Note that you cannot start with a literal string, since a literal string is invariant.

For example:

```
| resultString |  
resultString := String new.  
resultString add: 'String '  
              add: 'con';  
              add: 'catenation'.  
  
resultString  
%  
String concatenation
```

### Converting between String classes and encodings

To convert between UTF-8 encoded bytes and the various kinds of string classes, there are a number of methods:

- ▶ Instances of `Symbols` and `Traditional strings` can be converted to the lowest-storage type of `Unicode string` using `asUnicodeString`.
- ▶ Instances of `Symbols` and `Unicode strings` can be converted to the lowest-storage type of `Traditional strings` using `asString`.
- ▶ A `Traditional string` that is composed of raw UTF-8 encoded bytes can be decoded to a `Unicode string` using `decodeFromUTF8ToUnicode`, or to another `Traditional string` with decoded bytes, using `decodeFromUTF8ToString`.
- ▶ A `Traditional string` can be encoded into a `String` containing the raw UTF-8 encoded bytes, using `encodeAsUTF8IntoString`.

- ▶ To convert from a ByteArray containing UTF-8 or from a Utf8 to a Unicode string, use `decodeFromUTF8ToUnicode`, or to convert to a Traditional string, use `decodeFromUTF8ToString`.
- ▶ Instances of ByteArray and Utf8 may be converted to a Traditional string without decoding by using `bytesIntoString`.
- ▶ All kinds of strings can be encoded to an instance of Utf8 by using `encodeAsUTF8`.

## String Transformations

CharacterCollection and its subclasses define messages that let you perform various conversions.

Strings can be converted in case:

- ▶ `asUppercase` creates a new instance with all uppercase letters
- ▶ `asLowercase` creates a new instance with all lowercase letters
- ▶ `asTitlecase` creates a new instance with the first letter of each word capitalized, the remaining letters lowercase.
- ▶ `asFoldcase` returns a new instance in “fold case”, which is case-free for comparison, and usually is similar to the lowercase.

For example:

```
'abcde' asUppercase
  ABCDE
```

You can remove leading and/or trailing whitespace separators using methods such as `trimSeparators`. There are a number of variants; see the image for details.

For example:

```
' abcde ' trimSeparators
  'abcde'
```

Strings can be split using the `subStrings: method`, which allows you to specify one or more characters to use as markers.

For example, to split a text into lines with `/`:

```
'owa/tagu/siam' subStrings: '/'
  anArray( 'owa', 'tagu', 'siam')
```

Strings can be converted to numbers and other types of objects as well. For example:

```
'15' asFloat
  15.0
```

Note that not all Strings can be converted to all kinds of other objects; if the String does not contain the representation of a number, for example, it’s meaningless to convert it to an Integer, so this will return an error.

## Equality and Identity

Traditional strings are equal to each other if they contain the exact same Characters in the same case; equality is case-sensitive.

Unicode strings compared using `=` follow the ICU library comparison rules for equality, which are similar, although any non-whitespace control characters (such as null) are ignored for the comparison.

As mentioned above, Traditional strings and Unicode strings cannot be compared to each other for equality using `=`, when the repository is in Legacy String Comparison Mode. To compare traditional and Unicode strings in any combination, use `compareTo:collator:`, specifying `nil` for the collator to indicate the default collator.

Strings can be compared for case-insensitive equality using the methods `isEquivalent:` or `equalsNoCase:`.

### Identity in Literal vs. nonliteral

Literal and nonliteral Strings behave differently in identity comparisons. Each nonliteral String (created, for example, with `new`, `withAll:`, or `asString`) has a unique identity. That is, two Strings that are equal are not necessarily identical.

```
| nonlitString1 nonlitString2 |
nonlitString1 := String withAll: #($a $b $c).
nonlitString2 := String withAll: #($a $b $c).
(nonlitString1 == nonlitString2)
false
```

However, literal strings that contain the same character sequences and are compiled at the same time are both equal and identical:

```
| litString1 litString2 |
litString1 := 'abc'.
litString2 := 'abc'.
(litString1 == litString2)
true
```

This distinction can become significant in building sets. If you add both `litString1` and `litString2` to the same `IdentitySet`, the set will contain only one instance of `'abc'`; however, an `IdentitySet` would include both `nonlitString1` and `nonlitString2`.

### Searching and Pattern matching

`CharacterCollection` and its subclasses define methods that can tell you whether a string contains a particular sequence of characters and, if so, where the sequence begins. This search can be case sensitive, case insensitive, and may include wild cards.

Below are some common methods; see the image for further methods.

**Table 5.1 Search and Pattern Match Protocol**

Case-sensitive Search	Case-insensitive Search	Description
	<code>includesString: subString</code>	Return true if the receiver includes <code>subString</code> .
<code>findString: subString startingAt: anIndex</code>	<code>findStringNoCase: subString startingAt: anIndex</code>	Return the index of <code>subString</code> if it exists within the receiver at <code>anIndex</code> or above, otherwise zero (0).

Table 5.1 Search and Pattern Match Protocol

Case-sensitive Search	Case-insensitive Search	Description
matchPattern: <i>patternArray</i>		Return true if the receiver matches the specifications in <i>patternArray</i>
findPattern: <i>patternArray</i> startingAt: <i>anIndex</i>	findPatternNoCase: <i>patternArray</i> startingAt: <i>anIndex</i>	Return the index of a substring in the receiver that matches the specifications in <i>patternArray</i> at <i>anIndex</i> or above, otherwise zero (0).

### Pattern Matching Wild Cards

Pattern matching arguments (*patternArray*) consist of an Array containing combinations of Strings and the wildcard characters \$\* and \$?. The character \$? matches any single character in the receiver, and \$\* matches any sequence of characters in the receiver.

This is an example of the use of wildcard characters in pattern matching.

```
'weimaraner' matchPattern: #('w' $* 'r')
true
```

Since \$\* is interpreted as “any sequence of characters”, this returns true.

Similarly, The following example returns the index at which a sequence of characters beginning and ending with \$r occurs in the receiver.

```
'weimaraner' findPattern: #('r' $* 'r') startingAt: 1
6
```

If a wildcard character \$\* or \$? occurs in the receiver or within a string in the argument array, it is interpreted literally.

The following expressions illustrate what happens when the \* is within the string and interpreted literally:

```
'w*r' matchPattern: #('weimaraner')
false

'weimaraner' findPattern: #('w*r') startingAt: 1
0
```

## 5.3 String Sorting and Collation

While strings clearly have a natural sort order (collation), the details of that order are complex. Different languages may sort the same set of strings differently, according to the particular rules in that language. Even within one language, different applications may want to order string data differently. To complicate matters, some languages may treat certain sequences of characters as a unit when sorting strings.

Collation depends on the results of a comparison between two strings, which in turn depends on how the Characters within the string are collated. While this simple view breaks down with some sorting requirements and linguistic rules, basic string comparison is adequate for many uses and is faster than the more complete external collation.

## Comparison Mode

The Comparison Mode of a repository controls the way comparisons are done between instance of Traditional strings. The modes are:

- ▶ **Legacy String Comparison Mode**, the default for new applications.
- ▶ **Unicode Comparison Mode**, enabled in all GsDevKit-based applications.

In Legacy String Comparison Mode, Traditional strings and symbols cannot be compared to Unicode strings without using special protocol. Collation of Traditional strings and symbols is using character-based collation.

In Unicode Comparison Mode, Traditional strings and Symbols use ICU string-based collation, and can interoperate easily with Unicode strings.

A new repository can be easily switched to Unicode Comparison Mode. Since the collation rules may be subtly different, and affect system operations such as looking up class names in SymbolDictionaries, changing the mode for existing applications should be done with great care and thorough testing. To be safe, all indexes and sorted collections should be rebuilt, and all hashed collections re-hashed. The mode of a repository must be managed as part of System Administration, not by individual developers on a shared repository.

## StringConfiguration

The Comparison Mode is controlled by the Global #StringConfiguration. By default, StringConfiguration is set to String, and the repository is therefore in Legacy String Comparison Mode.

To enable Unicode Comparison Mode, as SystemUser, execute:

```
StringConfiguration enableUnicodeComparisonMode
```

This returns the previous setting for Unicode Comparison Mode. Note that this commits, but the current session is not affected; the new mode will take effect for all subsequent logins.

To enable Legacy String Comparison Mode, as SystemUser, execute:

```
StringConfiguration disableUnicodeComparisonMode
```

Again, note that this operation commits, but the change does not affect the current session; the new mode will take effect for all subsequent logins.

To verify the mode in this repository, execute:

```
StringConfiguration isInUnicodeComparisonMode
```

## Auto-conversion

When you create or update a kind of Traditional or Unicode string with a Character that requires more bits than the specific class of string can hold, it is transparently auto-converted to the appropriate class.

For example, if you add if you add the Euro character (code point 8364) to an instance of String, which can only hold codePoints up to 255, it will auto-convert to an instance of DoubleByteString. Likewise, if you add an Yen symbol (codePoint 177) to an instance of Unicode7, which can only hold codePoints up to 127, it is auto-converted to an instance of Unicode16.

When the repository is in Unicode Comparison Mode, an instance of String that would otherwise auto-convert to an instance of DoubleByteString is converted to an instance of Unicode16, for improved comparison and collation since comparisons use the ICU libraries.

## Legacy String Comparison Mode for Traditional Strings

Traditional strings (String, DoubleByteString, and QuadByteString) and symbols (Symbol, DoubleByteSymbol, and QuadByteSymbol) are collated, in Legacy String Comparison Mode, by individual character. The comparison of characters with values up to 255 are done according to the Default Unicode Collation Element Table (DUCET), and Character 256 and above are sorted by codePoint, the Unicode numeric value.

Legacy applications may have installed non-default internal character tables, which modified the character-based collation. This is no longer recommended; if the default character-based collation is not sufficient for your application, you should integrate the ICU string-based collation.

Enabling Unicode Comparison Mode (see “Comparison Mode” on page 82) causes Traditional strings and symbols to collate following the same rules as Unicode strings. This section only applies when in Legacy String Comparison Mode, not in Unicode Comparison Mode.

String ordering using <= (as well as <, >, and >=) is not case-sensitive. When instances of String, DoubleByteString, and QuadByteString are compared using <= or related operations, the comparison first is done case-insensitive. If they are found to be equal other than with respect to case – if the only difference is case – then they are collated according to the Character Data Table, which specifies uppercase comes before lowercase.

For example:

```
#( 'MM' 'c' 'Mm' 'mb' 'mM' 'x' 'mm' )
  sortAscending
  anArray( 'c' 'mb' 'MM' 'Mm' 'mM' 'mm' 'x' )
```

Since ordering is by character, with only case being excluded, the default ordering is sensitive to accents and other diacritical marks on characters. Characters with diacritical marks are not related to the base character.

For example, all words beginning with 'Co' and 'co' would sort before all words beginning with 'Có' and 'có':

```
#('Cór' 'COz' 'Coa' 'cóa')
  sortAscending
  anArray( 'Coa', 'COz', 'cóa', 'Cór')
```

## Unicode Comparison Mode and ICU Collation

Unicode strings, and all strings when in Unicode Comparison Mode, use the ICU (International Components for Unicode) libraries to provide string-based collation. The ICU libraries are a widely-used, open-source implementation of language-specific sorting and collation.

For a complete explanation of the features and subtleties of language-specific collation, you should refer to documentation on the ICU website, <http://icu-project.org/>.

The classes IcuLocale and IcuCollator provide an interface to the ICU libraries. Unicode strings (instance of Unicode7, Unicode16, and Unicode32) and instances of Utf8 use

IcuCollator and IcuLocale to perform sorting operations using the ICU libraries. The collation is performed by considering the entire string, not on a character-by-character basis, and requires a specific language and locale to determine the rules for the comparison. In addition to specific language rules, ICU sorting is highly configurable for other application-specific sorting requirements.

While collation will vary according to specific language and locale, in general ICU collation orders characters with diacritical marks with the base character, and sorts lowercase before uppercase.

For example, using the sorting examples in the previous section and the default collator for the US, a different sort ordering is produced from that of legacy collation:

```
#( 'MM' 'c' 'Mm' 'mb' 'mM' 'x' 'mm' )
  sortAscending
  anArray( 'c', 'mb', 'mm', 'mM', 'Mm', 'MM', 'x' )

#('Cór' 'COz' 'Coa' 'cóa')
  sortAscending
  anArray( 'Coa', 'cóa', 'Cór', 'COz' )
```

This is the default US collation; by configuring the IcuCollator, however, many other orderings may be produced.

## IcuLocale

Instances of IcuLocale represent a specific language, country, and language variant. The available IcuLocales are in the shared library and can be listed using `IcuLocale class >> availableLocales`.

A default instance of IcuLocale is instantiated on first reference, and stored in session state. The default IcuLocale is based on the operating system locale setting for the gem. The default IcuLocale affects collation, so some care should be taken in configuring the operating system locale for the gem processes. In applications with distributed locales, it may be safer to set a default IcuLocale on login, using `UserProfile >> loginHook:` (see the *System Administration Guide*).

To set a specific default IcuLocale, use the method `IcuLocale class >> default:`. This sets the default locale for the session executing this code. While the instance of IcuLocale can be made persistent, the default IcuLocale does not persist from session to session.

To determine what IcuLocale is currently in use, use the method `IcuLocale >> default`.

```
IcuLocale default
  IcuLocale en_US
```

## IcuCollator

An IcuCollator encapsulates the rules involved in collation for a specific IcuLocale. A default instance of IcuCollator is instantiated on first reference, based on the default IcuLocale, and stored in session state.

When comparing instances of Unicode string classes, the comparison always uses an IcuCollator, using the method `compareTo:collator:`. If an IcuCollator is not specified,

such as when Unicode string classes are compared using `>`, the `IcuCollator` default is used; which in turn uses `IcuLocale` default.

You can also create an instance of `IcuCollator` for a specific locale, if you need to use specific collation rules other than the default. You can do this using `IcuCollator` class methods `forLocale: anIcuLocale` or `forLocaleNamed: aString`. For example, to create an `IcuCollator` for the German language as used in Germany:

```
IcuCollator forLocaleNamed: 'de_DE'
```

The actual string comparison is done by the ICU libraries, and follows the ICU comparison rules for that locale. Collation rules are similar in most western languages, but there are differences in specific languages.

For example, in the Hungarian language, 'cs' is considered a single letter, so words that start with 'cs' are sorted together and follow other words beginning with 'c'. The following example sets up a collection that is sorted according to Hungarian rules:

### Example 5.1 Sorting in Hungarian IcuLocale

```
| hungarianWords collator |
collator := IcuCollator forLocaleNamed: 'hu_HU'.
hungarianWords := IcuSortedCollection newUsingCollator: collator.
hungarianWords
  add: 'csak' asUnicodeString;
  add: 'cukor' asUnicodeString;
  add: 'comb' asUnicodeString.
hungarianWords
a IcuSortedCollection
  sortBlock          a ExecBlock2
  collator           a IcuCollator
  #1 comb
  #2 cukor
  #3 csak
```

## Customizing Sort

`IcuCollator` includes a number of attributes that can be used to customize the sort. These attributes work within the specific language rules of the associated `IcuLocale`.

Keep in mind that while the default values and the descriptions listed in Table 5.2 apply to most locales, particularly with non-Western scripts, the defaults may be different in different locales, and the attribute may have different behaviors.

See the ICU site, particularly the pages under <http://userguide.icu-project.org/collation>, for more precise descriptions and more detailed documentation.

**Table 5.2 IcuCollator Attributes**

Attribute name	Allowed values	Default	
alternateHandling	true   false	false	When true, allows space and punctuation characters within the string to be ignored.

Table 5.2 IcuCollator Attributes (Continued)

Attribute name	Allowed values	Default	
caseFirst	'off', 'upperFirst', or 'lowerFirst'	'off'	When comparing case, determines if upper or lowercase is sorted first. Most locales sort lowercase first when caseFirst is 'off' as well as when 'lowerFirst'.
caseLevel	true   false	false	When true, considers case in the comparison, even if the strength would normally not consider case.
frenchCollation	true   false	false	When true, sorts secondary differences (e.g. differences in diacritical marks) in reverse order. This is the collation rule for French.
normalization	true   false	false	Determines whether to normalize input strings. Useful if input data may not be -normalized, but impacts performance.
numericCollation	true   false	false	When true, sorts numeric sequences within the string by numerical rather than string comparison; e.g. sort '100' after '2'.
strength	PRIMARY - 0 SECONDARY - 1 TERTIARY - 2 QUARTENARY - 4, or IDENTICAL - 15	TERTIARY	Determines the level of collation factors to consider, such as diacritical marks and case. See discussion below for more details.

Strength allows degrees of sort, to consider or not consider things like accent characters and case when performing the sort. The default strength is TERTIARY for most locales (the main exception being Japanese). The following are the sort strengths:

- ▶ PRIMARY sorts by primary differences, ignoring secondary and later differences. The base letter represents a primary difference, so for example 'a' and 'b'.
- ▶ SECONDARY sorts by primary and secondary differences, ignoring tertiary and later differences. An example of a secondary difference is diacritical differences on the same base letter, for example 'o' and 'ó'.
- ▶ TERTIARY sorts by primary, then secondary, then tertiary differences. Uppercase vs. lowercase is a tertiary differences. TERTIARY is the default sort order for most locales.
- ▶ QUATERNARY is used in Japanese, where it distinguishes between Japanese Katakana and Hiragana, and can be used to break ties among separator characters when alternateHandling is true.

- ▶ IDENTICAL sorts by the specific character, by codepoints in the NFD (Normalization Form Canonical Decomposition) form. There is a performance impact with this strength.

The default sort strength is TERTIARY. As an example, when two strings are compared using TERTIARY strength, characters in the strings are compared first by the base character, ignoring any case or diacritical marks. If the base characters are the same, they are compared by diacritical mark, ignoring case. If both base characters and diacritical marks are the same, then case is considered. Note that unlike GemStone's Strings or ASCII ordering, the default sorts places lowercase before uppercase.

Keep in mind that with lower sort strengths, when a factor such as case is not used, the relative position in the results of similar strings is not deterministic; the strings compare as the same, and so their position will depend on the order of the input.

By using the IcuCollator sort attributes, you have a great deal of control over your specific sorting.

For example, using the alternative handling example, you can sort strings that include spaces, dashes and other punctuation without considering the punctuation characters when doing the comparison:

### Example 5.2 Sort ignoring punctuation

```
| blues collator|
collator := IcuCollator forLocale: IcuLocale default.
collator alternateHandling: true.
blues := IcuSortedCollection newUsingCollator: collator.
blues add: (Unicode7 withAll: 'blue berry').
blues add: (Unicode7 withAll: 'blue moon').
blues add: (Unicode7 withAll: 'bluebird').
blues add: (Unicode7 withAll: 'blue bird').
blues add: (Unicode7 withAll: 'blue-bird').
blues add: (Unicode7 withAll: 'bluetooth').
blues
%
a IcuSortedCollection
  sortBlock          a ExecBlock2
  collator           a IcuCollator
  #1 blue berry
  #2 bluebird
  #3 blue bird
  #4 blue-bird
  #5 blue moon
  #6 bluetooth
```

### IcuSortedCollection

An IcuSortedCollection is a specialized subclass of SortedCollection for which you do not set the sortBlock. An IcuSortedCollection may only hold instances of subclasses of CharacterCollection. It is associated with a IcuCollator, which in turn is associated with an IcuLocale, and the sorting behavior is specific to the configuration of these instances. IcuSortedCollections rely on the open-source ICU libraries to perform the comparisons and produce correctly collated results.

Using `IcuSortedCollection` is recommended if you will have sorted collections containing Unicode strings. This avoids lookup failures if a different collator is used to lookup than was used to sort the elements in the collection.

## ICU libraries and versioning

### ICU and Unicode versioning

The Unicode Consortium periodically releases new versions of the Unicode Standard, with (usually minor) changes in collation and the addition of new characters. The ICU organization then periodically releases new versions of their libraries reflecting these changes in the standard. Major GemStone releases include the latest version of the ICU libraries.

The indexing structures depend on collation encodings from ICU that may change between versions, even if the collation changes would not otherwise affect the application. So even in cases where the Unicode differences are minor, the ICU library version loaded in an application must match the ICU version used to build indexes.

To accommodate the (generally) low value of upgrading to a new ICU library, and the potentially high cost of rebuilding structures in your application that depend on collation, GemStone preserves the existing ICU library version over upgrade.

### IcuLibraryVersion

The version of the ICU library that is used in a repository is stored under (`Globals at: #IcuLibraryVersion`). This is a string, which must correspond to one of the versions of the ICU libraries in the product distribution. When a session logs in, it will select the ICU shared libraries to load based on the `IcuLibraryVersion` value.

As with `StringConfiguration`, `IcuLibraryVersion` is a global, repository-wide setting that can be only changed by `SystemUser`, to avoid the risk of lookup failures and incorrect query results. It should be managed as part of System Administration, not by individual developers on a shared repository.

### Updating IcuLibraryVersion

To update the version of ICU libraries in your repository, you will need to follow this procedure:

1. Ensure no other users are on the system
2. Login as `SystemUser` and execute

```
Globals at: #IcuLibraryVersion put: newVersionString
Commit and logout.
```
3. Shut down and restart the Stone.
4. Login as `DataCurator`, or a user with the appropriate object access rights. If you are using a linked session, you may need to restart the application to allow the new version of the ICU shared library to be loaded
5. Update any persistent data structures that may be affected. This involves dropping and rebuilding indexes that involve Unicode strings, resorting `SortedCollections`, and resorting any application data structures that depend on Unicode string collation.

6. When this is complete and all changes have been committed, other users may be allowed to login.

## 5.4 Encrypting Strings

There are times when you may wish to encrypt strings in your repository or for transmittal to other systems. GemStone provides an interface to Advanced Encryption Standard (AES) encryption/decryption, provided by the OpenSSL open source libraries included with GemStone.

The AES specification is available at:  
<http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.197.pdf>.

All encryptions/decryptions are in cipher block chaining (CBC) mode; see the AES specification document for further details.

Encryption and decryption API methods are provided for 128-bit/16-byte keys, 192-bit/24-byte keys, and 256-bit/32-byte keys, using the following methods.

Encryption can be done on instances of `ByteArray` or `Utf8`, or subclasses of `CharacterCollection`. For encryption, you must provide a key that is a `ByteArray` of the appropriate size (16, 24, or 32 bytes) containing key bytes, and a salt that is a 16-byte `ByteArray` containing salt values.

The following methods encrypt or decrypt using the specified key and salt, return the encrypted or decrypted result:

```
aesEncryptWith128BitKey: aKey salt: aSalt
aesDecryptWith128BitKey: aKey salt: aSalt
```

```
aesEncryptWith192BitKey: aKey salt: aSalt
aesDecryptWith192BitKey: aKey salt: aSalt
```

```
aesEncryptWith256BitKey: aKey salt: aSalt
aesDecryptWith256BitKey: aKey salt: aSalt
```

These methods place the encrypted or decrypted result into `aByteObjOrNil`, starting at offset 1, and resizing if necessary. If `aByteObjOrNil` is nil, a new instance of the same class as the receiver will be created containing the results.

```
aesEncryptWith128BitKey: aKey salt: aSalt into: aByteObjOrNil
aesDecryptWith128BitKey: aKey salt: aSalt into: aByteObjOrNil
```

```
aesEncryptWith192BitKey: aKey salt: aSalt into: aByteObjOrNil
aesDecryptWith192BitKey: aKey salt: aSalt into: aByteObjOrNil
```

```
aesEncryptWith256BitKey: aKey salt: aSalt into: aByteObjOrNil
aesDecryptWith256BitKey: aKey salt: aSalt into: aByteObjOrNil
```

You may use `ByteArray withRandomBytes: N` to produce pseudo-random key and salt values for encryption. For example:

### Example 5.3 String Encryption

---

```
topaz 1> printit
| key salt encrypted |
key := ByteArray withRandomBytes: 32.
salt := ByteArray withRandomBytes: 16.
encrypted := 'My secret string' aesEncryptWith256BitKey: key
           salt: salt.
encrypted aesDecryptWith256BitKey: key salt: salt.
%
My secret string
```

---

# Numeric Classes

---

This chapter describes GemStone's Numeric and Time-related classes.

This includes Numbers, such as Integers, floating point, fractions. Most numbers can be specified as literals within your code, and most numbers can be used in expressions with, or converted to, other types of numbers.

Also included are Time-related classes, such as Date, Time, DateTime, and DateAndTime.

**Integers** (page 91)

Describes classes that represent whole numbers: SmallInteger and LargeInteger.

**Binary Floating Point** (page 92)

Describes classes for binary floating point numbers: SmallDouble and Float.

**Other Rational Numbers** (page 96)

Describes classes for other rational numbers with different ranges and precisions, including Fraction, FixedPoint, ScaledDecimal, and DecimalFloat.

**Dates and Times** (page 101)

Describes the classes that represent times.

**Internationalizing** (page 107)

How to control the display of decimal points.

**Random Numbers** (page 109)

Information on the set of random number generator classes, providing random numbers of various purposes.

## 6.1 Integers

Integers in GemStone are composed of SmallIntegers and LargeIntegers. Most Integers you are likely to use will be SmallIntegers, in the range of  $-2^{60}$  to  $2^{60} - 1$ . Integers outside this range are represented by LargeIntegers. Operations that result in a value outside the SmallInteger range transparently result in LargeIntegers, and vice-versa

The literal syntax for Integer will create either a SmallInteger or LargeInteger.

Integers can be specified using radix notation, using the `r` or `#` characters.

For example, to specify the hex `SmallInteger` value `FF`, the following are all valid:

```
FFr16
FF#16
Number fromString: 'FFr16'
'ff#16' asNumber
```

## SmallInteger

`SmallIntegers` are special (immediate) objects, that is, the number itself is encoded in the OOP, making instances of this class both small (since no further storage is required) and fast. They are also unique, so `SmallIntegers` of the same value are always identical (`==`) as well as equal (`=`).

`SmallIntegers` have a range from  $-2^{60}$  to  $2^{60} - 1$ . Values outside this range must be represented as `LargeIntegers`.

## LargeInteger

`LargeIntegers` are not special objects; they require an OOP.

Each instance of `LargeInteger` is stored as an array of bytes, where every 4 bytes represents a base 4294967296 digit. The first 4 bytes are the sign digit (0 or 1), the next 4 bytes in that array constitute the least significant base 4294967296 digit, and the last 4 bytes are the most significant base 4294967296 digit.

Instances of `LargeInteger` have a maximum size of 4067 digits plus the sign. The maximum absolute value for a `LargeInteger` is  $(2^{130144} - 1)$ . Attempting to create a `LargeInteger` that exceeds this maximum will fail with an Integer overflow error.

## Printing Integers

Integers are printed by default, using `Integer >> asString`, in base 10. You may print using other bases by invoking `printStringRadix:` or `printStringRadix:showRadix:.`

For example,

```
1234 printStringRadix: 2
%
10011010010

-1234 printStringRadix: 16 showRadix: true
%
-16r4D2
```

## 6.2 Binary Floating Point

Floating point values in GemStone are composed of `SmallDoubles` and `Floats`. The most commonly used floating points will be `SmallDoubles`. While both `SmallDouble` and `Float` represents 8-byte binary floating point numbers, as defined in IEEE standard 754, `SmallDoubles` have a reduced exponent range. Some floating point values therefore can

only be represented by instances of `Float`, rather than `SmallDouble`. Similarly to `SmallInteger` and `LargeInteger`, GemStone operations return one or the other as needed.

The numerical behavior of instances of `Float` is implemented by the mathematics package of the vendor of the machine on which the Gem process is running. There are slight variations in results with different platform's implementation of the IEEE-754 standard.

You can get the components of a floating point value using the methods `signBit`, `exponent`, and `mantissa`.

## SmallDouble

`SmallDoubles` are special objects; as with `SmallIntegers`, the number itself is encoded in the OOP, making instances small and fast. They are also unique, so `SmallDoubles` of the same value are identical (`==`) as well as equal (`=`).

Each `SmallDouble` contains a 61 bit value, in IEEE format but with reduced exponent range. There is 1 sign bit, 8 bits of exponent and 52 bits of fraction. `SmallDoubles` are always in big-endian format (both on disk and in memory).

`SmallDoubles` can represent C doubles that have value zero or that have exponent bits in range 0x381 to 0x3ff, which corresponds to about 5.0e-39 to 6.0e38; approximately the range of C 4-byte floats.

## Float

Floats are not special objects; they require an OOP.

Each `Float` contains a 64 bit value in IEEE format, with 1 sign bit, 11 bits of exponent and 52 bits of mantissa. Floats are in cpu-native byte order when in memory, and the byte order of the extent when on disk.

In addition to the finite numbers, the IEEE standard defines floating point formats to include Infinity (positive and negative) and NaNs (not a Number), which can be quiet or signaling. NaNs results from an operations whose result is not a real number, such as:

```
-23 sqrt
%
PlusQuietNaN
```

Infinity results from operations that return a value outside the range of representation, such as:

```
32.0 / 0
%
PlusInfinity
```

`ExceptionalFloats` are named, unique instances of `Float`, not of `SmallDouble`. `ExceptionalFloats` include:

```
PlusInfinity
MinusInfinity
PlusQuietNaN
MinusQuietNaN
PlusSignalingNaN
MinusSignalingNaN
```

Since the sign of NaNs is not defined, GemStone operations return only positive NaNs; they do not return `MinusQuietNan` or `MinusSignalingNan`.

An unusual quality of NaNs is that they are not equal to themselves. This means that NaNs can cause problems if used as keys of hashed equality-based collections.

```
PlusQuietNaN = PlusQuietNaN
%
false
```

## Signalling Exception rather than returning Exceptional Float

When performing operations on Floats, an ExceptionalFloat may not always be an appropriate result.

You can determine if a number is an ExceptionalFloat using the message `#isExceptionalFloat`.

You can configure your system to signal an exception, rather than return an ExceptionalFloat. The following are the types of Floating point error conditions that may arise:

- ▶ `#divideByZero`
- ▶ `#overflow`
- ▶ `#underflow`
- ▶ `#invalidOperation`
- ▶ `#inexactResult`

FloatingPointError has protocol to configure signalling for all or none of these error conditions, or any subset. For example,

```
FloatingPointError enableAllExceptions.
FloatingPointError enableExceptions: { #divideByZero }
```

After enabling exceptions, exceptional conditions will signal errors, rather than returning an exceptional Float, for the duration of that session.

### Example 6.1 Enabling floating point exceptions

```
topaz 1> printit
3 / 0.0
%
PlusInfinity
topaz 1> printit
FloatingPointError enableAllExceptions.
%
0
topaz 1> printit
3 / 0.0
%
ERROR 2724 , a FloatingPointError divideByZero
(FloatingPointError)
```

## Literal Floats

Literal numbers in evaluated code that include a decimal point by default create a `SmallDouble` or `Float`. If the value is in the `SmallDouble` range, a `SmallDouble` will be created, otherwise a `Float` will be created.

Literal floats may be specified using exponential notation. For example, `5.1e3` and `5.1e-3` are valid `SmallDouble` literals.

ANSI specifies that float values may have exponents `e`, `d`, or `q`. These exponents, as well as `E` and `D`, are legal in GemStone, but have the same result: a `SmallDouble` or `Float`. Likewise, the ANSI class names `FloatE`, `FloatD`, and `FloatQ` can be used in code, but all resolve to `Float` class.

Note that using a plus sign before the exponent is not allowed in literal floats, although it can be used to create floating points from strings (using `Float fromString:`). This avoids ambiguity with Smalltalk dialects that would interpret this as the addition operator. For example, `5.1E+3`, which historically GemStone would interpret as the same as `5.1E3`, is disallowed; code must either omit the `+`, or include white space to clarify the addition operator.

## Printing Binary Floating Points

`SmallDoubles` and `Floats` are printed by default using `asString` or `printString`, in the notation equivalent to the C `printf` expression `%.16g`. This provides a maximum of 16 significant digits, rounding the fractional portion and changing to exponent notation if the whole number portion has more than 16 digits.

You can use `asStringUsingFormat:` to control the details of how floating point numbers are formatted when printing. `asStringUsingFormat:` accepts an Array of three elements:

- ▶ an Integer between -1000 and 1000, specifying a minimum number of Characters in the result String. Negative arguments pad with blanks to the left, positive arguments pad to the right. Note that if the value of this element is not large enough to completely represent the `Float`, a longer String will be generated.
- ▶ an Integer between 0 and 1000, specifying the number of digits to display to the right of the decimal point. If the printed representation of the float requires fewer characters, the result is padded with blanks on the right. If the value is insufficient to completely specify the float, the value is rounded to fit.
- ▶ A Boolean indicating whether or not to display the magnitude using exponential notation. If true, exponential notation is used; if false, decimal notation.

For example:

```
12.3456 asString
%
12.3456

12.3456 asStringUsingFormat: #(-8 2 false)
%
12.35

12.3456 asStringUsingFormat: #(4 10 true)
%
1.2345600000e01
```

## 6.3 Other Rational Numbers

For some application, binary floating points are problematic, since there are common decimal values that cannot be expressed exactly in binary floating point; for example, 5.1 does not have a precise binary floating point representation. This can make computation results incorrect. For example:

```
5.1 * 100000
%
509999.9999999999
```

There are several options to avoid this: Fraction, FixedPoint, ScaledDecimal, and DecimalFloat. These classes are independent of each other, and each provides different qualities of precision and range.

### Fractions

Fractions precisely represent rational numbers. Fractions are composed of an integer numerator and an integer denominator. As the ratio of two Integers, fractions can represent any rational number to an unbounded level of precision.

The display of fractions is as the numerator and denominator separated by the `$/` character, which is also the division binary method. Fractions have no literal representation. An expression such as `1/3`, which performs a division of two Integers, will return a fraction if the result is not an Integer.

```
(1/3) printString
%
1/3
```

Any expression, not just division expressions, that could result in fractions will be reduced automatically, to the lowest fraction or to an Integer.

```
(5/6) + (1/6)
%
1
```

## SmallFraction

SmallFractions are special objects, in which the OOP itself encodes the value. As with SmallDouble and Float, creating a fraction will result in either an instance of SmallFraction or Fraction, depending on the specific value.

SmallFractions can hold objects with numerators between -536870912 and 536870911, and denominators from 1 to 134217727.

## Fraction

If the numerator or denominator is outside the SmallFraction range, an instance of Fraction is created. These are not special objects.

## ScaledDecimals

ScaledDecimals represent a decimal number to the precision of a fixed number of fractional digits. ScaledDecimals are composed of an integer mantissa and a power-of-10 scale.

Literal ScaledDecimal or SmallScaledDecimals can be created using the s notation; for example, 1.53s2. This is not an exponential notation; the 2 here is the scale, and mantissa is resized appropriately. The values 1.53s2, 1.53s3, and 1.53s4 are all equal.

The number of fractional digits must not be greater than the scale.

For returned values from mathematical operations, ANSI does not precisely specify the scale of a returned ScaledDecimal. The following rules are used:

- ▶ For unary messages, the scale of the result equals the scale of the receiver.
- ▶ For a one-argument message, the scale of the result is the greater of the scale of the receiver and argument. An integer receiver or argument coerced to a ScaledDecimal should effectively have a scale of zero, meaning the result will have the scale of the non-coerced ScaledDecimal argument or receiver.

For some mathematical operations, the returned value type is a ScaledDecimal, but the returned value cannot always be exactly represented as a ScaledDecimal with the correct scale. In these cases, the results are rounded using the following rules:

- ▶ Following the example of IEEE754 float rounding, the ScaledDecimal that is answered is selected as though we computed the numerically exact value and then chose the closest representable ScaledDecimal of the scale specified by the rules. If the numerically exact value falls exactly halfway between two adjacent representable ScaledDecimal values of the scale specified by the rules, the ScaledDecimal with an even least significant digit is answered.

## SmallScaledDecimal

SmallScaledDecimals are special objects, in which the OOP itself encodes the value. As with SmallDouble and Float, creating a ScaledDecimal will result in either an instance of SmallScaledDecimal or ScaledDecimal, depending on the specific value.

SmallScaledDecimals can hold values with a scale  $\leq 31$ , and a mantissa in the range  $-2^{50}$  (-1125899906842624) to  $2^{50}$  (1125899906842623).

## ScaledDecimal

If the scale or mantissa is outside the SmallScaledDecimal range, an instance of ScaledDecimal is created. These are not special objects. The maximum scale of a ScaledDecimal is 30000.

While ScaledDecimals represent decimal fractions to the precision specified, most floating point values cannot be represented exactly by ScaledDecimals. The system will error if there is an attempt to create a ScaledDecimal with more than 39177 decimal digits.

## FixedPoints

FixedPoints, like Fractions, represents rational numbers, but also include information on how they should be displayed. A FixedPoint is composed of an integer numerator, integer denominator, and an integer scale. Like Fraction, this allows rational numbers to be represented with unbounded precision, and since fractional arithmetic is used in calculations, numerical results do not lose precision. The scale provides automatic rounding when representing the FixedPoint as a String.

FixedPoint is a legacy format provided for backwards compatibility. New code should use ScaledDecimal or Fraction.

FixedPoint uses a literal notation using p, such as 1.23p2. This is not an exponential notation; the 2 here specifies scale. The values 1.23p2, 1.23p3, and 1.23p4 are all equal.

## DecimalFloat

DecimalFloats represent base 10 floating point numbers, per IEEE standard 854-1987.

Literal DecimalFloats can be specified in exponential notation using the f or F character; for example, 5.432F2 creates a DecimalFloat equivalent to 543.2.

Objects of class DecimalFloat have 20 digits of precision, with an exponent in the range -15000 to +15000. The first byte encodes the sign and kind of the floating-point number. Bit 0 is the sign bit. The values in bits 1 through 3 indicate the kind of DecimalFloat:

```
001x = normal
010x = subnormal
011x = infinity
100x = zero
101x = quiet NaN
110x = signaling NaN
```

Bytes 2 and 3 encode the exponent as a biased 16-bit number (byte 2 is more significant). The actual exponent is calculated by subtracting 15000. Bytes 4 through 13 form the mantissa of the number. Each byte holds two BCD digits, with bits 4 through 7 of byte 4 containing the most significant digit.

Similarly to Float, operations that would not result in a real number, or that produce a result outside the representable range, result in Exceptional numbers:

```
DecimalPlusInfinity
DecimalMinusInfinity
DecimalPlusQuietNaN
DecimalMinusQuietNaN
DecimalPlusSignalingNaN
DecimalMinusSignalingNaN
```

You can determine if a number is an ExceptionalFloat using the message `#isExceptionalFloat`.

## Summary of literal syntax

The following table lists the notations that may appear in a literal number.

#	radix notation
d, D	SmallDouble/Float exponential notation
e, E	SmallDouble/Float exponential notation
f, F	DecimalFloat exponential notation
p	FixedPoint notation
q	SmallDouble/Float exponential notation
s	ScaledDecimal notation
r	radix notation

## Custom numeric literals

You can instruct the compiler to understand a new numerical literal format by sending a message to your customized subclass of Number to register that format.

The following method provides this registration:

```
Number >> parseLiterals: aCharacter exponentRequired: aBoolean
```

Once this is sent to an instance of a subclass of Number, when the compiler encounters a numeric value using *aCharacter*, it will send `fromString:` to that class.

- ▶ The subclass of Number must implement `fromString:` in such a way as to be able to read the new literal format, and create the new instance.
- ▶ *aCharacter* must be an alphabetic Character with `codePoint <= 127`, and may not be an existing numeric literal character as listed in the table on page 99.
- ▶ *aBoolean* indicates if digits following the exponent are required or not.

For example, say you have defined a class `ComplexNumber`. For the literal format, you wish to use `NiM`, where *N* represents the real part and *M* represents the imaginary part. So for example, `4.5+5i` would be specified using the literal form `4.5i5`.

First, you would define the `ComplexNumber >> fromString:` method, which will parse a string of the form `NiM` and return the new instance of `ComplexNumber`.

Then, to allow the literals to be included in code, send the following message.

```
ComplexNumber parseLiterals: $i exponentRequired: true
```

Now, assuming you have implemented the behavior appropriately, the compiler can evaluate expressions of the form:

```
(3.5i5 + 7.1i3) asString
%
10.6i8.0
```

Once invoked, the new literal format will be recognized until the session logs out.

Note that for subsequent logins, compiled references to that literal will continue to be valid, but unless the method is invoked again, methods with that literal cannot be recompiled. Including the invocation of `parseLiterals:exponentRequired:` in session initialization code (such as using `loginHook:`) is recommended.

To uninstall a custom literal without logging out, use the same method, passing in `nil` for *aBoolean*. For example,

```
ComplexNumber parseLiterals: $i exponentRequired: nil
```

## 6.4 Result classes, Conversion, and Rounding

When a mathematical operation is performed on numeric values, the class of the result depends on the specific classes of the argument and receiver.

By default, the result will be according to ANSI, per the following table.

**Table 6.2 Default result class for mathematical operations**

Receiver class	Argument Integer	Argument Float	Argument ScaledDecimal	Argument Fraction
Integer	Integer	Float	ScaledDecimal	Fraction (may reduce to Integer)
Float	Float	Float	Float	Float
ScaledDecimal	ScaledDecimal	Float	ScaledDecimal	Fraction (may reduce to Integer)
Fraction	Fraction (may reduce to Integer)	Float	Fraction (may reduce to Integer)	Fraction (may reduce to Integer)
FixedPoint, DecimalFlat	Integer	Float	ScaledDecimal	Fraction

So, for example, the result of  $1 * 0.5$  is a `SmallDouble`, and the result of  $1/3 - 1.50s2$  is a `SmallFraction`.

Note that the legacy classes `FixedPoint` and `DecimalFloat` do not reliably conform to this standard; this maintain historic behavior for legacy applications.

### Conversion

Conversion is supported between kinds of numbers and from kinds of numbers to `String`, using `asInteger`, `asFloat`, `asFraction`, and so on; see the image for specifics. Note that converting to a `ScaledDecimal` or `FixedPoint` requires a scale argument.

In cases when floating points are involved, the precision limits may result in a converted value that is slightly different than you might expect.

For example, when printing the following `SmallDouble` using `asStringLegacy` (which uses exponential notation to print all digits of precision), the printed value varies slightly at the maximum digits of precision.

```
1.111222233334444 asStringLegacy
%
1.1112222333344439E00
```

## Truncation and Rounding

The methods `roundTo:` and `truncateTo:`, and related selectors, allow you to reduce unnecessary fractional parts.

For display, this can also be done using `asStringUsingFormat:` methods that allow you to specify the number of digits of a float, for example.

`truncateTo:` returns the multiple of the argument that is closest to the receiver towards zero.

`roundTo:` returns the multiple of the argument that is closest to the receiver in either direction; if the receiver is exactly between two multiples, the multiple farther from zero is returned. `roundedHalfToEven:` allows you to round, and if the receiver is exactly halfway between multiples, return the even one.

The result of `truncateTo:` and `roundTo:` follows the ANSI standard, and conforms to table on Table 6.2 on page 100. However, for truncation and rounding it is useful to be get a result that is the class of the argument; since, for example, `3.123 roundTo: 1` returns a `SmallDouble`, `3.0`.

The methods `truncated` and `rounded` perform `truncateTo:/roundTo:` to 1, and return an `Integer`. When truncating or rounding other than to 1, you can use the methods `truncateAndCoerceTo:` and `roundAndCoerceTo:`, which return an instance of the class of the argument.

## 6.5 Dates and Times

GemStone supports the date-time classes `Date`, `Time`, and `DateTime` and `DateAndTime`. Using date-time classes requires taking into account the `TimeZone`, the current local offset from UTC (GMT).

### Date

An instance of `Date` describes a month, day and year in time.

`Date` supports dates after December 31, 1900. While `Dates` earlier than this can be created and operated on, they are not tested and not officially supported. Technically legal years are in the range  $-2^{31}$  (-2147483648) to  $2^{31} - 1$  (2147483647).

### SmallDate

An instance of `SmallDate` is a special (that is, the OOP encodes the value), encoding the day of year and the year within the OOP. Newly created `Date` instances are instances of `SmallDate`.

All possible supported Dates can be represented as SmallDates. All newly created instances of Date are returned as instances of SmallDate, but existing instances of Date remain in upgraded repositories. Date and SmallDate are interoperable with each other, and the following discussion uses Date to refer to either Date or SmallDate.

## Instance Creation

Date instances can be created from formatted strings or streams, or by using instance creation methods that allow you to specify the properties. See the image for the available options. For example:

```
Date today
Date fromString: '21/07/2023'
Date newDay: 21 month: 'July' year: 2023
```

## Formatting for Instance Creation and Printing

Date supports printing methods and reading methods that use an explicit string-formatting Array. For Date, this is a 6-element Array that is passed to Date >> asStringUsingFormat:, or fromString: usingFormat: and related instance creation methods.

- ▶ Elements 1, 2 and 3 are integers that determine the position of the day (1), month (2), and year (3). So year, month, day would be 3, 2, 1.
- ▶ Element 4 is a character for the date separator, such as \$/.
- ▶ Element 5 is an integer that determines if the month format is printed as a number (1), three-letter abbreviation (2), or the entire name (3).
- ▶ Element 6 is an integer that determines if the year format is the entire number (1), or only the last two digits (2).

The default is DD/MM/YYYY, equivalent to #(2 1 3 \$/ 1 1). Note that the day is first, not the US convention of month first.

```
(Date fromString: '21/07/2023') monthName
  'July'

(Date today) asStringUsingFormat: #(3 2 1 $. 1 1)
  '2023.07.21'
```

## Time

An instance of Time describes a time of day.

### SmallTime

An instance of SmallTime is a special (that is, the OOP encodes the value), in microseconds resolution, using 56 bits of the OOP. Since there are 86400000000 microseconds in a day, all possible Times can be represented as SmallTimes. All newly created instances of Time are returned as instances of SmallTime, but existing instances of Time remain in upgraded repositories. Time and SmallTime are interoperable with each other, and the following discussion uses Time to refer to either Time or SmallTime.

## Instance Creation

Time instances can be created from formatted strings or streams, or by using instance creation methods that allow you to specify details. See the image for the available options. For example:

```
Time now
Time fromString: '17:23:05'
Time fromSeconds: 34234
```

## Formatting for Instance Creation and Printing

Explicit string-formatting specifications take the form of a 3-element Array that is passed to `Time >> asStringUsingFormat:`, or `fromString:usingFormat:` and related instance creation methods. In this array:

- ▶ Element 1 is a character for the date separator, such as \$:
- ▶ Element 2 is a boolean; if true, seconds are included, or false then they are omitted.
- ▶ Element 3 is a boolean; if true, then the time is in 12-hour format with am or pm; if false then the time is in 24-hour format. When using format to create an instance from a String or Stream, there must be a space before the am/pm.

The default is #(\$: true false), equivalent to HH:MM:SS.

For example

```
Time fromString: '23.34' usingFormat: #( $ . false false)
23:34:00

Time now asStringUsingFormat: #($: false true)
'02:30 PM'
```

## Time offset

Time instances themselves have no `TimeZone`. Class creation methods, and methods that print or display Times, that have "gmt" in the selector, interpret the time as GMT and apply the offset from the repository setting for `TimeZone current`.

For example, in the `America/LosAngeles` `TimeZone`:

```
Time fromStringGmt: '11:30:00'
04:30:00

(Time fromStringGmt: '11:30:00') asStringGmt
'11:30:00'
```

## DateTime

`DateTime` represents a point in time with millisecond resolution. `DateTime` supports points in time after midnight on December 31, 1900 at midnight GMT. While Dates earlier than this can be created and operated on, they are not tested and not officially supported.

A `DateTime` includes the year, day of year, milliseconds, and a `TimeZone`.

## Instance Creation

Date instances can be created from formatted strings or streams, or by using instance creation methods that allow you to specify the properties. See the image for the available options. For example,

```
DateTime now
DateTime fromString: '21/07/2023 23:15:26'
DateTime newWithYear: 2023 month: 7 day: 21 hours: 15
    minutes: 26 seconds: 0
```

## Formatting for Instance Creation and Printing

Explicit string-formatting specifications take the form of a Array of elements that is passed to `DateTime >> asStringUsingFormat:`, or `fromString:usingFormat:` and related instance creation methods.

This Array must have at least 8 elements; 4 additional elements are optional.

In this array:

- ▶ Elements 1, 2 and 3 are integers that determine the position of the day (1), month (2), and year (3). So year, month, day would be 3, 2, 1.
- ▶ Element 4 is a character for the date separator, such as \$/.
- ▶ Element 5 is an integer that determines if the month format is printed as a number (1), three-letter abbreviation (2), or the entire name (3).
- ▶ Element 6 is an integer that determines if the year format is the entire number (1), or only the last two digits (2).
- ▶ Element 7 is a character for the time separator, such as \$:
- ▶ Element 8 is a boolean, if true the time is included; if false, the time is omitted and values for elements 7, 9, and 10 are ignored. Elements following this are optional.
- ▶ Element 9 is a boolean, if true seconds are included, or false then they are omitted.
- ▶ Element 10 is a boolean, if true the time is in 12-hour format with am or pm; if false then the time is in 24-hour format. When using format to create an instance from a String or Stream, there must be a space before the am/pm.
- ▶ Element 11 is a boolean, if true the time zone information is included; if false it is omitted.
- ▶ Element 12 is a boolean, if true then the print is printed per the time zone in which the DateTime is created, not in the local time zone. If false then the DateTime is printed according to the local time zone (TimeZone current).

The default is (1 2 3 \$/ 1 1 \$: true true false false false), equivalent to (DD/MM/YYYY HH:MM:SS).

For example, to read a non-default formatted string, and create an instance of DateTime, which then by default prints in the default format:

```
DateTime fromString: 'January 7 2023 9:45 pm' usingFormat:
    #(2 1 3 $ 3 1 $: true false true false)
    07/01/2023 21:45:00
```

## DateAndTime

DateAndTime is the ANSI compliant class that is equivalent to DateTime, and represents a point in time. DateAndTime instances know their microseconds (plus or minus) since 00:00:00 on January 1 2001, UTC (not counting leap seconds), and the TimeZone offset in seconds.

There is no restriction on the year range for DateAndTime. DateAndTimes are in the Gregorian calendar (the common calendar adopted in Europe in 1582). Date times prior to the adoption of this calendar are given in the retrospective astronomical Gregorian calendar. The year 1 A.D. is astronomical Gregorian year 1, the year 1 B.C. is astronomical Gregorian year 0, the year 2 B.C. is astronomical Gregorian year -1, and so on.

DateAndTime instances are always created with a resolution of 6 decimal places, equivalent to microsecond resolution.

### SmallDateAndTime

An instance of SmallDateAndTime is a special (that is, the OOP encodes the value), that represents DateAndTimes in the range of from January 1, 2001 to May 10, 2072, that have timezone offsets in integer hours, in the range -16..15. Creating a DateAndTime will return an instance of SmallDateAndTime if it is within the range, otherwise it will return an instance of DateAndTime.

### Instance Creation

DateAndTime instances can be created from formatted strings or streams, or by using instance creation methods that allow you to specify details. See the image for the available options. For example:

```
DateAndTime now
DateAndTime fromString: '2023-07-21T15:26:00-08:00'
DateAndTime year: 2023 month: 07 day: 21 hour: 15
              minute: 26 second: 0 offset: (Duration seconds: -28800)
```

### Formatting for Instance Creation and Printing

The format for reading and printing DateAndTime is

```
[ - ]YYYY-MM-DDTHH:MM:SS[ .SSSSSS ]±hh:mm
```

where:

- ▶ - is a minus sign if the receiver represents a time with a year less than zero, and is optional.
- ▶ DD is the number of complete days, with leading zeros to fill two places
- ▶ HH is the number of complete hours, with leading zeros to fill two places
- ▶ MM is the number of complete minutes, with leading zeros to fill two places
- ▶ SS is the number of complete seconds, with leading zeros to fill two places
- ▶ ..SSSSSS is the fractional part of the number of seconds, and is optional. 1 to 6 digits of fractional seconds are supported; additional digits may be used but are rounded to 6 digits.

- ▶  $\pm$  a + or - is required
- ▶ hh is the number of complete hours of the offset, with leading zeros to fill two places
- ▶ mm is the number of complete minutes of the offset, with leading zeros to fill two places

Sending `printString` returns all digits of the `DateAndTime` seconds, while `asString` does not print fractional seconds.

For example:

```
(DateAndTime fromString: '2023-07-21T13:06:15-00:00')
printString
'2023-07-21T13:06:15+00:00'

(DateAndTime fromString: '2023-07-21T13:06:15.12356789-07:00')
printString
'2023-07-21T13:06:15.123568-07:00'

(DateAndTime fromString: '2023-07-21T13:06:15.12356789-07:00')
asString
'2023-07-21T13:06:15-07:00'
```

## TimeZone

Each instance of `DateTime` includes a reference to a `TimeZone` object, which handles the conversion from the internally stored Greenwich Mean Time (GMT)/Coordinated Universal Time(UTC) and the local time. `TimeZones` are also used to determine the local time offset for instances of `DateAndTime` now.

`TimeZones` encapsulate the daylight savings time (DST) rules, so a given GMT/UTC time is adjusted to local time based on `TimeZone` and the specific date.

Each session has a current `TimeZone`, which is used to display times, and to create instance of date and time classes when using methods that do not explicitly specify the `TimeZone`.

The current `TimeZone` (`TimeZone current`) is set from the default `TimeZone` (`TimeZone default`) during login. The default `TimeZone` is persistent, and can only be updated by `SystemUser`. The GemStone distribution comes with the `America/Los_Angeles` `TimeZone` installed as the default `TimeZone`. This is described in the *GemStone/S 64 Bit Installation Guide*, and instructions for updating the current `TimeZone` are in the *System Administration Guide*.

`DateTime` method selectors and comments refer to GMT, which is the older term; `DateAndTime` refers to UTC, which is the current preferred term.

## 6.6 Internationalizing

The convention for expressing numbers and dates varies in different countries. GemStone supports several ways to handle customization from the default display.

### Dates in GemStone log files

The timestamps printed in the log headers and in log messages are formatted according to the current system locale. You can override this using the `GS_CFTIME` environment variable. If this is set in the environment for the process, then the setting is used to control printing in log headers and log messages.

The setting for `GS_CFTIME` must be a valid `strftime` format string, and **must** contain fields for:

- ▶ Month: `%m` or `%b` or `%B` or `%h`
- ▶ Day: `%d`
- ▶ Hour: `%H`, or `%I` and `%p`, or `%I` and `%P`
- ▶ Minutes: `%M`
- ▶ Seconds: `%S`

If these criteria are not met, the default date format based on the system's `LOCALE` is used, or otherwise the US-centric date format. See the man page for `strftime`, for details on the formats.

For example,

```
EXPORT GS_CFTIME="%Y-%m-%d %H:%M:%S %Z"
```

Will include lines such as the following in the Gem log header:

```
| PROCESS ID: 2120954 DATE: 2023-07-21 16:20:53 PDT (UTC -7:00)
...
| DATEFORMAT: %Y-%m-%d %H:%M:%S %Z
```

### Internationalizing Decimal Points using Locale

The class `Locale` allows you to obtain operating system locale information and use or override it in GemStone. GemStone currently only uses the `decimalPoint` setting, to provide localized reading and writing of numbers involving decimal points. Updates to `Locale` are stored in session state, and only persist for the lifetime of the session. They are not affected by `commit` or `abort`.

To override the operating system locale information, use the following message:

```
Locale class >> setCategory: categorySymbol locale: LocaleString
```

Note that the `LocaleString` passed to `setCategory:locale:` must be defined on the host machine. If the given locale is not found, this method will return `nil`. You can use the UNIX command `locale -a` to get a list of all available `LocaleStrings`. To check the decimal point, the following method returns the `decimalPoint` setting for the current `Locale`:

```
Locale decimalPoint
%
```

While there are a number of Locale category symbols, the only ones that are of use in this release are `#LC_NUMERIC` and `#LC_ALL`, either of which will set the category that affects the decimal point.

For example, To use decimal localization appropriate for Germany:

```
Locale setCategory: #LC_NUMERIC locale: 'de_DE'.
```

To reset to UNIX default value, using period:

```
Locale setCategory: #LC_ALL locale: 'C'.
```

### **Smalltalk code requires the period separator**

Note that Smalltalk syntax requires the use of "." as the decimal point separator, so expressions involving literal floating point numbers within Smalltalk code require use of the period, regardless of Locale.

Smalltalk code that includes literal floating point values with a comma separator will fail to compile, regardless of the Locale settings.

In order to be able to export and input numerical values regardless of the Locale of a particular session, methods whose printed form includes the decimal point provide the following set of methods:

```
(instance method) asStringLocaleC  
(class method) fromStringLocaleC:
```

These methods use a period as a decimal separator, regardless of the current Locale.

## 6.7 Random Numbers

### Universally unique identifiers (UUIDs)

Universally unique identifiers (UUIDs) provide a practically unique 128-bit number that does not rely on a registry or a shared resource. UUIDs are generated randomly using the secure OpenSSL random number generator.

The class `GsUuidV4` implements a version 4 UUID as specified in RFC 4122, A Universally Unique Identifier (UUID) URN Namespace.

Instances of `GsUuidV4` are invariant and cannot be modified.

To create an instance of UUID:

```
GsUuidV4 class >> new
```

```
GsUuidV4 class >> fromString: aUuidString
```

*aUuidString* must be a valid UUID version 4 string in the following format:

```
xxxxxxxx-xxxx-4xxx-Vxxx-xxxxxxxxxxxx
```

where x is any valid lower-case hex digit and V is one of 8, 9, a or b.

### Random Number Generator

The class `Random` and its subclasses provide random number generation.

There are two types of random number generation, which correspond to separate subclass hierarchies. The `SeededRandom` subclasses provide random numbers generated within GemStone code, using a starting seed value. The `HostRandom` subclass provides access to the host operating system's `/dev/urandom` random number generator.

The class hierarchy of the `Random` classes are:

```
Object
  Random (abstract)
    HostRandom
    SeededRandom (abstract)
      Lag1MwcRandom
      Lag25000CmwcRandom
```

#### Random

The `Random` class is an abstract superclass for the random number generators. It also can be used to create an instance of a default random number generator class.

`Random new` will return an instance of `HostRandom`, the most basic kind of generator based on host OS `/dev/urandom`.

`Random seed:` will return an instance of `Lag1MwcRandom`. `HostRandom` does not support seeds.

While an instance of `Lag25000CmwcRandom` takes some time to create, it can produce a more fair and longer-period series of random numbers that are generated much more quickly than is done by the other `Random` subclasses.

Once you have an instance of a concrete subclass of `Random`, you can generate random numbers or collections of random numbers with the following range and type specifications:

```
float - a random Float in the range [0,1)
```

`floats: n` - a collection of  $n$  random floats in the range  $[0,1)$

`integer` - a random non-negative 32-bit integer, in the range  $[0,2^{32}-1]$

`integers: n` - a collection of  $n$  random non-negative integers in the range  $[0,2^{32}-1]$

`integerBetween: l and: h` - a random integer in the range  $[l,h]$ .  $l$  and  $h$  should be less than approximately  $2^{31}$ .

`integers: n between: l and: h` - a collection of  $n$  random integers in the range  $[l,h]$ .  $l$  and  $h$  should be less than approximately  $2^{31}$ .

`smallInteger` - Answer a random integer in the SmallInteger range,  $[-2^{60},2^{60}-1]$

Subsequent calls to the same instance will generate new random numbers.

You should create an instance of a Random subclass and retain that to generate many random numbers, rather than creating new instances of a Random subclass.

### HostRandom

HostRandom allows access to the host operating system's `/dev/urandom` random number generator.

HostRandom is much slower to generate numbers than the other subclasses of Random, but does not have the overhead of creating an instance. On some platforms, `/dev/urandom` may be intended to be a cryptographically secure random number generator, which none of the other subclasses are. It also has the advantage of not needing an initial seed, and so is good for generating random seeds for other Random subclasses.

HostRandom uses a shared singleton instance, which is accessed by sending `#new` to the class HostRandom. Sending `#new` has the side effect of opening the underlying file `/dev/urandom`. This file normally remains open for the life of the session, but if you wish to close it you can send `#close` to the instance, and later send `#open` to reopen it. If you store a persistent reference to the singleton instance the underlying file will not be open in a new session and you must send `#open` to the instance before asking for a random number.

Since HostRandom is a service from the operating system, it cannot be seeded, and should not be used when a repeatable random sequence of numbers is needed.

### SeededRandom

SeededRandom is an abstract superclass for classes that generate sequences of random numbers that can be generated repeatedly by giving the same initial seed to the generator.

In addition to creating new instances using the class methods `new` and `seed:`, the following instance methods allow repeatable sequences to be generated:

`seed: aSmallInteger`

Sets the seed of the receiver from the given seed, which can be any SmallInteger. The subsequent random number sequence generated will be the same as if this generator had been created with this seed.

`fullState, fullState: stateArray`

The internal state of a generator is more than can be represented by a single SmallInteger. These messages allow you to retrieve the full state of a generator at any time, and to restore that state later. The random number sequence generated after the restoration of the state will be the same as that generated after the retrieval

of the state. You might, for instance, allow a generator to get its initial state from `/dev/urandom`, then save this state so the random sequence can be repeated later.

### Lag1MwcRandom

Lag1MwcRandom is faster to create than Lag25000CmwcRandom, since it can be seeded by a single 61-bit SmallInteger, rather than a seed of more than 800000 bits as required by Lag25000CmwcRandom. After creation, however, it is slower, and it is not perfectly fair, and has a shorter period. It can be used when a small number of seeded random numbers are needed.

### Lag25000CmwcRandom

Lag25000CmwcRandom is a seedable random generator with a period of over  $10^{240833}$ . It is a lag-25000 generator using the complementary multiply-with-carry algorithm to generate random numbers. Its period is so long that every possible sequence of 24994 successive 32-bit integers appears somewhere in its output, making it suitable for generating random n-tuples where  $n < 24994$ . Its output is fair in that the number of 0 bits and 1 bits in the full sequence are equal.

While this generator is recommended for most uses, it is **not** cryptographically secure, so for applications such as key generation you should consider using HostRandom, once you satisfy yourself that HostRandom is secure enough on your operating system.

You can also allow the seed bits to be initialized from the HostRandom, then retrieve that state by sending `#fullState`. That state can later be restored by sending the retrieved state as an argument to `#fullState`.



# Files and Directories

---

Many GemStone applications will need to read and write files. GemStone provides two interfaces; the older GsFile interface, and the new FileSystem interface.

This chapter explains how to use GsFile and FileSystem to read, write, and manage files and directories.

**Accessing Files using GsFile** (page 113)

describes the protocol provided by class GsFile to open and close files, read their contents, and write to them.

**FileSystem** (page 121)

describes FileSystem, a set of classes ported and adapted from Pharo to support a wide range of File and Directory operations.

## 7.1 Accessing Files using GsFile

The class GsFile provides basic protocol to create and access operating system files. Most functions of GsFile class can also be performed using FileSystem.

This section provides a few examples of the more common operations using GsFile. For more information, see the GsFile methods in the image.

Instances of GsFile understand most protocol common to Streams.

### Client vs. Server files

Many of the methods in GsFile distinguish between files on the client and those on the server machine. In this context, the term *client* refers to the machine on which the GCI interface is executing, and the *server* refers to the machine on which the Gem is executing. This terminology is historic, and may be misleading, since the server in this case does not mean the machine the Stone is running on, if the Gem is remote from the Stone.

In the case of a linked interface, the interface and the Gem execute as a single process, so the client machine and the server machine are the same.

In the case of an RPC interface, the interface and the Gem are separate processes, and the client machine may be different from the server machine. When using GBS or topaz on Windows, GsFile client methods can operate on files in the Windows file system.

Methods that include "OnServer" operate on server files, accessible in the file system of the machine that the Gem is running on. Methods that do not mention "OnServer" operate on files on the machine that the client is running on.

GsFile is implemented in UserActions (for historic reasons), and thus client-side file access is not supported in external sessions using GsTsExternalSession or GsExternalSession, nor within code invoked via GciNb\* functions.

## Specifying Files

Many of the methods in the class GsFile take as arguments a *file specification*, which is any string that constitutes a legal file specification in the operating system under which GemStone is running. Wildcard characters and environment variables are legal in a file specification.

### Specifying Files Using Environment Variables

If you supply an environment variable instead of a full path when using the methods described in this chapter, the way in which the environment variable is expanded depends upon whether the process is running on the client or the server machine.

- ▶ If you are running a linked interface or you are using methods that create processes on the server, the environment variables accessed by your GemStone Smalltalk methods are those defined in the shell under which the Gem process is running.
- ▶ If you are running an RPC interface and using methods that create processes on a separate client machine, the environment variables are instead those defined by the remote user account on the client machine on which the application process is running.

## Creating a File

You can create a new operating system file from GemStone Smalltalk, using class methods to open a GsFile for write or append; these methods will create a file if it does not exist, as well as returning the open file.

Example 7.1 creates a file named aFileName in the current directory on the server.

### Example 7.1

```
myFilePath := 'aFileName'.
myFile := GsFile openWriteOnServer: myFilePath.
"Here would go code to write data to the file"
myFile close
```

Example 7.2 creates a file named aFileName in the current directory on the client machine.

**Example 7.2**


---

```
myFilePath := 'aFileName'.
myFile := GsFile openWrite: myFilePath.
"Here would go code to write data to the file"
myFile close
```

---

These methods return the instance of GsFile that was created, or nil if an error occurred. See “GsFile Errors” on page 120 for more on how to find the details of why a GsFile operation returned nil.

**Opening a File**

GsFile provides a wide variety of protocol to open files. For a complete set of methods, see the image. These methods return the GsFile instance if successful, or nil if an error occurs..

**Table 7.1 GsFile Class Methods to Open Files**

Method	Description
openReadOnServer: openReadOnServerCompressed:	Opens a file on the server for reading.
openWriteOnServer: openWriteOnServerCompressed:	Opens a file on the server for writing. Creates a new file if one does not exist, or truncates an existing file to 0.
openAppendOnServer:	Opens a file on the server for reading, appending the new contents instead of replacing the existing contents. Creates the file if it does not exist.
openUpdateOnServer:	Opens a file on the server for both reading and writing. Creates the file if it does not exist.
openRead: openReadCompressed:	Opens a file on the client machine for reading.
openWrite: openWriteCompressed:	Opens a file on the client machine for writing. Creates a new file if one does not exist, or truncates an existing file to 0.
openAppend:	Opens a file on the client machine for writing, appending the new contents instead of replacing the existing contents. Creates the file if it does not exist.
openUpdate:	Opens a file on the client machine for both reading and writing. Creates the file if it does not exist.

## Closing a File or Files

The following methods close the current instance, or multiple open files:

**Table 7.2 GsFile Method Summary**

Method	Description
GsFile class >> closeAllOnServer	Closes all open GsFile instances on the server except stdin, stdout, and stderr.
GsFile >> close	Closes the receiver.
GsFile class >> closeAll	Closes all open GsFile instances on the client machine except stdin, stdout, and stderr.

Your operating system limits the number of files a process can concurrently access. Instances of GsFile automatically have their C state closed when the instance is garbage collected or when a persistent instance drops out of memory. However, you should still close files as soon as you are done using them.

## Writing to a File

After you have opened a file for writing, you can add new contents to it in several ways.

For example, the instance methods `addAll:` and `nextPutAll:` take strings as arguments and write the string to the end of the file specified by the receiver. The method `add:` takes a single character as argument and writes the character to the end of the file. And various methods such as `cr`, `lf`, and `ff` write specific characters to the end of the file – in this case, a carriage return, a line feed, and a form feed character, respectively.

The write methods return the number of bytes that were written to the file, or `nil` if an error occurs.

For example, the following code writes the two strings specified to the file `myFile.txt`, separated by end-of-line characters.

### Example 7.3

```
myFile := GsFile openWrite: 'myFile.txt'.
myFile nextPutAll: 'All of us are in the gutter,'.
myFile cr.
myFile nextPutAll: 'but some of us are looking at the stars.'.
myFile close.
```

If the text you wish to write contains characters outside the ASCII range (that is, with `codePoints` greater than 127), then there are additional considerations.

Text with Characters with `codePoints` greater than 255 require more than one byte to represent. These must be encoded before they can be written to a GsFile. Encoding to UTF-8 is done by using `GsFile >> nextPutAllUtf8:` or by passing an instance of `Utf8` to the write methods.

For example, the Euro character € has the Unicode value U+20AC.

---

**Example 7.4 Writing the Extended Character € to a File**

---

```
| myfile str |
myfile := GsFile openWrite: 'extendedCharacterExample.txt'.
str := String new.
str add: 'How to write a Euro character '.
str add: (Character codePoint: 16r20AC).
str add: ' to a file'; lf.
myfile nextPutAllUtf8: str.
myfile close.
```

---

Text with Characters with codePoints in the range of 128..254 have ambiguity between the legacy output and UTF-8 encoding. Traditionally, GsFiles write Characters as byte data without encoding or decoding. Most modern systems encode files as UTF-8, and expect files to be encoded as UTF-8. However, to ensure legacy uses of GsFile do not create invalid Strings, the GsFile write methods continue to write data as bytes.

While unlikely, it is possible that Characters with codePoints in the range 128..254 could be either a portion of a UTF-8 encoding or an 8-bit character. Most often, specifying to default to the incorrect format will result in a badly formed UTF-8 error, or un-decoded bytes.

Since for ASCII Characters (codePoints in the 7-bit range), the legacy output and the UTF-8 encoding are the same, encoding all writes (and decoding all reads) is an effective way to remove ambiguity.

## Reading from a File

Instances of GsFile can be accessed in many of the same ways as instances of Stream subclasses. Like streams, GsFile instances also include the notion of a position, or pointer into the file. When you first open a file, the pointer is positioned at the beginning of the file. Reading or writing elements of the file ordinarily repositions the pointer as if you were processing elements of a stream.

A variety of methods allow you to read some or all of the contents of a file from within GemStone Smalltalk. For example, the `contents` method (at the end of Example 7.3) returns the entire contents of the specified file and positions the pointer at the end of the file.

In Example 7.5, `next: into:` takes the 12 characters after the current pointer position and places them into the specified string object. It then advances the pointer by 12 characters.

---

**Example 7.5**

---

```
| result |
result := String new.
myfile := GsFile openRead: 'myFileName'.
myfile next: 12 into: result.
myfile close
result.
```

---

To read a file containing data encoded in UTF-8, you may read the file as usual, and then send `decodeFromUTF8ToString` or `decodeFromUTF8ToUnicode` to decode the results. Alternatively, you may use the method

```
GsFile >> contentsAsUtf8
```

which you can then decode from the instance of `Utf8` similarly using `decodeToString` or `decodeToUnicode`.

Note that when reading files whose contents logically contain Characters with codePoints larger than 127, you must be aware of the whether the file is encoded in order to decode appropriately. `GsFile` reads the bytes and does not distinguish between encoded or un-encoded contents. A UTF-8 encoded file when read in using a `GsFile` and not explicitly decoded will be garbled, and a file written as 8-bit characters that you attempt to decode will almost always result in a badly formed UTF-8 error.

If you are writing a file in topaz format, for example source code, you may include a header line in the output file, either:

```
fileformat utf8  
fileformat 8bit
```

to instruct topaz of the file encoding.

## Positioning

You can also reposition the pointer without reading characters, or peek at characters without repositioning the pointer. The method:

```
GsFile peek
```

allows you to view the next character in the file without advancing the pointer.

To advance the pointer without reading the intervening characters, use:

```
GsFile skip: anInteger
```

To determine the position or set the position specifically, use:

```
GsFile position  
GsFile position: anInteger
```

## Testing Files

The class `GsFile` provides a variety of methods that allow you to determine facts about a file.

To test for existence of a file, use:

```
GsFile exists: aFileNameString  
GsFile existsOnServer: aFileNameString
```

These methods returns true if the file exists, false if it does not, and nil if an error occurred.

## Renaming Files

Files on the client or server can be renamed or moved. For example:

```
GsFile rename: '/tmp/myfile.txt' to: '/tmp/newname.txt'.  
  
GsFile renameFileOnServer: '$GEMSTONE/data/system.conf'  
to: '/users/david/mysystem.conf'.
```

## Removing Files

To remove a file from the client machine, use an expression of the form:

```
GsFile closeAll.  
GsFile removeClientFile: mySpec.
```

To remove a file from the server machine, use the method `removeServerFile:` instead. These methods return the receiver or `nil` if an error occurred.

## Examining a Directory

To get a list of the names of files in a directory, send GsFile the message `contentsOfDirectory: aFileSpec onClient: aBoolean`. This message acts very much like the UNIX `ls` command, returning an array of file specifications for all entries in the directory.

If the argument to the `onClient:` keyword is `true`, GemStone searches on the client machine. If the argument is `false`, it searches on the server instead.

For example:

### Example 7.6

---

```
GsFile contentsOfDirectory: '$GEMSTONE/examples/admin' onClient:  
false  
%  
an Array  
#1 /dbf/gsadmin/GS6437/examples/admin/.  
#2 /dbf/gsadmin/GS6437/examples/admin/..  
#3 /dbf/gsadmin/GS6437/examples/admin/onlinebackup.sh  
#4 /dbf/gsadmin/GS6437/examples/admin/archivelogs.sh
```

---

If the argument is a directory name, this message returns the full pathnames of all files in the directory, as shown in Example 7.6. However, if the argument is a filename, this message returns the full pathnames of all files in the current directory that match the filename. The argument can contain wildcard characters such as `*`. The following example shows a different use of this message.

```
GsFile contentsOfDirectory: '$GEMSTONE/ver*' onClient: false  
%  
an Array  
#1 /dbf/gsadmin/GS6432/version.txt
```

If you wish to distinguish between files and directories, you can use the message `contentsAndTypesOfDirectory: onClient:` instead. This method returns an array of pairs of elements. After the name of the directory element, a value of `true` indicates a file; a value of `false` indicates a directory. For example:

### Example 7.7

---

```
GsFile contentsAndTypesOfDirectory: '$GEMSTONE/uilib' onClient:  
false  
%  
a Array
```

---

```
#1 /dbf/gsadmin/GS6433/uilib/.  
#2 false  
#3 /dbf/gsadmin/GS6433/uilib/..  
#4 false  
#5 /dbf/gsadmin/GS6433/uilib/liboraapi23-643.so  
#6 true
```

---

All the above methods, like most GsFile methods, return nil if an error occurs.

## GsFile Errors

GsFile operations return nil in cases where an error occurs during the operation. For this reason, most GsFile operations should check for nil return. There are separate methods to check for errors within file operations on server files and client files.

To check for errors in an operation on a server file, the method is `GsFile >> serverErrorString`. It is nil if no error has occurred. This error is available until the next GsFile operation is executed.

### Example 7.8

---

```
| myFile |  
myFile := GsFile openReadOnServer: 'nonexistentfile'.  
myFile  
  ifNil: [GsFile serverErrorString]  
  ifNotNil: ['Successfully opened'].  
%  
No such file or directory : nonexistentfile
```

---

To check for similar errors for a client file, use the method `lastErrorString`. For example:

### Example 7.9

---

```
| myFile |  
myFile := GsFile openRead: 'privatefile'.  
myFile  
  ifNil: [GsFile lastErrorString]  
  ifNotNil: ['Successfully opened'].  
%  
Permission denied : privatefile
```

---

## Logging to stdout

GsFile class protocol allows you to write messages to stdout of either the Gem (server) or the client. Note that for clients running without a console (as may be the case, for example, using GBS on Windows), linked output may not be accessible.

## On the server (Gem)

```
GsFile gciLogServer: aString
```

`gciLogServer`: writes to stdout of the Gem. For an RPC login, this is the Gem log file. For a linked login on topaz, this is the console or (when using topaz) an output file as controlled by an **output push** command.

## On the client (GCI client)

```
GsFile gciLogClient: aString
```

`gciLogClient`: writes to stdout of the GCI application. This is the console for both linked and RPC logins (or on topaz, an output file as controlled by an **output push** command). If it is not possible to perform GsFile client writes (e.g, if within an external session or a non-blocking GCI execution), it will fall back to executing `gciLogServer`.

## 7.2 FileSystem

FileSystem is a set of classes, ported and adapted from Pharo, that support operations on files and directories. FileSystem provides a much more flexible and feature-rich environment than GsFile.

The term 'FileSystem' can be used to refer to the entire set of classes that support file and directory operations, or to the specific class named FileSystem. The specific class **FileSystem** represents the underlying file environment. Most operations within the general File System environment use the class **FileReference**, which represents a file or directory.

Unlike GsFile, FileSystem is implemented using FFI, which makes it easily extensible to support specific behavior under different OS platforms. However, FileSystem does not support operations on the client (that is, on the node on which the GCI client is running); only operations on the server (the node on which the Gem is running, which may or may not be the node on which the Stone is running). If you are running a client application on Windows, for example, you will need to continue to use GsFile to access files on the Windows client. FileSystem is also not supported on AIX.

FileSystem is still under refinement, and may be missing features or contain unexpected behavior. The low level support classes in particular may be refactored and/or have protocol modifications.

### FileReference

FileReference is the primary entry point for file and directory operations. A FileReference represents a file or directory, which may or may not exist on disk.

File references can be created from Strings or using FileReference or FileSystem class methods. For example:

```
'/gshost/test/foo.txt' asFileReference
```

See "Specifying a FileReference" on page 123 for other options.

Using the FileSystem environment, you do not normally "open" a file; instead, you create a read or write stream on a file to perform file operations on the contents. Unlike GsFile, FileReference instances themselves do not understand stream protocol.

FileReference includes many methods to get information about a file or directory, access it for read and write, decompose the filename and path, and file parent and child directories and files. Many of these are inherited from its superclass, AbstractFileReference.

## FileLocator

FileLocator is a sibling class of FileReference, which provides much the same file and directory behavior inherited from AbstractFileReference. It is explicitly designed to allow an environment-independent specification of paths. The actual physical file or directory is resolved according to the environment at runtime. This allows you to move code between environments without requiring explicit management of the paths.

FileLocator provides a number of common environments; these can be listed using

```
FileLocator class >> supportedOrigins
```

Class methods are available for the various supportedOrigins, including home, cache, temp, userData, tranlog, preferences, extent1Directory, extent1, documents, desktop, and workingDirectory.

To find the resolved value of the FileLocator for your current environment, you can send the message absolutePath; this returns an instance of a kind of Path.

```
FileLocator extent1 absolutePath printString  
Path / 'gshost' / 'GemStone3.7' / 'data' / 'extent0.dbf'
```

## FileSystem

While the primary entry point in the FileSystem environment is FileReference, the class FileSystem provides additional behavior.

FileSystem includes support for both ordinary disk based file systems (FsDiskFileSystem) and in-memory file systems (FsMemoryFileSystem). These can be retrieved using class methods. For example:

```
FileSystem disk  
returns a disk-based FileSystem, which can be used with @ or / to create a  
FileReference, e.g. FileSystem disk / '/gshost/test/foo.txt'.
```

```
FileSystem memory  
returns an in memory-based FileSystem, which can be used with @ or / to create a  
FileReference, e.g. FileSystem memory / 'foo' / 'bar'.
```

```
FileSystem workingDirectory  
creates a FileReference to the current working directory (disk-based).
```

FileSystem instances should not be used to perform file and directory operations; these instance methods are subject to change or removal in a future version. Obtaining instances of FileReference or changing the working directory are valid operations for FileSystem instances.

## Specifying a FileReference

There are many ways to specify a FileReferences to a particular file, using methods on FileReference, FileSystem, or by creating a FileReference from a String.

The following are all equivalent:

```

'/gshost/test/foo.txt' asFileReference
FileReference / '/gshost/test/foo.txt'
FileReference / 'gshost' / 'test' / 'foo.txt'
FileReference disk @ '/gshost/test/foo.txt'
FileSystem disk / '/gshost/test/foo.txt'
'/gshost/test/' asFileReference / 'foo.txt'

```

These all resolve to a FileReference that is printed as:

```
FileReference disk @ '/gshost/test/foo.txt'
```

The FileReference and FileSystem / operator and FileSystem @ operator create a new instance of FileReference with its argument string interpreted as a path or file relative to the path of the receiver. If the argument string is an absolute path (that is, includes a leading / in the string), then the new FileReference has that absolute path.

## Working Directory and other environment-independent paths

The working directory is a special case of environment-independent path. You can use working directory as a file location without needing to use using FileLocator.

Using FileSystem >> workingDirectory, you can create a FileReference can be used to access or create a file in the directory that is the current working directory when the code is executed.

For example,

```
FileSystem workingDirectory / 'myLogFile.txt'
```

The workingDirectory can be modified using:

```
FileSystem disk setWorkingDirectory: aFileReference
```

The working directory is also accessible from FileLocator, and other environment-independent root paths can be specified using FileLocator. For example,

```
FileLocator home / 'output.txt'.
FileLocator temp / 'testLogs' / 'performance.log'
```

## Reading from and Writing to a File

To read from or write to a file, open the file for read or write using:

```

aFileReference readStream
aFileReference writeStream

```

These methods return a kind of ZnStream, which understands standard stream protocol.

```

rdStream := '/gshost/test/foo.txt' asFileReference readStream.
[rdStream atEnd] whileFalse:
    [report add: rdStream nextLine; lf.].
rdStream close.

```

These methods have a number of variants. For example, the readStreamIfAbsent: and writeStreamIfPresent: provide easy checking for some common error conditions.

## Closing file streams

Streams that have been opened should be closed when they are no longer needed.

When closing a writestream, you should call `flush` to ensure all data is written to the file.

```
wrStream := '/gshost/test/foo.txt' asFileReference writeStream.
wrStream nextPutAll: SystemRepository fileSizeReport; lf.
wrStream flush; close.
```

The `readStreamDo:` and `writeStreamDo:` variants close the file after the `Do:` block is complete, avoiding the need for an explicit close.

```
'/gshost/test/foo.txt' asFileReference readStreamDo:
[:str |
[str atEnd] whileFalse: [
report add: str nextLine; lf.]
```

## File encodings

The above methods read the disk file with the default encoding, UTF-8. You can write instances of kinds of Strings or Unicode strings with Characters outside the ASCII range to a file, and you read instances of these classes from a file, without having to explicitly convert them to or from the raw file bytes.

FileReference also supports reading files using GemStone's legacy 8-bit encoding. To do this, use the methods that include `Encoded:`, which takes an instance of a kind of `ZnCharacterEncoder`. For example,

```
rdStream := '/gshost/test/foo.txt' asFileReference
readStreamEncoded: '8bit' asZnCharacterEncoder.
[rdStream atEnd] whileFalse:
[report add: rdStream nextLine; lf.].
rdStream close.
```

## Binary Streams

You can also read the files as binary, which allows you to do your own processing of the results. The `binary*Stream` methods return instance of `FsBinaryFileStream`, which reads and writes instances of `ByteArray`.

```
rdStream := '/gshost/test/foo.txt' asFileReference
binaryReadStream.
[rdStream atEnd] whileFalse:
[report add: rdStream contents decodeFromUTF8ToString; lf].
rdStream close.
```

## Operations on Files and Directories

There are a number of options for operating on files and directories; see the image for the full set of methods.

### Create

```
createFile
Create the file; signal an exception if the parent does not exist.
```

`ensureCreateFile`

Create the file, if it does not exist, including parents if needed.

`createDirectory`

Create the directory; signal an exception if the parent does not exist.

`ensureCreateDirectory`

Create the directory, if does not exist, including parents if needed.

## Delete

`delete`

Delete the file or directory. If the file or directory does not exist, or the directory is not empty, signal an exception.

`ensureDelete` - like `delete`, but does not signal an exception if the file or directory does not exist.

`deleteAll`

Delete the file or directory, including all files and directories under the directory. If the file or directory does not exist, signal an exception

`ensureDeleteAll`

Similar to `deleteAll`, but does not signal an error if the file or directory does not exist.

## Move

`moveTo: anotherFileReference`

Moves the file at location for the receiver to the location of the argument.

## Copy

`copyTo: anotherFileReference`

Copies the receiver's file or directory to the location specified by the argument.

`copyAllTo: anotherFileReference`

Copies the receiver's file or directory, and recursively all children of the directory, to the location specified by the argument.

## Rename

`renameTo: newName`

Renames the file or directory

## Information queries

`FileReference` contains many methods providing information about a file or directory. The following methods are available for `FileReference` and `FileLocator`.

### Status of a file or directory

`FileReference` methods allow you to test the status of a file or directory.

Some available testing methods are:

```
exists
isReadable
isWriteable
isFile
isDirectory
isExecutable
```

Operations can also be performed based on file status. For example,

```
ifAbsent: absentBlock
ifExists: existBlock
ifExists: existBlock ifAbsent: absentBlock
ifFile: fileBlock ifDirectory: directoryBlock ifAbsent: absentBlock
```

## Information about a file

FileReference returns information about a file or directory. There are many methods for file attributes; see the image for available methods.

```
exec '/gshost/test/foo.txt' asFileReference modificationTime %
2023-02-27 23:17:32.768

exec '/gshost/test/foo.txt' asFileReference size %
1823
```

## File permissions

File permissions use another class, FileSystemPermission, to represent a file or directory's permissions. FileSystemPermission includes methods to query for read and write permission.

```
exec '/gshost/test/foo.txt' asFileReference permissions %
a FileSystemPermission
posixPermission 420

exec '/gshost/test/foo.txt' asFileReference permissions
printString %
rw-r--r--

exec '/gshost/test/foo.txt' asFileReference permissions
ownerExecute %
false
```

## Decomposition of filename and path

Operations on files often need to filter out one segment of a path or filename. FileReference has a rich set of options to access segments of a path and filename.

```
exec '/gshost/test/foo.txt.gz' asFileReference basename %
foo.txt.gz
```

```
exec '/gshost/test/foo.txt.gz' asFileReference
  pathSegments printString %
anArray( 'gshost', 'test', 'foo.txt.gz')
```

```
exec '/gshost/test/foo.txt.gz' asFileReference
  extensions printString %
anOrderedCollection( 'txt', 'gz')
```

```
exec '/gshost/test/foo.txt.gz' asFileReference
  basenameWithoutExtension: 'gz' %
foo.txt
```

## Listing Directories and Files

FileReference uses the term children for the files and directories that are under the current FileReference.

For example, the directory /gshost/test/ contains two files and one subdirectory. The following reports all children, and only the children that are files.

```
exec '/gshost/test' asFileReference children printString %
anArray( File @ /gshost/test/foo.txt, File @ /gshost/test/logs,
File @ /gshost/test/bar.txt)
```

```
exec '/gshost/test' asFileReference files printString %
anArray( File @ /gshost/test/foo.txt, File @
/gshost/test/bar.txt)
```

Methods such as allChildren, allFiles and allDirectories will recursively return all files and/or directories underneath the receiver.

## Supporting Classes

There are a number of classes supporting FileSystem; many of these are for internal use in supporting FileSystem features. The following are some important subsystems of FileSystem.

### FsFileDescriptor

FsFileDescriptor represents the file itself. This is a lower level support class, however, it is possible to use this directly.

Within the FileSystem environment, you do not normally "open" a file; instead, you create a read or write stream (a kind of ZnStream) on a file to perform file operations on the contents.

FileReference open: methods return an instance of FsFileDescriptor. Instances of FsFileDescriptor read and write ByteArrays, rather than strings.

## Zinc Stream Classes

Subclasses of ZnObject, including ZnBuffered\*Stream, ZnEncoded\*Stream, and Zn\*Encoder, provide read and writestream support for FileSystem. Zn\*Encoder streams handle the encoding/decoding from UTF-8 and 8-bit (GemStone legacy) encoded files, and the creation of Legacy or Unicode String instances.

## Error Classes

FsError and its subclasses represent errors associated with FileSystem. FileSystem exception represent errors related to the FileSystem, FileException and its subclasses represent FileReference/FileLocator errors.

Other errors such as FsUnixError and its subclasses represent low level UNIX file errors, which are generally handled by public API.

## Path classes

Path, RelativePath, and Absolute Path encapsulate a path. A Path can be obtained from a FileReference and vice versa. While Path is an abstract class, you can use it to create instances, e.g. Path from: '/gshost/test/' will return an instance of AbsolutePath.

## Opening options

FsFileOpeningOptions and its subclasses provide the operating-system specific options for opening files.

## Store classes

FileSystemStore and subclasses represent the OS file system or memory file system, with differences for specific operating systems.

## Resolver classes

FileSystemResolver and subclasses support resolving various FileLocator origin options, which vary between operating systems.

## FFI support classes

FsLibcInterface and FsCStruct and their subclasses provide the FFI interface to the operating system file commands.

# Indexes and Querying

---

This chapter describes GemStone Smalltalk's indexing and querying mechanism, a system for efficiently retrieving elements of large collections.

**Overview** (page 130)

Reviews the concept of relations.

**Defining Queries** (page 137)

Describes the structure of query predicates, the types of queries, and how to construct a query.

**Creating Indexes** (page 139)

Discusses GemStone Smalltalk's facilities for creating indexes on collections.

**Results of Executing a GsQuery** (page 146)

How to execute a query and the options for working with the results.

**Enumerated and Set-valued Indexes** (page 151)

Describes how to create enumerated and collection-valued indexes and queries.

**Managing Indexes** (page 152)

How to perform index management: find out about indexes in your system, remove existing indexes, handle errors, and audit indexes.

**Indexing and Performance** (page 157)

Additional factors that can impact the performance of your queries.

**Historic Indexing API differences** (page 158)

The older indexing API, using `UnorderedCollection` methods and `select` blocks.

## 8.1 Overview

Most applications use one or more databases containing business data, which may be very large. Individual records in these databases may be added, removed, and/or updated, and need to be queried in multiple ways for different purposes. All these operations must be performed quickly and efficiently.

### Business Objects

In GemStone, a database is represented as an instance of a collection that holds instances of business objects. You may have thousands or millions of objects in a collection, and these objects may be complex composite objects holding many individual strings, dates, number and other basic data types.

The following example shows simple employee data in table form:

**Table 8.1 Employees**

First Name	Job	Age	Address
Fred	clerk	40	22313 Main, Dexter, OR
Sophie	bus driver	24	540 E. Sixth, Renton, WA
Conan	librarian	40	999 Walnut, Hilt, CA
Moppet	intern	18	17 SW Oak #6, Portland, OR

In Smalltalk, this can be represented as an Employee class, with instance variables `firstName`, `job`, `age`, and `address`; and an Address class, with `street`, `city`, and `state` instance variables.

### Database Collection

The collection itself may be an instance of a number of different types of Collection subclasses. For scaling, and to support indexes, a subclass of `UnorderedCollection` is recommended. Hashed collections such as dictionaries may become unbalanced if too many elements hash to the same value, and as a collection grows, may require the entire collection to be rebuilt. Indexed collections such as `Array` have limitations on adding and removing elements without affecting the entire collection. `UnorderedCollections`, particularly `IdentityBag`, `IdentitySet`, `RcIdentityBag`, and `RcIdentitySet`, use an optimized internal tree structure to hold the elements and are the recommended Collection classes for use for large databases. Collection classes are described in Chapter 4.

To make it easy to associate behavior with your set of Employees, it is often useful to define a class `SetOfEmployees` that is a subclass of `IdentitySet`. An instance of `SetOfEmployees` then can contain instances of `Employee`, with a reference from `UserGlobals` or from a class variable.

## Queries

Since `UnorderedCollections` aren't ordered, lookup is by value. For example, to find a particular `Employee`, you use `select:`, `detect:` or similar messages. For example,

```
MyEmployees select: [:ea | ea address state = 'OR']
MyEmployees detect: [:ea | ea firstName = 'Sophie']
```

These iterative messages may not scale well. For example, for the above `select:` expression, for each employee in the collection, the employee object and the address object must be faulted into memory, and the messages `address`, `state`, and `=` are sent. While this doesn't matter for small collections, it can become unreasonably slow for very large collections; particularly if objects in the collection are not in the shared page cache, and need to be read from disk.

## GemStone Indexes and Queries

### Indexes

Indexes and indexed queries provide a way to locate specific objects in a collection by value. Indexes are created on specific named instance variables, either by identity or by equality. Creating an index on a collection (e.g. on the instance variable `firstName`), creates parallel internal structures which provide a mapping from the indexed value (such as the `firstName` 'Sophie') to the root object in the collection (the employee). Using this index, only a few message sends are needed to lookup the collection element that is the same as, or less or greater than, a particular value.

**Identity indexes** support queries that are looking for identical values, while **equality indexes** support queries that compare using equality, or greater or less than, a particular value.

Indexes are created on objects based on instance variables, not on message sends; since the instance variable relationships are known by the system, indexes can be updated automatically as elements are added and removed from the collection, and when references on the path are changed. There are some exceptions to this which require manually updating the indexes.

Indexes may only be created for instance of subclasses of `UnorderedCollection`.

### GsQueries

To take advantage of an index you have built on your collection, you must perform the query using `GsQuery` syntax, rather than `select:` or similar iteration methods. A query performed using `GsQuery` will use indexes, as long as an index exists for the particular instance variable involved in the query. If an index does not exist, then the `GsQuery` will be performed iteratively, with performance similar to the comparable `select:` or `detect:` operation.

When the collection is properly indexed, `GsQueries` can return results without having to iterate the collection, fault the intermediate objects into memory, or send messages to each object.

`GsQueries` can be used on most kinds of `Collection`, not only `UnorderedCollection`. However, the performance benefit only appears on instances of subclasses of `UnorderedCollection` for which the appropriate index or indexes exist.

## Deciding what to optimize

As with any kind of optimization, it's important to consider the application's performance profile, performance requirements, and the entire context, rather than automatically creating indexes on all possible paths.

The process of creating indexes creates overhead. The additional internal objects created use some space, and building an index may take some time. As the data in the repository changes, including objects added to and removed from the collection itself as well as changes in actual values, the mappings in the index structures need to be updated. Periodically, indexes should be audited to ensure integrity, and rebuilt if necessary; rebuilds are required for some system upgrades. Indexes must be specifically removed when the collection is removed, to ensure the internal infrastructure is cleaned up.

While most collections with more than a few thousand objects will see better performance using indexed queries, it is wise to consider indexes with this overhead in mind. Before going through the trouble of creating an index, you should determine that the index provides value. There are a number of factors that strongly influence queries, both iterative queries and indexed queries. These factors interact with each other and there are other factors, such as caching, that also influence performance.

- ▶ The size of the collection. With smaller collections, iterative performance is fast enough that indexing provides little benefit. Iterative performance grows linearly with collection size, while indexed performance increases slowly.
- ▶ The length of the path. Longer paths require more lookups and more infrastructure, and take longer to complete. For longer paths, it is more efficient to cache the value higher within the object structure.
- ▶ The size of the result set. If you have a query that returns a very large number of results, creating the result set reduces performance; this is particularly so for indexed queries.

## Overview of the steps in creating and using indexed queries

In order to take advantage of efficient indexed queries on your collection, the following steps need to be done:

- a. Determine the queries that can benefit from optimization, and describe them using query syntax. Query syntax is described starting on page 133.

For example, to query for employees under 21 who live in Oregon, the query string might be:

```
(each.age < 21) & (each.address.state = 'OR')
```

2. Create one or more indexes on the collection, that specify the particular instance variable path on which you will perform the query. Creating indexes is described starting on page 140.

To support the above query, you may want to create two indexes, for example:

```
GsIndexSpec new
  equalityIndex: 'each.age' lastElementClass: SmallInteger;
  equalityIndex: 'each.address.state' lastElementClass:
    String;
  createIndexesOn: myEmployees.
```

- Execute the query on that indexed collection, using query protocol. How to define and execute queries is described starting on page 138.

For example:

```
(GsQuery fromString: '(each.age < 21) & (each.address.state =
    'OR')')
on: myEmployees;
queryResult
```

## Managing Indexes

In addition to creating indexes and queries, you will also need to do some management on your indexes and queries. For example, you should evaluate your indexes for performance, remove indexes that are no longer needed, and audit indexes to ensure the structures are correct. Many of these indexing tasks are handled by `IndexManager`.

## Special Syntax for Indexing

GemStone indexing uses several syntactical elements that are either specific to, or primarily used for, index creation and indexed queries.

### Path-dot syntax

Indexes are created, and queries formed, using special syntactic structure called a *path*, which designates variables for indexing and describes certain features of the index. Path syntax uses a period to represent the object/instance variable name relationship.

For example, given a collection of `Employees`, in which each employee has an `address` instance variable, which refers to an `Address` that has a `state` instance variable, the path is:

```
address.state
```

A longer path is

```
account.order.address.state
```

In the simplest case, a path on an instance variable on the collection elements, this is just the instance variable name. For example:

```
firstName
```

You may also specify an empty path, meaning the elements of the root collection itself.

Each instance variable name on the path is a *pathTerm*. In the above example, `address` and `state` are each *pathTerms*. Paths can contain a long string of *pathTerms*, if the elements of the collection represent a deeply nested tree of objects.

Path-dot syntax can be used anywhere in GemStone code; it is required in index creation and queries, for which message sends are not allowed.

### Initial each

An initial `'each.'`, where `each` represents the elements of the collection, is recommended but optional for `GsIndexSpec` index creation, and required for `GsQueries`. For example:

```
each.address.state
```

### Enumerated pathTerms

A vertical bar `|` in the path indicates the presence of two alternate instance variables that will be indexed together, as if they were a single variable.

For example, you might want to search on both name and nickname in a single operation. This might look like this:

```
account.name|nickname
```

### Set-value path terms

An asterisk \* in the path indicates a collection, which must be an instance of an indexable class (an instance of a subclass of `UnorderedCollection`). A set-valued path term may not be the first term in the path.

For example, if the instance variable `children` contains an `IdentityBag` of instances of `Child`, and a child has the instance variable `age`:

```
children.*.age
```

### Historic indexing syntax

The `GsIndexSpec`/`GsQuery` classes provide the general purpose indexing interface. An older syntax using `UnorderedCollection` methods to create indexes, and selection blocks with curly braces to define queries, is an alternate way to use indexes. This older syntax remains fully supported in order to ensure upgraded applications do not require changes. However, new features are not available using this historic API.

See section 8.6 on page 152 for information specific to the historic API.

## Last Element Class

Creating an equality index creates an internal btree that contains the ordered values of the instance variable that is indexed. For example, an index on `firstName` creates a btree containing 'Conan', 'Fred', and so on. This allows fast lookup of a position in this btree when performing the query, and values that are equal or greater or less than can be returned in order as needed.

Building this btree and providing predictable lookup requires that the values be comparable in well-known and efficient ways. When building indexes, there are choice to make in balancing the restrictivity of the indexed values vs. the impact of comparison on query performance.

Performing an identity query creates no such restrictions on the index, since the comparison is by identity (OOP), and any two objects can be compared this way.

To provide the definition of comparison, equality indexes require specifying the `lastElementClass`. This generally restricts the indexed values to instance of this class or of subclasses of this class, although string classes have some special handling.

### Optimized classes

The following classes, and subclasses of these classes, are optimized for indexes. In most cases, the final element you will create an index on will be one of the following. For legacy indexes, the index structures encode the value; for `btreePlusIndexes`, they can perform optimized comparisons. These classes are subclasses of `Magnitude` or `CharacterCollection`.

```
Character, SmallInteger, SmallDouble, SmallFraction,  
String, DoubleByteString, QuadByteString,  
Unicode7, Unicode16, Unicode32,  
Symbol, DoubleByteSymbol, QuadByteSymbol,
```

Time, Date, DateTime, DateAndTime,  
LargeInteger, Float, DecimalFloat, ScaledDecimal, FixedPoint, Fraction

Boolean is a special case; it is a special, and so does not require looking in legacy indexes. However, it does not support optimizedComparison.

## Using other classes

You can create indexes where the indexed values are instances of classes other than the above, including classes you have defined yourself.

Identity indexes on instances of your own classes require no extra work, since they compare on the identity of the objects.

If you wish to create an index where the values that are instance of application classes that do not subclasses of basic classes, you must ensure these classes implement comparison operators, as described on page 136.

## Comparing data types

Some cases of data type comparison have special handling in indexes.

- ▶ It may be useful to mix strings and symbols, but there is additional cost. While a string and a symbol can be ordered using `<=`, a string and a symbol that contain the same characters are not equal. There are two solutions: using alternate comparison methods which reduce performance; or optimizing the comparison operators and not mixing symbols and strings.
- ▶ NaN (not a number) are specialized kinds of Float that are not equal to themselves. As with strings, special handling is required to accommodate NaNs, at the cost of performance; or NaNs may be disallowed in Float indexes.
- ▶ The indexed comparison mechanism considers only the first 900 characters of each string operand, so two strings that differ only beginning at the 901st character are considered equal.
- ▶ nil is a special case of object that can be compared to any other object. They also require special handling in indexes. Since the appearance of nil signifies a value that is not there, less than and greater than comparison results will not include nil values. Since accommodating nil requires special protocol, nil may also be disallowed.

A nil along the path to an indexed slot is a different issue; such missing sections of a reference tree are allowed without special handling.

## Strings in indexes

Indexing on strings has complications, due to the different collation orders it is possible to configure. For more on collation, see Chapter 5.

To summarize, strings come in two "flavors":

- ▶ Traditional strings (String, DoubleByteString and QuadByteString, which are interchangeable based on the maximum Character codePoint size). Traditional strings, in Legacy String Comparison Mode, use character-based collation.

Symbols (Symbol, DoubleByteSymbol and QuadByteSymbol) follow the same collation rules as Traditional strings.

- ▶ Unicode strings (Unicode7, Unicode16, and Unicode32) always use ICU string-based collation.

A repository in Legacy String Comparison Mode disallows compare between Unicode strings and Traditional strings or symbols, to avoid unpredictable results. In this mode, you cannot mix Traditional and Unicode strings; it is difficult to avoid errors when using Unicode strings in Legacy String Comparison Mode.

A repository in Unicode Comparison Mode uses Unicode collation for all flavors of strings and symbols. In this mode, you can use Traditional strings and Unicode strings interchangeably.

Constraining the indexed variables using `lastElementClass` is not effective for strings, since Traditional string, symbol and Unicode string classes inherit by codePoint range rather than by collation or other behavior. It is allowed, but not recommended, to specify `CharacterCollection` (the superclass of all kinds of Strings and Symbols), since (depending on the mode and index type) it may create an ambiguous indexes.

In both Comparison Modes, specifying a `lastElementClass` of any of the following will create an index that includes a cached collator:

Unicode7, Unicode16, Unicode32

In Legacy String Comparison Mode, the `lastElementClass` of any of the following will permit instance of any of the classes:

String, DoubleByteString, QuadByteString,  
Symbol, DoubleByteSymbol, QuadByteSymbol

In Unicode Comparison Mode, the `lastElementClass` of any of the following will permit instance of any of the classes:

String, DoubleByteString QuadByteString,  
Symbol, DoubleByteSymbol, QuadByteSymbol  
Unicode7, Unicode16, Unicode32

Note that some optimized indexes disallow mixing Symbols with any kinds of Strings.

## Redefining Comparison Messages

If you create an index on values that are instances of your application classes, these classes must implement the basic comparison operators, at least `=`, `>`, `<`, and `<=`. You can redefine one or more of these in terms of another.

The operators must be defined to conform to the following rules:

- ▶ If  $a < b$  and  $b < c$ , then  $a < c$ .
- ▶ Exactly one of these is true:  $a < b$ , or  $b < a$ , or  $a = b$ .
- ▶  $a <= b$  if  $a < b$  or  $a = b$ .
- ▶ If  $a = b$ , then  $b = a$ .
- ▶ If  $a < b$ , then  $b > a$ .
- ▶ If  $a >= b$ , then  $b <= a$ .

While the indexing subsystem does not use hashing itself, note that redefining `=` does requires attention to the `hash` method to be consistent with the new definition of equality. Object that are equal must return the same hash value to ensure they behave in a consistent and logical manner in all use cases.

## 8.2 Defining Queries

Before you can define indexes on your collection, you need to determine the ways in which you will need to search your collection to retrieve elements. The queries you need determine the details of the indexes to create.

At its simplest, a query consists of the specification of an instance variable common to all the objects in the collection, a comparison operator, and a literal to which the value is compared. For example, if you wish to be able to find all employees 21 and older, your query formula could be something like this:

```
each.age >= 21
```

In this example, every object in the collection (*each*) has an instance variable *age*, which is specified using dot-path notation. The value of that instance variable is compared, greater than or equal, to the literal `SmallInteger 21`.

While this formula is simple, you can formulate queries based on multiple instance variable values, operators, and constants, and combine them using boolean logic. However, using this query syntax, you cannot include message sends; the indexes are based on structural relationships using instance variable names.

For performance and clarity, it is an advantage to use short and simple queries. However, it may be valuable to compose your queries based on the statement of business logic. This may mean creating a complicated query that is not in its most efficient form. The final query will be automatically optimized to a logically equivalent form that is more efficient for GemStone to execute. See “Formulating queries and performance” on page 158.

### Query Predicate Syntax

A query contains a *predicate* expression, which is a Boolean expression that, when evaluated with the elements of the collection, returns true or false. In a query, the expression usually compares an instance variable on the collection objects with another instance variable or with a constant.

A predicate contains one or more *predicate terms* – the expressions that specify comparisons.

#### Predicate Terms

A *term* is a Boolean expression containing an operand and usually a comparison operator followed by another operand. For example, in

```
each.age >= 18
```

`each.age` and `18` are operands, while `>=` is a comparison operator. The only time you would not have a comparison operator is if the operand is itself a Boolean (true or false).

#### Predicate Operands

An operand can be a path (*each.age*, in this case), a variable name, or a literal (18, in this example). All GemStone Smalltalk literals except arrays are acceptable as operands.

#### Predicate Operators

Predicate operators are `==`, `~~`, `=`, `~=`, `<`, `<=`, `>` and `>=`. No other operators are permitted in a GsQuery or selection block query.

## Combining Predicates using Boolean Logic

If you want retrieval of an element to be contingent on the values of two or more of its instance variables, you can join several terms using a conjunction operator `&` (logical AND) or disjunction operator `|` (logical OR).

The conjunction operator, `&`, makes the predicate true if and only if the terms it connects are true. The disjunction operator, `|`, makes the predicate true if either one, or both, of the terms it connects are true.

You may also negate individual predicate terms using `not`.

Each predicate term must be parenthesized.

For example, the following are legal queries.

```
(each.name = 'Conan') & (each.job = 'librarian')
(each.age <= 40) | (each.job = 'librarian') not
```

## Combining Range Predicates

Queries that use less than or greater than, such as `each.age >= 18`, define a starting (or ending) point in a range query. Specifying both a starting point and ending point creates a range query. For example,

```
(18 <= each.age) & (each.age <= 65)
```

These two terms can be combined into single range predicate.

```
18 <= each.age <= 65
```

Range specifications such this can only be defined with this syntax if the operands and comparison operators truly define a range.

## Creating a GsQuery

GsQuery is a programmatic way to define a query, allowing you to easily abstract, store and reuse various aspects of the query.

To create a GsQuery, you create an instance of GsQuery using query predicate syntax. The most simple way to create a GsQuery is by passing in a string. For example:

```
GsQuery fromString: 'each.age >= 18'
```

Since the `fromString:` protocol requires a string, if the query includes literal strings, you must include two single quotes within the string. For example:

```
GsQuery fromString: 'each.firstName = ''Fred'''
```

This message will return an instance of GsQuery. Before it can be executed, it must be bound to a collection:

- ▶ Create the GsQuery using `fromString:on:`: creates a GsQuery that is bound to a particular collection.
- ▶ Bind the query before executing using the `on:` method.

## Query Variables

The strings used to define GsQuery instances may contain variables—any element of a predicate that is are not a literal or path-dot expressions. This allows your query to be stored and executed later using different values.

For example, for a query such as

```
GsQuery fromString: '18 <= each.age <= 65'
```

This can be generalized to a query with variables:

```
GsQuery fromString: 'min <= each.age <= max'.
```

The resulting formula in the GsQuery includes 'min' and 'max' as variables. These must be bound to specific values before the query can be executed. Binding is done by sending the `bind:to:` message to the query. For the above example, to execute the query:

```
aQuery := GsQuery fromString: 'min <= each.age <= max'.
aQuery
  bind: 'min' to: 18;
  bind: 'max' to: 65;
  on: myEmployees;
  queryResult
```

Note that the “max” and “min” in the query formula are string elements, and are not affected by any temporary or instance variables named max or min in the scope of the code being executed. The only way to resolve max and min are by binding variables.

## 8.3 Creating Indexes

Queries can be executed without an associated index, but there is no performance benefit. To execute a query efficiently, you need to also create an index on the instance variables for the query. These indexes provide a mapping from the specific key values that you are interested in to the results (the objects in the collection).

The path you provide when creating an index provides the key that is needed to lookup the value during a query. These keys are the values of a specific instance variables within the elements of a collection, or the elements of the collection itself. For example, given a collection of Employees, and the path `each.address.state`, the objects at the `state` instance variable (perhaps two-character Strings) would be the keys.

The values for these keys are the objects in the collection itself, which are the results of the query using that index. For our example, the values are the instances of Employee in AllEmployees. When you make an indexed query for Employees with addresses in a given state, that `state` key is used to lookup the matching elements (instance of Employee).

### Equality and Identity Indexes

Indexes fall into two main types: **Equality Indexes** and **Identity Indexes**. Equality indexes support equality-based queries, including `>`, `>=`, `<`, `<=`, `=`, and `~=`. Identity indexes support queries containing identity comparisons, `==` and `~~`.

When creating an index, you specify whether an equality or identity index is created. Since identity comparisons are done by OOP, not by the object's contents, they are faster, and the `lastElementClass` does not matter; any two objects can be compared for identity.

If you only have an identity index on a variable, but form your query using an equality operator, the query will not have an index to use (and thus, will iterate the collection).

You may create both equality and identity indexes on the same path.

## Btree and Legacy Indexes

GemStone supports two different internal structures; the legacy structures, which includes a btree and an index dictionary; and the btreePlus structures, which use a btree+ and does not require the dictionary. The query results are the same for each, of course, but the performance profile is different.

The decision of which to use impacts your indexing work.

- ▶ The best query performance is with btreePlusIndexes with optimizedComparison. However, optimizedComparison places restrictions on lastElementClass data types, such that, for example, Strings and Symbols cannot be mixed, and nils and NaN floats may not be present.
- ▶ If your data does not conform to the data type restrictions, using legacy indexes is recommended.

With a legacy identity index, the index dictionary provides a identity-based lookup for the key. In a btreePlus identity index, the keys are in a btree. This allows you to stream over the results of a identity query only when using a btreePlus index.

The index structure you use can be specified for each index, otherwise it relies on the system or configured default. Since structures are shared between indexes on a collection, all indexes on a specific collection must use the same internal structure.

Note this is entirely distinct from the historic indexing API (using UnorderedCollection methods to create indexes); creating indexes using the historic API may create either kind of internal structure, depending on the current default.

See page 143 for details on how to configure each index type.

## Creating the Index

Creating an index involves creating an instance of GsIndexSpec and sending messages to define the index and the parameters and options for that index, then use this spec to create indexes on a specific collection.

Before creating an index, you must know:

- ▶ the paths for the instance variables that you will query on.
- ▶ The classes of the values of these instance variables, and if these instances are homogenous.
- ▶ If your queries will be by equality or identity

To create an index using GsIndexSpec, do the following:

### 1. Create the instance of GsIndexSpec

This is done by executing `GsIndexSpec new`

### 2. Define one or more indexes on the spec

To define an index, send an index creation message to the GsIndexSpec, including the path you want indexed, the class of the last element (for equality indexes), and options (if used).

The most general index creation methods include:

```
equalityIndex: lastElementClass:  
identityIndex:
```

While these methods can be used to create indexes on strings, there are additional index creation methods are specific to various kinds of string indexes. These methods have variants that allow you to specify the index options.

### 3. Create the index on a specific collection

To actually create the index, send the message `createIndexesOn:`, providing the specific collection on which you want to create the indexes.

To put this all together, for example:

```
GsIndexSpec new  
  identityIndex: 'each.userId';  
  equalityIndex: 'each.age' lastElementClass: SmallInteger;  
  equalityIndex: 'each.address.state' lastElementClass: String;  
  createIndexesOn: myEmployees.
```

This creates an identity index on `userId`, an equality index on `age`, and another equality index on `address.state`, all on the collection `myEmployees`.

You can view the indexes by recreating the specification from the indexed collection, using `indexSpec`. For example:

```
run  
myEmployees indexSpec printString  
%  
GsIndexSpec new  
  identityIndex: 'each.userId';  
  equalityIndex: 'each.age'  
    lastElementClass: SmallInteger;  
  equalityIndex: 'each.address.state'  
    lastElementClass: String;  
  yourself.
```

## Equality Indexes on strings

Equality indexes on strings present a variety of options and restrictions, depending on:

- ▶ If the indexed elements will be Traditional strings, Unicode strings, Symbols, or a mix.
- ▶ If you are using the `GsIndexOptions optimizedComparison` feature, which is strongly recommended with `btreePlus` indexes and disallowed with legacy indexes.
- ▶ If the application is in Unicode or Legacy String Comparison Mode.

The following methods can be used to create equality indexes on strings and/or symbols. Note that each has a variants that allow you to specify the index options.

```
equalityIndex:lastElementClass:  
unicodeIndex:  
unicodeIndex:collator:  
stringOptimizedIndex:  
symbolOptimizedIndex:  
symbolOptimizedIndex:collator:  
unicodeStringOptimizedIndex:  
unicodeStringOptimizedIndex:collator:
```

Which one you should use, and the rules allowing comparisons between different kinds of data, are different for repositories in Legacy String Comparison Mode or in Unicode Comparison Mode.

Comparison Modes are described on on page 82.

## Repositories in Legacy String Comparison mode

In Legacy String Comparison mode, it is disallowed to compare Traditional and Unicode strings, so it's not possible for the indexed variables to contain a mix of Unicode strings and Traditional strings or Symbols.

### Legacy indexes

To create a legacy index on Traditional strings, symbols, or a mix of the two, use a `equalityIndex:*` method specifying a `lastElementClass` of String.

If you are using Unicode strings in Legacy String Comparison Mode, use a `unicodeIndex:*` method.

### optimizedComparison (btreePlus) index

You cannot create an `optimizedComparison` index on a mix of types.

If your indexed elements are all Traditional strings, use a `stringOptimizedIndex:*` method.

If your indexed elements are all Unicode strings, use a `unicodeStringOptimizedIndex:*` method.

If your indexed elements are all Symbols, use a `symbolOptimizedIndex:*` method.

## Repositories in Unicode Comparison Mode

In Unicode Comparison Mode, Traditional strings are collated exactly like Unicode strings, and indexes make no distinction between them.

Symbols are also collated like Unicode strings, but due to the definition of equality, `optimizedComparison` indexes do make a distinction between strings and symbols.

### Legacy indexes

To create a legacy index in Unicode Comparison Mode on Traditional strings, Unicode strings, symbols, or any mix, use a `unicodeIndex:*` method, to ensure the collator is persisted with the index.

### optimizedComparison (btreePlus) index

optimizedComparison indexes may mix Traditional and Unicode strings, but may not mix strings and symbols.

If your indexed elements are all Traditional or Unicode strings, use the method `unicodeStringOptimizedIndex:*`.

If your indexed elements are all Symbols, use the method `symbolOptimizedIndex:*`.

## Implicit Indexes

With legacy indexes, the indexing internal structures include a dictionary. This dictionary, as a side effect, provides de facto identity indexes with some equality indexes: specifically, for non-terminal pathTerms, and where the lastElementClass is a Special (such as Character, SmallInteger, etc; see “Special” on page 35 for specific classes), in which equality and identity are the same). Such indexes are referred to as implicit indexes.

Since with btreePlusIndexes there is no dictionary, there are also no implicit indexes defined.

For clarity, and to avoid dependency on side-effects of the internal structures, it is recommended to explicitly define any identity indexes that you require. There is no risk in explicitly creating an identity index that would exist as a implicit index.

## GsIndexOptions

An instance of GsIndexOptions specifies features that will be used when creating a particular index on a collection. GsIndexSpec index definition methods all have variants that accept an instance of GsIndexOptions, although some override certain settings. If no GsIndexOptions is explicitly provided, the session or repository default is used.

The GsIndexOptions defines if the index is a legacy index or a btreePlus index, as well as other important indexing features. The options available for GsIndexOptions are:

`GsIndexOptions class >> legacyIndex`  
defines a legacy index structure, and disables btreePlusIndex and optimizedComparison.

`GsIndexOptions class >> btreePlusIndex`  
defines a btreePlus index structure, and disables legacyIndex.

`GsIndexOptions class >> optimizedComparison`  
adding optimizedComparison is only allowed with btreePlusIndex.

`GsIndexOptions class >> reducedConflict`  
Instructs the index to create the internal structures as reduced-conflict, recommended when indexing on a reduced-conflict collection.

`GsIndexOptions class >> optionalPathTerms`  
Instructs the index to allow objects that do not include an indexed instance variables to be present in the indexed collection.

These options are described in more detail starting on page 144.

## Combining options

GsIndexOptions can be combined using the plus operator and removed using the minus or not operators, with the caveat that not all options are compatible with each other. For example:

```
GsIndexOptions legacyIndex + GsIndexOptions reducedConflict
GsIndexOptions btreePlusIndex + GsIndexOptions
    optimizedComparison not
```

If you combine two options that conflict, the later one has precedence.

## Default options

Creating an instance of GsIndexOptions, using class methods such as GsIndexOptions >> legacyIndex, begins with the default, repository-wide GsIndexOptions.

The specific value requested by the class method (such as legacyIndex) overwrites the default only for that setting and its dependents.

For example, using GsIndexOptions legacyIndex will return a GsIndexOptions instance with legacyIndexes on and both btreePlusIndex and optimizedComparison disabled, regardless of the default. However, the default GsIndexOptions setting for other values, such as reducedConflict, will be retained.

The initial default GsIndexOptions is:

```
GsIndexOptions btreePlusIndex + GsIndexOptions optimizedComparison.
```

In an upgraded application, the system default is set instead to:

```
GsIndexOptions legacyIndex
```

to ensure that the behavior does not change from previous releases.

You can manually set the repository-wide default, as SystemUser, by executing GsIndexOptions class >> default:. Do this with care, since it may affect all indexes that are created in the future that do not explicitly set all the GsIndexOptions values.

For example, if you have an upgraded application and want to default to btreePlusIndexes and optimizedComparison, execute

```
GsIndexOptions default: (GsIndexOptions legacyIndex +
    GsIndexOptions reducedConflict)
```

You may also set a session-wide default that applies only to your session and only until you log out, using GsIndexOptions class >> sessionDefault:.

## The Options in GsIndexOptions

The options btreePlusIndex, optimizedComparison, and legacyIndex are used to specify the index type.

- ▶ GsIndexOptions legacyIndex enables the classic legacy btree and disables btreePlusIndex. legacyIndex is not compatible with optimizedComparison.
- ▶ GsIndexOptions btreePlusIndex enables the btreePlus structures and disables legacyIndex. For performance, this is normally used with the optimizedComparison option. btreePlusIndexes without optimizedComparison are somewhat less performant than legacy indexes in most cases.

The following table describes the three combinations:

<code>GsIndexOptions btreePlusIndex + GsIndexOptions optimized comparison</code>	Provides the best query performance, with somewhat slower update performance. There are restrictions on the contents of indexed instance variables; nil is not allowed, they cannot mix strings and symbols, and cannot mix floats and NaNs.
<code>GsIndexOptions legacyIndex</code>	Provides good performance. Data type restrictions are less strict.
<code>GsIndexOptions btreePlusIndex</code>	Data type restrictions are less strict, but the performance is not as good as legacyIndex.

Using `optimizedComparison`, it is disallowed to use a mix of certain kinds of objects in the collection. The following rules when using `optimizedComparison`:

- ▶ values must be a kind of the last element class.
- ▶ nil is not allowed as a value.
- ▶ For Float last element class, NaN floats are not allowed as a value.
- ▶ For String last element class, Symbols are not allowed as a value.
- ▶ For Symbol last element class, Strings are not allowed as a value.

When using the "Optimized" index specification methods to define an index, it overrides the settings for these three options in the default or argument `GsIndexOptions`.

## Reduced-Conflict

In a multi-user system, reduced-conflict collection classes may help avoid transaction conflicts if multiple users simultaneously add or remove objects from the collection; for more on this problem, see "Classes That Reduce the Chance of Conflict" on page 185. For example, using an `RcIdentityBag` rather than an `IdentityBag` allows concurrent updates to the collection itself.

If there are concurrent updates of the same indexed instance variable for different objects in the collection (for example, the addresses associated with two different customer objects are both changed), there is not an application object conflict, since the objects are independent. However, there may be a transaction conflict due to the indexes, since both addresses are keys in the same indexing structure.

This doesn't apply to legacy identity indexes, which are always reduced-conflict.

To avoid transaction conflicts from the indexing internal structures, specify that the indexes are `reducedConflict`, using `GsIndexOptions reducedConflict`.

For example:

```
GsIndexSpec new
  equalityIndex: 'each.address'
  options: (GsIndexOptions reducedConflict)
```

## Optional pathTerms

A homogenous collection is one in which each element in the indexed collection defines the instance variable described by the index, for each `pathTerm` in the indexed path. By default, indexes require that the collection be homogeneous. If any element does not have

the given instance variable, it will raise an error when the element is added to the collection.

If you want to create an index on a non-homogenous collection, you can define the indexes with optional pathTerms. For example:

```
GsIndexSpec new
  equalityIndex: 'each.nickName'
  options: (GsIndexOptions optionalPathTerms)
```

When creating an optional pathTerm index, it is not an error when the objects in the collection do not implement an instance variable specified by the index. For a multi-pathTerm index, that includes each pathTerm; objects with missing instance variable definitions for any of the pathTerms in the indexed path are not considered when creating query results.

Note that this option bypasses some error detection. If you create an index using an instance variable that does not exist at all (perhaps due to a typing error), then the index is created correctly and does not report an error, even if it does not create the index you might have intended to create.

## 8.4 Results of Executing a GsQuery

Once you have defined your query, created the GsQuery, and bound it to a collection, there are further options in how to access the results of the query.

To simply get the results, you can send `queryResult` to the instance of GsQuery.

`GsQuery >> queryResult` will, like selection block queries, return a new instance of collection of the same class as the base collection, unless protocol such as `asArray` are used to specify the class of the results.

Also similarly to selection block queries, queries on instances of reduced-conflict (Rc) collections, return the equivalent non-Rc collection.

The collection returned from a query has no index structures. Indexes belong to specific instances of collections, rather than the classes. If you want to perform indexed selections on the new collection, you must build the necessary indexes on the new collection.

### GsQuery's Collection protocol

GsQuery accepts other Collection protocol, and, provided the query has bound to a collection and to query variables, the GsQuery instance responds to as if the GsQuery was a collection of the results of the query. This means that rather than having to put the results of a query into a temporary variable for further processing, GsQuery can respond directly to the kinds of message you are likely to send to the query results.

You can convert the type of collection, for example, using `asArray` or `asIdentityBag`:

```
(GsQuery fromString: 'each.address.state = ''OR''
  on: Employees) asArray
```

Or fetch a single instance from the results:

```
(GsQuery fromString: 'each.firstName = ''Sophie''
  on: Employees) any
```

Performing one of the collection operations that are provided for GsQuery simplifies your code, since you may not have to put results in temporary variables. It may or may not allow you to avoid creating query result objects.

Enumeration methods also allows you to perform code while the query is executing, rather than waiting for the results.

### Caching Query Results

While GsQuery responds to messages as if it was a collection, the results of a query are not a static collection. By default, each time you execute any GsQuery collection protocol, the query is performed again. So, for example, sending `isEmpty` to a GsQuery before sending `asArray` will execute the underlying query twice.

You can cache the results of your GsQuery using GsQueryOptions `cacheQueryResult`. By default, it is false. Using this option allows the `resultSet` of the GsQuery to be cached. Note that this cache will **not** reflect changes in the root collection that occurred after the query was executed; you are responsible for re-running the query if current results are required.

To create an instance of GsQueryOptions with `cacheQueryResult` true, use this expression:

```
GsQueryOptions cacheQueryResult
```

And use this instance with GsQuery methods that includes the `options:` keyword.

For example:

```
query := (GsQuery fromString: 'each.address.state = 'OR''
         options: (GsQueryOptions cacheQueryResult)
         on: Employees).
query isEmpty ifTrue: [^'no results'].
report := self createReportingStructure.
query do: [:ea | report updateDataWith: ea].
...
```

### GsQuery enumeration methods accepting blocks

Among the collection protocol that GsQuery understands are the methods `do:`, `select:`, `reject:`, `collect:`, `detect:` and `detect:ifNone:`. These may look similar to iterative queries on the root collection, but since the actual query is already provided by the GsQuery, the action is quite different.

With GsQuery, these will operate on the *result set* of the initial query. In essence, you are adding an additional, non-indexed search criteria to the indexed query. This additional code will be executed for each element in the collection for which the indexed query matches, at the time that the index query is examining that result element.

For example, if you have an index on Employee age, and a query such as:

```
(GsQuery fromString: 'each.age <= 18' on: Employees)
```

Using this query, you can add an additional search criteria using `select:`, so that only Employees who live in Oregon are returned.

```
(GsQuery fromString: 'each.age <= 18' on: Employees) select:
[:each | each address state = 'OR']
```

This will return a result set that includes Employees under 18 who live in Oregon.

The `address` message is only sent to the elements (Employees) who are under 18, it is not executed for every element in the collection. Also note that the state comparison does not use an index; these are message sends.

### Order of results

Provided there is an index on the query path, the enumeration block operates on each object in the result set in the order specified by the index. However, if you wish to use the result of the `select:` or other enumeration method, the result will necessarily be a kind of `UnorderedCollection`, and the objects in the returned collection will be not be ordered.

You can still use the enumeration protocol to produce results that are ordered according to the index, by adding each element to a temporary Array. However, for ordered results, you may want to stream over the results instead.

### Efficiency of query vs. enumeration

It is more efficient to perform an indexed query with multiple predicates using `GsQuery`, than to add additional criteria using enumeration methods.

For example, the following code returns a collection of all employees who are 26 or younger, and who respond false to `hasOtherHealthInsurance`.

```
GsQuery fromString: 'each.age <= 26' on: myEmployees)
  reject: [:each | each hasOtherHealthInsurance]
```

This may be useful if you have predicates that require message sends. However, if you can formulate the second statement as an indexable predicate, it would be more efficient as a query. If `hasOtherHealthInsurance` was actually an instance variable, you could write this as:

```
(GsQuery fromString: '(each.age <= 26) &
  (each.hasOtherHealthInsurance) not' on: myEmployees)
  queryResults
```

### Early exit from execution

Since the code in the block provided to `select:` (and similar methods) is executed for each element that the indexed query itself would return, this provides a way to exit the indexed query early. In this block, you can execute any code (as long as it does not modify the collection or the objects in the collection, in ways that would change the result set). If it's no longer useful to continue the search, you can exit the block and potentially save a lot of time.

For example, say you have a collection of purchase orders, and you are generating a report of all open purchase orders. If a new order arrives during the period you are executing this operation, you might want not want to bother producing the already-obsolete report.

```
(GsQuery fromString: 'each.isOpen' on: MyOrders) do:
  [:anOrder |
  report add: anOrder description.
  self checkForNewOrders ifTrue: [^'report canceled']
  ]
```

## Query results as Streams

It may be more useful to return the result of an equality query as a stream, instead of a collection, especially if the result set is large. Returning the result as a stream not only is faster, is also avoids the need to have all the result objects in memory simultaneously.

You can stream on an identity query only when using a `btreePlusIndex`. You cannot stream on the results of an identity `legacyIndex`.

Streaming on index results return the results in order that is defined by the index, so you can iterate over the elements that are returned in the order defined by the index, with no extra effort.

To get the results as a stream, use the message `GsQuery >> readStream` or `GsQuery >> reversedReadStream`.

These methods return an instance of a specialized subclass of `Stream` that understand a limited number of `ReadStream` protocol. Legal messages to an index stream are:

```
atEnd
do:
next
reversed
size
skip:
```

Streams do not automatically save the resulting objects. If you do not save them as you read them, the results of the query are lost. You should not modify the objects in the base collection while streaming, nor add or remove objects; doing so can cause an error or corrupt the stream.

For example, suppose your company wishes to send a congratulatory letter to anyone who has worked there for thirty years or more. Once you have sent the letter, you have no further use for the data. Assuming that each employee has an instance variable called `lengthOfService`, and there is an index on this, you can use a stream to formulate the query as follows:

```
oldTimers := (GsQuery fromString: 'each.lengthOfService >= 30'
               on: myEmployees) readStream.
[ oldTimers atEnd ] whileFalse: [
  | anEmployee |
  anEmployee := oldTimers next.
  anEmployee sendCongratuations. ].
```

### Limitations on streamable queries

Streams on query results have certain limitations; for example, the predicate in the query must be logically streamable. The following restrictions apply:

- ▶ It takes a single predicate only; no conjunction of predicate terms is allowed. The exception is range predicates, which can be combined into a single predicate. For

example `(each.age > 18) & (each.age <= 65)` is legal, since it can be reformulated as a single range predicate, `(18 < each.age <= 65)`.

- ▶ The predicate can contain only one path.
- ▶ The collection you are streaming over must have an equality index on the path specified in an equality predicate; or have an identity `btreePlusIndex` on the path specified by an identity predicate.

## 8.5 Enumerated and Set-valued Indexes

### Enumerated path terms in indexes and queries

Enumerated path terms allow you query over more than one instance variable value in a single query. This is specified using the vertical bar | in the path term, between the instance variable names.

The instance variables are treated as alternate choices; if any one of the specified instance variables matches the search criteria, the predicate evaluates to true.

For example, you might want to search on both first name and nickname in a single operation. The query might look like this:

```
(GsQuery fromString: 'each.firstName|nickName = ''Freddie''  
on: MyEmployees) queryResult
```

When this is executed, the results will include all instances that have either the firstName equal to 'Freddie', or the nickName 'Freddie', or both.

In order to optimize this query with an index, you need to create an index on the specific enumeration, e.g. 'each.firstName|nickName'. An enumerated path term query will not use an index on the individual instance variables that are enumerated.

### Restrictions on predicates with enumerated pathTerms

The semantics of enumerated pathTerms do not allow multiple conjoined predicates using the same enumerated pathTerm, since each predicate is evaluated separately. (conjoined predicates are those connected using &).

### Indexes and Queries with collections on the path

Your business objects may themselves contain collections; for example, an employee may contain a collection of children; and you may want to search based on some criteria of the objects in that collection. As long as this collection is itself indexable, indexes and queries can include all elements within these contained collections.

Index paths that include collections, and the queries that use these indexes, are generally referred to as Set-valued indexes and queries for historical reasons, although any kind of indexable collection, not just Sets, may be used.

When you wish to specify a path containing an instance of a subclass of UnorderedCollection, the collection is represented by an asterisk \*. This syntax may be used to create indexes and perform queries. Only GsQuery may be used to perform set-valued queries.

For example, suppose you want to know which of your employees has children of age 18 or younger. To facilitate such queries, each of your employees has an instance variable named *children*, which is implemented as a set. This set contains instances of a class that has an instance variable named *age*.

To create the index:

```
GsIndexSpec new  
equalityIndex: 'each.children.*.age'  
lastElementClass: SmallInteger;  
createIndexesOn: myEmployees.
```

## Set-valued query results

When you execute a set-valued query, the results you get will follow the particular semantics of Set-valued queries. Since there are potentially multiple “true” query results for a given element in the base collection, the result of a set-valued query such as this can be larger than the original collection.

For example, consider the following query, using the index created above:

```
(GsQuery fromString: 'each.children.*.age <= 18'  
  on: myEmployees) queryResult
```

In this example, if the root collection `myEmployees` is a `Bag` or `IdentityBag` (rather than a `Set` or `IdentitySet`), and an employee has two children that are under 18, then that employee will appear in the results (a `Bag` or `IdentityBag`) twice. Employees with three minor children appear in the results three times, and so on. The resulting collection may be several times as large as the original collection, depending on the details of the query and data.

If the root collection `myEmployees` is a `Set`, which does not allow multiple instances of the same object, this potential source of confusion does not occur.

## Restrictions on predicates in set-valued queries

The semantics of set-valued indexes do not allow multiple conjoined predicates that use the same set-valued `pathTerm`, since each predicate is evaluated separately. (conjoined predicates are those connected using `&`).

In general, it is recommended to avoid using multiple- set-valued predicate queries, although some multiple-predicate set-valued queries can be optimized, or avoid the problem cases, and are safe and therefor allowed.

## 8.6 Managing Indexes

You may need to find out about all the indexes in your system, and to remove selected indexes or clean up indexes that were not successfully created. This functionality is provided by the class `IndexManager`.

`IndexManager` has a single instance which provides much of the functionality, accessible via `IndexManager current`.

This instance is lazy initialized, and stored in the `IndexManager` class instance variable after it is created. Any configuration you do on `IndexManager current`, therefore, will be used by all affected operations, if you commit after making the change.

### While Indexes are Being Created

Indexing a large collection will take some amount of time to create the infrastructure and tracking for each indexed object.

The message `progressOfIndexCreation` returns a description of the current status for an index as it is created.

## Queries during index creation

While the index is being created, the index is write-locked. Any query that would normally use the index is performed directly on the collection, by brute force. If a concurrent user modifies an object that is actively participating in the index at the same time, index creation is terminated with an error.

## Auto-commit

Creating or removing an index creates and/or modifies many objects related to the internal structures that support indexes. These modifications are uncommitted changes that must be kept in the session's memory until these changes are committed. Many uncommitted changes place a large demand on memory and creates a risk of out of memory conditions. Chapter 9, "Transactions and Concurrency Control", explains uncommitted objects and transactions in more detail, while Chapter 16, "Performance and Optimization" includes information on object memory use.

To avoid problems during index creation, it is often necessary to set the IndexManager to autoCommit. When IndexManager is set to autoCommit, it will commit the partially created index, rather than risk running out of resources and failing the index operation.

By default, autoCommit is false. When you send the following message:

```
IndexManager autoCommit: true
```

it configures your IndexManager such that the current transaction is committed during an indexing operation, whenever any of the following occur:

- ▶ The current session receives a signal indicating temporary object memory is almost full.
- ▶ The percentage of temporary object memory in use reaches the IndexManager's setting for percentTempObjSpaceCommitThreshold.

The default is 60. This threshold can be changed using `IndexManager >> percentTempObjSpaceCommitThreshold: anInt`

- ▶ The current session receives a signal to FinishTransaction. This occurs when the commit record backlog is larger than STN\_SIGNAL\_ABORT\_CR\_BACKLOG, and this session is holding the commit record.
- ▶ The number of modified objects in the current transaction reaches the IndexManager's setting for dirtyObjectCommitThreshold.

The default is SmallInteger maximum value, which means this limit is effectively disabled. This limit can be changed using `IndexManager >> dirtyObjectCommitThreshold: anInt`

When autoCommit is true, a transaction will be started (if necessary) before the indexing operation begins, and the IndexManager will commit at the completion of the indexing operation. Note that this means that, even if you are in manual transaction mode and not in a transaction, index operations will cause changes to be committed to the repository without you explicitly beginning a transaction.

If you want to enable autoCommit only for the current session, not for all index creation, you can use

```
IndexManager sessionAutoCommit: true
```

## Indexes on temporary collections

You may create indexes on temporary collections containing temporary and persistent objects. However, on abort, any indexes on temporary collections are removed.

## Inquiring About Indexes

For a full description of the indexes on a particular collection, send `indexSpec` to the collection. This produces a string containing the `GsIndexSpec` code that would recreate the same indexes, and provides useful documentation on those indexes.

For example,

```
myEmployees indexSpec printString
%
GsIndexSpec new
  equalityIndex: 'each.age'
  lastElementClass: SmallInteger;
  equalityIndex: 'each.address.state'
  lastElementClass: String;
  options: GsIndexOptions reducedConflict;
  identityIndex: 'each.userId';
  yourself.
```

The following `IndexManager` messages allow you to inquire about all indexes in the repository.

▶ `getAllNSCRoots`

Returns a collection of all `UnorderedCollections` in the repository that have indexes.

▶ `usageReport`

Returns a report on all indexes on all `UnorderedCollections` in the repository.

## Removing Indexes

There are a number of ways to remove indexes.

Since indexing internal structures create references to the indexed collection and to objects in the collection, before dereferencing a collection, you should be sure to remove all indexes on the collection. This allows the collection to be garbage collected.

### To remove indexes based on a `GsIndexSpec`

As you can create indexes based on an instance of `GsIndexSpec`, you can also use that specification to remove these indexes.

```
GsIndexSpec >> removeIndexesFrom: aCollection
```

This method removes the indexes described by the `GsIndexSpec` from the collection `aCollection`. If any of the indexes do not exist, they are not removed and no error is returned.

This is most useful in combination with the method that creates the spec from the existing collection. For example:

```
(MyEmployees indexSpec)
  removeIndexesFrom: MyEmployees.
```

To remove a single index, you may edit the specification code printed by `indexSpec`, or create a simple `GsIndexSpec` with information to remove a single index:

```
(GsIndexSpec new
  equalityIndex: 'each.age' lastElementClass: Object)
  removeIndexesFrom: MyEmployees.
```

## To remove indexes using IndexManager

`IndexManager`, which provides a system-wide view of all the indexes in the repository, provides a number of methods to remove indexes both individually, by collection, and globally.

`IndexManager >> removeEqualityIndexFor: aCollection on: aPathString`

Removes an equality index from the collection *aCollection* with the indexed path described by *aPathString*. If the path specified does not exist, this method returns an error. Implicit indexes are not removed.

`IndexManager >> removeIdentityIndexFor: aCollection on: aPathString`

Removes the identity index from the collection *aCollection* with the indexed path described by *aPathString*. If the path specified does not exist, this method returns an error. Implicit indexes are not removed.

`IndexManager >> removeAllIndexesOn: aCollection`

Removes all explicitly created indexes from the collection *aCollection*. Implicit indexes that were created by these elements participating in other indexed collections are not removed.

`IndexManager >> removeAllIndexes`

Removes all indexes on all `UnorderedCollections`, including all implicit and partial indexes.

`IndexManager >> removeAllTracking`

Removes all indexes on all `UnorderedCollections`, and all object tracking. While this is the fastest way and most complete way to remove indexing infrastructure, if you are using modification tracking for any other purpose, that tracking will be removed as well.

## Rebuilding Indexes

When objects that participate in an index are modified, the related indexing infrastructure must be updated. This causes some overhead. If you are performing an operation that will modify a large number of objects that participate in multiple indexes, such as a large migration, it may be more efficient to remove some or all of the indexes on the collection before performing the migrate, and rebuild those indexes after the migration is complete.

It is also sometimes required to remove and rebuild indexes as part of a GemStone upgrade; certain changes in GemStone kernel classes require you to either rebuild specific kinds of, or all, indexes. Any requirement to do this will be included in upgrade instructions in the *Installation Guide* for the version of GemStone to which you are upgrading.

To remove and rebuild indexes, you can extract and save the `GsIndexSpec`, and reuse that after the operation is complete.

For example:

```
| mySpec |  
mySpec := myCollection indexSpec.  
mySpec removeAllIndexesFrom: myCollection.  
<perform migration or other operation>  
mySpec createIndexesOn:myCollection
```

Using `IndexManager >> getAllNSCRoots`, you may extend this example to retrieve the `GsIndexSpec` for each collection in the repository, which will allow you to remove and rebuild the indexes.

## Indexing Errors

To ensure that indexing structures are consistent, some kinds of errors that may occur during index creation will disable commits. Before creating an index, it is advisable to commit any work in progress, to avoid losing any work if an indexing error does occur.

For example, if you create an index on a collection and one or more of the objects that participate in the index do not implement the instance variable on the path, it will raise an error (unless using optionalPathTerms, as described on page 145).

If an error occurs partly through index creation, and the `autoCommit` status (see “Auto-commit” on page 153) means that some portion of the index creation was committed, a collection may have unusable partial indexes. These indexes must be manually removed.

The following `IndexManager` instance methods allow you to remove incomplete indexes, while not affecting any complete, usable indexes:

```
IndexManager current removeAllIncompleteIndexes
```

Removes all incomplete indexes on all `UnorderedCollections`.

```
IndexManager current removeAllIncompleteIndexesOn: anNSC
```

Removes all incomplete indexes on the specified `UnorderedCollection`.

If you modify objects that participate in an index, try to commit your transaction, and your commit operation fails, query results can become inconsistent. If this occurs, abort the transaction and try again.

## Auditing Indexes

Indexes should be audited regularly, as part of your regular application maintenance, to ensure there are no problems.

You can audit the internal indexing structures for a particular collection by executing:

```
aCollection auditIndexes
```

This audits all the indexes, explicit and implicit, on the given collection. If indexes are correct, this method returns 'Indexes are OK' or 'Indexes are OK and the receiver participates in one or more indexes.'. If there are no indexes on the collection, a message such as 'No indexes are present.' is returned.

In the case of failure, a list of specific problems is returned.

You can audit all indexes in the entire repository at once using:

```
IndexManager current nscsWithBadIndexes
```

which will return an IdentitySet containing all collections that fail auditIndexes. Depending on the number of indexed collections in your system, this may take a considerable time to run.

In the rare case of a problem reported, the usual way to resolve the problem is to remove and rebuild the affected indexes. In some cases, removing all indexes on the collection may succeed even if the internal problems prevent a single index being removed.

## 8.7 Indexing and Performance

The value of Indexes is to improve performance, of course. It is always recommended to perform tests to verify performance improvements.

Indexing improves query performance dramatically (in most cases), but does have a negative impact on updating the indexed data, since the indexes must be kept up to date.

### Type of index

The performance characteristics of btreePlus and legacy indexes are quite different.

btreePlus indexes without optimized comparison are usually slower than other kinds of indexes. If your desired index cannot support optimizedComparison, you should use a legacyIndex.

btreePlus optimizedComparison indexes are usually considerably faster than a legacy index, but they create a somewhat larger negative impact on data updates.

### Data updates

As your application is in use and the data in the indexed collection changes, the index must be updated. While normally indexing a large collection speeds up queries performed on that collection and has little effect on other operations, there are cases in which maintaining the index can cause a performance bottleneck.

For example, you may notice slower than acceptable performance if you are making a great many modifications to the instance variables of objects that participate in an index, and more than one of the following is true:

- ▶ the path of the index is long;
- ▶ the object occurs many times within the indexed IdentityBag or Bag
- ▶ the object participates in many indexes

Even so, indexing a large collection is still likely to improve performance unless more than one of these circumstances holds true. If you do experience a performance problem, you can work around it in one of two ways:

If you have created relatively few indexes but are modifying many indexed objects, it may be worthwhile to remove the indexes, modify the objects, and then re-create the indexes.

If you are making many modifications to only a few objects, or if you have created a great many indexes, it is more efficient to commit frequently during the course of your work. That is, modify a few objects, commit the transaction, modify a few more objects, and commit again.

## Formulating queries and performance

The most efficient queries are the ones in which the first predicate will return the smallest result set. This is sometimes easy for a human to determine, but the query cannot predict this without actually running the query. Queries should be manually reviewed for these kinds of domain-specific optimizations.

For example, you might want to query for current orders for a particular customer.

```
(each.status = #current) & (each.customer.name = 'Smith')
```

If your application is likely to have only a few current orders, then this is more efficient. However, if you are likely to have many current orders, but only a few customers named Smith, it would be more efficient for you to write the formula in reverse order.

### Auto-optimize

Queries, by default, are optimized before execution; for example, the not operator is transformed into the logical equivalent by changing the comparison operator.

In addition, the predicates are reordered as follows, from left to right:

1. predicates involving indexed paths.
2. predicates with identity comparisons on paths without indexes.
3. predicates with equality comparisons on paths without indexes.

Auto-optimize can be disabled using the instance of `GsQueryOptions` that is associated with each query. The `GsQueryOptions` instance controls optimization and other query features. In addition to the various specific optimizations performed, `GsQueryOptions` controls if automatic query optimization is done; the default is to do auto-optimization.

## 8.8 Historic Indexing API differences

In older versions of GemStone/S and GemStone/S 64 Bit, indexes and queries used a more limited API based on `UnorderedCollection` methods and a block-like query syntax. This API remains fully supported and interoperates with the `GsIndexSpec/GsQuery` API, with some limitations. A number of features are not supported by the older API.

### Index creation using `UnorderedCollection` protocol

`UnorderedCollection` provides protocol to create indexes. This creates the same index structures as `GsIndexSpec`, but does not provide access to some index features.

The following index creation methods are defined on `UnorderedCollection`:

```
createIdentityIndexOn:  
createEqualityIndexOn:withLastElementClass:
```

The path argument is the same as the path used to create a `GsIndexSpec` index, however you may not include the initial "each".

For example, the following three statements create the same indexes that were created on page 141.

```
myEmployees createIdentityIndexOn: 'userId'.
myEmployees
    createEqualityIndexOn: 'age'
    withLastElementClass: SmallInteger.
myEmployees
    createEqualityIndexOn: 'address.state'
    withLastElementClass: String.
```

Enumerated and set-value indexes and queries are not supported using historic API.

## Internal legacy vs. btreePlus indexing structures

The used of legacyIndex or btreePlusIndex/optimizedComparison is based on the default GsIndexOptions. Whatever the session or system default is will determine the type of index being created

## String and Unicode Equality Indexes

Indexes on various kinds of strings follow the same rules as GsIndexSpec string indexes, with the exception that the optimized indexes cannot be created this way.

To create unicode indexes, specify a lastElementClass of any Unicode string class (Unicode7, Unicode 16, or Unicode32). Since no collator can be specified, the index will be created using the current default IcuCollator.

## Reduced-conflict Equality Indexes

An Rc Equality Index is a type of Equality Index in which internal indexing structures are reduced-conflict. This avoids some transaction conflicts when creating an index on a reduced-conflict (RC) collection, such as RcIdentityBag. Reduced-conflict classes are described in “Indexes and Concurrency Control” on page 172. Rc Equality indexes are described under “Reduced-Conflict” on page 145.

Using UnorderedCollection index creation protocol to create an index, the message is:

```
createRcEqualityIndexOn:withLastElementClass:
```

## Queries using Selection Blocks

*Selection blocks* are a kind of block specialized for queries, using curly braces instead of brackets. The compiler understands this syntax and creates the selection block instance when the code or method is compiled.

A selection block query might be written like this:

```
{:each | each.address.state = 'OR'}
```

Selection blocks are quite restrictive:

- ▶ A selection block has exactly one argument
- ▶ Message sends are not allowed in a selection block; you can only use the dot syntax to specify instance variables of the argument.
- ▶ The code inside the block is limited to predicates as described under “Query Predicate Syntax” on page 137, with additional limitations below.

- ▶ Set valued and enumerated syntax are not allowed in a selection block
- ▶ Range predicate syntax are not allowed in a selection block, although you may specify the same operation by conjoining two separate predicates.
- ▶ Selection block queries do not allow the `|` (disjunction operator), nor the `not` operator.
- ▶ Selection block can only be used as arguments to the methods `select:`, `reject:`, `detect:`, `detect:ifNone:`, or `selectAsStream:`.
- ▶ Selection block queries are not optimized.

In selection block queries, you can reference temporary, instance or other variables within the block, and these are resolved at runtime as in ordinary blocks.

## Executing Selection Block Queries

A selection block is used with `select:`, `reject:`, `detect:`, `detect:ifNone:`, or `selectAsStream:` to perform the query over a collection.

For example:

```
Employees select: { :each | each.address.state = 'OR' }
```

These have the same semantics as with standard blocks executed on a collection. For example, `reject:` will return a result set that includes all elements for which the block evaluation would return false. The results are in a collection the same class as the base collection (unless `species` or `speciesForSelect` specifies a different class, as with the RC classes).

The collection returned from a query has no index structures. If you want to perform indexed selections on the new collection, you must build the necessary indexes on the new collection.

## Results as a stream

To get the results as a stream, use `UnorderedCollection >> selectAsStream:`. This returns an instance of `RangeIndexReadStream`, which understands the following messages:

`next`

Returns the next value on a stream of range index values.

`atEnd`

Returns true if there are no more elements to return through the logical iteration of the stream.

`reversed`

Create a `ReversedRangeIndexReadStream` based on the receiver, allowing you to stream over the results from last to first.

## Creating a GsQuery from a selection block

If you have existing code that includes selection block queries, you can use those selection blocks to create the instances of `GsQuery`.

For example,

```
GsQuery fromSelectBlock: { :each | each.address.state = 'OR' }
```

This can be bound using `on:`, or created using `fromSelectBlock: on:`, similar to how you create and bind a `GsQuery` from a string.

## Managing indexes

### Information about indexes

Sending `indexSpec` to the collection provides a complete description of the indexes on a collection, and can be used for information without using the `GsIndexSpec` API; the extra details provided by `indexSpec` can be ignored.

You can also send messages to the collection that will return quick information on indexed paths.

`equalityIndexedPaths` and `identityIndexedPaths`

Returns, respectively, the equality indexes and the identity indexes on the receiver's contents. Each message returns an array of strings representing the paths in question.

For example, the following expression returns the paths into `myEmployees` that bear equality indexes:

```
myEmployees equalityIndexedPaths
%
anArray( 'age', 'address.state')
```

`kindsOfIndexOn: aPathNameString`

Returns information about the kind of index present on an instance variable within the elements of the receiver. The information is returned as one of these symbols: `#none`, `#identity`, `#equality`, `#identityAndEquality`.

`equalityIndexedPathsAndConstraints`

Returns an array in which the odd-numbered elements are the elements of the path, and the even-numbered elements are the constraints specified when creating an index using the keyword `withLastElementClass:`.

### Removing Indexes

Removing indexes can be done using the `GsIndexSpec`

You may send methods to the indexed collection directly to remove one or all indexes.

`UnorderedCollection >> removeEqualityIndexOn: aPathString`

Removes an equality index from the path indicated by `aPathString`. If the path specified does not exist, this method returns an error. Implicit indexes are not removed.

`UnorderedCollection >> removeIdentityIndexOn: aPathString`

Removes the identity index on the specified path. If the path specified does not exist, this method returns an error. Implicit indexes are not removed.

`UnorderedCollection >> removeAllIndexes`

Removes all explicitly created indexes from the receiver. Implicit indexes that were created by these elements participating in other indexed collections are not removed.



# Transactions and Concurrency Control

---

GemStone users can share code and data objects by maintaining common dictionaries that refer to those objects. However, if operations that modify shared objects are interleaved in any arbitrary order, inconsistencies can result.

This chapter describes how GemStone manages concurrent sessions to prevent inconsistencies resulting from multiple concurrent updates.

**GemStone's Conflict Management** (page 163)

introduces the concept of a transaction and describes how it interacts with each user's view of the repository.

**How GemStone Detects and Manages Conflict** (page 169)

describes how commit conflicts are detected and reported and how to handle and avoid conflicts.

**Controlling Concurrent Access with Locks** (page 175)

discusses the kinds of lock you can use to prevent conflict.

**Classes That Reduce the Chance of Conflict** (page 185)

describes the classes that help reduce the likelihood of a conflict.

## 9.1 GemStone's Conflict Management

GemStone prevents conflict between users by encapsulating each session's operations (computations, stores, and fetches) in units called *transactions*. The operations that make up a transaction act on what appears to you to be a private *view* of GemStone objects. When you tell GemStone to *commit* the current transaction, GemStone tries to merge the modified objects in your view with the shared object store.

### Note on Terminology

GemStone's terminology differs from commonly used DBMS terminology. A GemStone "view" is a consistent representation of all objects in the repository at a given point in time; essentially a snapshot. This document uses the term "snapshot view" for clarity.

GemStone uses the term "commit record" for the internal structure that tracks what composes a view or snapshot. This term may be used equivalently to view; for example, the oldest commit record is the oldest snapshot view of the repository.

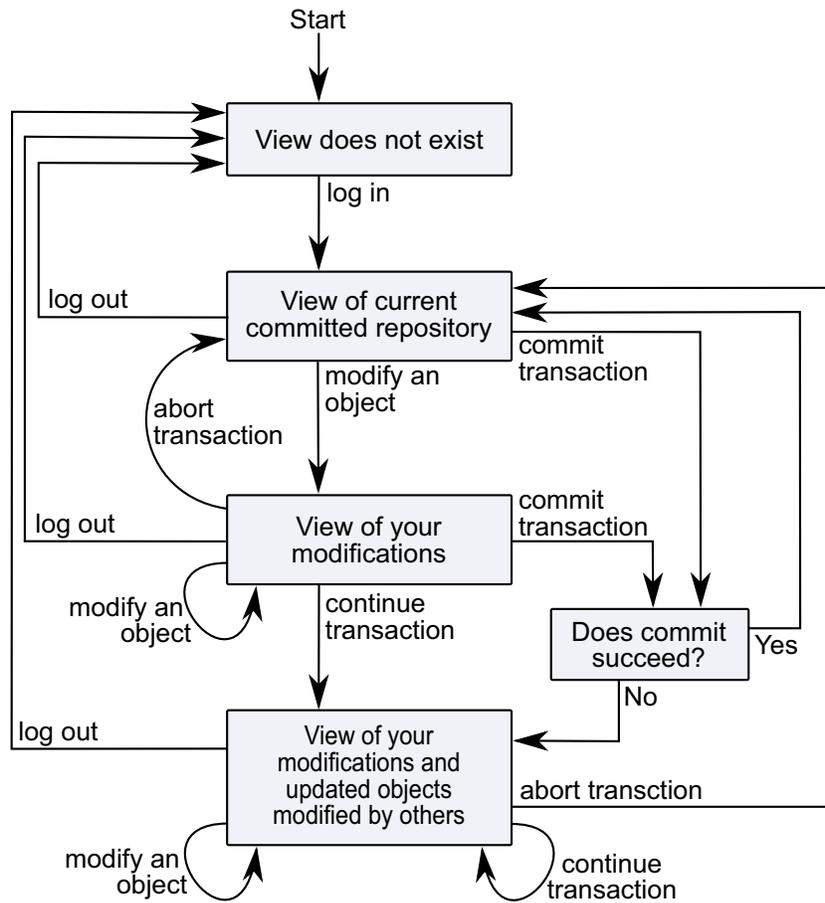
## Snapshot Views and Transactions

As shown in Figure 9.1, every user session maintains its own consistent snapshot view of the repository state. Objects that the repository contained at the beginning of your session are preserved in your snapshot view, even if you are not using them – and even if other users' actions have rendered them obsolete. The storage that those objects are using cannot be reclaimed until you commit or abort your transaction. Depending upon the characteristics of your particular installation (such as the number of users and the commit frequency), this burden can be trivial or significant.

When you log in to GemStone, you get a snapshot view of the latest repository state. After login, you may start a transaction automatically or manually, or remain outside of transaction. The repository snapshot view you get on login is updated when you begin a transaction or abort. When you commit a transaction, your changes are merged with other changes to the shared data in the repository, and your snapshot view is updated. When you obtain a new snapshot view of the repository, by commit, abort, or continuing, any new or modified objects that have been committed by other users become visible to you.

The transaction mode controls if a transaction is automatically started, or if you must manually begin a transaction. For details, see "Committing Transactions" on page 170.)

Figure 9.1 View States



## Transaction State and Transaction Modes

A GemStone session is always either in a transaction or not in a transaction. When in transaction, changes can be committed to the repository. When not in transaction, you can make changes to objects in the repository, but these changes cannot be committed.

A session that is in transaction may be in one of a number of transaction levels, depending on if nested transactions are involved.

When not in transaction, the session may merely be not in transaction, or it may be in the specialized transactionless mode. In transactionless mode, the session is not in transaction, but its snapshot view may be updated automatically at any time. Transactionless mode is primarily for idle sessions that do not need consistent repository data, since objects may change at any time; the topics that this chapter discusses, for the most part, do not apply to transactionless mode sessions.

The transaction modes provide different behavior with respect to starting new transactions. When in automatic transaction mode, the session is always in transaction. When in manual transaction mode, you may be in transaction or not in transaction, depending on specific messages your session sends.

The following are the GemStone transaction modes:

### Automatic transaction mode

In this mode, GemStone begins a transaction when you log in, and starts a new one after each commit or abort message. In this default mode, you are in a transaction the entire time you are logged into a GemStone session. Use caution with this mode in busy production systems, since your session will not receive the signals that your snapshot view is causing a strain on system resources.

This is the default transaction mode on login.

To change to automatic transaction mode, send the message:

```
System transactionMode: #autoBegin
```

This aborts the current transaction and starts a new transaction.

### Manual transaction mode

In this mode, you can be logged in and be outside of a transaction. You explicitly control whether your session starts a transaction, makes changes, and commits. Although a transaction is started for you when you log in, you can set the transaction mode to manual, which aborts the current transaction and leaves you outside a transaction. You can subsequently start a transaction when you are ready to start making changes that you want to commit. Manual transaction mode provides a method of minimizing the transactions, while still managing the repository for concurrent access.

In manual transaction mode, you can browse classes, data, and other objects in the repository, and make computations based upon object values. You cannot, however, make your changes permanent, nor can you add any new objects you may have created while outside a transaction. You can start a transaction at any time during a session; you can carry temporary results that you may have computed while outside a transaction into your new transaction, where they can be committed, subject to the usual constraints of conflict-checking.

To change to manual transaction mode, send the message:

```
System transactionMode: #manualBegin
```

This aborts the current transaction and leaves the session not in transaction.

To begin a transaction, execute

```
System beginTransaction
```

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

### Transactionless mode

In transactionless mode, you remain outside a transaction. This mode is intended primarily for idle sessions. If all you need to do is browse objects in the repository, transactionless mode can be a more efficient use of system resources. However, you are at risk of obtaining inconsistent repository data.

To change to transactionless transaction mode, send the message:

```
System transactionMode: #transactionless
```

This aborts the current transaction and leaves the session out of transaction.

### Determining transaction mode and transaction state

To determine the transaction mode you are in, send the message:

```
System transactionMode
```

To determine the transaction level you are at, send the message:

```
System transactionLevel
```

A transaction level of 1 or more means your session is in transaction, with values greater than 1 indicating the number of levels of transaction. A transaction level of 0 is not in transaction, while -1 indicates transactionless.

You can determine whether you are currently in a transaction by sending the message:

```
System inTransaction
```

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

## Reading and Writing in Transactions

GemStone considers the operations that take place in a transaction (or outside of a transaction) as *reading* or *writing* objects. Any operation that sends a message to an object, or accesses any instance variable of an object, is said to *read* that object. An operation that stores something in one of an object's instance variables is said to *write* the object. While you can read without writing, writing an object always implies reading it. GemStone must read the internal state of an object in order to store a new value in the object.

Operations that fetch information about an object also read the object. In particular, fetching an object's size, class, or security policy reads the object. An object also gets read in the process of being stored into another object.

The following expression sends a message to obtain the name of an employee and so reads the object:

```
theName := anEmployee name. "reads anEmployee"
```

The following example reads aName in the same operation that anEmployee is written:

```
anEmployee name: aName. "writes anEmployee, reads aName"
```

Some less common operations cause objects to be read or written. For example, assigning an object to a new object security policy, using the message `assignToObjectSecurityPolicy:`, writes the object and reads both the old and the new `GsObjectSecurityPolicy`. Modifying an object that participates in an index may write support objects built and maintained as part of the indexing mechanism.

For the purposes of detecting conflict among concurrent users, GemStone keeps separate sets of the objects you have written during a transaction and the objects you have only read. These sets are called the *write set* and the *read set*; the read set is always a superset of the write set.

## Reading and Writing Outside of Transactions

Outside of a transaction, reading an object is accomplished precisely the same way. You can write objects in the same way as well, but you cannot commit these changes to make them a permanent part of the repository.

## When Should You Commit a Transaction?

Most applications create or modify objects in logically separate steps, combining trivial operations in sequences that ultimately do significant things. To protect other users from reading or using intermediate results, you want to commit after your program has produced some stable and usable results. Changes become visible to other users only after you've committed.

Your chance of being in conflict with other users increases with the time between commits.

## Nested In-memory Transactions

Within a transaction, GemStone allows you to group units of work into logical transactions, which can be committed or aborted within the given session. These logical transactions can be nested with up to 16 levels of nesting (including the outer level actual transaction). When the full set of changes are ready to be committed, committing the outer transaction will make the changes persistent and detect any conflicts.

While the same protocol is used to commit the actual (outer) transaction and the nested transactions, the semantics are different. A commit of a nested transaction does not detect conflicts with changes by other users, does not update current session state, and does not make the changes persistent if the session exits unexpectedly or recoverable on system shutdown. Abort of a nested transaction returns the session to the state it was in at the beginning of the nested transaction, without updating the session's snapshot view with any changes by other users.

When transactions are discussed, unless specified otherwise, it only refers to an outer level actual transaction, not to a nested transaction.

To begin a nested transaction, use

```
System beginNestedTransaction
```

You should be already in transaction when executing this method.

Executing `commit`, `commitTransaction`, `abort`, or `abortTransaction` when in a nested transaction preserve or discard in-memory changes and return to the parent level of transaction. The same protocol is used at the outer level, actual transaction to perform the commit or abort.

`continueTransaction` cannot be used when in a nested transaction.

You can commit or abort all levels of nested transactions at once, including performing the outer level actual commit or abort, using the messages:

```
System commitAll
System abortAll
```

## 9.2 How GemStone Detects and Manages Conflict

GemStone detects conflict by comparing your write set with those of all other transactions that committed since your transaction began. The following conditions signal a possible concurrency conflict:

- ▶ An object in your write set is also in the write set of another transaction – a *write-write conflict*. Write-write conflicts can involve only a single object.
- ▶ An object in your write set is also in another session's dependency list – a *write-dependency conflict*. An object belongs to a session's *dependency list* if the session has added, removed, or changed a dependency (index) for that object. For details about how GemStone creates and manages indexes on collections, see Chapter 8, "Indexes and Querying".

If a write-write or write-dependency conflict is detected, then your transaction cannot commit; you must abort, and try again. The following section describes some approaches to handling this kind of situation.

### Concurrency Management

As the application designer, you determine your approach to concurrency control.

- ▶ Using the *optimistic* approach to concurrency control, you simply read and write objects as if you were the only user. The object server detects conflicts with other sessions only at the time you try to commit your transaction. Your chance of being in conflict with other users increases with the time between commits and the size of your write set.

Although easy to implement in an application, this approach entails the risk that you might lose the work you've done if conflicts are detected and you are unable to commit.

- ▶ Using the *pessimistic* approach to concurrency control, you detect and prevent conflicts by explicitly requesting *locks* that signal your intentions to read or write objects. By locking an object, other users are unable to use the object in a way that conflicts with your purposes. If you are unable to acquire a lock, then someone else has already locked the object and you cannot use the object. You can then abort the transaction immediately instead of doing work that can't be committed.
- ▶ Using *reduced-conflict (RC) classes* in places where write-write conflicts are likely. RC classes use internal structures and additional logic to allow a commit to succeed in spite of a write-write conflict, when the changes do not actually conflict with each other.

The GemStone reduced-conflict classes include: RcCounter, RcIdentityBag, RcIdentitySet, RcArray, RcPipe, RcQueue, and RcKeyValueDictionary. See "Classes That Reduce the Chance of Conflict" on page 185.

## Committing Transactions

Committing a transaction has two effects:

- ▶ It makes your new and changed objects visible to other users as a permanent part of the repository.
- ▶ It makes visible to you any new or modified objects that have been committed by other users in an the most recent state of the repository.

When you tell GemStone to commit your transaction, the object server performs these actions:

1. Checks whether other concurrent sessions have committed transactions that modify an object that you modified during your transaction.
2. Checks to see whether other concurrent sessions have added, removed, or changed indexes on an object that you have modified during your transaction.
3. Checks for locks set by other sessions that indicate the intention to modify objects that you have read.

If none of these conditions is found, GemStone commits the transaction.

The methods `System class >> commit` or `System class >> commitTransaction` commit the current transaction.

`System commitTransaction` returns true if GemStone commits your transaction and false if the commit fails and it cannot commit.

`System commit` performs the same commit, but returns true if GemStone commits your transaction and signals an error if it fails to commit.

## Handling Commit Failure in a Transaction

If GemStone refuses to commit your transaction, the transaction read or wrote an object that another user modified and committed to the repository (or involved in indexing operations) since your transaction began. Because you can't undo a read or a write operation, simply repeating the attempt to commit will not succeed.

You must abort the transaction in order to get a new snapshot view of the repository and, along with it, an empty read set and an empty write set. A subsequent attempt to run your code and commit can succeed. If the competition for shared data is heavy, subsequent transactions can also fail to commit. In this situation, locking objects that are frequently modified by other transactions gives you a better chance of committing.

## Transaction Conflicts

To find why your transaction failed to commit, you can send the message:

```
System transactionConflicts
```

This method returns a symbol dictionary that contains an Association whose key is `#commitResult` and whose value is one of the following symbols:

#commitResult value	Meaning
#readOnly	There were no modified objects to commit, so the commit did not do writes. In this case, <code>commitTransaction</code> returns true.

#success	Commit was successful.
#rcFailure	The replay of changes to instances of Rc classes failed.
#dependencyFailure	Commit failed, concurrency conflict on dependencyMap.
#failure	Commit failed.
#retryFailure	Commit failed, and the previous commit attempt failed with an rcFailure.
#commitDisallowed	Commits were disallowed for other errors.
#retryLimitExceeded	Up to 15 retry attempts are allowed.
#symbolFailure	Commit failed due to a Symbol creation failure.
#lockFailure	Commit failed to do a lock failure.

The remaining Associations in the dictionary, if any, are used to report the conflicts found. Each Association's key indicates the kind of conflict detected; its associated value is an Array of OOPs for the objects that are conflicting.

Table 9.1 lists the possible keys for the conflict.

**Table 9.1 Transaction Conflict Keys**

Key	Meaning
#'Read-Write'	StrongReadSet and WriteSetUnion conflicts, with the RcReadSet subtracted.
#'Write-Write'	WriteSet and WriteSetUnion conflicts.
#'WriteWrite_minusRcReadSet'	the same as #'Write-Write', but with the RcReadSet subtracted.
#'Write-Dependency'	WriteSet and DependencyChangeSetUnion conflicts.
#'Write-WriteLock'	WriteSet and WriteLockSet conflicts.
#'Write-ReadLock'	WriteSet and ReadLockSet conflicts.
#'Rc-Write-Write'	Logical Write-Write conflict on instances of a reduced conflict class.
#'RcReadSet'	The RcReadSet
#'Synchronized-Commit'	Details of the synchronized commit failure.

If there are no conflicts for the transaction, the returned symbol dictionary has no additional Associations.

Conflict sets are cleared at the beginning of a commit or abort and thus can be examined until the next commit, continue, or abort.

*NOTE*

*If you save a reference to the conflict set, be sure to clear this references to avoid making the conflict set persistent.*

To determine whether the current transaction has write-write conflicts, you can send the following message before attempting to commit the transaction:

```
System currentTransactionHasWWConflicts
```

Similarly, to determine whether the current transaction has write-dependency conflicts, you can send this message:

```
System currentTransactionHasWDConflicts
```

If the above message returns true, you can send the appropriate message to obtain a list of write-write (or write-dependency) conflicts in the current transaction:

```
System currentTransactionWWConflicts (write-write)
```

or:

```
System currentTransactionWDConflicts (write-dependency)
```

## More details about transaction conflicts

The information provided by `transactionConflicts` lets you know the objects that were committed by another session, but does not help in tracking down which session or user committed the changes that were the cause of the conflict. You can enable tracking in your session that lets you collect these details, but this must be enabled before your commit performs a commit that fails. Determining the session whose commit caused the conflict has performance overhead, so it is not recommended that you run this way by default, but it can be useful when you have ongoing conflicts that are difficult to track down.

To enable tracking, set the runtime-only configuration parameter `GemCommitConflictDetails` to true. For example,

```
System gemConfigurationAt: #GemCommitConflictDetails put: true
```

Then, once the commit fails, execute the method `System class >> detailedConflictReportString`, which returns a string containing information about the conflicting other commit. For example,

```
System detailedConflictReportString
%
Commit failed , failure
Attempt to commit at: 2023-05-14 15:39:55.959
1 Write-Write Conflicts
  ( 12200193(a SymbolDictionary))
  1 commits by other sessions
    session 7 at 2023-05-14 15:39:45.607 userId DataCurator
      ( 12200193(a SymbolDictionary))
```

## Indexes and Concurrency Control

It is also possible that you can encounter conflict on the internal indexing structures used by GemStone. For example, if two transactions modify the salaries of different employees that participate in the same indexed set, it is possible that both transactions will modify the same internal indexing structure and therefore conflict, despite the fact that neither transaction has explicitly accessed an object written by the other transaction. It is true even if the collection itself is an Rc collection and does not encounter transaction conflicts.

To check this possibility, examine the dictionary returned by evaluating `System transactionConflicts` (described on page 171). If that dictionary includes any Associations whose key is `#'Write-Dependency'`, you have experienced a conflict on some portion of an indexing structure. In that case, you can abort the transaction and try the modification again.

If you encounter conflicts in the internal indexing structures, you can create a reduced-conflict index. See “Reduced-Conflict” on page 145.

## Aborting Transactions

If GemStone refuses to commit your modifications, your snapshot view remains intact with all of the new and modified objects it contains. However, your repository now also includes other users' modifications to objects that are visible to you, but that you have not modified.

You must take some action to save the modifications in your session or in a file outside GemStone.

Then you need to *abort* the transaction. This discards all of the modifications from the aborted transaction, and gives you a new snapshot view containing all shared, committed objects. Depending on the activities of other users, you can repeat your operations using the new values and commit the new transaction without encountering conflicts.

The messages `abort` or `abortTransaction` discard the modified objects in your snapshot view. If you are in automatic transaction mode, these messages also begin a new transaction.

Aborting a transaction discards any changes you have made to shared objects during the transaction. However, work you have done within your own object space is not affected by an `abortTransaction`. GemStone gives you a new snapshot view of the repository that does not include any changes you made to permanent objects during the aborted transaction—because the transaction was aborted, your changes did not affect objects in the repository. The new snapshot view, however, does include changes committed by other users since your last transaction started.

Objects that you have created in the GemBuilder for Smalltalk object space, outside the repository, remain until you remove them or end your session.

## Updating the View Without Committing or Aborting

The message `System continueTransaction` gives you a new, up-to-date snapshot view of other users' committed work, without discarding the objects you have modified in your current session.

The message `continueTransaction` returns true if a commit on your transaction would succeed, or false if a commit would fail. After `continueTransaction` returns false, you may examine the results of `System transactionConflicts` to see what objects have conflicts.

Unlike `commitTransaction` and `abortTransaction`, `continueTransaction` does not end your transaction. It has no effect on object locks, and it does not discard any changes you have made or commit any changes. Objects that you have modified or created do not become visible to other users.

Work you have done locally within your own interface is not affected by a `continueTransaction`. Objects that you have created in your own application remain. Similarly, any execution that you have begun continues, unless the execution explicitly depends upon a successful commit operation.

Note that if you were unable to commit your transaction due to conflicts, you cannot use `continueTransaction` until you abort the transaction.

## Being Signaled To Abort

As mentioned earlier, being in a transaction incurs certain costs. When you are in a transaction, GemStone waits until you commit or abort before it attempts to reclaim obsolete objects in your snapshot view. While you are in a transaction, your session will not be signalled to abort, nor is it subject to losing its snapshot view of the repository or being terminated as a result of sigAbort mechanisms. A session in transaction may cause your repository to grow until it runs out of disk space.

When you are outside of a transaction, GemStone warns you when your snapshot view is outdated, and this is imposing a burden on the system, by sending your session the `TransactionBacklog` notification. You are allowed a certain amount of time to abort your current snapshot view, as specified in the `STN_GEM_ABORT_TIMEOUT` parameter in your configuration file. When you abort your current snapshot view (by sending the message `System abort` or `System abortTransaction`), GemStone can reclaim storage and you get a fresh snapshot view of the repository.

If you do not respond within the specified time period, the object server sends your session the exception `RepositoryViewLost` and then terminates the Gem.

Work that you have done locally (such as references to objects within your application) is retained, and you still cannot commit work to the repository when running outside of a transaction. However, you must read again those objects that you had previously read from the repository, and recompute the results of any computations performed on them, because the object server no longer guarantees that the application values are valid.

Your GemStone session controls whether it is signalled to abort by receiving the `TransactionBacklog` notification when it is out of transaction. To enable receiving it, send the message:

```
System enableSignaledAbortError
```

To disable receiving it, send the message:

```
System disableSignaledAbortError
```

To determine whether receiving this notification is currently enabled or disabled, send the message:

```
System signaledAbortErrorStatus
```

This method returns true if the notification is enabled, and false if it is disabled. By default, GemStone sessions disable receiving this notification. The `GemBuilder` interfaces may change this default. If you wish to be notified, then you must explicitly enable the signaled abort error, and re-enable it after each time the signal is received.

## Being Signaled to continueTransaction

As described earlier, when you are in a transaction, GemStone does not signal the session to abort, nor are you subject to losing your snapshot view of the repository. This entails a risk that your repository may grow until it runs out of disk space.

To avoid this problem, you can enable your GemStone session to receive the `TransactionBacklog` notification when you are in transaction. This prompts your session that it is now holding the oldest snapshot view of the repository, and potentially causing your repository to grow. When your session receives this signal, it may execute a `continueTransaction`, or abort or commit its changes.

Your GemStone session controls whether it receives the `TransactionBacklog` notification when in transaction. To enable receiving it, send the message:

```
System enableSignaledFinishTransactionError
```

To disable receiving it, send the message:

```
System disableSignaledFinishTransactionError
```

To determine whether receiving this error message is currently enabled or disabled, send the message:

```
System signaledFinishTransactionErrorStatus
```

This method returns true if the notification is enabled, and false if it is disabled. By default, GemStone sessions disable receiving this notification. If you wish to be notified, then you must explicitly enable it after each time the signal is received.

## Handlers for abort or continueTransaction notifications

Not only do you need to enable the receipt of the notification to abort or continueTransaction, you must also set up a signal handler to take the appropriate action. Sending `enableSignaledAbortError` and `enableSignaledFinishTransactionError` control whether you receive the `TransactionBacklog` notification when you are not in transaction or when you are in transaction, respectively. The handler for the `TransactionBacklog` notification needs to take both possible situations into account.

## 9.3 Controlling Concurrent Access with Locks

If many users are competing for shared data in your application, or you can't tolerate even an occasional inability to commit, then you can implement pessimistic concurrency control by using locks.

Locking an object is a way of telling GemStone (and, indirectly, other users) your intention to read or write the object. Holding locks prevents transactions whose activities would conflict with your own from committing changes to the repository. Unless you specify otherwise, GemStone locks persist across aborts as well as commits. If you lock on an object and then abort, your session still holds the lock after the abort. Aborting the current transaction (and starting another, if you are in manual transaction mode) gives you an up-to-date value for the locked object without removing the lock.

Remember, locking improves one user's chances of committing only at the expense of other users. Use locks sparingly to prevent an overall degradation of system performance.

### Lock Types

GemStone provides two kinds of locks you may use on any objects: *read* and *write*. A session may hold only one kind of lock on an object at a time.

GemStone also provides another type of lock, *applicationWriteLock*, which is limited to a single unique lock object; it provides similar locking, but is used to provide a mutex. While the *applicationWriteLock* behaves similarly to read and write locks, it is used differently is discussed separately starting on page 184.

## Read Locks

Holding a read lock on an object means that you can use the object's value, and then commit without fear that some other transaction has committed a new value for that object during your transaction. Another way of saying this is that holding a read lock on an object guarantees that other sessions cannot:

- ▶ acquire a write lock on the object, or
- ▶ commit if they have written the object.

To understand the utility of read locks, imagine that you need to compute the average age of a large number of employees. While you are reading the employees and computing the average, another user changes an employee's age and commits (right at the end of a birthday party). You have now performed the computation using out-of-date information. You can prevent this frustration by read-locking the employees at the outset of your transaction; this prevents changes to those objects.

Multiple sessions can hold read locks on the same object. A maximum of 1 million read locks can be held concurrently. Because locking incurs a cost at commit time, you should keep the aggregate number of locked objects as small as possible.

### NOTE

*If you have a read lock on an object and you try to write that object, your attempt to commit that transaction will fail.*

## Write Locks

Holding a write lock on an object guarantees that you can write the object and commit. That is, it ensures that you won't find that someone else has prevented you from committing by writing the object and committing it before you, while your transaction was in progress. Another way of looking at this is that holding a write lock on an object guarantees that other sessions cannot:

- ▶ acquire either a read or write lock on the object, or
- ▶ commit if they have written the object.

Write locks are useful, for example, if you want to change the addresses of a number of employees. If you write-lock the employees at the outset of your transaction, you prevent other sessions from modifying one of the employees and committing before you can finish your work. This guarantees your ability to commit the changes.

Write locks differ from read locks in that only one session can hold a write lock on an object. In fact, if a session holds a write lock on an object, then no other session can hold any kind of lock on the object. This prevents another session from receiving the assurance implied by a read lock; that the value of the object it sees in its snapshot view will not be out of date when it attempts to commit a transaction.

## Acquiring Locks

The kernel class `System` is the receiver of all lock requests. The following statements request one lock of each kind:

### Example 9.1

---

```
System readLock: SharedDictionary.  
System writeLock: myEmployees.
```

---

When locks are granted, these messages return `System`.

Commits and aborts do not necessarily release locks, although locks can be set up so that they will do so. Unless you specify otherwise, once you acquire a lock, it remains in place until you log out or remove it explicitly. (Subsequent sections explain how to remove locks.)

When a lock is requested, GemStone grants it unless one of the following conditions is true:

- ▶ You do not have suitable authorization. Read locks require read authorization; write locks require write authorization.
- ▶ The object is an instance of `SmallInteger`, `Boolean`, `Character`, `SmallDouble`, or `nil`. Trying to lock these special objects is meaningless.
- ▶ The object is already locked in an incompatible way by another session (remember, only read locks can be shared).

Variants of the `readLock:` and `writeLock:` messages allow you to lock collections of objects en masse. For details, see “Locking Collections of Objects Efficiently” on page 179.

## Lock Denial

If you request a lock on an object and another session already holds a conflicting lock on it, then GemStone denies your request; GemStone does not automatically wait for locks to become available.

If you use one of the simpler lock request messages (such as `readLock:`), lock denial generates an error. If you want to take some automatic action in response to the denial, use a more complex lock request message, such as this:

```
System readLock: anObject  
    ifDenied: [block1]  
    ifChanged: [block2].
```

A lock denial causes GemStone to execute the block argument to `ifDenied:`. The method in Example 9.2 uses this technique to request a lock repeatedly until the lock becomes available.

**Example 9.2**


---

```

Object subclass: #Dummy
  instVarNames: #()
  classVars: #()
  classInstVars: #()
  poolDictionaries: #()
  inDictionary: UserGlobals
  options: #()
%
method: Dummy
getReadLockOn: anObject tries: numTries
  "This method tries to lock anObject. If the lock is denied,
  it tries again, making up to numTries attempts."
| n |
n := 1.
[n <= numTries] whileTrue: [
  System readLock: anObject
    ifDenied: [System sleep: 1.]
    ifChanged: [System abortTransaction.].
  n := n + 1].
^(System myLockKind: anObject) = #read
%
UserGlobals at: #testObject put: Object new.
System commitTransaction.
%

Dummy new getReadLockOn: testObject tries: 3
%
```

---

**Deadlocks**

You may never succeed in acquiring a lock, no matter how long you wait. Furthermore, because GemStone does not automatically wait for locks, it does not attempt deadlock detection. It is your responsibility to limit the attempts to acquire locks in some way. For example, you can write a portion of your application in such a way that there is an absolute time limit on attempts to acquire a lock. Or you can let users know when locks are being awaited and allow them to interrupt the process if needed.

**Dirty Locks**

If another user has written an object and committed the change since your transaction began, then the value of the object in your snapshot view is out of date. Although you may be able to acquire a lock on the object, it is a *dirty lock* because you cannot use the object and commit, despite holding the lock.

This condition is trapped by the argument to the `ifChanged:` keyword following read lock request message:

```
System readLock: anObject
  ifDenied: [block1]
  ifChanged: [block2].
```

Like its simpler counterpart, this message returns `System` if it acquires a lock on *anObject* without complications. It generates an error if the user has no authorization for acquiring the lock, or selects one of the blocks passed as arguments and executes that block, returning the block's value.

For example, if a conflicting lock is held on *anObject*, this message executes the block given as an argument to the keyword `ifDenied:`. Similarly, if *anObject* has been changed by another session, it executes the argument to `ifChanged:`. The following sections provide some suggestions about the code such blocks might contain. For example:

### Example 9.3

---

```
System readLock: anObject
  ifDenied: []
  ifChanged: [System abortTransaction]
```

---

To minimize your chances of getting dirty locks, lock the objects you need as early in your transaction as possible. If you encounter a dirty lock in the process, you can keep track of the fact and continue locking. After you finish locking, you can abort your transaction to get current values for all of the objects whose locks are dirty. See Example 9.4.

### Example 9.4

---

```
| dirtyBag |
dirtyBag := IdentityBag new.
myEmployees do: [:anEmp |
  System readLock: anEmp
    ifDenied: []
    ifChanged: [ dirtyBag add: anEmp ] ].
dirtyBag isEmpty
  ifTrue: [ ^true ]
  ifFalse: [ System abortTransaction ].
```

---

Your new transaction can then proceed with clean locks.

## Locking Collections of Objects Efficiently

In addition to the locking request messages for single objects, GemStone provides messages to request locks on an entire collection of objects. If the objects you need to lock are already in collections, or if they can be gathered into collections without too much work, it is more efficient to use the collection-locking methods than to lock the objects individually.

The following statements request locks on each of the elements of two different collections:

---

**Example 9.5**

---

```
UserGlobals at: #myArray put: Array new;
  at: #myBag put: IdentityBag new.

System readLockAll: myArray.
System writeLockAll: myBag.
```

---

The messages in Example 9.5 are similar to the simple, single-object locking-request messages (such as `readLock:`) that you've already seen. If a clean lock is acquired on each element of the argument, these messages return `System`. If you lack the proper authorization for any object in the argument, GemStone generates an error and grants no locks.

The difference between these methods and their single-object counterparts is in the handling of other errors. The system does not immediately halt to report an error if an object in the collection is changed, or if a lock must be denied because another session has already locked the object. Instead, the system continues to request locks on the remaining elements, acquiring as many locks as possible. When the method finishes processing the entire collection, it generates an error. In the meantime, however, all locks that you acquired remain in place.

You might want to handle these errors from within your GemStone Smalltalk program instead of letting execution halt. For this purpose, class `System` provides collection-locking methods that pass information about unsuccessful lock requests to blocks that you supply as arguments. For example:

```
System writeLockAll: aCollection ifIncomplete: aBlock
```

The argument *aBlock* that you supply to this method must take three arguments. If locks are not granted on all elements of *aCollection* (for any reason except authorization failure), the method passes three arrays to *aBlock* and then executes the block.

- ▶ The first array contains all elements of *aCollection* for which locks were denied.
- ▶ The second array contains all elements for which dirty locks were granted.
- ▶ The third array is empty, and is there for compatibility with previous versions of GemStone.

You can then take appropriate actions within the block. See Example 9.6.

**Example 9.6**


---

```

classmethod: Dummy
handleDenialOn: deniedObjs
^ deniedObjs
%
classmethod: Dummy
getWriteLocksOn: aCollection
System writeLockAll: aCollection
    ifIncomplete: [:denied :dirty :unused |
        denied isEmpty ifFalse: [self handleDenialOn: denied].
        dirty isEmpty ifFalse: [System abortTransaction] ]
%
System readLockAll: myEmployees
%
Dummy getWriteLocksOn: myEmployees
%
```

---

**Upgrading Locks**

On occasion, you might want to *upgrade* a read lock to a write lock. For example, you might initially intend to read an object, only to discover later that you must also write the object.

However, if you have a read lock on an object, you cannot successfully write that object. If you attempt to do so, your attempt to commit that transaction will fail.

GemStone currently provides no built-in support for upgrading locks. However, to ensure your ability to commit, you can remove the read lock you currently hold on an object and then immediately request a write lock.

It is important to request the upgraded lock immediately, because between the time that the lock is removed, and the time that the upgraded lock is requested, another session has the opportunity to lock the object, or to write it and commit.

**Locking and Indexed Collections**

When indexes are present, locking can fail to prevent conflict. The reasons are similar to those discussed in the section “Indexes and Concurrency Control” on page 172. Briefly, GemStone maintains indexing structures in your snapshot view and does not lock these structures when an indexed collection or one of its elements is locked. Therefore, despite having locked all of the visible objects that you touched, you can be unable to commit.

Specifically, this means that:

- ▶ *if* an object is either an element of an indexed collection, or participates in an index (meaning it is a component of an element bearing an index);
- ▶ *and* another session can access the object, an indexed collection of which the object is a member, or one of its predecessors along the same indexed path;
- ▶ *then* locking the object does not guarantee that you can commit after reading or writing the object.

Therefore, don’t rely on locking an object if the object participates in an index.

## Removing or Releasing Locks

Once you lock an object, its default behavior is to remain locked until you either log out or explicitly remove the lock; unless you specify otherwise, locks persist through aborts and commits. In general, remove a lock on an object when you have used the object, committed the resulting values to the repository, and no longer foresee an immediate need to maintain control of the object.

Class System provides the following messages for removing locks:

`System removeLock: anObject`

Removes any lock you might hold on a single object. If *anObject* is not locked, GemStone does nothing. If another session holds a lock on *anObject*, this message has no effect on the other session's lock.

`System removeLockAll: aCollection`

Removes any locks you might hold on the elements of a collection.

If you intend to continue your session, but the next transaction is to work on a different set of objects, you might wish to remove all the locks held by your session. Class System provides two mechanisms for doing so.

`System commitTransaction; removeLocksForSession`

Attempts to commit the present transaction and removes all locks it holds, even if the commit does not succeed.

`System commitAndReleaseLocks`

Attempts to commit your transaction and release all the locks you hold in a single operation. If your transaction fails to commit, all locks are held instead of released.

## Releasing Locks Upon Aborting or Committing

After you have locked an object, you can add it to either of two special sets. One set contains objects whose locks you wish to release as soon as you commit your current transaction. The other set contains objects whose locks you wish to release as soon as you either commit or abort your current transaction. Executing `continueTransaction` does not release the locks in either set.

The following statement adds a locked object to the set of objects whose locks are to be released upon the next commit:

`System addToCommitReleaseLocksSet: aLockedObject`

The following statement adds a locked object to the set of objects whose locks are to be released upon the next commit or abort:

`System addToCommitOrAbortReleaseLocksSet: aLockedObject`

The following statement adds the locked elements of a collection to the set of objects whose locks are to be released upon the next commit:

`System addAllToCommitReleaseLocksSet: aLockedCollection`

The following statement adds the locked elements of a collection to the set of objects whose locks are to be released upon the next commit or abort:

`System addAllToCommitOrAbortReleaseLocksSet: aLockedCollection`

### NOTE

*If you add an object to one of these sets and then request an updated lock on it, the object is removed from the set.*

You can remove objects from these sets without removing the lock on the object. The following statement removes a locked object from the set of objects whose locks are to be released upon the next commit:

```
System removeFromCommitReleaseLocksSet: aLockedObject
```

The following statement removes a locked object from the set of objects whose locks are to be released upon the next commit or abort:

```
System removeFromCommitOrAbortReleaseLocksSet: aLockedObject
```

The following statement removes the locked elements of a collection from the set of objects whose locks are to be released upon the next commit:

```
System removeAllFromCommitReleaseLocksSet: aLockedCollection
```

The following statement removes the locked elements of a collection from the set of objects whose locks are to be released upon the next commit or abort:

```
System removeAllFromCommitOrAbortReleaseLocksSet: aLockedCollection
```

You can also remove all objects from either of these sets with one message. The following statement removes all objects from the set of objects whose locks are to be released upon the next commit:

```
System clearCommitReleaseLocksSet
```

The following statement removes all objects from the set of objects whose locks are to be released upon the next commit or abort:

```
System clearCommitOrAbortReleaseLocksSet
```

The statement `System commitAndReleaseLocks` also clears both sets if the transaction was successfully committed.

## Inquiring About Locks

GemStone provides messages for inquiring about locks held by your session and other sessions. Most of these messages are intended for diagnostic use, but some may be useful to ordinary applications.

The message `sessionLocks` gives you a complete list of all the locks held by your session. This message returns a three-element array. The first element is an array of read-locked objects; the second is an array of write-locked objects, and the third is "deferred unlocks", objects that have been unlocked, but the request is waiting for another session to release the commit token.

For example, the following code uses this information to remove all write locks held by the current session:

```
System removeLockAll:  
(System sessionLocks at: 2)
```

Other useful messages `systemLocksQuick`, `systemLocks`, `systemLocksReport`, and `systemLocksDetailedReport`, which report locks on all objects held by all sessions currently logged in to the repository. Note that these methods do not report on locks that other sessions are holding on their temporary objects – that is, objects that they have never committed to the repository. These objects are not visible to your session, so they are unlikely to be a cause of commit conflict.

Another lock inquiry message, `lockOwners: anObject`, is useful if you've been unable to acquire a lock because of conflict with another session. This message returns an array of `SmallIntegers` representing the sessions that hold locks on `anObject`. The method in

Example 9.7 uses `lockOwners:` to build an array of the userIDs of all users whose sessions hold locks on a particular object.

### Example 9.7

---

```

classmethod: Dummy
getNamesOfLockOwnersFor: anObject
| userIDArray sessionArray |
sessionArray := System lockOwners: anObject.
userIDArray := Array new.
sessionArray do:
    [:aSessNum | userIDArray add:
        (System userProfileForSession: aSessNum) userId].
^userIDArray
%

Dummy getNamesOfLockOwnersFor:
    (myEmployees detect: {:e | e.name = 'Conan' })
%
```

---

You can test to see whether an object is included in either of the sets of locked objects whose locks are to be released upon the next abort or commit operation. The following statement returns true if *anObject* is included in the set of objects whose locks are to be released upon the next commit:

```
System commitReleaseLocksSetIncludes: anObject
```

The following statement returns true if *anObject* is included in the set of objects whose locks are to be released upon the next commit or abort:

```
System commitOrAbortReleaseLocksSetIncludes: anObject
```

For information about the other lock inquiry messages, see the methods on System class in the image.

## Application Write Locks

Unlike read and write locks, application write locks can only be placed on a single object per lock queue (there are ten lock queues available). The object can be any persistent non-special object; the first time an application lock write is invoked on a lock queue, the object that is locked is registered for that lock queue, and all subsequent uses of that lock queue can only lock this particular object until the next Stone restart.

This allows it to be used as a mutex, or simplifies serializing modifications to a single critical object, such as a collection.

The call to acquire an application write lock also does not return until the lock is acquired, or the lock wait times out. This frees you from having to repeatedly request a lock if it is not immediately available. The timeout is controlled by the configuration parameter `STN_OBJ_LOCK_TIMEOUT`.

To set an application write lock on an object, send the message:

```
System waitForApplicationWriteLock: lockObject
    queue: lockIdx
    autoRelease: aBoolean
```

*lockIdx* must be a SmallInteger between 1 and 10, depending on which lock queue is being used. If *aBoolean* is true, the lock is released automatically on commit or abort, otherwise you must manually remove the lock when you are done.

This method errors if you attempt to lock a temporary object or AllSymbols, otherwise returns an integer code, one of the following:

- 1 - lock granted
- 2071 - undefined lock (*lockIdx* out of range)
- 2074 - dirty; the lock object written by other session since start of this transaction
- 2075 - lock denied (*lockObject* is an invalid object)
- 2418 - lock not granted, deadlock
- 2419 - lock not granted, wait for lock timed out

## 9.4 Classes That Reduce the Chance of Conflict

Often, concurrent access to an object is structural, but not semantic. GemStone detects a conflict when two users access the same object, even when respective changes to the objects do not collide.

For example, when two users both try to add something to a bag they share, GemStone perceives a write-write conflict on the second add operation, although there is really no reason why the two users cannot both add their objects. As human beings, we can see that allowing both operations to succeed leaves the bag in a consistent state, even though both operations modify the bag. A situation such as this causes commit conflicts that could potentially be avoided.

GemStone provides a number of reduced-conflict classes that you can use instead of their regular counterparts in applications that might otherwise experience too many unnecessary conflicts. Using these classes allows a greater number of transactions to commit successfully, but “reduced conflict” does not mean “no conflict.” For example, while two users should be able to add different objects to a shared collection, the code can’t be expected to resolve the problem of two users attempting to remove the same object.

When a conflict does occur - for example, two users attempting to remove the same object - this is a normal conflict. The second user will see a commit failure with a transaction conflict. When the commit fails, the user loses all changes made to the Rc object during the current transaction, and the persistent state remains in the state left by the earlier user who made the conflicting changes.

Reduced-conflict classes are not always appropriate; some of them require more storage, and may require maintenance under some usage conditions, or may cause commits to take longer to complete under some usage conditions.

### RcCounter

The class RcCounter can be used instead of a simple number in order to keep track of the amount of something. It allows multiple users to increment or decrement the amount at the same time without experiencing conflicts.

The class RcCounter is not a kind of number. It encapsulates a number – the counter – but it also incorporates other intelligence; you cannot use an RcCounter to replace a number anywhere in your application. It only increments and decrements a counter.

For example, imagine an application to keep track of the number of items in a warehouse bin. Workers increment the counter when they add items to the bin, and decrement the counter when they remove items to be shipped. This warehouse is busy; if each concurrent increment or decrement operation produced a conflict, work slows unacceptably.

Furthermore, the conflicts are mostly unnecessary. Most of the workers can tolerate a certain amount of inaccuracy in the value of the bin count at any time. They do not need to know the exact number of items in the bin at every moment; they may not even worry if the bin count goes slightly negative from time to time. They may simply trust that the data is not completely up-to-date, and that their fellow workers have added to the bin in the time since the data was last refreshed. For such an application, an RcCounter is helpful.

Instances of RcCounter understand messages such as `increment`, `decrement`, and `value`. For additional protocol, see the image.

For example, assuming that `binCount` refers to an instance of RcCounter, the following operations can take place concurrently from different sessions without causing a conflict:

### Example 9.8

---

```
!session 1
binCount incrementBy: 36.
System commitTransaction.
%
!session 2
binCount incrementBy: 24.
System commitTransaction.
%
!session 3
binCount decrementBy: 48
    ifLessThan: 0
        thenExecute: ['^Not enough widgets to ship today.'].
System commitTransaction.
%
```

---

This can result in some variable behavior, depending on the timing of the operations.

For example, if the starting `binCount` is 0, and these operations happen concurrently, then session 3 will not perform the decrement, and the final `binCount` will be 60.

However, if session 1 and 2 have committed their increment operations, and session 3 updated its snapshot view prior to executing the code, then session 3 will perform the decrement and the final `binCount` will be 12.

## Reduced-Conflict Collection Classes

GemStone provides a variety of reduced-conflict collection classes:

- ▶ RcArray, providing Array-like semantics
- ▶ RcKeyValueDictionary, producing dictionary-like semantics
- ▶ RcIdentityBag, RcLowMaintenanceIdentityBag, and RcIdentitySet, providing IdentityBag- and IdentitySet-like semantics.
- ▶ GsPipe, RcPipe, and RcQueue, providing queue semantics

In addition to varying collection semantics, individual classes have specific types of conflicts they are designed to avoid, and the amount of internal infrastructure or the cost of resolving a conflict varies. Selection of an RC class should consider the demands of the application and also the costs of the automatic conflict resolution.

RcArray, GsPipe, RcPipe, RcKeyValueDictionary, RcIdentitySet and RcLowMaintenanceIdentityBag provide reduced-conflict by automatic replay; when performing specific supported operations, if conflict occurs, the changes can be replayed, slowing down the commit by the second session but allowing the commit to occur.

In cases where there are likely to be many concurrent updates, there is a risk of developing a backlog of sessions replaying the operations; application in which a high degree of concurrent operations are expected may benefit by using an RcQueue or RcIdentityBag. RcIdentityBag and RcQueue provide add and remove sets for each session. This avoids the risk of conflict between sessions at the cost of additional time required to access elements, and some use patterns may require periodic manual cleanup.

## RcArray

The class RcArray provides much of the same functionality as Array. However, no conflict occurs on instances of RcArray with:

- ▶ Multiple producers: any number of users are adding objects to the array.

If a conflict with other update operations on the RcArray occur, the add is replayed so that the commit can succeed. Only methods that add elements at the end of the RcArray support concurrent updates. During conflict resolution, commit order determines the order of the elements in the RcArray.

Because implementation relies on the replay of the adds when there are conflicts, high levels of concurrency have a risk of creating a backlog, when a convoy of sessions are all trying to commit their additions to the RcArray. For applications with expected high rates of concurrency, consider using an RcQueue to accumulate the additions, and have a single gem process remove elements from the RcQueue, and put them in an RcArray.

## RcIdentityBag

The class RcIdentityBag provides much of the same functionality as IdentityBag. No conflict occurs on instances of RcIdentityBag with:

- ▶ Multiple producers: any number of users are adding objects .
- ▶ Limited multiple consumers: one user removes an object from the bag, or multiple users remove objects but only one tried to remove the last or only occurrence of an object.

When multiple sessions remove different occurrences of the same object, it may take a little longer to commit the second transaction.

RcIdentityBag uses per-session add and remove subcollections to avoid conflict. Each session adds to its individual subcollection, and removals of these items are tracked in a parallel bag.

If you create an index on an RcIdentityBag, you should also create a reduced-conflict index, otherwise the underlying index structure may have a conflict. However, even an indexed instance of RcIdentityBag reduces the possibility of a transaction conflict, compared to an instance of IdentityBag, indexed or not.

## RcLowMaintenanceIdentityBag and RcIdentitySet

The class `RcLowMaintenanceIdentityBag` and `RcIdentitySet` provide much of the same functionality as `IdentityBag` and `IdentitySet`. No conflict occurs on instances of `RcLowMaintenanceIdentityBag` or `RcIdentitySet` with:

- ▶ Multiple producers: any number of users are adding objects.
- ▶ Single consumer: one user removes objects.

`RcLowMaintenanceIdentityBag` and `RcIdentitySet` avoid conflict by performing a selective abort and replay of adds. If more than one user removes objects, they are likely to experience a commit failure with a transaction conflict. Instances of `RcLowMaintenanceIdentityBag` and `RcIdentitySet` may have indexes on their contents. It is recommended to create a reduced-conflict index.

## RcKeyValueDictionary

The class `RcKeyValueDictionary` provides the same functionality as `KeyValueDictionary`. No conflict occurs on instances of `RcKeyValueDictionary` with:

- ▶ Limited multiple producers: any number of users add keys and values to the dictionary, as long as the keys do not already exist in the dictionary.
- ▶ Limited multiple consumers: any number of users remove keys from the dictionary, as long as only one user removes the same key at a time.

`RcKeyValueDictionary` avoids conflict by performing a selective abort and replay of the modifications to the dictionary. A session that would otherwise have a commit failure due to a transaction conflict may take slightly longer to complete the commit.

## GsPipe

The class `GsPipe` implements a first-in-first-out queue with a single producer and a single consumer. No conflict occurs on instances of `GsPipe` with:

- ▶ Single producer: only one user at a time adds objects to the pipe.
- ▶ Single consumer: only one user at a time removes an object from the pipe.

`GsPipe` avoids commit conflict between adds and removes by the nature of its implementation, since modifying the head or tail of a linked list doesn't cause conflict.

## RcPipe

The class `RcPipe` implements a first-in-first-out queue with multiple producers and a single consumer. No conflict occurs on instances of `RcPipe` with:

- ▶ Multiple producers: any number of users are adding objects to the `RcPipe`.
- ▶ Single consumer: only one user at a time removes an object from the `RcPipe`.

`RcPipe` avoids conflict by performing a selective abort and replay of adds to the pipe. If more than one user removes objects from the pipe, they are likely to experience a commit failure with a transaction conflict. When the commit fails, the user loses all changes made to the pipe during the current transaction, and the pipe remains in the state left by the earlier user who made the conflicting changes.

## RcQueue

The class RcQueue approximates the functionality of a first-in-first-out queue. RcQueues are multiple-producer, single-consumer. No conflict occurs with:

- ▶ Multiple producers: any number of users are adding objects to the queue.
- ▶ Single consumer: only one user at a time removes an object from the queue.

RcQueue uses per-session add and remove subcollections to avoid conflict. Each session's modifications are only to its own add and remove subcollections; the RcQueue calculates the next element based on the contents of the individual session subcollections.

RcQueue approximates a first-in-first-out queue, but it cannot implement such functionality exactly because of the nature of repository snapshot views during transactions.

An object added to an RcQueue is ordered in the queue according to the time it is added to the RcQueue, but it only becomes visible to other sessions when the session commits. If objects in the RcQueue are "consumed" as soon as they appear, then it is possible for more recently created elements to be consumed before older ones that were not yet committed.

For example, suppose one user adds object A at 10:20, but waits to commit until 10:50. Meanwhile, another user adds object B at 10:35 and commits immediately. A third user viewing the queue at 10:30 will see neither object A nor B. At 10:35, object B will become visible to the third user. At 10:50, object A will also become visible to the third user, and will furthermore appear earlier in the queue, because it was created first.



# Object Security and Authorization

---

This chapter explains how to set up object security policies to restrict read and write access to application objects.

It covers:

**How GemStone Security Works** (page 191)

describes the Gemstone object security model.

**Assigning Objects to Security Policies** (page 195)

summarizes the messages for reporting your current security policy, changing your current policy, and assigning a policy to simple and complex objects.

**Application Example** (page 202) and **Development Example** (page 206)

provides examples for defining and implementing object security for your projects.

**Privileged Protocol for Class GsObjectSecurityPolicy** (page 214)

defines the system privileges for creating or changing security policy authorization.

## 10.1 How GemStone Security Works

GemStone provides security at several levels:

- ▶ Login authorization keeps unauthorized users from gaining access to the repository;
- ▶ Privileges limit ability to execute special methods affecting the basic functioning of the system (for example, the methods that reclaim storage space); and
- ▶ Object level security allows individual users, specific groups of users, and all users to have read, write, or no access to each object in the repository.
- ▶ Object Filtering provides further object level security, by allowing you to control the transmission of each object to a remote cache. This features is only available with X509 logins, which include additional security for Gems running on insecure remote nodes. See the *GemStone/S 64 Bit X509-Secured GemStone System Administration Guide* for more information.

## Login Authorization

You log into GemStone through any of the interfaces provided: GemBuilder for Smalltalk, GemBuilder for Java, Topaz, or the C interface (see the appropriate interface manual for details). Whichever interface you use, GemStone requires the presentation of a *user ID* (a name or some other identifying string) and a password. If the user ID and password pair match the user ID and password pair of someone authorized to use the system, GemStone permits interaction to proceed; if not, GemStone severs the logical connection.

The GemStone system administrator, or someone with equivalent privileges (see below), establishes your user ID and (depending on the login authentication used) your password, when he or she creates your *UserProfile*. The GemStone system administrator can also configure a GemStone system to monitor failures to log in, and to note the attempts in the Stone log file after a certain number of failures have occurred within a specified period of time. A system can also be configured to disable a user account after a certain number of failed attempts to log into the system through that account. See the *GemStone System Administration Guide* for details.

### The UserProfile

Each instance of *UserProfile* is created by the system administrator. The *UserProfile* is stored with a set of all other *UserProfiles* in a set called *AllUsers*. The *UserProfile* contains:

- ▶ Your UserID and Password.
- ▶ The *SymbolList* used for resolving symbols when compiling, including *SymbolDictionaries* such as *Globals* and *UserGlobals*. Chapter 3, “Resolving Names and Sharing Objects”, discusses these topics.
- ▶ The groups to which you belong
- ▶ The privileges you may have.
- ▶ A default *GsObjectSecurityPolicy* to assign your session at login, or nil.

See the *System Administration Guide* for instructions on creating *UserProfiles*, defining groups, and assigning users to groups.

## System Privileges

Actions that affect the entire GemStone system are tightly controlled by *privileges* to use methods or access instances of the *System*, *UserProfile*, *GsObjectSecurityPolicy*, and *Repository* classes, and to modify code. Privileges are given to individual *UserProfile* accounts to access various parts of GemStone or perform important functions such as storage reclamation.

The privileged messages for the *System*, *UserProfile*, *GsObjectSecurityPolicy* and *Repository* Classes are described in the image, and their use is discussed in the *System Administration Guide*.

## Object-level Security

GemStone object-level security allows you to:

- ▶ abstractly group objects;
- ▶ specify who owns the objects;

- ▶ specify who can read them; and
- ▶ specify who can write them.

Each site designs a custom scheme for its data security. Objects can be secured for selective read or write access by a group or individual users. Objects can also be left unsecured, so any user can read or modify them. Not restricting access will improve performance for sites with fewer security requirements.

The GemStone class `GsObjectSecurityPolicy` facilitates this security.

## GsObjectSecurityPolicy

Each object's header includes a 16-bit unsigned security policy Id that specifies the `GsObjectSecurityPolicy` to which the object has been assigned. (In previous releases, object security policies were known as Segments; references to Segment now mean `GsObjectSecurityPolicy`).

All objects assigned the same security policy have exactly the same protection. That is, if you can read or write one object assigned to a certain policy, you can read or write them all.

There are several ways that access to objects is controlled by the security policy:

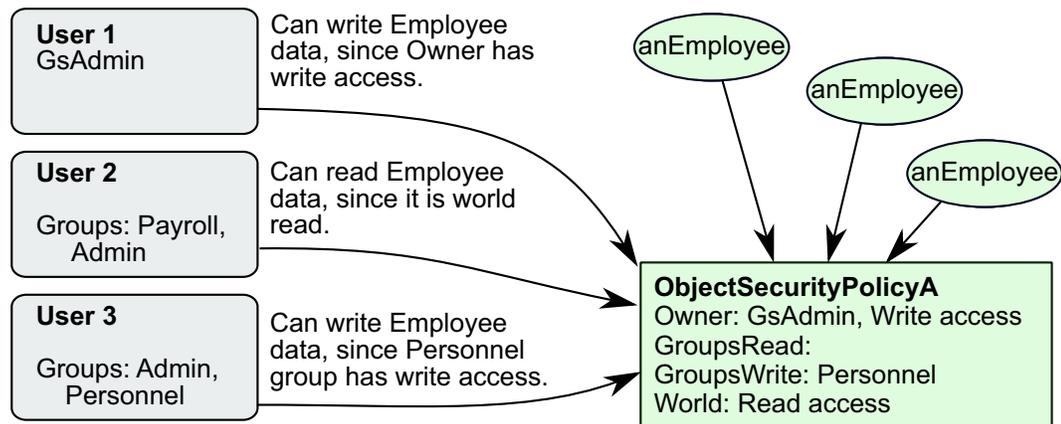
- ▶ Each policy is owned by a single user, and all objects assigned the same security policy have the same owner. The owner has write and read access to all objects associated with the security policy.
- ▶ A security policy has a setting for world; this allows every authorized GemStone user to have read, write, or no access to all the objects associated with a security policy
- ▶ Groups of users can be defined, and these groups can be configured to have read, write, or no access to all the objects associated with a security policy.

In addition, an object may also have no security policy, in which case its security policy Id is zero. This means that there are no restrictions on access to this object; any logged-in user can read and write this object.

Whenever an application tries to access an object, GemStone compares the object's authorization attributes in the security policy associated with the object with those of the user whose application is attempting access. If the user is appropriately authorized, the operation proceeds. If not, GemStone returns an error notification.

The user's group membership and security policy authorization control access to objects, as shown by Figure 10.1.

Figure 10.1 User Access to Application ObjectSecurityPolicyA



Three users access this application:

- ▶ The System Administrator, **GsAdmin**, owns `ObjectSecurityPolicyA` and can read and write the objects assigned with it.
- ▶ **User3** belongs to the Personnel group, which authorizes read and write access to `ObjectSecurityPolicyA`'s objects.
- ▶ **User2** doesn't belong to a group that can access `ObjectSecurityPolicyA`, but can still read those objects, because `ObjectSecurityPolicyA` gives read authorization to all GemStone users.

Because security policies are objects, access to a `GsObjectSecurityPolicy` object is controlled by the security policy it is assigned to, exactly like access to any other object.

`GsObjectSecurityPolicy` instances are usually assigned to the `DataCuratorObjectSecurityPolicy`. The authorization information stored in the `GsObjectSecurityPolicy` instance, which controls access to the objects assigned with that security policy, does not control access to the policy object itself.

Objects do not "belong" to an security policy. It is more correct to say that objects are associated with a security policy. Although objects know which policy they are assigned to, security policies do not know which objects are assigned to them. Security policies are not meant to organize objects for easy listing and retrieval. For those purposes, you must turn to symbol lists, which are described in Chapter 3, "Resolving Names and Sharing Objects".

## 10.2 Assigning Objects to Security Policies

For security policy authorizations to have any effect, you must assign some objects to the security policies whose authorizations you have set up.

### Default Security Policy and Current Security Policy

In your UserProfile, you may be assigned a *default* security policy, or this may be left empty. When you login to GemStone, your Session uses this default security policy as your current security policy. Any objects you create are assigned to your current security policy; if you do not have a current security policy, the new objects do not have a security policy, and so have world read and write access.

Class UserProfile has the message `defaultObjectSecurityPolicy`, which returns your default GsObjectSecurityPolicy (or nil). Sending the message `currentObjectSecurityPolicy:` to System changes your current security policy:

#### Example 10.1

---

```
| aPolicy myPolicy |
myPolicy := System myUserProfile
    defaultObjectSecurityPolicy.
aPolicy := GsObjectSecurityPolicy new.
System commitTransaction.
"change my current security policy to aPolicy"
System currentObjectSecurityPolicy: aPolicy
```

---

Only committed instances of GsObjectSecurityPolicy can be used.

If you commit after changing the security policy, the new GsObjectSecurityPolicy remains your current security policy until you change the security policy again or log out. If you abort after changing your current security policy, your current security policy is reset from your UserProfile's default security policy.

Unnamed GsObjectSecurityPolicies are often stored in a UserProfile, but named GsObjectSecurityPolicies are stored in symbol dictionaries like other named objects. Private security policies are typically kept in a user's UserGlobals dictionary; security policies for groups of users are typically kept in a shared dictionary.

#### Example 10.2

---

```
| myPolicy |
"get default security policy"
myPolicy := System myUserProfile defaultObjectSecurityPolicy.
"compare with current"
myPolicy = System currentObjectSecurityPolicy
%
true
```

---

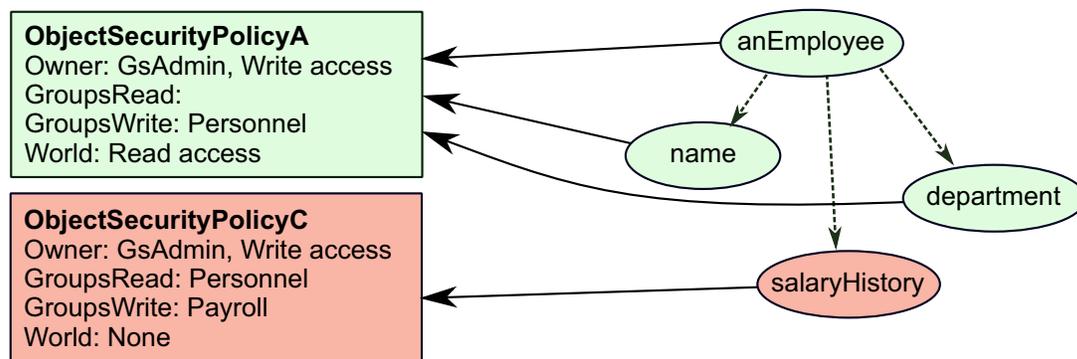
## Objects and Security Policies

GemStone object security is defined for objects. Your security scheme must be defined to protect sensitive data in separate objects, either by itself or as a member object of a customer class. Since each object has separate authorization, each object must be assigned separately.

### Compound Objects

Usually, the objects you are working with are compound, and each part is an object in its own right, with its own security policy assignment. For example, look at anEmployee in Figure 10.2. The contents of its instance variables (name, salary, and department) are separate objects that can be assigned to different security policies. Salary is assigned to ObjectSecurityPolicyC, which enforces more restricted access than ObjectSecurityPolicyA.

**Figure 10.2 Multiple Security Policy Assignments for a Compound Object**



### Collections

When you assign collections of objects to security policies, you must distinguish the container from the items it contains. Each of the items must also be assigned to the proper policy. Distinguishing between a collection and the objects it contains allows you to create collections most elements of which are publicly accessible, while some elements are sensitive.

## Configuring Authorization for an Object Security Policy

Object security policies store authorization information that defines what a particular user or group member can do to the objects with that policy. Three levels of authorization are provided:

- ▶ **write** — A user can read and modify any of the objects with that security policy and create new objects associated with the policy.
- ▶ **read** — A user can read any of the objects with that security policy, but cannot modify (write) them or add new ones.
- ▶ **none** — A user can neither read nor write any of the objects with that security policy.

By assigning a security policy to an object, you give the object the access information associated with that policy. Thus, all objects with a security policy have exactly the same

protection; that is, if you can read or write one object with to a certain policy, you can read or write them all.

Controlling authorizations at the security policy level rather than storing the information in each object makes them easy to change. Instead of modifying a number of objects individually, you just modify one security policy object. This also keeps the repository smaller, eliminating the need for duplicate information in each of the objects.

## How GemStone Responds to Unauthorized Access

GemStone immediately detects an attempt to read or write without authorization and responds by stopping the current method and issuing an error. When you successfully commit your transaction, GemStone verifies that you are still authorized to write in your current security policy. If you are no longer authorized to do so, GemStone issues an error, and your default security policy once again becomes your current security policy. If you are no longer authorized to write in your default security policy, GemStone terminates your session, and you are unable to log back in to GemStone. If this happens, see your system administrator for assistance.

## Owner, Group, and World Authorization

A `GsObjectSecurityPolicy` controls what access a user has to associated objects. Access can be separately assigned for:

- ▶ a security policy's *owner*
- ▶ *groups* of users (by name)
- ▶ the *world* of all GemStone users

Whenever a program tries to read or write an object, GemStone compares the object's authorization attributes with those of the user who is attempting to do the reading or writing. If the user has authorization to perform the operation, it proceeds. If not, GemStone returns an error notification.

These categories overlap. The owner of a security policy is also in the world of all GemStone users, and may also be in one or more groups that have other access authorization. When determining a user's authorization, the most permissive or generous authorization will be allowed and other, more restrictive authorizations, will be ignored. Thus, if world authorization is `#read`, but the user is a member of a group with `#write` authorization, then the world authorization will be ignored.

## Owner Authorization

Each `GsObjectSecurityPolicy` has an owner. The owner of a policy may be assigned read, write, or no access in the security policy, and therefore to the objects associated with this security policy. Usually, the owner of a policy has write authorization, but this isn't required (unless this is the default security policy for that user). Users may own more than one security policy.

The message `GsObjectSecurityPolicy >> ownerAuthorization:` *anAuthorizationSymbol* is used to set and clear authorization for the owner of the security policy. The message `GsObjectSecurityPolicy >> ownerAuthorization` returns the authorization for the owner of the security policy.

## Group Authorization

Groups are an efficient way to ensure that a number of GemStone users all will share the same level of access to objects in the repository, and all will be able to manipulate certain objects in the same ways.

Groups are typically organized as categories of users who have common interests or needs; for example, Payroll or Personnel.

The global collection AllGroups, a collection of instances of UserProfileGroup, defines all groups in the system. Membership in a group is granted either by adding the UserProfile to the UserProfileGroup, by adding the group to the user's UserProfile groups.

The message GsObjectSecurityPolicy>>authorizationForGroup: *group* returns the rights for users in that group.

The message GsObjectSecurityPolicy>>groupsWithAuthorization: *anAuthSymbol* returns the names of groups that have a particular level of access (#read, #write, or #none) for the receiver security policy.

To set group access, use the message GsObjectSecurityPolicy>>group: *groupNameString* authorization: *anAuthSymbol*. For example, to set the group authorization as shown in Example 10.3, use the following:

```
anObjectSecurityPolicy group: 'Managers' authorization: #read
```

## World Authorization

In addition to storing authorization for its owner and for groups, a security policy can also be told to authorize or to deny access by all GemStone users (*the world*.)

The message GsObjectSecurityPolicy>>worldAuthorization returns the rights for all users. A corresponding message,

GsObjectSecurityPolicy>>worldAuthorization: *anAuthSymbol*, sets the authorization for all GemStone users. For example:

```
anObjectSecurityPolicy worldAuthorization: #none
```

## Predefined GsObjectSecurityPolicies

The initial GemStone repository has eight GsObjectSecurityPolicies, with the following Ids:

### 1. SystemObjectSecurityPolicy

This security policy is defined in the Globals dictionary, and is owned by the SystemUser. All GemStone users, represented by world access, are authorized to read, but not write, objects associated with this security policy. The group #System is authorized to write to objects in this policy.

### 2. DataCuratorObjectSecurityPolicy

This security policy is defined in the Globals dictionary, and is owned by the DataCurator. All GemStone users, represented by world access, are authorized to read, but not write, objects associated with this security policy. The group #DataCuratorGroup is authorized to write in this security policy.

Objects in the DataCuratorObjectSecurityPolicy include the Globals dictionary, the SystemRepository object, all GsObjectSecurityPolicy objects, AllUsers (the set of all

GemStone UserProfiles), AllGroups (the collection of groups authorized to read and write objects in GemStone security policies), and each UserProfile object.

*NOTE:*

*When GemStone is installed, only the DataCurator is authorized to write in this security policy. To protect the objects in DataCuratorObjectSecurityPolicy from unauthorized modification, only administrative users should have write access.*

**3. GsTimeZoneObjectSecurityPolicy**

The initial repository does not use this Id. Repositories that have been converted from earlier GemStone/S server products use this for the **GsTimeZoneObjectSecurityPolicy**.

**4. GsIndexingObjectSecurityPolicy**

This security policy is used by the indexing subsystem.

**5. SecurityDataObjectSecurityPolicy**

This security policy is used by the system for passwords for UserProfiles, and other highly protected information.

**6. PublishedObjectSecurityPolicy**

This security policy is used for objects in the Published symbol dictionary.

**7. GcUserObjectSecurityPolicy**, the GsObjectSecurityPolicy for the GcUser user.

**8. NamelessObjectSecurityPolicy**, the GsObjectSecurityPolicy for the Nameless user.

**9. CodeLibrarianUserObjectSecurityPolicy**, the GsObjectSecurityPolicy for the CodeLibrarianUser.

**10. HostAgentUserObjectSecurityPolicy**, the GsObjectSecurityPolicy for the HostAgentUser.

**11. ObjectFiltersObjectSecurityPolicy**

The GsObjectSecurityPolicy that contains ObjectFilters, which are used for further object protection with X509 logins.

For repositories that have been converted from certain earlier versions, there may also be GsObjectSecurityPolicy with id 20, with world write.

## **GsObjectSecurityPolicy names**

Most of the system security policies have Symbolic references in Globals and can the name can be used in code. SecurityPolicies may have names (which are stored in a dynamic instance variable); most system security policies have names set. When defining a new object security policy, you may send name: to make it more easily identifiable. This does not itself create a global reference.

## **Changing the Security Policy for an Object**

If you have the authorization, you can change the accessibility of an individual object by assigning a different security policy to it.

The message Object >> objectSecurityPolicy returns the security policy that protects that receiver, or nil if the receiver does not have an associated security policy:

**Example 10.3**


---

```
UserGlobals objectSecurityPolicy
DataCuratorObjectSecurityPolicy(#2 in Repository SystemRepository,
Owner DataCurator write, Group DataCuratorGroup write, World read)
```

---

The message `Object >> objectSecurityPolicy: anObjectSecurityPolicy` assigns *anObjectSecurityPolicy* as the security policy for the receiver. You also use this method to remove the security policy, so the receiver object has world read and write access. You must have write authorization for both security policies, that of the receiver and the argument. Assuming the necessary authorization, this example assigns a new security policy to class `Employee`:

```
Employee objectSecurityPolicy: aPolicy.
```

You may override the method `objectSecurityPolicy:` for your own classes, especially if they have several components.

For objects having several components, such as collections, you may assign all the component objects to a specified security policy when you reassign the composite object. You can implement the message `objectSecurityPolicy:` to perform these multiple operations. Within the method `objectSecurityPolicy:` for your composite class, send the message `assignToObjectSecurityPolicy:` to the receiver and each object of which it is composed.

For example, an `objectSecurityPolicy:` method for the class `Menagerie` might appear as shown in Example 10.4. The object itself is assigned to another security policy using the method `assignToObjectSecurityPolicy:`. Its component objects, the animals themselves, have internal structure (names, habitats, and so on), and therefore call `Animal's objectSecurityPolicy:` method, which in its turn sends the message `assignToObjectSecurityPolicy:` to each component of an `Animal`, ensuring that each animal is properly and completely reassigned to the new security policy.

**Example 10.4**


---

```
Array subclass: 'Menagerie'
  instVarNames: #()
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals
%
method: Menagerie
objectSecurityPolicy: aPolicy
  self assignToObjectSecurityPolicy: aPolicy.
  self do: [:eachAnimal |
    eachAnimal objectSecurityPolicy: aPolicy]
%
```

---

Special objects — `SmallInteger`, `SmallDouble`, `Character`, `Boolean`, and `nil` — are assigned the `SystemObjectSecurityPolicy` and cannot be assigned another security policy.

## Security Policy Ownership

Each `GsObjectSecurityPolicy` has an owner—by default, the user who created it. An security policy's owner always has control over who can access the security policy's objects. As a security policy's owner, you can alter your own access rights at any time, even forbidding yourself to read or write objects with that security policy.

You might not be the owner of your default security policy. To find out who owns a security policy, send it the message `owner`. The receiver returns the owner's `UserProfile`, which you may read, if you have the authorization:

### Example 10.5

```
"Return the userId of the owner of the default security policy for
the current Session."
| aUserProf myDefaultPolicy |
"get default security policy"
myDefaultPolicy := System myUserProfile
    defaultObjectSecurityPolicy.
myDefaultPolicy ifNotNil:
    ["return its owner's UserProfile"
     aUserProf := myDefaultPolicy owner.
     "request the userId"
     aUserProf userId]
%
user1
```

Every security policy understands the message `owner`: `aUserProfile`. This message assigns ownership of the receiver to the person associated with `aUserProfile`. The following expression, for example, assigns the ownership of your default security policy to the user associated with `aUserProfile`:

```
System myUserProfile defaultObjectSecurityPolicy owner:
aUserProfile
```

In order to reassign ownership of a security policy, you must have write authorization for the `DataCuratorObjectSecurityPolicy`. Because of the way separate authorizations for owners, groups and world combine, changing access rights for the any one of them may or may not alter a particular user's rights to a security policy.

#### CAUTION

*Do not, under any circumstances, attempt to change the authorization of the `SystemObjectSecurityPolicy`.*

## Revoking Your Own Authorization: a Side Effect

You may occasionally want to create objects and then take away authorization for modifying them.

#### CAUTION

*Do not remove your write authorization for your default security policy or your current security policy. If you lose write authorization for your default security policy, you will not be able to log in again.*

## Finding Out Which Objects Are Protected by a Security Policy

It may be useful for you to be able to find all the objects that are protected by a particular security policy. An expression of the form:

```
SystemRepository listObjectsInObjectSecurityPolicies: anArray
```

takes as its argument an array of security policy IDs, and returns an array of arrays. Each inner array contains all objects whose security policy ID is equal to the corresponding security policy ID element in the argument *anArray*. Instances to which you lack read authorization are omitted without notification.

Note that this method aborts the current transaction and scans the object header of each object in the repository.

If the result set is very large, there is a risk of out of memory errors. To avoid the need to have the entire result set in memory, the following methods are provided:

```
Repository >> listObjectsInObjectSecurityPolicyToHiddenSet:
anObjectSecurityPolicyId
```

This method puts the set of all objects in the specified security policy in the ListInstancesResult hidden set. (a hidden set is an internal memory structure that, while not an object, is treated as one).

To enumerate the hidden set, you can use this method:

```
System >> hiddenSetEnumerate: hiddenSetId limit: maxElements
```

using a *hiddenSetId* of 1, which is the number of the "ListInstancesResult" hidden set in GemStone/S 64 Bit v3.7. This hidden set number is subject to change in new releases; to determine which hidden sets are in a particular release, use the GemStone Smalltalk method `System Class >> HiddenSetSpecifiers`. For more on hidden sets, see "Other Optimization Hints" on page 328.

You can also list objects that are protected by a particular security policies to an external binary file, which can later be read into a hidden set. To do this, use the method:

```
Repository >> listObjectsInObjectSecurityPolicies: anArray
toDirectory: aString
```

This method scans the repository for the instances protected by the security policies in *anArray* and writes the results to binary bitmap files in the directory specified by *aString*. Binary bitmap files have an extension of `.bm` and may be loaded into hidden sets using class methods in System.

Bitmap files are named:

```
objectSecurityPolicy<ObjectSecurityPolicyId>-objects.bm
```

where *ObjectSecurityPolicyId* is the security policy ID.

The result is an Array of pairs. For each element of the argument *anArray*, the result array contains *ObjectSecurityPolicyId*, *numberOfInstances*. The *numberOfInstances* is the total number written to the output bitmap file.

## 10.3 Application Example

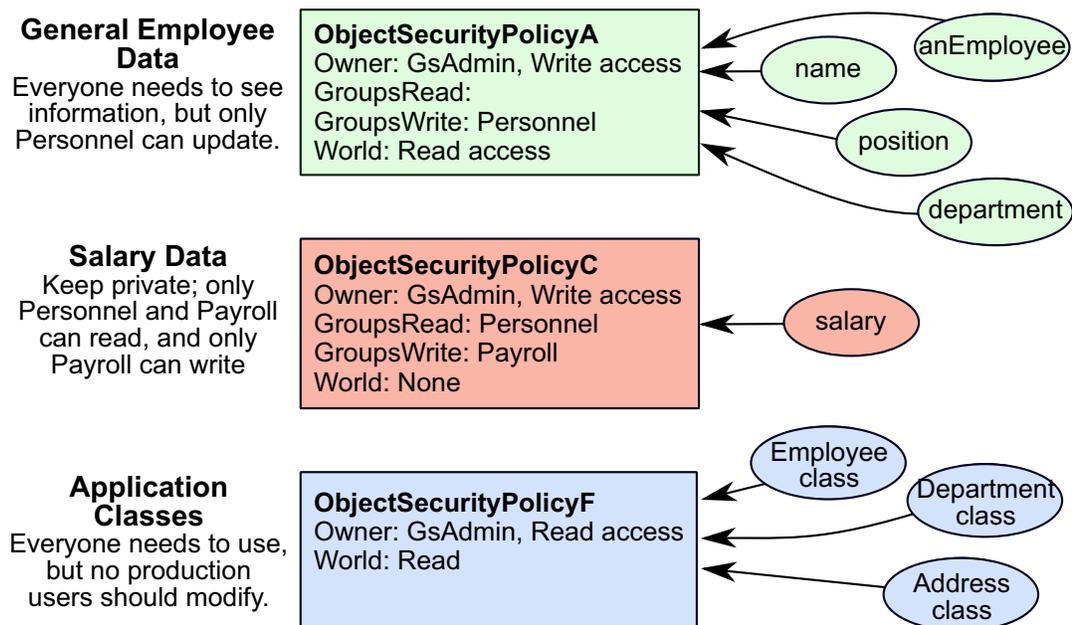
The structure of the user community determines how your data is stored and accessed. Regardless of their job titles, users generally fall into three categories:

- ▶ *Developers* define classes and methods.
- ▶ *Updaters* create and modify instances.
- ▶ *Reporters* read and output information.

When you have a group of users working with the same GemStone application, you need to ensure that everyone has access to the objects that should be shared, such as the application classes, but you probably want to limit access to certain data objects.

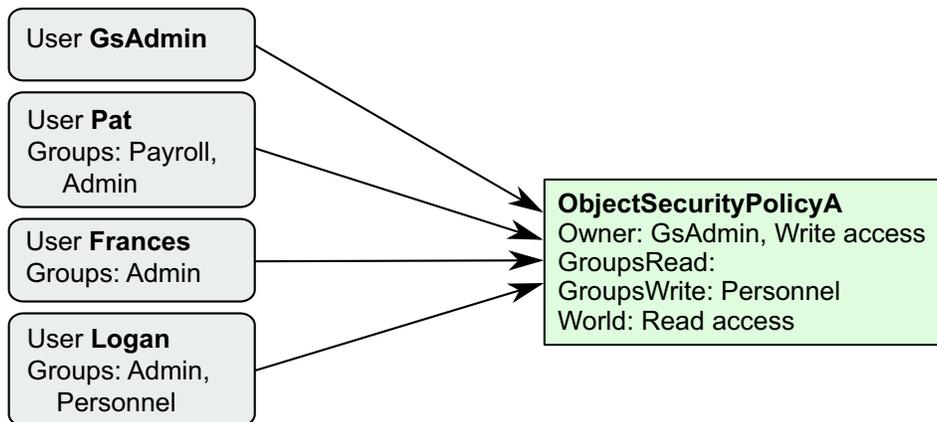
Figure 10.3 shows a typical production situation. In this example, all the application users need access to the data, but different users need to read some objects and write others. So most data goes into ObjectSecurityPolicyA, which anyone can look at, but only the Personnel group or owner can change. ObjectSecurityPolicyC is set up for sensitive salary data, which only the Payroll group or owner can change, and only they and the Personnel group can see. You don't want anyone to accidentally corrupt the application classes, so they go into ObjectSecurityPolicyF, which no one can change.

**Figure 10.3 Application Objects Assigned with Three Security Policies**



Given a set of users with different roles in the application, Figure 10.4 and Figure 10.5 indicate how group membership and security policy authorization control access to application objects:

Figure 10.4 User Access to Application ObjectSecurityPolicyA



Four users access this application:

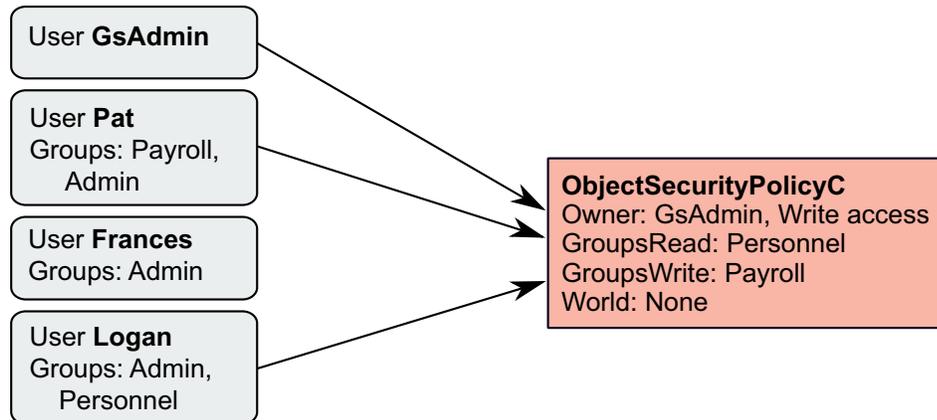
- ▶ The System Administrator, **GsAdmin**, owns both security policies and can read and write the objects assigned with them.
- ▶ **Logan** belongs to the Personnel group, which authorizes her to read and write objects associated with ObjectSecurityPolicyA, and read objects associated with ObjectSecurityPolicyC.
- ▶ **Pat** can read and write the objects assigned with ObjectSecurityPolicyC, because he belongs to the Payroll group. He doesn't belong to a group that can access ObjectSecurityPolicyA, but he can still read those objects, because ObjectSecurityPolicyA gives read authorization to all GemStone users.
- ▶ **Frances** does not belong to a group that can access either security policy. She can read the objects assigned with ObjectSecurityPolicyA, because it allows read access to all GemStone users. She has no access at all to ObjectSecurityPolicyC.

Pat and Logan are sometimes updaters and sometimes reporters, depending on the type of data. Frances is strictly a reporter.

---

**Figure 10.5 User Access to Application ObjectSecurityPolicyC**

---



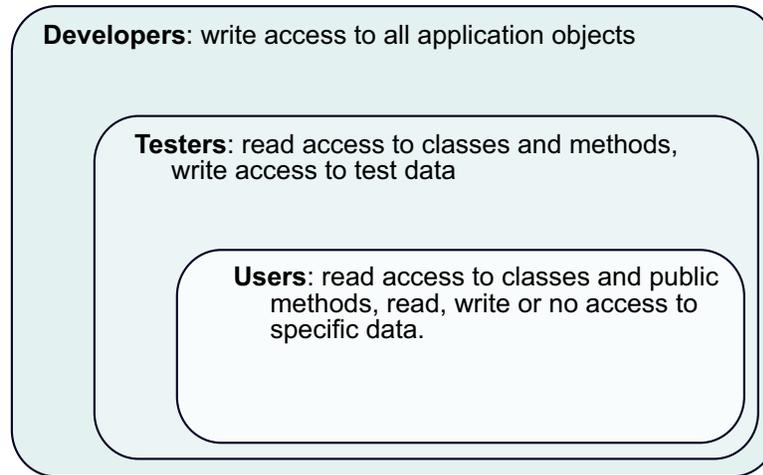
ObjectSecurityPolicyB is associated with the classes and methods for the application. These are world read, so all users can read these objects. No one, not even GsAdmin, can modify the classes.

## 10.4 Development Example

Up to now, this discussion has been limited to applications in a production environment, but issues of access and security arise at each step of application development. During the design phase you need to consider the security policies needed for the application life cycle: development, testing, and production.

The access required at each stage is a subset of the preceding one, as shown in Figure 10.6.

Figure 10.6 Access Requirements During an Application's Life Cycle



## 10.5 Planning Security Policies for User Access

As you design your application, decide what kind of access different end users will need for each object.

### Protecting the Application Classes

All the application users need read access to the application classes and methods, so they can execute the methods. To prevent accidental damage to them, however, you probably want to limit write access. The CodeModification privilege is required to create or modify classes and methods. You can further limit write access using security policies. You may even want to change the owner's authorization to read, until changes are required.

Like other objects, classes and their methods are assigned to security policies on an object-by-object basis. You may keep separate subsections of your application in different security policies, with different write authorizations, if you want.

### CodeModification privilege

All application developers will need to have CodeModification privilege. This is in addition to the ability to read and write the appropriate security policies. Without CodeModification privilege, you cannot compile methods or classes, add new methods,

add a Class to a SymbolDictionary, or perform other operations required for application development.

Application users, on the other hand, should not have CodeModification privilege, since they will not be modifying methods or classes. This allows you to protect the application code for inadvertent (or intentional) damage or modification, even if you do not want to implement object level security.

## Planning Authorization for Data Objects

Authorization for data objects means protecting the instances of the application's classes, which will be created by end users to store their data. You can begin the planning process by creating a matrix of users and their required access to objects. Table 10.1 shows part of such a matrix, which maps out access to instances of the class Employee and some of its instance variables.

Security is easier to implement if it is built into the application design at the beginning, not added later. In the following sections, planning for the third stage, end user access, comes first. Following the planning discussion comes the implementation instructions, which explain how to set up security policies for the developers, extend the access to the testers, and finally move the application into production.

Remember that in effect you have four options, shown on the matrix as:

**W** – need to write (also allows reading)

**R** – need to read, must not write

**N** – must not read or write

**blank** – don't need access, but it won't hurt

**Table 10.1 Access for Application Objects Required by Users**

Objects	Users						
	System Admin.	Human Resource	Employee Records	Payroll	Mktg	Sales	Customer Support
anEmployee	W	W	W	R	R	R	R
name	W	W	W	R	R	R	R
position	W	W	W	R		R	
dept.	W	W	W	R		R	
manager	W	W	W	R		R	
dateHired	W	W	W	R	N	R	N
salary	W	R	R	W	N	N	N
salesQuarter	W	R	R	R	N	W	N
salesYear	W	R	R	R	N	W	N
vacationDays	W	W	W	N	N	N	N
sickDays	W	W	W	N	N	N	N

### World Access

To begin analyzing your access requirements, check whether the objects have any Ns. For objects that do, world authorization must be set to none.

If you have people who need read access to nonsensitive information, give world read authorization to those objects. In this example, world can have read access to anEmployee, name, position, dept., and manager. The objects can still be protected from casual browsing by storing them in a dictionary that does not appear in everyone’s symbol list. This does not absolutely prevent someone from finding an object, but it makes it difficult. For more information, see Chapter 3, “Resolving Names and Sharing Objects”.

### Owner

By default, the owner has write access to the objects protected by a security policy. To choose an owner, look for a user who needs to modify everything. In terms of the basic user categories described earlier, the owner could be either an administrator or an updater. This depends on the type of objects that will be assigned to the security policy.

In Table 10.1 the system administrator is the user who needs write access. So the system administrator is made the owner, with full control of all the objects. The DataCurator and SystemUser logins are available to the system administrator. The DataCurator is not automatically authorized to read and write all objects, however. Like any other user account, it must be explicitly authorized to access objects in security policies it does not own. Although the SystemUser can read and write all objects, it should not be used for these purposes.

### Planning Groups

The rest of the access requirements must be satisfied by setting up groups. The thing to remember about groups is that they do not reflect the organization chart; they reflect differences in access requirements. Because the number of possible authorization combinations is limited, the number of groups required is also limited.

First look at the existing access to anEmployee, name, position, dept., and manager, as shown in Table 10.2. By making the system administrator the owner with write authorization and assigning read authorization to world, you have already satisfied the needs of five departments.

**Table 10.2 Access to the First Five Objects Through Owner and World Authorization**

Objects	Users						
	System Admin.	Human Resource	Employ. Records	Payroll	Mktg	Sales	Customer Support
Employee	W	W	W	R		R	
name	W	W	W	R		R	
position	W	W	W	R		R	
dept.	W	W	W	R		R	
manager	W	W	W	R		R	

write access as owner or no access as world

You still need to provide authorization for the Human Resources and Employee Records departments. In every case, they need the same access (see Table 10.1) so you only have to create one group for the two departments. This group, named Personnel, requires write authorization for the objects in Table 10.2.

Now look at the existing access to the rest of the objects. These objects store more sensitive information, so access requirements of different users are more varied. Assigning write authorization to owner and none to world has completely satisfied the needs of three departments, as shown in Table 10.3.

**Table 10.3 Access to the Last Six Objects Through Owner and World Authorization**

Objects	Users						
	System Admin.	Human Resource	Employ. Records	Payroll	Mktg	Sales	Customer Support
dateHired	W	W	W	R	N	R	N
salary	W	R	R	W	N	N	N
salesQuarter	W	R	R	R	N	W	N
salesYear	W	R	R	R	N	W	N
vacationDays	W	W	W	N	N	N	N
sickDays	W	W	W	N	N	N	N

write access as owner or no access as world

Two more departments, Human Resources and Employee Records, are already set up to access as the Personnel group. As shown in Table 10.4, this group needs write authorization to dateHired, vacationDays, and sickDays, which they must be able to read and modify. They need read authorization to salary, salesQuarter, and salesYear, which they must read but cannot modify.

**Table 10.4 Access to the Last Six Objects Through the Personnel Group**

Objects	Users						
	System Admin.	Human Resource	Employ. Records	Payroll	Mktg	Sales	Customer Support
dateHired	W	W	W	R	N	R	N
salary	W	R	R	W	N	N	N
salesQuarter	W	R	R	R	N	W	N
salesYear	W	R	R	R	N	W	N
vacationDays	W	W	W	N	N	N	N
sickDays	W	W	W	N	N	N	N

read or write access as Personnel group

Now the Payroll and Sales departments still require access to the objects, as shown in Table 10.3. Because these departments' needs don't match anyone else's, they must each have a separate group.

**Table 10.5 Access to the Last Six Objects Through the Payroll and Sales Groups**

Objects	Users						
	System Admin.	Human Resource	Employ. Records	Payroll	Mktg	Sales	Customer Support
dateHired	W	W	W	R	N	R	N
salary	W	R	R	W	N	N	N
salesQuarter	W	R	R	R	N	W	N
salesYear	W	R	R	R	N	W	N
vacationDays	W	W	W	N	N	N	N
sickDays	W	W	W	N	N	N	N

 read or write access as Payroll or Sales group

In all, this example only requires three groups: Personnel, Payroll, and Sales, even though it involves seven departments.

### Planning Security Policies

When you have been through this exercise with all your application's prospective objects and users, you are ready to plan the security policies. For easiest maintenance, use the smallest number of security policies that your required combinations of owner, group, and world authorizations allow. You don't need different security policies with duplicate functionality to separate particular objects, like the application classes and data objects. Remember that symbol lists, not security policies, are used to organize objects for listing and retrieval.

In this example you need six security policies, as shown in Figure 10.7. Notice that each one has different authorization.

### Developing the Application

During application development you implement two separate schemes for object organization: one for sharing application objects by the development team and one controlling access by the end users. In addition, you may need to allow access for the testers, who may need different access to objects.

Once you have planned the security policies and authorizations you want for your project, you can refer to procedures in the *System Administration Guide* for implementing that plan.

### Setting Up Security Policies for Joint Development

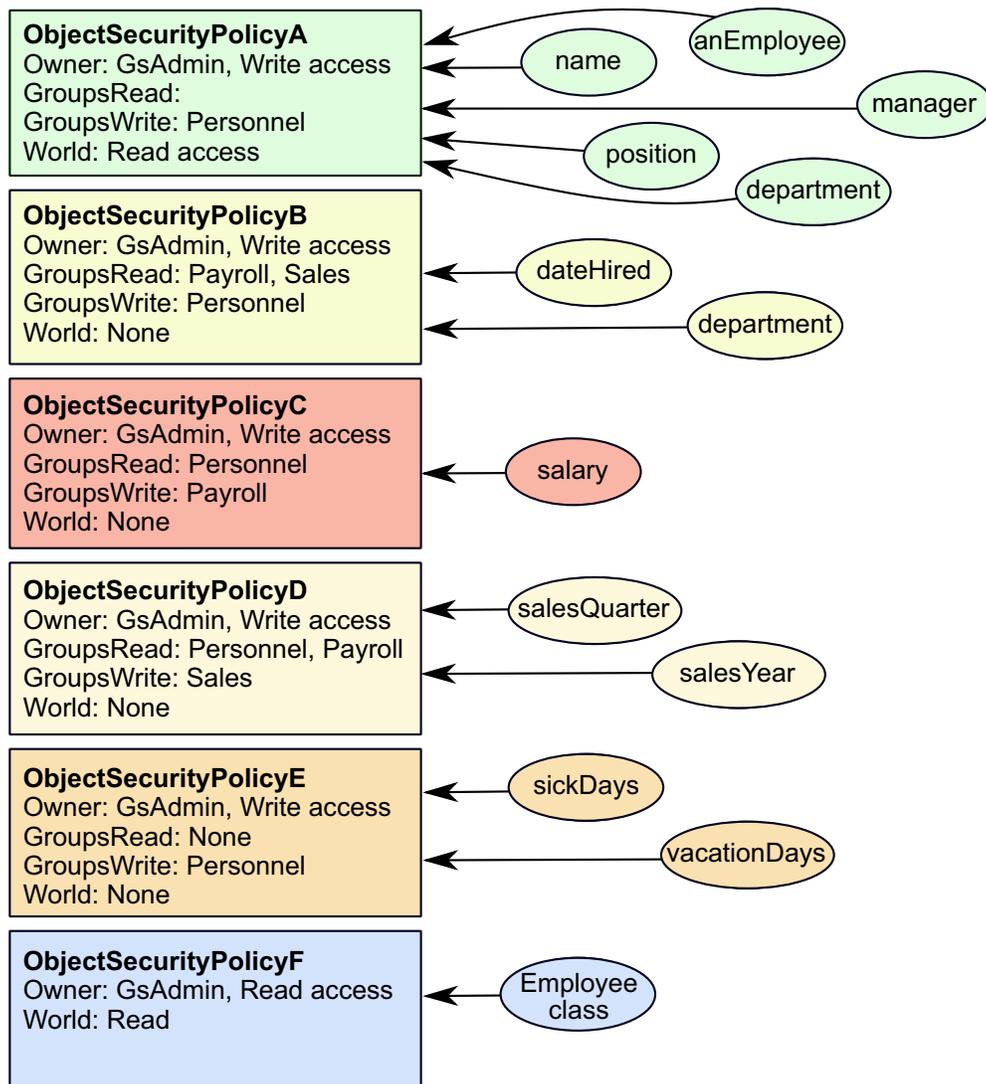
To make joint development possible, you need to set up authorization and references so that all the developers have access to the classes and methods that are being created. Create a new symbol dictionary for the application and put it in everyone's symbol list; make sure it includes references to any shared security policies. If only developers are using the

repository, you can give world access to shared objects, but if other people are using the repository, you must set up a group for developers.

You can organize security policy assignments in various ways:

- ▶ **Full access to all personal security policies.** Give all the developers their own default security policies to work in. Give everyone in the team write access to all the security policies. Because the objects you create are typically assigned to your default security policy, this method may be the simplest way to organize shared work.
- ▶ **Read access to all personal security policies.** Set up the same as above, except give everyone read access to the security policies. If each developer is doing a separate module, read access may be enough. Then everyone can use other people's classes, but not change them. This has the advantage of enforcing the line between application and data.
- ▶ **Full access to a shared security policy.** Give all developers the same default security policy, writable by everyone. This is an easy, informal way to share objects.
- ▶ **Full access to a shared security policy plus private security policies.** Developers work in their own default security policies and reassign their objects to the shared security policy when they are finished. This lets you share a collection, for example, but keep the existing elements private, so that other developers could add elements but not modify the elements you have already created. To share a collection this way, assign the collection object itself to the accessible security policy. The collection has references to many other objects, which can be associated with other security policies. Everyone has the references, but they get errors if they try to access objects with non-readable security policies. You might also choose to share an application symbol dictionary, so that other developers can put objects in it, without making the objects themselves public.

Figure 10.7 Security Policies Required for User Access to Application Objects



### Making the Application Accessible for Testing

Testers need to be able to alternate between two distinct levels of access:

- ▶ **Full access.** As members of the development team, they need read access to all the classes and methods in the application, including the private methods. Testers also need write access to their test data.
- ▶ **User-level access.** They need a way to duplicate the user environment, or more likely several environments created for different user groups.

This can be done by setting up a tester group and one or more sample user groups during the development phase. For testing the user environment, the application must already be set up for multi-user production use, as explained in the following section.

## Moving the Application into a Production Environment

When you have created the application, it is time to set it up for a multi-user environment. A GemStone application is developed in the repository, so all you have to do to install an application is to give other users access to it. This means implementing the rest of your application design, in roughly the reverse order of the planning exercise. To give other users authorization to use the objects in the application:

1. Create the security policies.
2. Create the necessary user groups specified in up-front development, if they don't exist.
3. Assign the required owner, world, and group authorizations to the security policies.
4. Assign testers to the user groups and complete multi-user testing.
5. Assign any end users that need group authorization to the user groups.
6. Assign the application's objects to the security policies you created.

You also have to give users a reference to the application so they can find it. An application dictionary is usually created with references to the application objects, including its security policies. A reference to this dictionary usually must appear in the users' symbol lists. For more information on the use of symbol dictionaries, see the discussion of symbol resolution and object sharing in Chapter 3, "Resolving Names and Sharing Objects".

## Security Policy Assignment for User-created Objects

Because security policy assignment is on an object-by-object basis, it is important to know how objects are assigned. When the objects are being created by end users of an application, as in this example, you may want to partially or fully automate the process of security policy assignment. Depending on the needs of the local site, you can implement various mechanisms to ensure data security, prevent accidental damage to existing data, or simply avoid misplaced data.

### Assign a Specified Security Policy to the User Account

Set up users with the proper security policy by default. This is a simple way to assure that someone who creates objects in a single security policy doesn't misplace them. To make it impossible to change security policies, rather than just unlikely, you also have to close write access for group and world to all the other security policies.

This solution would work for the Sales and Payroll groups in the example (Figure 10.7). They need read access to several security policies, but they only write in one.

The drawback of this solution is that the user can only use one security policy.

### Develop the Application to Create the Data Objects

Your best choice is to create objects in the correct security policy, using the `GsObjectSecurityPolicy>>setCurrentWhile:` method. With this method, the application stores data objects in the proper security policies. This provides the most protection. Besides guaranteeing that the objects end up in the proper security policy, this prevents users from accidentally modifying objects they have created. It also prevents

them from reading the data that other users enter, even when everyone is creating instances of the same classes.

## 10.6 Privileged Protocol for Class GsObjectSecurityPolicy

Privileges stand apart from the security policy and authorization mechanism. *Privileges* are associated with certain operations: they are a means of stating that, ordinarily, only the DataCurator or SystemUser is to perform these privileged operations. The DataCurator can assign privileges to other users at his or her discretion, and then those users can also perform the operations specified by the particular privilege.

### NOTE

*Privileges are more powerful than security policy authorization. Although the owner of a security policy can always use read/write authorization protocol to restrict access to objects protected by a security policy, the DataCurator can override that protection by sending privileged messages to change the authorization scheme.*

The following message to GsObjectSecurityPolicy always requires special privileges:

```
new (class method)
newInRepository: (class method)
```

You can always send the following messages to the security policies you own, but you must have special privileges to send them to other security policies:

```
group:authorization:
ownerAuthorization:
worldAuthorization:
```

For changing privileges, UserProfile defines two messages that also work in terms of the privilege categories described above. The message addPrivilege: *aPrivString* takes a number of strings as its argument, including the following:

```
'DefaultObjectSecurityPolicy'
'ObjectSecurityPolicyCreation'
'ObjectSecurityPolicyProtection'
```

For a full list of privileges, see the *System Administration Guide* chapter on User Management.

To add security policy creation privileges to your UserProfile, for example, you might do this:

```
System myUserProfile addPrivilege: 'ObjectSecurityPolicyCreation'.
```

This gives you the ability to execute GsObjectSecurityPolicy new.

A similar message, privileges:, takes an array of privilege description strings as its argument. The following example adds privileges for security policy creation and password changes:

```
System myUserProfile privileges:
#('ObjectSecurityPolicyCreation' 'UserPassword')
```

To withdraw a privilege, send the message `deletePrivilege: aPrivString`. As in preceding examples, the argument is a string naming one of the privilege categories. For example:

```
System myUserProfile deletePrivilege:  
    'ObjectSecurityPolicyCreation'
```

Because UserProfile privilege information is typically protected by a security policy that only the data curator can modify, you might not be able to change privileges yourself. You must have write authorization to the DataCuratorObjectSecurityPolicy, or be a member of DataCuratorGroup, in order to do so.

For direction and information about configuring user accounts, adding user accounts and assigning security policies to those accounts, and checking authorization for user accounts, see the *System Administration Guide*.



# Class versions and Instance Migration

---

Although you designed your schema with care and thought, after using it for a while you will probably find a few things you would like to improve. Furthermore, even if your design was perfect, real-world changes usually require changes to the schema sooner or later.

This chapter discusses the mechanisms GemStone Smalltalk provides to allow you to make changes in your schema and manage the migration of existing objects to the new schema.

**Versions of Classes** (page 217)

defines the concept of a class version and describes two different approaches you can take to specify one class as a version of another.

**ClassHistory** (page 219)

describes the GemStone Smalltalk class that encapsulates the notion of class versioning.

**Migrating Objects** (page 221)

explains how to migrate either certain instances, or all of them, from one version of a class to another while retaining the data that these instances hold, including transforming the data as needed.

**Multi-threaded instance migration** (page 232)

describes how to perform simple migrations more quickly by creating mappings and performing the migration in a single operation.

## 11.1 Versions of Classes

In order to create instances of a class, the class must be invariant, and invariant classes cannot be modified, except in some specific ways. While you defined your schema to be as complete as you could at the time you created the classes, inevitably further changes are needed. You may now have instances of invariant classes populating your database and a need to modify your schema by redefining certain of these classes.

To support this schema modification, GemStone allows you to define different versions of classes. Every class in GemStone has a class history – an object that maintains a list of all versions of the class – and every class is listed in at least one class history, the class history for the class itself. You can define as many different versions of a class as required, and declare that the different versions belong to the same class history. You can migrate some or all instances of one version of a class to another version when you need to. The values of the instance variables of the migrating instances are retained if you have defined the new version to do so.

## Defining a New Version

In GemStone Smalltalk classes have *versions*. Each version is a unique and independent class object, but the versions are related to each other through a common class history. The classes need not share a similar structure, nor even a similar implementation. The classes need not even share a name, although it is less confusing if they do, or if you establish and adhere to some naming convention.

If you define a new class in a SymbolDictionary that already contains an existing class with the same name, it automatically becomes a new version of the previously existing class. This is the most common way of creating new class versions. Instances that predate the creation of the new version remain unchanged, and continue to access the old class's methods, although tools such as GemBuilder may provide options to automatically migrate instances to the new class. Instances created after the redefinition have the new class's structure and access to the new class's methods.

When you define a class, the class creation protocol includes an option to specify the existing class of which the new class is a version. See the keyword `newVersionOf:`.

## New Versions and Subclasses

When you create a new version of a class – for example, `Animal` – subclasses of the old version of `Animal` still point to the old version of `Animal` as their superclass (unless you are using a tool which provides the option to automatically version and recompile subclasses). If you wish these classes to become subclasses of the new version, you need to recompile the subclass definitions to make new versions of the subclasses, specifying the new version of `Animal` as their superclass.

One way to do this is to file in the subclasses of `Animal` after making the new version of `Animal` (assuming the new version of the superclass has the same name).

## New Versions and References in Methods

When you create a new version of a class (such as `Animal`) you typically want your existing code to use the new version rather than the old version. That is, without being recompiled, existing methods containing code like the following should create an instance of the new version rather than of the old version of `Animal` class:

```
pet := Animal new.
```

As long as the new class version replaces an existing class in the same SymbolDictionary, then references from existing methods will be automatically updated to the new class version.

This works because a compiled method does not directly reference a global (e.g., the class `Animal`), but references a `SymbolAssociation` in a `SymbolDictionary`. When you originally compile the method, it resolves the name using an expression similar to the following:

```
System myUserProfile resolveSymbol: #theClassName
```

The compiled method includes the resulting `SymbolAssociation`, whose key is the name of the global and whose value is the class (or other object). The value can be updated at any time, for example when you create a new version of a class.

This tiny performance penalty is what allows global variables to vary. If you have a global that you know will be constant, then you can reference the value directly from a compiled method by making the `SymbolAssociation` invariant before compiling the method.

While the `SymbolAssociation` is updated with the new value by versioning the class within the same `SymbolDictionary`, keep in mind that under some circumstances you may have a `SymbolAssociation` that does not reference the latest version, or the version you expect. If you have a newer class with the same name in a different `SymbolDictionary`, or if you delete and recreate the class, the `SymbolAssociation` will continue to point to the older class.

## Class Variables and Class Instance Variables

Adding a Class Variable does not require a new version of your class, but adding a class instance variable does.

When you create a new version of a class, the values in any Class variables or Class Instances variables in the old class are referenced by the new class as well. By default, all versions of a class refer to the same objects referenced from Class or Class instance variables.

## Class versioning and Class options

When you define a class with the same name and variables, but with a different set of options (passed in via the `options:` keyword to the class creation method), it does not always need to create a new version of the class.

See page 37 for a description of the class options and how they can be used for specific class behavior.

Class options are automatically inherited by new version or by an updated version, unless the first element in `options:` array is `#noInheritOptions`.

If the new definition includes `#instancesNonPersistent`, then the existing class will be modified to add `#instancesNonPersistent`.

If the new definition includes `#subclassesDisallowed`, `#disallowGciStore`, `#traverseByCallback`, `#dbTransient`, or `#instancesInvariant`, and if `#noInheritOptions` is not the first element, the options will not be removed.

To remove an unwanted option, you may need to include `#noInheritOptions` as the first element of the `options:` array.

## 11.2 ClassHistory

In GemStone Smalltalk, every class has a class history, represented by the system as an instance of the class `ClassHistory`. A class history is an array of classes that are meant to be

different versions of each other. While they often have the same class name, this is not a requirement; you can rename classes as well as change their structure.

## Defining a Class as a new version of an existing Class

When you define a new class in the same symbol dictionary as an existing class with the same name, it is by default created as the latest version of the existing class and shares its class history.

When you define a new class by a name that is new to a symbol dictionary, the class is by default created with a unique class history. If you use a class creation message that includes the keyword `newVersionOf:`, you can specify an existing class whose history you wish the new class to share. This is useful if you want to create a version of a class with a different name or in a different symbol dictionary. If the new class version has the same name and is in the same symbol dictionary, it is not necessary to use `newVersionOf:`, since the new class will become a version of the existing class automatically.

For example, suppose your existing class `Animal` was defined like this:

### Example 11.1

---

```
Object subclass: 'Animal'  
  instVarNames: #('habitat' 'name' 'favoriteFood' 'predator')  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: {}  
  inDictionary: UserGlobals.  
  
Animal compileMissingAccessingMethods.
```

---

Example 11.2 creates a class named `NewAnimal` and specifies that the class shares the class history used by the existing class `Animal`.

### Example 11.2

---

```
Object subclass: 'NewAnimal'  
  instVarNames: #('name' 'diet' 'predator' 'species')  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: {}  
  inDictionary: UserGlobals  
  newVersionOf: Animal  
  description: nil  
  options: #().  
  
NewAnimal compileMissingAccessingMethods.
```

---

If you wish to define a new class `Animal` with its own unique class history – in other words, the new class `Animal` is **not** a version of the old class `Animal` – you can add it to a different

symbol dictionary, and specify the argument *nil* to the keyword `newVersionOf:`. However, this can easily create confusion, and make it difficult to diagnose problems. It is only recommended when the two symbol dictionaries will not normally both be loaded and in use at the same time.

If you try to define a new class with the same name as an existing class that you did not create, you will most likely get an error, because you are trying to modify the class history of that class — an object which you are probably not permitted to modify.

## Accessing a Class History

You can access the class history of a given class by sending the message `classHistory` to the class. For example, the following expression returns the class history of the class `Employee`:

```
Animal classHistory
```

This is a collection that includes all older versions of the class, in order, with the most recent version being the last one on the collection. For example:

```
Animal classHistory last == NewAnimal
true
Animal classHistory first == Animal
true
```

In the usual case, in which all classes in the class history have the same name, the compiler will resolve that name as the last one in the classHistory.

In some cases (such as in `GemBuilder` for Smalltalk) a class with a class history size that is larger than one these will be displayed with the position, for example:

```
Animal [1]
NewAnimal [2]
```

## Assigning to a Class History

You can assign a class to a class history by sending the message `addNewVersion:` to the class whose class history you wish to amend; the argument to this message is the class whose history is to be reassigned.

For example, suppose that we created `NewAnimal` using the regular class creation protocol, and did not use the method with the keyword `newVersionOf:`. To later specify that it is a new version of `Animal`, execute the following expression:

```
Animal addNewVersion: NewAnimal
```

## 11.3 Migrating Objects

Once you have defined a new version of your class, you may want to migrate your existing instances from the old class version to the new version. Migration in GemStone Smalltalk is a flexible, configurable operation.

- ▶ Instances of any class can migrate to any other. The two classes need not be similarly named, or, indeed, have much else in common, although it will usually make more sense if they represent the same conceptual object.

- ▶ Not all instances of a class need to migrate at the same time; you can migrate a single instance, a specific sets of instances, or all instances of the original class. Other instances need never migrate. However, instances that are versions of these older classes will not understand new methods that are added to the new version of the class, and may require special consideration during some GemStone upgrades, if you manage code using fileout, or if you use passivation.
- ▶ During migration, the value of instance variables in the new instance are initialized from values in the old instance. The mapping from the old instance variables to the new instance variables is under your control, and can be arbitrarily complex. A default mapping mechanism is provided for simple migrations, such as when adding new instance variables.

## Migration Destinations

You can establish the default destination class for migration, for use in later instance migration operations. To do so, send a message of the form:

```
OldClass migrateTo: NewClass
```

This configures the old class so it knows to migrate its instances to become instances of the new class. Migration does **not** occur as a result of sending the above message; this only sets the destination of migration.

It is not necessary to set a migration destination ahead of time; other protocol will allow you to specify the migration class for a specific instance migration. If you use methods that includes a specific migration destination class, the default destination is ignored.

Once you have set the migration destination, you can migrate a single specific instance, using a message of the form:

```
anInstanceOfOldClass migrate
```

Provided the object is an instance of a class for which a migration destination has been defined, the object becomes an instance of the new class. If no destination has been defined, no change occurs.

The following series of expressions, for example, creates a new instance of `Animal`, sets `Animal`'s migration destination to be `NewAnimal`, and then causes the new instance of `Animal` to become an instance of `NewAnimal`.

### Example 11.3

---

```
| aLemming |  
aLemming := Animal new.  
Animal migrateTo: NewAnimal.  
aLemming migrate.
```

---

Other instances of `Animal` remain unchanged, until such time as they receive the message to migrate.

## Bypassing the Migration Destination

You can bypass the migration destination, or migrate instances of classes for which no migration destination has been specified. To do so, specify the destination directly in the message that performs the migration.

*OldClass* migrateInstances: setOfInstancesOfOldClass to: NewClass  
 Migrate the specific instances in *setOfInstancesOfOldClass*, which are instances of *OldClass*, to become instances of *NewClass*. Ignores any existing migration destination, and does not set the default migration destination.

*OldClass* migrateInstancesTo: NewClass  
 Migrate all instances of *OldClass* to become instances of *NewClass*. Ignores any existing migration destination, and does not set the default migration destination.

Example 11.4 uses `migrateInstances:to:` to migrate all instances of all versions of a class, except the latest version, to the latest version.

For more on listing instances, see “Finding Instances” on page 227.

#### Example 11.4

```
| animalHist allAnimals |
animalHist := Animal classHistory.
allAnimals := SystemRepository listInstances: animalHist.
1 to: animalHist size-1 do: [:index |
  (animalHist at: index)
  migrateInstances:(allAnimals at: index)
  to: Animal currentVersion].
```

When you migrate a set of instances, it is possible that not all instances in the set can be successfully migrated.

The migration methods `migrateInstancesTo:` and `migrateInstances:to:` return an array of collections:

1. The set of objects for which the current user does not have read permission.
2. The set of objects for which the current user has read permission but not write permission.
3. The set of objects that could not be migrated due to index incompatibilities.
4. The set of objects whose class was not identical to the receiver (presumably, incorrectly gathered instances) and therefore were not migrated; this will always be empty if using `migrateInstancesTo:`.
5. The set of object that failed migration by signalling a `MigrationError`. This error may be signalled, for example, by customized migration code that encounters an object that cannot be migrated.

If all these collections are empty, all requested migrations have occurred.

## Migrating Instances that Participate in an Index

If an instance participates in an index (for example, because it is part of the path on which that index was created), then the indexing structure can, under certain circumstances, cause migration to fail. Three scenarios are possible:

- Migration succeeds. In this case, the indexing structure you have made remains intact. Commit your transaction.

- ▶ GemStone examines the structures of the existing version of the class and the version to which you are trying to migrate, and determines that migration is incompatible with the indexing structure. In this case, GemStone raises an error notifying you of the problem, and migration does not occur.

You can commit your transaction, if you have done other meaningful work since you last committed, and then follow these steps:

1. Remove the index in which the instance participates.
  2. Migrate the instance.
  3. Modify the indexing code as appropriate for the new class version and re-create the index.
  4. Commit the transaction.
- ▶ In the final case, GemStone fails to determine that migration is incompatible with the indexing structure, and so migration occurs and the indexing structure is corrupted. In this case, GemStone raises an error notifying you of the problem, and you will not be permitted to commit the transaction. Abort the transaction and then follow the steps explained above.

## Default Instance Variable Mappings

The purpose of migration is to retain the data contained in the old object, while updating the object to a possibly entirely new structure.

In most cases, it makes sense for the object contents at particular instance variable names to have the same values as before the migration, and this is the default behavior.

There are cases in which you might want to change the use of a particular instance variable name, and relocate the data to a new instance variable slot. This is described under “Customized Instance Variable Mappings” on page 225.

The default migration behavior:

- If the new class has an instance variable with the same name as the old class, data at that instance variable in the instance of the old class is moved to the instance variable in the migrated instance.
- If the new class has an instance variable that was not present in the old class, the migrated instance has that instance variable set to nil
- If the new class does not have an instance variable that was present in the old class, the value at that instance variable is dereferenced. The data it represents is no longer accessible from this object
- Values of dynamic instance variables in the instance of the old class remain as dynamic instance variables in the migrated instance.

Suppose, for example, you create two instances of class `Animal` and initialize their instance variables as shown in Example 11.5.

---

**Example 11.5**

---

```
| aLemming |
aLemming := Animal new
    name: 'Leopold';
    favoriteFood: 'grass';
    habitat: 'tundra';
    predator: 'owl';
    yourself.
UserGlobals at: #aLemming put: aLemming.
```

---

You then decide that class `Animal` needs an additional instance variable, `predator`. You create the class called `NewAnimal` (as described in Example 11.2, with four instance variables: `name`, `favoriteFood`, `habitat`, and `predator`). You then migrate `aLemming`.

Example 11.6 performs the migration, and then shows the results of printing the values of the instance variables.

---

**Example 11.6**

---

```
Animal migrateInstances: { aLemming } to: NewAnimal.

aLemming name.
Leopold

aLemming predator.
owl

aLemming diet.
nil
```

---

## Customized Instance Variable Mappings

To initialize an instance variable with the value of a variable that has a different name, you must provide an explicit mapping from the instance variable names of the older class to the instance variable names of the migration destination.

This can be done by overriding the implementation of the method `migrateFrom:instVarMap:` in the destination class.

When you migrate an object, `migrateFrom:instVarMap:` is sent to the new instance with the argument of the old instance. The `instVarMap:` argument is a mapping structure that can be further customized, but is usually left at the default.

In our example, the class `Animal` has the instance variables: `habitat`, `name`, `favoriteFood`, and `predator`, and `NewAnimal` has variables: `name`, `diet`, `predator`, and `species`.

When instances of `Animal` migrate to `NewAnimal`, the value of `diet` ought to be initialized with the value presently held in `favoriteFood`.

Also, we will keep `habitat`, by making it a dynamic instance variable on the new instance.

To accomplish this, `NewAnimal` implements `migrateFrom:instVarMap:`, as in Example 11.7.

#### Example 11.7 `NewAnimal` class >> `instVarMappingTo:`

---

```
method NewAnimal
migrateFrom: anOldInstance instVarMap: aMap
    super migrateFrom: anOldInstance instVarMap: aMap.
    self diet: anOldInstance favoriteFood.
    self dynamicInstVarAt: #habitat put: anOldInstance habitat.
%

Animal migrationDestination: NewAnimal.
aLemming migrate
%
a NewAnimal
  name          Leopold
  diet          grass
  predator      owl
  species       nil
  t1 habitat
  t2 tundra
```

---

## Transforming Variable Values

Another kind of customization is required when the format of data changes. For example, suppose that you have a class named `Point`, which defines two instance variables  $x$  and  $y$ . These instance variables define the position of the point in Cartesian two-dimensional coordinate space.

Suppose that you define a class named `NewPoint` to use polar coordinates. The class has two instance variables named *radius* and *angle*. Obviously the default mapping strategy is not going to be helpful here; migrating an instance of `Point` to become an instance of `NewPoint` loses its data (the actual position) completely. Nor is it correct to map  $x$  to *radius* and  $y$  to *angle*. Instead, what is needed is a method that implements the appropriate trigonometric function to transform the point to its appropriate position in polar coordinate space.

In this case, the method to override is `migrateFrom:instVarMap:`, which you implement as an instance method of the class `NewPoint`. Then, when you request an instance of `Point` to migrate to an instance of `NewPoint`, the migration code that calls `migrateFrom:instVarMap:` executes the method in `NewPoint` instead of in `Object`.

**Example 11.8 Point >> migrateFrom:instVarMap:**

```

Object subclass: #OldPoint
  instVarNames: #(x y )
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals.

OldPoint compileMissingAccessingMethods.

Object subclass: #Point
  instVarNames: #(radius angle )
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals.

Point compileMissingAccessingMethods.

method: Point
migrateFrom: oldPoint instVarMap: aMap
  | x y |
  super migrateFrom: oldPoint instVarMap: aMap.
  x := oldPoint x.
  y := oldPoint y.
  radius := ((x*x) + (y*y)) asFloat sqrt.
  angle := (y/x) asFloat arcTan.
  ^self

OldPoint migrationDestination: Point.
(OldPoint new x: 123; y: 456) migrate.
%
a Point
  radius          472.2975756871932
  angle           1.307329785759979

```

If you may be migrating instances from a completely separate version of class Point, that does not have the instance variables x and y, nor use the Cartesian coordinate system, then you should add behavior to the old class to make it polymorphic to the new class.

**Finding Instances**

Preparing the set of objects that needs to be migrated can be done in a number of ways. You may, for example, have application collections of the instances that need to be migrated.

Alternatively, there are several methods available to allow you to find instances of one or more classes.

Finding instances requires scanning the entire repository, which can take significant time for very large repositories. Likewise, in a large repository there may be many instances in

the result set, potentially more than can fit into memory. The choice of methods to use to locate objects for migration depends on the size of the repository and the number of instances.

The following tables list the methods and some considerations for use. Other methods are available; see the image for details.

**Table 11.9 Finding instances**

Expression	Return value	Utility
SystemRepository <b>listInstances:</b> <i>anArrayOfClasses</i> SystemRepository <b>fastListInstances:</b> <i>anArrayOfClasses</i>	Returns an Array of Arrays; each contains all instances whose class is equal to the corresponding element in <i>anArrayOfClasses</i> . <i>anArrayOfClasses</i> may include multiple classes with the same class name, such as versions of the same class.	Performs one repository scan finding instances of all the specified classes. The result set objects are in-memory. Both temporary instances and persistent instances are included in the results.
SystemRepository <b>allInstances:</b> <i>classOrCollOfClasses</i> SystemRepository <b>fastAllInstances:</b> <i>classOrCollOfClasses</i>	Returns an instance of GsBitmap, or a set of class/GsBitmap instances corresponding to the elements of the argument.	Performs one repository scan finding instances of all the specified classes. Only committed objects are included in the results. GsBitmap and its content objects do not consume object memory, but you will still need to bring each instance into memory in order to migrate.

## Tuning migration and managing memory

Your session is configured to have a certain amount of object memory, defined by the configuration parameter `GEM_TEMPOBJ_CACHE_SIZE`. The collection containing the instances you are going to migrate must fit into memory, as well as the instances themselves. If the number of instances you are going to migrate is large, you will likely not have enough memory to hold everything.

For a large migration, you will want to:

- ▶ Increase the amount of temporary object memory by tuning `GEM_TEMPOBJ_CACHE_SIZE`. See the *System Administration Guide* for details.
- ▶ Commit your transaction after migrating some number of instances. This allows that memory to be reused.
- ▶ Use GsBitmaps to hold the collections of objects to be migrated, and bring them into memory (fault them in) in manageable chunks. See Example 11.10.

- ▶ When finding instances, consider the time it will take to scan your repository. Repository scans can be tuned; a fast variant may be more appropriate, or a reduced-impact scans that can run in the background.
- ▶ Dividing the work between multiple gem sessions allows migrate to complete more quickly.
- ▶ Migration generates shadowed objects, and large migrations should tune reclaim to make sure these shadow objects do not cause the repository to grow excessively. See the *System Administration Guide* for how to manage reclaim.
- ▶ In some cases, migrating in page order can provide benefit. See Example 11.11 below.

## Using GsBitmaps to manage memory for large result sets

For large result sets, it may be helpful to use `Repository >> allInstances:`, which returns one or more instances of `GsBitmap`. A `GsBitmap` uses heap memory, not object memory, and the result can be arbitrarily large.

Once you have the `GsBitmap` or `GsBitmaps`, you may enumerate them using `do:`, or retrieve objects from them using methods such as `removeCount:`, and perform the migration. For more on how to use `GsBitmaps`, see the section “`GsBitmap`” on page 66.

While a `GsBitmap` cannot be committed, it can be saved to a file and reloaded later. However, note that if any of the objects become dereferenced in the period after this file is written, and becomes garbage collected, this cannot be detected. When you read in the `GsBitmap` file, the `oopNumber` may reference an entirely different object. It is important to read and process the file as soon as possible after it is created.

## Tuning system resource use when finding instances

Repository-wide scans such as the ones described above use a multi-threaded scan that can be tuned to use more or less resources of the system, thereby impacting performance of anything else running on this system to a greater or lesser degree.

The regular methods use a conservative amount of system resources, while the "fast" variants allow the scan to complete faster and use a greater percentage of system resources. If you are performing a scan in an offline or single-user system, the fast variants may be more appropriate. It is also possible to tune the scan to use fewer resources.

For details on tuning the multi-threaded scan, see the *System Administration Guide*.

## Committing the migration in chunks

You will, of course, always commit or abort immediately before starting the migration and commit at the end, to ensure all your migrations are committed successfully.

For larger migrations, you will likely need to perform periodic commits to avoid running out of memory, since the migrated objects must be kept in memory until they are committed.

Example 11.10 brings instances into memory and commits in chunks of 10000 objects. The appropriate chunk size will depend on your memory and result set size; much larger chunks may be more appropriate.

---

**Example 11.10 Migration using GsBitmap in chunks**

---

```
| searchResults startIndex limit |
searchResults := (SystemRepository fastAllInstances: { Animal } ).
searchResults do: [:entry | | bm |
    bm := entry last.
    [bm isEmpty] whileFalse:
        [ | chunk |
            chunk := bm removeCount: 10000.
            entry first migrateInstances: chunk to: NewAnimal.
            System commitTransaction
                ifFalse: [self error: 'commit failed'].
        ]
    ]
]
```

---

**Migrating instances in Page Order**

For the most efficient migration of large sets of objects of multiple classes, you should perform the migration in page order – the same order as the objects are stored on disk. This allows multiple objects of several different classes on the same page in the repository to be migrated at the same time.

If the repository is in active use, the objects will move from page to page (reclaim moves the live objects off of a page, in order to reclaim the page). If there has been movement such as this between the time the file was written and read, then the page-order efficiency will be lost.

Getting the collection of objects to migrate in page order uses GsBitmap file protocol:

```
GsBitMap >> writeToFileInPageOrder: aFileName
GsBitMap >> readFromFile: aFileName withLimit: int startingAt: startIndex
```

For example, to migrate all instances of Animal in page order, in chunks of 2000;

---

**Example 11.11 Page order migration**

---

```
| searchResults bm startIndex limit |
searchResults := (SystemRepository allInstances: { Animal })
first.
searchResults last writeToFileInPageOrder: 'animal_instances.bm'.
limit := bm size.
startIndex := 1.
Animal migrationDestination: NewAnimal.
[ startIndex <= limit ] whileTrue:
    [bm := GsBitmap new.
    (bm readFromFile: 'animal_instances.bm' withLimit: 2000
    startingAt: startIndex).
    bm do: [:ea | ea migrate].
    startIndex := startIndex + 2000.
    System commitTransaction
    ifFalse: [self error: 'commit failed']].
].
```

---

## 11.4 Multi-threaded instance migration

For simple migrations, the migration can be considerably faster using multi-threaded instance migration. This allows you to migrate all instances of up to 2000 classes in a single operation.

Only simple migration operations on non-collection classes are supported.

Supported:

- ▶ mapping instance variable values to differently named instance variables
- ▶ set values to nil
- ▶ preserve or remove dynamic instance variables

Not supported:

- ▶ migrating indexable or NSC objects (that is, collections, with or without instance variables)
- ▶ modifications to dynamic instance variables
- ▶ mapping between dynamic and named instance variables
- ▶ modifying instance variable values

The operation commits if successful, so this should be done when no other users would be committing, to avoid concurrency conflicts.

Multithreaded migration requires the MigrateObjects privilege.

Multi-threaded migration is configured using an instance of `InstVarMappingArray`, which maps an old class instance to the migration destination new class instance, and specifies the changes to be made during migration.

To ensure that your migration produces the results you expect and there are no unforeseen issues, it is recommended that prior to the actual migration, you test the migration with a more limited number of test objects. `Repository >> testMigrateMt:with:` allows you to perform the same migration as `Repository >> migrateMt:`, on a limited set of objects, and does not commit; after running this method commits are disallowed, so you must abort after running this test.

### InstVarMappingArray

The input to multithreaded migration is an array of instance of `InstVarMappingArray`. You will create an individual instance of `InstVarMappingArray` for each class that you want to migrate. This instance specifies migration of instance variables defined on that class and inherited instance variables in the same way.

Create the instance using `mappingFrom:to:`,

```
InstVarMappingArray class >> mappingFrom: originalClass to: newClass
```

By default, instance variables values and dynamic instance variable values are preserved in the new instance. You may configure some specific changes using the following methods:

```
InstVarMappingArray >> mapInstVarNamed: originalClassInstVarName
to: newClassInstVarName
```

```
InstVarMappingArray >> mapInstVarToNil: originalClassInstVarName
```

```
InstVarMappingArray >> preserveDynamic: aBoolean
```

## Example

For example, to create a mapping from `OrigClass` to `NewClass` that moves the value of the instance variable `oldName` to the location of the instance variable `newName` in the `NewClass`:

```
| im |  
im := InstVarMappingArray mappingFrom: OrigClass to: NewClass.  
im mapInstVarNamed: #oldName to: #newName.  
im mapInstVarToNil: #oldName.  
im preserveDynamic: false.  
SystemRepository migrateMt: { im }
```

Note that if you do not specify to set `#oldName` to `nil`, the value at `oldName` will be retained in the migrated instance (now a `NewClass`) in `#oldName`, as well as being set in the instance variable `#newName`.



# Encryption and Validation

---

Many modern systems require encryption and authentication, to ensure that sensitive material remains private, and that no questions can arise about where the material originated. OpenSSL is a robust, widely used toolkit for the Transport Layer Security (TLS) and Secure Sockets Layer (SSL) protocols, and contains support for general-purpose cryptography.

The GemStone distribution includes OpenSSL executables and libraries, which are used internally within GemStone. For many OpenSSL functions, you may use your own OpenSSL installation if you prefer.

This chapter describes the tools and features that GemStone provides that make working with SSL keys easier.

**Overview for SSL keys and certificates** (page 236)

How to create example keys, and the GemStone classes that encapsulated keys and certificates.

**Checksums and HMAC** (page 238)

Ways to verifying the consistency of text.

**Symmetric-Key Encryption** (page 239)

Encrypting text

**Digital Signatures** (page 240)

Digitally signing text.

**Digital Envelopes** (page 241)

How to create and use digital envelopes to transmit information with both encryption and signing.

## 12.1 Overview for SSL keys and certificates

GemStone provides the OpenSSL executable and shared libraries, which can be used to generate keys and support the GemStone features that authenticate SSL.

A full explanation of SSL key algorithms and the choice of which to use, and the details of generating keys, is outside the scope of this document, and SSL and security algorithms are under active development. Simple examples are provided here to illustrate the use of the tools. You should review your application security requirements with the latest security recommendations and ensure that your security is sufficiently strong. GemStone provides access to the OpenSSL security tools; you are responsible for using these tools correctly such that you are using encryption and authentication that is appropriate for your needs.

For applications that are not limited to internal, secure networks, security is a critical. The usage in these examples is NOT sufficient to provide complete security.

The `openssl` executable enables you to generate a private key. A public key can always be extracted from a private key. The following lines will create a private key and extract the public key, providing the public/private key pairs used in the examples.

### Example 12.1 Generate RSA public and private keys

```
unix> $GEMSTONE/bin/openssl genrsa -out privKey.pem 2048
unix> $GEMSTONE/bin/openssl rsa -in privKey.pem -pubout -out pubKey.pem
```

An X509 certificate contains a public SSL key as well as validation information, so the public key can be extracted from an X509 certificate file.

## GsTlsCredential

The classes `GsTlsCredential`, with subclasses `GsTlsPrivateKey`, `GsTlsPublicKey`, and `GsX509Certificate`, encapsulate TLS (SSL) private keys, public keys, and X509 certificates, respectively. Instances contain a hidden reference to C pointer to the OpenSSL representation of the TLS object.

Instances of `GsTlsPublicKey` and `GsTlsPrivateKey` are used for digital signing and digital envelopes, as well as in X509-Secured external sessions.

## Creating a GsTlsCredential

Instances are created by reading the PEM from a file or String. The class methods

```
newFromPemFile:
newFromPemString:
```

Private keys, which may have an optional (but generally recommended) passphrase, have additional methods:

```
GsTlsPrivateKey >> newFromPemFile:withPassphrase:
GsTlsPrivateKey >> newFromPemFile:withPassphraseFile:
GsTlsPrivateKey >> newFromPemString::withPassphrase:
GsTlsPrivateKey >> newFromPemString::withPassphraseFile:
```

These methods can be used to instantiate subclasses of `GsTlsCredential` from string or disk file arguments. For example,

```
GsTlsPublicKey newFromPemFile: 'pubKey.pem'
```

You can get the public key from an instance of `GsTlsPrivateKey` or `GsX509Certificate` by sending `asPublicKey`.

## Verifying public/private key pairs

The following method can be used to verify that two keys or X509 certificates form a matching public/private key pair:

```
matches: anotherKey
```

Return true if the receiver and *anotherKey* match each other as a valid public key-private key or certificate pair.

RSA and DSA key pairs match if both keys use the same modulus. Elliptic curve key pairs match if both keys use the same curve and the same point on that curve.

## Encryption and signing algorithms

OpenSSL 1.1.1 supports a number of encryption and signing algorithms. To query an instance of a kind of `GsTlsCredential` for the algorithm, the following methods are available:

```
algorithm
```

Answers a Symbol indicating the type of high-level PKI (Public Key Infrastructure) algorithm the receiver uses. The high-level PKI algorithms supported are:

```
#RSA - Rivest-Shamir-Adleman
```

```
#DSA - Data Signature Algorithm
```

```
#EC - Elliptic Curve Cryptography
```

All high-level algorithms have various sub-types. Use the `sslAlgorithm` method to obtain information about the specific PKI algorithm of the receiver.

```
sslAlgorithm
```

Answers a Symbol indicating the SSL type of PKI algorithm the receiver uses. See the image comment for details.

## 12.2 Checksums and HMAC

### Checksums

A checksum in the form of a `BigInteger` or hexadecimal `String` can be computed for a `String` or `ByteArray`, with the following algorithms:

- ▶ `md5`                      ▶ `sha3-224`
- ▶ `sha1`                      ▶ `sha3-256`
- ▶ `sha256`                   ▶ `sha3-384`
- ▶ `sha512`                   ▶ `sha3-256`

The methods are of the form ***algorithm***`Sum` to compute the numeric checksum, and as ***algorithm***`String` to compute the checksum as a string. See the image for the full list of methods on classes `ByteArray` and `CharacterCollection`.

For example,

```
aByteArray md5Sum
aString asSha1String
```

### HMAC (Hash-based message authentication codes)

For `Strings` or `ByteArrays`, a hash-based authentication code can be computed. This can be returned as a `BigInteger` or as a string of hex digits and letters.

The HMAC can be computed with:

- ▶ `md5`                      ▶ `sha3-224`
- ▶ `sha1`                      ▶ `sha3-256`
- ▶ `sha256`                   ▶ `sha3-384`
- ▶ `sha512`                   ▶ `sha3-256`

HMAC uses a secret key, which must be a single byte object: a `String`, `ByteArray` or `Unicode7`.

See the image for the full list of methods on classes `ByteArray` and `Character`.

For example:

```
'Fourscore and seven years ago our fathers brought forth on
this continent a new nation, conceived in liberty and dedicated
to the proposition that all men are created equal'
  asMd5HmacWithKey: 'Lincoln'
%
198049353258401870159700056328212954477

'It was the best of times, it was the worst of times, it was
the age of wisdom, it was the age of foolishness, it was the
epoch of belief, it was the epoch of incredulity'
  asSha3_224HmacStringWithKey: 'Dickens'
%
'a80a577c5ca6d29455a7ad168d92b91b59242907f9a254696e8675ce'
```

## 12.3 Symmetric-Key Encryption

Symmetric key encryption uses the same key to perform the encryption and the decryption (in contrast to asymmetric-key encryption, which use separate public and private keys).

AES encryption/decryption (Advanced Encryption Standard) is a block symmetric cipher, while ChaCha20 is a stream symmetric cipher.

AES supports several modes:

- ▶ CBC is an acronym for cipher block chaining.
- ▶ OCB is an acronym for Offset Cookbook Mode.
- ▶ GCM is an acronym for Galois Counter Mode.

ChaCha20 supports:

- ▶ Poly1305

OCB, GCM and Poly1305 are Authenticated Encryption with Associated Data (AEAD) modes. AEAD provides data authenticity, confidentiality, and integrity.

AEAD also supports Additional Authenticated Data (AAD). AAD is not encrypted and therefore not confidential, but its authenticity and integrity are guaranteed. If AAD is used, it is not included in the encrypted payload, but must be provided in order to decrypt the data. The additional data is optional, so the argument for this may be nil, in which cases it is not needed for decryption.

### Encryption

Both kinds of CharacterCollection (Strings and Unicode Strings) and ByteArrays can be encrypted.

**Algorithm**EncryptWith**NNN**BitKey: *aKey* salt: *aSalt* into: *destObjOrNil*  
tag: *aTag* extraData: *eData*

**Algorithm**DecryptWith**NNN**BitKey: *aKey* salt: *aSalt* into: *destObjOrNil*  
tag: *aTag* extraData: *eData*

These methods encrypt or decrypt the receiver, respectively, using **NNN** bits and the algorithm **Algorithm**.

**Algorithm** may be:

- ▶ aesOcb for AES-OCB.
- ▶ aesGcm for AES-GCM.
- ▶ chacha20Poly1305 for CHACHA20-Poly1305.

**NNN** may be:

- ▶ 128. In this case *aKey* must be a ByteArray of size 16.
- ▶ 192. In this case *aKey* must be a ByteArray of size 24.
- ▶ 256. In this case *aKey* must be a ByteArray of size 32.

See the image for the specific methods that are available. The example below uses aesOcbEncryptWith256BitKey:salt:into:tag:extraData: and aesOcbDecryptWith256BitKey:salt:into:tag:extraData:

*destObjOrNil* must be nil or an instance of a byte object (that is not invariant). If *destObjOrNil* is nil, the result of the operation will be placed into a new instance of `ByteArray` (encryption) or `String` (decryption); otherwise the result will be placed into the given byte object starting at offset 1.

The size of *destObjOrNil* will be modified to correctly contain all encrypted or decrypted data, and may differ from the size of the receiver due to the automatic addition or removal of padding by the cipher algorithm.

*aSalt* must be a `ByteArray` of size 12. *aKey* must be a `ByteArray` with the appropriate size for the method. The same key and salt must be used to decrypt as were used to encrypt.

During AEAD encryption, a tag is generated which is used during decryption to ensure data integrity. The tag data will be stored into the *aTag* argument, which must be an instance of a byte object. The extra data *eData* must be nil or a byte object with a character size of one (a `ByteArray`, `String`, or `Unicode7`) containing additional data to be used in generating the tag value. On decryption, the tag argument *aTag* must be the bytes generated during encryption, and the same bytes of *eData* must be provided, or nil.

When encrypting a receiver that a character size greater than one, data is placed into big-endian byte order before encryption. On decryption into a *destObjOrNil* object that has a character size greater than one, data is converted to big-endian byte order after decrypting.

## Example

The following code encrypts the `String` in *textToBeEncrypted* using AES-OCB with 256 bits, which requires a 32-byte key.

```
| textToBeEncrypted myKey mySalt myTag encoded decoded |
textToBeEncrypted := 'In sooth, I know not why I am so sad'.
myKey := ByteArray withRandomBytes: 32.
mySalt := ByteArray withRandomBytes: 12.
myTag := ByteArray new.
encoded := textToBeEncrypted aesOcbEncryptWith256BitKey: myKey
    salt: mySalt into: nil tag: myTag extraData: 'shakespeare'.
decoded := encoded aesOcbDecryptWith256BitKey: myKey
    salt: mySalt into: nil tag: myTag extraData: 'shakespeare'.
```

## 12.4 Digital Signatures

Digitally signing allows data (in the form of a `String` or `ByteArray`) to be transmitted to a remote destination, and the receiver to be certain that the data originated with the receiver and has not been altered during transmission. Signing data does not encrypt it.

The originating server signs the data using their private key. The originating server's public key can be used by any recipient to verify the signature.

A number of signing algorithms are supported:

Ec (Elliptic Curve key)	Sha3_256AndRsaPss (PSS padding)
Sha1AndDsa (PKCS1 padding)	Sha3_256AndRsa (PKCS1 padding)
Sha1AndRsaPss (PSS padding)	Sha3_384AndRsaPss (PSS padding)
Sha1AndRsa (PKCS1 padding)	Sha3_384AndRsa (PKCS1 padding)
Sha256AndDsa (PKCS1 padding)	Sha3_512AndRsaPss (PSS padding)
Sha256AndRsaPss (PSS padding)	Sha3_512AndRsa (PKCS1 padding)
Sha256AndRsa (PKCS1 padding)	Sha512AndDsa (PKCS1 padding)
Sha3_224AndRsaPss (PSS padding)	Sha512AndRsaPss (PSS padding)
Sha3_224AndRsa (PKCS1 padding)	Sha512AndRsa (PKCS1 padding)

Data is signed on the originating server with `ByteArray` or `CharacterCollection` instance methods of the form:

```
signWithAlgorithmAndKeyTypePaddingPrivateKey: aGstlsPrivateKey
into: aByteArrayOrNil
Hashes the receiver using Algorithm and signs the resulting hash with the KeyType
and Padding. Returns a ByteArray containing the resulting signature.
```

Specific methods include, for example:

```
signWithSha3_256AndRsaPssPrivateKey:into:
signWithSha512AndDsaPrivateKey:into:
```

After transmission to the destination, it is verified with the matching `ByteArray` or `CharacterCollection` method of the form:

```
verifyWithAlgorithmAndKeyTypePaddingPublicKey: aGstlsPublicKey
signature: aByteArray
Hashes the receiver using Algorithm and verifies the resulting hash using the
KeyType and Padding. Returns true if the signature is correct.
```

Specific methods include, for example:

```
verifyWithSha3_256AndRsaPublicKey:signature:
verifyWithSha512AndDsaPublicKey:signature:
```

## 12.5 Digital Envelopes

`GsDigitalEnvelopes` allow both encryption and signing in a single operation. A `GsDigitalEnvelope` has the following security features:

- ▶ **Confidentiality** -- the message is encrypted using a randomly generated AES encryption session key and initialization vector. The session key is then encrypted with the provided public encryption key.
- ▶ **Integrity** -- authenticated encryption guarantees the cipher text has not been altered. The digital signature guarantees the encrypted key, initialization vector and tag have not been altered.
- ▶ **Authentication** -- the receiver of the envelope is assured that the sender signed the envelope with the private key matching the public key used to successfully verify the signature.

NOTE: In order to guarantee authentication, the receiver must confirm that the public verification key actually belongs to the sender. Normally this confirmation is done by verifying that an X509 certificate containing the sender's public key has been signed by a reputable certificate authority.

GsDigitalEnvelope does NOT do this public key/certificate signature verification; it is up to the envelope's recipient to ensure that the public key or X509 certificate used to verify the signature is trustworthy.

## Creating the GsDigitalEnvelope

Digital Envelopes are created using the following methods:

```
GsDigitalEnvelope class >> encryptMessage: messageBytes
    withPublicKey: publicEncryptionKey
    cipherId: cipherOpCode
    withPrivateSigningKey: privateSigningKey
```

```
GsDigitalEnvelope class >> encryptMessage: messageBytes
    withPublicEncryptionKeys: arrayOfPublicEncryptionKey
    cipherId: cipherOpCode
    withPrivateSigningKey: privateSigningKey
```

With the following arguments:

*messageBytes*

The message to be encrypted may be any kind of String or ByteArray. The name of the class is stored in the envelope so when the envelope is opened, an instance of the same kind of object is returned. Internally, if the message is a multi-byte String, the string is converted into big-endian form, but this is transparent to the user. It is an error if the class of object that was encrypted into the envelope is not resolvable at the destination.

*publicEncryptionKey* and *privateSigningKey*

The message is encrypted using a public encryption key, and signed using a private signing key, which should be distinct. The envelope can only be "opened" (decrypted) by a recipient with the private key that matches the public encryption key, and the public key that matches the private signing key.

Providing an array of multiple encryption keys allows a message to be encrypted by multiple public keys into multiple envelopes in one operation. Each envelope can be opened by the single matching private key, and the public signing key.

The encryption and signing keys can be instances of GsTlsPublicKey or GsX509Certificate. Encryption keys must be RSA; signing keys may be DSA, RSA, Ed25519, ECDSA, EC, or Ed448.

For signing algorithms that require a message digest algorithm, the SHA2-256 message digest is used. Signatures generated with RSA keys are padded using the RSA\_PKCS1\_PSS\_PADDING scheme.

*cipherOpCode*

The `cipherId`: argument accepts a numeric code specifying the AEAD (Authenticated Encryption with Additional Data) of the cipher; legal values are in the following table.

opCode	Cipher	Mode	bits/Bytes
4	AES	OCB	128/16
5	AES	OCB	192/24
6	AES	OCB	256/32
7	AES	GCM	128/16
8	AES	GCM	192/24
9	AES	GCM	256/32
10	CHACHA20	Poly1305	256/32

Ciphers corresponding to 6 and 10 are considered the most secure and are recommended for most applications.

## Using the GsDigitalEnvelope

GsDigitalEnvelopes can be passivated into string form, for ease of transmission to the destination.

Once a GsDigitalEnvelope object is created, it can be converted to a passive object (String) for transmission to the recipient. Decrypting ("opening") the given envelope requires the private key matching the argument public key, and the public signing key matching the private key used to sign the envelope. The recipient activates the GsDigitalEnvelope object, and then opens it using the matching keys.

A detailed example is provided in class methods in GsDigitalEnvelope. The example follows the following form, first creating the envelope and then passivating it on the originating server, and activating and decryption on the destination.

### On the originating server

```
origEnvelope := GsDigitalEnvelope
  encryptMessage: messageText
  withPublicKey: aGsTlsPubEncrKey
  cipherId: 10
  withPrivateKey: aGsTlsPrivSigKey.
aStream := WriteStream on: String new.
PassiveObject passivate: origEnvelope toStream: aStream.
aPassiveString := aStream contents.
```

**On the destination server**

```
destEnvelope := (PassiveObject newWithContents: aPassiveString)
    activate.
decryptedMsg := destEnvelope
    decryptWithPrivateKey: aGsTlsPrivEncrKey
    withPublicVerificationKey: aGsTlsPubSigKey.
```

# Operating System Access

---

GemStone applications interact with their operating system environment in many ways. The ability to directly execute operating system commands provides a flexible and powerful way to accomplish many tasks. Sockets provide another powerful way to communicate between processes.

This chapter also includes information on serialized GemStone data using `PassiveObject` and `JSON`.

**Executing Operating System Commands** (page 246)

how to execute operating system commands from GemStone.

**Setting environment variables** (page 249)

Setting environment variables from the GemStone image.

**Creating and Using Sockets** (page 249)

describes the protocol provided by class `GsSocket` and `GsSecureSocket` to create operating system sockets and exchange data between two independent interface processes.

**ssh and sftp using OpenSSL** (page 258)

how to use the specialized sockets for `ssh` and `sftp`.

**Serializing Data** (page 261)

describes the mechanism that GemStone provides for creating serialized versions of the objects that represent your data.

## 13.1 Executing Operating System Commands

GemStone can execute arbitrary operating system command line expressions, using `System performOnServer:`, or an instance of the `GsHostProcess` class.

### Privileges and limiting OS access

Performing server operations gives your application broad powers, and you may wish to limit the specific server operations that are permitted.

GemStone user privileges, which are described in the *System Administration Guide*, may restrict who can execute OS commands, and which OS commands can be executed.

GemStone UserProfiles that have the (reverse) privilege `NoPerformOnServer` cannot execute any OS commands using `performOnServer:` or using `GsHostProcess`.

Users with `NoPerformOnServer` may be allow to execute specific named commands that are added to the allowlist for their UserProfile. An administrator with `OtherPassword` privilege can add the full path of a command to the allowlist for a specific user, or remove it, using:

```
UserProfile >> addPerformOnServerCommand: aString
```

```
UserProfile >> removePerformOnServerCommand: aString
```

See the *System Administration Guide* for more information, or review methods in the image.

### performOnServer: for simple commands

System also understands the message `performOnServer: aString`, which causes the UNIX shell commands given in *aString* to execute in a subprocess of the current GemStone process. The output of the commands is returned as an instance of `String`. For example:

```
System performOnServer: 'date'
%
Fri Aug 04 13:10:21 PDT 2023
```

The commands in the argument to `performOnServer:` can have exactly the same form as a shell script; for example, new lines or semicolons can separate commands; you may embed single quotes or double quotes within the argument starting, to allow arguments with spaces; and the character “\” can be used as an escape character.

For example:

```
System performOnServer: 'cat ''a b.txt'''
contents of file
```

```
System performOnServer: 'export NEWENVVAR=44
echo $NEWENVVAR'
%
44
```

The string returned is whatever an equivalent shell command writes to *stdout*. If the operating system command reports an error, you must interpret the string to determine the correct action. For example:

```
System performOnServer: 'cp'
%
[40738561 size:63 String] cp: missing file operand
Try 'cp --help' for more information.
```

The String returned from the `performOnServer:` command is always composed of 8-bit Characters. If the operating system command produces UTF-8 encoded results, then the Smalltalk String will be UTF-8 encoded. You will need to send `decodeFromUTF8ToString` or `decodeFromUTF8ToUnicode` to decode the results. `performOnServer:` may also return results as 8-bit non-encoded, extended ASCII, if that is what was returned by the operating system commands that were executed.

## Using other shells

By default, `performOnServer:` executes using `/bin/sh`. To use other shells, specify the full path to the shell executable using `System class >> performOnServer: aString withShell: aShellOrNil`.

Not all shells will work as an argument to `performOnServer:withShell:.` `bash` (`/bin/bash`), Korn Shell (`/bin/ksh`), Z Shell (`/bin/zsh`) and are known to work. C Shell (`/bin/csh`) and its variants (`/bin/tcsh`) are known to not work.

## GsHostProcess for more complex interactions

`System >> performOnServer:` can execute arbitrary OS code on the server, but only operates synchronously; Smalltalk will block until the command has completed.

To interact with an operating system command, perform code asynchronously, and to allow Smalltalk to read from *stdout* or write to *stdin* or to redirect *stdout* or *stderr* to files, you can use the class `GsHostProcess`.

To use this, use one of the class methods `execute:`, `execute:input:`, or `fork:`, passing the command line you wish to execute. `execute:` will execute the command and return the results; `fork:` will return immediately with an instance of `GsHostProcess` with sockets on *stdin*, *stdout*, and *stderr*. You can use socket protocol to read from or write to these sockets.

For example, to fetch the date from the operating system:

```
GsHostProcess execute: '/bin/date'.
%
Fri Aug 04 13:10:21 PDT 2023
```

Note that pathname resolution is **not** provided. You must fully qualify executable paths. Build-in shell commands cannot be invoked with `GsHostProcess`.

As with `performOnServer:`, `GsHostProcess` accepts a string, which can include spaces to separate arguments and single or double quotes.

```
GsHostProcess execute: '/bin/cat helloworld.txt'
GsHostProcess execute: '/bin/cp /gshost/f1.txt /gshost/f2.txt'
```

However, `GsHostProcess` does not simply execute a shell command, so you cannot pass multiple commands using newlines, since newlines within the command are treated as argument separators; nor can you use other shell syntax.

When an operating system command executed by `GsHostProcess` fails, it will signal a `ChildError` with the error status and message from `stderr`. For example,

```
GsHostProcess execute: '/bin/cp'
%
ERROR 2710 , a ChildError occurred (error 2710), status 1,
/bin/cp: missing file operand
Try '/bin/cp --help' for more information.
```

## Using `execute:input:` to pass data to `stdin`

The method `GsHostProcess >> execute:input:` allows you to pass in data to the `stdin` of the command being executed.

For example:

```
GsHostProcess execute: '/bin/grep wo' input: 'hello
world'
%
world
```

## Using `fork:` for interactive commands

Creating a `GsHostProcess` using the class method `fork:` allows you to interact with the `GsHostProcess` instance. For example:

```
| hostprocess |
hostprocess := GsHostProcess fork: '/bin/date'.
hostprocess stdout read: 1024
%
Fri Aug 04 13:10:21 PDT 2023
```

`GsHostProcess >> readOutErr` will read from both `stdout` and `stderr`.

When using `fork:`, the `GsHostProcess` can report its status using `childStatus`; you may kill the child process using `killChild` or `killChild: timeoutSeconds`.

## Connecting `stdin`, `stdout`, and `stderr`

When using `GsHostProcess >> fork` or `fork:`, rather than `execute:*`, it is recommended to specify one or more of:

```
GsHostProcess>>stderrPath:
GsHostProcess>>stdinPath:
GsHostProcess>>stdoutPath:
```

to redirect standard files to the filesystem. Otherwise there is a risk the child may hang when trying to read or write to a pipe connected to the parent process, if the parent is not executing the data reads and writes equivalent to those implemented in `GsHostProcess>>_executeWithInput:`.

To redirect both stdout and stderr to the same file, use `stdoutPath:`, followed by `GsHostProcess >> redirectStderrToStdout`.

## 13.2 Setting environment variables

You can ask about and set environment variables in the server or client environments using the following methods:

```
System class >> clientEnvironmentVariable: varName
System class >> clientEnvironmentVariable: varName put: valueStr
System class >> gemEnvironmentVariable: varName
System class >> gemEnvironmentVariable: varName put: valueStr
```

These methods cannot be used if the `NoGsFileOnServer` privilege is set for the `UserProfile`.

For example,

```
run
System gemEnvironmentVariable: 'GS_MY_INDEX_KEY' put: '47'.
System clientEnvironmentVariable: 'GS_MY_KEY' put: '84'.
System performOnServer: 'env | grep GS_MY'
%
GS_MY_KEY=84
GS_MY_INDEX_KEY=47
```

When methods that accept a filename, such as `GsFile` methods, include an environment variable, environment variables set in this way are used in the expansion.

## 13.3 Creating and Using Sockets

Sockets open a connection between two processes, allowing a two-way exchange of data. The class `GsSocket` provides a mechanism for manipulating operating system sockets from within GemStone Smalltalk.

Methods in the class `GsSocket` do not use the terms *client* and *server* in the same way as the methods in class `GsFile`. Instead, these terms refer to the roles that two processes play with respect to the socket: the server process creates the socket, binds it to a port number, and listens for the client, while the client connects to an already created socket. Both client and server are processes created (or spawned) by a Gem process.

In addition to standard sockets created by `GsSocket`, you can create secure SSL sockets using the class `GsSecureSocket`. `GsSecureSocket` is a subclass of `GsSocket` that adds protocol to specify certificates and require authentication.

Both `GsSocket` and `GsSecureSocket` contain a number of examples as class methods, in the method category `Examples`.

`GsSocket` class methods such as `clientExample` and `serverExample`. These methods provide examples of how to create a socket connection between two sessions. The example methods work together; they require two separate sessions running from two independently executing interfaces, one running the server example and one running the client example. You can execute these methods from `Topaz` or from `GemBuilder` for

Smalltalk, but note that `serverExample`, which should be started first, will take control of the interface until the `clientExample` completes the socket connection.

`GsSecureSocket` has a more extensive collection of examples that include single-method examples, as well as pairs of examples. See the image for the available examples.

## GsSocket and GsSignallingSocket

`GsSocket` and `GsSignallingSocket` represent a basic socket. Methods on `GsSocket` generally return `nil` on error, while methods on `GsSignallingSocket` signal an `Error`.

### Establishing the connection

To setup a socket connection, you create instances of `GsSocket` in both the client and server processes.

1. On the server side, create an instance of `GsSocket`, and call `makeServerAtPort`: This creates a listening socket on the given port.

To have the operating system select a port, use a wildcard bind using `makeServer:`, or pass `nil` as the port argument. You will then need to determine the port that the client should connect at using the `port` method.

2. On the client side, create an instance of a `GsSocket` and call one of the following:

- ▶ `connectTo`: for a connection to a process on the same host
- ▶ `connectTo:on`: if the server is on a different machine
- ▶ `connectTo:on:timeoutMs`: to specify a timeout for the connection

Provided there was a listening server socket setup as in step 1, this will initiate the connection to the server.

3. The server then does an `accept`, or `acceptTimeoutMs`: (to specify a timeout) . This returns a new instance of `GsSocket` for the client connection.

Note that the server side has two sockets; a listening socket and the established socket with the client.

### Communication on the socket

Each process can write and read to the socket using protocol such as `write:` and `read:`. See the image methods in the categories Reading and Writing for specific methods.

Writes and reads are of byte objects such as `String` or `ByteArray`. Read operations are for a specified number of bytes, and return the actual number of bytes read if fewer bytes were available (if fewer bytes were written to the socket by the peer). A return value of `nil` means an error occurred, and for read operations, a return value of 0 means the socket EOF was reached.

### Closing the socket

When completed, the client should close its socket and the server close the listening and established sockets. This is done by simply sending `close` to the sockets.

## Socket Configuration

Socket configuration can be done using the method

```
GsSocket >> option:put:
```

See the comments in this method for details on socket configuration.

The most common option is blocking.

## Blocking

Sockets can be made blocking or non-blocking, and the blocking status checked, using the following methods:

```
GsSocket >> makeNonBlocking
```

```
GsSocket >> makeBlocking
```

```
GsSocket >> isNonBlocking
```

```
GsSocket >> isBlocking
```

## GsSecureSocket

GsSecureSocket creates a secure socket Secure Sockets Layer (SSL), providing access to the open-source OpenSSL library. This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>).

The protocol is Transport Layer Security (TLS), the successor to SSL; GemStone continues to use the term "SSL" in most cases, since SSL remains commonly used to refer to both.

To create a secure socket, you create instances of the GsSecureSocket class and first establish the connection as a regular socket. Then, further protocol authenticates the connection to make the socket secure.

GsSecureSocket class protocol includes setup of certificates and types of authentication, which can be set for general socket operations, so each new socket does not have to be separately configured.

## Certificates, keys, and passphrases

GsSecureSocket instances must be configured with the CA certificates, private key files, and passphrases, to allow them to complete the secure handshake. This can be done using class methods that apply to all new instances, or you can send instance methods to configure individual settings.

The required certificates, CA certificates, private key files and passphrases may be provided by your organization. The GemStone distribution includes example certificates, keys, and passphrases to verify your code, under the directories:

```
$GEMSTONE/examples/openssl/certs/  
directory containing example public keys
```

```
$GEMSTONE/examples/openssl/private/  
directory containing example private keys and passphrases
```

The distribution includes script that will allow you to generate certificates:

```
$GEMSTONE/examples/openssl/create_ca.sh
```

```
$GEMSTONE/examples/openssl/create_new_certs.sh
```

And the openssl executable, matching the version that GemStone uses:

```
$GEMSTONE/bin/openssl
```

Using this openssl executable, rather than any version that may be present on your system, is recommended. For details on the openssl interface, see <http://www.openssl.org/docs/apps/openssl.html>.

## Enable or disable verifying CA Certificate

Certificate Authority (CA) certificates should be setup prior to creating instance of GsSecureSocket. An example CA certificate file is provided here:

```
$GEMSTONE/examples/openssl/certs/cacert.pem
```

By default, client sockets verify the certificates presented by the server, and are configured with the CA certificate. By default, servers do not verify certificates. This is the usual case, and if this is your preferred behavior you do not need to explicitly enable or disable verification.

Class methods that configure verification must be executed prior to the creation of the client or server GsSecureSocket instance.

### Client Sockets

By default, client sockets verify connections from the server. You will need to specify the client certificate file, or the directory containing CA certificates, using one of the following methods.

```
GsSecureSocket class >> useCACertificateFileForClients: certfile
GsSecureSocket class >> useCACertificateDirectoryForClients: aDir
```

A CA certificate file must be in PEM format. It may contain more than one certificate. The certificate has effect only for the life of the session; the state is not committed to the repository.

To disable validation on the client, use the methods:

```
GsSecureSocket disableCertificateVerificationOnClient
aGsSecureClientSocket disableCertificateVerification
```

Parallel methods exist to re-enable validation after validation is disabled.

### Server Sockets

By default, server sockets do not verify the certificate from the client, so you do not need to specify the CA certificate file or directory.

Setting the CA certificate file or directory also enables verification. To enable, execute one of the following:

```
GsSecureSocket class >> useCACertificateFileForServers: certfile
GsSecureSocket class >> useCACertificateDirectoryForServers: aDir
```

Parallel methods exist to disable validation after validation is enabled.

The server accepts additional verification options. Using the methods

```
GsSecureSocket setCertificateVerificationOptionsForServer:
aGsSecureServerSocket setCertificateVerificationOptions:
```

The following options may be set:

```
#SSL_VERIFY_FAIL_IF_NO_PEER_CERT
    if the client did not return a certificate, the TLS/SSL handshake is immediately
    terminated with a 'handshake failure' alert.

#SSL_VERIFY_CLIENT_ONCE
    only request a client certificate on the initial TLS/SSL handshake. Do not ask for a
    client certificate again in case of a renegotiation.
```

## Set certificate, private key, and passphrase

The certificate, private key, and private key passphrase can be setup by class methods to apply to all instances, or by sending messages to the instance of GsSecureSocket.

You can pass the certificate and private key either by using pathnames to the files, or by the string. If a string is used, it must exactly match the contents of the corresponding certificate file (including white space, line feeds, etc.), or the strings will not be accepted.

Both certificate and private key must be in PEM format, and the private key must match the certificate. The same file may be specified for the certificate file and the private key file.

The certificate may contain a certificate chain or a single certificate.

If the private key requires a passphrase, it must be specified as a string. If the private key does not require a passphrase, the argument is expected to be nil.

Example certificates, public and private keys, and passphrases are under the `$GEMSTONE/examples/openssl/` directory. A number of examples are provided, for example:

```
$GEMSTONE/examples/openssl/certs/server_1_servercert.pem
$GEMSTONE/examples/openssl/private/server_1_serverkey.pem
$GEMSTONE/examples/openssl/private/server_1_server_passwd.txt
```

## Class setup for server sockets

To specify the server certificates, private key file, and passphrase (if required), that will be used for all secure server sockets that are created after these methods are invoked, use the method:

```
GsSecureSocket class >>
    useServerCertificateFile: certFilePathAndName
    withPrivateKeyFile: privateKeyFilePathAndName
    privateKeyPassphrase: passphraseStringOrNil
```

A variant is available that allows you to pass in the strings containing the certificate and key as string; see the image methods. If a string is used, it must exactly match the contents of the corresponding certificate file (including white space, line feeds, etc.), or the strings will not be accepted.

## Class setup for client sockets

By default, certificates are not verified on the client, so the certificate and private key do not need to be setup for clients.

If you need to enable client validation, the follow methods specify the client certificates, private key file, and passphrase, that will be used to validate server connections for secure client sockets that are created after these methods are invoked:

```
GsSecureSocket class >>
  useClientCertificateFile: certFilePathAndName
  withPrivateKeyFile: privateKeyFilePathAndName
  privateKeyPassphraseFile: passphraseStringOrNil
```

A variant that accepts the content strings instead of filenames is available.

### Instance setup for client or server sockets

You can specify the certificate, private key, and passphrase for a single specific instance of GsSecureSocket (either a server socket or a client socket), using instance method:

```
GsSecureSocket >>
  useCertificateFile: certFilePathAndName
  withPrivateKeyFile: privateKeyFilePathAndName
  privateKeyPassphraseFile: passphraseStringOrNil
```

Again, a variant that accepts the content strings instead of filenames is available.

### Setup the Cipher list

The list of ciphers that are acceptable to use can be configured, either on the class side for servers and clients, or for specific instances of GsSecureSocket client or server sockets.

The cipher list is specified as a formatted string. See <http://www.openssl.org/docs/apps/ciphers.html> for details on the format of this string (as well as other information on ciphers).

For example, to use all ciphers except NULL ciphers and anonymous Diffie-Hellman (DH), and sort by strength, use the following string:

```
'ALL:!ADH:@STRENGTH'
```

To configure the cipher list for all instances of GsSecureSocket, use the following methods. These methods return true if the specification finds one or more usable ciphers, false if no usable ciphers match the specification.

```
GsSecureSocket class >>
  setClientCipherListFromString: aString
GsSecureSocket class >>
  setServerCipherListFromString: aString
```

To configure the cipher list for a specific instance of a server socket or client socket, the ciphers must be set before `secureConnect:/secureAccept` are executed. This method returns true if the specification finds one or more usable ciphers, false if no usable ciphers match the specification, and nil if the operation has no affect because the receiver is already connected.

```
GsSecureSocket class >>
  setCipherListFromString: aString
```

Once an instance of GsSecureSocket is successfully connected, you can fetch the cipher in use using:

```
GsSecureSocket >> fetchCipherDescription
```

## Establishing the connection

Rather than creating instances using `GsSocket class >> new`, with `GsSecureSocket` sockets are instantiated using `newClient` and `newServer`.

To establish the socket connection, as with regular `GsSocket`,

1. The server creates the socket using `newServer`, and calls `makeServerAtPort :` on an unused port, to create the server listener socket on that port.
2. The client creates the socket using `newClient`, and calls `connectTo :`, specify the same port as in Step 1.
3. The server socket calls `accept`, or `acceptTimeoutMs :`, which creates the connected socket on the given port.

This establishes the standard socket, but the connection is not secure. Another client-server interaction is required to make this a secure socket.

At this point, you can setup specific certificates and ciphers that will apply to these sockets only, as described in the preceding sections. This is needed if you have not previously set up certificates and ciphers that apply to all `GsSecureSocket` connections.

Then continue with the process that makes the socket secure:

4. The client socket calls `secureConnect`.
5. The server socket calls `secureAccept`, or `secureAcceptTimeoutMs :`.

If these methods return true, then the connection is secure. To determine if you have a secure connection, use the method:

```
GsSecureSocket >> hasSecureConnection
```

## Communication on the socket

At this point reads and writes are done as for standard sockets.

## Closing the socket

You can either close the socket connection entirely, or close the secure connection and remain connected for normal (not secure) communication.

To close the socket entirely, use

```
GsSecureSocket >> close
```

Which performs both the secure close and the regular close.

Note that the secure `close` requires a handshake. If the socket is blocking, and the peer does not respond, then the close will hang. To close the socket, we recommend first making it non blocking:

```
mySecureSocket makeNonBlocking.  
mySecureSocket close.
```

To close only the secure socket and leave the connection available for non-secure communication, you can use the method

```
GsSecureSocket >> secureClose
```

Which must be executed by both sockets on the connection. You can then call `close` later, to close the connection entirely.

## HTTPS connection

The image includes an example of using a `GsSecureSocket` to setup an HTTPS connection to a well known search engine URL, verifying of the server certificate, and performing a simple GET request. See the image for the example method:

```
GsSecureSocket class >> httpsClientExample
```

## Error handling

### GsSocket

Many `GsSocket` operations return nil if an error occurs.

On a nil return, you can query the system for the details on the error that occurred using the following methods. These methods are implemented both for the class and instance of `GsSocket`. For errors in `GsSocket` class methods, use the class side error methods, and for errors in `GsSocket` instance methods, use the instance methods

```
lastErrorString
```

Returns a String containing information about the last error or nil if no error has occurred. Clears the error information.

```
lastErrorCode
```

Returns an integer representing the last OS error or nil if no error has occurred. Does not clear the error information

```
lastErrorSymbol
```

Returns a Symbol representing the last OS error or nil if no error has occurred. Does not clear the error information.

### Signalling errors

You can send `raiseExceptionOnError:` to a socket, in which case a `SocketError` is signalled in cases where a nil was returned from the C code.

### GsSecureSocket

Most `GsSecureSocket` operations signal a `SocketError` or a `SecureSocketError` if an error occurs. Ordinary socket operations will signal a `SocketError`; errors in the security configuration will signal a `SecureSocketError`.

You can query for the last error from an instance operation using the instance methods:

```
GsSecureSocket >> lastErrorString
```

```
GsSecureSocket >> fetchLastIoErrorString
```

These methods fetch and clear the error string from a call to SSL functions for connect, accept, read, or write.

On the class side, the following methods return error strings for any SSL function call errors:

```
GsSecureSocket class >> fetchErrorStringArray
```

Returns an Array of error strings generated by the OpenSSL package. The errors returned are cleared from the SSL error queue. The array is ordered from oldest to newest error.

```
GsSecureSocket class >> lastErrorString
  Returns a string describing the elements in the SSL error queue, based on the contents
  of fetchErrorStringArray. The errors are cleared from the error queue.

GsSecureSocket class >>
  fetchLastCertificateVerificationErrorForClient
GsSecureSocket class >>
  fetchLastCertificateVerificationErrorForServer
  These methods fetch and clear a string representing the last certificate verification error
  logged by, respectively, a client SSL socket or a server SSL socket.

To clear the error queue, use the method
GsSecureSocket class >> clearErrorQueue
```

## GsSecureSocket pre-shared keys

Pre-shared keys are symmetric keys shared in advance among the communicating parties. Pre-shared keys are set up using the following method:

```
GsSecureSocket >> setPreSharedKey: aByteArray
  Sets the Pre-Shared Key (psk) for the connection to be the bytes contained in
  aByteArray. Raises an exception if the receiver is already connected or listening
  for a connection. aByteArray must have a size of at least 8 bytes, not more than 64
  bytes and be a multiple of 8 (16, 24, 32 etc). A key size of at least 16 bytes is
  recommended.
```

## TLS version handling

Support for pre-shared keys requires callback protocols that are dependent on the TLS version. The following methods allow you to specify the version, and limit the maximum and minimum TLS version:

```
GsSecureSocket >> tlsActualVersion
GsSecureSocket >> tlsMaxVersion
GsSecureSocket >> tlsMaxVersion: verStringOrSymbol
GsSecureSocket >> tlsMinVersion
GsSecureSocket >> tlsMinVersion:
GsSecureSocket class >> tlsClientMaxVersion
GsSecureSocket class >> tlsClientMaxVersion: verStringOrSymbol
GsSecureSocket class >> tlsClientMinVersion
GsSecureSocket class >> tlsClientMinVersion: verStringOrSymbol
GsSecureSocket class >> tlsServerMaxVersion
GsSecureSocket class >> tlsServerMaxVersion: verStringOrSymbol
SGsSecureSocket class >> tlsServerMinVersion
GsSecureSocket class >> tlsServerMinVersion: verStringOrSymbol
```

## Example in image

```
preSharedKeyTlsClientExample: logToGciClient usingPort: portNum on: host
  Pre-shared key transport layer security (PSK-TLS) requires the client and server to
```

both know a pre-shared secret key before the connection is initiated. The key must be at least 8 bytes in size (16 hex digits) and be stored in a ByteArray.

```
GsSecureSocket class >> preSharedKeyTlsServerExample: logToGciClient
usingPort: portNum on: host
```

Pre-shared key transport layer security (PSK-TLS) requires the client and server to both know a pre-shared secret key before the connection is initiated. The key must be at least 8 bytes in size (16 hex digits) and be stored in a ByteArray.

## 13.4 ssh and sftp using OpenSSL

The libssh libraries include support for ssh and sftp. GemStone supports commonly used functions with the following classes:

```
GsSshSocket
GsSftpSocket
GsSftpRemoteFile
SshSocketError (error 2759)
```

Examples of using these classes are provided in class side methods.

### SSH with GsSshSocket

#### Creating the SSH connection

To create a SSH connection, create a client GsSshSocket using methods inherited from GsSocket, and connect using sshConnect.

For example, the following code connects to the host *hostnameOrIP*:

```
sshSock := GsSshSocket newClient.
sshSock connectToHost: hostnameOrIP timeoutMs: 2000.
sshSock userId: userName; password: userPassword.
sshSock disableHostAuthentication.
sshSock sshConnect.
```

As with all sockets, you should close the socket when you no longer need the connection.

```
sshSock close
```

#### SSH operations

Commands are executing using executeRemoteCommand:. For example,

```
result := sshSock executeRemoteCommand: 'uptime'.
```

The commands can be executed non-blocking using the following methods:

```
nbExecuteRemoteCommand:
hasCommandInProgress
nbRemoteCommandResult
nbRemoteCommandResultReady
```

## SFTP with GsSftpSocket

### Creating Sftp connection

Creating an SFTP connection is just like creating an SSH connection. To create a SFTP connection, create a client `GsSftpSocket` using methods inherited from `GsSocket`, and connect using `sshConnect`.

For example, the following code connects to the host *hostnameOrIP*:

```
sshFtpSock := GsSftpSocket newClient.
sshFtpSock connectToHost: hostnameOrIP timeoutMs: 2000.
sshFtpSock userId: userName; password: userPassword.
sshFtpSock disableHostAuthentication.
sshFtpSock sshConnect.
```

As with all sockets, you should close the socket when you no longer need the connection.

```
sshFtpSock close
```

### Sftp operations

Once the `GsSftpSocket` instance is created and connected, there are a number of methods supported, including:

```
GsSftpSocket >> remoteFileExists: aFileOrDirectory
GsSftpSocket >> currentRemoteDirectory
GsSftpSocket >> createRemoteDirectory: dirname
GsSftpSocket >> createRemoteDirectory: dirname mode: modeInt
GsSftpSocket >> contentsOfRemoteDirectory: dirname
GsSftpSocket >> contentsOfRemoteDirectory: dirname matchPattern:
    aString
GsSftpSocket >> renameRemoteFile: oldName to: newName
GsSftpSocket >> removeRemoteFile: filename
GsSftpSocket >> removeRemoteDirectory: dirname
```

### GsSftpRemoteFile for operations on remote files

The class `GsSftpRemoteFile` represents an instance of a Remote File that can be used for reading and writing. You must have an established `GsSftpSocket` connection, which is an argument to the instance creation methods.

```
GsSftpRemoteFile class >> openRemoteFile: fileName mode: openMode
    permissions: permInt errorIfExists: aBoolean withSftpSocket:
    aGsSftpSocket
GsSftpRemoteFile class >> openRemoteFileAppend: fileName
    withSftpSocket: aGsSftpSocket
GsSftpRemoteFile class >> openRemoteFileReadOnly: fileName
    withSftpSocket: aGsSftpSocket
GsSftpRemoteFile class >> createNewRemoteFile: fileName
    withSftpSocket: aGsSftpSocket
GsSftpRemoteFile class >> createOrOverwriteRemoteFile: fileName
    withSftpSocket: aGsSftpSocket
```

Once an instance of `GsSftpRemoteFile` has been created, use the instance protocol to read, write, and fetch information about the remote file. See the image for specific methods.

For example, to download a remote file using SFTP:

```
sshFtpSock := GsSftpSocket newClient.  
sshFtpSock connectToHost: hostnameOrIP timeoutMs: 2000.  
sshFtpSock userId: userName; password: userPassword.  
sshFtpSock disableHostAuthentication.  
sshFtpSock sshConnect.  
remoteSftpFile := GsSftpRemoteFile  
    openRemoteFileReadOnly: remoteFileName  
    withSftpSocket: sftpSock.  
localFile := GsFile openWriteOnServer: localFileName.  
bytes := remoteSftpFile readAllInto: localFile.  
localFile close.  
sftpFile close.
```

## 13.5 Serializing Data

There are several options to serialize GemStone data. `PassiveObject` preserves the classes and object relationships, but may not work over versions or if activated in a different repository. JSON is highly portable for simple data types.

### PassiveObject

To serialize your data to a file or for transmitting over a socket, you can *passivate* the objects themselves. Objects representing your data are stored into a serialized, text-based form by the GemStone class `PassiveObject`. `PassiveObject` starts with a root object and traces through its instance variables, and their instance variables, recursively until it reaches special objects (instances of `SmallInteger`, `Character`, `Boolean`, `SmallDouble`, or `UndefinedObject`), or classes that can be reduced to special objects (strings and numbers that are not integers), creating a representation of the object that preserves all of the values required to re-create it. The resulting *network* of object descriptions can be written to a file, stream, or string.

Each file can hold only one network—you cannot append additional networks to an existing passive object file, stream, or string.

A few objects and aspects of objects are not preserved:

- ▶ Instances of `UserProfile` cannot be preserved in this way, for obvious security reasons.
- ▶ `SystemRepository` cannot be preserved.
- ▶ In-memory `FileSystems` are not preserved.
- ▶ Blocks that refer to globals or other variables outside the scope of the block cannot be reactivated correctly.
- ▶ Blocks associated with objects (such as the sort block in `SortedCollections`) have the block's source code preserved, and recompiled on activate, which may affect behavior depending on the activation environment's `symbolList`.
- ▶ Any indexes you have created on the object are lost.
- ▶ Identities (OOPs) are not preserved.

Dynamic instance variables are preserved. This can be customized for specific classes, by overriding the method `shouldWriteDynamicInstVar: instVarName`, with the specific dynamic instance variable that you do not want included in the passivated form.

The relationship between two objects is conserved only so long as they are described in the same network. Similarly, if two separate objects A and B both refer to the same third object C, then making A and B passive in two separate operations will result in duplicating the object C, which will be represented in both A's and B's network. Because the resulting network of objects can be quite large anyway, you want to avoid such unnecessary duplication. For this reason, it is usually a good idea to create one collection to hold all the objects you wish to preserve before invoking one of the `PassiveObject` methods.

In addition, since object identity is not preserved, behavior that depends on identity may not work as expected. For example, for objects that implement `=` using `==`, the re-activated object will not be `=` to the original.

The class `PassiveObject` implements the method `passivate: anObject toStream: aGsFileOrStream` to write objects out to a stream or a file. To write the object `AllEmployees`

out to the file *allEmployees.obj* in the current directory, execute an expression of the form shown in Example 13.1.

### Example 13.1

---

```
| empFile |
empFile := GsFile openWriteOnServer: 'allEmployees.obj'.
PassiveObject passivate: AllEmployees toStream: empFile.
empFile close.
```

---

The class `PassiveObject` implements the method `newOnStream: aGsFileOrStream` to read objects from a stream or file into a repository. The method `activate` then restores the object to its previous form.

The following example reads the file *allEmployees.obj* into a GemStone repository:

### Example 13.2

---

```
| empFile passivatedEmployees |
empFile := GsFile openReadOnServer: 'allEmployees.obj'.
passivatedEmployees := PassiveObject newOnStream: empFile.
AllEmployees := passivatedEmployees activate.
empFile close.
```

---

## JSON

For transmitting data between sockets, GemStone supports serializing objects in JSON format. JSON allows you to serialize the data only in a language-independant human-readable form.

For example, using the `Employee` examples from `Employee` class and `AllEmployees` (page 70):

```
AllEmployees asJson
%
' [{"firstName": "Lee", "lastName": "Smith", "job": "librarian",
  "age": 40}, {"firstName": "Kay", "lastName": "Adams", "job":
  "clerk", "age": 24}, {"firstName": "Al", "lastName": "Jones", "job":
  "busdriver", "age": 40}]'
```

Unlike passivation, the default JSON serialized form does not preserve class names, object identity, dynamic variables, or other objects attributes. A pointer object is represented using `{}` dictionary syntax, with key-value pairs of instance variable name-value at that instance variable; indexable objects are represented using `[]` arrays syntax. Other attributes, such as `sortBlocks` and dynamic instance variables are omitted.

By implementing `printJsonOn:`, you can customize the output of your own classes to provide additional information, as long as you write corresponding code to understand the resulting JSON string.

The `JsonParser` allows you to parse a JSON string, either a string you have created based on a `GemStone` object, or external JSON code. The `Json` parser returns a `Dictionary` or `Array` containing the equivalent of the JSON string.

```
JsonParser parse: AllEmployees asJson
%
anArray( aDictionary( 'firstName'->'Lee', 'age'->40, 'job'->
  'librarian', 'lastName'->'Smith'), aDictionary( 'firstName'
->'Kay', 'age'->24, 'job'->'clerk', 'lastName'->'Adams'),
aDictionary( 'firstName'->'Al', 'age'->40, 'job'
->'busdriver', 'lastName'->'Jones'))
```

To recreate the original object, you can write custom code to set the instance variable values.

For example, implement the following method

```
Employee class >> fromJson: aJsonString
| json n |
n := super new.
json := JsonParser parse: aJsonString.
json keysAndValuesDo: [:key :val |
  n perform: key, ':' with: val
].
^n
```



# Signals and Notifiers

---

This chapter discusses how to communicate between one session and another, and between one application and another.

**Communicating Between Sessions** (page 265)

introduces two ways to communicate between sessions.

**Object Change Notification** (page 266)

describes the process used to enable object change notification for your session.

**Gem-to-Gem Signaling** (page 274)

describes one way to pass signals from one session to another.

**Other Signal-Related Issues** (page 279)

describes performance, signal buffer overflow, and other signal related considerations.

## 14.1 Communicating Between Sessions

Applications that handle multiple sessions often find it convenient to allow one session to know about other sessions' activities. GemStone provides two ways to send information from one current session to another:

▶ **Object change notification**

Reports the changes recorded by the object server. You set your session to be notified when specific objects are modified. Once enabled, notification is automatic, but a signal is not sent until the changed objects are committed.

▶ **Gem-to-Gem signaling**

Reports events that happen independent of the transaction space. Currently logged-in users signal to send messages to each other. Gems can also pass information that is not necessarily visible to users, such as the name of a queue that needs servicing. Sending a signal requires a specific action by the other Gem; it happens immediately.

Object change notification and Gem-to-Gem signals only reach logged-in sessions. For applications that need to track processes continuously, you can create a Gem that runs

independently of the user sessions and monitors the system. See the instructions on creating a custom Gem in the *GemBuilder for C* manual.

## 14.2 Object Change Notification

Object change notifiers are signals that can be generated by the object server to inform you when specified objects have changed. You can request that the object server inform you of these changes by adding objects to your *notify set*.

When a reference to an object is placed in a notify set, you receive notification of all changes to that object (including the changes you commit) until you remove it from your notify set or end your GemStone session. The notification you receive can vary in form and content, depending on which interface to GemStone you are running and how the notification action was defined.

Your application can respond in several ways:

- ▶ Prompt users to abort or commit for an updated image.
- ▶ Log the information in an object change report.
- ▶ Use the notifiers to trigger another action. For example, a package for managing investment portfolios might check the stock that triggered the notifier and enter a transaction to buy or sell if the price went below or above preset values.

To set up a simple notifier for an object:

1. Create the object and commit it to the object server.
2. Add the object to your session's notify set with one of the messages:

```
System addToNotifySet: aCommittedObject  
System addAllToNotifySet: aCollectionOfCommittedObjects
```

3. Define how to receive the notifier with either a notifier message or by polling.
4. Define what your session will do upon receiving the notifier.

The following section describes each of these steps in detail.

### Setting Up a Notify Set

GemStone defines a notify set for each user session to which you add or remove objects. Except for a few special cases discussed later, any object you can refer to can be added to a notify set.

Notify sets persist through transactions, living as long as the GemStone session in which they were created. When the session ends, the notify set is no longer in effect. If you need notification regarding the same objects for your next session, you must once again add those objects to the notify set.

### Adding an Object to a Notify Set

To add an object to your notify set, use an expression of the form:

```
System addToNotifySet: aCommittedObject
```

When you add an object to the notify set, GemStone begins monitoring changes to it immediately.

Most GemStone objects are composite objects, made up of a root object and a few subobjects. Usually you can just ignore the subobjects. However, there are circumstances in which the both the root object and subobjects must appear in the notify set. For details, see "Special Classes" on page 272.

Example 14.1 creates a collection of stock holdings and then creates a notify set for the stocks in the collection. Finally, the session is set to automatically receive the notifier.

#### Example 14.1

```
"Create a Class to record stock name, number and price"
Object subclass: #Holding
  instVarNames: #('name' 'number' 'price')
  classVars: #()
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: Published.

"Compile accessing methods"
Holding compileAccessingMethodsFor: Holding instVarNames.

"Add a Collection for Holdings to UserGlobals dictionary"
UserGlobals
  at: #MyHoldings put: IdentityBag new.

"Add some stocks to my collection"
MyHoldings add:
  (Holding new name: #USSteel; number: 1000; price: 50.00).
MyHoldings add:
  (Holding new name: #VMware; number: 50000; price: 95.00).
MyHoldings add:
  (Holding new name: #ATT; number: 100000; price: 30.00).

"Add the collection object to the notify set"
System addToNotifySet: MyHoldings.
(System notifySet) includesIdentical: MyHoldings.

"Enable receipt of signals"
System enableSignaledObjectsError.
```

#### Objects That Cannot Be Added

Not every object can be added to a notify set. Objects in a notify set must be visible to more than one session; otherwise, other sessions could not change them. So, objects you have created for temporary use or have not committed cannot be added to a notify set. GemStone responds with an error if you try to add such objects to the notify set.

You also receive an error if you attempt to add special objects, such as true, false, nil, and instances of Character, SmallInteger and SmallDouble.

## Adding a Collection to a Notify Set

To add a collection of objects to your notify set, use an expression like this:

```
System addAllToNotifySet: aCollectionOfCommittedObjects
```

This expression adds the elements of the collection to the notify set.

You don't have to add the collection object itself, but if you do, use `addToNotifySet:` rather than `addAllToNotifySet:`. When a collection object is in the notify set, adding elements to the collection or removing elements from it trigger notification. Modifications to the elements do not trigger notification on the collection object; if you want to know when the elements change, you must add them to the notification set.

Example 14.2 shows the notify set containing both the collection object and the elements in the collection.

### Example 14.2

---

```
"Add the stocks in the collection to the notify set"
System addAllToNotifySet: MyHoldings.
System notifySet.
%
an Array
#1 a Holding
#2 a Holding
#3 a Holding

"Add the collection object itself to the notify set"
System addToNotifySet: MyHoldings.
System notifySet.
%
an Array
#1 a Holding
#2 a Holding
#3 a Holding
#4 an IdentityBag
```

---

## Very Large Notify Sets

You can register any number of objects for notification, but very large notify sets can degrade system performance. GemStone can handle thousands of objects without significant impact. Beyond that, test whether the response times are acceptable for your application.

If performance is a problem, you can set up a different system of change recording:

1. Have each session maintain its own list of the last several objects updated (a modify list). The list is a collection written only by that session.
2. Create a global collection of collections that contains every session's list of changes.
3. Put the global collection and its elements in your notify set, so you receive notification when a session commits a modified list of changed objects. Then you can check for changes of interest.

If the modify lists are ordered, this preserves the order of the additions, so that the new objects can be serviced in the correct order. Using the `notifySet`, notification on a batch of changed objects is received in OOP order.

### Listing Your Notify Set

To determine the objects in your notify set, execute:

```
System notifySet
```

### Removing Objects From Your Notify Set

To remove an object from your notify set, use an expression of the form:

```
System removeFromNotifySet: anObject
```

To remove a collection of objects from your notify set, use an expression of the form:

```
System removeAllFromNotifySet: aCollection
```

This expression removes the elements of the collection. If the collection object itself is also in the notify set, remove it separately, using `removeFromNotifySet: .`

To remove all objects from your notify set, execute:

```
System clearNotifySet
```

## Notification of New Objects

In a multi-user environment, objects are created in various sessions, committed, and immediately open to modification. It may not be sufficient to receive notifiers on the objects that existed at the beginning of your session. You may also need notification concerning new objects.

You cannot put unknown objects in your notify set, but you can create a collection for those kinds of objects and add that collection to the notify set. Then when the collection changes, meaning that objects have been added or removed, you can stop and look for new objects. For example, to receive notification when the price of any stock in your portfolio changes, you can perform the following steps:

1. Create a globally known collection (for example, `MyHoldings`) and add your existing stock holdings (instances of class `Holding`) to it.
2. Place all of these stocks in your notify set:

```
System addAllToNotifySet: MyHoldings
```

3. Place the collection `MyHoldings` in your notify set, so that you receive notification that the collection has changed when a stock is bought or sold:

```
System addToNotifySet: MyHoldings
```

4. Place new stock purchases in `MyHoldings` by adding code to the instance creation method for class `Holding`.
5. When you receive notification that the contents of `MyHoldings` have changed, compare the new `MyHoldings` with the original.
6. When you find new stocks, add them to your notify set, so that you will be notified if they are changed.

Example 14.3 shows one way to do steps 5 and 6.

---

**Example 14.3**

---

```
"Make a temporary copy of the set."  
  
| tmp newObjs |  
tmp := MyHoldings copy.  
  
"Refresh the snapshot view (commit or abort)."  
System commitTransaction.  
  
"Get the difference between the old and new sets."  
newObjs := (MyHoldings - tmp).  
  
"Add the new elements to the notify set."  
newObjs size > 0 ifTrue: [System addAllToNotifySet: newObjs].
```

---

You can also identify objects to remove from the notify set by doing the opposite operation:

```
tmp - MyHoldings
```

This method could be useful if you are tracking a great many objects and trying to keep the notify set as small as possible.

Note that only IdentityBag and its subclasses understand "-" as a difference operator.

### Receiving Object Change Notification

After a commit, each session's snapshot view is updated. The object server also updates its list of committed objects. This list of objects is compared with the contents of the notify set for each session, and a set of the changed objects for each notify set is compiled.

You can receive notification of committed changes to the objects in your notify set in two ways:

- ▶ Enabling automatic notification, which is faster and uses less CPU
- ▶ Polling for changes

## Automatic Notification of Object Changes

For automatic notification, you enable your session to receive the exception `ObjectsCommittedNotification`. By default, `ObjectsCommittedNotification` is disabled (except in `GemBuilder` for `Smalltalk`, which enables the signal as part of `GbsSession>>notificationAction:`).

To enable the event signal for your session, execute:

```
System enableSignaledObjectsError
```

To disable the event signal, send the message:

```
System disableSignaledObjectsError
```

To determine whether this error message is enabled or disabled for your session, send the message:

```
System signaledObjectsErrorStatus
```

This method returns true if the signal is enabled, and false if it is disabled.

This setting is not affected by commits or aborts. It remains until you change it, you end the session, or you receive the signal. The signal is automatically disabled when you receive it so that the exception handler can take appropriate action.

The receiving session handles the notification with an exception handler. Your exception handler is responsible for reading the set of signaled objects (by sending the message `System class>>signaledObjects`) as well as taking the appropriate action.

`ObjectsCommittedNotification addDefaultHandler:`

```
[ :ex |
  | changes |
  changes := System signaledObjects.
  "do something with the changed objects"
  System enableSignaledObjectsError].
```

## Reading the Set of Signaled Objects

The `System class>>signaledObjects` method reads the incoming changed object signals. This method returns an array, which includes all the objects in your notify set that have changed since the last time you sent `signaledObjects` in your current session. The array contains objects changed and committed by all sessions, including your own. If more than one session has committed, the OOPs are OR'd together. The elements of the array are arranged in OOP order, not in the order the changes were committed. If none of the objects in your notify set have been changed, the array is empty.

Use a loop to call `signaledObjects` repeatedly, until it returns an empty collection. The empty collection guarantees that there are no more signals in the queue.

Also see the discussion on "Frequently Changing Objects" on page 272.

## Polling for Changes to Objects

You also use `System class>>signaledObjects` to poll for changes to objects in your notify set.

Example 14.4 uses the polling method to inform you if anyone has added objects to a set or changed an existing one. Notice that the set is created in a dictionary that is accessible to other users, not in `UserGlobals`.

---

**Example 14.4**

```
System disableSignaledObjectsError;
    signaledObjectsErrorStatus.
%

"Create a set."
Published at: #Changes put: IdentitySet new.
System commitTransaction.

System addToNotifySet: Changes.
%

"Login a separate session to perform the following"
Changes add: 'here is a change'.
System commitTransaction
%

"In the original session, see the signal"
| mySignaledObjs count |
System abortTransaction.
count := 0 .
[ mySignaledObjs := System signaledObjects.
mySignaledObjs size = 0 and:[ count < 50]
]
whileTrue: [
    System sleep: 10 .
    count := count + 1
].
^ mySignaledObjs.
%
```

---

## Troubleshooting

Notification on object changes may occasionally produce unexpected results. The following sections outline areas of concern.

### Frequently Changing Objects

If users are committing many changes to objects in your notify set, you may not receive notification of each change. You might not be able to poll frequently enough, or your exception handler might not process the errors it receives fast enough. In such cases, you can miss some intermediate values of frequently changing objects.

### Special Classes

Most GemStone objects are composite objects, but for the purposes of notification you can usually ignore this fact. They are almost always implemented so that changes to subobjects affect the root, so only the root object needs to go into the notify set.

Common operations that trigger notification on the root object include:

- ▶ Assignment to an instance variable:

```
name := 'dowJones'
```

- ▶ Updating the indexable portion of an object:

```
self at: 3 put: 'active'.
```

- ▶ Adding to a collection:

```
self add: 3.
```

In a few cases, however, the changes are made only to subobjects. For the following GemStone kernel classes, both the object and the subobjects must appear in the notification set:

- ▶ RcQueue
- ▶ RcIdentityBag
- ▶ RcCounter
- ▶ RcKeyValueDictionary

You can also have the problem with your own application classes. Wherever possible, you should implement objects so that changes modify the root object. You must also balance the needs of notification with potential problems of concurrency conflicts.

If you are not being notified of changes to a composite object in your notify set, look at the code and see which objects are actually modified during common operations such as `add:` or `remove:`. When you are looking for the code that actually modifies an object, you may have to check a lower-level method to find where the work is performed.

Once you know the object's structure and have discovered which elements are changed, add the object and its relevant elements to the notify set. For cases where elements are known, you can add them just like any other object:

```
System addToNotifySet: anObject
```

Example 14.5 shows a method that creates an object and automatically adds it to the notify set in the process.

#### Example 14.5

---

```
method: SetOfHoldings
add: anObject
    System addToNotifySet: anObject.
    ^super add: anObject
%
```

---

## Methods for Object Notification

Methods related to notification are implemented in class `System`. Browse the class `System` and read about these methods:

```
addAllToNotifySet:  
addToNotifySet:  
clearNotifySet  
disableSignaledObjectsError  
enableSignaledObjectsError  
notifySet  
removeAllFromNotifySet:  
removeFromNotifySet:  
signaledObjects  
signaledObjectsErrorStatus
```

See Chapter 15 for more on handling Exceptions such as `ObjectsCommittedNotification`.

## 14.3 Gem-to-Gem Signaling

`GemStone` enables you to send a signal from your Gem session to any other current Gem session. The exception class `InterSessionSignal` implements several methods for communicating between two sessions. Unlike object change notification, inter-session signaling operates on the event layer and deals with events that are not being recorded in the repository. Signaling happens immediately, without waiting for a commit.

An application can use signals between sessions for situations like a queue, when you want to pass the information quickly. Signals can also be a way for one user who is currently logged in to send information to another user who is logged in.

*NOTE*

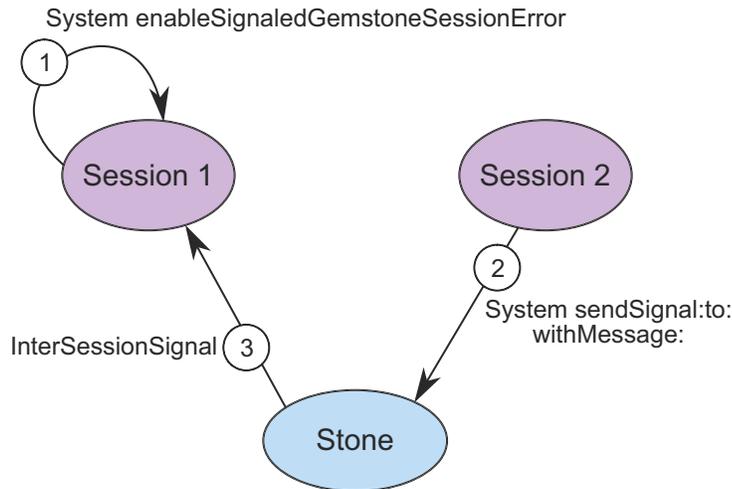
*A signal is not an interrupt, and it does not automatically awaken an idle session. The signal can be received only when your session is actively executing Smalltalk code.*

You can receive a signal from another session by polling for the signal or by receiving automatic notification.

As an example of Gem-to-Gem signaling, Figure 14.1 shows the following sequence of events:

1. `session1` enables event signals from other Gem sessions. (For details, see “Receiving a Notification” on page 278.)
2. `session2` sends a signal to `session1`. (See “Sending a Signal” on page 275.)
3. The Stone sends the exception `InterSessionSignal` to `session1`. The receiving session processes the signal with an exception handler. For details, see Chapter 15, “Handling Exceptions”.

Figure 14.1 Communicating from Session to Session



## Sending a Signal

To communicate, one session must send a signal and the receiving session must be set up to receive the signal.

### Identifying the Receiving Session

To send a signal to another Gem session, you must know its session ID.

To see a description of sessions that are currently logged in, execute the following method:

```
System currentSessions
```

This message returns an array of SmallIntegers representing session IDs for all current sessions. Example 14.6 shows how you might use this method to find the session ID for user1 and send a message.

### Sending the Message

When you have the session ID, you can use the method

```
InterSessionSignal class
  sendSignal: aSignalNumber
  to: sessionId
  withMessage: aMessage.
```

- ▶ *aSignalNumber* is determined by the particular protocol you arranged at your site and the specific message you wish to send. Sending the integer "1," for example, doesn't convey a lot unless everyone has agreed that "1" means "Ready to trade." Your application may use these for specific types of signals, to indicate priority, or whatever other needs your application has; meaning can also be encoded in the message argument. You may create an application-level symbol dictionary of meanings for the different signal numbers.
- ▶ *sessionId* is the sessionId of the session to which the signal will be sent.

▶ *aMessage* is a String object with up to 1023 characters.

### Example 14.6 Sending a Signal

```
| sessionId serialNum otherSession signalToSend |
sessionId := System currentSessions
  detect:[:each |(((System descriptionOfSession: each) at: 1)
  userId = 'TradeUser') ]
  ifNone: [nil].
sessionId ifNotNil: [
  InterSessionSignal
  sendSignal: 4
  to: sessionId
  withMessage: 'reinvest form is here'.
]
```

### Legacy protocol to send signals

While `InterSessionSignal` is recommended, you may also use

```
System sendSignal: aSignalNumber to: sessionId withMessage: aMessage
```

There are a number of legacy/deprecated ways to signal sessions. These use the serial number of a session rather than the `sessionId`, or a `GsSession` instance.

To determine the serial number of a session based on the `sessionId`, use the following expression:

```
GsSession serialOfSession: sessionId
```

To loop up a `GsSession`:

```
GsSession sessionWithSerialNumber: sessionSerialNumber
```

Signalling another session can be done using expressions such as:

- ▶ `GsInterSessionSignal signal: sessionSerialNumber message: aMessage`
- ▶ `signalToSend sendToSession: otherGsSession`
- ▶ `otherGsSession sendSignalObject: signalToSend`

No matter how the message is sent, the other session needs to receive it, as shown in Example 14.7.

### Receiving a Signal

Once the signal is sent, the other session needs to retrieve it. For example, after sending the signal as described in Example 14.6:

#### Example 14.7 Receiving a signal

```
| theSignal |
theSignal := InterSessionSignal poll.
theSignal ifNotNil: [
  theSignal sentInt asString, '-', theSignal sentMessage].
%
```

---

4-reinvest form is here

---

You can receive a signal from another session in either of two ways: you can poll for such signals, or you can enable notification from GemStone. Signals are queued in the receiving session in the order in which they were received. If the receiving session has inadequate heap space for an incoming signal, the contents of the signal is written to *stdout*, whether the receiving session has enabled receiving such signals or not.

The method `InterSessionSignal poll` reads the incoming signals, whether you poll or receive a signal. If there are no pending signals, the array is empty.

Use a loop to call `InterSessionSignal poll` repeatedly, until it returns a nil. This guarantees that there are no more signals in the queue. If signals are being sent quickly, you may not receive a separate `InterSessionSignal` for every signal. Or, if you use polling, signals may arrive more often than your polling frequency.

## Polling

To poll for signals from other sessions, send the following message as often as you require:

```
InterSessionSignal >> poll
```

This method returns nil if there is no signal waiting. Example 14.8 shows how to poll for Gem-to-Gem signals. If the polling process finds a signal, it returns; otherwise it sleeps for 10 seconds and tries again. This example includes a limit such that it stops after 10 minutes, if no signal arrives; polling, of course, must continue running.

### Example 14.8 Polling

---

```
| theSignal loopLimit |
loopLimit := 0.
[ theSignal := InterSessionSignal poll.
  theSignal notNil and:[ loopLimit < 60 ]
] whileTrue: [
  System sleep: 10.
  loopLimit := loopLimit + 1
].
^theSignal
```

---

## Polling and Signalling

Rather than fetching a signal and then handling it, you can setup a poll that sends a signal when a signal arrives. For example

### Example 14.9 Polling and Signalling

```
[ 60 timesRepeat: [
  InterSessionSignal pollAndSignal.
  System sleep: 10 ]
] on: InterSessionSignal
  do: [:ex |
  "perform handling for the signal"
  ex resume]
```

## Legacy Polling protocol

Older application may use this deprecated method to poll for signals:

```
System signalFromGemStoneSession
```

If no signal has been sent, this method returns an empty array. If a signal has been sent, this method returns a four-element array containing:

- ▶ An instance of GsSession representing the session that sent the signal. (equivalent to GsSession sessionWithSerialNumber: (anInterSessionSignal sendingSessionSerial).
- ▶ The signal value (a SmallInteger) (equivalent to anInterSessionSignal sentInt)
- ▶ The string containing the signal message. (equivalent to anInterSessionSignal sentMessage)
- ▶ The number of additional pending signals.(equivalent to anInterSessionSignal numPending)

## Receiving a Notification

To use the exception mechanism to receive signals from other Gem sessions without polling, you must enable receipt of the InterSessionSignal notification. By default, the InterSessionSignal notification is disabled, except in the GemBuilder for Smalltalk interface, which enables the error as part of GbsSession>>gemSignalAction:.

When enabled or disabled, the setting is not affected by commits or aborts. It remains until you change it, you end the session, or you receive the error.

To enable this exception, execute:

```
InterSessionSignal enableSignalling
```

To disable the exception, send the message:

```
InterSessionSignal disableSignalling
```

To determine whether receiving this exception is presently enabled or disabled, send the message:

```
InterSessionSignal signallingEnabled
```

If you have enabled signalling on `InterSessionSignal`, it is automatically disabled when you receive it, so that the exception handler can take appropriate action without further interruption. You must re-enable it afterwards; usually this is the final step in the block that handles the signal.

The following legacy protocol performs the same actions:

```
System enableSignaledGemStoneSessionError
System disableSignaledGemStoneSessionError
System signaledGemStoneSessionErrorStatus
```

## 14.4 Other Signal-Related Issues

GemStone notifiers and Gem-to-Gem signals use the same underlying implementation. The following performance and other considerations apply when using either mechanism.

### Inactive Gem

Receiving the signal can be delayed, if the session is inactive. GemStone is not an interrupt-driven application programming interface. It is designed to make no demands on the application until the application specifically requests service. Therefore, Gem-to-Gem signals and object change notifiers are not implemented as interrupts, and they do not automatically awaken an idle session. They can be received only when GemBuilder is running, not when you are running client code, sitting at the Topaz prompt, waiting for activity on a socket, or waiting on a semaphore (as for a child process to complete). The signals are queued up and wait until you read them, which can create a problem with signal overflow if the delay is too long and the signals are coming rapidly.

You can receive signals at reliable intervals by regularly performing some operation that activates GemBuilder. For example, in a GemStone Smalltalk application, you could set up a polling process that periodically sends out `GbsSession>>pollForSignal`. The `pollForSignal` method causes GemBuilder for Smalltalk to poll the repository. GemBuilder for C also provides a function **GciPollForSignal**.

### Dealing With Signal Overflow

Gem-to-Gem signals and object change notification signals are queued separately in the receiving session. The queues maintain the order in which the signals are received.

#### NOTE

*For object change notification, the queue does not preserve the order in which the changes were committed to the repository. Each notification signal contains an array of OOPs, and these changes are arranged in OOP order. See "Receiving Object Change Notification" on page 270.*

Each session has a signal buffer that will accommodate 50 signals. Signals remain in the signal buffer until they are received and read by the receiving session. If the receiving session does not read the signals, or if it does not read them fast enough to keep up with signals that are being sent, the signal buffer will fill up. In this case, further signals will cause the Exception `SignalBufferFull` to be signalled on the sender. Set your application so that the sender gracefully handles this error. For example, the sender might try to send the signal five times, and finally display a message of the form:

```
Receiver not responding.
```

The most effective way to prevent signal overflow is to keep the session in a state to receive signals regularly, using the techniques discussed in the preceding section. When you do receive signals, make sure you read all the signals off the queue. Repeat `signaledObjects` or `InterSessionSignal poll` until it returns a nil.

## Sending Large Amounts of Data

If you want to pass large amounts of data between sessions, sockets are more appropriate than Gem-to-Gem signals. Chapter 13, “Operating System Access”, describes the GemStone interface to TCP/IP sockets. That solution does not pass data through the Stone, so it does not create system overload when you send a great many messages or very long ones.

## Maintaining Signals and Notification When Users Log Out

Object change notification and Gem-to-Gem signals only reach logged-in sessions. For applications that need to track processes continuously, you can create a Gem that runs independently of the user sessions and monitors the system. For example, such a Gem can monitor a machine and send a warning to all current sessions when something is out of tolerance. Or it might receive the information that all the users need and store it where they can find it when they log in.

Example 14.10 shows some of the code executed by an error handler installed in a monitor Gem. It traps Gem-to-Gem signals and writes them to a log file.

### Example 14.10

---

```
| theSignal logString |
theSignal := InterSessionSignal poll.
logString := String new.
logString add: '-----'; lf;
    add: 'Signal ';
    add: (theSignal sendInt) asString;
    add: ' received from session ';
    add: (theSignal sendingSessionId) asString;
    add: ' with message: ';
    addAll: (theSignal sentMessage); lf.
(GsFile openWriteOnServer: '$GEMSTONE/gemmessage.txt')
    addAll: logString; close.
```

---

# Handling Exceptions

---

GemStone Smalltalk implements the ANSI exception handling protocols, with provisions for signaling that an exception has occurred and for defining handlers for signaled exceptions.

**The Exception Class Hierarchy** (page 281)

describes the exception class hierarchy, listing the subclasses that correspond to events that you may want to handle.

**Signaling Exceptions** (page 283)

describes the mechanism whereby an application can signal that a some notable event occurred. The class of the signaled exception determines which handler(s) will be invoked. A handler might halt execution and report an error to the user.

**Handling Exceptions** (page 284)

describes how to define handlers in your application to cope with signaled exceptions. Depending on the type of the exception, your application might be able to handle the exception gracefully, possibly even without the user being informed of the exception.

**The Legacy Exception Handling Framework** (page 290)

describes the legacy exception handling framework.

## 15.1 The Exception Class Hierarchy

GemStone/S 64 Bit supports the ANSI Exception framework. The ANSI Exception framework defines subclasses to match the granularity of errors that you may want to handle.

GemStone also supports a Legacy Exception framework, for compatibility with earlier versions of Gemstone. This can be used to signal and handle ANSI exceptions. The Legacy Exception framework is described under “The Legacy Exception Handling Framework” on page 290.

Figure 15.1 shows some of the classes in the ANSI exception handler class hierarchy.

**Figure 15.1 (Partial) Exception Class Hierarchy**

---

```
AbstractException
  Exception
    ControlInterrupt
      Break
      Breakpoint
      ClientForwarderSend
      Halt
      TerminateProcess
    Error
      AlmostOutOfMemoryError
      AlmostOutOfStackError
      CompileError
      EndOfStream
      ExitClientError
      ExternalError
        CryptoError
        IOError
          SocketError
            SecureSocketError
      GciError
      ImproperOperation
        ArgumentError
        ArgumentTypeError
        CannotReturn
        LookupError
        OffsetError
        OutOfRange
          FloatingPointError
      IndexingErrorPreventingCommit
      InternalError
        GciTransportError
      LockError
      MigrationError
      NameError
        MessageNotUnderstood
      NumericError
        ZeroDivide
      RepositoryError
      SecurityError
      SignalBufferFull
      TransactionError
      UncontinuableError
      UserDefinedError
    Notification
      Admonition
        AlmostOutOfMemory
        AlmostOutOfStack
        RepositoryViewLost
      Deprecated
      FloatingPointException
      InterSessionSignal
      TransactionBacklog
      Warning
        CompileWarning
```

---

For the complete set of Exception classes, refer to the image.

## 15.2 Signaling Exceptions

ANSI Exceptions are *class-based*: you use a class in the Exception hierarchy to describe errors and other exceptions in your GemStone Smalltalk programs.

You can extend the built-in exception types by defining new subclasses. You can also change your new exception's default behavior by adding method overrides to the new class (for example, `defaultAction` and `isResumable`).

The ANSI exception handling framework provides for zero or more dynamic (stack-based) handlers and a list of zero or more default handlers, ordered in the sequence they were installed.

When an application sends a message of the form:

```
Exception signal: aString
```

GemStone Smalltalk creates an *instance* of the signaled class and performs the following search for a suitable handler:

1. Search the stack for a handler associated with the exception class. In a dynamic (stack-based) handler (described on page 284), you explicitly identify a block of application code that might signal an exception to which you wish to respond.
2. Search the default (static) handlers. A default handler (described on page 288) is invoked if a dynamic handler is not found or if the last dynamic handler passes the exception.
3. Search the exception class for an implementation of the instance method `defaultAction`. Some exception classes redefine this method, thereby establishing a handler to use in the case that there is no suitable dynamic or default handler or if the last such handler passes the exception. For example, with `Notification`, the default action is to ignore the exception.

If the exception class does not override the implementation of `defaultAction` in class `AbstractException`, halt the GemStone Smalltalk interpreter and pass the exception back to the client to be handled (by `Topaz`, `GemBuilder`, or another application) as an error.

### Example 15.1

---

```
topaz 1> method: Employee
age: anInt
(anInt between: 15 and: 65)
  ifFalse: [Error signal: 'Employee age out of range'].
age := anInt.
%
```

---

## 15.3 Handling Exceptions

Other than a few fatal errors, most signaled exceptions can be handled in your GemStone Smalltalk application. To do so, you identify the type of exception that might be signaled (Exception or, more often, a subclass of Exception) and provide GemStone Smalltalk code to handle the exception.

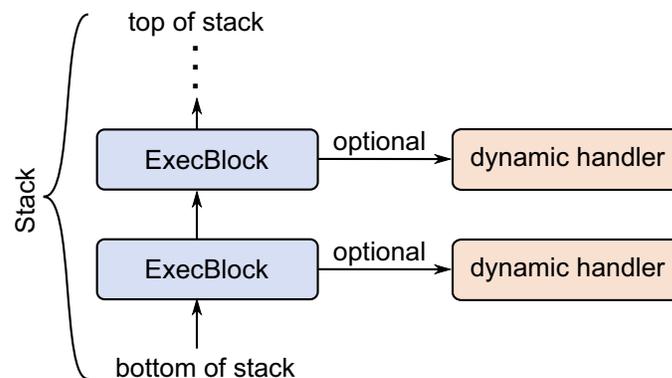
GemStone Smalltalk allows you to define two kinds of exception handlers: *dynamic (stack-based) handlers* and *default (static) handlers*.

### Dynamic (Stack-Based) Handlers

A dynamic (stack-based) handler is associated with an executable block (instance of ExecBlock) and the associated state in which the GemStone Smalltalk virtual machine is presently executing. These handlers live and die with their associated blocks – when the block is exited, the handler is gone.

A dynamic handler is associated with exactly one ExecBlock and applies as long as the ExecBlock is being executed. Because an ExecBlock can be embedded in another ExecBlock (either directly or via another method), multiple dynamic handlers can be active at one time. Figure 15.2 illustrates this relationship.

Figure 15.2 ExecBlock and Associated Handlers



To define a dynamic handler for an ExecBlock, send the `on:do:` message to the block. Example 15.2 defines an `averagePay` method for the `Employee` class. The method calculates an average by dividing two values. If the division signals a `ZeroDivide` exception, the exception handler returns zero as the result of the method. In this implementation, the method will never result in a “division by zero error” being seen by the user. (Of course, there are other ways you might write this particular method. This example simply serves to highlight the `on:do:` exception handling approach.)

---

**Example 15.2**

---

```
method: Employee
averagePay

[
  ^self totalPay / self yearsOfService.
] on: ZeroDivide do: [:ex |
  ^0.
].

%
```

---

The first argument to the `on:do:` method specifies what types of exception the handler should catch. The argument can be a class in the Exception hierarchy, or it can be an `ExceptionSet` made up of one or more classes in the Exception hierarchy.

The second argument specifies a one-argument `ExecBlock` that will be invoked when the specified exception is signaled. The one argument is the newly-created instance of the class of the exception that was signaled, and can contain additional information about the exception (including the string that was passed to the `signal:` method). For example, an instance of the `ZeroDivide` error can be queried for the dividend (obviously, the divisor is zero). Similarly, an instance of the `MessageNotUnderstood` error can be queried for the receiver and message (selector and arguments).

## Selecting a Handler

When an exception is signaled, GemStone starts at the top of the current process's stack, searching down the stack for a handler that handles the exception. Each exception handler in the stack is examined to see if it was installed (using the `on:do:` message) as a handler for the signaled exception's class. If a handler is found but it does not handle the signaled exception, it is passed over and the search continues down the stack.

A handler for a superclass will handle subclass exceptions. That is, an exception handler for the class `Error` will be invoked for an exception of its subclass `ZeroDivide`, and an exception handler for the class `Notification` will be invoked for an exception of its subclass `Warning`.

A subclass does not, however, handle a superclass exception. This means that an exception handler for the class `MessageNotUnderstood` will not be invoked for an exception of its superclass `Error`.

Example 15.3 contains six blocks, three protected blocks and three handler blocks. Each of the three `on:do:` messages creates a new stack frame that has an associated handler block.

---

**Example 15.3**

```

method: Employee
doStuff

| a b c |
a := [
  self doStuffA.
  b := [
    self doStuffB.
    c := [
      self doStuffC.
      self doStuffD.
    ] on: ZeroDivide do: [:zdx |
      self handleZeroDivide: zdx.
      ^self.
    ].
    self doStuffE.
  ] on: Warning do: [:wEx |
    self handleWarning: wEx.
    wEx resume: #ok.
  ].
  self doStuffF.
  #good.
] on: Error do: [:erEx |
  self handleError: erEx.
  erEx return: #bad.
].

%

```

---

As shown in Figure 15.3, the handler for Error is installed first, and catches any Error or subclass exception signaled during the block that begins with `self doStuffA`.

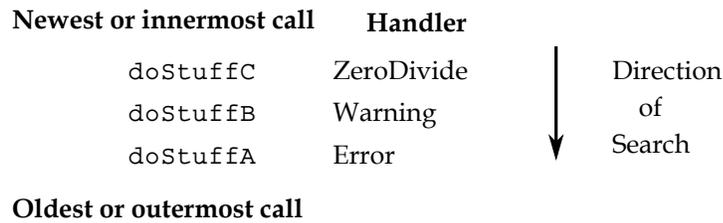
The handler for Warning is installed next, and catches any Warning or subclass exception signaled during the block that begins with `self doStuffB`.

If a ZeroDivide error is signaled during `doStuffB`, it is handled by the Error handler, not by the ZeroDivide handler (which is not yet installed).

The handler for ZeroDivide is installed last, and catches any ZeroDivide error or subclass exception signaled during the block that begins with `self doStuffC`.

If a MessageNotUnderstood error were signaled during `doStuffC`, it would not be handled by either the ZeroDivide or Warning handler, even though they were installed more recently. Those handlers are not of the proper class; MessageNotUnderstood does not inherit from ZeroDivide or Warning. Instead, a MessageNotUnderstood error would be handled by the Error handler associated with the block that begins with `self doStuffA`.

Figure 15.3 Selecting a Handler



## Flow of Control

Once control is passed by sending `value:` to the handler block with the exception instance as an argument, the handler block can attempt to address the situation.

Keep in mind that a dynamic handler is just an `ExecBlock` that is defined in a method and passed as an argument during a message send (like a block sent with a `select: message`). As such, the dynamic handler has access to the method context in which it is defined, including method temporaries and block variables in its scope, as well as the object in which the method is defined (including instance variables). The handler may, of course, send messages to any object to which it has access.

In particular, the dynamic handler may return from the method containing the dynamic handler. In Example 15.3 on page 286, the `ZeroDivide` handler returns `self`. If a `ZeroDivide` exception were signaled during `doStuffC`, then the `doStuff` method would return and other messages would never be sent (`doStuffD`, `doStuffE`, and `doStuffF`).

### Messages That Alter the Flow of Control

In addition to an explicit return from the containing method, a dynamic handler can send the following messages to the exception instance to cause other changes in the flow of control. Sending one of these messages is similar to a method return in that there is no return from these messages (except for `outer`, which might return).

`resume: anObject`

Causes *anObject* to be returned as the result of the `signal:` message that triggered the exception. Sending `resume:` to a non-resumable exception is an error.

In Example 15.3, the `Warning` handler returns `#ok` as the result of the `signal:` message.

`resume`

Causes `nil` to be returned as the result of the `signal:` message. Sending `resume` to a non-resumable exception is an error.

`return: anObject`

Causes *anObject* to be returned as the result of the `on:do:` message to the protected block. In Example 15.3, the `Error` handler returns `#bad` to the local variable 'a' as the result of the `on:do:` message. If no `Error` occurred during the protected block, then the `on:do:` method would return `#good` as the result of evaluating the protected block.

`return`

Causes `nil` to be returned as a result of the `on:do:` message.

`retry`

Unwinds the stack and re-evaluates the protected block (by sending the `on:do:` message again).

`retryUsing: aBlock`

Unwinds the stack and evaluates the replacement block as the protected block, sending it the `on:do:` message.

`pass`

Exits the current handler and searches for the next handler. In Example 15.3, if the `ZeroDivide` handler sends `pass` to the `ZeroDivide` exception instance, control passes to the `Error` handler as if the `ZeroDivide` handler didn't exist (except that any side effects of its operation up to the `pass` message are preserved).

`outer`

Similar to `pass`, except that if the outer handler sends `resume:` or `resume` to the exception instance, control returns to the inner handler from the `outer` message.

`resignalAs: replacementException`

Sending this message causes GemStone Smalltalk to start searching for an exception handler for `replacementException` at the top of the stack as if the original `signal:` message had been sent to `replacementException` instead of the receiver.

#### NOTE

*If none of the above messages are sent to alter the flow of control, the value of the last expression in the block will be returned as the result of the `on:do:` message. (For clarity, you could make this behavior explicit by using the `return:` message.)*

## Default Handlers

As described above, a dynamic (stack-based) handler protects a particular block of code that exists in the same method as the handler. This is appropriate when you only want to handle a particular exception during execution of the protected code. When the protected block finishes executing, the handler is no longer in effect.

There are, however, other exceptions that could happen at any time for reasons entirely unrelated to your code — for example, being notified that the disk is full (`RepositoryError`) or that another Gem is sending you a signal (`InterSessionSignal`). For such exceptions, you can establish a default (or static) handler.

Since ANSI does not provide a direct API for adding and removing default handlers at runtime, GemStone provides the following methods to deal with default handlers in the context of the ANSI framework.

`Exception class >> addDefaultHandler: aOneArgumentBlock`

Returns a `GsExceptionHandler` that understands the message `remove` and adds the new handler to the beginning of the `defaultHandlers` list. After `aOneArgumentBlock` (equivalent to the second argument to `on:do:`) is invoked, the argument (an instance of `Exception` or one of its subclasses) responds appropriately to `pass` and `outer` seamlessly between stack-based and default handlers.

`AbstractException class >> defaultHandlers`

Returns a `SequenceableCollection` (or subclass) of `GsExceptionHandler` instances that will catch instances of the receiver (typically, a subclass of `AbstractException`). The result does not include any legacy static handlers (as discussed on page 291). This collection may be empty and typically is a subset of the installed default (static) handlers.

`GsExceptionHandler >> remove`

Since a default handler is not tied to a specific block of code, once installed it remains in effect until explicitly removed (or until the session logs out). This method removes (and returns) the default handler if it is found. If it is not found, returns `nil`.

For example, the following code sets a default handler for an `AlmostOutOfMemoryError`. Note that this will handle the almost out of memory by performing a commit, which could commit unexpected, when the asynchronous signal occurs; and that commit may not free up memory and avoids entirely running out of memory.

---

#### Example 15.4 Set a default handler for almost out of memory

---

```
AlmostOutOfMemoryError addDefaultHandler: [:ex |
    GsFile gciLogClient: ex messageText.
    System commit.
    AlmostOutOfMemoryError enable.
    ex resume].
AlmostOutOfMemoryError enable.
%
```

---

See also the table on page 292, and the example on page 324.

## Default Actions

The third line of defense for an exception (after dynamic and default handlers) occurs when the virtual machine sends the message `defaultAction` to the signaled exception. Because `defaultAction` is implemented in `AbstractException`, every exception will eventually be handled. The ultimate default action (in `AbstractException`) is to stop the GemStone Smalltalk interpreter and pass the exception back to the client (to be handled by Topaz, GemBuilder, or another application).

Exception subclasses can override this method to provide alternate behavior. For example, the default action for `Notification` is to ignore the notification and return `nil` from the `signal: message`. For `Deprecated`, the default action is to log information; for `MessageNotUnderstood`, the default action is to retry the original action.

To define a default handler for a new exception, add a `defaultAction` method to your new exception class.

## 15.4 The Legacy Exception Handling Framework

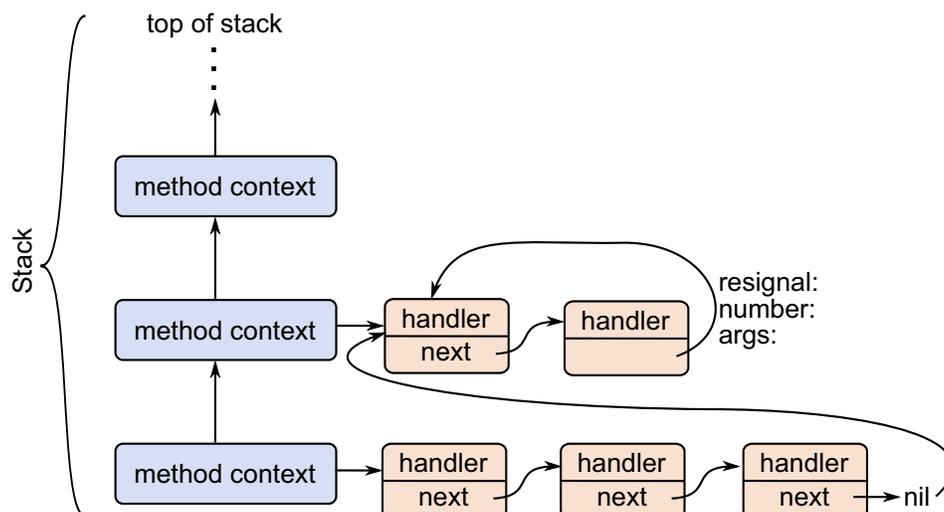
ANSI exception handling, as described previously, is the primary mechanism for dealing with errors in your programs. The legacy handler protocol is deprecated, and all exceptions are now raised as ANSI exceptions. While we strongly encourage the use of ANSI protocol, legacy protocol may be used to raise and handle ANSI exceptions.

### Dynamic (Stack-Based) Exception Handler

In ANSI, a dynamic (stack-based) exception handler is associated with an ExecBlock. By contrast, a dynamic legacy exception handler is associated with a method being executed. These exception handlers live and die with their associated method contexts—when the method returns, control is passed to the next method and the exception handler is gone.

Each exception handler is associated with one method context, but each method context can have a stack of associated exception handlers. The relationship is diagrammed in Figure 15.4.

Figure 15.4 Method Contexts and Associated Handlers



### Installing a Dynamic (Stack-Based) Exception Handler

To define a legacy dynamic (stack-based) handler for an exception, use the class method `Exception category:number:do:`.

- ▶ The argument to the `category:` keyword is ignored.
- ▶ The argument to the `number:` keyword is the specific error number you wish to catch, which can be `nil` (to catch all exceptions).
- ▶ The argument to the `do:` keyword is a four-argument block you wish to execute when the error is raised.
  - ▶ The first argument to the four-argument block is the instance of `Exception` that was signaled.

- ▶ The second argument to the four-argument block is always GemStoneError.
- ▶ The third argument to the four-argument block is an error number.
- ▶ The fourth argument to the four-argument block is the data passed in when invoking the error.

If your exception handler does not specify an error number (an error number of nil), then it receives control in the event of any exception.

The exception handler in Example 15.5 catches the GemStone exception ZeroDivide and returns either PlusInfinity or MinusInfinity, depending on the sign of the dividend.

### Example 15.5

```
| a b c |
a := 0.
Exception
  category: GemStoneError
  number: 2026
  do: [:ex :cat :num :args |
    "Return a value as a result of the #'/' message"
    ex dividend * 1.0e0 / 0].

"This might give a ZeroDivide error,
depending on the value of a"
b := -10 / a.
c := b * 3.
c
```

#### NOTE

*Keep the handler as simple as possible, because you cannot receive any additional errors while the handler executes. Normally your handler should never terminate the ongoing activity and change to some other activity.*

## Default (Static) Exception Handlers

A *default (static) exception handler* is a final line of defense—if you define one, it will take control in the event of any error for which no other handler has been defined. A static exception handler executes without changing in any way the stack, or the return value of the method that called it. Static exception handlers are therefore useful for handling errors that appear at unpredictable times, such as the errors listed in Table 15.1. You can use a static exception handler as you would an interrupt handler, coding it to change the value of some global variable, perhaps, so that you can determine that an error did, in fact, occur.

### Installing a Default (Static) Exception Handler

To define a default (static) exception handler, use the Exception class method `installStaticException:category:number:.`

- ▶ The argument to the `installStaticException:` keyword is the block you wish to execute when the error is raised.
- ▶ The argument to the `category:` keyword is ignored.

- ▶ The argument to the `number` keyword is the specific error number you wish to catch.

The following exception handler, for example, handles the error `#abortErrLostOtRoot`:

### Example 15.6

```
UserGlobals at: #tx3 put:
( "Handle lost OT root"
  Exception
    installStaticException: [:ex :cat :num :args |
      System abortTransaction.
    ]
    category: nil
    number: 3031
    subtype: nil
).
```

To remove the handler, execute:

```
self removeExceptionHandler: (UserGlobals at: #tx3).
```

## GemStone Event Exceptions

The errors in Table 15.1 are sometimes called *event exceptions*. Although they are not true errors, their implementation is based on the GemStone error mechanism. For examples that use these event exceptions, also called signals, see Chapter 14, “Signals and Notifiers”.

In Table 15.1, the legacy error symbol (and number) is listed along with the corresponding current exception class.

### NOTE

*The array `LegacyErrNumMap` (in `Globals`) describes the mapping of legacy (pre-3.0) error numbers to ANSI exception classes (as described in Chapter 15, “Handling Exceptions”).*

**Table 15.1 Common GemStone Event Exceptions**

Exception class Legacy symbol (and number)	Description
TransactionBacklog  <code>#rtErrSignalAbort</code> (6009) <code>#rtErrSignalFinishTransaction</code> (6012)	When <code>System inTransaction</code> returns false (running outside a transaction), Stone requested Gem to abort. This error is generated only if you have executed either <code>System enableSignaledAbortError</code> or <code>TransactionBacklog enableSignalling</code> .  When <code>System inTransaction</code> returns true (the session is in transaction), Stone has requested the session to commit, abort, or continue (with <code>continueTransaction</code> ) the current transaction. This error is received only if you have executed either <code>System enableSignaledFinishTransactionError</code> or <code>TransactionBacklog enableSignalling</code> .

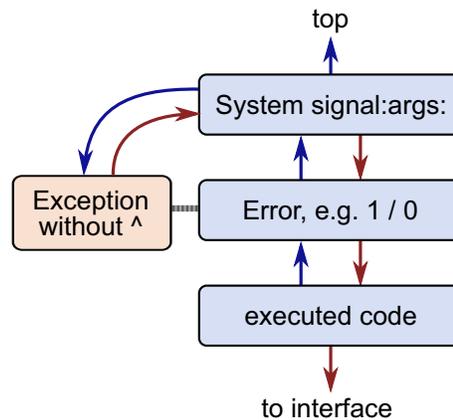
Table 15.1 Common GemStone Event Exceptions

Exception class Legacy symbol (and number)	Description
ObjectsCommittedNotification  <i>#rtErrSignalCommit (6008)</i>	An element of the notify set was committed and added to the signaled objects set. This error is received only if you have executed either <code>System enableSignaledObjectsError</code> or <code>ObjectsCommittedNotification enableSignalling</code>
InterSessionSignal  <i>#rtErrSignalGemStoneSession (6010)</i>	Your session received a signal from another GemStone session. This error is received only if you have executed either <code>System enableSignaledGemstoneSessionError</code> or <code>InterSessionSignal enableSignalling</code> . InterSessionSignal arguments: 1. The session ID of the session that sent the signal. 2. An integer representing the signal. 3. A message string.
AlmostOutOfMemory  <i>#rtErrSignalAlmostOutOfMemory (6012)</i>	Temporary object memory for the session is almost full. The error is deferred if in user action or index maintenance. This error is enabled by default, but has no effect; you must also execute <code>AlmostOutOfMemory enableNotification</code> , or set a default handler. After a signal is received, it must be reenabled using <code>AlmostOutOfMemory enable</code> .
AlmostOutOfMemoryError  <i>(6013)</i>	Temporary object memory for the session is almost full. The error is deferred if in user action or index maintenance. This error is not enabled by default. After a signal is received, it must be reenabled using <code>AlmostOutOfMemoryError enable</code> .
RepositoryError  <i>#rtErrTranlogDirFull (2339)</i>	All available transaction log directories or partitions are full. This error is received if you are DataCurator or SystemUser, otherwise only if you have executed <code>System enableSignalTranlogsFull</code> .
RepositoryViewLost  <i>#abortErrLostOtRoot (3031)</i>	While running outside a transaction, Stone requested Gem to abort. Gem did not respond in the allocated time, and Stone was forced to revoke access to the object table.

## Flow of Control

Exception handlers with no explicit return operate like interrupt handlers – they return control directly to the method from which the exception was raised. You must write all default (static) exception handlers this way, because the stack usually changes by the time they catch an error. Dynamic (stack-based) exception handlers can also be written to behave that way, like the one in Example 15.5 on page 291. See Figure 15.5.

Figure 15.5 Default Flow of Control in Legacy Exception Handlers



Sometimes, however, this is not useful behavior – the application may simply have to raise the same error again. In dynamic (stack-based) exception handlers, it can be useful instead to return control to the method that defined the handler.

You can accomplish this by defining an explicit return (using the return character `^`) in the block that is executed when the exception is raised. For example, the method in Example 15.7 redefines how the GemStone exception `#ZeroDivide` is to be handled.

#### Example 15.7

```

| a b c |
a := 0.
Exception
  category: GemStoneError
  number: 2026
  do: [:ex :cat :num :args |
    "Return from this method with a String"
    ^'zero divide'
  ].

```

"When a is zero, the error will be caught and the method will return without assigning any value to b or c"

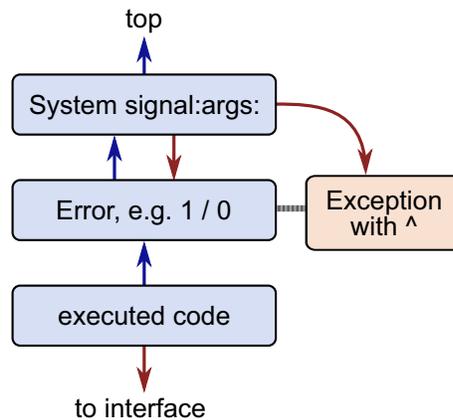
```

b := -10 / a.
c := b * 3.
c

```

Figure 15.6 shows the flow of control in Example 15.7.

Figure 15.6 Dynamic (Stack-Based) Exception Handler with Explicit Return



### Signaling Other Exception Handlers

Under certain circumstances, your exception handler can choose to pass control to a previously defined exception handler, one that is below the present exception handler on the stack. To do so, your exception handler can send the message `resignal: number: args:.`

- ▶ The argument to the `resignal:` keyword is ignored.
- ▶ The argument to the `number:` keyword is the specific error number you wish to signal.
- ▶ The argument to the `args:` keyword is an array of information you wish to pass to the exception handler. This is the array whose elements might be used to build the error message.

### Removing Exception Handlers

You can define an exception so that it removes itself after it has been raised, using the Exception instance method `remove`. In conjunction with the `resignal:` mechanism described in the previous section, `remove` allows you to set up your application so that successive occurrences of the same error (or category of errors) are handled by successively older exception handlers that are associated with the same context.

For example, suppose we execute the following code:

**Example 15.8**


---

```

| x y |
Exception
  category: GemStoneError
  number: 2026
  do: [:ex :cat :num :args | ex remove. 'first result'].
Exception
  category: GemStoneError
  number: 2026
  do: [:ex :cat :num :args | ex remove. 'second result'].
x := 1 / 0. "handled by the second (most recent) handler"
y := 2 / 0. "handled by the first handler; the second was removed"
Array with: x with: y.
%
anArray( 'second result', 'first result')

```

---

The first occurrence of the error executes the most recent exception defined. The exception then removes itself, so that the next occurrence of the same error executes the exception handler stacked previously within the same method context. This exception handler returns an array of two strings, as shown here.

**Recursive Errors**

If you define an exception handler broadly to handle many different errors, and you make a programming mistake in your exception handler, the exception handler may then raise an error that calls itself repeatedly. Such infinitely recursive error handling eventually reaches the stack limit. The resulting stack overflow error is received by whichever interface you are using.

If you receive such an error, check your exception handler carefully to determine whether it includes errors that are causing the problem.

**Raising Exceptions**

Legacy methods for raising exceptions can be used, but raise ANSI exceptions.

To raise an exception, use the class method `System signal:args:signalDictionary:`.

- ▶ The argument to the `signal:` keyword is the specific error number you wish to signal.
- ▶ The argument to the `args:` keyword is an array of information you wish to pass to the exception handler. This is the array whose elements are passed to the handler.
- ▶ The argument to the `signalDictionary:` keyword is ignored.

To raise the generic exception defined for you in `ErrorSymbols` as `#genericError`, use the class method `System genericSignal:text:args:`, or one of its variants.

- ▶ The argument to the `genericSignal:` keyword is an object you can define to further distinguish between errors, if you wish. Alternatively, it can be `nil`.

- ▶ The argument to the `text : keyword` is a string you can use for an error message. It will appear in GemStone's error message when this error is raised. It can be `nil`.
- ▶ The argument to the `args : keyword` is an array of information you wish to pass to the exception handler, as described above.

Other variants of this message are `System genericSignal:text:arg:` for errors having only one argument, or `System genericSignal:text:` for errors having no arguments.

## ANSI Integration

The ANSI and legacy frameworks should work together so that signaling an ANSI exception is caught by a legacy exception handler. Example 15.9 shows a sample use of a legacy handler to catch signaled ANSI exceptions.

### Example 15.9

---

```

method: Employee
  legacyMethod

    self doA.
    "Install a legacy handler"
    Exception
      category: nil
      number: nil
      do: [:ex :cat :num :args |
          self handlerCode.
          self shouldReturn ifTrue: [
              ^self returnValue.
          ].
          self continueValue.
      ].
    self doB.
    "Signal an ANSI error"
    instVar1 := Error signal: 'something bad happened!'.
    self doC.
    ^instVar2.

%
```

---

When this method is invoked, it calls `doA` before installing the exception handler. After the exception handler is installed, the method calls `doB`. If any exception is signaled during the execution of `doB`, the handler is invoked.

Next, an explicit error is invoked, using the ANSI protocol. This signaled ANSI exception is caught by the legacy exception handler installed earlier in the method. After evaluating the `handlerCode`, the handler decides whether to return from the method or continue. If it returns, the result of `returnValue` is returned. If it continues, the result of `continueValue` is stored in `instVar1`, and the method proceeds with `doC` and finally returns `instVar2`.



# Performance and Optimization

---

GemStone Smalltalk includes several tools to help you tune your applications for faster performance.

**Profiling Smalltalk Execution** (page 300)

Profiling tools that allow you to pinpoint the problem areas in your application code.

**Clustering Objects for Faster Retrieval** (page 311)

How to cluster objects that are often accessed together so that many of them can be found in the same disk access.

**Modifying Cache Sizes for Better Performance** (page 320)

How to increase or decrease the size of various caches in order to minimize disk access and storage reclamation.

**Managing VM Memory** (page 322)

Issues to consider when managing temporary object memory, and presents techniques for diagnosing and addressing OutOfMemory conditions.

**NotTranloggedGlobals** (page 327)

Optimize certain operations by avoiding writing tranlog entries.

**Other Optimization Hints** (page 328)

Allow operations on large collections without using temporary object memory.

## 16.1 Profiling Smalltalk Execution

Many things impact performance, and cache size and disk access often have the largest impact on application performance. However, your GemStone Smalltalk code can also affect the speed of your application. There are a number of tools to help you identify issues and optimize your code.

### Time to execute a block

If you simply want to know how long it takes a given block to return its value, you can use GemStone Smalltalk methods that execute a block and return a number.

#### CPU Time

The familiar method `System class >> millisecondsToRun:` takes a zero-argument block as its argument and returns the time in milliseconds required to evaluate the block.

```
topaz 1> printit
System millisecondsToRun: [
  System performOnServer: 'ping -c1 gemtalksystems.com']
%
0
```

For microseconds resolution use the parallel `microsecondsToRun:`

```
topaz 1> printit
System microsecondsToRun: [
  System performOnServer: 'ping -c1 gemtalksystems.com']
%
484
```

#### Elapsed Time

`Time class >> millisecondsElapsedTime:` works similarly, but returns the elapsed rather than the CPU time required.

```
topaz 1> printit
Time millisecondsElapsedTime: [
  System performOnServer: 'ping -c1 gemtalksystems.com']
%
20
```

To get further resolution, use `Time class >> secondsElapsedTime:`, which returns a float with system-dependent resolution. For example, to get a result in microseconds:

```
topaz 1> printit
((Time secondsElapsedTime: [
  System performOnServer: 'ping -c1 gemtalksystems.com']) *
  1000000) asInteger
%
19961
```

## ProfMonitor

The ProfMonitor class allows you to sample the methods that are executed in a given block of code and analyze the percentage of total execution time represented by each method. When an instance starts profiling, it will take a method call stack at specified intervals for a specified period of time. When it is done, it collects the results and returns them in the form of a string formatted as a report.

ProfMonitorTree is a subclass of ProfMonitor, that by default returns an execution tree report or reports, in addition to the reports generated by ProfMonitor. By specifying the desired reports in arguments to ProfMonitor, these tree reports can be also generated from ProfMonitor.

### Sample intervals

ProfMonitor, by default, will take a sample every millisecond (1 ms). You can specify the interval at which ProfMonitor takes samples using the instance methods `interval:` or `intervalNs:`, or class methods with these keywords. `interval:` specifies milliseconds, while `intervalNs:` specifies the interval in nanoseconds (a nanosecond is a billionth of a second). The minimum interval is 1000 nanoseconds.

It may be convenient to refer to Table 16.1 when determining the sample interval and reading the results:

**Table 16.1 Subsecond time conversions**

seconds	milliseconds ms	microseconds μs	nanoseconds ns
1	1000	1,000,000	1,000,000,000
	1	1000	1,000,000
		1	1000

### Reporting limits

By default, ProfMonitor reports every method it found executing. It is usually useful to limit the reporting of methods to the ones that appear more frequently, to reduce clutter in the results and allow you to focus on what is taking the most time.

To limit the reporting results, set the lower limit using the instance method `reportDownTo: limit` or methods with the keyword `downTo:`. Each result at the limit or larger is included in the report.

These methods accept a *limit* of either an integer, which is an absolute number of samples, or a SmallDouble, which defines a percentage of the total number of samples.

For example, a `downTo:` of 50 would specify that the reports include information for every method that was sampled at least 50 times, regardless of whether the number of samples was 100 or 1000. A `downTo:` of 0.50 would specify that the reports include information for methods that were sampled 50% of the time or more; if the total number of samples is 100, this would be 50 actual samples, for a sample set size of 1000, this would be 500 samples.

## Reports

ProfMonitor provides profiling results in the form of a string, containing up to six individual reports that analyze the profiling raw data in different ways. The desired reports to be output can be specified using methods with the reports: keyword. By specifying reports, you can also enable object creation tracking.

Available reports include:

- `#samples`—sample counts report, labeled STATISTICAL SAMPLING RESULTS.
- `#stackSamples`—stack sampling report, labeled STATISTICAL STACK SAMPLING RESULTS.
- `#senders`—method senders report, labeled STATISTICAL METHOD SENDERS RESULTS.
- `#objCreation`—object creation report, labeled OBJECT CREATION REPORT. Including this in the reports: argument enables object tracking.
- `#tree`—method execution tree report, labeled STACK SAMPLING TREE RESULTS. Including this in the reports: argument causes ProfMonitorTree to be used for profiling.
- `#objCreationTree`—object creation tree report, labeled OBJECT CREATION TREE REPORT. Including this in the reports: argument enables object tracking and causes ProfMonitorTree to be used for profiling.

The default reports that are provided depend on the initial class specified;

- ▶ ProfMonitor defaults to { `#samples . #stackSamples . #senders` }
- ▶ ProfMonitorTree defaults to { `#samples . #stackSamples . #senders . #tree` }

## Temporary results file

ProfMonitor stores its results temporarily in a file with the default filename `/tmp/gempid.tmp`. You can specify a different filename by using ProfMonitor's instance creation method `newWithFile:` and variants. This file is deleted by profiling block methods, `profileOff`, and `reportAfterRun*` methods. Note that if the Gem that is executing profiling terminates abnormally, it may leave this file behind; such files must be manually deleted.

## Real vs. CPU time

Profiling operates by taking samples of the stack at intervals specified by the `interval:` or `intervalNs:` arguments. Generally, this specifies that samples are taken at the given intervals in CPU time, which provides information about the relative performance of operations based on how much CPU time they use.

It is also possible to profile based on real time, in which case samples are taken after the specified interval of real time has elapsed. This can detect performance issues that are not based on CPU execution, such as a sleep:, and expose the performance impact of disk access and other performance issues external to the code executing.

Sampling time is one of the options that is defined by the `setOptions:` keyword, either the convenience profiling methods or the instance method. You may include `#real` or `#cpu` in this array.

## Profiling Code

### Convenience Profiling of a Block of Code

ProfMonitor provides several methods that allow you to profile a block of code and report the results with a single class method.

The following profiling methods are available:

```
monitorBlock:
monitorBlock:reports:
monitorBlock:intervalNs:
monitorBlock:intervalNs:options:
monitorBlock:downTo:
monitorBlock:downTo:interval:
monitorBlock:downTo:intervalNs:
monitorBlock:downTo:intervalNs:options:
monitorBlock:downTo:intervalNs:reports:
```

The following defaults apply:

- ▶ When no `downTo:` keyword is provided, the default is 1; each method sampled is reported.
- ▶ When no `interval:` or `intervalNs:` keyword is provided, the default is 0.5ms (500 microseconds).
- ▶ When no `options:` are provided, #CPU is used.
- ▶ When no `reports:` are specified, the default is `{#samples . #stackSamples . #senders}` for ProfMonitor, and `{#samples . #stackSamples . #senders . #tree}` for ProfMonitorTree.

For example, to take samples every millisecond, and only report methods that were sampled at least 10 times:

```
ProfMonitor
  monitorBlock: [ 100 timesRepeat:
    [ System myUserProfile dictionaryNames ]]
  downTo: 10
  interval: 1
```

For a more detailed report, you could take samples every 1/10 of a millisecond; this interval is 100000 nanoseconds. This creates many more samples; to make it easier to control the reporting limit we'll use a percent, and only include methods whose number of samples was 20% or more of the total.

```
ProfMonitor
  monitorBlock: [ 100 timesRepeat:
    [ System myUserProfile dictionaryNames ]]
  downTo: 0.2
  intervalNs: 100000
```

These two reports will give you similar results, but since there are many more samples, the effect of chance sampling error will be less. The choice of sampling interval and report limit

depends on the specific code you are profiling. You may need to run a number of iterations, starting with a more coarse-grained profile and refining for subsequent runs.

## Background Profiling

To sample blocks of code, the quick profiling methods are sufficient. You can also explicitly start and stop profiling, allowing you to profile any arbitrary sequence of GemStone Smalltalk statements.

To start and stop profiling, use the class method `profileOn`, which create an instances of `ProfMonitor` and starts profiling; when you are done, the instance method `profileOff` stops profiling and reports the results.

For example:

```
run
UserGlobals at: #myMonitor put: ProfMonitor profileOn.
%

run
100 timesRepeat: [ System myUserProfile dictionaryNames ].
%

run
(UserGlobals at: #myMonitor) profileOff.
%
```

## Manual Profiling

You can also create and configure the instance of `ProfMonitor`. To profile in this way, perform the following steps:

- Step 1.** Create instance using `ProfMonitor new`, `newWithFile:`, `newWithFile:interval:`, or `newWithFile:intervalNs:`.
- Step 2.** Configure it as desired, using instance methods including `interval:`, `intervalNs:`, `setOptions:`, and `traceObjectCreation:`.
- Step 3.** start profiling using the instance method `startMonitoring`.
- Step 4.** execute your code.
- Step 5.** stop profiling using the instance method `stopMonitoring`.  
Steps 3, 4 and 5 can also be done using `runBlock:`.
- Step 6.** gather results and report, using `reportAfterRun` or `reportAfterRunDownTo:`.

For example:

```
| aMonitor |
aMonitor := ProfMonitor newWithFile:
    '$GEMSTONE/data/profMon.dat'.
aMonitor interval: 2.
aMonitor setOptions: {#objCreation}.
aMonitor startMonitoring.
100 timesRepeat: [ System myUserProfile dictionaryNames ].
aMonitor stopMonitoring.
aMonitor reportAfterRun.
```

## Saving a ProfMonitor for later analysis

ProfMonitor raw data is written to a disk file, and as long as the disk file is available, you may save the instance of ProfMonitor and it will reopen its file to perform the analysis later or in a different session.

To ensure that the file is saved, use methods such as `runBlock:`, which do not automatically create the report and delete the file.

For example:

```
run
UserGlobals at: #aProfMon put:
    (ProfMonitor runBlock: [
        200 timesRepeat: [System myUserProfile dictionaryNames]
    ]).
%
commit
logout
login
run
aProfMon reportAfterRun
%
```

## The Profile Report

The profiling methods discussed in the previous sections return a string formatted as a report. The following example shows a sample run and the resulting report.

### Example 16.1

```
topaz 1> printit
ProfMonitor
  monitorBlock:[
    200 timesRepeat:[ System myUserProfile dictionaryNames ] ]
  reports: { #samples . #stackSamples . #senders . #tree}
%
=====
STATISTICAL SAMPLING RESULTS
elapsed CPU time:    90 ms
monitoring interval: 1.0 ms
report limit threshold: 2 hits / 2.2%
0 pageFaults 2061 objFaults 0 gcMs 824413 edenBytesUsed

tally      %   class and method name
-----
 23  24.21  Array                >> _at:
 22  23.16  IdentityDictionary   >> associationsDo:
 18  18.95  block in SymbolList  >> names
 18  18.95  AbstractDictionary   >> _at:
 11  11.58  block in AbstractDictionary >> associationsDetect:ifNone:
  2   2.11  Object                >> _basicSize
  1   1.05  11 other methods
 95 100.00  Total

=====
STATISTICAL STACK SAMPLING RESULTS
elapsed CPU time:    90 ms
monitoring interval: 1.0 ms
report limit threshold: 2 hits / 2.2%
0 pageFaults 2061 objFaults 0 gcMs 824413 edenBytesUsed

total      %   class and method name
-----
 95 100.00  GsNMethod class      >> _gsReturnToC
 95 100.00  executed code
 95 100.00  ProfMonitor class    >> monitorBlock:downTo:
 95 100.00  ProfMonitor          >> monitorBlock:
 94  98.95  block in executed code
 94  98.95  UserProfile          >> dictionaryNames
 94  98.95  SymbolList           >> namesReport
 94  98.95  SymbolList           >> names
 94  98.95  AbstractDictionary   >> associationsDetect:ifNone:
 94  98.95  IdentityDictionary   >> associationsDo:
 29  30.53  block in AbstractDictionary >> associationsDetect:ifNone:
 23  24.21  Array                >> _at:
 18  18.95  block in SymbolList  >> names
 18  18.95  AbstractDictionary   >> _at:
  2   2.11  Object                >> _basicSize
```

```

1    1.05  2 other methods
95  100.00 Total

```

```
=====
```

```
STATISTICAL METHOD SENDERS RESULTS
```

```
elapsed CPU time: 90 ms
```

```
monitoring interval: 1.0 ms
```

```
report limit threshold: 2 hits / 2.2%
```

	% self Time	% total Time	total ms	local %	Parent Method Child
=	0.0	100.0	90.0	0.0	GsNMethod class >> _gsReturnToC executed code
=	0.0	100.0	90.0	0.0	GsNMethod class >> _gsReturnToC executed code
=	0.0	100.0	90.0	0.0	ProfMonitor class >> monitorBlock:downTo: executed code
=	0.0	100.0	90.0	0.0	ProfMonitor class >> monitorBlock:downTo: ProfMonitor >> monitorBlock:
=	0.0	100.0	90.0	0.0	ProfMonitor class >> monitorBlock:downTo: ProfMonitor >> monitorBlock: block in executed code
=	0.0	100.0	89.1	98.9	ProfMonitor >> startMonitoring
=	0.0	98.9	89.1	0.0	ProfMonitor >> monitorBlock: block in executed code
=	0.0	98.9	89.1	0.0	UserProfile >> dictionaryNames
=	0.0	98.9	89.1	0.0	block in executed code UserProfile >> dictionaryNames
=	0.0	98.9	89.1	0.0	SymbolList >> namesReport
=	0.0	98.9	89.1	0.0	UserProfile >> dictionaryNames SymbolList >> namesReport
=	0.0	98.9	89.1	0.0	SymbolList >> names SymbolList >> names
=	0.0	98.9	89.1	0.0	SymbolList >> namesReport SymbolList >> names
=	0.0	98.9	89.1	0.0	SymbolList >> namesReport AbstractDictionary >> associationsDetect:ifNone:
=	0.0	98.9	89.1	0.0	SymbolList >> names

```

=   0.0  98.9   89.1   0.0  AbstractDictionary >> associationsDetect:ifNone:
      89.1 100.0      IdentityDictionary >> associationsDo:
-----

=  23.2  98.9   89.1 100.0   AbstractDictionary >> associationsDetect:ifNone:
      89.1 23.4 IdentityDictionary >> associationsDo:
      1.9  2.1   Object      >> _basicSize
      27.5 30.9   block in AbstractDictionary >> associationsDetect:ifNone:
      21.8 24.5   Array       >> _at:
      17.1 19.1   AbstractDictionary >> _at:
-----

=  11.6  30.5   27.5 100.0   IdentityDictionary >> associationsDo:
      27.5 37.9 block in AbstractDictionary >> associationsDetect:ifNone:
      17.1 62.1   block in SymbolList >> names
-----

=  24.2  24.2   21.8 100.0   IdentityDictionary >> associationsDo:
      21.8 100.0 Array       >> _at:
-----

=  18.9  18.9   17.1 100.0   block in AbstractDictionary >> associationsDetect:ifNone:
      17.1 100.0 block in SymbolList >> names
-----

=  18.9  18.9   17.1 100.0   IdentityDictionary >> associationsDo:
      17.1 100.0 AbstractDictionary >> _at:
-----

=   2.1   2.1   1.9 100.0   IdentityDictionary >> associationsDo:
      1.9 100.0 Object      >> _basicSize
-----

```

=====

#### STACK SAMPLING TREE RESULTS

elapsed CPU time: 90 ms  
monitoring interval: 1.0 ms  
report limit threshold: 2 hits / 2.2%

```

100.0% (95) executed code      [UndefinedObject]
  100.0% (95) ProfMonitor class  >> monitorBlock:downTo: [ProfMonitor class]
    100.0% (95) ProfMonitor      >> monitorBlock: [ProfMonitor]
      98.9% (94) block in executed code [ExecBlock0]
        | 98.9% (94) UserProfile      >> dictionaryNames
        |   98.9% (94) SymbolList      >> namesReport
        |     98.9% (94) SymbolList    >> names
        |       98.9% (94) AbstractDictionary >> associationsDetect:ifNone: [SymbolDictionary]
        |         98.9% (94) IdentityDictionary >> associationsDo: [SymbolDictionary]
        |           30.5% (29) block in AbstractDictionary >> associationsDetect:ifNone: [ExecBlock1]
        |             | 18.9% (18) block in SymbolList >> names [ExecBlock1]
        |             | 24.2% (23) Array >> _at: [IdentityCollisionBucket]
        |             | 18.9% (18) AbstractDictionary >> _at: [SymbolDictionary]
        |             | 2.1% (2) Object >> _basicSize [IdentityCollisionBucket]

```

As you can see, the report is in four sections, corresponding to the requested reports:

- ▶ #samples: STATISTICAL SAMPLING RESULTS
- ▶ #stackSamples: STATISTICAL STACK SAMPLING RESULTS
- ▶ #senders:STATISTICAL METHOD SENDERS RESULTS
- ▶ #tree: STACK SAMPLING TREE RESULTS

Each section includes the same set of methods that the profile monitor encountered when it checked the execution stack every millisecond; the report is presented to give different views of this data.

Keep in mind that these numbers are based on sampling, and depending on the size and number of samples, may not exactly reflect the actual percentage of time spent in each method and will likely vary from run to run. Also, if you make external calls to the OS, to user actions or other C libraries, this will also distort results for the invoking method.

## Profiling Beyond Performance

Profiling as previously described is focused on the performance of a block of code. ProfMonitor provides additional options that let you track other things that are going on, alongside your code execution, that can impact application performance. These are:

- ▶ number of persistent and temporary objects created
- ▶ number of object faults by this Gem
- ▶ number of page faults by this Gem
- ▶ space used in the temporary object "eden" space
- ▶ time spent in in-memory garbage collection

To profile these attributes, use profiling methods with the `setOptions:` keyword, specify the option that you want to profile.

When profiling these options, you get one or more standard time-based reports along with the specific attribute profile or profiles. This allows you to correlate with the basic performance over a specific execution run.

Type of Profile	Options Keyword	Units	Default time profiling
Object creation	#objCreation		#cpu
Object faults	#objFaults	faults	#real
Page faults	#pageFaults	faults	#real
GC operations	#gcTime	milliseconds	#cpu
Temporary memory eden space	#edenUsage	bytes	#real

## Object Creation Tracking

Object creation tracking is enabled in a number of ways:

- ▶ #objCreation in the setOptions: array;
- ▶ Specifying #objCreation or #objCreationTree in the reports: array; or by
- ▶ using the ProfMonitor instance method traceObjectCreation:.

When object creation tracking is enabled, after the standard report sections, an additional section is included to report the count and object creation.

For example:

### Example 16.2 Object creation report

```
OBJECT CREATION REPORT:
elapsed CPU time:    40 ms
monitoring interval: 2.0 ms

tally  class of created object
       call stack
-----  -----

    600  String class
-----  -----
        500  SmallInteger >> asString
          500  SymbolList  >> namesReport
            500  UserProfile >> dictionaryNames
              500  executed code
                500  GsNMethod class >> _gsReturnToC
-----  -----
        100  String class >> new
          100  SymbolList  >> namesReport
            100  UserProfile >> dictionaryNames
              100  executed code
                100  GsNMethod class >> _gsReturnToC
-----  -----

    100  Array class
-----  -----
        100  SymbolList  >> names
          100  SymbolList  >> namesReport
            100  UserProfile >> dictionaryNames
              100  executed code
                100  GsNMethod class >> _gsReturnToC
```

## Memory Use Profiling

While object creation is the most important way to profile a Gem's use of memory, ProfMonitor provides several other options to allow you to track the impact of the operations by methods on the Gem's memory use.

Using the following keys in the `setOptions:` argument changes the profiling to the specific kind of profile. Only one of these can be used at a time.

- ▶ `#objFaults` – profiles object faults
- ▶ `#pageFaults` – profiles page faults
- ▶ `#edenSpace` – profiles eden space, the area of temporary objects memory into which new temporary objects are put.
- ▶ `#gcTime` – profiles the time spent in in-memory garbage collection, as tracked by the cache statistic `TimeInScavenges`.

By default, for object faults, page faults, and eden space, the sampling frequency is calculated in real time, rather than CPU time. This can be changed by also including `#cpu` in the `setOptions:` array. Garbage collection time is sampled by default in CPU time.

When using these option, the first two reports describe faults, rather than milliseconds. The third report provides millisecond performance information to allow correlation with the faulting data.

## 16.2 Clustering Objects for Faster Retrieval

As you've seen, GemStone ordinarily manages the placement of objects on the disk automatically – you're never forced to worry about it. Occasionally, you might choose to group related objects on secondary storage to enable GemStone to read all of the objects in the group with as few disk accesses as possible.

Because an access to the first element usually presages the need to read the other elements, it makes sense to arrange those elements on the disk in the smallest number of disk pages. This placement of objects on physically contiguous regions of the disk is the function of class `Object's clustering` protocol. By clustering small groups of objects that are often accessed together, you can sometimes improve performance.

Clustering a group of objects packs them into disk pages, each page holding as many of the objects as possible. The objects are contiguous within a page, but pages are not necessarily contiguous on the disk.

### Will Clustering Solve the Problem?

Clustering objects solves a specific problem – slow performance due to excessive disk accessing. However, disk access is not the only factor in poor performance. In order to determine if clustering will solve your problem, you need to do some diagnosis. You can use GemStone's VSD utility to find out how many times your application is accessing the disk. VSD allows you to chart system statistics over time to better understand the performance of your system. See the *VSD User's Guide* for more information on using VSD.

The following statistics are of interest:

- ▶ `PageReads` – how many pages your session has read from the disk since the session began
- ▶ `PageWrites` – how many pages your session has written to the disk since the session began

You can examine the values of these statistics before and after you commit each transaction to discover how many pages it read in order to perform a particular query, and to determine the number of disk accesses required by the process of committing the transaction.

It is tempting to ignore these issues until you experience a problem such as an extremely slow application, but if you keep track of such statistics on a regular (even if intermittent) basis, you will have a better idea of what is “normal” behavior when a problem crops up.

## Cluster Buckets

You can think of clustering as writing the components of their receivers on a stream of disk pages. When a page is filled, another is randomly chosen and subsequent objects are written on the new page. A new page is ordinarily selected for use only when the previous page is filled, or when a transaction ends. Sending the message `cluster` to objects in repeated transactions will, within the limits imposed by page capacity, place its receivers in adjacent disk locations. (Sending the message `cluster` to objects repeatedly within a transaction has no effect.)

The stream of disk pages used by `cluster` and its companion methods is called a *bucket*. GemStone captures this concept in the class `ClusterBucket`.

If you determine that clustering will improve your application’s performance, you can use instances of the class `ClusterBucket` to help. All objects assigned to the same instance of `ClusterBucket` are to be clustered together. When the objects are written, they are moved to contiguous locations on the same page, if possible. Otherwise the objects are written to contiguous locations on several pages.

Once an object has been clustered into a particular bucket and committed, that bucket remains associated with the object until you specify otherwise. When the object is modified, it continues to cluster with the other objects in the same bucket, although it might move to another page within the same bucket.

## Using Existing Cluster Buckets

By default, a global array called `AllClusterBuckets` defines seven instances of `ClusterBucket`. Each can be accessed by specifying its offset in the array. For example, the first instance, `AllClusterBuckets at: 1`, is the default bucket when you log in. This bucket is invariant—you cannot modify it.

The second, third, and seventh cluster buckets in the array can be used for whatever purposes you require and can all be modified.

The GemStone system makes use of the fourth, fifth, and sixth buckets of the array `AllClusterBuckets`:

- ▶ `AllClusterBuckets at: 4` is the bucket used to cluster the methods associated with kernel classes.
- ▶ `AllClusterBuckets at: 5` is the bucket used to cluster the strings that define source code for kernel classes.
- ▶ `AllClusterBuckets at: 6` is the bucket used to cluster other kernel objects such as globals.

You can determine how many cluster buckets are currently defined by executing:

```
System maxClusterBucket
```

A given cluster bucket's offset in the array specifies its *clusterId*. A cluster bucket's *clusterId* is an integer in the range of 1 to (`System maxClusterBucket`).

*NOTE*

*For compatibility with previous versions of GemStone, you can use a *clusterId* as an argument to any keyword that takes an instance of *ClusterBucket* as an argument.*

You can determine which cluster bucket is currently the system default by executing:

```
System currentClusterBucket
```

You can access all instances of cluster buckets in your system by executing:

```
ClusterBucket allInstances
```

You can change the current default cluster bucket by executing an expression of the form:

```
System clusterBucket: aClusterBucket
```

## Creating New Cluster Buckets

You are not limited to the predefined instances of *ClusterBucket*. You can create new instances of *ClusterBucket* with the simple expression `ClusterBucket new`.

This expression creates a new instance of *ClusterBucket* and adds it to the array `AllClusterBuckets`. You can then access the bucket in one of two ways. You can assign it a name:

```
UserGlobals at: #empClusterBucket put: (ClusterBucket new)
```

You could then refer to it in your application as `empClusterBucket`. Alternatively, you can use the offset into the array `AllClusterBuckets`. For example, if this is the first cluster bucket you have created, you could refer to it this way:

```
AllClusterBuckets at: 8
```

(Recall that the first seven elements of the array are predefined.)

You can determine the *clusterId* of a cluster bucket by sending it the message `clusterId`. For example:

```
empClusterBucket clusterId
8
```

You can access an instance of *ClusterBucket* with a specific *clusterId* by sending it the message `bucketWithId:`.

You can create and use as many cluster buckets as you need; up to thousands, if necessary.

*NOTE*

*For best performance and disk space usage, use no more than 32 cluster buckets in a single session.*

## Cluster Buckets and Concurrency

Cluster buckets are designed to minimize concurrency conflicts. As many users as necessary can cluster objects at the same time, using the same cluster bucket, without experiencing concurrency conflicts. Cluster buckets do not contain or reference the objects clustered on them -- the objects that are clustered keep track of their bucket. This also avoids problems with authorizations.

However, creating a new instance of *ClusterBucket* automatically adds it to the global array `AllClusterBuckets`. Adding an instance to `AllClusterBuckets` causes a concurrency

conflict when more than one transaction tries to create new cluster buckets at the same time, since all the transactions are all trying to write the same array object.

To avoid concurrency conflicts, you should design your clustering when you design your application. Create all the instances of ClusterBucket you anticipate needing and commit them in one or few transactions.

To facilitate this kind of design, GemStone allows you to associate descriptions with specific instances of ClusterBucket. In this way, you can communicate to your fellow users the intended use of a given cluster bucket with the message `description:`. For example:

### Example 16.3

```
UserGlobals at: #empClusterBucket put: (ClusterBucket new).
empClusterBucket description: 'Use this bucket for
    clustering employees and their instance variables.'
```

As you can see, the message `description:` takes a string of text as an argument.

Changing the attributes of a cluster bucket, such as its description or `clusterId`, writes that cluster bucket and thus can cause concurrency conflict. Only change these attributes when necessary.

#### NOTE

*For best performance and disk space usage as well as avoiding concurrency conflicts, create the required instances of ClusterBucket all at once, instead of on a per-transaction basis, and update their attributes infrequently.*

## Cluster Buckets and Indexing

Indexes on instance of subclasses of `UnorderedCollection` are created and modified using the cluster bucket associated with the specific collection, if any. To change the clustering of an indexed collection:

1. Remove its index.
2. Recluster the collection.
3. Re-create its index.

## Clustering Objects

Class `Object` defines several clustering methods. One method is simple and fundamental. Another method is more sophisticated and attempts to order the receiver's instance variables as well as writing the receiver itself.

### The Basic Clustering Message

The basic clustering message defined by class `Object` is `cluster`. For example:

```
myObject cluster
```

This simplest clustering method simply assigns the receiver to the current default cluster bucket; it does not attempt to cluster the receiver's instance variables. When the object is next written to disk, it will be clustered according to the attributes of the current default cluster bucket.

If you wish to cluster the instance variables of an object, you can define a special method to do so.

**CAUTION**

*Do not redefine the method `cluster` in the class `Object`, because other methods rely on the default behavior of the `cluster` method. You can, however, define a `cluster` method for classes in your application if required.*

Suppose, for example, that you defined class `Name` and class `Employee` as shown in Example 16.4.

#### Example 16.4

---

```
Object subclass: 'Name'
  instVarNames: #('first' 'middle' 'last')
  classVars: #( )
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals.
Object subclass: 'Employee'
  instVarNames: #('name' 'job' 'age' 'address')
  classVars: #( )
  classInstVars: #()
  poolDictionaries: {}
  inDictionary: UserGlobals.
```

---

The following clustering method might be suitable for class `Employee`. (A more purely object-oriented approach would embed the information on clustering first, middle, and last names in the `cluster` method for `Name`, but such an approach does not exemplify the breadth-first clustering technique we wish to show here.)

#### Example 16.5

---

```
method: Employee
clusterBreadthFirst
  self cluster.
  name cluster.
  job cluster.
  address cluster.
  name first cluster.
  name middle cluster.
  name last cluster.
  ^false
%

| Lurleen |
Lurleen := Employee new name: (Name new first: #Lurleen);
  job: 'busdriver'; age: 24; address: '540 E. Sixth'.
Lurleen clusterBreadthFirst
```

---

The elements of byte objects such as instances of `String` and `Float` are always clustered automatically. A string's characters, for example, are always written contiguously within disk pages. Consequently, you need not send `cluster` to each element of each string stored in `job` or `address`; clustering the strings themselves is sufficient. Sending `cluster` to individual special objects (instances of `SmallInteger`, `Character`, `Boolean`, `SmallDouble`, or `UndefinedObject`) has no effect. Hence no clustering message is sent to `age` in the previous example.

After sending `cluster` to an `Employee`, the `Employee` is clustered as follows:

```
anEmp aName job address first middle last
```

`cluster` returns a Boolean value. You can use that value to eliminate the possibility of infinite recursion when you're clustering the variables of an object that can contain itself. Here are the rules that `cluster` follows in deciding what to return:

- ▶ If the receiver has already been clustered during the current transaction or if the receiver is a special object, `cluster` declines to cluster the object and returns true to indicate that all of the necessary work has been done.
- ▶ If the receiver is a byte object that has not been clustered in the current transaction, `cluster` writes it on a disk page and, as in the previous case, returns true to indicate that the clustering process is finished for that object.
- ▶ If the receiver is a pointer object that has not been clustered in the current transaction, `cluster` writes the object and returns false to indicate that the receiver might have instance variables that could benefit from clustering.

## Depth-First Clustering

`clusterDepthFirst` differs from `cluster` only in one way: it traverses the tree representing its receiver's instance variables (named, indexed, or unordered) in depth-first order, assigning each node to the current default cluster bucket as it is visited. That is, it writes the receiver's first instance variable, then the first instance variable of that instance variable, then the first instance variable of that instance variable, and so on to the bottom of the tree. It then backs up and visits the nodes it missed before, repeating the process until the whole tree has been written.

After sending `clusterDepthFirst` to an `Employee`, the `Employee` is clustered as follows:

```
anEmp aName first middle last job address
```

## Assigning Cluster Buckets

Both `cluster` and `clusterDepthFirst` use the current default cluster bucket. If you wish to use a specific cluster bucket instead, you can use the method `clusterInBucket:.` For example, the following expression clusters `aBagOfEmployees` using the specific cluster bucket `empClusterBucket`:

```
aBagOfEmployees clusterInBucket: empClusterBucket
```

In order to determine the cluster bucket associated with a given object, you can send it the message `clusterBucket`. For example, after executing the example above, the following example would return the value shown below:

```
aBagOfEmployees clusterBucket
empClusterBucket
```

## Clustering and Memory Use

Clustering tags objects in memory so that when the next successful commit occurs, the objects are clustered onto data pages according to the method specified. After an object has been clustered, it is considered to be “dirty”. If you cluster a large number of objects, you may need to increase temporary object memory to avoid running out of session memory. See “Managing VM Memory” on page 322.

## Using Several Cluster Buckets

When you want to write a loop that clusters parts of each object in a group into separate pages, it is helpful to have multiple cluster buckets available.

Suppose that you had defined class `SetOfEmployees` and class `Employee` as in Example 16.4 on page 315. Suppose, in addition, that you wanted a clustering method to write all employees contiguously and then write all employee addresses contiguously.

With only one cluster bucket at your disposal, you would need to define your clustering method as shown in Example 16.6. In this approach, each employee is fetched once for clustering, then fetched again in order to cluster the employee’s address.

### Example 16.6

```
method: SetOfEmployees
clusterEmployees
    self do: [:n | n cluster].
    self do: [:n | n address cluster].
%
myEmployees clusterEmployees
```

## Clustering Class Objects

Clustering provides the most benefit for small groups of objects that are often accessed together – for example, a class with its instance variables. Those instance variables of a class that describe the class’s variables are often accessed in a single operation, as are the instance variables that contain a class’s methods. Therefore, class `Behavior` defines the following special clustering methods for classes:

**Table 16.2 Clustering Protocol**

<code>clusterBehavior</code>	Clusters in depth-first order the parts of the receiver required for executing GemStone Smalltalk code (the receiver and its method dictionary).
<code>clusterDescription</code>	Clusters in depth-first order those instance variables in the receiver that describe the structure of the receiver’s instances. (Does not cluster the receiver itself.) The instance variables clustered are <i>instVarNames</i> , <i>classVars</i> , <i>categories</i> , and <i>class histories</i> .

Table 16.2 Clustering Protocol (Continued)

<pre>clusterBehaviorExceptMethods:   aCollectionOfMethodNames</pre>	<p>This method can sometimes provide a better clustering of the receiving class and its method dictionary by omitting those methods that are seldom used. This omission allows frequently used methods to be packed more densely.</p>
---	---

The code in Example 16.7 clusters class Employee's structure-describing variables, then its class methods, and finally its instance methods.

### Example 16.7

```
| behaviorBucket descriptionBucket |
behaviorBucket := AllClusterBuckets at: 4.
descriptionBucket := AllClusterBuckets at: 5.
System clusterBucket: descriptionBucket.
Employee clusterDescription.
System clusterBucket: behaviorBucket.
Employee class clusterBehavior.
Employee clusterBehavior.
```

The following clusters all of class Employee's instance methods except for address and address:

```
Employee clusterBehaviorExceptMethods: #(#address #address:).
```

## Maintaining Clusters

Once you have clustered certain objects, they do not necessarily stay clustered in the same way forever. If you edit some of the objects in the data structure, the edited object will be placed on a new page in the same clusterBucket. The performance benefit of clustering is that the objects are on the same page, but since the clusterBucket will span multiple pages, the objects may be in the same clusterBucket but not on the same page.

You may therefore wish to check an object's location, especially if you suspect that such declustering is causing your application to run more slowly than it used to.

### Determining an Object's Location

To enable you to check your clustering methods for correctness, Class Object defines the message `page`, which returns an integer identifying the disk page on which the receiver resides. For example:

```
anEmp page
2539
```

Disk page identifiers are returned only for temporary use in examining the results of your custom clustering methods—they are not stable pointers to storage locations. The page on which an object is stored can change for several reasons, as discussed in the next section.

For special objects (instances of `SmallInteger`, `Character`, `Boolean`, `SmallDouble`, or `UndefinedObject`), the page number returned is 0.

## Why Do Objects Move?

The page on which an object is stored can change for any of the following reasons:

- ▶ A clustering message is sent to the object or to another object on the same page.
- ▶ The current transaction is aborted.
- ▶ The object is modified.
- ▶ Another object on the page with the object is modified.

As your application updates clustered objects, new values are placed on secondary storage using GemStone's normal space allocation algorithms. When objects are moved, they are automatically reclustered within the same clusterId. If a specific clusterId was specified, it continues to be used; if not, the default clusterId is used.

If, for example, you replace the string at position 2 of the clustered array ProscribedWords, the replacement string is stored in a page separate from the one containing the original, although it will still be within the same clusterId. Therefore, it might be worthwhile to recluster often-modified collections occasionally to counter the effects of this fragmentation. You'll probably need some experience with your application to determine how often the time required for reclustered is justified by the resulting performance enhancement.

## 16.3 Modifying Cache Sizes for Better Performance

As code executes in GemStone, committed objects must be fetched from disk or from cache, and temporary objects must be managed. This is handled transparently by the GemStone repository monitor. The performance of your application can be affected both by the tuning of the caches, and the structure and usage patterns of your application.

### GemStone Caches

GemStone uses two kinds of caches: temporary object space and the shared page cache.

#### Temporary Object Space

The *temporary object space* cache is used to store temporary objects created by your application. Each Gem session has a temporary object memory that is private to the Gem process and its corresponding session. When you fault persistent (committed) objects into your application, they are copied to temporary object memory.

Some of the temporary objects in the cache may ultimately become permanent and reside on the disk, but probably not all of them. Temporary objects that your application creates merely in order to do its work reside in temporary object space until they are no longer needed, when the Gem's garbage collector reclaims the storage they use.

It is important to provide sufficient temporary object space. At the same time, you must design your application so that it does not create an infinite amount of reachable temporary objects. Temporary object memory must be large enough to accommodate the sum of live temporary objects and modified persistent objects. If that sum exceeds the allocated temporary object memory, the Gem can encounter an OutOfMemory condition and terminate.

The amount of memory allocated for temporary object space is primarily determined by the `GEM_TEMPOBJ_CACHE_SIZE` configuration option. You should increase this value for applications that create a large number of temporary objects – for example, applications that make heavy use of the reduced conflict classes or sessions performing a bulk load.

You will probably need to experiment somewhat before you determine the optimum size of the temporary object space for the application. The default of 50000 (50 MB) should be adequate for normal user sessions. For sessions that place a high demand on the temporary object cache, such as upgrade, you may wish to use a much larger cache size.

For a more exhaustive discussion of the issues involved in managing the size of temporary object memory, and a general discussion of garbage collection, see the “Garbage Collection” chapter of the *System Administration Guide*.

For details about how to set the size of `GEM_TEMPOBJ_CACHE_SIZE` in the Gem configuration file, see the “GemStone Configuration Options” appendix of the *System Administration Guide*.

#### Shared Page Cache

The *shared page cache* is used to hold the *object table* – a structure containing pointers to all the objects in the repository – and copies of the disk pages that hold the objects with which users are presently working. The system administrator defines the size of the shared page cache in the configuration file for the Stone. The single active Stone process per repository has one shared page cache per host machine.

Whenever the Gem needs to read an object, it reads the entire page on which an object resides into the shared page cache. If the Gem then needs to access another object, GemStone first checks to see if the object is already in the shared page cache. If it is, no further disk access is necessary. If it is not, it reads another page into the shared page cache.

For acceptable performance, the shared page cache should be large enough to hold the entire object table. To get the best possible performance, make the shared page cache as large as possible, ideally large enough to contain all objects in the repository.

You can determine the size of the object table by examining the results of a **pageaudit** (see the *System Administration Guide*), which includes the number and size required of various kinds of pages in the repository, including the kinds of object table pages.

The amount of memory allocated for the shared page cache is determined by the `SHR_PAGE_CACHE_SIZE_KB` configuration parameter (in the Stone configuration file). This is described in the *System Administration Guide* (Appendix A, GemStone Configuration Options).

## Getting Rid of Non-Persistent Objects

As discussed in Chapter 4, you can create instances of `KeySoftValueDictionary` to enable your session to free up temporary object memory as needed. The entries in a `KeySoftValueDictionary` are *non-persistent*; that is, they cannot be committed to the database. When there is a demand on memory, you can configure GemStone to clear non-persistent entries as needed during a VM mark/sweep garbage collection.

The action taken during mark/sweep depends on two configuration parameters, along with `startingMemUsed` – the percentage of temporary object memory in-use at the beginning of the VM mark/sweep.

Case 1: `GEM_SOFTREF_CLEANUP_PERCENT_MEM < startingMemUsed < 80%`

If `startingMemUsed` is greater than `GEM_SOFTREF_CLEANUP_PERCENT_MEM` but less than 80%, the VM mark/sweep will attempt to clear an internally determined number of least recently used `SoftReferences` (non-persistent entries). Under rare circumstances, you might choose to specify a minimum number (`GEM_KEEP_MIN_SOFTREFS`) that will not be cleared.

Case 2: `startingMemUsed < GEM_SOFTREF_CLEANUP_PERCENT_MEM`

No `SoftReferences` will be cleared.

Case 3: `startingMemUsed > 80%`

VM mark/sweep will attempt to clear all `SoftReferences`.

For more about these and other configuration parameters, see the “GemStone Configuration Options” appendix of the *System Administration Guide*.

Several cache statistics may also be of interest: `NumSoftRefsCleared`, `NumLiveSoftRefs`, and `NumNonNilSoftRefs`. For more about these statistics, see the “Monitoring GemStone” chapter of the *System Administration Guide*.

## 16.4 Managing VM Memory

As mentioned earlier in this chapter, each Gem session has a temporary object memory that is private to the Gem process and its corresponding session. When you fault persistent (committed) objects into your application, they are copied to temporary object memory.

It is important to provide sufficient temporary object space. At the same time, you must design your application so that it does not create an infinite amount of reachable temporary objects. Temporary object memory must be large enough to accommodate the sum of live temporary objects and modified persistent objects. If that sum exceeds the allocated temporary object memory, the Gem can encounter an `OutOfMemory` condition and terminate.

There is a limit on how large a transaction can be, either in terms of the total size of previously committed objects that are modified, or of the total size of temporary objects that are transitively reachable from modified committed objects. For large applications, you may need to commit incrementally, rather than waiting to commit all at once.

The remainder of this chapter discusses issues to consider when allocating and managing temporary object memory, and presents techniques for diagnosing and addressing `OutOfMemory` conditions. This section assumes you have read the general discussion of memory organization in the “Managing Memory” chapter of the *System Administration Guide*.

### Large Working Set

If your application requires a large working set of committed objects in memory, you can configure the `pom` area to be large (compared to other object spaces) without having an adverse effect on in-memory garbage collection. To do this, increase the setting for the configuration parameter `GEM_TEMPOBJ_POMGEN_SIZE`. For details on how to do this, see the *System Administration Guide*, Appendix A.

### Class Hierarchy

If your application references a very deep class hierarchy, you may need to adjust the memory configuration accordingly to allow a larger temporary object memory. When an object is in memory, its class is also faulted into the `perm` area of temporary object memory, along with the class’s superclass, extending up through the hierarchy all the way to `Object`. While this approach provides for significantly faster message lookups, it also increases the consumption of temporary object memory.

For example, the default configuration provides 1 MB for the `perm` area. Each class consumes about 400 bytes (including the metaclass). Thus, the default configuration can accommodate about 2500 classes in memory at once.

### UserAction Considerations

*NOTE*

*Do not compact the `code` region of temporary object memory while a `UserAction` is executing.*

When using `GemBuilder` for C, you may encounter an `OutOfMemory` error within an `UserAction` in either of the following situations:

- ▶ The UserAction faults in a large number of methods via **GciPerform**.
- ▶ The UserAction compiles a large number of anonymous methods via **GciExecute**.

## Exported Set

The Export Set is a collection of objects for which the Gem process has handed out its OOP to one of the interfaces (GCI, GBS, objects returned from topaz run commands). Objects in the export set are prevented from being garbage collected by any of the garbage collection processes (that is, by a Gem's in-memory collection of temporary objects, markForCollection, or the epoch garbage collection). The export set is used to guarantee referential integrity for objects only referenced by an application, that is, objects that have no references to them within the Gem.

The application program is responsible for timely removal of objects from the export set. The contents of the export set can be examined using hidden set methods defined in class System.

In general, the smaller the size of the export set, the better the performance is likely to be. There are several reasons for this relationship. The export set is one of the root sets used for garbage collection. The larger the export set, the more likely it is that objects that would otherwise be considered garbage are being retained. One threshold for performance is when the size of the export set exceeds 16K objects. When its size is smaller than 16K objects, the export set is a small object in object memory. When its size is larger than 16K, the export set becomes a large object, implemented as a tree of small objects in memory.

The configuration parameter `#GemDropCommittedExportedObjs` will allow committed object to be removed from the export set when memory is low, at the expense of having to re-fault these object when they are needed.

## Debugging out of memory errors

If you find that your application is running out of temporary memory, you can set several GemStone environment variables to help you identify which parts of your application are triggering OutOfMemory conditions. These environment variables allow you to obtain multiple Smalltalk stack printouts and other useful information before your application runs out of temporary object memory. For example, it displays how many objects of each class are in temporary memory.

Details on these environment variables are provided in the *System Administration Guide*, and they are listed in the `$GEMSTONE/sys/gemnetdebug` file, which is a debug version of the `gemnetobject` script. `gemnetdebug` enables some, but not all, available memory related environment variables. By using `gemnetdebug` instead of `gemnetobject` in your RPC login parameters, you can generate memory logging information. For help with analysis, contact GemTalk Technical Support.

Once you've identified the cause/s of the problem, you can modify your application to reduce the demand on memory, or adjust your GemStone configuration options to provide a larger amount of memory.

## Signal on low memory condition

When a session runs low on temporary object memory, there are actions it can take to avoid running out of memory altogether; for example, the session may commit or abort, or discard temporary objects.

By enabling handling for `AlmostOutOfMemoryError`, an application can take appropriate action before memory is entirely full. This exception is asynchronous, so may be received at any time memory use is greater than the threshold the end of an in-memory markSweep. However, if the session is executing a user action, or is in index maintenance, the error is deferred and generated when execution returns.

After an `AlmostOutOfMemoryError` notification is delivered, the handling is automatically disabled. Handling must be reenabled each time the signal occurs. Handling this signal is enabled by executing either of the following:

```
AlmostOutOfMemoryError enable
```

The default the default threshold (the amount of memory that is used) is 90%. You can specify a different threshold using,

```
AlmostOutOfMemoryError enableAtThreshold: integerBetween1And125
```

When you enable this error, when your temporary object memory is 90% full, your session will get an error; but the session will not terminate with an out of memory error. You can take such steps as are needed to recover temporary object space (this will likely happen automatically as temporary variables in your executing code will be available for reclaim).

You may also set an handler block to take specific action, such as commit or abort, on an `AlmostOutOfMemoryError`. The following example shows how you can catch an `AlmostOutOfMemoryError`, automatically commit your work, and resume execution.

#### Example 16.8 Commit on out of memory

```
AlmostOutOfMemoryError enable.
[ code ]
  on: AlmostOutOfMemoryError
  do: [:ex |
    GsFile gciLogClient: ex messageText.
    System commit.
    AlmostOutOfMemoryError enable.
    ex resume.
  ].
```

`AlmostOutOfMemory` is similar to `AlmostOutOfMemoryError`, but is a notification rather than an error. The semantics are similar; however, to allow the `AlmostOutOfMemory` to be effective, you must also enable signalling by sending `AlmostOutOfMemory enableNotification`.

## Methods for Computing Temporary Object Space

To find out how much space is left in the `old` area of temporary memory, the following methods in class `System` (category `Performance Monitoring`) are provided:

```
System _tempObjSpaceUsed
  Returns the approximate number of bytes of temporary object memory being used to store objects.
```

```
System _tempObjSpaceMax
  Returns the size of the old area of temporary object memory; that is, the approximate maximum number of bytes of temporary object memory that are
```

usable for storing objects. When the `old` area fills up, the Gem process may terminate with an `OutOfMemory` error.

`System _tempObjSpacePercentUsed`

Returns the approximate percentage of temporary object memory that is being used to store temporary objects. This is equivalent to the expression:

```
(System _tempObjSpaceUsed * 100) //
System _tempObjSpaceMax.
```

Note that it is possible for the result to be slightly greater than 100%. Such a result indicates that temporary memory is almost completely full.

To measure the size of complex objects, you might create a known object graph containing typical instances of the classes in question, and then execute the following methods at various points in your *test* code to get memory usage information:

*CAUTION*

*Do not execute this sequence in your production code!*

### Example 16.9

```
System _vmMarkSweep.
System _tempObjSpaceUsed.
```

## Statistics for monitoring memory use

You can monitor the following statistics to better understand your application's memory usage. The statistics are grouped here with related statistics, rather than alphabetically.

**Table 16.3 Statistics Related to the Objects Copied into Memory**

<code>ObjectsRead</code>	The number of committed objects copied into VM memory since the start of the session.
<code>ClassesRead</code>	The number of classes copied into the <code>perm</code> generation area of VM memory since the start of the session.
<code>MethodsRead</code>	The number of <code>GsNMethods</code> copied into the <code>code</code> generation area of VM memory since the start of the session.
<code>ObjectsRefreshed</code>	The number of committed objects in VM memory that have been re-read from the shared page cache after transaction boundaries, since the start of the session.

**Table 16.4 Statistics Related to Mark/Sweeps and Scavenges**

<code>NumberOfMarkSweeps</code>	The number of mark/sweeps executed by the in-memory garbage collector.
<code>NumberOfScavenges</code>	The number of scavenges executed by the in-memory garbage collector. Only updated at mark/sweeps.
<code>TimeInMarkSweep</code>	The real time (in milliseconds) spent in in-memory garbage collector mark/sweeps.

**Table 16.4 Statistics Related to Mark/Sweeps and Scavenges**

TimeInScavenge	The real time (in milliseconds) spent in in-memory garbage collector scavenges. Only updated at mark/sweeps.
----------------	--

**Table 16.5 Statistics Related to Object Memory Regions**

CodeCacheSizeBytes	Total size in bytes of copies of GsNMethods that are in the code generation area and ready for execution, as of the end of mark/sweep.
NewGenSizeBytes	The number of used bytes in the new generation at the end of mark/sweep.
OldGenSizeBytes	The number of used bytes in the old generation at the end of mark/sweep.
PomGenSizeBytes	The number of used bytes in the pom generation area at the end of mark/sweep. Pom generation holds clean copies of committed objects.
PermGenSizeBytes	The number of used bytes in the perm generation area at the end of mark/sweep. Perm generation holds copies of Classes.
MeSpaceUsedBytes	The number of bytes occupied by the remembered set (remSet), in-memory oopMap, and in-use map entries.
MeSpaceAllocatedBytes	The number of bytes allocated for the remembered set (remSet), in-memory oopMap, and map entries.

**Table 16.6 Statistics Related to Stubbing**

NumRefsStubbedMarkSweep	The number of in-memory references that were stubbed (converted to a POM objectId) by in-memory mark/sweep.
NumRefsStubbedScavenge	The number of in-memory references that were stubbed (converted to a POM objectId) by in-memory scavenge.

**Table 16.7 Statistics Related to Garbage Collection**

CodeGenGcCount	The number of times the code generation area has been garbage collected.
PomGenScavCount	The number of times scavenge has thrown away the oldest pom generation space.

### Symbol Creation

When a new symbol is needed (which may just be from evaluating a code snippet that includes a symbol), it is created by the SymbolGem. The SymbolGem process runs in the background and is responsible for creating all new Symbols, based on session requests that

are managed by the Stone. You can examine the following statistics to track the effect of symbol creation activity on temporary object memory.

**Table 16.8 Statistics Related to Symbol Creation**

NewSymbolRequests	The number of symbol creation requests by a session to the symbol creation gem.
NewSymbolsCount	The number of symbol creation requests by a session that did not resolve to an already committed symbol.
TimeWaitingForSymbols	Cumulative elapsed time (in milliseconds) waiting for symbol creation requests to be processed.

**Table 16.9 Other Statistics**

ExportedSetSize	The number of objects in the ExportSet (see page 323).
TrackedSetSize	The number of objects in the Tracked Objects Set, as defined by the GCI. You can use <code>GciReleaseObjs</code> to remove objects from the Tracked Objects Set. For details, see the <i>GemStone/S 64 Bit GemBuilder for C</i> manual.
DirtyListSize	The number of modified committed objects in the temporary object memory dirty list.
WorkingSetSize	The number of objects in memory that have an <code>objectId</code> assigned to them; approximately the number of committed objects that have been faulted in plus the number that have been created and committed.
TempObjSpacePercentUsed	The approximate percentage of temporary object memory for this session that is being used to store temporary objects. If this value approaches or exceeds 100%, sessions will probably encounter an <code>OutOfMemory</code> error. This statistic is only updated at the end of a mark/sweep operation.  Compare with <code>System _tempObjSpacePercentUsed</code> , which is computed whenever the primitive is executed.

## 16.5 NotTranloggedGlobals

All changes to the repository are written to the transaction logs when the transaction is committed, to ensure these changes are recoverable in case of unexpected shutdown, and to allow these changes to be applied to warm standby copies of the repository. However, you may have data that you will be committing changes to, but that does not need to be recovered in case of system crash or corruption. For this kind of data, you can avoid the overhead of writing each change to the transaction logs, and the disk space required for the transaction logs to archive large amounts of non-critical data.

For objects that are intended to be persistent, but not log changes in the transaction logs, there must be no reference from persistent objects, and the reference should be from the variable `NotTranloggedGlobals`. This is in the `Globals SymbolDictionary`.

For example:

```
NotTranloggedGlobals at: #perfLog put: PerformanceLogger new.
```

If the object in `NotTranloggedGlobals` is reachable from `AllUsers` (the regular root for all persistent objects), it will generate an error on commit.

On system crash or unexpected shutdown, the state of the objects reachable from `NotTranloggedGlobals` will be as was recorded in the most recent checkpoint prior to the shutdown; changes made after that checkpoint will be lost. If the repository is restored from backup, and transaction logs applied, the state of these objects will be as of the time the backup was taken; all changes made since the backup was taken are lost.

## 16.6 Other Optimization Hints

While optimization is an application-specific problem, we can provide a few ideas for improving application performance:

- ▶ Arrays tend to be faster than sets. If you do not need the particular semantics that a set affords, use an array instead.
- ▶ The following Number classes are listed in decreasing order of performance:
  - SmallInteger
  - SmallDouble
  - Float
  - LargeInteger
  - ScaledDecimal
  - DecimalFloat
- ▶ Avoid coercing integers to floating point numbers. Although GemStone Smalltalk can easily handle mixing integers and floating point numbers in computations, the coercion required can be time-consuming.
- ▶ If you create an instance of a Dictionary class (or subclass) that you intend to load with values later, create it to be approximately the final required size in order to avoid rehashing, which can significantly slow performance.
- ▶ Prefer methods that invoke primitives, if possible, or methods that cause primitives to be invoked after fewer intermediate message-sends. (For information on writing your own primitive methods, see the *GemBuilder for C* manual.)
- ▶ Prefer message-sends over path notation, where possible. (This is not possible in indexed queries, however.)
- ▶ Prefer simpler blocks to more complex blocks. The most efficient blocks refer only to one or more literals, global variables, pool variables, class variables, local block arguments, or block temporaries; they also do not include a return statement.  
  
Less efficient blocks include a return statement and can also refer to one or more of the pseudovariables *super* or *self*, instance variables of *self*, arguments to the enclosing

method, temporary variables of the enclosing method, block arguments, or block temporaries of an enclosing block.

The least efficient blocks enclose a less efficient block of the kind described in the above paragraph.

Blocks provided as arguments to the methods `ifTrue:`, `ifFalse:`, `ifTrue:ifFalse:`, `ifFalse:ifTrue:`, `whileFalse:`, and `whileTrue:` are specially optimized. Unless they contain block temporary variables, you need not count them when counting levels of block nesting.

- ▶ Use optimized selectors whenever possible. For example, iterations using `to:do` are specially optimized; using `to:do:` instead of another collection iteration method avoids a message send and a level of block nesting, possibly avoiding the cost of using a block altogether. A list of optimized selectors is under “Reserved and Optimized Selectors” on page 399.

In the same way, for fastest performance in iterating over Collections, use the `to:do:` or `to:by:do:` methods to iterate, rather than `do:` or other collection iteration methods

- ▶ Append to rather than concatenate strings. `String >>` , creates a new string that combines the receiver and argument, while `String >> add:` modifies the receiver. This is much more efficient in memory use, although otherwise performance is similar.
- ▶ If you have a choice between a method that modifies an object and one that returns a modified copy, use the method that modifies the object directly if your application allows it. This creates fewer temporary objects whose storage will have to be reclaimed.
- ▶ Avoid generating temporary objects whose storage will need to be reclaimed. Storage reclamation can slow your application significantly.
- ▶ Keep repository files on a disk reserved for their use, if possible. Particularly avoid putting repository files on the disk used for swapping.
- ▶ For large applications, you may need to commit incrementally, rather than waiting to commit all at once. There is a limit on how large a transaction can be, either in terms of the total size of previously committed objects that are modified, or of the total size of temporary objects that are transitively reachable from modified committed objects.
- ▶ Consider trade-offs in indexing. While indexes can improve query performance on large collections, there is overhead. If the collection has fewer than about 2000 objects, the extra overhead in internal objects and index maintenance may not be worth negligible performance gain in queries.



# Working with Classes and Methods

---

An object responds to messages defined and stored with its class and its class's superclasses. The classes named `Object`, `Class`, and `Behavior` are superclasses of every class. Although the mechanism involved may be a little confusing, the practical implication is easy to grasp — every class understands the instance messages defined by `Object`, `Class`, and `Behavior`.

This chapter provides an overview of the `Behavior` methods that are inherited by all classes, and so can be used to programmatically create and access methods, categories, pool dictionaries and variables for your classes.

**Creating and Removing Methods** (page 331)

describes the protocol in class `Behavior` for adding and removing methods.

**Information about Class and Methods** (page 335)

describes the protocol in class `Behavior` for examining the method dictionary of a class.

**ClassOrganizer** (page 337)

describes the protocol in class `Behavior` for examining, adding, and removing method categories.

**Handling Deprecated Methods** (page 339)

How to locate and clean up references to methods that have been deprecated.

## 17.1 Creating and Removing Methods

Class `Behavior` defines messages for creating methods and removing methods.

### Defining Simple Accessing and Updating Methods

Class `Behavior` provides an easy way to define simple methods for establishing and returning the values of instance variables. For each instance variable named by a symbol in the argument array, the message `compileAccessingMethodsFor: arrayOfSymbols` creates one method that sets the instance variable's value and another method that returns

it. These methods are added to the categories “Accessing” (return the instance variable’s value) and “Updating” (set its value).

For example, this invocation of the method:

```
Animal compileAccessingMethodsFor: #(#name)
```

has the same effect as the following topaz:

```
category: 'Accessing'
method: Animal
name
    ^name
%
category: 'Updating'
method: Animal
name: aName
    name := aName
%
```

You can also use `compileAccessingMethodsFor:` to define class methods for accessing class, class instance and pool variables, by sending `compileAccessingMethodsFor:` to the *class* of the class that defines the variables of interest.

The similar method `compileMissingAccessingMethods` will create accessing methods for any instance variables for which accessor methods with the standard selector do not already exist.

## Compiling Methods

Class Behavior defines the basic method for compiling a new method for a class and adding the method to the class’s method dictionary.

An invocation of the method has this form:

```
aClass compileMethod: sourceString
    dictionaries: arrayOfSymbolDicts
    category: aCategoryNameString
    environmentId: environmentId
```

The first argument, *sourceString*, is the text of the method to be compiled, beginning with the method’s selector. The second argument, *arrayOfSymbolDicts*, is an array of `SymbolDictionaries` to be used in resolving the source code symbols in *sourceString*. Under most circumstances, you will probably use your symbol list for this argument. The third argument names the category to which the new method is to be added.

*environmentId* specifies one of potentially multiple compile environments, provided for applications using multiple environments; it is normally 0. Unless you are using multiple environments, use 0; or omit this keyword, and it will default to an *environmentId* of 0.

You may use a simplified form,

```
aClass compileMethod: sourceString
```

which defaults to using the current symbol list, the category '(as yet unclassified)', and an environment Id of 0.

The following code compiles an accessor method named `habitat` for the class `Animal`, adding it to the category "Accessing":

```
Animal
  compileMethod:
    'habitat
    "Return the value of the habitat instance variable"
    ^habitat'
  dictionaries: (System myUserProfile symbolList)
  category: 'Accessing'
  environmentId: 0
```

When you write methods for compilation in this way, remember to double each apostrophe within the source string.

If `compileMethod: . . .` executes successfully, it adds the new method to the receiver. If the source string contains errors, this method signals a `CompileError`, with details on the specific causes of the failure.

## Removing Methods

You can remove a method by sending `removeSelector:` to a class or metaclass.

The following examples remove instance and class methods, respectively:

```
Animal removeSelector: #habitat
```

```
Animal class removeSelector:#newWithName:favoriteFood:habitat:
```

To remove all methods in a method category, as well as the category itself, use `removeCategory: categoryName`. For example,

```
Animal removeCategory: 'Accessing'
```

## Pragmas

A pragma is a literal selector or keyword message pattern that occurs between angle brackets at the start of a method after any temporaries. Pragmas are useful to provide metadata about methods.

Specifying a pragma is done using unary or keyword method selector syntax, and may include arguments that are literals - they may not include variables. For example,

```
<foo: 123 >
<foo: 5 bar: 'update'>
<bar>
```

While they follow method selector syntax, they are symbol literals within the method, not message sends. A primitive directives in GemStone looks like a pragma, but is not; `primitive:` is a reserved word in the first pragma in a method.

You may include multiple pragmas on one or multiple lines.

---

**Example 17.1 Compiling a method with pragmas**

---

```
Animal
  compileMethod:
    'checkHabitat
    <version: 2.1> <AnimalManagement>
    ^self habitat isPreferred'
  dictionaries: (System myUserProfile symbolList)
  category: 'Operations'
  environmentId: 0
```

---

**Pragma class**

The Pragma class provides a way to find out information about pragmas. An instance of Pragma references the method that defines it, and the keyword and argument or arguments.

Pragma class methods provide search capabilities. `Pragma >> allNamed:in:` returns a collection of all Pragmas with the given keyword in methods in the given class.

For example:

```
(Pragma allNamed: #AnimalManagement in: Animal)
  first method
  GsNMethod Animal>>checkHabitat
```

Sending `GsNMethod >> pragmas` will return an array of instances of Pragma.

## 17.2 Information about Class and Methods

Classes Behavior and Class define messages that let you discover information about a class, such as the class's instance variables, selectors, and categories. The class ClassOrganizer provides searching over methods in the image.

For full protocol, see the image.

### Information about the Class

Protocol in Class provides listing of superclasses and subclasses:

```
Class >> allSubclasses
Class >> allSuperclasses
Class >> allInstances
```

Each class also has a class comment and a category. This information can be accessed and updated using:

```
Class >> comment
Class >> comment: aString
Class >> category
Class >> category: aString
```

### Information about Instance, Class, and Shared Pool variables

Protocol in Behavior allows you to discover the class variables names, instance variable names, and shared pools defined for a given class, or for that class and all its superclasses.

```
Behavior >> classVarNames
Behavior >> allClassVarNames
Behavior >> instVarNames
Behavior >> allInstVarNames
Behavior >> sharedPools
Behavior >> allSharedPools
```

### Information about Method Selectors

Protocol in Behavior allows you to discover the selectors for the methods in a class, or in that class and its superclasses, and query on particular selectors.

```
Behavior >> selectors
Behavior >> allSelectors
Behavior >> includesSelector: aSelector
Behavior >> canUnderstand: aSelector
Behavior >> whichClassIncludesSelector: aSelector
```

## Accessing and Managing Method Categories

The methods in a class are associated with a method category, which is used to organize and document the method but does not affect execution. Method categories can be managed programmatically using the following methods in Behavior:

```
Behavior >> categoryNames
Behavior >> selectorsIn: categoryName
Behavior >> categoryOfSelector: selector
Behavior >> addCategory: categoryName
Behavior >> removeCategory: categoryName
Behavior >> renameCategory: categoryName to: newCategoryName
Behavior >> moveMethod: aSelector toCategory: categoryName
```

## Specific Methods

Each method is compiled into an instance of GsNMethod. You can query a class for its methods, and get source code and other information about the method.

To get the source code for a method, use:

```
Behavior >> sourceCodeAt: aSelector
```

To retrieve the compiled method itself, use:

```
Behavior >> compiledMethodAt: aSelector
```

This returns an instance of GsNMethod, from which you can then get source code. For example,

```
(Animal compiledMethodAt: #habitat) sourceString
```

Some GsNMethod methods that may be particularly useful are:

```
GsNMethod >> sourceString
GsNMethod >> sourceStringToFirstComment
GsNMethod >> selector
GsNMethod >> recompileFromSource
```

## 17.3 Transient Methods

Defining a method creates a persistent method on that class, and in most cases, the classes you are using are shared with other users. You can modify a method in an individual transaction, and you will be able to use your modification within that transaction, but any commit will make that change persistent and visible to other users.

For cases in which private modifications to methods in shared classes are useful, using special protocol allows you to compile methods that are only visible to a single session for the life of that session, and are unaffected by commit or abort.

The method must already exist, you cannot create an entirely new method this way. You must also have write permission for the object security policy of the class of the method; this avoids creating a security hole.

To create a transient method, use:

```
Behavior >> compileTransientMethod: dictionaries: environmentId:
```

This method is similar to `compile: dictionaries: environmentId:`, except that the `compiledMethod` is installed into the transient method dictionary for the receiver. The method created this way will exist in your image until logout. If you wish to remove them sooner, use one of the following:

```
Behavior >> removeTransientSelector: environmentId:
```

```
Behavior >> removeTransientSelector: environmentId: ifAbsent:
```

## 17.4 ClassOrganizer

ClassOrganizer provides useful methods to analyze your repository and perform operations such as searching for senders, receivers, or implementors, and string searches over method source. While usually you would perform these operations using GBS or another Smalltalk IDE, ClassOrganizer provide the ability to do customized analysis and reporting.

ClassOrganizer provides both reporting methods, which return formatted Strings, and query methods, which return collections of symbols or instances of GsNMethods that can be used for further analysis and reporting. See the image for the full set of protocol that ClassOrganizer understands.

For example, to get a report of all the senders of the `#dayOfWeek: selector:`

### Example 17.2 Senders of #dayOfWeek:

---

```
ClassOrganizer new sendersOfReport: #dayOfWeek:
%
Date >> next:
Date >> previous:
```

---

This example shows methods that are part of the GemStone kernel. In an application environment, it would also include your application senders. Application senders are usually of more interest than GemStone kernel class senders.

You can perform a search over the senders that are not from Globals using:

```
ClassOrganizer newExcludingGlobals sendersOfReport: #dayOfWeek:
```

If you want to perform more analysis on the methods or add additional reporting, send `sendersOf:`, which will return two arrays, the first an array of GsNMethods, the second the offset into the source code. For example:

### Example 17.3 Senders of #dayOfWeek: with source offsets

---

```
ClassOrganizer new sendersOf: #dayOfWeek:
%
anArray( anArray( GsNMethod Date>>next:, GsNMethod
Date>>previous:), anArray( 102, 137))
```

---

There are two senders of #dayOfWeek:; the method Date >> next:, at source code offset 102, and method Date >> previous:, at source code offset 137.

## Examples

The following example relates to the upgrade to 3.7, in which the GsHostProcess class is renamed to ObsoleteGsHostProcess and moved to ObsoleteClasses. This code looks up the class object, finds all references to the class (excluding kernel methods), and recompiles each method from source code.

### Example 17.4 Recompile methods referencing GsHostProcess

```
((ClassOrganizer newExcludingGlobals) referencesToObject:
 (ObsoleteClasses at: #ObsoleteGsHostProcess))
  do: [:aGsNMethod | aGsNMethod recompileFromSource].
```

The following code looks for all methods that are sent the message subclassResponsibility:, and makes sure all subclasses override that implementation. This example will return false positives, however, since it does not distinguish abstract classes.

### Example 17.5 Verify subclasses override subclassResponsibility:

```
| clsOrg meths report |
clsOrg := ClassOrganizer new.
report := String new.
meths := (clsOrg sendersOf: #subclassResponsibility:) at: 1.
meths do:
  [:srMeth |
  (clsOrg subclassesOf: srMeth inClass) do:
    [:subcls |
    ((srMeth selector ~= #subclassResponsibility) and:
    [(subcls whichClassIncludesSelector: srMeth selector) =
    srMeth inClass]) ifTrue: [
      report
        add: subcls name asString;
        add: ' does not override ';
        add: srMeth inClass asString;
        add: '>>';
        add: srMeth selector asString;
        lf
      ].
    ]
  ].
report
```

## 17.5 Handling Deprecated Methods

As GemStone features change, some methods may no longer be appropriate, or the method names may be incorrect or misleading. To allow obsolete methods to continue to function and provide a gentle transition to new methods, these obsolete methods may be deprecated.

Deprecated methods may be removed in future major releases, although some deprecated methods may remain in the image for longer periods for the convenience of existing applications.

Usually, deprecated methods will continue to work exactly as they did in the previous releases. However, in some cases the old behavior may not be meaningful in a new version; the deprecated method will continue to work as similarly as possible, but there may be differences.

Behavior may also change for existing methods. With any new release, you should review the Release Notes for changes in behavior as well as for newly deprecated methods.

Deprecated methods in GemStone are indicated by:

- ▶ Officially deprecated method include a call to `deprecated:`.
- ▶ Deprecated methods are in method category with a name including 'Deprecated'.
- ▶ Deprecation may be mentioned in the method comment. This may indicate an intention to deprecate.

Private methods, in a category with a name including 'Private', or which begin with an underscore, or which the method comment says private, may or may not be deprecated prior to removal. It is strongly recommended to avoid calling private methods.

Kernel methods that call `deprecated:` provide a string, which will generally include the class and selector, the version in which this method was deprecated, and the method that replaces it or some other indications of alternate action.

Since deprecated methods are subject to removal in major releases, it is important to keep your application updated so that no deprecated methods are called.

### Deprecated handling

By default, nothing happens when a deprecated method is called; the call to `deprecated:` has no action. This is most convenient when you first upgrade or convert to a new release of GemStone.

After you have updated your application references to deprecated methods, you can enable Deprecation handling, which can be configured to error or to log all calls to any deprecated methods. By running with this setting, you can locate and fix calls you may have missed, or confirm that you have indeed fixed all calls.

Changing deprecation handling can only be done by a user with write permission for the DataCurator object security policy. Once committed, the setting affects all users of the repository.

There are several levels of action that can be taken when a deprecated method is called:

- ▶ **Do nothing** — calls to deprecated methods are execute the same as any other method. This is the default.

To turn off any action on deprecation that you have previously enabled, execute:

```
Deprecated doNothingOnDeprecated
```

- ▶ **Raise an exception** – calls to deprecated methods signal an exception.

To enable this, execute:

```
Deprecated doErrorOnDeprecated
```

- ▶ **Log the call** – when a call to a deprecated method occurs, the call to the deprecated method is logged to the deprecation log file, and execution continues. There is no impact on the application, other than performance.

To enable this, execute:

```
Deprecated doLogOnDeprecated
```

- ▶ **Log the call stack** – when a call to a deprecated method occurs, the call to the deprecated method and the call stack are logged to the deprecation log file, and execution continues. There is no impact on the application, other than performance.

To enable this, execute:

```
Deprecated doLogStackOnDeprecated
```

## Deprecation log

When deprecations are configured to write to a log, a file named `DeprecatedPID.log` is created in the same location as a the gem log for an RPC login.

This file continues to grow and must be manually deleted. Logging methods or call stacks consumes resources and can noticeably affect performance, and use significant disk space. Methods called repeatedly, such as calls from within sort blocks, are particularly likely to impact the application.

## Listing deprecated methods

You can find all currently deprecated methods in a particular version by executing :

```
ClassOrganizer new sendersOfReport: #deprecated:
```

## Determining senders of deprecated methods

For each deprecated method, you can use development tools to determine if you have any senders within your application. In addition to GBS or other IDE tools, you can use `ClassOrganizer` methods.

For example, having determined that `setSegmentId:` has been deprecated, you can execute the following to find all senders of that selector within your application:

```
ClassOrganizer sendersOfReport: #setSegmentId:
```

Since deprecation only applies to a method associated with a specific class, and this search looks for all senders of the selector, you will have to examine the list to determine if the call is actually deprecated. This is the consequence of how typing is handled in Smalltalk. For example, `String >>+` is deprecated, but `Integer >>+` is not.

Also note that this technique will not find methods that are symbols sent to `perform:` statements, in code in client applications that is sent to GemStone for execution, or in topaz scripts.

## 17.6 File In and File Out of GemStone code

To archive your application or transfer GemStone classes to another repository you can *file out* GemStone Smalltalk source code for classes and methods to a text file. To port your application to another repository, you can *file in* that text file, and the source code for your classes and methods is immediately available in the new repository.

### Fileout

Methods in behavior allow you to file out a class, category, or method. For example, to file out a single class named Customer:

```
| myFile |
myFile := GsFile openWriteOnServer: 'CustomerClassFileout.gs'.
myFile ifNil: [^GsFile serverErrorString].
Customer fileOutClassOn: myFile.
myFile close.
%
```

Using ClassOrganizer, you can file out all the classes and methods in a SymbolDictionary, ordered correctly for filein. For example, to file out UserGlobals:

```
| myFile |
myFile := GsFile openWriteOnServer: 'UserGlobalsFileout.gs'.
myFile ifNil: [^GsFile serverErrorString].
ClassOrganizer new fileOutClassesAndMethodsInDictionary:
    UserGlobals on: myFile.
myFile close.
%
```

File out can also be done using the topaz command **fileout**. See the *Topaz User's Guide* for more information.

### Filein

File in is done using topaz **input** command, facilities provided by GBS, or the GemStone class GsFileIn.

For example, to file in the fileout of UserGlobals from the previous example:

```
topaz 1> input UserGlobalsFileout.gs
```

### GsFileIn

GsFileIn supports filing topaz-format GemStone source files into the image, without the use of topaz.

To use, send one of the *from\** class methods, specifying the path and file; this reads the file and also performs the filein, and returns itself.

```
GsFileIn fromGemHostPath: '/mynode/users/gadmin/myCode.gs'
GsFileIn fromGciClientPath: 'C:\sourceCode\myCode.gs'
```

The Gem Host is the server; this name clarifies that it is the host that the Gem is on, not the host for the Stone. The Gci Client is the path on the node on which the client is running on, including a Windows path when running topaz or a GBS Smalltalk environment on Windows.

Note that not all topaz commands are supported. The GsFileIn input file may include any of the commands that are normally present in GemStone's fileout results, and a few others. Input files may also be hand-coded provided they adhere to these limitations.

The following are supported and perform much the same function as in topaz:

**doit printit run nbrun send**

**input**

**category:**

**classmethod classmethod:**

**commit abort**

**env**

**fileformat**

**method method:**

**removeallmethods removeallclassmethods**

**set** specific options: **compile\_env class category sourcestring class enableremoveall project package** (other set commands are ignored)

The following are acceptable in a file, but are ignored:

**expectvalue expecterror iferr iferr\_clear iferr\_list errorcount**

**fileout output**

**display omit level limit**

**time remark status**

**login logout list**

Other topaz commands will error.

# GemStone System Features

---

This chapter describes some features and internal structures that provide specialized behavior.

These structures are intended for use by experienced GemStone programmers.

**Hidden Sets** (page 343)

Describes HiddenSets, a non-persistent way to manage objects using bitmaps.

**SessionTemps and access to Session State** (page 345)

Ways to keep session-temporary data available for the life of a session.

**Shared Counters** (page 345)

Integer counters that can be shared between sessions. Both non-persistent and persistent counters are available.

**GsEventLog** (page 347)

Easily track errors, trace code and record objects in a multi-session environment.

## 18.1 Hidden Sets

GemStone's internal bitmap structures that hold OOPs efficiently and with minimal memory demands are encapsulated by instances of GsBitmap. Specific internal bitmaps are accessible to the image via the hidden set API in System class.

For historic reasons, some repository scan methods are available that write results temporarily to a hidden set. These results can be converted to a GsBitmap for processing, with the advantages of avoiding overloading temporary object memory. While there are replacement methods that return GsBitmaps directly, you may still encounter methods that put the results in a hidden set, rather than returning a GsBitmap

For example, while the method `allInstances:` returns a GsBitmap, the equivalent method `listInstancesToHiddenSet:` puts the results of a `listInstances` operation into a specific hidden set (in this case, hidden set 1) The following code shows the call to this method, and how to use hidden set protocol to migrate each object:

**Example 18.1**


---

```

SystemRepository listInstancesToHiddenSet: MyClass.
bm := GsBitmap newForHiddenSet: (GsBitmap hiddenSetIdAsSymbol: 1).
[ bm size > 0 ]
  whileTrue:
    [
      | resultBatch |
      resultBatch := bm removeCount: 1024.
      resultBatch do: [:aMyClass | aMyClass migrate].
      System commitTransaction.
    ].

```

---

For more on working with instances of GsBitmap, see “GsBitmap” on page 66.

The specific set symbolic names that are accessible from GsBitmap are listed in the method `GsBitmap class >> hiddenSetSpecifiers`. Many of these are read-only.

Note that the method `System >> HiddenSetSpecifiers` lists integer indexes of hidden set IDs. This list is ordered differently than the symbolic names in `GsBitmap specifiers`; the index offsets are not interchangeable. You can look up a GsBitmap identifier based on a hidden set integer id using the method `GsBitmap class >> hiddenSetIdAsSymbol::`

**Sets still accessed via System methods**

The following internal sets may be accessed using System methods and respond to some hidden set protocol:

**NotifySet**

The #NotifySet is System hidden set 25.

The #NotifySet has an interface in class System with methods in the Notification category; `clearNotifySet`, `addToNotifySet::`, etc.

**ExportedDirtyObjs and TrackedDirtyObjs**

#ExportedDirtyObjs is System hidden set 22.

#TrackedDirtyObjs is System hidden set 23.

The methods `gciDirtyObjsInit` and `gciTrackedObjsInit` enable dirty and tracked object sets, respectively. Use `getAndClearGciDirtySet:into:` to retrieve the results.

**PureExportSet and GciTrackedObjs**

#PureExportSet is System hidden set 39

#GciTrackedObjs is System hidden set 40.

Methods in class System, category ‘Gci Set Support’ allow you to add and remove objects from these hidden sets. Note you must initialize the dirty set using `gciDirtyObjsInit` before adding or removing from `GciTrackedObjs`.

## 18.2 SessionTemps and access to Session State

Data in GemStone is either temporary or persistent. While most temporary data is only retained for as long as the method is executing, or until the session updates its commit record by committing or aborting, you may sometimes want data that is not persistent and not shared, so does not risk transaction conflicts, but remains unaffected by transaction status.

Session-specific data of this kind can be put into SessionTemps. `SessionTemps current` provides access to a kind of SymbolDictionary; elements in the SessionTemps dictionary remain until the session logs out or exits, are not affected by commit or abort, and are not visible outside of the session.

For example, if you wish to open a log file and leave it open:

```
SessionTemps current at: #Log put:
  (GsFile openAppend: 'myFile.log')
```

Actual code, of course, would do more error checking. To write to the file, use code similar to this:

```
(SessionTemps current at: #Log)
  nextPutAll: 'a message for the log file'.
```

Objects in SessionTemps use temporary object memory, and the objects cannot be removed from memory by in-memory garbage collection. While there is no limit on how much data can be stored in SessionTemps, if your session reaches the memory limit and exits, that data will be lost.

### SessionState

SessionTemps uses a slot in the internal Session State structure, which is primarily provided for use by the kernel. Access to customer-available SessionState slots is provided primarily for legacy uses, but may be useful depending on application requirements.

SessionState is accessed by integer index, with slots 1 to 1994 available for use. The SessionState array is variable size, and will grow as needed.

The following methods can be used to read and update SessionState:

```
System >> sessionStateAt: anIndex
System >> sessionStateAt: anIndex put: anObject
System >> sessionStateSize
```

## 18.3 Shared Counters

There are two types of Shared Counters available; AppStat Shared Counters and Persistent Shared Counters.

AppStat Shared Counters provide a way for sessions on the same shared cache to read and update a set of counters. These counters are stored in the shared cache and are not persistent across cache restart. They are not visible to sessions on different shared page caches - that is, a session on the stone's cache and a session on a remote cache cannot access the same Shared Counters. Values are also not recoverable from tranlogs.

Persistent shared counters are stored in the repository, and are visible to all sessions on all shared caches. On repository recovery or restore, the values of persistent shared caches are restored.

## AppStat Shared Counters

Shared counters allow multiple sessions on the same SPC to read and update a common counter value.

Shared counters are indexed from 0 to (System numSharedCounters - 1), which is set by the configuration parameter SHR\_PAGE\_CACHE\_NUM\_SHARED\_COUNTERS. The default value for SHR\_PAGE\_CACHE\_NUM\_SHARED\_COUNTERS is 1900. Each counter is protected by a unique spinlock. The index of the first counter is 0.

Shared counters may be set to any signed 64 bit integer value, in the range:

$-2^{63}$  (-9223372036854775808) to  $2^{63} - 1$  (9223372036854775807)

If you increment or decrement so that the result would be outside the range of a signed 64-bit integer, the value will be set to the minimum or maximum; directly setting an out of range value will result in an error.

Shared counters are transient, that is, they do not persist across cache restart.

Shared counter values are recorded by statmonitor when using the -n option and recorded as AppStats.

The following methods may be used to read and update shared counters. For details, see the method comments in the image.

```
System class >> numSharedCounters
System class >> sharedCounter: index
System class >> sharedCounter: index setValue: value
System class >> sharedCounter: index incrementBy: amount
System class >> sharedCounter: index decrementBy: amount
System class >> sharedCounter: index decrementBy: amount
    withFloor: floorValue
System class >> sharedCounterFetchValuesFrom: firstCounter
    to: lastCounter
```

## Persistent Shared Counters

Persistent shared counters allow all sessions in a repository to read and update a set of counters. Persistent shared counters are globally visible to all sessions on all shared page caches.

There are 1536 persistent shared counters, numbered from 1 to 1536. The index of the first counter is 1.

Persistent shared counters may be set to any signed 64 bit integer value, in the range:

$-2^{63}$  (-9223372036854775808) to  $2^{63} - 1$  (9223372036854775807)

No limit checks are done when incrementing or decrementing a counter. If you increment or decrement so that the result would be outside the range of a signed 64-bit integer, the value will “rollover” and the overflow bits will be lost. Directly setting an out of range value will result in an error.

Values of all persistent shared counters are stored in the repository and in tranlog records. They are persistent through Stone restart, and recovered on Stone crash, restore from backup, and restore from tranlog.

Persistent shared counters are independent of database transactions. Updates to counters are visible immediately and not affected by aborts.

Each update to a persistent shared counter causes a roundtrip to the Stone; but reading the value is handled by the gem (and the page server, if remote), and does not cause a roundtrip to the stone.

The following methods may be used to read and update persistent shared counters. For details, see the method comments in the image.

```
System class >> numberOfPersistentSharedCounters
System class >> persistentCounterAt: index put: value
System class >> persistentCounterAt: index
System class >> persistentCounterAt: index incrementBy: amount
System class >> persistentCounterAt: index decrementBy: amount
```

## 18.4 GsEventLog

GsEventLog is a shared, reduced-conflict logging tool that allows multiple sessions to:

- ▶ record errors and informational messages to a single location
- ▶ flexibly include objects as well as text
- ▶ easily and flexibly query and sort by particular attributes.

A GsEventLog class variable holds a instance of GsEventLog, which in turn holds a collection of log entries, which are instances of kinds of GsEventLogEntry. It is allowed for user applications to create custom subclasses of GsEventLogEntry.

Each event has a priority. Built in priorities are fatal, error, warning, info, debug, and trace; these are stored internally as integers.

The entries are stored in an instance of RcArray, allowing concurrent writes of log entries. Note that the order of elements is based on the order in which the commits occurred, while entry timestamps reflect the time at which the entry was created.

### Adding events

GsEventLog may hold both application (user) events and system events. User entries can be added in two ways: class convenience methods such as `logError:`, `logInfo:`, etc., or by creating an instance of GsEventLogEntry and sending `addToLog`. A commit is required to make the log entry persistent.

System events should be added only by GemStone code. In this release, GemStone code does not write System events.

### Querying and reporting

To create a string containing text representation of the entire contents, send

```
GsEventLog current report
```

To search for a subset of entries with particular attributes

```
(GsEventLog current entriesSatisfying: aBlock) report
```

## Deleting events

GsEventLog is not reduced conflict for delete. It is recommended to lock the log using GsEventLog >> acquireGsEventLogLock. The lock is cleared automatically on commit.

You can clear all events using GsEventLog current deleteAllEntries. By querying for specific methods, you can delete those methods using deleteEntry: or deleteEntries:.

It is possible to restrict modifying or removing events. To do this, execute

```
GsEventLog entriesUnmodifiable
```

After this is executed, new entries to the log are made invariant and the standard delete methods will not delete them. However, they are not protected from delete using private delete protocol.

System events are also protected from modification or delete, other than using private delete protocol.

### Example 18.2 Debugging code using GsEventLog

```
topaz 1> printit
[GsEventLog logDebug: 'About to perform divide by zero'.
 1 / 0.
 GsEventLog logDebug: 'After divide by zero'.
 true]
  on: Error
  do: [:ex | GsEventLog logObject: ex. ex resume].
%
true
topaz 1> printit
GsEventLog current report
%
2023-08-15 16:40:21.7112    3440742 Debug    About to perform divide by
zero
2023-08-15 16:40:21.7115    3440742 Trace    a ZeroDivide occurred
(error 2026), reason:numErrIntDivisionByZero, attempt to divide 1 by
zero
2023-08-15 16:40:21.7116    3440742 Debug    After divide by zero
```

# The Foreign Function Interface

---

This chapter describes the Foreign Function Interface (FFI) classes and methods, and how you can use them to build and interface to an existing C library.

**Overview of the Foreign Function Interface (page 349)**

The purpose and use of the FFI.

**Using the FFI (page 350)**

Describes the specific steps and classes involved in using the FFI.

**Example using Zlib (page 359)**

Simple example using Zlib.

**Example using CCalloutStructs and variable arguments with RabbitMQ® (page 366)**

An example illustrating the use of CCalloutStructs and variable arguments.

## 19.1 Overview of the Foreign Function Interface

For certain applications, you may need to provide functionality that is not readily available within GemStone Smalltalk. This may include functionality that can be provided by third-party products that provide a C interface, such as the zlib compression library, data encryption, Oracle, MySQL, or other databases, and so on.

To interact with third-party products such as these, you can use the FFI to make C library calls from within GemStone Smalltalk. Using the FFI, you can access C functions in external libraries without the need to write and compile UserAction libraries.

The FFI is composed of a number of classes: CLibrary, CHeader, CDeclaration, CCallout, CCalloutStructs, CByteArray, and others.

These classes can be used to hand-craft access to a single or a few simple functions; they also allow you to parse the C header file to create new classes for C libraries, with Smalltalk methods that access each C function, and new classes for complex C data types.

Parsing C header files requires that you have gcc/g++ installed in your development environment. Once you have created C library interface classes for your application, the C compiler is not needed to execute the FFI code.

The GemStone image includes the class `GciLibrary`, which is an FFI interface class that allows calls to GemStone's C interface (GCI). You may examine the class `GciLibrary` as an example of the results of creating an FFI wrapper class.

## 19.2 Using the FFI

Making a C function call involves defining an instance of `CCallout` for the specific function, and then invoking that `CCallout` with the appropriate arguments at runtime.

The `CCallout` definition includes the specifications of the various argument and return types, and the name of the compiled library that supports that function. The defined `CCallout` instances can be persistent and are reusable.

The FFI supports parsing the header file that accompanies C libraries, to simplify and automate the process of defining `CCallouts` and making these easy to use from GemStone Smalltalk.

### CLibrary: defining the compiled C or C++ library

To make a call to a third party or custom C or C++ library, you will need the compiled C library file (`.so` or `.dyn`). The file is loaded into the gem or linked process by a primitive which calls `dlopen`; See "man `dlopen`" for the operating system details of library search order. An absolute path can be used, or the environment variable `LD_LIBRARY_PATH` can be used to specify a search path.

Instances of `CLibrary` are created using:

```
CLibrary class >> named:libraryName
```

passing in the path and name of the C shared library to be loaded. The platform-specific extension (such as `.so`) is optional.

To generate support classes and methods, the FFI interface uses the header file, which is also located by name. These are normally made available along with compiled libraries.

### CCallout: function definitions in GemStone

Functions in a C library are accessed via instances of `CCallout` or a subclass, including `CCalloutStructs`; there is one `CCallout` instance for each function you wish to call directly from the FFI. The `CCallout` is created with the function name and declarations for the result and argument types.

For example,

```
myCCallout := CCallout
  library: (CLibrary named: 'myLib.so')
  name: 'DoSomething'
  result: #'uint64'
  args: #(#'char*' #'uint64').
```

You will need a `CCalloutStructs` if there is a structure (passed by value) either as an Argument or a result. These are handled somewhat differently and have more restrictions than `CCallouts`; see "CCalloutStructs: Using C Structures passed by value" on page 355.

Creating `CCallouts` by hand is tedious and error prone, especially if variable arguments or structures are involved. In general, you should parse the C header file, using `CHeader` class wrapper: \* methods, to create the `CCallout` definitions.

The CCallout definitions include the C types of all the arguments (except variable arguments, which are declared when the function is called), and the type of the result. For the supported argument type symbols, see “C type symbols” on page 357.

To ensure you have the correct version of the function, you may send `version:` to the CCallout; this string will be used as third argument to `dlsym()`.

Note that while UserActions creation performs checks of your code against function prototypes of the external library, the FFI does not do this checking. You are responsible for ensuring that the CCallout definition matches the C call.

## CHeader: Parsing the C header file

The class CHeader allows you to parse the C header file/s that corresponds to the C library you wish to call. You can examine the functions in that library, and define a class and instance methods to invoke the C functions.

Header files are parsed using the methods:

```
CHeader class >> path: headerFileOrPath
```

You may pass in the full path and name of a header file, or just name of the header file. If the *headerFileOrPath* does not fully specify a file, will look in the system search path, including the current directory, `/usr/include/`, and `/usr/local/include/`. These paths are also used to locate any files that are parsed due to include statements. You may lookup the current search path using: `CPreprocessor new allSearchPaths`

```
CHeader class >> path: headerFileOrPath searchPath: aPath
```

Search for *headerFileOrPath* and its include files by looking first in *aPath*, and then in the current directory and the system search path.

```
CHeader class >> path: headerFileOrPath searchPaths: collOfPaths
```

Search for *headerFileOrPath* and its include files by searching first in the directories in *collOfPaths* in order, then in the current directory and the system search path.

Parsing the header file - that is, creating the instance of CHeader - does not give you the location of the actual C library file that you will be calling. To define or invoke the call to an FFI function, you will also need to specify the compiled C library name.

Parsing the file does not itself create a class or methods.

Cheader parsing is also used to create classes and methods for struct data types; see “CByteArray: Allocating memory for data” on page 352.

## CDeclaration: Finding details on a C function

The instance of CHeader file that is returned by the parsing methods contains CDefinitions for the functions, data types, and other C elements in the header file, and all included header files. The specific function calls are presented as instances of CDeclaration, and you may view the source code and the processed representation of the function prototype.

For example:

```
((CHeader path: 'myLib.hf') functions at: #DoSomething)
  source
%
extern uint32 DoSomething (char * __buf, size_t __size)
  __attribute__ ((__nothrow__ , __leaf__));

((CHeader path: 'myLib.hf') functions at: #DoSomething)
  printString
%
extern int32 DoSomething(uint8 *dest, uint64 *destLen, const
  uint8 *source, uint64 sourceLen)
```

## Creating a Class and Methods from the CHeader

The instance of CHeader can be used to create classes and methods for the CDeclaration types and functions that were parsed. The resulting class contains methods that allow you to invoke the C function from a Smalltalk method. The class will include reusable persistent instances of CCallout and CCalloutStructs.

- ▶ The name of the class to create is supplied as an argument. Method variants exist that create a default name, which is the name without 'lib'; see the image for details.
- ▶ When creating this class, the library name is an additional argument, and is stored in a class variable. Method variants exist that allow to you include an expression to execute to lookup the library name; see the image for details.
- ▶ The instance of CHeader includes definitions for every function in the header file and in every included header file. You may limit the set of functions that are made accessible in the Smalltalk class using a select block. The select block should evaluate to a Boolean that indicates whether or not to include the particular function.

For example, the following code parses a C header file and creates a class named MyLibrary, that contains methods corresponding to all functions for which the first letter is uppercase.

```
| wrapperClass |
wrapperClass := (CHeader path: 'myLib.hf')
  wrapperNamed: 'MyLibrary'
  forLibraryAt: 'myLib.so'
  select: [:each | each name first isUppercase].
wrapperClass initializeFunctions.
UserGlobals at: wrapperClass name put: wrapperClass.
```

The instance methods created by this code invoke a persistent instance of CCallout or CCalloutStructs for the specified library name, so you can invoke the method with only the arguments that are specific to that function.

## CByteArray: Allocating memory for data

The FFI, in general, allows you to work with C objects without needing to yourself malloc and free memory.

A CByteArray represents an allocation of C memory. When objects such as pointers or strings are passed to or from C functions, creating a CByteArray, using a method that

invokes `CByteArray gcMalloc:`, ensures that the memory will be valid following the call; and that once the object is dereferenced, the associated heap memory will be freed.

When creating a String, for example, using `CByteArray class>>withAll:`, the operation mallocs memory for the string, fills the `CByteArray` from the string contents, and adds the terminal 0 byte.

You may allocate memory using `CByteArray >> malloc:`, which is not automatically freed; this may be useful for some interfaces that require long-lived data types.

You can add to and retrieve Smalltalk objects from a `CByteArray` using data type specific methods. For example:

```
CByteArray >> int64At: zeroBasedOffset
CByteArray >> int64At: zeroBasedOffset put: anInteger
```

See the image for the many other available access methods, including methods to add and extract Arrays and adding and extracting strings as UTF-8 bytes.

## Complex data types

For struct and class data types, the `CHeader` that results from parsing C header files can be used to create subclasses of `CByteArray`, using `CHeader>>wrapperForTypeNamed:`. This allows creating instances of the correct size for the struct, and retrieving individual elements using instance methods rather than offset into the `CByteArray`.

For example, the GCI includes a struct type `GciFetchObjInfoArgsSType`. To see the `CDeclaration` for this struct:

```
((CHeader path: '$GEMSTONE/include/gci.hf') types at:
  'GciFetchObjInfoArgsSType') source
%
typedef struct { int64 startIndex; int64 bufSize;
  int64 numReturned; GciObjInfoSType *info; ByteType *buffer;
  int retrievalFlags; BoolType isRpc; } GciFetchObjInfoArgsSType;
```

A GemStone data class can be created for this definition by executing:

```
UserGlobals at: #GciFetchObjInfoArgsSType put:
  ((CHeader path: '$GEMSTONE/include/gci.hf')
   wrapperForTypeNamed: 'GciFetchObjInfoArgsSType')
```

The resulting class `GciFetchObjInfoArgsSType` is a subclass of `CByteArray` that includes instance creation methods that `gcMalloc` 48 bytes for the new instance,

```
GciFetchObjInfoArgsSType class >> new
  ^self gcMalloc: 48
```

and has (logical) getter and setter instance methods for `startIndex`, `bufSize`, and so on, that access the correct offsets within the `CByteArray`.

```
GciFetchObjInfoArgsSType >> bufSize
  ^ self int64At: 8
```

## Making FFI calls

To call the C function via the FFI, you will create an instance of `CCallout` that describes the C function, and send `#callWith:` to that instance of `CCallout`, passing in an array containing the argument values.

These arguments may include instances of CByteArray or CPointer, which handle allocating memory and ensure that the memory remains allocated while needed.

## Executing a CCallout created by hand

For example, if the definition was created by hand:

```
| doSomethingCallout |
doSomethingCallout := CCallout
    library: (CLibrary named: 'myLib.so')
    name: 'DoSomething'
    result: #'uint64'
    args: #(#'char*' #'uint64').
```

Then it can be invoked:

```
result := doSomethingCallout callWith: { 'string1' . 22 }
```

## Execute CHeader-generated definition

When using CHeader to generate the CCallouts and store them in a class, you will invoke the generated instance method; each argument is separate. The instance method names that are generated use `_:` to indicate each argument, in the order required by the C header file's function definition.

To execute the function DoSomething in the generated MyLibrary class, for example:

```
result := MyLibrary new DoSomething_: 'string2' _: 33
```

Internally, when using `wrapperNamed:forLibraryAt:select:` to generate the MyLibrary class, the following methods pairs were created, one pair for each function that matched the select block:

```
MyLibrary class >> initializeFunction_DoSomething_inLibrary:
    Function_DoSomething := CCallout
        library: (CLibrary named: 'myLib.so')
        name: 'DoSomething'
        result: #'uint64'
        args: #(#'char*' #'uint64').
```

This method is executed by `initializeFunctions`, which stores the CCallout in a Class variable. The CCallout is accessed via the corresponding instance method. The method that was called in the example is defined as:

```
MyLibrary >> DoSomething_: theString_: theNumber
    Function_DoSomething callWith: { theString . theNumber }
```

## Arguments and Return Values

When a CCallout is configured, the arguments are described as an array of C type symbols. Then, when the CCallout is used to make a C call, the values for these arguments are provided as an array that matches the CCallout's configuration in type, size, and order.

for example, in the CCallout instance creation:

```
... args: #(#'char*' #'uint64').
```

And when this CCallout is invoked:

```
myCallout callWith: { 'string1' . 22 }
```

See "C type symbols" on page 357 for the supported argument type symbols.

## String type arguments

For a function argument of type `char*`, `const char*`, `UChar*` or `const UChar*`, an instance of `CCallout` will accept an argument of the appropriate kind of string, and copy the body of that string to C memory. However that C memory data is only valid for the duration of the call; after the function completes, the reference to memory is no longer valid.

If the string will be updated by the function or needs to be referenced later, a `CByteArray` should be created for the string arguments. Creating these via `CByteArray class >> withAll:` allocs memory, adds the terminal nul, and performs tracking so that this memory is later freed.

For example:

```
( CByteArray withAll: 'abc' ) printString
%
aCByteArray size=4 gcFree=0x3 dead=false address=0x561d9150ade0
```

In this result, note the size includes the terminal zero byte, and the `gcFree` indicates that `free()` will be called when the instance is garbage collected.

## Variable Arguments

Functions may also have variable arguments, as indicated by `,...` in the C function header. When making a call to a function with variable arguments, the `callWith:` array should include, in addition to any regular arguments (which are only the value), two elements for each variable argument; the type definition for the argument, followed by the argument value.

For example, if the C header function is defined as:

```
myfunct( const char* arg1, int arg2,...)
```

than to pass in two string variable arguments, along with the two fixed arguments:

```
myCCallout callWith: {
    'string1' .
    22 .
    #'const char*' . 'string2' .
    #'const char*' . 'string3' }
```

Functions using `varArgs` generally may have a maximum of 20 variable arguments. This limit is lower if native code is disabled for this session, as described under “Limitations with native code disabled” on page 359.

For an example of a function that accept variable arguments, see under “Example using `CCalloutStructs` and variable arguments with `RabbitMQ@`” on page 366.

## CCalloutStructs: Using C Structures passed by value

The class `CCalloutStructs` is a kind of `CCallout`, that allows you to use C Structures passed by value (rather than by reference). `CCalloutStructs` can be used on Linux and Mac only, on x86 or ARM platforms.

The C structs that are used in a `CCalloutStructs` are in the form of a `CByteArray`, which must be the correct size to accommodate the data types within the struct. These should be created using `CHeader >> wrapperForTypeNamed:.` This ensures that the size of the `CByteArray` is correct, and allows you to set and fetch the individual struct fields by name rather than offset.

While you can hand-build the definition for a CCalloutStructs, it is recommended to generate these using CHeader wrapper\* methods, to ensure the definition and function matches correctly. The CHeader wrapper\* methods that generate CCallouts will generate a CCallout or CCalloutStructs, and the instance method that invokes it, depending on the declaration in the header file.

When the C function returns a struct, there are special considerations; these do not apply when an argument is a struct, but the return value is a simple data type.

- ▶ the generated instance method has an initial argument, in the first position, to hold the struct result. The argument passed in must be an instance of a kind of CByteArray of the correct size. It is strongly recommended that this be an instance of a data type subclass of CByteArray that was generated for the specific struct type that is returned
- ▶ If invoking this by hand, the method CCalloutStructs callWith:structResult:errno: must be used.

For an example of a function that return a struct, see under “Example using CCalloutStructs and variable arguments with RabbitMQ®” on page 366.

## CPointer

CPointer encapsulates a C pointer that does not have auto-free semantics. New instances are created by CFunction calls with result type #ptr, and are also used for certain arguments of CFunctions.

## Errno Handling

C functions may set errno following execution. To access this value, instead of using callWith: to call the function, use callWith:errno:, and pass in a 1-element array (not invariant). The errno for the function execute will be placed in the first position in this array.

If the element in the first position of the 1-element array is a SmallInteger, errno will be initialized to this value before making the C call.

For example:

```
| errnoArray myErr myCCallout |
errnoArray := { 0 }.
myCCallout callWith: #( 3 0 ) errno: myErr.
myErr := errnoArray at: 1
```

## CCallin: Creating and invoking callbacks

A CCallin represents a signature for a C function to be called by C code. The resulting CCallin may be used as a type within the *argumentTypes* array when defining a CCallout.

For example, with the following C definitions compiled into the library myLib.so,

```
typedef double (*CbdfType)(double p1, int64 p2);

extern "C" double myCallBack( int64 a1, double a2, CbdfType a3)
{
    double d = (*a3)(a2, a1); return d;
}
```

The following examples creates the callback and invokes it:

### Example 19.1 Callback example

```

topaz 1> run
| myCallin myCallback myResult |
myCallin := CCallin
    name: 'doMath'
    result: #double
    args: { #double . #int64 }.

myCallback := CCallout
    library: (CLibrary named:'myLib.so')
    name: 'myCallBack'
    result: #double
    args: { #int64 . #double . myCallin } .

myResult := myCallback callWith: {
    115 .
    338.0 .
    [:a1 :a2 | a1 + a2 ]
}.
myResult
%
453.0

```

## C type symbols

Table 19.1 lists the symbols used for creating *resType* (result type) and *argumentTypes* arguments when creating CCallouts. See the comments in the method CCallout class >> library:name:result:args: for details.

**Table 19.1 C Type**

	Return type	Argument type
#int64	Integer	Integer representable as a C int64
#uint64	Integer	Integer representable as a C uint64
#int32	SmallInteger	Integer representable as a C 32-bit int
#uint32	SmallInteger	Integer representable as a C 32-bit uint
#int16	SmallInteger	Integer representable as a C 16-bit short
#uint16	SmallInteger	Integer representable as a C 16-bit ushort
#int8	SmallInteger	Integer representable as a C 8-bit signed char
#uint8	SmallInteger	Integer representable as a C 8-bit uchar
#bool	true or false	Boolean
#double	SmallDouble or Float	SmallDouble or Float

Table 19.1 C Type (Continued)

	Return type	Argument type
#float	SmallDouble or Float	SmallDouble or Float
#'char*'¹	nil or a String	a non-nil String or Unicode7. The body is copied to C memory before the call, and a zero byte appended; and copied from C memory (and possibly grown/shrunk) after call. C memory will not be valid after the call finishes. A CByteArray can be used if the argument should be available after the call.
#'const char*'¹		nil (to pass NULL), String, Unicode7, or Utf8. The body is copied to C memory before the call, and a zero byte appended. C memory will not be valid after the call finishes. A CByteArray can be used if the argument should be available after the call.
#void	nil	
#'UChar*'¹		Unicode7, Unicode16, or Unicode32. The body is copied to C memory in UTF16 encoding before the call, and a zero byte appended; and copied from C memory (and possibly grown/shrunk) after call, which may cause the class to change. C memory will not be valid after the call finishes. A CByteArray can be used if the argument should be available after the call.
#'const UChar*'¹		nil (to pass NULL), or Unicode7, Unicode16, or Unicode32. The body is copied to C memory in UTF16 encoding before the call, and a zero byte appended. C memory will not be valid after the call finishes. A CByteArray can be used if the argument should be available after the call.
#'ptr'¹	nil or a CPointer	One of: <ul style="list-style-type: none"> <li>▶ nil, in which case C NULL is passed.</li> <li>▶ a CByteArray, in which case the address of the body is passed.</li> <li>▶ a CPointer, in which case the encapsulated pointer is passed.</li> </ul>
#' &ptr'¹		a CPointer. The CPointer's value will be passed and updated on return.
#'struct*'¹	a CByteArray (CCalloutStructs only)	a CByteArray (CCalloutStructs only)

## Limitations with native code disabled

If the generation of native code is disabled, there are further limitations:

- ▶ Functions using `varArgs` may have a maximum of four fixed and 10 total arguments.
- ▶ Functions not using `varArgs` are limited to a maximum of 15 total arguments.
- ▶ Arguments and results of C type float are not supported.
- ▶ Functions with one or more args of C type double are limited to a maximum of four arguments.
- ▶ `CCallin` cannot be used

Native code generation is on by default, but may be configured to be disabled or becomes disabled when breakpoints are set. See the *System Administration Guide* for more information on native code generation.

## 19.3 Example using Zlib

The following examples use `zlib`, a commonly available software library for data compression that is available on many platforms. The examples are based on `zlib v1.2.8` on Linux; with other versions of `zlib.h` or on other platforms, you may need to experiment. Documentation on `zlib` is available at <http://zlib.net/manual.html>.

### Parsing the header file

For this example, the `ZLib` header file is installed in `/usr/include/zlib.h`, which is on the machine search path.

The C library is install at `/lib/x86_64-linux-gnu/libc.so.6`.

The following example analyzes a a header file and stores the result in a variable in `UserGlobals`:

#### Example 19.2 Create a CHeader for zlib.h

---

```
topaz 1> printit
UserGlobals at: #'ZLibHeader' put:
  (CHeader path: '/usr/include/zlib.h').
%
```

---

Given the location of `zlib.h` on the C compiler search path `/usr/include/`, the following expressions can also be used to lookup the header file:

```
(CHeader path: 'zlib.h')
(CHeader path: 'zlib.h' searchPath: '/usr/include/')
(CHeader path: 'zlib.h' searchPaths: {'/usr/include/'})
(CHeader path: 'include/zlib.h' searchPath: '/usr/')
(CHeader path: 'include/zlib.h' searchPaths: {'/usr/'})
```

Once you have a `CHeader` object, you can get information about the various things defined in the header file and those it includes.

**Example 19.3 CDeclaration for compress()**


---

```

topaz 1> printit
(ZLibHeader functions at: 'compress')
%
a CDeclaration
  header          a CHeader
  name            compress
  storage         extern
  linkageSpec     nil
  type            int32
  count           nil
  pointer         0
  fields          nil
  parameters     a Array
  enumTag         nil
  isConstant      false
  includesCode    false
  isVaryingArgCount false
  bitCount        nil
  source          extern int compress (Bytef *dest, uLongf
                 *destLen, const Bytef *source, uLong sourceLen)
  file            /usr/include/zlib.h
  line            1060

```

---

While the `compress()` function is directly in `zlib.h`, this isn't necessarily the case for all functions in `ZLibHeader`. Functions that are defined in any header file that is `#included` in the parsed header file also will have definitions in the instance of `CHeader`.

For example, on Linux the `zlib.h` file `#includes` `unistd.h`, so functions such as `getcwd()` also have definitions in the instance of `CHeader`:

```

topaz 1> printit
(ZLibHeader functions at: 'getcwd') file.
%
/usr/include/unistd.h

```

On other platforms, `zlib.h` may not `#include` `unistd.h`. In this case, the definition is not included in `ZLibHeader`. In this case (if you wanted to access these functions from GemStone), you could create a separate instance of `CHeader` for `unistd.h`:

```

topaz 1> printit
UserGlobals at: #'UnistdLibHeader'
  put: (CHeader path: '/usr/include/unistd.h').
%

```

**Simple function call – `getcwd()`**

To take an example that is in `unistd.c`, viewing the source for the `getcwd()` function declaration will let us see the argument declarations.

The CHeader allows you to view the method source, and a version with resolved types:

```
topaz 1> printit
(ZLibHeader functions at: 'getcwd') source
%
extern char *getcwd (char *__buf, size_t __size) __attribute__
    ((__nothrow__ , __leaf__));

topaz 1> printit
(ZLibHeader functions at: 'getcwd') printString
%
extern uint8 *getcwd(uint8 *__buf, typedef uint64 size_t
    __size)
```

This tells us that the function takes two arguments, a pointer to a string and an integer, and returns a pointer to a string. Knowing that the function defined by this header is in `libc`, and the library path and filename is `/lib/x86_64-linux-gnu/libc.so.6`, we can manually create a call to this function:

#### Example 19.4 CCallout to invoke `getcwd()`

```
topaz 1> printit
| string ccallout_getcwd result |
string := String new: 200.
ccallout_getcwd := CCallout library:
    (CLibrary named: '/lib/x86_64-linux-gnu/libc.so.6')
    name: 'getcwd'
    result: #'char*'
    args: #('char*' #'uint64').
result := ccallout_getcwd callWith:
    (Array with: string with: string size).
%
```

It's important to note the way arguments are defined, since C handles memory differently from Smalltalk. The temporary string that is created as an argument to the function must be created with a size larger than the expected result. This is required for heap space to be allocated for the C function; if it is not large enough, the function will error. Also keep in mind that it's very important that the specified size of the string in the second argument **not** be larger than the actual size of the string. The C function will write results to memory limited by the second argument.

`getcwd()` updates the argument as well as returning a value; both contain the same string, but different instances. In both cases, after execution the String's size is the actual size; the String is truncated from the original size of 200.

#### More complex function call – `compress()`

A more complex example is the ZLib function `compress()`. This is defined in `zlib.h` as follows:

```
ZEXTERN int ZEXPORT compress OF((Bytef *dest, uLongf *destLen,
    const Bytef *source, uLong sourceLen));
```

You can view a simplified and clarified definition using CHeader `printString`:

```
topaz 1> printit
(ZLibHeader functions at: 'compress') printString
%
extern int32 compress(uint8 *dest, uint64 *destLen, const uint8
*source, uint64 sourceLen)
```

This tells us that `compress()` takes four arguments:

- ▶ a pointer to a destination buffer
- ▶ a pointer to the length of the destination buffer
- ▶ a pointer to the source data
- ▶ the length of the source data

The function compresses the source data and places the result in the destination buffer. The destination length is updated with the space actually used. The function returns a flag indicating success or the type of error experienced.

We can manually create a call to this function using the core classes described in 16.1:

```
CCallout
  library: (CLibrary
    named: '/lib/x86_64-linux-gnu/libz.so.1')
  name: 'compress'
  result: #'int32'
  args: #'ptr' #'ptr' #'const char*' #'uint64'.
```

This creates an object that can be used to call the `compress()` function in the library. The constructor takes four arguments: (1) an instance of `CLibrary`; (2) the name of the function; (3) the result type; and (4) a list of the types of the arguments.

In order to call the function from Smalltalk we need to create the arguments. The source string and the source length are easy; they are just instances of a Smalltalk String and Integer. The destination and destination length are a bit more complex. They are both pointers to memory locations where the function will retrieve information (`destLen` starts as the available length of the destination buffer) as well as return information (`dest`, where the result is placed, and `destLen`, the amount of `dest` actually used).

In general, C libraries cannot deal directly with Smalltalk objects since the format is different and objects can move in memory with various garbage collection operations. As part of making the C function call, the virtual machine converts the Smalltalk objects to C data and constructs a C stack before making the C library call. For many objects this works fine; as we saw in the `getcwd()` example above, simple String and Integer objects are handled properly. But when an argument is a pointer to a chunk of memory in which the C library will place arbitrary data, we need to explicitly allocate that space and pass a pointer to it.

The class `CByteArray` represents a chunk of memory that is outside the Smalltalk object space (it is on the "heap"), and when an instance of `CByteArray` is passed as a `#'ptr'` type, the virtual machine puts a pointer to the space on the stack before making the function call. There are methods in `CByteArray` to place various Smalltalk objects in the allocated memory and to retrieve Smalltalk objects from the memory.

To allocate memory for the destination buffer, we can do the following:

```
destination := CByteArray gcMalloc: 100.
```

The `gcMalloc` constructor says to create space on the heap (outside of Smalltalk's object memory) and create a Smalltalk object (in object memory) that references the external memory. The heap memory will be automatically freed when the Smalltalk object is garbage collected. We don't need to put anything into the memory since the `compress()` function will not retrieve anything from the buffer. We pick a size that is enough to hold the expected result (we made an educated guess for this example; in real use we could get a better estimate by calling `compressBound()` with the source length).

To allocate memory for the destination size, and put a value in the location, we can do the following:

```
dest_size := CByteArray gcMalloc: 8.
dest_size uint64At: 0 put: destination size.
```

This allocates 8 bytes in the heap and puts the integer 100 (or whatever size we have allocated for the destination buffer) in that memory location (starting at a zero-based offset of 0). When we call the function we will pass a pointer to the number, not the number itself. This is so we provide a place for the function to tell us the amount of the destination buffer actually used (reusing the memory we allocated). After we make the call we can get the size back from the memory location:

```
used := dest_size uint64At: 0.
```

Once we know the amount of the destination actually used, we can extract the zip data. Note that the zip data is generic binary data, not a string, and may include bytes with a value of 0 (so cannot be treated as a C-string). Note that we are again dealing with zero-based offsets since our underlying structures are C memory:

```
compressed := destination byteArrayFrom: 0 numBytes: used.
```

We can put this all together and pass a source string to be compressed:

### Example 19.5 CCallout to invoke `compress()`

```
topaz 1> printit
| ccallout_compress source destination dest_size result used
compressed |
ccallout_compress := CCallout
  library: (CLibrary
    named: '/lib/x86_64-linux-gnu/libz.so.1')
  name: 'compress'
  result: #'int32'
  args: #(#'ptr' #'ptr' #'const char*' #'uint64').
source := 'The quick brown fox jumped over the lazy dog'.
destination := CByteArray gcMalloc: 100.
dest_size := CByteArray gcMalloc: 8.
dest_size uint64At: 0 put: destination size.
result := ccallout_compress callWith:
  { destination . dest_size . source . source size }.
result == 0 ifFalse:[ Error signal:'compress failed' ].
used := dest_size uint64At: 0.
compressed := destination byteArrayFrom: 0 numBytes: used.
%
```

If the result is zero (`Z_OK`), then the function executed successfully, and `compressed` will reference a `ByteArray` that contains the compressed data.

## Using CHeader wrapper methods to create a class

The `CHeader` object can be used to create a new Smalltalk class and automatically generate methods to invoke the C functions.

The method `CHeader >> wrapperForLibraryAt:` can be used to create a Smalltalk class with default name and methods for each function. The default name is the library name without the 'lib', so for `zlib.h`, the resulting class name is simply `"Z"`.

When creating Smalltalk methods that allow arguments to be passed to the C function in the generated interface methods, each function argument is represented with `"_."`. So for example for the `getcwd()` function, which has two arguments, the equivalent Smalltalk method is:

```
getcwd_: buffer_: size
```

To generate a wrapper class for the `zlib` library, in the most simple case you could use the following code:

### Example 19.6 Create wrapper class using default

```
topaz 1> printit
| header wrapperClass |
header := CHeader path: '/usr/include/zlib.h'.
wrapperClass := header wrapperForLibraryAt:
    '/lib/x86_64-linux-gnu/libz.so.1'.
wrapperClass initializeFunctions.
UserGlobals at: wrapperClass name put: wrapperClass.
%
```

After this is executed, you can use a code browser to view the class-side methods that create the `CCallout` instances, and the instance-side methods that call the functions.

As mentioned earlier, the header file may include many functions beyond that provided in the library – all the functions that are defined in the referenced include files. And we can call any of these functions through this library, due to the way the C function lookup occurs.

For example, the function `getpid()` is defined to take no arguments and return a 32-bit number. This makes it very easy to call once we have defined a wrapper class:

### Example 19.7 Invoke Z function getpid

```
topaz 1> printit
Z new getpid
%
22753
```

We probably don't want to allow the `Z` class to have access to every function that is included – for example, it might be better not to have access to `sethostid()`, which changes the current machine's Internet number. It's better to be more selective about what

functions to include in the wrapper. It's also desirable to have a more descriptive name for the library wrapper class.

The method `CHeader>> wrapperNamed:forLibraryAt:select:` allows you to specify the name and a select block to determine the specific functions to include. The select block should evaluate to a Boolean that indicates whether or not to include the particular function.

For example, to create a wrapper for various compress functions, you could do the following:

---

### Example 19.8 Create wrapper class specifying name and functions

---

```
topaz 1> printit
| header class |
UserGlobals removeKey: #'ZLib' ifAbsent: [].
header := CHeader path: '/usr/include/zlib.h'.
class := header
    wrapperNamed: 'ZLib'
    forLibraryAt: '/lib/x86_64-linux-gnu/libz.so.1'
    select: [:each |
        each name includesString: 'compress'].
class initializeFunctions.
UserGlobals at: class name put: class.
%
```

---

This code creates a wrapper class, `ZLib`, that contains only four functions: `compress()`, `uncompress()`, `compress2()`, and `compressBound()`, all the ones that happen to include the string "compress". The select block may be considerably more complex, depending on which specific functions you want to include.

To invoke `compress` using the `Zlib` class rather than manually creating a `CCallout`:

---

### Example 19.9 Invoke Zlib function compress()

---

```
topaz 1> printit
| source destination dest_size result used compressed |
source := 'The quick brown fox jumped over the lazy dog'.
destination := CByteArray gcMalloc: 100.
dest_size := CByteArray gcMalloc: 8.
dest_size uint64At: 0 put: destination size.
result := ZLib new
    compress_: destination
    _: dest_size
    _: source
    _: source size.
used := dest_size int64At: 0.
compressed := destination byteArrayFrom: 0 to: used - 1.
compressed
%
```

---

## 19.4 Example using CCalloutStructs and variable arguments with RabbitMQ®

When interfacing to a library that includes functions that have arguments or return values that are structs, you will use CCalloutStructs. There are some differences from using CCallout, especially when the function returns a struct.

The following example calls a function in the open-source RabbitMQ messaging library, `amqp_login`, which includes variable arguments and returns a struct.

Note that for simplicity of illustration, the examples here use `(CPointer newNull)` for the argument that requires an established connection from an earlier `amqp` call. The methods as described can be executed, but obviously not function correctly.

### Example 19.10 Source for the `amqp_login` function

---

```
((CHeader path: '/usr/include/amqp.h')
  functions at: 'amqp_login') source
%
__attribute__((visibility("default")))amqp_rpc_reply_t
amqp_login(amqp_connection_state_t state, char const *vhost,
  int channel_max, int frame_max, int heartbeat,
  amqp_sasl_method_enum sasl_method, ...); amqp_sasl_method_enum
  sasl_method, ...);
```

---

As you can determine, this function has six regular arguments, includes variable arguments, and returns a struct of type `amqp_rpc_reply_t`.

To generate a data type class for the return struct `amqp_rpc_reply_t`, and save that class in `UserGlobals`:

### Example 19.11 Create wrapper: `CByteArray` class for data type

---

```
"create wrapper for the struct"
| wrapper |
wrapper := (CHeader path: '/usr/include/amqp.h')
  wrapperForTypeNamed: 'amqp_rpc_reply_t'.
UserGlobals at: wrapper name put: wrapper.
```

---

It is strongly recommended that the CCalloutStructs instances also be created using CHeader wrapper: methods, to avoid error-prone hand construction. The generated instances contain the same code as a hand-construct call, and you can use the information in the generated methods to build your own class and methods.

## Hand creation of CCalloutStructs

The following example shows the creation of a hand-constructed CCalloutStructs, and how it is invoked.

In this example, note that the argument arrays for the `args: keyword`, `structSizes:`, and `callWith:` have different sizes, as they include or omit the result and the variable arguments):

- ▶ the first element in the `structSizes:` argument is for the return value; this must be the size of the return value. In this case, since the function returns an instance of `amqp_rpc_reply_t` (a subclass of `CByteArray`), that is the required size.
- ▶ The function takes two variable arguments; strings that represent the userid and password. Variable arguments require passing in both type and value in the array of arguments. In the example, the arguments and variable arguments are composed separately and appended to make the call.

### Example 19.12 Hand-constructed CCalloutStructs

```
topaz 1 > run
| theResult myCO fArgs varArgs |
theResult := amqp_rpc_reply_t new.

"create the CCalloutStructs"
myCO := CCalloutStructs
  library: (CLibrary named: '$GEMSTONE/testlib/testmq')
  name: 'amqp_login'
  result: #struct
  args: #( #'ptr' #'const char*' #'int32' #'int32' #'int32'
           #'int32' )
  varArgsAfter: 6
  structSizes: { (theResult size) . nil . nil . nil . nil .
                 nil . nil }.
  fArgs := { (CPointer newNull) . '/' . 0 . 131072 . 0 . 0 }.
  varArgs := #( #'const char*' 'guest' #'const char*' 'guest' ).

"invoke the CCalloutStructs"
myCO
  callWith: (fArgs, varArgs)
  structResult: theResult
  errno: nil.
theResult
%
```

## Using CHeader wrapper methods to create CCalloutStructs

Defining the CCalloutStructs correctly, including the structSizes and varArgs, is challenging, so it is strongly recommended to use the CHeader utilities to create these methods.

---

### Example 19.13 Creating and using the Smalltalk class with the amqp\_login function

---

```
"create the Wrapper class MyAmq and the CCalloutStructs"
topaz 1> run
| class wrapper |
wrapper := (CHeader path: '/usr/include/amqp.h')
  wrapperNamed: 'MyAmq'
  forLibraryAt: '$GEMSTONE/testlib/testmq'
  select: [:each | each name = 'amqp_login'].
wrapper initializeFunctions.
UserGlobals at: wrapper name put: wrapper.
%

"call the Method on the wrapper class"
topaz 1> run
| theResult |
theResult := MyAmq new
  amqp_login_: (amqp_rpc_reply_t new)
  _: (CPointer newNull)
  _: '/'
  _: 0
  _: 131072
  _: 0
  _: 0
  varArgs: #(#'const char*' 'guest' #'const char*' 'guest').
theResult
%
```

---

## Further refinement: creating a convenience method

The generated method does not have meaningful method names, and includes arguments for both important arguments, and arguments that may be static for your application. You can create a further method for ease of use, as in the following example.

### Example 19.14 Convenience method

---

```
topaz 1> method MyAmq
amqpLoginTo: conn userId: userId password: pass
| result varArgsArray |
result := amqp_rpc_reply_t new.
varArgsArray := Array
    with: #'const char*' with: userId
    with: #'const char*' with: pass.
self
    amqp_login_: result
    _: conn _: '/'
    _: 0 _: 131072 _: 0 _: 0
    varArgs: varArgsArray.
^result
%
topaz 1> commit

" execute the convenience method "
topaz > run
MyAmq new
    amqpLoginTo: (CPointer newNull)
    userId: 'guest'
    password: 'guest'
%
```

---



# External Sessions

---

GemStone/S 64 Bit incorporates a number of classes that facilitate spawning and managing external sessions. This chapter describes these classes and how to use them

**External Sessions** (page 372)

How to create and use external sessions

**NRS and Login Parameter Support** (page 377)

describes how to setup the login parameters and define NRS strings

**Special Cases of External Sessions** (page 379)

Special cases of external sessions, such as solo sessions.

## 20.1 Overview

External sessions allow you to execute Smalltalk code in separate Gems, which may run on different servers and log in as different users to different repositories, including repositories running different versions of GemStone. This allows you to do things such as partitioning work among multiple gems or managing separate repositories.

GemStone includes:

- ▶ **GsTsExternalSession**, using `GciTsLibrary` based on the thread-safe library `libgcits*` described by `gcits.hf` and `GciTsLibrary`. This is the recommended external session class. Logging into repositories running a different version of GemStone is only supported with `GsTsExternalSession`.
- ▶ **GsExternalSession**, using `GciLibrary` based on `libgci*` described by `gci.hf`.
- ▶ **GsLegacyExternalSession**, which provides a limited subset of external session behavior using the `GciInterface` class. This is needed for environment that do not support the FFI, such as AIX.

While the API for these classes is similar, there are some differences, which are mentioned in the text.

## 20.2 External Sessions

To use an external session, you must create an instance of the external session class, `GsTsExternalSession`, `GsExternalSession`, or a subclass; with the appropriate login parameters. Sending login to this instance creates a new Gem session that is logged into the Stone specified by the login parameters. This may be the same stone as the originating Gem or a different Stone.

You can then send messages to the instance of external session, which will be executed on the remote Gem.

After you have executed operations on the remote Gem, you must logout, to ensure the remote Gem is terminated and does not continue to use resources.

You cannot persist instances of kinds of `AbstractExternalSession` in the repository.

### Setup the External Session

The login parameters you configure are the same as when logging in via topaz or other interfaces: the Stone's NRS, the Gem's NRS, and the `userId` and password that the external session will login as. If your login requires host username and host password, these are also required and are provided as part of the NRS arguments.

The Stone and Gem NRS may be provided as strings using GemStone's standard NRS syntax, described in the *System Administration Guide* appendix C. For convenience, you may also use instances of the class `GsNetworkResourceString`, which makes it easier to compose complicated NRS Strings. See section 20.3 on page 377 for classes that make composing NRS easier.

### Creating the External Session

To create the external session, create an instance and specify instances of `GsNetworkResourceString` or NRS strings. The following examples show equivalents using `GsNetworkResourceString` and NRS strings.

#### Default -- logging in to the same stone as the originating gem

`GsTsExternalSession class >> newDefault` creates a new instance of `GsTsExternalSession`, based on the login parameters of the originating Gem (with the exception of the password, which defaults to 'swordfish'). You then only have to apply any parameters that are different.

For example:

```
GsTsExternalSession newDefault
  username: 'Newton';
  password: 'gravity';
  yourself
```

## Using NRS strings

GsTsExternalSession class >> gemNRS:stoneNRS:username:password: allows you to specify the require login parameters. These can be strings or instances of GsNetworkResourceString.

```
GsTsExternalSession
  gemNRS: '!'@santiam#netldi:gs64ldi!gemnetobject'
  stoneNRS: '!'@santiam!gs64stone'
  username: 'Newton'
  password: 'gravity'
```

## With GsNetworkResourceString

```
GsTsExternalSession
  gemNRS: (GsNetworkResourceString
    gemNRSForNetLDI: 'gs64ldi' onHost: 'santiam')
  stoneNRS: (GsNetworkResourceString
    stoneNRSForStoneName: 'gs64stone' onHost: 'santiam')
  username: 'Newton'
  password: 'gravity'
```

## When logging into a different version of GemStone

Logging into a Stone that is running a different version of GemStone is not supported with GsExternalSession, only with GsTsExternalSession, and is only supported with a limited number of recent GemStone versions. Logging into a different version of GemStone also requires that you have gcc/g++ installed on the originating node.

You must use a method that specifies the GciTsLibrary version to be used, for example:

```
myExternalSession := GsTsExternalSession
  parameters: nil
  library: (GciTsLibrary
    newForVersion: '3.6.6'
    product: '/lark/users/gsadmin/GemStone64_366').
myExternalSession
  stoneNRS: '!'#netldi:ldi_366!stone366';
  username: 'DataCurator';
  password: 'swordfish';
  gemNRS: '!'#netldi:ldi_366!gemnetobject'.
```

## Avoiding passwords in code

Using a one-time password allows you to login a second session without including the password in code; a temporary password is granted based on the current logged-in user's authentication. This feature is described in detail in the *System Administration Guide*.

One-time password requires that the logged-in user have the OneTimePassword privilege. It is possible to login as a different user than the currently-logged in user; however, this requires that the other userId be on an allowlist.

The temporary password is created with a time limit, and then immediately passed to the login code. The one-time password is only usable for a single login.

To use this password to login an external session, for example:

```
tempPW := GsCurrentSession currentSession
        createOnetimePasswordValidForSeconds: 30.
myExternalSession := (GsTsExternalSession newDefault)
        onetimePassword: tempPW;
        yourself.
```

## Log in the External Session

To login, send #login to the configured external session:

```
myExternalSession login
```

Login creates a remote Gem session that is logged in and in transaction in the specified Stone, either the same Stone as the originating Gem session, or a different Stone. If the remote Gem is logged into a Stone that is in active use, you must manage the Gem appropriately to avoid creating a commit record backlog in that Stone; avoid leaving remote Gems logged in and idle, and ensure that the code you execute commits or aborts regularly.

To logout, send #logout to the logged-in external session:

```
myExternalSession logout.
```

## Executing Code

Code to be executed by the remote Gem can be passed as strings or blocks. These can be executed synchronously or asynchronously.

### Code in Strings

To synchronously execute code contained in a string, use the method `executeString:.`

For example:

```
myExternalSession executeString:
'SystemRepository fullBackupTo: '
'/backups/gs/bkup23-08-08.dat'''.
```

### Code in Blocks

What is actually sent to the remote Gem is always in the form of a String, but methods are provided that accept blocks containing the code to execute in the remote Gem. The source strings for these block will be passed to the remote Gem. This allows Smalltalk tools to manage the source, detect senders, and so on, which is not possible with strings.

To use blocks, the blocks must be able to compile in both the originating Gem and the remote Gem in which you intend them to execute, although the block's code is not necessarily meaningful in the originating session. Any variable resolution, etc. in the blocks will be resolved again in the environment of the remote Gem when the block is compiled after being transmitted as a string, and if the variables cannot be resolved in the remote Gem, it will result in an error.

Code in block can also be executed synchronously or asynchronously.

To synchronously execute code contained in a block, use:

```
executeBlock: aNoArgBlock
executeBlock: aOneArgBlock with: aValue
executeBlock: aTwoArgBlock with: aValue with: anotherValue
executeBlock: aBlock withArguments: aCollectionOfValues
```

These methods execute the source code contained in the given block, and return the result of executing that code.

When passing arguments to the block, the arguments values must be objects for which the `printString` allows the correct object state to be recreated in the remote session. This is true for all objects, including specials, strings, integers and floats; use caution to avoid unexpected conversion or loss of information, as well as errors.

## Return Values

After code is executed in the remote Gem, the result is returned to the originating session.

If the result of the expression is a special (Character, Boolean, SmallInteger, SmallFloat, etc.), or a String, Symbol, or ByteArray, the results are converted into the appropriate object in the originating Gem.

Expressions that return another type of object will return an Array containing the OOP of the result, the OOP of the class of the result, and the size of the result (for `GsExternalSession`, only the OOP of the result is returned). This should be avoided, except when performing additional remote operations on returned OOPs. The returned OOP is for the value of the result in the remote Gem, which may not exist or be resolvable in the originating Gem; and OOP lookup has an inherent risk of unexpected results.

When the result is not a special, then the OOP of the result is placed in the `ExportSet` of the remote Gem. See "Important caution on Export Set of remote session" on page 376.

Since the evaluation is done in a separate Gem process, any transient changes in the remote Gem are not visible in the originating Gem. In order for persistent changes in the remote Gem to be visible to the originating Gem, the remote Gem must commit the changes, and the originating Gem must abort.

## Asynchronous Execution

The `executeString:` and `executeBlock:` methods block the originating Gem until execution completes. To execute the remote code asynchronously and return control immediately to the originating Gem, the following equivalent methods are available:

```
forkString: aString
forkBlock: aNoArgBlock
forkBlock: aOneArgBlock with: aValue
forkBlock: aTwoArgBlock with: aValue with: anotherValue
```

When you execute asynchronously, an external call is in progress, and the methods you can invoke on the remote session are limited:

```
isResultAvailable
    Check whether the current call in progress has finished and save the result if it has.

lastResult
    Answer the result received when the last isResultAvailable answered true, which includes after a waitForResult operation completed.

waitForResult
    Wait for the external Gem to complete the current operation.
```

`waitForResultForSeconds: numSeconds`

Wait up to *numSeconds* seconds for the external Gem to complete the current operation.

`waitForResultForSeconds: numSeconds otherwise: aBlock`

Wait up to *numSeconds* seconds for the external Gem to complete the current operation. If the operation does not complete within that time, answer the result of evaluating *aBlock*.

### Operations on remote objects

If you perform a remote operation that returns an OOP, you can send specific selectors to that remote object by OOP.

`send: selector to: anOop`

Send the given selector to the object in the external session with the OOP *anOop*.

`send: selector to: anOop withArguments: anArrayOfValues`

Send the given selector to the object represented by the given OOP, which is an OOP in the external session, and pass the Array of arguments.

The OOPs of the arguments are passed to the remote Gem. These arguments must be specials, or persistent objects that exist on both the calling and remote sessions, otherwise it will result in an error.

### Important caution on Export Set of remote session

For objects other than specials (Integers, Characters, etc.) that are returned by the remote Gem, the remote Gem adds these objects to its export set. This includes Strings and other byte collections, Exceptions returned by the remote Gem, and other objects that are returned as OOPs.

Strings and similar byte-format results and exceptions are converted into new String (or appropriate) instances in the calling Gem (with a new OOP); the OOP of the original String on the remote Gem remains in the external Gem's export set, for logins using `GsExternalSession`. When using `GsTsExternalSession`, in many cases the OOPs are released from the export set automatically.

Otherwise, these OOPs remain in the export set of the remote Gem, and will not be garbage collected, until the remote Gem is logged out.

These OOPS can be removed manually from the export set using `GsTsExternalSession >> releaseOop: or releaseOops:.` The export set can be examined using `(GsBitmap newForHiddenSet: #PureExportSet)` on the remote Gem.

### Exceptions

The class `GciError` and `GciLegacyError` are provided to represent errors during remote execution. If the code being executed on the remote session encounters an exception, this is raised as a `GciError` in the calling session, or a `GciLegacyError` if using `GsLegacyExternalSession`.

Since remote debugging is not possible with this interface, the stack of the error is included with the error description.

For example, given the following code which triggers an error on the remote session:

```
result := [myExternalSession executeString: '1/0']
  on: GciError
  do: [:ex | ex description].
```

The result in the calling session will be:

```
GciError: a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by zero
1 AbstractException >> _signalWith: @5 line 25
  receiver size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
  inCextensionArg nil
  res nil
  .t1 size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
2 AbstractException >> signal @2 line 47
  receiver size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
  .t1 size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
  .t2 nil
3 Number >> _errorDivideByZero @6 line 7
  receiver 1
  .t1 size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
  .t2 size:0 a ZeroDivide occurred (error 2026),
reason:numErrIntDivisionByZero, attempt to divide 1 by z
4 SmallInteger >> / @6 line 7
  receiver 1
  aNumber 0
  .t1 1
5 Executed Code @2 line 1
  receiver nil
  .t1 1
  .t2 0
6 GsNMethod class >> _gsReturnToC @1 line 11
  receiver oop:144897 GsNMethod
```

## 20.3 NRS and Login Parameter Support

GemStone logins require a number of parameters. The Stone name, and the Gem service, are specified in the form of a Network resource string, or NRS. NRS includes a number of features; the NRS syntax is documented in the *System Administration Guide*, Appendix C.

While you may compose strings in standard NRS syntax for your external session logins, if you are unfamiliar with NRS syntax, the `GsNetworkResourceString` class provides a way to compose sophisticated NRS strings by specifying the individual building blocks.

The NRS Strings to specify a Stone are different than the strings used to specify a Gem service. The class `GsNetworkResourceString` supports variables for both sets of arguments; you will need to pass in the correct arguments for a Stone or Gem service.

## For a Stone

- ▶ `body`: the name of the Stone.
- ▶ `node`: not needed for a remote Stone running on localhost. Otherwise, the name or IP address of the remote Stone's node.
- ▶ `netldi`: not needed if the NetLDI serving the remote Stone is running as `gs64ldi`. Otherwise the name or port of the NetLDI serving the remote Stone.

For example:

```
(GsNetworkResourceString new
  node: 'santiam';
  netldi: '51234';
  body: 'gs64stone';
  yourself) printString
'!@santiam#netldi:51234!gs64stone'
```

## For a Gem

- ▶ `body`: the name of the gem service (e.g. 'gemnetobject' or 'gemnetdebug'). Gem service arguments such as `gemnetobject -C` can be included here.
- ▶ `node`: not needed if the Gem will run on localhost. Otherwise, the name or IP address of the node that the remote Gem will run on.
- ▶ `netldi`: not needed if the NetLDI serving the remote Stone is running as `gs64ldi`. Otherwise the name or port of the NetLDI serving the remote Stone.
- ▶ `authorization`: if the remote Netldi is not running in guest mode, the host `userId` and host password of the account that will own the remote Gem.
- ▶ `log`: optional, specifies a log file for the remote Gem.
- ▶ `dir`: optional, specifies a working directory for the remote Gem.

For example:

```
(GsNetworkResourceString new
  node: 'santiam';
  netldi: '51234';
  authorization: 'user@passwd';
  body: 'gemnetobject';
  yourself) printString
'!@santiam#auth:user@passwd#netldi:512345!gemnetobject'
```

## Convenience methods for common arguments

There are a number of class methods that allow you to pass in specific required elements of an NRS. In addition to these specific listed methods, see the image for other variants.

`gemNRS`

Creates a Gem NRS that uses the default Gem script, which will expect to find `gs64ldi` on localhost; the equivalent of 'gemnetobject'

`gemNRSForNetLDI: nameOrPort onHost: gemhostname`

Creates a Gem NRS to run the Gem process on the host named `gemhostname` and

with the NetLDI of *nameOrPort* with the standard gem script *gemnetobject*; the equivalent of `'!@gemhostname#netldi:nameOrPort!gemnetobject'`

`stoneNRS`

Creates a Stone NRS for a Stone named *gs64stone*, which should be running on localhost.

`stoneNRSForStoneName: aStoneName onHost: stoneHostName`

Creates a Stone NRS for *aStoneName*, which should be running on *stoneHostName*; the equivalent of `'!@stoneHostName!aStoneName'`.

## 20.4 Special Cases of External Sessions

### Solo external sessions

Using a "solo" session, you can perform most GemStone operations without accessing a Stone. An external session can be logged in solo either on the same host or another host and does not have a login to any stones, regardless of a Stone name pass into the NRS for gemstone.

By default, solo sessions use the read-only empty distribution extent; this can be configured to use another extent file, provided the following are true:

- ▶ The extent must not part of a multi-extent repository
- ▶ The extent file must either have read-only file permissions, or it will be exclusive-locked by the solo session.
- ▶ The extent repository must have been previously cleanly shutdown by the Stone

To login solo, the only thing necessary is to send `loginSolo` rather than `login`. Note that the login parameters include the stone name, but this is ignored.

The configuration parameter `GEM_SOLO_EXTENT` specifies the extent file to be used by a Solo session. This defaults to the clean, read-only extent within the distribution, `$GEMSTONE/bin/extent0.dbf`. You can pass in another extent file by specifying the `-C` argument to `gemnetobject`. For example:

#### Example 20.1 Passing a custom extent to a solo login

```
| myExternalSession |
myExternalSession := GsTsExternalSession
  gemNRS: 'gemnetobject -C GEM_SOLO_EXTENT=$GEMSTONE/archive.dbf'
  stoneNRS: nil
  username: 'DataCurator'
  password: 'swordfish'.
myExternalSession loginSolo.
```

## Primitive-based external sessions on AIX

GsExternalSession uses the FFI GciLibrary to implement its functionality. On platforms on which the FFI GCI library cannot be used, instances of GsLegacyExternalSession provides external sessions. AIX is the only known environment where this class must be used.

## X509 external sessions

The classes GsTsX509ExternalSession and GsX509ExternalSession, along with GemStoneX509Parameters, allow you to create an external session login using X509-Secured GemStone. See the *GemStone/S 64 Bit X509-Secured GemStone System Administration Guide* for details on how to set this up.

# Testing and SUnit

---

SUnit is a framework that supports the creation of automated unit and application tests. This chapter discusses the role of repeatable unit tests and how to easily create these tests using SUnit.

This chapter was originally based on “SUnit Explained” by Stéphane Ducasse (<http://www.iam.unibe.ch/~ducasse/Programmez/OnTheWeb/Eng-Art8-SUnit-V1.pdf>) and is used by permission.

**Application Testing** (page 381)

describes the general goals of automated testing, and the benefits of the SUnit framework.

**SUnit example: ExampleSetTest** (page 384)

presents a step-by-step example that illustrates the use of SUnit.

**The SUnit Framework and Implementation** (page 387)

describes the core classes of the SUnit framework.

**Running a Single Test** (page 388)

explores key aspects of the implementation by following the execution of a test and test suite.

## 21.1 Application Testing

### The value of testing

Many traditional development methodologies include testing as a step that follows coding, and this step is often cut short when time pressures arise. Development of automated tests in parallel with application development can save time in ensuring that each part of the application that is written is working as intended, that new development does not break existing functionality, and allow making changes in the application with much higher confidence.

The Smalltalk community has a long tradition of testing, due to the incremental development supported by its programming environment. Testing is commonly done during development: a typical practice is to compile a method and then, from a workspace, write a small expression to test it, perhaps even keeping the expression in method comments. However, these examples can be difficult to understand after some time passes, and there is no easy way to keep track of them and to automatically run them. Unfortunately, tests that you cannot automatically run are less likely to be run. Moreover, having a code snippet to run in isolation often does not readily indicate the expected result. Automated tests play several roles:

- ▶ Automated tests are an active and *always synchronized* documentation of the functionality they cover. Testing can demonstrate how the application classes and methods are meant to be used together, and are a valuable training resource.
- ▶ By writing tests at the same time or even before writing code, it forces you to think about the functionality you want to design. You are required to clearly state the context in which your functionality will run, the way it will interact with other code, and, more important, the expected results. Moreover, when you are writing tests, you are your first client and your code will naturally improve.
- ▶ The application lifecycle includes initial development, and also further development, bug fixes, and new features. With automated testing, you can have confidence that the changes you introduce are not breaking existing functionality, or easily locate and address any issues that arise.

By using “test-first” or “test-driven” development, eXtreme Programming proposes to write tests even before writing code. While this is counter-intuitive to the traditional “design-code-test” mindset, it can have a powerful impact on the overall result. Test-driven development can improve the design by helping you to discover the needed interface for a class and by clarifying when you are done (the tests pass!).

## Writing good tests

It is clear that we cannot test all the aspects of an application. Covering a complete application is simply impossible and should not be the goal of testing. Even with a good test suite, defects can creep into the application and be left hidden waiting for an opportunity to damage your system. While there are a variety of test practices that can address these issues, the goal of *regression* tests is to ensure that a previously discovered and fixed defect is not reintroduced into a later release of the product.

Writing good tests is a technique that can be easily learned by practice. Let us look at the properties that tests should have to get a maximum benefit:

- ▶ Repeatable. We should be able to easily repeat a test and get the same result each time.
- ▶ Automated. Tests should be run without human intervention. You should be able to run them during the night.
- ▶ Tell a story. A test should cover one aspect of a piece of code. A test should act as a specification for a unit of code.
- ▶ Resilient. You should not have to change many tests for every application change. One way to achieve this property is to write tests based on the interfaces of the tested functionality.

In addition, for test suites, the number of tests should be somehow proportional to the bulk of the tested functionality. For example, changing one aspect of the system might break *some* tests, but it should not break *all* the tests. This is important because having 100 tests broken should be a much more important message for you than having 10 tests failing.

## Why SUnit?

SUnit is a mature public-domain framework, widely available for most dialects of Smalltalk. SUnit is easy to use for basic testing, and also includes support for more sophisticated testing requirements.

The value of SUnit is that it provides a code framework to describe the context of your tests as well as to run them automatically. With a set of SUnit tests, you can quickly add tests that become part of an automated test suite. This represents a vast improvement over writing small code snippets in an ephemeral workspace.

SUnit was developed originally by Kent Beck and was extended by Joseph Pelrine and others. SUnit is not limited to Smalltalk; SUnit inspired testing frameworks in many other languages, such as JUnit for Java, RUnit for R, and so on. The language-independent framework is called XUnit.

*"There are simpler ways of running automated tests, and many, many more complicated ways, but this architecture seems to hit a sweet spot."*

## Portable SUnit classes and GemStone's SUnit classes

The SUnit framework is available in many Smalltalk dialects. Application tests that are written using base SUnit classes (TestCase, TestResult, TestSuite, etc), allows your tests to be portable to other Smalltalk dialects.

The GemStone distribution also includes a GemStone-specific refinement of SUnit, via subclasses GsTestCase, GsTestResult, and GsTestSuite. These classes provides additional test methods and support for GemStone-specific behavior in testing.

There are some differences in behavior; when using GsTestCase, results are written to log files, including SUnit.log and SUnitDefects.log. Also, GsTestResults do not preserve the actual TestCases, as does TestResult; instead, the printString is stored.

## Terminology: test errors vs. failures

SUnit recognizes two kinds of defects: not getting the correct answer (a **failure**) and not completing the test (an **error**).

Usually, a failure indicates that there is a problem in the code you are testing; the test is working correctly to uncover defects in your code. Of course, it may also be the case that the test itself includes incorrect logic.

An error indicates just that: an unhandled code error occurred, such as a divide by zero, and the test could not continue running. This may indicate a problem in the SUnit test code itself. For example, your SUnit class may send a method to an application class that is not implemented, and a message not understood occurs. You should either remove that message send from the test, or setup your test to expect that particular error.

## 21.2 SUnit example: ExampleSetTest

Before going into the details of SUnit, let's look at a simple example. This example tests the class `Set`. The example test class, `ExampleSetTest`, is included in the GemStone distribution, so that you can read the code directly in the image and execute the tests.

### Define the new testing Class

To start off, look at the definition of the class `ExampleSetTest`, which is a subclass of `TestCase`.

#### Example 21.1

---

```
TestCase subclass: 'ExampleSetTest'  
  instVarNames: #( full empty)  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: #()  
  inDictionary: Globals
```

---

The class `ExampleSetTest` here is intended to contain all tests related to the class `Set`.

It establishes the context of all the tests that we will specify. Here, the context is described by specifying two instance variables, `full` and `empty`, that represent a full and empty set, respectively.

This example is a subclass of `TestCase`; you may also subclass `GsTestCase` for testing that involves GemStone-specific behavior.

### Define setUp and tearDown methods

Testing often requires a context in which to run the tests, which is established by the method `setUp`. This method is automatically invoked before the execution of each test method defined in a subclass of `TestCase`.

Here, we initialize the instance variable `empty` to refer to an empty set, and the `full` instance variable to refer to a set containing two elements.

#### Example 21.2

---

```
ExampleSetTest>>setUp  
  empty := Set new.  
  full := Set with: 5 with: #abc.
```

---

The complementary method `tearDown` allows you to clear or release any resources that were used during the tests, so you can restore state to exactly as it was prior to testing.

It is not required to have a `setUp` or `tearDown` method, if the testing code does not need to take any action. `ExampleSetTest` does not need or have a `tearDown` method.

## Define the actual testing methods

We start off by defining three methods to test Set class. Each method represents one test.

The selectors for methods that represent tests, normally start with 'test'. Methods that match this pattern are automatically collected by SUnit for testing. You may need to implement support methods that are not themselves tests; these methods should not begin with 'test'.

### Example 21.3

---

```
ExampleSetTest>>testIncludes
  self assert: (full includes: 5).
  self assert: (full includes: #abc).

ExampleSetTest>>testOccurrences
  self assert: (empty occurrencesOf: 0) = 0.
  self assert: (full occurrencesOf: 5) = 1.
  full add: 5.
  self assert: (full occurrencesOf: 5) = 1.

ExampleSetTest>>testRemove
  full remove: 5.
  self assert: (full includes: #abc).
  self deny: (full includes: 5).
```

---

Each method starts executing with an initial state defined by the `setUp` method. So the `testIncludes` method tests that after running the `setUp` method in Example 21.2, the full instance variable include both 5 and #abc.

While the `testRemove` method removes an element from full, this has no effect on any other tests that might run later, since the initial state for each test comes from the `setUp` method.

## Execute the tests

Tests can be executed by executing code, in Topaz, GBS, or other tools.

To run a test, for example, execute the following code:

```
ExampleSetTest run: #testRemove.
```

To run all tests on the class, execute

```
ExampleSetTest suite run
```

These methods return an instance of the `TestResult` class, showing the tests that errored, failed, and passed. You must examine the object returned from these methods to determine if the test passed or not.

During test development, it is much more useful to run the test and bring up a walkback if an error or failure occurs. To do this, execute:

```
ExampleSetTest debug: #testRemove.
```

## Assertions and other claims in test methods

SUnit provides a number of methods to support verifying the results from code execution.

### Checking result value with `assert:` and `deny:`

The method `TestCase>>assert:` requires a single argument, a boolean that represents the value of a tested expression. When the argument is true, the expression is considered to be correct, and we know that the test is valid. When the argument is false, then the test failed.

The method `deny:` is the negation of `assert:`. Hence

```
self deny: Set new size > 0.
```

is equal to

```
self assert: (Set new size > 0) not.
```

There are other methods such as `assert:description:` and `assert:equals:`, see the image for details.

The method `should:` takes a block. For example,

```
self should: [Set new size > 0].
```

### Checking for exceptions with `should:raise:`

In addition to testing for correct results, your application may need testing for how it responds to abnormal input or unexpected conditions by raising an error. This is supported by the `should:raise:` and `shouldnt:raise:` methods.

#### Test for general error

Because SUnit runs on a variety of Smalltalk dialects, the SUnit framework factors out the variant parts (such as the name of the exception). This pattern is important if you plan to write tests that are intended to be cross-dialect.

This will catch `Error`, or any subclass of `Error`.

```
self should: [empty at: 5] raise: TestResult error.
```

#### Specific named error

You can also test for specific exception classes. This code may be portable provided that the target Smalltalk dialect implements the specific `Error` class. For example, `MessageNotUnderstood` is a fairly standard error. However, `GemStone` and other Smalltalk dialects will include many dialect-specific error classes that will not be portable.

```
self should: [empty foo] raise: MessageNotUnderstood.
```

#### GemStone-specific legacy error numbers

`GemStone`'s legacy error handling used numbers, rather than classes, and some ANSI error classes map to multiple more specific error numbers. For example, `[empty at: 5]` throws ANSI class `Error`, with legacy number `2007/#rtErrShouldNotImplement`.

A test for `Error` or `TestResult error` will catch `MessageNotUnderstood` as well as `Error`. Using a subclass of `GsTestCase`, rather than of `TestCase`, you may use the legacy error numbers in place of the error class to test for more specific errors.

In a subclass of `GsTestCase`:

```
self should: [empty at: 5] raise: 2007.
```

This will catch `#rtErrShouldNotImplement`, but will not catch `MessageNotUnderstood` (legacy error 2010).

With a subclass of `GsTestCase`, you may use both ANSI exception classes and GemStone legacy error numbers interchangeably. However, testing for legacy error numbers will not be portable to other Smalltalk dialects.

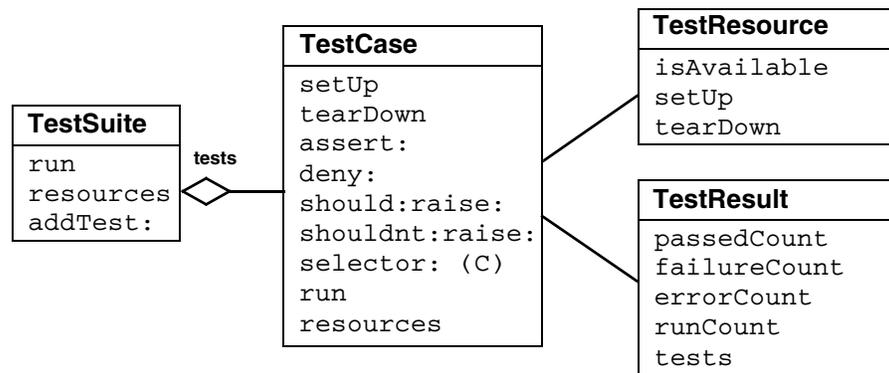
## 21.3 The SUnit Framework and Implementation

While you can use SUnit without understanding the implementation, it may be useful when you need to customize your test suite.

### SUnit Core Classes

SUnit is implemented by four main classes: `TestSuite`, `TestCase`, `TestResult`, and `TestResource`. See Figure 21.1. (Note that this is an object composition diagram, not a class hierarchy diagram.)

Figure 21.1 The SUnit Core Classes



#### TestSuite

The class `TestSuite` represents a collection of tests. A `TestSuite` contains subclasses of `TestCase` and other instances of `TestSuite`. The classes `TestSuite` and `TestCase` form a composite pattern in which `TestSuite` is the composite and `TestCase` is the leaf.

#### TestCase

The class `TestCase` represents a family of tests that share a common context. The context is specified by instance variables on a subclass of `TestCase` and by the specialization method `setUp`, which initializes the context in which each of the tests will be executed. The method `setUp` is invoked before the execution of every test.

The class `TestCase` also defines the method `tearDown`, which is responsible for cleanup, including releasing the objects allocated by `setUp`. The method `tearDown` is invoked after the execution of every test.

## TestResult

The class `TestResult` represents the results of a `TestSuite` execution. This includes a description of which tests passed, which failed, and which had errors.

## TestResource

Recall that the `setUp` method is used to create a context in which the test will run. Often that context is quite inexpensive to establish, as in Example 21.2 seen earlier, which creates two instances of `Set` and adds two objects to one of those instances.

At times, however, the context may be comparatively expensive to establish. In such cases, the prospect of re-establishing the context for each run of each test might discourage frequent running of the tests. To address this problem, SUnit introduces the notion of a *resource* that is shared by multiple tests.

The class `TestResource` represents a resource that is used by one or more tests in a suite, but instead of being set up and torn down for each test, it is established once before the first test and reset once after the last test. By default, an instance of `TestSuite` defines as its resources the list of resources for the `TestCase` instances that compose it.

As shown in Example 21.4, a resource is identified by overriding the class method `resources`. Here, we define a subclass of `TestResource` called `MyTestResource`. We associate it with `MyTestCase` by overriding the class method `resources` to return an array of the test classes to which it is associated.

### Example 21.4

---

```
MyTestCase class>>resources
  "associate a resource with a testcase"
  ^ Array with: MyTestResource.
```

---

As with a `TestCase`, we use the method `setUp` to define the actions that will be run during the setup of the resource.

## Running a Single Test

To execute a single test, we evaluate the expression

```
(TestCase selector: aSymbol) run.
```

The method `TestCase>>run` creates an instance of `TestResult` to contain the result of the executed tests, and then invokes the method `TestCase>>run:`, which in turn invokes the method `TestResult>>runCase:`. See Figure 21.2.

**Figure 21.2 TestCase instance methods run and run: (source code)**


---

```

TestCase>>run
  | result |
  result := TestResult new.
  [self run: result]
    ensure: [TestResource resetResources: self resources].
  ^result.

TestCase>>run: aResult
  aResult runCase: self.

```

---

The runCase: method (Figure 21.5) invokes the method TestCase>>runCase, which executes a test. Without going into the details, TestCase>>runCase pays attention to the possible exception that may be raised during the execution of the test, invokes the execution of a TestCase by calling the method runCase, and counts the errors, failures, and passed tests.

**Example 21.5 TestResult instance method runCase: (source code)**


---

```

TestResult>>runCase: aTestCase
  [aTestCase runCase.
  self addPass: aTestCase]
    on: self class failure , self class error
    do: [:ex | ex sunitAnnounce: aTestCase toResult: self]

```

---

As shown in Figure 21.6, the method TestCase>>runCase calls the methods setUp and tearDown.

**Example 21.6 TestCase instance method runCase (source code)**


---

```

TestCase>>runCase
  self resources do: [:each | each availableFor: self].
  [self setUp.
  self performTest]
    ensure: [ tornDown ifNil:[
      tornDown := true .
      self tearDown ]
    ]

```

---

## Running a TestSuite

To execute more than a single test, we invoke the method TestSuite>>run on a TestSuite (see Figure 21.7). The class TestCase provides the functionality to build a test suite from its methods. The expression MyTestCase suite returns a suite containing all the tests defined in the class MyTestCase.

The method TestSuite>>run creates an instance of TestResult, verifies that all the resource are available, then invokes the method TestSuite>>run: to run all the tests that compose the test suite. All the resources are then reset.

---

**Example 21.7 TestSuite instance methods run and run: (source code)**

---

```
TestSuite>>run
| result |
result := TestResult new.
[self run: result]
ensure: [TestResource resetResources: self resources].
^result

TestSuite>>run: aResult
self tests do: [:each |
self sunitChanged: each.
each run: aResult]
```

---

The class TestResource and its subclasses use the class method current to keep track of their currently created instances (one per class). This instance is cleared when the tests have finished running and the resources are reset. The resources are created as needed. See Figure 21.8.

---

**Example 21.8 TestResource class methods isAvailable and current (source code)**

---

```
TestResource class>>isAvailable
^self current notNil

TestResource class>>current
current ifNil: [current := self new].
^current
```

---

# GemStone Smalltalk Syntax

---

This appendix outlines the syntax for GemStone Smalltalk and introduces the important kinds of GemStone Smalltalk objects.

**GemStone and ANSI Smalltalk** (page 391)

The relationship between GemStone Smalltalk and the ANSI Smalltalk standard.

**GemStone Smalltalk** (page 392)

An overview of the Smalltalk Syntax.

**Blocks** (page 405)

How blocks are used in GemStone Smalltalk.

**GemStone Smalltalk BNF** (page 410)

GemStone Smalltalk BNF

## A.1 GemStone and ANSI Smalltalk

GemStone's programming language, GemStone Smalltalk, is a dialect of the Smalltalk programming language. The Smalltalk language standard is defined by an ANSI Smalltalk standard. While GemStone follows this standard, there are places where either for historical reasons or by choice, GemStone Smalltalk does not follow the ANSI standard.

Some known places in which GemStone Smalltalk does not conform to the ANSI standard:

- ▶ GemStone Array sizes are not fixed; an Array may increase in size by 1 when an operation assign to an index not more than 1 larger than the current size.
- ▶ GemStone Array constructors using {} are not part of the standard.
- ▶ `Integer >> asInteger` truncates rather than rounds.
- ▶ The Fixed point literal syntax with 'p' is not part of the standard.

## GemStone and ANSI limits

The limits in GemStone, such as the maximum length of an identifier, are generally higher than the ANSI minimum requirement. While this makes working in GemStone more flexible, adhering to the ANSI limits, or the limits for other Smalltalk dialects, makes it easier to port GemStone code to other platforms.

**Table 1.1 Implementation Limits**

	<b>GemStone limit</b>	<b>ANSI minimum upper limit</b>
Length of identifiers	1024	200
Length of binary selectors	1024	2
Total length of keyword selectors (including colons)	1024	500
Number of named instance variables per object (including inherited)	2030	127
Number of class variables per class	no limit	1000
Number of methods per behavior	no limit	1000
Number of arguments per method or block	255	15
Number of temporary variables per method or block	1013	15
Total number instance variables (system-wide)	no limit	65535

## A.2 GemStone Smalltalk

Every object is an instance of a class, taking its methods and its form of data storage from its class. Defining a class thus creates a kind of template for a whole family of objects that share the same structure and methods. Instances of a class are alike in form and in behavioral repertoire, but independent of one another in the values of the data they contain.

Classes are much like the data types (string, integer, etc.) provided by conventional languages; the most important difference is that classes define actions as well as storage structures. In other words, Algorithms + Data Structures = Classes.

Smalltalk provides a number of predefined classes that are specialized for storing and transforming different kinds of data. Instances of class Float, for example, store floating-point numbers, and class Float provides methods for doing floating-point arithmetic. Floats respond to messages such as +, -, and reciprocal.

Instances of class `Array` store sequences of objects and respond to messages that read and write array elements at specified indices.

The Smalltalk classes are organized in a treelike hierarchy, with classes providing the most general services nearer the root, and classes providing more specialized functions nearer the leaves of the tree. This organization takes advantage of the fact that a class's structure and methods are automatically conferred on any classes defined as its subclasses. A subclass is said to inherit the properties of its parent and its parent's ancestors.

GemStone Smalltalk is case-sensitive; that is, names such as "SuperClass," "superclass," and "superClass" are treated as unique items by the GemStone Smalltalk compiler.

## How to Create a New Class

Classes are created using a number of class creation methods, defined on the class `Class`. For example, following message expression makes a new subclass of class `Object`, the class at the top of the class hierarchy:

```
Object subclass: 'Animal'  
  instVarNames: #()  
  classVars: #()  
  classInstVars: #()  
  poolDictionaries: {}  
  inDictionary: UserGlobals
```

This subclass creation message establishes a name ('`Animal`') for the new class and installs the new class in a `Dictionary` called `UserGlobals`. The `String` used for the new class's name must follow the general rule for variable names — that is, it must begin with an alphabetic character and its length must not exceed 1024 characters. Installing the class in `UserGlobals` makes it available for use in the future — you need only write the name `Animal` in your code to refer to the new class. For more on class creation, see Chapter 2.

## Statements

The basic syntactic unit of a Smalltalk program is the *statement*. A lone statement needs no delimiters; multiple statements are separated by periods:

```
a := 2.  
b := 3.
```

In a group of statements to be executed en masse, a period after the last statement is optional.

A statement contains one or more *expressions*, combining them to perform some reasonable unit of work, such as an assignment or retrieval of an object.

## Comments

GemStone Smalltalk usually treats a string of characters enclosed in quotation marks as a *comment* — a descriptive remark to be ignored during compilation. Here is an example:

```
"This is a comment."
```

A quotation mark does *not* begin a comment in the following cases:

- ▶ Within another comment. You cannot nest comments.
- ▶ Within a string literal (see page 396). Within a GemStone Smalltalk string literal, a "comment" becomes part of the string.

- ▶ When it immediately follows a dollar sign (\$). GemStone Smalltalk interprets the first character after a dollar sign as a data object called a character literal (see page 395).

A comment terminates tokens such as numbers and variable names. For example, GemStone Smalltalk would interpret the following as two numbers separated by a space (by itself, an invalid expression):

```
2" this comment acts as a token terminator"345
```

## Expressions

An expression is a sequence of characters that Smalltalk can interpret as a reference to an object. Some references are direct, and some are indirect.

Expressions that name objects directly include both variable names and literals such as numbers and strings. The values of those expressions are the objects they name.

An expression that refers to an object indirectly by specifying a message invocation has the value returned by the message's receiver. You can use such an expression anywhere you might use an ordinary literal or a variable name. This expression:

```
2 negated
```

has the value (refers to) -2, the object that 2 returns in response to the message `negated`.

The following sections describe the syntax of GemStone Smalltalk expressions and tell you something about their behavior.

A GemStone Smalltalk expression can contain a combination of the following:

- ▶ a literal
- ▶ a variable name
- ▶ an assignment
- ▶ a message expression
- ▶ an array constructor
- ▶ a path
- ▶ a block

The following sections discuss each of these kinds of expression in turn.

## Literals

A *literal expression* is a representation of some object such as a character or string whose value or structure can be written out explicitly. The kinds of GemStone Smalltalk literals are:

- ▶ numbers
- ▶ characters
- ▶ strings
- ▶ symbols
- ▶ arrays of literals

## Numeric Literals

In Smalltalk, literal numbers look and act much like numbers in other programming languages. Like other Smalltalk objects, numbers receive and respond to messages. Most of those messages are requests for arithmetic operations. In general, Smalltalk numeric expressions do the same things as their counterparts in other programming languages. For example:

```
5 + 5
```

returns the sum of 5 and 5.

A literal floating point number must include at least one digit after the decimal point:

```
5.0
```

You can express very large and very small numbers compactly with scientific notation. To raise a number to some exponent, simply append the letter "e" and a numeric exponent to the number's digits. For example:

```
8.0e2
```

represents 800.0. The number after the e represents an exponent (base 10) to which the number preceding the e is to be raised. The result is always a floating point number. Here are more examples:

```
1e-3 represents 0.001
```

```
1.5e0 represents 1.5
```

The literal numeric type GemStone/S 64 Bit supports are:

- ▶ "e", "E", "d" and "D" for floating point literals (SmallDouble or Float)
- ▶ "f" and "F" for DecimalFloat literals
- ▶ "s" for ScaledDecimal literals
- ▶ "p" for FixedPoint literals

For details, see "GemStone Smalltalk Lexical Tokens" on page 412.

To represent a number in a nondecimal base literally, write the number's base (in decimal), followed by the radix "r" or character "#", and then the number itself. Here, for example, is how you could write octal 23 and hexadecimal FF:

```
8#23
```

```
16rFF
```

The largest radix available is 36.

## Character Literals

A Smalltalk character literal represents a character, such as one of the symbols of the alphabet. To create a character literal, write a dollar sign (\$) followed by the character's alphabetic symbol. Here are some examples:

```
$b $B $4 $? $$
```

If a nonprinting ASCII character such as a tab or a form feed follows the dollar sign, Smalltalk creates the appropriate internal representation of that character.

GemStone Smalltalk interprets this statement, for example, as a representation of ASCII character 32:

```
$ . "Creates the character representing a space (ASCII 32)"
```

In this example, the period following the space acted as a statement terminator. If no space had separated the dollar sign from the period, GemStone Smalltalk would have interpreted the expression as the character literal representing a period.

## String Literals

Literal strings represent sequences of characters. They are instances of the class `String`, described in Chapter 4, “Collection and Stream Classes”. A literal string is a sequence of characters enclosed by single quotation marks. These are literal instances of `String`:

```
'Intellectual passion drives out sensuality.'  
'A difference of taste in jokes is a great strain  
  on the affections.'
```

When you want to include apostrophes in a literal string, double them:

```
'You can''t make omelettes without breaking eggs.'
```

GemStone Smalltalk faithfully preserves control characters when it compiles literal strings. The following example creates a `String` containing a line feed (ASCII 10), the end-of-line character:

```
'Control characters such as line feeds  
  are significant in literal strings.'
```

Strings may hold characters with values up to 255, that is, characters that can be representing in a single byte. Characters themselves may have values much higher. If a string includes any characters larger than 255, it is converted to a `DoubleByteString`. If any of the characters require more than two bytes, it becomes a `QuadByteString`. For example, this is a `DoubleByteString`:

```
' koda'
```

## Symbol Literals

A literal `Symbol` is similar to a literal `String`. It is a sequence of characters preceded by a pound sign (`#`). For example:

```
#stuff  
#nonsense  
#may_24_thisYear
```

Literal `Symbols` specified in this way must be legal identifiers or keywords: they must begin with a letter, contain only alphanumeric characters, underscore, and colon. A `Symbol` can contain other characters, or start with a number: in this case, they must be preceded by a pound sign (`#`) and must also be delimited by single quotation marks. For example:

```
#'Gone With the Wind'
```

As with strings that contain characters that require more than a byte to represent, `DoubleByteSymbol` and `QuadByteSymbol` are used for symbol literals that include characters with values over 255.

## Array Literals

Arrays can hold objects of any type, and they respond to messages that read and write individual elements or groups of elements.

A literal `Array` can contain only other literals—Characters, Strings, Symbols, other literal Arrays, and three “special literals” (`true`, `false`, `nil`). The elements of a literal `Array` are

enclosed in parentheses and preceded by a pound sign (#). White space must separate the elements.

Here is an Array that contains two Strings, a literal Array, and a third String:

```
#('string one' 'string two' #('another' 'Array') 'string3')
```

The following Array contains a String, a Symbol, a Character, a Number, and a Boolean:

```
#('string one' #symbolOne $c 4 true)
```

ByteArray literals are similar, but may only hold SmallIntegers in the range 0 to 255, and use square brackets instead of parenthesis.

For example:

```
#[99 97 116]
```

Besides Array literals, you may also specify Array constructors in your code, which are used similarly, but follow quite different rules. For a discussion of array constructors, see page 403.

## Variables and Variable Names

A variable name is a sequence of characters of either or both cases. A variable name must begin with an alphabetic character or an underscore (“\_”), but it can contain numerals. Spaces are not allowed, and the underscore is the only acceptable punctuation mark. Here are some permissible variable names:

```
zero
relationalOperator
Top10SolidGold
A_good_name_is_better_than_precious_ointment
```

Most Smalltalk programmers begin local variable names with lowercase letters and global variable names with uppercase letters. When a variable name contains several words, Smalltalk programmers usually begin each word with an uppercase letter (sometimes called “camelcase”). You are free to ignore either of these conventions, but remember that Smalltalk is case-sensitive. The following are all different names to Smalltalk:

```
VariableName
variableName
variablename
```

Variable names can contain up to 1024 characters.

## Declaring Temporary Variables

GemStone Smalltalk requires you to declare new variable names (implicitly or explicitly) before using them. The simplest kind of variable to declare is the temporary variable. Temporary variables are so called because they are defined only for one execution of the set of statements in which they are declared.

To declare a temporary variable, you surround it with vertical bars as in this example:

```
| myTempVariable |
myTempVariable := 2.
```

You can declare at most 1013 temporary variables for a set of statements. Once declared, a variable can name objects of any kind.

To store a variable for later use, or to make its scope global, you must put it in one of GemStone's shared dictionaries that GemStone Smalltalk uses for symbol resolution. For example:

```
| myTempVariable |  
myTempVariable := 2.  
UserGlobals at: #MyPermanentVariable put: myTempVariable.
```

Subsequent references to `MyPermanentVariable` return the value 2.

## Pseudovariables

You can change the objects to which most variable names refer simply by assigning them new objects. However, five GemStone Smalltalk variables have values that cannot be changed by assignment; they are therefore called *pseudovariables*. They are:

### **nil**

Refers to an object representing a null value. Variables not assigned another value automatically refer to **nil**. **nil** is an instance of `UndefinedObject`.

### **true**

Refers to the object representing logical truth. **true** is an instance of `Boolean`.

### **false**

Refers to the object representing logical false. **false** is an instance of `Boolean`.

### **self**

Refers to the receiver of the message, which differs according to the context. **self** may be used anywhere a method argument or method temporary would be used, except **self** is not allowed on the left side of an assignment. When **self** is used in code that is not part of a method, it resolves to **nil**.

### **super**

Refers to the receiver of the message, but the search for the method to execute will start in the superclass of the class in which the sending method was compiled. **super** may only be used as the receiver of a message send, in code within a method.

## Assignment

Assignment statements in Smalltalk look like assignment statements in many other languages. The following statement assigns the value 2 to the variable `MightySmallInteger`:

```
MightySmallInteger := 2.
```

The next statement assigns the same String to two different variables (C programmers may notice the similarity to C assignment syntax):

```
nonmodularity := interdependence := 'No man is an island'.
```

## Message Expressions

Smalltalk objects communicate with one another by means of messages. Most of your effort in Smalltalk programming will be spent in writing expressions in which messages are passed between objects. This subsection discusses the syntax of those message expressions.

You have already seen several examples of message expressions:

```
2 + 2
5 + 5
```

In fact, the only GemStone Smalltalk code segments you have seen that are not message expressions are literals, variables, and simple assignments:

```
2                "a literal"
variableName     "a variable"
MightySmallInteger := 2.    "an assignment"
```

The ubiquity of message-passing is one of the hallmarks of object-oriented programming.

## Messages

A message expression consists of:

- ▶ an identifier or expression representing the object to receive the message,
- ▶ one or more identifiers called *selectors* that specify the message to be sent, and
- ▶ (possibly) one or more arguments that pass information with the message (these are analogous to procedure or function arguments in conventional programming). Arguments can be written as message expressions.

## Reserved and Optimized Selectors

GemStone represents selectors internally as symbols, and almost all symbols that conform to the unary, binary, or keyword selector patterns are acceptable as selectors. For details on legal selectors, see the BNF on page 410.

There are a few selectors that have been reserved for the sole use of the GemStone kernel classes. The compiler will not allow you to compile methods with reserved selectors. To see the list of reserved and optimized selectors, execute:

```
GsNMethod optimizedSelectors
```

These selectors are reserved:

```
==                _isScaledDecimal    ifTrue:ifFalse:
__inProtectedMode  _isSmallInteger      isKindOf:
_and:             _isSymbol            or:
_downTo:by:do:    _leaveProtectedMode repeat
_downTo:do:       _or:                timesRepeat:
_gsReturnNoResult _stringCharSize     to:by:do:
_gsReturnNothingEnableEvents ~~ to:do:
_isArray          and:                untilFalse
_isExceptionClass ifFalse:           untilFalse:
_isExecBlock     ifFalse:ifTrue:      untilTrue
_isFloat         ifNil:              untilTrue:
_isInteger       ifNil:ifNotNil:    whileFalse
_isNumber        ifNotNil:           whileFalse:
_isOneByteString ifNotNil:ifNil:    whileTrue
_isRange         ifTrue:             whileTrue:
```

In addition, the following methods are optimized or inlined in the class `SmallInteger`:

```
+      -      *      =      ~=      <      <=      >      >=
```

## Conditionally Optimized Selectors

The selectors `isNil`, `notNil`, and `yourself` are optimized by default, so methods with these selectors cannot be compiled.

This optimization is configurable, and can be removed. To see the list of conditionally optimized selectors, execute `GsNMethod configurableOptimizedSelectors`.

### Un-optimizing conditionally optimized selectors

To remove the optimization, as `SystemUser` you must reset `GsNMethods`'s class variable `OptimizedSelectors` to the desired list; note that symbols other than those in the defined optimized selectors will be ignored. This can be done using an expression such as `GsNMethod _classVars at: #OptimizedSelectors put: anArray`.

Then you must run **upgradeImage** or otherwise recompile GemStone source code, and recompile all application methods, so existing compiled methods are updated.

## Messages as Expressions

In the following message expression, the object 2 is the receiver, `+` is the selector, and 8 is the argument:

```
2 + 8
```

When 2 sees the selector `+`, it looks up the selector in its private memory and finds instructions to add the argument (8) to itself and to return the result. In other words, the selector `+` tells the receiver 2 what to do with the argument 8. The object 2 returns another numeric object 10, which can be stored with an assignment:

```
myDecimal := 2 + 8.
```

The selectors that an object understands (that is, the selectors for which instructions are stored in an object's instruction memory or "method dictionary") are determined by the object's class.

### Unary Messages

The simplest kind of message consists only of a single identifier called a unary selector. The selector `negated`, which tells a number to return its negative, is representative:

```
7 negated  
-7
```

Here are some other unary message expressions:

```
9 reciprocal. "returns the reciprocal of 9"  
myArray last. "returns the last element of Array myArray"  
DateTime now. "returns the current date and time"
```

### Binary Messages

Binary message expressions contain a receiver, a single selector consisting of one or two nonalphanumeric characters, and a single argument. You are already familiar with binary

message expressions that perform addition. Here are some other binary message expressions (for now, ignore the details and just notice the form):

```
8 * 8 "returns 64"
4 < 5 "returns true"
myObject = yourObject "returns true if myObject and
yourObject have the same value"
```

## Keyword Messages

Keyword messages are the most common. Each contains a receiver and up to 15 keyword and argument pairs. In keyword messages, each keyword is a simple identifier ending in a colon.

In the following example, 7 is the receiver, `rem:` is the keyword selector, and 3 is the argument:

```
7 rem: 3 "returns the remainder from the division of 7 by 3"
```

Here is a keyword message expression with two keyword-argument pairs:

```
| arrayOfStrings |
arrayOfStrings := Array new: 4.
arrayOfStrings at: (2 + 1) put: 'Curly'.
"puts 'Curly' at index position 3 in the receiver"
```

In a keyword message, the order of the keyword-argument pairs (`at: arg1` `put: arg2`) is significant.

## Combining Message Expressions

In a previous example, one message expression was nested within another, and parentheses set off the inner expression to make the order of evaluation clear. It happens that the parentheses were optional in that example. However, in GemStone Smalltalk as in most other languages, you sometimes need parentheses to force the compiler to interpret complex expressions in the order you prefer.

Combinations of unary messages are quite simple; GemStone Smalltalk always groups them from left to right and evaluates them in that order. For example:

```
9 reciprocal negated
```

is evaluated as if it were parenthesized like this:

```
(9 reciprocal) negated
```

That is, the numeric object returned by `9 reciprocal` is sent the message `negated`.

Binary messages are also invariably grouped from left to right. For example, GemStone Smalltalk evaluates:

```
2 + 3 * 2
```

as if the expression were parenthesized like this:

```
(2 + 3) * 2
```

This expression returns 10. It may be read: "Take the result of sending + 3 to 2, and send that object the message \* 2."

All binary selectors have the same precedence. Only the *sequence* of a string of binary selectors determines their order of evaluation; the identity of the selectors doesn't matter.

However, when you combine unary messages with binary messages, the unary messages take precedence. Consider the following expression, which contains the binary selector `+` and the unary selector `negated`:

```
2 + 2 negated
0
```

This expression returns the result 0 because the expression `2 negated` executes before the binary message expression `2 + 2`. To get the result you may have expected here, you would need to parenthesize the binary expression like this:

```
(2 + 2) negated
-4
```

Finally, binary messages take precedence over keyword messages. For example:

```
myArrayOfNums at: 2 * 2
```

would be interpreted as a reference to `myArrayOfNums` at position 4. To multiply the number at the second position in `myArrayOfNums` by 2, you would need to use parentheses like this:

```
(myArrayOfNums at: 2) * 2
```

## Summary of Precedence Rules

1. Parenthetical expressions are always evaluated first.
2. Unary expressions group left to right, and they are evaluated before binary and keyword expressions.
3. Binary expressions group from left to right, as well, and take precedence over keyword expressions.
4. GemStone Smalltalk executes assignments after message expressions.

## Cascaded Messages

You will often want to send a series of messages to the same object. By *cascading* the messages, you can avoid having to repeat the name of the receiver for each message. A cascaded message expression consists of the name of the receiver, a message, a semicolon, and any number of subsequent messages separated by semicolons.

For example:

```
| arrayOfPoets |
arrayOfPoets := Array new.
(arrayOfPoets add: 'cummings'; add: 'Byron'; add: 'Rimbaud';
yourself)
```

is a cascaded message expression that is equivalent to this series of statements:

```
| arrayOfPoets |
arrayOfPoets := Array new.
arrayOfPoets add: 'cummings'.
arrayOfPoets add: 'Byron'.
arrayOfPoets add: 'Rimbaud'.
arrayOfPoets
```

You can cascade any sequence of messages to an object. And, as always, you are free to replace the receiver's name with an expression whose value is the receiver.

## Array Constructors

Most of the syntax described in this chapter so far is standard Smalltalk syntax. However, GemStone Smalltalk also includes a syntactic construct called a *Array constructor*. An Array constructor is similar to a literal array, but its elements can be written as nonliteral expressions as well as literals. GemStone Smalltalk evaluates the expressions in an Array constructor at run time.

Array constructors look a lot like literal Arrays; the differences are that array constructors are enclosed in braces and have their elements delimited by periods.

The following example shows an Array constructor whose last element, represented by a message expression, has the value 4.

```
"An Array constructor"
{'string one' . #SymbolOne . $c . 2+2}
```

NOTE

*The Array constructor is not part of the Smalltalk standard. You should avoid its use in any code that might be ported to an other Smalltalk dialect. Instead, use a message send constructor such as `Array class >> #with:`, such as `Array with: 'string one' with: $c with: 2+2`.*

Because any valid GemStone Smalltalk expression is acceptable as an array constructor element, you are free to use variable names as well as literals and message expressions:

```
| aString aSymbol aCharacter aNumber |
aString := 'string one'.
aSymbol := #symbolOne.
aCharacter := $c.
aNumber := 4.
{aString . aSymbol . aCharacter . aNumber}
```

The differences in the behavior of array constructors versus literal arrays can be subtle. For example, the literal array:

```
#{123 huh 456}
```

is interpreted as an array of three elements: a `SmallInteger`, a `Symbol`, and another `SmallInteger`. This is true even if you declare the value of `huh` to be a `SmallInteger` such as 88, as shown in this example:

```
| huh |
huh := 88.
#{ 123 huh 456 )
[20176897 sz:3 cls: 66817 Array] an Array
#1 [986 sz:0 cls: 74241 SmallInteger] 123 == 0x7b
#2 [27086593 sz:3 cls: 110849 Symbol] huh
#3 [3650 sz:0 cls: 74241 SmallInteger] 456 == 0x1c8
```

The same declaration used in an array constructor, however, produces an array of three SmallIntegers:

```
| huh |
huh := 88.
{ 123 . huh . 456 }
[20192001 sz:3 cls: 66817 Array] an Array
#1 [986 sz:0 cls: 74241 SmallInteger] 123 == 0x7b
#2 [706 sz:0 cls: 74241 SmallInteger] 88 == 0x58
#3 [3650 sz:0 cls: 74241 SmallInteger] 456 == 0x1c8
```

## Path Expressions

With the exception of Array constructors, most of the syntax described in this chapter so far is standard Smalltalk syntax. GemStone Smalltalk also includes a syntactic construct called a *path*. A path is a special kind of expression that returns the value of an instance variable.

A path is an expression that contains the names of one or more instance variables separated by periods; a path returns the value of the last instance variable in the series. The sequence of the names reflects the order of the objects' nesting; the outermost object appears first in a path, and the innermost object appears last. The following path points to the instance variable name, which is contained in the object anEmployee:

```
anEmployee.name
```

The path in this example returns the value of instance variable name within anEmployee.

If the instance variable name contained another instance variable called last, the following expression would return the value of last:

```
anEmployee.name.last
```

### NOTE

*Use paths only for their intended purposes. Although you can use a path anywhere an expression is acceptable in a GemStone Smalltalk program, paths are intended for specifying indexes, formulating queries, and sorting. In other contexts, a path returns its value less efficiently than an equivalent message expression. Paths also violate the encapsulation that is one of the strengths of the object-oriented data model. Using them can circumvent the designer's intention. Finally, paths are not standard Smalltalk syntax. Therefore, programs using them are less portable than other GemStone Smalltalk programs.*

## Returning Values

Previous discussions have spoken of the "value of an expression" or the "object returned by an expression." Whenever a message is sent, the receiver of the message returns an object. You can think of this object as the message expression's value, just as you think of the value computed by a mathematical function as the function's value.

You can use an assignment statement to capture a returned object:

```
| myVariable |
myVariable := 8 + 9.      "assign 17 to myVariable"
myVariable              "return the value of myVariable"
17
```

You can also use the returned object immediately in a surrounding expression:

```
"puts 'Moe' at position 2 in arrayOfStrings"
| arrayOfStrings |
arrayOfStrings := Array new: 4.
(arrayOfStrings at: 1+1 put: 'Moe'; yourself) at: 2
```

And if the message simply adds to a data structure or performs some other operation where no feedback is necessary, you may simply ignore the returned value.

## A.3 Blocks

A GemStone Smalltalk block is an object that contains a sequence of instructions. The sequence of instructions encapsulated by a block can be stored for later use, and executed by simply sending the block the unary message value. Blocks find wide use in GemStone Smalltalk, especially in building control structures.

A literal block is delimited by brackets and contains one or more GemStone Smalltalk expressions separated by periods. Here is a simple block:

```
[3.2 rounded]
```

To execute this block, send it the message value.

```
[3.2 rounded] value
3
```

When a block receives the message value, it executes the instructions it contains and returns the value of the last expression in the sequence. The block in the following example performs all of the indicated computations and returns 8, the value of the last expression.

```
[89*5. 3+4. 48/6] value
8
```

You can store a block in a simple variable:

```
| myBlock |
myBlock := [3.2 rounded].
myBlock value.
3
```

or store several blocks in more complex data structures, such as Arrays:

```
| factorialArray |
factorialArray := Array new.
factorialArray at: 1 put: [1];
           at: 2 put: [2 * 1];
           at: 3 put: [3 * 2 * 1];
           at: 4 put: [4 * 3 * 2 * 1].
(factorialArray at: 3) value
6
```

Because a block's value is an ordinary object, you can send messages to the value returned by a block.

```
| myBlock |
myBlock := [4 * 8].
myBlock value / 8
4
```

The value of an empty block is nil.

```
[ ] value
nil
```

Blocks are especially important in building control structures. The following section discusses using blocks in conditional execution.

## Blocks with Arguments

You can build blocks that take arguments. To do so, precede each argument name with a colon, insert it at the beginning of the block, and append a vertical bar to separate the arguments from the rest of the block.

Here is a block that takes an argument named *myArg*:

```
[ :myArg | 10 + myArg]
```

To execute a block that takes an argument, send it value: *anArgument*. For example:

```
| myBlock |
myBlock := [ :myArg | 10 + myArg].
myBlock value: 10.
20
```

The following example creates and executes a block that takes two arguments. Notice the use of the two-keyword message value:value:.

```
| divider |
divider := [:arg1 :arg2 | arg1 / arg2].
divider value: 4 value: 2
2
```

A block assigns actual parameter values to block variables in the order implied by their positions. In this example, *arg1* takes the value 4 and *arg2* takes the value 2.

Variables used as block arguments are known only within their blocks; that is, a block variable is local to its block. A block variable's value is managed independently of the values of any similarly named instance variables, and GemStone Smalltalk discards it after the block finishes execution. This example illustrates this:

```
| aVariable |
aVariable := 1.
[:aVariable | aVariable ] value: 10.
aVariable
1
```

You cannot assign to a block variable within its block. This code, for example, would elicit a compiler error:

```
"The following expression attempts an invalid assignment
to a block variable."
[:blockVar | blockVar := blockVar * 2] value: 10
```

## Blocks and Conditional Execution

Most computer languages, GemStone Smalltalk included, execute program instructions sequentially unless you include special flow-of-control statements. These statements specify that some instructions are to be executed out of order; they enable you to skip some instructions or to repeat a block of instructions. Flow of control statements are usually conditional; they execute the target instructions if, until, or while some condition is met.

GemStone Smalltalk flow of control statements rely on blocks because blocks so conveniently encapsulate sequences of instructions. GemStone Smalltalk's most important flow of control structures are message expressions that execute a block if or while some object or expression is true or false. GemStone Smalltalk also provides a control structure that executes a block a specified number of times.

## Conditional Selection

You will often want GemStone Smalltalk to execute a block of code only if some condition is true or only if it is false. GemStone Smalltalk provides the messages `ifTrue: aBlock` and `ifFalse: aBlock` for that purpose. This example contains both of these messages:

```
5 = 5 ifTrue: ['yes, five is equal to five'].  
yes, five is equal to five
```

```
5 > 10 ifFalse: ['no, five is not greater than ten'].  
no, five is not greater than ten
```

In the first of these examples, GemStone Smalltalk initially evaluates the expression `(5 = 5)`. That expression returns the value `true` (a Boolean), to which GemStone Smalltalk then sends the selector `ifTrue:`. The receiver (`true`) looks at itself to verify that it is, indeed, the object `true`. Because it is, it proceeds to execute the block passed as an argument to `ifTrue:`, and the result is a String.

The receiver of `ifTrue:` or `ifFalse:` must be Boolean; that is, it must be either `true` or `false`. In the above example, the expressions `(5 = 5)` and `(5 > 10)` returned `true` and `false`, respectively, because GemStone Smalltalk numbers know how to compute and return those values when they receive messages such as `=` and `>`.

## Two-Way Conditional Selection

You will often want to direct your program to take one course of action if a condition is met and a different course if it isn't. You could arrange this by sending `ifTrue:` and then `ifFalse:` in sequence to a Boolean (`true` or `false`) expression. For example:

```
2 < 5 ifTrue: ['two is less than five'].  
two is less than five
```

```
2 < 5 ifFalse: ['two is not less than five'].  
nil
```

However, GemStone Smalltalk lets you express the same instructions more compactly by sending the single message `ifTrue: block1 ifFalse: block2` to an expression or object that has a Boolean value. Which of that message's arguments GemStone Smalltalk executes depends upon whether the receiver is `true` or `false`. In this example, the receiver is `true`:

```
2 < 5 ifTrue: ['two is less than five']  
ifFalse: ['two is not less than five'].  
two is less than five
```

## Conditional Repetition

You will also sometimes want to execute a block of instructions repeatedly as long as some condition is true, or as long as it is false. The messages `whileTrue: aBlock` and `whileFalse: aBlock` give you that ability. Any block that has a Boolean value responds

to these messages by executing *aBlock* repeatedly while it (the receiver) is true (`whileTrue:`) or false (`whileFalse:`).

Here is an example that repeatedly adds 1 to a variable until the variable equals 5:

```
| sum |
sum := 0.
[sum = 5] whileFalse: [sum := sum + 1].
sum
5
```

The next example calculates the total payroll of a miserly but egalitarian company that pays each employee the same salary.

```
| totalPayroll numEmployees salariesAdded standardSalary |
totalPayroll := 0.00.
salariesAdded := 0.
numEmployees := 40.
standardSalary := 5000.00.
[salariesAdded < numEmployees] whileTrue:
    [totalPayroll := totalPayroll + standardSalary.
     salariesAdded := salariesAdded + 1].
totalPayroll
200000.0
```

Blocks also accept two unary conditional repetition messages, `untilTrue` and `untilFalse`. These messages cause a block to execute repeatedly until the block's last statement returns either true (`untilTrue`) or false (`untilFalse`).

The following example uses `untilTrue` (rather than `whileFalse:`).

```
| sum |
sum := 0.
[sum := sum + 1. sum = 5] untilTrue.
sum
%
5
```

When GemStone Smalltalk executes the block initially (by sending it the message `value`), the block's first statement adds one to the variable `sum`. The block's second statement asks whether `sum` is equal to 5; since it isn't, that statement returns false, and GemStone Smalltalk executes the block again. GemStone Smalltalk continues to reevaluate the block as long as the last statement returns false (that is, while `sum` is not equal to 5).

The descriptions of classes `Boolean` and `Block` in the image describe these flow of control messages and others.

## Formatting Code

GemStone Smalltalk is a free-format language. A space, tab, line feed, form feed, or carriage return affects the meaning of a GemStone Smalltalk expression only when it separates two characters that, if adjacent to one another, would form part of a meaningful token.

In general, you are free to use whatever spacing makes your programs most readable. The following are all equivalent:

```
{'string one'.2+2.'string three'.$c.9*arglebargle}  
  
{ 'string one' . 2+2 . 'string three' . $c . 9*arglebargle }  
  
{ 'string one'.  
  2 + 2.  
  'string three'.  
  $c.  
  9 * arglebargle }
```

## A.4 GemStone Smalltalk BNF

This section provides a complete BNF description of GemStone Smalltalk. Here are a few notes about interpreting the grammar:

A = expr

This defines the syntactic production 'A' in terms of the expression on the right side of the equals sign.

B = C | D

The vertical bar '|' defines alternatives. In this case, the production "B" is one of either "C" or "D".

C = '<'

A symbol in accents is a literal symbol.

D = F G

A sequence of two or more productions means the productions in the order of their appearance.

E = [ A ]

Brackets indicate zero or one optional productions.

F = { B }

Braces indicate zero or more occurrences of the productions contained within.

G = A | (B|C)

Parentheses can be used to remove ambiguity.

In the GemStone Smalltalk syntactic productions in Figure A.1, white space is allowed between tokens. White space is required before and after the '\_' character.

**Figure A.1 GemStone Smalltalk BNF**

```

AExpression = Primary [ AMessage { ';' ACascadeMessage } ]
ABinaryMessage = [ EnvSpecifier ] BinarySelector Primary [ UnaryMessages ]
ABinaryMessages = ABinaryMessage { ABinaryMessage }
ACascadeMessage = UnaryMessage | ABinaryMessage | AKeywordMessage
AKeywordMessage = [ EnvSpecifier | RubyEnvSpecifier ]
    AKeywordPart { AKeywordPart }
AKeywordPart = Keyword Primary UnaryMessages { ABinaryMessage }
AMessage = [ UnaryMessages ] [ ABinaryMessages ] [ AKeywordMessage ]
AnyTerm = Operand [ Operator Operand ]
Array = '(' { ArrayItem } ')'
ArrayBuilder = '#[' [ AExpression { ',' AExpression } ] ']'
    (exists only if System configurationAt:#GemConvertArrayBuilder is true)
ArrayItem = Number | SymbolArrayItem | SymbolLiteral | StringLiteral |
    CharacterLiteral | Array | ArrayLiteral | ByteArrayLiteral
ArrayLiteral = '#' Array
Assignment = VariableName ':=' Statement | VariableName '_' Statement
BinaryMessage = [ EnvSpecifier | RubyEnvSpecifier ] BinarySelector Primary
    [ UnaryMessages ]
BinaryMessages = BinaryMessage { BinaryMessage }

```

```

BinaryPattern = BinarySelector VariableName
Block = '[' [ BlockParameters ] [ Temporaries ] Statements '''
BlockParameters = { Parameter } '|'
ByteArrayLiteral = '#' '[' [ Number { Number } ] '''
    (exists only if System configurationAt:#GemConvertArrayBuilder is false)
CascadeMessage = UnaryMessage | BinaryMessage | KeywordMessage
CurlyArrayBuilder = '{' [ AExpression { '.' AExpression } ] '''
Expression = Primary [ Message { ';' CascadeMessage } ]
KeywordMessage = [ EnvSpecifier | RubyEnvSpecifier ] KeywordPart {
    KeywordPart }
KeywordPart = Keyword Primary UnaryMessages { BinaryMessage }
KeywordPattern = Keyword VariableName { Keyword VariableName }
KeywordPragma = PragmaPair [ PragmaPair ]
Literal = Number | NegNumber | StringLiteral | CharacterLiteral |
    SymbolLiteral | ArrayLiteral | SpecialLiteral | ByteArrayLiteral
Message = [ UnaryMessages ] [ BinaryMessages ] [ KeywordMessage ]
MessagePattern = UnaryPattern | BinaryPattern | KeywordPattern
Method = MessagePattern [ Primitive ] MethodBody
MethodBody = [ Pragmas ] [ Temporaries ] [ Statements ]
NegNumber = '-' Number
Operand = Path | Literal | Identifier
Operator = '=' | '==' | '<' | '>' | '<=' | '>=' | '~=' | '~~'
ParenStatement = '(' Statement ')'
ParenTerm = '(' AnyTerm ')'
Pragma = '< PragmaBody >'
PragmaBody = UnaryPragma | KeywordPragma
Pragmas = Pragma [ Pragma ]
Predicate = ( AnyTerm | ParenTerm ) { '&' Term }
Primary = ArrayBuilder | CurlyArrayBuilder | Literal | Path | Block |
    SelectionBlock | ParenStatement | VariableName
Primitive = '<' [ 'protected' | 'unprotected' ] [ 'primitive:' Digits ] '>'
PragmaPair = ( KeywordNotPrimitive | BinarySelector ) PragmaLiteral
UnaryPragmaIdentifier is any Identifier except 'protected', 'unprotected',
    'requiresVc'
PragmaLiteral = Number | NegNumber | StringLiteral | CharacterLiteral |
    SymbolLiteral | SpecialLiteral
SelectionBlock = '{' Parameter '|' Predicate '''
Statement = Assignment | Expression
Statements = { [ Pragmas ] { Statement '.' } } [ Pragmas ] [ ['^']
    Statement ['.' [ Pragmas ] ] ]
Temporaries = '|' { VariableName } '|'
Term = ParenTerm | Operand
UnaryMessage = [ EnvSpecifier ] Identifier
UnaryMessages = { UnaryMessage }
UnaryPattern = Identifier
UnaryPragma = SpecialLiteral | UnaryPragmaIdentifier

```

GemStone Smalltalk lexical tokens are shown in Figure A.2. No white space is allowed within lexical tokens.

Figure A.2 GemStone Smalltalk Lexical Tokens

---

```

BinaryExponent = ( 'e' | 'E' | 'd' | 'D' | 'q' ) ['-'] Digits
BinarySelector = SelectorCharacter [SelectorCharacter]
Character = Any Ascii character with ordinal value 0..255
CharacterLiteral = '$' Character
Comment = '"' { Character } '"'
DecimalExponent = ( 'f' | 'F' ) ['-'] Digits
Digit = '0' | '1' | '2' | ... | '9'
Digits = Digit {Digit}
EndOfSource = the end of the method source string
Exponent = BinaryExponent | DecimalExponent | ScaledDecimalExponent |
    FixedPointExponent
EnvSpecifier = '@env' Digits ':'
    (no white space before or after Digits)
FractionalPart = '.' Digits [Exponent]
FixedPointExponent = 'p' [ ['-'] Digits ]
Identifier = SingleLetterIdentifier | MultiLetterIdentifier
KeyWord = Identifier ':'
KeyWordNotPrimitive is any KeyWord other than 'primitive:'
Letter = 'A' | 'B' | ... | 'Z' | 'a' | 'b' | ... | 'z' | '_'
MultiLetterIdentifier = Letter ( Letter | Digit ) { Letter | Digit }
Number = RadixedLiteral | NumericLiteral
Numeric = Digit | 'A' | 'B' | ... | 'Z'
NumericLiteral = Digits ( [FractionalPart] | [Exponent] )
Numerics = Numeric { Numeric }
Parameter = ':' VariableName
    (white space allowed between : and variableName)
Path = Identifier '.' PathIdentifier { '.' PathIdentifier }
PathIdentifier = Identifier | '*'
RadixedLiteral = Digits ( '#' | 'r' ) Numerics
ScaledDecimalExponent = 's' [ ['-'] Digits ]
ScdExponTerminator = '"' | WhiteSpace | ',' | ')' | ']' | '}' | '.' |
    ';' | EndOfSource
SelectorCharacter = '+' | '-' | '\' | '*' | '~' | '<' | '>' | '='
    | '|' | '/' | '&' | '@' | '%' | ',' | '?' | '!'
SingleLetterIdentifier = Letter
SpecialLiteral = 'true' | 'false' | 'nil' | '_remoteNil'
StringLiteral = '"' { Character | '"' } '"'
Symbol = Identifier | BinarySelector | ( Keyword { Keyword } )
SymbolArrayItem = Identifier | ( KeyWord { KeyWord } )
SymbolLiteral = '#' ( Symbol | StringLiteral )
VariableName = Identifier

```

---