
GemStone®

*System Administration
Guide
for GemStone/S 64 Bit™*

Version 3.2

April 2014



GEMTALK™
SYSTEMS

GEMSTONE S™ 64

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.

This documentation, or any part of it, may not be reproduced, displayed, photocopied, transmitted, or otherwise copied in any form or by any means now known or later developed, such as electronic, optical, or mechanical means, without express written authorization from GemTalk Systems.

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 © 1986-2014 GemTalk Systems, LLC. All rights reserved by GemTalk Systems.

PATENTS

GemStone software is 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". GemStone software may also be covered by one or more pending United States patent applications.

TRADEMARKS

VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

Sun, **Sun Microsystems**, and **Solaris** are trademarks or registered trademarks of Oracle and/or its affiliates. **SPARC** is a registered trademark of SPARC International, Inc.

HP, **HP Integrity**, and **HP-UX** are registered trademarks of Hewlett Packard Company.

Intel, **Pentium**, and **Itanium** are registered trademarks of Intel Corporation in the United States and other countries.

Microsoft, **MS**, **Windows**, **Windows XP**, **Windows 2003**, **Windows 7**, **Windows Vista** and **Windows 2008** are registered trademarks of Microsoft Corporation in the United States and other countries.

Linux is a registered trademark of Linus Torvalds and others.

Red Hat and all Red Hat-based trademarks and logos are trademarks or registered trademarks of Red Hat, Inc. in the United States and other countries.

SUSE is a registered trademark of Novell, Inc. in the United States and other countries.

AIX, **POWER5**, **POWER6**, and **POWER7** are trademarks or registered trademarks of International Business Machines Corporation.

Apple, **Mac**, **Mac OS**, **Macintosh**, and **Snow Leopard** are trademarks of Apple Inc., in the United States and other countries.

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

About This Manual

This manual tells how to perform day-to-day administration of your GemStone/S 64 Bit repository.

Prerequisites

This manual is intended for users that are at least somewhat familiar with using Smalltalk and the Topaz programming environment to execute GemStone Smalltalk code. It also assumes some familiarity with UNIX.

You should have the GemStone system installed correctly on your host computer, as described in the *GemStone/S 64 Bit Installation Guide* for your platform.

How This Manual is Organized

This manual is organized in three parts: initial configuration, day-to-day administration, and appendixes.

Part 1: System Configuration

- ▶ Chapter 1, “Configuring the GemStone Server,” tells how to adapt the GemStone central repository server to the needs of your application.
- ▶ Chapter 2, “Configuring Gem Session Processes,” tells how to configure the GemStone processes that provide the services to individual application clients.
- ▶ Chapter 3, “Connecting Distributed Systems,” explains the additional steps necessary to run GemStone in a networked environment. It includes examples of how to set up common configurations.

Part 2: System Administration

- ▶ Chapter 4, “Running GemStone,” tells how to start and stop the GemStone system and how to troubleshoot startup problems and unexpected shutdowns.

- ▶ Chapter 5, “Monitoring GemStone,” explains where the system logs are located, how to audit the repository, and how to monitor the performance of the GemStone server and its clients using GemStone Smalltalk methods.
- ▶ Chapter 6, “User Accounts and Security,” provides details on how to log in to the repository, how to create, modify, and remove GemStone user accounts, and how to configure GemStone login security.
- ▶ Chapter 7, “Managing Repository Space,” gives procedures for managing the repository itself: checking free space, adding space, and controlling its growth. It also how to recover from disk-full conditions.
- ▶ Chapter 8, “Managing Transaction Logs,” gives procedures for setting up the optional full incremental logging, managing log space, and archiving the log files.
- ▶ Chapter 9, “Making and Restoring Backups,” gives procedures for making a GemStone full backup while the repository is in use, and for using backups and transaction logs to restore the repository.
- ▶ Chapter 10, “Warm and Hot Standbys,” describes how to setup a standby system for failover in case of unexpected shutdown.
- ▶ Chapter 11, “Managing Memory,” discusses process memory, how to configure it, and how to diagnose problems.
- ▶ Chapter 12, “Managing Growth,” presents the main concepts underlying garbage collection in GemStone and tells when, how, and why to invoke the garbage collection mechanisms.

Part 3: Appendixes

- ▶ Appendix A, “GemStone Configuration Options,” explains how GemStone uses configuration files and describes each configuration option.
- ▶ Appendix B, “GemStone Utility Commands,” describes each of the GemStone-supplied commands defined for use by the GemStone system administrator.
- ▶ Appendix C, “Network Resource String Syntax,” lists the syntax for network resource strings, which allow you to specify the host machine for a GemStone file or process.
- ▶ Appendix D, “GemStone Kernel Objects,” lists the GemStone-supplied objects that are present in your repository after the GemStone system has been successfully installed.
- ▶ Appendix E, “Environment Variables,” lists all environment variables used by GemStone, including those that are reserved.
- ▶ Appendix F, “Object State Change Tracking,” describes how to analyze tranlogs to track the history of object modifications.

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 (|).

Other GemStone Documentation

You may find it useful to look at documents that describe other GemStone system components:

- ▶ *Topaz Programming Environment* – describes Topaz, a scriptable command-line interface to GemStone Smalltalk. Topaz is most commonly used for performing repository maintenance operations.
- ▶ *Programming Guide for GemStone/S 64 Bit* – a programmer’s guide to GemStone Smalltalk, GemStone’s object-oriented programming language.
- ▶ *GemBuilder for Smalltalk Users’s Guide* – describes GemBuilder for Smalltalk, a programming interface that provides a rich set of features for building and running client Smalltalk applications that interact transparently with GemStone Smalltalk.
- ▶ *GemBuilder for C* – describes GemBuilder for C, a set of C functions that provide a bridge between your application’s C code and the application’s database controlled by GemStone.
- ▶ *VSD User’s Guide* – describes VSD, a graphical tool to examine statistics data files generated by the GemStone/S server

In addition, each release of GemStone/S 64 Bit includes *Release Notes*, describing changes in that release, and platform-specific *Installation Guides*, providing system requirements and installation and upgrade instructions.

A description of the behavior of each GemStone kernel class is available in the class comments in the GemStone Smalltalk repository. Method comments include a description of the behavior of methods.

Technical Support

GemStone Website

<http://gemtalksystems.com/techsupport>

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

- ▶ **Documentation** for released versions of all GemTalk products, in PDF form.

- ▶ **Product downloads**, including past and current versions of GemTalk software.
- ▶ **Bugnotes**, identifying performance issues or error conditions that you may encounter when using a GemTalk product.
- ▶ **TechTips**, providing information and instructions that are not in the documentation.
- ▶ **Compatibility matrices**, listing supported platforms for GemTalk product versions.

This material is updated regularly; we recommend checking this site on a regular basis.

Help Requests

You may need to contact Technical Support directly, if your questions are not answered in the documentation or by other material on the Technical Support site. Technical Support is available to customers with current support contracts.

Requests for technical assistance may be submitted online, by email, or by telephone. We recommend you use telephone contact only for more serious 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

Email: techsupport@gemtalksystems.com

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

When submitting a request, please include the following information:

- ▶ Your name and company name.
- ▶ The versions of all related GemTalk products, and of any other related products, such as client Smalltalk products.
- ▶ The operating system, version, and operating system platform you are using.
- ▶ A description of the problem or request.
- ▶ Exact error message(s) received, if any, including log files 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 get immediate response 24 hours a day, 7 days a week, 365 days a year, for issues impacting a production system. For more details, contact GemTalk Support Renewal.

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.

Chapter 1. Configuring the GemStone Server	21
1.1 Configuration Overview	22
The Server Configuration File	23
Example Configuration Settings	24
Recommendations About Disk Usage	26
Why Use Multiple Drives?	26
When to Use Raw Partitions	26
Developing a Failover Strategy	27
1.2 How To Establish Your Configuration	27
Gathering Application Information	27
Planning Operating System Resources	28
Estimating Memory Needs	28
Estimating Swap Space Needs	28
Estimating File Descriptor Needs	29
Reviewing Kernel Tunable Parameters	29
Checking the System Clock	30
/opt/gemstone/locks and /opt/gemstone/log	30
To Set the Page Cache Options and the Number of Sessions	31
Shared Page Cache	31
Procedure	33
Stone's Private Page Cache	33
Diagnostics	34
To Configure the Repository Extents	34
Estimating Extent Size	34
Choosing the Extent Location	35
Setting a Maximum Size for an Extent	35
Pregrowing Extents to a Fixed Size	36
Pregrowing Extents to the Maximum Size	36
Allocating Data to Multiple Extents	37

To Configure the Transaction Logs	42
Choosing a Logging Mode	42
Estimating the Log Size	42
Choosing the Log Location and Size Limit	43
To Configure Server Response to Gem Fatal Errors	44
To Set File Permissions for the Server	45
Using the Setuid Bit	45
Alternative: Use Group Write Permission	46
File Permissions for the Files and Directories	46
1.3 How To Set Up a Raw Partition	47
Sample Raw Partition Setup	48
Changing Between Files and Raw Partitions	48
Moving an Extent to a Raw Partition	48
Moving an Extent to the File System	49
Moving Transaction Logging to a Raw Partition	49
Moving Transaction Logging to the File System	49
1.4 How To Access the Server Configuration at Run Time	50
To Access Current Settings at Run Time	50
To Change Settings at Run Time	50
1.5 How To Tune Server Performance	52
Tuning the Shared Page Cache	52
Adjusting the Cache Size	53
Matching Spin Lock Limit to Number of Processors	53
Clustering Objects That Are Accessed Together	53
Reducing Swapping	53
Controlling Checkpoint Frequency	53
Tuning Page Server Behavior	55
To Add AIO Page Servers	55
Free Frame Page Servers	55
Process Free Frame Caches	56
1.6 How To Run a Second Repository	56

Chapter 2. Configuring Gem Session Processes **57**

2.1 Overview	57
Linked and RPC Applications	58
The Session Configuration File	59
2.2 How To Configure Gem Session Processes	59
Gathering Application Information	59
Planning Operating System Resources	59
Estimating Memory Needs	60
Estimating Swap Space Needs	60
Estimating File Descriptor Needs	60
Reviewing Kernel Tunable Parameters	61
To Set Ownership and Permissions for Session Processes	61
To Set Access for Linked Applications	62

To Set Access for All Other Applications	62
To Set Access to Other Files	62
To Configure for Remote Shared Page Caches	63
2.3 Set the Gem Configuration Options	63
Configure Temporary Object Space	63
Configure SSL for remote Gem sessions	63
2.4 How To Access the Configuration at Run Time	64
To Access Current Settings at Run Time	64
To Change Settings at Run Time	64
2.5 Tuning Gem Performance	66
Private Page Cache	66
Native Code	66
2.6 How To Install a Custom Gem	67

Chapter 3. Connecting Distributed Systems **69**

3.1 Overview	69
GemStone NetLDIs	71
NetLDI Ports and Names	72
Stone and Shared Page Cache Monitor	72
GemStone Page Servers	72
GemStone Network Objects	73
Shared Page Cache in Distributed Systems	73
Disrupted Communications	74
3.2 How To Arrange Network Security	75
Running as Root with Authentication	75
Setting host Username and Password	75
Authentication Levels	76
Running in Guest Mode with Captive Account	76
3.3 How To Use Network Resource Strings	78
To Set a Default NRS	78
3.4 How To Set Up a Remote Session	79
To Duplicate the GemStone Installation	80
To Share a GemStone Directory	80
Configuration Examples	80
To Run a Linked Application on a Remote Node	81
To Run the Gem Session Process on the Stone's Node	83
To Run the Gem and Stone on Different Nodes	84
To Run the Application, Gem, and Stone on Three Nodes	86
To Run a Distributed System with a Mid-Level Cache	88
3.5 Troubleshooting Remote Logins	90
If You Still Have Trouble	91
Check NetLDI Log Files	92

Chapter 4. Running GemStone **93**

4.1 How To Start the GemStone Server	93
To Start GemStone	94
To Troubleshoot Stone Startup Failures	95
Missing or Invalid Key File	95
Shared Page Cache Cannot Be Attached	95
Extent Missing or Access Denied	96
Extent Open by Another Process.	96
Extent Already Exists	96
Other Extent Failures	97
Transaction Log Missing	97
Other Startup Failures.	98
4.2 How To Start a NetLDI.	98
To Troubleshoot NetLDI Startup Failures	99
4.3 To List Running Servers	100
4.4 How To Start a GemStone Session	100
To Define a GemStone Session Environment	100
To Start a Linked Session.	101
To Start an RPC Session	102
To Troubleshoot Session Login Failures	103
4.5 How To Identify Sessions Logged In	104
4.6 How To Shut Down the Object Server and NetLDI.	105
4.7 How To Recover from an Unexpected Shutdown.	106
Normal Shutdown Message	107
Disk Failure or File System Corruption.	107
Shared Page Cache Error	108
Fatal Error Detected by a Gem.	108
Some Other Shutdown Message.	108
No Shutdown Message	109
4.8 How To Bulk-Load Objects	109
4.9 Considerations for Large Repositories	110
Disk Space and Commit Record Backlogs	110
Handling signals indicating a commit record backlog	111

Chapter 5. Monitoring GemStone **113**

5.1 GemStone System Logs	113
GemStone Server Logs	114
Log file deletion on shutdown	114
Stone Log	115
Admin Gem Logs	116
Reclaim Gem Log	116
Shared Page Cache Monitor Log.	116
Free Frame Page Server Log	117
AIO Page Server Log	117

Page Manager Log117
Symbol Gem Log117
Logs Related to Gem Sessions.118
NetLDI Logs119
Logsender and logreceiver logs.119
Localizing timestamps in log files119
Programmatically adding messages to logs120
5.2 How To Audit the Repository120
To Perform a Page Audit120
To Perform an Object Audit and Repair122
Performing the Object Audit.123
Error Recovery123
5.3 Profiling Repository Contents125
5.4 Monitoring Performance127
Statmonitor and VSD.127
Programmatic Access to Cache Statistics127
Host Statistics132
Host Statistics for processes132
Host Statistics for OS132

Chapter 6. User Accounts and Security 135

6.1 GemStone Users135
UserProfiles135
AllUsers136
Special System Users136
6.2 Creating and Removing Users137
UserProfile Data137
User ID.138
Password.138
Default Object Security Policy.138
Privileges.138
Groups141
Symbol Lists141
Creating Users143
Removing Users143
6.3 Administering Users.145
List Existing Users145
Modifying the UserId145
Modifying Password145
Modifying defaultObjectSecurityPolicy146
Modifying Groups148
Modifying Privileges.149
Modifying SymbolLists151
Disable and Enable User Logins152
Disable and Enable Commits by User154

6.4 Password Authentication	155
GemStone Authentication	155
UNIX Authentication	155
LDAP Authentication.	156
Note on anonymous binds to LDAP server.	157
Determining an Account's Authentication Scheme	158
6.5 Configuring GemStone Login Security.	159
Limiting Choice of Passwords	159
Disallowing Particular Passwords.	160
Disallowing Reuse of Passwords	161
Password Aging – Require Periodic Password Changes	162
Account Aging – Disable Inactive Accounts	164
Enabling Account Aging and lastLoginTime.	165
Limit Logins Until Password Is Changed	166
Limit Concurrent Sessions by a Particular UserId.	166
Limit Login Failures.	167
6.6 Tracking User Logins.	168
Login logging	168
Login Hook	168

Chapter 7. Managing Repository Space **169**

7.1 Repository Growth	169
7.2 How To Check Free Space	170
7.3 How To Add Extents	171
To Add an Extent While the Stone is Running.	171
Possible Effects on Other Sessions	171
Repository>>createExtent:	172
Repository>>createExtent:withMaxSize:	173
7.4 How To Remove an Extent	173
7.5 How To Reallocate Existing Objects Among Extents	173
To Reallocate Objects Among a Different Number of Extents	174
To Reallocate Objects Among the Same Number of Extents	174
7.6 How To Shrink the Repository	175
7.7 How To Check Page Fragmentation	178
7.8 How To Recover from Disk-Full Conditions.	178
Repository Full	179

Chapter 8. Managing Transaction Logs **181**

8.1 Overview	181
Logging Modes	182
Recovering from an Unexpected Shutdown	184
Restoring Transactions to a Backup.	184
How the Logs Are Used.	184

8.2 How To Manage Full Logging 185
 To Archive Logs 186
 To Add a Log at Run Time 187
 To Force a New Transaction Log 188
 To Initiate a Checkpoint 188
 To Change to Partial Logging 189
 8.3 How To Manage Partial Logging. 189
 To Change to Full Logging 189
 8.4 How To Recover from Tranlog-Full Conditions 190
 Transaction Log Space Full 190

Chapter 9. Making and Restoring Backups 191

9.1 Overview 191
 9.2 Types of Backups. 192
 9.3 How To Make an Extent Snapshot Backup 193
 Extent Snapshot Backup when the Repository is shutdown. 193
 Extent Snapshot Backup when the Repository is running 194
 9.4 How To Make a Smalltalk Full Backup 196
 Monitoring and Performance 198
 Compressed Backups 199
 Verifying a Backup is Readable. 199
 Checking Backup Start and Completion 200
 9.5 How to Restore from Backup 200
 Restoring from an Extent Snapshot Backup 203
 Restoring from a Full Backup 204
 9.6 How to Restore Transaction Logs 207
 9.7 Special Cases and Errors in Tranlog Restore 209
 Precautions When Restoring a Subset of Transaction Logs 209
 Restoring Logs up to a Specific Log 210
 Restoring Logs to a Point in Time 211
 Errors While Restoring Transaction Logs 212
 9.8 Recovering from File System Problems 214
 9.9 Version Compatibility 215
 9.10 Warm and Hot Standby Systems 215

Chapter 10. Warm and Hot Standbys 217

10.1 Warm Standby 218
 Setup and run the warm standby. 218
 Activate the warm standby in case of failure in the primary. 219
 10.2 Hot Standby 220
 Hot standby processes 220
 logsender. 220
 logreceiver. 221

Continuous Restore Mode	221
Transaction Record Transmittal	221
Multiple standby repositories	222
To setup and run the hot standby	222
Activate the hot standby in case of failure in the primary	223
Planned failovers	224
Connecting using SSL Mode	225
Chapter 11. Managing Memory	227
11.1 Memory Organization	227
11.2 Configuring Temporary Memory Usage	228
Configuration Options	228
Methods for Computing Temporary Object Space	229
Debugging out-of-memory errors	230
Signal on low memory condition	232
Chapter 12. Managing Growth	233
12.1 Basic Concepts	233
Shadow or Dead?	235
What Happens to Garbage?	238
Admin and Reclaim Gems	239
GemStone's Garbage Collection Mechanisms	240
Marking	240
Reclaiming	240
GcLock	241
Symbol Garbage Collection	241
12.2 MarkForCollection	242
Impact on Other Sessions.	243
Scheduling markForCollection	243
12.3 The FDC/MGC Process.	244
Run markGcCandidatesFromFile:.	245
Impact on Other Sessions.	246
12.4 Epoch Garbage Collection	247
Running Epoch Garbage Collection.	247
Tuning Epoch	248
Cache Statistics	252
12.5 Reclaim	253
Tuning Reclaim	254
Reclaim Configuration Parameters	254
Reclaim Commit Frequency	255
Controlling the impact of reclaim	255
Speeding up reclaim.	255
Avoiding disk space issues	255

Cache Statistics255
12.6 Running Admin and Reclaim Gems256
Configuring Admin Gem256
Configuring Reclaim Gem256
Starting GcGems257
Stopping GcGems257
Adjusting the number of Reclaim sessions258
12.7 Further Tuning Garbage Collection.259
Multi-Threaded Scan.259
Tuning Multi-Threaded Scan259
Memory Impact260
Identifying Sessions Holding Up Voting261
Tuning Write Set Union Sweep261
Identifying Sessions Holding Up Page Reclaim.261
Finding large objects that are using excessive space262
Identify Large Objects in the Repository262
Finding References to an Object that prevent garbage collection263

Appendix A. GemStone Configuration Options **265**

A.1 How GemStone Uses Configuration Files266
Search for a System-Wide Configuration File266
Search for an Executable Configuration File.268
Creating or Using a System Configuration File269
Creating an Executable Configuration File269
Naming Executable Configuration Files269
Naming Conventions for Configuration Options.271
A.2 Configuration File Syntax271
Errors in Configuration Files272
Syntax Errors272
Option Value Errors273
A.3 Configuration Options274
DBF_ALLOCATION_MODE274
DBF_EXTENT_NAMES274
DBF_EXTENT_SIZES274
DBF_PRE_GROW275
DBF_SCRATCH_DIR276
DUMP_OPTIONS276
GEM_ABORT_MAX_CRIS276
GEM_FREE_FRAME_CACHE_SIZE.276
GEM_FREE_FRAME_LIMIT276
GEM_FREE_PAGEIDS_CACHE277
GEM_GCI_LOG_ENABLED277
GEM_HALT_ON_ERROR277
GEM_KEEP_MIN_SOFTREFS277
GEM_MAX_SMALLTALK_STACK_DEPTH.278

GEM_NATIVE_CODE_ENABLED	278
GEM_PGSRV_COMPRESS_PAGE_TRANSFERS	278
GEM_PGSRV_FREE_FRAME_CACHE_SIZE	279
GEM_PGSRV_FREE_FRAME_LIMIT	279
GEM_PGSRV_UPDATE_CACHE_ON_READ	279
GEM_PRIVATE_PAGE_CACHE_KB	280
GEM_REPOSITORY_IN_MEMORY	280
GEM_RPCGCI_TIMEOUT	280
GEM_RPC_KEEPALIVE_INTERVAL	280
GEM_RPC_USE_SSL	281
GEM_SOFTREF_CLEANUP_PERCENT_MEM	281
GEM_TEMPOBJ_AGGRESSIVE_STUBBING	281
GEM_TEMPOBJ_CACHE_SIZE	282
GEM_TEMPOBJ_MESPACE_SIZE	282
GEM_TEMPOBJ_OOPMAP_SIZE	283
GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE	283
GEM_TEMPOBJ_POMGEN_SCAVENGE_INTERVAL	283
GEM_TEMPOBJ_POMGEN_SIZE	284
GEM_TEMPOBJ_SCOPES_SIZE	284
GEM_TEMPOBJ_START_ADDR	284
KEYFILE	284
LOG_WARNINGS	285
SHR_NUM_FREE_FRAME_SERVERS	285
SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_POLICY	285
SHR_PAGE_CACHE_LOCKED	286
SHR_PAGE_CACHE_NUM_PROCS	286
SHR_PAGE_CACHE_NUM_SHARED_COUNTERS	286
SHR_PAGE_CACHE_PERMISSIONS	287
SHR_PAGE_CACHE_SIZE_KB	287
SHR_SPIN_LOCK_COUNT	287
SHR_TARGET_FREE_FRAME_COUNT	288
SHR_WELL_KNOWN_PORT_NUMBER	288
STN_ADMIN_GC_SESSION_ENABLED	288
STN_ALLOCATE_HIGH_OOPS	288
STN_ALLOW_NFS_EXTENTS	289
STN_CACHE_WARMER	289
STN_CACHE_WARMER_SESSIONS	289
STN_CHECKPOINT_INTERVAL	289
STN_COMMIT_QUEUE_THRESHOLD	290
STN_COMMIT_RECORD_QUEUE_SIZE	290
STN_COMMIT_TOKEN_TIMEOUT	290
STN_COMMITS_ASYNC	290
STN_CR_BACKLOG_THRESHOLD	291
STN_DISABLE_LOGIN_FAILURE_LIMIT	
STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT	291
STN_DISKFULL_TERMINATION_INTERVAL	292
STN_EPOCH_GC_ENABLED	292

STN_EXTENT_IO_FLAGS	292
STN_FREE_FRAME_CACHE_SIZE	293
STN_FREE_SPACE_THRESHOLD	293
STN_GEM_ABORT_TIMEOUT	293
STN_GEM_LOSTOT_TIMEOUT	294
STN_GEM_TIMEOUT	294
STN_HALT_ON_FATAL_ERR	294
STN_LISTENING_ADDRESSES	295
STN_LOG_IO_FLAGS	295
STN_LOG_LOGIN_FAILURE_LIMIT	
STN_LOG_LOGIN_FAILURE_TIME_LIMIT	296
STN_LOGIN_LOG_ENABLED	296
STN_LOOP_NO_WORK_THRESHOLD	297
STN_MAX_AIO_RATE	297
STN_MAX_AIO_REQUESTS	297
STN_MAX_GC_RECLAIM_SESSIONS	298
STN_MAX_LOGIN_LOCK_SPIN_COUNT	298
STN_MAX_REMOTE_CACHES	298
STN_MAX_SESSIONS	299
STN_MAX_VOTING_SESSIONS	299
STN_NUM_AIO_WRITE_THREADS	299
STN_NUM_GC_RECLAIM_SESSIONS	299
STN_NUM_LOCAL_AIO_SERVERS	300
STN_OBJ_LOCK_TIMEOUT	300
STN_PAGE_MGR_COMPRESSION_ENABLED	300
STN_PAGE_MGR_MAX_WAIT_TIME	301
STN_PAGE_MGR_PRINT_TIMEOUT_THRESHOLD	301
STN_PAGE_MGR_REMOVE_MAX_PAGES	301
STN_PAGE_MGR_REMOVE_MIN_PAGES	302
STN_PRIVATE_PAGE_CACHE_KB	302
STN_REMOTE_CACHE_PGSRV_TIMEOUT	302
STN_REMOTE_CACHE_TIMEOUT	302
STN_SHR_TARGET_PERCENT_DIRTY	303
STN_SIGNAL_ABORT_CR_BACKLOG	303
STN_SYMBOL_GC_ENABLED	303
STN_TRAN_FULL_LOGGING	303
STN_TRAN_LOG_DEBUG_LEVEL	304
STN_TRAN_LOG_DIRECTORIES	304
STN_TRAN_LOG_LIMIT	304
STN_TRAN_LOG_PREFIX	305
STN_TRAN_LOG_SIZES	305
STN_TRAN_Q_TO_RUN_Q_THRESHOLD	305
STN_WELL_KNOWN_PORT_NUMBER	306
A.4 Runtime-only Configuration Parameters	306
GemConvertArrayBuilder	306
GemDropCommittedExportedObjs	306
GemExceptionSignalCapturesStack	306

LogOriginTime	307
SessionInBackup.	307
StnCurrentTranLogDirId	307
StnCurrentTranLogNames.	307
StnLogFileName.	307
StnLogGemErrors	307
StnLoginsSuspended	307
StnMaxReposSize	308
StnMaxSessions	308
StnSunsetDate	308
StnTranLogOriginTime.	308

Appendix B. GemStone Utility Commands **309**

copydbf	310
gslis	314
pageaudit	316
pstack.	317
removedbf	318
startcachewarmer	319
startlogreceiver	320
startlogsender	322
startnetldi	324
startstone.	326
statmonitor	327
stoplogreceiver	329
stoplogsender	330
stopnetldi	331
stopstone.	332
topaz	333
vsd	334
waitstone.	335

Appendix C. Network Resource String Syntax **337**

Overview	337
Defaults	338
Notation	338
Syntax	339

Appendix D. GemStone Kernel Objects **343**

Non-Numeric Constants	343
---------------------------------	-----

Numeric Constants343
Repository and GsObjectSecurityPolicies344
Global Variables and Collections.345
Current TimeZone348
Zoneinfo349
Utilities350

Appendix E. Environment Variables **353**

Public Environment Variables353
System Variables Used by GemStone357
Reserved Environment Variables357

Appendix F. Object State Change Tracking **359**

F.1 Overview359
F.2 Tranlog Analysis Scripts360
Script Prerequisites360
Output360
Tranlog Assumptions360
Filter Criteria361
printlogs362
Examples.362
searchlogs363
Examples.363
F.3 Tranlog Structure363
Tranlog Entries364
Tranlog Entry Types365
Very Large Objects366
Full vs. Normal Mode367
F.4 Example of Tranlog Analysis368
Tracking Changes to an Employee369
Changed vs. new objects370
Details of Changes to an Employee.371
F.5 Further Analysis373
Class Operations373
Deleted Objects373
Managing Volume374

Index **375**

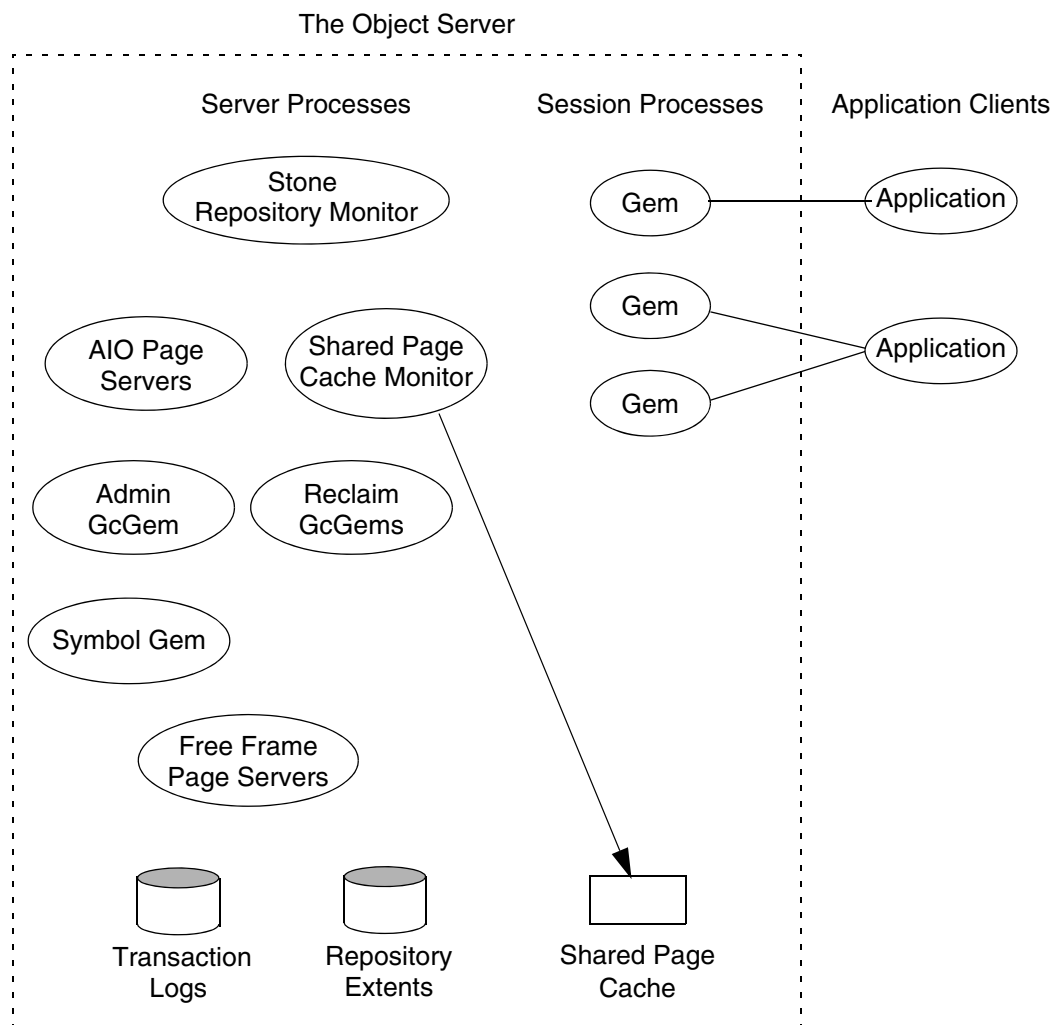
Configuring the GemStone Server

Figure 1.1 shows the basic GemStone/S 64 Bit architecture as seen by its administrator. The object server can be thought of as having two active parts. The *server processes* consist of the Stone repository monitor and a set of subordinate processes. These processes provide resources to individual Gem *session processes*, which are servers to application clients.

This chapter tells you how to configure the GemStone server processes, repository, transaction logs, and shared page cache to meet the needs of a specific application. For information about configuring session processes for clients, refer to Chapter 2.

The elements shown in Figure 1.1 can be distributed across multiple nodes to meet your application's needs. For information about establishing distributed servers, refer to Chapter 3.

Figure 1.1 The GemStone Object Server



1.1 Configuration Overview

Figure 1.1 shows the key parts that define the server configuration:

- ▶ The *Stone repository monitor* process acts as a resource coordinator. It synchronizes critical repository activities and ensures repository consistency.
- ▶ The *shared page cache monitor* creates and maintains a *shared page cache* for the GemStone server. The monitor balances page allocation among processes, ensuring that a few users or large objects do not monopolize the cache. The size of the shared page cache is configurable and should be scaled to match the size of the repository and the number of concurrent sessions; a larger size often provides better performance.

- ▶ The *AIO page servers* perform asynchronous I/O for the Stone repository monitor. Their primary tasks are to update the extents periodically and to pre-allocate (grow) the extents at startup when that feature is enabled. The default configuration uses one AIO page server, but additional ones should be specified for systems having several extents.
- ▶ The *Admin Gem* is a Gem server process that is dedicated to performing the administrative garbage collection tasks under supervision of the Stone. Each repository can have up to one Admin Gem process running.
- ▶ The *Reclaim Gem* performs page reclaim operations on both shadow objects and dead objects. Each repository can have up to one Reclaim Gem process running. The Reclaim Gem controls multiple sessions performing the reclaim task.
- ▶ The *Symbol Gem* is a Gem server background process that is responsible for creating all new Symbols, based on session requests that are managed by the Stone.
- ▶ The *Free Frame Page Servers* are Gem server processes that are dedicated to the task of adding free frames to the free frame list, from which a Gem can take as needed. The default configuration uses one free frame page server, but you can configure as many as 30 free frame page server processes.
- ▶ Objects are stored on the disk in one or more *extents*, which can be files in the file system, data in raw partitions, or a mixture. The location of each extent is configurable.
- ▶ *Transaction logs* permit recovery of committed data if a system crash occurs, and in *full logging mode* allows transaction logs to be used with GemStone backups for full recovery of committed transactions in the event of media failure.

The transaction logs should reside on a different disk drive from the extents, and neither should be on a drive that contains the operating system swap space (sometimes called page space).

The Server Configuration File

At start-up time, GemStone reads a system-wide configuration file. By default this file is `$(GEMSTONE)/data/system.conf`, where `GEMSTONE` is an environment variable that points to the directory in which the GemStone software is installed.

Appendix A, “GemStone Configuration Options,” tells how to specify an alternate configuration file and how to use supplementary files to adjust the system-wide configuration for a specific GemStone executable. The appendix also describes each of the configuration options.

Here is a brief summary of important facts about the configuration file:

- ▶ Lines that begin with `#` are comments. Settings supplied as comments are the same as the default values. You can easily change the configuration by altering the option value and moving the `#` symbol to the line previously in force.
- ▶ Options that begin with “`GEM_`” are read only by Gem session processes at the time they start. Chapter 2, “Configuring Gem Session Processes,” describes their use.
- ▶ Options that begin with “`SHR_`” are read both by the Stone repository monitor and by the first Gem session process to start on a node remote from the Stone. These options configure the local shared page cache.

- ▶ Most other options (those not beginning with “GEM_” or “SHR_”) are read by the Stone repository monitor. If another GemStone process needs that information, it is exchanged through a TCP/IP connection with the Stone.
- ▶ If an option is defined more than once, only the last definition is used. Certain run-time configuration changes, such as the addition of an extent, cause the repository monitor to append new configuration statements to the file. Be sure to check the end of a configuration file for possible entries that override earlier ones.
- ▶ Most configuration parameters have default values. If you do not specify a setting for the parameter, or if there is an error in your settings such as an out of range value, the default is used. After making changes, check the log file headers to ensure that your configuration setting is correct.

Example Configuration Settings

The configuration file that is provided in the GemStone distribution, in `$GEMSTONE/data/system.conf`, provides default values that are suitable as an initial configuration, for small systems.

This section describes possible sets of configuration parameter setting changes that may be useful for larger systems. These provides a starting point; the actual values will need to be adjusted for your particular hardware and application requirements. They give some sense of how some of the more important configuration parameters scale relative to each other.

Large systems will almost certainly require additional tuning for optimal performance. GemStone Professional Services can provide expert assistance in establishing your configuration, tuning configurations for performance and to accommodate growth.

More information about these settings is provided in the detailed instructions for establishing your own configuration, beginning on page 27. For details on the specific configuration parameters, see Appendix A.

Gemservers for Large Configurations

Depending on your hardware, there is an upper limit to the number of processes that you can run before performance becomes unacceptable. For very large configurations, it may be useful to establish a separate Gem server machine with a remote shared page cache set up specifically to run Gem sessions, with a high bandwidth connection between the repository server and the Gem server.

Table 1 Settings for Sample Configurations

Characteristic or Configuration Option	Sample Server Configuration		
	Small	Medium	Large
Application Characteristics			
Maximum number of user sessions	12	250	1000
Repository size	100 MB	10 GB	500 GB
System Requirements			
Typical number of CPUs	2	4	8+
RAM	4 GB	8 GB	128 GB
Kernel shared memory	2 GB	6 GB	100 GB
Number of disk drives	4	8	12
Configuration Settings			
STN_MAX_SESSIONS	20	150	1200
SHR_PAGE_CACHE_SIZE_KB	500 MB	4 GB	96 GB
STN_CR_BACKLOG_THRESHOLD	40	500	3000
STN_SIGNAL_ABORT_CR_BACKLOG	20	400	2800
STN_NUM_LOCAL_AIO_SERVERS	1	4	8
Approximate Memory Usage			
Stone repository monitor (MB)	20	40	100
Each Gem session process ^a (MB)	50	50	50
Extents			
DBF_EXTENT_NAMES	(2 files)	(4 files)	(8 files)
DBF_EXTENT_SIZES	(Unlimited)	(unlimited)	(unlimited)
DBF_ALLOCATION_MODE	10, 10	10,10,10,10	10,10,10,...
Transaction Logs			
STN_TRAN_LOG_DIRECTORIES	(2 directories)	(4 directories)	(8 directories)
STN_TRAN_LOG_SIZES	100	499 ^b each	1995 ^b each
^a Depends on the value of GEM_TEMPOBJ_CACHE_SIZE (default=50 MB but a larger setting is often required).			

Recommendations About Disk Usage

You can enhance server performance by distributing the repository files on multiple disk drives. If you are using a SAN or a disk using some RAID configurations, or have configured multiple logical disks on the same physical spindle, the follow discussion is not entirely applicable.

Why Use Multiple Drives?

Efficient access to GemStone repository files requires that the server node have at least three disk drives (that is, three separate spindles or physical volumes) to reduce I/O contention. For instance:

- ▶ One disk for swap space and the operating system (GemStone executables can also reside here).
- ▶ One disk for the repository extent, perhaps with a lightly accessed file system sharing the drive.
- ▶ One disk for transaction logs and possibly user file systems if they are only lightly used for non-GemStone purposes.

When developing your own configuration, bear in mind the following guidelines:

1. Keep extents and transaction logs separate from operating system swap space. Don't place either extents or logs on a disk that contains a swap partition; doing so drastically reduces performance.
2. Place the transaction logs on a disk that does not contain extents. Placing logs on a different disk from extents increases the transaction rate for updates while reducing the impact of updates on query performance. You can place multiple logs on the same disk, since only one log file is active at a time.
3. To benefit from multiple extents on multiple disks, you must use weighted allocation mode. If you use sequential allocation, multiple extents provide little benefit. For details about weighted allocation, see "Allocating Data to Multiple Extents" on page 37.
4. More than one AIO page server is not of value unless you have multiple extents on multiple disks.

When to Use Raw Partitions

Each raw partition (sometimes called a raw device or raw logical device) is like a single large sequential file, with one extent or one transaction log per partition. The use of raw disk partitions can yield better performance, depending on how they are used and the balancing of system resources.

Placing transaction logs on raw disk partitions is likely to yield better performance.

Usually, placing extents on file systems is as fast as using raw disk partitions. It is possible for this to yield better performance, if doing so reduces swapping; but it is not recommended to use configurations in which swapping occurs. Sufficient RAM should be made available for file system buffers and the shared page cache.

The use of raw partitions for transaction logs is useful for achieving the highest transaction rates in an update-intensive application because such applications primarily are writing sequentially to the active transaction log. Using raw partitions can improve the

maximum achievable rate by avoiding the extra file system operations necessary to ensure that each log entry is recorded on the disk. Transaction logs use sequential access exclusively, so the devices can be optimized for that access.

Because each partition holds a single log or extent, if you place transaction logs in raw partitions, you must provide at least two such partitions so that GemStone can preserve one log when switching to the next. If your application has a high transaction volume, you are likely to find that increasing the number of log partitions makes the task of archiving the logs easier.

For information about using raw partitions, see “How To Set Up a Raw Partition” on page 47.

Developing a Failover Strategy

In choosing a failover strategy, consider the following needs:

- ▶ Applications that cannot tolerate the loss of committed transactions should mirror the transaction logs (using OS-level tools) and use full transaction logging. A mirrored transaction log on another device allows GemStone to recover from a read failure when it starts up after an unexpected shutdown. Full logging mode allows transactions to be rolled forward from a GemStone backup to recover from the loss of an extent without data loss.
- ▶ Applications that require rapid recovery from the loss of an extent (that is, without the delay of restoring from a backup) may wish to replicate all extents on other devices through hardware means, in addition to mirroring transaction logs. Restoring a large repository (many GB) from a backup may a significant time.
- ▶ Setting up a warm or hot standby system can allow fast failover. See “Warm and Hot Standbys” on page 217.

1.2 How To Establish Your Configuration

Configuring the GemStone object server involves the following steps:

1. Gather application specifics about the size of the repository and the number of sessions that will be logged in simultaneously.
2. Plan the operating system resources that will be needed: memory and swap (page) space.
3. Set the size of the GemStone shared page cache and the number of sessions to be supported.
4. Configure the repository extents.
5. Configure the transaction logs.
6. Set GemStone file permissions to allow necessary access while providing adequate security.

Gathering Application Information

When you begin configuring GemStone, be sure to have the following information at hand:

- ▶ The number of simultaneous sessions that will be logged in to the repository (in some applications, each user can have more than one session logged in).
- ▶ The approximate size of your repository. It's also helpful, but not essential, to know the approximate number of objects in the repository.

This information is central to the sizing decisions that you must make.

Planning Operating System Resources

GemStone needs adequate memory and swap space to run efficiently. It also needs adequate kernel resources—for instance, kernel parameters can limit the size of the shared page cache or the number of sessions that can connect to it.

Estimating Memory Needs

The amount of memory required to run your GemStone server depends mostly on the size of the repository and the number of users who will be logged in to active GemStone sessions at one time. These needs are in addition to the memory required for the operating system and other software.

- ▶ The Stone and related processes need memory as shown in the configurations in Table 1 (on page 25). Note that this amount of memory is only for the server processes.
- ▶ The shared page cache should be increased in proportion to the overall size of your repository. Typically it should be at least 10% of the repository size to provide adequate performance, and the larger the better. The best performance is when the entire repository can fit into memory and cache warming is used to load all pages into memory.

On a node that is dedicated to running GemStone, we recommend in general that you allocate approximately one-third to one-half of your total system RAM to the shared page cache. If it is not a dedicated node, you may need to reduce the size to avoid swapping.

- ▶ Each Gem session process needs at least 30 MB of memory on the node where it runs—more may be needed, depending on the setting for `GEM_TEMPOBJ_CACHE_SIZE`. (See the discussion of memory needs for session processes on page 60.) Each Gem process that runs on a remote (client) node also needs about 0.25 MB on the server node for a GemStone page server process that accesses the repository extents.

Estimating Swap Space Needs

To provide reasonable flexibility, the total swap space on your system (sometimes called page space) in general should be at least equal to the system RAM. Preferably, swap space should be twice as much as system RAM. For example, a system with 4 GB of RAM should have at least 4 GB of swap space. The command to find out how much swap space is available (`swap`, `swapinfo`, `pstat`, or `lsvg`) depends on your operating system. Your *GemStone/S 64 Bit Installation Guide* contains an example for your platform.

Swap space should not be on a disk drive that contains any of the GemStone repository extent files. In particular, do not use operating system utilities like `swap` or `swapon` to place part of the swap space on a disk that also contains the GemStone extents or transaction logs.

If you want to determine the additional swap space needed just for GemStone, use the memory requirements derived in the preceding section, including space for the number of sessions you expect. These figures will approximate GemStone's needs beyond the swap requirement for UNIX and other software such as the X Window System.

Estimating File Descriptor Needs

When they start, most GemStone processes attempt to raise their file descriptor limit from the default (soft) limit to the hard limit set by the operating system. In the case of the Stone repository monitor, the processes that raise the limit this way are the Stone itself and two of its child processes, the AIO page server and the Admin Gem. The Stone uses file descriptors this way:

- 9 for stdin, stdout, stderr, and internal communication
- 2 for each user session that logs in
- 1 for each local extent or transaction log within a file system
- 2 for each extent or transaction log that is a raw partition
- 1 for each extent or transaction log that is on a remote node

You can cause the above processes to set a limit less than the system hard limit by setting the `GEMSTONE_MAX_FD` environment variable to a positive integer. A value of 0 disables attempts to change the default limit.

The shared page cache monitor always attempts to raise its file descriptor limit to equal its maximum number of clients plus five for stdin, stdout, stderr, and internal communication. The maximum number of clients is set by the `SHR_PAGE_CACHE_NUM_PROCS` configuration option (page 286), which is computed by the `STN_MAX_SESSIONS` configuration option (page 299).

Reviewing Kernel Tunable Parameters

Operating system kernel parameters limit the interprocess communication resources that GemStone can obtain. It's helpful to know what the existing limits are so that you can either stay within them or plan to raise the kernel limits. There are four parameters of primary interest:

- ▶ The maximum size of a shared memory segment (typically `shmmax` or a similar name) limits the size of the shared page cache for each repository monitor.
For information about platform-specific limitations on the size of the shared page cache, refer the *GemStone/S 64 Bit Installation Guide* for your platform.
- ▶ The maximum number of semaphores per semaphore id limits the number of sessions that can connect to the shared page cache, because each session uses two semaphores. (Typically this parameter is `semmsl` or a similar name, although it is not tunable under all operating systems.)
- ▶ The maximum number of users allowed on the system (typically `maxusers` or a similar name) can limit the number of logins and sometimes also is used as a variable in the allocation of other kernel resources by formula. In the latter case, you may need to set it somewhat larger than the actual number of users.
- ▶ The hard limit set for the number of file descriptors can limit the total number of logins and repository extents, as described previously.

How you determine the existing limits depends on your operating system. If the information is not readily available, proceed anyway. A later step shows how to verify that the shared memory and semaphore limits are adequate for the GemStone configuration you chose.

Checking the System Clock

The system clock should be set to the correct time. When GemStone opens the repository at startup, it compares the current system time with the recorded checkpoint times as part of a consistency check. A system time earlier than the time at which the last checkpoint was written may be taken as an indication of corrupted data and prevent GemStone from starting. The time comparisons use GMT. It is not necessary to adjust GemStone for changes to and from daylight savings time.

`/opt/gemstone/locks` and `/opt/gemstone/log`

GemStone requires access to two directories under `/opt/gemstone/`:

- ▶ `/opt/gemstone/locks` is used for lock files, which among other things provide the names, ports, and other important data used in interprocess communication and reported by `gslis`.

Under normal circumstances, you should never have to directly access files in this directory. To clear out lock files of processes that exited abnormally, use `gslis -c`.

- ▶ `/opt/gemstone/log` is the default location for NetLDI log files, if `startnetldi` does not explicitly specify a location using the `-l` option.

If `/opt/gemstone/` does not exist, GemStone may use `/usr/gemstone/` instead.

Alternatively, you can use the environment variable `GEMSTONE_GLOBAL_DIR` to specify a different location. Since the files in this location control visibility of GemStone processes to one another, all GemStone processes that interact must use the same directory.

Host Identifier

`/opt/gemstone/locks` (or an alternate directory, as described above) is also the location for a file named `gemstone.hostid`, which contains the unique host identifier for this host. This file is created by the first GemStone process on that host to require a unique identifier, by reading eight bytes from `/dev/random`. This unique `hostId` is used instead of host name or IP address for GemStone inter-process communication, avoiding issues with multi-homed hosts and changing IP address.

You can access the host identifier for the machine hosting the gem session using the method `System class >> hostId`.

GemStone shared libraries

Your GemStone installation includes shared library files as well as executables. Access to these shared library files is required for the GemStone executables. In the standard installation of the GemStone software, these shared libraries are located in the `$GEMSTONE/lib` and `$GEMSTONE/lib32` directories.

For installations that do not include a full server, such as remote nodes running gems, or GemBuilder clients, these libraries may be put in a directory other than this standard. See the *Installation Guide* for more information.

To Set the Page Cache Options and the Number of Sessions

You will configure the shared page cache and the Stone's private page cache according to the size of the repository and the number of sessions that will connect to it simultaneously. Then use a GemStone utility to verify that the OS kernel will support this configuration.

Shared Page Cache

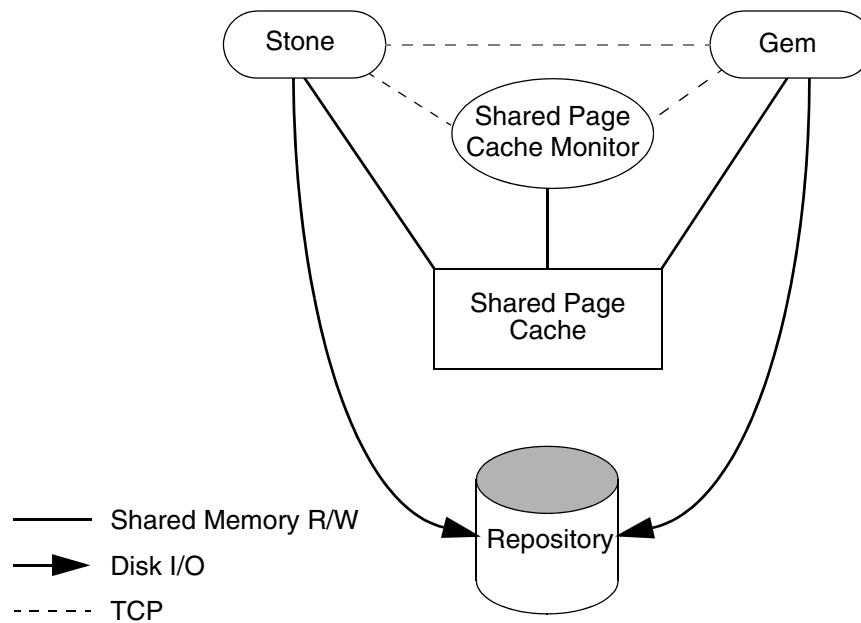
The GemStone shared page cache system consists of two components, the shared page cache itself and a monitor process (`shrpcmonitor`). Figure 1.2 shows the connections between these two and the main GemStone components when GemStone runs on a single node. There is no direct connection between the shared page cache and the repository.

The shared page cache resides in a segment of the operating system's virtual memory that is available to any authorized process. When the Stone repository monitor or a Gem session process needs to access an object in the repository, it first checks to see whether the page containing that object is already in the cache. If the page is already present, the process reads the object directly from shared memory. If the page is not present, the process reads the page from the disk into the cache, where all of its objects also become available to other processes.

The name of the shared page cache monitor ordinarily is derived from the name of the Stone repository monitor and a host-specific ID; for instance, `gs64stone~d7e2174792b1f787`. This host-specific ID is created by the first GemStone process to use a particular host, and remains the same for anything running on that host.

Each shared page cache is associated with exactly one Stone process and repository. A Stone may never have more than one shared page cache on the same node. The Stone spawns the shared page cache automatically during startup, and the shared page cache monitor creates the shared memory region and allocates the semaphores for the system. If other Gem session processes on the same node need to access that repository, they must connect to the same shared page cache and monitor process to ensure cache coherency.

Figure 1.2 Shared Page Cache Configuration



Estimating the Size of the Shared Page Cache

The goal in sizing the shared page cache is to make it large enough to hold the application's working set of objects, thereby reducing disk I/O, while not inducing swapping at the operating system level. Three factors are important in estimating the size:

1. The number of objects in the repository.
For best performance, the entire object table should be in shared memory. At a minimum, we recommend that you allow room for one-third to one-half of the object table in the cache. Each object uses five bytes in the table.
2. The size of the repository. We recommend that you keep at least 25% of the repository in the cache. Although this factor is application-dependent, increasing the amount in the cache will improve performance.
3. The number of users and the size of their transactions. Again, this factor is specific to your application, depending (among other things) on how frequently users will be committing their transactions.

These rules of thumb provide a basic guideline. Ultimately, the cache size needed depends on the working set size (the number of objects and transactional views) that need to be maintained in the cache to provide adequate performance.

Once your application is running, you can tune the cache size by monitoring the free space. See "Monitoring Performance" on page 127, especially the statistics `NumberOfFreeFrames`, `FramesFromFreeList`, and `FramesFromFindFree`.

Procedure

To configure the shared page cache, follow these steps:

Step 1. Set the `SHR_PAGE_CACHE_SIZE_KB` configuration option (page 287), using Table 1 (on page 25) or your own estimate derived above (remember to convert to KB). For instance, for a 1.5 GB cache:

```
SHR_PAGE_CACHE_SIZE_KB = 1500000;
```

Step 2. If the number of sessions will be greater than 40, increase the `STN_MAX_SESSIONS` configuration option (page 299) accordingly.

Make sure the `SHR_PAGE_CACHE_NUM_PROCS` option (page 286) is set to its default (-1), which causes GemStone to calculate a value based on `STN_MAX_SESSIONS`.

For instance:

```
STN_MAX_SESSIONS = 50;
SHR_PAGE_CACHE_NUM_PROCS = -1;
```

Step 3. Use GemStone's `shmem` utility to verify that your OS kernel supports the chosen cache size and number of processes. The command line is

```
$GEMSTONE/install/shmem existingFile cacheSizeKB numProcs
```

where

- ▶ `$GEMSTONE` is the directory where the GemStone software is installed.
- ▶ *existingFile* is the name of any writable file, which is used to obtain an id (the file is not altered).
- ▶ *cacheSizeKB* is the `SHR_PAGE_CACHE_SIZE_KB` setting.
- ▶ *numProcs* is the value calculated by `SHR_PAGE_CACHE_NUM_PROCS`.

For instance, for the values used in Steps 1 and 2, with *numProcs* calculated by computing `SHR_PAGE_CACHE_NUM_PROCS` based on `STN_MAX_SESSIONS`:

```
% touch /tmp/shmem
% $GEMSTONE/install/shmem /tmp/shmem 1500000 62
% rm /tmp/shmem
```

If `shmem` is successful in obtaining a shared memory segment of sufficient size, no message is printed. Otherwise, diagnostic output will help you identify the kernel parameter that needs tuning. The total shared memory requested includes cache overhead for cache space and per-session overhead. The actual shared memory segment in this example would be 104865792 bytes (your system might produce slightly different results).

Stone's Private Page Cache

As the Stone repository monitor allocates resources to each session, it stores the information in its private page cache. The size of this cache is set by the `STN_PRIVATE_PAGE_CACHE_KB` configuration option (page 302). The default size of 2 MB is sufficient in most circumstances. If you think you might need to adjust this setting, contact GemStone Technical Support.

Diagnostics

The shared page cache monitor creates or appends to a log file, *gemStoneNamePIDpcmon.log*, in the same directory as the log for the Stone repository monitor. The *PID* portion of the name is the monitor's process id. In case of difficulty, check for this log.

The operating system kernel must be configured appropriately on each node running a shared page cache. If **startstone** or a remote login fails because the shared cache cannot be attached, check *gemStoneName.log* and *gemStoneNamePidpcmon.log* for detailed information.

The following configuration settings are checked at startup:

- ▶ The kernel shared memory resources must be enabled and sufficient to provide the page space specified by `SHR_PAGE_CACHE_SIZE_KB` plus the cache overhead.

The kernel semaphore resources must also be sufficient to provide an array of size $(\text{SHR_PAGE_CACHE_NUM_PROCS} * 2) + 1$ semaphores.

Use the **shmem** utility to test the settings (see Step 3 on page 33). If multiple Stones are being run concurrently on the same node, each Stone requires a separate set of semaphores and separate semaphore id.

- ▶ Sufficient file descriptors must be available at startup to provide one descriptor for each of the `SHR_PAGE_CACHE_NUM_PROCS` processes plus an overhead of five. Compare the value computed by `SHR_PAGE_CACHE_NUM_PROCS` configuration setting to the operating system file descriptor limit per process. Some operating systems report the descriptor limit in response to the Bash built-in command **ulimit -a**.

On operating systems that permit it, the shared page cache monitor attempts to raise the descriptor soft limit to the number required. In some cases, raising the limit may require superuser action to raise the hard limit or to reconfigure the kernel.

To Configure the Repository Extents

Configuring the repository extents involves these primary considerations:

- ▶ Providing sufficient disk space
- ▶ Minimizing I/O contention
- ▶ Providing fault tolerance

Estimating Extent Size

When you estimate the size of the repository, allow 10 to 20% for fragmentation. Also allow at least 0.5 MB of free space for each session that will be logged in simultaneously. In addition, while the application is running, overhead is needed for objects that are created or that have been dereferenced but are not yet removed from the extents. The amount of extent space required for this depends strongly on the particular application and how it is used.

If there is room on the physical disks, and the extents are not at their maximum sizes as specified using `DBF_EXTENT_SIZES`, then the extents will grow automatically when additional space is needed.

Example 1.1 Extent Size Including Working Space

Size of repository	=	10 GB
Free-Space Allowance .5 MB * 20 sessions	=	10 MB
Fragmentation Allowance 10 GB * 15%	=	1500 MB
Total with Working Space	=	11.6 GB

If the free space in extents falls below a level set by the `STN_FREE_SPACE_THRESHOLD` configuration option, the Stone takes a number of steps to avoid shutting down. For information, see “How To Recover from Disk-Full Conditions” on page 178. (The default setting for `STN_FREE_SPACE_THRESHOLD` is 1 MB; see page 293.)

For planning purposes, you should allow additional disk space for making GemStone backups and for making a copy of the repository when upgrading to a new release. A GemStone full backup may occupy 75% to 90% of the total size of the extents, depending on how much space is free in the repository at the time.

Choosing the Extent Location

You should consider the following factors when deciding where to place the extents:

- ▶ Keep extents on a spindle different from operating system swap space.
- ▶ Where possible, keep the extents and transaction logs on separate spindles.

Specify the location of each extent in the configuration file. The following example uses two raw disk partitions (your partition names will be different):

```
DBF_EXTENT_NAMES = /dev/rdisk/c1t3d0s5, /dev/rdisk/c2t2d0s6;
```

Extent disk configuration

Extents benefit from efficiency of both random access (16 KB repository pages) and sequential access. Don't optimize one by compromising the other. Sequential access is important for such operations as garbage collection and making or restoring backups. Use of RAID devices or striped file systems that cannot efficiently support both random and sequential access may reduce overall performance. Simple disk mirroring may be give better results.

Setting a Maximum Size for an Extent

You can specify a maximum size in MB for each extent through the `DBF_EXTENT_SIZES` configuration option (page 274). When the extent reaches that size, GemStone stops allocating space in it. If no size is specified, which is the default, GemStone continues to allocate space for the extent until the file system or raw partition is full.

NOTE

For best performance using raw partitions, the maximum size should be slightly smaller than the size of the partition, so that GemStone can avoid having to handle system errors. For example, for a 2 GB partition, set the size to 1995 MB.


Each size entry is for the corresponding entry in `DBF_EXTENT_NAMES` (page 274). Use a comma to mark the position of an extent for which you do not want to specify a limit. For example, the following settings are for two extents of 500 MB each in raw partitions.

```
DBF_EXTENT_NAMES = /dev/rdisk/c1t3d0s5, /dev/rdisk/c2t2d0s6;
DBF_EXTENT_SIZES = 498, 498;
```

Pregrowing Extents to a Fixed Size

Allocating disk space requires a system call that introduces run time overhead. Each time an extent is expanded (Figure 1.3), your application must incur this overhead and then initialize the added extent pages.

Figure 1.3 Growing an Extent on Demand



You can increase I/O efficiency while reducing file system fragmentation by instructing GemStone to allocate an extent to a predetermined size (called pregrowing it) at startup.

You can specify a pregrow size for each extent through the `DBF_PRE_GROW` configuration option (page 275). When this is set, the Stone repository monitor allocates the specified amount of disk space when it starts up with an extent that is smaller than the specified size. The extent files can then grow as needed up to the limit of `DBF_EXTENT_SIZES`, if that is set, or to the limits of disk space.


Pregrowing extents avoids repeated system calls to allocate and initialize additional space incrementally. This technique can be used with any number of extents, and with either raw disk partitions or extents in a file system.

The disadvantages of pregrowing extents are that it takes longer to start up GemStone, and unused disk space allocated to pregrown extents is unavailable for other purposes.

Pregrowing Extents to the Maximum Size

You may pregrow extents to the maximum sizes specified in `DBF_EXTENT_SIZES` by setting `DBF_PRE_GROW` to `True`, rather than to a list of pregrow sizes.

Pregrowing extents to the maximum size provides a simple way to reserve space on a disk for a GemStone extent. Since extents cannot be expanded beyond the maximum specified size, the system should be configured with sufficiently large extent sizes that the limit will not be reached, to avoid running out of space.

Figure 1.4 Pregrowing an Extent

Two configuration options work together to pregrow extents. `DBF_PRE_GROW` (page 275) enables the operation, and optionally sets a minimum value to which to size that extent. When `DBF_PRE_GROW` is set to `True`, the Stone repository monitor allocates the space specified by `DBF_EXTENT_SIZES` (page 274) for each extent, when it creates a new extent or starts with an extent that is smaller than the specified size. It may also be set to a list of sizes, which sets the pregrow size individually for each extent to a value that is smaller than `DBF_EXTENT_SIZES`.

For example, to pregrow extents to the maximum size of 1 GB each:

```
DBF_EXTENT_SIZES = 1000, 1000, 1000;
DBF_PRE_GROW = TRUE;
```

To pregrow extents to 500M, but allow them to later expand to 1 GB if GemStone requires additional space, and that disk space is available:

```
DBF_EXTENT_SIZES = 1000, 1000, 1000;
DBF_PRE_GROW = 500, 500, 500;
```

Allocating Data to Multiple Extents

If your application is query-intensive, you should consider dividing the repository into multiple extents and placing each extent on a separate spindle so that accesses can overlap. When GemStone schedules disk writes, it assumes that you have done so. Because several extents can be active at once, putting them on the same spindle limits the maximum update rate and causes updating transactions to have unexpected impact on the response time for queries.

The `DBF_ALLOCATION_MODE` configuration option (page 274) determines whether GemStone allocates new disk pages to multiple extents by filling each extent sequentially or by balancing the extents using a set of weights you specify. If you have placed each extent on a separate disk drive as recommended, the weighted allocation yields better performance because it distributes disk accesses.

Sequential Allocation

By default, the Stone repository monitor allocates disk resources sequentially by filling one extent to capacity before opening the next extent. (See Figure 1.5 on page 39.) For example, if a logical repository consists of three extents named A, B, and C, then all of the disk resources in A will be allocated before any disk resources from B are used, and so

forth. Sequential allocation is used when the `DBF_ALLOCATION_MODE` configuration option is set to `SEQUENTIAL`.

Weighted Allocation

For weighted allocation, you use `DBF_ALLOCATION_MODE` to specify the number of extent pages to be allocated from each extent on each allocation request. The allocations are positive integers in the range 1..40 (inclusive), with each element corresponding to an extent of `DBF_EXTENT_NAMES`. For example:

```
DBF_EXTENT_NAMES = a.dbf, b.dbf, c.dbf;  
DBF_ALLOCATION_MODE = 12, 20, 8;
```

You can think of the total weight of a repository as the sum of the weights of its extents. When the Stone allocates space from the repository, each extent contributes an allocation proportional to its weight.

NOTE

We suggest that you avoid using very small values for weights, such as "1,1,1". It's more efficient to allocate a group of pages at once, such as "10,10,10", than to allocate single pages repeatedly.

One reason for specifying weighted allocation is to share the I/O load among a repository's extents. For example, you can create three extents with equal weights, as shown in Figure 1.6 (on page 40).

Figure 1.5 Sequential Allocation

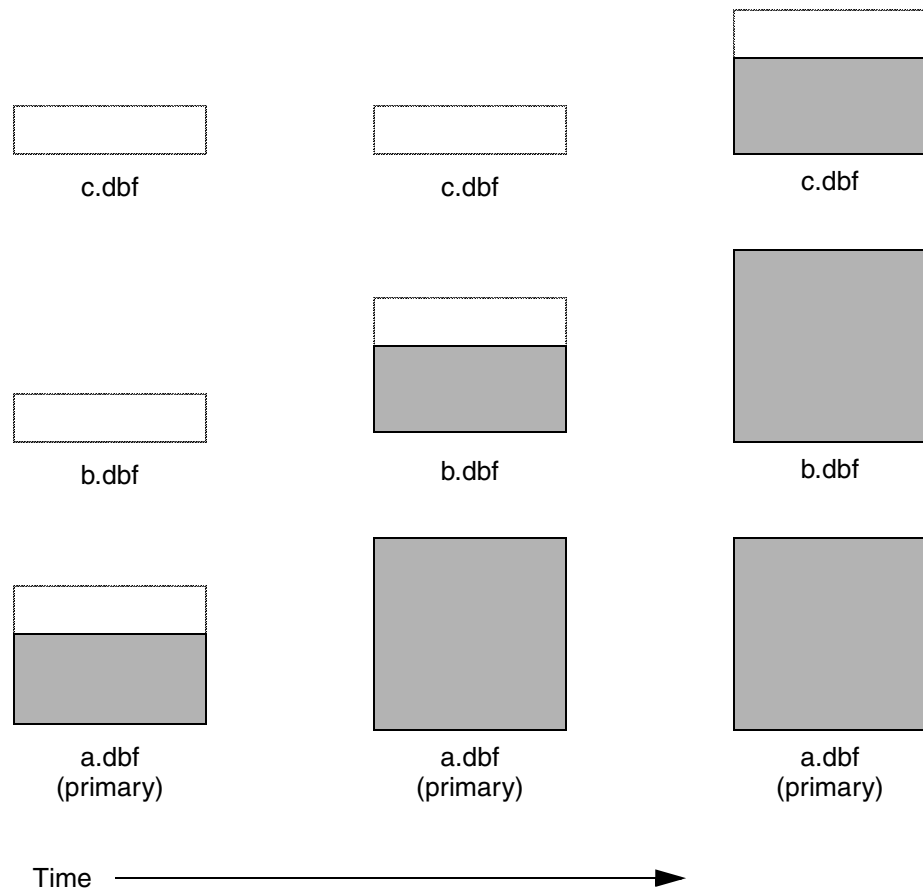



Figure 1.6 Equally Weighted Allocation


```
DBF_ALLOCATION_MODE = 10,10,10;
```



Although equal weights are most common, you can adjust the relative extent weights for other reasons, such as to favor a faster disk drive. For example, suppose we have defined three extents: A, B, and C. If we defined their weights to be 12, 20, and 8 respectively, then for every 40 disk units (pages) allocated, 12 would come from A, 20 from B, and 8 from C. Another way of stating this formula is that because B's weight is 50% of the total repository weight, 50% of all newly-allocated pages are taken from extent B. Figure 1.7 shows the result.

Figure 1.7 Proportionally Weighted Allocation

```
DBF_ALLOCATION_MODE = 12,20,8;
```



You can modify the relative extent weights by editing your GemStone configuration file and modifying the values listed for `DBF_ALLOCATION_MODE`. You can also change `DBF_ALLOCATION_MODE` to `SEQUENTIAL` without harming the system. The new values you specify take effect the next time you start the GemStone system.

Effect of Clustering on Allocation Mode

Explicit clustering of objects using instances of `ClusterBucket` that explicitly specify an `extentId` takes precedence over `DBF_ALLOCATION_MODE`. For information about clustering objects, refer to the *GemStone/S 64 Bit Programming Guide*.

Weighted Allocation for Extents Created at Run Time

Smalltalk methods for creating extents at run time (`Repository>> createExtent:` and `Repository>>createExtent:withMaxSize:`) do not provide a way to specify a weight for the newly-created extent. If your repository uses weighted allocation, the Stone repository monitor assigns the new extent a weight that is the simple average of the repository's existing extents. For instance, if the repository is composed of three extents with weights 6, 10, and 20, the default weight of a newly-created fourth extent would be 12 (36 divided by 3).

To Configure the Transaction Logs

Configuring the transaction logs involves considerations similar to those for extents:

- ▶ Choosing a logging mode
- ▶ Providing sufficient disk space
- ▶ Minimizing I/O contention
- ▶ Providing fault tolerance, through the choice of logging mode

Choosing a Logging Mode

GemStone provides two modes of transaction logging:

- ▶ *Full logging*, the default mode, provides real-time incremental backup of the repository. Deployed applications should use this mode. All transactions are logged regardless of their size, and the resulting logs can be used in restoring the repository from a GemStone backup.
- ▶ *Partial logging* is intended for use during evaluation or early stages of application development. Partial logging allows a simple operation to run unattended for an extended period and permits automatic recovery from system crashes that do not corrupt the repository. Logs created in this mode cannot be used in restoring the repository from a backup.

CAUTION

The only backups to which you can apply transaction logs are those made while the repository is in full logging mode. If you change to full logging, be sure to make a GemStone backup as soon as circumstances permit.

Changing the logging mode from full to partial logging requires special steps. See “To Change to Partial Logging” on page 189. To re-enable full transaction logging, change the configuration setting to True and restart the Stone repository monitor

For general information about the logging mode and the administrative differences, see “Logging Modes” on page 182.

Estimating the Log Size

How much disk space does your application need for transaction logs? The answer depends on several factors:

- ▶ The logging mode that you choose
- ▶ Characteristics of your transactions
- ▶ How often you archive and remove the logs

If you have configured GemStone for full transaction logging (that is, `STN_TRAN_FULL_LOGGING` is set to True), you must allow sufficient space to log all transactions until you next archive the logs.

CAUTION

If the Stone exhausts the transaction log space, users will be unable to commit transactions until space is made available.

You can estimate the space required from your transaction rate and the number of bytes modified in a typical transaction. Example 1.2 provides an estimate for an application that expects to generate 4500 transactions a day.

At any point, the method `Repository>>oldestLogFileIdForRecovery` identifies the oldest log file needed for recovery from the most recent checkpoint, if the Stone were to crash. Log files older than the most recent checkpoint (the default maximum interval is 5 minutes) are needed only if it becomes necessary to restore the repository from a backup. Although the older logs can be retrieved from archives, you may want to keep them online until the next GemStone full backup, if you have sufficient disk space.

Example 1.2 Space for Transaction Logs Under Full Logging

Average transaction rate	= 5 per minute
Duration of transaction processing	= 15 hours per day
Average transaction size	= 5 KB
Archiving interval	= Daily
Transactions between archives	
5 per minute * 60 minutes * 15 hours	= 4500
Log space (minimum)	
4500 transactions * 5 KB	= 22 MB

If GemStone is configured for partial logging, you need only provide enough space to maintain transaction logs since the last repository checkpoint. Ordinarily, two log files are sufficient: the current log and the immediately previous log. (In partial logging mode, transaction logs are used only after an unexpected shutdown to recover transactions since the last checkpoint.)

Choosing the Log Location and Size Limit

The considerations in choosing a location for transaction logs are similar to those for extents:

- ▶ Keep transaction logs on a different disk than operating system swap space.
- ▶ Where possible, keep the extents and transaction logs on separate disks — doing so reduces I/O contention while increasing fault tolerance.
- ▶ Because update-intensive applications primarily are writing to the transaction log, storing raw data in a disk partition (rather than in a file system) may yield better performance.

WARNING

Because the transaction logs are needed to recover from a system crash, do NOT place them in directories such as /tmp that are automatically cleared during power-up.

Transaction logs use sequential access exclusively, so the devices can be optimized for that access.

With raw partitions, or when in partial transaction logging mode, GemStone requires at least two log locations (directories or raw partitions) so it can switch to another when the current one is filled. In full transaction mode, logging to transaction logs on the file system, one directory may be used, in which case all transaction logs are created in that directory.

When you set the log locations in the configuration file, you should also check their size limit.

Although the size of 100 MB provided in the default configuration file is adequate in many situations, update-intensive applications should consider a larger size to limit the frequency with which logs are switched. Each switch causes a checkpoint to occur, which can impact performance.

NOTE

For best performance using raw partitions, the size setting should be slightly smaller than the size of the partition so GemStone can avoid having to handle system errors. For example, for a 2 GB partition, set it to 1998 MB.

The following example sets up a log in a 2 GB raw partition and a directory of 100 MB logs in the file system. This setup is a workable compromise when the number of raw partitions is limited. The file system logs give the administrator time to archive the primary log when it is full.

```
STN_TRAN_LOG_DIRECTORIES = /dev/rdisk/c4d0s2, /user3/tranlogs;  
STN_TRAN_LOG_SIZES = 1998, 100;
```

All of the transaction logs must reside on Stone's node.

To Configure Server Response to Gem Fatal Errors

The Stone repository monitor is configurable in its response to a fatal error detected by a Gem session process. If configured to do so, the Stone can halt and dump debug information if it receives notification from a Gem that the Gem process died with a fatal error. By stopping both the Gem and the Stone at this point, the possibility of repository corruption is minimized.

In the default mode, the Stone does not halt if a Gem encounters a fatal error. This is usually preferable for deployed production systems.

During application development, it may be helpful to know exactly what the Stone was doing when the Gem went down. It may in some cases be preferred to absolutely minimize the risk of repository corruption, at the risk of system outage. To configure the Stone to halt when a fatal gem error is encountered, change the following in the Stone's configuration file:

```
STN_HALT_ON_FATAL_ERR = TRUE;
```

To Set File Permissions for the Server

The primary consideration in setting file permissions for the Server is to protect the repository extents. All reads and writes should be done through GemStone repository executables: the executables that run the Stone, Page Servers, and Gems .

For the tightest security, you can have the extents and executables owned by a single UNIX account, using the setuid bit on the executable files, and making the extents writable only by that account. When setuid is set, the processes started from that executable are owned by the owner you specify for the file, regardless of which user actually starts them.

Alternatively, you can make the extents writable by a particular UNIX group and have all users belong to that group. This has the advantage that linked sessions that perform fileouts and other I/O operations will be done using the individual user's id instead of the single *gsadmin* account.

Using the Setuid Bit

When all extents and executables are owned and can only be written by a single UNIX account, it provides the strongest security. By setting the setuid bit, the processes started from that executable are owned by the owner you specify for the file.

Table 1 shows the recommended file settings. In this table, *gsadmin* and *gsgroup* can be any ordinary UNIX account and group (do NOT use the root account for this purpose). The person who starts the Stone must be logged in as *gsadmin* or have execute permission.

Table 1 Recommended File and Executable Permissions for the Server

Resource or Process	Filename	Protection Mode	File Owner	File Group	Process Runs As
Repository extents	(default) data/extent*.dbf	-rw-----	gsadmin	gsgroup	
Stone	sys/stoned	-r-sr-xr-x	gsadmin	gsgroup	gsadmin
AIO and FreeFrame Page Servers	sys/pgsvrmain	-r-sr-xr-x	gsadmin	gsgroup	gsadmin
Gem	sys/gem	-r-sr-xr-x	gsadmin	gsgroup	gsadmin

Ownership and permissions for the netluid executable depend on the authentication mode chosen, and are discussed in Chapter 3.

If you are logged in as root when you run the GemStone installation program, it offers to set file protections in the manner described in Table 1. To set them manually, do the following as root:

```
# cd $GEMSTONE/sys
# chown gsadmin gem pgsvr pgsvrmain stoned
# chmod u+s gem pgsvr pgsvrmain stoned
# cd $GEMSTONE/data
# chown gsadmin extent0.dbf
# chmod 600 extent0.dbf
```

The protection mode for the shared memory segment is set by the configuration parameter `SHR_PAGE_CACHE_PERMISSIONS`.

You must take similar steps to provide access for repository clients, which are presented in Chapter 2. See “To Set Ownership and Permissions for Session Processes” on page 2-61.

Alternative: Use Group Write Permission

For sites that prefer not to use the `setuid` bit, the alternative is to make the extents writable by a particular UNIX group and have all users belong to that group. That group must be the primary group of the person who starts the Stone (that is, the one listed in `/etc/passwd`). Do the following, where `gsgroup` is a group of your choice:

```
% cd $GEMSTONE/data
% chmod 660 extent0.dbf
% chgrp gsgroup extent0.dbf
```

Sites that run the linked sessions may also prefer to use this protection so that fileouts and other I/O operations that do not read or write the repository will be done using the individual user’s id instead of the single `gsadmin` account.

File Permissions for the Files and Directories

GemStone creates log files and other special files in several locations. In a multi-user environment, the protection of these resources must be such that the appropriate file can be created or updated in response to actions by several users.

<code>/opt/gemstone</code>	All users should be able to read files in the directory <code>/opt/gemstone/locks</code> on each node (or an equivalent location, as discussed on page 30). Users who will start a Stone or NetLDI process require read, write and execute access to <code>/opt/gemstone/locks</code> and (if used for logging) <code>/opt/gemstone/log</code> .
<code>system.conf</code>	The Stone must be able to write as well as read its primary configuration file. If certain configuration changes are made while the Stone is running, the Stone updates its configuration file; for example, <code>Repository>>createExtent</code> : updates the configuration file, so that subsequent restart will be correct. By default, this file is <code>\$GEMSTONE/data/system.conf</code> . The user who owns the Stone process must have write permission to the configuration file.

1.3 How To Set Up a Raw Partition

WARNING

Using raw partitions requires extreme care. Overwriting the wrong partition destroys existing information, which in certain cases can make data on the entire disk inaccessible.

The instructions in this section are incomplete intentionally. You will need to work with your system administrator to locate a partition of suitable size for your extent or transaction log. Consult the system documentation for guidance as necessary.

You can mix file system-based files and raw partitions in the same repository, and you can add a raw partition to existing extents or transaction log locations. The partition reference in `/dev` must be readable and writable by anyone using the repository, so you should give the entry in `/dev` the same protection as you would use for the corresponding type of file in the file system.

The first step is to find a partition (raw device) that is available for use. Depending on your operating system, a raw partition may have a name like `/dev/rdisk/c1t3d0s5`, `/dev/rsd2e`, or `/dev/vg03/r1vol1`. Most operating systems have a utility or administrative interface that can assist you in identifying existing partitions; some examples are **prtvtoc** and **vgdisplay**. A partition is available if all of the following are true:

- ▶ It does not contain the root (`/`) file system (on some systems, the root volume group).
- ▶ It is not on a device that contains swap space.
- ▶ Either it does not contain a file system or that file system can be left unmounted and its contents can be overwritten.
- ▶ It is not already being used for raw data.

When you select a partition, make sure that any file system tables, such as `/etc/vfstab`, do not call for it to be mounted at system boot. If necessary, unmount a file system that is currently mounted and edit the system table. Use **chmod** and **chown** to set read-write permissions and ownership of the special device file the same way you would protect a repository file in a file system. For example, set the permissions to 600, and set the owner to the GemStone administrator.

If the partition will contain the primary extent (the first or only one listed in `DBF_EXTENT_NAMES`), initialize it by using the GemStone **copydbf** utility to copy an existing repository extent to the device. The extent must not be in use when you copy it. If the partition already contains a GemStone file, first use **removedbf** to mark the partition as being empty.

Partitions for transaction logs do not need to be initialized, nor do secondary extents into which the repository will expand later.

Sample Raw Partition Setup

The following example configures GemStone to use the raw partition `/dev/rsd2d` as the repository extent.

Step 1. If the raw partition already contains a GemStone file, mark it as being empty. (The `copydbf` utility will not overwrite an existing repository file.)

```
% removedbf /dev/rsd2d
```

Step 2. Use `copydbf` to install a fresh extent on the raw partition. (If you copy an existing repository, first stop any Stone that is running on it.)

```
% copydbf $GEMSTONE/bin/extent0.dbf /dev/rsd2d
```

Step 3. As root, change the ownership and the permission of the partition special device file in `/dev` to what you ordinarily use for extents in a file system. For instance:

```
# chown gsAdmin /dev/rsd2d
# chmod 600 /dev/rsd2d
```

You should also consider restricting the execute permission for `$GEMSTONE/bin/removedbf` and `$GEMSTONE/bin/removeextent` to further protect your repository. In particular, these executable files should not have the setuid (S) bit set.

Step 4. Edit the Stone's configuration file to show where the extent is located:

```
DBF_EXTENT_NAMES = /dev/rsd2d;
```

Step 5. Use `startstone` to start the Stone repository monitor in the usual manner.

Changing Between Files and Raw Partitions

This section tells you how to change your configuration by moving existing repository extent files to raw partitions or by moving existing extents in raw partitions to files in a file system. You can make similar changes for transaction logs.

Moving an Extent to a Raw Partition

To move an extent from the file system to a raw partition, do this:

Step 1. Define the raw disk partition device. Its size should be at least 1 to 2 MB larger than the existing extent file.

Step 2. Stop the Stone repository monitor.

Step 3. Edit the repository's configuration file, substituting the device name of the partition for the file name in `DBF_EXTENT_NAMES` (page 274).

Set `DBF_EXTENT_SIZES` for this extent to be 1 to 2 MB smaller than the size of the partition.

Step 4. Use `copydbf` to copy the extent file to the raw partition. (If the partition previously contained a GemStone file, first use `removedbf` to mark it as unused.)

Step 5. Restart the Stone.

Moving an Extent to the File System

The procedure to move an extent from a raw partition to the file system is similar:

- Step 1.** Stop the Stone repository monitor.
- Step 2.** Edit the repository's configuration file, substituting the file pathname for the name of the partition in `DBF_EXTENT_NAMES`.
- Step 3.** Use `copydbf` to copy the extent to a file in a file system, then set the file permissions to the ones you ordinarily use.
- Step 4.** Restart the Stone.

Moving Transaction Logging to a Raw Partition

To switch from transaction logging in the file system to logging in a raw partition, do this:

- Step 1.** Define the raw disk partition. If you plan to copy the current transaction log to the partition, its size should be at least 1 to 2 MB larger than current log file.
- Step 2.** Stop the Stone repository monitor.
- Step 3.** Edit the repository's configuration file, substituting the device name of the partition for the directory name in `STN_TRAN_LOG_DIRECTORIES` (page 304). Make sure that `STN_TRAN_LOG_SIZES` for this location is 1 to 2 MB smaller than the size of the partition.
- Step 4.** Use `copydbf` to copy the current transaction log file to the raw partition. (If the partition previously contained a GemStone file, first use `removedbf` to mark it as unused.)

You can determine the current log from the last message "Creating a new transaction log" in the Stone's log. If you don't copy the current transaction log, the Stone will open a new one with the next sequential fileId, but it may be opened in another location specified by `STN_TRAN_LOG_DIRECTORIES`.

- Step 5.** Restart the Stone.

Moving Transaction Logging to the File System

The procedure to move transaction logging from a raw partition to the file system is similar:

- Step 1.** Stop the Stone repository monitor.
- Step 2.** Edit the repository's configuration file, substituting a directory pathname for the name of the partition in `STN_TRAN_LOG_DIRECTORIES`.
- Step 3.** Use `copydbf` to copy the current transaction log to a file in the specified directory. The `copydbf` utility will generate a file name like `tranlognnn.dbf`, where `nnn` is the internal fileId of that log.
- Step 4.** Restart the Stone.

1.4 How To Access the Server Configuration at Run Time

GemStone provides several methods in class `System` that let you examine, and in certain cases modify, the configuration parameters at run time from Smalltalk.

To Access Current Settings at Run Time

Class methods in category Configuration File Access let you examine the system- Stone configuration. The following access methods all provide similar server information:

`stoneConfigurationReport`

Returns a `SymbolDictionary` whose keys are the names of configuration file parameters, and whose values are the current settings of those parameters in the repository monitor process.

`configurationAt: aName`

Returns the value of the specified configuration parameter, giving preference to the current session process if the parameter applies to a Gem.

`stoneConfigurationAt: aName`

Returns the value of the specified configuration parameter from the Stone process, or returns `nil` if that parameter is not applicable to a Stone.

(The corresponding methods for accessing a session configuration are described on page 64.)

Here is a partial example of the Stone configuration report:

```
topaz 1> printit
System stoneConfigurationReport asReportString
%
#'StnEpochGcEnabled'    false
#'StnPageMgrRemoveMinPages'  40
#'STN TRAN LOG SIZES'    100
#'StnTranLogDebugLevel'  0
...
```

Keys in mixed capitals and lowercase, such as `StnEpochGcEnabled`, are internal run-time parameters.

To Change Settings at Run Time

The class method `System class>>configurationAt: aName put: aValue` lets you change the value of the internal run-time parameters in Table 1, if you have the appropriate privileges.

Parameters with names in all uppercase are read-only parameters; for parameters that can be changed at runtime, the name is in mixed case.

CAUTION

Avoid changing configuration parameters unless there is a clear reason for doing so. Incorrect settings can have serious adverse effects on performance. For additional guidance about run-time changes to specific parameters, see Appendix A, "GemStone Configuration Options."

Table 1 Server Configuration Parameters Changeable at Run Time

Configuration File Option	Internal Parameter
SHR_SPIN_LOCK_COUNT	#SpinLockCount ^a
STN_ADMIN_GC_SESSION_ENABLED	#StnAdminGcSessionEnabled ^b
STN_CHECKPOINT_INTERVAL	#StnCheckpointInterval ^a
STN_COMMIT_QUEUE_THRESHOLD	#StnCommitQueueThreshold ^a
STN_CR_BACKLOG_THRESHOLD	#StnCrBacklogThreshold ^a
STN_DISABLE_LOGIN_FAILURE_LIMIT	#StnDisableLoginFailureLimit ^c
STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT	#StnDisableLoginFailureTimeLimit ^c
STN_DISKFULL_TERMINATION_INTERVAL	#StnDiskFullTerminationInterval ^a
STN_EPOCH_GC_ENABLED	#StnEpochGcEnabled ^b
STN_FREE_SPACE_THRESHOLD	#StnFreeSpaceThreshold
STN_GEM_ABORT_TIMEOUT	#StnGemAbortTimeout ^a
STN_GEM_LOSTOT_TIMEOUT	#StnGemLostOfTimeout ^a
STN_GEM_TIMEOUT	#StnGemTimeout ^a
STN_HALT_ON_FATAL_ERR	#StnHaltOnFatalErr ^a
STN_LOG_LOGIN_FAILURE_LIMIT	#StnLogLoginFailureLimit ^c
STN_LOG_LOGIN_FAILURE_TIME_LIMIT	#StnLogLoginFailureTimeLimit ^c
STN_LOOP_NO_WORK_THRESHOLD	#StnLoopNoWorkThreshold ^a
STN_MAX_AIO_RATE	#StnMntMaxAioRate ^a
STN_MAX_LOGIN_LOCK_SPIN_COUNT	#StnMaxLoginLockSpinCount
STN_MAX_VOTING_SESSIONS	#StnMaxVotingSessions ^a
STN_NUM_GC_RECLAIM_SESSIONS	#StnNumGcReclaimSessions ^b
STN_OBJ_LOCK_TIMEOUT	#StnObjLockTimeout
STN_PAGE_MGR_COMPRESSION_ENABLED	#StnPageMgrCompressionEnabled ^a
STN_PAGE_MGR_MAX_WAIT_TIME	#StnPageMgrMaxWaitTime
STN_PAGE_MGR_PRINT_TIMEOUT_THRESHOLD	#StnPageMgrPrintTimeoutThreshold ^a
STN_PAGE_MGR_REMOVE_MAX_PAGES	#StnPageMgrRemoveMaxPages ^a
STN_PAGE_MGR_REMOVE_MIN_PAGES	#StnPageMgrRemoveMinPages ^a
STN_REMOTE_CACHE_PGSRV_TIMEOUT	#StnRemoteCachePgsvrTimeout ^d
STN_REMOTE_CACHE_TIMEOUT	#StnRemoteCacheTimeout ^a
STN_SHR_TARGET_PERCENT_DIRTY	#ShrPcTargetPercentDirty ^a

Table 1 Server Configuration Parameters Changeable at Run Time (Continued)

Configuration File Option	Internal Parameter
STN_SIGNAL_ABORT_CR_BACKLOG	#StnSignalAbortCrBacklog ^b
STN_SYMBOL_GC_ENABLED	#StnSymbolGcEnabled ^b
STN_TRAN_LOG_DEBUG_LEVEL	#StnTranLogDebugLevel ^d
STN_TRAN_LOG_LIMIT	#StnTranLogLimit ^a
STN_TRAN_Q_TO_RUN_Q_THRESHOLD	#StnTranQToRunQThreshold ^a
(none)	#StnLoginsSuspended ^d

^a Can be changed only by SystemUser.

^b Requires GarbageCollection privilege.

^c Requires OtherPassword privilege.

^d Requires SystemControl privilege.

The following example first obtains the value of #StnAdminGcSessionEnabled. This value can be changed at run time by a user with GarbageCollection privilege:

```
topaz 1> printit
System configurationAt: #StnAdminGcSessionEnabled
%
true
topaz 1> printit
System configurationAt: #StnAdminGcSessionEnabled put: false
%
false
```

For more information about these methods, see the comments in the image.

1.5 How To Tune Server Performance

There are a number of configuration options by which you can tune the GemStone server. These options can help make better use of the shared page cache, reduce swapping, and control disk activity caused by repository checkpoints.

Tuning the Shared Page Cache

Two configuration options can help you tailor the shared page cache to the needs of your application: SHR_PAGE_CACHE_SIZE_KB and SHR_SPIN_LOCK_COUNT.

You may also want to consider object clustering within Smalltalk as a means of increasing cache efficiency.

Adjusting the Cache Size

Adjust the `SHR_PAGE_CACHE_SIZE_KB` configuration option (page 287) according to the total number of objects in the repository and the number accessed at one time. For proper performance, the entire object table should be in shared memory.

In general, the more of your repository you can hold in your cache, the better your performance will be, provided you have enough memory to avoid swapping.

You should review the configuration recommendations given earlier (“Estimating the Size of the Shared Page Cache” on page 32) in light of your application’s design and usage patterns. Estimates of the number of objects queried or updated are particularly useful in tuning the cache.

You can use the shared page cache statistics for a running application to monitor the load on the cache. In particular, the statistics `FreeFrameCount` and `FramesFromFindFree` may be useful, as well as `FramesFromFreeList`.

Matching Spin Lock Limit to Number of Processors

The `SHR_SPIN_LOCK_COUNT` configuration option (page 287) specifies the number of times a process should attempt to obtain a lock in the shared page cache using the spin lock mechanism before resorting to setting a semaphore and sleeping. We recommend you leave `SHR_SPIN_LOCK_COUNT` set to `-1` (the default), which causes GemStone to determine whether multiple processors are installed and set the parameter accordingly.

Clustering Objects That Are Accessed Together

Appropriate clustering of objects by the application can allow a smaller shared page cache by reducing the number of data pages in use at once. For general information about clustering objects, see the *GemStone/S 64 Bit Programming Guide*.

Reducing Swapping

Be careful not to make the shared page cache so large that it forces swapping. You should ensure that your system has sufficient RAM to hold the configured shared page cache, with extra space for the other memory requirements.

Controlling Checkpoint Frequency

At a checkpoint, any committed changes that have not yet been written to the extents are written out, and the repository updated to a new consistent committed state. Between checkpoints, each commit is written to the transaction logs, and updates that are written to the extents are not part of the repository’s consistent committed state.

If checkpoints interferes with other GemStone activity, you may want to adjust their frequency.

- ▶ In full transaction logging mode, most checkpoints are determined by the `STN_CHECKPOINT_INTERVAL` configuration option, which by default is five minutes (see page 289). A few Smalltalk methods, such as `Repository>>fullBackupTo: ,` force a checkpoint at the time they are invoked. A checkpoint also is performed each time the Stone begins a new transaction log.

- ▶ In partial logging mode, checkpoints also are triggered by any transaction that is larger than `STN_TRAN_LOG_LIMIT`, which sets the size of the largest entry that is to be appended to the transaction log (see page 304). The default limit is 1 MB of log space. If checkpoints are too frequent in partial logging mode, it may help to raise this limit. Conversely, during bulk loading of data with large transactions, it may be desirable to lower this limit to avoid creating large log files.

A checkpoint also occurs each time the Stone repository monitor is shut down gracefully, as by invoking `stopstone` or `System class>>shutDown`. This checkpoint permits the Stone to restart without having to recover from transaction logs. It also permits extent files to be copied in a consistent state.

While less frequent checkpoints may improve performance in some cases, they may extend the time required to recover after an unexpected shutdown. In addition, since checkpoints are important in the recycling of repository space, less frequent checkpoints can mean more demand on free space (extent space) in the repository.

Suspending Checkpoints

You can call the method `System class>>suspendCheckpointsForMinutes:` to suspend checkpoints for a given number of minutes, or until `System class>>resumeCheckpoints` is executed. (To execute these Smalltalk methods, you must have the required GemStone privilege, as described in Chapter 6, “User Accounts and Security”.)

Generally, this approach is used only to allow online extent backups to complete. For details on how to suspend and resume checkpoints, see “How To Make an Extent Snapshot Backup” on page 193.

Tuning Page Server Behavior

GemStone uses page servers for three purposes:

- ▶ To write dirty pages to disk.
- ▶ To transfer pages from the Stone host to the shared page cache host, if different.
- ▶ To add free frames to the free frame list, from which a Gem can take as needed.

Page servers referred to as *AIO page servers* perform all three functions. By default, at least one AIO page server is running at all times, although larger systems are likely to need additional page servers.

In addition, by default, one *free frame page server* is running. Free frame page servers are dedicated only to the third task listed above: adding free frames to the free list. In some cases, increasing the number of free frame page servers can improve overall system performance. For example, if Gems are performing many operations requiring writing pages to disk, the AIO page server may have to spend all its time writing pages, never getting a chance to add free frames to the free list. Alternatively, if Gems are performing operations that require only reading, the AIO page server will see no dirty frames in the cache – the signal that prompts it to take action. In that case, it may sleep for a second, even though free frames are present in the cache and need to be added to the free list.

To Add AIO Page Servers

By default the Stone spawns a single page server process on its local node to perform asynchronous I/O (AIO) between the shared page cache and the extents. This page server ordinarily is the process that updates extents on the local node during a checkpoint. (In some cases, the Stone may use additional page servers temporarily during startup to pre-grow multiple extents.)

If your configuration has multiple extents on separate disk spindles, you should generally increase the number of AIO page servers. You can do this by changing the `STN_NUM_LOCAL_AIO_SERVERS` configuration option (page 300).

For multiple page servers to improve performance, they must be able to execute at the same time and write to disk at the same time. If you have only one CPU, or your extents are on a single disk spindle, multiple AIO page servers will not be able to write pages out faster than a single page server.

Free Frame Page Servers

A Gem can get free frames either from the free list (the quick way), or, if sufficient free frames are not listed, by scanning the shared page cache for a free frame instead. (What constitutes sufficient free frames is determined by the configuration parameter `GEM_FREE_FRAME_LIMIT`; for details, see page 276.)

To assist the AIO page servers in adding frames back to the free list, the stone spawns one or more free frame page servers.

By default, when you start the Stone, it spawns the same number of free frame page server processes as AIO page server processes. This is recommended so that the distribution of free pages and used pages is balanced over the repository.

Process Free Frame Caches

There is a communication overhead involved in getting free frames from the free frame list for scanning. To optimize this, you can configure the Gems and their remote page servers to add or remove multiple free frames from a free frame cache to the free frame list in a single operation.

When using the free frame cache, the Gem or remote page server process removes enough frames from the free list to refill the cache in a single operation. When adding frames to the free list, the process does not add them until the cache is full.

You can control the size of the Gem and remote page server free frame caches by setting the configuration parameters `GEM_FREE_FRAME_CACHE_SIZE` and `GEM_PGSRV_FREE_FRAME_CACHE_SIZE`, respectively.

The default behavior depends on the size of the shared page cache; if the shared page cache is 100MB or larger, a page server free frame cache size of 10 is used, so ten free frames are acquired in one operation when the cache is empty. For shared page cache sizes less than 100MB, the Gem or remote page server acquires the frames one at a time.

1.6 How To Run a Second Repository

You can run more than one repository on a single node—for example, separate production and development repositories. There are several points to keep in mind:

- ▶ Each repository requires its own Stone repository monitor process, extent files, transaction logs, and configuration file. Each Stone will start its own shared page cache monitor and a set of other processes, as described on page 22.
- ▶ Multiple Stones that are running the same version of GemStone can share a single installation directory, provided that you create separate repository extents, transaction logs, and configuration files. If performance is a concern, the first step should be to isolate each Stone's data directory and the system swap space on separate drives. Then, review the discussion "Recommendations About Disk Usage" on page 26.
- ▶ You must give each Stone a unique name at startup. That name is used to identify the server and to maintain separate log files. Users will connect to the repository by specifying the Stone's name.
- ▶ A single NetLDI serves all Stones and Gem session processes on a given node, provided both Stones are running the same version. If you are running two different versions of GemStone, you will need two NetLDIs.

Configuring Gem Session Processes

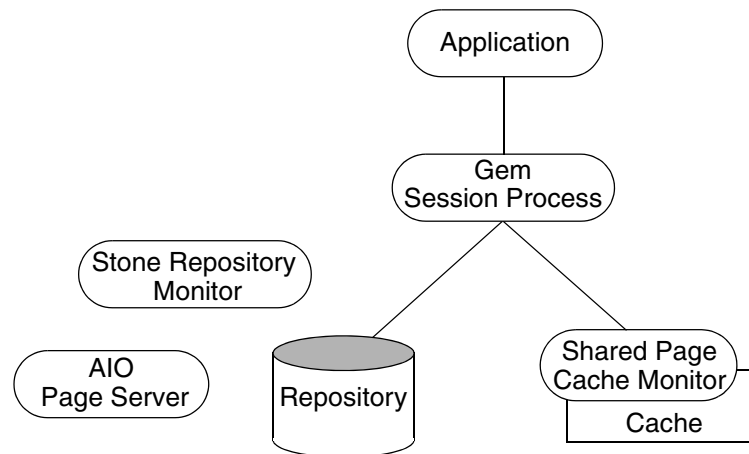
This chapter tells how to configure the GemStone/S 64 Bit session processes for your application. For additional information about running session processes on a node remote from the Stone repository monitor, refer also to Chapter 3.

2.1 Overview

As shown in Figure 2.1, a GemStone session involves the following components in a client-server relationship:

- ▶ The user application
- ▶ A session manager process (Gem), which acts as a server for a particular application
- ▶ The Stone repository monitor
- ▶ The shared page cache monitor and cache
- ▶ The Stone's AIO page server
- ▶ The repository itself

Figure 2.1 GemStone Session Elements



The Gem session process provides the bulk of the repository capabilities as seen by the application. From the viewpoint of the application, the Gem *is* the object server:

- ▶ It logs in to the repository through the Stone repository monitor, and it obtains object locks, free object identifiers, and free pages from the repository monitor.
- ▶ It presents the application with a consistent view of the repository during a transaction and tracks which objects the application accesses.
- ▶ It executes Smalltalk methods within the repository.
- ▶ It reads the repository as the application accesses objects, and (with the help of the AIO page server) it updates the repository when the application successfully commits a transaction.

Linked and RPC Applications

The Gem session process can be run as a separate process (as in Figure 2.1) or integrated with the application into a single process, in which case the application is called a *linked* application.

When the Gem runs as a separate process, it responds to Remote Procedure Calls (RPCs) from the application; this is called an *RPC* application. Applications that use a separate Gem process start that process automatically (from the user's viewpoint) while logging in to the repository.

Either type of application can be used on a single node or across a network.

An application can have only one linked login, since only one Gem can be integrated with the process. It may have multiple RCP logins, or one linked and multiple RPC logins.

GemStone provides both linked and RPC versions of topaz for repository administration, and shared client libraries that allow linked and RPC. The linked version allows both linked and RPC logins, but the RPC version allows only RPC logins.

C programmers should always use an RPC version during development and debugging to protect Gem data structures from possible corruption.

The Session Configuration File

At start-up time, each Gem session process looks for a configuration file, which by default is the same system-wide configuration file sought by the repository monitor when it starts. However, there are three important differences:

- ▶ The session configuration file is optional. If one is not found, the session process uses system defaults.
- ▶ All session processes read those configuration options that begin with “GEM_” and the few that are used by both Stones and Gems (such as DUMP_OPTIONS and LOG_WARNINGS). Other settings that the Gem needs are obtained from the repository monitor by network protocol and are the same for all sessions logged in to that Stone.
- ▶ The first session process on a node remote from the Stone and extents uses the shared page cache configuration options (SHR_), which determine the configuration of the cache on that node.

It can be useful for certain sessions to use a different configuration. Appendix A, “GemStone Configuration Options,” tells how to specify an alternate configuration file and how to use supplementary files to adjust the system-wide configuration for a specific session process. Appendix A describes each of the configuration options.

2.2 How To Configure Gem Session Processes

To configuring a Gem session process, you perform the following steps:

1. Gather application specifics about the number of sessions that will be logged in to the repository simultaneously from this node.
2. Plan the operating system resources that will be needed: memory and swap (paging) space.
3. Set the Gem configuration options.

If this node is remote from the repository monitor, enable a local GemStone shared page cache. Gem session processes running on a server node always use the Stone’s shared page cache.

4. Set GemStone file permissions to allow session processes access while providing adequate security.

Gathering Application Information

System resources needed for session processes primarily depend on the number of sessions that will be logged in to a particular repository from this node. Remember that in some applications each user can have more than one session logged in.

Planning Operating System Resources

GemStone session processes need adequate memory and swap space to run efficiently. In addition, kernel parameters can limit the number of sessions that can connect to the shared page cache.

Estimating Memory Needs

Two factors determine the memory needs for session processes:

- ▶ The size of the shared page cache on a node remote from the Stone and extents will depend on the configuration of the Gem that starts the cache. (There is only one cache on each node for a particular repository; session processes running on the server node attach to the Stone repository monitor's cache.)
- ▶ The amount of memory required by a Gem session is dependent on how it is configured, as determined by the system requirements. To avoid out-of-memory conditions, Gems must be configured with an adequate temporary object cache.

Assuming the default of 50 MB for the Gem's temporary object cache, the first Gem session process on a node ordinarily requires about 80 MB of memory, of which 5 MB is for code that can be shared by other session processes. Each additional session process requires about 65 MB. The requirement is the same for Gems linked with an application. If you tune the cache size for Gems, add any increase to the amount given here.

In addition to the memory needs for session processes, you must also allocate memory for object server processes. For details, see "Planning Operating System Resources" on page 28.

For Gem session processes running on machines that are remote from the object server, there are additional memory needs on the server. For information about this, see "Estimating Memory Needs" on page 28.

Estimating Swap Space Needs

Swap (paging) space on machines remote from the Stones should follow the same general guidelines given for servers on page 28. In determining the additional swap space needed for GemStone session processes, use the memory requirements derived in the preceding section ("Estimating Memory Needs"), including space for the number of sessions you expect. The resulting figures will approximate the client's needs, and are in addition to the swap requirement for the object server and non-GemStone processes.

Estimating File Descriptor Needs

When a Gem session process starts, it attempts to raise the file descriptor limit from the default (soft) limit to the hard limit set by the operating system. GemBuilder applications and page servers do the same. Gem session processes use file descriptors this way:

- 7 for stdin, stdout, stderr, and internal communication
- 1 for each connection between the Gem and an RPC application
- 1 for each local extent within a file system
- 2 for each local extent that is a raw partition
- 1 for each extent on a remote node

GemBuilder applications that start a large number of RPC Gems need a correspondingly large number of file descriptors.

You can override the default behavior of raising the file descriptor limit to the hard limit by setting the `GEMSTONE_MAX_FD` environment variable to a positive integer. A lower limit may be desirable in some cases to reduce the amount of virtual memory used by the process. A value of 0 disables attempts to raise the default limit.

The value of `GEMSTONE_MAX_FD` in the environment of a NetLDI (Network Long Distance Information) server is passed to its child processes.

Reviewing Kernel Tunable Parameters

The kernel parameter of primary relevance to GemStone session processes is the maximum number of semaphores per semaphore id (typically `semms1` or a similar name, although it is not tunable under all operating systems). This parameter limits the number of sessions than can connect to the shared page cache, because each session uses two semaphores.

How you determine the existing limits depends on your operating system. If the information is not readily available, proceed with setup. A later step shows how to verify that the limits are adequate for the GemStone configuration you set up.

To Set Ownership and Permissions for Session Processes

The primary consideration in setting file ownership permissions for client access is to make sure the Gem session process can read and write both the extents and the shared page cache. You must also take into account the following factors:

- ▶ Is the Gem session process linked or RPC?
- ▶ Does the Gem session process runs on the server or on a node remote from the Stone.?
- ▶ Does the server uses `setuid` bit and protection mode 600 for the extents (as recommended on page 45), or does it use the alternative of group write permission?

Extents

The extents may be protected as read-write only by their owner (protection 600) if you use the `setuid` (S) bit for repository executables as recommended on page 45. Otherwise, the extents must be writable by a group to which the GemStone users belong (protection 660).

Shared Page Cache

The shared memory and semaphore resources used by GemStone are created and owned by the user account under which the Stone repository monitor is running and have the same group membership. Access for the shared page cache is set by the configuration parameter `SHR_PAGE_CACHE_PERMISSIONS`; by default, it is read-write for the owner and read for the group (the equivalent of file protection 640). You can inspect the cache ownership and permissions by using the `ipcs` command. The configuration parameter is discussed on page 287.

For a session to log in using a shared page cache, the Operating System user account of the linked application or Gem session process must either be the same as that of the Stone (such as the `gsadmin` account) or be one that belongs to the same group as the Stone. The same requirement applies to page server processes, which are discussed in Chapter 3, “Connecting Distributed Systems.”

If the `setuid` bit is set on repository executables as recommended in Table 1 on page 45, the Stone process and shared page cache will belong to the owner you specify for those files (such as `gsadmin`).

To Set Access for Linked Applications

For linked applications *on the server*, we recommend you try using the `setuid` bit on the application's executable file. Have the file owned by `gsadmin` as it is defined on page 45. This works well for `topaz -l`. The `installgs` script offers to set the file ownership and permissions for you. To do it manually, do this while logged in as root:

```
# cd $GEMSTONE/bin
# chmod u+s topaz
# chown gsadmin topaz
```

You may prefer not to use the `setuid` bit with linked applications that do not distinguish between real and effective user IDs. GemStone's Topaz executable performs repository reads and writes as the *effective* user (the account that owns the executable's file), but performs other reads and writes as the *real* user (the one who invoked it). Linked applications that do not make this distinction, such as a third-party Smalltalk used with GemBuilder, are likely to perform *all* I/O as the effective user, or `gsadmin`. If this result is unsatisfactory, remove the `S` bit on that executable and add group write permission to the extents.

To Set Access for All Other Applications

All RPC applications require a GemStone NetLDI service to start a separate Gem session process and/or a page server. For these sessions, we recommend that the Gem session process and page server always be owned by (run as) the `gsadmin` account. That arrangement ensures that the Gem will be able to read and write both the extents and the shared page cache. The ownership and protection of the application executables themselves is not a factor.

To Set Access to Other Files

GemStone creates log files and other special files for session processes in several locations, which are described below. In a multi-user environment, the protection of these resources must be such that the appropriate file can be created or updated in response to actions by several users.

<code>\$HOME</code>	GemStone ordinarily creates log files for spawned processes (such as RPC Gem session processes and page servers) in the home directory of the user or the NetLDI captive account. In situations where the home directory cannot be writable, the environment variable <code>GEMSTONE_NRS_ALL</code> can be used to specify an alternative location; see "To Set a Default NRS" on page 78.
<code>/opt/gemstone</code>	<p>All users should have read/write/execute access to the directories <code>/opt/gemstone/log</code> and <code>/opt/gemstone/locks</code> on each host (or an equivalent location, as discussed on page 30).</p> <p>By default, NetLDI (Network Long Distance Information) and hot standby processes create log files in the <code>log</code> directory.</p> <p>Repository-wide lock files are created in the <code>locks</code> directory, and other processes use this location to look up access information for them.</p>

To Configure for Remote Shared Page Caches

The configuration file used by Gem sessions will also be used, on remote nodes, to configure the associated remote shared page cache monitor.

The Stone always creates a shared page cache monitor and cache on its own node, based on parameters in the Stone's configuration file. This cache is always used by Gems that run on this node.

On other nodes, when the first Gem session process logs in, the Stone starts a shared page cache monitor and cache on that node. Parameters in the configuration file used by the first Gem that logs in determine the size of the cache and the number of processes that can attach to it (`SHR_PAGE_CACHE_SIZE_KB` and `SHR_PAGE_CACHE_NUM_PROCS`, respectively). All subsequent sessions that log in from that remote node will use the same cache.

2.3 Set the Gem Configuration Options

There are a number of configuration options for Gem session processes.

Configure Temporary Object Space

Gems use temporary object memory for a number of purposes, described in detail in Chapter 11. It is important to provide sufficient temporary object space. If temporary object memory is exhausted, the Gem can encounter an out-of-memory condition and terminate.

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 need to use a larger value.

The configured value must be large enough to accommodate the memory needs of any gem using that configuration. The amount of temporary object memory required may be different for different application gems, depending on the specific tasks of the gem; you may wish to set up special configuration files for application sessions that have a particularly high demand on memory.

The full amount of `GEM_TEMPOBJ_CACHE_SIZE` is not allocated on gem login, but it is reserved and will impact memory usage per user.

You will probably need to experiment somewhat before you determine the optimum size of the temporary object space.

Memory management is discussed in greater detail in Chapter 11, "Managing Memory", on page 227.

Configure SSL for remote Gem sessions

During the RPC login process, a Secure Socket Layer (SSL) socket is used to establish the login. After that, for gem processes on the same node as their client application, communication reverts to a normal socket connection. For Gem processes that are running on a different node than their client, communication continues to use SSL.

This results in a slightly slower connection, due to the overhead of encrypting and decrypting data. To configure the gem to use normal socket communication for remote connections, set `GEM_RPC_USE_SSL` to false.

2.4 How To Access the Configuration at Run Time

GemStone provides several methods in class `System` that let you examine, and in certain cases modify, the session configuration parameters at run time.

To Access Current Settings at Run Time

`System` class methods let you examine the configuration of your current Gem session process. There are three access methods for session processes:

`gemConfigurationReport`

Returns a `SymbolDictionary` whose keys are the names of configuration file parameters, and whose values are the current settings of those parameters in the current session's Gem process.

`gemConfigurationAt: aName`

Returns the value of the specified configuration parameter from the current session, or returns `nil` if that parameter is not applicable to a session process.

`configurationAt: aName`

Returns the value of the specified configuration parameter, giving preference to the current session process if the parameter applies to a Gem.

To Change Settings at Run Time

The class method `System class >>configurationAt: aName put: aValue` lets you change the value of the internal run-time parameters in Table 1, provided you have the appropriate privileges.

Parameters read from the configuration file at process startup are in all uppercase, and are read-only. Parameters that can be changed at runtime have similar, but not necessarily identical, names that are in mixed case without underscores.

CAUTION

Do not change configuration parameters unless there is a clear reason for doing so. Incorrect settings can have serious adverse effects on GemStone performance. Appendix A provides additional guidance about run-time changes to specific parameters.

Table 1 Session Configuration Parameters Changeable at Run Time

Configuration File Option	Internal Parameter
GEM_ABORT_MAX_CRIS	#GemAbortMaxCrs
(none)	#GemConvertArrayBuilder
(none)	#GemDropCommittedExportedObjs
(none)	#GemExceptionSignalCapturesStack
GEM_FREE_FRAME_LIMIT	#GemFreeFrameLimit
GEM_FREE_PAGEIDS_CACHE	#GemFreePageIdsCache
GEM_HALT_ON_ERROR	#GemHaltOnError
GEM_KEEP_MIN_SOFTREFS	#GemKeepMinSoftRefs
GEM_NATIVE_CODE_ENABLED	#GemNativeCodeEnabled
GEM_PGSRV_COMPRESS_PAGE_TRANSFERS	#GemPgsvrCompressPageTransfers
GEM_PGSRV_UPDATE_CACHE_ON_READ	#GemPgsvrUpdateCacheOnRead
GEM_REPOSITORY_IN_MEMORY	#GemRepositoryInMemory
GEM_SOFTREF_CLEANUP_PERCENT_MEM	#GemSoftRefCleanupPercentMem
GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE	#GemPomGenPruneOnVote
GEM_TEMPOBJ_POMGEN_SCAVENGE_INTERVAL	#GemTempObjPomgenScavengeInterval

The following example changes the value of the configuration option #GemFreeFrameLimit:

```
topaz 1> printit
System configurationAt: #GemFreeFrameLimit put: 4000
%
4000
```

For more information about the parameters that can be changed at run time, see Appendix A, "GemStone Configuration Options."

2.5 Tuning Gem Performance

Private Page Cache

The configuration option `GEM_PRIVATE_PAGE_CACHE_KB` sets the size (in KB) of the Gem's private page cache. The default value of this option is 1000; in most cases, this value is acceptable, and you do not need to do any further tuning.

When monitoring the Gem, if the value of the statistic `LocalCacheOverflowCount` is non-zero, you may want to increase the setting for `GEM_PRIVATE_PAGE_CACHE_KB` as needed.

Native Code

By default, generation of native code for Smalltalk methods is enabled. This is configured using the configuration parameter `GEM_NATIVE_CODE_ENABLED`. When native code is disabled, execution is interpreted; behavior will be identical but somewhat slower.

Native code generation can be disabled using the runtime parameter `#GemNativeCodeEnabled`, but it cannot then be re-enabled for that session's lifetime.

If any breakpoints are set in any methods, native code is disabled. Native code is re-enabled when all breakpoints have been removed.

To determine if native code is in use by the currently executing session, execute:

```
GsProcess usingNativeCode
```

For more information on this parameter, see page 278.

Under some configurations on x86, in particular with Darwin, the 32-bit offset limit may be exceeded in some cases with a very large temporary object cache. If this occurs, native code is disabled.

Note that the Foreign Function Interface (FFI) feature has additional limitations when native code is disabled. FFI is discussed in the *Programming Guide for GemStone/S 64 Bit*.

2.6 How To Install a Custom Gem

The *GemBuilder for C* manual explains how to create a custom Gem session executable containing your own C functions to be called from Smalltalk. One way to make this custom Gem available to all users is to perform the following steps as system administrator:

Step 1. Copy the shell script `gemnetobject` from `$GEMSTONE/sys` to your working directory. This shell script is used to start Gem session processes under the UNIX shell. You will modify this script to start your custom Gem executable instead of the standard one.

Step 2. In your copy of `gemnetobject`, find the section labeled `User-definable symbols`. In that section, replace `gem` in the line

```
gemname="gem"
```

with the name of the new Gem executable. For example:

```
gemname="MyGem"
```

Step 3. Rename your modified copy of the shell script `gemnetobject` so that it has a distinct filename. For example:

```
% mv gemnetobject MyGemnetObject
```

Step 4. Copy the new shell script to `$GEMSTONE/sys`. Make sure that all GemStone users have read and execute (**r-x**) permission for the script. For example:

```
-r-xr-xr-x 1 root 912 Feb 24 20:22 MyGemnetObject
```

If necessary, change the permissions:

```
% chmod 555 MyGemnetObject
```

Step 5. Add an entry for the new shell script to the services database, `$GEMSTONE/sys/services.dat`. A NetLDI checks that file to translate the name of a service to a command it can execute. For example:

```
MyGemnetObject $GEMSTONE/sys/MyGemnetObject
```

Step 6. Copy the new Gem executable to the GemStone system directory. For example:

```
% cp MyGem $GEMSTONE/sys
```

Step 7. Make sure that all GemStone users have read and execute (**r-x**) permission for the new Gem executable.

The custom Gem executable is now available for shared use.

Connecting Distributed Systems

This chapter tells how to set up GemStone/S 64 Bit in a distributed environment:

- ▶ *Overview* (page 69) – An introduction to the GemStone processes and network objects that facilitate distributed GemStone systems.
- ▶ *How To Arrange Network Security* (page 75) – Three ways to provide access to GemStone processes on other nodes.
- ▶ *How To Use Network Resource Strings* (page 78) – How to specify where distributed GemStone resources are located.
- ▶ *How To Set Up a Remote Session* (page 79) – Step-by-step examples for setting up typical distributed client-server configurations. It also contains troubleshooting tips.

3.1 Overview

A properly configured network system is nearly transparent to GemStone users, but it requires additional steps by the system administrator. Users must be given access to all the workstations that will run their GemStone processes. Pointers to network services must be set up, and file and process specifications must include the node name in addition to the file name and path. Because processes are running on different nodes, the log files are spread throughout the network, and troubleshooting may become more complicated.

The nodes in your system can be any combination of GemStone-supported platforms, as long as they are connected by means of TCP/IP.

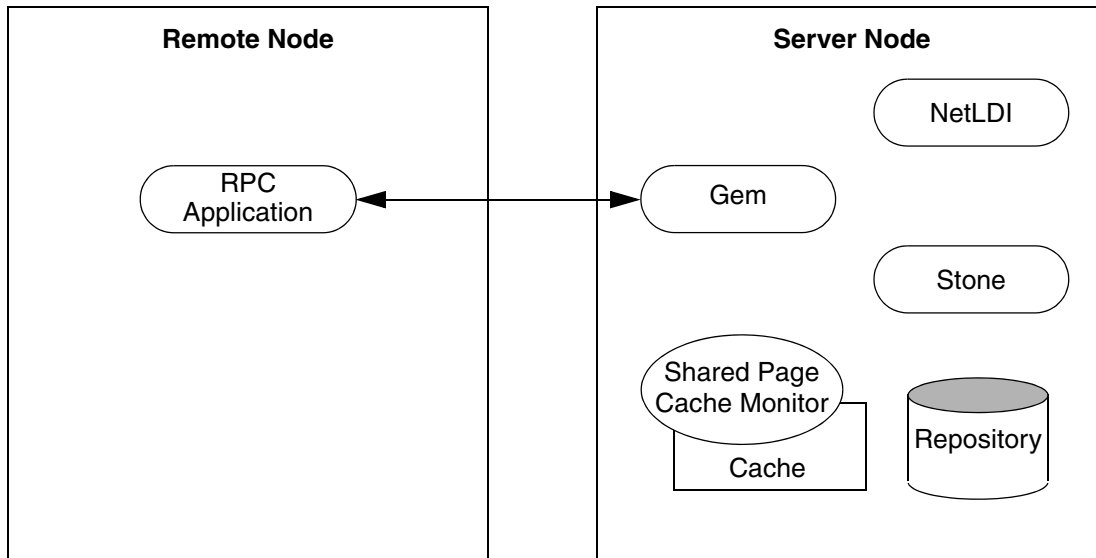
Although the Network File System (NFS) can be used to share executables, libraries, and configuration files, they are not recommended, and by default disallowed, for sharing repository extents and tranlogs.

Figure 3.1 and Figure 3.2 show two typical distributed configurations in which an application on a remote node is logged in to a repository and Stone repository monitor running on a server node.

In Figure 3.1, an application communicates with a Gem session process on the server node by way of RPC calls. This configuration lets the Gem execute Smalltalk code in the

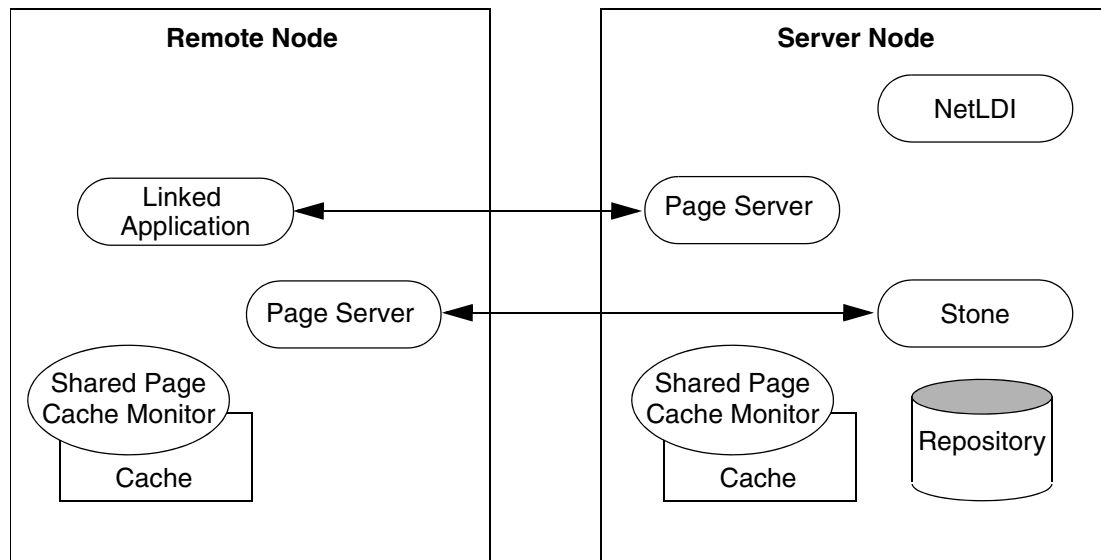
repository without first bringing complex objects across the network. The Gem can access the shared page cache that was started by the Stone repository monitor. For instructions on setting up this configuration, see “To Run the Gem Session Process on the Stone’s Node” on page 83.

Figure 3.1 Gem Session Process on Server Node



In Figure 3.2, the application and the Gem are linked in a single process that runs on the remote node. This configuration avoids the overhead of RPC calls, but in some applications it may increase network traffic substantially if large objects must be brought across the network. The Stone repository monitor starts a shared page cache on the remote node when the first user from that node logs in to the repository. The Stone and the Gem session process each use a GemStone page server to access data pages residing on the other node. For instructions on setting up this configuration, see “To Run a Linked Application on a Remote Node” on page 81.

Figure 3.2 Gem Session Process on Remote Node



GemStone NetLDIs

The GemStone network server process is called NetLDI (Network Long Distance Information). The NetLDIs are the glue holding a distributed GemStone system together. Each NetLDI reports the location of GemStone services on its node to remote processes that must connect to those services. It also spawns other GemStone processes on request.

In a distributed system, each node where a Stone repository monitor or Gem session process runs must have its own NetLDI. You do not need a NetLDI on nodes that are running only linked applications, or RPC applications with the gems on a different node.

You start a NetLDI directly by invoking the `startnetldi` command (page 324). The NetLDI, in turn, starts Gem session processes and page servers on demand. (See the following section for more about page servers.) These child processes belong by default to the user account of the process requesting the service – sometimes that account is a user logging in to GemStone, other times it is the account that started the repository monitor.

Because most operating systems only let the root account start processes that will be owned by other accounts, your NetLDI must be setup correctly if it is to serve more than one user. There are two ways to accomplish this:

- ▶ Set the owner and S bit for `$GEMSTONE/sys/netldid` so that the NetLDI runs as root.
- ▶ Run `netldi` in captive account mode, in which all processes started by the NetLDI belong to a single user account.

For more details on configuring NetLDI security, see “How To Arrange Network Security” on page 75.

NetLDI Ports and Names

The NetLDI listens for all requests on a single, fixed TCP/IP port, either as provided by the **startnetldi -P** option, or by querying the OS using `getservbyname()`. During the installation process, the NetLDI name is added to network services database (for a simple installation, the `/etc/services` file), and is assigned a port number.

The default name of the NetLDI process is `gs64ldi`. You may add your own NetLDI name instead, or in addition to this, and include multiple NetLDI names with different ports.

To access NetLDIs by name in a distributed system, the same name and port must be defined on every node. This includes all server nodes that will run Stones, Gems, or Caches, and also client nodes, Windows as well as UNIX, even if they do not run a NetLDI. This allows client nodes to reference a NetLDI by name during login.

To avoid the need to update the network services database, you can configure the well-known port by starting NetLDI using the **startnetldi -P** option, or by specifying a valid, unused port instead of a name in `startnetldi`. You must then reference the NetLDI using the port number, not the name; lookup by name requires an entry in `/etc/services`,

If you do not use the default NetLDI name `gs64ldi`, then you need to specify the NetLDI by name or port in NRS, using a line of the form `#netldi:netLdiName`. For example,

```
!@hostname.com#netldi:MyNetLDI!stoneName
```

To avoid entering the `netldi` with each NRS, you can add this information to the `GEMSTONE_NRS_ALL` environment variable for each user. For example:

```
$ GEMSTONE_NRS_ALL=#netldi:MyNetLDI
$ export GEMSTONE_NRS_ALL
```

For more information about `GEMSTONE_NRS_ALL`, see “How To Use Network Resource Strings” on page 78.

Stone and Shared Page Cache Monitor

Every named GemStone process – Stone, NetLDI, and Shared Page Cache Monitor – creates a `serviceName.LCK` at startup, in the directory `/opt/gemstone/locks` or `/usr/gemstone/locks`, or in a directory specified by `GEMSTONE_GLOBAL_DIR` (as discussed on page 30). The `.LCK` file holds the well-known port that the process is listening on. When the NetLDI (for example) contacts a particular Stone, it first looks up the well-known port for that Stone by locating the lock file for that Stone, then contacts the Stone on that port.

The ports used by the Stone and Shared Page Cache Monitor can be specified using the configuration parameters `STN_WELL_KNOWN_PORT_NUMBER` and `SHR_WELL_KNOWN_PORT_NUMBER`, respectively. These must be valid port numbers that are not already in use.

GemStone Page Servers

GemStone repository I/O is carried out by page server processes, running the executable file `pgsvrmain`. The Stone repository monitor creates one or more page servers at startup, to perform asynchronous I/O to the repository. There will also be Free Frame

page servers that are dedicated to added free frames to the free frame list. For more about page servers, see page 55.

Sessions on remote hosts also require page servers. For each process that connects to a repository extent across the network, the NetLDI service spawns a `pgsvrmain` on the stone's node.

GemStone Network Objects

GemStone uses the concept of *network objects* to encompass the services that a NetLDI can provide to a client. In addition to the page server, other network objects include the following services requested by the Stone at startup: the shared page cache monitor, SymbolGem, and garbage collection (GcGem) sessions.

The network object most visible to users is the Gem session process requested by an RPC application. This object can be `gemnetobject` or the name of a custom Gem. The request can be sent to the NetLDI on the same node to start a local session process, or (by using a network resource string) to a NetLDI on another node to start a process there.

The NetLDI first tries to map the requested object to the path of an executable by looking for an entry in `$GEMSTONE/sys/services.dat`. That file contains an entry for the standard Gem session process:

```
gemnetobject          $GEMSTONE/sys/gemnetobject
```

For example, when you enter "gemnetobject" as a session login parameter (such as for `gemnetid` in Topaz), the NetLDI uses the `services.dat` file to map the request to the script `$GEMSTONE/sys/gemnetobject`. Similarly, an object name can be entered while setting up a GemBuilder session (as Name of Gem Service) or other application. Application programmers provide the name as a parameter to `GciSetNet()`.

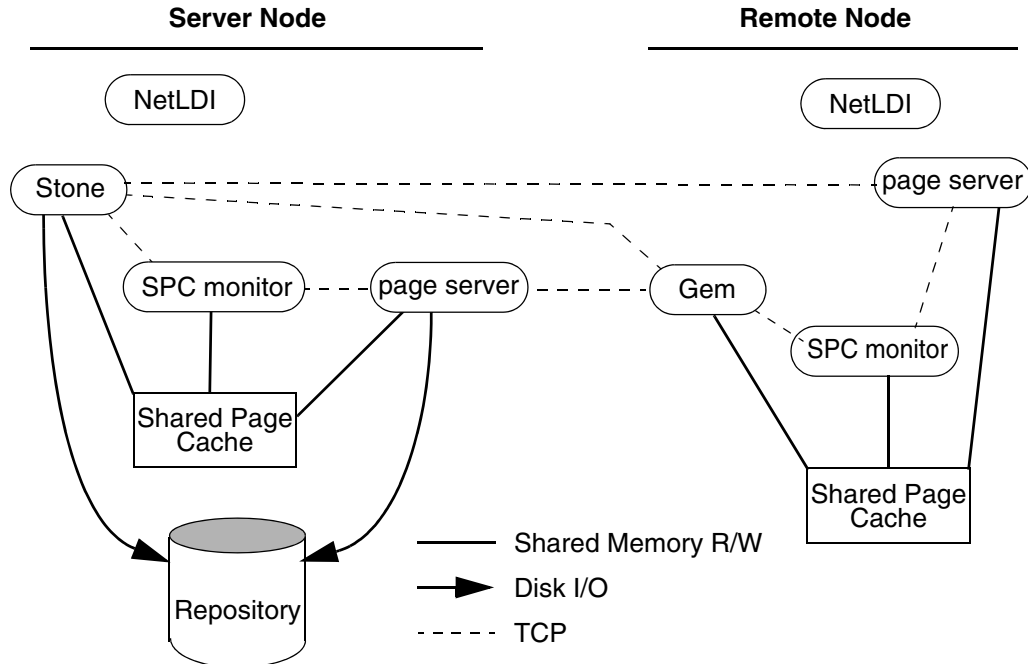
If your application uses a custom Gem executable, you can edit `services.dat` to include the appropriate mapping. For the procedure, see "How To Install a Custom Gem" on page 67.

If the NetLDI does not find the requested object in `services.dat`, it searches for an executable with that name in the user's `$HOME` directory. If you have a private Gem executable, place the executable in `$HOME` and then enter its name in place of `gemnetobject` during a GemStone login. Because of the search order, the private name must not be the same as that of an object in `services.dat`. The name must be the name of a file in `$HOME`, not a pathname.

Shared Page Cache in Distributed Systems

When the remote session logs in to the repository, the Stone repository monitor uses a NetLDI and page server (`pgsvrmain`) on the remote node to start a monitor process, and that monitor requests its local NetLDI to create a local shared page cache. When the remote Gem wants to access a page in the repository, it first checks the shared page cache on the remote node. If the page is not found, the Gem uses a `pgsvrmain` on the server node, checking in the shared cache on that node and then, if necessary, reading the page from the disk. See Figure 3.3.

Figure 3.3 Shared Page Cache with Remote Gem



Disrupted Communications

Several incidents can disrupt communications between the GemStone server and remote nodes in a distributed configuration. This can include node crashes, firewalls, and loss of the communications channel itself.

GemStone ordinarily depends on the network protocol *keepalive* option to detect that a remote process no longer exists because of an unexpected event. The *keepalive* interval is set globally by the operating system, typically at two hours. When that interval expires, the GemStone process tries to obtain a response from its partner. The parameters governing these attempts also are set by the operating system, with up to 10 attempts in 15 minutes being typical. If no response is received, the local GemStone process acts as if its partner was terminated abnormally. In some cases, you may wish to adjust the *keepalive* interval to allow for a longer timeout.

3.2 How To Arrange Network Security

This section describes the levels to which the system administrator can set the GemStone authentication requirement.

Running as Root with Authentication

When the NetLDI is started in default mode, it performs authentication for starting client processes. Since running the NetLDI in this mode from an ordinary user account requires that only that user may start processes, this mode is normally used with the NetLDI running as root, with the `setuid` bit set on the executables.

When running in this mode, the file permissions and ownership for the NetLDI executable should look similar to this:

```
-r-sr-xr-x 1 root gsadmin 674488 Mar 7 11:29 netldid
```

The UNIX host login name and password that will be used for authentication are set using the application's user interface, such as the `HostUserName` and `HostPassword` parameters in Topaz.

NOTE

Authentication is always done using the "real" user id, not the effective user id as set by the S bit on GemStone executables.

In this mode, PAM (Pluggable Authentication Modules) is used to authenticate, and the system should either configure a service with the name `gemstone.netldi`, or ensure that the default PAM authentication will allow logins. The actual means of authentication, such as LDAP, depend on how your system administrators have configured your UNIX installation.

Setting host Username and Password

Your application's login interface will generally let you specify a UNIX login name and password for the node on which you will be running an RPC Gem session process. For example, Topaz lets you set these as variables:

```
topaz> set hostusername yourLogin
topaz> set hostpassword yourPassword
```

GemBuilder for Smalltalk (GBS) provides similar fields in its login dialog, and other GemStone products and tools also provide ways to enter this information. .

You may specify the login name and password in NRS syntax. When you use topaz, GBS, or other GemStone tools, the fields for host user name and host password are used to build the NRS containing the authentication information. This is then passed to the NetLDI. For example, if you set the Topaz login parameters `HostUserName` and `HostPassword`, the application puts them in an NRS like the following:

```
'!@Server#auth:HostUserName@HostPassword!gemnetobject'
```

Although it is less convenient for ordinary use, this can be done manually, by entering the authorization modifier directly using the Topaz `GemNetId` parameter. For example:

```
topaz> set gemnetid !@Server#auth:name@password!gemnetobject
```

Authentication Levels

When running the NetLDI as root with `setuid`, there are two levels of authentication:

Default

In the default level of authentication, authentication is required each time a NetLDI attempts to start certain processes for a client, even if that process is to run on the node where the user is logged in. These situations always require authentication:

- ▶ Starting an RPC Gem session process, even on the same node.
- ▶ Creating or restoring a GemStone backup using a device on a node remote from the Gem performing the operation.
- ▶ Using `copydbf` between nodes.

Once a Stone or Gem is running, the NetLDI treats it as a trusted client and starts the page servers needed by a remote login without authentication. Simple network information requests, such as a request to look up a port number, also do not require authentication.

Secure Mode

`startnetldi -s` starts the NetLDI in *secure mode*. All accesses are authenticated, including simple requests to look up a server name. This mode affects the `waitstone` command and such user actions as connecting a session process to a remote Stone (a NetLDI is asked to look up the Stone's address). In secure mode, authentication is needed before a Gem or Stone can start a page server to access an extent or shared page cache on another node.

Running in Guest Mode with Captive Account

As an alternative, you may start the NetLDI with a captive account and in guest mode, which together provide the NetLDI security.

`startnetldi -aname` starts the NetLDI in *captive account mode*. All child processes created by the NetLDI will belong to the single, designated account *name*. This mode can be used when the NetLDI runs either as root or as *name*, but is normally combined with captive account.

The effect of captive account mode is much like setting the S bits on executables, but it only affects ownership of processes started by the NetLDI, not linked applications invoked directly by the user.

The captive account can be an ordinary user account or one created for that purpose, such as a GemStone administrative account. Log files by default will be in the captive account's home directory.

`startnetldi -g` starts the NetLDI in *guest mode*. No accesses are authenticated. Guest mode is not permitted if the NetLDI will run as root, so while it can be used without a captive account, it is normally combined with captive account mode.

Running the NetLDI in guest mode with captive account serves multiple users with the convenience of guest mode and with improved security, because the child processes no longer belong to accounts of individual users who request services.

The principal advantage of this combination is that the NetLDI can spawn processes on behalf of multiple users without being run as root. To make this capability possible, the captive account must own the `netldi` process. Change the file permissions and

ownership for the NetLDI executable to remove the setuid bit. The resulting permissions look like this:

```
-r-xr-xr-x 1 gsadmin gsadmin 674488 Mar 7 11:29 netldid
```

A disadvantage of the captive account for some applications is that the Gem session process will perform *all* I/O as that account, not as the account running the application – all file-ins, file-outs, and `System class >> performOnServer:.`

The captive account mode differs from the setuid method in that captive account mode affects *all* services started by the NetLDI, including any ad hoc processes, which are processes started from the user's home directory. (The NetLDI looks in the user's home directory if it cannot find a service listed in `$GEMSTONE/sys/services.dat.`) If you prefer, you can prohibit such ad hoc services by specifying the `-n` option when starting the NetLDI.

This is the procedure to configure guest and captive account:

- Step 1.** Create an OS account to own the GemStone distribution tree and serve as the captive account. We refer to this account as *gsadmin*.
- Step 2.** Make *gsadmin* the owner of the distribution tree, and set the setuid bit for any linked GemStone executables that run on the server node. Make the repository extents accessible only by *gsadmin* (mode 600). For instructions, see “To Set File Permissions for the Server” on page 45.

- Step 3.** Make sure that *gsadmin* has execute permission for `$GEMSTONE/sys/netldid.` The setgid bit should NOT be set on the `netldid` executable; it should look similar to this:

```
-r-xr-xr-x 1 gsadmin gsadmin 674488 Mar 7 11:29 netldid
```

- Step 4.** Log in as the captive account (such as *gsadmin*). Then start the NetLDI in guest mode and captive account mode. For instance:

```
% startnetldi -g -a gsadmin
```

You may specify other `startnetldi` options; for more details about the `startnetldi` command and its options, see page 324.

Table 1 shows how guest mode and captive account mode combinations affect NetLDI operation.

Table 1 NetLDI Guest and Captive Account Restrictions

NetLDI Options	Host passwords Required	Owner of Spawned Processes	Owner of NetLDI Process	Which Accounts Can Start Processes
(none)	Yes	Client's account	Ordinary user	Owner of NetLDI
			Root	Any user
-aname	Yes	Account <i>name</i> (must start the NetLDI)	Ordinary user (<i>name</i>)	Owner of NetLDI
			Root	Any user
-g	No	Client's account	Ordinary user	Owner of NetLDI
			Root – not allowed	
-aname -g	No	Account <i>name</i> (must start the NetLDI)	Ordinary user (<i>name</i>)	Any user
			Root – not allowed	

3.3 How To Use Network Resource Strings

Once you have chosen the remote and server nodes, network resource strings (NRS) allow you to specify the location of each part of the GemStone system. Use an NRS on a network system where you would use a process or file name on a single-node system. For example, suppose you want to know whether a Stone is running. If the Stone is on the local node, use this command:

```
% waitstone gemStoneName -1
```

If the Stone is on a remote node, use a command like this instead:

```
$ waitstone !@oboe!gemStoneName -1
```

where *oboe* is the Stone's node. You can also use an Internet address in "dot" form, such as 120.0.0.4, to identify the remote node. Note that each "!" must be preceded by a backslash (\) when your command will be processed by the C shell.

Appendix B, "GemStone Utility Commands," indicates which options of each UNIX-level GemStone command can be specified as an NRS. Besides location, an NRS can describe the network resource type so that GemStone can more accurately interpret the command line. Sometimes an NRS can also include your authorization to use that resource. For more information, see Appendix C, "Network Resource String Syntax."

To Set a Default NRS

You can set a default NRS header (the part between "! ... !") by setting the environment variable GEMSTONE_NRS_ALL. This variable determines which modifiers GemStone will use by default in each NRS it processes on your behalf. For instance, you can cause all Gem session process logs to be created with a specific name in a specific directory.

- ▶ If you set `GEMSTONE_NRS_ALL` before starting a NetLDI, which is a system-wide service, that setting is passed to all its children and becomes the default for all users of that service.
- ▶ If you set `GEMSTONE_NRS_ALL` before starting a Stone, an application, or a utility (such as `copydbf`), that setting applies only to your own processes and does not affect other users.

Because these settings are defaults, they take effect only if an explicit setting is not provided for the same modifier in a specific request.

Use the `#dir` modifier to set the current (working) directory for NetLDI child processes, such as `gemnetobject`. Without this setting, the default is the user's home directory. If the directory specified does not exist or is not writable at run time, an error is generated. For example:

```
$ GEMSTONE_NRS_ALL=#dir:/user2/apps/logs
$ export GEMSTONE_NRS_ALL
```

For further information about the modifiers and templates available, see Appendix C, "Network Resource String Syntax."

3.4 How To Set Up a Remote Session

Configuring a Gem session process on a remote node is much the same as configuring a session process on the server, which is described in Chapter 2, "Configuring Gem Session Processes." Keep the following points in mind:

- ▶ A remote node (on which a Gem is running) must have its kernel configured for shared memory similarly to how it is configured on the primary server node.
- ▶ Only nodes running a Stone need a GemStone key file, not nodes running remote sessions.
- ▶ If your site doesn't run NIS, add each node in the GemStone network to `/etc/hosts`.
- ▶ If your site doesn't run NIS, add the NetLDI entry to `/etc/services` on each node. Be sure to specify the same name and network port number each time.
- ▶ It's best if each node has its own directory for `/opt/gemstone/log` and `/opt/gemstone/locks` (or `/usr/gemstone/log` and `/usr/gemstone/locks`, or other location under `$GEMSTONE_GLOBAL_DIR`). If these directories are on an NFS-mounted partition, make sure that two nodes are not using the same directories. Each Stone and NetLDI needs a unique lock file. Shared log files may make it impossible to diagnose problems.
- ▶ Unless you run the NetLDIs in guest mode with a captive account, users must have an account on, and authorized network access to, all nodes that are part of the GemStone network for the repository they will be using. This includes the nodes on which the Stone Repository monitor and the user's Gem session process reside.
- ▶ Unless you run the NetLDIs in guest mode with a captive account, *the user who starts the Stone repository monitor* ordinarily needs an account on *every* node where a Gem session process will run.

You can either repeat the installation from the GemStone distribution media, as described in the next topic, or mount the directory on the server node that contains \$GEMSTONE (page 80).

To Duplicate the GemStone Installation

If you repeat the installation on the remote node, we recommend that you also run \$GEMSTONE/install/installgs. In particular, you should make the same selections regarding the ownership and group for the GemStone files as you did on the primary server node. You can save disk space later by deleting initial repositories (\$GEMSTONE/data/extent0.dbf and \$GEMSTONE/bin/*.dbf) and the complete upgrade (\$GEMSTONE/upgrade) and seaside (\$GEMSTONE/seaside) directories.

To Share a GemStone Directory

The following example prepares to run an application and Gem session process on a remote node using a shared software directory on the server. The GEMSTONE environment variable points to the shared installation directory, which is on the node *Server* and is NFS-mounted as */Server/users/g64stone*.

Step 1. Set the GEMSTONE environment variable to point to the NFS-mounted installation directory, and then invoke gemsetup:

```
(Bourne or Korn shell)
$ GEMSTONE=/Server/users/g64stone
$ export GEMSTONE
$ . $GEMSTONE/bin/gemsetup.sh
```

```
(C shell)
Remote% setenv GEMSTONE /Server/users/g64stone
Remote% source $GEMSTONE/bin/gemsetup.csh
```

Step 2. If they do not already exist, create the GemStone log and locks directories on the local node. (NetLDIs use this log directory.) You may need to have a system administrator do this for you as root.

```
# cd /opt
# mkdir gemstone gemstone/log gemstone/locks
# chmod 777 gemstone gemstone/log gemstone/locks
```

Configuration Examples

GemStone supports several configurations in which the application communicates with the Gem session process by using remote procedure calls (RPCs). Although the calls to network routines inevitably are time-consuming, they are essential when the application runs on a different node from the Gem, and they are desirable during code development because they isolate the application and Gem address spaces.

Use of RPC configurations for production repositories should be based on careful analysis of system loads and network traffic to select the most efficient configuration for a particular application. The RPC configuration may be desirable when the application accesses large or complex objects that would saturate the network if they were brought across it on a frequent basis.

This section presents examples that illustrate the following distributed applications:

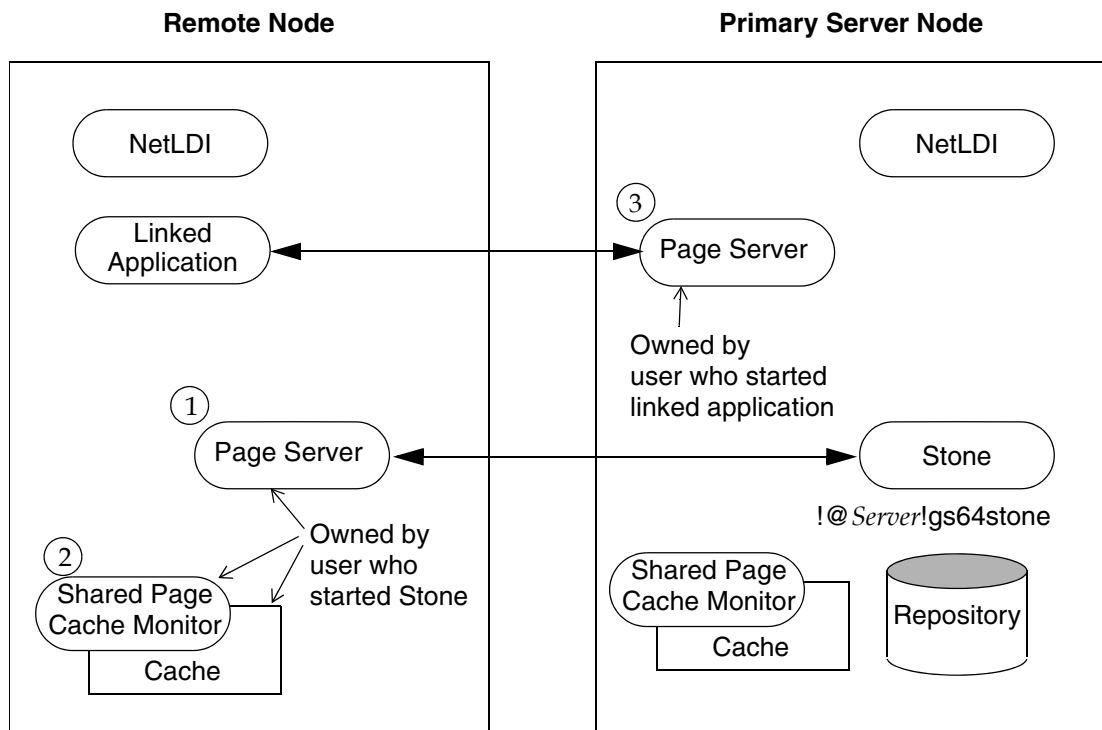
- ▶ A linked application connected to a Stone on another node (page 81).
- ▶ An RPC application with both the session process and the Stone on the server node (page 83).
- ▶ An RPC application with the session process on the application's node (page 84).
- ▶ An RPC application in which all three are on different nodes (page 86).

Two other examples show how to set up an extent on a node that is remote from the Stone, and how to use `copydbf` between nodes.

To Run a Linked Application on a Remote Node

Figure 3.4 shows how a linked application on a remote node communicates with a Stone and repository on the primary server node. This configuration typically is the best choice when you must offload some processes from a server node, especially when the application accesses relatively small objects or small groups of large objects.

Figure 3.4 Connecting a Linked Application to a Remote Server



Two NetLDIs and two page servers ordinarily are required. NetLDIs start the page servers on request of the Stone and the application. Numbers show the order in which these processes are started:

- ▶ One page server (1) lets the Stone start a shared page cache and monitor (2) on the remote node. The page server and monitor processes will be owned by the user who started the Stone (or by the captive account), so the owner must have an account on the remote node. The cache itself will have the same owner and group as the Stone. The linked application must have permission to access the cache, either through group membership or through an S bit on the application executable.
- ▶ The other page server (3) lets the Gem session process (the linked application) access the repository on the primary server. There will be one such page server process on the primary server node for each session logged in from a remote node; its owner (which may be a captive account) must have an account on the primary server. The page server process must have read-write permission for the repository, either through group membership or through an S bit on the `pgsvrmain` executable.

Because the shared page cache is readable and writable only by its owner and members of the same group (protection 660), the user running the application may need to belong to that group. See “To Set Ownership and Permissions for Session Processes” on page 61.

The following steps set up a linked application on the remote node. They use software in an NFS-mounted installation on the primary server node. (That directory is already mounted on the remote node as `/Server/g64stone`.)

Step 1. Set the GEMSTONE environment variable to point to the installation directory, and then invoke `gemsetup`:

```
(Bourne or Korn shell)
$ GEMSTONE=/Server/g64stone
$ export GEMSTONE
$ . $GEMSTONE/bin/gemsetup.sh
```

```
(C shell)
Remote% setenv GEMSTONE /Server/g64stone
Remote% source $GEMSTONE/bin/gemsetup.csh
```

Step 2. Verify that a Stone and NetLDI are running on the primary server node. One way to do this verification is to use the `gslis`t utility. For example:

```
Remote% gslis -m serverName
```

(The `-m` option tells `gslis`t to list only processes that are running on the specified node. For more about `gslis`t, see page 314.)

Step 3. Start a NetLDI on the remote node.

- ❑ To start the NetLDI for password authentication, make sure `$GEMSTONE/sys/netldid` is owned by root and has the S bit set. Issue this command (on some operating systems, you may have to issue it as root):

```
Remote% startnetldi
```

- ❑ To start the NetLDI in guest mode (authentication is not required), make sure `$GEMSTONE/sys/netldid` does NOT have the S bit set. Log in as the captive account `name`, then issue this command:

```
Remote% startnetldi -g -aname
```

Step 4. Start the linked application (for instance, Topaz) on the remote node, then set the *GemStone* login parameter to include the name of the primary server node in network resource syntax. For instance, to log in to Topaz as DataCurator:

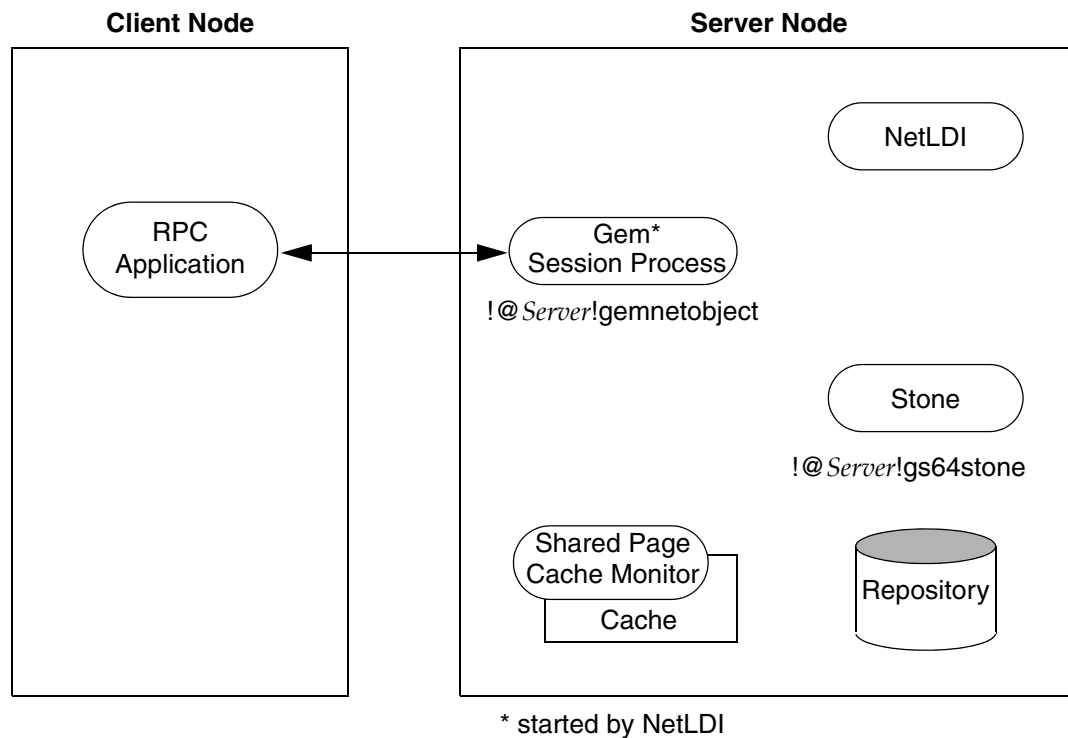
```
Remote% topaz -l
topaz> set gemstone !@Server!gs64stone
topaz> set username DataCurator
topaz> login
GemStone Password?
successful login
topaz 1>
```

To Run the Gem Session Process on the Stone's Node

If the Gem session process is going to run on the server node (as shown in Figure 3.5), an RPC application uses a NetLDI on that node to start a Gem session process. Unless the NetLDI is running in guest mode with a captive account, the application user must provide authentication to the NetLDI. You should also specify the Gem network object (*gemnetobject*) that matches your UNIX shell on the server. For more information about network objects and how to invoke them, see “GemStone Network Objects” on page 73.

The following procedure assumes that you are already set up to run GemStone applications, as described in Chapter 2. In particular, you must have defined the *GEMSTONE* environment variable and invoked `$GEMSTONE/bin/gemsetup.sh` or its equivalent.

Figure 3.5 Starting a Session Process on the Server Node



Step 1. Make sure that the NetLDI and Stone are running on the server. One way to do this is to use the `gslist` command. For example:

```
% gslist -m serverName
```

Step 2. Unless the NetLDI is running in guest mode, set the application login parameters, such as `HostUserName` and `HostPassword`, after you start the application. For example:

```
topaz> set hostusername yourLogin  
topaz> set hostpassword yourPassword
```

Step 3. Log in to the application node and start the RPC version of your application (for instance, `Topaz`), then set `UserName`. For example:

```
% topaz  
topaz> set username DataCurator
```

Step 4. Set `GemNetId` to `gemnetobject`. Because the session process is to run on the server, be sure to include the node name in the `GemNetId` NRS. (It's not necessary to set the `GemStone` login parameter when the `Stone` repository monitor runs on the same node as the `Gem`.) For example:

```
topaz> set gemnetid !@Server!gemnetobject
```

Step 5. Log in to the repository:

```
topaz> login  
GemStone Password?  
successful login  
topaz 1>
```

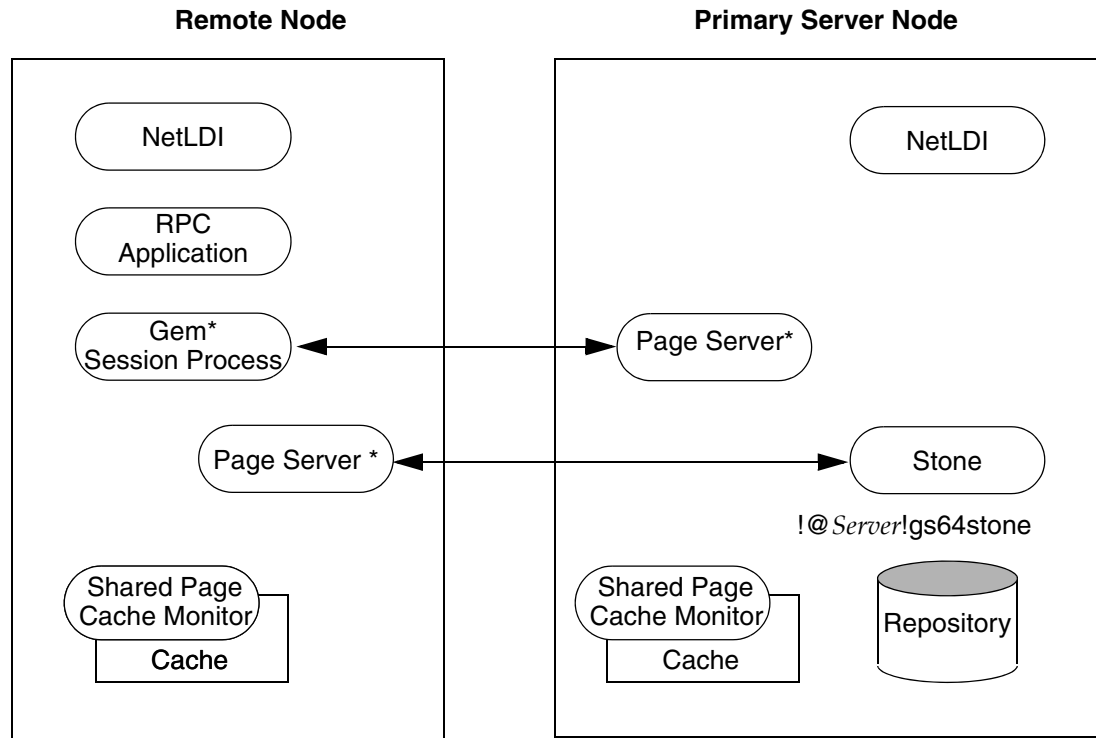
At this point, you are logged in to a `Gem` session process on the server node. That session process acts as a server to `Topaz` RPC and as a client to the `Stone`.

To Run the Gem and Stone on Different Nodes

The configuration shown in Figure 3.6 is unusual in that the RPC application and its session process are running on the same node. (While this configuration might be desirable during application development, a linked application, if it is available, probably would give better performance.)

The `NetLDIs` and page servers function similarly to those described for the linked application (see "To Run a Linked Application on a Remote Node" on page 81). In Figure 3.6, however, the `NetLDI` also starts the RPC `Gem` session process at the request of the application.

Figure 3.6 Starting the Session Process on a Remote Node



* started by NetLDI

Step 1. Unless the NetLDIs are running in guest mode, decide how you will provide access so that application can start a Gem session process on the remote node.

Set the application login parameters, such as HostUserName and HostPassword, after you start the application. For example:

```
topaz> set hostusername yourLogin
topaz> set hostpassword yourPassword
```

Step 2. Log in to the remote node and start a NetLDI.

- ▶ To start the NetLDI for password authentication, make sure that `$GEMSTONE/sys/netldid` is owned by root and has the S bit set. Issue this command (on some operating systems, you may have to issue it as root):


```
Remote% startnetldi
```
- ▶ To start the NetLDI in guest mode (authentication is not required), make sure `$GEMSTONE/sys/netldid` does NOT have the S bit set. Log in as the captive account *name*, then issue this command:


```
Remote% startnetldi -g -aname
```

Step 3. Start the RPC version of your application (for instance, Topaz):

```
Remote% topaz
```

Step 4. Set GemNetId to gemnetobject. This network object identifies scripts that start a session process. For example:

```
topaz> set gemnetid gemnetobject
```

Step 5. Set the GemStone name, using NRS syntax to specify its location on the primary server node. Then set the UserName and log in. For example:

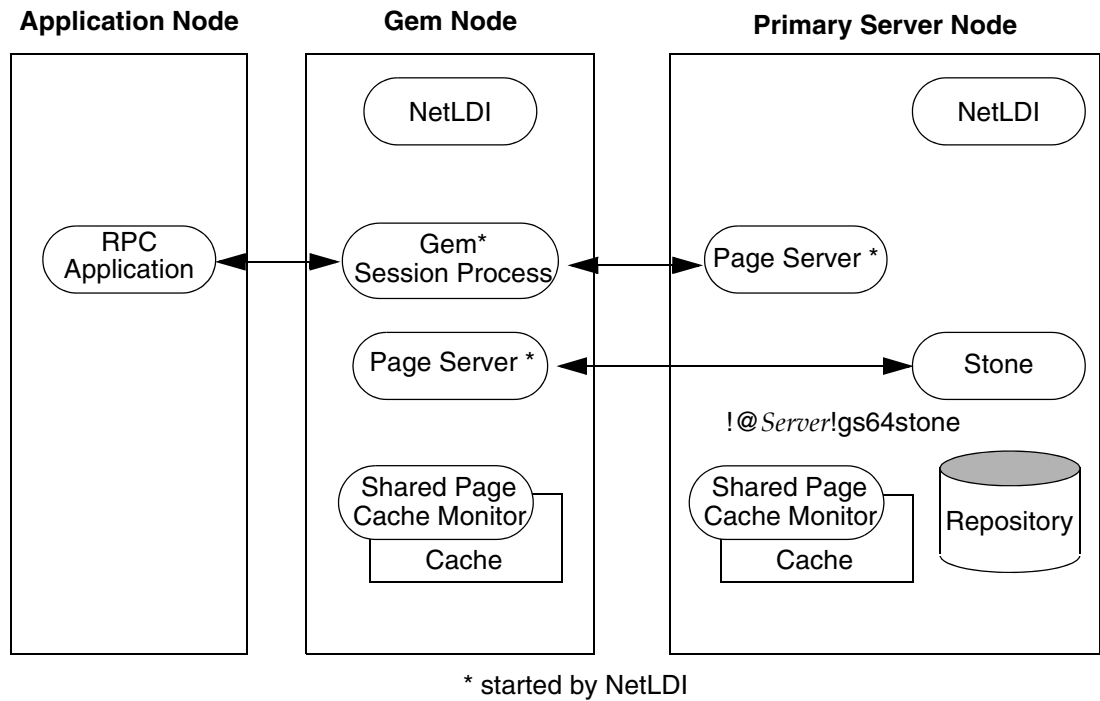
```
topaz> set gemstone !@Server!gs64stone
topaz> set username DataCurator
topaz> login
GemStone Password?
successful login
topaz 1>
```

To Run the Application, Gem, and Stone on Three Nodes

The RPC application, session process, and Stone can run on three separate nodes, as shown in Figure 3.7. The application runs on its node and connects to a Gem session process on the Gem's node. That session process communicates with the repository on the primary server node by way of a page server.

Again we see that a NetLDI must be running on each node where part of GemStone executes (but not necessarily on the application node, which runs only the RPC application).

Figure 3.7 Connecting an RPC Application, Three Nodes



The network access problem is similar to that in other RPC configurations: unless the NetLDI on the Gem node is running in guest mode, you must provide authentication to start the Gem session process.

Step 1. Unless the NetLDI on the Gem node is running in guest mode, set the application login parameters, such as `HostUserName` and `HostPassword`, after you start the application. For example:

```
topaz> set hostusername yourLogin
topaz> set hostpassword yourPassword
```

Step 2. Log in to the Gem's node and start the NetLDI.

- ▶ To start the NetLDI for password authentication, make sure that `$GEMSTONE/sys/netldid` is owned by root and has the S bit set. Issue this command (on some operating systems, you may have to issue it as root):

```
Remote% startnetldi
```

- ▶ To start the NetLDI in guest mode (authentication is not required), make sure `$GEMSTONE/sys/netldid` does NOT have the S bit set. Log in as the captive account `name`, then issue this command:

```
Remote% startnetldi -g -aname
```

Step 3. Log in to the application node. Start the RPC version of your application (for instance, Topaz):

```
Application% topaz
```

Step 4. Set GemNetId to gemnetobject, and include the location, *gemNode*, in the NRS. For example:

```
topaz> set gemnetid !@gemNode!gemnetobject
```

Step 5. Use NRS syntax to specify the location and name of the repository. Then set the username and log in. In Topaz, for example, set GemStone and UserName:

```
topaz> set gemstone !@Server!gs64stone
topaz> set username DataCurator
topaz> login
GemStone Password?
successful login
topaz 1>
```

At this point, your Topaz application on the application node has logged you in to a Gem session process on the Gem's node, and the session process has logged in to the repository on the server.

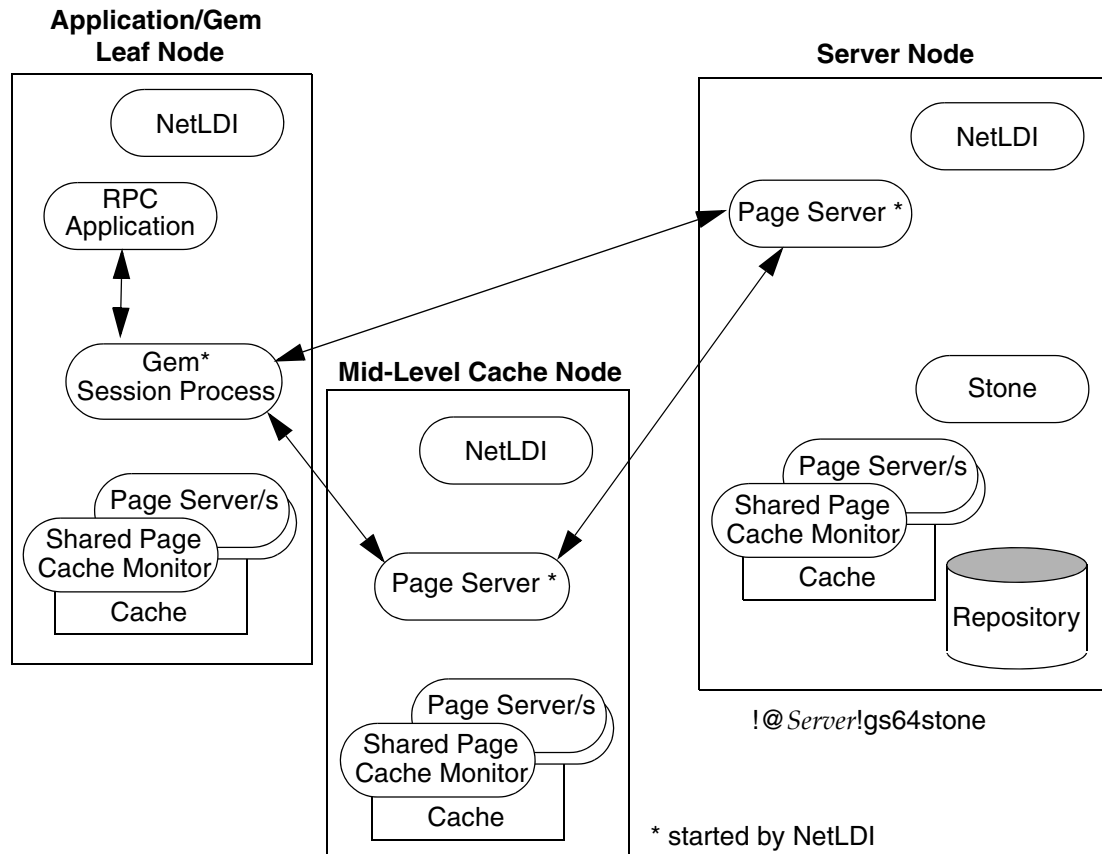
To Run a Distributed System with a Mid-Level Cache

In a distributed system over a Wide Area Network (WAN), with many remote nodes that are topographically distant from the Stone but close to each other, a mid-level cache can improve performance for the remote sessions. In this configuration, the RPC application and the session process may be on the same or different nodes, with the mid-level cache and Stone running on separate nodes. In Figure 3.8, the Gem session process (on the leaf node) connects to a mid-level cache (on the mid-level cache node). The session process communicates with the repository (on the primary server node) by way of a page server.

When the Gem needs a page but can't find it in its local cache, it first looks in the mid-level cache. If the Gem can't find the page in the mid-level cache, it then forwards the request to the page server on the Stone's host.

The page manager's pgsvr aggregates the responses from the pgsvrs on each of the gemSCs, and returns a combined response to the page manager. This reduces the number of round trips from the page manager to distant nodes.

If a mid-level cache is in use, then for each Gem process using the mid-level cache, all the shared caches to which the Gems are attached are subordinate to that mid-level cache.



Setting up a configuration with a mid-level cache requires that the session execute code following login to start a mid-level cache on a specified host, or to connect to an existing mid-level cache. Unlike the other remote configurations discussed in this chapter, the configuration is not established entirely by configuration settings and login arguments.

The Gem must be on a node that is remote from the Stone, and the request to connect to a mid-level cache must specify a node that is neither the Stone's nor the Gem's node.

If a Gem is running on the same machine as a mid-level cache, that Gem will use the mid-level cache as its local cache.

Connection Methods

System Class methods in the Shared Cache Management category allow you to connect to a mid-level cache.

`midLevelCacheConnect : hostName`

Attempts to connect to a mid-level cache on the specified host, if the cache already exists. The session's Gem process must be on a machine different from the machine running the Stone process.

`midLevelCacheConnect : hostName cacheSizeKB : aSize
maxSessions : nSess`

If a mid-level cache does not already exist on the specified host, and *aSize* > 0, attempts to start the cache and connect to it. If a cache is already running on the host, this method attempts to connect to the cache and ignores the other arguments.

The size of the mid-level cache is controlled by the method argument *aSize*, rather than by configuration parameters (as with other shared caches).

For example,

```
topaz> set gemstone !@Server!gs64stone
topaz> set username DataCurator
topaz> set gemnetid !@gemNode!gemnetobject
topaz> login
GemStone Password?
<details omitted>
successful login
topaz 1>
topaz 1> printit
    System midLevelCacheConnect: 'midLevelCacheNode'
        cacheSizeKB: 2048 maxSessions: 10
%
```

Reporting Methods

System Class methods in the Shared Cache Management category return lists of the shared caches on your system.

`remoteCachesReport`

Returns a String that lists all shared caches that the Stone process is managing, not including the cache on the Stone machine.

`midLevelCachesReport`

Similar to `remoteCachesReport`, but only includes the mid-level caches.

3.5 Troubleshooting Remote Logins

Logging into GemStone from a remote node requires proper system configuration of the remote node and frequently requires permission for network access from the primary server to the remote node as well as from the remote node to the primary server.

- ▶ The UNIX kernel on the remote node should meet shared memory and semaphore requirements similar to those for the server, although smaller sizes may be sufficient.
- ▶ Make sure that NetLDIs are running on all nodes that require them (see the figure for your configuration). Also make sure that the NetLDIs have the same port number in `/etc/services`. All nodes must be listed in `/etc/hosts`.
- ▶ If an RPC application is being started (that is, one with a separate Gem session process), make sure the user who starts the application has entered their host username and host password, or that NetLDI is running in guest mode with a captive account. The owner of the Gem process needs an account on the node where the Gem will run and needs write access to the Gem log, typically in `$HOME`. Ownership and permissions for `$GEMSTONE/sys/netldid` must be appropriate for the

authentication system in use (see pages 75 and 77), and the directories in `/opt/gemstone` must be writable.

- ▶ Make sure that the user who started the Stone has an account on the remote node. This user also must have write permission for `$HOME` so that log files for the remote node can be created, unless steps are taken to create the log files in another directory.
- ▶ Check any `GEMSTONE` environment variables for definitions that point to a previous version: `env | grep GEM`.

If You Still Have Trouble

If you still can't log in to GemStone from an application on a remote node, try logging in on the server node as the same UNIX user account. We suggest that you first try a linked application, such as `topaz -l`, and when that works, move on to an RPC application (such as `topaz` or the equivalent `topaz -r`), still on the server.

Try Linked Topaz on the Server

A linked application on the server offers the least complicated kind of login because the server's shared page cache is already running and no network facilities are used. Any problems are likely to involve access permission for the shared page cache or the repository extents, which can also block attempts to log in from a client node.

- ▶ Make sure that the owner of the topaz process (`$GEMSTONE/bin/topaz`) can access the shared page cache. Use the UNIX command `ipcs -m` to display permissions, owner, and group for shared memory; for example:

```
Server% ipcs -m
IPC status from <running system> as of Thu Mar 13 16:22:27 PDT
2014
T      ID      KEY          MODE          OWNER        GROUP
Shared Memory:
m      768 0x4c177155  --rw-rw----  gsadmin      pubs
```

Compare the owner and group returned by `ipcs` with the owner of the Topaz process. You can use the `ps` command to determine the owner; for example, `ps -ef | grep topaz`. (The switches may be different on your system.)

A typical problem arises when root owns the Stone process and the shared page cache because their group ordinarily will be a special one to which Topaz users do not belong. Related problems may occur with a linked GemBuilder session. The third-party Smalltalk may be installed without the S bits and therefore may rely on group access to the shared page cache and repository. For background information, see "To Set Ownership and Permissions for Session Processes" on page 61.

To correct a shared page cache access failure, either change the owner and group of the setuid files or have the Stone started by a user whose primary group is one to which other GemStone users belong. Unlike file permissions, the shared page cache permissions cannot be set directly.

- ▶ Make sure the owner of the Topaz process has read-write access to `$GEMSTONE/data/extent0.dbf`.

Try Topaz RPC on the Server

The next step should be to try running Topaz on the server with a separate Gem session process. This configuration relies on the NetLDI to start a Gem session process, and that process, not the application itself, must be able to access the shared page cache and repository extent.

- ▶ Make sure that a NetLDI is running on the server by invoking **gslist**. The default name is `gs64ldi`. If you need to start a NetLDI, the command is **startnetldi**.

GemStone uses the NetLDI to start a Gem session process that does repository I/O in this configuration. For the NetLDI to start processes for anyone other than its owner, it must be owned by root or it must be started in guest mode and captive account mode by someone logged in as the captive account.

The NetLDI writes a log file with the default name `/opt/gemstone/log/gs64ldi.log`; this may be overridden by the **startnetldi -l** option. Using the **gslist -x** command will provide the location of all log files. The log file contents may help you diagnose problems. (See the following discussion, “Check NetLDI Log Files.”)

- ▶ Make sure that the owner of the resulting Gem session process (`$GEMSTONE/sys/gem`) can access the shared page cache and `extent0.dbf` through group membership or S bits. The troubleshooting is the same as that given on page 91 for the `topaz` executable.

The user who starts `topaz` (or the NetLDI captive account when it is in use) must have write permission for `$HOME` so that the session process can create a log file there. (For a workaround for situations where write permission is not allowed, see “To Set a Default NRS” on page 78.)

Check NetLDI Log Files

Troubleshooting on a distributed GemStone system can be complicated. What looks like a hung process may actually be caused by incorrect NRS syntax or by another node on the network going down. The information for analyzing problems may be found in log files on all the nodes used by GemStone.

Where the log file messages include NRS strings, be sure to check their syntax. The problem may be as simple as an incorrect NRS or one that was not expanded by the shell as you intended.

If you can't identify the problem from the standard log messages, try running the NetLDI in debug mode, which puts additional information in the log. The command line is **startnetldi [netLdiName] -d**.

This chapter shows you how to perform some common GemStone/S 64 Bit system operations:

- ▶ Starting the GemStone Object Server (page 93)
- ▶ Starting Network Long Distance Information (NetLDI) servers (page 98)
- ▶ Identifying running servers (page 100)
- ▶ Logging in to a GemStone session (page 100)
- ▶ Identifying the current sessions (page 104)
- ▶ Shutting down the object server (page 105)

Additional topics explain how to:

- ▶ Recover from an unexpected shutdown (page 106)
- ▶ Load objects in bulk (page 109)
- ▶ Enter and use manual transaction mode, which we recommend that you use as often as possible (page 109)

4.1 How To Start the GemStone Server

In order to start a Stone repository monitor, the following must be identified through your operating system environment:

- ▶ Where GemStone is installed

The GEMSTONE environment variable must point to the directory where GemStone is installed, such as `/users/gemstone`. The directory `$GEMSTONE/bin` should be in your search path for commands.

- ▶ Which configuration parameters to use

The repository monitor must find a configuration file. The default is `$GEMSTONE/data/system.conf`. Other files can supplement or replace the default file; for information, see “How GemStone Uses Configuration Files” on page 266.

▶ Which repository to use

The configuration file must give the path to one or more repository files (extents) and to space for transaction logs. The default configuration file specifies `$GEMSTONE/data/extent0.dbf` for the extent file, and places transaction logs in `$GEMSTONE/data/`. You may want to move these files to other locations. For further information, see “Choosing the Extent Location” on page 35.

To Start GemStone

Follow these steps to start GemStone following installation or an orderly shutdown. (To recover from an abnormal shutdown, refer to “Do NOT use kill -9 or another uncatchable signal, which may not result in a clean shutdown or may cause the Stone repository monitor to shut down when you intended to kill only a Gem process. If for some reason you need to send kill -9 to a shared page cache monitor, use `ipcs` and `ipcrm` to identify and free the shared memory and semaphore resources for that cache. If you send kill -9 to a Stone, use `ipcs` to determine whether `ipcrm` should be invoked.” on page 106.)

Step 1. Set the `GEMSTONE` environment variable to the *full pathname* (starting with a slash) of the directory where GemStone is installed. Ordinarily this directory has a name like `GemStone64Bit3.2.0-x86_64.Linux` (depending on your platform). For example:

```
$ GEMSTONE=/users/GemStone64Bit3.2.0-x86_64.Linux
$ export GEMSTONE
```

If you have been using another version of GemStone, be sure you update or unset previous settings of these environment variables:

- ▶ `GEMSTONE`
- ▶ `GEMSTONE_SYS_CONF`
- ▶ `GEMSTONE_EXE_CONF`
- ▶ `GEMSTONE_NRS_ALL`

Step 2. Set your UNIX path. One way to do this is to use one of the `gemsetup` scripts. There is one version for users of the Bourne and Korn shells and another for users of the C shell. These scripts also set your man page path to include the GemStone man pages. Note that these scripts append to the end of your path or man path.

```
(Bourne or Korn shell)
$ . $GEMSTONE/bin/gemsetup.sh

or (C shell)
% source $GEMSTONE/bin/gemsetup.csh
```

Step 3. Start GemStone by using the `startstone` command:

```
% startstone [gemStoneName]
```

where *gemStoneName* is optional and is the name you want the repository monitor to have. The default name is `gs64stone`. For additional information about `startstone`, see page 326.

To Troubleshoot Stone Startup Failures

If the Stone repository monitor fails to start in response to a **startstone** command, it's likely that the cause is one of the following. Inspect the Stone log for clues (the default location is `$GEMSTONE/data/g64stone.log`).

- ▶ The GemStone key file is missing or invalid.
- ▶ The shared page cache cannot be attached.
- ▶ A problem with the extents: a missing extent or one that is in use by another process
- ▶ A problem with transaction logs: a log needed for recovery is missing, or the log directory or device does not exist
- ▶ The repository has become corrupted

Missing or Invalid Key File

The Stone repository monitor must be able to read the GemStone key file. This is by default in the filename and at the file path `$GEMSTONE/sys/gemstone.key`, but the location and filename can be configured by the `KEYFILE` configuration parameter.

Ordinarily, you create this file during installation from information provided by GemStone. Be careful to enter the information correctly. GemStone key files are platform-specific, and key files for earlier versions may not work with new releases.

If you do not have a valid key file, contact GemStone Technical Support as described on page 5.

Shared Page Cache Cannot Be Attached

The shared page cache monitor must be able to create and attach to the shared memory segment that will serve as the shared page cache. Several factors may prevent this from happening:

- ▶ On some platforms, shared memory is not enabled in the kernel by default, or its default maximum size is too small to accommodate the GemStone configuration. GemStone's default configuration requires a shared memory segment somewhat larger than 75 MB. For specifics about configuring shared memory, refer to the *GemStone/S 64 Bit Installation Guide* for your platform.
- ▶ If the size of the shared page cache has been increased, the operating system's limit on shared memory regions may need to be increased accordingly. Page 33 describes a utility (`$GEMSTONE/install/shmem`) that will help you check the configuration.
- ▶ The repository executables (the Stone, Gems, and page servers) must have permission to read and write the shared page cache. Ways to set up access are described in "To Set File Permissions for the Server" on page 45. In general, users must belong to the same group as the Stone repository monitor. If the Stone is running as root, it is unlikely that other users will be able to access the shared page cache.

Extent Missing or Access Denied

If the Stone repository monitor cannot access a repository extent file, it logs a message like the following:

```
GemStone is unable to open the file  
!TCP@pelican#dbf!/users/GemStone/data/extent0.dbf.  
reason = File = /users/GemStone/data/extent0.dbf  
DBF Op: Open; DBF Record: -1;  
Error: open() failure; System Codes: errno=2, ENOENT, The file or  
directory specified cannot be found
```

An error occurred opening the repository for exclusive access.

Stone startup has failed.

Examine the message for further clues. The extent file could be missing, the permissions on the file or directory could be set incorrectly, or there may be an error in the configuration file that points to the extents. Correct the problem, then try starting GemStone again.

Extent Open by Another Process

If another process has an extent file open when you attempt to restart GemStone, a message like the following appears in the Stone log (by default, `$/GEMSTONE/data/gs64stone.log`):

```
GemStone is unable to open the file  
!TCP@pelican#dbf!/users/GemStone/data/extent0.dbf.  
reason = File = /users/GemStone/data/extent0.dbf  
DBF Op: Open; DBF Record: -1;  
Error: exclusive open: File is open by another process.; System Codes:  
errno=11, EAGAIN, No more processes (due to process table full, user  
quotas, or insufficient memory)
```

An error occurred opening the repository for exclusive access.

Stone startup has failed.

Close any other Gem sessions (including Topaz sessions) that are accessing the repository you are trying to restart, or wait for a `copydbf` to complete. Use `ps -ef` (the options on your system may differ) to identify any `pgsvrmain` processes that are still running, and then use `kill processid` to terminate them. Try again to start GemStone.

Extent Already Exists

If GemStone attempts to recover from a system crash that occurred just after an extent was created, and GemStone was not able to write a checkpoint when the extent was added, you will find an error message like the following in the Stone log:

```
An error occurred in recovery for extentId 2:  
fileName= !TCP@pelican#dbf!/users/GemStone/data/extent1.dbf  
File already exists; you must delete it before recovery can succeed.
```


Check that an extent was being added to the repository at or shortly before the crash. If necessary, look for a message near the end of the Stone log file.

- ▶ If an extent was being added, there is no committed data in the extent file yet. Delete the specified file and do not replace it with anything. Try to start GemStone again. The recovery procedure will recreate the extent file.
- ▶ If an extent was NOT being added, it is possible that an existing extent has been corrupted. For instance, `extent0.dbf` of a multiple-extent repository may have been overwritten. Try to determine the cause and whether the action can be rectified. You may have to restore the repository from a backup.

Other Extent Failures

At startup, the GemStone system performs consistency checks on each extent listed in `DBF_EXTENT_NAMES`.

All extents must have been shut down cleanly with a repository checkpoint the last time the system was run. This consistency check is the only one for which GemStone attempts automatic recovery.

The following consistency checks, if failed, cause the startup sequence to terminate. These failures imply corruption of the disk or file system, or that the extents were modified at the operating system level (such as by `cp` or `copydbf`) outside of GemStone's control and in a manner that has corrupted the repository.

- ▶ Extents must be in proper sequence within `DBF_EXTENT_NAMES`.
- ▶ Extents must be properly sequenced in time.
- ▶ The last checkpoint must have occurred earlier than or at the same time as the current system time (in GMT).
- ▶ Extents must belong to the correct repository.

Transaction Log Missing

If GemStone cannot find the transaction log file for the period between the last checkpoint and an unexpected shutdown, it puts a message like this in the Stone log:

```
Extent 0 was not cleanly shutdown; recovery is needed.  
<Repository startup statistics>
```

```
Repository startup is from checkpoint = (fileId 6, blockId 3)
```

```
ERROR: cannot find log file(s) to recover repository.  
To proceed without tranlogs and lose transactions committed  
since the last checkpoint use "-N" switch on your startstone  
command.
```

```
An error occurred when attempting to start repository recovery.  
Waiting for aiowrites to complete
```

```
Stone startup has failed.
```

If the log file was archived and removed from the log directory, restore the file.

If the log file is no longer available, you can use **startstone -N** to restart from the most recent checkpoint in the repository. However, any transactions that occurred during the intervening period cannot be recovered.

NOTE

*When you use **startstone** with the **-N** option, any transactions occurring after the last checkpoint are permanently lost.*

Other Startup Failures

- ▶ Check `/opt/gemstone/locks` (or equivalent location, as discussed on page 30) and remove old files. On Solaris systems, also check `/tmp/gemstone` for `stoneName . . FIFO`.
- ▶ Certain unexpected shutdowns may leave UNIX interprocess communication facilities allocated, which can block attempts to restart the repository monitor. Use the command **ipcs** to identify the shared memory segments and semaphores allocated, then use **ipcrm** to free those resources allocated to a repository monitor that is no longer running. For information about **ipcs** and **ipcrm**, consult your operating system's documentation.
- ▶ If it takes more than 5 minutes for your cache to complete initialization, the startup timeout may be expiring. Set the environment variable `$GEMSTONE_SPCMON_STARTUP_TIMEOUT`; see page 354.
- ▶ Check your installation configuration and make sure that all required files and libraries are present and uncorrupted.
- ▶ Try to run **pageaudit** on the repository. (See "How To Audit the Repository" on page 120.)

If you are still unable to start GemStone or determine the reason that startup is failing, contact your local GemStone administrator or GemStone Technical Support.

If this is an existing GemStone repository and the problems reported on startup attempts indicate that the repository is corrupt, you may need to restore from backups, as described in Chapter 9. See "How to Restore from Backup" on page 200.

4.2 How To Start a NetLDI

You will usually need to start a GemStone Network Long Distance Information (NetLDI) server when starting a Stone repository monitor. NetLDI servers are needed to start up Gem processes for RPC logins, and for starting up caches on behalf of Gems that are on other nodes.

If you are running distributed configurations, you will need to perform these steps on each node that requires a NetLDI.

To start a NetLDI server, perform the following steps on the node where the NetLDI is to run:

- Step 1.** Set the GEMSTONE environment variable to the *full pathname* (starting with a slash) of the directory where GemStone is installed. Ordinarily this directory has a

name like `GemStone64Bit3.2.0-x86_64.Linux` (depending on the platform). For example:

```
$ GEMSTONE=/installDir/GemStone64Bit3.2.0-x86_64.Linux
$ export GEMSTONE
```

If you have been using another version of GemStone, be sure you update or unset previous settings of these environment variables:

- ▶ `GEMSTONE_NRS_ALL`

Step 2. Use one of the `gemsetup` scripts to set your UNIX path. There is one version for users of the Bourne and Korn shells and another for the C shell. These scripts also set your man page path to include the GemStone man pages.

(Bourne or Korn shell)

```
$ . $GEMSTONE/bin/gemsetup.sh
```

or (C shell)

```
% source $GEMSTONE/bin/gemsetup.csh
```

Step 3. Start the NetLDI by using the `startnetldi` command.

- ❑ To start the NetLDI for password authentication, make sure that `$GEMSTONE/sys/netldid` is owned by root and has the S bit set. Issue this command (on some operating systems, you may have to issue it as root):

```
% startnetldi
```

- ❑ To start the NetLDI in guest mode (authentication is not required), make sure that `$GEMSTONE/sys/netldid` does NOT have the S bit set. Log in as the captive account *name*, then issue this command:

```
% startnetldi -g -aname
```

For additional information about `startnetldi`, see the command description in Appendix B. For information about the authentication modes, see “How To Arrange Network Security” on page 75.

To Troubleshoot NetLDI Startup Failures

If the NetLDI service fails to start in response to a `startnetldi` command, it’s likely that the cause is one of the following:

- ▶ The NetLDI is to run as root but the guest mode option is specified. This combination is not allowed.
- ▶ The account starting the NetLDI does not have permission to create or append to its log file.
- ▶ The account starting the NetLDI does not have read and execute permission for `$GEMSTONE/sys/netldid`.

Check the NetLDI log for clues. By default, the NetLDI log (`netLdiName.log`) is located in `/opt/gemstone/log`. On some systems, this file may be located in `/usr/gemstone/log`, and may be overridden using the `-l` option to the `startnetldi` command, or by setting `$GEMSTONE_GLOBAL_DIR`.

4.3 To List Running Servers

The `gslist` utility lists all Stone repository monitors, shared page cache monitors, and NetLDIs that are running. The `gslist` command by itself checks the locks directory (`/opt/gemstone/locks`, `/usr/gemstone/locks`, or `$GEMSTONE_GLOBAL_DIR/locks`) for entries. The `-v` option causes it to verify that each process is alive and responding. For example:

```
% gslist -v
Status Version Owner      Started      Type  Name
-----
OK      3.2.0   gsadmin   Mar 11 12:02 cache  gs64stone-1c9fa07f0412665
OK      3.2.0   gsadmin   Mar 11 12:02 Stone  gs64stone
OK      3.2.0   gsadmin   Mar 11 10:13 Netldi  gs64ldi
```

By default, `gslist` lists servers on the local node. The `-m host` option performs the operation on node `host`, which must have a compatible NetLDI running.

4.4 How To Start a GemStone Session

This section tells how to start a GemStone session and log in to the repository monitor. The instructions apply to all logins from the node on which the Stone repository monitor is running.

- ▶ For additional information about the GemStone administrative logins, see Chapter 6, “User Accounts and Security.”
- ▶ For additional information about logging in from a remote node, see Chapter 3, “Connecting Distributed Systems.”

This section begins with a brief discussion of environmental variables, and then presents two examples. The first example starts a linked application and logs in to GemStone. The second example starts an RPC application, which in turn spawns a separate Gem session process that communicates with the GemStone server.

The examples use Topaz as the application because it is part of the standard GemStone Object Server distribution. Other applications may use different steps to accomplish the same purpose. Some users may prefer to make these steps part of an initialization file.

For an explanation of the difference between linked and RPC sessions, see “Linked and RPC Applications” on page 58.

To Define a GemStone Session Environment

In order to start a GemStone session, the following must be defined through your operating system environment:

- ▶ Where GemStone is installed

All GemStone users must have a `GEMSTONE` environment variable that points to the GemStone installation directory, such as `/installDir/GemStone64Bit3.2.0-x86_64.Linux` (depending on your platform). The directory `$GEMSTONE/bin` should be in your search path for commands. For an example, see the next topic, “To Start a Linked Session”.

- ▶ Which configuration parameters to use

Because each GemStone session can have its own configuration file, some users may need a second environmental variable, such as `GEMSTONE_EXE_CONF`. If no other file is found, the session uses system defaults. For further information, see “How GemStone Uses Configuration Files” on page 266.

To Start a Linked Session

The following steps show how to start a linked application (here, the linked version of Topaz). The steps for setting the `GEMSTONE` environment variable and the operating system path for a session are the same as those given on page 94 for starting a repository monitor. They are repeated here for convenience.

The procedure assumes that the Stone repository monitor has already been started and has the default name `gs64stone`.

Step 1. Set the `GEMSTONE` environment variable to the *full pathname* (starting with a slash) of the directory where GemStone is installed. Ordinarily this directory has a name like `GemStone64Bit3.2.0-x86_64.Linux` (depending on your platform). For example:

```
$ GEMSTONE=/installDir/GemStone64Bit3.2.0-x86_64.Linux
$ export GEMSTONE
```

If you have been using another version on GemStone, be sure you update or delete previous settings of these environment variables:

- ▶ `GEMSTONE`
- ▶ `GEMSTONE_SYS_CONF`
- ▶ `GEMSTONE_EXE_CONF`
- ▶ `GEMSTONE_NRS_ALL`

Step 2. Set your UNIX path. One way to do this is to use one of the `gemsetup` scripts. There is one version for users of the Bourne and Korn shells and another for users of the C shell. These scripts also set your man page path to include the GemStone man pages.

(Bourne or Korn shell)

```
$ . $GEMSTONE/bin/gemsetup.sh
```

or (C shell)

```
% source $GEMSTONE/bin/gemsetup.csh
```

Step 3. Start linked Topaz:

```
% topaz -l
```

Step 4. Set the `UserName` login parameter:

```
topaz> set username DataCurator
```

Step 5. Log in to the Gem session. It will query you for the password.

```
topaz> login
GemStone Password?
[Info]: LNK client/gem GCI levels = 860/860
[Info]: libssl-3.2.0-64.so: loaded
[Info]: User ID: DataCurator
[Info]: Repository: gs64stone
[Info]: Session ID: 6
[Info]: GCI Client Host: <Linked>
[Info]: Page server PID: -1
[Info]: Login Time: 03/14/2014 11:36:47.508 PDT
Gave this VM preference for OOM killer, Wrote to
/proc/6923/oom adj value 4
[03/14/2014 11:36:47.510 PDT]
gci login: currSession 1 linked session
successful login
topaz 1>
```

At this point, you are logged in to a Gem session process, which is linked with the application. The session process acts as a server to Topaz and as a client to the Stone. Information about Topaz is in the manual *GemStone Topaz Programming Environment*.

When you are ready to end the GemStone session, you can log out of GemStone and exit Topaz in one step by invoking the Topaz **exit** command:

```
topaz 1> exit
```

To Start an RPC Session

The following steps show how to start an RPC application (here, the RPC version of Topaz) on the server node. The procedure assumes that the Stone is running under the default name *gs64stone* and that you are already set up to run a GemStone session as described in Steps 1 and 2 of the previous example (“To Start a Linked Session”).

Sessions that login RPC use SRP (Secure Remote Password) and SSL to authenticate passwords for login. If the Gem is running on the server node, the connection reverts to normal socket communication after login completes.

The following steps demonstrate an RPC login from topaz:

Step 1. Use **gslist** to find out if a NetLDI is already running. The default name for the NetLDI is *gs64ldi*.

```
% gslist
Status Version  Owner      Started   Type  Name
-----
exists 3.2.0   gsadmin  Mar 11 12:02  cache  gs64stone~1c9fa07f0412665
exists 3.2.0   gsadmin  Mar 11 12:02  Stone  gs64stone
exists 3.2.0   gsadmin  Mar 11 10:13  Netldi  gs64ldi
```

If necessary, start a NetLDI following the instructions that start on page 98.

Step 2. Unless the NetLDI is running in guest mode with a captive account, set the application login parameters, such as HostUserName and HostPassword, after you start the application. For example:

```
topaz> set hostusername yourUnixId
topaz> set hostpassword yourPassword
```

Step 3. Start the RPC application (such as Topaz), then set the UserName.

```
topaz> set username DataCurator
```

Step 4. Set GemNetId (the name of the Gem service to be started) to gemnetobject. This script starts the separate Gem session process for you. For example:

```
topaz> set gemnetid gemnetobject
```

Step 5. Log in to the GemStone session.

```
topaz> login
GemStone Password?
[Info]: libssl-3.2.0-64.so: loaded
[03/14/2014 11:36:47.777 PDT]
_gci login: currSession 1 rpc gem processId 6943
successful login
topaz 1>
```

At this point, you are logged in through a separate Gem session process that acts as a server to Topaz RPC and as a client to the Stone repository monitor.

When you are ready to end the GemStone session, you can log out of GemStone and exit Topaz by in one step by invoking the Topaz **exit** command:

```
topaz 1> exit
```

To Troubleshoot Session Login Failures

Several factors may prevent successful login to the repository:

- ▶ Your GemStone key file may establish a maximum number of user sessions that can simultaneously be logged in to GemStone. (Note that a single user may have multiple GemStone sessions running simultaneously.) The limit itself is encoded in the keyfile used to start the stone (by default, `$GEMSTONE/sys/gemstone.key`), and reported in the stone log on startup. By default, the Stone log file is `$GEMSTONE/data/gemStoneName.log`. Look for a line like this:

```
SESSION MAX: The licensed concurrent session max is 10.
```

- ▶ The GemStone configuration option `STN_MAX_SESSIONS` (page 299) can restrict the number of logins to fewer than a particular key file allows. An entry in the Stone log file shows the maximum at the time the Stone started. By default, the Stone log file is `$GEMSTONE/data/gemStoneName.log`. Look for a line like this:

```
SESSION CONFIGURATION: The maximum number of concurrent
sessions is 41
```

- ▶ The GemStone configuration option `SHR_PAGE_CACHE_NUM_PROCS` (page 286) restricts the number of sessions that can attach to a particular shared page cache. This number can be different on each node, depending on the configuration file that is read

by the process that starts the cache. On the node where the Stone runs, one of this number is used by the Stone, the shared page cache monitor, each GcGem (garbage collection) session, each Stone AIO page server, the page manager, the SymbolGem, and each free frame page server. On other nodes, the Stone's page server and the shared page cache monitor each use one. For details, see "To Set the Page Cache Options and the Number of Sessions" on page 31. Check the Stone's log for warnings that the value requested for SHR_PAGE_CACHE_NUM_PROCS has been adjusted to match your system's configuration.

- ▶ The UNIX kernel must provide two semaphores for each session that wants to attach to the shared page cache. See "Reviewing Kernel Tunable Parameters" on page 29.
- ▶ The UNIX kernel file descriptor limit can restrict the number of sessions, and GemStone executables attempt to raise that limit. For information, see the discussions of "Estimating File Descriptor Needs" on page 29 (for the Stone) and page 60 (for Gems). On some operating systems, you can examine the kernel limit by invoking **limit**.
- ▶ The owner of the Gem or a linked application process must have write access to the extent file and to the shared page cache. Use the UNIX command **ipcs -m** to display permissions, owner, and group for shared memory. For example:

```
server% ipcs -m
IPC status from <running system> as of Mon March 10 16:21:08
PDT 2014
T      ID      KEY          MODE          OWNER        GROUP
Shared Memory:
m      25089    0x4c000ed5  --rw-rw----  gsadmin      users
```

Typical problems occur with linked applications, which may be installed without the S bit and therefore rely on group access to the shared page cache and the repository.

- ▶ If the session is using a separate (RPC) gem process, see "Troubleshooting Remote Logins" on page 90.

4.5 How To Identify Sessions Logged In

Privileges required: SessionAccess.

To identify the sessions currently logged in to GemStone, send the message `System class>>currentSessionNames`. This message returns an array of internal session numbers and the corresponding UserId. For example:

```
topaz 1> printit
System currentSessionNames
%
session number: 2    UserId: GcUser
session number: 3    UserId: GcUser
session number: 4    UserId: SymbolUser
session number: 5    UserId: DataCurator
```

The session number can be used with other System class methods to stop a particular session or to obtain its UserProfile. See `stopSession:aSessionId`,

`terminateSession:aSessionId timeout:seconds` and
`userProfileForSession:aSessionId`.

NOTE

Be aware that it may take as long as a minute for a session to terminate after you send `stopSession`. If the Gem is responsive, it usually terminates within milliseconds. However, if a Gem is not active (for example, sleeping or waiting on I/O), the Stone waits one minute for it to respond before forcibly logging it out. You can bypass this timeout by sending `terminateSession:timeout`:

The method `System class>>descriptionOfSession:aSessionId` returns an array of descriptive information by which you can trace the session name to a particular person: the second element shows the operating system process id (pid), and the third element shows the name of the node on which it is running. In this example, the DataCurator session is running on "node1" as pid 3010:

```
topaz 1> printit
System descriptionOfSession: 5
%
an Array
  #1 an UserProfile
  #2 3010
  #3 node1
  ...
```

For details about these methods and the information returned, see the class and method comments in the image.

4.6 How To Shut Down the Object Server and NetLDI

Privileges required: SystemAccess and SystemControl.

To shut down GemStone from UNIX, first make sure that all user sessions have logged out. One way to find out about other user sessions is to send the message `currentSessionNames` to System. For example, using Topaz:

```
topaz 1> printit
System currentSessionNames
%
session number: 2   UserId: GcUser
session number: 3   UserId: GcUser
session number: 4   UserId: SymbolUser
session number: 5   UserId: DataCurator
```

The SymbolUser and GcUser sessions are system session and will be shut down cleanly when the stone is shut down. The above example includes session 5, which is the user executing the example code.

After all user sessions have logged out, use the **stopstone** command, which performs an orderly shutdown in which all committed transactions are written to the extent files.

```
% stopstone [gemStoneName] [-i]
```

If you do not supply the name of the Stone repository monitor, **stopstone** prompts you for one. The default name during startup was `gs64stone`. If necessary, use **gslis** (page 314) to find the name.

The **-i** option aborts all current (uncommitted) transactions and terminates all active user sessions. If you do not specify this option and other sessions are logged in, GemStone will not shut down and you will receive a message to that effect.

stopstone prompts you to supply a GemStone username and password. The user must have the SystemControl privilege (initially, this privilege is granted to System-User and DataCurator). For details about user accounts and privileges, see Chapter 6.

There is a similar command to shut down the NetLDI network service.

```
% stopnetldi [netLdiName]
```

For more information about the **stopstone** and **stopnetldi** commands, refer to Appendix B, "GemStone Utility Commands."

If you are logged in to a GemStone session, you can invoke `System class>>shutDown`, which also requires the SystemControl privilege.

CAUTION

*If you must halt a specific Gem session process or GemStone server processes, be sure to use only **kill** or **kill -term** so that the Gem can perform an orderly shutdown.*

*Do NOT use **kill -9** or another uncatchable signal, which may not result in a clean shutdown or may cause the Stone repository monitor to shut down when you intended to kill only a Gem process. If for some reason you need to send **kill -9** to a shared page cache monitor, use **ipcs** and **ipcrm** to identify and free the shared memory and semaphore resources for that cache. If you send **kill -9** to a Stone, use **ipcs** to determine whether **ipcrm** should be invoked.*

4.7 How To Recover from an Unexpected Shutdown

GemStone is designed to shut down in response to certain error conditions as a way of minimizing damage to the repository. If GemStone stops unexpectedly, it probably means that one of the following situations has occurred:

- ▶ Disk failure
- ▶ Shared page cache monitor failure
- ▶ Fatal error detected by a Gem
- ▶ File system corruption
- ▶ Power failure
- ▶ Operating system crash

When GemStone shuts down unexpectedly, check the message at the end of the Stone log file to begin diagnosing the problem. Unless you have set up an environment variable, or specified another file on the **startstone** command line, the Stone log is `$(GEMSTONE)/data/gemStoneName.log`.

The \$GEMSTONE/data directory also contains log files for the Stone child processes. The child processes have log names formed from *gemStoneName*, the process id, and a descriptive abbreviation. For instance: Once the problem is identified, your recovery

gs64stone.log	Stone repository monitor
gs64stone_14033admingcgem.log	Admin GemAdmin Gem
gs64stone_2963pcmon.log	Shared page cache monitor
gs64stone_2967pgsvrff.log	Free Frame page server
gs64stone_2984pgsvraio.log	AIO page server
gs64stone_2987pagemanager.log	Page Manager
gs64stone_2992reclaimcgem.log	Reclaim GemReclaim Gem
gs64stone_2994symbolgem.log	SymbolGem

strategy should take into account the interdependence of GemStone system components. For instance, if an extent becomes unavailable, to restart the system and recover you may have to kill the Stone repository monitor if it is still running. The **stopstone** command won't work in this situation, since the orderly shutdown process requires the Stone to clean up the repository before it stops.

Normal Shutdown Message

If you see a shutdown message in the system log file, GemStone has stopped in response to a **stopstone** command or a Smalltalk System shutdown method:

```
--- 03/21/14 13:00:43 PDT ---
LoginsSuspended is set to 1 by DataCurator from session 5

SHUTDOWN command was received from user DataCurator session 5
gem processId 29188.
Waiting for aiowrites to complete
Waiting for NetRead thread to stop

Now stopping GemStone.
```

After a normal shutdown, restart GemStone in the usual manner. For instructions, see "How To Start the GemStone Server" on page 93 of this chapter.

Disk Failure or File System Corruption

GemStone prints several different disk read error messages to the GemStone log file. For example:

```
Repository Read failure,
fileName = !#dbf!/users/gemstone/data/extent0.dbf
PageId = 94
File = /users/gs64stone/data/extent0.dbf
too few bytes returned from read()
DBF Operation Read; DBF record 94, UNIX codes: errno=34,...
"A read error occurred when accessing the repository."
```

If you see a message similar to the above, or if your system administrator identifies a disk failure or a corrupted file system, try to copy your extents to another node or back them up immediately. The copies may be bad, but it is worth doing, just in case. If you're lucky, you may be able to copy them back after the underlying problem is solved and start again with the current committed state of your repository.

Otherwise, you may need to restore the repository. For details, see the restore procedures in Chapter 9.

Shared Page Cache Error

If you find a message similar to the following in the GemStone log, the shared page cache (SPC) monitor process (`shrpcmonitor`) died. The SPC monitor log, `$GEMSTONE/data/gemStoneName_pcmomnnnn.log`, may indicate the reason.

```
--- 03/10/14 15:07:19 PDT ---  
The stone's connection to the local shared cache monitor was  
lost.  
Error Text: 'Network partner has disconnected.'
```

The unexpected shutdown of a Gem process may, in rare cases, result in a “stuck spin lock” error that brings down the shared page cache monitor and the Stone. GemStone uses spin locks to coordinate access to critical structures within the cache. In most cases, the monitor can recover if a Gem dies while holding a spin lock, but not all spin locks can be recovered safely. Stuck spin locks may result from a Gem crash, but a typical cause is the use of `kill -9` to kill an unwanted Gem process. If you must halt a Gem process, be sure to use only `kill` or `kill -TERM` so that the Gem can perform an orderly shutdown.

Use `startstone` to restart GemStone. For instructions, see “How To Start the GemStone Server” on page 93.

Fatal Error Detected by a Gem

If a Gem session process detects a fatal error that would cause it to halt and dump a core image, the Stone repository monitor may do the same when it is notified of the event. This response on the part of the Stone is configurable through the `STN_HALT_ON_FATAL_ERR` option (page 294). When that option is set to True and a Gem encounters a fatal error, the Stone prints a message like this in its log file:

```
Fatal Internal Error condition in Gem  
when halt on fatal error was specified in the config file
```

By default, `STN_HALT_ON_FATAL_ERR` is set to False. That setting causes the Stone to attempt to keep running if a Gem encounters a fatal error; it is the recommended setting for GemStone in a production system. You can set `STN_HALT_ON_FATAL_ERR` to True during development and testing to provide additional checks for potential risks.

Some Other Shutdown Message

In the event of other shutdown messages in the GemStone log:

1. Consider whether the shutdown might have been caused by a disk failure or a corrupt file system, especially if you see an unexpected message such as `Object not found`.

If you suspect one of these conditions, start with a page audit of the repository file (see “How To Audit the Repository” on page 120).

If the page audit fails, refer to “Disk Failure or File System Corruption” on page 107 of this chapter, and consult your operating system administrator.

If the audit succeeds, continue to the next step.

2. If you don't suspect disk failure or a corrupt file system, try using **startstone** to restart GemStone. For instructions, see “How To Start the GemStone Server” on page 93.
3. If the restart fails, you may have to restore the repository. For details, see the restore procedures in Chapter 9.

No Shutdown Message

If the GemStone log doesn't contain a shutdown message, there has probably been a power failure or an operating system crash. In that event, the Stone repository monitor automatically recovers committed transactions the next time it starts. Use **startstone** to restart GemStone, as described on page 93. For information about the **startstone** command, see page 326.

4.8 How To Bulk-Load Objects

During bulk loading of objects into the repository, it may be desirable to make the following changes:

- ▶ You may need to commit incrementally, but if so, commit as seldom as possible without running out of memory. 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.

To address this concern, increase the `GEM_TEMPOBJ_CACHE_SIZE` configuration option. The size of each transaction (the number of 16 KB pages written) should be approximately 1/3 to 1/2 the size of `GEM_TEMPOBJ_CACHE_SIZE`, and no more than 1/4 to 1/2 the size of the shared page cache.

- ▶ Disable epoch garbage collection, if it was enabled (it is disabled by default). This step saves the CPU time ordinarily devoted to scanning for dereferenced objects. To do this, log in as GcUser and set `#epochGcEnabled` to `False`.
- ▶ Alternatively, you can increase performance during bulk loads by adding the following entries to your configuration file:

```
STN_TRAN_LOG_DIRECTORIES = /dev/null, /dev/null;
STN_TRAN_FULL_LOGGING = TRUE;
```

For information about these configuration file options, see pages 304 and 303, respectively.

NOTE

Be aware that using `/dev/null` for the tranlog directories will prevent you from being able to restore tranlogs in the event of a system failure.

4.9 Considerations for Large Repositories

GemStone/S 64 allows you to define a very large shared page cache, thereby enabling you to run very large repositories. This section presents special considerations that apply to large repositories.

Loading the object table at startup

When starting the repository, the object table is not loaded into memory, and initial accesses can take an excessively long time. If you encounter degraded application performance for a period after restarting the Stone, you may want to start up using cache warming. There is an initial heavy I/O load at Stone startup with cache warming, but subsequent application performance should be consistent with your normal application performance. You may choose to preload just the object table pages, which are most important for performance. Alternatively, you can also preload data pages, which will improve performance if you have a large cache with a relatively fixed working set of data pages.

There are two ways to perform cache warming on startup.

- ▶ The configuration parameters `STN_CACHE_WARMER` and `STN_CACHE_WARMER_SESSIONS` enable and configure cache warming, respectively. See page 289 for more information.
- ▶ Alternatively, you may run the `startcachewarmer` utility. For details about `startcachewarmer`, see page 319.

Making efficient use of remote caches

When running a system on which many users log in simultaneously, consider using remote caches so that you don't need to run all Gem processes on the same machine. There are a couple of ways to optimize this. The following configuration options are of particular interest:

- ▶ To allow Gems to make more efficient use of the large cache, set the `GEM_PGSRV_FREE_FRAME_CACHE_SIZE` configuration option (page 279) to increase the size of the Gem free frame cache. For example:

```
GEM_PGSRV_FREE_FRAME_CACHE_SIZE = 25;
```
- ▶ To improve performance on remote caches, set `GEM_PGSRV_UPDATE_CACHE_ON_READ` (page 279) to `True` so that remote Gem sessions will update their local caches. For example:

```
GEM_PGSRV_UPDATE_CACHE_ON_READ = TRUE;
```

Disk Space and Commit Record Backlogs

Sessions only update their view of the repository when they commit or abort. The repository must keep a copy of each session's view so long as the session is using it, even if other sessions frequently commit changes and create new views (commit records). Storing the original view and all the intermediate views uses up space in the repository, and can result in the repository running out of space. To avoid this problem, all sessions in a busy system should commit or abort regularly.

For a session that is not in a transaction, if the number of commit records exceeds the value of `STN_CR_BACKLOG_THRESHOLD`, the Stone repository monitor signals the session to abort by signaling `TransactionBacklog` (also called “sigAbort”). If the session does not abort, the Stone repository monitor reinitializes the session or terminates it, depending on the value of `STN_GEM_LOSTOT_TIMEOUT`.

Sessions that are in transaction are not subject to losing their view forcibly. Sessions in transaction enable receipt of the signal `TransactionBacklog`, and handle it appropriately, but it is optional. It is important that sessions do not stay in transaction for long periods in busy systems; this can result in the Stone running out of space and shutting down. However, sessions that run in automatic transaction mode are *always* in transaction; as soon as they commit or abort, they begin a new transaction. (For a discussion of automatic and manual transaction modes, see the “Transactions and Concurrency Control” chapter of the *GemStone/S 64 Bit Programming Guide*.)

To avoid running out of disk space, we recommend that you use *manual transaction mode* whenever possible. To enter manual transaction mode:

```
topaz> printit
System transactionMode: #manualBegin
%
```

At the point that this session needs to commit a change, begin a transaction manually, then make the changes:

```
topaz> printit
System beginTransaction .
AllUsers addNewUserWithId: #Jane password: 'gemstone' .
System commitTransaction
%
```

After you commit (or abort) the transaction, your session will return to waiting outside of a transaction.

Handling signals indicating a commit record backlog

Even in manual transaction mode, it is possible to cause a commit record backlog, depending on how your system is configured. Sessions should ensure that they commit or abort regularly, or set up sigAbort handlers to abort when requested by the Stone. A sigAbort handler may be as simple as this:

Example 4.1 sigAbort handler

```
Exception
installStaticException:
[ :exception :GSdictionary :errID :array |
  System abortTransaction.
  System enableSignaledAbortError).
```

Note that a session that is entirely idle does not become aware of the signal to abort, and may timeout and be terminated by the stone in spite of the handler. If your application may have idle sessions, we recommend setting up a timer that causes regular aborts when the session is otherwise idle.

Sessions that are in transaction, and therefore immune from the `sigAbort` mechanism, may also be signaled when there is a commit record backlog. When the number of commit records exceeds the value of `STN_CR_BACKLOG_THRESHOLD`, and the session holding the oldest commit record is in transaction, the Stone repository monitor signals the session by sending `TransactionBacklog`. The session then has the opportunity to perform a `continueTransaction` to update its view of unmodified objects. It may also commit or abort. Unlike `sigAbort`, the session can choose to ignore this message and will not receive further signals from the stone.

Example 4.2 finishTransaction handler

```
Exception
  installStaticException:
    [ :exception :GSdictionary :errID :array |
      System continueTransaction.
      System enableSignaledFinishTransactionError).
```

For more information on these signals, see the *Programming Guide* for GemStone/S 64 Bit.

Monitoring GemStone

This chapter tells you:

- ▶ Where to look for the log files that GemStone/S 64 Bit processes create
- ▶ How to audit the repository
- ▶ How to analyze the repository contents
- ▶ How to monitor the performance of the GemStone server and its clients using GemStone Smalltalk methods

If you decide to keep a GemStone session running for occasional use, be careful not to leave it in an active transaction. A prolonged transaction can cause an excessive commit record backlog and undesirable repository growth, until you either commit or abort.

NOTE

Monitoring on active repositories should be done in manual transaction mode or transactionless mode. For details on entering and using manual transaction mode, see page 110.

5.1 GemStone System Logs

In addition to transaction logs, GemStone creates three types of log files:

- ▶ Logs for GemStone object server processes (page 114)
- ▶ Logs for processes related to individual GemStone sessions (page 118)
- ▶ Logs for GemStone network server processes, NetLDIs (page 119), and for logsenders and logreceivers (page 119)

If a GemStone server is running, you can use the **gslist** utility to locate its logs. Use **gslist -x** to display the location of the current log file for Stones, NetLDIs, logsenders, logreceivers, and the shared page cache monitors.

The logs for the AIO page servers, free frame page servers, SymbolGem, Page Manager, and Admin and Reclaim Gems are in the same location as the corresponding Stone's log.

WARNING

The Stone writes several files to the `/opt/gemstone/locks` or equivalent directory (as discussed on page 30). These lock files are usually regenerated if deleted, but we recommend not removing them manually. Use `gslist -c` to clear out unnecessary lock files.

For more about the `gslist` command, see page 314.

GemStone Server Logs

The Stone repository monitor and its child processes each create a log file in a single location. By default, the files are in `$GEMSTONE/data` and have a name beginning with the Stone name. Table 1 shows typical log names for a Stone with the default name of `gs64stone`. Log names for child processes also include the process id and a descriptive suffix.

Table 1 Representative Log Names for GemStone Server Processes

<code>gs64stone.log</code>	Stone repository monitor
<code>gs64stone_14033admingcgem.log</code>	Admin Gem
<code>gs64stone_2963pcmon.log</code>	Shared page cache monitor
<code>gs64stone_2967pgsvrff.log</code>	Free frame page server
<code>gs64stone_2984pgsvraio.log</code>	AIO page server
<code>gs64stone_2987pagemanager.log</code>	Page Manager (Stone thread)
<code>gs64stone_2992reclaimcgem.log</code>	Reclaim Gem
<code>gs64stone_2994symbolgem.log</code>	Symbol Gem

Several factors can alter the name and location of these logs. The precedence is

1. A path and filename supplied by `startstone -l logFile.logfile` may be a filename, or a relative or absolute path and filename, to which the account starting the Stone has write permission. If `logFile` is a filename without a path, `logFile` is created in the current directory. Logs for the child processes in Table 1 are placed in the same directory.
2. A path and filename specified by the `GEMSTONE_LOG` environment variable. As with `startstone -l`, this may be set to a filename or to a relative or absolute path and filename. Child process log files are created in the same location.
3. `$GEMSTONE/data/gemStoneName.log`.

Log file deletion on shutdown

When a GemStone server shuts down, some server process log files are deleted. Other files are retained in case they are needed later for tuning or problem diagnosis. Table 2 details the specific behavior for each process's logs, and how to change that behavior in cases where the behavior is configurable.

Log files are never deleted if the process exits abnormally. In such cases, the log file may be requested when you contact GemStone Technical Support.

Since some logs are not deleted, if you restart GemStone often, you may need to manually remove log files periodically.

Table 2 Log file handling by process type

Stone	The same log file is appended to on restart. Log file is never deleted.
Shared Page Cache Monitor	A new log file is created on restart, including the process PID. Log file is never deleted on exit.
Page Manager	(Not an independent process) A new log file is created when the stone is restarted, including the Stone's PID. Log file is never deleted on exit.
Symbol Gem	A new log file is created on restart, including the process PID. Log file is never deleted on exit.
Free frame page servers, AIO page serv- ers	A new log file is created on restart, including the process PID. Log file is deleted by default on normal exit. To prevent this log from being deleted on normal exit, edit <code>\$GEMSTONE/sys/runpgsvrmain</code> to look like this: # unset \$GEMSTONE_KEEP_LOG. This will affect all page servers, free frame, AIO page servers, and page servers for remote gems.
Admin Gem	A new log file is created on restart, including the process PID. Log file is deleted by default on normal exit. To prevent this log from being deleted on normal exit, edit <code>\$GEMSTONE/sys/runadmingcgem</code> to look like this: # unset \$GEMSTONE_KEEP_LOG.
Reclaim Gem	A new log file is created on restart, including the process PID. Log file is deleted by default on normal exit. To prevent this log from being deleted on normal exit, edit <code>\$GEMSTONE/sys/runreclaimcgem</code> to look like this: # unset \$GEMSTONE_KEEP_LOG.

Stone Log

The log for the Stone repository monitor is always appended to, and is therefore cumulative across runs by default. This log is the first one you should check when a GemStone system problem is suspected. In addition to possible warnings and error messages, the log records the following useful information:

- ▶ The GemStone version.
- ▶ The configuration files that were read at startup and, if `DUMP_OPTIONS` is set to True (page 276), the resulting Stone configuration.
- ▶ Each startup and shutdown of the Stone, the reason for the shutdown, and whether recovery from transaction logs was necessary at startup.
- ▶ Each expansion of a repository extent and its current size.
- ▶ Each opening of a new transaction log.
- ▶ Each startup and shutdown of each GcGem session, and the corresponding processId.

- ▶ Each #abortErrLostOtRoot sent to a Gem.
- ▶ Each suspension and resumption of logins.
- ▶ Certain changes to the login security system.
- ▶ Each time a backup is started and when the backup is completed.

Admin Gem Logs

Each time the Stone repository monitor starts an administrative garbage collection session (Admin Gem) process, a new log is created in the same location as the Stone's log. The log name is formed using the pattern:

```
stoneName_PIDadminggem.log
```

where *stoneName* is the name of the Stone, and *PID* is the process Id of the Admin Gem process.

This log shows the startup value of the Admin Gem parameters that are stored in GcUser's UserGlobals, and any changes to them, and records other Admin Gem functions.

Reclaim Gem Log

Each time the Stone repository monitor starts a reclaim garbage collection session (Reclaim Gem) process, a new log is created in the same location as the Stone's log. The log name is formed using the pattern:

```
stoneName_PIDreclaimggem.log
```

where *stoneName* is the name of the stone, and *PID* is the process Id of the Reclaim Gem process.

This log shows the startup value of the Reclaim Gem parameters that are stored in GcUser's UserGlobals, and any changes to them, and records other Reclaim Gem functions.

Shared Page Cache Monitor Log

The log for the shared page cache monitor on the Stone's machine is located in the same directory as the Stone's log. This log file has a name of the form

```
stoneName_PIDpcmon.log
```

Check this log if other messages refer to a shared page cache failure.

When a session logs in from another node, a log is created for the shared page cache monitor on the remote node. This log is located by default in the home directory of the account that started the Stone, but this location can be modified by environment variable settings. The default name is of the form

```
startshrpcmonPIDNode.log
```

where *PID* is the process Id of the monitor process, and *Node* is the name of the remote node.

Among the items included in the log for the shared page cache monitor are:

- ▶ Its configuration (for remote nodes, this may be different from the configuration on the Stone's node).
- ▶ The number of processes that can attach (which can limit the number of logins).
- ▶ The UNIX identifiers for the memory region and the semaphore array (these identifiers are helpful in the event you must remove them manually using the `ipcrm` command).

Free Frame Page Server Log

One or more free frame page servers started up on repository startup. Each one has an individual log file, located in the same directory as the log for the shared page cache monitor. These log files have names of the form

```
stoneName_PIDpgsvrff.log
```

where *stoneName* is the name of the stone, and *PID* is the process Id of the free frame page server process. These logs ordinarily are not of interest unless they contain an error message.

AIO Page Server Log

One or more AIO page servers are started up on repository startup. Each one has an individual log file, located in the same directory as the log for the shared page cache monitor. These log files have names of the form

```
stoneName_PIDpgsvraio.log
```

where *stoneName* is the name of the stone, and *PID* is the process Id of the AIO page server process. These logs ordinarily are not of interest unless they contain an error message.

Page Manager Log

The Page Manager is a thread in the Stone, and is not a separate process, but it writes to a separate log for ease of maintenance. The Page Manager log is located in the same directory as the log for the shared page cache monitor. This log file has a name of the form

```
stoneName_PIDpagemanager.log
```

where *stoneName* is the name of the stone, and *PID* is the process Id of the Stone process. Ordinarily these logs are not of interest unless errors occur or tuning is required.

Symbol Gem Log

The Symbol Gem log is located in the same directory as the Stone's log. This log file has a name of the form

```
stoneName_PIDsymbolgem.log
```

where *stoneName* is the name of the stone, and *PID* is the process Id of the Symbol Gem process. Ordinarily these logs are not of interest unless an error has occurred.

Logs Related to Gem Sessions

Except for linked session that are running on the same node as the Stone, login depends on NetLDI services to spawn one or more supporting processes. In each case, the NetLDI creates a log file that includes in its name the identity of the node on which the process is running.

Linked logins do not have separate log files. Log file output is sent to stdout of the linked process.

All RPC logins spawn a Gem session process.

For RPC logins where the Gem is not on the same node as the Stone, or for linked logins that are not on the same node as the Stone, the following additional processes are also spawned:

- ▶ A page server (for the session) to access a repository extent on the server node.
- ▶ A page server (for the Stone) to start or access a shared page cache on the client's node.
- ▶ A shared page cache monitor (for the Stone) to manage the cache on the client's node.

By default, the log files for these processes are located in the home directory of the account that owns the corresponding process. For the Gem session process and the page server on the server node, that account ordinarily is the application user. For the shared page cache monitor and page server on the client node, that account is the one that invoked **startstone**.

You can change the default location by setting **#dir** or **#log** in the `GEMSTONE_NRS_ALL` environment variable for the NetLDI itself or for individual clients (see "To Set a Default NRS" on page 78). Alternatively, when you log in to GemStone, you can specify a different network resource string (NRS) in your login parameters.

Table 3 shows typical log names for session-related processes, given a Stone and repository on *node1* with a login from a Gem session process on *node2*.

Table 3 Typical Logs Supporting Gem Sessions

Typical Name	GemStone Process
<code>gemnetobject27853node2.log</code>	Gem session process on <i>node2</i> (serves an RPC session)
<code>pgsvrmain27819node2.log</code>	Page server on <i>node2</i> that the repository monitor uses to create and access its shared page cache on <i>node2</i>
<code>startshrpcmon27820node2.log</code>	Shared page cache monitor on <i>node2</i>
<code>pgsvrmain12397node1.log</code>	Page server on <i>node1</i> that the Gem session process uses to access the repository extents on <i>node1</i>

If a process shuts down normally, the log file may be automatically removed. (See Table 2 on page 115 for specific behavior by process type.) Log files are never deleted if the process shuts down abnormally. This way, the log files that remain may provide helpful diagnostic information.

If you want to retain the log even when a Gem session process exits normally, edit the scripts according to the instructions in Table 2 on page 115. If the NetLDI on the client node has a separate \$GEMSTONE directory, edit the appropriate scripts in the client's installation directory.

NetLDI Logs

Each NetLDI creates a log file (*netLdiName.log*) in `/opt/gemstone/log` (or an equivalent; see page 30) on the node on which it runs. This location and name can be overridden by including the option `-llogname` when starting the NetLDI. Each NetLDI you start with the same name appends to one log, so it's a good idea to remove outdated messages occasionally.

By default, the NetLDI log contains only configuration information and error messages. The configuration information reflects the environment at the time the NetLDI was started and the effect of any authentication switches specified as part of the `startnetldi` command.

In some cases it is helpful to log additional information by starting the NetLDI in debug mode (`startnetldi -d`). The debug log records each exchange between the NetLDI and a client. Because the log becomes much larger, you probably won't want to use this mode routinely.

Logsender and logreceiver logs

The logsender and logreceiver processes are started only if you are setting up a hot standby system.

Each logsender and logreceiver creates a log file in `/opt/gemstone/log` on the node on which it runs. The log file's name, by default, is `logsender_listeningPort.log` or `logreceiver_listeningPort.log`. This location and name can be overridden by including the option `-llogname` when starting the logsender or logreceiver. Each logsender and logreceiver you start with the same log file name and path — explicitly specified or the default — appends to one log, so you should periodically remove outdated messages.

Localizing timestamps in 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 the criteria are not met, the default date format based on the system's LOCALE is used, or otherwise the US-centric date format.

Programmatically adding messages to logs

You can write a message to the stone log using

```
System addAllToStoneLog: aString
```

To write a message to the gem log for the current session, use

```
GsFile gciLogServer: aString
```

The following writes a log message to the GCI client, such as the topaz console:

```
GsFile gciLogClient: aString
```

5.2 How To Audit the Repository

This section describes two levels of checks that you can perform on the repository.

- ▶ A *page audit* (page 120) typically is invoked to ensure page-level consistency after some kind of system failure, such as a read-write error or a cache coherency error. In these cases, a successful page audit indicates that the problem did not affect the committed repository. GemStone must be halted when you perform a page audit.
- ▶ An *object audit* (page 122) checks the consistency of the repository at the object level. An object audit can be performed as part of routine maintenance and is performed while GemStone is running.

To Perform a Page Audit

Page audits allow you to diagnose problems in the system repository by checking for consistency at the page level.

The **pageaudit** utility can be run only on a repository that is not in use.

Pageaudit scans the rootpages in a repository, pages used in the bitmap structures referenced by the rootpage, and all other pages (including data pages) to confirm page-level consistency. It does not check that the data on data pages is valid. For that, you need to run object audit; see “To Perform an Object Audit and Repair” on page 122.

To check for page-level problems, run **pageaudit** on the repository defined in your ordinary GemStone configuration by issuing this command at the operating system level:

```
% pageaudit [gemStoneName] [-e exeConfig] [-z systemConfig] [-f] [-d]
  [-l logfile] [-h]
```

where:

- ▶ *gemStoneName* is the name of the GemStone repository monitor.
- ▶ *systemConfig* is the system configuration file (page 266).
- ▶ *exeConfig* is the executable configuration file (page 268).
- ▶ *logfile* is the name of the output file.

All of these arguments are optional in a standard GemStone configuration. If these options are not supplied, **pageaudit** uses `gs64stone-audit` for `gemStoneName` and writes output to `stdout`.

You can get details of the available options for **pageaudit** by executing `pageaudit -h`, which returns **pageaudit** usage. **pageaudit** is also described on page 316.

In addition to the audit results, **pageaudit** prints status updates as it is running, and the output includes repository statistics to the screen. For example:

```

GemStone is starting a page audit of the Repository.
Finished auditing reserved pages in extent 0.
Finished pages past end
Begin auditing current checkpoint.
Finished auditing checkpoint bitmaps.
Finished auditing scavengable pages.
Start auditing object table pages.
Finished auditing object table pages.
Finished auditing commit records.
Finished auditing alloc pages shadowed.
Begin auditing data pages. Count=970
Finished auditing data pages

PAGE AUDIT STATISTICS paris sun4u (Solaris 2.10 Generic 141444-09)
- 03/12/2014 09:57:26.234 PDT

16384 bytes = 1 GemStone Page
1048576 bytes = 1 Mbytes
Repository Size                               53 Mbytes
Data Pages                                     6 Mbytes
Meta Information Pages                         1 Mbytes
Shadow Pages                                  0 Mbytes
Free Space in Repository                       44 Mbytes
**** Number of differences found in page allocation = 0.
Page Audit of Repository completed successfully.

```

The report contains the following statistics:

Repository Size

The total physical size of the repository; this is the same size that the operating system reports for an extent file.

Data Pages

This includes all pages referenced from the object table.

Meta Information Pages

Pages that contain only internal information about the repository, such as the object table.

Shadow Pages

Pages scheduled for scavenging by the reclaim task.

Free Space in Repository

Free space in the repository is computed as the number of free pages times the size of a page (16 KB). That value reflects the number of pages available for allocation to Gem session processes. It excludes space fragments on partially filled data pages.

If the page audit finds problems, the message to the screen ends with a message like this:

```
----- PAGE AUDIT RESULTS -----
**** NumberOfFreePages = 980 does not agree with audit
      results = 988

**** Problems were found in Page Audit.
**** Refer to recovery procedures in System Administrator's Guide.
```

If there are problems in the page audit, you will need to restore the repository file from backups. (See the section “How to Restore from Backup” on page 200.)

To Perform an Object Audit and Repair

Privileges required: SystemControl.

Object audits check the consistency of the repository at the object level. Starting with Object Table, each object is located and validated.

Object audit is performed using multiple threads (lightweight sessions), and can be configured to perform as quickly as possible using a large amount of system resources, or configured to use fewer resources and take longer to run.

Object audit should be run from linked Topaz, and on the same machine as the Stone.

Repository >> objectAudit

objectAudit is the normal way to perform the audit. You may have other sessions logged in and running simultaneously, but the audit will impact performance. This audit uses two threads and up to 90% of the CPU.

Repository >> fastObjectAudit

fastObjectAudit is like objectAudit, but is configured to use most or all system resources to complete as quickly as possible. This is useful when running an audit on offline systems.

Repository >> objectAuditWithMaxThreads: *maxThreads*

percentCpuActiveLimit: *aPercent*

This method allows you to specify the exact performance/impact parameters for the audit, if neither objectAudit nor fastObjectAudit is satisfactory for your requirements.

Performing the Object Audit

To perform an object audit:

Step 1. Log in to GemStone using linked Topaz (**topaz -l**).

Step 2. Send one of the audit messages to the repository. For example:

```
topaz 1> printit
SystemRepository objectAudit
%
```

The audit involves a number of checks and specific error messages. Checks include:

- ▶ Object corruption – The object header should contain valid (legal) information about the object's tag size, body size (number of instance variables), and physical size (bytes or OOPs).
- ▶ Object reference consistency – No object should contain a reference to a nonexistent object, including references to a nonexistent class.
- ▶ Identifier consistency – OOPs within the range in use (that is, up to the high-water mark) should be in either the Object Table or the list of free OOPs, and OOPs for objects existing in data pages should be in the Object Table.

If the repository is consistent and no errors are found, the audit will complete with the line:

```
Object Audit: Audit successfully completed; no errors were
detected.
```

Otherwise, the reasons for failure with the specific problems found are reported to standard output

Error Recovery

If an object audit reports errors, these issues should be addressed. You may want to contact GemStone Technical Support for advice.

The following are general approaches to errors from object audit.

Collect and reclaim garbage and retry

If errors are reported during the object audit, you may wish to perform a `markForCollection` and `reclaimAll` and repeat the object audit. This may clear up problems if the object (s) that is (are) corrupt are not referenced from any live objects. Whether this is useful will depend on the particular errors reported.

Restore from backup

The safest approach when you find object audit errors is to restore from backup. GemStone recommends that you make regular backups, run in full transaction logging mode, and archive transaction logs as needed to recover. This would allow you to recover at any time from unexpected problems such as repository corruption.

If you do not have the set of backups and transaction logs that would allow you to restore from a backup and recover later transactions, or if you are in partial transaction logging mode, you can still make and restore a backup. Backups made using `fullBackupTo`; when

restored, rebuild the internal data structures. Depending on the specific problems found in audit, this may clear up the problem.

Manual repair of invalid object references

Invalid object references can be repaired manually, if you know what the missing data should be, or if the referenced data is not important.

Use the Topaz object identity specification format *@identifier* to substitute nil or an appropriate reference for an invalid reference.

For example, given an instance of Array with the OOP 51369729, if the element at slot 3 is an object that does not exist, it can be repaired by setting the reference to nil using the following expression:

```
topaz 1> send @51369729 at: 3 put: nil
```

Repository repair

You can have GemStone attempt appropriate repairs during the re-scan by invoking `Repository>>repair`. The following repairs illustrate their nature:

- ▶ nil is substituted for an invalid object reference.
- ▶ Class String is substituted for an invalid class of a byte object, class Array for a pointer object, or class IdentitySet for a nonsequenceable collection object.
- ▶ Oops in the Object Table for which the referenced object does not exist are inserted into the list of free Oops. Oops for which an object exists but which are also in the list of free Oops are removed from the free list.

The repair audits the repository, keeping track of errors. After the initial audit completes, each error found is repaired. A descriptive message is displayed for each repair. The repair will commit periodically to avoid memory issues, and log off when it is complete. For example:

```
topaz 1> run
SystemRepository repair
%
Object [20897537] references class [27554561] which does not exist
Object [27551745] references class [27554561] which does not exist
In object [27553281] of class Widget [26374657], the logical size
42 is invalid
In object [27554049] of class Widget [26374657], the object format
1 disagrees with theclass format 0
Object [27554817] references class [27554561] which does not exist

Object Audit: 5 errors were found
Repairing error: BadClassId - Object [20897537] references class
[27554561] which does not exist
  Changing class to String [74753]
Repairing error: BadClassId - Object [27551745] references class
[27554561] which does not exist
  Changing class to IdentitySet [73985]
Repairing error: BadLogicalSize - In object [27553281] of class
Widget [26374657], the logical size 42 is invalid
```

```

    resetting logialSize to 8
Repairing error: BadFormat - In object [27554049] of class Widget
[26374657], the object format 1 disagrees with the class format 0
    Changing class to String [74753]
Repairing error: BadClassId - Object [27554817] references class
[27554561] which does not exist
    Changing class to Array [66817]

[Info]: Logging out at 03/12/2014 09:57:26.234 PDT
ERROR 4061 , The session is terminating abnormally, completed the
repair of 5 objects, forcing logout.

```

5.3 Profiling Repository Contents

Some questions – such as “what is using up all the space in my Repository?” – can only be answered by examining the types and numbers of objects in your repository. To find out this information, you can use methods on `GsObjectInventory`.

The methods in `GsObjectInventory` count all instances of all classes in the repository – or in any collection, or in a hidden set, or in a file of disconnected possible garbage objects – and report the results, ordered by the number of instances or by space consumed.

`GsObjectInventory` performs a multi-threaded scan of the repository, and thus should only be run in session on the same machine as the Stone. To tune the impact of the scan, additional protocol allows you to perform fast scans or to specify the impact levels. For details, see methods in the image.

The scans require the `GcLock`, and so cannot be run while any garbage collection operation is running, nor can garbage collection operations be started while a `GsObjectInventory` scan is going on.

The following code will report the number of instances and the space required for all Classes whose total space requirements are more than 10000 bytes.

```

topaz 1> run
GsObjectInventory profileRepository byteCountReportDownTo: 1000
%
    *** GsObjectInventory byteCountReport printed at: 25/02/2014 20:17:27
    ***
Hidden classes are included in this report.

```

Class	Instances	Bytes
String	22497	8126360
GsNMethod	15289	3005728
Array	18921	2945904
GsMethodDictionary	2570	1292696
Symbol	15146	658624
LargeObjectNode	30	456512
CanonStringBucket	2016	269984
Class	1254	195776
IdentityKeyValueDictionary	1271	172880

SymbolAssociation	3923	157416
ExecBlock	2312	148128
SymbolDictionary	766	141008
IdentityCollisionBucket	1336	131824
SymbolSet	3502	103808
GsClassDocumentation	464	48272
DepListBucket	751	42064
DateTime	634	35528
ClassHistory	625	30176
GsDocText	714	28624
LargeInteger	9	22976
TimeZoneTransition	313	17528
CanonSymbolDict	1	16176
WordArray	13	13920
EqualityCollisionBucket	128	13248

The same profiling with an instance count report is much shorter, since the number of instances, rather than the bytes of space used, limits the results.

```
topaz 1> run
GsObjectInventory profileRepository instanceCountReportDownTo: 10000
%
*** GsObjectInventory instanceCountReport printed at: 25/02/2014
20:17:27 ***
Hidden classes are included in this report.
```

Class	Instances	Bytes
String	22497	8126360
Array	18921	2945904
GsNMethod	15289	3005728
Symbol	15147	658680

These reports include instances of hidden classes - these are classes that are used to implement internal GemStone objects, which are invisible to the image. One such class is LargeObjectNode. Instances of LargeObjectNodes are used to implement the tree structures that underlie large collections. To avoid seeing hidden classes - which will include the space used by the hidden class within the root, public object, profile using the method `profileRepositoryAndSkipHiddenClasses` rather than `profileRepository`.

For more on GsObjectInventory, see the methods in the image.

5.4 Monitoring Performance

As part of your ongoing responsibilities, you may find it useful to monitor performance of the object server or individual session processes.

GemStone includes graphical tools to allow you to record statistics in file and analyze this data graphically. You can also programmatically access these statistics.

A full list of the statistics that are recorded and are available programmatically can be found in the *VSD User's Guide*.

Statmonitor and VSD

GemStone includes the statmonitor utility, which records statistics about GemStone processes to a disk file. You can configure the statistics recorded, how frequently the statistics are collected, and other details. See page 327 for more information on the statmonitor utility.

Both GemStone-specific and operating system statistics are collected. The operating system statistics include general host information as well as information specific to the individual GemStone processes.

We recommend running statmonitor at all times, as it provides a valuable record of many aspects of system behavior.

To view this data, VSD (Visual Statistics Display) graphically displays the statistics. For more details on using VSD, see the *VSD User's Guide*.

Programmatic Access to Cache Statistics

A set of methods on the System class provide a way for you to analyze performance by programmatically examining the statistics that are collected in the shared page cache. This is the same data that is visible using statmonitor and VSD, although statmonitor and VSD can collect additional OS level information. This additional OS level information is also available programmatically; see "Host Statistics" on page 132

A process can only access statistics that are kept in the shared page cache to which it is attached. Sessions that are running on a different node than the Stone use a separate shared cache on that remote node. This means that processes that are on a different node than the Stone, cannot access statistics for the Stone or for other server processes that are attached to the Stone's shared page cache.

Within the shared page cache, GemStone statistics are stored as an array of *process slots*, each of which corresponds to a specific process. Process slot 0 is the shared page cache monitor. On the Stone's shared page cache, process slot 1 is the Stone; on remote caches, slot 1 is the page server for the Stone that started the cache. Subsequent process slots are the page servers, Page Manager, Admin and Reclaim Gems, Symbol Gem, and user Gems. The order of these slots depends on the order in which the processes are started up, and is different on remote caches.

The specific set of statistics is different for each type of process that can attach to the shared page cache. The types of processes are numbered:

- 1 = Shared page cache monitor
- 2 = Stone
- 4 = Page server
- 8 = Gem (including Topaz, GBS, and other GCI applications).

Statistics by name

To obtain the value for a specific statistics for the Stone, the Stone's SPC monitor, or for the current session, use the following methods:

```
System class >> stoneCacheStatisticWithName:
System class >> primaryCacheMonitorCacheStatisticWithName:
System class >> myCacheStatisticWithName:
```

These methods will return the statistics value corresponding to the given name for that process. If the statistics name is not found, it returns nil.

For example, to retrieve the statistics named 'CommitRecordCount' for the Stone:

```
topaz 1> printit
System stoneCacheStatisticWithName: 'CommitRecordCount'.
%
23
```

To retrieve the current session's PageReads:

```
topaz 1> printit
System myCacheStatisticWithName: 'PageReads'.
%
548
```

All statistics for a process

The general way to retrieve statistics is as an array of values. To understand what the value at each index refers to, there are corresponding description methods to return an array of Strings. Matching the index of the statistic name to the index within the values locates the value for that statistic.

Since the statistics are different for the different types of processes, you will need to use corresponding methods to collect the statistics and the descriptions.

For the Stone, the Gem that is running the code, and the Stone's shared page cache monitor, no further information is needed to identify them within the cache, so the following pairs of methods can be used:

```
System cacheStatisticsDescriptionForGem.
System myCacheStatistics.

System cacheStatisticsDescriptionForStone.
System stoneCacheStatistics.

System cacheStatisticsDescriptionForMonitor.
System sharedPageCacheMonitorCacheStatistics.
```


For example, while you would normally use `stoneCacheStatisticForName:`, here is another possible way to get the `CommitRecordCount`:

```
topaz 1> printit
| index |
index := System cacheStatisticsDescriptionForStone indexOf: 'Com-
mitRecordCount'.
System stoneCacheStatistics at: index.
%
```

23

To collect statistics for other Gems, and for page servers, you need to determine the process Id, session Id, or slot of the specific Gem or page server, or the cache name of the Gem. There are a variety of ways you might determine this, but one way is to examine the results of:

```
System cacheStatisticsForAllSlotsShort
```

This method returns the name, process Id, session Id, statistics type, and process slot for each process currently attached to the cache. For example:

```
topaz 1> printit
(System cacheStatisticsForAllSlotsShort) collect:
[:ea | ea printString]
%
```

an Array

#1	anArray('ShrPcMonitor', 7722, 4294967295, 1, 0)
#2	anArray('gs64stone', 7721, 0, 2, 1)
#3	anArray('FreeFrmPgsvr2', 7725, 4294967294, 4, 2)
#4	anArray('AioPgsvr3', 7726, 4294967294, 4, 3)
#5	anArray('pagemgrThread', 7729, 1, 8, 4)
#6	anArray('GcAdmin5', 7734, 2, 8, 5)
#7	anArray('SymbolGem6', 7735, 3, 8, 6)
#8	anArray('GcReclaim6 7', 7733, 4, 8, 7)
#9	anArray('Gem26', 2271, 5, 8, 8)
#10	anArray('Gem27', 16924, 6, 8, 9)

Of course, a Gem may log out between the time you execute this and the time you collect statistics, so be sure that your code handles that condition gracefully.

The methods you use to get the statistics and the corresponding descriptions will depend on how you have determined the specific process you want information about.

By name:

```
System cacheStatisticsForProcessWithCacheName: aString
(You must manually determine the process type)
or
```

```
System cacheStatsForGemWithName: aString.
System cacheStatisticsDescriptionForGem.
```

By operating system Process Id (PID):

```
System cacheStatisticsProcessId: aPid.
System cacheStatisticsDescriptionAt:
    (System cacheSlotForProcessId: aPid).
```

By process slot:

```
System class >> cacheStatisticsAt: aProcessSlot
System class >> cacheStatisticsDescriptionAt: aProcessSlot
```

By session Id:

The page server for a Gem assumes the same sessionId as its Gem.

```
System gemCacheStatisticsForSessionId: aSessionId.
System cacheStatisticsDescriptionForGem.
or
```

```
System cacheStatsForPageServerWithSessionId: aSessionId
System cacheStatisticsDescriptionForPageServer
```

For example, to find an aggregate value for TimeInFramesFromFindFree of all Gems in the system:

```
topaz 1> printit
| gemPids index time |
gemPids := Array new.
System cacheStatisticsForAllSlotsShort do:
    [:anArray |
    (anArray at: 4) = 8 ifTrue:
        [gemPids add: (anArray at: 2)].
    ].
index := System cacheStatisticsDescriptionForGem indexOf:
    'TimeInFramesFromFindFree'.
time := 0.
gemPids do: [:aPid | | stats |
    stats := System cacheStatisticsProcessId: aPid.
    stats ifNotNil: [time := time + (stats at: index)].
].
time
%
0
```

Setting the name for the Gem in the cache

To make it easier for you to track cache statistics for specific Gems, you can explicitly give each Gem a unique name. The method

```
System cacheName: aString
```

sets the name for the current Gem session in the cache statistics, thus making it much easier to read the statistics in VSD.

Set the cache name soon after login. If you are collecting statistics information using statmonitor, information may be logged using the default name for the Gem, and you may have two separate lines of data for the same session.

Session Statistics

In addition to the system-generated statistics listed below, GemStone provides a facility for defining session statistics – user-defined statistics that can be written and read by each session, to monitor and profile the internal operations specific to your application.

There are 48 session cache statistic slots available, with names of the form SessionStat01...SessionStat47.

You can use the following methods to read and write the session cache statistics:

System class >> sessionCacheStatAt: *anIndex*

Returns the value of the statistic at the designated index. *anIndex* must be in the range -2 to 47. Negative values are reserved for internal use.

System class >> sessionCacheStatAt: *anIndex* put: *aValue*

Assigns a value to the statistic at the designated index and returns the new value. *anIndex* must be in the range -2 to 47. Negative values are reserved for internal use.

System class >> sessionCacheStatAt: *anIndex* incrementBy: *anInt*

Increment the statistic at the designated index by *anInt*, and returns the new value. *anIndex* must be in the range -2 to 47. Negative values are reserved for internal use.

System class >> sessionCacheStatAt: *anIndex* decrementBy: *anInt*

Decrement the statistic at the designated index by *anInt*, and returns the new value. *anIndex* must be in the range -2 to 47. Negative values are reserved for internal use.

System class >> sessionCacheStatsForProcessSlot: *aProcessSlot*

Return an array containing the 48 session statistics for the given process slot, or nil if the process slot is not found or is not in use.

System class >> sessionCacheStatsForSessionId: *aSessionId*

Return an array containing the 48 session statistics for the given session id, or nil if the session is not found or is not in use.

Global Session Statistics

In addition to the Gem session statistics, GemStone/S 64 Bit provides global session statistics – user-defined statistics that can be written and read by any Gem on any Gem server. Unlike *session* cache statistics, which are stored in the shared page cache of the machine that the Gem is running on, *global* session statistics are stored in the shared page cache of the Stone. Global session statistics are not transactional. For a given statistic, every session sees the same value, regardless of its transactional view.

There are 48 global cache statistic slots available, with names of the form GlobalStat01...GlobalStat47.

You can use the following methods to read and write the global cache statistics:

System class >> globalSessionStatAt: *aProcessSlot*

Returns the value of the statistic at the designated slot (must be in the range 0..47).

System class >> globalSessionStatAt: *aProcessSlot* put: *aValue*

Assigns a value to the statistic at the designated slot (must be in the range 0..47) and returns the new value. The value must be a SmallInteger in the range of -2147483648 to 2147483647.

System class >> `incrementGlobalSessionStatAt: aProcessSlot by: anInt`
 Increments the value of the statistic at the designated slot by *anInt* and returns the new value of the statistic. The value *anInt* must be a `SmallInteger` in the range of -2147483648 to 2147483647.

Host Statistics

Host Statistics for processes

Process-level statistics require an OS call, which can cause cache statistics to impact performance. These statistics are not part of the information returned by regular cache statistics interface methods. To get this information, use the following methods.

System class >> `hostProcessStatisticsNames`

Returns an array of `Strings` which are the names of the per-process statistics provided by this host.

System class >> `hostStatisticsForMyProcess`

Returns an array of `SmallIntegers` which represent the host statistics for this process. The names of each statistic are returned by the `#hostProcessStatisticsNames` method.

System class >> `hostStatisticsForProcess: processId`

Returns an array of `SmallIntegers` which represent the host statistics for the process with the given process ID. The names of each statistic are returned by the `#hostProcessStatisticsNames`

Specific methods are also available to return the host CPU statistics only:

System class >> `hostCpuStatsForProcessId: anInt`

Return an Array of two integers as follows:

- 1 - user mode CPU milliseconds
- 2 - system mode CPU milliseconds

Both array elements will be -1 if the process slot is out of range or not in use or if this method is not supported for the host architecture.

It is not required that the process with pid *anInt* is attached to the shared page cache or even is a GemStone process. The method will succeed for any process for which the Gem session executing the method has permission to view the target process' CPU usage statistics.

System class >> `hostCpuStatsForProcessSlot: anInt`

For the process using the cache process slot *anInt*, return an Array of two integers as follows:

- 1 - user mode CPU milliseconds used
- 2 - system mode CPU milliseconds used

Both array elements are set to -1 if the process slot is out of range or not in use, or if this method is not supported for the host architecture.

Host Statistics for OS

While most monitoring is of the object server and session processes, it is also useful to monitor the performance of the operating system that is running GemStone. On host

platforms that support it, the following methods return statistics provided by the operating system. This is the same information that is available via statmonitor; for details, see page 327.

System class >> fetchSystemStatNames

Return an array of Strings with the names of the available OS level statistics. The length is host-dependent. If the host system does not support system statistics, this method returns nil.

System class >> fetchSystemStats

Return an array of Numbers corresponding to the names returned by the #fetchSystemStatNames method. The length of the result array is host dependent. While most elements in the result array will be SmallIntegers, the result may also contain other types of Numbers such as SmallDoubles, Floats, LargeIntegers, etc. If the host system does not support system statistics, this method returns nil.

You can also monitoring specific CPU usage for the host using the following method:

System class >> hostCpuUsage

Returns an Array of 5 SmallIntegers with values between 0 and 100 which have the following meanings:

- 1 - Percent CPU active (user + system)
- 2 - Percent CPU idle
- 3 - Percent CPU user
- 4 - Percent CPU system (kernel)
- 5 - Percent CPU I/O wait

On hosts with multiple CPUs, these figure represent the average across all processors. The results of the first call to this method are invalid and should be discarded. Returns nil if the host system does not support collecting CPU statistics.

User Accounts and Security

This chapter also shows you how to perform some common GemStone user administration tasks:

- ▶ How to create and modify user accounts, including passwords, privileges, group memberships, and symbol resolution, and how to control the user's read-write access to objects through the use of *object security policies*.
- ▶ How to set up alternate ways to authenticate logins, such as authenticating by UNIX `userId` or LDAP.
- ▶ How to configure GemStone login security by restricting valid passwords, imposing password and account age limits, and monitoring intrusion attempts.
- ▶ Tracking user account logins and logouts

To perform most of these tasks you must have explicit *privilege* to execute a restricted Smalltalk method, and you may also need to be explicitly authorized to modify an affected `GsObjectSecurityPolicy`. This chapter introduces these concepts. For a full description of privileges and `GsObjectSecurityPolicies`, see the chapter of the *GemStone/S 64 Bit Programming Guide* that discusses security.

6.1 GemStone Users

This section provides background information about how GemStone stores user accounts, what accounts are predefined, and what determines an account's name space.

UserProfiles

Each GemStone user is associated with an instance of class `UserProfile`. This `UserProfile` object contains information describing objects that the user is allowed to examine or modify, privileges that the user has to perform certain operations, and security information.

Each `UserProfile` has the following information:

User ID A unique String that identifies the user to the GemStone system.

Password	The GemStone-specific password (an InvariantString) to use to validate logins. GemStone stores the password in encrypted form in a secure manner.
Default Object Security Policy	Either nil or an instance of GsObjectSecurityPolicy. This determines the default read and write authorizations for objects created by the user.
Privileges	A collection of symbols that allow the user to perform certain “privileged” system functions.
Groups	In conjunction with object security policies, group membership is used to allow access to restricted objects for specific categories of users.
SymbolList	The list of SymbolDictionaries that this user can see.
Login Hook	Method selector or block of code to execute on login.

These are discussed in more detail starting on page 137.

Other information related to the user account is stored in an instance of UserSecurityData; this includes data related to security features. Instances of UserSecurityData are private and protected, but some information, such as lastLoginTime, can be accessed via methods in UserProfile.

AllUsers

Each instance of UserProfile must be in the global collection, AllUsers. AllUsers is the single instance of UserProfileSet. AllUsers acts as the “root” for all objects in the repository; any object in the repository must be reachable from AllUsers, usually via the SymbolLists of the UserProfiles, otherwise it is subject to garbage collection.

Special System Users

When GemStone is first installed, AllUsers has UserProfiles already defined for the following users. These are the special system users. *You must never delete these users.* These users may not have privileges removed, cannot be disabled, must use GemStone authentication, and their accounts are not subject to password or account age limits.

You can determine if an account is a special system user by executing:

```
UserProfile isSpecialUserId: 'theUserId'
```

SystemUser

SystemUser is analogous to root in UNIX. SystemUser has all privileges, belongs to all predefined groups, and is authorized to read and write all objects regardless of GsObjectSecurityPolicy protection. These privileges cannot be taken away, so SystemUser can always write to all objects. This account is used only to perform GemStone system upgrades, modify some system configuration settings, and other special-purpose operations that must be highly restricted.

The SystemUser account is the owner of the SystemObjectSecurityPolicy, which contains the kernel classes.

WARNING

*Logging in to GemStone as SystemUser is like logging in to your workstation as root: an accidental modification to a kernel object can cause a great deal of harm. Use the DataCurator account for system administration functions except those that **require** SystemUser privileges, such as a repository upgrade.*

DataCurator

The DataCurator account is the account that is normally used for day-to-day administration tasks. Initially, DataCurator is granted all privileges and belongs to all predefined groups. All GemStone UserProfiles are protected by the DataCuratorObjectSecurityPolicy.

GcUser

The GcUser account is a special account that logs in to the repository automatically to perform garbage collection tasks. You normally would only login as GcUser in order to update configuration parameters stored in GcUser's UserGlobals.

SymbolUser

The SymbolUser account is a special account that is used to perform symbol creation tasks. Login as SymbolUser is disallowed. Directly accessing the AllSymbols collection, which is in the UserGlobals of the SymbolUser, is possible from another administrative session.

Nameless

The Nameless account is a special account for use only by other GemStone products. Do not use this account or change it unless instructed to do so by GemStone Technical Support.

6.2 Creating and Removing Users

Methods that create UserProfiles add the new UserProfile to AllUsers, the global collection (UserProfileSet) of users. UserProfiles that are not in AllUsers cannot log in.

In addition to creating the new UserProfile, you should also see that each user's UNIX environment is set up to provide access to GemStone. This is described in the *GemStone/S 64 Bit Installation Guide*.

Removing a user, in addition to removing the UserProfile from AllUsers, requires cleanup to ensure that objects that refer to the deleted user do not inadvertently prevent the deleted user from being garbage collected, and that objects that are referred to only by the deleted users are not inadvertently garbage collected. The methods to remove users perform this cleanup.

UserProfile Data

Each user profile has the following instance variable data, either explicitly specified during instance creation, or provided with default values. This information can be updated.

The requirements for updating this information vary. In many cases, the requirements are different between updating information for another user and for updating your own information. See the update methods for details.

User ID

Each `UserProfile` is created with a unique user Id String. Embedded spaces are permitted.

`UserId` may only be changed by `SystemUser` or by a user with the privilege `ChangeUserId`. The `userIds` of special system users cannot be changed.

Password

Each `UserProfile` is created with an initial password, an Invariant String, which must not be the same as the User Id and may not be longer than 1024 characters. The user supplies this password for identification purposes at login, unless another form of login authentication is used; see “Password Authentication” on page 155.

Users must have the explicit privilege `#UserPassword` to change their own passwords, and there are a number of ways to constrain the choice of passwords. See the section starting on page 159 for details.

To change the password of a user other than yourself, the `#OtherPassword` privilege is required, a different method is used, and password constraints do not apply.

Default Object Security Policy

A user's `defaultObjectSecurityPolicy` determines the default read and write authorizations for objects created by the user. When you add a new user to the GemStone system, you can either allow the default security policy to be nil, use protocol that creates a new `GsObjectSecurityPolicy`, or specify an existing `GsObjectSecurityPolicy` for the user.

For more information on how security policies are used to control read and write authorization for objects in the repository, see the chapter in the *GemStone/S 64 Bit Programming Guide* that discusses security.

A `defaultObjectSecurityPolicy` of nil means that objects created by that user, by default, have world read and write access; that is, are not restricted from being read or written by all other users. Not requiring authorization checks has the benefit of improved performance, if your application does not require object level security.

To specify a `defaultObjectSecurityPolicy` for a user, you may use an existing instance of `GsObjectSecurityPolicy`, or you must create and commit the `GsObjectSecurityPolicy` before it can be used.

To modify the `defaultObjectSecurityPolicy` of any user, you must have write access to the security policy of that `UserProfile`. In addition, to modify your own `defaultObjectSecurityPolicy`, you must have the `DefaultObjectSecurityPolicy` privilege.

Privileges

When you create a new `UserProfile`, you determine whether the new user may perform certain “privileged” system functions. For example, stopping another session or the repository itself requires a particular privilege to do so. Table 1 describes the types of functions that each privilege controls.

For developers, you must also specifically grant the privilege that allows the user to modify code.

Note that privileges are more powerful than security policy authorization (see page 146). Although the owner of a security policy can always use authorization protocol to restrict read or write access to objects in a policy, an administrator with appropriate privileges, such as DataCurator, can override that protection by sending privileged messages that let you change the authorization scheme.

Table 1 GemStone Privileges

Type of Privilege	Privileged Operations
SystemControl	SystemControl is required by methods that start or stop sessions, including operations that invoke the multi-threaded scan (see page 259); for methods that suspend or resume logins, send signals to other sessions, and manage checkpoints.
SessionAccess	SessionAccess privilege is required to find out information about sessions other than the current session, or to perform operations on other sessions.
UserPassword	Required to change your own password using <code>UserProfile>>oldPassword:newPassword:</code>
DefaultObjectSecurityPolicy	This privilege is required to set a UserProfile's default ObjectSecurityPolicy using <code>UserProfile>>defaultObjectSecurityPolicy:</code> or a method that invokes that. For compatibility with previous versions, the DefaultSegment privilege also resolves to this privilege
CodeModification	You must have CodeModification privilege to create or modify instances of GsNMethod, GsMethodDictionary, or Class. See the discussion following this table.
OtherPassword	You must have OtherPassword privilege to make any changes to a UserProfile. This includes adding or removing a SymbolDictionary to/from a SymbolList that is not your own. OtherPassword is also required to find out information about UserProfiles other than the currently logged in session. OtherPassword is required to make any changes to AllUsers, including creating a new user and configuring security requirements.
ObjectSecurityPolicy Creation	Required in order to creating a new GsObjectSecurityPolicy, using <code>GsObjectSecurityPolicy class>>new, newInRepository:</code> or any methods that invoke these. For compatibility with previous versions, the SegmentCreation privilege also resolves to this privilege.

Table 1 GemStone Privileges (Continued)

Type of Privilege	Privileged Operations
ObjectSecurityPolicyProtection	You must have ObjectSecurityPolicyProtection to update the authorizations of a GsObjectSecurityPolicy, other than one that is owned by the current session's user. This includes GsObjectSecurityPolicy>>group:authorization:, ownerAuthorization:, and worldAuthorization:. For compatibility with previous versions, the SegmentProtection privilege also resolves to this privilege.
FileControl	FileControl is required for operations that access external files, including operations related to backup, restore, transaction logs, and extents.
GarbageCollection	Required to perform any garbage collection operation, to start and stop Admin and Reclaim Gems, and force epoch or reclaim to run. Also required to audit and profile the repository.
NoPerformOnServer	If you have this privilege, you cannot execute System class>>performOnServer:
NoUserAction	If you have this privilege, you cannot execute System class>> loadUserActionLibrary:
NoGsFileOnServer	If you have this privilege, you cannot execute any GsFile operation which accesses a file on the server.
NoGsFileOnClient	If you have this privilege, you cannot execute any GsFile operation which accesses a file on the client.
SessionPriority	Required to modify the priority of any session, or to check the priority of a session other than the current session.
CompilePrimitives	Allow user to compile primitive methods, which is otherwise restricted to SystemUser.
ChangeUserId	Allow user to execute userId:password: to rename a user, which is otherwise restricted to SystemUser.

Code Modification Privilege

CodeModification privilege is required to execute any method that modifies Smalltalk code:

- ▶ You must have CodeModification privilege to create instances of GsNMethod, or to create or modify instances of GsMethodDictionary or Class. (You cannot modify a GsNMethod once it has been created.)
- ▶ You must have CodeModification privilege to add a Class to, or remove a Class from, a SymbolDictionary and its subclasses.
- ▶ You must have CodeModification privilege to add or remove a SymbolDictionary from your own SymbolList. (See page 151)

You cannot use GemBuilder for C to modify instances of the following classes (or their subclasses): GsNMethod, GsMethodDictionary, Class, SymbolDictionary, SymbolList, UserProfile.

Groups

GemStone uses group membership to facilitate access to objects that are protected by GsObjectSecurityPolicies. While certain objects must be protected from read or write access by other users in the system (the “world”), you may still need to grant access to specific individual users. By adding group authorization to the GsObjectSecurityPolicy that protects these objects, and adding the corresponding group membership to that user, you can provide a user with the appropriate access to these objects.

A GsObjectSecurityPolicy can authorize multiple groups and a user can be a member of multiple groups. There are several predefined groups, as shown in Table 2. By default, all new users become members of group Subscribers.

AllGroups

AllGroups is a global collection of Strings, that includes all groups defined for all users and all security policies.

Initially, AllGroups contains the following:

Table 2 GemStone Groups

Group name	Access
System	Members of this group have write access to objects protected by the GcUser’s object security policy.
DataCuratorGroup	Members of this groups have write access to objects protected by the DataCuratorObjectSecurityPolicy . This is useful if you wish to make a user other than DataCurator to be a system administrator, since many operations that update users require write access to DataCuratorObjectSecurityPolicy.
Publishers	Members of this group have write access to objects protected by PublishedObjectSecurityPolicy .
Subscribers	Members of this group have read access to objects protected by PublishedObjectSecurityPolicy .
SymbolUser	(Reserved).

Symbol Lists

As explained in the *GemStone/S 64 Bit Programming Guide*, the GemStone Smalltalk compiler follows a well-defined path in resolving objects named by source code symbols. First, the compiler considers the possibility that a variable name might be either local (a temporary variable or an argument) or defined by the class of the current method definition (an instance variable, a class variable, or a pool variable). If a variable is none of these, the compiler refers to an Array of SymbolDictionaries in the user’s UserProfile and current session state. That Array is called the user’s *symbol list*. The symbol list tells Smalltalk which of many possible GemStone SymbolDictionaries to search for an object named in a Smalltalk program.

For each user, a persistent instance of class SymbolList is stored in the repository and is referenced from the UserProfile associated with this user as the symbolList instance variable. In addition, a transient copy of that SymbolList is stored in the GsCurrentSession object for the logged-in session.

A session's transient copy can be modified without affecting (or causing concurrency conflicts with) either the persistent symbol list or the transient copies controlling other sessions. Changes to your own UserProfile's persistent symbol list also change the symbol resolution of your current session. However, changes to the persistent symbol list are likely to cause concurrency conflicts with other sessions logged in under the same userId.

For further information about symbol lists, refer to the *GemStone/S 64 Bit Programming Guide*.

New UserProfiles are created with the following SymbolDictionaries, in this order:

- UserGlobals** Each UserProfile has its own SymbolDictionary for the user's private symbols.
- Globals** The second element in each user's initial symbol list is a "system globals" SymbolDictionary, *Globals*. This dictionary contains all of the GemStone Smalltalk kernel classes (Object, Class, Collection, Integer, and so forth). Although users can read the objects in *Globals*, ordinarily they cannot modify objects in that Dictionary.
- Published** The third and final element in each user's initial symbol list is a SymbolDictionary for application objects that are "published" to all users. Users who are members of the group Publishers can place objects in this dictionary to make them visible to other users. 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 symbolList instance variable.

Although all users automatically share access to objects in *Globals*, sharing application objects between users requires that the objects be in a SymbolList that is visible to both users. There are three primary ways to do this:

- ▶ As a member of group Publishers, you can add the objects to the Published dictionary. This dictionary is already in each user's symbol list, so whatever you add becomes visible to users the next time they obtain a fresh transaction view of the repository. You may do this by sending the message `Published at: aKey put: aValue`.
- ▶ You can define a special SymbolDictionary, and add that to the user's SymbolList. The procedure is described under "Adding a SymbolDictionary to Someone Else's Symbol List" on page 151.
- ▶ The application itself can add the objects to the individual user's symbol list, either to the permanent symbol list in the UserProfile or to a transient symbol list for that session. For information about this approach, refer to the *GemStone/S 64 Bit Programming Guide*.

For more information, refer to the chapter on symbol resolution and object sharing in the *GemStone/S 64 Bit Programming Guide*.

Creating Users

Privileges required: OtherPassword, and write access to the DataCuratorObjectSecurityPolicy. You may also need ObjectSecurityPolicyCreation.

Simple User Creation

At minimum to create a new UserProfile, you must supply the new user's userId and password, each as a String. This creates the new user with no privileges, only the Subscribers group, and with a nil defaultObjectSecurityPolicy.

```
AllUsers addNewUserWithId: 'theUserId'
  password: 'thePassword' .
```

Simple User Creation with GsObjectSecurityPolicy Creation

To create a new user and specify that a new instance of GsObjectSecurityPolicy should be created for the new user, use the following method:

```
AllUsers addNewUserWithId: 'theUserId'
  password: 'thePassword'
  createNewObjectSecurityPolicy: true
```

User Creation With Privileges, Groups, and ObjectSecurityPolicy

Using the complete form allows you to assign privileges to the new user, add the user to groups, and specify an ObjectSecurityPolicy. The ObjectSecurityPolicy may be nil.

```
AllUsers addNewUserWithId: 'theUserId'
  password: 'thePassword'
  defaultObjectSecurityPolicy: anObjectSecurityPolicyOrNil
  privileges: anArrayOfPrivStrings
  inGroups: aCollectionOfGroupStrings
```

For example:

```
topaz 1> printit
AllUsers addNewUserWithId: 'Mary'
  password: 'herPasswd'
  defaultObjectSecurityPolicy: nil
  privileges: #( 'UserPassword' )
  inGroups: #( 'MarathonRunners' ).
"commit the new UserProfile"
System commitTransaction.
%
```

For additional user creation protocol, see the image. UserProfileSet instance methods and UserProfile class methods both create new UserProfiles and add the new user to AllUsers.

Removing Users

Privileges required: OtherPassword.

In addition to removing the UserProfile from AllUsers, removing a user requires that any GsObjectSecurityPolicies owned by the deleted user are moved to a new user.

In addition, to ensure that work being done by the user being deleted is not lost, the SymbolLists of the deleted user are moved to a separate location where they can be periodically reviewed and manually dereferenced.

DeletedUserProfile and AllDeletedUsers

When a user is removed, a new instance of DeletedUserProfile is created, with the userId and the SymbolLists of the user being removed, and the current DateTime.

This instance of DeletedUserProfile is moved to a collection #AllDeletedUsers, in Globals SymbolDictionary.

An administrator should review the classes and methods in the SymbolLists of the DeletedUserProfile, copy anything valuable elsewhere, and manually remove the DeletedUserProfile from AllDeletedUsers.

Remove User, with ObjectSecurityPolicies Going to the Current User

The following methods remove the user and reassign any GsObjectSecurityPolicies owned by the user to be removed to the UserProfile of the current user (the user executing this code, such as DataCurator).

This form passes in the instance of UserProfile for the user to be deleted:

```
AllUsers removeAndCleanup: aUserProfile.
```

While this form passes in the userId of the user to be deleted, and includes a block to execute if no UserProfile with the given userId exists in AllUsers:

```
AllUsers removeAndCleanupUserWithId: 'aUserId'
  ifAbsent: aBlock
```

Remove User, with ObjectSecurityPolicies going to another User

The following methods remove the user and reassign any GsObjectSecurityPolicies owned by that user to another specific UserProfile.

This form passes in the instance of UserProfile for the user to be deleted and the UserProfile to which to reassign the ObjectSecurityPolicies:

```
AllUsers removeAndCleanup: aUserProfile
  migrationSecurityPoliciesTo: anotherUserProfile.
```

While this form passes in the userId String of the user to be deleted and the userId String of the UserProfile to which to reassign the ObjectSecurityPolicies, and includes a block to execute if a UserProfile does not exist with either of the UserId Strings:

```
AllUsers removeAndCleanupUserWithId: 'aUserId'
  migrationSecurityPoliciesUserWithId: 'anotherUserId'
  ifAnyAbsent: aBlock
```

For example,

```
topaz 1> printit
AllUsers removeAndCleanup: (AllUsers userWithId: 'John')
  migrationSecurityPoliciesTo: (AllUsers userWithId: 'Ann')
System commitTransaction.
%
```


6.3 Administering Users

List Existing Users

Privileges required: None.

There is no direct method within GemStone Smalltalk to list only the names of existing accounts. The following example shows one way to obtain that information:

```
topaz 1> printit
(AllUsers collect: [:each | each userId ]) asArray.
%
#1 SystemUser
#2 DataCurator
#3 Nameless
#4 SymbolUser
#5 GcUser
```

Modifying the UserId

Privileges required: ChangeUserId.

Updating the userId requires resetting the password for that user. The new user ID and password will take effect when you commit the current transaction. The names of special system users cannot be changed.

To modify the user ID of a GemStone user, execute the following expression:

```
topaz 1> printit
(AllUsers userWithId: 'theUserId')
  userId: 'newId'
  password: 'newPassword'.
%
```

An error is raised if *newId* is the userId of an existing UserProfile.

Modifying Password

Users Changing Their Own Password

Privileges required: UserPassword.

In many cases, users set their own passwords and may be required to update them periodically. These users must be given the UserPassword privilege to do so, and use the following method to update their password:

```
UserProfile >> oldPassword:'oldPwd' newPassword: 'newPwd'
```

For example,

```
topaz 1> printit
System myUserProfile
  oldPassword: 'oldPasswordString'
  newPassword: 'newPasswordString'.
System commitTransaction
%
```

Password choice is constrained by login security that is configured for the repository; see the discussion on password restrictions starting on page 159.

A different method, requiring other privileges, is used by Administrators to update the password of another user.

The new password takes effect when you commit the current transaction.

Changing Another User's Password

Privileges required: OtherPassword.

To modify the password of any GemStone user, execute the following expression.

```
topaz 1> printit
(AllUsers userWithId: 'theUserId')
  password: 'newPasswordString' .
System commitTransaction
%
```

The new password takes effect when you commit the current transaction.

The password set by this method is not subject to the constraints described on page 159 because it can only be set by a user having the OtherPassword privilege. The password must not be the same as the UserId and must not be longer than 1024 characters.

Each password change of this type is noted in the GemStone security log, which currently is the Stone's log file. The entry includes the userId of the session making the change but not the new password.

Modifying defaultObjectSecurityPolicy

Each security policy maintains access authorization for its owner, the world, and an unlimited number of groups. There are three levels of authorization: none, read (read-only), and write (which includes read permission).

Determining Who Is Authorized to Read or Write in an Object Security Policy

Privileges required: read authorization for the security policy with which this policy is associated, such as the DataCuratorObjectSecurityPolicy.

You can find out who is authorized to read or write objects in an security policy by sending it the message `asString`. For instance:

```
topaz 1> printit
PublishedObjectSecurityPolicy asString
%
anObjectSecurityPolicy, Number 6 in Repository SystemRepository,
Owner SystemUser write, Group Subscribers read, Group Publishers
write, World none
```

Changing the Authorization of an Object Security Policy

Privileges required: ObjectSecurityPolicyProtection or be the security policy's owner, and write authorization to the DataCuratorObjectSecurityPolicy.

The new authorization will take effect for logins following the commit of the current transaction.

CAUTION

Do not attempt to change the authorization of SystemObjectSecurityPolicy.

To change the authorization for a security policy, execute any (or all) of the following expressions.

```
topaz 1> printit
theObjectSecurityPolicy ownerAuthorization: #anAuthorizationSymbol .
%
topaz 1> printit
theObjectSecurityPolicy worldAuthorization: #anAuthorizationSymbol .
%
topaz 1> printit
theObjectSecurityPolicy group: 'aGroupString'
authorization: #anAuthorizationSymbol .
%
```

NOTE

Exercise caution when changing the authorization for any security policy that a user may be using as his or her default or current security policy – whether or not the user owns the affected policy. If a user attempts to commit a transaction, but has created objects with a policy for which he or she no longer has write authorization, an error will be generated.

For example, to authorize the group Accounting to read (but not write) in user Eli's default security policy, you could execute the following expression:

```
topaz 1> printit
(AllUsers userWithId: 'Eli') defaultObjectSecurityPolicy
group: 'Accounting'
authorization: #read.
%
```

If the group 'Accounting' does not exist, GemStone will return an error. The discussion under “Adding a User to a Group” on page 148 explains how to create a new GemStone group.

Remove a Group from an Object Security Policy's Authorization List

Privileges required: ObjectSecurityPolicyProtection, and write authorization for the security policy.

To remove a group from a security policy's list of authorized groups, execute the following expression:

```
theSecurityPolicy group: 'aGroupString' authorization: #none
```

Change a User's Default Object Security Policy

Privileges required: DefaultObjectSecurityPolicy, and write authorization to the DataCurator ObjectSecurityPolicy.

Changes to another user's default security policy do not take effect until the next login.

To change a user's default security policy, execute the following expression:

```
(AllUsers userWithId: 'theUserId') defaultObjectSecurityPolicy:
aNewSecurityPolicy
```

NOTE

If you change any user's default security policy (including your own) to a security policy for which that user lacks write authorization, and you subsequently commit the transaction, the affected user will no longer be able to log in to GemStone.

Modifying Groups

Examining Group Memberships

No privileges are required for this operation.

To find out which groups a user belongs to, execute the following expression:

```
(AllUsers userWithId: 'theUserId') groups
```

This expression returns a Set of Strings indicating the groups to which the user belongs.

Adding a User to a Group

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

To add a user to a group, do the following:

```
topaz 1> printit
" If this is a new group, add it to AllGroups "
('MarathonRunners' in: AllGroups)
  ifFalse: [AllGroups add: 'MarathonRunners' ].
(AllUsers userWithId: 'theUserId') addGroup: 'MarathonRunners'.
System commitTransaction
%
```

This expression adds the user to the group MarathonRunners by adding the group name to the list of groups maintained in the UserProfile. (This action takes effect when you commit the current transaction.)

If the group MarathonRunners did not previously exist, this expression creates it in AllGroups (the “master list” of all group names). See page 141 for more information about AllGroups.

Removing a User from a Group

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

To remove a user from a group, execute an expression of the following form:

```
topaz 1> printit
(AllUsers userWithId: 'theUserId') removeGroup: 'Sprinters' .
System commitTransaction
%
```

This expression removes the designated group from the list of groups to which the user belongs. This action will take effect when you commit the current transaction. For more

information about groups, see the Security chapter of the *GemStone/S 64 Bit Programming Guide*.

Listing Members of a Group

No privileges are required for this operation.

To list all members of a user group, execute the following expression:

```
AllUsers membersOfGroup: aString
```

This expression returns an IdentitySet containing the userId for each member of the group.

Removing a Group

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

To remove a user group from the global object AllGroups, first remove each member from the group, then remove the group itself, using the expression

```
AllGroups removeGroup: aGroupString.
```

For example, to remove a group named MarathonRunners:

```
topaz 1> printit
  "Remove each member from the group "
  (AllUsers usersInGroup: 'MarathonRunners') do:
    [:aUserProfile| aUserProfile removeGroup: theGroup].
  "Now remove the group itself "
  AllGroups remove: 'MarathonRunners'.
  System commitTransaction
  %
```

Modifying Privileges

Examining a User's Privileges

No privileges are required for this operation.

GemStone provides messages that allow you to determine which privileged methods a GemStone user may execute, and to change the privileges of any user. Naturally, you need the appropriate privileges to use those methods.

To find out which privileged methods a given user is permitted to execute, send the following message to the desired user's UserProfile:

```
topaz 1> printit
  (AllUsers userWithId: 'theUserId') privileges
  %
```

This message returns an Array of Strings. Each String in the Array corresponds to one of the user's privileges. Table 1 (on page 139) lists the Smalltalk methods that correspond to each privilege.

Adding a Privilege

Privileges required: OtherPassword and write authorization to the ObjectSecurityPolicy of the user's UserProfile.

The new privileges will take effect when you commit the current transaction.

To add to a user's existing privileges, execute the following expression:

```
(AllUsers userWithId: 'theUserId') addPrivilege: aPrivilegeString .
```

Here's an example that assigns three new privileges to user Bob:

```
topaz 1> printit
(AllUsers userWithId: 'Bob')
  addPrivilege: 'SystemControl';
  addPrivilege: 'SessionAccess';
  addPrivilege: 'UserPassword' .
System commitTransaction
%
```

Revoking a Privilege

Privileges required: OtherPassword and write authorization to the security policy of the user's UserProfile.

The privileges will be revoked when you commit the current transaction.

To revoke one (or more) of a user's existing privileges, execute the following expression:

```
(AllUsers userWithId: 'theUserId') deletePrivilege: aPrivilegeString .
```

The following example revokes three of user Jane's privileges:

```
topaz 1> printit
(AllUsers userWithId: 'Jane')
  deletePrivilege: 'SystemControl';
  deletePrivilege: 'SessionAccess';
  deletePrivilege: 'UserPassword' .
System commitTransaction
%
```

Reassigning All Privileges

Privileges required: OtherPassword and write authorization to the security policy of the user's UserProfile.

The new privileges will take effect when you commit the current transaction.

To redefine the full set of a user's privileges, perhaps adding some and revoking others, execute the following expression:

```
(AllUsers userWithId: theUserId) privileges: anArrayOfStrings
```

This expression supersedes any previous privilege assignments. After the change is committed, only those privileges listed in the expression are valid for the user. Any privileges that were previously valid, but are not listed, are revoked.

For example:

```
topaz 1> printit
(AllUsers userWithId: 'Sam') privileges:
    #( 'UserPassword' ) .
System commitTransaction
%
```

Modifying SymbolLists

Adding a SymbolDictionary to Your Own Symbol List

Privileges required: CodeModification.

You can add a dictionary to a symbol list by sending the message `UserProfile>>insertDictionary: aSymbolDictionary at: anIndex`. The change does not affect the transient copy of the symbol list that is used by another currently logged in session until that session commits or aborts.

This example inserts dictionary `NewDict` (which already exists in the Published dictionary) into the user's own symbol list:

```
topaz 1> printit
System myUserProfile
    insertDictionary: NewDict at: 2.
%
```

Inserting the new dictionary at index 2, as in the example, places it between the `UserGlobals` and the `Globals` dictionaries in the search order. Because symbol resolution depends on the order of dictionaries in a user's symbol list, the index used in this example may not be appropriate for all situations.

Adding a SymbolDictionary to Someone Else's Symbol List

Privileges required: `OtherPassword` and write permission to the security policy of the other user's `SymbolDictionary`.

This example inserts dictionary `NewDict` (which already exists in the Published dictionary) into user Jerry's symbol list:

```
topaz 1> printit
(AllUsers userWithId: 'Jerry')
    insertDictionary: NewDict at: 7.
System commitTransaction
%
```

Removing a SymbolDictionary from Your Own Symbol List

Privileges required: CodeModification.

You can remove a dictionary from a symbol list by sending the message `UserProfile>>removeDictionary: aSymbolDictionary`.

The change does not affect the transient copy of the symbol list that is used by another currently logged in session, until that session commits or aborts.

This example removes dictionary OldDict from the user's own symbol list:

```
topaz 1> printit
System myUserProfile
    removeDictionary: OldDict.
System commitTransaction
%
```

Removing a SymbolDictionary from Someone Else's Symbol List

Privileges required: OtherPassword and write permission to the security policy of the other user's SymbolDictionary,

This example removes dictionary OldDict from user Jerry's symbol list:

```
topaz 1> printit
(AllUsers userWithId: 'Jerry')
    removeDictionary: OldDict.
System commitTransaction
%
```

Disable and Enable User Logins

A disabled account cannot log in. Disabling an account consists of setting its password to a non-enterable character. Once an account is disabled, a new password must be assigned for that account by an administrator, before the account can log in again.

Disable an Account Explicitly

Privileges required: OtherPassword. Logins cannot be disabled for special system accounts.

Users can be disabled using the method `disable` or `disableWithReason:`. This prevents the user from logging in, but does not affect current login sessions.

For example:

```
topaz 1> printit
(AllUsers userWithId: 'Sam')
    disableWithReason: 'on Sabbatical'.
System commitTransaction
%
```

Re-enable an Account

Privileges required: OtherPassword.

To re-enable a user account, an administrative user must login and reset the user's password. For example, to reset Sam's password and require him to change his password the first time he logs in:

```
topaz 1> printit
(AllUsers userWithId: 'Sam')
    password: 'AaBbCc';
    loginsAllowedBeforeExpiration: 1.
System commitTransaction
%
```


Find Out Which Accounts Have Been Disabled

Privileges required: OtherPassword.

The message `AllUsers findDisabledUsers` returns a `SortedCollection` of `UserProfiles` that are disabled explicitly or by one of the security precautions discussed in this chapter:

- ▶ The password expired (through aging or a login limit).
- ▶ The account remained inactive.
- ▶ There were repeated password failures.

For example:

```
topaz 1> level 1
topaz 1> printit
AllUsers findDisabledUsers
  collect: [:aUser | aUser userId ] .
%
an Array
  #1 qa2
  #2 qa3
```

DataCurator or another user with the OtherPassword privilege can reactivate an account by giving it a new password, as described on page 152.

Check If an Account Is Disabled

Privileges required: OtherPassword.

You can check if a particular account is disabled by sending the message `isDisabled` to the account's `UserProfile`. The method returns either `True` or `False`. This example inquires about account `qa2`:

```
topaz 1> printit
(AllUsers userWithId: 'qa2') isDisabled
%
true
```

Find Out Why an Account Was Disabled

Privileges required: OtherPassword.

You can find out why a particular account was disabled by sending the message `reasonForDisabledAccount` to the account's `UserProfile`. This example inquires about account `qa2`:

```
topaz 1> printit
(AllUsers userWithId: 'qa2') reasonForDisabledAccount
%
LoginsWithSamePassword
```

If the account was disabled using the method `disableWithReason:`, this method will return that argument. Otherwise if the account was disabled by login security, it will return one of these Strings: 'PasswordAgeLimit', 'StaleAccount', 'LoginsWithSamePassword', or 'LoginsWithInvalidPassword'.

Disable and Enable Commits by User

Commits can be disabled for particular users to ensure “read only” access to the GemStone repository. These users can still log in and view data for which they have read or write authorization, and can modify objects, but they cannot commit and make any changes permanent.

Disable Commits

Privileges required: OtherPassword. Commits cannot be disabled for special system users.

To disable commits for a user, execute the following expression:

```
(AllUsers userWithId: theUserId) disableCommits
```

This expression disables any commits for the given user, beginning with the next login of the user after the session making this change commits. If this user is currently logged in, it does not affect the user's transaction, if any.

For example:

```
topaz 1> printit  
(AllUsers userWithId: 'Sam') disableCommits.  
System commitTransaction  
%
```

Re-enable Commits for a User

Privileges required: OtherPassword.

To enable commits for a user, execute the following expression:

```
(AllUsers userWithId: theUserId) enableCommits
```

This expression enables commits for the given user, beginning with the next login after the session making this change commits.

Check If a User Can Commit

Privileges required: OtherPassword.

To test if commits have been disabled for a user account, execute the following expression:

```
(AllUsers userWithId: theUserId) isReadOnly
```

6.4 Password Authentication

When a user wants to log in to the GemStone repository, their login must be authenticated: they must present a password that is matched against stored information for their UserId, to verify that they are authorized to log in. This authentication can be done entirely within GemStone, or GemStone can use UNIX or LDAP to perform the authentication.

Performing the authentication entirely with GemStone – Gemstone authentication – is the initial default for all users. Other authentication schemes can be configured for individual UserProfiles, although special system users will always use GemStone authentication. The repository may contain UserProfiles using a mix of authentication schemes.

After the authentication scheme is modified for a user and the change is committed, it will take effect the next time the user logs in. Existing logins are not affected.

In addition to GemStone user login authentication, users may need to provide OS userId and password authentication, to execute processes at the OS level. This depends on how the system is configured; see Chapter 3 for details interprocess security and host authentication.

GemStone Authentication

Privileges required: OtherPassword

GemStone authentication is the default authentication, using the UserId and password store with the UserProfile of the account. In older versions, this was the only way of authentication within GemStone, and must still always be used by the special system users - SystemUser, DataCurator, GcUser, SymbolUser, and Nameless.

```
topaz 1> printit
(AllUsers userWithId: 'Mary') enableGemStoneAuthentication.
System commitTransaction.
%
```

UNIX Authentication

Privileges required: OtherPassword

When UNIX account authentication is enabled for a user, they will enter their UNIX account password instead of their GemStone password. Password management for this user then becomes the Unix password management; sending messages to change the GemStone password are not useful, and these accounts are not subject to GemStone's password aging mechanisms.

UNIX authentication uses PAM (pluggable authentication module) , and PAM must be configured on the system in order to use UNIX Authentication. Login will look for a module gemstone.gem. See the *Installation Guide* for your server platform for details.

The GemStone userId may be the same as the UNIX userId, or they may be different. When you enable UNIX authentication for an account, you may specify the UNIX userId associated with the GemStone UserProfile and this will be used for authentication. Using nil means to use the GemStone userId as the UNIX userId.

```
(AllUsers userWithId: GemStoneUserId)
  enableUnixAuthenticationWithAlias: UNIXUserIdOrNil
```

for example, if Mary's UNIX `userId` is `msmith`, you can use the first form if her GemStone `userId` is `Mary`. If her GemStone `userId` is also `msmith` you leave the argument `nil`.

```
topaz 1> printit
(AllUsers userWithId: 'Mary')
  enableUnixAuthenticationWithAlias: 'msmith'.
System commitTransaction.
%
topaz 1> printit
(AllUsers userWithId: 'msmith')
  enableUnixAuthenticationWithAlias: nil.
System commitTransaction.
%
```

To verify that the UNIX `userId` exists, execute:

```
System unixUserIdExists: UNIXUserId
```

LDAP Authentication

Privileges required: OtherPassword

An LDAP (Lightweight Directory Access Protocol) server consolidates and centralizes user authentication, and can be used to authenticate logins to the GemStone server. UNIX authentication, which uses PAM, may ultimately use LDAP if PAM is configured to use LDAP. When you configure GemStone to use LDAP authentication, it goes to LDAP directly, rather than using PAM.

To use LDAP for GemStone authentication, you must have an LDAP server available. When a `UserProfile` has been configured for LDAP authentication, on login, GemStone performs an LDAP bind to authenticate the `userId`.

In most cases, the LDAP bind requires a Distinguished Name (DN), which is the unique identifier for an entry. The DN includes both the `userId` and domain information, e.g. `'uid=msmith,ou=employees,dc=somecompany,dc=com'`. GemStone composes the DN based on the arguments provided when you configure LDAP authentication.

To configure a user to use LDAP for authentication, use the following method:

```
(AllUsers userWithId: GemStoneUserId)
  enableLDAPAuthenticationWithAlias: LDAPUserIdOrNil
  baseDn: baseDn
  filterDn: filterDn
```

UserId or Alias

As with UNIX authentication, if the LDAP `userId` is the same as the GemStone Id, you may pass in `nil` for the argument `UNIXUserIdOrNil`; otherwise, pass in the LDAP `userId`. The user will use their GemStone `userId` and the password for their LDAP account to log in.

Fully Qualified DN

You can configure the LDAP authentication for a particular user with the specific DN, with `'%s'` replacing the `userId`, for the `baseDn:` argument. In this case you should pass in `nil` for the `filterDn:` argument, which will disable the query.

For example, to configure authentication to use a specific fully qualified DN:

```
topaz 1> printit
(AllUsers userWithId: 'Mary')
  enableLDAPAuthenticationWithAlias: 'msmith'
  baseDn: 'uid=%s,ou=employees,dc=somecompany,dc=com'
  filterDn: nil.
System commitTransaction.
%
```

Search for DN

You can also configure authentication to include the base domain information in the baseDn: argument, and include a filter in the filterDn: argument. The filter must contain '%s' in the position for the userId. GemStone will perform a query to get the full DN given the base and filter information.

The base and filter information is provided at the time the authentication is configured, not at login time, so a search option is particularly useful if the LDAP structure is likely to be modified.

To configure authentication to do a search for the given user:

```
topaz 1> printit
(AllUsers userWithId: 'Mary')
  enableLDAPAuthenticationWithAlias: 'msmith'
  baseDn: 'dc=somecompany,dc=com'
  filterDn: '(uid=%s)'.
System commitTransaction.
%
```

Note on anonymous binds to LDAP server

GemStone login authentication requires that the LDAP server allow anonymous binds. If authentication is required to access the LDAP server for authentication, the GemStone login must be done using a different authentication.

Once logged in, you can perform an validation call to the LDAP server that passes in the LDAP authentication as well as the user authentication, using this method:

```
System class >> validatePasswordUsingLdapServers:baseDn:
  filterDn:userId:password:bindDn:bindPassword:
```

The additional bind arguments should be nil if authenticating in explicit mode, or with anonymous bind. For example:

Explicit mode

```
System validatePasswordUsingLdapServers:
  (Array with: 'ldaps://myldap.mydomain.com')
  baseDn: 'uid=%s,ou=Users,dc=mycompany,dc=com'
  filterDn: nil
  userId: 'MyUserId' password: 'swordfish'
  bindDn: nil bindPassword: nil
```

Search mode with anonymous bind

```
System validatePasswordUsingLdapServers:
  (Array with: 'ldaps://myldap.mydomain.com')
  baseDn: 'ou=Users,dc=mycompany,dc=com'
  filterDn: '(uid=%s)'
  userId: 'MyUserId' password: 'swordfish'
  bindDn: nil bindPassword: nil
```

Search mode with authenticated bind

```
System validatePasswordUsingLdapServers:
  (Array with: 'ldaps://myldap.mydomain.com')
  baseDn: 'ou=Users,dc=mycompany,dc=com'
  filterDn: '(uid=%s)'
  userId: 'MyUserId' password: 'swordfish'
  bindDn: 'LdapBindUser' bindPassword: 'LdapBindPassword'
```

For further details, see the method comments in the image.

Determining an Account's Authentication Scheme

Privileges required: OtherPassword

Your repository may contain a mix of authentication schemes, with some users (at least certainly the special system accounts) using GemStone authentication, others authenticating using UNIX accounts, and still other using LDAP.

You can determine what scheme an account is using by sending `authenticationScheme`. This will return the symbol `#GemStone`, `#UNIX`, or `#LDAP`. For example,

```
topaz 1> printit
(AllUsers userWithId: 'DataCurator') authenticationScheme.
%
GemStone
```

You may also send messages to check for specific schemes. The following methods return true or false.

```
(AllUsers userWithId: userId) authenticationSchemeIsUNIX
(AllUsers userWithId: userId) authenticationSchemeIsLDAP
(AllUsers userWithId: userId) authenticationSchemeIsGemStone
```

6.5 Configuring GemStone Login Security

GemStone provides several login security features. You can:

- ▶ Constrain the choice of passwords to a certain pattern, ban particular passwords altogether, or ban reuse of a password by the same account
- ▶ Require users to change their passwords periodically (password aging)
- ▶ Limit the number of logins under a temporary password
- ▶ Disable accounts that have not logged in for a specified interval (account aging)
- ▶ Limit the number of concurrent sessions by a particular account
- ▶ Monitor failed login attempts and, if necessary, disable further login attempts on that account

In all cases, the password must not be the same as the UserId and must not be longer than 1024 characters.

Additional methods let you determine which accounts have been disabled by one of these security features and why a particular account was disabled.

CAUTION

GemStone records certain administrative changes to these security features in the Stone log. You may want to restrict access to that file.

The special system users - SystemUser, DataCurator, SymbolUser, GcUser, and Nameless - are never disabled by the security features.

Limiting Choice of Passwords

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

You can constrain a user's choice of passwords in terms of pattern (such as the number of characters that repeat). Independently, you can establish a list of words that are disallowed as passwords, and you can keep a user from choosing the same password more than once.

The constraints described here apply only when a user changes his or her own password by using the message `UserProfile>>oldPassword:newPassword:` and only to password changes after the constraint is committed to the repository. That is, the constraints (other than the prohibition of `userId` as the password) do not apply to changing a user's user's password using the `password:` method (which requires OtherPassword privilege), and they do not invalidate existing passwords.

Table 1 lists the messages that you can use to set pattern constraints. You send these messages to the global object `AllUsers`. For example, to set the minimum password length to six characters, do this:

```
topaz 1> printit
AllUsers minPasswordSize: 6 .
System commitTransaction
%
```

The default setting in all cases is 0, which means there is no constraint on the pattern.

Table 1 Ways to Constrain the Password Pattern

Message to AllUsers	Comments
<code>minPasswordSize: aPositiveInteger</code>	Sets the minimum number of characters in a new password; 0 means no constraint.
<code>maxPasswordSize: aPositiveInteger</code>	Sets the maximum number of characters in a new password; 0 disables the constraint. (The password String itself must not be longer than 1024 characters.)
<code>maxRepeatingChars: aPositiveInteger</code>	Sets the maximum number of adjacent characters that can have the same value; for example, 1 allows 'aba' but not 'aa'. 0 means no constraint.
<code>maxConsecutiveChars: aPositiveInteger</code>	Sets the maximum number of adjacent characters that can be an ascending or descending sequence, such as '123' or 'zyx' based on a case-sensitive comparison. 0 means no constraint.
<code>maxCharsOfSameType: aPositiveInteger</code>	Sets the maximum number of adjacent characters that can be of the same type (alpha, numeric, or special); for example, 3 allows 'abc4de' but not 'abcde'. 0 means no constraint.
<code>passwordRequiresUppercase: aBoolean</code>	Set the requirement that the password contain at least one uppercase character. false means no constraint.
<code>passwordRequiresLowercase: aBoolean</code>	Set the requirement that the password contain at least one lowercase character. false means no constraint.
<code>passwordRequiresSymbol: aBoolean</code>	Set the requirement that the password contain at least one character that is not alphanumeric. false means no constraint.
<code>passwordRequiresDigit: aBoolean</code>	Set the requirement that the password contain at least one digit. false means no constraint.

Any user can inquire about the current setting of a password pattern constraint by sending its corresponding Accessing message (that is, without the colon or argument shown in Table 1).

For example, to determine the current minimum size for a password:

```
topaz 1> printit
AllUsers minPasswordSize
%
6
```

Disallowing Particular Passwords

Privileges required: OtherPassword and write authorization to the DataCuratorObjectSecurityPolicy.

You can create a list of disallowed passwords by adding Strings to the AllUsers instance variable disallowedPasswords. Any messages understood by class Set can be used. For instance:

```
topaz 1> printit
(AllUsers disallowedPasswords)
  addAll: #( 'Mother' 'apple_pie' ) .
System commitTransaction
%
```

The default is an empty set.

Additions to this list affect only new passwords requested after the additions are committed; that is, additions do not invalidate existing passwords. If a user attempts to change that account's password to one of the Strings in disallowedPasswords, a SecurityError is signaled.

Any user can examine the current list of globally disallowed passwords by sending the message AllUsers disallowedPasswords.

Disallowing Reuse of Passwords

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

You can prevent each user from choosing the same password more than once by setting AllUsers disallowUsedPasswords to true. By default, disallowUsedPasswords is false.

When reuse of passwords is disallowed, GemStone maintains a separate encrypted set of old passwords for each user. Each time a user invokes oldPassword:newPassword:, the new password is checked against the prior passwords for that account. If the new password matches a prior one, a SecurityError is signaled.

Disallow All Previously Used Passwords

To disallow password reuse, use the method:

```
AllUsers disallowUsedPasswords: aBoolean.
```

For example:

```
topaz 1> printit
AllUsers disallowUsedPasswords: true .
System commitTransaction
%
```

Disallow a Specific Number of Previous Passwords

Privileges required: OtherPassword and write authorization to DataCuratorObjectSecurityPolicy.

To disallow a fixed number of previously-used password, but allow earlier passwords, use the following:

```
AllUsers numberOfDisallowedPasswords: anInteger
```

anInteger must be a number between 0 and 65535; 0 means that the user may not reuse any previously-used passwords. The limit on the number of disallowed passwords has no effect if disallowUsedPasswords is false.

For example,

```
topaz 1> printit
AllUsers disallowUsedPasswords: true.
AllUsers numberOfDisallowedPasswords: 10.
System commitTransaction
%
```

Clearing a User's Disallowed Old Passwords

Privileges required: OtherPassword.

You can clear the set of old passwords so that they can be reused by sending the message `clearOldPasswords` to that user's `UserProfile`. As mentioned above, this set is maintained for each user when the `AllUsers` instance variable `disallowUsedPasswords` is set to `true`. The following example clears the remembered passwords for Mary's account:

```
topaz 1> printit
(AllUsers userWithId: 'Mary') clearOldPasswords.
System commitTransaction
%
```

Password Aging – Require Periodic Password Changes

You can configure your system so users must change their password periodically. This can be configured at the repository-wide level, or for individual users. Note that if you are not using GemStone login authorization, password aging does not apply.

Privileges required: OtherPassword and write authorization to `DataCuratorObjectSecurityPolicy`. The special GemStone accounts are not disabled by password aging.

Repository-Wide Password Aging

You can require users to change their password periodically by setting up a password age limit. This can be set for all users in the repository, and can be overridden for specific individual users.

To set a password age limit for all users in the repository, use the method `UserProfileSet>>passwordAgeLimit: numberOfHours`.

For example, to set the limit to 120 days:

```
topaz 1> printit
AllUsers passwordAgeLimit: 120 * 24 .
System commitTransaction
%
```

The `passwordAgeLimit` is added to the time the password was last changed to determine when the password will expire. A setting of 0 (the default) disables password aging.

Each time this method is invoked, the action is recorded in the Stone's log.

If a user does not change the account's password within the specified interval, the account is disabled, and attempts to log in result in an error.

`DataCurator` or another user with the `OtherPassword` privilege can reactivate the disabled account by giving it a new password; see page 152 for details.

Password Age Limits for Individual Users

To override the repository-wide setting for password aging, you can set a password age limit for a specific user, by using the method

```
UserProfile>>passwordAgeLimit: numberOfHours.
```

For example, to set the limit for a Mary to 5 days:

```
topaz 1> printit
(AllUsers userWithId: 'Mary') passwordAgeLimit: 5 * 24.
System commitTransaction
%
```

If repository-wide password aging is enabled, you can also override this for individual users, such as managers or administrators, either by setting the `passwordAgeLimit` for that `UserProfile` to 0, or by executing `setPasswordNeverExpires:` with a true argument.

For example, to prevent the password used by batch jobs from expiring, execute:

```
topaz 1> printit
(AllUsers userWithId: 'BatchUser')
  setPasswordNeverExpires: true.
System commitTransaction
%
```

Repository-Wide Password Expiration Warning

Privileges required: OtherPassword and write authorization to `DataCuratorObjectSecurityPolicy`. This does not apply to system users, and has no effect if password aging is not enabled.

You can provide an automatic warning to users repository-wide whose password is about to expire by sending the message

```
UserProfileSet>>passwordAgeWarning: numberOfHours.
```

For example, to warn users who log in within five days of the time their password will expire, do this:

```
topaz 1> printit
AllUsers passwordAgeWarning: 5 * 24.
System commitTransaction
%
```

Logins within *numberOfHours* prior to expiration cause a `SecurityError` to be signaled.

Per-User Password Expiration Warning

You can provide a similar automatic warning to specific user using

```
UserProfile>>passwordAgeWarning: numberOfHours.
```

For example, to warn Mary within three days of the time her password will expire, do this:

```
topaz 1> printit
(AllUsers userWithId: 'Mary') passwordAgeWarning: 3 * 24.
System commitTransaction
%
```

Finding Accounts with Password About to Expire

Privileges required: OtherPassword.

You can find out which accounts have a password within the warning period set by `passwordAgeWarning`. To do this, send the message `findProfilesWithAgingPassword` to `AllUsers`. For example:

```
topaz 1> printit
AllUsers findProfilesWithAgingPassword
  collect: [ :u | u userId] .
%
an OrderedCollection
#1 qa1
#2 qa2
#3 qa3
```

Finding Out When a Password Was Changed

Privileges required: OtherPassword.

You can find out the last time the password was changed for a particular `userId` by sending the message `lastPasswordChange` to that account's `UserProfile`. This example converts the `DateTime` returned to a particular pattern based on MM/DD/YY:

```
topaz 1> printit
(AllUsers userWithId: 'qa2') lastPasswordChange US12HrFormat
%
03/24/13 11:28 AM
```

Account Aging – Disable Inactive Accounts

You can configure your system so users must log in periodically, by disabling accounts for which there has been no login for a specified length of time. This can be configured at the repository-wide level, or for individual users.

Privileges required: OtherPassword and write authorization to `DataCuratorObjectSecurityPolicy`. The special GemStone accounts are not disabled by stale account aging.

Repository-Wide Stale Account Aging

To do this, send the message `staleAccountAgeLimit`: *numberOfHours* to `AllUsers`. This example disables accounts when they have not logged in for 30 days:

```
topaz 1> printit
AllUsers staleAccountAgeLimit: 30 * 24.
System commitTransaction
%
```

Each time this method is invoked, the action is recorded in the Stone log.

A setting of 0 (the default) disables account aging.

`DataCurator` or another user with the OtherPassword privilege can reactivate the disabled account by giving it a new password; see page 152 for details.

Per-User Stale Account Aging

You can override the repository-wide setting for account aging for a specific user using `UserProfile>>staleAccountAgeLimit: numberOfHours`.

For example, for the 'Auditor' account who may log in less frequently, you can set up a 180-day stale account age limit. This will override a repository-wide 30 day setting.

```
topaz 1> printit
(AllUsers userWithId: 'Auditor')
  staleAccountAgeLimit: 180 * 24.
System commitTransaction
%
```

Finding Out When an Account Last Logged In

Privileges required: OtherPassword.

If at least one age limit applies to an account, you find out when that account last logged in by sending the message `lastLoginTime` to that account's `UserProfile`. For example:

```
topaz 1> printit
(AllUsers userWithId: 'Mary') lastLoginTime US12HrFormat
%
03/24/13 01:40 PM
```

The time of the last login is maintained only if `loginsAllowedBeforeExpiration` is set in that `UserProfile` or if at least one of these instance variables is set in `AllUsers` or the specific `UserProfile`: `passwordAgeLimit`, `passwordAgeWarning`, or `staleAccountAgeLimit`. If none of these features are enabled, the `lastLoginTime` may be nil, the time of the account creation, or a time representing a login during an earlier period when one of these features was enabled. This is also true if a feature that enables `lastLoginTime` recording has been enabled on more recently than the last login of the user.

The data curator may explicitly set the time of the last login, using the method `UserProfile >> lastLoginTime:.`

Enabling Account Aging and lastLoginTime

Privileges required: OtherPassword.

The time of the last login is recorded for a `UserProfile` only if `loginsAllowedBeforeExpiration` is set in that `UserProfile` or if at least one of these instance variables is set in `AllUsers` or the specific `UserProfile`: `passwordAgeLimit`, `passwordAgeWarning`, or `staleAccountAgeLimit`. This is to avoid the commit during login, which is required to record the `lastLoginTime`.

As a result, you should use caution in enabling account aging on existing repositories. Enabling account aging may result in user accounts being disabled.

To avoid this, when you enable account aging, you can set the `lastLoginTime` to the current date, or to nil, for all affected `UserProfiles`. A nil setting disables account aging checks, allowing the aging period to begin with the next login.

```

topaz 1> printit
AllUsers passwordAgeLimit: 120 * 24.
AllUsers do: [:aUserProfile |
    aUserProfile lastLoginTime: DateTime now.
].
System commitTransaction.
%
```

Limit Logins Until Password Is Changed

Privileges required: OtherPassword.

When you assign a password to an account, you can make the password temporary by limiting the number of times it can be used. This limitation applies only to a specific account, that is, to the UserProfile that is the receiver of the message. It is intended for use with a new or reactivated account as a means of ensuring that the user changes the password. For example, the following limits the account "qa2" to two more logins under the current password:

```

topaz 1> printit
(AllUsers userWithId: 'Mary')
    loginsAllowedBeforeExpiration: 2.
System commitTransaction
%
```

A setting of 0 (the default) disables this feature.

The limit remains in effect until the user changes the password (see pages 145). Once the password is changed, the limit for that account is set to 0. The password will not expire again unless a new limit is set by repeating `loginsAllowedBeforeExpiration:`.

If the limit is exceeded before the password is changed, the system disables the account. DataCurator or another user with the OtherPassword privilege can reactivate the disabled account by giving it a new password; see page 152 for details.

The special user accounts are not disabled by this mechanism.

Limit Concurrent Sessions by a Particular UserId

Privileges required: OtherPassword.

You can limit the number of concurrent sessions logged in under a particular `userId` by sending the message `activeUserIdLimit: aPositiveInteger` to the UserProfile for that account.

For example, the following limits the `userId` "qa2" to four concurrent sessions:

```

topaz 1> printit
(AllUsers userWithId: 'qa2') activeUserIdLimit: 4.
System commitTransaction
%
```

A setting of 0 (the default) disables this feature.

If a user attempts to log in when the maximum number of sessions for that `userId` are already logged in, the login is denied and a SecurityError is signalled.

Limit Login Failures

Record Login Failures

The Stone repository monitor keeps track of login failures (incorrect passwords) for each account and can write that information to the Stone's log.

By default, messages are logged when the same account fails login attempts 10 or more times within ten minutes. You can change the default limits by setting the `STN_LOG_LOGIN_FAILURE_LIMIT` and `STN_LOG_LOGIN_FAILURE_TIME_LIMIT` configuration options (see page 296).

The log message gives the following information:

```
---Fri 10 May 2013 09:39:40 PDT ---  
_____  
GemStone user Mary has failed on 10 attempt(s)  
_____  
to log in within 1 minute(s).  
_____  
The last attempt was from user account writer1 on host  
name docs.
```

Disabling Further Login Attempts

If login failures continue, the Stone repository monitor can disable the account. By default, the account is disabled when the number of failures exceeds 15 within 15 minutes. You can change the default limits by setting the `STN_DISABLE_LOGIN_FAILURE_LIMIT` and `STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT` configuration options (page 291).

Subsequent attempts to login as that account result in the following error message:

```
Login failed: the GemStone userId/password combination is invalid  
or expired.
```

The special user accounts are not disabled by login failures.

DataCurator or another user with the OtherPassword privilege can reactivate the disabled account by giving it a new password; see page 152 for details.

6.6 Tracking User Logins

Login logging

It is sometimes useful to keep a log recording when each session logs in and out. You can configure your system to record login/logout events for each session in the repository. Other related operations, such as Stone startup and shutdown, are also recorded.

Tracking login/logout events is disabled by default. It is enabled by setting the configuration parameter `STN_LOGIN_LOG_ENABLED` to true in the configuration file used by the Stone, prior to Stone startup. See page 296 for details on this configuration parameter.

Once this is enabled for the repository, by default, all sessions that login to the stone will have logins and logouts logged. You can specify specific UserProfiles that will not have events logged, to avoid having the log cluttered with system logins. This is done with the method `UserProfile >> disableLoginLogging`. After this is executed, that UserProfile will not have logins or logouts recorded in the log file.

Logins and logouts are recorded to a text file named `stoneName_login.log`, in the same directory as the Stone log. Each log entry is on an individual line, with the following fields:

```
TimestampString TimeStampSeconds EventKind UserName SessionId
ProcessId RealUserID EffectiveUserID HostName GemIPAddress
ClientIPAddress NumCommits
```

For example:

```
"10/03/2013 16:42:48.720" 1380843768 STARTUP Stone 0 15270 631 631
kata.gemtalksystems.com 204.45.122.94 0.0.0.0 0
"10/03/2013 16:43:13.017" 1380843793 LOGIN DataCurator 3 15317 631
631 kata.gemtalksystems.com 204.45.122.94 10.94.141.15 0
"10/03/2013 16:43:23.488" 1380843803 SHUTDOWN Stone 0 15270 631
631 kata.gemtalksystems.com 204.45.122.94 0.0.0.0 0
```

Login Hook

The loginHook is an optional feature that allows you to specify code to be executed each time a specific userProfile logs in.

The argument to the method `UserProfile >> loginHook`: specifies either a symbol, which must be a unary selector of an instance method that was added to UserProfile, or a zero-argument block. When the userProfile logs in, if the `loginHook`: is not nil, than the method associated with the selector, or the block, is executed.

For example:

```
topaz 1> printit
(AllUsers userWithId: 'Mary')
  loginHook: [MyApplicationLogger logLogin: 'Mary']
System commitTransaction.
%
```


Managing Repository Space

The *repository* is the logical unit that represents the universe of shared objects that are stored within a GemStone/S 64 Bit system. Within GemStone Smalltalk, the repository is the single instance of Class Repository, with the name `SystemRepository`.

The logical repository maps to one or more physical *extent* files in the file system or to data on one or more raw disk partitions. Chapter 1 explains how this mapping is done through GemStone configuration options. Initially, the repository is contained in a single file, `$GEMSTONE/data/extent0.dbf`.

Whenever GemStone performs a *checkpoint*, it makes sure that transactions committed before the checkpoint have been written to the repository extents. The `STN_CHECKPOINT_INTERVAL` configuration option (page 289) sets the maximum time between checkpoints. (The default is five minutes, but various factors may cause a checkpoint to occur sooner.) The checkpoint limits the amount of time that is needed to recover from a system crash by guaranteeing that the data for the transaction is written to the extent and not just to the transaction log. For information, see “Controlling Checkpoint Frequency” on page 53.

This chapter explains how the repository grows, and tells you how to perform a number of administrative tasks related to the repository.

7.1 Repository Growth

Repository extents not only have to hold the data in your database, they also need to hold the changes all the users make, and coordinate the views of each user so the user has a consistent view of the data. All these activities require space in the extents.

The repository begins in the compact form of `$GEMSTONE/data/extent0.dbf`. As repository activity progresses, the extent file expands for a variety of reasons, in increments of 1 MB or more. Not only does your new application data require space, but space is required for internal structures that organize and manage the objects. Sessions, and their transactional views of the repository, also require space.

Garbage objects – objects that are no longer referenced, or the older versions of objects, also use space in the repository until they are garbage collected. To manage the size of the extents, you need to regularly perform garbage collection on your GemStone repository. The frequency can vary from monthly to daily, depending on the amount of activity in the repository. Without garbage collection, the repository will continue to grow and be filled with wasted space. See Chapter 12, “Managing Growth,” for a discussion of garbage collection in GemStone.

7.2 How To Check Free Space

Use the methods `Repository>>fileSize` and `Repository>>freeSpace` to obtain reports about the logical repository as a whole.

The result of the message `fileSize` is the total size of the repository in bytes, including all extent files. If the repository consists of a single extent file, it is ordinarily the same result as you would obtain by using the operating system commands to find file size. For example:

```
topaz 1> printit
SystemRepository fileSize
%
153092096
```

The result of `freeSpace` tells how much space (in bytes) is available for allocation within the repository at its current size. Free space is equal to the sum (for all extents in the repository) of the number of free pages in each extent multiplied by the page size. This space does not include fragments on partially filled data pages.

```
topaz 1> printit
SystemRepository freeSpace
%
26411008
```

Depending on the configuration options selected and the available disk space, the Stone repository monitor may be able to create additional free space by enlarging the repository.

If your configuration has more than one extent, use `Repository>>fileSizeReport` to generate statistics about each individual extent and also totals for the entire repository. (The heading “Extent #1” identifies the primary extent regardless of its file name, which initially is `extent0.dbf`.)

For example:

```
topaz 1> printit
SystemRepository fileSizeReport
%
Extent #1
-----
Filename = /users/extents/extent0.dbf

File size = 208.00 Megabytes
Space available = 3.69 Megabytes

Extent #2
```

```

-----
Filename = /users/extents/extent1.dbf

File size =      74.00 Megabytes
Space available =  4.77 Megabytes

Totals
-----
Repository size = 282.00 Megabytes
Free Space =      8.45 Megabytes

```

The amount of free space in the repository can also be determined from the cache statistic FreePages. To obtain the free space, multiply FreePages by the page size, 16384.

7.3 How To Add Extents

GemStone provides two ways to add extents:

- ▶ You can add new extents at startup by editing your GemStone configuration file and adding extent names and sizes to the DBF_EXTENT_NAMES and DBF_EXTENT_SIZES configuration options (page 274). Append the new values to the existing entries, just before the semicolon (;) delimiter. The new extents will be created the next time the Stone starts up.
- ▶ You can add extents while the Stone is running by invoking the Smalltalk methods described in this section. These methods are especially useful in avoiding or resolving low disk space conditions because the change takes effect immediately.

To Add an Extent While the Stone is Running

To prevent the repository from becoming full, you can dynamically add another extent to the Stone configuration. This section describes the Smalltalk methods that allow you to do this.

For general information about multiple extents, see “To Configure the Repository Extents” on page 34.

Possible Effects on Other Sessions

When a new extent is dynamically added to the logical repository through Smalltalk, sessions that are currently logged in must have access to the new extent. The possibility exists that an online session may terminate because it cannot open a new extent. Reasons for this condition could range from the inability to start a remote page server process to file permission problems.

CAUTION

The operating system creates the new extents with the ownership and permissions of the Stone repository monitor process. If these permissions are not the same as

for other extents, you should use operating system commands to modify them as soon as possible. Such changes can be made without stopping the Stone.

A session's view of which files make up the logical repository is updated whenever one of the following events occurs:

- ▶ Users commit or abort the session.
- ▶ The Stone repository monitor hands out disk resources to the session.

An explicit **commit** or **abort** may succeed but then cause the session to be terminated because of the inability to mount new extents immediately after the **commit** or **abort** operation.

Repository>>createExtent:

Privileges required: FileControl.

The Smalltalk method `createExtent :extentFilename` creates a new repository extent with the given extent file name (specified as a String). The new extent has no maximum size. The extent must be located on the machine running the Stone process; NFS-mounted disks are only allowed if your stone is configured to allow NFS mounted disks (see "STN_ALLOW_NFS_EXTENTS" on page 289). For example:

```
topaz 1> printit
SystemRepository createExtent: '$GEMSTONE/data/extent2.dbf'
%
```

You can execute this method when other users are logged in.

The Stone creates the new extent file, and it also appends the augmented extent list to your configuration file:

```
# DBF_EXTENT_NAMES written by Stone (user Bob) on 2/12/14 12:56:18
PST
DBF_EXTENT_NAMES = "$GEMSTONE/data/extent0.dbf",
"$GEMSTONE/data/extent1.dbf",
"!TCP@mozart#dbf!/users/gemstone/data/extent2.dbf;"
```

If the given file already exists, the method returns an error and the specified extent is not added to the repository.

Creating an extent with this method bypasses any setting you may have specified for the `DBF_PRE_GROW` option at system startup (page 275). Because extents created with this method have no maximum size and do not have an entry in a list setting for `DBF_PRE_GROW`, they cannot be pre-grown. If the repository is using weighted allocation, the new extent will be given a weight equal to the average weight of all other extents. (For a discussion of weighted allocation, see page 34.)

If this method is run from a session on a host remote from the Stone, `extentFilename` must include a Network Resource String (NRS) specifying the Stone host. The syntax is shown above in the excerpt from the augmented configuration file. For information about NRS syntax, see Appendix C.

Repository>>createExtent:withMaxSize:

Privileges required: FileControl.

The Smalltalk method `createExtent:extentFilename withMaxSize:aSmallInteger` creates a new repository extent with the specified *extentFilename* and sets the maximum size of that extent to the specified size. You can execute this method when other users are logged in.

The size must be a non-zero positive integer representing the maximum physical size of the file in MB.

If the specified extent file already exists, this method returns an error and the extent is not added to the logical repository.

If the configuration file option `DBF_PRE_GROW` is set to True (page 275), this method will cause the newly created extent to be pre-grown to the given size. If the pre-grow operation fails, then this method will return an error and the new extent will not be added to the logical repository.

7.4 How To Remove an Extent

The only way to remove an extent file is by first performing a Smalltalk full backup and restore to move the contents of that extent to other extents.

Privileges required: FileControl.

Reducing the number of existing extents requires special steps to ensure data integrity. If you must remove an extent file, follow this procedure:

Step 1. Back up your repository using the Smalltalk full backup procedure described on page 196.

You cannot use an online or offline extent backup to remove or shrink extents (obviously).

Step 2. Shut down the Stone repository monitor.

Step 3. Modify the `DBF_EXTENT_NAMES` configuration parameter (page 274) to show the new extent structure.

Step 4. Restore the repository from your Smalltalk full backup. Follow the GemStone restore procedure described on page 204.

7.5 How To Reallocate Existing Objects Among Extents

If you want to reallocate existing objects among two or more extents, the procedure depends in part on whether you are also changing the number of extents. Because changes to `DBF_ALLOCATION_MODE` directly affect only the subsequent allocation of pages for new or modified objects, additional steps are necessary.

To Reallocate Objects Among a Different Number of Extents

If you are increasing or decreasing the number of extents and want to change allocation of existing objects as part of that operation, perform a Smalltalk full backup, then restore the backup after setting appropriate weights in the `DBF_ALLOCATION_MODE` configuration option (page 274).

For example, suppose your existing repository contains 800 MB and you want to divide them about equally between the existing extent and a new one. To populate each extent with about 400 MB, follow this procedure:

Step 1. Back up your repository, using the Smalltalk full backup procedure described on page 196.

Step 2. Shut down the Stone repository monitor.

Step 3. Modify the `DBF_EXTENT_NAMES` configuration parameter to show the new extent structure.

```
DBF_EXTENT_NAMES = $GEMSTONE/data/extent0.dbf,  
$GEMSTONE/data/extent1.dbf;
```

Step 4. Edit the `DBF_ALLOCATION_MODE` configuration option to reflect the intended distribution of pages (see “Allocating Data to Multiple Extents” on page 37). For example:

```
DBF_ALLOCATION_MODE = 10, 10;
```

Step 5. Restore the repository from your Smalltalk full backup, using the procedure beginning on page 204.

If objects in the repository were explicitly clustered using instances of `ClusterBucket` that explicitly specified the first extent, those objects may tend to migrate back to that extent over time. You can prevent such migration by placing size limits on the existing extent, or by explicitly reclustered those objects in the new extent using a `ClusterBucket` that specifies either an `extentId` of `nil` or the `extentId` of the new extent. For information about clustering, refer to the *GemStone/S 64 Bit Programming Guide*.

To Reallocate Objects Among the Same Number of Extents

Changes to `DBF_ALLOCATION_MODE` (page 274) directly affect only the subsequent allocation of pages for new or modified objects. When you restore into a repository with the same number of extents, the distribution of the original repository will be used in the restored repository, regardless of the `DBF_ALLOCATION_MODE`.

To change the allocation of existing objects, you can either restore into a repository with a different number of extents (as discussed on page 174) or you can specify a maximum size on the extent files, to force objects to be distributed as you want them.

For example, suppose your existing repository has three extents, and that you are running in sequential allocation mode. The first extent has 600 MB, while the second and third extents are 1 MB each (the minimum size). You now want to redistribute the objects so they are spread evenly over all three extents. You cannot simply change the `DBF_ALLOCATION_MODE`, and restore a backup into three extents; existing objects

would be distributed according to the original allocation mode, that is, entirely in the first extent. Only new objects created after the restore would be created evenly over the three extents.

To populate the three extents evenly, you can follow this procedure:

Step 1. Back up your repository, using the Smalltalk full backup procedure described on page 196.

Step 2. Shut down the Stone repository monitor.

Step 3. Edit the `DBF_EXTENT_SIZES` configuration option (page 274) to limit the size of the first extent temporarily to 200 MB. For example:

```
DBF_EXTENT_SIZES = 200, , ;
```

Step 4. Edit the `DBF_ALLOCATION_MODE` configuration option (page 274) to reflect the intended distribution of pages (see “Allocating Data to Multiple Extents” on page 37). This setting determines the distribution of new or modified objects. For example:

```
DBF_ALLOCATION_MODE = 10, 10, 10;
```

Step 5. Restore the repository from your Smalltalk full backup, using the procedure beginning on page 204.

Step 6. If you want the first extent to grow beyond the temporary limit you set in Step 3, stop the Stone after you restore the repository. Edit the configuration file again, either specifying a higher limit or no limit. For example:

```
DBF_EXTENT_SIZES = , , ;
```

If objects in the repository were explicitly clustered using instances of `ClusterBucket` that explicitly specified the first extent, those objects may tend to migrate back to that extent over time. You can prevent such migration by maintaining the size limit set in Step 3, or by explicitly recluster those objects in the new extent using a `ClusterBucket` that specifies either an `extentId` of `nil` or the `extentId` of the new extent. For information, refer to the *GemStone/S 64 Bit Programming Guide*.

7.6 How To Shrink the Repository

Privileges required: `SystemControl`, `GarbageCollection`, and `FileControl`.

To shrink the repository to its minimum size, make a Smalltalk full backup. Then take the repository offline and restore the backup into a copy of the GemStone distribution repository. Use the following procedure, which compacts the repository into the minimum set of consecutive pages.

Step 1. Mark your repository for garbage collection, and wait for GemStone to complete the garbage collection and reclaim the space.

For example:

```
topaz 1> printit
SystemRepository markForCollection
%
```

The time required depends on several factors: the size of the repository, the number of Reclaim Gem sessions currently running, and (in multi-user mode) the status of other sessions.

If other users are logged in, space will not be reclaimed until all sessions have committed or aborted any transactions concurrent with the `markForCollection` process.

For further information on `markForCollection` and the garbage collection process, see Chapter 12, "Managing Growth", starting on page 233. Details on the `markForCollection` method are under "MarkForCollection" on page 242.

Step 2. Make a Smalltalk full backup of your repository by sending it the message `fullBackupTo:` or `fullBackupTo: MBytes:`.

For example:

```
topaz 1> printit
SystemRepository fullBackupTo: '/users/bk/march20'.
%
```

This example writes the backup to a single disk file. If you need to write multiple files, see "The `fullBackupTo: Methods`" on page 197.

Step 3. Take the repository offline:

```
topaz 1> printit
System shutDown
%
```

Step 4. Remove the existing repository extents. Obtain a copy of the distribution repository as the first (primary) extent by using the `copydbf` command. For example, assuming that all of your GemStone extents are in `$GEMSTONE/data`:

```
% cd $GEMSTONE/data
% rm extentNames
% copydbf $GEMSTONE/bin/extent0.dbf primaryExtentName
```

Use `chmod` to set the extent permission to what you ordinarily use for your repository.

Step 5. To put the repository back online, issue the `startstone` command:

```
% startstone gemStoneName
```

If you do not specify `gemStoneName`, `startstone` defaults to `gs64stone`.

Step 6. Log in to linked Topaz again.

NOTE

To perform the remaining parts of this procedure, you must be the only user logged in to GemStone. Logins will be disabled when you start the next step.

Step 7. Restore the repository by using the method

`Repository>>restoreFromBackup:fileOrDevice`, using the same file or device as in Step 2. Because it is being restored into a copy of the initial repository, the restored

repository will be compressed to the minimum space. This example restores the backup from a single disk file:

```
topaz 1> printit
SystemRepository restoreFromBackup: '/users/bk/march20'
%
```

If you need to restore multiple files, use the method `Repository>>restoreFromBackups :fileOrDeviceArray` instead:

```
topaz 1> printit
SystemRepository restoreFromBackups:
#( '%/backups/February_20.1'
    '%/backups/February_20.2' )
%
```

(For more information, see “Restoring from a Full Backup” on page 204.)

GemStone reads the backup(s) and rebuilds the repository in a “shadow” object space that is invisible to users at this time. If the restore succeeds, GemStone commits the restore and returns a summary in the form of a nonfatal error message like the following:

```
Restore from full backup completed with 616227 objects
restored.
```

Each restore operation, on completion, terminates its GemStone session. You will need to log in again before performing the next restore operation.

Step 8. Between the time the full backup was started (Step 2) and the time the repository was shut down (Step 3), there may have been transactions on your repository. To ensure that no work is lost, restore from transaction logs and commit the restore. For example:

```
topaz 1> printit
SystemRepository restoreFromCurrentLogs
%
Restore from transaction log(s) succeeded.
```

```
topaz> login
<details omitted>
successful login
```

```
topaz 1> printit
SystemRepository commitRestore
%
Restore from transaction log(s) succeeded. commitRestore
succeeded
```

7.7 How To Check Page Fragmentation

Space within the repository is managed in pages having a fixed size of 16 KB. It is possible for these pages to become *fragmented*—that is, only partially filled with objects. GemStone automatically schedules reclaim of pages with greater than 10% free space as part of its garbage collection activity.

You can inquire about the amount of fragmentation in the repository by executing the following expression.

```
SystemRepository pagesWithPercentFree:aPercentage
```

Typical values of *aPercentage* range from 10 to 25. This method returns an array containing the following statistics:

- ▶ The total number of data pages processed
- ▶ The sum (in bytes) of free space in all pages
- ▶ The page size (in bytes)
- ▶ The number of data pages having at least the specified percentage of free space
- ▶ The number of data pages having at least the specified percentage of free space, that contain only a single object
- ▶ The total number of pages in the repository that contain only a single object

`pagesWithPercentFree`: executes using the multi-threaded scan. See “Multi-Threaded Scan” on page 259 for details.

7.8 How To Recover from Disk-Full Conditions

The Stone repository monitor has two critical needs for disk space. It must be able to:

- ▶ Append to the transaction log as sessions commit changes.
- ▶ Expand the repository as necessary to allocate free pages to current sessions or to sessions logging in.

Whenever the Stone cannot log transactions or cannot find sufficient free space in the repository, it issues an error message to any session logged in as `DataCurator` or `SystemUser`. If users report that GemStone appears to be hung or that they get a disk-full error while logging in, you should check one of these administrative logins for such a message. The message is also written to the Stone’s log file.

The following topics explain the Stone’s actions in greater detail and describe steps you can take to provide sufficient space.

For details on how tranlog full conditions are handled, see “How To Recover from Tranlog-Full Conditions” on page 190.

Repository Full

The Stone takes a number of actions to prevent the repository from becoming completely full. If the free space remaining in the repository falls below the level set by the `STN_FREE_SPACE_THRESHOLD` configuration option (page 293) and the Stone cannot allocate more space in any extent, it takes the following actions to prevent a system crash:

1. It becomes more aggressive about disposing of commit records so that garbage collection can proceed. (If the Stone is very busy, a backlog of commit records can accumulate.)
2. It starts a checkpoint if there isn't one in progress and reduces the checkpoint interval to three minutes until the condition clears. (This checkpoint may free pages that have been reclaimed.)
3. It writes a message to the Stone log to indicate the condition.
4. It prevents new logins except for DataCurator and SystemUser accounts. It issues a disk-full error to other sessions attempting to log in.
5. It signals the Exception RepositoryError to any sessions logged in (or logging in) as DataCurator or SystemUser so that they know disk space is becoming critical.
6. It signals Gem session processes to return all except five free pages per extent. It responds to requests for additional pages by allocating only five pages at a time.
7. If the free space available drops below 400 KB (50 pages), the Stone stops responding to page requests from sessions that are not logged in as DataCurator or SystemUser. This action prevents users from acquiring all of the available space, which would cause the system to crash. Gem session processes appear to “hang” while they are waiting for pages. The unhonored page requests are granted when the free space goes back above the threshold.
8. If the previous steps do not solve the problem within the time specified by the `STN_DISKFULL_TERMINATION_INTERVAL` (page 292), then the Stone begins to terminate sessions holding on to the oldest commit record *even if the session is in a transaction*. This action applies to any user session, including logins as SystemUser and DataCurator. Allowing the Stone to dispose of the commit record allows additional garbage collection.

NOTE

You can configure the Stone to never terminate sessions by setting `STN_DISKFULL_TERMINATION_INTERVAL` to 0; however, doing so increases the risk of GemStone shutting down because of a lack of free space in the repository.

9. When the condition clears, another message is written to the Stone log and operation returns to normal.

If you see a message like the following while logged in as DataCurator or SystemUser or in the Stone log, disk space is becoming critical:

The repository is currently running below the freeSpaceThreshold.

When the system must dynamically expand the repository, it checks one extent at a time, in the order dictated by the allocation strategy, to see if that extent can be expanded to create more space. When no extents can be extended and the free space is below `STN_FREE_SPACE_THRESHOLD`, the Stone takes the actions described above.

Failure to expand an extent has two possible causes: either the disk containing the extent is full, or the extent has reached its maximum size as set by the `DBF_EXTENT_SIZES` configuration option.

There are a number of things you can do to create more space in an existing extent, or you can create a new extent. Each of these actions may create sufficient additional space for immediate needs:

- ▶ Warn the current users about the problem, and have them log out until enough space is made available.
- ▶ Remove any non-essential files to create enough space for expanding the repository.
- ▶ Create a new extent through Smalltalk with `Repository>>createExtent:` or a related method. If the Stone has stopped, you can create a new extent by editing the parameters in the configuration file before restarting it. See “How To Add Extents” on page 171.

Managing Transaction Logs

8.1 Overview

A transaction log contains the information necessary to redo transactions to the repository that have been committed by GemStone sessions since the last checkpoint or orderly shutdown. This log is used to recover from crashes such as those caused by a power failure, an operating system failure, or the killing of GemStone monitor processes.

If you need to restore the repository from a backup, transaction logs written in full-logging mode can be used to recreate all transactions committed since the most recent backup was written.

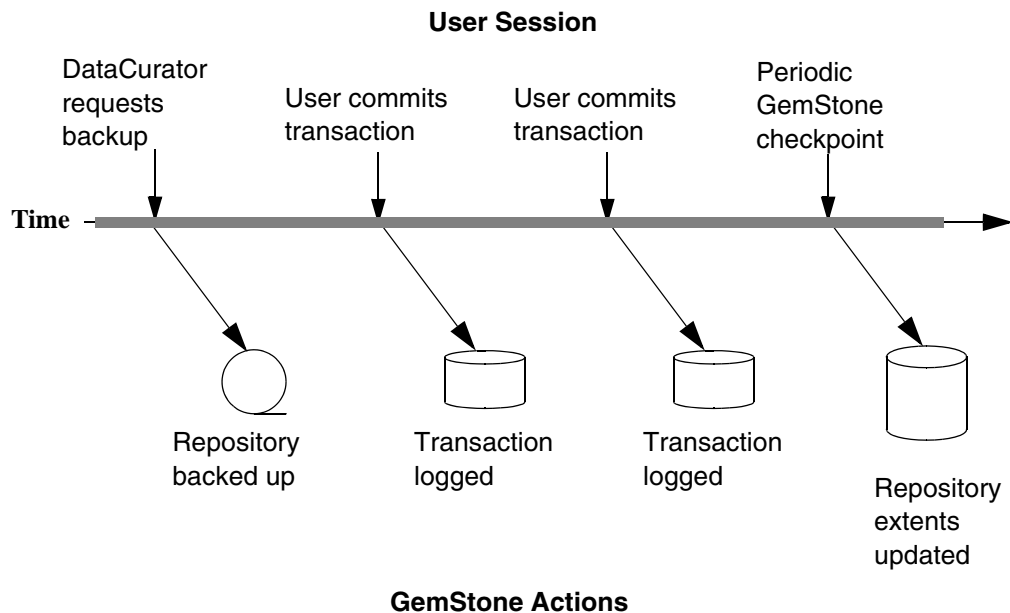
The transaction log is implemented as a sequence of files having names of the form `tranlog1.dbf ... tranlogNNN.dbf`. The numeric `fileId` starts at 0 when the Stone starts with a copy of the initial repository extent (`$GEMSTONE/bin/extent0.dbf`), and there are no existing transaction logs with that `tranlog id` in any transaction log directories. If the Stone starts on an existing repository without any logs present, the `fileId` will be one greater than when the repository was last shut down cleanly.

The filename prefix, by default “`tranlog`”, can be controlled by setting the `STN_TRAN_LOG_PREFIX` configuration option (page 305).

The transaction logs are written to a list of directories or raw partitions specified by the `STN_TRAN_LOG_DIRECTORIES` configuration option (page 304), which is treated as a circular list. Each log is limited to the size set for that directory or raw partition by `STN_TRAN_LOG_SIZES` (page 305). When one log is full, logging switches to the next directory or raw partition. (What happens when logs have been created in all directories is discussed in Table 1 on page 183.) Collectively, the transaction log files logically form an extremely large sequential file with a maximum size of 4×10^6 GB.

As users commit changes to the repository, GemStone writes the changes to a transaction log (Figure 8.1). Periodically, the stone repository monitor performs a checkpoint, at which point all committed changes are guaranteed to be written to the extents. During the periods between checkpoints, the transaction logs hold the record of the changes, so they are available if an unexpected shutdown occurs.

Figure 8.1 Normal Operation



Use ordinary UNIX utilities to backup the transaction logs in the file system. To backup a transaction log that is on a raw disk partition, use **copydbf** to copy it to a file system. You'll also need to use **removedbf** to clear the partition for reuse.

Logging Modes

GemStone provides two modes of transaction logging, selected by setting the `STN_TRAN_FULL_LOGGING` configuration option:

- ▶ To provide real-time incremental backup of the repository, set `STN_TRAN_FULL_LOGGING` to True. All transactions are logged regardless of their size. This mode is strongly recommended for deployed GemStone systems.
- ▶ To allow a simple operation without critical data to run unattended for an extended period, set `STN_TRAN_FULL_LOGGING` to False. This mode, known as *partial* logging, provides automatic recovery from system crashes that do not corrupt the repository, such as fatal errors or loss of power. However, if you do experience extent corruption and need to restore from backup, you cannot recover any changes made after the backup was taken.

Table 1 compares full and partial transaction logging.

Table 1 Comparison of Full and Partial Transaction Logging

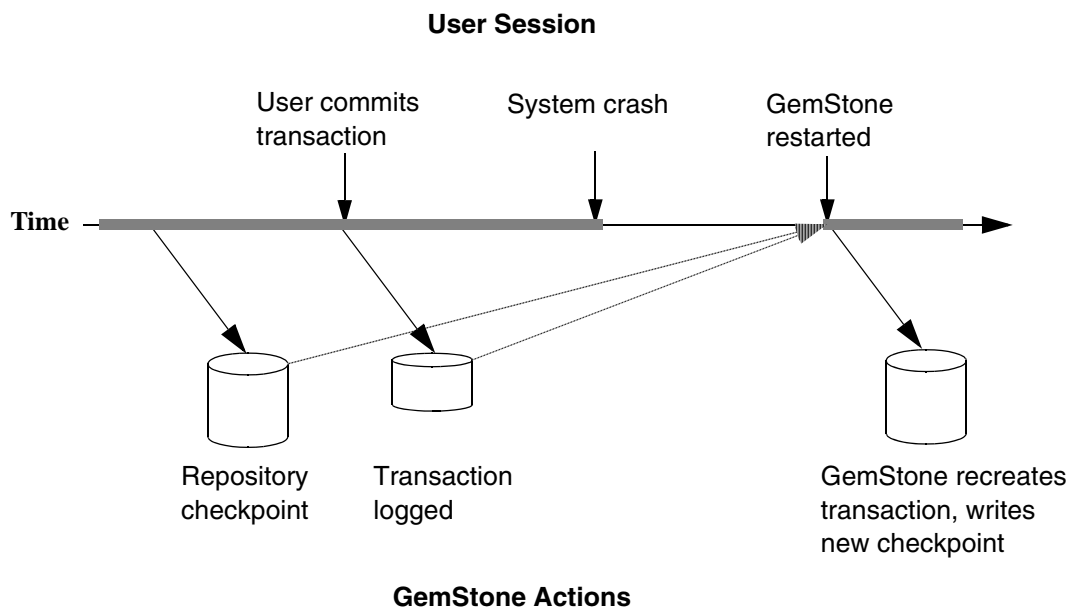
Characteristic	Full Logging STN_TRAN_FULL_LOGGING =TRUE	Partial Logging STN_TRAN_FULL_LOGGING =FALSE
Type of transaction logged	All transactions	Only those transactions smaller than STN_TRAN_LOG_LIMIT; successful commits of larger transactions cause an immediate checkpoint
Recovery from system crash (extents are OK)	Yes, automatic recovery during restart using checkpoint and log	Yes, automatic recovery during restart using checkpoint and log
Replay of transactions since last backup (as after media failure)	Yes – can carry forward GemStone backup by recreating subsequently committed transactions	No – cannot replay transactions since the backup
Action when current log is full	Logging moves to the next directory or to the head of the list. If it is a file system directory, the Stone opens a new log file there; existing transaction logs are retained. If it is a raw partition, a new log can be opened only if the previous one has been archived and removed. The maximum number of file system logs online at one time depends on disk space. The maximum number of raw partition logs depends on the number of partitions listed in STN_TRAN_LOG_DIRECTORIES.	Logging moves to the next directory or to the head of the list. The Stone removes the existing transaction log before opening a new one. The maximum number of logs on line at one time depends on the number of directories or raw partitions in the list.
Action when log space becomes full	The Stone pauses execution if it cannot find space in any of the specified directories or raw partitions.	The Stone deletes log files from the circular list of directories and keeps running.
Administrative task	Monitor log space; archive log files and delete them as necessary	None

Recovering from an Unexpected Shutdown

In the event of a system crash, GemStone can recover by automatically replaying transactions recorded in the transaction log, from the latest checkpoint before the crash to the end of the log at the moment of the crash (Figure 8.2). This allows you to restart from where you were at the time of the crash.

No special administrative action is required to perform this recovery. Every time the system does a normal startup, it checks to see if there are transactions that need to be recovered. Thus, if the extents and transaction logs are good, the system will automatically recover from a crash.

Figure 8.2 System Crash



Restoring Transactions to a Backup

An important use of transaction logs is to restore transactions that were committed between the last full backup and a system failure. If you experience system failure and your current extents are no longer usable, you can still recover all data, provided that the repository already is in full transaction logging mode and that the backup was made in that mode.

For details on restoring transactions to a backup, see “How to Restore Transaction Logs” on page 207.

How the Logs Are Used

You can only use transaction logs to restore your system from backup if you are in full logging mode.

After creating a backup, you need to retain all the transaction logs that are created during and after the backup. To determine the oldest transaction log that will be required, use the `copydbf` utility to query the backup file. For example:

```
unix> copydbf -i backup.dat
Source file: backup.dat
  file type: backup  fileId: 0
  byteOrder: Intel (LSB first)  compatibilityLevel: 844
  Full backup started from checkpoint at: 05/20/12 17:16:35
  PDT.
  Oldest tranlog needed for restore is fileId 4 ( tranlog4.dbf
  ).
  Backup was created by GemStone Version: 3.2.0 .
```

For this example, `tranlog4.dbf` and later must be available. If any of these transaction logs are deleted or lost, you will not be able to recover completely. These transaction logs may be archived elsewhere, as long as they can be readily be made available if you do need to restore from backup.

If you do need to restore your system, first restore from backup using the using the appropriate procedure, then replay transactions committed since the backup by reading the transaction logs in the order in which they were created.

For detailed procedures on restoring from backup and replaying transaction logs, see the instructions in Chapter 9, starting on page 191

NOTE

Restoring a repository from backup resets it - both the GemStone kernel and your application classes and data - to the state it was in at the time the backup was created. Anything that was done after that can be recovered only by replaying transaction logs in order, or by restoring a more recent backup.

8.2 How To Manage Full Logging

When the system is operating with the `STN_TRAN_FULL_LOGGING` configuration option set to `True`, you (as system administrator) should monitor the available log space. If the log space defined by `STN_TRAN_LOG_DIRECTORIES` becomes full, users will be unable to commit transactions to the repository until space is made available.

For transaction logs in file system directories, “full” means that there is no free space in the file systems containing those directories.

For transaction logs in raw partitions, “full” means that all partitions listed already contain a GemStone transaction log or other repository file. After archiving an existing log, you must invoke **`removedbf`** (page 318) before that partition can be reused.

There are two recovery situations to consider in managing transaction logs under full logging:

- ▶ Recovery from a system crash requires logs for all transactions committed since the last *checkpoint*. Because of the way GemStone logs changes involving large objects, parts of these transactions may be in earlier logs. The method `Repository>>oldestLogFileIdForRecovery` returns the `fileId` of the oldest

log that would be needed if a crash were to occur at that point. All logs needed for crash recovery should be kept online.

- ▶ Recovery from damaged extents, such as a media failure, requires all transaction logs since the last *backup*, and earlier logs may be needed if lengthy transactions were in progress at the time the backup started. Log files not needed for crash recovery may be archived offline, although restoring them will take longer.

To Archive Logs

Ordinary UNIX tools, such as **tar** and **cp**, can be used to move log files to other locations or to archival media. We recommend that you archive and free one complete log directory at a time, in the order listed in the `STN_TRAN_LOG_DIRECTORIES` configuration option (page 304).

NOTE

*If you must rename the log files, we recommend that you preserve the digits in the original filename as an aid to restoring the files in sequence should that become necessary. If your transaction logs are in raw disk partitions, **copydbf** adds the fileId when you copy a log to a file system directory.*

Two special commands are provided for working with raw disk partitions:

- ▶ The **copydbf** command copies a repository file (extent, transaction log, or full backup) to or from a raw disk partition. If the destination is a directory in the file system, **copydbf** generates a filename that includes the file type and its internal fileId.
- ▶ The **removedbf** command writes over a raw partition so that GemStone will no longer think it contains a repository file.

Both the **copydbf** and **removedbf** commands can be used with a remote node that is not running NFS, provided that a NetLDI is running on that node. For further information about **copydbf** and **removedbf**, see the command descriptions in Appendix B.

To determine the current size of a transaction log that is in a raw partition, use the method `Repository>>currentTranlogSizeMB`. This method returns the log size (in MB) as an Integer.

To determine the oldest transaction log that would be needed to recover from the most recent checkpoint, use the method `Repository>>oldestLogFileIdForRecovery`. This method returns the internal fileId, which is part of the filename for transaction logs in the file system. If a session was in a lengthy transaction at the time of a system crash, the oldest log needed during recovery may be one that was written before the last checkpoint occurred.

You can obtain similar information by applying **copydbf -i** to an extent. For example:

```
% copydbf -i extent0.dbf
Source file: extent0.dbf
  file type: extent  fileId: 0
  byteOrder: Sparc (MSB first) compatibilityLevel: 844
  Last checkpoint written at: 04/04/13 15:30:30 PDT.
  Oldest tranlog needed for recovery is fileId 5 ( tranlog5.dbf
  ).
  Extent was shutdown cleanly; no recovery needed.
```

To determine the oldest transaction log needed to roll forward from a backup, apply `copydbf -i` to the backup:

```
% copydbf -i backup.dat
Source file: backup.dat
file type: backup  fileId: 0
byteOrder: Sparc (MSB first)  compatibilityLevel: 844
Full backup started from checkpoint at: 04/04/13 15:28:37 PDT.
Oldest tranlog needed for restore is fileId 5 ( tranlog5.dbf ).
Backup was created by GemStone Version: 3.2.0 .
```

For an example script showing how to archive transaction logs out of raw partitions, see `$GEMSTONE/examples/archivelog.sh`. You will need to edit the script to conform to your own partition names and archive location, and then test it.

To Add a Log at Run Time

Privileges required: FileControl.

You can add a directory or a raw partition for transaction logs to the existing list without shutting down the Stone repository monitor. When you do this, the repository monitor also records the change in its configuration file so that the addition becomes permanent. Send the following message:

```
SystemRepository addTransactionLog: deviceOrDirectory size: aSize
```

For example:

```
topaz 1> printit
SystemRepository addTransactionLog: '/users/tlogs2' size: 100
%
```

The argument *aSize* sets the maximum log size (in MB) for *deviceOrDirectory*. It will be added to the list in `STN_TRAN_LOG_SIZES` (page 305).

You can use the method `System class>>stoneConfigurationAt:` to examine the contents of `STN_TRAN_LOG_DIRECTORIES` at run time. For information, see “How To Access the Server Configuration at Run Time” on page 50. The Repository methods in Table 1 return other information that is helpful in managing transaction logs.

Table 1 Repository Methods for Information About Transaction Logs

Method	Description
<code>currentLogDirectoryId</code>	Returns a positive <code>SmallInteger</code> , which is the one-based offset of the current log file into the list of log directory names.
<code>currentLogFile</code>	Returns a <code>String</code> containing the name of the transaction log file to which records currently are being appended.
<code>currentTranlogSizeMB</code>	Returns an <code>Integer</code> that is the size (in MB) of the currently active transaction log.
<code>allTranlogDirectories</code>	Returns an <code>Array of Strings</code> corresponding to the <code>STN_TRAN_LOG_DIRECTORIES</code> configuration.

Table 1 Repository Methods for Information About Transaction Logs

Method	Description
<code>allTranLogSizes</code>	Returns an Array of SmallIntegers corresponding to the <code>STN_TRAN_LOG_SIZES</code> configuration.
<code>logOriginTime</code>	Returns the log origin time of the receiver, the time when a new sequence of log files was started. For details, see the method comment in the image.
<code>oldestLogFileIdForRecovery</code>	Returns a positive SmallInteger, which is the internal fileId of the oldest transaction log needed to recover from the most recent checkpoint, if the Stone were to crash as of now.
<code>restoreStatusOldestFileId</code>	Returns a SmallInteger, which is the internal fileId of the oldest transaction log needed for the next restore from log operation.

To Force a New Transaction Log

Privileges required: FileControl.

You can force closure of the current log and opening of a new log at almost any time by using the method `Repository>>startNewLog`. The method performs the following sequential actions:

1. Starts a checkpoint.
2. Waits until the checkpoint completes.
3. Starts the new log.
4. Returns a SmallInteger, which is the `fileId` of the new log.

In the following example, the new transaction log file would be `tranlog9.dbf`.

```
topaz 1> printit
SystemRepository startNewLog
%
_
```

If a checkpoint is already in progress when you execute `startNewLog`, the method will fail and return `-1` instead. If you're using this method in an application, therefore, you need to accommodate the possibility of such a failure with code such as:

```
| id |
id := SystemRepository startNewLog.
[ id < 0 ] whileTrue: [
    System sleep: 1.
    id := SystemRepository startNewLog ].
```

To Initiate a Checkpoint

Privileges required: SystemControl.

System class provides two methods that you can use to start checkpoints manually. These methods do not commit, abort, or otherwise modify the current transaction.

```
startCheckpointSync
```

Starts a new checkpoint and returns when the new checkpoint has completed. If a checkpoint is already in progress, this method waits until the current checkpoint completes, then starts a new checkpoint. Returns true if a new checkpoint was successfully completed, returns false if a new checkpoint could not be started because transaction logs are full or checkpoints are suspended.

`startCheckpointAsync`

Starts a new checkpoint if a checkpoint is not already in progress and returns immediately, without waiting for the checkpoint to finish. Returns true if a checkpoint is in progress or was started; returns false if a checkpoint could not be started because transaction logs are full or checkpoints are suspended.

To Change to Partial Logging

Once the full transaction logging has been started on a repository, the `STN_TRAN_FULL_LOGGING` state of True persists regardless of later changes to the configuration file.

To terminate full logging, the procedure is:

Step 1. Make a Smalltalk full backup (not an extent snapshot), following the instructions starting on page 196.

Step 2. Edit the configuration file to set the `STN_TRAN_FULL_LOGGING` option to False.

Step 3. Stop the Stone repository monitor, and restore the backup following the instructions starting on page 204. into a copy of the empty distribution extent (`$GEMSTONE/bin/extent0.dbf`).

Step 4. Restart GemStone.

8.3 How To Manage Partial Logging

Partial logging provides ease of administration along with some protection against loss of data from system crashes. The Stone repository monitor treats the log directories as a circular list. If the file in the current directory reaches the limit set by `STN_TRAN_LOG_SIZES`, the Stone switches to the next directory in the list that does not contain a transaction log. In the process of creating log *n*, the Stone attempts to find and delete log (*n* - `size_of_STN_TRAN_LOG_DIRECTORIES`); for example, if the new log will be `tranlog7.dbf` and there are three elements in `STN_TRAN_LOG_DIRECTORIES`, the Stone searches all three in attempting to delete `tranlog4.dbf`.

If there is only one directory specified in by the `STN_TRAN_LOG_DIRECTORIES`, then the stone deletes the log file in this directory before starting a new log. This means that badly-timed crashes may not be recoverable. You should ensure that there are two directories specified and that there is sufficient disk space for at least two log files, so that one can be preserved when the next is opened.

To Change to Full Logging

To change a repository from partial to full logging, simply change the `STN_TRAN_FULL_LOGGING` setting to True (page 303) and restart the Stone repository monitor.

Be sure to make a new GemStone full backup in full-logging mode so that you will be able to restore from the transaction logs if necessary. Transaction logs cannot be restored to backups that were made in partial-logging mode.

8.4 How To Recover from Tranlog-Full Conditions

Transaction Log Space Full

If the space for transaction logs becomes full, the Stone stops processing commits or other requests that initiate a write to the transaction log. Sessions performing these operations are blocked until the condition is resolved and may appear to the user to be hung. The Stone writes a message like the following in its log:

The tranlog directories are full and the stone process is waiting for an operator to make more space available by either cleaning up the existing files (copying them to archive media and deleting them) or by adding a new tranlog directory.

Also, the `Exception RepositoryError` is signaled in any sessions that have enabled receipt of this error by sending `System enableSignalTranlogsFull`, and setting up a handler for this error.

Once enabled, you can disable receipt of this error by sending `System disableSignalTranlogsFull`, and determine the current status by `System signalTranlogsFullStatus`.

If the transaction log space is full, you have the following options:

- ▶ You can free space by taking some existing log files offline. Archive them using operating system utilities and then remove them. GemStone can reuse that slot in the circular list of log directories. (To archive and remove a log file from a raw partition, use `copydbf` and then `removedbf`.)
- ▶ You can increase the available log space by adding a raw partition or a directory on another disk drive to the `STN_TRAN_LOG_DIRECTORIES` configuration option (page 304). Add its maximum file size to `STN_TRAN_LOG_SIZES`. For information on how to make these changes while GemStone is running, see “To Add a Log at Run Time” on page 187.

When transaction log space becomes available, waiting sessions can complete operations that were blocked.

Making and Restoring Backups

This chapter describes how to make backups of your GemStone/S 64 Bit repository and how to use the backups and transaction logs to restore the repository.

This chapter includes the following topics:

- ▶ The different types of backups, and how to choose a backup strategy (page 192)
- ▶ How To Make an Extent Snapshot Backup (page 193)
- ▶ How To Make a Smalltalk Full Backup (page 196)
- ▶ How to Restore from Backup (page 200)
This includes both Restoring from an Extent Snapshot Backup (page 203) and Restoring from a Full Backup (page 204).
- ▶ How to Restore Transaction Logs (page 207)

9.1 Overview

To safeguard your repository, you should create a backup of your GemStone repository periodically, and store the backup in a safe place. Backups provide security in case of problem with power, operating system, disks, or other system corruption, and if used in combination with transaction logs, preserve all committed data against loss.

Making a backup of the GemStone repository captures the state of the system at a particular moment in time, and restoring that backup can return your system to the state it was in at the time the backup started. A GemStone backup is a backup of not only your application data, but also your application code and GemStone kernel code, and of user profiles and passwords and so on - everything in the repository. Because the backup includes kernel code, backups can only be restored into the same version of GemStone as that in which the backup was created; otherwise the kernel classes and methods may not be appropriate for that version.

Between these periodic backups, transaction logs capture all committed changes that occur in the repository (provided the repository is in full logging mode). By preserving

the backup and a set of transaction logs, you have the ability to recreate the system up to the last committed change in the transaction logs.

In partial tranlog mode, the transaction logs cannot be applied after restoring a backup. In this case, the transaction logs are useful when recovering from transient problems such as unexpected shutdown, but restoring a backup can only restore the system to the state it was in at the time of the backup. Later transaction logs in partial logging mode cannot be applied to recover work done after the time of the backup.

You should establish a regular backup process and schedule that fits your application requirements, and a system of managing and archiving the backup files and transaction logs that will allow you to recover smoothly after any problems.

In addition to regular backups, to ensure protection from disk failure, we recommend that you either use mirrored disks or operating system mirroring. For more information, see “Developing a Failover Strategy” on page 27.

9.2 Types of Backups

GemStone supports three types of backup:

- ▶ Offline extent snapshot backups
- ▶ Online extent snapshot backups
- ▶ Smalltalk full backups.

Extent snapshot backups consist of operating system copies of the extent files.

When the repository is offline, and was cleanly shut down, the extent files can be copied using regular OS copy functions with no further considerations.

To make extent file copies of a repository that is in use (online), checkpoints must be suspended for the duration of the extent copy. The extents are updated during checkpoints, so if a checkpoint occurs during extent file copy, it is likely the backup files will be corrupted and unusable.

Smalltalk full backups are made by executing backup methods in GemStone code. These can only be created when the system is running. Executing the backup methods will cause all live objects in the repository as of the time the backup method execution began to be written out to one or more operating system files. Dead objects, and internal structures such as the object table, are not written out, so these files typically are somewhat smaller than the repository extent size (excluding free space in extents).

Determining which type of backup to make depends on the size of your repository and the uptime requirements.

- ▶ Offline extent backups are the most simple, since nothing is needed beyond clean shutdown and file copy. However, since these must be taken when the repository is shut down, they are not suitable for systems that must be available 24x7.
- ▶ Full backups are convenient to run in highly-available systems that are not shut down regularly. However, backup execution places load on the system, and should be avoided during periods of heavy system use. Restore from full backup will also take much longer than offline or online extent backups

Backups made using full Backups methods have other value: restoring these backup files can be used to change the number of extents, redistribute objects among extents, or reduce the size of extent files.

- ▶ Online extent snapshots require the most effort to setup, since checkpoints must be suspended for the entire duration of the file copy. Since file copy is limited only by throughput of the physical disks, for large repositories that are in use 24x7, online extent backups will have the smallest impact on availability and on performance. They will also be much faster to restore than full backups.

Full vs. Partial Transaction logging

As described in “Logging Modes” on page 182, your repository may be run in partial transaction logging mode, or in full transaction logging mode.

In partial transaction logging mode, you cannot make online extent backups, since checkpoints cannot be suspended while you are this mode.

While you can make Smalltalk full backups, or offline extent copy backups, you cannot restore transaction logs into these backups. If you need to restore from backup, any work done after the start of the backup is permanently lost. For repositories with valuable data, we recommend that you run in full logging mode to avoid data loss in case of extent corruption.

Verify Backup Process

Creating a backup and archiving transaction logs is only useful if you can restore them successfully in case of a system failure. To make sure that your procedures for archiving and restoring backups is complete and correct, it is good practice to perform the restore operation into a non-production system, replay tranlogs, and audit the restored repository. Audit information can be found starting on page 120.

Performing this exercise ensures that if you do have an emergency situation, you will have the required files available and be familiar with the process of restore, and avoid the risk of losing data.

9.3 How To Make an Extent Snapshot Backup

Extent snapshot backups are file system copies of the repository extents. These copies can be made when the repository is not running (offline); or when the repository is running (online), provided you suspend checkpoints for the duration of the extent file copy.

WARNING

File system copies of the extents of a running GemStone repository that are taken during a period that includes a checkpoint will have inconsistent state, and not be usable for restore.

Extent Snapshot Backup when the Repository is shutdown

When the repository is shut down, you can safely perform a file system backup of the extents files. During the shutdown process, a checkpoint is performed in which all committed transactions are written to the extents. A copy of the extents after an orderly shutdown constitutes a complete operating system backup of the repository without requiring any transaction logs.

If GemStone was not shut down cleanly, file system copies of the extents are usable, but they will not include any transactions committed since the last completed checkpoint before the shutdown. In order to recover later work, you will also need one or more transaction logs.

This applies for both partial logging and full logging.

copydbf -i will report if the extents were cleanly shutdown and the oldest tranlog required for recovery if the extents were not cleanly shutdown.

Extent Snapshot Backup when the Repository is running

When the repository is running, you must suspend checkpoints before starting the extent file copy, and resume checkpoints when the file copy is complete.

You should not attempt to take online extent snapshot backups when the repository is in partial logging mode (`STN_TRAN_FULL_LOGGING = FALSE`), since checkpoints cannot be suspended in partial logging mode.

Three steps are involved in an online extent backup

1. Suspend checkpoints.

Checkpoints are not permitted while the extent file are being copied for the online backup. There must not be a checkpoint in progress when the first extent file copy starts, and no checkpoints are allowed to begin until the last extent file copy has completed. All other database operations (including commits, aborts, and the creation of new tranlogs) are permitted during the online extent snapshot backup.

To suspend checkpoints for a specified number of minutes, call `System class >> suspendCheckpointsForMinutes : .` If this method is called while a checkpoint is already in progress, it will block until the current checkpoint completes. On some systems under heavy load, checkpoints may take some time to complete; the period in which checkpoints are suspended does not begin until the previous checkpoint is complete.

If one session attempts to suspend checkpoints and is blocked while the current checkpoint completes, and then a second session attempts to suspend checkpoints, the second session fails and the method returns false.

If the system is shut down while checkpoints are suspended, checkpoints will be re-enabled and a final checkpoint will be written during the clean shutdown process. Any extent snapshot backups in progress during system shutdown must be discarded.

To query the current status of checkpoints, call `System class >> checkpointStatus`. This method returns an Array object containing a Boolean that indicates whether checkpoints are suspended, and an Integer giving the number of seconds remaining in the suspension.

For example:

```
topaz 1> printit
System checkpointStatus
%
an Array
  #1 false
  #2 0
```

```
topaz 1> printit
System suspendCheckpointsForMinutes: 15
%
true
```

```
topaz 1> printit
System checkpointStatus
%
an Array
  #1 true
  #2 900
```

We recommend using a value of *minutes* that is much larger than any possible anticipated time, taking into consideration the amount of time backups may take after future repository growth. If checkpoints resume before the extent/s copy is complete, the snapshot will not be usable.

It is preferable to have checkpoints suspended for as short a time as possible, but it is safer for the backup script to manually resume checkpoints after the file copies are completed, rather than relying on tuning the time out period.

2. Copy the repository extents.

Once checkpoints are suspended, the session requesting the suspension can log out from GemStone and start the extent copy, using operating system commands or copy-dbf.

3. Resume checkpoints.

Once the extent copy has completed, a session should log in to GemStone and request the Stone to resume checkpoints (`System class >> resumeCheckpoints`). The result of this method is `false` if checkpoints were not previously suspended before executing `System class >> suspendCheckpointsForMinutes:` (as in step 1), and `true` if they were previously suspended.

```
topaz 1> printit
System resumeCheckpoints
%
true
```

From this result, you can determine if the online extent backup was completed while checkpoints were still suspended. If the backup was completed in time, no further action is required and the backup is complete. If the backup did not complete before checkpoints were resumed, then the backup must be discarded and another online extent backup must be taken.

CAUTION

Make sure your backup code checks this result, since a false return value means that your backup is not usable.

An Example Script

The GemStone installation directory includes an example script `$GEMSTONE/examples/admin/onlinebackup.sh`. You can customize this script for your own system.

This script does not include code to make file system copies of the extents; you must add the necessary code to perform this task. This script provides a default checkpoint suspension of 15 minutes, which may or may not be sufficient time.

NOTE

The example script `onlinebackup.sh` is unsupported. It is provided here for your convenience, and is subject to change in future releases.

Be sure to review and test your script adequately to ensure the integrity of your backups.

9.4 How To Make a Smalltalk Full Backup

You can create a backup of the objects in your repository by performing Smalltalk full backups, using methods provided as part of the GemStone kernel. Smalltalk full backups are required if you want to reduce the number of extents in the repository or redistribute objects within the repository. During a Smalltalk full backup, dynamic internal data structures are not copied and will be rebuilt, which can, at least temporarily, improve the performance of such routine maintenance tasks as garbage collection.

In a Smalltalk full backup, the methods `Repository>>fullBackupTo:` or `fullBackupTo:MBytes:` save the most recently committed version of the repository in a way that is consistent from a transaction viewpoint. These methods force a checkpoint of the repository at the time the method is executed and then creates a backup from that checkpoint, copying all objects in the repository and arranging them in a compact form in one or more files.

You can make Smalltalk full backups while the repository is in use. Other sessions can continue to commit transactions, but those transactions are not included in the backup. Full backups require the GcLock, and so full backups cannot be made while other operations that hold the GcLock are running.

A Smalltalk full backup includes these three steps:

1. The Gem performing the backup scans the object table, building a list of objects to back up. This step runs in a transaction and can therefore cause a temporary commit record backlog in systems with high transaction rates. This step normally completes fairly quickly.
2. The Gem performing the backup next writes all shadow objects to the backup file. This step also runs in a transaction; furthermore, backing up shadow objects requires more disk I/O than backing up live objects, so the rate of objects backed up per second is slower in this step than in the next.

(For definitions of shadow and live objects, see “Basic Concepts” on page 233.)

3. In the final step, all remaining live objects are written to the backup file. This step is performed outside a transaction; if the Stone signals the session to abort, it will do so. This step takes the longest of the three.

The fullBackupTo: Methods

```
Repository>>fullBackupTo:filename
```

```
Repository>>fullBackupTo: arrayOfFileNames MBytes: mByteLimit
```

In these methods, *filename* or *arrayOfFileNames* specifies one or more files where the backup is to be created. You must specify the name of the files, not a directory name. You may include a relative or absolute path in addition to the file name.

If you use a relative path, the path is relative to the directory of the Gem process or linked session. For linked topaz sessions, this is the directory from which topaz was started. For RPC Gems, this is either specified by #dir: in the login parameters, or the home directory of the Gem's UNIX user.

You can create backups on a remote node by using a network resource string (NRS) to specify the node name as part of the file name, and ensuring a NetLDI is running on the remote node.

mByteLimit is either a single integer, or an array of integers with the same number of elements as *arrayOfFileNames*. This argument limits the maximum size of each file, except the last. If *mByteLimit* is one integer, each backup file will use that value; if it is an array of integers, each file will be limited by the matching entry. A value of 0 means the file sizes are unlimited.

In order to avoid running out of space for the backup, the last file is not limited, regardless of the size limit specified. If the number and size limit of *arrayOfFileNames* is too small to hold the entire backup, after each of the earlier files reaches its *mByteLimit*, the last file may grow significantly larger to contain the remainder of the backup.

WARNING

If there is not sufficient space to write the entire backup, the backup will return an error and deletes the incomplete backup files. Make sure you have sufficient disk space and the appropriate value for mByteLimit.

If you do not want to limit the size of the backup file, specify a *mByteLimit* of 0.

For example:

```
topaz 1> printit
"Create a full backup of the Repository"
SystemRepository
  fullBackupTo: {
    '/users/backups/bkup13.3.15-1' .
    '/users/backups/bkup13.3.15-2' .
    '/users/backups/bkup13.3.15-3'
  }
  MBytes: 0.
%
true
```

This writes the backup into three files, named `bkup13.3.15-1`, `bkup13.3.15-2`, and `bkup13.3.15-3`. Messages are written to the stone log indicating when the backup started and when it completed.

During the backup, after the initial period in transaction, the session is put into manual transaction mode so the backup won't interfere with ongoing garbage collection. When the backup completes, the session is left outside of a transaction. If you want to make changes to the repository after a backup, send `System beginTransaction` or `System transactionMode: #autoBegin`.

Backup fails to run or encounters an error

If the backup file already exists, a path cannot be found, or if any of the file names are empty strings, the method returns an error.

If another session is holding the GcLock, the backup will wait for up to 5 minutes for the other operations to complete and release the GcLock, otherwise it will fail and return an error. You can determine the session holding the GcLock by using:

```
System sessionIdHoldingGcLock
```

This method will return 0 if no session is holding the GcLock.

Backup (and restore) require at least one extra session be available, beyond the session that is starting the backup. If the number of users logged in is equal to the `STN_MAX_SESSIONS` setting, the backup will fail with an error.

If backup encounters an error, then any backup files that were created are automatically deleted.

Monitoring and Performance

The following performance and monitoring topics apply to full backups. The performance of on and offline extent snapshot backups depend on your operating system and disk performance and can be monitored and optimized outside of GemStone using OS level tools.

Shared Page Cache Size

You can often improve both backup and restore performance by increasing the size of the shared page cache.

Multi-threading

Full backups are written and restored multi-threaded to allow the reads and file writes to progress in parallel for faster performance.

The maximum number of threads is based on the number of extents in the repository and the number of backup files specified. For backup, threads may be deactivated and system impact reduced using the methods described in "Tuning Multi-Threaded Scan" on page 259. Restore is always done with maximum performance.

The multi-threaded algorithm uses one session per extent in your repository, up to a limit of 16 sessions. If there are not sufficient session slots available – if the number of users logged in is close to the `STN_MAX_SESSIONS` setting – then the backup or restore will use fewer sessions and performance will be slower. In this case, a message is printed to stdout (the topaz -l terminal) and to the stone log.

Cache Statistics

During the main part of a full backup, the statistic `ProgressCount` for the session performing the backup indicates the number of objects written to the backup file thus far. If you know the total number of objects in the repository, you can use this statistic to determine how far the backup has progressed.

Backups and Garbage Collection

NOTE

You will find it easier to understand the following discussion if you have first read and understood the section “Basic Concepts” on page 233.

Because shadow objects must be backed up, it is more efficient to run a Smalltalk full backup when there are few shadow objects. If possible, first check the statistic `PagesNeedReclaimSize`. If that statistic is high, run one or more Reclaim Gem sessions before performing the backup. (See “Admin and Reclaim Gems” on page 239.)

Dead objects waiting to be reclaimed (measured by the statistic `DeadNotReclaimedObjs`) are not backed up, as these objects are going to be deleted anyway.

Compressed Backups

It is possible to write and read full backup files in compressed mode.

Writing to, and reading from, a compressed file can be performed only to a local file system file or to a file system that is NFS-mounted.

Backup files written in compressed mode are automatically appended with the suffix `.gz` if you do not specify that suffix.

All restore methods automatically detect whether a file is compressed or not and read the file accordingly. Even a backup originally created in uncompressed mode, then later compressed externally with `gzip`, is readable by `restoreFromBackup:.`

The following class methods in `Repository` are provided to create compressed full backups:

`fullBackupCompressedTo: filename`

This method backs up the receiver to a single backup file in `gzip` format. The output file is written compressed in `gzip` format.

`fullBackupCompressedTo: arrayOfFileNames MBytes: mByteLimit`

This method is similar to `fullBackupTo:MBytes:` except that the output file is written compressed in `gzip` format.

Verifying a Backup is Readable

To verify that a backup file is readable, use the GemStone utility `copydbf`. You can conserve disk space and reduce disk activity by specifying `/dev/null` as the destination. For instance:

```
% copydbf /users/backup/bkup13.3.15-1 /dev/null
```

Checking Backup Start and Completion

The time a backup is started, and the time that it completes successfully, are written to the stone log. For multi-file backups, only the first filename is listed.

```
--- 02/20/14 10:19:52 PST ---
    Full backup of the repository has been started.
        Host: ip6-localhost          ProcessId: 2930
        User: DataCurator           SessionId: 5
--- 02/20/14 10:21:35 PST ---
    Full backup successfully completed by sessionId 5 to file:
bkup.dat
```

9.5 How to Restore from Backup

There are several circumstances under which you will want to restore from backup.

If you have disk errors or file corruption, or if you encounter object corruption in your repository, you will need to restore from backup and replay transaction logs to recover all work up to the time of the corruption.

Restoring from backup is also used to set up and refresh warm or hot standby systems, and to set up test environments that match production systems.

To make the repository smaller, or to redistribute objects among a different number of extents, or to change your system to use partial logging mode, you must restore from full backup. Restore from full backup may also improve space use and performance by recreating dynamic internal structures.

Note that if your intention is to redistribute objects over a different number of extents, if the number of extents during restore is the same as the number of extents when the backup was started, this takes precedence over the `DBF_ALLOCATION_MODE` configuration setting during restore. If the number of extents differs, then the `DBF_ALLOCATION_MODE` setting at the time of the restore controls the distribution of objects across extents.

The ability to restore from backup is critical to the reliability of your GemStone system. You should ensure that you regularly take backups, and from time to time, verify that the processes that you use to make the backups result in complete and usable backup files.

If you are concerned about losing work that is done between backups, ensure that you are in full transaction log mode. In this mode, the transaction logs record all commits in your repository and the transaction logs are not automatically deleted, so they can be replayed into a restored backup if they are needed.

There are two phases of restoring from backup:

Phase 1 - restore the backup. The process will vary depending on if you are restoring from an extent snapshot backup or from a full backup.

- ▶ To restore from extent snapshot backups, see “Restoring from an Extent Snapshot Backup” on page 203.
- ▶ For restore from a fullbackup, see “Restoring from a Full Backup” on page 204

Phase 2 - restore transaction logs. This phase is only possible in full transaction logging mode.

▶ “How to Restore Transaction Logs” on page 207

After the backup has been restored, the repository reflects its state at the time of the backup. All the objects are intact and ordinarily are clustered in a way similar, but not identical, to their organization in the original repository. This clustering reflects both explicit clustering of objects by the application and default clustering into the generic cluster bucket.

Restore Status

Before, during, and after restore from backup and from transaction logs, you can use the message `restoreStatus` to determine where you are in the process. This status is an attribute of the repository, not of the session, and persists across login sessions and stopping and restarting the Stone.

Not in restore mode

```
topaz 1> printit
SystemRepository restoreStatus
%
Restore is not active
```

During restore from transaction logs

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from transaction log files, restored to 03/20/12
10:15:07 PDT, nextFileId = 1, record = 409 oldest fileId = 1
```

Restore Overview

For a graphical overview of the process of restoring from backup and transaction logs, Figure 9.1 shows the process for restoring from fullbackup, while Figure 9.2 show the steps in restoring from extent snapshot backups.

Figure 9.1 System Timeline: Restoring from a Smalltalk Full Backup

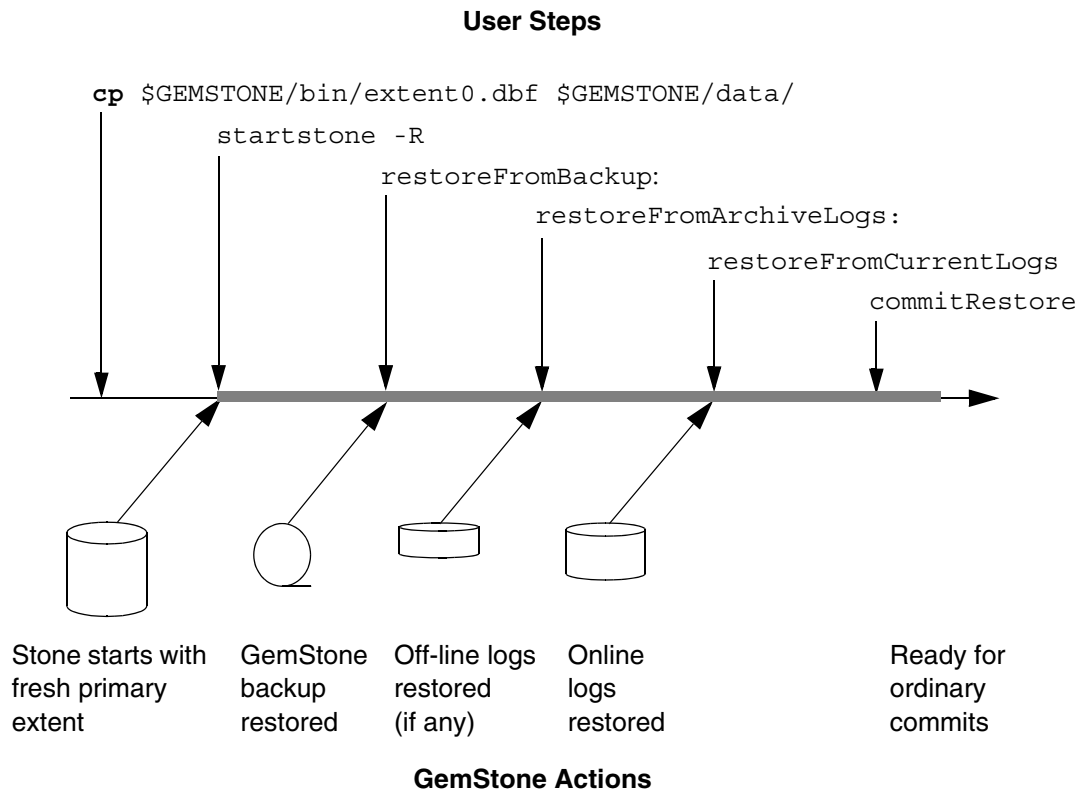
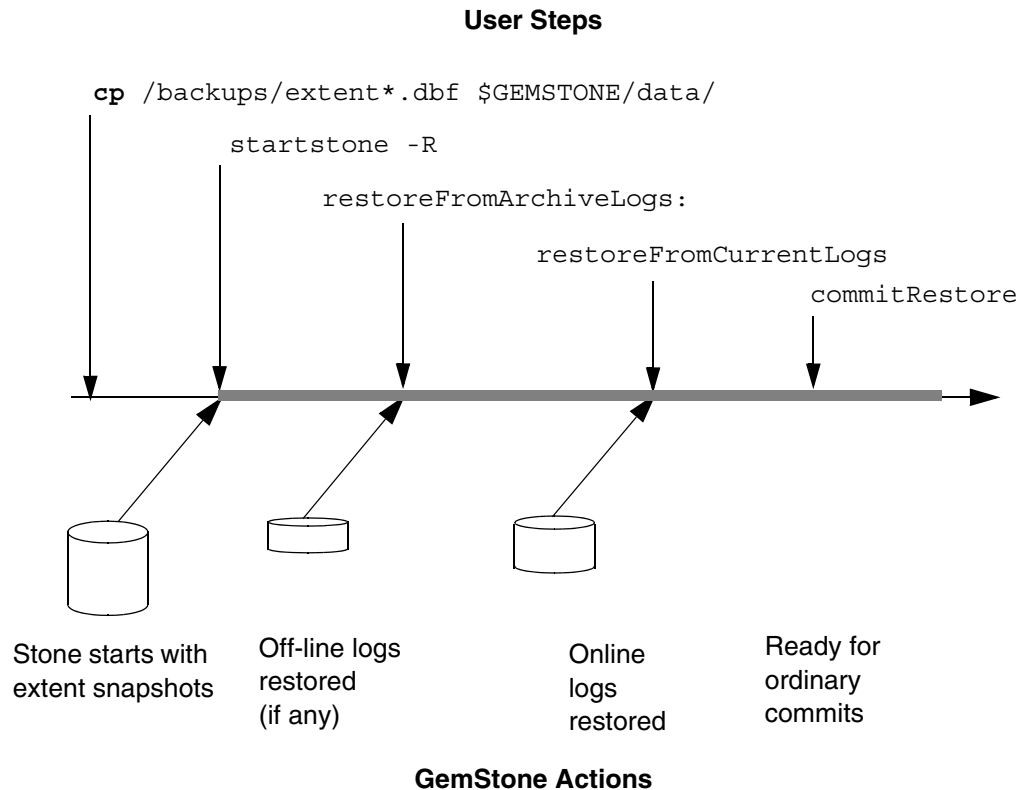


Figure 9.2 System Timeline: Restoring from a Extent Snapshot Backup



Restoring from an Extent Snapshot Backup

This section describes how to restore the repository from an operating system backup made using utilities such as `cp` to take a snapshot of the extent files. In order to recover, this backup must have been made while checkpoints were suspended for the entire time the copy was being made, or while the repository monitor was shut down.

If the backup consists of multiple extents, all extents must be available, along with all transaction logs written since the backup was started. One or more transaction logs from before the backup may also be required.

If the file system itself has been corrupted, not just the extent files, see the section “Disk Failure or File System Corruption” on page 107.

Step 1. If GemStone is still running, tell all users to log out and use `stopstone` to stop the repository monitor. Certain file system failures while the Stone is running may make it necessary to use `kill processid` to kill the Stone process.

Step 2. If you are restoring the repository because of a suspected GemStone failure, preserve a copy of the extents in case Technical Support wants to examine them.

WARNING

Do NOT delete the transaction log files – leave them online in their current locations.

Step 3. Delete all extent files specified by `DBF_EXTENT_NAMES` in your configuration file.

Step 4. Restore the operating system backup copies of the extent files to the locations specified by the `DBF_EXTENT_NAMES` configuration option.

Step 5. Ensure that there is space to create a log file. At least one of the directories specified by `STN_TRAN_LOG_DIRECTORIES` must have space available or one of the raw partitions must be empty. You may need to add entries to `STN_TRAN_LOG_DIRECTORIES` and `STN_TRAN_LOG_SIZES` in your configuration file.

Step 6. Start up the stone.

If partial transaction logging (`STN_TRAN_FULL_LOGGING = False`) was in effect at the time the backup was made, tranlogs are not restored. Restart Gemstone by invoking **startstone** in the usual manner. The restore process is now complete.

If full transaction logging (`STN_TRAN_FULL_LOGGING = True`) was in effect, start in restore mode to restore transaction logs. Use **startstone -R** to restart GemStone.

Step 7. If in full transaction logging, continue by restoring transaction logs (starting on page 207)

Restoring from a Full Backup

To begin, you need a file copy (*not* a GemStone backup) of a good repository. We recommend that you use a copy of the `extent0.dbf` that was shipped in `$GEMSTONE/bin`, although any extent file that is a complete, uncorrupted repository will work. If you are using the backup/restore process to reduce the size of your extent, the new extent file must be smaller than your current extent.

NOTE

Make sure that you have all backup files are complete. If the backup consists of multiple files, the complete set must be available.

The user restoring the backup must be the only user logged in to the server. The method that starts the restoration will suspend other logins.

NOTE

We recommend that you log in as `DataCurator` or `SystemUser` to restore the backup. If you start the restore as another user and that `UserProfile` disappears as a result of the restore, `Topaz` will see a fatal error.

To restore your repository from a Smalltalk full backup, perform the following procedure:

Step 1. If GemStone is still running, tell all users to log out and use **stopstone** to stop the system. Certain file system failures while the Stone is running may make it necessary to use **kill processid** to kill the Stone process.

Step 2. If you are restoring the repository because of a suspected GemStone failure, preserve a copy of the extents in case Technical Support wants to examine them.

Step 3. Delete all extent files specified in `DBF_EXTENT_NAMES` in your configuration file.

WARNING

Do NOT delete the transaction log files up to the time of the crash – leave them online in their current locations.

Step 4. Copy the distribution extent to the location of your primary extent, which is the extent listed first in DBF_EXTENT_NAMES.

We recommend that you use the GemStone **copydbf** command to create the copy, rather than using the UNIX **cp** command; **copydbf** must be used if you are copying to or from a raw partition.

Make sure there are no other extent files in that location. Do not copy any other extent files to the extent location. If you have more than one extent, the Stone repository monitor will create the new extents at startup.

Use **chmod** to give the copy the same permissions you ordinarily assign to your repository files.

For example:

```
% copydbf $GEMSTONE/bin/extent0.dbf \
  $GEMSTONE/data/extent0.dbf
% chmod 600 $GEMSTONE/data/extent0.dbf
```

Step 5. Ensure that there is space to create a log file. At least one of the directories specified by STN_TRAN_LOG_DIRECTORIES must have space available or one of the raw partitions must be empty. You may need to add entries to STN_TRAN_LOG_DIRECTORIES and STN_TRAN_LOG_SIZES in your configuration file.

Step 6. Use **startstone -R** to restart the Stone.

The **-R** option starts the stone in restore mode and avoids creating an orphan transaction log.

For optimal performance, your extent files should be pre-grown during startup, rather than growing incrementally during restore. See “Pregrowing Extents to a Fixed Size” on page 36.

Step 7. Log in to GemStone as DataCurator or SystemUser using linked Topaz (**topaz -l**). Remember that the password will be the original one supplied when you installed GemStone, not necessarily the one you have been using.

NOTE

To perform the following steps, you must be the only user logged in to GemStone. Once you start the next step, other logins will be suspended.

Step 8. Restore the most recent full backup to the new repository by sending the message **restoreFromBackup:** or **restoreFromBackups:**. These methods automatically detect whether a backup is compressed or not and reads it accordingly.

```
topaz 1> printit
SystemRepository restoreFromBackup: 'backup.gz'
%
```

To restore from a multi-file backup, you must specify all the files in the backup, in the order the backups were created.

```
topaz 1> printit
SystemRepository restoreFromBackups:
  #( '/backups/bkup13.3.15-1'
    '/backups/bkup13.3.15-2'
    '/backups/bkup13.3.15-3')
%
```

When restore from backup is complete, the session logs out.

```
[Info]: Logging out at 03/20/12 14:21:41 PDT
The restore from backup completed, with 97655 objects restored.
Ready for restore from transaction log(s).
If partial logging was in effect (STN_TRAN_FULL_LOGGING = false) at the time the
backup was made, the final status line reads:
```

```
Restore complete. (Backup made while in partial logging mode.)
This status means that transaction logs cannot be restored. The repository is ready for
ordinary use, and logins have been enabled.
```

If an error occurs during the restore, the system returns to the state it was in before restore. Determine the cause of the error and correct it, and return to Step 7.

Step 9. If full logging was in effect (STN_TRAN_FULL_LOGGING = true), the status line indicates the next step:

```
Ready for restore from transaction log(s).
Continue with "How to Restore Transaction Logs" on page 207.
```

Controlling Reclaim Activity When Restore Completes

During restore, pages that contain free space are by default not added to the scavengeable pages at the end of the restore. This avoids a load on the reclaim gems immediately after the commitRestore. In the normal course of operations as repository objects are operated on, pages with free space will tend to be reclaimed over time.

This can be controlled, so that pages with a specific percentage of free space are made scavengeable, so they will be reclaimed after the commitRestore. This will result in the largest amount of free space after this initial reclaim, at the expense of heavy reclaim load on the repository shortly after startup. Since reclaim requires pages, you should use some caution to avoid running out of free space before the newly reclaimed pages become available.

To explicitly specify the page free space percentage required to add pages to the scavengeable pages list, use the method `Repository >> restoreFromBackups: arrayOfFileNames scavengePagesWithPercentFree: aPercent`.

A *aPercent* value of 100 means no pages are added (the default), while 0 means pages with any free space at all are added.

For example,

```
topaz 1> printit
SystemRepository restoreFromBackups:
  #( '/backups/bkup13.3.15-1'
    '/backups/bkup13.3.15-2'
    '/backups/bkup13.3.15-3')
scavengePagesWithPercentFree: 90
%
```

9.6 How to Restore Transaction Logs

The second phase of restoring the repository is to roll forward from the state at the starting point of the last backup to the state of the last committed transaction. This action repeats the transactions in the order in which they were committed.

You can do this only if the STN_TRAN_FULL_LOGGING configuration option was set to True at the time the backup was made. You cannot restore transaction logs that are not part of a sequence of tranlogs that includes the backup. Since restore breaks this sequence, the transactions being restored cannot span a more recent restore.

Note that while backup files can be written in either uncompressed or compressed format, transaction logs are always written in uncompressed format. However, transaction logs may be compressed with **gzip** before archiving them. These compressed tranlogs can be restored directly, without having to manually run **gunzip** on them.

At this point, GemStone should be running and in restore mode, following a restore from either an extent snapshot backup or from a full backup. The following steps describe the most common case of restoring the transaction logs.

CAUTION

Ordinarily, you will restore transactions from all log files written since the backup. If for some reason you plan to omit one or more log files, refer to the section “Special Cases and Errors in Tranlog Restore” on page 209.

Step 1. Log in to GemStone as DataCurator or SystemUser using linked Topaz (**topaz -l**).

Step 2. Determine which transaction logs are needed for restore and their locations. The method `restoreStatus` identifies the earliest transaction log that is needed. In this example it is `tranlog6.dbf`:

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from Transaction Log files,
  restored to 03/02/13 13:26:31 PST
  next fileId = 6, record = 9.
```

Compare the `fileId` in the message with the names of the transaction log files in the directories specified in `STN_TRAN_LOG_DIRECTORIES`. For transaction logs in the file system, `fileId` forms the numeric portion of the file name, `tranlogNN.dbf`. For transaction logs in raw partitions, use **copydbf -i fileName** to display the `fileId`.

Transaction log files that are located in a directory specified in `STN_TRAN_LOG_DIRECTORIES` are “current”. If some required transaction logs have been moved to another location, they are “archive” logs, and are restored using a different method.

Step 3. Restore archive transaction logs, if any.

If any of the tranlogs to be restored are not in one of the current tranlog directories, collect the names of directories containing all these archive logs, and restore using `Repository>>restoreFromArchiveLogs`: or related methods.

You will have to login prior to running this step.

```
topaz 1> printit
SystemRepository restoreFromArchiveLogs:
  #( 'GS-archive' )
%
```

See the method comments in the image for details. A directory location can include an NRS for a remote node, but a NetLDI must be running on that node.

If you encounter a failure because of a truncated or corrupted transaction log, refer to “Errors While Restoring Transaction Logs” on page 212.

Step 4. Before continuing to restore tranlogs, you must log in again. Restore operations terminate the session when complete.)

Step 5. Restore transactions from the current log files by executing the method `Repository>>restoreFromCurrentLogs`. All the remaining log files must be in directories or raw partitions specified in `STN_TRAN_LOG_DIRECTORIES`.

```
topaz 1> printit
SystemRepository restoreFromCurrentLogs
%
Restore from transaction log(s) succeeded.
```

Step 6. If restoration from the transaction logs was successful, send the message `commitRestore` to tell the system that you are finished restoring. After this, no further logs can be restored, and normal user commits will be allowed.

You will have to login again prior to running this step.

```
topaz 1> printit
SystemRepository commitRestore
%
Restore from transaction log(s) succeeded. commitRestore
succeeded
```

If you send `commitRestore` prior to `restoreFromCurrentLogs`, a warning is issued because all previously committed transactions may not have been restored. However, this usage provides a way to recover as much as is available when a log file has been corrupted or lost.

Step 7. Make a new GemStone backup as soon as operational circumstances permit.

9.7 Special Cases and Errors in Tranlog Restore

If all transaction logs needed to restore up to the current time are available, transaction log restore is simple. However, in some cases transaction logs may be missing or corrupt, or you may wish to restore to an earlier point in time. This section describes these special cases and problems that you may encounter during transaction log restore.

Precautions When Restoring a Subset of Transaction Logs

When you determine the need restore an incomplete set of transaction logs, be aware of the likely consequences:

- ▶ Obviously, the omitted transactions will be lost. Presumably that is unavoidable or intentional.
- ▶ Less obviously, it may be impossible to reverse your action later and restore the omitted logs. Operations after the first `commitRestore` create a time fork in the repository, and attempting to reverse the course later results in inconsistent data and object audit errors. For a detailed example illustrating this, see the following discussion on “Fork-in-Time” Scenario.

If there is any chance that you may want to restore from the omitted transaction logs later, prior to restore archive the repository backup and all transaction logs required for complete restore to a separate location. The transaction logs should not be on any directory listed in `STN_TRAN_LOG_DIRECTORIES`.

Later, if you wish to perform a second restore, you can repeat the entire restore process, including restoring any omitted transaction logs.

Since any new work done in the partially restored system constitutes a “Fork-in-Time”, the work done after the partially restored system’s `commitRestore` cannot be restored to this second restored system. That work will be lost.

“Fork-in-Time” Scenario


In some cases, you may encounter problems with restoring from your most recent backup file and must restore from an earlier backup. This scenario presents a risk of transaction logs that are out of sequence due to a “fork-in-time.” Consider the following sequence of repository events:

1. Generate backup1.
2. Generate transaction logs `tranlog1`, `tranlog2`, `tranlog3`.
3. Generate backup2.
4. Generate transaction logs `tranlog4`, `tranlog5`, `tranlog6`.
5. Restore backup2.
6. `commitRestore` (without replaying transaction logs `tranlog4`, `tranlog5`, `tranlog6`).

The repository is now at same state as step 3.

7. Generate transaction logs `tranlog7`, `tranlog8`, `tranlog9`.
8. Restore backup1.
9. Replay transaction logs `tranlog1` through `tranlog9`.

In terms of the repository lifecycle, this scenario has two timelines, with a fork-in-time at the end of tranlog3:



If, at step 5, we also restored the transaction logs (tranlog4, tranlog5, tranlog6), the resulting sequence could be replayed without problems. The problem is caused when the continuity of the transaction log chain is broken.

After restoring backup1 in step 8, it would be possible to safely replay transaction logs tranlog1 through tranlog6 without problems, but any changes made in (tranlog7, tranlog8, tranlog9) would be lost.

During step 9, the replay of (tranlog7, tranlog8, tranlog9) is likely to produce problems. If any object changes made in (tranlog4, tranlog5, tranlog6) are logically inconsistent with those made in (tranlog7, tranlog8, tranlog9), possible errors are wide-ranging, including UTL_ASSERT/UTL_GUARANTEE errors or errors of the form:

```

recovery/restore: invalid operation XXXXXXXXXXXX
Transaction expected to abort.
non-empty invalidObjs in recover.c:commitTran
  
```

In the worst case, errors may not be written to the Stone log during transaction log replay, but the final repository may be corrupted in obscure ways. If the corruption is structural, it may be detected by an object audit (page 120). Otherwise, the corruption may go undetected unless picked up by application code.

If you are presented with a situation wherein you are forced to restore from an earlier backup, keep in mind the following:

1. Be aware of the fork-in-time phenomenon and avoid restore/replay operations that would create a fork.
2. When restoring into an ongoing transaction log sequence, only restore a backup file generated earlier within that same sequence, and then replay *all* transaction logs in that sequence generated since that backup.
3. If for some reason you cannot follow guideline 2, realize that you cannot restore from an earlier backup and replay transaction logs beyond the point of the initially restored backup.

Restoring Logs up to a Specific Log

To restore transaction logs, stopping at a specific log, execute
`Repository>>restoreToEndOfLog:fileId`. This restores all transaction logs up to and

including the specified transaction log. All tranlogs from the next tranlog required through the specified tranlog must be available. For example:

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from Transaction Log files,
  restored to 03/02/13 13:51:19 PST
  next fileId = 7, record = 0 oldest fileId = 7
topaz 1> printit
SystemRepository restoreToEndOfLog: 15
%
[Info]: Logging out at 03/24/13 14:37:07 PDT
Restore from transaction log(s) succeeded.
```

If the transaction logs to be restored are in a archive location, use the similar methods `restoreFromArchiveLogs:toEndOfLog:` or `restoreFromArchiveLogs:toEndOfLog:withPrefix:.`

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from Transaction Log files, restored to 03/02/13
13:51:19 PST, next fileId = 7, record = 0 oldest fileId = 7
topaz 1> printit
SystemRepository
  restorefromArchiveLogs: #(GS-archive)
  toEndOfLog: 15
%
[Info]: Logging out at 03/24/13 14:37:07 PDT
Restore from transaction log(s) succeeded.
```

Restoring Logs to a Point in Time

Ordinarily, the methods to restore one or more transaction logs restores each individual transaction within the log file. However, you can specify an earlier stopping point and restore only part of a transaction log, by sending one of the following messages:

```
restoreToPointInTime: aDateTime
restoreFromArchiveLogs: arrayOfDirSpec toPointInTime: aDateTime
restoreFromArchiveLogs: arrayOfDirSpec toPointInTime: aDateTime
  withPrefix: tranlogPrefix
```

Restoration will stop at the first repository checkpoint that originally occurred at or after *aDateTime*. This may be several minutes after *aDateTime*, depending on the checkpoint frequency in the transaction log.

To display the time a transaction log was started and the time of each checkpoint recorded in it, use `copydbf -I fileName`. By default, the interval between checkpoints is five minutes. For example:

```
% copydbf -I tranlog2.dbf

Source file: tranlog2.dbf
  file type: tranlog  fileId: 2
```

```

byteOrder: Sparc (MSB first) compatibilityLevel: 910
The file was created at:      03/25/13 14:55:59 PDT.
The previous file last recordId is 69.
Scanning file to find last checkpoint...
Checkpoint 1 started at: 03/25/13 14:55:59 PDT.
  oldest transaction references fileId -1 ( this file ).
Checkpoint 2 started at: 03/25/13 14:57:23 PDT.
  oldest transaction references fileId -1 ( this file ).
File size is 2.2 MBytes (4350 records).

```

The method to use to restore to a point in time depends on if the logs are archive (not in a directory on STN_TRAN_LOG_DIRECTORIES), OR ONLINE (IN A DIRECTORY USED FOR CURRENT TRANSACTION LOGS).

If the point in time that you wish to restore to occurs in an current/online transaction log, first restore any archives logs using `restoreFromArchiveLogs:`.

Then, restore all current logs up to a specified time. The following example restores the repository to the first checkpoint that would have included a commit on March 22, 2013 at 2:56:00 p.m.:

```

topaz 1> printit
SystemRepository restoreToPointInTime:
  (DateTime fromString: '22/03/2013 14:56:00').
%
```

To restore to a point in time that is in an archived tranlog, use the method `restoreFromArchiveLogs:toPointInTime:` or `restoreFromArchiveLogs:toPointInTime:withPrefix:`. This second method allows you to also specify alternate file prefixes, if you rename files as part of the archive process.

The following sequence restores the repository to the first checkpoint that would have included a commit on March 22, 2013 at 2:56:00 p.m.:

```

topaz 1> printit
SystemRepository restoreFromArchiveLogs:
  #( 'GS-archive' )
  toPointInTime:
    (DateTime fromString: '22/03/2013 14:56:00').
%
```

You can continue restoring past *aDateTime* by issuing further restore messages.

Errors While Restoring Transaction Logs

Missing Transaction Log File

If a transaction log file in the sequence is missing, the tranlog restore stops at that point, and reports an error if it detects the existence of later transaction logs.

For example, if you have tranlog1.dbf through tranlog10.dbf, but tranlog4.dbf is missing, restoreFromCurrentLogs stops after restoring from tranlog3.dbf.

```
topaz 1> printit
SystemRepository restoreFromCurrentLogs.
%
[Info]: Logging out at 03/01/12 16:27:03 PST
ERROR 4049 , Restore from transaction log failed,
      EndOfAllLogs reached after fileid 3 before last log for
recovery 10 found.
```

The tranlog after the one reported in the error is the one that is missing. You can also execute the method restoreStatus to identify the next log file explicitly. Locate the missing file or files, and then continue the restore process.

Truncated or Corrupt Transaction Log File

If a transaction log is truncated or corrupt, it may not be noticed until the next transaction log is restored. This may occur, for example, if you have an undetected disk full condition when copying a transaction log.

The truncated log may restore successfully, but when the next log is restored, the gap is detected and the error is reported.

In the following example, tranlog6.dbf is truncated, and restoreFromCurrentLogs reports an error.

```
topaz 1> printit
SystemRepository restoreFromCurrentLogs
%
[Info]: Logging out at 03/24/13 14:37:07 PDT
ERROR 4049 , Restore from transaction log failed
      Log with fileId 6 is truncated or corrupt, or log 7 is
corrupt.
```

Logging in again and checking the restore status confirms that tranlog6.dbf is incomplete:

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from Transaction Log files, restored to 03/02/13
13:26:31 PST, next fileId = 6, record = 4409 oldest fileId = 6
```

After locating a complete, uncorrupted copy of tranlog6.dbf, it is copied into the appropriate directory and the restore is done again:

```
topaz 1> printit
SystemRepository restoreFromCurrentLogs
%
[Info]: Logging out at 03/24/13 14:37:07 PDT
Restore from transaction log(s) succeeded.
```

You can verify that this and any later transaction logs were restored by logging in again and checking the restore status:

```
topaz 1> printit
SystemRepository restoreStatus
%
Restoring from Transaction Log files, restored to 03/02/13
13:51:19 PST, next fileId = 11, record = 4409 oldest fileId =
11
```

Since in this case all available transaction logs are now successfully restored, login again and commit the restored repository:

```
topaz 1> printit
SystemRepository commitRestore
%
Restore from transaction log(s) succeeded., commitRestore
succeeded
```

If you cannot find an undamaged copy of the transaction log, you cannot restore any further logs. Executing `commitRestore` will commit as much as has been restored. However, if there is any chance of a finding a good copy, see the discussion, "Precautions When Restoring a Subset of Transaction Logs".

9.8 Recovering from File System Problems

We recommend disk or operating system mirroring for applications that cannot tolerate the risk of data loss. In particular, recent transaction logs should be mirrored, or at minimum copied to an archive location on a frequent basis. In the case of a disk failure or a corrupt file system, if any of the transaction logs created since the last backup are corrupt or unusable, this recent work may be permanently lost.

In the case of disk failure or a corrupt file system, the file system must be repaired or restored. The most reliable strategy is to restore GemStone from backup, restoring copies of all transaction logs for which you have uncorrupted copies.

However, if you have important work that may be lost, you may want to attempt recovery of the existing repository. If each of these steps completes successfully, your repository is uncorrupted and you can resume normal operations.

Step 1. Perform page audit

Execute `pageaudit` per the instructions on page 120, to verify page-level integrity of the repository.

Step 2. Restart GemStone

Step 3. Perform object audit

Execute `objectAudit` per the instructions on page 122 to verify objects in the repository. This may take some time.

Some types of `objectAudit` failures indicate corruption in internal GemStone structures, which are rebuilt during restore of a full backup. If `objectAudit` reports errors, it

may be worthwhile to attempt to make a fullBackup of the repository. If this succeeds, restoring it may provide a uncorrupted repository.

9.9 Version Compatibility

It is not always possible to restore backups made by a previous version of GemStone into a new version. Since kernel classes and methods are also included in the full backup, restoring an older version will result in GemStone Smalltalk code that is not correct for the GemStone version.

If you archive backups of your GemStone repository over multiple upgrades of your GemStone installation, you should also archive the GemStone executables for each version.

While not supported, in cases where it is possible to restore the backup from an older version into a more recent version, you should follow the upgrade instructions in the *Installation Guide* to run `upgradeImage`, to ensure kernel code is updated.

9.10 Warm and Hot Standby Systems

GemStone's backup and restore mechanisms can be used to set up a secondary server, running in parallel with the primary server and ready to take over as quickly as possible in case of any failure of the primary system.

To do this, a backup of the primary server is restored into a separate location. This backup stays running in restore mode, and as transactions are generated on the primary server, they are restored into the standby system. In case of failure of the primary system, the standby can be quickly ready to use and in a state identical to the failed system.

For details on how to set up a warm or hot standby system, see Chapter 10, "Warm and Hot Standbys", on page 217.

For high-availability production systems, it is a serious problem if the repository has an unexpected error and has to be shut down, or possibly require restore from backup. While such problems are rare, critical systems must be prepared.

For such systems, a second GemStone system can be kept running in parallel, so it can be brought into use with minimal downtime. GemStone provides several options for standby systems.

- ▶ *Warm Standby* (page 218) – manually transmit entire transaction logs and restore into the standby system.
- ▶ *Hot Standby* (page 220) – set up processes to automatically transmit transaction log records as they are generated, and automatically restore the records into the standby.

This chapter discuss how to set up the standby server and the process for restoring logs, which differ between warm and hot standbys. For general information about restoring backups and transaction logs, see “How to Restore from Backup” on page 200. This discussion assumes you are familiar with that procedure.

Overview

Customers with critical, high-availability systems may want to keep a duplicate of a production GemStone server running almost in parallel as a standby system. This duplicate continually runs in restore mode, restoring transactions from the production server. If anything goes wrong with the primary production server, the standby can be brought into use very quickly.

The production system is referred to here as the primary or master system. The standby system is also referred to as the slave system.

Following a failover, these roles change; the standby system becomes the master and users log in to perform work. Often the system that was previously the master, after correcting the problem that caused failover, is updated to become the new standby.

To operate a standby, the primary system must be running in full logging mode, as described in “Choosing a Logging Mode” on page 42.

10.1 Warm Standby

With a warm standby, transaction logs from the primary system are manually copied to the standby system as each log is completed. The standby runs in restore mode, and restores each log as it is closed.

When failover is needed, the final transaction log from the primary is copied over and restored, and the standby system performs a `commitRestore` and is then available for use as the new primary server.

An important point to remember is that the transaction logs copied from the production server, called the *archive logs* here, must be kept separate from transaction logs created by the duplicate server. You can do that by using different log directories or different file name prefixes.

Setup and run the warm standby

Step 1. Install the duplicate server. It is best to do a complete GemStone installation on a second node.

Step 2. Decide on a naming convention or location that you will use on the duplicate server to keep the archive logs separate from those being created by the duplicate server itself. For instance, if both Stones use the default prefix of *tranlog*, you might copy `tranlog123.dbf` on the production server to `$GEMSTONE/data/prodtranlog123.dbf` on the duplicate server.

Step 3. Make an extent copy backup, or a full backup of the primary system. (Instructions begin on page 193.) You'll have to do this at least once, when you start this system; however, regular backups will simplify matters when you need to synchronize the primary and the standby systems.

Step 4. Copy the extent backups, or restore the full backup, into the duplicate server. If you use extent copies, use the `-R` option to start the Stone, which causes the Stone to enter restore mode. For instance, **startstone -R**.

Step 5. As each transaction log completes on the primary system, move the log to a file system accessible to the warm standby GemStone installation. Avoid moving these logs into the transaction log directory that the warm standby uses for its own transaction logs.

Wait several seconds after the new log is created before copying the old log, to ensure the completion of any asynchronous writes.

You can limit each transaction log to a tolerable amount of information either by limiting transaction log size or by starting a new log at regular time intervals:

- ▶ **Size-based:** Limit the transaction log size, as described in "Choosing the Log Location and Size Limit" on page 43. When a transaction log grows to the specified limit, GemStone starts a new transaction log.
- ▶ **Time-based:** On your primary system, run a script at regular intervals that terminates the current transaction log and starts a new one, using the method `System class >> startNewLog`.

Step 6. On the warm standby, restore the transaction logs as they are available.

Since the transaction logs for the primary are not in the STN_TRAN_LOG_DIRECTORIES of the standby, you will restore from an archive log directory for these tranlogs. You can restore from an archive log directory using one of the following methods:

```
Repository>>restoreFromArchiveLogs:
Repository>>restoreFromArchiveLogs:tranlogPrefix:
Repository>>restoreFromArchiveLogs:toEndOfLog:
Repository>>restoreFromArchiveLogs:toEndOfLog:tranlogPrefix:
The tranlogPrefix: argument allows you to use a different setting for
STN_TRAN_LOG_PREFIX on the production and standby systems.
```

Since restoring transaction logs terminates the session, you will need to login for each restore. For example:

```
topaz> login
<details omitted>
successful login
topaz 1> printit
System restoreFromArchiveLogs: {'GS-archive'}.
%
```

For more information about restoring logs, see page 207.

Step 7. Repeat Steps 5 and 6 as necessary.

You may find it necessary to shut down the standby from time to time. Ensure that you shut down the stone using stopstone. This does not affect the restore status.

Note that if the standby is not shut down cleanly (i.e. an unexpected shutdown), the restarted system has the status of its last checkpoint. You may need to restore tranlogs again that were previously restored.

Activate the warm standby in case of failure in the primary

Step 1. If the primary system fails, replay its latest transaction log on the standby system.

Step 2. On the standby server, send the message `Repository>>commitRestore` to terminate the restore process and enable logins.

Step 3. Client applications will have to reconnect to the standby system, which now becomes the primary system. Applications may have to perform their own failure recovery code as necessary, as well.

NOTE

Design your application and configuration so that, after a failure occurs and the standby is activated, client applications can reconnect to the new primary correctly.

Step 4. Correct the problem on the failed system and restart it.

Depending on how much time has elapsed since the standby system became the primary system, either make a full backup of the new primary system and restore it on the system that failed, or replay the new primary system's transaction logs on the system that failed. Maintain that system in restore mode as the new standby.

10.2 Hot Standby

A Hot Standby provides faster failover since it can always remain synchronized with the primary server. If failover is needed, all that is required is to stop tranlog transmittal and restore, perform the final commitRestore, and the repository is ready for use.

As with the warm standby, a backup of the primary (or master) repository is installed on a standby (or slave) system. Then, special processes run to transmit the transaction log records as they are generated on the primary system. Individual transaction records, rather than entire transaction logs, are transmitted in compressed form between the systems. The standby system runs in a special mode where it will continuously restore the transaction log records. As a result, the standby system can keep closely synchronized with the primary.

Precautions regarding tranlog sequences

Since the hotstandby process relies on a logical sequence of tranlogs, some care must be taken to avoid situations such as described under "Fork in Time" on page 209.

For example, if you restore from backup on the primary, you must make a new backup of the primary, and restore this into the hot standby. Likewise, you cannot restore a backup into both the primary and the standby, since this creates a fork in the logical sequence of tranlogs, regardless of the tranlog numbering. While in some cases the automated process that transmit and restore the transaction log records will detect this, there may be cases where the processes will appear to hang waiting for the correct logs.

Hot standby processes

logsender

The **logsender** process runs on the master system. It determines when new transaction records are available on the primary system, and sends these records in compressed form to the standby slave system.

The logsender process is started using the startlogsender utility command. When starting the logsender, you must specify:

- ▶ The address and port to listen on for connections from a logreceiver on a slave system.
- ▶ all directories or raw partitions containing transaction logs generated by the master system.
- ▶ the name of the master stone. While technically optional, without this the logsender will not be notified of new data to send.

There are other optional arguments. For more specific details on the startlogsender utility, see Appendix B, page 322.

A logsender continues to run when the associated master stone is shut down, and reconnects to the stone when the stone is restarted. Since it is continuously listening for connection requests from a logreceiver, the connection with the logreceiver can also be automatically reestablished after a disconnect, provided the logreceiver is running.

The logsender must be stopped explicitly using the stoplogsender utility command. For details on the stoplogsender command, see Appendix B, page 330.

logreceiver

The **logreceiver** process runs on the standby slave system. It receives transaction logs from the master system and writes them to a location where the slave stone can restore them.

The logreceiver process is started using the startlogreceiver utility command. When starting the logreceiver, you must specify:

- ▶ The address and port that the logsender on the master system is listening on.
- ▶ One or more directories to write incoming transaction logs from the master system. It is not recommended to use raw partitions.
- ▶ the name of the slave stone. While technically optional, without this the logreceiver cannot notify the slave stone that new data has arrived.

There are other optional arguments. For more specific details on startlogreceiver and the command line options, see Appendix B, page 320.

A logreceiver continues to run when the associated slave stone is shut down, and reconnects to the stone when the stone is restarted. If the connection to the logsender is lost, the logreceiver will attempt to reconnect.

The logreceiver must be stopped explicitly using the stoplogreceiver utility command. For details on the stoplogreceiver command, see Appendix B, page 329.

Continuous Restore Mode

A slave system in a hot standby runs in continuous restore mode. In this mode, it can restore individual transaction records as they become available.

To enter continuous restore mode, execute:

```
SystemRepository continuousRestoreFromArchiveLogs:  
anArrayOfRestoreDirectories
```

To exit continuous restore mode, execute:

```
SystemRepository stopContinuousRestore
```

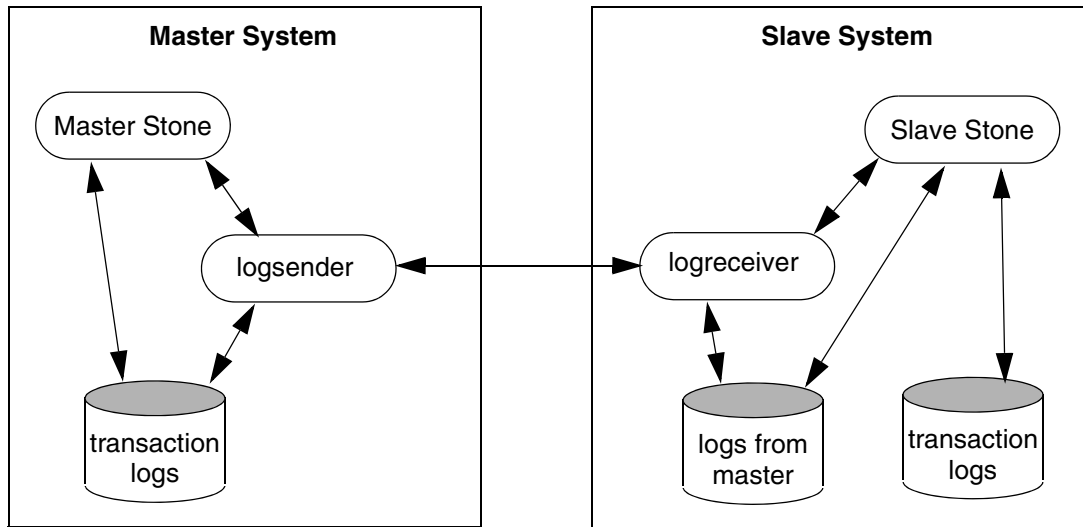
You must be in restore mode before you can enter continuous restore mode, and executing stopContinuousRestore leaves the stone still in restore mode.

You may not execute the commitRestore, which exits from restore mode, without first exiting continuous restore mode.

Transaction Record Transmittal

The process that transaction records follow from the master to the slave is described in this section.

Figure 10.1 Processes in a hot standby



As transactions are committed on the master stone, transaction log records are written to the transaction logs – this is standard behavior. The logsender process is logged into the master stone and is aware when new transaction records are generated. The logsender also has an established connection with the logreceiver process. The logsender transmits the transaction log records in compressed form to the logreceiver.

The logreceiver accepts the transaction log records and writes them to the slave system's directories for restore logs. The logreceiver is logged into the slave stone, and makes the slave stone aware that new transaction records are available for restore.

The slave stone is running in continuous restore mode, and restores the transaction records.

When a logsender or logreceiver is logged into a stone, or while a stone is in continuous restore mode, the stone cannot restore a full backup or restore tranlogs (using other restore methods), nor can it perform a commitRestore.

Multiple standby repositories

A given master stone may have multiple slave stones. Each slave system will have a separate logreceiver process running. The logsender on the master system can transmit data to up to 5 logreceiver processes.

To setup and run the hot standby

- Step 1.** Install the slave server. It is best to do a complete GemStone installation on a second node.
- Step 2.** Decide on a directory location that you will use on the slave server to keep for the logs transmitted from the master. This should be a separate directory or directories from the tranlog directories of the slave stone.

Step 3. Make an extent copy backup, or a programmatic full backup of the primary system. This must be a backup from the primary system; you cannot restore a backup into both the primary and standby, since restore creates a fork-in-time (see page 209).

You'll have to do this at least once, when you start this system, and after each restore from backup or upgrade on the primary system. Regular backups will simplify matters when you need to synchronize the primary and the standby systems.

Step 4. Copy the extent backups to the slave system and start the stone using the `startstone -R -N` option; or start the stone on a clean extent and restore the full backup. The slave stone must be running in restore mode.

Step 5. Start the logsender process on the master system using the `startlogsender` utility.

Before starting the logsender, you will need to determine the set of directories that contain transaction logs on the master system. This will include all the entries in the master stone's `STN_TRAN_LOG_DIRECTORIES`. If the master stone's transaction logs are copied to another directory as part of an archive process, these archive directories may also need to be specified.

You will also need to select an port number that is unused on the master and slave systems for the logsender to listen on.

For example:

```
startlogsender -P 57222 -A masterListeningAddress -T $GEMSTONE/data
-s masterStone.
```

Step 6. Start the logreceiver process on the slave system using the `startlogreceiver` utility.

You will use the same port as the logsender, and the directory or directories you determined in Step 2.

```
startlogreceiver -P 57222 -A masterNode
-T /gemstone/masterTransLogs -s slaveStone.
```

Step 7. Put the slave stone into continuous restore mode. To do this, log into the slave system and execute `Repository >> continuousRestoreFromArchiveLogs:`, passing in the list of directories you determined in Step 2. After entering continuous restore mode, this method will exit and you can log out.

Activate the hot standby in case of failure in the primary

In case of a failure in the master system, perform the following steps:

Step 1. Confirm that all tranlog records have been transmitted to the slave system, and that the slave system has restored all transaction log records.

To do this, provided the master system is operational, check the results of `restoreStatus` on the slave system, and compare this to the results of `copydbf -i` on the final tranlog on the master system.

Step 2. If the master system is operational, stop the logsender process on the master system using the `stoplogsender` utility command.

```
stoplogsender -P 57222
```

Step 3. On the slave server, stop the logreceiver process using the stoplogreceiver utility command.

```
stoplogreceiver -P 57222
```

Step 4. On the slave system, send the message `Repository>>stopContinuousRestore` to exit continuous restore mode, then `Repository>>commitRestore` to terminate the restore process and enable logins. The slave system is now ready for use.

Step 5. Client applications will have to reconnect to the slave system, which now becomes the primary system. Applications may have to perform their own failure recovery code as necessary, as well.

NOTE

Design your client applications so that, after detecting a failure, they can determine which system is the new primary and reconnect correctly.

Step 6. Correct the problem on the failed former-master system and restart this as the new slave system.

You will need to make a fresh backup of the new master system and restore this into the new slave (the former master), prior to setting up the hot standby again.

Planned failovers

When you have a planned failover, there are additional steps you can take to ensure that all records have been transmitted, and stop new work from taking place on the master during the failover itself.

Step 1. On the master system, execute:

```
SystemRepository suspendCommitsForFailover
```

which will suspend commits and perform a checkpoint. This is the failover timestamp. At this point, no changes can be made in the master that will affect the transaction logs: commits are suspended, including garbage collection and other activities. This state persists across repository shutdown, until you execute `SystemRepository resumeCommits`.

Step 2. On the slave system, when a transaction record of a checkpoint with a failover timestamp is received, the slave system will stop continuous restore. This ensures the slave has all records from the master, and no further work may be done on the master that could be lost.

At this point, you may perform the commit restore on the slave system and begin using that as your master system, and shut down the previous-master.

However, if the standby Stone was started up later than the transaction record containing a failover timestamp, it is not considered to be a failover scenario and the standby does not stop continuous restore. If you intend to do a planned failover, ensure that the slave stone is running before suspending commits for failover on the master system.

Connecting using SSL Mode

Generally, both master and slave nodes of a hot standby would be within your secure network, and benefit from the ease and performance of regular socket connections.

However, you may also configure the logsender-logreceiver connection to use SSL, if the network between your master and slave system is not secure.

When you run the hotstandby in SSL mode, you will need SSL credentials for both the logsender and logreceiver in order for them to connect, and you will also need these to stop a logsender or logreceiver running in SSL mode.

SSL standard TLS v1.2 is used.

The SSL-specific arguments are optional, and include:

- C *fileName* certificate in PEM format that will be sent to the peer upon request.
- J *fileName* certificate authority (CA) file in PEM format to use for peer certificate verification.
- K *fileName* private key in PEM format for the certificate (-C option).
- Q *string* private key passphrase. Required if the -K option is used and the private key is encrypted.
- S enable SSL mode. Must be specified to use any other SSL options and must also be specified when starting the peer process.
- V Disable verification of the peer's certificate.

For example, to setup a logsender/logreceiver with certificate verification fully enabled, using the example certificates provided with the GemStone/S 64 Bit distribution:

```
startlogsender -A masterListeningAddress -P 57222 -T $GEMSTONE/data
-s masterStone -S -C $GEMSTONE/examples/openssl/certs/server.pem
-K $GEMSTONE/examples/openssl/certs/server.pem
-J $GEMSTONE/examples/openssl/certs/serverCA.pem

startlogreceiver -A masterNode -P 57222 -s slaveStone
-T /gemstone/masterTransLogs -S
-C $GEMSTONE/examples/openssl/certs/server.pem
-K $GEMSTONE/examples/openssl/certs/server.pem
-J $GEMSTONE/examples/openssl/certs/serverCA.pem

stoplogsender -P 57222 -S
-C $GEMSTONE/examples/openssl/certs/server.pem
-K $GEMSTONE/examples/openssl/certs/server.pem
-J $GEMSTONE/examples/openssl/certs/serverCA.pem

stoplogreceiver -P 57222 -S
-C $GEMSTONE/examples/openssl/certs/server.pem
-K $GEMSTONE/examples/openssl/certs/server.pem
-J $GEMSTONE/examples/openssl/certs/serverCA.pem
```

Self signed certificates

Note that by default, self-signed certificates will be rejected, because the certificate cannot be verified by a known certificate authority. To use self-signed certificates, the signer must be added to the CA list in the CA file (-J flag), or certificate verification must be disabled with the -V flag. Using -V effectively tells OpenSSL to ignore certificate errors. In this mode, communications between the logsender and receiver are encrypted, but the identities of the logsender and/or logreceiver have not been verified.

Managing Memory

Executing your application code will naturally require access to previously committed objects in the repository. These objects are faulted into the Gem session's memory to be examined or updated. When new objects are created, they must reside in memory while they are being modified.

GemStone automatically garbage-collects temporary objects that are no longer referenced, and clears out the space used by persistent, committed objects, when memory is needed. However, the memory available for any session is finite. If you need to create large temporaries or modify many objects within a transaction, you may need to tune your application to increase the available memory, or to use memory more efficiently.

This chapter discusses the following topics:

- ▶ **Memory Organization** – How a GemStone session's memory is organized.
- ▶ **Configuring Temporary Memory Usage** – How to configure temporary object memory, and debug out-of-memory problems.

11.1 Memory Organization

Each Gem session has a temporary object memory that is private to the Gem process and its corresponding session. This local object memory is divided into the following regions:

- ▶ `new` – Young temporary objects; includes two subspaces named `eden` and `survivor`
- ▶ `old` – Older temporary objects
- ▶ `pom` – Unmodified faulted-in objects, divided into ten subspaces
- ▶ `perm` – Faulted-in or created Classes and Metaclasses
- ▶ `code` – Instances of `GsNMethod` being executed, or recently executed
- ▶ `mE` – `oopMap` entries that map `objId` to in-memory objects for committed objects

Temporary objects are created in the new area of local object memory. When the new area fills up, a scavenge occurs which throws away unreferenced objects in new. After an object has survived a number of scavenges, it is copied to the old area. After the old area has grown by some amount or is almost full, a mark/sweep takes place, finding all live objects and then compacting the new, old, perm, and code areas as needed to remove dead objects.

Committed objects referenced by the session are copied from the shared page cache into the pom, perm, or code areas at the point they are first referenced by interpreter execution or a GCI call. (This is called a "copy-on-read" design.)

If a committed object in the pom area has been modified, it is copied to the old area if a scavenge occurs before the change is committed. Objects that are sent to the GCI client, such as GBS, are also moved to the old area, whether or not they are modified.

When the pom area becomes full, the contents of its oldest subspace (that is, the oldest 10%) are discarded, and that subspace is reused to continue faulting-in committed objects. Before the oldest subspace is recycled, any objects in the subspace that have been modified, or that are currently referenced from the interpreter stack, are copied to the old area.

At transaction commit, any committed objects that have been modified, and any new objects transitively reachable from those modified objects, are copied to new data pages in the shared cache. A transaction conflict check is then performed. If the commit succeeds, the in-memory state of all new objects copied to the shared cache is changed to "committed". The newly committed objects are now eligible to be removed from temporary memory by a mark/sweep or scavenge if they are no longer directly referenced from temporary objects.

11.2 Configuring Temporary Memory Usage

You may encounter an OutOfMemory error if you create too large a graph of live temporary objects at any time, or if you try to modify too many committed objects in a single transaction. OutOfMemory is a fatal error that terminates the session.

Very large numbers of Classes can also fill up temporary object memory. Any class that is referenced by message send or iteration is loaded into the perm area, and its method dictionaries, classHistory, and so on are loaded into the old area.

Persistent objects that are in the export set for a GCI client, such as GemBuilder for Smalltalk, are also moved to the old area. This includes objects that are replicated but not modified.

If you find that your application is running out of temporary memory, you can use several GemStone environment variables to help you identify which parts of your application are triggering garbage collection. Once you've done that, you can set GemStone configuration options to provide the needed memory.

Configuration Options

The values for these options are set when the gem or topaz -l process is initialized. You cannot change these values without restarting the VM. For more about these options, see the descriptions that begin on page 282.

GEM_TEMPOBJ_CACHE_SIZE

The maximum size (in KB) of temporary object memory. (This limit also applies to linked Topaz sessions and linked GemBuilder applications.) When you only change this setting, and the other GEM_TEMPOBJ... configuration options use default values, then all of the various spaces remain in proportion to each other.

GEM_TEMPOBJ_MESPACE_SIZE

The maximum size (in KB) of the mE (Map Entries) space within temporary object memory. Unless you are trying to minimize the memory footprint on HP-UX or AIX, you should always leave GEM_TEMPOBJ_MESPACE_SIZE at its default value so that the system can calculate the appropriate value. Otherwise, you are at risk of premature OutOfMemory errors.

GEM_TEMPOBJ_OOPMAP_SIZE

The size of the hash table (that is, the number of 8-byte entries) in the objId-to-object map within temporary object memory. This option should normally be left at its default setting.

GEM_TEMPOBJ_POMGEN_SIZE

The maximum size (in KB) of the POM generation area in temporary object memory. The POM generation area holds unmodified copies of committed objects that have been faulted into a Gem, and is divided into ten subspaces. This option should normally be left at its default setting so that the system can calculate the value.

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 approximate maximum number of bytes of temporary object memory that are usable for storing objects.

System _tempObjSpacePercentUsed

Returns the approximate percentage of temporary object memory that is in use 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.

Sample Configurations

This section presents several sample configurations:

- ▶ Default configuration
- ▶ Larger old area, smaller pom
- ▶ Smaller old area, larger pom

These examples assume that you have already set the `GS_DEBUG_VMGC...` environment variables (page 228) to produce the resulting printouts. The examples shown are for Solaris and may vary on other platforms.

Default Configuration

A default value (10000) for `GEM_TEMPOBJ_CACHE_SIZE` (that is, 10 MB) produces a limit of about 7 MB of temporary plus modified-committed objects (`old`), space for a working set of about 8 MB of unmodified committed objects (`pom`), and a maximum memory footprint on the order of 25 MB.

The following example shows the printout for this configuration:

```
vmGc spaceSizes: eden init 1864K max 1864K, survivor init 320K
max 320K,
vmGc old max 7496K, code max 2000K, perm max 1000K, pom 10 *
840K=8400K,
vmGc remSet 216K, meSpace max 9592K oopMapSize 65536
```

(The internal structures `remSet`, `meSpace`, and `oopMapSize` are not of interest here.)

Larger old, Smaller pom

The following settings configure the application for a 5 MB working set of unmodified committed objects (smaller than the default), and a maximum of 25 MB of temporary plus modified objects (larger than the default).

```
GEM_TEMPOBJ_CACHE_SIZE = 25000;
GEM_TEMPOBJ_POMGEN_SIZE = 5000;
```

The following example shows the printout for this configuration:

```
vmGc spaceSizes: eden init 2000K max 4688K, survivor init 400K
max 784K
vmGc old max 18744K, code max 5000K, perm max 2504K, pom 10 *
504K=5040K
vmGc remSet 360K, meSpace max 16824K oopMapSize 65536
```

Smaller old Area, Larger pom

The following settings configure an application with a large working set of committed objects and small temporary object space.

```
GEM_TEMPOBJ_CACHE_SIZE = 7000;
GEM_TEMPOBJ_POMGEN_SIZE = 100000;
```

The following example shows the printout for this configuration:

```
vmGc spaceSizes: eden init 1304K max 1304K , survivor init 224K
max 224K
vmGc old max 5248K, code max 1400K, perm max 704K, pom 10 *
10000K=100000K
vmGc remSet 1008K, meSpace max 48880K oopMapSize 524288
```

Debugging out-of-memory errors

When any of the following environment variables are set to a positive non-zero value, they have the effect described here for each Gem or linkable Topaz (`topaz -l`) process that

you subsequently start. For all of these environment variables, the printout goes to the "output push" file of a linkable Topaz (topaz -l) session, for use in testing your application. If that file is not defined, the printouts go to standard output of the session's Gem or topaz -l process.

To set any of these environment variables, you must first uncomment them in the \$GEMSTONE/sys/gemnetdebug file. Be aware that the contents of gemnetdebug are subject to change at any time. For the most current information about these and other variables, examine the gemnetdebug file.

GS_DEBUG_COMPILE_TRACE

Trace method compiles. The following are valid values:

- 0 - no tracing
- 1 - one line (class,selector) of each method compiled
- 2 - in addition to above, bytecode disassembly
- 3 - in addition to above, native code assembly listing

GS_DEBUG_VMGC_MKSW_MEMORY_USED_SOFT_BREAK

At the end of each mark/sweep, if the percent of memory used is greater than the threshold specified by this variable, a SoftBreak (error 6003) is generated, and the threshold is raised by 5 percent. We suggest a setting of 75%.

GS_DEBUG_VMGC_MKSW_PRINT_STACK

The mark/sweep count at which to begin printing the Smalltalk stack at each mark/sweep.

GS_DEBUG_VMGC_MKSW_PRINT_C_STACK

The mark/sweep count at which to begin printing the C stack at each mark/sweep. This variable is very expensive, consuming two seconds plus the cost of fork() for each printout.

GS_DEBUG_VMGC_PRINT_MKSW

The mark/sweep count at which to begin printing mark/sweeps.

GS_DEBUG_VMGC_PRINT_MKSW_MEMORY

The mark/sweep count at which to begin printing detailed memory usage (20 lines) for each mark/sweep.

GS_DEBUG_VMGC_PRINT_MKSW_MEMORY_USED

Specifies when Smalltalk stack printing starts as the application approaches OutOfMemory conditions. At the end of each mark/sweep, if the percent of memory used is greater than the threshold specified by this variable, the mark/sweep is printed, the Smalltalk stack is printed, and the threshold is raised by 5 percent. In a situation producing an OutOfMemory error, you should get several Smalltalk stacks printed in the Gem log file before the session dies.

GS_DEBUG_VMGC_PRINT_SCAV

The scavenge count at which to begin printing scavenges. Once this takes effect, all mark/sweeps will also be printed. Be aware that printing scavenges can produce large quantities of output.

GS_DEBUG_VM_PRINT_TRANS

Print transaction boundaries (begin/commit/abort) in the log file.

GS_DEBUG_VMGC_SCAV_PRINT_STACK

The scavenge count at which to begin printing the Smalltalk stack at each scavenge. Be aware that this print activity can produce large quantities of output.

GS_DEBUG_VMGC_SCAV_PRINT_C_STACK

The scavenge count at which to begin printing the C stack at each scavenge. This variable is very expensive, consuming 2 seconds plus the cost of fork() for each printout.

GS_DEBUG_VMGC_VERBOSE_OUTOFMEM

Automatically call the primitive for `System class>>_vmPrintInstanceCounts:0` when an OutOfMemory error occurs, and also print the Smalltalk stack. (For details about this method, see the comments in the image.) This applies to each Gem or linkable Topaz (topaz -l) process that you subsequently start.

GS_DEBUG_VMGC_VERIFY_MKSW

The mark/sweep count at which to begin verifying object memory before and after each mark/sweep.

GS_DEBUG_VMGC_VERIFY_SCAV

The scavenge count at which to begin verifying object memory before and after each scavenge. Once this takes effect, `GS_DEBUG_VMGC_VERIFY_MKSW` will also be in effect. Be aware that this activity uses significant amounts of CPU time.

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. By enabling handling for the signal `AlmostOutOfMemory`, an application can take appropriate action before memory is entirely full. This signal 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.

When performing index operations, such as creating indexes for large collections, the `IndexManager` can be configured to use this facility to automatically commit when memory is low. Committing objects allows them to be removed from memory, since they can be re-loaded as needed from the persistent object. See the *Programming Guide for GemStone/S 64 Bit* for details on `IndexManager autoCommit`.

For more information on handling `AlmostOutOfMemory`, see the *Programming Guide for GemStone/S 64 Bit*.

In the course of everyday operations, your GemStone/S 64 Bit repository will grow. Some of this growth will be the result of new data in your repository, but some will represent unreferenced or outdated objects. These objects, no longer needed, must be removed to prevent the repository from growing arbitrarily large. The process of removing unwanted objects to reclaim their storage is referred to as *garbage collection*.

This chapter describes GemStone's garbage collection mechanisms and explains how and when to use them.

This chapter discusses the following topics:

- ▶ **Basic Concepts** (page 233) – The main concepts underlying garbage collection.
- ▶ **Garbage Collection Operations** – `markForCollection` (page 242), FDC/MGC (page 244), epoch garbage collection (page 247), and `reclaim` (page 253).
- ▶ **GcGems** (page 256) – How to configure, start, and stop the Admin Gems and the Reclaim Gem.
- ▶ **Further Tuning Garbage Collection** (page 259) – Tuning multi-threaded scan operations, and other special issues affecting Garbage Collection.

12.1 Basic Concepts

Smalltalk execution can produce a number of objects needed only for the moment. In addition, normal business operations can cause previously committed objects to become obsolete. To make the best use of system resources, it is desirable to reclaim the resources these objects use as soon as possible.

Different Types of Garbage

Garbage collection mechanisms vary according to *where* garbage collection occurs – temporary (scratch) memory or permanent object space – and *how* it occurs – automatically, or in response to an administrator's action.

Each Gem session has its own private memory intended for scratch space, known as *local object memory*. The Gem session uses local object memory for a variety of temporary objects, which can be garbage-collected individually. This type of garbage collection is handled automatically by the session and is (for the most part) not configurable, although memory can be configured for specific gem requirements. These issues are covered in Chapter 11, “Managing Memory” on page 227.

Permanent objects are organized in units of 16 KB called *pages*. Pages exist in the Gem’s private page cache, the Stone repository monitor’s private page cache, the shared page cache, and on disk in the extents. When first created, each page is associated with a specific transaction; after its transaction has completed, GemStone does not write to that page again until all its storage can be reclaimed.

Objects on pages are not garbage-collected individually. Instead, the presence of a shadow object or dead object triggers reclaim of the page on which the object resides. Live objects on this page are copied to another page.

The Process of Garbage Collection

Removing unwanted objects is a two-phase process:

1. Identify – *mark* – superfluous objects.
2. *Reclaim* the resources they consume.

Together, *marking* and *reclaiming* unwanted objects is *collecting garbage*.

Complications ensue because each Gem in a transaction is guaranteed a consistent view of the repository: all visible objects are guaranteed to remain in the same state as when the transaction began. If another Gem commits a change to a mutually visible object, both states of the object must somehow coexist until the older transaction commits or aborts, refreshing its view. Therefore, resources can be reclaimed only after all transactions concurrent with marking have committed or aborted.

Older views of committed, modified objects are called *shadow objects*.

Garbage collection reclaims three kinds of resources:

- ▶ The storage occupied by dead objects
- ▶ The storage occupied by shadow objects
- ▶ Object identifiers (OOps) for dead objects

Live objects

GemStone considers an object *live* if it can be reached by traversing a path from AllUsers, the root object of the GemStone repository. By definition, AllUsers contains a reference to each user’s UserProfile. Each UserProfile contains a reference to the symbol list for a given user, and those symbol dictionaries in these lists in turn point to classes and instances created by that user’s applications. Thus, AllUsers is the root node of a tree whose branches and leaves encompass all the objects that the repository requires at a given time to function as expected.

Transitive closure

Traversing such a path from a root object to all its branches and leaves is called *transitive closure*.

Dead objects

An object is *dead* if it cannot be reached from the AllUsers root object. Other dead objects may refer to it, but no live object does. Without living references, the object is visible only to the system, and is a candidate for reclaim of both its storage and its OOP.

Shadow objects

A *shadow object* is a committed object with an outdated value. A committed object becomes shadowed when it is modified during a transaction. Unlike a dead object, a shadow object is still referenced in the repository because the old and new values share a single object identifier. The shadow object must be maintained as long as it is visible to other transactions on the system; then the system can reclaim only its storage, not its OOP (which is still in use identifying the committed object with its current value).

Commit records

Views of the repository are based on *commit records*, structures written when a transaction is committed. Commit records detail every object modified (*the write set*), as well as the new values of modified objects. The Stone maintains these commit records; when a Gem begins a transaction or refreshes its view of the repository, its view is based on the most recent commit record available.

Each session's view is based on exactly one commit record at a time, but any number of sessions' views can be based on the same commit record.

NOTE

The repository must retain each commit record and the shadow objects to which it refers as long as that commit record defines the transaction view of any session.

Commit record backlog

The list of commit records that the Stone maintains in order to support multiple repository views is the *commit record backlog*.

Shadow or Dead?

The following example illustrates the difference between dead and shadow objects. In Figure 12.1, a user creates a SymbolAssociation in the SymbolDictionary Published. The SymbolAssociation is an object (oop 27111425) that refers to two other objects, its instance variables key (#City, oop 20945153), and value ('Beaverton', oop 27110657).

The Topaz command “display oops” causes Topaz to display within brackets ([]) the identifier, size, and class of each object. This display is helpful in examining the initial SymbolAssociation and the changes that occur.

Figure 12.1 An Association Is Created and Committed

```

topaz 1> display oops
topaz 1> printit
Published at: #City put: 'Beaverton'.
Published associationAt: #City
%
[27111425 sz:2 cls: 111617 SymbolAssociation] a SymbolAssociation
key [20945153 sz:4 cls: 110849 Symbol] City
value [27110657 sz:9 cls: 74753 String] Beaverton
topaz 1> commit
Successful commit

```




Figure 12.2 shows a second Topaz session that logs in at this point. Notice that the Topaz prompt identifies the session by displaying a digit. Because Session 1 committed the SymbolAssociation to the repository, Session 2 can see the SymbolAssociation.

Figure 12.2 A Second Session Can See the Association

```

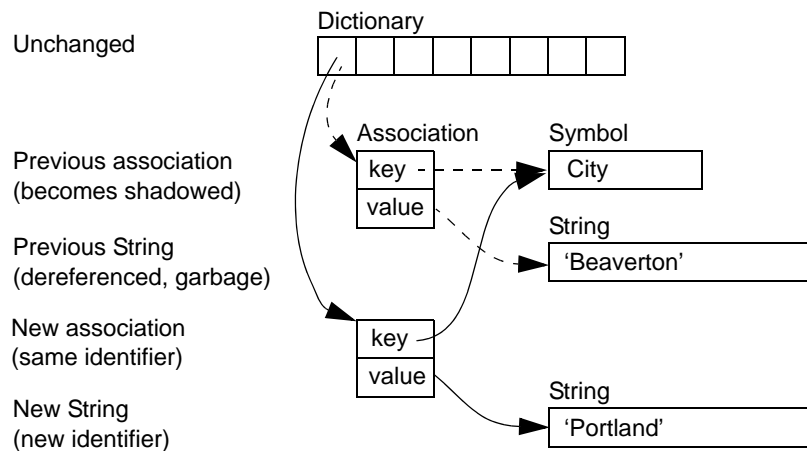
topaz 2> display oops
topaz 2> printit
Published associationAt: #City .
%
[27111425 sz:2 cls: 111617 SymbolAssociation] a SymbolAssociation
key [20945153 sz:4 cls: 110849 Symbol] City
value [27110657 sz:9 cls: 74753 String] Beaverton

```

Now Session 1 changes the *value* instance variable, creating a new SymbolAssociation (Figure 12.3). Notice in the oops display that the new SymbolAssociation object has the same identifier (27111425) as the previous Association.

Figure 12.3 The Value Is Replaced, Changing the Association

```
topaz 1> printit
City := 'Portland'.
Published associationAt: #City.
%
[27111425 sz:2 cls: 111617 SymbolAssociation] a SymbolAssociation
key [20945153 sz:4 cls: 110849 Symbol] City
value [27109121 sz:8 cls: 74753 String] Portland
topaz 1> commit
Successful commit
```



- ▶ The SymbolAssociation is now *shadowed*. Because the shadow SymbolAssociation was part of the committed repository and is still visible to other transactions (such as that of Session 2), it cannot be overwritten. Instead, the new SymbolAssociation is written to another page, one allocated for the current transaction.
- ▶ The previous value (oop 27110657) is no longer referenced in the repository. For now, this object is considered *possibly dead*; we cannot be sure it is dead because, although the object has been dereferenced by a committed transaction, other, concurrent transactions might have created a reference to it.

Even though Session 1 committed the change, Session 2 continues to see the original SymbolAssociation and its value (Figure 12.4). Session 2 (and any other concurrent sessions) will not see the new SymbolAssociation and value until it either commits or aborts the transaction that was ongoing when Session 1 committed the change.

Figure 12.4 Session 2 Sees Change After Renewing Transaction View of Repository

```

topaz 2> printit
Published associationAt: #City.
%
[27111425 sz:2 cls: 111617 SymbolAssociation] a SymbolAssociation
  key [20945153 sz:4 cls: 110849 Symbol] City
  value [27110657 sz:9 cls: 74753 String] Beaverton

topaz 2> abort
topaz 2> printit
Published associationAt: #City.
%
[27111425 sz:2 cls: 111617 SymbolAssociation] a SymbolAssociation
  key [20945153 sz:4 cls: 110849 Symbol] City
  value [27109121 sz:8 cls: 74753 String] Portland

```

Only when all sessions with concurrent transactions have committed or aborted can the shadow object be garbage collected.

What Happens to Garbage?

This section describes the steps involved in garbage collection. Specific garbage collection mechanisms will follow these steps, although the details will vary when using different garbage collection mechanisms.

The basic garbage collection process encompasses nine steps:

1. Find all the live objects in the system by traversing references, starting at the system root AllUsers. This step is called *mark/sweep*.
2. The Gem that performed mark/sweep now has a list of all live objects. It also knows the universe of all possible objects: objects whose OOPs range from zero to the highest OOP in the system. It can now compute the *set of possible dead objects* as follows:
 - a. Subtract the live objects from the universe of possible objects.
 - b. Subtract all the unassigned (free) OOPs in that range.

This step is called the *object table sweep* because the Gem uses the object table to determine the universe of possible objects and the unassigned OOPs.

3. The Gem performing this work now has a list of *possibly dead* objects. We can't be sure they're dead because, during the time that the mark/sweep and object table sweep were occurring, other concurrent transactions might have created references to some of them.

The Gem sends the Stone the *possible dead set* and returns.

4. Now, in a step called *voting*, each Gem logged into the system must search its private memory to see if it has created any references to objects in the possible dead set. When it next commits or aborts, it votes on every object in the possible dead set. Objects referenced by a Gem are removed from the possible dead set.

NOTE

Gems do not vote until they complete their current transaction. If a Gem is sleeping or otherwise engaged in a long transaction, the vote cannot be finalized

and garbage collection pauses at this point. Commit records accumulate, garbage accumulates, and a variety of problems can ensue.

5. Because all the previous steps take time, it's possible that some Gems were on the system when the mark/sweep began, created a reference to an object now in the possible dead set, and then logged out. They cannot vote on the possible dead set, but objects they've modified are in the write sets of their commit records. The *Admin Gem*, a process dedicated to administrative garbage collection tasks, scans all these write sets (the *write set union*), and votes on their behalf. This is called the *write set union sweep*.
6. After all voting is complete, the resulting set now holds nothing but unreferenced objects. The Stone now promotes the objects from possibly dead to dead.
7. the *Reclaim Gem* reclaims pages: it copies live objects on the page onto a new page, thereby compacting live objects in page space. The page now contains only recyclable objects and perhaps free space.
8. The Reclaim Gem commits. The reclaimed OOPs are returned to their free pool.
9. The Reclaim Gem's commit record is disposed of. The reclaimed pages are returned to their free pool.

Admin and Reclaim Gems

It is useful to understand the distinction between the Admin Gem and the Reclaim Gem:

- ▶ The Admin Gem finalizes the vote on possibly dead objects (page 239, Step 5), and performs the write set union sweep. The Admin Gem also performs epoch garbage collection (page 242), if enabled.
- ▶ The Reclaim Gem is dedicated to the task of reclaiming shadowed pages and dead objects repository-wide, along with their OOPs.

The Reclaim Gem includes a master session and multiple reclaim sessions, each being a thread within the Reclaim Gem process. This allows reclaim to occur in parallel.

By default, the Admin Gem and the Reclaim Gem with one reclaim session are configured to run, and are started automatically when the Stone is started. By default, epoch is disabled.

- ▶ We recommend that you leave the Admin Gem running at all times, although it is required only following a `markForCollection` or `markGcCandidatesFromFile:`, or epoch garbage collection operation. (Subsequent sections of this chapter describe these operations in detail.) If the Admin Gem is not running following one of these operations, the garbage collection process cannot complete, and garbage can build up in your system.
- ▶ We recommend that you have the Reclaim Gem running at all times, to reclaim shadow objects.

Both the Admin and Reclaim Gems are run from the GcUser account, a special account that logs in to the repository to perform garbage collection tasks. This account is used to set configuration values for the GcGems. To modify configuration parameter that apply to either GcGem, log in as GcUser (GcUser's default password is the same as DataCurator's or SystemUser's default password) and send the message at `:aKey put :aValue` to UserGlobals.

For example, to set `#reclaimMinPages` to 100:

```
topaz> set user GcUser password thePassword
login
...
topaz 1> printit
UserGlobals at: #reclaimMinPages put: 100.
System commitTransaction
%
```

Specific configuration parameters and how to apply them are discussed in detail later in this chapter.

GemStone's Garbage Collection Mechanisms

GemStone provides the following mechanisms that together mark and reclaim garbage, thereby helping you to control repository growth.

Marking

Repository-wide marking — To prevent the repository from growing large enough to cause problems on a regular basis, you can run `Repository >> markForCollection` (page 242). This method combines a full sweep of all objects in the repository and the marking of each possible dead object in a single operation.

FDC/MGC — This two-step process allows the first step to be performed on an offline copy of the repository, reducing the impact on the production system. In the first step (FDC), the method `Repository >> findDisconnectedObjectsAndWriteToFile:` explicitly finds disconnected objects in the repository and identifies a list of possible targets for garbage collection, and writes the OOPS of the possibly dead objects to a file. In the second step (MGC), the method `Repository >> markGcCandidatesFromFile:` examines the set of candidate garbage objects, ensures that they are not referenced by live objects, and then marks them as dead for subsequent garbage collection. For details about the FDC/MGC process, see page 244.

Epoch garbage collection — If enabled, the Admin Gem periodically examines all transactions written since a specific, recent time (the beginning of this *epoch*) for objects that were created and then dereferenced during that period. However, epoch garbage collection cannot reclaim objects that are created in one epoch but dereferenced in another. In spite of its name, epoch garbage collection only marks; it does not reclaim. You can configure various aspects to maximize its usefulness. Epoch garbage collection is disabled by default. For details about epoch garbage collection, see page 242.

Reclaiming

Reclaim — Once you've run `markForCollection`, FDC/MGC, or epoch garbage collection, the Reclaim Gem will reclaim pages that contain either dead or shadow objects. When there are a high number of objects needing to be reclaimed, you may increase the number of sessions under the Reclaim Gem. For details about reclaiming pages, see page 253.

GcLock

Many garbage collection process, such as mark/sweep operations, should not be run concurrently. To prevent this, there is a shared internal lock called the GcLock. Garbage collection processes that cannot run concurrently, such as markForCollection, get the GcLock, which prevents another one from starting up. The GcLock is also held by the Admin Gem at certain periods.

In addition to garbage collection, some repository-wide operations such as GsObjectInventory (page 125) also hold the GcLock while they are running.

If another task that requires the GcLock is in progress at the time you try to do markForCollection or findDisconnectedObjects..., they will not execute, but report an error similar to that shown below.

```
-- Request for MFC gclock by session 10 denied, reason: vote state  
is voting, sessionId not voted 2  
ERROR 2501 , a Error occurred (error 2501), Request for gcLock  
timed out. 'Request for MFC gclock by session 10 denied, reason:  
vote state is voting, sessionId not voted 2' (Error)
```

The cause of the conflict may be:

- ▶ Another operation that requires the GcLock is in progress in another session; this includes epoch, MFC, and MGC, and also operations such as GsObjectInventory.
- ▶ A previous epoch, markGcCandidatesFromFile: or markForCollection completed the mark phase, but voting on possibly dead objects has not completed.

For voting to complete, the Admin Gem must be running. Also, any long-running session that neither aborts nor commits will prevent the vote from completing.

Symbol Garbage Collection

Symbols in GemStone are a special case of Object, since they must always have a unique OOP across all sessions. To ensure this, symbol creation is managed by the SymbolUser, who creates all new Symbols. Symbols are stored in the AllSymbols dictionary, and are not removed, to avoid any risk of creating duplicate symbols.

However, there are cases where a large number of unimportant symbols are created, perhaps inadvertently. To reclaim this space and to manage the size of AllSymbols, you can configure GemStone to collect unreferenced symbols in a multi-step process that ensures that symbols in use are not collected.

By default, Symbol garbage collection is not enabled. It can be enabled using the configuration parameter, STN_SYMBOL_GC_ENABLED, or by the runtime equivalent, #StnSymbolGcEnabled. If enabled, symbol garbage collection is performed automatically in the background and requires no management.

When enabled, unused symbols are located and put in a possibleDeadSymbols collection as part of a markForCollection,. These symbols are hidden, to remove references from AllSymbols but retain the OOPs until the voting, union, and finalization is done. Any lookups on the hidden symbol will return the existing hidden symbol and restore it to the AllSymbols dictionary.

Once voting and write-set union sweep are done, the symbols that are otherwise unreferenced are removed from the possibleDeadSymbols, so they will be collected by the next markForCollection.

12.2 MarkForCollection

Privileges required: GarbageCollection.

The method `Repository>>markForCollection` sweeps the entire repository and marks as live all objects that can be reached through a transitive closure on the symbol lists in `AllUsers`, as described on page 234. The remaining objects become the list of possible dead objects.

`markForCollection` only provides a set of possible dead objects for voting and eventual reclaiming as described under “What Happens to Garbage?” on page 238. It does not reclaim the space or OOPs itself – the Reclaim Gem does that, as described beginning on page 253.

To mark unreferenced GemStone objects for collection, log in to GemStone and send your repository the message `markForCollection`, as in the following example:

```
topaz 1> printit  
SystemRepository markForCollection  
%
```

If you are performing `markForCollection` on a large production repository, consider the steps described under “Impact on Other Sessions” on page 243.

This method aborts the current transaction and runs `markForCollection` inside a transaction, but monitors the commit record backlog so it can abort as necessary to prevent the backlog from growing. When `markForCollection` completes, the session reenters a transaction, if it was in one when this method was invoked.

When `markForCollection` completes successfully, the Gem that started it displays a message such as the one below:

```
Warning: a Warning occurred (notification 2515), markForCollection  
found 110917 live objects, 3496 dead objects(occupying approx  
314640 bytes)
```

If another garbage collection task is in progress at the time you try to do `markForCollection`, this method will retry for a fixed period, reporting status. If the other operation does not complete within the timeout period, it reports an error indicating it could not get the `GcLock`. See “`GcLock`” on page 241 for more details.

Before issuing the error, the `markForCollection` method waits up to a minute for the other operation to complete. To have the `markForCollection` wait for a longer period, use `markForCollectionWait: waitTimeSeconds`. To wait as long as necessary for the other garbage collection to finish, pass the argument `-1`. Do so with caution, however; under certain conditions, the session could wait forever. To avoid this:

- ▶ Make sure that other sessions are committing or aborting, which allows voting on possible dead to complete.
- ▶ Make sure that the Admin Gem is running to complete processing of dead objects once the vote is completed.

Impact on Other Sessions

The `markForCollection` operation uses multi-threaded scan. For more details on this, see “Multi-Threaded Scan” on page 259.

By default, `markForCollection` limits its use of CPU resources if the CPU load on the system reaches 90%. It starts the operation with two threads and a page buffer size of 128. If the CPU limit is reached, the code automatically causes threads to sleep until the load is less than 90%. Depending upon the I/O required, the system may never reach this limit.

To enable `markForCollection` to complete as quickly as possible, you can use:

```
SystemRepository fastMarkForCollection
```

This uses higher settings (95% of CPU, and a number of threads based on the current hardware) to use as many system resources as possible. The performance of anything else running on the same system may be heavily degraded.

For maximum control, use the method

```
SystemRepository markForCollectionWithMaxThreads: threadsCt
  waitForLock: seconds
  pageBufSize: pageBufSize
  percentCpuActiveLimit: percentLimit
```

This allows you to specify the precise limits.

Starting `markForCollection` with these limits provides a specification for the trade-off you wish to make between speed to complete and the impact on other sessions. The desired trade-off may vary over time; for example, if your `markForCollection` extends over both business hours and non-business hours, you may accept greater impact during these periods of light load. The Multi-threaded scan parameters can be changed at runtime, as described on page 259.

After the `markForCollection` has completed, there may be additional impact on other sessions, since it is likely that dead objects that require reclaim were identified. After the remaining Garbage Collection steps have completed, the Reclaim Gem Sessions may become busy reclaiming the dead objects.

Scheduling markForCollection

To invoke `markForCollection` using the `cron` facility, create a three-line script file similar to the Topaz example on page 242 by entering everything except the prompt. Use this script as standard input to `topaz`, and redirect the standard output to another file:

```
topaz < scriptName > logName
```

Make sure that `$GEMSTONE` and any other required environment variables are defined during the `cron` job. Either create a `.topazini` file for a user who has GarbageCollection privilege, or insert those login settings at the beginning of the script. For information about using `cron`, refer to your operating system documentation.

12.3 The FDC/MGC Process

As discussed in the previous section, `markForCollection` combines a full sweep of all objects in the repository and marking of each possible dead object in a single operation. In a large production system, this can have a significant impact on performance. To lessen this impact, you can perform garbage collection as a two-step process:

`findDisconnectedObjectsAndWriteToFile:` (FDC) and `markGcCandidatesFromFile:` (MGC). These two steps together perform the same function as running `markForCollection`.

In this two-step process, the possibly dead objects are first written to a file, thus allowing you to run the FDC in a copy of the repository rather than in the production system. The OOPs of the possibly dead objects are then transferred to the production system, where they are marked as dead and subsequently garbage-collected.

Running in a copy of the repository avoids the performance impact of this processing on the production repository, and permits operations on the repository that would not be possible during a long-running repository sweep; for example, the production repository can be shut down.

If you prefer to run garbage collection in the production system rather than in a copy of the repository, `markForCollection` would be a better option.

For offline FDC/MGC to be effective, Epoch garbage collection must be disabled on the production repository during the period from the time the production copy is made, until the `markGcCandidates` completes. If the possibly dead objects that are passed into the MGC operation contain objects that have already been garbage collected and reused, performance degrades to the point of uselessness. Epoch garbage collection is disabled by default.

Step 1. Prepare for FDC.

FDC produces a large file of encoded OOPs; the size of the file (in bytes) is slightly more than eight times the number of dead objects. On the server that will run the backup, ensure that there is enough space for the resulting file, and that the filename you will use does not already exist.

Step 2. Create a copy of the production repository.

Make a backup of your production repository, and restore the backup into another location. You may use online extent file backup, offline extent file backup, or Smalltalk `fullBackup`; see Chapter 9, starting on page 233, for details on backup and restore

This restored repository that will run the FDC, while your production repository continues to operate normally.

Step 3. Run `findDisconnectedObjectsAndWriteToFile:....`

On the running backup repository, execute the FDC method:

```
SystemRepository
  findDisconnectedObjectsAndWriteToFile: filename.txt
  pageBufferSize: 128
  saveToRepository: true
```

This method does the following:

- ▶ Runs code similar to the `markForCollection` sweep phase on the backup repository.
- ▶ Stores the resulting list of candidate dead objects to the named file (*filename.txt*), as encoded OOPs. If a relative filename is specified, the file is created relative to the current working directory. Note that for RPC sessions the working directory is the home directory of the gem process' UNIX user.
- ▶ If `saveToRepository:` is true, stores the list of encoded OOPs to an internal array as well. Set this argument to true if you want to be able to subsequently examine and analyze the results of the FDC operation; it is not needed as part of the garbage collection process.

If `saveToRepository:` is true, you can subsequently use the array of dead objects to recreate the FDC file in the event that the file is deleted or lost. See the following discussion, "If You Need to Recreate the FDC File."

Also note that a larger `pageBufferSize` may improve performance, provided you have adequate memory. See "Memory Impact" on page 260.

Step 4. At this point, you can shut down the backup repository; it is no longer needed for the MGC/FDC process.

Step 5. Run `markGcCandidatesFromFile:` (page 245) on the production system.

If You Need to Recreate the FDC File

If your FDC file should inadvertently be deleted or become lost, you may still be able to recreate the FDC file for subsequent use by `markGcCandidatesFromFile:` (MGC) – provided that `saveToRepository:` was true when you ran FDC.

You can execute this method to write the array of candidate dead objects found by the last FDC operation (for which `saveToRepository:` was true) to a file:

```
SystemRepository writeFdcArrayToFile: aFileNameString
```

NOTE

Before running this method, ensure that you have sufficient disk space to hold the entire file, whose size in bytes will be approximately four times the size of the array.

This method expects Globals at: `#FdcResults` to contain the array of candidate dead objects, in order and sorted by OOP. You must specify the file as a String, using the full path and filename of a file that does not yet exist.

The method returns true if successful, false if the entire file could not be written, or nil if the results of the last FDC could not be found.

Run `markGcCandidatesFromFile:`

Privileges required: `GarbageCollection`.

`markGcCandidatesFromFile:` (MGC) performs a multi-threaded sweep of the repository to identify any objects in the FDC output file that still have references. If the only references to candidate objects are from other candidates, the objects are marked for subsequent reclaim.

The MGC operation scans all *non-candidate* objects to ensure that there are no references to candidate objects. If there are no such references, the MGC operation is complete. If, however, there are any references to candidate objects, the MGC operation performs a transitive closure on all *former candidate* objects to identify any other candidate objects that cannot be collected.

To minimize the impact on performance, run the MGC operation at a time when the load on the production repository is low, such as during off-hours.

On the production system, execute:

```
SystemRepository markGcCandidatesFromFile: aFileNameString
where aFileNameString is a String representing the full path to the file of encoded
OOps that was created during the FDC (page 244, Step 3). Make sure that this file is
available from the production location.
```

`markGcCandidates` requires the `GcLock`. If another garbage collection (`markForCollection`) is in progress at the time you try to do `markGcCandidatesFromFile:`, or if voting has not completed, it may report a conflict error similar to that shown on page 241. Try `markGcCandidatesFromFile:` again after a few minutes.

Recall that marking only provides a set of possible dead objects for voting and eventual reclaiming, as described under “What Happens to Garbage?” on page 238. Marking does not by itself reclaim space or identifiers—the Reclaim Gem does that, as described beginning on page 253.

As mentioned elsewhere in this chapter, the Admin Gem must be running to complete processing of dead objects following an MGC operation.

In some cases, there could be irrelevant errors that would prevent MGC from completing. To override such errors, execute the following method:

```
SystemRepository markGcCandidatesFromFile: aFileNameString
forceOnError: aBoolean
```

This method is virtually identical to `markGcCandidatesFromFile:`, except that if *aBoolean* is true, the MGC will run even if the given file contains objects that fail validation, or there are references from live objects to some objects in the candidate collection. The live objects in the candidate collection are not marked as dead.

Impact on Other Sessions

The FDC operation uses multi-threaded scan. For more details on this, see “Multi-Threaded Scan” on page 259.

Since FDC is designed to be run in an offline copy of the repository, by default, FDC uses very aggressive settings, 95% of CPU and two threads per CPU core. It may consume all CPU and disk I/O host system resources.

If you want to run the FDC with a reduced impact on other sessions, you can use the basic method:

```
basicFindDisconnectedObjectsAndWriteToFile: aFileNameString
pageBufferSize: pageBufSize saveToRepository: saveToRep
withMaxThreads: maxThreads maxCpuUsage: aPercentage
```

which allows you to include the specific thread and CPU limits.

12.4 Epoch Garbage Collection

Privileges required: `GarbageCollection`.

Epoch garbage collection operates on a finite set of recent transactions: the *epoch*. Using the write set that the Stone maintains for each transaction, the Admin Gem examines every object created during the epoch. If an object is unreferenced by the end of the epoch, it is marked as garbage and added to the list of possible dead objects.

Epoch collection is efficient because:

- ▶ It's faster and easier to perform a transitive closure on a few recent transactions than on the entire repository.
- ▶ Most objects die young, especially in applications characterized by numerous small transactions updating a few previously committed objects. An epoch of the right length can collect most garbage automatically.

Although epoch collection identifies a lot of dead objects, it cannot replace `markForCollection` because it will never detect objects created in one epoch and dereferenced in another.

By default, epoch garbage collection is disabled. You can enable it in either of two ways:

- ▶ Before you start the Stone, set the `STN_EPOCH_GC_ENABLED` configuration parameter to `TRUE`. For details about this parameter, see page 292.
- ▶ Execute the method `System class >> enableEpochGc`. You may also manually disable epoch garbage collection using `System class >> disableEpochGc`. Using these methods updates the system configuration file.

After your installation has been operating for a while, and you've had the chance to collect operational statistics, consider this: epochs of the wrong length can be notably inefficient. An in-depth discussion of the performance trade-offs of short or long epochs starts on page 248.

Running Epoch Garbage Collection

When epoch garbage collection is enabled, it will run automatically according to the `GcUser` configuration parameters `#epochGcTimeLimit` and `#epochGcTransLimit`.

You can force an epoch garbage collection to begin using `System class >> forceEpochGc`. `forceEpochGc` will return `false`, and not start an epoch garbage collection, if any of the following are true:

- ▶ Checkpoints are suspended.
- ▶ Another garbage collection operation is in progress.
- ▶ Unfinalized possible dead objects exist (that is, `System voteState` returns a non-zero value).
- ▶ The system is in restore mode.
- ▶ The Admin Gem is not running.
- ▶ Epoch garbage collection is disabled (that is, `STN_EPOCH_GC_ENABLED = FALSE`).
- ▶ The system is performing a `reclaimAll`.

- ▶ A previous `forceEpochGc` operation was performed and the epoch has not yet started or completed.

Tuning Epoch

Epoch Configuration Parameters

The following configuration parameters are available to control the performance of epoch garbage collection, and are stored in `GcUser's UserGlobals`.

<code>#epochGcTimeLimit</code>	The maximum frequency of epoch garbage collection (in seconds). Default: one hour (3600 seconds). This value should be at least 1800 (30 minutes), since the aging of objects faulted into Gem memory uses 5 minute aging for each of 10 subspaces of the POM generation.
<code>#epochGcTransLimit</code>	The minimum number of transactions required to trigger epoch garbage collection. Default: 5000.
<code>#epochGcPercentCpuActiveLimit</code>	Limit active epoch threads when system <code>percentCpuActive</code> is above this limit. Default: 90.
<code>#epochGcPageBufferSize</code>	Size in pages of buffer used for epoch GC (must be power of 2). Default: 64.
<code>#epochGcMaxThreads</code>	The <code>MaxThreads</code> used for next <code>epochGc</code> .

Epoch garbage collection uses the multi-threaded scan (see page 259) and can be tuned to complete more quickly with more system performance and resources impact, or take longer and use fewer system resources. The parameters `#epochGcPercentCpuActiveLimit`, `#epochGcPageBufferSize`, and `#epochGcMaxThreads` are used to tune epoch garbage collection's multi-threaded scan impact.

Determining the Epoch Length

Epoch garbage collection's ability to identify unreferenced objects depends on the relationship between three variables:

- ▶ The rate of production R of short-lived objects.
- ▶ The lifetime L of these objects.
- ▶ The epoch length E .

The only variable under your direct control is epoch length. Although you cannot specify it explicitly, the following configuration parameters jointly control the length of an epoch:

- ▶ `#epochGcTimeLimit`
- ▶ `#epochGcTransLimit`

Epoch garbage collection occurs when:

```
(the time since last epoch > epochGcTimeLimit ) AND
(transactions since last epoch > epochGcTransLimit)
```


The following discussion assumes that the epoch is determined by the minimum time interval (`#epochGcTimeLimit`) because other threshold is always met.

Figure 12.5 shows the effect of the epoch on the number of items marked. If $L = E$, for example, five minutes, every object's lifetime spans epochs (top part of graph), and none are collected.

When the epoch is longer than an average object's lifetime, however, some objects live and die within the same epoch, and can be marked. The lower part of Figure 12.5 shows an example where $E = 3L$ and objects are created at a uniform rate. Objects created during the first two-thirds of the interval die before its end and are marked. Only those created during the final third survive to the next epoch.


The results shown in Figure 12.5 can be expressed as:

$$\begin{aligned}\text{Objects Missed by EpochGC} &= R \times L \\ \text{Objects Recovered by EpochGC} &= R(E - L)\end{aligned}$$

For example, assume $R = 1000$ objects per minute, $L = 5$ minutes, and $E = 15$ minutes. Then, for each epoch:

$$\begin{aligned}\text{Objects Missed} &= 1000 \times 5 = 5000 \\ \text{Objects Recovered} &= 1000 (15 - 5) = 10000\end{aligned}$$

Figure 12.5 Effect of Collection Interval on Epoch Garbage Collection



Therefore:

- ▶ Set `#epochGcTimeLimit` $E >$ lifetime L of short-lived objects.


Figure 12.6 graphs the effect of the epoch. When $E = L$, epoch garbage collection is in effect disabled – all objects survive into the next epoch; the number of unmarked yet dead objects in the repository grows at the creation rate. These dead objects remain unidentified until you run `markForCollection`.

When the epoch is extended so that $E = 3L$, each epoch garbage collection marks those objects both created and dereferenced during that interval. This ratio causes the sawtooth pattern in the graph. If the creation rate is uniform, two-thirds of the dead objects are marked $((E - L)/E)$, and one-third are missed (L/E) . Consequently, the repository grows at one-third the rate of the case $E = L$.

This configuration trades short bursts of epoch garbage collection activity for:


- ▶ moderate growth in the repository, and
- ▶ the need to run `markForCollection` often enough to mark dead objects that survive between epochs.

Figure 12.6 Repository Growth with Short Epoch



Suppose we extend the epoch to $E = 12L$. The result is shown in Figure 12.7, superimposed on part of the previous figure.

Figure 12.7 Effect of Longer Epoch on Repository Growth



Although the longer epoch allows many more dead objects to accumulate, the growth rate of the repository is substantially less—25% of the previous case.

This configuration trades a slower growth rate for:

- ▶ a need for greater headroom on the disk, and
- ▶ longer bursts of epoch garbage collection activity.

Certain cases have needed an epoch as long as several hours, or even a day.

Cache Statistics

Several cache statistics include information about the epoch garbage collection process. These are visible by using statmonitor data viewed in VSD (the visual statistics display tool). You may also access methods in System to get the values programmatically; see “Programmatic Access to Cache Statistics” on page 127 for more information.

The following statistics may be useful in monitoring epoch:

EpochGcCount	The number of times that the epoch garbage collection process was run by the Admin Gem since the Admin Gem was started. For a system in steady state, look for uniform periods between runs or a uniform run rate.
EpochNewObjs	The number of new objects that were created during the last epoch.
EpochPossibleDeadObjs	The number of possible dead objects found by the last epoch garbage collection.
EpochScannedObjs	The number of objects scanned by the last epoch garbage collection.

12.5 Reclaim

The Reclaim Gem is responsible for reclaiming both dead and shadowed objects (see “Shadow or Dead?” on page 235 for the difference between these types of garbage).

Shadowed objects are created naturally as your application modifies existing objects, so it is a good idea to always have the Reclaim Gem running to avoid shadowed objects accumulating. Some operations, such as migration, create a very large number of shadowed objects that need to be reclaimed.

After a mark/sweep operation – MFC, FDC/MGC, or epoch – completes, there will be a number of dead objects that need to be reclaimed.

Although it is objects that are dead or shadowed, reclaim is done in pages. Pages that contain dead or shadowed objects may also contain some live objects; these live objects are copied to fresh pages, and the resulting page may then be reclaimed.

Reclaim is performed multi-threaded. Each thread within the Reclaim Gem is similar to a session, but runs within the Reclaim Gem process.

When the Reclaim Gem is running, its sessions examine pages marked reclaimable because they contain either dead or shadow objects, and reclaim fragments of space left by transactions that did not fill an entire page. This occurs in the background, with no specific action required.

Although it is recommended to allow the background processes to perform the reclaim, you can explicitly invoke it by executing:

```
SystemRepository reclaimAll
```

Reclaimed space does not appear as free space in the repository until other sessions have committed or aborted all transactions concurrent with the reclaim transaction, and the Stone has disposed the commit record. If other users are logged in and holding up this process, you can determine which sessions are viewing the oldest commit record, thereby impeding reclaim. See the discussion, “Further Tuning Garbage Collection,” on page 259.

Tuning Reclaim

Reclaim Configuration Parameters

The following configuration parameters are available to control the reclaim task, and are stored in GcUser's UserGlobals.

`#deadObjsReclaimedCommitThreshold`

The maximum number of dead objects to reclaim in a single transaction, including dead objects reclaimed when reclaiming shadow pages. The default is 20000.

`#deferReclaimCacheDirtyThreshold`

If the primary shared page cache (the shared cache on the stone's machine) is more than this percentage dirty, then reclaim gems will wait until the cache is less than 5% below this threshold before resuming reclaims. The default is 75%.

`#maxTransactionDuration`

The maximum length (in seconds) of a GcGem transaction. The transaction will be committed once this time is exceeded. Must be ≥ 10 ; the default is 300, maximum is `SmallInteger maximumValue`.

`#objsMovedPerCommitThreshold`

The approximate maximum number of live objects to move in a reclaim transaction. Must be ≥ 100 ; the default is 20000, maximum is `SmallInteger maximumValue`.

`#reclaimDeadEnabled`

A Boolean indicating whether or not to reclaim dead objects; the default is true.

`#reclaimMinFreeSpaceMb`

Minimum repository free space which must be available in order for reclaims to proceed. Reclaims will be temporarily suspended if the repository free space drops below this threshold. The default value of 0 specifies a limit computed as the current size of the repository divided by 1000, with a minimum value of 5MB. Default and minimum 0, maximum 65536.

`#reclaimMinPages`

The minimum number of pages to process in a single reclaim operation (reclaiming does not start until this threshold is reached). Must be ≥ 1 ; the default is 30 pages, maximum is `SmallInteger maximumValue`.

`#sleepTimeBetweenReclaimMs`

The minimum amount of time in milliseconds that the process will sleep between reclaims, even when work is scheduled. The default is 0 milliseconds, maximum 3600000.

Reclaim Commit Frequency

A Reclaim Gem session will commit reclaim changes as soon as any one of the following conditions is met:

- ▶ Number of live objects moved exceeds `#objsMovedPerCommitThreshold`.
- ▶ Duration of the transaction exceeds `#maxTransactionDuration`.
- ▶ Number of dead objects reclaimed exceeds `#deadObjsReclaimedCommitThreshold`.

Controlling the impact of reclaim

Reclaim, particularly with a larger number of sessions configured for the Reclaim Gem, can perform quickly but place a large load on your system. If you are likely to be doing reclaim during periods where users will also need to use the system, you may wish to slow down reclaim. This can be done in a number of ways:

- ▶ Reduce the number of reclaim sessions using `System class >> changeNumberOfReclaimGemSessions:` with a argument of 1 or 0.
- ▶ Set `#sleepTimeBetweenReclaimMs` to ensure that reclaim Gem sessions pause between reclaim operations.

Speeding up reclaim

You can also setup your system to run reclaim with the maximum impact during off-hours. If you have a large amount of reclaim to perform, this allows the reclaim to finish more quickly. You can increase the number of Reclaim session to the maximum using:

```
System class >> startMaxReclaimGemSessions
```

This will start the number of sessions specified by `STN_MAX_GC_RECLAIM_SESSIONS`.

Avoiding disk space issues

Reclaim requires pages from the repository in order to copy non-dead objects. There are further steps that the stone must complete, before the space on the reclaimed pages is available again. So initially, reclaim will cause the amount of free space in the repository to drop.

Depending on overhead required by your system and the largest amount of reclaim that needs to be done at any time, you may want to configure a larger `#reclaimMinFreeSpaceMb`. This will ensure that reclaim pauses before your repository becomes dangerously low in free space.

Cache Statistics

Several cache statistics provide information about reclaim. These are visible by using `statmonitor` data viewed in VSD (the visual statistics display tool). You may also access methods in `System` to get the values programmatically; see “Programmatic Access to Cache Statistics” on page 127 for more information.

The following Stone statistics may be useful in monitoring reclaim:

`DeadNotReclaimedObjs`

The number of objects known to be dead but not yet reclaimed.

DeadObjsReclaimedCount	The total number of dead objects reclaimed since the Stone repository monitor process was last started.
GcVoteState	Indicates the current phase of garbage collection: Gems voting, voting complete, Possible Dead Write-Set Union Sweep (PDWSUS) in progress, or PDWSUS complete.
PagesNeedReclaimSize	The amount of work waiting for the reclaim task.
PossibleDeadObjs	The number of objects marked as dereferenced but not yet declared to be dead.
ReclaimCount	The number of times the reclaim process has been run.
ReclaimedPagesCount	The number of scavenged pages.

12.6 Running Admin and Reclaim Gems

Admin Gem Privileges required: GarbageCollection

The initial configuration for the Admin and Reclaim Gems are provided in the system configuration file for the stone; by default, `$GEMSTONE/data/system.conf`. These settings determine what is started automatically when the stone starts up. During runtime, you can start and stop the Admin Gem and change the number of Reclaim sessions that are running.

Configuring Admin Gem

The Admin Gem is enabled or disabled by the configuration parameter `STN_ADMIN_GC_SESSION_ENABLED`, described in detail on page 288. By default, this is enabled, and normally you should leave this enabled. You can stop and restart the Admin Gem at runtime as needed.

Configuring Reclaim Gem

The number of Reclaim sessions is set by the configuration parameter `STN_NUM_GC_RECLAIM_SESSIONS`, described in detail on page 299. By default, this is one, and you should normally keep at least one Reclaim session running. Most systems will benefit from increasing the number of Reclaim sessions. In general, we recommend running one Reclaim session for between 5 and 10 extents. You may need to experiment to find the correct balance for your system. The number of Reclaim sessions can be changed at runtime as needed.

To ensure that Reclaim sessions do not impact the number of user sessions, a separate parameter, `STN_MAX_GC_RECLAIM_SESSIONS`, configures the maximum number of Reclaim sessions you will be running. This parameter is described in detail on page 298. By default, this is set to the number of extents on your system. This parameter cannot be changed without restarting the stone. The upper limit for the number for the number of Reclaim sessions that can be run under any configuration is 255.

While the number of Reclaim sessions should normally be less than or equal to `STN_MAX_GC_RECLAIM_SESSIONS`, it is possible to start a larger number of Reclaim sessions. However, this will reduce the number of user sessions that can login to this

Stone. If your system does not have excess unused user sessions, you should be careful to configure `STN_MAX_GC_RECLAIM_SESSIONS` high enough that you will never want to run a larger number of Reclaim sessions.

Starting GcGems

You can ensure all configured GcGems are running using:

```
System startAllGcGems
```

If the Admin Gem is not running, start it. If the Reclaim Gem is not running, start it with the configured number of Reclaim sessions. Return true if the Admin Gem and at least one Reclaim sessions are started.

or by executing both:

```
System startAdminGem
```

If the Admin Gem is not running, start it. Return true if the Admin Gem is running, false if the Admin Gem could not be started.

```
System startReclaimGem
```

If the Reclaim Gem is not running, start it with the configured number of Reclaim sessions. Return the number of Reclaim sessions that will be running. If the Reclaim Gem is already running, has no effect and returns the number of Reclaim sessions already running.

It may take a little time for the GcGems to complete login. The above methods do not block; they initiate the startup and return immediately. To wait a given period of time for the GcGems to start up:

```
System waitForAllGcGemsToStartForUpToSeconds: anInt
```

If the Admin Gem is not running, start it. If the Reclaim Gem is not running, start it with the configured number of Reclaim sessions. If all the GcGems have not started up within that time, return false. However, this does not necessarily mean that any GcGems have failed to start; on a slow system with a short timeout, this method may return false, even though all GcGems eventually start correctly.

To confirm that the GcGems are running,:

```
System hasMissingGcGems
```

Returns false if either the Admin Gem or the Reclaim Gem is not running.

To determine the number of Reclaim sessions that are currently running:

```
System reclaimGcSessionCount
```

Returns the total number of Reclaim sessions that are running.

Stopping GcGems

To ensure that the Admin Gem and all Reclaim sessions are stopped:

```
System stopAllGcGems
```

or you may execute both:

```
System stopAdminGem
```

```
System stopReclaimGem
```

Adjusting the number of Reclaim sessions

You can adjust the number of Reclaim sessions that are running during the course of operation of your application. When there is a large amount of reclaim and little other load on your system, running a large number of Reclaim sessions will allow the reclaim work to complete more quickly. During normal operation, reducing the number of Reclaim sessions avoids using too many system resources and impacting users.

To set the number of Reclaim sessions that are running:

```
System changeNumberOfReclaimGemSessions: targetReclaimSessionCount  
  Start the ReclaimGem, if it is not running, with targetReclaimSessionCount Reclaim  
  sessions.
```

targetReclaimSessionCount should be a number less than or equal to the value for STN_MAX_GC_RECLAIM_SESSIONS. Using a larger argument does not error, but may have consequences for user logins; see the discussion on page 256.

Return the new target number of Reclaim sessions; Reclaim sessions will be started or stopped to reach this number. This method does not block, so it may take a little time before the correct number of Reclaim sessions is actually running.

Using this method only changes the currently running number of Reclaim sessions, but does not affect the configured number. After stopping the ReclaimGem, on restart the regular configured number of sessions will be started.

To change the default number of Reclaim sessions that will be started by default when the ReclaimGem starts up:

```
System configurationAt: #StnNumGcReclaimSessions  
  put: targetReclaimSessionCount.
```

This does not effect the number of Reclaim sessions that are currently running, if any. Changes to the runtime parameter do not persist if the Stone is restarted. For a permanent change, you should edit the configuration parameters in the configuration file used by the stone: STN_NUM_GC_RECLAIM_SESSIONS, and if necessary, STN_MAX_GC_RECLAIM_SESSIONS.

12.7 Further Tuning Garbage Collection

Multi-Threaded Scan

For large systems, it can take a considerable amount of time to scan the entire repository, as is required by a mark/sweep operation (or other operations such as listInstances). To allow these scans to complete faster, operations that scan the entire repository use multiple threads running in parallel. There is a trade-off between how fast the operation completes and how much of the system resources it uses. Obviously, the faster the scan completes, the less of anything else can be done on that system during that period.

The trade-off can be configured per operation and can be changed as the operation proceeds. You can choose to run it to complete as quickly as possible but use all the system resources, or with minimal impact on the rest of your application, but taking much longer to complete. You can switch between these approaches as often as you need to.

Multi-threaded scan operations must be performed by Gems running locally to the Stone. When executed from a Gem that is running on a machine remote from the Stone, the scan will be performed by one thread, regardless of the number of threads requested. You can use `System class>>sessionIsOnStoneHost` to check that the gem is local.

Tuning Multi-Threaded Scan

Each session has two variables that control the impact of the multi-threaded operation it is running:

<code>MtThreadsLimit</code>	The upper limit on the number of threads can be activated.
<code>MtPercentCpuActiveLimit</code>	The total CPU load level at which the scan starts to deactivate threads.

These variables are arguments to all scan operations, although most scan operations have variants that use default values.

While these variables are passed in during operation startup, you can also update them while the scan is running. This enables you, for example, to reduce impact during working hours, while allowing more resources to be used during off hours.

Since the scan is running, of course, you need to update these variables from a second session, using the `sessionId` of the session that is running the scan.

One way to determine the session Id of the session that is running a scan operation is by checking the session holding the `GcLock`. However, while only one session can be holding the `GcLock` at a time, and `markForCollection` and `FDC` require the `GcLock`, other operations, such as `GsObjectInventory`, also may have the `GcLock`.

To access the upper limit on the number of threads:

```
System mtThreadsLimit: aSessionId
```

To update the upper limit on the number of threads:

```
System mtThreadsLimit: aSessionId setValue: anInt
```

To access the CPU load limit:

```
System mtPercentCpuActiveLimit: aSessionId
```

To update the CPU load limit:

System mtPercentCpuActiveLimit: *aSessionId* setValue: *anInt*

Both of these variables are used in tuning, but they have somewhat different uses. The primary way you will tune the impact on your system is by setting MtPercentCpuActiveLimit. The operation then controls its impact by activating or deactivating threads, up to a limit of MtThreadsLimit. The operation will proceed, using more or less resources at any particular time depending on what else is executing on your system. Note that the CPU load includes non-GemStone process running on this same machine, so if a machine is heavily used by non-GemStone processes, the operation may make little progress even if the GemStone repository itself is idle.

MtThreadsLimit acts as a ceiling on the impact as well. Since this limit is of more relevance within GemStone, on heavily loaded machines you may want to pay more attention to this limit to control the impact within the repository. This limit is also useful when you want to pause the scan. Setting the MtThreadsLimit to 0 means that the scan cannot perform work, but does not stop executing, it waits until a non-zero limit is set.

Cache Statistics

The following cache statistics are important for tuning multi-threaded scans. These are visible by using statmonitor data viewed in VSD (the visual statistics display tool); see *VSD User's Guide*. You may also access methods in System to get the values programmatically; see "Programmatic Access to Cache Statistics" on page 127 for more information.

MtThreadsLimit	The upper limit on the number of threads that can be running at any one time.
MtPercentCpuActiveLimit	The upper limit on percent of CPU that can be active before threads are deactivated.
percentCpuActive	The current percentage of CPU that is active.
MtActiveThreads	The current number of active threads

Memory Impact

Multi-threaded operations may require considerable heap memory. This memory requirement is **not** part of temporary object cache memory. You can configure your GEM_TEMPOBJ_CACHE_SIZE according to other application Gem requirements, or even configure the sessions performing repository scan operations with a very small temporary object cache size.

The amount of memory space that is needed depends primarily upon the current oopHighWater value, the number of threads, and the page buffer size. markForCollection uses a pageBufferSize of 128, epoch and writeSetUnionSweep use a size of 64, and it is an explicit argument to FDC.

The overhead associated with the oopHighWater value can be computed as:

$$(\text{stnOopHighWater} + 10\text{M}) / 2$$

The memory cost per thread is:

$$50\text{K} + (180\text{K} * \text{pageBufSize})$$

For example, a system with an `oopHighWater` mark of 500M running eight threads with a page buffer size of 128 would require a minimum of about 440 MB of free memory.

Identifying Sessions Holding Up Voting

Voting is phase 4 of garbage collection, as described on page 238. During this phase, each logged-in gem must vote on possibly dead objects. Sessions perform this vote on the next abort or commit that they execute, or on logout. If there are idle sessions that do not commit or abort, voting will not be able to complete.

You may find these sessions using:

```
System class >> notVotedSessionNames
```

The method `System class >> descriptionOfSession:` can help in tracking down such sessions. The array returned by this method includes the not voted status, as element 20. For details, see the comment in the image.

Tuning Write Set Union Sweep

The write set union sweep is phase 5 of garbage collection, as described on page 239. It is performed by the Admin Gem.

The write set union sweep is performed using the multi-threaded scan (page 259), and can be tuned using the following `GcGem` parameters:

```
#sweepWsUnionPercentCpuActiveLimit
    Limit active wsUnion threads when system
    percentCpuActive is above this limit. Default: 90.

#sweepWsUnionPageBufferSize
    Size (in pages) of buffer used for wsUnion sweep. Must be a
    power of 2. Default: 64. Minimum: 8. Maximum: 1024.

#sweepWsUnionMaxThreads
    The maximum threads used for next wsUnion sweep. By
    default, use one thread.
```

Identifying Sessions Holding Up Page Reclaim

Reclaiming pages can proceed only up to those pages currently providing some session's transaction view of the repository — that is, only up to the oldest commit record. When other sessions are logged in, reclaim stops at that point until all sessions using that commit record either commit or abort their transaction.

It can be helpful to identify which sessions are holding on to the oldest commit record. The method `System class>>sessionsReferencingOldestCr` returns an array of session IDs, which can be mapped to GemStone logins through `System class>>currentSessionNames` or `System class>>descriptionOfSession:aSessionId`. For example:

```
topaz 1> printit
System sessionsReferencingOldestCr
%
an Array
  #1 5
```

```
topaz 1> printit
System currentSessionNames
%
session number: 2    UserId: GcUser
session number: 3    UserId: GcUser
session number: 4    UserId: SymbolUser
session number: 5    UserId: DataCurator
```

The method `descriptionOfSession:` is particularly useful in that it returns an array of descriptive information. The second element is the operating system process ID (pid), and the third element is the name of the node on which the process is running. For details, see the comment in the image.

Finding large objects that are using excessive space

If you know that you have large objects that are no longer needed, another way to free space is to explicitly remove references to them. To remove such objects, you must first identify them. Then you can find all references to them and remove those references.

Identify Large Objects in the Repository

The following expression causes GemStone to look through the symbol list for each user in `AllUsers` and gather information on any named objects larger than the `SmallInteger aSize`.

```
topaz 1> printit
AllUsers findObjectsLargerThan: aSize limit: aSmallInt
%
```

This method locates large collections or strings stored directly in the Symbol Lists, for example, as created by an expression such as

```
UserGlobals at: #aCollection put: IdentityBag new
```

It will not locate collections stored within the class variables of classes or stored in instances of classes.

It returns an Array of up to `aSmallInt` elements, each of the form:

```
{ { aUserId . aKey . anObject } }
```

where `anObject` is an object larger than `aSize` defined in the symbol list of `aUserId`, and `aKey` is the Symbol associated with that object.

If any references to `anObject` are protected by a `GsObjectSecurityPolicy` for which you do not have read authorization, the last element of the result is the String 'Read Authorization Error encountered'.

With the information from this method, you may be able to track down temporary collections that were inadvertently committed, by checking the `SymbolLists` of the given user and locating an object referenced by the given key. However, if there are large objects that are not stored directly in symbol lists, you can send the same message to `System` to perform a global search.

The following method which will return any large objects, regardless of where they are stored:

```
topaz 1> printit
System findObjectsLargerThan: aSize limit: aSmallInt
%
```

This returns an Array of all objects in the repository larger than the SmallInteger *aSize*, whether they are named in a user's symbol list or not. As above, the Array is limited to a maximum of *aSmallInt* elements.

Again, if any references to an Object reside in ObjectSecurityPolicies for which you do not have read authorization, the last element of the result is the String 'Read Authorization Error encountered'.

This method performs the repository wide search using multi-threaded scan; see page 259 for a description. `findObjectsLargerThan:limit:` uses two threads and a CPU limit of 100%, so it may have some impact, particularly on smaller host systems, and may take some time to complete.

To complete the operation as quickly as possible, using as much of the system resources as necessary, you can use the similar method `System class >> fastFindObjectsLargerThan:limit:.`

Finding References to an Object that prevent garbage collection

Full reference path

To find a complete reference path to a particular object, you may use the method `Repository >> findReferencePathToObject:.` This method reports the complete reference path from a root object to the argument.

NOTE

This method runs in transaction, may take a considerable time to run. Avoid using it in production systems.

The result of this method is an array, with the first two elements consisting of the search object and true or false indicating if a reference path was found. No reference path means that the object would be garbage collected by the next MFC or FDC/MGC cycle. If a reference path was found, following this in the array are the list of objects that connect a root object to the search object, with the search object again appearing at the end, the reference to the search object in the second to last position, and so on.

For example,

```
topaz 1> run
SystemRepository findReferencePathToObject: AllDeletedUsers
%
<RefPathScan information>
an Array
  #1 an IdentitySet
  #2 true
  #3 a SymbolDictionary
  #4 an IdentityCollisionBucket
  #5 a SymbolAssociation
```

#6 an IdentitySet

AllDeletedUsers is an IdentitySet (#1 and #6). It is referred to in a SymbolDictionary (#3), using internal implementation objects (an IdentityCollisionBucket and a SymbolAssociation) that actually have the references.

This method returns the first path from a root object to the argument object that is found, but there may be multiple paths. You may search for multiple reference paths to the object using the method `Repository >> findAllReferencePathsToObject: .` This returns an Array of Arrays containing comparable information. See the method comments in the image for more information.

Note that you cannot find references to a Class or Metaclass using these methods.

Once you have found the references to the unwanted object, set those references to `nil`. This allows the object to be removed during normal garbage collection.

All References

You can search the repository for multiple references to an object by sending the following message:

```
topaz 1> printit
anObject findReferencesWithLimit: aSmallInt
%
```

This returns an Array of objects in the repository that reference *anObject*. If an object contains multiple references to *anObject*, that object will appear only once in the resulting Array. The Array is limited to a maximum of *aSmallInt* elements, to limit the duration and memory requirements of the result.

The resulting Array contains only those references that are protected by `GsObjectSecurityPolicies` for which you have read authorization. If any references to *anObject* are protected by `GsObjectSecurityPolicies` for which you do not have read authorization, the last element of the result is the String 'Read Authorization Error encountered'.

`findReferencesWithLimit:` uses multi-threaded scan (see page 259 for details). It uses a default of two threads and up to 90% of the CPU.

GemStone Configuration Options

A GemStone/S 64 Bit configuration file is a file containing information that, when read at startup time, can control the configuration, behavior, and functionality of the system at run time. Some of these configuration settings can be modified dynamically by sending messages in GemStone Smalltalk.

The Stone, Gem, and linked applications (collectively, the repository executables) are able to read two different types of configuration files: system-wide configuration files and executable-dependent configuration files.

- **System-wide configuration files** allow the GemStone system administrator to set options pertaining to all GemStone executables on a system- or network-wide basis. This file is required for a Stone to start.

System-wide configuration files are found in a default location, passed by argument, or located by a `GEMSTONE_SYS_CONF` environment variable.

- **Executable-dependent configuration files** can be used by individual users to control their own running copy of the GemStone system. Options contained in executable-dependent configuration files override the options specified in a system-wide configuration file.

Executable-dependent configuration files are found by name, passed by argument, or located by a `GEMSTONE_EXE_CONF` environment variable.

These environment variables can be set in the usual way. For example:

```
$ GEMSTONE_EXE_CONF=$HOME/myFile.conf
$ export GEMSTONE_EXE_CONF
```

Both `GEMSTONE_SYS_CONF` and `GEMSTONE_EXE_CONF` can be defined to point to either a file or a directory.

GemStone executables such as **startstone**, **pageaudit**, and **topaz** accept command line arguments to point to configuration files:

- The **-z** option sets the system configuration file.
- The **-e** option sets the executable-dependent configuration file.

A.1 How GemStone Uses Configuration Files

At startup time, GemStone repository executables attempt to find and read both a system-wide and an executable-dependent configuration file, searching for these files in the following manner.

Search for a System-Wide Configuration File

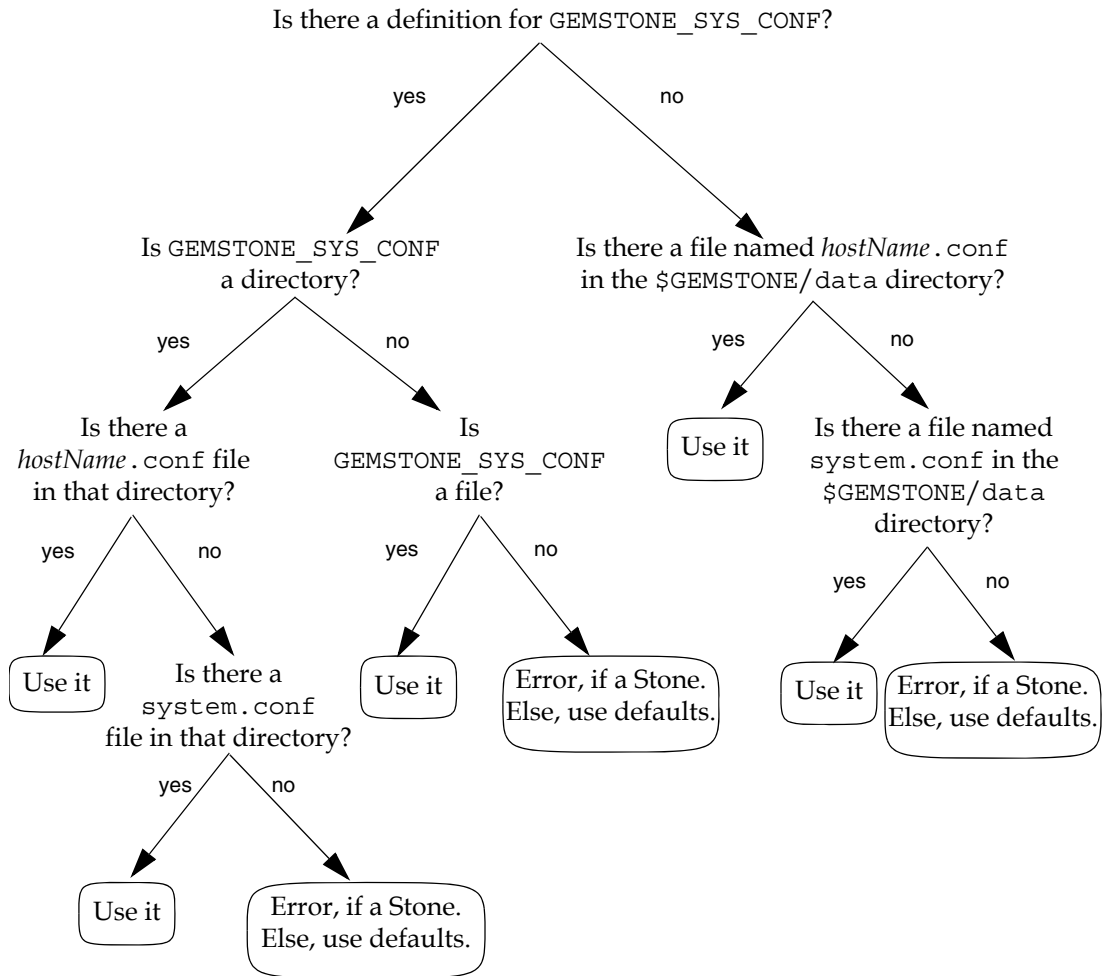
GemStone repository executables begin by attempting to find a system-wide configuration file.

1. As shown in Figure A.1, GemStone first checks to see if there is an environment variable defined for `GEMSTONE_SYS_CONF`.
2. If `GEMSTONE_SYS_CONF` is *not* defined, GemStone looks for a file named `hostName.conf` in `$GEMSTONE/data` and uses that file. `hostName` must match the results of executing the `hostname` command on the machine on which the executables are running.
3. If no such file exists, it looks for a file named `system.conf` in `$GEMSTONE/data` and uses that.
4. If neither of those files exist, the system defaults are used, unless the executable is Stone, in which case an error is generated.
5. If `GEMSTONE_SYS_CONF` is defined, GemStone checks to see if it points to a directory.
6. If `GEMSTONE_SYS_CONF` points to a directory, GemStone looks for a file named `hostName.conf` in that directory. If it finds such a file, it uses it. `hostName` must match the results of executing the `hostname` command on the machine on which the executables are running. If no `hostName.conf` is found, it looks in that directory for a file named `system.conf` and uses that. If neither of those files exist, the system defaults are used, unless the executable is Stone, in which case an error is generated.

If the `GEMSTONE_SYS_CONF` environment variable points to a file instead of a directory, GemStone just uses that file.

Within each file, if an option is listed more than once, then the value it is given the last time it is specified is used as its true value at executable run time. This rule also applies between the two types of configuration files. If the same option is given a value in both the system-wide and executable-dependent configuration files, the value in the executable-dependent configuration file overrides the system-wide configuration file's value.

Figure A.1 Search Path for a System-Wide Configuration File

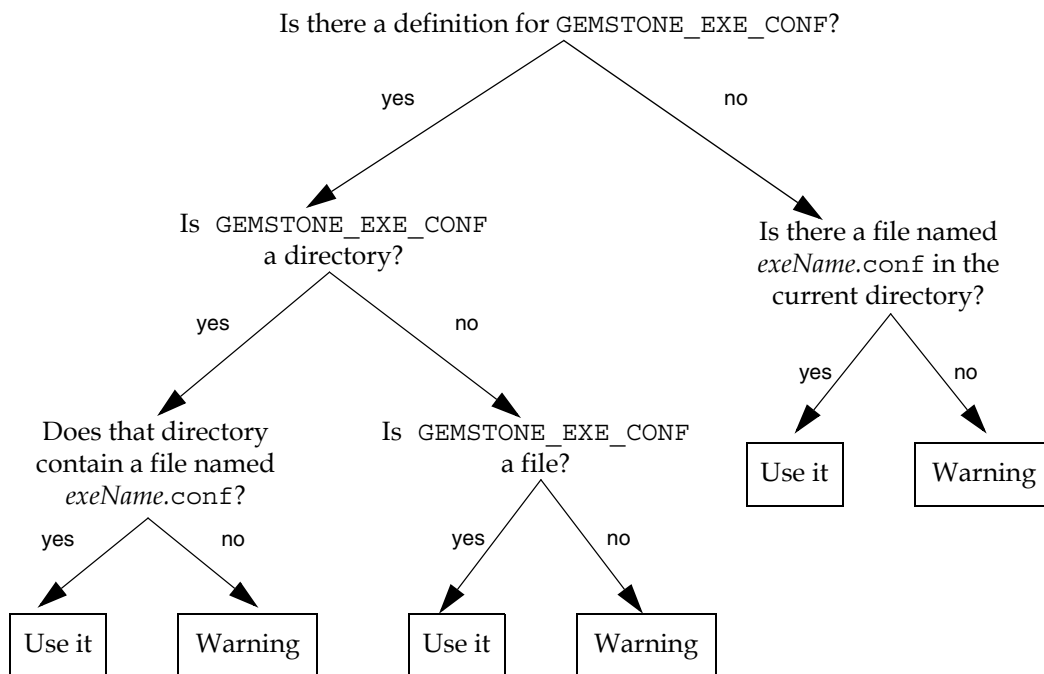


Search for an Executable Configuration File

Ordinarily, GemStone repository executables next try to find an executable-dependent configuration file. (The exception is a Stone repository monitor that failed to find its system-wide configuration file – it exits with an error.)

1. As shown in Figure A.2, GemStone begins by checking to see if there is an environment variable defined for GEMSTONE_EXE_CONF.
2. If GEMSTONE_EXE_CONF is not defined, GemStone tries to find a file called *exeName.conf* in the current working directory. (For information about the naming conventions, see “Naming Executable Configuration Files” on page 269.)
3. If it succeeds at finding such a file, it uses that file. If such a file does not exist, it generates a warning and relies solely on the system-wide configuration file for configuration parameters.

Figure A.2 Search Path for an Executable-Dependent Configuration File



If GEMSTONE_EXE_CONF is defined, GemStone first looks to see if it points to a directory.

- If GEMSTONE_EXE_CONF points to a directory, GemStone looks for a file named *exeName.conf* in that directory. If such a file exists, it uses it; if not, a warning is generated and GemStone relies on the system-wide configuration file for configuration parameters.
- If GEMSTONE_EXE_CONF points to a file, rather than to a directory, GemStone simply uses that file.

- If GEMSTONE_EXE_CONF points to a directory or file that doesn't exist, a warning is generated and GemStone defaults to using the system-wide configuration file for configuration parameters.

Creating or Using a System Configuration File

If you are satisfied with the standard options and the defaults, the simplest thing to do is to just use the configuration file provided in `$GEMSTONE/data/system.conf`. You can either copy this file and set the GEMSTONE_SYS_CONF environment variable to point to your new file, or you can do nothing and let GemStone use `$GEMSTONE/data/system.conf` itself.

Creating an Executable Configuration File

There are two ways to create a configuration file for a specific executable:

- You can copy the entire system-wide configuration file to a new file, name it appropriately, and change selected parameters.
- You can create a new file, give it an appropriate name, and include only those parameters that you want to differ from the default.

To make sure that GemStone is able to find and use your executable-dependent configuration file, you can set the GEMSTONE_EXE_CONF environment variable to point to your file. GEMSTONE_EXE_CONF can be either a file name or a directory name. If you set the environment variable to a directory name, be sure to name the configuration file `exeName.conf` so GemStone can find it at startup. (Information about the naming conventions for configuration files is just ahead.)

If you don't set the GEMSTONE_EXE_CONF environment variable, GemStone looks for a file named `exeName.conf` in the current working directory at startup. If it doesn't find one, it uses the configuration parameters set in the system-wide configuration file, or it uses the system defaults.

NOTE

Make sure your executable-dependent file is both readable and writable by the Stone process, which will update options by writing to it if you make certain configuration changes at run time.

Naming Executable Configuration Files

The default name of an executable configuration file generally is determined from the name of the executable itself.

Application Gems

Stand-alone (RPC) Gems look for a file named `gem.conf` in the current working directory unless GEMSTONE_EXE_CONF is defined. The working directory by default is the user's home directory, unless the `gemnetobject` script has been customized. The file `$GEMSTONE/sys/gemnetobject` is a script that a NetLDI invokes to start a GemStone session process. This script can be edited to define the name of the Gem to execute, the directory where the Gem resides, and the GEMSTONE_SYS_CONF and GEMSTONE_EXE_CONF environment variables.

System Gems

It is sometimes useful to change the parameters in a configuration file specific to a system Gem, such as the Admin or Reclaim Gem. This allows customized configuration settings that remain in effect if the system is stopped and restarted.

To do so:

Step 1. Copy `$GEMSTONE/data/system.conf`.

Step 2. Edit the copy, setting the values you want.

Step 3. Save your changes, renaming the file appropriately (for example, `admingcgem.conf`, `reclaimcgem.conf`, etc., depending on which system Gem the new configuration file is for). Place the file in GemStone's `sys` directory.

Step 4. Make a copy of the appropriate script/s.

<code>runadmingcgem</code>	Starts the Admin Gem.
<code>runcachewarmergem</code>	Starts the cache warmer Gems.
<code>runotcachewarmergem</code>	Starts the Object Table cache warmer Gems.
<code>runreclaimcgem</code>	Starts the Reclaim Gem.
<code>runsymbolgem</code>	Starts the SymbolGem.

These scripts are located in `$GEMSTONE/sys/`. Name the copy appropriately; for example, `$GEMSTONE/sys/myrunreclaimcgem`

Step 5. Edit `myrunreclaimcgem` to specify the customized configuration file.

For example, to use a Reclaim Gem configuration file named `$GEMSTONE/sys/reclaimcgem.conf`, locate the lines:

```
exeConfig=""
```

and change to:

```
exeConfig="$GEMSTONE/sys/reclaimcgem.conf"
```

Step 6. Edit the file `$GEMSTONE/sys/services.dat`. Comment out the existing line:

```
runreclaimcgem      $GEMSTONE/sys/runreclaimcgem
```

and add an entry specifying the new script to be executed for the `runreclaimcgem` service.

```
#runreclaimcgem    $GEMSTONE/sys/runreclaimcgem
runreclaimcgem     $GEMSTONE/sys/myrunreclaimcgem
```

Stone

Stone looks for a file named `stoneName.conf` in the current working directory.

Linked Topaz

The linked version of Topaz looks for the configuration file `gem.conf` so, by default, Gem and Topaz can share the same options. The default location of the file is the user's home directory (`$HOME`).

Linked GBS

Linked GBS logins by default look for a file named `gbs.conf`. The default location of the file is the user's home directory (`$HOME`).

Linkable GemBuilder for C Applications

Linkable GemBuilder for C applications look for a file named `gci.conf` in the current working directory unless the application has provided a different name by calling `GciInitAppName()`.

Naming Conventions for Configuration Options

The prefix "GEM_" indicates that the option is processed directly by Gems. Unless indicated otherwise by the phrase "used by all executables," most other options are processed only by the Stone, which passes the information to executables as needed through network connections. Exceptions are the shared page cache configuration options ("SHR_"). The first Gem session process on a node remote from the Stone and extents reads these options, which determine the configuration of the shared page cache on that node.

All executables (that is, the Stone and Gems) understand the standard options used in the file `$GEMSTONE/data/system.conf` as shipped. The GemStone executables generate a warning message whenever they encounter an option that is not in the standard list.

NOTE:

If the DUMP_OPTIONS option is set to true, then once the system-wide and executable-dependent configuration files have been processed, the values of all the options understood by the executable are displayed. You can access the configuration parameters from Smalltalk by using the methods described starting on page 50.

A.2 Configuration File Syntax

The following section describes the rules of grammar to be used in editing configuration files.

White space	Leading white space is ignored in the parsing of configuration files. Trailing white space is ignored if it follows the statement termination symbol (;).
New lines	New lines within a statement are allowed only after an equal sign or after a comma within a list of values.
Comments	The comment symbol for GemStone configuration files is the pound sign (#). Comments can be embedded in a configuration file using the following rules: <ul style="list-style-type: none"> • by starting a line with the comment symbol • by placing any text after the statement termination symbol (;)
Lists	Lists are separated by commas; list elements can be empty, for example:

```
DBF_EXTENT_SIZES = 1999, , 1999;
```

Within lists of values, leading and trailing white space is ignored.

Strings

Strings are encased in quotes. An empty string is acceptable in the grammar, and may be expressed by either two double quotes (""), or by no value at all (for instance, `OPTION = ;`).

Within strings, the escape character is the backslash (\). It can be used as follows:

<u>To generate:</u>	<u>Use the sequence:</u>
backslash (\)	\\
quote (")	\"
statement termination symbol (;)	\;
list separation character (,)	\,
control characters	\ followed by decimal representation of the character as a zero-padded 3-digit decimal number. For example, the string control-N would read \014, because control-N is ASCII 14.

Case-sensitivity String option values are case-sensitive; boolean option names are *not* case-sensitive.

Maximum Sizes The maximum number of characters allowed for a GemStone configuration option name is 64. The maximum length of a string option is 1024 characters. There is no limit on the number of elements within a list.

Use of Environment Variables In Options

Options that are either file names or directories may have environment variables as the first part of their value or the entire value. For instance, `$GEMSTONE/data/extent0.dbf`.

Errors in Configuration Files

At startup, each GemStone executable reads the configuration files. If any error is detected, information about the error is written to the standard output. This information indicates the file and line containing the error and the error's severity.

Two kinds of errors can be generated by the processing of configuration files: syntax errors and option value errors.

Syntax Errors

Syntax errors are generated whenever a grammatical error is detected in the configuration file. All syntax errors are warnings; they do not cause execution to terminate. These errors include:

- End-of-line or end-of-file detected before expected

- Invalid starting character for an option name or invalid character within an option name
- Equals or semicolon sign expected
- Invalid 3-digit escape sequence
- Invalid escape character
- Terminating quote missing in a quoted string

Option Value Errors

Option value errors are generated when the value assigned to an option has no meaning or is of the wrong type. For example, an option value error is generated when an option defined to need a boolean for its value has been set to an integer.

Option value errors vary in severity. Some options, such as not specifying the list of files that make up a logical repository, will necessarily terminate execution. Other option value errors, such as a invalid cache size, might only generate warnings. When a warning is issued, the executable ignores the given value and use the option's default value.

A.3 Configuration Options

The system-wide configuration file contains the following standard configuration options. In this discussion, *default* refers to the value that results when an option is not explicitly set by a statement in the configuration file. *Initial setting* refers to an explicit setting in the initial `system.conf` file that differs from the default.

Some configuration options have an internal runtime parameter that can be changed while GemStone is running. Where such a parameter exists, its name is given as part of the entry. For more information, see “To Change Settings at Run Time” on page 50.

The `$GEMSTONE/bin` directory contains a write-protected file named `initial.config` that is an exact replicate of `$GEMSTONE/data/system.conf` as it was originally shipped. After modifying `system.conf`, you can always recover its original condition.

DBF_ALLOCATION_MODE

`DBF_ALLOCATION_MODE` describes the space allocation heuristic to be used when filling repository extents.

Permissible values are either Sequential or a series of allocation weights, separated by commas. Under sequential allocation, each extent has its full resources used before the next extent’s resources are used. Under weighted allocation, those extents with a larger weight will have proportionally more of their disk resources allocated than those with smaller weights. Each weight applies to the corresponding extent in the series of extents specified in `DBF_EXTENT_NAMES`, and the number of elements must match. Extent allocation weights must be integers in the range 1..40 (inclusive).

Default: Sequential

DBF_EXTENT_NAMES

`DBF_EXTENT_NAMES` list of all repository extents, in order, primary extent first, separated by commas. Taken together, all of the listed file resources make up the logical repository. This option is required, and must contain at least one entry, the name of the primary extent. The maximum number of extents is 255.

An extent name can be a file name or the device name for a raw disk partition. The name can have an environment variable as its first component.

Default: EMPTY. The system will not run unless you define an extent list.

Initial setting: `$GEMSTONE/data/extent0.dbf`

DBF_EXTENT_SIZES

`DBF_EXTENT_SIZES` sets the maximum sizes of all repository extents, in order, primary extent first, separated by commas. Each size applies to the corresponding extent in the series of extents specified in `DBF_EXTENT_NAMES`.

A size entry may be empty, which indicates that the corresponding extent has no fixed maximum size. This setting allows the extent to grow until it fills the disk containing it.

When an extent is on a raw partition, for optimal performance the corresponding setting in `DBF_EXTENT_SIZES` should be slightly smaller than the size of the partition so GemStone

can avoid having to handle system errors. For example, set it to about 1995 MB for a 2 GB partition.

You can modify the size of an existing extent under these conditions:

- If the original maximum size was unlimited, the new maximum size must be larger than the current physical size of the extent.
- If the original maximum size was limited, the new maximum size must be larger than the original maximum size.

The Stone repository monitor is the only executable allowed to change `DBF_EXTENT_SIZES`. At GemStone system startup, the maximum size of each extent is written to the system log.

If no units are specified, the value is in MB (1 Megabyte = 1048576 bytes). You may also specify units as KB, MB, or GB.

Default Units: MB

Min: 1 (MB)

Max: 33554432 (subject to disk, operating system and platform limits)

Default: EMPTY (no maximum sizes)

DBF_PRE_GROW

If `DBF_PRE_GROW` is set to `TRUE` and there are extents for which a size is specified in `DBF_EXTENT_SIZES`, then on repository startup, each extent with a size in `DBF_EXTENT_SIZES` larger than the current size will be pregrown to the specified size. If the grow fails, the extent is reset to its original size and startup fails.

If `DBF_PRE_GROW` is set to `TRUE` and a new extent is added programmatically with a size specified, it will be pregrown to that size. If the extent cannot be grown to the maximum size because of disk capacity problems, then extent creation will fail.

The default value for `DBF_PRE_GROW` is `FALSE`. This setting indicates that extents will grow only when new space is needed. An extent without a maximum size is never pregrown.

The value of `DBF_PRE_GROW` may also be a list of integer sizes, which may include blanks for specific extents that will not be pregrown. Extents with an integer size specified in `DBF_PRE_GROW` will be pregrown to this size if needed. Extents without an entry in `DBF_PRE_GROW` will not be pregrown.

It is an error if the value for an extent's `DBF_PRE_GROW` size is larger than the corresponding `DBF_EXTENT_SIZES` size; if one but not both are empty, it is not an error and the extent will not be pregrown.

Elements of `DBF_PRE_GROW` may be blank to specify pregrow sizes for some but not all extents, such as:

```
DBF_PRE_GROW = 1000, , 1000;
```

If no units are specified, the value is in MB (1 Megabyte = 1048576 bytes). You may also specify units as KB, MB, or GB.

Default Units: MB

Min: 1

Max: 33554432 (limited by `DBF_EXTENT_SIZES` values)

Default: `FALSE`

DBF_SCRATCH_DIR

DBF_SCRATCH_DIR specifies a scratch directory that the Stone process can use to create “scratch” repositories for use during **pageaudit**. The file name is appended to the directory name *without* an intervening delimiter, so a trailing delimiter is necessary here.

Default: \$GEMSTONE/data/

DUMP_OPTIONS

If DUMP_OPTIONS is set to true, dumps a summary of all configuration options as part of the process log file headers.

Default: true

GEM_ABORT_MAX_CRIS

When a Gem is not in a transaction and aborts, GEM_ABORT_MAX_CRIS specifies the maximum number of commit records to analyze to compute the writeSetUnion since the last time this session aborted. If the number of commit records would exceed this limit, the abort is treated similar to a LostOt and all in-memory copies of committed objects are marked invalid and will be re-read as needed during subsequent execution.

A value of 0 (zero) means no limit on number of commit records to analyze.

Runtime parameter: #GemAbortMaxCris

Min: 0

Max: 2147483647

Default: 0

GEM_FREE_FRAME_CACHE_SIZE

GEM_FREE_FRAME_CACHE_SIZE specifies the size of the Gem’s free frame cache. When using the free frame cache, the Gem removes enough frames from the free frame list to refill the cache in a single operation. When adding frames to the free list, the Gem does not add them until the cache is full.

A value of 0 disables the free frame cache (the Gem acquires frames one at a time). A value of -1 means use the default value: 0 for caches less than 100 MB and 10 for caches of 100 MB or greater.

Cache Statistic: FreeFrameCacheSize (Gem)

Units: frames

Min: -1

Max: 63

Default: -1 (cache size=0 for caches less than 100 MB and 10 for caches of 100 MB or greater)

GEM_FREE_FRAME_LIMIT

When the number of free frames in the shared page cache is less than GEM_FREE_FRAME_LIMIT, the Gem session process scans the cache for a free frame rather than using one from the free frame list. This action is desirable for performance reasons so the remaining frames in the list are available for use by the Stone repository monitor.

If the value of `GEM_FREE_FRAME_LIMIT` is `-1`, the free frame limit is set to one of the following default values:

- For primary shared page cache that is 800 MB or smaller: 10% of the number of frames in the cache
- For primary shared page cache greater than 800 MB: 5000 frames
- For a remote shared page cache: 0

Runtime parameter: **#GemFreeFrameLimit**

Cache Statistic: FreeFrameLimit (Gem)

Min: 1

Max: 65536

Default: `-1` (see above discussion)

GEM_FREE_PAGEIDS_CACHE

`GEM_FREE_PAGEIDS_CACHE` specifies the maximum number of free pageIds to be cached in Gem. Larger values reduce number of calls to Stone, at a cost of needing more free space within the extents.

Runtime parameter: **#GemFreePageIdsCache**

Min: 40

Max: 3500

Default: 200

GEM_GCI_LOG_ENABLED

This option has no effect in customer executables.

Default: `false`

GEM_HALT_ON_ERROR

`GEM_HALT_ON_ERROR` causes a Gem to halt and dump core if an error with the specified GemStone error number occurs. The value `0` means “never halt”. Ordinarily this option is used only to assist Technical Support in diagnosing problems.

Runtime parameter: **#GemHaltOnError**

Default: `0`

GEM_KEEP_MIN_SOFTREFS

`GEM_KEEP_MIN_SOFTREFS` determines the minimum number of most recently used SoftReferences that will not be cleared by VM markSweep if *startingMemUsed* – the percentage of temporary object memory in use at the beginning of a VM mark/sweep – is greater than `GEM_SOFTREF_CLEANUP_PERCENT_MEM` but less than 80%.

In most cases, the default (`0`) is appropriate and should not be changed.

Runtime parameter: **#GemKeepMinSoftRefs**

Min: `0`

Max: 10000000

Default: `0`

GEM_MAX_SMALLTALK_STACK_DEPTH

GEM_MAX_SMALLTALK_STACK_DEPTH determines the size of the GemStone Smalltalk execution stack space that is allocated when the Gem logs in. The unit is the approximate number of method activations in the stack. This setting causes heap memory allocation of approximately 64 bytes per activation. Exceeding the stack depth results in generation of the error RT_ERR_STACK_LIMIT.

Min: 100
Max: 1000000
Default: 1000

GEM_NATIVE_CODE_ENABLED

GEM_NATIVE_CODE_ENABLED enables or disables generation of native code. This is set to an integer 0, 1, or 2. For compatibility with config files from earlier versions, it may be also set to TRUE or FALSE.

Breakpoints in methods disable native code. Also, a session with a very large GEM_TEMPOBJ_CACHE_SIZE, on the Mac, may disable native code for internal reasons.

The runtime parameter #GemNativeCodeEnabled can be used to disable native code, and can be used to control whether subsequent GsProcesses start execution using interpreted or native code. Enabling generation of native code for methods as they are loaded for execution can only be controlled by the value of GEM_NATIVE_CODE_ENABLED in the configuration files at process startup.

- 0 or FALSE disables native code generation
- 1 or TRUE enables native code generation
- 2 enables native code generation, with inlining of some SmallInteger math primitives.

Runtime parameter: #GemNativeCodeEnabled
Minimum: 0
Maximum: 2
Default: 2

GEM_PGSRV_COMPRESS_PAGE_TRANSFERS

If GEM_PGSRV_COMPRESS_PAGE_TRANSFERS is true, use compress2() from zlib library with default compression level to compress page transfers between the page server on Stone's machine and Gem or mid-cache page server.

For the first Gem to login on a remote machine, that Gem's configuration file value of GEM_PGSRV_COMPRESS_PAGE_TRANSFERS is propagated to the page manager, and is used to configure the page manager's communication to the page manager's pgsvr on the new remote cache.

When a Gem triggers creation of a mid-level cache via the method **midLevelCacheConnect:cacheSizeKB:maxSessions:**, that Gem's current runtime value of GEM_PGSRV_COMPRESS_PAGE_TRANSFERS is propagated to the page manager, and is used to configure the page manager's communication to the page manager's pgsvr on the new mid-level cache.

Runtime parameter: #GemPgsvrCompressPageTransfers
Default: FALSE

GEM_PGSRV_FREE_FRAME_CACHE_SIZE

GEM_PGSRV_FREE_FRAME_CACHE_SIZE specifies the size of the free frame cache used by the Gem's remote page server. This configuration option has no effect for Gems that are local to the repository extents (which have a page server).

When using the free frame cache, the page server removes enough frames from the free frame list to refill the cache in a single operation. When adding frames to the free list, the page server does not add them until the cache is full.

A value of 0 disables the free frame cache (the page server acquires frames one at a time). A value of -1 means use the default value: 0 for caches less than 100 MB and 10 for caches of 100 MB or greater.

Cache Statistic: FreeFrameCacheSize (Page Server)

Units: frames

Min: -1

Max: 63

Default: -1 (cache size=0 for caches less than 100 MB and 10 for caches of 100 MB or greater)

GEM_PGSRV_FREE_FRAME_LIMIT

GEM_PGSRV_FREE_FRAME_LIMIT determines the free frame limit used by the Gem's remote page server. It has no effect for Gems local to the repository extents (which do not have a page server). For a description of free frames, see the GEM_FREE_FRAME_LIMIT configuration option (page 276).

If the value of GEM_PGSRV_FREE_FRAME_LIMIT is -1, the free frame limit is set to one of the following default values:

- For primary shared page cache that is 800 MB or smaller: 10% of the number of frames in the cache
- For primary shared page cache greater than 800 MB: 5000 frames

To tune the free frame limit of a page server at runtime, use the method `System class>>changeCacheSlotFreeFrameLimit: aSlot to: aValue`.

Cache Statistic: (Page Server) FreeFrameLimit

Min: -1

Max: 65536

Default: -1 (see above discussion)

GEM_PGSRV_UPDATE_CACHE_ON_READ

GEM_PGSRV_UPDATE_CACHE_ON_READ determines the read behavior of the Gem's remote page server when pages are read from disk. If this option is set to true, pages read from disk are also added to the shared page cache on the page server's host. If this option is false, pages read are not added to the page server's shared cache.

This option has no effect for Gems that are local to the repository extents, which do not have page servers, nor on mid-level caches.

Runtime parameter: `#GemPgsvrUpdateCacheOnRead`

Default: false

GEM_PRIVATE_PAGE_CACHE_KB

GEM_PRIVATE_PAGE_CACHE_KB sets the size of the Gem's private page cache. This setting also applies to linked Gems.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Default units: KB
Min: 128
Max: 524288
Default: 1000

GEM_REPOSITORY_IN_MEMORY

Determines the performance behavior of the gem for certain operations that scan the entire repository. If set to TRUE, the gem assumes most or all of the data pages in the repository have been previously loaded into the shared page cache. If set to FALSE, the gem assumes most or all of the data pages in the repository are not in the shared page cache and must be read from disk.

This setting affects performance only. All operations affected by this setting will succeed and produce the same results.

Repository instance methods affected by this setting are:

findReferencePathToObject: (and related methods)
findAllReferencePathsToObjects: (and related methods)
pagesWithPercentFree: (and related methods)

Runtime parameter: **#GemRepositoryInMemory**
Default: FALSE

GEM_RPCGCI_TIMEOUT

GEM_RPCGCI_TIMEOUT specifies the time in minutes after which lack of an Rpc command will cause a Gem to terminate. Negative timeouts are not allowed. Resolution of timeouts is one-half the specified timeout interval.

Min: 0
Default: 0 (Gem waits forever)

GEM_RPC_KEEPALIVE_INTERVAL

GEM_RPC_KEEPALIVE_INTERVAL is the Interval in seconds for the RPC GCI client to send a packet to the gem to ensure the network connection is kept alive.

With the traditional GCI interface (gci.hf), this controls how often keep-alive packets are sent during the GciPollForSignal() calls. The application needs to be calling GciPollForSignal at regular intervals at least as often as the configured value.

In the thread safe GCI (gcits.hf), the gem will send an interrupt byte periodically, and the application must be calling GciTsWaitForEvent at least as often as this config value.

Min: 0
Max: 7200
Default: 0 (disabled)

GEM_RPC_USE_SSL

GEM_RPC_USE_SSL controls whether a remote RPC gem uses a secure socket layer (SSL) connection to converse with its RPC client. RPC sessions always establish a secure connection during the login sequence. This parameter controls whether the gem and its remote RPC client continue using the SSL connection. Otherwise, a standard TCP/IP socket connection is used.

This option has no effect for linked gems and local RPC gems (i.e., a gem running on the same host its client). Local gems always revert to a standard TCP/IP socket after login.

Secure sockets are slightly slower than insecure sockets due to the overhead of encrypting and decrypting data.

Default: TRUE

GEM_SOFTREF_CLEANUP_PERCENT_MEM

GEM_SOFTREF_CLEANUP_PERCENT_MEM controls the cleanup of SoftReferences.

If *startingMemUsed* – the percentage of temporary object memory in-use at the beginning of a VM mark/sweep – is less than the value of this option, no SoftReferences will be cleared.

If *startingMemUsed* is greater than the value of this option and less than 80%, the VM mark/sweep will attempt to clear an internally determined number of least recently used SoftReferences. Under rare circumstances, you might choose to specify a minimum number (GEM_KEEP_MIN_SOFTREFS) that will not be cleared.

If *startingMemUsed* is greater than 80%, VM mark/sweep will attempt to clear all SoftReferences.

Also see the statistics NumSoftRefsCleared, NumLiveSoftRefs, and NumNonNilSoftRefs.

Runtime parameter: **#GemSoftRefCleanupPercentMem**

Min: 10

Max: 80

Default: 50

GEM_TEMPOBJ_AGGRESSIVE_STUBBING

GEM_TEMPOBJ_AGGRESSIVE_STUBBING controls stubbing in in-memory garbage collection. If instance variable X in object A references object B, and X contains a memory pointer to B, then the reference is *stubbed* by storing into instance variable X the objectId of object B.

When this option is TRUE (the default), references from temporary objects to in-memory copies of committed objects are stubbed whenever possible, during both scavenge and mark/sweep. Also, references from not-dirty in-memory copies of committed objects to other committed objects are stubbed whenever possible. This reduces the number of committed objects forced to stay in-memory, but can slow down garbage collection and subsequent execution.

When this option is FALSE, references from temporary objects to in-memory copies of committed objects are never stubbed. References from not-dirty in-memory copies of committed objects to other committed objects are stubbed after the number of objects

flushed during commits reaches a threshold, or if almost OutOfMemory. Performance may be faster, but there is a greater risk of OutOfMemory errors.

Stubbing is always disabled when a commit attempt is in progress, regardless of the setting of this parameter. Certain objects private to the object manager are always immune from stubbing, and so are references stored into Session State by using `System class >> _sessionStateAt:put:.`

Also see the statistics NumRefsStubbedMarkSweep and NumRefsStubbedScavenge.

Default: true

GEM_TEMPOBJ_CACHE_SIZE

GEM_TEMPOBJ_CACHE_SIZE sets the maximum size of the Gem's temporary object memory. This limit also applies to memory in linked Topaz sessions and linked GemBuilder applications. This value is set when the VM is initialized and cannot be changed without restarting the VM. When you only change this setting, and the other GEM_TEMPOBJ* configuration options use default values, then all of the various spaces remain in proportion to each other.

This setting defines the maximum memory size. The initial memory allocated will be smaller, and as the actual space required for objects grows, the VM requests and allocates virtual memory as needed. As the limit is approached, in-memory garbage collection becomes more aggressive; if the limit is reached, the Gem will exit.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Statistic: (Gem) GemTempObjCacheSizeKb

DefaultUnits: KB

Min: 2000

Max: 64000000

Default: 50000

GEM_TEMPOBJ_MESPACE_SIZE

GEM_TEMPOBJ_MESPACE_SIZE sets the maximum size of the Map Entries space within the Gem's temporary object memory. This value is set when the VM is initialized and cannot be changed without starting the VM.

One Map Entry is required for each faulted-in committed object, or for any temporary object that might become committed, referenced from an IdentityBag, or exported to the GCI. One Map Entry occupies approximately 24 bytes.

If a Map Entry is needed and the Map Entries Space is full, an OutOfMemory occurs, terminating the session.

Unless you are trying to minimize the memory footprint on HP-UX or AIX, you should always leave GEM_TEMPOBJ_MESPACE_SIZE at its default value (0) so that the system can calculate the size of the Map Entries space based on other memory sizes. Otherwise, you are at risk of premature OutOfMemory errors.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Default units: KB
Min: 1000
Max: 1000000
Default: 0

GEM_TEMPOBJ_OOPMAP_SIZE

GEM_TEMPOBJ_OOPMAP_SIZE sets the size of the hash table (that is, the number of 8-byte entries) in the objId-to-object map within the Gem's temporary object memory. This value is set when the VM is initialized and cannot be changed without starting the VM.

The value specified is rounded up to the next higher power of 2.

This option should normally be left at its default value (0) so that the system can calculate the size of the map based on other memory sizes.

Min: 16384
Max: 524288000
Default: 0

GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE

GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE sets the percent of POM generation area to be thrown away when voting on possible dead objects.

If the value is 0, no subspaces of POM generation are cleared; if the value is 100, all subspaces are cleared. For values greater than 0 and less than 100, the number of spaces that are in use and older than 5 minutes is computed, and the specified parameter is the percentage, rounded down, of these subspaces that are discarded.

Runtime parameter: **#GemPomGenPruneOnVote**
Default: 50
Min: 0
Max: 100

GEM_TEMPOBJ_POMGEN_SCAVENGE_INTERVAL

GEM_TEMPOBJ_POMGEN_SCAVENGE_INTERVAL is the interval in seconds in which the oldest POM generation subspace will be discarded. Lower values may reduce Gem memory usage but may also cause objects to be re-read. Larger values may result in higher Gem memory usage and may reduce disk I/O. Setting this value to zero disables scheduled POM generation scavenges. In this case, POM generation will only be scavenged when all subspaces become full.

Runtime parameter: **#GemTempObjPomgenScavengeInterval**
Units: seconds
Min: 0
Max: 86400
Default: 1800

GEM_TEMPOBJ_POMGEN_SIZE

GEM_TEMPOBJ_POMGEN_SIZE sets the maximum size of the POM generation area in the Gem's temporary object memory. This value is set when the VM is initialized and cannot be changed without starting the VM.

The POM generation area holds unmodified copies of committed objects that have been faulted into a Gem, and is divided into ten subspaces.

This option should normally be left at its default value (0) so that the POM generation area is allocated to the default, which is approximately 0.8 times the size of the GEM_TEMPOBJ_CACHE_SIZE.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Default units: KB
Min: 1000
Max: 1000000
Default: 0

GEM_TEMPOBJ_SCOPES_SIZE

GEM_TEMPOBJ_SCOPES_SIZE is the size of the scopes stack in Gem temporary object garbage collection. This value is set when the VM is initialized and cannot be changed without starting the VM.

The scopes stack consumes (8 bytes * GEM_TEMPOBJ_SCOPES_SIZE) of C heap memory.

The primary user-visible effect of this setting is maximum depth of nested expressions that can be compiled by the method compiler. The default setting is sufficient for expression nesting of about 200, such as in depth of nested parenthesized expressions.

Min: 1000
Max: 10000000
Default: 2000

GEM_TEMPOBJ_START_ADDR

GEM_TEMPOBJ_START_ADDR applies for AIX only. If the default mmap of temp obj memory fails, this value is used to define the starting address at which to attempt to mmap temporary object memory using MAP_FIXED at fixed addresses and munmap to simulate MAP_NORESERVE.

A config file value of zero results in an internal default of 0xA000000000000000 for AIX 7, and 0x7000000000000000 for AIX 6. A non-default value must be coded as an exact address and may be affected by use of mmap by other shared libraries.

Default: 0

KEYFILE

KEYFILE sets the location of GemStone licensing keyfile.

Default: \$GEMSTONE/sys/gemstone.key

LOG_WARNINGS

If LOG_WARNINGS is set to true, warnings are printed for invalid configuration options.

Default: true

SHR_NUM_FREE_FRAME_SERVERS

SHR_NUM_FREE_FRAME_SERVERS specifies the number of free frame page server processes that will be started when the shared page cache is created. A value of -1 means the default value should be used. On the primary shared page cache (the cache to which the stone attaches), the default value is equal to the value used for the STN_NUM_LOCAL_AIO_SERVERS parameter. On a remote shared page cache, the default is 1.

Min: -1

Max: 255

Default: -1

SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_POLICY

Specifies whether large memory pages will be used when creating the shared page cache. Enabling large memory pages can result in significant performance gains when large shared page caches are used. The improvement is due to a reduction in translation lookaside buffer (TLB) cache misses. The TLB is an internal structure used by the operating system to manage memory address translation.

Large memory page support is an operating system and hardware dependent feature. Currently, GemStone supports large memory pages on AIX and Linux only. This configuration option is silently ignored on all other platforms.

Three policies are available on supported operating systems:

0 - Disabled: No large memory page support.

1 - Advisory: Large memory pages are requested when the shared page cache is created. If the operating system denies the request, a warning is printed in the SPC monitor log file and the cache is started without large memory pages. NOT RECOMMENDED ON AIX!

2 - Mandatory: Large memory pages are requested when the shared page cache is created. If the operating system denies the request, an error is printed in the SPC monitor log file and the shared page cache fails to start.

Both Linux and AIX require operating system kernel changes in order to enable large memory pages. Refer to the release notes for more information.

Default: 0

Minimum: 0

Maximum: 2

SHR_PAGE_CACHE_LOCKED

SHR_PAGE_CACHE_LOCKED specifies whether the shared page cache should be locked in memory. On systems that permit a portion of memory to be dedicated to GemStone, this option may provide higher performance.

On Solaris 10, GemStone uses Intimate Shared Memory for the shared page cache, making setting this variable unnecessary.

Other specific operating systems may restrict this action to processes running as root or may require special privileges (such as MLOCK on HP-UX) for this option to take effect; for further information, check the shared page cache monitor log for error messages and consult your operating system documentation.

Default: false

SHR_PAGE_CACHE_NUM_PROCS

SHR_PAGE_CACHE_NUM_PROCS sets the maximum number of processes allowed to attach to the shared page cache. This parameter is used to allocate space in the shared page cache for session information and cache statistics. This cache space is in addition to extent page space allocated by SHR_PAGE_CACHE_SIZE_KB.

The value for SHR_PAGE_CACHE_NUM_PROCS must accommodate the GcGems and various background GemStone processes, as well as user Gem and Topaz session processes. If the value is too small, sessions might be unable to login because they can't attach to the cache. If the value is too large, space in the cache may be wasted.

It is recommended to leave this at the default value. When the default setting of -1 is specified, the value of this parameter is calculated as:

```
STN_MAX_SESSIONS
+ 8 (for system logins)
+ STN_MAX_GC_RECLAIM_SESSIONS
+ SHR_NUM_FREE_FRAME_SERVERS
+ STN_NUM_LOCAL_AIO_SERVERS
```

Cache Statistic: (SPC Monitor) SlotsTotalCount
Min: 15, or the number extents + 3, whichever is larger
Max: determined by STN_MAX_SESSIONS or file descriptor limits
Default: -1

SHR_PAGE_CACHE_NUM_SHARED_COUNTERS

SHR_PAGE_CACHE_NUM_SHARED_COUNTERS specifies the number of shared counters available in the shared page cache. On most platforms, each counter consumes 128 bytes of shared memory. On AIX, each counter consumes 256 bytes of shared memory. Shared memory used for shared counters is in addition to the shared memory size specified in SHR_PAGE_CACHE_SIZE_KB.

Cache Statistic: (SPC Monitor) NumSharedCounters
Min: 0
Max: 500000
Default: 1900

SHR_PAGE_CACHE_PERMISSIONS

SHR_PAGE_CACHE_PERMISSIONS specifies the UNIX permission settings of the shared page cache, expressed as an octal number. The first two digits are constant and must always be 06. The 0 indicates an octal constant and 6 indicates the UNIX user which created the cache has read/write permissions.

The last two digits specifies the group and other permissions respectively. Each of the last two digits must be one of the following:

- 6 - read/write
- 4 - read only
- 0 - no access

By default, the shared page cache is created with group read/write permission but no access for other users.

Default: 0660

SHR_PAGE_CACHE_SIZE_KB

SHR_PAGE_CACHE_SIZE_KB sets the base size of the shared page cache. Additional shared memory is used for overhead, so the actual size of the memory segment will be somewhat larger.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Default Units: KB
Min: 16000
Max: Limited by system memory and kernel configurations
Default: 75000

NOTE

For information about platform-specific limitations on the size of the shared page cache, refer to Chapter 1 of your GemStone/S Installation Guide.

SHR_SPIN_LOCK_COUNT

SHR_SPIN_LOCK_COUNT specifies the number of tries to get a spin lock before the process sleeps on a semaphore. Semaphores involve a relatively time-consuming call to the operating system. Spin locks involve busy-wait loops. Efficient locking may require a combination of these methods.

In single-processor architectures, this value should always be 1 since there is no value in spinning (the lock won't change until the process holding the lock gets scheduled). On multiple-processor architectures, a value of 5000 is recommended.

We recommend that you leave this option set to the default value of -1, which causes GemStone to use a value of either 1 or 5000, based upon the number of CPUs detected.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: **#SpinLockCount**
Default: -1 (use either 1 or 5000, based on the number of CPUs detected)

SHR_TARGET_FREE_FRAME_COUNT

SHR_TARGET_FREE_FRAME_COUNT specifies the target number of free frames to keep in the shared cache at all times. The free frame page server process(es) will attempt to keep the number of free frames in the cache equal to or greater than this value.

If the value of the parameter is -1, the target free frame count is set to a percentage of the total frames in the shared cache. For the main shared cache (the cache to which the stone attaches), the default is 1/8 the frames in the cache. For remote caches, the default is 1/100 the frames in the cache.

For best performance, keep this setting greater than GEM_FREE_FRAME_LIMIT.

If the value of SHR_TARGET_FREE_FRAME_COUNT is -1, the target number of free frames is set to one of the following default values:

- For primary shared page cache that is 800 MB or smaller: 12.5% of the number of frames in the cache
- For primary shared page cache greater than 800 MB: 7000 frames
- For a remote shared page cache: 1% of the number of frames, or 2000 frames, whichever is smaller

Min: -1

Max: 65536

Default: -1 (see above discussion)

SHR_WELL_KNOWN_PORT_NUMBER

SHR_WELL_KNOWN_PORT_NUMBER specifies the port number that the shared page cache monitor will use as its well-known port. The well-known port is used by all Gems and page servers on this host to connect to the cache monitor.

If the specified port is in use by another process, the monitor process will not start and exits with an error.

A value of zero indicates that the port number will be selected by the system.

Min: 1

Max: 65535

Default: 0

STN_ADMIN_GC_SESSION_ENABLED

STN_ADMIN_GC_SESSION_ENABLED determines whether the Admin Gem is started when the Stone is started. (The Admin Gem performs administrative garbage collection functions such as write set union sweeps; the Reclaim Gem performs dead object and page reclaim.)

Runtime parameter: **#StnAdminGcSessionEnabled**

Default: true

STN_ALLOCATE_HIGH_OOPS

If true, the Stone skips the first 16 million object identifiers and begins to allocate object identifiers (GCI OopTypes) for non-special objects at 16r100000001.

This option is designed for testing conversion of GCI applications and user actions. Do not set this option in a production environment.

Default: FALSE

STN_ALLOW_NFS_EXTENTS

If TRUE, stone will startup using extents and tranlogs which are on NFS-mounted filesystems. This is less reliable and less performant than locally mounted filesystems, or filesystems on storage arrays which appear as local mounts. This variable cannot be changed at runtime.

Default: FALSE

STN_CACHE_WARMER

Specifies if the cache warmer should be run when the Stone is started and whether to load just the object table pages or both the object table and the data pages.

STN_CACHE_WARMER has the following possible values:

- 0 - Disabled, the stone does not start the cache warmer. This is the default.
- 1 - Start the cache warmer and load only the object table pages.
- 2 - Start the cache warmer and load the object table and data pages.

Default: 0

STN_CACHE_WARMER_SESSIONS

Specifies the number of worker sessions (threads) to use by the cache warmer Gem to perform cache warming on startup. In addition to the specified number of threads, there is one additional "master" session allocated. The warmer will exit with an error if not enough sessions are available.

Cache Statistic: NumCacheWarmers (Stone)

Units: sessions

Min: 0

Max: 256

Default: 0 - the warmer is run with numberOfCpus + numberOfExtents sessions

STN_CHECKPOINT_INTERVAL

STN_CHECKPOINT_INTERVAL sets the maximum interval between checkpoints. Checkpoints may be written more often, depending on other factors. The unit is seconds.

This can be changed at runtime only by SystemUser.

Runtime parameter: **#StnCheckpointInterval**

Units: seconds

Min: 5

Max: 1800

Default: 300

STN_COMMIT_QUEUE_THRESHOLD

STN_COMMIT_QUEUE_THRESHOLD determines whether the Stone defers the disposal of commit records, based on the number of sessions in the commit queue and the run queue. If the size of either of these queues exceeds this threshold, the Stone defers commit record disposal until all queues have sizes less than or equal to the value.

This setting is ignored if the commit record backlog exceeds the value of STN_CR_BACKLOG_THRESHOLD.

Runtime parameter: **#StnCommitQueueThreshold**
Default: -1 (never defer commit record disposal)
Min: -1
Max: 1024

STN_COMMIT_RECORD_QUEUE_SIZE

STN_COMMIT_RECORD_QUEUE_SIZE determines the size of the Stone's internal commit record cache. The Stone will keep copies of up to this many commit records in heap memory. Stone is able to dispose commit records more quickly when a copy of the commit record is found in this cache.

When the default value of -1 is specified, Stone sets this value to be twice the value of the STN_SIGNAL_ABORT_CR_BACKLOG option.

Units: Commit Record Pages
Default: -1
Min: 16
Max: 1000000

STN_COMMIT_TOKEN_TIMEOUT

STN_COMMIT_TOKEN_TIMEOUT sets the maximum interval (in seconds) that a session may possess the commit token. If the session possesses the token for longer than this period, the session will be logged off the system and an error message written to the Stone log. GcGems of all types are exempted from this timeout.

Default: 0 (Stone waits forever)
Min: 0
Max: 86400

STN_COMMITS_ASYNC

If STN_COMMITS_ASYNC is set to TRUE, it causes the stone to acknowledge each commit or persistent shared counter update to the requesting session without waiting for the tranlog writes for that commit to complete.

Default: FALSE

STN_CR_BACKLOG_THRESHOLD

STN_CR_BACKLOG_THRESHOLD sets the size of the commit record backlog above which the Stone aggressively disposes of commit records. This setting overrides the deferral of commit record disposal provided by the STN_COMMIT_QUEUE_THRESHOLD parameter.

The default setting of -1 causes the Stone to use a setting equal to $(2 * STN_MAX_SESSIONS)$. A setting of 0 disables this threshold.

Runtime parameter: **#StnCrBacklogThreshold**
Default: -1
Min: -1
Max: 500000

STN_DISABLE_LOGIN_FAILURE_LIMIT STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT

These options control when a user account is disabled because the user exceeded the STN_DISABLE_LOGIN_FAILURE_LIMIT of failed login attempts within the time in minutes specified by STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT. When an account exceeds these limits, the user account is disabled (the system changes the password on the account to one that is invalid) and a record of the event is written to the Stone log file. The user account can only be restored by another user with OtherPassword privileges.

Changes to the runtime parameters require the OtherPassword privilege.

STN_DISABLE_LOGIN_FAILURE_LIMIT:
Runtime parameter: **#StnDisableLoginFailureLimit**
Units: login attempts
Default: 15
Min: 0
Max: 65536

STN_LOG_LOGIN_FAILURE_TIME_LIMIT:
Runtime parameter: **#StnDisableLoginFailureTimeLimit**
Units: Minutes
Default: 15
Min: 1
Max: 1440 (24 hours)

STN_DISKFULL_TERMINATION_INTERVAL

STN_DISKFULL_TERMINATION_INTERVAL specifies how soon (in minutes) the Stone should start terminating sessions holding on to the oldest commit record when the repository free space is below the value set for STN_FREE_SPACE_THRESHOLD. Such sessions are sent the fatal diskfull error.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: **#StnDiskFullTerminationInterval**

Units: Minutes

Min: 0 (no sessions are terminated)

Max: 1440 (24 hours)

Default: 3

STN_EPOCH_GC_ENABLED

STN_EPOCH_GC_ENABLED determines if epoch garbage collection can be run on the system. Leave this value set to FALSE unless you plan to run epoch garbage collection on the system. Setting this to TRUE adds a small amount of overhead to commit processing.

Runtime parameter: **#StnEpochGcEnabled**

Default: FALSE

STN_EXTENT_IO_FLAGS

STN_EXTENT_IO_FLAGS specifies what (if any) special I/O flags will be used to open the database extents. Two kinds of special I/O are supported: direct I/O and concurrent I/O.

Direct I/O tells the operating system avoid caching extent data in the file system cache. Enabling direct I/O tells the operating system to treat the database extents as if they were on raw partitions. Direct I/O may greatly improve database performance in some cases. Concurrent I/O is only available to extents running on the Enhanced JFS file system (aka JFS2) on AIX. Setting this flag has no effect on other operating systems.

STN_EXTENT_IO_FLAGS has the following possible values:

0 - no special I/O flags. This is the default.

1 - enable Direct I/O on all extents on file systems.

2 - enable concurrent I/O (AIX only)

If the requested I/O mode is not available, the stone will fail to start and an error message will be printed in the stone log. If this happens, change this option back to zero and restart the stone.

STN_EXTENT_IO_FLAGS has no effect on extents which reside on raw partitions.

Once the stone starts, all processes which open the database extents (gems and page servers) will open the extents using the same I/O flags. This behavior is required by some operating systems.

Default: 0

Min: 0

Max: 2 (AIX only), 1 (All Others)

STN_FREE_FRAME_CACHE_SIZE

STN_FREE_FRAME_CACHE_SIZE specifies the size of the Stone's free frame cache. When using the free frame cache, the Stone removes enough frames from the free frame list to refill the cache in a single operation.

Units: frames

Default: 1 (disables the free frame cache; Stone acquires frames one at a time)

Min: 1

Max: 1% of the frames in the cache

STN_FREE_SPACE_THRESHOLD

STN_FREE_SPACE_THRESHOLD sets the minimum amount of free space to be available in the repository. If the Stone cannot maintain this level by growing an extent, it begins to take action to prevent the shutdown of the system. If the amount of free space remains below this level for more than the number of minutes specified by STN_DISKFULL_TERMINATION_INTERVAL, the stone will start terminating sessions. For more information, see "Repository Full" on page 179.

The default value of 0 specifies a varying STN_FREE_SPACE_THRESHOLD that is computed when needed as the current size of the repository divided by 1000, with a minimum value of 5 MB.

If no units are specified, the value is in MB. You may also specify units as KB, MB, or GB.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: **#StnFreeSpaceThreshold**

Default units: MB

Default: 0 (system computes based on repository size)

Min: 0

Max: 65536

STN_GEM_ABORT_TIMEOUT

STN_GEM_ABORT_TIMEOUT sets the time in minutes that the Stone will wait for a Gem running outside of a transaction to abort (in order to release a commit record), after Stone has signaled that Gem to do so. If the time expires before the Gem aborts, the Stone sends the Gem the error ABORT_ERR_LOST_OT_ROOT, and then either stop the Gem or force it to completely reinitialize its object caches, depending on the value of the related configuration parameter STN_GEM_LOSTOT_TIMEOUT. Negative timeouts are not allowed. Resolution of timeouts is one half the specified time-out interval.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: **#StnGemAbortTimeout**

Default: 1

Min: 1

Max: 1440

STN_GEM_LOSTOT_TIMEOUT

STN_GEM_LOSTOT_TIMEOUT sets the time in seconds that the Stone will wait after signaling the `Exception RepositoryViewLost`, before retracting the Gem's commit record and forcibly stopping the session. Negative timeouts other than -1 are not allowed. Resolution of timeouts is one half the specified timeout interval.

If the value is -1, the Stone does not stop the Gem; it immediately retracts the session's commit record, forcing the Gem to completely reinitialize its object caches.

CAUTION

A value of -1 entails a slight risk that the sleeping Gem will reactivate and begin writing to the shared page cache before it responds to the ABORT_ERR_LOST_OT_ROOT error, thus corrupting the shared page cache.

Because of these risks, design your application to minimize the chances of receiving the ABORT_ERR_LOST_OT_ROOT error.

The runtime parameter can be changed only by `SystemUser`.

Runtime parameter: **#StnGemLostOfTimeout**
Default: 60
Min: -1
Max: 5000000

STN_GEM_TIMEOUT

STN_GEM_TIMEOUT sets the time in minutes after which lack of interaction with the Gem causes the Stone to terminate the session. Negative timeouts are not allowed. Resolution of timeouts is one half the specified time-out interval. If non-zero, this timeout is also the maximum time allowed for a Gem to complete processing of its login to the Stone. If this timeout is 0, the maximum time for login processing is set to five minutes.

The runtime parameter can be changed only by `SystemUser`.

Runtime parameter: **#StnGemTimeout**
Min: 0
Default: 0 (Stone waits forever)

STN_HALT_ON_FATAL_ERR

If STN_HALT_ON_FATAL_ERR is set to true, the Stone will halt and dump core if it receives notification from a Gem that the Gem died with a fatal error that would cause Gem to dump core. By stopping the Stone at this point, the possibility of repository corruption is minimized. true is the recommended setting for systems during development.

If STN_HALT_ON_FATAL_ERR is set to false, the Stone will attempt to keep running if a Gem encounters a fatal error. false is the recommended setting for systems in production use.

Internally, the setting 0 = false, 1 = true. The runtime parameter can be changed only by `SystemUser`.

Runtime parameter: **#StnHaltOnFatalErr**
Default: false

STN_LISTENING_ADDRESSES

A list of 0 to 10 addresses upon which stone should listen for login connections. If list is empty, the default address ":::" is used, which means listen on any active network interfaces, plus the loopback ("::1") interface.

Each element of the list may be a name or a numeric IPv6 address. Each named address must resolve via `getaddrinfo()` to at least one address legal to listen on, i.e. resolve to the loopback or wildcard address, or to an address assigned to a network interface on this machine. Each numeric address must be an address legal to listen on.

Numeric IPv6 addresses may be any form recognized by `inet_pton(AF_INET6, address, etc.)` or by `inet_pton(AF_INET, address, etc.)` on the host operating system. Per RFC 2373 this includes these forms:

- IPv4 dotted-decimal format, d.d.d.d
- IPv6 hex format x:x:x:x:x:x where x is a 16 bit hexadecimal number
- IPv4-mapped IPv6 ::FFFF:d.d.d.d where d is an 8 bit decimal number

IPv6 format may contain at most one :: which is a contiguous group of zeros. The loopback address 0:0:0:0:0:0:1 can be written as ::1. The wildcard address 0:0:0:0:0:0:0 can be written as ::.

If the list contains the wildcard address ::, the other elements of the list are ignored.

If the list does not contain ::, then the loopback addresses ::1 and 127.0.0.1 are always listened on, even if not explicitly in the list, to support logins from system gems.

See public documents RFC 4291 and RFC 4038 for more information on IPV6 addressing.

Default: An empty list, ":::"

STN_LOG_IO_FLAGS

STN_LOG_IO_FLAGS specifies what (if any) special I/O flags will be used to open the database transaction logs.

This configuration is only applicable for Solaris.

Direct I/O tells the operating system avoid caching extent data in the file system cache. Enabling direct I/O tells the operating system to treat the database tranlogs as if they were on raw partitions. Direct I/O may greatly improve database performance in some cases.

STN_LOG_IO_FLAGS has no effect for tranlogs on raw partitions.

STN_LOG_IO_FLAGS has the following possible values:

- 0 - no special I/O flags. This is the default.
- 1 - enable Direct I/O on all tranlogs on file systems (Solaris only).

If the requested I/O mode is not available, the stone will fail to start and an error message will be printed in the stone log. If this happens, change this option back to zero and restart the stone.

STN_LOG_IO_FLAGS has no effect on tranlogs which reside on raw partitions.

Min: 0
 Max: 1 (Solaris only), 0 (All Others)
 Default: 0

STN_LOG_LOGIN_FAILURE_LIMIT

STN_LOG_LOGIN_FAILURE_TIME_LIMIT

If a user has greater than or equal to the `STN_LOG_LOGIN_FAILURE_LIMIT` number of login failures within the time in minutes specified by `STN_LOG_LOGIN_FAILURE_TIME_LIMIT`, a message is written to the Stone log file.

Changes to the runtime parameters require the `OtherPassword` privilege.

`STN_LOG_LOGIN_FAILURE_LIMIT`:

Runtime parameter: **#StnLogLoginFailureLimit**

Units: login attempts

Min: 0

Max: 65536

Default: 10

`STN_LOG_LOGIN_FAILURE_TIME_LIMIT`:

Runtime parameter: **#StnLogLoginFailureTimeLimit**

Units: Minutes

Min: 1

Max: 1440 (24 hours)

Default: 10

STN_LOGIN_LOG_ENABLED

Enable the logging of all session login and logout events to a separate log file owned by the stone. The file will be named `stoneName_login.log` and will be placed in the same directory as the stone log.

When this feature is enabled, logins and logouts are recorded for all sessions by default. Logging may be disabled for a `UserProfile` by sending the `#disableLoginLogging` message to a `UserProfile` instance and committing the transaction.

The login log file is a text file that contains one line per event. Fields within a line are separated by spaces; the Timestamp String is quoted. The fields logged in each line are:

- Timestamp String - time in human-readable form
- Timestamp Seconds - seconds from the epoch (January 1, 1970, 00:00 UTC)
- Event Kind - one of STARTUP, SHUTDOWN, LOGIN, LOGIN_FAIL, LOGOUT, or COMMIT_RESTORE.
- UserName - the `UserProfile`'s `userId`, or "Stone" for the stone process.
- SessionId
- ProcessId
- Real UNIX user ID - numeric value; always 0 on Windows.
- Effective UNIX user ID - numeric value; always 0 on Windows.
- Host Name - node name where the gem process is running.
- Gem IP Address - IP Address of the gem.
- Client IP Address - IP Address of the gem's client.
- NumCommits - number of commits performed by the session.

STARTUP and SHUTDOWN records are written to indicate when the stone was started and stopped and do not indicate a session login or logout.

Login failures are written for non-exempt sessions that fail a login attempt, usually due to specifying a bad password.

Default: FALSE

STN_LOOP_NO_WORK_THRESHOLD

STN_LOOP_NO_WORK_THRESHOLD indicates the maximum number of times the stone will continue executing its main service loop when there is no work to do. If the stone loops more than this number of times and finds no work, the stone will sleep for up to one second. The stone will immediately wake up when there is any work to be done.

Setting this value to zero disables this feature. Setting this value to a non-zero setting, in addition to causing the above behavior, will also cause the stone to not sleep whenever any of the following conditions are true and the no work threshold has not been exceeded:

- a session holds the commit token
- one or more sessions are waiting in the commit queue
- one or more sessions are waiting in the run queue.

Setting this parameter to a non-zero value will always cause the stone to consume more CPU.

Runtime parameter: **#StnLoopNoWorkThreshold**

Default: 0

Min: 0

Max: 536870911

STN_MAX_AIO_RATE

STN_MAX_AIO_RATE specifies the maximum I/O rate that each AIO page server is allowed when performing asynchronous writes. Since the I/O rate specified is applied to each page server, the total maximum I/O rate on the disk system is this value multiplied by STN_NUM_LOCAL_AIO_SERVERS.

The page servers use this maximum I/O rate for both dirty page and checkpoint writes.

Runtime parameter: **#StnMntMaxAioRate**

Min: 20

Max: 1000000

Default: 3000

STN_MAX_AIO_REQUESTS

STN_MAX_SESSIONS specifies the maximum number of asynchronous write requests the stone can have pending. If more than this number of asynchronous writes are requested, the stone will wait (sleep) until one or more of the pending requests have completed. Asynchronous write requests are only used to write to the current tranlog.

The maximum value allowed depends on the maximum allowed by the UNIX kernel. The maximum value for this parameter allowed by GemStone is the value of `_SC_AIO_MAX`

or 4096, whichever is lower. On some systems (such as Solaris), it is not possible to determine the value of `_SC_AIO_MAX`. In that case, GemStone will impose a maximum value of 128. Otherwise the maximum is 4096 or `_SC_AIO_MAX`, whichever is lower.

For further information on the `_SC_AIO_MAX` kernel parameter, refer to the UNIX documentation for your system and/or the UNIX man page for the `sysconf()` call.

Min: 100

Max: the minimum of 4096 or `_SC_AIO_MAX`; or 128 (see above)

Default: 128

STN_MAX_GC_RECLAIM_SESSIONS

The maximum number of page reclaim garbage collector sessions which are expected to be used on the system. When the default is specified, the actual value used is the number of extents defined in `DBF_EXTENT_NAMES` configuration.

Default: 0

Minimum: 0

Maximum: 256

STN_MAX_LOGIN_LOCK_SPIN_COUNT

Maximum number of times a session will attempt to write lock its user security data object at login time before raising a fatal error and failing the login. The session will sleep for 100 milliseconds between retries.

Enabling certain UserProfile security features (password aging, etc) causes each session to update its user security data object at login time and commit. A write lock must be acquired on this object to guarantee the commit succeeds.

Each UserProfile has a unique user security data object. Lock retries may be required when 2 sessions attempt to login with the same user ID at nearly the same instant. Simultaneous logins that use different user IDs never require lock retries.

Repositories that do not enable UserProfile security features are not affected by this parameter because the write-lock and commit described above are not required at login time.

Runtime parameter: `#StnMaxLoginLockSpinCount`

Default: 100

Minimum: 1

Maximum: 36000

STN_MAX_REMOTE_CACHES

`STN_MAX_REMOTE_CACHES` specifies the maximum number of remote shared page caches that the system may have.

Min: 0

Max: 10000

Default: 255

STN_MAX_SESSIONS

STN_MAX_SESSIONS limits the number of simultaneous sessions (number of Gem logins to Stone). The actual value used by Stone is the value of this parameter or the number of sessions specified by the software license key file, whichever is less. Using a value that is larger than needed will result in wasted space in the cache.

The number of logins may also be limited by changes in SHR_PAGE_CACHE_NUM_PROCS, or by the setting for the maximum number of file descriptors per process (imposed by the operating system kernel) .

Recommended: 40 or more, depending on your requirements

Min: 1

Max: 10000

Default: 40

STN_MAX_VOTING_SESSIONS

STN_MAX_VOTING_SESSIONS specifies the maximum number of sessions that can simultaneously vote on possible dead objects, at the end of a markForCollection or epoch garbage collection. To help prevent the voting on possible dead objects from causing large increases in response time of the system, set this to a value substantially lower than STN_MAX_SESSIONS.

Runtime parameter: **#StnMaxVotingSessions**

Min: 1

Max: 1000000

Default: 100

STN_NUM_AIO_WRITE_THREADS

STN_NUM_AIO_WRITE_THREADS specifies the number of native threads the Stone will start to perform writes to the tranlog. In commit-intensive systems, this should be increased to 8 or 16.

Cache Statistic: StnAioNumWriteThreads (Stone)

Min: 4

Max: 256

Default: 4

STN_NUM_GC_RECLAIM_SESSIONS

STN_NUM_GC_RECLAIM_SESSIONS specifies the number of page reclaim garbage collector sessions (Reclaim Gem sessions) that will be started when the Stone starts. This value must be less than or equal to STN_MAX_GC_RECLAIM_SESSION.

Runtime parameter: **#StnNumGcReclaimSessions**

Default: 1

Min: 0

Max: 256

STN_NUM_LOCAL_AIO_SERVERS

STN_NUM_LOCAL_AIO_SERVERS is the approximate number of page server processes to start as local asynchronous I/O servers for the shared page cache on the node where the Stone runs. The number of extents known to the Stone at startup is divided by the value of STN_NUM_LOCAL_AIO_SERVERS to compute the internal configuration parameter `StnRDbfMaxFilesPerServer`. The latter parameter is the approximate number of extent files to be serviced by each AIO page server.

For instance, if your configuration has six extents, setting STN_NUM_LOCAL_AIO_SERVERS to 3 causes each AIO page server to service $(6 \div 3) = 2$ extents.

Under certain circumstances, multiple AIO page servers can help you achieve the maximum possible commit rate. A value greater than one is recommended only if there are two or more extents, the host has multiple CPUs (to allow parallel execution), and the disk drive hardware allows concurrent writes to disk (the extents are on separate spindles, or the equivalent).

Min: 1
Max: 256
Default: 1

STN_OBJ_LOCK_TIMEOUT

STN_OBJ_LOCK_TIMEOUT specifies the time in seconds that a session is allowed to wait to obtain one of the special single object write locks. For more information, see the methods `System >> waitForRcWriteLock:` and `System >> waitForApplicationWriteLock:queue:autoRelease:.`

Runtime parameter: **#StnObjLockTimeout**
Min: 0
Max: 86400
Default: 0 (stone waits forever)

STN_PAGE_MGR_COMPRESSION_ENABLED

STN_PAGE_MGR_COMPRESSION_ENABLED determines if the page manager will compress the list of pages that it sends to remote shared page caches for removal. If set to TRUE, all lists of pages larger than 50 will be compressed before transmission. The same compressed list is used to send to all remote shared page caches; i.e., the compression operation is performed no more than once for each list of pages to be sent. Has no effect on systems that do not use remote shared page caches.

Runtime parameter: **#StnPageMgrCompressionEnabled**
Cache Statistic: PageMgrCompressionEnabled (Stone)
Default: FALSE

STN_PAGE_MGR_MAX_WAIT_TIME

Maximum time the stone will defer servicing the page manager thread because the stone is busy with other tasks. Normally the stone services the page manager whenever it has idle time and no session is performing a commit. If the time the page manager has been waiting for service exceeds this value, the stone will service the page manager unconditionally and increment the cache statistic PageManagerStarvedCount.

Runtime parameter: **#StnPageMgrMaxWaitTime**
Units: Milliseconds
Default: 200
Min: 1
Max: 1000

STN_PAGE_MGR_PRINT_TIMEOUT_THRESHOLD

STN_PAGE_MGR_PRINT_TIMEOUT_THRESHOLD is the threshold in real seconds used by the page manager to determine if a slow response from a remote shared page cache should be printed to the page manager log file. If a remote cache takes longer than this number of seconds to respond to the page manager, the page manager will print a message to the log file. If a remote cache takes less than this number of seconds to respond, no message is printed.

Note that this value controls the writing of log messages only. The connection to the remote cache will not be terminated by page manager unless STN_REMOTE_CACHE_PGSRV_TIMEOUT is exceeded.

Runtime parameter: **#StnPageMgrPrintTimeoutThreshold**
Cache Statistic: PageMgrPrintTimeoutThreshold (Stone)
Min: 0
Max: 3600
Default: 5

STN_PAGE_MGR_REMOVE_MAX_PAGES

STN_PAGE_MGR_REMOVE_MAX_PAGES sets the maximum batch size for the Page Manager gem. This is the maximum number of pages in a single request to the stone. Must be greater than or equal to STN_PAGE_MGR_REMOVE_MIN_PAGES

Runtime parameter: **#StnPageMgrRemoveMaxPages**
Cache Statistic: PageMgrRemoveMaxPages (Stone)
Min: 1
Max: 16384
Default: 16384

STN_PAGE_MGR_REMOVE_MIN_PAGES

STN_PAGE_MGR_REMOVE_MIN_PAGES sets the minimum batch size for the Page Manager gem. When the number of pages waiting to be processed by the Page Manager is greater than this value, then the Page Manager will request the pages from the stone and process them. Otherwise the Page Manager will wait until this threshold is exceeded before requesting pages from the stone. Must be less than or equal to STN_PAGE_MGR_REMOVE_MAX_PAGES

Runtime parameter: **#StnPageMgrRemoveMinPages**

Cache Statistic: PageMgrRemoveMinPages

Min: 0

Max: 1792

Default: 40

STN_PRIVATE_PAGE_CACHE_KB

STN_PRIVATE_PAGE_CACHE_KB sets the default size of the Stone page cache.

If no units are specified, the value is in KB. You may also specify units as KB, MB, or GB.

Default units: KB

Min: 128

Max: 524288

Default: 2000

STN_REMOTE_CACHE_PGSRV_TIMEOUT

STN_REMOTE_CACHE_PGSRV_TIMEOUT specifies the maximum time in seconds that the page manager session will wait for a response from a page server on a remote shared page cache. If no response is received within the timeout period, all Gems attached to that cache are logged off and a message is written to the Stone and page manager logs. Negative timeouts are not allowed. A timeout value of zero causes the page manager to wait forever.

Runtime parameter: **#StnRemoteCachePgsvrTimeout**

Cache Statistic: PageMgrRemoteCachePgsvrTimeout (Stone)

Min: 0

Max: 3600

Default: 15

STN_REMOTE_CACHE_TIMEOUT

STN_REMOTE_CACHE_TIMEOUT sets the time in minutes after the last active process on a remote node logs out before the Stone shuts down the shared page cache on that node.

A value of 0 causes the Stone to shut down the remote cache as soon as possible.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: **#StnRemoteCacheTimeout**

Min: 0

Max: 5000000

Default: 5

STN_SHR_TARGET_PERCENT_DIRTY

STN_SHR_TARGET_PERCENT_DIRTY specifies the maximum percentage of the Stone's shared page cache that can contain dirty pages without AIO page servers increasing their IO rates.

Runtime parameter: `#StnShrPcTargetPercentDirty`

Min: 1

Max: 90

Default: 20

STN_SIGNAL_ABORT_CR_BACKLOG

STN_SIGNAL_ABORT_CR_BACKLOG sets the number of old transactions (commit records) above which the Stone will start to generate SignalAbort messages to Gems. The Gem receives these as a TransactionBacklog exception.

If the Gem is not in transaction, this is received if the Gem has enabled receipt of sigAborts by invoking `System >> enableSignaledAbortError`. If the Gem that is out of transaction does not respond within the time allowed by `STN_GEM_ABORT_TIMEOUT`, the Gem will receive a `ABORT_ERR_LOST_OT_ROOT`.

If the Gem is in transaction, it will receive `finishTransaction` if it has invoked `System >> enableSignaledFinishTransactionError`. No further signals are sent to a Gem that is in transaction, whether or not it responds to the signal.

You may need to tune this option according to your application's commit rate and repository free space.

Changing the runtime parameter requires GarbageCollection privilege.

Runtime parameter: `#StnSignalAbortCrBacklog`

Default: 20

Min: 2

Max: 65536

STN_SYMBOL_GC_ENABLED

STN_SYMBOL_GC_ENABLED determines if symbol garbage collection is allowed to run on the system. Setting this value to true enables symbol garbage collection.

Updating the runtime parameter requires the GarbageCollection privilege.

Runtime parameter: `#StnSymbolGcEnabled`

Default: FALSE

STN_TRAN_FULL_LOGGING

If `STN_TRAN_FULL_LOGGING` is set to true, all transactions are logged, and log files are not deleted by the system. This is full transaction logging mode. In this mode, the transaction logs are providing real-time incremental backup of the repository. If no disk space is available for logs, Gem session processes may appear to "hang" until space becomes available.

If `STN_TRAN_FULL_LOGGING` is set to false, only transactions smaller than `STN_TRAN_LOG_LIMIT` are logged; larger transactions cause a checkpoint, which updates

the extent files. This is partial transaction logging mode. Log files are deleted by the system when the circular list of log directories wraps around. This setting allows a simple installation to run unattended for extended periods of time, but it does **not** provide real-time backup.

Once you have started the Stone on a repository with `STN_TRAN_FULL_LOGGING = true`, then the true state will persist in the repository; any subsequent changes to this parameter in the configuration file are ignored. To change the repository back to partial logging, you must do a Smalltalk full backup and then restore the backup into a copy of `$GEMSTONE/bin/extent0.dbf`.

For further information, see Chapter 8, "Managing Transaction Logs."

Default: true.

STN_TRAN_LOG_DEBUG_LEVEL

This option is only for GemStone internal use. Customers should not change the default setting unless directed to do so by GemStone Technical Support.

Runtime parameter: `#StnTranLogDebugLevel`

Default: 0

STN_TRAN_LOG_DIRECTORIES

`STN_TRAN_LOG_DIRECTORIES` lists the directories or raw disk partitions to be used for transaction logging. This list defines the maximum number of log files that will be online at once. Each entry must be a directory or a raw disk partition. Directories may appear multiple times in the list. A given raw disk partition may appear only once. If raw partitions are used or if `STN_TRAN_FULL_LOGGING` is false, at least two entries should be included.

Min: 1 entry

Max: 2^{31} entries

Default: Empty (the system will not run without at least one entry)

Initial setting: `$GEMSTONE/data/`

STN_TRAN_LOG_LIMIT

`STN_TRAN_LOG_LIMIT` sets the maximum transaction log entry size limit in KB. Successful commits of transactions consuming more than this amount of log file space when `STN_TRAN_FULL_LOGGING` is set to false will cause a checkpoint. This option has no effect when `STN_TRAN_FULL_LOGGING` is set to true.

The runtime parameter can be changed only by SystemUser.

Runtime parameter: `#StnTranLogLimit`

Min: 25

Max: 1000

Default: 1000

STN_TRAN_LOG_PREFIX

STN_TRAN_LOG_PREFIX sets file name prefixes for transaction log files. A sequence number and “.dbf” is added to the prefix; for example, “tranlog” produces “tranlog0.dbf, tranlog1.dbf, ...”. You can set this configuration option to permit multiple repository monitors to share a log directory without conflict.

Default: tranlog

STN_TRAN_LOG_SIZES

STN_TRAN_LOG_SIZES sets the maximum sizes (in MB) of all log files, in order and separated by commas. Each size applies to a corresponding log file specified in STN_TRAN_LOG_DIRECTORIES, and the number of entries must match. The sizes also apply to the corresponding entries in STN_REPL_TRAN_LOG_DIRECTORIES when that list is not empty.

For transaction logs on raw partitions, if the size specified is larger than the physical size of the corresponding raw partition, the STN_TRAN_LOG_SIZES VALUE is adjusted appropriately.

Min: 1

Max: 16384

Default: Empty (the system will not run unless sizes are specified)

Initial setting: 100

STN_TRAN_Q_TO_RUN_Q_THRESHOLD

STN_TRAN_Q_TO_RUN_Q_THRESHOLD specifies the number of sessions in the commit queue (waiting for the commit token) above which the stone will allow the remaining sessions in the queue to process unions (read old commit records) while waiting for the commit token.

For example, if this parameter is set to 6 and there are 9 sessions in the commit queue, the 3 last sessions will be allowed to process unions while waiting for the token. If there are 6 or fewer sessions in the queue, no sessions will process unions.

The first session in the commit queue never processes unions since it will receive the token when the current commit completes.

Runtime parameter: **#StnTranQToRunQThreshold**

Min: 1

Max: 1024

Default: 6

STN_WELL_KNOWN_PORT_NUMBER

STN_WELL_KNOWN_PORT_NUMBER is the port number that the Stone will use as its well-known port. The well-known port is used by all Gems while establishing their initial connection to the Stone during the login sequence.

If the specified port is in use by another process, the Stone will not start and exits with an error.

A value of zero indicates the port number will be selected by the system.

Min: 1

Max: 65535

Default: 0

A.4 Runtime-only Configuration Parameters

The parameters described in this section are similar to the configuration parameters above, but can only be read or modified at runtime.

The process for modifying is similar to that for the runtime parameter equivalents for the configuration parameters listed in the configuration files.

The runtime parameters are read using the method `System class>>configurationAt:`, and updating using `System class>>configurationAt:put:`.

GemConvertArrayBuilder

If True, allows old style Array Builder syntax `#[a, b]` to be parsed correctly. The compiler converts this syntax to the new form `{ a . b }`, and updates method source as well as compiled code.

Default: false

GemDropCommittedExportedObjs

If this configuration parameter is true, clean, committed objects may be dropped from RAM. This reduces demand on memory in the Gem, but there is the small cost of an additional bitmap lookup when the object is faulted, to detect if this object is in the Pure Export Set.

Default: false

GemExceptionSignalCapturesStack

If this configuration parameter is true, invocations of `AbstractException>>_signalWith:` fill in the `gsStack` instance variable of the receiver, allowing subsequent calls to `Exception >> stackReport`.

Default: false

LogOriginTime

#LogOriginTime is the time the current sequence of Stone logs was started. It is the same value returned by `Repository>>logOriginTime`. For information about when a new sequence is started, see the method comment for `Repository>>commitRestore` in the image.

This should not be modified.

SessionInBackup

#SessionInBackup is the GemStone session number of the session performing a full backup, or -1 if a backup is not in progress.

This should not be modified by the user.

StnCurrentTranLogDirId

#StnCurrentTranLogDirId is the one-based offset of the current transaction log into the list of log directory names, `STN_TRAN_LOG_DIRECTORIES`. It is the same value returned by `Repository>>currentLogDirectoryId`.

This should not be modified.

StnCurrentTranLogNames

#StnCurrentTranLogNames is an Array containing up to two Strings: the name of the transaction log to which records currently are being appended, and the name of the current replicated log. These are the same values returned by `Repository>>currentLogFile` and `currentLogReplicate`, respectively.

StnLogFileName

The stone log file path and name.

This is a read-only value.

StnLogGemErrors

#StnLogGemErrors is intended for internal debugging use. When it is set to 1, the Stone logs error messages it sends to Gems.

StnLoginsSuspended

#StnLoginsSuspended ordinarily has the values 0 (false) and 1 (true) as set by `System class>>suspendLogins` and `resumeLogins`.

Changing this parameter requires the `SystemControl` privilege.

StnMaxReposSize

#StnMaxReposSize is the maximum size of the repository for your GemStone license, as set by the GemStone keyfile. The value of this parameter is 0 when not limited by the keyfile.

This should not be modified.

StnMaxSessions

#StnMaxSessions is the maximum allowed number of sessions for your GemStone license, as set by the GemStone keyfile. It is not related to the STN_MAX_SESSIONS configuration option. The value of this parameter is 0 when not limited by the keyfile.

This should not be modified.

StnSunsetDate

#StnSunsetDate is the sunset date for your GemStone license, as set by the GemStone keyfile. The value of this parameter is 0 when not limited by the keyfile.

This should not be modified.

StnTranLogOriginTime

#StnTranLogOriginTime is the time when the current transaction log was started.

This should not be modified.

GemStone Utility Commands

The GemStone/S 64 Bit utility commands in this appendix are provided in the `$GEMSTONE/bin` directory.

- `copydbf` (page 310) – copy an extent, transaction log or backup.
- `gslist` (page 314) – list running GemStone server processes.
- `pageaudit` (page 316) – offline audit of repository pages.
- `pstack` (page 317) – get C-level and Smalltalk stack traces.
- `removedbf` (page 318) – delete an extent, transaction log or backup.
- `startcachewarmer` (page 319) – start cache warmers.
- `startlogreceiver` (page 320) – start logreceiver for hot standby system.
- `startlogsender` (page 322) – start logsender for hot standby system.
- `startnetldi` (page 324) – start a NetLDI.
- `startstone` (page 326) – start a Stone.
- `statmonitor` (page 327) – start logreceiver for hot standby system.
- `stoplogreceiver` (page 329) – stop logreceiver for hot standby system.
- `stoplogsender` (page 330) – stop logsender for hot standby system.
- `stopnetldi` (page 331) – shut down a NetLDI.
- `stopstone` (page 332) – shut down a Stone.
- `topaz` (page 333) – command line scripting tool for GemStone. See also the separate manual *Topaz Users's Guide*.
- `vsd` (page 334) – graphic tool to analyze GemStone statistics files. See also the separate manual *VSD Users's Guide*.
- `waitstone` (page 335) – verify status of a Stones

copydbf

copydbf	<i>sourceNRS destinationNRS</i> [-C -c] [-E] [-filePrefix] [-netLdiName] [-ppgsvrId] [-sMbytes] [-l -m -P] [-h -v]
copydbf	<i>sourceNRS</i> [-i -I] [-netLdiName] [-ppgsvrId] [-h -v]
<i>sourceNRS</i>	The source file or raw partition (containing an extent, a transaction log, or a backup file created by <code>fullBackupTo:</code>) as a GemStone network resource string.
<i>destinationNRS</i>	The destination file, directory, or raw partition as a GemStone network resource string. If the destination is a file system directory (the trailing <code>/</code> is optional), a file name is generated and appended to <i>destinationNRS</i> based on the type and internal <code>fileId</code> of the source. Use of <code>/dev/null</code> as the destination is supported only for files as a means of verifying that the file is readable.
-C	Compress output. The output must be a filesystem file. Write the output compressed, in gzip format. The output file name will have the suffix <code>.gz</code> appended to it, if it does not end in <code>.gz</code> .
-c	Compress transaction log with record-level compression. This option only applies to transaction logs. The resulting log does not have the <code>.gz</code> option.
-E	Ignore disk read errors. If the disk read error occurs while reading an extent root page, the copy will fail. Otherwise, pages of the source file that cannot be read will be replaced with an invalid-page-kind page in the destination file. The destination file may not be usable. This option only applies to extents, not to transaction logs or backup files.
-filePrefix	If <i>destinationNRS</i> is a file system directory, then <i>filePrefix</i> overrides the filename prefix that would be generated based on the contents of <i>sourceNRS</i> . If <i>destinationNRS</i> is other than a file system directory, this option has no effect.
-netLdiName	The name of the GemStone network server; the default is <code>gs64ldi</code> .
-ppgsvrId	The name of a specific <code>runpgsvr</code> (similar to <code>gemnetid</code>).
-sMB	The size (in MB) to pre-allocate the destination file. For instance, <code>-s10</code> allocates at least 10 MB to the created file. If you do not specify the <code>-s</code> option, the output file is made as short as possible.
-h	Displays a usage line and exits.
-i	Information only. When this option is present without <i>destinationNRS</i> , information about <i>sourceNRS</i> is printed without performing a file copy. If both <code>-i</code> and <i>destinationNRS</i> are present, an error message is printed.

- I** Full information. The same information is printed as for **-i**. In addition, if the file is a transaction log, all checkpoint times found are listed instead of only the last one.
- l** Least-significant-byte ordering for the *destinationNRS*. This byte ordering is the native byte ordering for Intel processors.
- m** Most-significant-byte ordering for the *destinationNRS*. This byte ordering is the native byte ordering for Solaris SPARC, AIX POWER, and HP-UX PA-RISC and Itanium processors.
- P** Preserve byte ordering. This option creates the destination file using the byte ordering found in the source file. The default is to write the file using the host's native byte ordering.
- v** Print version and exit

To make copies of extent files or transaction logs, the user executing **copydbf** must have read permission to the file. If you attempt to copy extent files that are in use, if checkpoints are not suspended the resulting repository may be corrupt and unusable. See "How To Make an Extent Snapshot Backup" on page 193 for more information.

GemStone repository files on the UNIX file system can usually be copied using the ordinary **cp** command. However, if you are copying between operating systems and the source and destination have different byte ordering, you should use **copydbf**. You can use the first form of **copydbf** for disk-to-disk copies between machines (without NFS) or copies to and from raw disk partitions. To use **copydbf** between remote nodes, you must have a NetLDI running on both nodes, and authentication must be configured to allow access.

You must give an NRS (network resource string) for both the source file and the destination. A local machine file specification is a valid NRS.

If the destination is a directory in a file system, **copydbf** generates a filename based on the type of file. The generated name includes a prefix (*extent*, *tranlog*, or *backup*), a fileId representing an internal sequence number that starts at 0, and the extension *.dbf*. You can use the **-f** option to change the prefix.

A message describing the source and destination files is printed to stderr before starting the copy. The size of the destination file is printed to stderr after the copy is completed.

For example:

```
% copydbf $GEMSTONE/bin/extent0.dbf .
```

```
Source file: /users/GemStone/data/extent0.dbf
  File type: extent  fileId: 0
  ByteOrder: Intel (LSB first)  compatibilityLevel: 844
  Last checkpoint written at: 02/24/2014 14:37:59 PST.
Destination file: ./extent0.dbf
  ByteOrder: Intel (LSB first)
  Clean shutdown, no tranlog needed for recovery.
  Last tranlog written to had fileId 5 ( tranlog5.dbf ).
  File contains 4608 records occupying 75.5 MBytes.
```

To obtain the same source file information (but not the size) without making a copy, use the second form of the command: **copydbf -i sourceNRS**. In this usage, you do not specify a *destinationNRS*.

The following **copydbf -i** example displays information for an extent, and indicates the oldest transaction log that would be needed to recover from a system crash:

```
% copydbf -i extent0.dbf
```

```
Source file: extent0.dbf
  File type: extent  fileId: 0
  ByteOrder: Intel (LSB first)  compatibilityLevel: 844
  Last checkpoint written at: 03/07/2014 19:03:46 PST.
  Oldest tranlog needed for recovery/restore is fileId 5
  ( tranlog5.dbf ).
  Extent was shutdown cleanly; no recovery needed.
```

The next example displays information for a backup, and indicates the oldest transaction log that would be needed to restore subsequent transactions.

```
% copydbf -i backup.dat
```

```
File type: backup  fileId: 0
  ByteOrder: Intel (LSB first)  compatibilityLevel: 844
  The file was created at: 02/25/2014 20:34:01 PST.
  Full backup started from checkpoint at: 02/25/2014 20:34:00 PST.
  Oldest tranlog needed for restore is fileId 2 ( tranlog2.dbf ).
  Backup was created by GemStone Version: 3.2.0 .
```

To obtain the size of a repository file in a raw partition, use this form:

```
% copydbf sourceNRS destinationNRS
```


For a listing of all checkpoints recorded in a transaction log, use **copydbf -I sourceNRS**. This information is helpful in restoring a GemStone backup to a particular point in time. For example:

```
% copydbf -I tranlog5.dbf
```

```
Source file: tranlog5.dbf
  File type: tranlog  fileId: 5
  ByteOrder: Intel (LSB first)  compatibilityLevel: 911
  The file was created at: 02/27/2014 11:49:05 PST.
  The previous file last recordId is 36 .
  Scanning file to find last checkpoint...
  Checkpoint 1 started at: 02/27/14 11:54:35 PST.
    oldest transaction references fileId -1 ( this file ).
  Checkpoint 2 started at: 02/27/14 11:56:17 PST.
    oldest transaction references fileId -1 ( this file ).
  Checkpoint 3 started at: 02/27/14 11:58:30 PST.
    oldest transaction references fileId -1 ( this file ).
  Checkpoint 4 started at: 02/27/14 12:03:01 PST.
    oldest transaction references fileId -1 ( this file ).
  File contains 69 records occupying 35328 bytes.
```

To pre-allocate disk space in the destination file, use the **-s** option. For instance, **-s50** would allocate at least 50 MB to the created file. The output file is made as short as possible by default.

In the following example, the local GemStone repository file "extent0.dbf" is copied to a remote machine using a full *destinationNRS*. In this example, the repository file is copied to a remote machine named "node," using remote user account "username" and "password," with a remote filespec of "/users/path/extent0.dbf_copy," via the standard GemStone network server *gs64ldi* using TCP protocol:

```
% copydbf $GEMSTONE/data/extent0.dbf
!@node#auth:username@password#dbf!/users/extent0.dbf_copy
```

The next example copies a fresh repository extent to an existing raw disk partition. If the raw partition already contains a repository file or backup, use **removedbf** first to mark it as being empty.

```
% copydbf $GEMSTONE/bin/extent0.dbf /dev/rsd3h
```

Copying a transaction log from a raw partition to a file system directory generates a file name having the same form as if that transaction log had originated in a file system. If the internal fileId is 43, this example would name the destination file `/dsk1/tranlogs/tranlog43.dbf`:

```
% copydbf /dev/rsd3h /dsk1/tranlogs/
```

gslis

gslis	<code>[-h] [-V] [-l -p -x] [-c] [-q] [-v] [-t secs] [-u user] [-m host]+ [[-n] name]+</code>
-c	Removes locks left by servers that have been killed.
-h	Prints a usage line and exits.
-l	Prints a long listing (includes pid and port).
-m host	Only list servers on machine <i>host</i> ; default is '.' which represents the local host. If not the local host, the machine <i>host</i> must be running a compatible version of netldi with the default netldi name for the environment in which <i>gslis</i> is being executed.
[-n] name	Only list the server <i>name</i> .
-p	Prints only the pid (process id), or 0 if the server does not exist.
-q	(Quiet.) Don't print any extra information; intended for use when the output will be processed by some other program.
-t secs	Wait <i>secs</i> seconds for server to respond (only with -v); default is 2 seconds.
-u user	Only list servers started by <i>user</i> .
-v	Verify the status of each server.
-V	Print the version information and exit.
-x	Prints an exhaustive listing, with each item on a separate line.

The **gslis** command prints information about GemStone servers. The default listing prints the following server attributes in columns:

Status	One of the following: OK server is accepting clients (-v only) frozen server is not responding (-v only) restoreLogs servers is in restore mode (-v only) full server can't accept any more clients (-v only) exists server process exists but is not verified killed server process does not exist startup server process is not yet accepting clients
Version	GemStone version of the server.
Owner	The account name of the user who created the server.
Started	The date and time that the server was started.
Type	One of the following: Netldi, Stone, cache (shared page cache monitor), logsender, logreceiver
Name	The server's name.

When you include the **-l** or **-x** option, the following additional columns are printed:

Pid	The process id of the server's main process.
Port	The port number of the server's listening socket.
The -x option prints the preceding attributes on separate lines, and adds lines for the following as appropriate:	
options	Options used when the server was started.
logfile	Full path of server's log file, if it exists.
sysconf	The GemStone system configuration file (see page 266).
execonf	The GemStone executable configuration file (see page 268).
GEMSTONE	Root of the product tree used by the server.

If many servers are reported as **frozen**, try increasing `-t secs`.

By default, status is returned for all servers on the current host. To specify a particular server, use the `-n` switch, or just include the server name on the command line (since the `-n` is optional). To specify multiple server names, include the `-n name` option for each server.

The `-m` option allows you to list servers on a remote host. To specify more than one remote host, include the `-m host` option for each host. Names on remote hosts are prefixed by `host :`, where *host* is the name of the remote machine.

To list servers on a remote host, **gslist** must be able to contact a NetLDI running on the remote machine. The remote NetLDI must be named according to the default NetLDI for the environment in which **gslist** is running, either "gs64ldi" or as specified by `$GEMSTONE_NRS_ALL`. In addition, the GemStone version of the remote NetLDI must be compatible with **gslist**.

The date and time that the process started is normally printed in a format specific to **gslist** and fitting into the table display. To get a parseable but potentially less readable format, you may use the environment variable `GS_GSLIST_TIME_FORMAT` to specify a UNIX-style date format string,

The exit status has the following values:

0	Operation was successful.
1	No servers were found.
2	A stale lock was removed (in response to <code>-c</code> switch).
3, 4	An error occurred.

gslist is available for Windows, in addition to regular server platforms, as part of the Windows Client installation. **gslist** on Windows must be used with the `-m` option, specifying a hostname for a host that is running Gemstone.

pageaudit

pageaudit	<i>[gemStoneName]</i> [-e <i>exeConfig</i>] [-z <i>systemConfig</i>] [-f] [-d] [-l <i>logfile</i>] [-h]
<i>gemStoneName</i>	Name of the GemStone repository monitor; the default is <code>gs64stone-audit</code> . Network resource syntax is not permitted.
-e <i>exeConfig</i>	The GemStone executable-dependent configuration file (see page 268).
-z <i>systemConfig</i>	The GemStone system configuration file (see page 266).
-f	Keeps running beyond the first error, if possible.
-d	Disable audit of data pages. Only audit Object Table pages, bitmaps, and other non-data pages.
-l <i>logfile</i>	Write output to the file with the given name. The file is created if it does not exist. If there is an existing file with this name, the pageaudit output is appended to the end of this file.
-h	Displays a usage line and exits.

Audit the pages in a GemStone repository, which must not be in use. **pageaudit** opens the repository specified by the relevant configuration files. The arguments **-e** *exeConfig*, and **-z** *systemConfig* determine which configuration files **pageaudit** reads. Other options for this command generally are not needed.

By default, all pages in the repository are verified; this includes data pages as well as Object Table, bitmap, and other pages containing internal information. Audit of data pages can be disabled using the **-d** option.

An error is returned if another Stone is running as *gemStoneName* or has opened the same repository.

When you include the **-f** switch, **pageaudit** prints all errors possible. Without **-f**, the default is to stop after the first error is found.

This utility can take a long time to run, so it is best to run it as a background job.

For additional information about **pageaudit** and a description of its output, see "To Perform a Page Audit" on page 120.

pstack

pstack **[-b]** *processPid*

processPid The PID of a running GemStone process.

-b If specified, outputs a brief stack rather than both brief and full stacks.

The pstack command attaches to the process with the given pid, outputs the stack of that process to stdout, and detaches. The Smalltalk execution stack is printed to the Gem log or topaz linked session.

Execution of the running process briefly pauses while gbd is attaching, but the process will subsequently continue running normally.

pstack is similar to functions provided with some OS platforms.

The -b option allows you to specify brief format stacks. Without the -b option, stacks with one line per frame for each thread, and complete stacks are generated. Using the -b option, only the one-line per frame stack traces are output.

removedbf

removedbf	<i>dbfNRS</i> [- n <i>netLdiName</i>] [- p <i>pgsvrNetId</i>] [- h]
<i>dbfNRS</i>	The GemStone repository filename or device, as a network resource string, for the repository to be removed.
- n <i>netLdiName</i>	The name of the GemStone network server; default is <i>gs64ldi</i> .
- p <i>pgsvrNetId</i>	The name of a specific page server to use.
- h	Displays a usage line and exits.

This command removes (erases) a GemStone repository file. It is provided primarily for erasing an extent or transaction log from a raw partition, but it also works on files in the file system and on remote machines. You must provide the NRS (network resource string) of the file to be removed. (A local machine filespec is a subset of an NRS.)

If you specify a file in the file system, this command is equivalent to the **rm** command. If you specify a raw disk partition, GemStone metadata in the partition is overwritten so that GemStone will no longer think there is a repository file on the partition.

This command does not disconnect an extent from the logical repository. To alter the configuration by disconnecting an existing extent, see "How To Remove an Extent" on page 173.

The options for this command generally are not needed for a standard GemStone configuration.

startcachewarmer

startcachewarmer	[-d -D] [-h] [-l <i>limit</i>] [-L <i>path</i>] [-n <i>numSessions</i>] [-s <i>stone</i>] [-w <i>delayTime</i>] [-W]
-d	Reads data pages into the cache (default: only object table pages are read).
-D	Reads data pages into the UNIX file buffer cache and not the shared page cache.
-h	Displays a usage line and exits.
-L <i>path</i>	Path to a writable log file directory (default: current directory)
-l <i>limit</i>	Stops cache warming if the number of free frames in the cache falls below the specified <i>limit</i> . If <i>limit</i> is -1 (the default), have the system compute the actual limit based on the size of the shared cache. If <i>limit</i> is 0, force cache warming to continue even if the shared cache is full.
-n <i>numSessions</i>	Number of worker sessions to start. The default is the number of CPUs + number of Extents, plus 1 additional master session. The cachewarmer will exit if not enough sessions are available.
-s <i>stone</i>	Name of the running Stone (default: <code>gs64stone</code>).
-w <i>delayTime</i>	Wait <i>delayTime</i> seconds between spawning Gems (default: 1)
-W	Wait for cache warming Gems to exit before exiting this script. By default, this script spawns Gems in the background and exits immediately.

The **startcachewarmer** command warms up the shared page cache on startup, by preloading object and data tables into the cache. This allows the overhead of initial page loading to occur in a controlled way on system startup, rather than more gradually as the repository is in use.

Cache warming runs in one or two phases:

- In the first phase, the object table is loaded into the shared page cache. During this phase, if any data pages will be loaded later, the data pages that are referenced by the object table lookups are recorded for use in phase 2.
- In the second phase, data pages remembered from the first phase are loaded either into the shared page cache or the file buffer. This phase runs only if data pages will be loaded into the shared page cache or the file buffer.

For greater efficiency, you can set the configuration file parameters `STN_CACHE_WARMER` and `STN_CACHE_WARMER_SESSIONS` to perform cache warming on startup. See Appendix A for more details on these parameters.

startlogreceiver

- startlogreceiver** **-P** *listeningPort* **-A** *listeningAddress* **-T** *tranlogDir*+ **[-I** *logFile*
[-s *stoneName*] **[-p** *alternatePort*] **[-d]** **[-v]** **[-h]** **[-C** *certFileName*]
[-J *certAuthFileName*] **[-K** *keyFileName*] **[-Q** *passphrasestring*] **[-S]** **[-V]**
- A** *listeningAddress* Address that will be used to attempt connection to a logsender.
- d** Print debug tracing of tranlog read and write operations to log file. The log file will be much larger.
- I** *logFileOrDir* The path and filename or directory for the logged output of the logreceiver process. The default is `/opt/gemstone/log/logreceiver_listeningPort.log`.
- h** Displays a usage line and exits.
- P** *listeningPort* Port or named service that will be used to attempt connection to a logsender.
- p** *alternatePort* The logreceiver listens on localhost on this port for stoplogreceiver commands. If logsender and logreceiver are run on same machine, then **-p** must be used to specify a port different than *listeningPort* specified with the **-P** argument.
- s** *stoneName* The name of the slave stone. *stoneName* is required for the logreceiver to be able to notify a stone in continuous restore mode that new log records have arrived. Without *stoneName*, logreceiver will just write tranlogs to the file system.
- T** *tranlogDir* Directory(s) where logreceiver writes files received from the logsender, and continuousRestoreFromArchiveLogs: will read from. At least one is required, up to 20 **-T** may be specified.
- v** Print version and exit.
- Additional arguments for SSL:
- C** *certFileName* Certificate in PEM format that will be sent to the peer upon request.
- J** *certAuthFileName* Certificate authority (CA) file in PEM format to use for peer certificate verification.
- K** *keyFileName* Private key in PEM format for the certificate (**-C** option).
- Q** *passphrasestring* Private key passphrase. Required if the **-K** option is used and the private key is encrypted.
- S** Enable SSL mode. Must be specified to use any other SSL options and must also be specified when starting the peer process.
- V** Disable verification of the peer's certificate.

This utility starts up a logreceiver process. As part of a hot standby setup, the logreceiver process runs on the slave system to receive transaction log records, as they are generated, from a logsender process that is running on a master system.

The received transaction log records are written to files in a specified directory on the slave system. The slave system can run `continuousRestoreFromArchiveLogs`: to restore these transaction log records.

The logreceiver writes logging output to a file with the path and name or in the directory specified by the `-l` argument. If this file is a directory, the file name is `logreceiver_listeningPort.log`. If no `-l` argument is used, it writes to `/opt/gemstone/log/logreceiver_listeningPort.log`. If a file with the specified name exists, it is appended to. It is not deleted on process exit.

The logreceiver process will continue running until explicitly stopped using the **stoplogreceiver** command. If connection to the stone or the logsender process is lost, it will attempt to reconnect.

gslist will report running logreceiver processes.

To start the logsender using SSL to establish a secure connection between logsender and log receiver, both `startlogsender` and `startlogreceiver` can be provided with SSL credentials. When SSL arguments are provided, `startlogreceiver` will start the logreceiver in SSL mode.

startlogsender

- startlogsender** **-P** *listening port* **-A** *listeningAddress*+ **-T** *tranlogDir*+ [**-I** *logFile*] [**-s** *stoneName*] [**-u**] [**-d**] [**-v**] [**-h**] [**-C** *certFileName*] [**-J** *certAuthFileName*] [**-K** *keyFileName*] [**-Q** *passphrasestring*] [**-S**] [**-V**]
- A** *listeningAddress* The address (es) on which the logsender will listen for connections from logreceivers. At least one is required, up to 4 may be specified.
- d** Print debug tracing of tranlog read and write operations to log file. The log file will be much larger.
- I** *logFileOrDir* The path and filename or directory for the logged output of the logsender process. The default is `/opt/gemstone/log/logsender_listeningPort.log`
- h** Displays a usage line and exits.
- P** *listeningPort* The port or named service on which on which the logsender will listen for connections from logreceivers.
- s** *stoneName* The name of the master stone. *stoneName* is required for the logsender to detect that stone has written new data to the master tranlogs. Without stone name, logsender will transmit tranlogs that it finds on the file system when it starts up.
- T** *tranlogDir* Directory(s) where the master stone's transaction logs are located. Normally the same as stone's STN_TRAN_LOG_DIRECTORIES but may also include archive directories. logsender will examine these directories for new files to send. At least one is required, up to 20 -T may be specified.
- u** Do not compress tranlog files at record level before transmission.
- v** Print version and exit.
- Additional arguments for SSL:
- C** *certFileName* Certificate in PEM format that will be sent to the peer upon request.
- J** *certAuthFileName* Certificate authority (CA) file in PEM format to use for peer certificate verification.
- K** *keyFileName* Private key in PEM format for the certificate (-C option).
- Q** *passphrasestring* Private key passphrase. Required if the -K option is used and the private key is encrypted.
- S** Enable SSL mode. Must be specified to use any other SSL options and must also be specified when starting the peer process.
- V** Disable verification of the peer's certificate.

This utility starts up a logsender process. As part of a hot standby setup, the logsender process runs on the master system to transmit transaction log records, as they are generated, to a logreceiver process that is running on a slave system.

The logsender writes logging output to a file with the path and name or in the directory specified by the `-l` argument. If this file is a directory, the file name is `logsender_listeningPort.log`. If no `-l` argument is used, it writes to `/opt/gemstone/log/logsender_listeningPort.log`. If a file with the specified name exists, it is appended to. It is not deleted on process exit.

The logsender process will continue running until explicitly stopped using the **stoplogsender** command. If connection to the stone is lost, it will attempt to reconnect.

gslist will report running logsender processes.

To start the logsender using SSL to establish a secure connection between logsender and log receiver, both `startlogsender` and `startlogreceiver` can be provided with SSL credentials. When SSL arguments are provided, `startlogsender` will start the logsender in SSL mode.

startnetldi

startnetldi	<i>[netLdiName]</i> [-l logFile] [-t timeout] [-a name] [-P portNumber] [-A addresses]+ [-n] [-g -s] [-d] [-h] [-v]
<i>netLdiName</i>	The name or port number of the GemStone network server. If <i>netLdiName</i> is a numeric value equal to or less than 65535, it is interpreted as a port number. If the given port is in use, it will result in an error. Other values are interpreted as the netldi name. If the -P argument is omitted, this name is looked up in the network services database to determine the port number. Network resource syntax is not permitted. The default is <code>gs64ldi</code> .
-a name	Captive account; all child processes created by the NetLDI will belong to the account named <i>name</i> . By default, child processes belong to the client's account.
-A addresses	Address to listen on. Up to 10 arguments are accepted. If no -A arguments are provided, listening is on the default wildcard address <code>::</code> . If this default wildcard address is included, then other -A arguments are ignored. If the -A entry arguments do not include <code>::</code> or <code>::1</code> , then <code>::1</code> is also listened on.
-d	Debug mode; inserts more extensive messages in the log file.
-g	Guest mode; no accesses are authenticated. This option is not allowed if the NetLDI's effective user id is the root account.
-h	Displays a usage line and exits.
-l logFile	The logged output of the NetLDI; the default is <code>/opt/gemstone/log/NetLdiName.log</code>
-P portNumber	The well-known port number that netldi will use.
-n	Do not allow any <i>ad hoc</i> processes to be created (ad hoc processes are ones not listed in <code>\$GEMSTONE/sys/services.dat</code>).
-s	Secure; require authentication for all NetLDI accesses.
-t timeout	Seconds to wait for a spawned client to start, such as a Gem; default is 30 seconds.
-v	Print version and exit

This command starts a GemStone network server with the specified *netLdiName* and *timeout* (given in seconds). The server spawns GemStone processes in response to login requests from remote applications and requests for remote repository access from Stone repository monitor processes. If your machine is slow or heavily loaded, and RPC logins time out before completing, specify a larger timeout value.

The NetLDI listens for requests, including RPC login requests, on a specified or configured port and optionally on a specific address. During the RPC login process, it uses a separate set of ports to establish communication between the Gem and its client. If you are running

over a firewall, you can specify the port using the **-P** options, and configure your firewall to permit access via this port.

To start the NetLDI for password authentication, make sure that `$/GEMSTONE/sys/netldid` is owned by root and has the S bit set. Issue this command (on some operating systems, you may have to issue it as root):

```
% startnetldi
```

Alternatively, to start the NetLDI in guest mode (authentication is not required), make sure `$/GEMSTONE/sys/netldid` does NOT have the S bit set. Log in as the captive account *name*, then issue this command:

```
% startnetldi -g -aname
```

For information about authentication modes, see “How To Arrange Network Security” on page 75.

For assistance with startup failures, refer to “To Troubleshoot NetLDI Startup Failures” on page 99.

startstone

startstone	<i>[gemStoneName]</i> [- I <i>logFile</i>] [- e <i>exeConfig</i>] [- z <i>systemConfig</i>] [- h] [- R][- N][- v]
<i>gemStoneName</i>	Name of the GemStone repository monitor, default is <code>gs64stone</code> . Network resource syntax is not permitted.
-I <i>logFile</i>	The location (or filename) for the logged output of the stone; the default is (1) the setting of the <code>GEMSTONE_LOG</code> environment variable, if defined; (2) <code>\$(GEMSTONE)/data/<i>gemStoneName</i>.log</code>
-e <i>exeConfig</i>	The GemStone executable-dependent configuration file (see page 268).
-z <i>systemConfig</i>	The GemStone system configuration file (see page 266).
-h	Displays a usage line and exits.
-R	Start up from the most recent checkpoint and go into restore mode. This allows transaction logs to be restored. Used when restoring from backup.
-N	Do not replay transaction logs as part of startup. Unless used with the -R option, and transaction logs are replayed manually, work may be lost.
-C	startup for conversion; only applies when starting up a version that requires conversion. Refer to the Installation Guide for the specific version for details.
-v	Print version and exit

This command opens the GemStone repository specified by the relevant configuration files. The three arguments *gemStoneName*, **-e** *exeConfig*, and **-z** *systemConfig* determine which configuration files **startstone** reads. (For more information, see “How GemStone Uses Configuration Files” on page 266.)

Legal characters in a Stone name are A-Z, a-z, 0-9, and the characters period, underscore, and dash (. _ -). Stone names with other characters are disallowed.

The **-N** option is intended for use when the repository needs recovery, but the transaction logs specified in the configuration file cannot be found or have become corrupted. The **-N** forces the repository to start up anyway, even though transactions committed since the last checkpoint will be lost. A new transaction log will be initialized as part of the startup.

The **-R** option is used when restoring from backup. This starts up the repository at the point of the most recent checkpoint in the extents, and puts the repository into restore mode. You can then invoke Topaz to restore from transaction logs and commit the restored state.

If the extents against which Stone is being started require recovery, or if you are restoring from online extent backups, then you must specify the **-N** option along with the **-R** option; otherwise, an error will result and the repository monitor will not start.

For startup failures, refer to “To Troubleshoot Stone Startup Failures” on page 95.

statmonitor

statmonitor *stoneName* **-f** *fileName* [*options*]

<i>stoneName</i>	Required; the name of the Stone to monitor.
-A	Collect all available system statistics.
-C	Collect system statistics for each individual CPU.
-D	Collect system statistics for all disks and partitions
-f <i>fileName</i>	The output file name. By default, the output filename is statmon <i>N</i> .out, where <i>N</i> is the process ID. If this file already exists, statmonitor will not start. To send output to stdout instead of a file, specify -f stdout .
-i <i>interval</i>	The interval in seconds (default: 20). Select either -i or -I .
-I <i>intervalMs</i>	The interval in milliseconds (default 20000; minimum 100). Select either -i or -I .
-J	Sample the Stone, shared cache monitor, and page manager only.
-h <i>hours</i>	The number of hours (default: forever).
-m <i>stoneHostName</i>	The Stone host name (default: localHost).
-n <i>numAppStats</i>	The number of application statistics (default: 0).
-N	Collect system statistics for all network interfaces.
-p <i>sessionId</i>	A GemStone sessionId to monitor. You may specify multiple sessions. (Default: monitor all sessions.)
-P	Sample the Stone, shared cache monitor and all AIO page servers only.
-q	Quiet mode. Only print messages if an error occurs.
-Q <i>pid1,pid2,...</i>	Record statistics for a list of process IDs.
-r	Restart a new output file when the current one completes. Each file will be given a unique name. This option may only be used with the -h or -t switches, which control when a restart is done.
-s0	(deprecated) Same as -Y
-s1	(ignored)
-s2	(deprecated) Same as -D
-s3	(deprecated) Same as -D -N
-s4	(deprecated) Same as -A
-S	Sample only the Stone and shared cache monitor.

- t** *times* The maximum number of samples to collect before starting a new output file (default is forever). Select either **-h** or **-t**.
- u** *seconds* The maximum number of seconds to wait before flushing the cached information to the output file (default: 60). If 0 then the flush will be done every interval.
- v** Print version and exit
- W** Collect system statistics for system memory pages.
- U***uid1,uid2,...* Record statistics for all processes owned by one or more UNIX user IDs.
- X** Run in host-only mode. Sample host system statistics only and do not attach to a shared page cache. May be combined with other flags EXCEPT: **-m**, **-n**, **-p**, **-P**, **-S**, or **-Y**.
- Y** Disable collection of all system statistics, including per-process data.
- z** Write the output in compressed gzip format.

Record statistics for a repository and/or the operating system to a disk file. **statmonitor** runs in the background, sampling specified data at a specified interval and recording the data to a text file. The data in this file can be viewed by the graphical application VSD. For more information on VSD, see page 334 and the *VSD Users Guide*.

Statistics are collected from the shared page cache. Only data for GemStone processes attached to the shared page cache on the host on which statmonitor is running are collected. To monitor Gems on systems with remote Gem servers, you must run statmonitor both on the Stone's machine and on the Gem server machine.

The *stonename* argument is used to specify the cache to monitor, both for caches that are local to the Stone and caches that are remote from the Stone with that name. Configurations in which a single machine is hosting remote caches for multiple stones that are running with the same Stone name are ambiguous and will not work correctly; this configuration is strongly discouraged.

stoplogreceiver

- stoplogreceiver** **-P** *listening port* [**-v**] [**-h**] [**-C** *certFileName*] [**-J** *certAuthFileName*] [**-K** *keyFileName*] [**-Q** *passphrasestring*] [**-S**] [**-V**]
- h** Displays a usage line and exits.
- P** *port* The port that the logreceiver is listening on for **stoplogreceiver** connections. If **-p** was used in the **startlogreceiver** command, that port must be specified here; otherwise the port specified by the **-P** argument to **startlogreceiver**.
- v** Print version and exit.

Additional arguments for SSL:

- C** *certFileName* Certificate in PEM format that will be sent to the peer upon request.
- J** *certAuthFileName* Certificate authority (CA) file in PEM format to use for peer certificate verification.
- K** *keyFileName* Private key in PEM format for the certificate (**-C** option).
- Q** *passphrasestring* Private key passphrase. Required if the **-K** option is used and the private key is encrypted.
- S** Enable SSL mode. Must be specified to use any other SSL options and must also be specified when starting the peer process.
- V** Disable verification of the peer's certificate.

This utility stops a logreceiver process that was started by a previous **startlogreceiver** command. This may only be executed on the same node as the logreceiver is running.

If the logreceiver is operating in SSL mode, the SSL credentials (certificates, etc) are required to stop the logreceiver.

stoplogsender

- stoplogsender** **-P** *listening port* **[-v]** **[-h]** **[-C** *certFileName* **]** **[-J** *certAuthFileName* **]** **[-K** *keyFileName* **]** **[-Q** *passphrasestring* **]** **[-S]** **[-V]**
- h** Displays a usage line and exits.
- P** *port* The port on which this logsender is listening. It must be the same port as specified in the **-P** argument of the **startlogsender**.
- v** Print version and exit.
- Additional arguments for SSL:
- C** *certFileName* Certificate in PEM format that will be sent to the peer upon request.
- J** *certAuthFileName* Certificate authority (CA) file in PEM format to use for peer certificate verification.
- K** *keyFileName* Private key in PEM format for the certificate (**-C** option).
- Q** *passphrasestring* Private key passphrase. Required if the **-K** option is used and the private key is encrypted.
- S** Enable SSL mode. Must be specified to use any other SSL options and must also be specified when starting the peer process.
- V** Disable verification of the peer's certificate.

This utility stops a logsender process that was started by a previous **startlogsender** command. This may only be executed on the same node as the logsender is running.

If the logsender is operating in SSL mode, the SSL credentials (certificates, etc) are required to stop the logsender.

stopnetldi

stopnetldi *[netLdiName]* [-h][-v]

netLdiName The name of the GemStone network server; the default is gs64ldi.
Network resource syntax is not permitted.

-h Displays a usage line and exits.

-v Print version and exit

Gracefully stop a GemStone network server. The argument to this command is optional, and are not needed for a standard GemStone configuration.

stopstone

stopstone	<i>[gemStoneName [gemStoneUserName [gemStonePassword]]]</i> [-i] [-t] [-h]
<i>gemStoneName</i>	The name of the GemStone repository monitor, commonly <code>gs64stone</code> . Network resource syntax is not permitted.
<i>gemStoneUserName</i>	A privileged GemStone user account name, such as “DataCurator” or “SystemUser”.
<i>gemStonePassword</i>	The GemStone password for the specified <i>gemStoneUserName</i> .
-i	Causes any current GemStone sessions to be terminated immediately.
-h	Displays a usage line and exits.
-t	Specifies how long to wait for other processes to detach from the cache. Default is -1, wait forever.

Gracefully shut down a GemStone repository monitor. In the process, a checkpoint is performed in which all committed transactions are written to the extents. If you specify the **-i** (immediate) option, the repository monitor is shut down even if there are GemStone users logged in. The **-t** option specifies how long to wait for all session to detach from the cache before returning; if omitted, **stopstone** will wait forever. If you do not supply the *gemStoneName*, *gemStoneUserName*, and *gemStonePassword* on the command line, this command will prompt you for them.

Because this command uses a Gem session process to connect to *gemStoneName*, it fails if the Gem is unable to connect for any reason, such as inability to attach a shared page cache. For assistance, refer to “To Troubleshoot Session Login Failures” on page 103.

topaz

topaz [-r]	[-i] [-n <i>netLdiName</i>] [-q] [-I <i>topazini</i>] [-u <i>useName</i>] [-h]
topaz -l	[-i] [-n <i>netLdiName</i>] [-e <i>exeConfig</i>] [-z <i>systemConfig</i>] [-T <i>tocSizeKB</i>] [-q] [-I <i>topazini</i>] [-u <i>useName</i>] [-h]
-e <i>exeConfig</i>	The GemStone executable-dependent configuration file (applies only to linked sessions). See page 268.
-h	Displays a usage line and exits
-i	Ignore the initialization file, <i>.topazini</i> .
-I <i>topazini</i>	Specify a complete path and file to a topazini initialization files, and use this rather than any <i>.topazini</i> in the default location.
-l	Invoke the linked version of Topaz.
-n <i>netLdiName</i>	The name of the GemStone network server; the default is (1) the setting of the GEMSTONE_NRS_ALL environment variable; (2) <i>gs641di</i>
-q	Start Topaz in quiet mode, suppressing printout of the banner and other information.
-r	Invoke the RPC (remote procedure call) version of Topaz.
-T <i>tocSizeKB</i>	The GEM_TEMPOBJ_CACHE_SIZE that will be used. Overrides any settings provided in configuration files passed as arguments with the -e or -z options. Only applies to linked sessions.
-u <i>useName</i>	The value is not used by the topaz executable, but may be useful in identifying processes in OS utilities such as <i>top</i> or <i>ps</i> .
-z <i>systemConfig</i>	The GemStone system configuration file (applies only to linked sessions). See page 266.
-v	Print version and exit
-w	On Windows client only; forces terminal behavior regardless of I/O device.

This command invokes various forms of topaz. The default is to invoke the remote procedure call (RPC) version of topaz; to do this explicitly, use the **-r** option. To invoke the linked version of topaz, which is recommended or required for some maintenance operations, use the **-l** option. Several topaz arguments only apply in linked topaz. The arguments **-e***exeConfig* and **-z***systemConfig* determine the configuration files that **topaz -l** reads. For more information about this, see “How GemStone Uses Configuration Files” on page 266.

Settings within topaz can allow linked topaz to perform RPC logins.

Topaz is available for Windows, in addition to regular server platforms, as part of the Windows Client installation. Options are more limited since linked logins are not possible.

For further information, see the *Topaz Programming Environment*.

vsd

vsd [**-h**] [*statmonDataFile*]

-h Displays a usage line and exits

statmonDataFile Specify the name of a data file to load into the VSD that is started.

This command starts VSD, the graphical tool to view and analyze statmonitor data. The default is to open an empty instance of VSD, into which you can then load one or multiple statmonitor data files. After loading data files, you may select one or multiple GemStone processes, and view charts on one or multiple statistics for these processes.

VSD is available for Windows, in addition to regular server platforms, as part of the Windows Client installation.

For a complete description of VSD, see the *VSD User's Guide*.

waitstone

waitstone	[<i>gemStoneName</i> <i>netLdiName</i>] [<i>timeout</i>]
<i>gemStoneName</i>	The name of the GemStone repository monitor.
<i>netLdiName</i>	The name of the GemStone netldi service.
<i>timeout</i>	How many minutes to wait for GemStone to initialize before reporting a problem. The default (0) means wait forever; -1 means don't wait, try once and return the result. Only valid when specifying either <i>gemStoneName</i> or <i>netLdiName</i> .

This command reports whether the GemStone repository *gemStoneName* is ready to accept logins or whether *netLdiName* is ready to accept requests. If neither *gemStoneName* nor *netLdiName* are specified, **waitstone** will report on the default stone name, *gs64stone*.

During the first 10 seconds, **waitstone** tests every two seconds to see if the Stone or NetLDI is ready; thereafter, a new connection is attempted once every 10 seconds. When the service is ready, **waitstone** issues a message to `stdout`. If the service is not ready by the time the specified number of minutes have elapsed, **waitstone** reports an error.

You may specify the *gemStoneName* and *netLdiName* arguments as a GemStone network resource string. (See Appendix C, "Network Resource String Syntax.")

This command returns 0 exit status if the operation is successful; otherwise, it returns a non-zero value.

waitstone does not have a `help` argument, but you can type:

```
man waitstone
```

to display the online manual page.

Network Resource String Syntax

This appendix describes the syntax for network resource strings. A network resource string (NRS) provides a means for uniquely identifying a GemStone file or process by specifying its location on the network, its type, and authorization information. GemStone utilities use network resource strings to request services from a NetLDI.

Overview

One common application of NRS strings is the specification of login parameters for a remote process (RPC) GemStone application. An RPC login typically requires you to specify a GemStone repository monitor and a Gem service on a remote server, using NRS strings that include the remote server's hostname. For example, to log in from Topaz to a Stone process called "gs64stone" running on node "handel", you would specify two NRS strings:

```
topaz> set gemstone !@handel!gs64stone
topaz> set gemnetid !@handel!gemnetobject
```

Many GemStone processes use network resource strings, so the strings show up in places where command arguments are recorded, such as the GemStone log file. Looking at log messages will show you the way an NRS works. For example:

```
Opening transaction log file for read,
filename = !@oboe#dbf!/user1/gemstone/data/tranlog0.dbf
```

An NRS can contain spaces and special characters. On heterogeneous network systems, you need to keep in mind that the various UNIX shells have their own rules for interpreting these characters. If you have a problem getting a command to work with an NRS as part of the command line, check the syntax of the NRS recorded in the log file. It may be that the shell didn't expand the string as you expected.

NOTE

Before you begin using network resource strings, make sure you understand the behavior of the software that will process the command.

See each operating system's documentation for a full discussion of its own rules about escaping certain characters in NRS strings that are entered at a command prompt.

If there is a space in the NRS, you can replace the space with a colon (:), or you can enclose the string in quotes (" "). For example, the following network resource strings are equivalent:

```
% waitstone !@oboe#auth:user@password!gs64stone
% waitstone "!@oboe#auth user@password!gs64stone"
```

Defaults

The following items uniquely identify a network resource:

- Communications protocol— such as TCP/IP
- Destination node— the host that has the resource
- Authentication of the user— such as a system authorization code
- Resource type— such as server, database extent, or task
- Environment— such as a NetLDI, a directory, or the name of a log file
- Resource name— the name of the specific resource being requested.

A network resource string can include some or all of this information. In most cases, you need not fill in all of the fields in a network resource string. The information required depends upon the nature of the utility being executed and the task to be accomplished. Most GemStone utilities provide some context-sensitive defaults. For example, the Topaz interface prefixes the name of a Stone process with the **#server** resource identifier.

When a utility needs a value for which it does not have a built-in default, it relies on the system-wide defaults described in the syntax productions in "Syntax" on page 339. You can supply your own default values for NRS modifiers by defining an environment variable named GEMSTONE_NRS_ALL in the form of the *nrs-header* production described in the Syntax section. If GEMSTONE_NRS_ALL defines a value for the desired field, that value is used in place of the system default. (There can be no meaningful default value for "resource name.")

A GemStone utility picks up the value of GEMSTONE_NRS_ALL as it is defined when the utility is started. Subsequent changes to the environment variable are not reflected in the behavior of an already-running utility.

When a client utility submits a request to a NetLDI, the utility uses its own defaults and those gleaned from its environment to build the NRS. After the NRS is submitted to it, the NetLDI then applies additional defaults if needed. Values submitted by the client utility take precedence over those provided by the NetLDI.

Notation

Terminal symbols are printed in boldface. They appear in a network resource string as written:

#server

Nonterminal symbols are printed in italics. They are defined in terms of terminal symbols and other nonterminal symbols:

username ::= *nrs-identifier*

Items enclosed in square brackets are optional. When they appear, they can appear only one time:

address-modifier ::= [*protocol*] [*@ node*]

Items enclosed in curly braces are also optional. When they appear, they can appear more than once:

nrs-header ::= ! [*address-modifier*] {*keyword-modifier*} !

Parentheses and vertical bars denote multiple options. Any single item on the list can be chosen:

protocol ::= (**tcp** | **serial** | **default**)

Syntax

nrs ::= [*nrs-header*] *nrs-body*

where:

nrs-header ::= ! [*address-modifier*] {*keyword-modifier*} [*resource-modifier*]!

All modifiers are optional, and defaults apply if a modifier is omitted. The value of an environment variable can be placed in an NRS by preceding the name of the variable with "\$". If the name needs to be followed by alphanumeric text, then it can be bracketed by "{" and "}". If an environment variable named `foo` exists, then either of the following will cause it to be expanded: `$foo` or `${foo}`. Environment variables are only expanded in the *nrs-header*. The *nrs-body* is never parsed.

address-modifier ::= [*protocol*] [*@ node*]

Specifies where the network resource is.

protocol ::= (**tcp** | **serial** | **default**)

Supports heterogeneous connections by predicating address on a network type. If no protocol is specified, `GCI_NET_DEFAULT_PROTOCOL` is used. On UNIX hosts, this default is **tcp**.

node ::= *nrs-identifier*

If no node is specified, the current machine's network node name is used. The identifier may also be an Internet-style numeric address. For example:

```
!@120.0.0.4#server!cornerstone
```

nrs-identifier ::= *identifier*

Identifiers are runs of characters; the special characters `!`, `#`, `$`, `@`, `^` and white space (blank, tab, newline) must be preceded by a "^". Identifiers are words in the UNIX sense.

keyword-modifier ::= (*authorization-modifier* | *environment-modifier*)

Keyword modifiers may be given in any order. If a keyword modifier is specified more than once, the latter replaces the former. If a keyword modifier takes an argument, then the keyword may be separated from the argument by a space or a colon.

authorization-modifier ::= (**#auth** | **#encrypted**) [:] *username* [@ *password*])

#auth specifies a valid OS user name on the target network. A valid OS user password is needed only if the resource type requires authentication. **#encrypted** is used by GemStone utilities. This type of authorization is the default.

username ::= *nrs-identifier*

If no OS user name is specified, the default is the current OS user. See the earlier discussion of *authorization-modifier*.

password ::= *nrs-identifier*

Only needed if the resource type requires authentication; see the earlier discussion of *authorization-modifier*.

environment-modifier ::= (**#netldi** | **#dir** | **#log**) [:] *nrs-identifier*

#netldi causes the named NetLDI to be used to service the request. If no NetLDI is specified, the default is `gs64ldi`. When you specify the **#netldi** option, the *nrs-identifier* (page 339) is either the name of a NetLDI service or the port number at which a NetLDI is running.

#dir sets the default directory of the network resource. It has no effect if the resource already exists. If a directory is not set, the pattern “%H” (home directory) is used. (See the definition of *nrs-identifier* on page 339.)

#log sets the name of the log file of the network resource. It has no effect if the resource already exists. If the log name is a relative path, it is relative to the working directory. If a log name is not set, the pattern “%N%P%M.log” (defined below) is used. (See the definition of *nrs-identifier* on page 339.)

The argument to **#dir** or **#log** can contain patterns that are expanded in the context of the created resource. The following patterns are supported:

%H	home directory
%M	machine's network node name
%N	executable's base name
%P	process pid
%U	user name
%%	%

resource-modifier ::= (**#server** | **#spawn** | **#task** | **#dbf** | **#monitor** | **#file**)

Identifies the intended purpose of the string in the *nrs-body*. An NRS can contain only one resource modifier. The default resource modifier is context sensitive. For instance, if the system expects an NRS for a database file, then the default is **#dbf**.

#server directs the NetLDI to search for the network address of a server, such as a Stone or another NetLDI. If successful, it returns the address. The *nrs-body* is a network server name. A successful lookup means only that the service has been defined; it does not indicate whether the service is currently running. A new process will not be started. (Authorization is needed only if the NetLDI is on a remote node and is running in secure mode.)

#task starts a new Gem. The *nrs-body* is a NetLDI service name (such as “gemnetobject”), followed by arguments to the command line. The NetLDI creates the named service by looking first for an entry in `$GEMSTONE/sys/services.dat`, and then in the user's home directory for an executable having that name. The NetLDI

returns the network address of the service. (Authorization is needed to create a new process unless the NetLDI is in guest mode.) The **#task** resource modifier is also used internally to create page servers.

#dbf is used to access a database file. The *nrs-body* is the file spec of a GemStone database file. The NetLDI creates a page server on the given node to access the database and returns the network address of the page server. (Authorization is needed unless the NetLDI is in guest mode).

#spawn is used internally to start the garbage collection and other service Gem processes.

#monitor is used internally to start up a shared page cache monitor.

#file means the *nrs-body* is the file spec of a file on the given host (not currently implemented).

nrs-body ::= unformatted text, to end of string

The *nrs-body* is interpreted according to the context established by the *resource-modifier*. No extended identifier expansion is done in the *nrs-body*, and no special escapes are needed.

GemStone Kernel Objects

This appendix describes the predefined objects that are located in a freshly installed GemStone/S 64 Bit repository.

Non-Numeric Constants

The following non-numeric constants are defined in the Globals dictionary and protected by the SystemObjectSecurityPolicy:

- **true** (an instance of Boolean)
- **false** (an instance of Boolean)
- **nil** (an instance of UndefinedObject)

Numeric Constants

Floating point constants are instances of class Float or class DecimalFloat. They are defined in the Globals dictionary and are protected by the SystemObjectSecurityPolicy. Refer to IEEE standards 754-1987 and 854-1987 for more information regarding their meanings in floating-point calculations.

DecimalPlusInfinity
DecimalMinusInfinity
DecimalPlusQuietNaN
DecimalMinusQuietNaN
DecimalPlusSignalingNaN
DecimalMinusSignalingNaN
PlusInfinity
MinusInfinity
PlusQuietNaN
MinusQuietNaN
PlusSignalingNaN
MinusSignalingNaN

Repository and GsObjectSecurityPolicies

SystemRepository. This single instance of Repository is defined in the Globals dictionary. Repository is a subclass of Collection, and the indexable part of SystemRepository contains references to all the security policies (all the instances of GsObjectSecurityPolicy) in GemStone.

The SystemRepository object initially contains eight security policies, three of which are public and named: the SystemObjectSecurityPolicy (owned by the SystemUser), the DataCuratorObjectSecurityPolicy (owned by the DataCurator), and the PublishedObjectSecurityPolicy (owned by SystemUser). The SystemRepository object itself is protected by the DataCuratorObjectSecurityPolicy. New GsObjectSecurityPolicies may be created and added to the SystemRepository object, using the methods `new`, `newInRepository:`, or by some methods that create new users.

For more on GsObjectSecurityPolicies, see the *Programming Guide for GemStone/S 64 Bit*.

SystemObjectSecurityPolicy. This security policy is defined in the Globals dictionary. For backwards compatibility, the key `#SystemSegment` also refers to this security policy, and may be used in upgraded repositories. The SystemObjectSecurityPolicy object itself is protected by the DataCuratorObjectSecurityPolicy.

The SystemObjectSecurityPolicy is the default security policy for its owner, the SystemUser (who has write authorization for any of the objects in this security policy). The “world” (that is, the set of all GemStone users) is authorized to read, but not write, the objects in this security policy. In addition, the group `#System` is authorized to write in this security policy.

DataCuratorObjectSecurityPolicy. This security policy is defined in the Globals dictionary. For backwards compatibility, the key `#DataCuratorSegment` also refers to this security policy, and may be used in upgraded repositories. The DataCuratorObjectSecurityPolicy object itself is protected by the DataCuratorObjectSecurityPolicy.

The DataCuratorObjectSecurityPolicy is the default security policy for its owner, the DataCurator (who has write authorization for any of the objects in this security policy). The “world” (that is, the set of all GemStone users) is authorized to read, but not write, the objects in this security policy. In addition, the `#DataCuratorGroup` is authorized to write in this security policy.

Objects in the DataCuratorObjectSecurityPolicy include the Globals dictionary, the SystemRepository object, most instances of GsObjectSecurityPolicy, AllUsers (the set of all GemStone UserProfiles), AllGroups (the collection of groups authorized to read and write objects in security policies), and each UserProfile object.

PublishedObjectSecurityPolicy. This security policy is defined in the Globals dictionary. For backwards compatibility, the key `#PublishedSegment` also refers to this security policy, and may be used in upgraded repositories. The PublishedObjectSecurityPolicy object itself is protected by the DataCuratorObjectSecurityPolicy.

The PublishedObjectSecurityPolicy is owned by the SystemUser. The group `#Subscribers` is authorized to read in this security policy. The group `#Publishers` is authorized to read and write in this security policy. The “world” is not authorized to read or write the objects in this security policy.

DbfHistory. DbfHistory is a String object that contains information regarding conversions and updates applied to the repository.

NativeLanguage. This Symbol is not used in this version, but may exist in upgraded repositories.

Global Variables and Collections

AllGroups. This CanonicalStringDictionary is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. Each Symbol in AllGroups corresponds to a group of users. When GemStone is first installed, AllGroups contains the symbols #System, #Publishers, #Subscribers, #DataCuratorGroup, and #SymbolUser.

AllUsers. The AllUsers object (a UserProfileSet) is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. AllUsers contains the UserProfiles of all GemStone users. When GemStone is first installed, AllUsers contains five UserProfiles: SystemUser, DataCurator, GcUser, SymbolUser, and Nameless.

For more information on the Special system accounts, see page 136.

AllDeletedUsers. This collection is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. AllDeletedUsers contains DeletedUserProfiles representing GemStone users that have been removed from AllUsers. When GemStone is first installed, AllDeletedUsers is empty.

AllClusterBuckets. This ClusterBucketArray is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. AllClusterBuckets contains instances of ClusterBucket, which group objects on extent pages to improve performance. When GemStone is first installed, AllClusterBuckets contains the following predefined cluster buckets (listed by cluster id):

1. A generic bucket whose extent is "don't care". This bucket, the current default after session login, is invariant and may not be modified.
2. A generic bucket whose extent is "don't care".
3. A generic bucket whose extent is "don't care".
4. The kernel classes "behaviorBucket", extent 1.
5. The kernel classes "descriptionBucket", extent 1.
6. The kernel classes "otherBucket", extent 1.
7. A generic bucket whose extent is "don't care".

ConfigurationParameterDict. This SymbolKeyValueDictionary is defined in the Globals dictionary, and is protected by the SystemObjectSecurityPolicy. Its keys list the names of the configuration parameters available to a session. Its values are only used internally in GemStone, to locate the values of the parameters themselves for an individual session.

DbfHistory. This String describes the history of this repository, from the version in which it was first created, and each subsequent upgrade.

- DbfOrigin.** This SmallInteger identifies the version in which this repository was first created.
- DeprecationEnabled.** This is nil, or a keyword. When deprecation is enabled, it will be set to the configured handling instructions when a deprecated method is encountered. Deprecation is described in the *Programmer's Guide for GemStone/S 64 Bit*.
- ErrorSymbols.** This SymbolDictionary is defined in the Globals dictionary, and is protected by the SystemObjectSecurityPolicy. It maps mnemonic symbols to error numbers.
- GciStructsMd5String** and **GciTsStructsMd5String.** These constants encode the GCI structures that this version of GemStone uses. On upgrade, they can be compared to determine if there are changes in the GCI structures that would impact C code using GemBuilder for C.
- GemStoneError.** This SymbolDictionary is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. Each key is a Symbol representing a native language, and is associated with an Array of error messages in that language. Initially, this dictionary contains the single key #English.
- GemStone_Legacy_Streams.** This SymbolDictionary contains classes implementing the legacy GemStone PositionableStream interface.
- GemStone_Portable_Streams.** This SymbolDictionary contains classes implementing the PositionableStream interface that is ANSI-complaint and portable to other Smalltalk dialects.
- InstancesDisallowed.** This IdentitySet is defined in the Globals dictionary, and is protected by the SystemObjectSecurityPolicy. This collection is used for error reporting for some cases where instance creation is disallowed.
- NotTranloggedGlobals.** This SymbolDictionary is defined in the Globals dictionary, and is protected by the DataCuratorObjectSecurityPolicy. This collection holds objects for which changes are committed but not recorded in the transaction logs. When GemStone is first installed, NotTranloggedGlobals is empty.
- SourceStringClass.** This is the default class used to instantiate literal strings, and to control the behavior of String comparison, including String equality. Changing the value should only be done on new repositories, since it may break existing collections of Strings and other usages that depend on string comparison.
- Transcript.** An instance of TranscriptStreamPortable.

Table D.1 Initial Contents of the Globals Dictionary

	Key	The object's class
Numeric Constants	#DecimalMinusInfinity	DecimalFloat
	#DecimalMinusQuietNaN	DecimalFloat
	#DecimalMinusSignalingNaN	DecimalFloat
	#DecimalPlusInfinity	DecimalFloat
	#DecimalPlusQuietNaN	DecimalFloat
	#DecimalPlusSignalingNaN	DecimalFloat
	#MinusInfinity	Float
	#MinusQuietNaN	Float
	#MinusSignalingNaN	Float
	#PlusInfinity	Float
	#PlusQuietNaN	Float
	#PlusSignalingNaN	Float
Non-Numeric Constants	#false	Boolean
	#nil	UndefinedObject
	#true	Boolean
Repository and instances of GsObjectSecurityPolicy	#DataCuratorObjectSecurityPolicy, #DataCuratorSegment	GsObjectSecurityPolicy
	#DbfHistory	String
	#DbfOrigin	SmallInteger
	#PositionableStream_position	String
	#PublishedObjectSecurityPolicy, #PublishedSegment	GsObjectSecurityPolicy
	#SystemRepository	Repository
Collections	#AllClusterBuckets	ClusterBucketArray
	#AllDeletedUsers	IdentitySet
	#AllGroups	CanonicalStringDictionary
	#AllUsers	UserProfileSet
	#ConfigurationParameterDict	SymbolKeyValueDictionary
	#ErrorSymbols	SymbolDictionary
	#GemStoneError	SymbolDictionary
	#GemStone_Portable_Streams	SymbolDictionary
	#GemStone_Legacy_Streams	SymbolDictionary
	#Globals	SymbolDictionary
	#InstancesDisallowed	IdentitySet
	#LegacyErrNumMap	Array
	#NotTranloggedGlobals	SymbolDictionary

Table D.1 Initial Contents of the Globals Dictionary (Continued)

	Key	The object's class
Other Customer-usable Globals	#DeprecationEnabled	Symbol
	#GciStructsMd5	String
	#GciTsStructsMd5	String
	#SourceStringClass	Class
	#Transcript	TranscriptStreamPortable
GemStone Internal Objects	#AsciiCollatingTable	ByteArray
	#ConversionReservedOopMap	Array
	#ConversionStatus	Array
	#DoubleByteAsciiCollatingTable	DoubleByteString
	#FdcResults	UndefinedObject
	#GcCandidates	UndefinedObject
	#GcCandidatesCount	UndefinedObject
	#GcHints	UndefinedObject
	#GcWeakReferences	Array
	#GemStoneRCLock	Object
	#GsCompilerClasses	SymbolDictionary
	#GsIndexingObjectSecurityPolicy, #GsIndexingSegment	GsObjectSecurityPolicy
	#ImageVersion	SymbolDictionary
	#ObsoleteClasses	SymbolDictionary
	#QuadByteAsciiCollatingTable	QuadByteString
	#RcBTreeNode	UndefinedObject
	#_remoteNil	UndefinedObject
	#SecurityDataObjectSecurityPolicy, #SecurityDataSegment	GsObjectSecurityPolicy
	#SharedDependencyLists	DepListTable
#VersionParameterDict	SymbolKeyValueDictionary	
plus all kernel classes		

Current TimeZone

Each instance of `DateTime` includes a reference to a `TimeZone` object, which handles the conversion from the internally stored Greenwich Mean Time (GMT, also referred to as UTC or Coordinated Universal Time) and the local time. `TimeZones` encapsulate the daylight savings time (DST) rules, so a given GMT time is adjusted to local time based on `TimeZone` and the specific date. `TimeZones` are also used to calculate the internal stored GMT for newly created `DateTime` instances.

Each session has a current `TimeZone` and a default `TimeZone`, which are used to display times, and in `DateTime` creation when methods that do not explicitly specify the `TimeZone` are used. These are installed as part of application installation or configuration; by default, the GemStone distribution has the `America/Los_Angeles` `TimeZone` installed. This is described in the *GemStone/S 64 Bit Installation Guide*.

GemStone uses the public domain **zoneinfo** database to create TimeZone, loading the information from platform and language independent source files. If the rules change for the TimeZone that your application uses, you must recreate the TimeZone instance from the source files. Depending on the nature of the rules change, you may also need to update references from DateTime instances to the new TimeZone instance, or possibly update the DateTime internal offsets.

There are a number of ways to create TimeZone instances for your application:

- **From the OS on Solaris or Linux.** On these operating systems, you can create the TimeZone instance based on the current machine configuration using:

```
newTZ := TimeZone fromOS
```

- **GemStone's time zone database.** Using the interactive script `tzselect`, you can determine the correct time zone descriptor name for your local time zone. With this, you can create the new TimeZone instance using the time zone database provided with GemStone.

```
newTZ := TimeZone fromGemPath: '$GEMSTONE/pub/timezone/etc/zoneinfo/Europe/Zurich'
```

or, if GemStone's time zone database is installed in the default location:

```
newTZ := TimeZone named: 'Europe/Zurich'
```

- **Your own time zone database.** With the time zone descriptor name for your TimeZone, you can specify the full path to the time zone information.

```
newTZ := TimeZone fromGemPath: yourPath, '/Europe/Zurich'.
```

You must then install this TimeZone instances as the current and default time zone.

Zoneinfo

The widely used public-domain time zone database, **ZoneInfo** or **tz**, contains code and data that records time zone information for locations worldwide. It is updated periodically when boundaries or rules change in any of the represented locations.

Each record in the tz database represents a location where all clocks are kept on the same time as each other throughout the year, coordinating any time adjustments such as DST, and have done so for many years. Locations are identified by continent (or ocean, for islands) and name, which is usually the largest city within the region. For example, `America/Los_Angeles`, `Europe/London`, etc.

tz is provided as text files, which may be compiled into binary files using tz's compilers. GemStone's TimeZone implementation uses the compiled binary form, which is also used by the Solaris and Linux operating systems. GemStone's files are based on `tzdata2006p.tar.gz`. To get updated source files, download from:

```
http://www.iana.org/time-zones
```

The timezone sources may be compiled using the `zic` timezone compiler, which GemStone provides as a convenience (see "zic" on page 350).

Utilities

tzselect, **zdump** and **zic** are public domain, open source utilities that are useful in working with the zoneinfo database. These utilities are provided with the Solaris and Linux operating systems; for the convenience of users on other operating systems, these utilities are provided along with the other zoneinfo database files.

NOTE

These are not GemStone utilities. Support for their use is not provided by GemStone.

You may download the source code for these utilities here:

<http://www.iana.org/time-zones>

Shipped files are based on `tzcode2006p.tar.gz`. Documentation for these utilities is provided as man pages. To read the man pages, add the directory `$GEMSTONE/pub/timezone/man` to the `MANPATH`.

To run these, you may wish to add `$GEMSTONE/pub/timezone/etc` to the executable path.

tzselect

`tzselect` allows you to interactively select a time zone. The interactive script asks you a series of questions about the current location and outputs the resulting time zone description to standard output. The output is suitable as a value for the `TZ` environment variable and GemStone scripts.

You may need to set the environment variable `$TZDIR` to `$GEMSTONE/pub/timezone/etc/zoneinfo` (or the path to your zoneinfo database, for this script to work correctly. You may also need to set the environment variable `$AWK`, to any POSIX compliant `awk` program.

For further details on using `tzselect` see the man page.

zdump

```
zdump [-v] [-c cutoffyear] [zonename...]
```

`zdump` prints time zone information. It prints the current time for each time zone (`zonename`) listed on the command line.

Specifying an invalid zone name to `zdump` does NOT return an error; instead, it returns the `zdump` output for GMT. This reflects the same behavior of the time routines in `libc`.

The `-v` option will display the entire contents of the time zone database for the given time zone name.

For further details on using `zdump`, including the command line options, see the man page.

zic

```
zic [-s] [-v] [-l localtime] [-p posixrules] [-d directory]
    [-y yearistype] [filename...]
```

`zic` compiles time zone source files. It reads input text in files named on the command line, and creates the time zone binary files.

To create files in a specific location, rather than the standard platform directory (on Solaris, `/usr/share/lib/zoneinfo`), use the `-d` *directory* option.

For example, to recompile sources on Solaris to the GemStone timezone database, execute the following:

```
zic -d $GEMSTONE/pub/timezone/etc/zoneinfo/  
/usr/share/lib/zoneinfo/src/northamerica
```

For further details on using `zic`, including the command line options and the structure of the source code files, see the man page for `zic`.

This appendix lists the environment variables used by GemStone/S 64 Bit. The list has two parts: variables intended for public use, and variables that are reserved for internal use.

Public Environment Variables

The following environment variables are intended for use by customers. The variable GEMSTONE is required; the others may be useful in particular situations.

GEMSTONE

The location of the GemStone Object Server software, which must be a full path, beginning with a slash, such as `/user3/GemStone-hppa.hpux`.

GEMSTONE_ADMIN_GC_LOG_DIR

The directory location for Admin Gem logs. By default, Admin Gem log files are created in the same directory as the Stone log, which is `$GEMSTONE/data`. You can set this environment variable to specify that Admin Gem log files are created in a different directory.

GEMSTONE_EXE_CONF

The location of an executable-dependent configuration file; see “Creating an Executable Configuration File” on page 269.

GEMSTONE_GLOBAL_DIR

The location for the global GemStone logs and locks file, overriding the default `/opt/gemstone/`. This directory must already exist.

This directory controls visibility between GemStone processes, and must be the same for all GemStone processes that may want to interact with a given repository, including `stone`, `gems`, `topaz`, `statmonitor`, `netldi`, `gslis`, etc. GemStone processes that do not share a common location for `/gemstone/locks` – either `/opt/gemstone`, `/usr/gemstone`, or a directory specified by `$GEMSTONE_GLOBAL_DIR` – operate as if they are in independent name spaces.

GEMSTONE_KEEP_LOG

To keep a process’s log from being deleted when the process terminates normally,

unset this variable in the appropriate script, such as
`unset GEMSTONE/sys/gemnetobject.`

GEMSTONE_LIB

Specifies the directory for the gem and gem dynamic libraries. This is primarily of use in debugging low-level problems, and is used by the `gemnetdebug` script to specify the slow gem and gem dynamic libraries.

GEMSTONE_LOG

The location of system log files for the Stone repository monitor and its child processes. For further information, see “GemStone System Logs” on page 113.

GEMSTONE_MAX_FD

Limits the number of file descriptors requested by a GemStone process. For further information, see “Estimating File Descriptor Needs” on page 29.

GEMSTONE_NRS_ALL

Sets a number of network-related defaults, including the type of user authentication that GemStone expects. For further information, see “To Set a Default NRS” on page 78.

GEMSTONE_PAGE_MGR_LOG_DIR

The directory location for Page Manager Gem logs. By default, Page Manager Gem log files are created in the same directory as the Stone log, which is `$GEMSTONE/data`. You can set this environment variable to specify that Page Manager Gem log files are created in a different directory.

GEMSTONE_RECLAIM_GC_LOG_DIR

The directory location for all Reclaim Gem logs. By default, Reclaim Gem log files are created in the same directory as the Stone log, which is `$GEMSTONE/data`. You can set this environment variable to specify that Reclaim Gem log files are created in a different directory.

GEMSTONE_SPCMON_STARTUP_TIMELIMIT

Internally, GemStone waits for five minutes for the shared page cache to start up and initialize. This environment variable overrides this timeout, and specifies the time, in seconds, that the Stone will wait for the shared page cache to startup before giving up.

GEMSTONE_SSL_LIB_DIR

A directory location in which to look for the SSL shared library.

GEMSTONE_SYMBOL_GEM_LOG_DIR

The directory location for SymbolGem logs. By default, Symbol Gem log files are created in the same directory as the Stone log, which is `$GEMSTONE/data`. You can set this environment variable to specify that Symbol Gem log files are created in a different directory.

GEMSTONE_SYS_CONF

Location of a system-wide configuration file; see “How GemStone Uses Configuration Files” on page 266.

GS_CFTIME

If defined, it should contain a date and time format string that overrides the GemStone LOCALE-based default. See “Localizing timestamps in log files” on page 119.

GS_CORE_TIME_OUT

If `GS_WRITE_CORE_FILE` is defined, this is the number of seconds to wait before a

catastrophically failing GemStone/S process writes a core file and terminates – by default, 60 seconds. To determine the cause of a problem, GemStone/S Technical Support needs a stack trace, which is usually written to the process log file prior to the process shutdown.

If you need to derive a stack trace directly from a failing (but not yet terminated) process by attaching a debugger to it, you can set this variable to increase the time available to attach the debugger.

GS_DEBUG_PAM

If this variable is set to any value, PAM debugging information will be printed to stdout.

GS_DEBUG_SHARED_MEM

If this variable is set to any value, the shared page cache monitor process will print extra debugging to its log file.

GS_DEBUG_SSL_LOG_DIR

If this variable is set to a directory, a process that logs in RPC will write output of SSL calls made during to a file named `GsSSLDebug_<pid>.log` in the specified directory. This file may get very large.

GS_DEBUG_VMGC_MKSW_MEMORY_USED_SOFT_BREAK

At the end of each mark/sweep, if the percent of memory used is greater than the threshold specified by this variable, a SoftBreak (error 6003) is generated, and the threshold is raised by 5 percent. We suggest a setting of 75%.

GS_DEBUG_VMGC_MKSW_PRINT_STACK

The mark/sweep count at which to begin printing the Smalltalk stack at each mark/sweep.

For this and all other `GS_DEBUG_VMGC_*` environment variables, the printout goes to the "output push" file of a linkable Topaz (topaz -l) session, for use in testing your application. If that file is not defined, the printouts go to standard output of the session's gem or topaz -l process.

GS_DEBUG_VMGC_MKSW_PRINT_C_STACK

The mark/sweep count at which to begin printing the C stack at each mark/sweep. This variable is very expensive, consuming 2 seconds plus the cost of `fork()` for each printout.

GS_DEBUG_VMGC_PRINT_MKSW

The mark/sweep count at which to begin printing mark/sweeps.

GS_DEBUG_VMGC_PRINT_MKSW_MEMORY

The mark/sweep count at which to begin printing detailed memory usage (20 lines) for each mark/sweep.

GS_DEBUG_VMGC_PRINT_MKSW_MEMORY_USED

Specifies when Smalltalk stack printing starts as the application approaches OutOfMemory conditions. At the end of each mark/sweep, if the percent of memory used is greater than the threshold specified by this variable, the mark/sweep is printed, the Smalltalk stack is printed, and the threshold is raised by 5 percent. In a situation producing an OutOfMemory error, you should get several Smalltalk stacks printed in the Gem log file before the session dies. The default setting is 75%.

GS_DEBUG_VMGC_PRINT_SCAV

The scavenge count at which to begin printing scavenges. Once this takes effect, all mark/sweeps will also be printed. Be aware that printing scavenges can produce large quantities of output.

GS_DEBUG_VMGC_PRINT_TRANS

Print transaction boundaries (begin/commit/abort) in the log file.

GS_DEBUG_VMGC_SCAV_PRINT_STACK

The scavenge count at which to begin printing the Smalltalk stack at each scavenge. Be aware that this print activity can produce large quantities of output.

GS_DEBUG_VMGC_SCAV_PRINT_C_STACK

The scavenge count at which to begin printing the C stack at each scavenge. This variable is very expensive, consuming 2 seconds plus the cost of fork() for each printout.

GS_DEBUG_VMGC_VERBOSE_OUTOFMEM

Automatically call the primitive for `System class>>_vmPrintInstanceCounts:0` when an OutOfMemory error occurs, and also print the Smalltalk stack. (For details about this method, see the comments in the image.) This applies to each Gem or linkable Topaz (topaz -l) process that you subsequently start.

GS_DEBUG_VMGC_VERIFY_MKSW

The mark/sweep count at which to begin verifying object memory before and after each mark/sweep.

GS_DEBUG_VMGC_VERIFY_SCAV

The scavenge count at which to begin verifying object memory before and after each scavenge. Once this takes effect, `GS_DEBUG_VMGC_VERIFY_MKSW` will also be in effect. Be aware that this activity uses significant amounts of CPU time.

GS_DISABLE_CHARACTER_TABLE_LOAD

Any value disables the loading of any customized Extended Character Set Character Data Tables. This feature is deprecated; see the *Programming Guide* for details.

GS_DISABLE_KEEPALIVE

A non-empty string disables the network keepalive facility. For further information about keepalive, see “Disrupted Communications” on page 74

.GS_DISABLE_WARNING

A non-empty string disables a warning that GemStone is using `/opt/gemstone` instead of `/usr/gemstone` for log and lock files when both directories exist. Use of `/usr/gemstone` is only for compatibility with previous releases; the default location is `/opt/gemstone`.

GS_DISABLE_SHMDT

Disables the system call to `shmdt()` made on cache detach during logout.

GS_DISABLE_SIGNAL_HANDLERS

When this environment variable is enabled in the environment for a gem process, by setting it to any value, this gem sessions will not attempt to handle a SIGSEGV, SIGBUS or SIGILL signal, but will shut down immediately. It will not generate a core nor print C stacks. This avoids side effects with user action or client Smalltalk code.

GS_GSLIST_TIME_FORMAT

This can be set to UNIX-style date format string, to allow the output of the gslist utility to be displayed in a parseable format.

GS_PAGE_MGR_PRINT_REMOTE_STACKS

If this variable is set, if a remote cache page server becomes stuck, the page manager will request that the remote cache page server print its call stack to its log file.

GS_WRITE_CORE_FILE

By default, core files are not created when a fatal error occurs. (The C level stack trace is written to the process log file prior to the process shutdown.) You can set this environment variable if you need a core file.

upgradeLogDir

The location for log files produced during the upgrade of a repository for a new version of GemStone.

System Variables Used by GemStone

GemStone uses the following system variables that exist for other purposes:

EDITOR	Used by Topaz to determine which editor to invoke.
PATH	The search path of locating executable files.
SHELL	Used to determine what shell to use for an exec, such as by <code>System class>>performOnServer:.</code>

Reserved Environment Variables

The following environment variables are reserved for internal use. Customers should not define these variables for use with GemStone unless specifically instructed to do so. Please refrain from using these variables for other purposes.

GCIRTL_BASELIBNAME**GEMSTONE***

All environment variable names beginning with "GEMSTONE" other than those above are reserved.

GS_*

All environment variable names beginning with "GS_" other than those above are reserved.

NT_PARENT_PID

gs64ldi

runpgsvr

Object State Change Tracking

F.1 Overview

GemStone transaction logs (tranlogs), which provide a way to recover all changes made to the repository in the case of repository crashes, or allow warm standbys to apply changes made to a primary repository, include detailed records of all changes in an encoded and compressed form. Converting the tranlog information to human-readable form, and analyzing this output, provides invaluable information for debugging and testing. It allows us to determine exactly what changes have been made to any of the objects in the repository over time. The tools used to perform this have been used internally to GemStone for many years, and have been provided to customers on occasion when needed to analyze specific application problems.

The ability to track changes to objects may be useful for customers who need to identify details on changes that have been made to application data objects. For this reason, we are making these scripts available as part of the GemStone product. Additional information has been added to the tranlogs to allow tracking of the specific user or machine that originated the object changes.

The scripts used to perform tranlog analysis are located in the `$GEMSTONE/bin/` directory and are named:

`printlogs` — to output the complete contents (optionally filtered) of selected tranlogs in human-readable form.

`searchlogs` — to search all tranlogs in a directory for selected OOPs (Object Oriented Pointers, or Object Ids), and output the matching entries (optionally filtered) in human readable form.

A GemStone repository performs many automatic operations, including things like garbage collection and checkpoints, that are transparent to end users. The tranlogs, of course, must contain records of any changes made by these operations. Complete details on everything that tranlogs may contain is beyond the scope of this documentation. The information provided here is intended to allow the use of the tranlog analysis scripts to locate and identify the details of changes to application objects.

Object oriented design's principle of encapsulation allows you to hide internal object complexity. However, to understand the data recorded in the tranlogs, you must have a detailed understanding of the actual structure of the objects. This includes both your own application classes, and the classes that are part of GemStone Smalltalk.

Also, note that since the tranlog analysis scripts are general purpose, used for a wide variety of purposes in which a detailed record of internal repository operation is required, the scripts may output much more information than is necessary for tracking object state changes. You may need to ignore this extraneous information as you perform your analysis.

F.2 Tranlog Analysis Scripts

Script Prerequisites

The environment variable `$GEMSTONE` must be set, and the `$GEMSTONE/bin/` directory must be in your executable path.

The scripts perform the analysis on tranlogs that are in the current working directory. If you are using raw partitions for your tranlogs, locate disks on the file system with adequate space for both the tranlogs themselves, and the script output files, which may be larger than the original tranlogs. Use `copydbf` to copy the tranlogs from the raw partition to the file system. For more information on the `copydbf` utility, see Appendix B, "GemStone Utility Commands."

Output

Output from the scripts goes directly to `stdout`. To preserve the output and allow it to be used in later steps of analysis, redirect this; for example:

```
$> printlogs tranlog1.dbf > tranlog1.out
```

Note that these scripts can produce very large amounts of output, so make sure that you have adequate disk space.

In some cases the resulting files may be too large for unix text editors such as `vi` or `emacs` to open. You may find it necessary to use the unix `split` utility to break up very large output files into more manageable chunks.

Tranlog Assumptions

By default, the tranlogs are assumed to be named using the GemStone convention `tranlogNNN.dbf`. If you are using a different naming convention, you can override this by setting the environment variable `$GS_TRANLOG_PREFIX` to the prefix you are using.

The scripts use the creation date of the tranlog file to determine the order in which the tranlogs are analyzed. If you copy or manipulate the tranlogs in a way that changes the creation date, this may cause the analysis to be done out of order. The output will warn of this with the message:

```
*** Warning: scanning tranlogs out of order
```


Filter Criteria

The scripts both allow you to filter the results, to locate entries that are related to a particular `UserId`, `GemHost`, or `ClientIP`.

UserId – The `userId` (user name) of the `UserProfile` associated with this session: `DataCurator`, `SystemUser`, etc. The filter keyword is `user`.

GemHost – The name or IP address of the host machine running the gem process. For a linked session, which links the gem into the client, this is the same machine as the client.

If the gem is running on the same machine as the stone, use the name of the host machine. Otherwise, if the gem is on a different machine than the stone, use the IP address of the remote machine.

The filter keyword is `host`.

ClientIP – The IP address of the host machine running the client process.

For an RPC session, this is the machine running the client application. Clients may be `topaz -r`, `GemBuilder` for `Smalltalk`, or `GemBuilder` for C applications. For a linked session, this is the machine running the linked client/gem (the same machine as the `GemHost`). However, the `ClientIP` is always the IP address, even if it is on the same machine as the stone.

The filter keyword is `client`.

Effective UNIX user ID – The integer that is the effective UNIX user id of the gem process.

The filter keyword is `euid`.

UNIX user ID – The integer that is the real UNIX user ID of the gem process.

The filter keyword is `ruid`.

effective UNIX user name – The effective UNIX user name of the gem process.

The filter keyword is `euidstr`.

UNIX user name – The real UNIX user name running the gem process.

The filter keyword is `ruidstr`.

process ID – The integer process id (PID) of the gem process.

The filter keyword is `gempid`.

session ID – The integer session id of the session within GemStone.

The filter keyword is `sessionid`.

WARNING

*Information about `UserId`, `GemHost`, and `ClientIP` are derived from a **Login** tranlog entry created when a session first logs in. This entry associates the `UserId/GemHost/ClientIP` with a particular `sessionID`, which is then used as a key for subsequent tranlog entries. If you start analysis from a later tranlog which does not include this **Login** entry, these fields will be left blank for that session,*

*and printouts/searches using filters based on these fields will not locate any results. Likewise, scanning through tranlogs out of order may result in the wrong **Login** entry being associated with a given sessionID. This would set UserId/GemHost/ClientIP incorrectly for that particular session, and produce incorrect results when filtering.*

printlogs

This script prints out the contents of one or more tranlogs in the current working directory in a human-readable form.

Warning

This script produce a very large amount of output, which (unfiltered) will exceed the size of original tranlog/s, and depending on the contents may be twice as large as the original tranlogs. Consider disk space, the use of filters, and restricting the set of tranlogs before running this script. Use caution in including the `full` keyword.

Usage:

```
printlogs [<filters>] [full] [all] [<tlogA> ... <tlogZ>]
```

If `<filters>` are specified, only print out the tranlog entries that match the specified criteria. Filters may be combined. Possible filters are:

```
user <userId> - Filter by the userId (the user name) of the GemStone UserProfile
host <hostnameOrIP> - Filter by gem/topaz process host or IP address
client <N.N.N.N> - Filter by client IP Address
eid <integer> - Filter by gem's effective UNIX user ID
ruid <integer> - Filter by gem's real UNIX user ID
eidstr <string> - Filter by gem's effective UNIX user name
ruidstr <string> - Filter by gem's real UNIX user name
gempid <integer> - Filter by gem's process ID
sessionid <integer> - Filter by gem's session ID
```

`full` – Produce more detailed information. WARNING: this produces much larger output results.

`all` – Print out the contents of all tranlogs in the current working directory.

Examples

To print out the entire contents of all tranlogs in this working directory:

```
printlogs all
```

To print out all entries in a selected number of tranlogs (note that tranlogs in the sequence must be contiguous):

```
printlogs tranlog5.dbf tranlog6.dbf tranlog7.dbf
```

To print out all tranlog entries for the user DataCurator in any tranlog:

```
printlogs user DataCurator all
```

To print out detailed information for all entries in tranlog5.dbf for the user DataCurator:

```
printlogs full user DataCurator tranlog5.dbf
```

searchlogs

This script prints out tranlog entries associated with particular OOP values, according to the arguments in the command line. All tranlogs in the current working directory are scanned.

Usage:

```
searchlogs [<filters>] <oopA> ... <oopB>]
```

If <filters> are specified, only print out the tranlog entries that match the specified criteria. Filters may be combined. Possible filters are:

```
user <userId> – Filter by the userId (the user name) of the GemStone UserProfile
host <hostnameOrIP> – Filter by gem/topaz process host or IP address
client <N.N.N.N> – Filter by client IP Address
euid <integer> – Filter by gem's effective UNIX user ID
ruid <integer> – Filter by gem's real UNIX user ID
euidstr <string> – Filter by gem's effective UNIX user name
ruidstr <string> – Filter by gem's real UNIX user name
gempid <integer> – Filter by gem's process ID
sessionid <integer> – Filter by gem's session ID
```

When using more than one filter, you must list the filters in the listed order.

Examples

To print out all entries involving OOP 1234 and OOP 5678:

```
searchlogs 1234 5678
```

To print out all entries involving OOP 1234 performed by DataCurator:

```
searchlogs user DataCurator 1234
```

To print out all entries involving OOP 1234 and OOP 5678 performed by DataCurator while logged in from client machine 10.20.30.40:

```
searchlogs user DataCurator client 10.20.30.40 1234 5678
```

F.3 Tranlog Structure

In order to make sense of the output from the scripts, you need a basic understanding of how tranlogs are structured.

GemStone transaction logs consist of a sequence of **tranlog records**. Tranlog records are written on **physical pages** of 512 bytes (note that this is different from the larger page size used for extents); a tranlog record may extend over more than one page, but its size is always a multiple of 512 bytes.

Each tranlog record contains one or more **tranlog entries** (sometimes referred to internally as logical records). The tranlog entries contain the information of interest - the actual changes to objects in the repository (and any other repository operations).

Output from the scripts will include header information for the tranlog record (see Example F.1), followed by data from each of the tranlog entries held by that tranlog record.

Example F.1 Tranlog Record Header Output

```
Dump of record 11 hdr.pageId 42949672962 , Kind=(16)Tran Log Record
thisRecordId:(file:2 rec:11) previousRecordId:(file:2 rec:10)
recordSize: 1024 numLogicalRecs: 2 fileCreationTime: 1297200655
```

Tranlog records are identified by the pageId that they begin on. Since a tranlog record may extend over multiple pages, there may be a gap in the sequence of record ids in the output. For example, the first tranlog record in Example F.2, record 7, has a recordSize of 1536 (three 512-byte physical pages), and so the next tranlog record will be 10.

Example F.2 Gap in Tranlog Record Sequence

```
Dump of record 7 hdr.pageId 25769803778 , Kind=(16)Tran Log Record
thisRecordId:(file:2 rec:7) previousRecordId:(file:2 rec:6)
recordSize: 1536 numLogicalRecs: 2 fileCreationTime: 1297200655
```

```
2.7.0 BeginData session: 4 beginId:(100979 1)
  numObjs:3  pad1:0  pad2:0
2195969  240385  20423937
```

```
2.7.1 Commit session: 4 beginId:(100979 1)  crPage: 1039  commitSeq: 0.158
timeWritten: 02/08/11 13:31:52 PST
  seqType:normal beginLogRecord:(file:2 rec:7)
-----
```

```
Dump of record 10 hdr.pageId 30064771074 , Kind=(16)Tran Log Record
thisRecordId:(file:2 rec:10) previousRecordId:(file:2 rec:7)
recordSize: 512 numLogicalRecs: 3 fileCreationTime: 1297200655
```

Tranlog Entries

Each tranlog entry contains a unique identifying code, a descriptive entry type, and information on the session that originated the tranlog entry.

The identifying code consists of three numbers in the form:

```
AAA.BBB.CCC
```

where:

```
AAA - the fileId of the tranlog holding the entry
BBB - the pageId for the tranlog record holding the entry
CCC - the entryId: the number of the entry within the tranlog record
```

Each tranlog entry is of a particular type, according to the event that it represents and the information it contains. Types are indicated by short descriptive terms such as **Login**, **Abort**, **Commit**, and **Data**. There are a large number of entry types, most of which are not important for tracking object state changes and can be ignored (for example, a **Checkpoint** entry. The ones that are important are documented below.

Each tranlog entry is associated with an Integer sessionID. SessionsIDs are unique to a specific session at any point in time. However, when a session logs out, the sessionID becomes available again for reuse by a new session logging in, so over a period of time, a sessionId may be used by a number of different sessions.

A sessionID of zero is used for tranlog entries generated by the stone.

If the entry was not originated from the stone (that is, the tranlog entries sessionID is not 0), the tranlog entry header includes more information about the session: the UserId, the GemHost, and the ClientIP address. These are described in more detail under “Filter Criteria” on page 361.

Example F.3 Example Tranlog Entry

```
2.3.4 Login session: 4 beginId:(100973 1) userProfileOop: 208385
timeWritten: 02/08/11 13:30:56 PST userId: SymbolUser gemhost: myhost
clientIP: 10.20.30.40 processId: 12663 rUid: 531 eUid: 531 realU: gsadmin
effU: gsadmin
```

```
2.3.5 StartSymbolGem session: 4 beginId:(100973 5)
```

Example F.3 shows that is in the tranlog with fileId 2 (by the default naming convention, tranlog2.dbf), it is in the third physical page and in tranlog record 3, and these are the fourth and fifth tranlog entries in tranlog record 3.

These entries are of the tranlog entry types Login and StartReclaimGcGem. The session is logged in as the user named SymbolUser; the gem session is executing on the same machine as the stone, a machine named myhost; the client is executing on a machine with the IP address 10.20.30.40; and the OS process is owned by the OS userid gsadmin. (beginId contains transaction tracking information that you can ignore).

Other information will follow this basic data, depending on the type of entry.

Tranlog Entry Types

There are a large number of tranlog entry types. Most of these are not relevant to tracking object state history - they record internal operations of the system, such as garbage collection or checkpoints. These tranlog entry types are not documented, although their functions can often be deduced by their names.

Below are the entries important for tracking object state history:

Login

A new session is logging into GemStone. As mentioned earlier, this entry logs the UserId, GemHost, ClientIP and other data for this particular sessionID. If you start analysis on a tranlog that follows this event, these fields will be left blank for that session.

For example, if session 7 logs in while tranlog4 is in effect, and logs out when tranlog5 is in effect, and you begin analysis on tranlog5, entries for this session will not contain any session detail information. If sessionID 7 is reused by a new session logging in later during tranlog5, that login will be recorded in tranlog5, and subsequent entries for this new session will be displayed properly.

AbortLogout

This entry is written when a session logs out or the Stone detects that the session has been killed. This entry is not flushed to the transaction log until a commit occurs.

BeginData

Data

BeginStoreData

StoreData

GemStone uses the above four entry types for recording new or changed object data. The basic entry information is followed by additional information about the changes, including the OOP values of the changed objects. Using the optional “full” flag in the `printlogs` script will cause the output to include additional information on the exact changes made.

Commit

All the changes recorded in earlier **BeginData**, **Data**, **BeginStoreData**, or **StoreData** entries are now officially committed to the repository.

Abort

Any changes recorded earlier in this transaction are discarded.

BreakSerialization

This entry indicates that a transaction conflict occurred during an attempt to commit. Any changes recorded earlier in this transaction are not yet permanent. This event is usually followed by an **Abort** entry, although it's possible that the next event seen for this sessionID is a **Login**, if the original session logged out and a new session reuses the sessionID.

Very Large Objects

GemStone is designed so that all objects will fit on a single page of 16384 bytes. This means that a byte object can be no larger than about 16000 bytes (since page header information uses some space), and pointer objects can only have about 4000 elements. GemStone internally represents objects larger than this as a tree structure, where the root node object references 2 or more leaf node objects, which then reference the actual elements of the collection object. Extremely large objects, such as large collections, may have internal branch nodes, if the number of leaf objects needed exceeds 4000.

This internal structure is usually transparent to the user. So, for example, you may create and manipulate an Array containing 20,000 elements as if it was a single large object, while the actual representation is a root object that references five leaf objects, each containing a 4000-element chunk of the array. While this makes application development with GemStone much simpler, the entries in the tranlogs reflect the actual implementation; you will need to be aware of this to understand tranlog output relating to collections larger than ~4000 object references or ~16000 bytes. Adding an element to the large Array, for example, may mean the tranlog entry includes a change to an instance of `LargeObjectNode` (the leaf node object), rather than a change to the Array itself.

Full vs. Normal Mode

When using the printlogs script, you can optionally specify “full” mode to get more details on the changes made to objects in the repository. But this will greatly increase the size of the resulting log files. For example, using normal mode on our test tranlogs generated a log file that contained an entry that looked like this:

Example F.4 Tranlog entry, normal mode

```
2.11.0 BeginData session: 5 beginId:(100978 1)
      numObjs:3 pad1:0 pad2:0
      25567233 8532993 25880577
```

which tells you that three objects were created or modified during this event, with oops 25567233, 8532993, and 25880577.

Using “full” mode will produce a much more detailed listing for this event:

Example F.5 Tranlog entry, full mode

```
2.11.0 BeginData session: 5 beginId:(100978 1)
      numObjs:3 pad1:0 pad2:0

objId 25567233 class 114177 segId:9 len 14 gcSz 14 psize 136 bits 0x1 page
1152
Oop values: 50(sI 6) 8532993 5042177 25881601 17979137 26620161
20272897 25801985
8: 20322817 25559553 20324097 25559297 20423937 25869569

objId 8532993 class 111361 segId:9 len 14 gcSz 14 psize 136 bits 0x1 page
1152
Oop values: 162(sI 20) 122(sI 15) 20002(sI 2500) 42(sI 5) 20 18224129 20
25567233
8: 20 25566977 20 25801729 20 25566721

objId 25880577 isNewObj=1 class 73985 segId:9 len 23 gcSz 23 psize 208
bits 0x241 page 1152
Oop values: 18(sI 2) 18(sI 2) 20 2(sI 0) 25870081 25871105 20 20
8: 20 20 20 20 20 20 20 20
16: 20 20 20 20 20 20 20
```

Note that there is now a description of each individual created or modified object, containing these fields:

objId: The OOP of this object

isNewObj: A Boolean indicating if the object is a newly created object.

class: The OOP of the class of this object

segId: The id of the GsObjectSecurityPolicy (formerly Segment) associated with this object
len: The logical size of this object (in bytes for a byte object, OOPs for an OOP object)
gcSz: The logical size of this object (in bytes for a byte object, OOPs for an OOP object)
psize: The physical size of this object on disk in bytes (including 20 bytes of object header information)
bits: The format bits for this object (internal GS use)
page: The extent page that this object is written on
Bytes: The actual bytes that make up this object (if a byte object)
 [or]
Oop values: The actual OOPs that make up this object (if an OOP object)

If the Oop values contain more than eight elements, they are broken into lines of eight items, each of which is prefixed by a counter. For example, an Array of 30 items might look like this:

```

objId 39056641 isNewObj=1 class 66817 segId:9 len 30 gcSz 30
psize 264 bits 0x201 page 1005
Oop values: 39056385 39056129 39055873 39055617 39055361
39055105 39054849 39054593
8: 39054337 39054081 39053825 39053569 39053313 39053057
39052801 39052545
16: 39052289 39052033 39051777 39051521 39051265 39051009
39050753 39050497
24: 39050241 39049985 39049729 39049473 39049217 39048961
  
```

Bytes are broken up similarly, but the sections are 60 bytes rather than 8. For example, the source code string for the name: method might look like this:

```

objId 38711809 isNewObj=1 class 74753 segId:9 len 91 gcSz 0 psize
120 bits 0x280 page 1097
Bytes: name: newValue

"Modify the value of the instance variabl
 61: e 'name'."
name := newValue
  
```

F.4 Example of Tranlog Analysis

For this example, we created a simple database containing some Employee information and performed some simple operations creating and modifying these Employee objects.

The structure of classes associated with the Employee data is as follows:

Employee:

name - a Name object
 age - a SmallInteger
 address - an Address object

Name:

last - a String object
 first - a String object
 middle - a String object

Address:

addr1 - a String object
 addr2 - a String object
 city - a String object
 state - a String object
 zip - a SmallInteger

After having User1 create five Employee objects (with associated Name and Address objects), we then had User1 and User2 make some minor changes to one of the Employees:

- User2 incremented the age after a birthday.
- User1 changed the address after a move.

When completed, we had two tranlogs, tranlog2.dbf and tranlog3.dbf.

Tracking Changes to an Employee

Let's say we want to examine the change history of a particular Employee. Using the method #asOop, we find the OOP of the Employee object of interest is 38808321.

```
topaz 1> printit
| myEmployee |
myEmployee := <code to locate employee object>.
myEmployee asOop.
%
38808321
```

The data composing an Employee is contained in subobjects (such as address), as well as directly (such as age). So, we will also need to track changes to these subobjects. Again using #asOop, we find that the OOP of the Name object associated with this Employee is 155073, and that the OOP for the Address object is 155069.

```
myEmployee name asOop
38808577
myEmployee address asOop
38808833
```

You can use #asOop on any persistent object in the repository. For example,

```
73 asOop
586
nil asOop
20
```

We can now search our tranlogs for any events involving these objects. Using the command:

```
$> searchlogs 38808321 38808577 38808833
```

results in the following output:

Example F.6 searchlogs output for Employee

Searching for oops: 38808321 38808577 38808833

Searching tranlogs:

```
tranlog2.dbf
tranlog3.dbf
```

... header material omitted ...

```
2.118.0 BeginData session: 5 beginId:(101033 0) newobj 38808321 cls
39017217 onPage 1081, newobj 38808577 cls 38726401 onPage 1081, newobj
38808833 cls 38715649 onPage 1081
```

```
2.123.0 Commit session: 5 beginId:(101033 0) timeWritten: 02/08/11
14:43:42 PST
```

```
2.215.0 BeginData session: 6 beginId:(101664 1) object 38808321 cls
39017217 onPage 871,
```

```
2.215.1 Commit session: 6 beginId:(101664 1) timeWritten: 02/08/11
16:14:48 PST
```

```
3.22.0 BeginData session: 8 beginId:(101779 1) object 38808833 cls
38715905 onPage 900,
```

```
3.22.1 Commit session: 8 beginId:(101779 1) timeWritten: 02/08/11
16:32:58 PST
```

From this output, we can see that in tranlog entry 2.118.0, Session 5 made changes to all three objects (in this case, when the Employee and associated subobjects were first created). In entry 2.215.0, session 6 made a change to the Employee, and then later in entry 3.22.0, session 8 made a change to 38808833, the Address object.

Note that the **BeginData** entries are each followed by a **Commit**. You should always confirm that a **BeginData/Data/BeginStoreData/StoreData** of interest is followed by a **Commit**. If it doesn't then the reported event was not made persistent in the repository.

Changed vs. new objects

In the above example, while the field of an Address object changed, the Address object itself was the same (had the same OOP). Depending on how the Smalltalk application is written, this may not always be the case. If application that was initiating these changes created a new Address object, and assigned the Employee's address instance variable to

this new object, then the Employee object would reference a new OOP, rather than OOP 38808833. This would make the analysis somewhat different. For example, in the initial stage of the analysis when you look up the OOP of the Address object in your application, you would find the new OOP rather than the original OOP. Looking back in time, you would see when this Address object was created and assigned to the Employee instance.

Details of Changes to an Employee

Having used the `searchlogs` script to get a general idea of which tranlogs are of interest, you can now use the `printlogs` script to get more details.

For example, you might want more details on the creation of Employee object 38808321 and its associated subobjects 38808577 and 38808833 in entry 2.118.0. The "2" in "2.118.0" indicates that `tranlog2.dbf` is the tranlog of interest. The command:

```
$> printlogs tranlog2.dbf
```

will generate a condensed listing of all events in `tranlog2.dbf`. By searching the resulting file for the entry number 2.118.0 you can find the relevant entry:

Example F.7 Employee modification

```
3.61.0 BeginData session: 2 user: User1 gemhost: merlin clientIP:
10.20.30.40 beginId:(134 0)
  clusterId: 1, extentId: 0 numObjs:35

145021 155041 155045 155049 155053 155069 155073 155077 155089
155093 155097 155101 155121 155273 155285 155317 156589
156597 156609 156613 156661 156689 156729 156733 156749
156825 156829 156833 156837 156857 156861 156885 179629
180045 180653
```

This shows the oops of *all* objects created during this event. If you want to see more details on the actual changes made, use the "full" argument in the `printlogs` command:

```
$> printlogs full tranlog3.dbf
```

This will produce a more detailed listing of all events. For tranlog entry 3.61.0, you'll find:

Example F.8 Example Employee modifications, output in full mode

```
2.118.0 BeginData session: 5 beginId:(101033 0)
  numObjs:11 pad1:0 pad2:0

[details for other objects omitted]

objId 38808321 isNewObj=1 class 39017217 segId:9 len 3 gcSz 3 psize 48
bits 0x201 page 1081
Oop values: 38808577 306(sI 38) 38808833

objId 38808577 isNewObj=1 class 38726401 segId:9 len 3 gcSz 3 psize 48
bits 0x201 page 1081
Oop values: 38809089 38809601 20
```

```

objId 38808833 isNewObj=1 class 38715649 segId:9 len 5 gcSz 5 psize 64
bits 0x201 page 1081
Oop values: 38809857 20 38810113 38810369 777714(sI 97214)

objId 38809089 isNewObj=1 class 74753 segId:9 len 7 gcSz 0 psize 32 bits
0x280 page 1081
Bytes: Patrick

objId 38809601 isNewObj=1 class 74753 segId:9 len 5 gcSz 0 psize 32 bits
0x280 page 1081
Bytes: Ohara

objId 38809857 isNewObj=1 class 74753 segId:9 len 13 gcSz 0 psize 40 bits
0x280 page 1081
Bytes: 2556 Fir Blvd

objId 38810113 isNewObj=1 class 74753 segId:9 len 7 gcSz 0 psize 32 bits
0x280 page 1081
Bytes: Ashford

objId 38810369 isNewObj=1 class 74753 segId:9 len 2 gcSz 0 psize 32 bits
0x280 page 1081
Bytes: OR

```

So, for the Employee object 38808321, we find:

```

objId 38808321 isNewObj=1 class 39017217 segId:9 len 3 gcSz 3
psize 48 bits 0x201 page 1081
Oop values: 38808577 306(sI 38) 38808833

```

Indicating that this is an instance of class 39017217, the Employee class, which has three instance variables: name, age, and address. By position, we can identify the data in the instance variables:

- name: 38808577 - the OOP of an instance of Name, found later in the entry
- age: 306 - the OOP of the SmallInteger (sI) 38
- address: 38808833 - the OOP of an instance of Address, found later in the entry

Looking at the Name object 38808577 we find:

```

objId 38808577 isNewObj=1 class 38726401 segId:9 len 3 gcSz 3
psize 48 bits 0x201 page 1081
Oop values: 38809089 38809601 20

```

Indicating that this is an instance of the class with OOP 38726401 (Name). Name contains three instance variables, last, first, and middle. By position, we see the data is:

- last: 38809089 - the OOP of a String, described below
- first: 38809601 - the OOP of a String, described below
- middle: 20 (the OOP of nil) - in this example, no middle name was set

For the last name object 38809089, we find:

```
objId 38809601 isNewObj=1 class 74753 segId:9 len 5 gcSz 0 psize
32 bits 0x280 page 1081
Bytes: Ohara
```

Indicating the last name is the string "Ohara".

By a similar process you can follow the trail of objects to examine the structure of other subobjects in the Name and Address objects.

F.5 Further Analysis

Class Operations

To find all objects created or modified that belong to a particular class, first generate `printlogs` output in full mode of the tranlogs of interest. Each time an object of that class is created or modified, the full tranlog entry includes the line

```
class <OOP>
```

Use the unix `grep` command to find all references to the OOP of the class.

For example, to find all creation or modification of any instance of our example class `Employee` in the `printlogs` full output.

Since the `Employee` class is OOP 39017217, execute the `grep` command:

```
$> grep "class: 39017217" tranlogfull.txt
```

this will produce output of the form:

```
objId 38804481 isNewObj=1 class 39017217 segId:9 len 3 gcSz 3
psize 48 bits 0x201 page 1087
objId 38808321 isNewObj=1 class 39017217 segId:9 len 3 gcSz 3
psize 48 bits 0x201 page 1081
objId 40062465 isNewObj=1 class 39017217 segId:9 len 3 gcSz 3
psize 48 bits 0x201 page 1317
objId 38808321 class 39017217 segId:9 len 3 gcSz 3 psize 48 bits
0x1 page 871
objId 38808321 class 39017217 segId:9 len 3 gcSz 3 psize 48 bits
0x1 page 900
```

This gives the first line from the entry creating/modifying the object belonging to that class. From this, you can use other commands to search for and/or track the history of these objects.

Deleted Objects

An object-oriented system doesn't actually delete objects; objects cease to be referenced and are eventually garbage-collected. Noting the removal of an object requires examining the references to that object (such as from a collection) and identifying when the referencing object was modified in such a way that the object of interest is no longer referenced. Meanwhile, as you examine the `printlogs` output, you may find references to the OOP of a dereferenced object in garbage collection tranlog entries.

Managing Volume

As noted above, the `printlogs` produce a very large amount of output. GemStone tranlogs may be multiple GB in size. The output of `printlogs` in normal mode will be somewhat larger than the original tranlog (the `printlogs` output, being human readable, is less dense). The output from this script in full mode is much larger.

To manage the volume:

- Avoid configuring your system with very large tranlogs.
- Ensure that you have plenty of disk space available before beginning analysis.
- Print only the tranlogs containing data you need. Use the `searchlogs` script to identify exactly where the required information is located.
- Make sure that only the relevant tranlogs are in the current directory; move the unneeded ones elsewhere. However, you must retain a continuous set of tranlogs without gaps in sequence, and you must include the tranlog with the original log entry, in order to have the `UserId` and other information provided.
- Once you have printed the output, use the UNIX utility `grep -n` to locate the lines of interest, and the UNIX utility `split -l` to break the resulting file up into more manageable size chunks.

Index

A

- ad hoc processes 77, 324
- adding
 - a new user 143
 - a user to a group 148
 - extent 171
 - extent, while running 172
 - symbolList to user 151
 - transaction log at runtime 187
 - user privileges 150
- Admin Gem 239
 - configuring 256
- AIO page server 55
 - defined 23
 - log files 107, 114, 117
- AllClusterBuckets 345
 - see also the *Programming Guide*
- AllDeletedUsers 144, 345
- AllGroups 345
 - adding a group to 148
- allocation
 - of extents, weighted 38
 - of objects to new extents 173, 174
- AllUsers 345
- application, linked vs. RPC 58
- assigning privileges to a user 143
- audit
 - object audit 123
 - page audit 120
- auth modifier, NRS 75
- authentication scheme
 - setting to GemStone 155
 - setting to LDAP 156
 - setting to UNIX 155
- authentication scheme, determining 158
- authorization
 - and object security policies 138, 146
 - list of an object security policy, removing a

- group from 147
- automatic garbage collection 227–252

B

- backups
 - checkpoint at start of 196
 - compressed 199
 - errors in full backup 198
 - for warm or hot standby 217
 - full backup 192
 - GcLock 198
 - interactions with garbage collection 198, 199
 - making full backup 196–199
 - making offline extent snapshot 193–194
 - making online extent snapshot 194–196
 - monitoring full backup 199
 - of transaction logs 186
 - offline extent snapshot 192
 - online extent snapshot 192
 - performance of full backup 198
 - restoring
 - to point in time 211
 - restoring from extent snapshot 203–204
 - restoring from full backup 204–207
 - restoring transaction logs into 207–208
 - transaction mode and 198
 - using copydbf 193
 - using operating system facilities 193
 - verifying readability 199
 - with repository online 196
- baseDn: (login authentication) 156
- beginTransaction (System) 111

C

- cache statistics
 - programmatically access 127
 - see also under *statistics*

- setting gem name in 130
- statistics names 128
- user defined session 131
- user-defined global session 131
- cache warming 289, 319
- cacheName
 - setting name in cache for gem 130
- captive account mode, NetLDI
 - guest mode with 76
- ChangeUserId (privilege) 140
- checkpoint
 - defined 169
 - determining status 194
 - frequency of 53
 - identifying in transaction logs 313
 - operating system backup and 193
 - resuming 195
 - starting 188
 - Stone shutdown and 54
 - suspending 194
- clock, system 30
- clustering
 - new extents and 174, 175
 - restoring backups and 201
 - see also the *Programming Guide*
- code
 - region of temporary object memory 227
- CodeModification (privilege) 139, 140
- commands
 - copydbf** 310
 - gslis**t 314
 - pageaudit** 316
 - pstack** 317
 - removedbf** 318
 - startcachewarmer** 319
 - startlogreceiver** 221, 223, 320
 - startlogsender** 220, 223, 322
 - startnetldi** 324
 - startstone** 326
 - statmonitor** 327
 - stoplogreceiver** 221, 224, 329
 - stoplogsender** 220, 223, 330
 - stopnetldi** 331
 - stopstone** 332
 - topaz** 333
 - vsd** 334
 - waitstone** 335
- commit
 - disable user 154
 - re-enable user 154
 - test if user is enabled to 154
- commit record
 - defined 235
- commit record backlog 110, 303
 - defined 235
- committed objects
 - in local object memory 227
- communications, disrupted 74
- CompilePrimitives (privilege) 140
- compiler and symbol resolution 141
- compressed backups
 - creating 199
- concurrent I/O 292
- configuration
 - access at run time
 - Gem 64
 - Stone 50
 - extent locations 35
 - file descriptors 34
 - Gem session processes 59
 - kernel resources 29
 - memory needs for server 28
 - multiple extents 37
 - raw partitions 47
 - shared page cache 31
 - single-host 58
 - Stone private page cache 33
 - swap space needs 28
 - system resources 28
 - transaction logs 42
 - tuning 52
- configuration files
 - examining parameters from Smalltalk 50, 64
 - executable 265
 - for server (Stone) 23
 - for sessions 59
 - naming options 271
 - option value errors in 272
 - options, warning messages and 271
 - printing summary of all options 276
 - searching for 266
 - syntax errors in 272
 - syntax of 271
 - system-wide 265
 - topaz and 270
 - used by GemStone 266
- configuration options
 - and ConfigurationParameterDict system object 345
 - DBF_ALLOCATION_MODE 38, 274
 - DBF_EXTENT_NAMES 97, 274
 - DBF_EXTENT_SIZES 37, 180, 274

- DBF_PRE_GROW 37, 275
DBF_SCRATCH_DIR 276
DUMP_OPTIONS 276
GEM_ABORT_MAX_CRG 276
GEM_FREE_FRAME_CACHE_SIZE 276, 56
GEM_FREE_FRAME_LIMIT 55, 276
GEM_FREE_PAGEIDS_CACHE 277
GEM_GCI_LOG_ENABLED 277
GEM_HALT_ON_ERROR 277
GEM_KEEP_MIN_SOFTREFS 277
GEM_NATIVE_CODE_ENABLED 278
GEM_PGSRV_COMPRESS_PAGE_TRANSFERS
278
GEM_PGSRV_FREE_FRAME_CACHE_SIZE
279, 56
GEM_PGSRV_FREE_FRAME_LIMIT 279
GEM_PGSRV_UPDATE_CACHE_ON_READ 65,
279
GEM_PRIVATE_PAGE_CACHE_KB 280
GEM_REPOSITORY_IN_MEMORY 280
GEM_RPCGCI_TIMEOUT 280
GEM_RPC_KEEPALIVE_INTERVAL 280
GEM_RPC_USE_SSL 281
GEM_SMALLTALK_STACK_DEPTH 278
GEM_SOFTREF_CLEANUP_PERCENT_MEM
281
GEM_TEMPOBJ_AGGRESSIVE_STUBBING
281
GEM_TEMPOBJ_CACHE_SIZE 63, 229, 282
GEM_TEMPOBJ_MESPACE_SIZE 229, 282
GEM_TEMPOBJ_OOPMAP_SIZE 229, 283
GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE
283
GEM_TEMPOBJ_POMGEN_SCAVENGE_INTER
VAL 283
GEM_TEMPOBJ_POMGEN_SIZE 229, 284
GEM_TEMPOBJ_SCOPES_SIZE 284
KEYFILE 284
LOG_WARNINGS 285
SHR_NUM_FREE_FRAME_SERVERS 285
SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_
POLICY 285
SHR_PAGE_CACHE_LOCKED 286
SHR_PAGE_CACHE_NUM_PROCS 33, 63, 286,
29
SHR_PAGE_CACHE_NUM_SHARED_COUNTER
S 61, 286
SHR_PAGE_CACHE_PERMISSIONS 287
SHR_PAGE_CACHE_SIZE_KB 33, 63, 287
SHR_SPIN_LOCK_COUNT 53, 287
SHR_TARGET_FREE_FRAME_COUNT 288
SHR_WELL_KNOWN_PORT_NUMBER 72, 288
STN_ADMIN_GC_SESSION_ENABLED 288
STN_ALLOCATE_HIGH_OOPS 288
STN_ALLOW_NFS_EXTENTS 289
STN_CACHE_WARMER 289
STN_CACHE_WARMER_SESSIONS 289
STN_CHECKPOINT_INTERVAL 53, 289
STN_COMMIT_QUEUE_THRESHOLD 290
STN_COMMIT_RECORD_QUEUE_SIZE 290
STN_COMMITS_ASYNC 290
STN_COMMIT_TOKEN_TIMEOUT 290
STN_CR_BACKLOG_THRESHOLD 291
STN_DISABLE_LOGIN_FAILURE_LIMIT
167, 291
STN_DISABLE_LOGIN_FAILURE_TIME_LI
MIT 167, 291
STN_DISKFULL_TERMINATION_INTERVAL
179, 292
STN_EPOCH_GC_ENABLED 247, 292
STN_EXTENT_IO_FLAGS 292
STN_FREE_FRAME_CACHE_SIZE 293
STN_FREE_SPACE_THRESHOLD 293, 179
STN_GEM_ABORT_TIMEOUT 293
STN_GEM_LOSTOT_TIMEOUT 294
STN_GEM_TIMEOUT 294
STN_HALT_ON_FATAL_ERR 294, 44
STN_LISTENING_ADDRESSES 295
STN_LOGIN_LOG_ENABLED 296
STN_LOG_IO_FLAGS 295
STN_LOG_LOGIN_FAILURE_LIMIT 167, 296
STN_LOG_LOGIN_FAILURE_TIME_LIMIT
167, 296
STN_LOOP_NO_WORK_THRESHOLD 297
STN_MAX_AIO_RATE 297
STN_MAX_AIO_REQUESTS 297
STN_MAX_GC_RECLAIM_SESSIONS 298
STN_MAX_LOGIN_LOCK_SPIN_COUNT 298
STN_MAX_REMOTE_CACHES 298
STN_MAX_SESSIONS 33, 299
STN_MAX_VOTING_SESSIONS 299
STN_NUM_AIO_WRITE_THREADS 299
STN_NUM_GC_RECLAIM_SESSIONS 299
STN_NUM_LOCAL_AIO_SERVERS 55, 300
STN_OBJ_LOCK_TIMEOUT 300
STN_PAGE_MGR_COMPRESSION_ENABLED
300
STN_PAGE_MGR_MAX_WAIT_TIME 301
STN_PAGE_MGR_PRINT_TIMEOUT_THRESH
OLD 301
STN_PAGE_MGR_REMOVE_MAX_PAGES 301
STN_PAGE_MGR_REMOVE_MIN_PAGES 302

- STN_PRIVATE_PAGE_CACHE_KB 33, 302
 - STN_REMOTE_CACHE_PGSRV_TIMEOUT 302
 - STN_REMOTE_CACHE_TIMEOUT 302
 - STN_SHR_TARGET_PERCENT_DIRTY 303
 - STN_SIGNAL_ABORT_CR_BACKLOG 303
 - STN_SYMBOL_GC_ENABLED 303
 - STN_TRAN_FULL_LOGGING 42, 182, 303
 - STN_TRAN_LOG_DEBUG_LEVEL 304
 - STN_TRAN_LOG_DIRECTORIES 44, 304
 - STN_TRAN_LOG_LIMIT 54, 304
 - STN_TRAN_LOG_PREFIX 305
 - STN_TRAN_LOG_SIZES 43, 44, 305
 - STN_TRAN_Q_TO_RUN_Q_THRESHOLD 305
 - STN_WELL_KNOWN_PORT_NUMBER 72, 306
 - configuration parameters
 - garbage collection 254
 - configurationAt: (System) 50, 64, 306
 - configurationAt:put: (System) 306
 - ConfigurationParameterDict 345
 - continuous restore mode, for hot standby 221, 223
 - copydbf** command 310
 - archiving transaction logs 186
 - example
 - with raw partition 48
 - using with raw partitions 47
 - createExtent: (Repository) 172
 - createExtent:withMaxSize: (Repository) 173
 - creating
 - executable configuration files 269
 - extents 41
 - new user group 148
 - transaction logs 42, 182
 - current object security policy, exercise caution
 - when changing 147
 - currentLogDirectoryId (Repository) 187
 - currentLogFile (Repository) 187
 - currentSessionNames (System) 104, 105, 261
 - currentTranlogSizeMB (Repository) 186, 187
 - custom gem executable, installing 67
- D**
- DataCurator
 - and AllUsers system object 345
 - and DataCuratorObjectSecurityPolicy object 344
 - described 137
 - DataCuratorGroup 141
 - DataCuratorObjectSecurityPolicy 141, 344
 - Date and time formatting 119
 - DBF_ALLOCATION_MODE
 - (configuration option) 174
 - adding extents and 174, 175
 - effect when restoring backups 200
 - DBF_ALLOCATION_MODE (configuration option) 38
 - DBF_ALLOCATION_MODE (configuration option) 174, 274
 - DBF_EXTENT_NAMES (configuration option) 97, 171, 173, 174, 175, 274
 - DBF_EXTENT_SIZES (configuration option) 37, 171, 180, 181, 274
 - DbfHistory 345
 - DbfOrigin 346
 - DBF_PRE_GROW (configuration option) 37, 172, 173, 275
 - DBF_SCRATCH_DIR (configuration option) 276
 - dead object
 - contrasted with shadow object 235–238
 - defined 235
 - DeadNotReclaimedObjs 255
 - #deadObjsReclaimedCommitThreshold (GcUser parameter) 254
 - DeadObjsReclaimedCount 256
 - debugging
 - out-of-memory errors 230
 - DecimalMinusInfinity 343
 - DecimalMinusQuietNaN 343
 - DecimalMinusSignalingNaN 343
 - DecimalPlusInfinity 343
 - DecimalPlusQuietNaN 343
 - DecimalPlusSignalingNaN 343
 - default object security policy
 - changing a user's 147
 - defined 136, 138
 - described 136
 - exercise caution when changing 147
 - privilege required in UserProfile 147
 - DefaultObjectSecurityPolicy (privilege) 139
 - #deferReclaimCacheDirtyThreshold (GcUser parameter) 254
 - DeletedUserProfile 144
 - deleting user privileges 150
 - DeprecationEnabled 346
 - descriptionOfSession (System) 105, 261, 262
 - dictionaries
 - Globals 142
 - Published 142
 - Direct I/O 292
 - directory, current, of child process 79
 - disableCommits:

- (UserProfile) 154
 - disaster recovery 326
 - disk drives
 - I/O among multiple extents 38
 - multiple drives recommended 26
 - raw partitions 47
 - usage recommendations 26
 - disk failure 107
 - disk or repository full error 178
 - disk space
 - disk-full errors 178
 - distinguished name, and LDAP authentication 156
 - distributed system
 - and mid-level cache 88
 - DN (distinguished name), and LDAP 156
 - DUMP_OPTIONS (configuration option) 276
- ## E
- enableCommits:
 - (UserProfile) 154
 - enableEpochGc (System) 247
 - enableSignaledAbortError 111
 - enableSignaledFinishTransactionError 112
 - environment variables 353
 - GEMSTONE 353
 - GEMSTONE_ADMIN_GC_LOG_DIR 353
 - GEMSTONE_EXE_CONF 265, 268, 269, 270, 353
 - GEMSTONE_GLOBAL_DIR 353, 30
 - GEMSTONE_KEEP_LOG 353
 - GEMSTONE_LIB 354
 - GEMSTONE_LOG 114, 354
 - GEMSTONE_MAX_FD 354, 29, 60
 - GEMSTONE_NRS_ALL 354, 79
 - GEMSTONE_PAGE_MGR_LOG_DIR 354
 - GEMSTONE_RECLAIM_GC_LOG_DIR 354
 - GEMSTONE_SPCMON_STARTUP_TIMELIMIT 354
 - GEMSTONE_SYMBOL_GEM_LOG_DIR 354
 - GEMSTONE_SYS_CONF 265, 266, 354
 - GS_CFTIME 354
 - GS_CORE_TIME_OUT 354
 - GS_DEBUG_PAM 355
 - GS_DEBUG_SHARDE_MEM 355
 - GS_DEBUG_SSL_LOG_DIR 355
 - GS_DISABLE_CHARACTER_TABLE_LOAD 356
 - GS_DISABLE_KEEPALIVE 356
 - GS_DISABLE_SIGNAL_HANDLER 356
 - GS_DISABLE_WARNING 356
 - GS_PAGE_MGR_PRINT_REMOTE_STACKS 357
 - GS_WRITE_CORE_FILE 357
 - reserved names 357
 - upgradeLogDir 357
 - used in options 272
 - epoch garbage collection 240, 247–252
 - benefits of 247
 - determining epoch length 248
 - forcing 247
 - limitations of 247
 - parameters for tuning 248
 - when 248
 - EpochGcCount (cache statistic) 253
 - #epochGcMaxThreads (GcUser parameter) 248
 - #epochGcPageBufferSize (GcUser parameter) 248
 - #epochGcPercentCpuActiveLimit (GcUser parameter) 248
 - #epochGcTimeLimit (GcUser parameter) 248
 - #epochGcTransLimit (GcUser parameter) 248
 - EpochNewObjs (cache statistic) 253
 - EpochPossibleDeadObjs (cache statistic) 253
 - EpochScannedObjs (cache statistic) 253
 - error messages
 - native language 345
 - errors
 - after restoring backup 204
 - disk full, diagnosing 178
 - extent already exists 96
 - extent already open 96
 - extent missing or access denied 96
 - fatal 44
 - in configuration files, option value 272
 - in configuration files, syntax 272
 - invalid password 167
 - key file 95
 - object audit 123
 - shared page cache not attached 95
 - Stone response to Gem fatal error 44
 - stuck spin lock 108
 - tranlog directories full 190
 - transaction log missing 97
 - ErrorSymbols 346
 - /etc/services file 79
 - examining
 - user privileges 141, 149
 - executable configuration files 270
 - creating 269
 - defined 265

- names 269
- search for 268
- setting permission 269
- expanding extents 180
- extent files 180
 - see also *repository*
 - allocating space 274
 - allocation mode 37
 - checkpoint defined 169
 - creating new 41, 172
 - defined 23
 - disk full condition 179
 - disk space, managing 180
 - free space in 170
 - group access to 46
 - location 35
 - moving to raw partition 48
 - naming 274
 - permissions for dynamically added 171
 - pre-growing 36, 275
 - reallocating objects in 173, 174
 - removing 173
 - specifying size of 274
 - using multiple extents 37

F

- failed login messages in log 167
- failover to standby system 219, 223
- false 343
- fastObjectAudit (Repository) 122
- FDC (findDisconnectedObjects)
 - and MGC 240, 244
 - fast 245
- file descriptors 29, 34, 60
- file permissions 45
- FileControl (privilege) 140
- files
 - permissions for Gem processes 61
- fileSize (Repository) 170
- fileSizeReport (Repository) 170
- findDisconnectedObjects 240, 244
 - reducing impact of 246
- findObjectsLargerThan:limit:(Object) 262
- findReferencePathToObject: (Repository) 263
- findReferences: (Object) 264
- findReferencesWithLimit: (Object) 264
- finishTransaction (signal) 112, 303
- forceEpochGc (System) 247
- fork-in-time scenario 209

- free frame cache
 - specifying the size of 276
- Free Frame page server
 - log files 107
- free frame page server 55
 - defined 23
 - log files 114, 117
- free list page server
 - benefits of 55
- free space
 - in repository 170
- freeing repository space
 - when 239
- freeSpace (Repository) 170
- full backup 192
- full logging
 - how to manage 185
- fullBackupCompressedTo: (Repository) 199
- fullBackupCompressedTo:MBytes: (Repository) 199
- fullBackupTo: (Repository) 196

G

- garbage collection
 - automatic mechanisms 227–252
 - backups and 198, 199
 - concepts 233
 - conflicts between mechanisms 241, 242, 246
 - dead object, defined 235
 - determining epoch length 248
 - epoch collection 247–252
 - findDisconnectedObjects 244
 - GcUser configuration parameters 240
 - identifying garbage 234
 - live object, defined 234
 - local object memory 234
 - markForCollection 242
 - mark/sweep, defined 238
 - object table sweep, defined 238
 - overview 238–239
 - pages, defined 234
 - parameters for tuning 248, 254, 261
 - possible dead objects, defined 238
 - reclaim 234, 240
 - resources reclaimed 234
 - shadow object, defined 235
 - transitive closure, defined 234
 - tuning reclaim 254
 - two-step process 244
 - voting, defined 238

- write set union sweep, defined 239
- GarbageCollection (privilege) 140
- GcGems
 - starting up all 257
 - tasks of 239
- GciStructsMd5String 346
- GciTsStructsMd5String 346
- GcLock 241
 - and markForCollection 242
 - backups and 198
- GcUser
 - changing parameters for 240
 - described 137
- GcVoteState 256
- Gem session process
 - configuration 59
 - file 59
 - run time access to 64
 - tuning 63
 - custom executable 73
 - custom executable, installing 67
 - defined 58
 - file ownership and permissions for 61
 - linked and RPC 58
 - linked, setting up access 62
 - log files related to 118
 - remote from stone 79
 - RPC or remote, setting up access 62
 - starting 100
 - linked session 101
 - RPC session 102
 - troubleshooting 103
 - system resources for 59
 - temporary object space, tuning 63
 - tuning configuration 63
- GEM_ABORT_MAX_CRIS (configuration option) 276
- gem.conf file 269
- gemConfigurationAt: (System) 64
- gemConfigurationReport (System) 64
- GemConvertArrayBuilder (internal parameter) 306
- GemDropCommittedExportedObjs (internal parameter) 306
- GemExceptionSignalCapturesStack (internal parameter) 306
- GEM_FREE_FRAME_CACHE_SIZE (configuration option) 276
- GEM_FREE_FRAME_CACHE_SIZE (configuration option) 56
- GEM_FREE_FRAME_LIMIT (configuration option) 55, 276
- GemFreeFrameLimit (internal parameter) 65, 277
- GEM_FREE_PAGEIDS_CACHE (configuration parameter) 277
- GEM_GCI_LOG_ENABLED (configuration option) 277
- GEM_HALT_ON_ERROR (configuration option) 277
- GEM_KEEP_MIN_SOFTREFS (configuration option) 277
- GEM_MAX_SMALLTALK_STACK_DEPTH (configuration option) 278
- GEM_NATIVE_CODE_ENABLED (configuration option) 278
- GemNetId for remote Stone 84
- gemnetobject** executable
 - for custom Gem executable 67
 - mapping 73
 - modifying for custom Gem executable 67
 - RPC session and 103
- GEM_PGSRV_COMPRESS_PAGE_TRANSFERS (configuration option) 278
- GEM_PGSRV_FREE_FRAME_CACHE_SIZE (configuration option) 279
- GEM_PGSRV_FREE_FRAME_CACHE_SIZE (configuration option) 56
- GEM_PGSRV_FREE_FRAME_LIMIT (configuration option) 279
- GEM_PGSRV_UPDATE_CACHE_ON_READ (configuration option) 65, 110, 279
- GEM_PRIVATE_PAGE_CACHE_KB (configuration option) 280
- GEM_REPOSITORY_IN_MEMORY (configuration option) 280
- GEM_RPCGCI_TIMEOUT (configuration option) 280
- GEM_RPC_KEEPALIVE_INTERVAL (configuration option) 280
- GEM_RPC_USE_SSL (configuration option) 281
- gemsetup.csh 82, 94, 99
- gemsetup.sh 82, 94, 99
- GEM_SOFTREF_CLEANUP_PERCENT_MEM (configuration option) 281
- GemStone
 - see also *Stone repository monitor* and *Gem session process*
 - actions, user-defined 67
 - component overview 57
 - configuration files used in 266
 - privileges required for system administration tasks 135
 - privileges, defined 138
 - service name 67
 - shutting down repository 105
 - avoid **kill -9** 106

- starting repository monitor 94
- SymbolDictionaries, used in symbol resolution 141
- system logs, examining 113
- user ID, defined 135
- GEMSTONE (environment variable) 353
 - setting 101
- GemStone login authentication 155
- GEMSTONE_ADMIN_GC_LOG_DIR (environment variable) 353
- GemStoneError 346
- GEMSTONE_EXE_CONF (environment variable) 101, 265, 268–269, 270, 353
- GEMSTONE_GLOBAL_DIR (environment variable) 30, 353
- gemstone.hostid (host identifier) 30
- GEMSTONE_KEEP_LOG (environment variable) 353
- GemStone_Legacy_Streams 346
- GEMSTONE_LIB (environment variable) 354
- GEMSTONE_LOG (environment variable) 114, 354
- GEMSTONE_MAX_FD (environment variable) 29, 60, 354
- GEMSTONE_NRS_ALL (environment variable) 72, 79, 354
- GEMSTONE_PAGE_MGR_LOG_DIR (environment variable) 354
- GemStone_Portable_Streams 346
- GEMSTONE_RECLAIM_GC_LOG_DIR (environment variable) 354
- GEMSTONE_SPCMON_STARTUP_TIMELIMIT (environment variable) 354
- GEMSTONE_SYMBOL_GEM_LOG_DIR (environment variable) 354
- GEMSTONE_SYS_CONF (environment variable) 265, 266–267, 269, 354
- GEM_TEMPOBJ_AGGRESSIVE_STUBBING (configuration option) 281
- GEM_TEMPOBJ_CACHE_SIZE (configuration option) 229, 282
 - and bulk loading of objects 109
 - tuning 63
- GEM_TEMPOBJ_CACHE_SIZE (configuration parameter)
 - and multi-threaded operations 260
- GEM_TEMPOBJ_MESPACE_SIZE (configuration option) 229, 282
- GEM_TEMPOBJ_OOPMAP_SIZE (configuration option) 229, 283
- GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE (configuration option) 283
- GEM_TEMPOBJ_POMGEN_SCAVENGE_INTERVAL (configuration option) 283
- GemTempObjPomgenScavengeInterval
 - (internal parameter) 283
- GEM_TEMPOBJ_POMGEN_SIZE (configuration option) 229, 284
- GEM_TEMPOBJ_SCOPES_SIZE (configuration option) 284
- global session statistics 131
- Globals (system globals dictionary) 142
 - initial contents of 347
- group:authorization: (GsObjectSecurityPolicy) 147
- groups
 - access to extents 46
 - adding a new user to 143, 148
 - and AllGroups system object 345
 - and object security policy authorization 138
 - creating new 148
 - defined 141
 - removing a user from 148
 - removing from an object security policy's authorization list 147
- GS_CFTIME (environment variable) 119, 354
- GS_CORE_TIME_OUT (environment variable) 354
- GS_DEBUG_COMPILE_TRACE (environment variable) 231
- GS_DEBUG_PAM 355
- GS_DEBUG_SHARDE_MEM 355
- GS_DEBUG_SSL_LOG_DIR (environment variable) 355
- GS_DEBUG_VMGC_MKSW_MEMORY_USED_SOFT_BREAK (environment variable) 231, 355
- GS_DEBUG_VMGC_MKSW_PRINT_C_STACK (environment variable) 231, 355
- GS_DEBUG_VMGC_MKSW_PRINT_STACK (environment variable) 231, 355
- GS_DEBUG_VMGC_PRINT_MKSW (environment variable) 231, 355
- GS_DEBUG_VMGC_PRINT_MKSW_MEMORY (environment variable) 231, 355
- GS_DEBUG_VMGC_PRINT_MKSW_MEMORY_USED (environment variable) 231, 355
- GS_DEBUG_VMGC_PRINT_SCAV (environment variable) 231, 356
- GS_DEBUG_VMGC_PRINT_TRANS (environment variable) 356
- GS_DEBUG_VMGC_SCAV_PRINT_C_STACK (environment variable) 232, 356
- GS_DEBUG_VMGC_SCAV_PRINT_STACK (environment variable) 232, 356
- GS_DEBUG_VMGC_VERBOSE_OUTOFMEM (environment variable) 232, 356
- GS_DEBUG_VMGC_VERIFY_MKSW (environment variable) 232, 356
- GS_DEBUG_VMGC_VERIFY_SCAV (environment variable) 232, 356

- GS_DEBUG_VM_PRINT_TRANS (environment variable) 231
- GS_DISABLE_CHARACTER_TABLE_LOAD (environment variable) 356
- GS_DISABLE_KEEPALIVE (environment variable) 356
- GS_DISABLE_SIGNAL_HANDLERS (environment variable) 356
- GS_DISABLE_WARNING (environment variable) 356
- gslist** command 100, 314
using to find log locations 113
- GsObjectInventory
and GcLock 241
and repository scan 125
- GS_PAGE_MGR_PRINT_REMOTE_STACKS (environment variable) 357
- GS_WRITE_CORE_FILE (environment variable) 357
- guest mode, NetLDI
captive accounts with 76
- ## H
- high-availability systems, setting up standbys for 217
- host identifier 30
- hot standby 220–224, 320, 322
- ## I
- identifying garbage 234
- identifying large objects in the repository 262
- initialization
executing code on gem login 168
- instance creation
and InstancesDisallowed 346
- InstancesDisallowed 346
- invalid password error 167
- IPv6 295
- ## K
- keepalive, network option 74
- kernel requirements
for Gem session processes 61
for Stone repository monitor 29
- KEYFILE (configuration option) 284
- killing Gem or Stone processes 106
- ## L
- large objects, identifying in the repository 262
- large repositories
special considerations 110
- lastLoginTime 136
- LDAP authentication 156
- Legal characters in a Stone name 326
- licensing keyfile
setting the location of 284
- linked vs. RPC application 58
- live object
defined 234
- local object memory
defined 234
organization 227
- lock files 30
- log files 113–119
AIO page server 107, 117
and process type 114
for child processes 62
for RPC Gems 62
free frame page server 107, 117
garbage collection session 107
Gem session process 118
NetLDI 119
Page Manager 107, 117
shared page cache monitor 34, 107
Stone repository monitor 107
Symbol Gem 117
SymbolGem 107
write access for 46, 62
- login
executing code on gem 168
- login authentication
GemStone 155
LDAP 156
UNIX 155
- login object security policy, described 138
- loginHook
168
- LogOriginTime (read-only runtime configuration parameter) 307
- logOriginTime (Repository method) 188
- logreceiver 221, 222, 223, 224, 320, 329
- logsender 220, 222, 223, 322, 330
- LOG_WARNINGS (configuration option) 285
- low memory
signal 232

M

manual transaction mode 111, 113, 198
 markForCollection
 conflicts with other garbage collection 242
 fast version 243
 scheduling with cron 243
 markForCollection
 conflicts with other garbage collection 241
 scheduling 243
 markForCollection (Repository) 240
 markForCollectionWithMaxThreads:...
 243
 markGcCandidatesFromFile:
 conflicts with other garbage collection 246
 markGcCandidatesFromFile: (Repository)
 240, 244
 marking garbage repository-wide 240
 marking objects for garbage collection 242
 mark/sweep
 defined 238
 master repository (in hot standby) 217, 220
 #maxTransactionDuration (GcUser
 parameter) 254
 mE
 region of temporary object memory 227
 memory
 and multi-threaded operations 260
 Gem session processes 60
 local object 227
 server needs 28
 signalling on low 232
 temporary object 228
 memory space
 for multi-threaded operations
 calculating 260
 MGC (markGcCandidates)
 and FDC 240, 244
 mid-level cache
 connection methods 89
 in distributed system 88
 reporting methods 90
 MinusInfinity 343
 MinusQuietNaN 343
 MinusSignalingNaN 343
 modes
 allocation 25, 37
 debugging (NetLDI) 324
 file protection 46, 61
 full logging 23, 27, 42, 181, 182
 guest (NetLDI) 324
 manual transaction 113

 partial logging 43, 182, 189
 transaction logging 42
 modifying
 another user's ID 145
 another user's password 145, 146
 monitoring
 cache statistics, programmatic access 127
 MtActiveThreads 260
 MtPercentCpuActiveLimit 259, 260
 MtThreadsLimit 259, 260
 multi-threaded operations
 backup 198
 memory impact of 260
 monitoring 260
 session use and 198
 tuning 259
 multi-threaded scan 259
 defined 259
N
 name, stone legal name limits 326
 Nameless (predefined system account) 137
 naming
 configuration file options 271
 executable configuration files 269
 extent files 274
 Native Code 66
 NativeLanguage 345
 NetLDI
 debug mode 92
 default mode 76
 default name 72
 defined 71
 guest mode 76
 in GEMSTONE_NRS_ALL 72
 list of 100
 log files 30, 119
 permissions for executable netldid 75, 77
 ports 72
 secure mode 76
 shutting down 105
 starting 98, 324
 troubleshooting 99
 well-known port 324
 network
 authentication, when required 76
 disrupted communications 74
 /etc/services file 79
 Gem session process on Stone's machine 83
 GemStone network objects (gemnetobject) 73

- guest mode with captive account 76
 - keepalive option 74
 - linked application on remote machine 81
 - log files for spawned processes 62
 - NetLDI 71
 - page server processes 72
 - remote sessions, setting up 79
 - resource string (NRS) 78
 - syntax 337
 - setting up remote sessions 79
 - shared GemStone directory 80
 - shared page cache and 73
 - shell scripts, modifying for custom Gem
 - executable 67
 - Stone and RPC Gem on different machines 84
 - troubleshooting remote logins 90
 - typical configurations 69
 - network resource string
 - #auth modifier 75
 - new
 - region of temporary object memory 227
 - nil 343
 - NoGsFileOnClient (privilege) 140
 - NoGsFileOnServer (privilege) 140
 - NoPerformOnServer (privilege) 140
 - NotTranloggedGlobals 346
 - NoUserAction (privilege) 140
 - NRS (network resource string) 78
 - GEMSTONE_NRS_ALL 79
 - syntax 337
- O**
- object audit 122
 - repairing errors 124
 - object memory
 - organization of 227
 - object security policy
 - changing a user's default 147
 - changing the authorization of a 147
 - unit of authorization 138
 - used in read/write authorization 146
 - Object Server, see *Stone repository monitor*
 - object table
 - loading at startup 110
 - object table sweep
 - defined 238
 - objectAudit (Repository) 122, 123
 - objectAuditWithMaxThreads:percentCpuActiveLimit: (Repository) 122
 - objects, large, identifying in the repository 262
 - ObjectSecurityPolicies
 - predefined 344
 - ObjectSecurityPolicy
 - predefined instances 344
 - ObjectSecurityPolicyCreation (privilege) 139
 - ObjectSecurityPolicyProtection (privilege) 140
 - #objsMovedPerCommitThreshold (GcUser parameter) 254
 - offline extent snapshot backup 192, 193
 - old
 - region of temporary object memory 227
 - oldestLogFileIdForRecovery (Repository) 185
 - oldestLogFileIdForRecovery (Repository) 186, 188
 - online extent snapshot backup 192
 - onlinebackup.sh
 - file in GemStone examples directory 196
 - oopHighWater value
 - and multi-threaded operations 260
 - /opt/gemstone/ directory
 - file access permission 62
 - /opt/gemstone/ used for GemStone files 30, 46, 114
 - OtherPassword (privilege) 139
 - OutOfMemory error 228
- P**
- page audit 120
 - Page Manager
 - log files 107, 114, 117
 - page server process
 - AIO page server 55
 - free frame page server 55
 - tasks of 55
 - pageaudit command 316
 - pages
 - defined 234
 - PagesNeedReclaimSize 256
 - password
 - configuring authentication 155
 - constraining choices 159–166
 - modifying another user's 145, 146
 - modifying your own 145
 - percentCpuActive 260
 - perm
 - region of temporary object memory 227
 - permission, setting for executable configuration files 269
 - permissions
 - file 45

- for Gem session processes 61
 - for shared page cache 287
 - PlusInfinity 343
 - PlusQuietNaN 343
 - PlusSignalingNaN 343
 - pom
 - region of temporary object memory 227
 - ports
 - NetLDI 72
 - shared page cache monitor 72
 - Stone 72
 - POSIX 326
 - possible dead objects
 - defined 238
 - PossibleDeadObjs 256
 - pre-growing repository extents 36, 275
 - prerequisites 3
 - primary repository (in standby systems) 217
 - primitives, user-defined 67
 - printing configuration options 276
 - privileges
 - adding to a user's 141, 150
 - assigning to a new user 143
 - defined 138
 - deleting a user's 150
 - examining a user's 141, 149
 - redefining a user's 141, 150
 - removing from user 150
 - required for system administration tasks 135
 - process slots
 - for cache statistics 127
 - process type
 - and log files 114
 - profiling
 - repository contents 125
 - pstack** command 317
 - PublishedObjectSecurityPolicy 344
 - Publishers (predefined group) 141
 - purging unneeded objects 234
- ## R
- RAID devices 35
 - raw partitions
 - changing to and from 48
 - removing old contents 318
 - setting up 47
 - use recommended 26
 - read/write authorization 146
 - and object security policies 138
 - reclaim
 - configuration parameters affecting 254
 - garbage collection task 240
 - parameters for tuning 254
 - tuning 259
 - Reclaim Gem 239
 - and reclaim 253
 - configuring 254, 256
 - ReclaimCount 256
 - #reclaimDeadEnabled (GcUser parameter) 254
 - ReclaimedPagesCount 256
 - reclaiming system resources 234
 - #reclaimMinFreeSpaceMb (GcUser parameter) 254, 255
 - #reclaimMinPages (GcUser parameter) 254
 - recovery
 - after full disk error 178
 - after NetLDI startup failure 99
 - after Stone startup failure 95
 - after unexpected shutdown 106
 - using offline extent backup 203
 - redefining a user's privileges 141, 150
 - references to repository objects
 - searching for 264
 - remote logins
 - troubleshooting 90
 - remoteCachesReport (System) 90
 - removedbf** command 318
 - archiving transaction logs 186
 - removing
 - a user from a group 148
 - a user from the system 143
 - a user's privileges 150
 - extent 173
 - privilege from a user 150
 - symbolList from user 151
 - repair (Repository) 124
 - repository
 - see also *extent files* and *transaction logs*
 - audit at object level 122
 - audit at page level 120
 - backups, see *backups*
 - bulk loading of 109
 - checkpoint frequency 53
 - continuous restore mode 221
 - disaster recovery 326
 - disk full condition 179
 - error when below free space threshold 179
 - free space in 170
 - growth of 169, 233
 - identifying large objects in 262
 - marking garbage in 240

- object references, searching for 264
- oldest log needed for recovery 186
- page fragmentation 178
- profiling 125
- running multiple 56
- running warm or hot standby 217
- shrinkage, when 239
- shrinking to minimum size 175
- shutting down 105
- single instance of 344
- standby systems 217
- starting monitor 94
- transaction logs, defined 23
- updating views of extents 172
- when free space appears in 253
- resolving symbols, `symbolList` used in 141
- restore mode 221
- `restoreFromBackup`: (Repository) 205
- `restoreStatus` (Repository) 201, 207
- `restoreStatusOldestFileId` (Repository) 188
- restoring the GemStone repository to a point in time 211
- RPC Gems
 - configuration file name 269
- RPC vs. linked application 58
- S**
- scan
 - multi-threaded 259
- scheduling `markForCollection` 243
- search for
 - executable configuration files 268
 - references to repository objects 264
 - system-wide configuration file 266
- secure mode, NetLDI 76
- security 159
 - disabling inactive accounts 164
 - finding disabled accounts 153
 - last login by account 165
 - limiting concurrent sessions by user 166
 - login failures
 - disabling further 167
 - logging 167
- passwords
 - aging 162
 - clearing disallowed list of 162
 - constraining choice of 159
 - disallowing certain 161
 - disallowing reuse of 161
 - login limit under a 166
 - warning of expiration 163
 - when last changed 164
 - see also *passwords*
- service name, GemStone 67
- `services.dat`
 - file in GemStone system directory 73
- session statistics 131
- `SessionAccess` (privilege) 139
- `SessionInBackup` (internal parameter) 307
- `SessionPriority` (privilege) 140
- sessions
 - current session names 105
 - determining session holding up reclaim 261
 - find who is logged in 104
 - finding process id of 105, 261, 262
 - identifying current 261
- setting
 - full transaction logging 303
 - permission, executable configuration files 269
- setuid bit
 - for Gem session processes 61
 - on executable files 45
- shadow object
 - contrasted with dead object 235–238
 - defined 235
- shared memory
 - access by Gems 61
 - utility to check system 33
- shared page cache
 - cleanup after **kill -9** 106
 - configuration 31
 - disconnect error 108
 - enabling 63
 - for remote Gem session processes 60, 63
 - maximum processes 33, 286
 - monitor process 22, 31, 108
 - log file 114, 116, 117
 - on remote machine 82
 - sessions on remote hosts and 73
 - size 32, 287
 - spin lock attempts 53
 - stuck spin lock 108
 - timeout of remote 302
- shared page cache monitor

- log files 107
- shared system objects, in Globals dictionary 142
- shrinking the repository 175
- SHR_NUM_FREE_FRAME_SERVERS (configuration option) 285
- SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_POLICY (configuration option) 285
- SHR_PAGE_CACHE_LOCKED (configuration option) 286
- SHR_PAGE_CACHE_NUM_PROCS (configuration option) 29, 63, 103, 286
 - adjusting to number of users 33
- SHR_PAGE_CACHE_NUM_SHARED_COUNTERS (configuration option) 286
- SHR_PAGE_CACHE_PERMISSIONS (configuration option) 61, 287
- SHR_PAGE_CACHE_SIZE_KB (configuration option) 33, 63, 287
- shrpcmonitor 31
- SHR_SPIN_LOCK_COUNT (configuration option) 287
- SHR_TARGET_FREE_FRAME_COUNT (configuration option) 288
- SHR_WELL_KNOWN_PORT_NUMBER (configuration option) 72, 288
- shutdown message 107, 108
- sigAbort 111, 303
 - configuration parameter controlling 303
 - handler 111
- signal
 - sent on commit record backlog 111
 - sent on low memory 232
- slave repository (in hot standby) 217, 220
- #sleepTimeBetweenReclaimMs (GcUser parameter) 254
- Smalltalk
 - compiler, and symbol resolution 141
 - kernel classes, and Globals dictionary 142
 - methods, and GemStone privileges 138
- Smalltalk full backups 196
- SourceStringClass 346
- space
 - determining usage in repository 125
- SpinLockCount (internal parameter) 287
- standalone Gems 269
- standby systems 217
- startcachewarmer** command 319
 - when to use 110
- startCheckpointSync (System) 196
- starting a checkpoint 188
- starting a NetLDI 76
- starting GemStone 93
- startlogreceiver** command 221, 223, 320
- startlogsender** command 220, 223, 322
- startnetldi** command 71, 76, 324
- startNewLog (Repository) 188
- startstone** command 94, 326
 - when transaction logs are missing 98
- statistics
 - global 131
 - obtaining value by name 128
 - operating system 132
 - programmatic access to process CPU 132
 - see also under *cache statistics*
 - session 131
 - shared page cache 127
- statmonitor 127
- statmonitor** command 327
- StnAdminGcSessionEnabled (internal parameter) 288
- STN_ADMIN_GC_SESSION_ENABLED(configuration option) 288
- STN_ALLOCATE_HIGH_OOPS(configuration option) 288
- STN_ALLOW_NFS_EXTENTS (configuration option) 289
- STN_CACHE_WARMER (configuration option) 110, 289
- STN_CACHE_WARMER_SESSIONS (configuration option) 110, 289
- STN_CHECKPOINT_INTERVAL (configuration option) 53, 169, 289
- STN_COMMIT_QUEUE_THRESHOLD (configuration option) 290
- StnCommitQueueThreshold (internal parameter) 290
- STN_COMMIT_RECORD_QUEUE_SIZE (configuration option) 290
- STN_COMMITS_ASYNC (configuration option) 290
- STN_COMMIT_TOKEN_TIMEOUT (configuration option) 290
- STN_CR_BACKLOG_THRESHOLD (configuration option) 111, 291
- StnCurrentTranLogDirId (internal parameter) 307
- StnCurrentTranLogNames (internal parameter) 307
- STN_DISABLE_LOGIN_FAILURE_LIMIT (configuration option) 167, 291
- STN_DISABLE_LOGIN_FAILURE_TIME_LIMIT (configuration option) 167, 291
- StnDisableLoginFailureTimeLimit (internal parameter) 291
- STN_DISKFULL_TERMINATION_INTERVAL (configuration option) 179, 292
- StnDiskFullTerminationInterval (internal

- parameter) 292
- STN_EPOCH_GC_ENABLED (configuration option) 247, 292
- STN_EXTENT_IO_FLAGS (configuration option) 292
- STN_FREE_FRAME_CACHE_SIZE (configuration option) 293
- STN_FREE_SPACE_THRESHOLD (configuration option) 179
- STN_FREE_SPACE_THRESHOLD (configuration option) 293
- StnFreeSpaceThreshold (internal parameter) 293
- STN_GEM_ABORT_TIMEOUT (configuration option) 293
- STN_GEM_LOSTOT_TIMEOUT (configuration option) 111, 294
- STN_GEM_TIMEOUT (configuration option) 294
- StnGemTimeout (internal parameter) 294
- STN_HALT_ON_FATAL_ERR (configuration option) 44, 108, 294
- STN_LISTENING_ADDRESSES (configuration option) 295
- StnLogFileName (internal parameter) 307
- StnLogGemErrors (internal parameter) 307
- STN_LOGIN_LOG_ENABLED (configuration option) 296
- StnLoginsSuspended (internal parameter) 307
- STN_LOG_IO_FLAGS (configuration option) 295
- STN_LOG_LOGIN_FAILURE_LIMIT (configuration option) 167, 296
- StnLogLoginFailureLimit (internal parameter) 296
- STN_LOG_LOGIN_FAILURE_TIME_LIMIT (configuration option) 167, 296
- StnLogLoginFailureTimeLimit (internal parameter) 296
- STN_LOOP_NO_WORK_THRESHOLD (configuration option) 297
- STN_MAX_AIO_RATE (configuration option) 297
- STN_MAX_AIO_REQUESTS (configuration option) 297
- STN_MAX_GC_RECLAIM_SESSIONS (configuration option) 298
- STN_MAX_LOGIN_LOCK_SPIN_COUNT (configuration option) 298
- STN_MAX_REMOTE_CACHES (configuration option) 298
- StnMaxReposSize (internal parameter) 308
- STN_MAX_SESSIONS 198
- STN_MAX_SESSIONS (configuration option) 33, 103, 299
- StnMaxSessions (internal parameter) 308
- STN_MAX_VOTING_SESSIONS (configuration option) 299
- StnMaxVotingSessions (internal parameter) 299
- StnMntMaxAioRate (internal parameter) 297
- STN_NUM_AIO_WRITE_THREADS (configuration option) 299
- STN_NUM_GC_RECLAIM_SESSIONS (configuration option) 299
- StnNumGcReclaimSessions (internal parameter) 299
- STN_NUM_LOCAL_AIO_SERVERS (configuration option) 55, 300
- STN_OBJ_LOCK_TIMEOUT (configuration option) 300
- StnObjLockTimeout (internal parameter) 300
- stnOopHighWater and multi-threaded operations 260
- STN_PAGE_MGR_MAX_WAIT_TIME (configuration option) 301
- StnPageMgrMaxWaitTime (internal parameter) 51, 301
- STN_PAGE_MGR_PRINT_TIMEOUT_THRESHOLD (configuration option) 301
- STN_PAGE_MGR_REMOVE_MAX_PAGES (configuration option) 301
- StnPageMgrRemoveMaxPages (internal parameter) 301
- STN_PAGE_MGR_REMOVE_MIN_PAGES (configuration option) 302
- StnPageMgrRemoveMinPages (internal parameter) 302
- STN_PRIVATE_PAGE_CACHE_KB (configuration option) 33, 302
- STN_REMOTE_CACHE_TIMEOUT (configuration option) 302
- STN_REMOTE_CACHE_PGSRV_TIMEOUT (configuration option) 302
- STN_SHR_TARGET_PERCENT_DIRTY (configuration option) 303
- STN_SIGNAL_ABORT_CR_BACKLOG (configuration option) 303
- StnSignalAbortCrBacklog (internal parameter) 303
- StnSunsetDate (internal parameter) 308
- STN_SYMBOL_GC_ENABLED (configuration option) 303
- STN_TRAN_FULL_LOGGING (configuration option) 42, 182, 185, 189, 303
- STN_TRAN_LOG_DEBUG_LEVEL (configuration option) 304
- StnTranLogDebugLevel (internal parameter) 304
- STN_TRAN_LOG_DIRECTORIES (configuration option) 181, 183, 185, 187, 189, 304 and bulk loading of objects 109

- STN_TRAN_LOG_LIMIT (configuration option) 54, 183, 304
- StnTranLogLimit (internal parameter) 304
- StnTranLogOriginTime (internal parameter) 308
- STN_TRAN_LOG_PREFIX (configuration option) 181, 305
- STN_TRAN_LOG_SIZES (configuration option) 43, 44, 181, 187, 189, 305
- STN_TRAN_Q_TO_RUN_Q_THRESHOLD (configuration option) 305
- StnTranQToRunQueueThreshold (internal parameter) 305
- STN_WELL_KNOWN_PORT_NUMBER (configuration option) 72, 306
- Stone private page cache tuning 33
- Stone repository monitor
 - AIO page servers 55
 - configuration
 - file 23
 - run time access to 50
 - configuring server 27
 - defined 22
 - disk usage 26
 - extents 34
 - on raw partitions 48
 - file descriptors for 29
 - file permissions for 45
 - Gem fatal errors, response to 44
 - identifying configuration file in use 315
 - identifying sessions logged in 104
 - kernel parameters for 29
 - listing of 100
 - log files 107, 114, 115
 - memory for 28
 - private page cache 33
 - raw partitions 26
 - using 47
 - recovery 106
 - disk error 107
 - disk full condition 178
 - fatal error by Gem 108
 - shared page cache error 108
 - removing stale locks 314
 - running multiple servers 56
 - security, *see security*
 - shared page cache 31
 - diagnostics for 34
 - tuning of 52
 - shutting down 105
 - starting 93
 - troubleshooting 95
 - status of 314
 - swap (paging) space for 28
 - swapping, excessive 53
 - transaction logs 42
 - on raw partitions 49
- stone.conf file 270
- stoneConfigurationAt: (System) 50, 187
- stoneConfigurationReport (System) 50
- stoplogreceiver** command 221, 224, 329
- stoplogsender** command 220, 223, 330
- stopnetldi** command 331
- stopSession: (System) 105
 - delay for inactive sessions 105
- stopstone** command 332
 - shutting down repository 105
- stuck spin lock error 108
- Subscribers (predefined group) 141
- suspending checkpoints 194
- swap space
 - system needs for server 28
- #sweepWsUnionMaxThreads (GcGem parameter) 261
- #sweepWsUnionPageBufferSize (GcGem parameter) 261
- #sweepWsUnionPercentCpuActiveLimit (GcGem parameter) 261
- Symbol Gem
 - defined 23
 - log files 117
- symbol list
 - adding to user 151
 - removing from user 151
- symbol list, and UserProfile 141
 - removing from 152
- symbol resolution 141
- SymbolGem
 - log files 107, 114
- SymbolUser (predefined group) 141
- syntax
 - configuration files 271
 - errors, in configuration files 272
- system
 - GemStone logs 114
 - objects, in Globals dictionary 142
 - shutdowns, diagnosing 106, 178
- System (predefined group) 141
- system clock 30
- SystemControl (privilege) 139
- SystemObjectSecurityPolicy 344

- SystemRepository 344
 - SystemUser
 - and AllUsers system object 345
 - and SystemObjectSecurityPolicy 344
 - described 136
 - system-wide configuration files
 - defined 265
 - search for 266
- T**
- temporary object memory
 - error on low memory 228
 - temporary object space
 - finding out currently used 229
 - temporary objects 227
 - terminateSession:timeout: (System) 105
 - time, changes in system 30
 - Timeout, on SPC startup 354
 - Timestamp formatting 119
 - timeToRestoreTo: (Repository) 211
 - TimeZone 348–351
 - topaz**
 - configuration files and 270
 - command 333
 - tranlog directories full (error message) 190
 - tranlogXXX.log (transaction log) 181
 - transaction logging
 - comparison of full, partial 182
 - enabling 42, 182, 303
 - partial logging checkpoint threshold 42, 182, 304
 - transaction logs
 - adding online 187
 - archiving 186
 - backups for 186
 - current log directory 187
 - current log file 187
 - current log size 187
 - disk full condition 190
 - disk space, managing 190
 - hot standby, use in 221
 - identifying checkpoints in 313
 - limiting size 218
 - log directories 181, 304
 - log not found error 97
 - log origin time 188
 - log size limit 181, 305
 - missing 212
 - moving to raw partition 49
 - oldest log needed for recovery 186
 - out of sequence 209
 - restarting stone without 98
 - restoring a subset of 209
 - restoring into backup 207–208
 - warm standby, use in 218
 - transaction mode, manual 113, 198
 - transaction record backlog 303
 - TransactionBacklog 111, 112, 303
 - transactionMode: (System) 111
 - transactions
 - checkpoint
 - starting 188
 - transitive closure
 - defined 234
 - troubleshooting
 - NetLDI startup 99
 - remote sessions 90
 - session login 103
 - Stone startup 95
 - true 343
 - tuning reclaim 259
 - tz (TimeZone database) 349
 - tzselect (TimeZone utility) 350
- U**
- UNIX
 - file system corruption 107
 - kernel configuration 34
 - UNIX authentication (for GemStone login) 155
 - upgradeLogDir (environment variable) 357
 - user actions, initializing 67
 - user groups
 - adding users to 148
 - and AllGroups system object 345
 - assigning a new user to 143
 - defined 141
 - removing a user from 148
 - removing from an object security policy's
 - authorization list 147
 - used in object security policy authorization 138
 - UserPassword (privilege) 139
 - UserProfile
 - and AllUsers system object 345
 - described 135
 - userProfileForSession: (System) 105
 - users
 - current sessions 105
 - default object security policy 136, 138

- dictionaries 141
 - removing 151
- disabling inactive accounts 164
- environment variables 353
- finding disabled accounts 153
- group membership 136, 148
- limiting concurrent sessions by same 166
- listing all 145
- password 136, 138
- predefined users 136
- privileges 136, 138
 - changing 150
- security, *see security*
- sessions holding up reclaim 261
- symbol list 141
- userId 138
 - changing 145
 - when last logged in 165
 - why account disabled 153
- UserSecurityData 136
- user-written C functions, calling from Smalltalk 67

V

- verification of backups 199
- voting
 - defined 238
- VSD 127
- vsd** command 334

W

- waitstone** command 78, 335
- warm standby 218–219
- warming, cache 289, 319
- weighted allocation of extents 38
- write authorization 146
- write set union sweep
 - defined 239
 - tuning 261
- writeFdcArrayToFile: (Repository) 245

Z

- zdump (TimeZone utility) 350
- zic (TimeZone utility) 350
- zoneinfo (TimeZone database) 349