


GemStone/S 64 BitTM **EARLY DRAFT Release** **Notes**

Version 3.7

This document is a EARLY DRAFT version of 3/8/23, and may
change substantially before release.
You must review the final version at time of release for
updates and corrections.



INTELLECTUAL PROPERTY OWNERSHIP

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

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

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

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

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

COPYRIGHTS

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

PATENTS

GemStone software has been covered by U.S. Patent Number 6,256,637 "Transactional virtual machine architecture" (1998-2018), Patent Number 6,360,219 "Object queues with concurrent updating" (1998-2018), Patent Number 6,567,905 "Generational garbage collector with persistent object cache" (2001-2021), and Patent Number 6,681,226 "Selective pessimistic locking for a concurrently updateable database" (2001-2021).

TRADEMARKS

GemTalk, **GemStone**, **GemBuilder**, **GemConnect**, and the GemTalk logo are 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.

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

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

Microsoft, **Windows**, and **Windows Server** 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.

Ubuntu is a registered trademark of Canonical Ltd., Inc., in the U.S. and other countries.

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

AIX, **POWER7**, **POWER8**, **POWER9** and **VisualAge** are trademarks or registered trademarks of International Business Machines Corporation.

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

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

Raspberry Pi is a trademark of the Raspberry Pi Foundation

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 LLC
15220 NW Greenbrier Parkway
Suite 240
Beaverton, OR 97006

Preface

About This Documentation

These release notes describe changes in the GemStone/S 64 Bit™ version 3.7 release. Read these release notes carefully before you begin installation, upgrade, or development with this release.

For information on installing or upgrading to this version of GemStone/S 64 Bit, please refer to the *GemStone/S 64 Bit Installation Guide* for version 3.7.

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 LLC, previously GemStone Systems, Inc. and a division of VMware, Inc.

Technical Support

Support Website

gemtalksystems.com

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

- ▶ **Documentation** for the current and for previous released versions of all GemTalk products, in PDF form.
- ▶ **Product download** for the current and selected recent versions of GemTalk software.
- ▶ **Bugnotes**, identifying performance issues or error conditions that you may encounter when using a GemTalk product.

- ▶ **Supplemental Documentation** and **TechTips**, providing information and instructions that are not in the regular documentation.
- ▶ **Compatibility matrices**, listing supported platforms for GemTalk product versions.

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

Help Requests

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

Website: techsupport.gemtalksystems.com

Email: techsupport@gemtalksystems.com

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

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

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

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

24x7 Emergency Technical Support

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

Training and Consulting

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



Table of Contents

Chapter 1. DRAFT Release Notes for 3.7

Overview	13
New keyfiles required	13
Supported Platforms.	14
Platforms for Version 3.7	14
GemBuilder for Smalltalk (GBS) Versions.	14
VSD Version.	15
GemBuilder for Java (GBJ) Version.	15
Rowan	15
Sparkle.	15
Upgrade	15
Documentation Changes	16
1.1 Library and Distribution changes	16
Updated Library Versions.	16
RPM distributions available.	16
clientLibs distributions.	17
Distribution Changes	17
\$GEMSTONE/projects moved	17
Further changes in permissions within the distribution	17
Example scripts for setting up GemStone in systemd	17
1.2 Optimizations	18
New native code generation and compiler optimizations	18
Native code not reimplemented on AIX	18
Older CPU machines cannot run v3.7	18
Process Scheduling and Debugger Support	18
Scheduler now preemptive	18
Breakpoint Handling and Debugger Support	19
Interrupting sockets	19
Process and debugging bugs fixed	19

Handling asynchronous events in a designated process	19
Added and renamed Process and related methods	20
GsProcess "Step" methods renamed	20
Signaling the GCI	20
Other added methods	21
Removed Methods	21

Chapter 2. Changes in Administration

2.1 Multithreaded instance migration	24
InstVarMappingArray	25
Example	25
2.2 One-time password	26
One-time passwords allow login without further authentication	26
Privilege to create one-time passwords	26
One-time passwords for other UserProfiles	26
Creating and Using a one-time password	26
New topaz command	27
Tracking one-time passwords	27
Login log	27
Added Cache Statistics	27
Added methods	28
Changes to descriptionOfSession:	28
Additional fields in Login log	29
New Privilege	29
2.3 Changes in Users and ObjectSecurityPolicies	29
CreateOnetimePassword privilege	29
MigrateObjects privilege	29
2.4 Changes in Backup and Restore	29
Full control over compression levels	29
Consistent default for Lz4 to level 0	29
Specifying level in existing methods	30
New methods including level argument	30
Repository scan operations now allowed in solo sessions	30
In restore mode, compiler now generates arguments and temporaries as Strings	30
2.5 Hot Standby Changes	31
suspendCommitsForFailover no longer supported	31
Hot Standby with failover and multiple slaves	31
Repository >> failOverStatus improved report	31
Full details in JSON from gslist	32
Upgrade via hot standby	33
2.6 Garbage Collection Changes	34
markForCollection performance improvement for repositories with long object chains	34
Number of threads configured for AdminGem	34
Renamed method to configure number of reclaim threads	34

Reclaim GcGem Logging	35
Determining sessions that have not voted	35
2.7 Other Changes related to administrative operations	35
Cache Warming.	35
Cache warming during upgrade	35
Deprecated cache warming configuration parameters removed	35
Determining status of cache warming	36
Changes in abort behavior in solo session	36
Changes in session information for GcLock	36
Methods added to determine recursive size of an object tree	37
Fast variants of findAllReferences	37
2.8 Changes in Errors	38
Added Errors	38
Added Error and Exception Classes	38
Removed Errors	38
Cannot return now not resumable	38
OS-origin error messages changed	38
2503 changed from FloatingPointError to NumericError	38
AbstractException >> signalToGci	38
Print GCI trace records on protocol error	39
MFC Warning now includes possibleDeadSymbols	39
2.9 Changes in Utilities	39
Many utilities faster on Linux.	39
largememorypages improved and usable for remote caches	39
Changes in argument handling	39
secure_backup_extract added.	40
Topaz changes	40
Strings containing now print hex value rather than space	40
LIMIT added option LEV2OOPS	41
SET added option ONETIMEPASSWORD	41
Improvements to debugging using DEBUGGEM	41
System waitForDebug	41
System cancelWaitForDebug	41
Added Topaz commands	41
2.10 Configuration Parameter Changes	42
Removed configuration parameters	42
STN_NUM_GC_RECLAIM_SESSIONS now applies to Admin Gem	42
GEM_NATIVE_CODE_ENABLED.	42
STN_GEM_TIMEOUT	42
2.11 Changes in Statmonitor and Cache Statistics	43
Performance improvement for Shared Page Cache Monitor.	43
Improved lz4 compression and new VSD required	43
Statmonitor -z -Z arguments now accept compression level	43
Statistics Types reorganized.	43
Pgsvr split	43
Stone, Page Manager, and Stone's Restore Threads split.	44

Gems split	44
Added API for PageManager statistics	44
Changes in memory statistics on Linux.	45
Per-process statistics	45
Add statistic for number of memory mapped regions in use.	45
Added statistics for smaps thread	46
New restrictions for per-process Statistics for other users	46
Other changes in Existing Statistics	47
Details for GarbageCollectionState and GcHighWaterPage	47
ReposAtMaxSize changed	47
Changes in Session*WithGcLock and GcLockKind	48
Removed Statistics	48
Added Statistics	48
"Time Since" statistics	48
Other Added Cache Statistics	49

Chapter 3. Changes in GemStone Smalltalk

3.1 Changes in Classes, compilation, and code management	53
SystemUser SymbolList Changes	53
ClassOrganizer >> newExcludingGlobals	54
GsNMethod >> recompileFromSource	54
Added method Behavior >> compileMethod:category:environmentId:	54
#logCreation and logging class creation	54
Improvements to filein using GsFileIn	54
3.2 FileSystem	55
Differences from GsFile.	55
Primary API classes	55
Supporting Classes	57
Zinc Stream Classes	57
Opening options.	57
Error Classes	57
Store classes	58
Resolver classes	58
FFI support classes	58
3.3 Support for ssh and sftp using OpenSSL.	58
SSH with GsSshSocket	58
SFTP with GsSftpSocket	59
3.4 Changes in Magnitude classes.	60
Float printing improvements	60
Time support for microseconds	61
Converting Non-specials in special range	61
highBit now ANSI-compliant	61
Other added method	61
3.5 Stream and String related changes	62
String compare to ByteArray now returns false for matching bytes.	62

New ReadByteStream classes optimized for Strings	62
readStream reimplemented to return ReadByteStream	62
Support for Lz4 encode/decode added to String and ByteArray classes	62
ByteArray read/write of binary numbers in native format	63
Utf8 now disallows #readStream.	63
codePointAt:put: now performs auto-conversion if needed.	63
Other Changes	64
atOrNil: added	64
Stream classes consistently return objects of their collection class	64
3.6 JsonParser and JsonPetitParser	64
Upgrade impacts	64
3.7 External Session Changes	65
Added methods for GsTsExternalSession	65
Return value from GsTsExternalSession >> waitForReadReadyTimeOut:	65
Easier to use GEM_HALT_ON_ERROR in external sessions	65
3.8 Other Changes and Enhancements.	65
ProfMonitor improved formatting for object creation tree report.	65
GsSecureSocket example updates	65
isValidIdentifier, validateIsIdentifier optimized	66
Other Added methods	66
3.9 Deprecated and removed classes and methods.	66
Deprecated Methods	67
AbstractCharacter removed	67
Removed Public Methods	67
Removed Previously-deprecated Methods	67
Removed Private Methods	68

Chapter 4. Changes in the GCI Interfaces

4.1 GCI changes	69
GciDirtyTrackedObjs functionality removed	69
Removed GCI calls	69
Removed Image Methods related to TrackedObjects.	69
GsBitmap hiddenSetSpecifiers	69
Added Lz4 compression functions	70
GciCompressUsingLz4.	70
GciUncompressUsingLz4	70
Removed Mnemonics	70
4.2 GCI thread-safe changes.	70
Added Thread-safe function	71
GciTsNbPoll.	71
4.3 Foreign Function Interface (FFI) related changes.	71
errno: unsafe; use alternate methods to access to CCallout errno.	71
In-memory GC now responsive to memory use by FFI	71
CPreprocessor failed to cleanup temporary files	71
4.4 Updated Compile and Link Information	72

Linux Compile and Link Information.	72
AIX Compile and Link Information.	73
DARWIN Compile and Link Information	73
Windows Compile and Link Information	74

Chapter 5. Bug Fixes

Checkpoint when tranlogs full can result in Stone shutdown	75
Shrpcmon crash during crashed client frame lock recovery	75
Cache Warming Issues	75
Cache warming could have caused commit record backlog	75
Cache warming could hang	75
Race condition results in Gem SEGV in copyFrom:to:	75
Login related issues	76
A stuck login with UserSecurityData locked may block further logins	76
Slow timeout failure on login when network connection table is full	76
Improved reporting on client disconnect errors	76
Gems in login during stone shutdown may fail to exit	76
Login log did not record failed logins.	76
When commits suspended, symbol creation causes SymbolGem failure.	76
SecurityError on objectSecurityPolicy: for a class with no instance variables	76
Bugs in Configuration Parameters	77
Improved distribution over extents for DBF_ALLOCATION_MODE	77
Computation of SHR_PAGE_CACHE_NUM_PROCS too small.	77
Backup and Restore issues	77
Backup and restore failed to check OOP upper bound	77
Backup to a file with a nonexistent directory produced unclear error message 77	
Commit as restoreFromBackup starts could cause Stone shutdown.	77
Finding references failed to report objects in large IdentityBags	77
GsFile primitives not interrupted by SIGTERM	77
Upgrade fails with nonstandard decimalPoint Locale with GsPackagePolicy	77
Gem crash recovery may cause Stone to SEGV	78
After nbLogin, GsTsExternalSession wait methods SIGSEGVed or errorred.	78
Memory leak in GciTsNbLogin	78
TransactionBacklog signaled incorrectly	78
Object >> subclassResponsibility error was not reported correctly	78
stoned exited with code 3 on clean shutdown	78
Gems may be slow to shutdown due to memory checks	78
Hot standby related issues	79
Starting hotstandby continuous restore with obsolete tranlogs may crash	79
Hot standby logreceiver errors if there are other logreceivers for the same logsender	79
Problems after circular failOverToSlave	79
Checking for vote state incorrect	79
\$GEMSTONE paths containing symlinks that change	79

SymbolGem printed excessive messages to its log on GC 79

Unicode Compares causes strings to not be returned in ExportedDirtyList. . . . 80

FFI structures may fault out of memory and lose C data 80

performOnServer: failed for configurations with client s-bit set 80

Unnecessary commit conflicts. 80

 With Indexing operations 80

 On RcIdentitySet/Bag 80

Numeric and Time related issues. 80

 A Time did not compare equal to equivalent SmallTime 80

 SmallDateAndTime passivate activate failed for years 2035-2072. 80

 Number class >> parseLiterals:exponent: handling nil character argument. 80

 Division by SmallInteger minimumValue threw error 81

 ScaledDecimal, Fraction >> asFloat could return incorrect results 81

 CharacterCollection >> isDigits incorrect for empty string 81

 Float >> asDecimalFloat results not as precise as possible. 81

Protocol error on incomplete LGC_PAD_N read over socket. 81

topaz set cachename did not work for linked topaz 81

Class subclassed from nil did not have default doesNotUnderstand handling . . 81

Bitmap written by #saveWriteSetUnionToFile not readable 81

GsTestCase assert:isEquivalentTo: failed to catch errors. 81

Transaction log debug level change 82

PositionableStreamLegacy positionW, positionW: did not signal Warning 82

PPPParser >> child included a send to #assert: 82

DRAFT Release Notes for 3.7

Overview

GemStone/S 64 Bit™ 3.7 is a new version of the GemStone/S 64 Bit object server. Version 3.7 includes a number of important new features, including compiler optimizations, changes to process scheduling and debugging, new FileSystem support, multithreaded instance migration, and other features and bug fixes.

These Release Notes include changes between the previous version of GemStone/S 64 Bit, v3.6.5, and v3.7. If you are upgrading from a version prior to 3.6.5, review the release notes for each intermediate release to see the full set of changes.

For details about installing GemStone/S 64 Bit 3.7 or upgrading from earlier versions of GemStone/S 64 Bit, see the *GemStone/S 64 Bit Installation Guide* for v3.7 for your platform.

New keyfiles required

The keyfiles for v3.6.x and earlier cannot be used with v3.7; new keyfiles are required for this release. To obtain a new keyfile for GemStone/S v3.7, write to keyfiles@gemtalksystems.com. In your request, include your license information, platform and any updates to contact information.

Please contact GemTalk Technical Support if you have issues or questions.

Supported Platforms

Platforms for Version 3.7

GemStone/S 64 Bit version 3.7 is supported on the following platforms:

- ▶ Red Hat-compatible Linux 7.9, 8.7, and 9.1, and Ubuntu 20.04 and 22.04, on x86; Ubuntu 20.04 on ARM (Ubuntu on ARM is supported for development only)
GemStone is tested on a mixture of Red Hat, CentOS, Rocky and Alma; these are all considered fully certified platforms. Any reference to Red Hat applies to any Red Hat-compatible distribution.
- ▶ AIX 7.1 and 7.2
- ▶ OSX 13.1 (Ventura) with Darwin 22.2.0 kernel, and OSX 12.6 (Monterey) with Darwin 21.6.0 kernel on x86; and OSX 11.6 (Big Sur) with Darwin 20.6.0 kernel on Apple silicon.
(Mac is supported for development only)

Note that GemStone/S 64 Bit v3.7 was built using compilation features that are not available on older CPUs; v3.7 will not run on these CPUs, regardless of the Linux OS version. See page 18 for more details.

Distributions for Solaris/x86 and Solaris/SPARC are no longer available.

For more information and detailed requirements for each supported platforms, please refer to the *GemStone/S 64 Bit v3.7 Installation Guide* for that platform.

X509-Secured GemStone feature is fully tested and supported on Linux platforms only.

GemBuilder for Smalltalk (GBS) Versions

The following versions of GBS can be used with GemStone/S 64 Bit version 3.7:

GBS/VW version 8.6

VisualWorks 9.1.1 32-bit and 64-bit
<ul style="list-style-type: none"> ▶ Windows 10 ▶ RedHat ES 7.9, 8.7, and 9.1; Ubuntu 20.04 and 22.04

GBS/VA version 5.4.6

VAST Platform 11.0.1	VAST Platform 10.0.2
▶ Windows Server 2016 and Windows 10	▶ Windows Server 2016 and Windows 10

For more details on GBS and client Smalltalk platforms and requirements, see the *GemBuilder for Smalltalk Installation Guide* for that version of GBS.

VSD Version

The GemStone/S 64 Bit v3.7 distribution includes VSD version 5.6. The previous version of GemStone/S 64 Bit, v3.6.5, included VSD v5.5.4.

VSD version 5.6 includes a number of new features and bug fixes. For details on the changes, see the [Release Notes for VSD v5.6](#).

With GemStone/S 64 Bit v3.7, **statmonitor** now uses an additional file to improve lz4 compression, which is used by VSD when reading an lz4-compressed statmonitor data file. Further details on this change can be found on page 43.

As a result, lz4-compressed statmonitor files from v3.7 cannot be read by versions of VSD earlier than v5.6. VSD 5.6 can read statmonitor files generated in older versions of GemStone/S 64, 32-bit GemStone, and GBS, as well as those generated by GemStone/S 64 Bit v3.7.

VSD 5.6 is included with the GemStone distribution, and can also be downloaded as a separate product from <https://gemtalksystems.com/vsd/>.

GemBuilder for Java (GBJ) Version

The most recent version of GBJ, v3.1.3, can be used with GemStone/S 64 Bit v3.7.

GBJ 3.1.3 has not been tested on AIX with v3.7, although it is expected to work, and the distribution on AIX does not include the GBJ shared library. Contact GemTalk Technical Support if you require GBJ on AIX.

Rowan

The GemStone/S v3.7 distribution includes Rowan v 2.5.

Sparkle

Sparkle, the Pharo IDE for GemStone, provides GemStone development tools in the Pharo client Smalltalk environment. GemStone/S 64 Bit v3.7 is required with Sparkle.

Sparkle is under active development, and there are ongoing changes in GemStone server code to support Sparkle.

See the *Installation guide for Sparkle* in the Documentation subdirectory of the project, under <https://github.com/GemTalk/Sparkle/>.

Upgrade

Upgrade is supported from all 3.4.x, 3.5.x, and 3.6.x versions. Upgrade for 3.3.x is not disallowed, but is not tested or supported. To upgrade from earlier versions, upgrade to a supported upgrade origin version, and then upgrade to 3.7.

Note that **JsonParser** has changed definition in v3.7, and **GsHostProcess** changed definition in 3.6.5. If you have subclasses of these classes, or methods that reference either class directly or by name, you may need to recompile methods and/or edit source code, to ensure these methods work as intended. See the *Installation Guide for v3.7* for details.

Note that several obsolete or deprecated configuration parameters have been removed. These will now error when read, rather than being ignored. You will need to update legacy configuration files if you have not done so previously. See “Removed configuration parameters” on page 42.

For large systems with many symbols, in which upgrade performance is critical, we are now recommending to run System clusterAllSymbols in the old environment, prior to upgrade.

Documentation Changes

Documentation has been revised for this release, with modifications to incorporate new and changed features, as well as corrections and improvements.

In addition to the maintenance changes, and addition of new information and features, the following improvements have been made:

- ▶ The *Topaz User's Guide* has an additional chapter describing scripting, including information on superDoit scripts.
- ▶ The *System Administration Guide* section on Hot Standbys has been updated

The *GemStone/S 64 Bit X509-Secured GemStone System Administration Guide* has not been updated for this release.

1.1 Library and Distribution changes

Updated Library Versions

The version of OpenSSL has been updated to 3.0.8. This is a major version update.

The version of OpenLDAP has been updated to 2.6.4

The version of libssh has been updated to 0.10.4

The version of Kerberos has been updated to 1.20.1

The version of lz4 has been updated to 1.9.4

The version of prometheus has been updated to 1.0.1

The version of aws-sdk-cpp has been updated to 1.9.362

The version of zoneinfo has been updated to 2022d

The version of zlib has been updated to 1.2.12

A new library is now included, double-conversion, v3.2.1.

RPM distributions available

RPM (Redhat Package Manager) distributions are now available.

Since RPM installs as a read-only shared product tree, there are some limitations on locating configuration and log files as well as other GemStone files. See the new *Installation Guide for Linux RPM* for details.

When using an RPM installation, you must specify the configuration file using `-e`, `-z`, or set `$GEMSTONE_EXE_CONF` or `$GEMSTONE_SYS_CONF`.

To support simple installations, a new template configuration file, `$GEMSTONE/bin.default.conf`, has been added, which references an environment variable, `$GEMSTONE_DATADIR`, rather than `$GEMSTONE`.

clientLibs distributions

Along with the GemStone distribution, GemStone now includes a clientlibs directory tree. This provides a way to flexibly handling multiple versions of client library files. The structure of this file is:

```
clientlibs
  versionNum
    64bit
      specific library files
    32bit
      specific library files
```

When a new version of the server is released, the *newVersionNum* tree from that release can be added to the clientlibs directory without perturbing other uses.

Distribution Changes

\$GEMSTONE/projects moved

Earlier 3.6.x releases included the projects directory under \$GEMSTONE/upgrade/. This directory includes rowan-format source code for projects such as superDoit.

This directory has been moved up and is now a base directory, \$GEMSTONE/projects/.

Further changes in permissions within the distribution

The GemStone distribution included subdirectories and files with incorrect permissions; primarily non-executable files with execute permission; and some cases of subdirectories with owner write. These permissions have been corrected.

\$GEMSTONE/data and \$GEMSTONE/uilib continue to have owner and group write permission.

Example scripts for setting up GemStone in systemd

Under \$GEMSTONE/examples/admin/, the following directory and files have been added:

```
systemd
  netldi.service
  gemstone.env
  netldi.env
  gemstone.service
```

These scripts include the details needed to set up the GemStone and NetLDI to run as services managed by systemd.

1.2 Optimizations

New native code generation and compiler optimizations

The native code generator has been extensively improved for performance, and the compiler has been further optimized.

Some operations, such as integer arithmetic, have 2x-5x performance improvements, overall a 10% improvement can be expected.

Note that operations that are not CPU-bound, i.e. that are I/O bound waiting on disk reads or writes, will not observe performance improvements.

Native code not reimplemented on AIX

In this release, AIX runs in interpreted mode only.

Native code is available on Linux on ARM 64-bit and Darwin on Apple silicon, as well as Linux and Mac on x64_64.

Older CPU machines cannot run v3.7

GemStone has been compiled using optimizations that require CPU features that are not available on older CPUs (more than about 10 years old).

The minimum required for 3.7:

- ▶ Intel CPUs with architecture Sandy Bridge or newer (excluding Atom CPUs)
- ▶ AMD CPUs with architecture Bulldozer version 1 or newer

v3.7 cannot run on older CPUs. Attempting to run any executable will result in an error such as:

```
[Error]: Program cannot start because the CPU is missing the
following feature(s):
[Error]: avx pclmul
[Error]: For more information, please contact GemTalk technical
support.
```

Process Scheduling and Debugger Support

GemStone/S 64 Bit v3.7 includes additional infrastructure and changes in process scheduling and management, breakpoints, and stepping from v3.6.x. These differences includes both additional methods, and behavior changes in specific methods.

Scheduler now preemptive

The scheduling of a higher priority process previously required an explicit yield or wait in GemStone; now, GemStone does preemptive scheduling, similar to the behavior in Pharo and other Smalltalk dialects.

For example `ExecBlock >> forkAt` : no longer requires a subsequent yield, to allow a high priority `GsProcess` to start running. `GsProcess >> resume` and `GsProcess >> priority`: will automatically yield, if the receiver is or becomes higher priority than the current process.

Code with coordinating multiple processes may see differences in behavior in this version of GemStone.

Breakpoint Handling and Debugger Support

The GsProcess, Behavior and GsNMethod API that supports breakpoints and debugging has been redesigned, to improve and facilitate debugger implementation. Additional exception codes have been added to distinguish signaling to the GCI vs. signaling to Smalltalk. These changes are likely to affect behavior in corner cases, but should be generally transparent to end users. If you have implemented a debugger or made extensive customizations, contact GemTalk Engineering for more information.

Interrupting sockets

When a socket is waiting on activity (such as read or accept) in one GsProcess, while another GsProcess is running, previously the socket that is waiting could not get scheduled when the wait was complete and it was ready for further execution.

Now, a socket can be configured as interrupting. The following methods have been added:

```
GsSocket >> interrupting
GsSocket >> interrupting: aBoolean
```

After executing *aSocket* `interrupting: true`, then a GsProcess waiting for that socket to be ready for read or write in an instance method of GsSocket or a subclass of GsSocket will be scheduled to run when that socket is ready.

Process and debugging bugs fixed

A number of bugs have been fixed in the underlying code. These issues may or may not have been in previous releases.

- Step over in recursive method stopped in a frame of the same selector, which may be incorrect (#49539)
- Process terminate may be ignored if unwind block continuously runs. (#49464)
- GsProcess >> suspend did not suspend a process waiting for a delay or on a semaphore (#49526)
- GsProcess >> resume resumed a process that was waiting on a delay but not suspended (#49802)
- Terminated GsProcess does not allow ensure block in progress to complete, risking stuck semaphore (#48271)
- After AlmostOutOfStack, attempting to execute further encountered error clearing stack (#47373)
- GsProcess terminate on a process waiting on a socket without a timeout was not terminated correctly. (#47196)

Handling asynchronous events in a designated process

In a multi-process application, you may now configure a specific process to receive and handle these signals. The following methods have been added. For an example of how these can be used, see `GsSignalingSocket class >> _readNotificationExample`.

`GsSignalingSocket class >> newForAsyncExceptions: anArray`
 Returns a kind of `GsSignalingSocket` which has been registered to receive data representing asynchronous Notifications. The socket returned represents one end of an underlying UNIX domain socket pair, the other end of that pair is written to by the C thread that is the source of the Notification.

anArray must contain one or more of the classes `InterSessionSignal`, `ObjectsCommittedNotification`, `TransactionBacklog`, and `GcFinalizeNotification`, to specify the type of Notification to receive. Use `GsSignalingSocket >> readNotification` to receive the Notifications.

Do not use `InterSessionSignal class >> enableSignalling`, `ObjectsCommittedNotification class >> enableSignalling`, or `TransactionBacklog class >> enableSignalling` in conjunction with this method.

Only one socket can be registered in a session to receive asynchronous Notifications. `GsSignalingSocket class >> disableAsyncExceptions` must be used to cancel a previous registration before the next execution of `newForAsyncExceptions::`

`GsSignalingSocket >> readNotification`
 Read a notification from the receiver, which must be the result of executing `GsSignalingSocket class >> newForAsyncExceptions::`

`GsSignalingSocket class >> disableAsyncExceptions`
 Cancel a registration for receiving asynchronous notifications.

Added and renamed Process and related methods

GsProcess "Step" methods renamed

The methods `GsProcess >> step*FromLevel:*` have been renamed to more accurately describe their function; the new names are `setStep*BreaksAtLevel:*`. The earlier methods are still present and usable, but deprecated.

Signaling the GCI

The following methods have been added:

`AbstractException >> signalToGci`
 Receiver is signaled and it will not be trappable by any exception handler. An Error will be return to the GCI.

`AbstractException class >> signalToGci`
 An instance of the concrete subclass is created and signaled; not be trappable by any exception handler. An Error will be return to the GCI.

Other added methods

```
GsProcess >> abandonUnwindAndTrimStackToLevel:
GsProcess >> breakpointLevel
GsProcess >> breakpointLevel:
GsProcess >> terminateTries:eachTimeoutMs:
GsProcess >> trimStackToLevel:
GsProcess >> unsafeTerminate
GsProcess >> unsafeTrimStackToLevel:
GsProcess class >> current
GsMethod >> setBreakAtStepPoint:breakpointLevel:
ProcessorScheduler class >> highestPriority
ProcessorScheduler class >> lowestPriority
```

Removed Methods

The following related methods have been removed:

```
Breakpoint class >> trappable:
ExecBlock >> valueNowOrOnUnwindDo:
```

Other removed methods are listed under “Deprecated and removed classes and methods” on page 66.

Changes in Administration

This chapter describes changes and new features that apply to the administration of a GemStone Repository, and changes in the GemStone environment, configuration, and tools, including:

Multithreaded instance migration	24
One-time password	26
Changes in Users and ObjectSecurityPolicies	29
Changes in Backup and Restore	29
Hot Standby Changes	31
Garbage Collection Changes	34
Other Changes related to administrative operations	35
Changes in Errors	38
Changes in Utilities	39
Configuration Parameter Changes	42
Changes in Statmonitor and Cache Statistics	43

2.1 Multithreaded instance migration

Faster instance migration is now possible using the new migration API. This allows you to specify up to 2000 classes; every instance in the repository, for each of these classes, is migrated in a single operation within C code. Only simple migration operations on non-collection classes is supported. The primitive that performs this operation commits if successful, so this should be done when no other users would be committing, to avoid concurrency conflicts.

Multithreaded migration requires the MigrateObjects privilege (see “MigrateObjects privilege” on page 29).

This initial version of multi-threaded migration supports a limited set of migration.

Supported:

- ▶ mapping instance variable values to differently named instance variables
- ▶ set values to nil
- ▶ preserve or remove dynamic instance variables

Not supported:

- ▶ migrating indexable or NSC objects (that is, collections; including those with instance variables)
- ▶ modifications to specific dynamic instance variables
- ▶ mapping between dynamic and named instance variables

For migrations that are not supported by the multi-threaded migration function, you can continue to use the object based message-send migration options.

Multi-threaded migration is configured using the class `InstVarMappingArray`, which was previously a support class for instance migration that was not directly used. An instance of this class is the input to the multi-threaded migration operations.

To ensure that your migration has the results you expect and there are no unforeseen issues, it is recommended that prior to the actual migration, you test the migration with a more limited number of test objects. `Repository >> testMigrateMt:with:` allows you to perform the same migration as `Repository >> migrateMt:`, on a limited set of objects, and does not commit; after running this method commits are disallowed, so you must abort after running this test.

The following methods have been added:

```
Repository >> migrateMt: arrayOfMappings
```

Performs a multi threaded scan of the repository migrating instances of the classes found in the *arrayOfMappings* to their corresponding new classes. Up to 2000 classes can be migrated in a single operation. The classes must be committed. An error is generated if there are any modified persistent objects in the temporary object memory at the time it is executed.

The method will commit all of the migrated objects before it returns. It is possible that the commit may fail due to concurrency conflicts, so to avoid that the operation should be performed while there is no other activity on the system.

The *arrayOfMappings* is Array of instances of `InstVarMappingArray` which defines the classes that are to be migrated. For example:

```
SystemRepository migrateMt:
    {(InstVarMappingArray mappingFrom: oldClass to: newClass)}
```


Before executing this method you may wish to execute `testMigrateMt:with:` to validate that the migration does what is intended.

`Repository >> fastMigrateMt: arrayOfMappings`

Like `Repository >> migrateMt:`, but performs the migration aggressively, using more system resources, in order to complete more quickly.

`Repository >> testMigrateMt: arrayOfMappings with: testObjs`

This method is similar to `migrateMt:` and can be used to test the `migrateMt:` operation on a limited number of objects before performing a full repository scan and migrating all objects.

It performs a multi threaded scan of only the objects listed in the test array, migrating the objects in the same way that `migrateMt:` would, but it does NOT commit them. The session is left in a state with commits disallowed. Once the objects have been inspected and have been verified that the migration performed as expected the user should abort the current transaction.

InstVarMappingArray

The existing `InstVarMappingArray` class been modified and updated, and an instance can be created and customized directly for use by multi-threaded migration.

Note that ordinary object based migration continues to use instances of `InstVarMappingArray` without specific customization.

The following methods have been added:

`InstVarMappingArray >> mapInstVarNamed: oldName to: newName`

The value at the instance variable named `oldName` in the old class should be migrated to the instance variable named `newName` in the new class.

`InstVarMappingArray >> mapInstVarToNil: newName`

The migration should not retain the value of any instance variable `newName` present in the instance of the old class. After migration; the object will have a nil value at this instance variable.

The origin and target classes are now in instance variables, and methods have been added to access these:

`InstVarMappingArray >> oldClass`

`InstVarMappingArray >> newClass`

Example

For example, to create a mapping from `OrigClass` to `NewClass` that moves the value of the instance variable `oldName` to the location of the instance variable `newName` in the `NewClass`:

```
| im |
im := InstVarMappingArray mappingFrom: OrigClass to: NewClass.
im mapInstVarNamed: #oldName to: #newName.
im mapInstVarToNil: #oldName.
im preserveDynamic: false.
SystemRepository migrateMt: { im }
```

Note that if you do not specify to set `#oldName` to `nil`, the value at `oldName` will be retained in the migrated instance (now a `NewClass`) in `#oldName`, as well as being set in the instance variable `#newName`.

2.2 One-time password

One-time passwords allow login without further authentication

A new feature has been added, to allow a session to create a one-time, time-limited password for a specific `UserProfile`, which will allow a login as that user without further authentication. This allows, for example, to write code that uses `GsExternalSession` to login other sessions to divide tasks over multiple sessions, without needing to hard-code the GemStone password in code or query for it on the command line. The current session's authentication is "carried over" for a one-time authentication without re-entering a password.

Privilege to create one-time passwords

Only `userIds` with the new privilege `#CreateOnetimePassword` can create one-time passwords. This is granted by default to `SystemUser` and `DataCurator`.

Temporary passwords cannot be created to login as `SystemUser`.

One-time passwords for other UserProfiles

In addition to being able to login a second session as your own `UserProfile`, one-time passwords can also be created for other `UserProfiles`. Creating a one-time password for another `UserProfile` requires that the other `UserProfile` be on this session's allowlist of `UserProfiles`.

Adding a `UserProfile` to the allowlist requires `#OtherPassword` privilege. Once the whitelist is updated, `#OtherPassword` is not needed to create a temporary password, only `#CreateOnetimePassword` is needed.

Creating and Using a one-time password

Methods to create a one-time passwords are in `GsCurrentSession`. The one-time password is registered in the Stone, but not persisted (no commit is needed). To login the secondary session, the `UserId` and this one-time password are used; this allows a single login, within the specified timeout. After the successful login or the timeout has elapsed, the one-time password is no longer usable.

The GemStone GCI login functions `GciLogin()` and `GciTsLogin()` include a new flag, `GCI_ONETIME_PASSWORD`, that must be set in order to perform a login using a one-time password. This can be set using new methods in `GsTsExternalSession` and `GsExternalSession`, and a new command in topaz. The one-time password can be used to login using `GsTsExternalSession` or `GsExternalSession` (the expected use), and can also be used to login using topaz (linked or RPC). Other tools can also be used, provided that they support the `GCI_ONETIME_PASSWORD` flag.

Example 2.1 Using GsExternalSession to perform a login as SubordinateUser

The following example is executed as DataCurator, and presumes the existence of a GemStone UserProfile with the userId SubordinateUser.

First, DataCurator adds SubordinateUser to their allowlist:

```
System myUserProfile addOnetimePasswordUserId:  
  'SubordinateUser'.
```

Then, a temporary password is created:

```
tempPW := GsCurrentSession currentSession  
  createOnetimePasswordForUserId: 'SubordinateUser'  
  validForSeconds: 30.
```

To use this password to login an external session:

```
sess := (GsTsExternalSession newDefault)  
  username: 'SubordinateUser';  
  onetimePassword: tempPW;  
  yourself.  
sess login.
```

New topaz command

To use one-time passwords in topaz, a new command has been added.

set onetimepassword onOrOff

Determines if the password is a normal GemStone password or a special one-time password. When OFF, the default, the password field is interpreted as a normal GemStone password. When ON, the password field is interpreted as a onetime password.

Tracking one-time passwords

Login log

The login log feature allows you to enable logging of logins and logouts for a repository, using the STN_LOGIN_LOG_ENABLED configuration parameter.

The file generated when login logging is enabled now includes two additional fields. A login record for a session that was logged in using a one-time password include the OOP of the UserProfile that requested the one-time password, and the pid of the session that requested it. Other entries in the login log include 20 (nil) and -1 in these slots.

Added Cache Statistics

The follow cache statistics have been added:

OnetimePasswordsActive (Stn)

Number of onetime passwords in the stone's hash map.

OnetimePasswordsCreated (Gem)

Number of onetime passwords this session has created.

OnetimePasswordsExpired (Stn)

Number of expired onetime passwords automatically removed by stone.

OnetimePasswordsNotValidated (Stn)

Number of failed onetime password validations.

OnetimePasswordsTotal (Stn)

Number of onetime passwords ever created by any session since the stone was started.

OnetimePasswordsValidated (Stn)

Number of successful onetime password validations.

OnetimePasswordUsed (Gem)

A boolean with value 1 if the session used a onetime password to login, otherwise 0.

Added methods

```

GsCurrentSession >> createOnetimePasswordForUserId: aUserId
    validForSeconds: seconds
GsCurrentSession >> createOnetimePasswordForUserProfile: aUserProfile
    validForSeconds: seconds
GsCurrentSession >> createOnetimePasswordValidForSeconds: seconds
GemStoneParameters >> onetimePassword: pword
GemStoneParameters >> onetimePasswordFlag
GemStoneParameters >> onetimePasswordLoginFlags
GemStoneParameters >> setLoginWithOnetimePassword
GsExternalSession >> onetimePassword: aString
GsTsExternalSession >> onetimePassword: aString
UserProfile >> addOnetimePasswordUserId: aUserId
UserProfile >> addOnetimePasswordUserProfile: aUserProfile
UserProfile >> addOnetimePasswordUserProfiles: anArray
UserProfile >> onetimePasswordUserProfiles
UserProfile >> removeAllOnetimePasswordUserProfiles
UserProfile >> removeOnetimePasswordUserId: aUserId
UserProfile >> removeOnetimePasswordUserProfile: aUserProfile
UserProfile >> removeOnetimePasswordUserProfiles: anArray

```

Changes to descriptionOfSession:

The results of `System >> descriptionOfSession:` now includes additional fields (note that a new entry 25 has also been added; see page 36):

26. UserProfile which created the onetime password for the session, or nil if no onetime password was used.

27. Process ID of the gem that created the onetime password for the session, or -1 if no onetime password was used.

Additional fields in Login log

The Login log (enabled by the configuration parameter `STN_LOGIN_LOG_ENABLED`), reports two additional fields for each entry.

ParentUserProfile - OOP of the UserProfile that created the token used for login or 20 (OOP_NIL) if no token was used.

ParentPid - process ID of the parent session that created the token used for login or -1 if no token was used.

One-time logins that fail to login due to a mismatch in UserId and one-time password, or an invalid or expired password are also reported.

New Privilege

The privilege `#CreateOnetimePassword` has been added. This privilege is granted to DataCurator and SystemUser, but existing and newly created users will not have this privilege.

2.3 Changes in Users and ObjectSecurityPolicies

CreateOnetimePassword privilege

This new privilege allows creating a one-time password that allow another session to login; see “One-time password” on page 26 for details.

MigrateObjects privilege

This privilege was present in earlier versions, but not implemented. This privilege is required for the new optimized instance migration; see “Multithreaded instance migration” on page 24.

2.4 Changes in Backup and Restore

Full control over compression levels

GemStone programmatic `fullBackup` and `secureFullBackup` allow you to compress backups using `gzip` (`zlib`) or `lz4`. Generally speaking, `lz4` backups are faster but do not compress as much.

Both `gzip` and `lz4` support different compression levels, with lower compression level values providing faster performance with less compression. `gzip` supports compression levels 1 to 9, while `lz4` supports 0 to 12. For GemStone backups, the upper ranges provide little addition compression and take much longer.

Consistent default for Lz4 to level 0

In previous releases, the default for `lz4`-compressed `fullBackup` was level 3. Since the reason for using `lz4` rather than `gz` is generally performance, this has been changed to

default to level 0. This is the same compression level as the default for secure lz4 compressed backup.

Specifying level in existing methods

The existing methods:

```
Repository >> fullBackupTo: fileNames MBytes: mByteLimit
                compressKind: anInt bufSize: count
Repository >> secureFullBackupTo: fileNames MBytes: mByteLimit
                compressKind: anInt bufSize: count encryptKind: encKind ...
```

now accept for *anInt* not only 0, 1 (for gzip), and 2 (for lz4), but also additional values to indicate compression level. 101... 109 indicate various gzip compression levels, and 200--212 indicate the various lz4 compression levels.

New methods including level argument

New methods have been also added to allow you to specify the particular level of compression you wish to use. These methods are variants of existing methods with the added `compressLevel`: argument.

```
Repository >> fullBackupGzCompressedTo: fileNames MBytes: mByteLimit
                compressLevel: cLevel
Repository >> fullBackupLz4CompressedTo: fileNames MBytes:
                mByteLimit compressLevel: cLevel
Repository >> secureGzFullBackupTo: fileNames MBytes: mByteLimit
                compressLevel: cLevel bufSize: count encryptKind: encKind
                publicKeyCerts: anArrayOfString signatureHashKind: hashKind
                signingKey: signingKeyFn signingKeyPassphrase: aPassphrase
                numThreads: numThreads
Repository >> secureLz4FullBackupTo: fileNames MBytes: mByteLimit
                compressLevel: cLevel bufSize: count encryptKind: encKind
                publicKeyCerts: anArrayOfString signatureHashKind: hashKind
                signingKey: signingKeyFn signingKeyPassphrase: aPassphrase
                numThreads: numThreads
```

Repository scan operations now allowed in solo sessions

A number of repository scan operations that were previously disallowed in solo sessions can now be used.

This includes Repository methods for `allInstances*`, `listInstances*`, `allReferences*`, `listReferences*`, and related methods; and `GsObjectInventory` scans.

In restore mode, compiler now generates arguments and temporaries as Strings

When a repository is in restore mode, it cannot create new Symbols, since Symbols are canonicalized and require the `SymbolGem` to commit, which would create a state difference between the restoring repository and the repository it is restoring from. This

created problems executing code that includes arguments or temporary variables with names that did not already exist in the repository, since variable names are created as symbols for efficiency.

Now, when the repository is in restore mode, temporary variable names in compiled code are created as Strings, which will not require the SymbolGem commit. (#49984)

2.5 Hot Standby Changes

suspendCommitsForFailover no longer supported

Repository >> suspendCommitsForFailover should no longer be used; instead, use the alternate method Repository >> failOverToSlave.

suspendCommitsForFailover allowed you to keep the master system running as a master, while you created a forked duplicate system on the slave node. After suspendCommitsForFailover to setup the former slave as a master, you have two masters that are not compatible, since they have separate, incompatible tranlog sequences.

With suspendCommitsForFailover, to setup the former master as a slave of the new master, you had to restore a new backup of the former slave, new master, into the former master.

If you do want to continue running the master as an independent fork along with the system running on the former-slave, execute commitRestore on the former master after executing failOverToSlave.

Hot Standby with failover and multiple slaves

A number of bugs have been fixed relating to hot standby; see “Hot standby related issues” on page 79.

As in previous releases, it is possible for multiple slave stones to be restoring from a single master stone. While only one logsender (the **primary logsender**) can attach to the master stone, the primary logsender may service up to 5 slave logreceivers. Additional logsenders can run on the master node; these read from the tranlogs, but are not attached to the Stone, so the master Stone is not aware of them, and they are not notified if, for example, another tranlog directory is added.

The first Slave system’s logreceiver to connect to the Master’s primary logsender is the **primary logreceiver**, and its Stone is the **primary Slave**. This is the system for which failOverStatus is reported.

A failOverToSlave operation prepares all slaves to become the new master; whichever slave performs the commitRestore will be the one that is the new master.

Repository >> failOverStatus improved report

The failOverStatus method results now include more details about the status of the master Stone executing it, and the primary slave Stone being restored into. This includes the most recent tranlog record as well as the most recent checkpoint, and other critical information on the master repository. This allows you to verify the exact restore status of the primary slave and compare to the status of the current master.

For example:

```
This repository last tranlog commit file 6 record 206
  last checkpoint file 6 record 15
Other repository restored to commit file 6 record 77
  last restored checkpoint file 6 record 15
This repository commits allowed
```

Note that this may be misleading if you have multiple slaves, since only the primary slave's details are reported. If you have more than one slave, you may check on intended failover slave itself, or use `gslist -j -v logsender`, as described in the next section.

Full details in JSON from `gslist`

logsenders and logreceivers write keys into the locks directory, which allows them to be queried by `gslist`. This provides a way to see the status of all logsenders or logreceivers.

`gslist` has a new features which allows you to query for information on a logsender and all the logreceivers attached to it; using `gslist -j -v` (both options are required).

For example, the following is the output on the Stone's node, for a system running on host benton, with a primary slave on host fossa and a secondary slave on host santiam, on the stone's node. The "statusInMasterCache:true" indicates that this is the primary slave, whose details are reported by `failOverStatus`.

```
unix > gslist -v -j logsender_57222
{"GemStoneServers": [
  {
    "Name": "logsender_57222",
    "Host": "benton",
    "HostId": "69621bb0476b1937",
    "Ip": "10.94.441.15",
    "Status": "OK",
    "Type": "Logsender",
    "Version": "3.7.0",
    "Creator": "lalmarod",
    "Started": "2023-01-12T18:13:29.000-08:00",
    "Pid": 15613,
    "Port": 57222,
    "Options": {
      "-d": null,
      "-P": "57222",
      "-A": "benton.gemtalksystems.com",
      "-l": "/benton/admin/logs/logsender57222.log",
      "-s": "bentonstone"
    },
    "LogName": "/benton/admin/logs/logsender57222.log",
    "Sysconf": null,
    "Execonf": null,
    "GEMSTONE": "/benton/admin/GS6437",
    "Exe": "/benton/admin/GS6437/sys/gem",
```



```

    "Logreceivers":[
      {
        "statusInMasterCache":"true",
        "peerIp":"10.94.441.127",
        "peerHost":"fossa.gemtalksystems.com",
        "processId": 3637081,
        "replayedCommit":{"file": 4,"record": 5235},
        "replayedCheckpoint":{"file": 4,"record": 5235}
      },
      {
        "statusInMasterCache":"false",
        "peerIp":"10.94.441.104",
        "peerHost":"santiam.gemtalksystems.com",
        "processId": 791089,
        "replayedCommit":{"file": 4,"record": 5235},
        "replayedCheckpoint":{"file": 4,"record": 5235}
      }
    ]

```

Upgrade via hot standby

It is now possible to perform a GemStone upgrade of a hot standby system, by upgrading the slave, failing over to make this the master, and allowing the former master, new slave to be upgraded via transaction logs.

This does not require changes in 3.6.x; the changes to support this feature are added in v3.7, and this process can be used to upgrade from 3.6.5 to 3.7.

Note that this process is newly developed and may need refinement, and has not been verified in a multiple slave system.

To upgrade via hot standby:

1. Stop the slave system's logreceiver and the slave stone. Restart both using the v3.7 binaries, and restart the continuous restore. Do not perform the **upgradeimage** yet.

The slave will continue to restore transaction from the master, although logins will report version mismatch errors. Login as SystemUser is allowed. These errors are temporary until the upgrade process is complete.

2. On the master, which is still on the original version, perform `failOverToSlave`. This stops further commits, checkpoints, and puts the now former master into restore mode.

Commits are now disallowed on the master.

3. Wait for the master's failover transaction to be replayed on the slave. You can determine this using `SystemRepository restoreStatusInfo at: 13`, which returns a non-zero timestamp when the failover transaction is replayed. On the slave system, execute:

```
[ 0 == (SystemRepository restoreStatusInfo at: 13) ]
  whileTrue: [ System _sleepMs: 500 ].
```

Once the failover transaction has been replayed, stop the master stone and the logsender. The former-master stone may be left running to service read-only operations, until Step 6.

4. On the slave, stop the logreceiver, and perform the `stopContinuousRestore` and do `commitRestore`. This makes the now former slave into the new master.
5. On the new master (the former slave), perform the `upgradeImage` and any other required upgrade steps.
The new master is now available for logins.
Start the logsender on this new master.
6. On the former master, now slave system restart the stone and logreceiver using the 3.7 binaries, and start continuous restore (`continuousRestoreFromArchiveLogs`). Do not perform `upgradeImage`.
As the transactions coming from the new master are replayed on the new slave, the new slave will be upgraded to v3.7. There will be a lag after the transactions are restored before next checkpoint completes, and the version is updated. Until this time, logins to the new slave system will see a version mismatch error.

2.6 Garbage Collection Changes

markForCollection performance improvement for repositories with long object chains

When a repository has very long object chains, performance of `markForCollection` can become slow. The code was improved substantially in v3.6.5 ([bug 50025](#)). v3.7 includes additional improvements; thread sleep time has adjustments to moderately improve performance under these circumstances.

Number of threads configured for AdminGem

The configuration file value for `STN_NUM_GC_RECLAIM_SESSIONS` is now used to ensure that the AdminGem uses a number of threads appropriate for the repository configuration to perform write set union sweep and epoch.

In 3.7, the Stone passes the value for `STN_NUM_GC_RECLAIM_SESSIONS` to the AdminGem. The larger of the AdminGem's setting for `#epochGcMaxThreads` or `#epochGcMaxThreads` and this argument used to determine the number of AdminGem threads used to perform the operation.

Renamed method to configure number of reclaim threads

A new method has been added:

```
System class >> changeNumberOfReclaimThreads:
```

Which has identical behavior as the existing method `System class >> changeNumberOfReclaimGemSessions:`. The older method continues to be usable and is not deprecated. The new method is preferred; the name reflects the current reclaim implementation.

Reclaim GcGem Logging

The ReclaimGem configuration option `#verboseLogging` now has multiple levels. Higher levels may produce a very large amount of output for debugging errors and evaluating performance, and are not recommended for general use.

For compatibility with earlier versions of GemStone, in which only booleans were expected true and false are accepted; false = 0, true = 1

These are the definitions for the levels:

- 0 - only summary information every 5 minutes in slow, 15 minutes in a fast
- 1 - log sigAbort, gcGemAlert, gcHwPage, etc
- 2 - numDead from stone, numDead processed per thread
- 3 - summary: commitCount, stayInTrans, crPage, numLive, numDead, numReclPages
- 4 - abort, start phases
- 5 - conflict set
- 6 - num live and dead found per page
- 7 - summary of info added to ot slot
- 8 - deltas in rootPage
- 9 - push live or dead

The default ReclaimGem logging (with level 0) is now on one line, making the log easier to read. Note at this level, summaries of reclaim operations are printed every 15 minutes.

Determining sessions that have not voted

The following method has been added:

```
System class >> notVotedSessionsReport
    Returns a String describing sessions that have not voted since the last epochGc or
    markForCollection.
```

2.7 Other Changes related to administrative operations

Cache Warming

Cache warming during upgrade

During the upgrade process, after the stone has started but before `upgradeImage` has completed, logins may fail with a bad Gem version error, which prevents cache warming from happening. Since cache warming may result in faster upgrade, it is now allowed for cachewarmers to login before `upgradeImage` has completed.

Deprecated cache warming configuration parameters removed

The deprecated configuration parameters `STN_CACHE_WARMER` and `STN_CACHE_WARMER_SESSIONS` have been removed in v3.7. To run cache warming automatically on stone startup, you must use the configuration parameter `STN_CACHE_WARMER_ARGS`.

Determining status of cache warming

The following methods have been added, which report the SessionId and other information for active cache warmers. This allows you to detect when startup cache warming is complete.

```
System class >> cacheWarmerSessions
    Returns an Array of sessionIds of cache warmer sessions.

System class >> cacheWarmerSessionsReport
    Returns a String describing cache warmer sessions.
```

Changes in abort behavior in solo session

A solo session (a GemStone login that does not connect to a Stone process) has no transactional behavior nor ability to make persistent changes. `System class >> commitTransaction` has always be disallowed in solo sessions.

When converting existing scripts to solo, calls to `System class >> abortTransaction` or `System class >> beginTransaction` that had been present within the scripts previously, could result in losing changes to persistent objects that cannot be committed in a solo session (such as adding classes to `UserGlobals`).

To make code conversion more reliable, `abortTransaction` now signals an error if invoked in a solo session where changes to persistent objects would be lost. This allows you to find and remove the abort.

A new method, `System class >> soloAbort`, can be used within a solo session to perform a logical abort.

`System class >> beginTransaction` is also disallowed, if in a solo session and changes to persistent objects would be lost.

Changes in session information for GcLock

Up to five sessions can hold the GcLock. The method `System class >> sessionIdHoldingGcLock` and the statistic **SessionWithGcLock** reported only the last session to get a lock, and was cleared whenever a session released its lock, even if other sessions held a lock.

`sessionsHoldingGcLock` remains, but is deprecated.

The following methods have been added:

```
System class >> sessionsHoldingGcLock
    Returns an Array of sessionIds of sessions holding a garbage collection or
    repository scan lock.

System class >> gcLocksCount
    Returns a SmallInteger which is the number of sessions holding a garbage
    collection or repository scan lock, or 0 if the lock is free.

System class >> gcLocksReport
    Returns a string report on the GcLocks.
```

System class >> descriptionOfSession: now includes an additional field providing the session lock held by the give session:

25. gcLockKind. 0 or the type of gcLock or repository scan lock held by the session.

The associated cache statistics **SessionWithGcLock** is replaced by **SessionsWithGcLock**, and **GcLockKind** is now a per-session stat, describing the kind of lock that this session is holding.

Methods added to determine recursive size of an object tree

The following methods have been added to support repository analysis.

Object >> recursiveSize

Returns an Array of the form { *size* . *numberOfObjects* } where *size* is approximate physical size on disk of an object and all of the objects referenced by it and *numberOfObjects* is the count of those objects. Excludes referenced objects in the symbolList, referenced Symbols and referenced special objects.

Object >> recursiveSizeInMemory

Returns an Array of the form { *size* . *numberOfObjects* } where *size* is approximate physical size on disk of an object and all of the objects referenced by it and *numberOfObjects* is the count of those objects. Excludes committed objects in memory unless those objects have uncommitted changes. Excludes referenced objects in the symbolList, referenced Symbols and referenced special objects.

Object >> recursiveSizeInMemoryReport

Return the result of recursiveSizeInMemory as a readable String.

Object >> recursiveSizeReport

Return the result of recursiveSize as a readable String.

Fast variants of findAllReferences

The following methods have been added:

Object >> fastAllReferences

Same as Object >> findAllReferences, but uses more threads.

Object >> fastPrivateReferences

Same as Object >> findAllReferences, but uses more threads and returns references from internal nodes of large objects and Nscs.

2.8 Changes in Errors

Added Errors

The following exceptions have been added:

- RT_ERR_STEP_St 6023
Single-step breakpoint signalled to Smalltalk.
- RT_ERR_CODE_BREAKPOINT_St 6024
Method breakpoint signalled to Smalltalk.
- RT_ERR_STACK_BREAKPOINT_St 6025
Stack breakpoint signalled to Smalltalk.

Added Error and Exception Classes

The **SshSocketError** class, associated with the new feature “Support for ssh and sftp using OpenSSL”, starting on page 58, was added to support the error ERR_SshSocketError 2759. ERR_SshSocketError was present, but reserved, in previous releases.

The error **GcFinalizeNotification**, which is private to GemStone internals, has been added; it is used in finalization of Ephemeron.

A number of Exception classes have been also been added as part of FileSystem.

Removed Errors

The following errors have been removed:

- RT_ERR_TERMINATE_PROCESS 6018
- RT_ERR_TRACKED_OBJS_NEEDS_INIT 2381

Cannot return now not resumable

The **CannotReturn** error returned true for `isResumable`, although this error could not actually be resumed. (#49677)

OS-origin error messages changed

Previously, GemStone provided it's own strings to describe errors that originate from the OS, for example, "ENOENT, The file or directory specified cannot be found".

Now, the wording provided by the OS (on Unix) is used, which may be different; for example, "No such file or directory".

2503 changed from FloatingPointError to NumericError

This error number is now associated with the class NumericError.

AbstractException >> signalToGci

The following method has been added:

- `AbstractException >> signalToGci`
Receiver is signaled and it will not be trappable by any exception handler. An Error will be return to the GCI.

Print GCI trace records on protocol error

To aid in analysis of GCI protocol errors, recent GCI call history is now automatically printed to the gem log when a GCI protocol error occurs. GCI call trace records are maintained in a wraparound buffer, and the most recent 100 records will be printed to the log.

MFC Warning now includes possibleDeadSymbols

The warning signaled by `markForCollection` now includes the `possibleDeadSymbols` count, as well as live and dead object counts. This change was present in 3.6.4, but not documented.

2.9 Changes in Utilities

Many utilities faster on Linux

The `startstone`, `stopstone`, `stopnetldi`, and `waitstone` utilities have been tuned and optimized for modern hardware, and require much less overhead time than in previous releases. In practice, the amount of performance improvement will depend on the specific `startstone` tasks for an application.

largememorypages improved and usable for remote caches

The calculations of the number of large memory pages was not entirely accurate, and could produce too low values for very large caches. (#50216).

In addition, the `-r` option has been added to `largememorypages`, which calculates the page requirements for a remote cache, which is somewhat smaller than for the Stone's cache.

Changes in argument handling

When using a configuration file as an argument to `largememorypages`, the configuration file must have large pages enabled. If `SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_POLICY` is not enabled in a `-z` or `-e` argument, the `largememorypages` output includes cache configuration information, but no large page size recommendations.

In addition, the `-p` or `SHR_PAGE_CACHE_LARGE_MEMORY_PAGE_SIZE_MB` in a `-z` or `-e` argument no longer defaults; you must now specify the page size on Linux and AIX.

secure_backup_extract added

The **secure_backup_extract** utility provides a way to verify the signature of a secure backup, that can be used on secure backup files from any earlier version (the secure backup feature was added in v3.4), without requiring a GemStone environment. This utility is available on little-endian hosts only (Linux and Mac).

Usage:

```
secure_backup_extract (sig | cert) secureBackupFile outputFile
where:
  sig - extract the digital signature and write to outputFile
  cert - extract the public signing key and write to
        outputFile
  secureBackupFile - a valid GemStone secure backup
  outputFile - location to write script results; this file
               must not already exist
```

secure_backup_extract extracts the public signing cert or the digital signature from a secure backup file, and writes it to a new file. While \$GEMSTONE is not needed, you must have openssl and perl available.

This script produces the same result as **copydbf**, using **copydbf -X** to extract the public signing cert, or **copydbf -Y** to extract the digital signature. To verify the digital signature, compare the binary files use the unix utility **cmp**.

For example, if you have a secure backup file *backup.sdbf* that was signed using GemStone's example private cert *backup_sign_2_clientkey.pem*, either of the following will extract the public signing cert:

```
secure_backup_extract cert backup.sdbf signingCert.secbkext
copydbf -X signingCert.copydbf backup.sdbf
```

This can be compared with each other, or against the signing cert used to make the backup; for example:

```
cmp -l signingCert.secbkext $GEMSTONE/examples/openssl/certs/backup_s
ign_2_clientcert.pem
```

To extract the digital signature itself:

```
secure_backup_extract sig backup.sdbf digitalSig.secbkext
copydbf -Y digitalSig.copydbf backup.sdbf
```

which can be compared for validation:

```
cmp -l digitalSig.secbkext digitalSig.copydbf
```

Topaz changes

Strings containing now print hex value rather than space

Character codePoint: 16rA0 is the non-blocking space . If this character is included in a string, previously topaz displayed this the same as a standard space, code point 16r20. Now, non-blocking spaces will be displayed as '\ua0' to avoid confusion.

White space characters outside of the ASCII range, which includes codePoint 16rA0 and others, are not considered whitespace by the GemStone compiler.

LIMIT added option LEV2OOPS

A new option has been added to the **limit** command:

limit lev2oops *aNum*

When **level** is set to 2 or greater, or **obj2** is used, the **lev2oops** limit controls how many oops to display of instVar values.

SET added option ONETIMEPASSWORD

An option has been added to support the one-time password feature (see page 26).

set onetimepassword *onOrOff*

Determines if the password is a normal GemStone password or a special one-time password. When OFF, the default, the password field is interpreted as a normal GemStone password. When ON, the password field is interpreted as a onetime password.

Improvements to debugging using DEBUGGEM

The topaz **debuggem** command allows you to attach topaz to an executing process for debugging. v3.7 has new features that make this easier.

System waitForDebug

The added method `System class >> waitForDebug` enables `GEM_LISTEN_FOR_DEBUG`; it is not necessary to use the configuration parameter, nor to invoke `listenFoeDebugConnection`. It then immediately waits, so another session can attach. When using this method, the **debuggem** command is the same, but there is no need to perform the **stack set** in the debugging session, as was necessary when using **topazwaitfordebug**.

System cancelWaitForDebug

Executing the added method `System class >> cancelWaitForDebug` in the executing session cancels the effect of the waiting part of the `waitForDebug` method. On logout, the executing session will resume running. This is not needed if using the new **detach** topaz command.

Added Topaz commands

kill

In the debugging topaz session (the session created by **debuggem**), execute a `kill -TERM` against the executing gem or topaz -l (the session being debugged).

detach

In the debugging topaz session (the session created by **debuggem**), disconnects from the executing gem or topaz -l (the session being debugged) and logs out. The `GsProcess` being debugged will be continued after the logout, and `System class >> waitForDebug` will return to the caller.

2.10 Configuration Parameter Changes

Removed configuration parameters

The following obsolete and nonfunctional parameters have been removed:

GEM_PRIVATE_PAGE_CACHE_KB

STN_PRIVATE_PAGE_CACHE_KB

GEM_RPCGCL_TIMEOUT

The following legacy parameters have also been removed; applications should use the parameter STN_CACHE_WARMER_ARGS for automatic cache warming on startup.

STN_CACHE_WARMER

STN_CACHE_WARMER_SESSIONS

STN_NUM_GC_RECLAIM_SESSIONS now applies to Admin Gem

The number of sessions used by the Admin GcGem now may be determined using the setting for reclaim sessions, which is calculated based on the setting for STN_MAX_GC_RECLAIM_SESSIONS. See “Number of threads configured for AdminGem” on page 34.

GEM_NATIVE_CODE_ENABLED

Previously, a setting of 0 disabled native code, 1 enabled native code, and 2 enabled native code with inlining some math primitives.

Now, settings of 1 and 2 both produce pure native code. A setting of 1 is reserved for future use and interpreted as a 2.

STN_GEM_TIMEOUT

If the value of this parameter is 0, the login timeout used was previously 5 minutes, and is now 1 minute.

2.11 Changes in Statmonitor and Cache Statistics

The information written on cache statistics has been extensively revised and reorganized, and new Process Types added to reflect the different responsibilities of different GemStone processes or process threads. In particular, the general Pgsvr process type has been divided up into the different specific page servers (cache, free frame, AIO, etc.), and the PageManager has been divided from the Stone. A number of obsolete statistics have been removed, and new statistics have been added.

Performance improvement for Shared Page Cache Monitor

Statistics collection in the Shared Page Cache Monitor has been adjusted so expensive per-frame statistics are collected less frequently, not more than once per second, to reduce overhead for large systems.

Improved lz4 compression and new VSD required

Compressing statmonitor data with lz4 now has improved compression (between 3% and 25%) by using a C dictionary; this dictionary is in the distribution as `$GEMSTONE/bin/lz4.statmon.dict.bin`, and is used by VSD to read lz4-compressed statmonitor data files in VSD v5.6 and later.

Statmonitor files generated by 3.7 and compressed using lz4 can only be read using VSD v5.6 or later. VSD 5.6 can read lz4-compressed statmonitor files generated by earlier versions of GemStone, and older versions of VSD can read statmonitor files generated by GemStone v3.7 that are not compressed or that are compressed using gz.

To uncompress an lz4-compressed statmonitor data file on the command line, you must use an additional `-D $GEMSTONE/bin/lz4.statmon.dict.bin` argument; for example:

```
unix> lz4 -D $GEMSTONE/bin/lz4.statmon.dict.bin stats.out.lz4
```

Statmonitor -z -Z arguments now accept compression level

The **statmonitor -z** (gzip) and **-Z** (lz4) compression arguments now accept an optional argument to specify the compression level. Higher levels provide greater compression, at the expense of performance.

-z accepts 1...9; the default, if no level is included, is 6.

-Z accepts 1...12; the default, if no level is included, is 3.

Statistics Types reorganized

Several Process Types have been split up. For statistics that applied to the previous general process type, these are now associated only with the relevant process types. For example, the PageMgr* stats that were previously of type Stn, since the PageManager is a thread in the Stone, are now specifically of type PageMgrThr.

Pgsvr split

The general type Pgsvr no longer exists. It is replaced by:

StnAioThr (AIO PageServer threads in Stone)

FreeFrameThr (Free Frame PageServer threads in Stone)

CachePgsvr (Cache PageServer)

GemPgsvrMain, GemPgsvrThr (Gem PageServer)

Statistics that were previously of type Pgsvr will now apply to one or more of these new types.

Stone, Page Manager, and Stone's Restore Threads split

The PageManager has been a thread in the Stone for some number of releases, but statistics for PageManager were previously recorded under the same statistics type as the Stone. Now, the PageManager has its own Type, **PageMgrThr**.

Note that since the PageManager is a thread, not a process, host process stats are not collected for the PageManager. To see host process stats, you must look at the Stone's statistics.

The restore threads used during restoreFromBackup have also been made a new process type, **RestoreThr**.

Gems split

For threads within the Gem, the **GemThr** type has been added, to distinguish from the **Gem** type. Host system stats are not collected for **GemThr**.

A new method has been added to list statistics that apply to Gem threads:

```
System class >> cacheStatisticsDescriptionForGemThreads
  Returns an Array of Strings describing cache statistics applicable to threads started
  by gems for certain repository-wide operations such as markForCollection, full
  backup, etc.
```

Added API for PageManager statistics

```
System class >> cacheStatisticsDescriptionForPageManager
  Returns an Array of Strings describing cache statistics applicable to the stone's
  page manager thread.
```

```
System class >> pageManagerCacheStatistics
  Return the cache statistics for the page manager thread if this session is located on
  the same host as the stone process. Only cache statistics applicable to the page
  manager thread are returned.
```

```
System class >> pageManagerCacheStatisticWithName: aString
  Return the value of the cache stat with the given name, which must match an
  element of the Array returned by the #cacheStatisticsDescription method
  applicable to the page manager thread. The UserTime and SysTime statistics
  cannot be accessed using this method.
```

```
System class >> pageManagerProcessSlot
  Answer the cache slot for the page manager thread or -1 if the slot was not
  found. Page manager exists only on the shared page cache used by the stone
  process.
```

Changes in memory statistics on Linux

The existing statistics to capture memory use on Linux did not reliably capture important information, and in many cases were misleading.

Now, on Linux, **statmonitor** collects memory use related statistics from `/proc/<pid>/smaps` instead of `/proc/<pid>/maps`.

There is some performance impact, since it is more expensive to read `smaps` than `maps`, particularly on large systems with many processes and short sampling intervals. A new thread in `statmonitor` is dedicated to reading `smaps`; this thread is designed to consume a upper maximum of 25% of 1 core, although it will normally take less; much less on smaller systems.

Per-process statistics

The following per-process statistics have been removed:

- ImageKBytes**
- MaxImageSize**
- MaxRSS**
- RSSData**
- RSSKBytes**
- RSSLib**
- RSSStack**
- RSSText**

The following are the new per-process statistics. Note that these are only collected for the main thread within the Gem.

PrivateClean (All)

The amount of memory in kilobytes mapped to the process which has not been modified since it was mapped.

PrivateDirty (All)

The amount of memory in kilobytes mapped to the process which has been modified since it was mapped.

PSSSize (All)

The proportional set size of the process in kilobytes. PSS includes memory counted in `PrivateClean` and `PrivateDirty` as well as a portion of the memory shared by this process and with one or more other processes.

Swap (All)

The amount of memory in kilobytes used by this process which has been swapped out to disk.

SwapPss (All)

The amount of memory in kilobytes swapped out to disk. `SwapPss` includes memory counted in `Swap` as well as a portion of the memory swapped out which is shared by this process and with one or more other processes.

Add statistic for number of memory mapped regions in use

MemMapRegions (CachePgsvr Gem GemPgsvrMain Linux_Process Netldi Shrpc Statmon Stn)

The number of memory mapped regions the process is using.

Added statistics for smaps thread

The following statistics allow monitoring performance of the new statmonitor smaps thread itself.

SmapsCollectCount (Linux System)

Number of times the statmonitor smaps thread has collected smaps statistics for all processes it is monitoring.

SmapsLastActualSleepTime (Linux System)

The number of real milliseconds the statmonitor smaps thread actually slept in its most recent sleep.

SmapsLastReqSleepTime (Linux System)

The number of real milliseconds the statmonitor smaps thread requested to sleep in its most recent sleep.

SmapsLastTimeCollectingStats (Linux System)

The number of real microseconds the statmonitor smaps thread spent collecting smaps statistics for all monitored processes for its most recent sample.

SmapsPidsMonitoredCount - The number of processes the statmonitor smaps thread is currently monitoring.

SmapsProcFileFailedReads (Linux System)

Number of times the statmonitor smaps thread received an error when attempting to read the smaps file for a process. Usually indicates the process has exited.

SmapsSharedTableBuckets (Linux System)

Number of collision buckets present in the smaps shared process id table.

SmapsSharedTableSize (Linux System)

Number of keys present in the smaps shared process id table.

New restrictions for per-process Statistics for other users

Access to `/proc/pid/smaps` is restricted. Statmonitor and gem primitives which collect per-process statistics cannot collect smaps statistics for processes owned by other users.

Attempting to collect the smaps statistics does not raise an error, but the values for those specific statistics will always be zero (all other statistics will show up normally for other users)

To enable collecting per-process memory statistics for other users, you must give the executable the `cap_sys_ptrace` capability, by executing the following:

statmonitor executable:

```
sudo setcap cap_sys_ptrace=pe $GEMSTONE/bin/statmonitor
```

executables for Gem processes:

```
sudo setcap cap_sys_ptrace=pe $GEMSTONE/bin/topaz
sudo setcap cap_sys_ptrace=pe $GEMSTONE/sys/gem
```

To verify the executable has the correct capability, execute the following.

```
unix> setcap -v cap_sys_ptrace=pe $GEMSTONE/bin/statmonitor
./statmonitor: OK
```

Other changes in Existing Statistics

Details for **GarbageCollectionState** and **GcHighWaterPage**

These statistics report numeric values that represent internal state. These have been verified and the statistics definition updated for 3.7.

GarbageCollectionState (Gem)

Indicates the phase of a garbage collection task or other multi-threaded operation.

Values are defined as follows:

- 0 - Inactive
- 1 - Warm Dependency Map
- 2 - Scan Object Table
- 3 - Scan Data Pages
- 4 - Scan Shadowed (scavengable) Data Pages
- 5 - Scan Remaining Data Pages
- 6 - Backup Scavengable Pages
- 7 - Backup Remaining Pages
- 8 - Restore Data Pages
- 9 - Rescan Shadowed Pages (for listRefToInstOfClasses)
- 10 - Rescan Remaining Pages
- 11 - Audit Scavengable Pages
- 12 - Audit Remaining Pages
- 13 - Audit Rescan Object Table (rescan to find references to non existentobjects)
- 14 - Audit Rescan Data
- 15 - All Symbols Sweep
- 16 - Mark Sweep
- 17 - Write Set Union Sweep
- 18 - Find Connected Objects Sweep
- 19 - Migrate Merge Deltas

GcLockKind (Gem)

Indicates the state of the garbage collection lock and why it is being held:

- 0 - Free
- 1 - MarkForCollection
- 2 - FindDisconnectedObjects
- 3 - Epoch GC
- 4 - Write-set union sweep
- 5 - Reclaim All
- 6 - Backup
- 7 - Repository Scan (listInstances, listReferences, etc)
- 8 - Conversion

ReposAtMaxSize changed

Previously, **ReposAtMaxSize** was a boolean. Now, it is an integer indicating the current state of the extents:

- 0 - extents are not at maximum configured size and not in a disk full state.
- 1 - all extents have file system full state.

2 - all extents are currently at their configured size limit.

Changes in Session*WithGcLock and GcLockKind

As described on page 36, the GcLock can be held by up to five sessions, which makes existing protocol misleading.

SessionWithGcLock has been removed, replaced by the following statistic:

SessionsWithGcLock

Number of sessions holding a Gc Lock or repository Scan lock

The Stone's statistics no longer include **GcLockKind**. **GcLockKind** is now a per-session stat, describing the kind of lock that this session is holding.

Removed Statistics

The following statistics have been removed (in addition to Linux per-process stats listed above, and **SessionWithGcLock**):

- AbortInProgress**
- ClientAbortInProgress**
- CodeGenFullCount**
- ObjectMemoryGrowCount**
- PreemptedBitmapPages**
- PreemptedCommitRecordPages**
- PreemptedDataPages**
- PreemptedObjectTablePages**
- PreemptedOtherPages**
- SpinLockNewSymSleepCount**
- TimeSleepingMs**
- TrackedSetSize**
- TteCrPageFreeCount**

Added Statistics

"Time Since" statistics

The following statistics provide the time since a particular event occurred; this avoids the need to manually determine the delta between a given timestamp and the event timestamp. 0 means the event just occurred, -1 means the event has not occurred in the lifetime of the process.

SecondsSinceAbort (Gem Stn)

Stone: number of seconds since any session has aborted. Gem: number of seconds since this session has aborted.

SecondsSinceCheckpoint (Stn)

Number of seconds since a checkpoint was started.

SecondsSinceCommit (Gem Stn)

Stone: number of seconds since any session has committed. Gem: number of seconds since this session has committed.

SecondsSinceCrDisposal (Stn)

Number of seconds since a commit record was disposed.

SecondsSinceExtentGrow (Stn)

Number of seconds since any extent was grown.

SecondsSinceFailedCommit (Gem Stn)

Stone: number of seconds since any session attempted a commit which ultimately failed.

Gem: number of seconds since this session attempted a commit which ultimately failed.

SecondsSinceLogin (Gem Stn)

Stone: number of seconds since any session logged in.

Gem: number of seconds since this session logged in.

SecondsSinceLogout (Stn)

Number of seconds since any session logged out.

SecondsSinceLowFreespace (Stn)

Number of seconds since repository low free space was detected.

SecondsSinceNewTranlog (Stn)

Number of seconds since a new tranlog was started.

SecondsSinceReclaim (Stn)

Number of seconds since a reclaim gem committed.

SecondsSinceReclaimDead (Stn)

Number of seconds since a reclaim gem committed and reclaimed object identifiers.

SecondsSinceSigAbort (Gem)

Number of seconds since the session detected a SigAbort.

SecondsSinceSigLostOt (Gem)

Number of seconds since the session detected a SigLostOtRoot.

SecondsSinceStartLongPrim (Gem)

Number of seconds since the session started a long primitive operation.

SecondsSinceStartStone (Stn)

Number of seconds since the stone has been available for logins.

SecondsSinceTranlogFull (Stn)

Number of seconds since the a tranlog full condition was detected.

Other Added Cache Statistics

BitmapPagesPreempted (CachePgsvr FreeFrameThr PageMgrThr)

Number of bitmap pages removed from the cache by the process.

BitmapPageWrites (StnAioThr)

Number of bitmap written to disk by the process.

CacheIdle (Shrpc)

A boolean indicating if the shared page cache is idle.

CachePgsvrRemoveFrameId (CachePgsvr)

The frameId that the cache page server is attempting to recycle.

CachePgsvrRemovePageId (CachePgsvr)

The pageId that the cache page server is attempting to remove from the remote shared page cache.

CacheRegionNumFrames (FreeFrameThr StnAioThr)

The number of frames in the cache region for the process.

CommitRecordPagesPreempted (CachePgsvr FreeFrameThr PageMgrThr)

Number of commit record pages removed from the cache by the process.

CommitRecordPageWrites (StnAioThr)

Number of commit record pages written to disk by the process.

DataPagesPreempted (CachePgsvr FreeFrameThr PageMgrThr)

Number of data pages removed from the cache by the process.

DataPageWrites (StnAioThr)

Number of data pages written to disk by the process.

FfiGcHeapBytes (Gem)

Bytes allocated by `CByteArray class >> gcMalloc:` and other C data allocations for instances of `GsSocket`, etc, and not yet freed by in-memory GC.

FfiHeapBytes (Gem)

Cumulative bytes allocated by `CByteArray class >> malloc:`, that the VM GC will not free, not decremented when application FFI code calls free.

FreeFrameListOkLimit (FreeFrameThr StnAioThr)

The ok free frame limit for the process.

FreeFrameLowerLimit (FreeFrameThr StnAioThr)

The lower free frame limit for the process. When the number of free frames is below this value, the process will aggressively add frames to the free list (free frame threads) or write dirty pages to disk (Aio threads).

FreeFrameUpperLimit (FreeFrameThr StnAioThr)

The upper free frame limit for the process.

GemThreadsInCacheCount (Shrpc)

The number secondary gem threads currently attached to the cache.

LowFreespaceCount (Stn)

Number of times stone detected a low free space condition.

NumInSecurityDataQueue (Stn)

Number of sessions waiting for `SymbolGem` to update security data of their `UserProfile`.

ObjectTablePagesPreempted (CachePgsvr FreeFrameThr PageMgrThr)

Number of object table pages removed from the cache by the process.

ObjectTablePageWrites (StnAioThr)

Number of object table pages written to disk by the process.

OnetimePasswordsActive (Stn)

Number of onetime passwords in the stone's hash map.

OnetimePasswordsCreated (Gem)

Number of onetime passwords this session has created.

OnetimePasswordsExpired (Stn)

Number of expired onetime passwords automatically removed by stone.

OnetimePasswordsNotValidated (Stn)

Number of failed onetime password validations.

OnetimePasswordsTotal (Stn)

Number of onetime passwords ever created by any session since the stone was started.

OnetimePasswordsValidated (Stn)

Number of successful onetime password validations.

OnetimePasswordUsed (Gem)

A boolean with value 1 if the session used a onetime password to login, otherwise 0.

OtherPagesPreempted (CachePgsvr FreeFrameThr PageMgrThr)

Number of pages removed from the cache by the process that were not object table, data, commit record or bitmap pages.

OtherPageWrites (StnAioThr)

The number of pages written by the process that were not object table, data, commit record or bitmap pages since the process was started.

RecovNumPendingCheckpoints (Stn)

Number of checkpoints pending replay from the tranlog during recovery.

RecovSkippedCheckpointCount (Stn)

Number of checkpoints skipped during recovery.

RecovTimeWaitingForCheckpoints (Stn)

Total amount of real milliseconds the stone recovery thread spent waiting for checkpoints.

ScanLimitNumFrames (FreeFrameThr StnAioThr)

Maximum number of cache frames scanned by the process before sleeping.

TimeInCheckpoint (Stn)

Cumulative number of real seconds the stone has spent writing checkpoints.

TimeInLowFreespace (Stn)

Cumulative number of real seconds the stone has spent in low free space mode.

TranlogFullCount (Stn)

Number of times stone detected a tranlog full condition.

TteCrPageInUse (Stn)

Number of commit record page entries in use in stone's internal commit record cache.

Changes in GemStone Smalltalk

This chapter describes changes and new features important for programmers using GemStone Smalltalk, including:

Changes in Classes, compilation, and code management ..	53
FileSystem	55
Support for ssh and sftp using OpenSSL	58
Changes in Magnitude classes	60
Stream and String related changes	62
JsonParser and JsonPetitParser	64
External Session Changes	65
Other Changes and Enhancements	65
Deprecated and removed classes and methods	66

3.1 Changes in Classes, compilation, and code management

SystemUser SymbolList Changes

In v3.6.x, SystemUser's SymbolList (visible when logged in as SystemUser), included a number of additional SymbolDictionaries, including GsCompilerClasses, ObsoleteClasses, RowanKernel, and in some versions GemStone_Portable_Streams and GemStone_Legacy_Streams. This facilitated work, or were artifacts of, the transition to Rowan-based code management.

For 3.7, all these dictionaries have been removed; the classes are located in subdictionaries within Globals, as needed.

ClassOrganizer >> newExcludingGlobals

This method allows you to perform ClassOrganizer queries that do not return results exclusive to Globals. This can be used, for example, to find application methods that reference GsHostProcess or JsonParser class (which may need to be recompiled, as described in the *Installation Guide* upgrade instructions), without including kernel methods that do not need recompile.

ClassOrganizer >> newExcludingGlobals replaces the private method `_newExcludingGlobals`, which was available in some recent releases. `_newExcludingGlobals` did not reliably exclude methods on classes in Globals.

GsNMethod >> recompileFromSource

This method has been added for convenient recompilation of a method from source code, such as the recompile needed for GsHostProcess and JsonParser. Existing recompile methods were designed for recompiling after bytecode changes; since they relied on the existing literal references, they did not update references to the obsolete class.

Note that if the symbolList of the user that is executing the recompile resolves literal names to different objects than the original compile, this may result in unexpected changes in behavior. Depending on the complexity of your application's use of users' symbolLists, you may wish to examine the source code before recompiling.

Added method Behavior >> compileMethod:category:environmentId:

The following method has been added, allowing you to compile a method using the current symbol list:

```
Behavior >> compileMethod: sourceString category: aCategoryString
environmentId: environmentId
```

#logCreation and logging class creation

The option `#logCreation` can be included in the options: of a class creation message, which results in the class creation being logged (using `GsFile gciLogServer:`). The primary purpose was for tracking class creation during build/upgrade filein. This option is not persistent and not printed (reported by `Class >> definition`), and thus was not retained if a class was filed out and in.

Now, you may set the key `#GsClass_logCreation` in `SessionTemps`, to force logging of all class creations for the lifetime of the session, using an expression such as:

```
SessionTemps current at: #GsClass_logCreation put: true
```

The legacy behavior is unchanged.

Improvements to filein using GsFileIn

GemStone supports filein of topaz-format source code using the class `GsFileIn`. This provides a functionality similar to `topaz input`, but permits only a subset of the full set of topaz commands. `GsFileIn` allows you to filein GemStone code programmatically, entirely from within GemStone code.

`GsFileIn` has been extensively revised. With v3.7, `GsFileIn` now supports source files that include extended characters encoded as Utf8; and performance has been greatly improved.

While instances of `GsFileIn` are not normally persistent, on upgrade, existing instances become instances of `ObsoleteGsFileIn`.

The `GsFileIn` class >> `from*`: API has been streamlined. While `fromServerPath:` and `fromClientPath:` may be used, the preferred methods (methods whose names are more accurate) are `fromGemHostPath:` and `fromGciHostPath:`, respectively. The method `GsFileIn` >> `fromStream:` has been added, to allow filein from streams.

For more details on `GsFileIn`, see class comment and methods in the image.

3.2 FileSystem

`FileSystem` is a set of classes, ported and adapted from Pharo, that support operations on files and directories. `FileSystem` provides a much more flexible and feature-rich environment than `GsFile`. `FileSystem` performance is comparable to `GsFile`.

In these notes, the term 'FileSystem' can be used to refer to the entire set of classes that support file and directory operations, or to the specific class named `FileSystem`. The specific class `FileSystem` represents the underlying file environment. Most operations within the general File System environment use the class `FileReference`, which represents a file or directory.

`FileSystem` is still under active development and may be missing features or contain unexpected behavior.

Differences from GsFile

`FileSystem` supports a superset of the functions provided by `GsFile` for server files, and the API is significantly different.

Some differences:

- ▶ `FileSystem` is implemented using multiple classes to provide a flexible and portable API, while `GsFile` is implemented as a single class with fixed behaviors.
- ▶ `FileSystem` is implemented based on an FFI interface to the OS file system functions, while `GsFile` is implemented with user actions (for historic reasons) and C primitives.
- ▶ On error, `FileSystem` signals an appropriate error; `GsFile` functions return nil and require you to query for the error details.
- ▶ `FileSystem` does not support operations on the GCI client (for example, reading files on a Windows client), while `GsFile` does.

`FileSystem` is supported with Linux on x86 and ARM, and macOS on x86 and Apple Silicon. `FileSystem` is not supported on AIX; on AIX, you must continue to use `GsFile`.

Primary API classes

There are a few main entry points for `FileSystem`.

■ FileReference

`FileReference` is the primary entry point for file and directory operations. You may create files using `FileReference` >> `/`, `aString asFileReference`, or by using

FileSystem class protocol. FileReferences can also be created from another instance of FileReference. The following are all equivalent:

```

'/gshost/test/foo.txt' asFileReference
FileReference / '/gshost/test/foo.txt'
FileReference / 'gshost' / 'test' / 'foo.txt'
FileSystem disk / '/gshost/test/foo.txt'
'/gshost/test/' asFileReference / 'foo.txt'

```

The FileReference that is created prints as `File @ /gshost/test/foo.txt`

There are a great many supported operations, including both information about the file or directory, and operations on that file or directory. Note that much of the API is on the superclass, AbstractFileReference. The methods include:

- ▶ report if the receiver is a file or directory, and for directories, report details on the contained children (files or directories).
- ▶ return details of the path, filename, and extension
- ▶ return information about a file, such as size, permissions, modificationTime, etc.
- ▶ open a read or write stream on the contents, including encoded and binary streams.

Operations on the contents of a FileReference will normally use streams. Methods such as `readStream` and `writeStream` return instances of Zinc-based streams, in this case `ZnCharacterReadStream` or `ZnCharacterWriteStream`. For example:

```

rdStream := 'gshost/test/foo.txt' asFileReference readStream.
[rdStream atEnd] whileFalse:
    [report add: rdStream nextLine; lf.].
rdStream close.

```

Variants such as `readStreamDo:` and `writeStreamDo:` automatically close the file after performing the block operation.

By default, these streams provide automatic UTF-8 encoding, although method variants allow you to specify an encoder to support legacy 8-bit encoding.

Binary read and write streams created using `binaryReadStream` and `binaryWriteStream` create instance of `FsBinaryFileStream`, allowing byte-based operations with no encoding.

■ FileLocator

FileLocator provides similar file operations to FileReference, but the file or directory does not need to have an explicit full path. The location for a FileLocator instance can be relative to your environment; for example, the working directory, the directory containing the extent files or the directory containing gem logs. This allows you to move code between environments without requiring explicit management of the paths. For example the following code does not need to be changed if executed in different environment:

```

outFile := FileLocator home / 'output.txt'.
outFile
    writeStreamDo: [:stream | stream nextPutAll: 'test results']
    ifPresent: [FileAlreadyExistsException signalWith: outFile]

```


You can see the available environment-dependent origins using `FileLocator class >> supportedOrigins`; class methods are available for all, including `home`, `cache`, `temp`, `userData`, `tranlog`, `preferences`, `extent1Directory`, `extent1`, `documents`, `desktop`, and `workingDirectory`.

To find the resolved value of the `FileLocator`, you can send the message `absolutePath`; this returns an instance of a kind of `Path`.

```
FileLocator extent1 absolutePath printString
  Path / 'gshost' / 'GemStone3.7' / 'data' / 'extent0.dbf'
```

`FileLocator` understands most of the methods available for `FileReference`.

■ **FileSystem**

While the primary entry point is `FileReference`, the class `FileSystem` can provide a way to define a `FileReference`. `FileSystem` supports ordinary disk based file systems and in-memory file systems.

`FileSystem` has a number of useful methods to defining `FileReferences`. For example:

```
FileSystem disk
  returns a disk-based FileSystem, which can be used with / to create a FileReference
  to a disk file, e.g. FileSystem disk / '/gshost/test/foo.txt'.
```

```
FileSystem memory
  returns an in memory-based FileSystem, which can be used with / to create a
  FileReference, e.g. FileSystem memory / 'foo' / 'bar'.
```

```
FileSystem workingDirectory
  creates a FileReference to the current working directory (disk-based).
```

`FileSystem` allows you to perform a number of file and directory operations, although the preferred way is to use `FileReference`.

■ **FsFileDescriptor**

`FsFileDescriptor` represents the file itself. `FileReference open :` methods return an instance of this class, which can be used to perform basic file operations.

`FsFileDescriptors` read and write `ByteArrays` rather than strings.

Supporting Classes

Zinc Stream Classes

Subclasses of `ZnStream`, including `ZnBuffered*Stream`, `ZnEncoded*Stream`, and `Zn*Encoder`, provide read and write streams for `FileSystem`. `Zn*Encoder` streams handle the encoding/decoding from UTF-8 and 8-bit (GemStone legacy) encoded files, and the creation of `Legacy` or `Unicode String` instances.

Opening options

`FsFileOpeningOptions` provide the operating-system specific options for opening files.

Error Classes

`FsError` and its subclasses represent `FileSystem` errors, `FileException` and its subclasses represent `FileReference/FileLocator` errors. Other errors such as `FsUnixError` and its

subclasses represent low level UNIX file errors, which are generally handled by public API.

Path classes

Path, RelativePath, and AbsolutePath encapsulate a path. A Path can be obtained from a FileReference and vice versa. While Path is an abstract class, you can use it to create instances, e.g. Path from: '/gshost/test/' will return an instance of AbsolutePath.

Store classes

FileSystemStore and subclasses represent the OS file system or memory file system, with differences for specific operating systems.

Resolver classes

FileSystemResolver and subclasses support resolving various FileLocator origin options, which vary between operating systems.

FFI support classes

FsLibcInterface and FsCStruct and their subclasses provide the FFI interface to the operating system file commands.

3.3 Support for ssh and sftp using OpenSSL

The OpenSSL libraries includes support for ssh and sftp. These commonly used functions are now available programmatically using GemStone classes.

The following classes have been added:

- GsSshSocket
- GsSftpSocket
- GsSftpRemoteFile
- SshSocketError (error 2759)

Examples of using these classes are provided in class side methods.

SSH with GsSshSocket

Creating the SSH connection

To create a SSH connection, create a client socket using methods inherited from GsSocket, and connect using sshConnect.

For example, the following code connects to the host *hostnameOrIP*:

```
sshSock := GsSshSocket newClient.
sshSock connectToHost: hostnameOrIP timeoutMs: 2000.
sshSock userId: userName; password: userPassword.
sshSock disableHostAuthentication.
sshSock sshConnect.
```

As with all sockets, you should close the socket when you no longer need the connection.

```
sshSock close
```

SSH operations

Commands are executing using `executeRemoteCommand`. For example,

```
result := sshSock executeRemoteCommand: 'uptime'.
```

The commands can be executed non-blocking using the following methods:

```
nbExecuteRemoteCommand:  
hasCommandInProgress  
nbRemoteCommandResult  
nbRemoteCommandResultReady
```

SFTP with GsSftpSocket

Creating Sftp connection

Creating an SFPT connection is just like creating an SSH connection. To create a SSH connection, create a client socket using methods inherited from `GsSocket`, and connect using `sshConnect`.

For example, the following code connects to the host *hostnameOrIP*:

```
sshFtpSock := GsSshSocket newClient.  
sshFtpSock connectToHost: hostnameOrIP timeoutMs: 2000.  
sshFtpSock userId: userName; password: userPassword.  
sshFtpSock disableHostAuthentication.  
sshFtpSock sshConnect.
```

As with all sockets, you should close the socket when you no longer need the connection.

```
sshFtpSock close
```

Sftp operations

Once the `GsSftpSocket` instance is created and connected, there are a number of methods supported, including:

```
GsSftpSocket >> remoteFileExists: aFileOrDirectory  
GsSftpSocket >> currentRemoteDirectory  
GsSftpSocket >> createRemoteDirectory: dirname  
GsSftpSocket >> createRemoteDirectory: dirname mode: modeInt  
GsSftpSocket >> contentsOfRemoteDirectory: dirname  
GsSftpSocket >> contentsOfRemoteDirectory: dirname matchPattern:  
    aString  
GsSftpSocket >> renameRemoteFile: oldName to: newName  
GsSftpSocket >> removeRemoteFile: filename  
GsSftpSocket >> removeRemoteDirectory: dirname
```

GsSftpRemoteFile for operations on remote files

The class GsSftpRemoteFile represents an instance of a Remote File that can be used for reading and writing. You must have an established GsSftpSocket connection, which is an argument to the instance creation methods.

```
GsSftpRemoteFile class >> openRemoteFile: fileName mode: openMode
  permissions: permInt errorIfExists: aBoolean withSftpSocket:
  aGsSftpSocket

GsSftpRemoteFile class >> openRemoteFileAppend: fileName
  withSftpSocket: aGsSftpSocket

GsSftpRemoteFile class >> openRemoteFileReadOnly: fileName
  withSftpSocket: aGsSftpSocket

GsSftpRemoteFile class >> createNewRemoteFile: fileName
  withSftpSocket: aGsSftpSocket

GsSftpRemoteFile class >> createOrOverwriteRemoteFile: fileName
  withSftpSocket: aGsSftpSocket
```

Once an instance of GsSftpRemoteFile has been created, use the instance protocol to read, write, and fetch information about the remote file. See the image for specific methods.

For example, to download a remote file using SFTP:

```
sshFtpSock := GsSftpSocket newClient.
sshFtpSock connectToHost: hostnameOrIP timeoutMs: 2000.
sshFtpSock userId: userName; password: userPassword.
sshFtpSock disableHostAuthentication.
sshFtpSock sshConnect.
remoteSftpFile := GsSftpRemoteFile
  openRemoteFileReadOnly: remoteFileName
  withSftpSocket: sftpSock.
localFile := GsFile openWriteOnServer: localFileName.
bytes := remoteSftpFile readAllInto: localFile.
localFile close.
sftpFile close.
```

3.4 Changes in Magnitude classes

Float printing improvements

GemStone now includes the library double-conversion, which allows more control over floating point display. Float and SmallDoubles such as 9.45, that previously displayed as the technically correct, but not optimal, '9.449999999999999', now display as '9.45'. (#47003)

In addition, Doubles may have up to 17 significant decimal digits, but the earlier code limited this to 16 decimal places. For example, a float created from '7.6000000000000001' was previously printed as 7.6000000000000001, and is now 7.6000000000000005. (#46855)

Time support for microseconds

In previous releases, Time has included microsecond resolution; however, methods to support using microsecond times were limited.

The following methods have been added:

`asStringUs`

Returns a String that expresses the receiver in local time in the format HH:MM:SS.ssssss where sss are microseconds.

`addMicroseconds: anInteger`

Returns a Time that describes a time of day anInteger microseconds later than that of the receiver.

`subtractMicroseconds: anInteger`

Returns a Time that describes a time of day anInteger milliseconds earlier than that of the receiver.

Converting Non-specials in special range

A special object is one in which the OOP encodes the value, and thus requires no repository space; this is a canonical form of that value. A number of magnitude classes have subclasses with the prefix 'Small', which provide a special/canonical form of a subset of that class's potential values. While new values are created in canonical form, it is possible to have non-canonical instances that have an exactly equivalent canonical value. For example, an upgraded application may have an instance of Fraction with the value 1/3, rather than a SmallFraction with the value 1/3.

To convert a non-canonical to the canonical equivalent, the method `asCanonicalForm` has been added. This either returns self, or the canonical object exactly equivalent to the receiver. Specific implementations apply to Float, Date, DateAndTime, Time, Fraction, LargeInteger, and ScaledDecimal.

Note that this does not perform a conversion in place; you must modify the object referring to the non-canonical value to refer to the canonical value.

highBit now ANSI-compliant

The method `Integer`, `LargeInteger`, and `SmallInteger` >> `highBit` previously returned nil for a zero receiver. Now, this method returns 0 for a zero receiver, as specified by the ANSI spec.

Other added method

The following method has been added:

`SmallInteger` >> `minimum32bitInteger`

Return the minimum value representable by a 32bit signed integer

3.5 Stream and String related changes

String compare to ByteArray now returns false for matching bytes

If a String is compared to a ByteArray using `=`, and they both contain the same bytes (e.g. `' ' = #[32]`), in previous releases, the comparison returned true; while a similar comparison between a ByteArray and a String (e.g. `#[32] = ' '`) returned false. (#49031) Both cases now return false for this comparison.

New ReadByteStream classes optimized for Strings

The class ReadByteStream, implemented by ReadByteStreamPortable and ReadByteStringLegacy classes, has been added. This provides a ReadStream that is optimized for performance, and limited to holding kinds of String, MultiByteString, or ByteArray.

The API of ReadByteStream is largely the same as the existing ReadStream classes.

readStream reimplemented to return ReadByteStream

The method readStream, which was previously inherited from SequencableCollection to return an instance of ReadStream, has been reimplemented in MultiByteString, String, and ByteArray (and inherited by respective subclasses), to return an instance of ReadByteStream.

Support for Lz4 encode/decode added to String and ByteArray classes

The following methods have been added, which allow Lz4 compression into a ByteArray, and decompression from Lz4-compressed data in a ByteArray into a byte-format object; normally but not necessarily a kind of String, MultiByteString, or ByteArray.

Compressing a multi-byte character object converts it to a little endian format before compressing, and decompression assumes the data is in little endian format.

You cannot compress, nor decompress into, a Symbol, DoubleByteSymbol or QuadByteSymbol.

Note that the compressed form does not distinguish the original class; for example, String or DoubleByteString; you must specify the correct class in the argument to the decompress method.

```
ByteArray, String, MultiByteString >> compressWithLz4: opCode
  Compress the receiver with lz4 compression, and return an instance of ByteArray
  containing the compressed bytes. Compression modes include 0/frame mode and
  1/blockmode; frame mode should normally be used.
```

```
ByteArray >> decompressWithLz4IntoNewInstanceOf: aClass
  decompressedSize: aSize
  Decompress the receiver using lz4 and store the resulting bytes into a new instance
  of aClass. aClass must be a byte format Class. The lz4 compression mode (block or
  frame) used to compress the data is automatically detected by this method. aSize is
  a SmallInteger which indicates the size (in characters, not bytes) of the data when
  decompressed. If the contents were compressed using lz4 frame mode, aSize may
```

be set to 0 indicating the decompressed size is not known. On success, returns a new instance of *aClass* containing the decompressed data.

ByteArray read/write of binary numbers in native format

The following methods have been added to allow you to read and write numeric values as raw bytes into a ByteArray in the current native byte order. Values are written in native format, and read as if in native format: no swizzling is done.

Note that there are existing methods in ByteArray to write and read numeric values; these methods read and write in big endian format, regardless of platforms.

```
ByteArray >> unsigned16AsNativeAt: startIndex
ByteArray >> unsigned32AsNativeAt: startIndex
ByteArray >> unsigned64AsNativeAt: startIndex
ByteArray >> signed16AsNativeAt: startIndex
ByteArray >> signed32AsNativeAt: startIndex
ByteArray >> signed64AsNativeAt: startIndex
ByteArray >> doubleAsNativeAt: startIndex
ByteArray >> floatAsNativeAt: startIndex
ByteArray >> signed16At: startIndex putAsNative: anInteger
ByteArray >> signed32At: startIndex putAsNative: anInteger
ByteArray >> signed64At: startIndex putAsNative: anInteger
ByteArray >> unsigned16At: startIndex putAsNative: anInteger
ByteArray >> unsigned32At: startIndex putAsNative: anInteger
ByteArray >> unsigned64At: startIndex putAsNative: anInteger
ByteArray >> doubleAt: startIndex putAsNative: aDouble
ByteArray >> floatAt: startIndex putAsNative: aFloat
```

The existing ByteArray methods `at:put:signed:width:` and `at:signed:width:` now accept additional `width:` arguments of 512, 1024, or 2048, which specify native format with 2, 4, or 8 bytes, respectively.

Utf8 now disallows #readStream

The `#readStream` method, previously inherited from `SequenceableCollection`, has been disallowed in `Utf8`. `Utf8` is intended to be only used in the encoded form; you should decode this prior to creating a `readStream` on it.

codePointAt:put: now performs auto-conversion if needed

Previously, sending `codePointAt:put:` to a `String` or `DoubleByteString`, with a `codePoint` is out of range for that class, resulted in an `OutOfRangeException` error.

Now, the receiver will be transparently converted to the class that can contain the given `codePoint`; that is, convert from a `String` or `DoubleByteString` to a `DoubleByteString` or `QuadByteString`).

Other Changes

atOrNil: added

This method has been added to `ByteArray`, `String` and `String`'s subclasses `DoubleByteString` and `QuadByteString`, as part of the optimizations for `ReadStream`.

`atOrNil: anIndex`

Returns the Character at `anIndex`, or nil if `anIndex` is beyond end of the receiver

This is disallowed for `Utf8` and `Utf16`.

Stream classes consistently return objects of their collection class

Some methods could return an empty instance of `String`, rather than an instance of the collection class.

3.6 JsonParser and JsonPetitParser

Previous releases includes a `PetitParser`-based JSON Parser, which could read a JSON string and create the equivalent Smalltalk object structure. Using this class required some experience with `PetitParser`.

In this version, a new JSON parser has been added. This is a faster and simpler recursive descent parser, not related to `PetitParser`.

The class name 'JsonParser' is now applied to the new, simple JSON Parser. The old JSON parser, formerly 'JsonParser', is now 'JsonPetitParser'. `JsonPetitParser` is deprecated.

The API for the new `JsonParser` is much smaller; like `JsonPetitParser`, the new `JsonParser` is also accessed using the class method `parse:`. The new `JsonParser` signals errors on `parse:` of invalid JSON, unlike `JsonPetitParser`.

Upgrade impacts

Compiled methods that contain references to the `JsonParser` class in an earlier version, will refer to `JsonPetitParser` after upgrade and continue to work as before. Provided you are only using the `JsonParser` to parse JSON, it is optional to get the improved performance by recompiling methods that reference the `JsonParser` class.

If you are using `JsonParser` as part of a `PetitParser` customization, references to the class by name must be updated to refer to 'JsonPetitParser' instead of 'JsonParser'. It is recommended to edit the source of methods that reference `JsonParser` class to reference `JsonPetitParser`, and edit the superclass of subclasses of `JsonParser`.

If you have subclasses of `JsonParser` that must be file in, these will be compiled and refer to the new `JsonParser` class. If this subclass directly references inherited instance variables, compilation of those methods will fail; and depending on the specific modifications may or may not be functional. In most cases before file in, you should modify the code on disk to refer to `JsonPetitParser`.

See the *Installation Guide for v3.7*, Upgrade chapter, for recompilation details

3.7 External Session Changes

Added methods for GsTsExternalSession

The following methods have been added:

`GsTsExternalSession >> newUtf8String: aUtf8 toUnicode: aBoolean`
Create a new instance of `Utf8` (if `aBoolean` is false) or a `Unicode7`, `Unicode16` or `Unicode32` (if `aBoolean` is true), in the external session, containing the contents of `aUtf8`. Returns the oop of the object that was created in the external session.

`GsTsExternalSession >> releaseOop: anOop`
Release an `Oop` (an oop in the external session) from the `ExportSet` of the current session, allowing the object to be garbage collected in the external environment.

`GsTsExternalSession >> releaseOops: anArray`
Release the oops in an `Array` (oops in the external session) from the `ExportSet` of the current session, allowing the objects to be garbage collected in the external environment.

`GsTsExternalSession >> send: aSelector to: rcvrOop withOops: anArray`
The same as `send:to:withArguments:`, but the arguments are already in the form of oops.

Return value from GsTsExternalSession >> waitForReadReadyTimeout:

Previous, the method `GsTsExternalSession >> waitForReadReadyTimeout: msToWait` returned `self` for success, and signaled an error if the timeout expired or an error occurred.

Now, this method returns `true` on success, `false` if there was a timeout, and signals an error if there was an error.

Easier to use GEM_HALT_ON_ERROR in external sessions

The default for `haltOnErrNum` has been changed to `-1`, to allow the Gem's configuration file setting for `GEM_HALT_ON_ERROR` to be in effect for external sessions.

3.8 Other Changes and Enhancements

ProfMonitor improved formatting for object creation tree report

The tree report produced when doing object creation tracing in `ProfMonitor` or `ProfMonitorTree` has improved formatting.

GsSecureSocket example updates

The `GsSecureSocket` class method

```
httpsClientExampleForHost:certificateDirectory:withSniName:
```

has been renamed to

```
httpsClientExampleForHost:withSniName:.
```

The following methods has been added:

GsSecureSocket class >> setCaCertLocation

On linux: try to find the trusted CA cert file on this host. If we find it, use it. If we do not, disable cert verification so the examples work correctly.

isValidIdentifier, validateIsIdentifier optimized

The following methods are now implemented in primitives:

CharacterCollection >> isValidIdentifier

Object >> validateIsIdentifier

Other Added methods

GsFile >> beforeEnd

Answer true if the receiver is not positioned at the end, false otherwise

GsFile >> print: *anObject*

Writes the print representation of the argument to the GsFile receiver instance.

GsHostProcess >> readOutErr

Assuming both out and err are sockets, attempt to read from both and return combined result.

GsNetworkResourceString >> gemConfig: *aString*

Include a -C configuration definition for the receiver. *aString* should be a valid -C argument; for example, 'GEM_TEMPOBJ_CACHE_SIZE=100MB;'.

Array >> copyNotNilFrom: *startIndex* to: *stopIndex*

Return an Array containing the non-nil elements that are within the specified offsets of the receiver.

Array >> indexOfNotNil: *startOffset* to: *endOffset*

If *startOffset* <= *endOffset*, returns the first offset within in the specified range of a non-nil element of the receiver. If *startOffset* > *endOffset*, returns the last offset within the specified range of a non-nil element of the receiver. Returns zero if all are nil. Intended for use only on Arrays of size <= 2000. Performance on larger arrays will be slow.

TestAsserter >> assert: *anObject* identical: *otherObj*

Assert that two objects are identical.

TestAsserter >> fail

Cause the test to report failure unconditionally

TestAsserter >> fail: *descriptionString*

Cause the test to report failure unconditionally, with the given description.

3.9 Deprecated and removed classes and methods

See “GciDirtyTrackedObjs functionality removed” on page 69 for additional removed methods related to the remove Tracked Dirty Objects set.

Deprecated Methods

The following methods are newly deprecated. See the deprecation message for the replacement.

```
GsProcess >> stepIntoFromLevel:  
GsProcess >> stepOverFromLevel:  
GsProcess >> stepThroughFromLevel:  
GsProcess >> terminateTimeoutMs:  
GsProcess >> threadRubyEnvId  
System >> sessionIdHoldingGcLock
```

AbstractCharacter removed

The abstract class AbstractCharacter has been removed; Character superclass is now Magnitude. The AbstractCharacter definition is now in ObsoleteClasses.

Removed Public Methods

The following methods have been removed, associated with changed mentioned elsewhere in these release notes:

```
CPreprocessor >> defaultSparcSolarisDefinitions  
CPreprocessor >> defaultX86SolarisDefinitions  
Breakpoint class >> trappable:  
GsSecureSocket class >> httpsClientExampleForHost:certificateDi  
rectory:withSniName:  
GsTestCase >> assert:isEquivalentTo:  
Object >> removeObjectFromBtrees
```

Removed Previously-deprecated Methods

The following methods were previously deprecated, and have been removed:

```
ExecBlock >> valueNowOrOnUnwindDo:  
Repository >> shrinkExtents  
System class >> cacheStatisticsDescription
```

Removed Private Methods

```

AbstractException >> _signalAsync
AbstractException >> _signalFromPrimitive:
AbstractException >> _signalGcFinalize
AbstractException >> _signalTerminateProcess
AbstractException >> _signalTimeout
AbstractException >> _signalWith:
Behavior >> _transientSessionMethodBehaviorsCacheName
CCallout class >> _errno:
CharacterCollection >> _validIdentifier:
ClassOrganizer class >> _newExcludingGlobals
Delay >> _activate
ExecBlock >> _valueNowOrOnUnwindDo:
ExecBlock >> _valueNowOrOnUnwindDoPrim:
GsFileIn >> compileMethodIn:
GsFileIn >> fileStream:
GsFileIn class >> fromNestedClientPath:to:
GsFileIn class >> fromNestedPath:on:to:
GsFileIn class >> fromNestedServerPath:to:
GsFileIn class >> _fromStream:
GsFileIn class >> _fromStream:to:
GsNMethod >> _breakOperation:forStepPoint:
GsNMethod >> _setBreakAtIp:operation:frame:process:
GsNMethod >> __setBreakAtIp:operation:frame:process:
GsNMethod class >> _setBreakAtIp:operation:frame:process:
GsProcess >> _activate
GsProcess >> _continue
GsProcess >> _finishTermination
GsProcess >> _isRubyThread
GsProcess >> _serviceTerminationInterrupt
GsProcess >> _stepOverFromLevel:
GsProcess >> _stepOverInFrame:mode:replace:tos:
GsSocket >> _waitingProcessesInto:
GsSocket >> _waitingProcessesInto:inGroup:
GsTsExternalSession >> _getObjInfo:buffer:
IdentityBag >> _basicAdd:
IdentityBag >> _rcIncludes:
IdentityBag >> _rcIncludesValue:
IndexingErrorPreventingCommit >> _signalWith:
Integer >> _floatParts
Object >> addObjectToBtreesWithValues:
Object >> _respondsTo:private:flags:
Object >> _errorsDuringPreventCommit:
RcIdentityBag >> _thisSessionRemovalIndex
RcKeyValueDictionary >> _keyCollisionOk
SmallInteger >> _floatParts
System class >> _deprecatedCacheStatisticsDescription
System class >> _disallowSubsequentCommits:
System class >> _primitiveAbort
System class >> _sessionsReportExcluding:
System class >> _hostWaitForDebuggerIfSlow

```

Changes in the GCI Interfaces

4.1 GCI changes

GciDirtyTrackedObjs functionality removed

The Dirty Tracked Objects set is no longer needed, and has been removed.

Removed GCI calls

```
GciDirtyTrackedObjs  
GciReleaseAllTrackedOops  
GciReleaseTrackedOops  
GciSaveAndTrackObjs  
GciTrackedObjsFetchAllDirty  
GciTrackedObjsInit  
GCI_JIS_CHAR_TO_OOP  
GCI_OOP_TO_JIS_CHAR
```

Removed Image Methods related to TrackedObjects

The GciLibrary functions associated with the removed GCI calls are no longer included.

In addition, this method has been removed:

```
System class >> gciTrackedObjsInit
```

GsBitmap hiddenSetSpecifiers

```
#TrackedDirtyObjs is now #TrackedDirtyObjs_NotImplemented  
#GciTrackedObjs is now #GciTrackedObjs_NotImplemented
```

Added Lz4 compression functions

GciCompressUsingLz4

```
(int) GciCompressUsingLz4(
    const char* src,
    char* dst,
    int srcSize,
    int dstCapacity
);
```

Wrapper for the Lz4 function `LZ4_compress_default()`.

Compress *srcSize* bytes from buffer *src* into already allocated *dst* buffer of size *dstCapacity*. Compression is faster, and guaranteed to succeed, if *dstCapacity* \geq `LZ4_compressBound(srcSize)`.

If the function cannot compress *src* into a more limited *dst* buffer size, compression stops, function result is zero, and *dst* content is undefined (invalid).

GciUncompressUsingLz4

```
(int) GciUncompressUsingLz4(
    const char* src,
    char* dst,
    int compressedSize,
    int dstCapacity
);
```

Wrapper for `LZ4_decompress_safe()`.

Decompress the compressed string at **src* and place the result in **dst*. **dst* must be already allocated. The exact size of the compressed string at **src*, and the maximum size for the destination must be provided; if *dstCapacity* is insufficient, decompression stops and the function returns a negative result.

Returns the number of bytes decompressed into **dst*.

Removed Mnemonics

The definitions for the obsolete classes `OOP_CLASS_EUC_STRING`, `OOP_CLASS_INVARIANT_EUC_STRING`, `OOP_CLASS_EUC_SYMBOL`, `OOP_CLASS_JIS_CHARACTER`, and `OOP_CLASS_JIS_STRING` have been renamed to include an additional `_`.

4.2 GCI thread-safe changes

The thread-safe GCI is provided in `gcits.hf`.

Added Thread-safe function

GciTsNbPoll

```
(int) GciTsNbPoll(  
    GciSession session,  
    int timeoutMs,  
    GciErrSType *gciErr  
);
```

Function results:

1 - result or error is ready, call GciTsNbResult to get the result or error.

0 - result is not ready

-1 - error, (invalid session, no NB call in progress, peer disconnected), details in *gciErr*.

Argument *timeoutMs* may be:

0 - do not block, return immediately

-1 - block forever until the Nb call finishes or an error occurs.

> 0 block for up to *timeoutMs* ms before returning the result

4.3 Foreign Function Interface (FFI) related changes

errno: unsafe; use alternate methods to access to CCallout errno

Since CCallout class methods `errno` and `errno:` access a value stored in the `GsProcess`, there is a risk of that two FFI calls (within a single `GsProcess`) may access this, and set/retrieve incorrect values. (#49432)

The methods `errno` and `errno:` now signal an error and should not be used. You may still use `callWith:` if you do not need access to the `errno` information.

The method `CCallout >> callWith:errno:`, which was added in v3.6.4, can be used when the `errno` information is needed. See the method image comments for details.

Generated library code such as `GciLibrary` methods now invoke `callWith:errno:` rather than `callWith:`.

In-memory GC now responsive to memory use by FFI

In-memory garbage collection scans are now also responsive to memory pressure resulting from heavy use of `gcMalloc:`. This avoids excessive `CHeap` growth with intensive use of FFI with limited pressure on memory from `Smalltalk` objects.

CPreprocessor failed to cleanup temporary files

CPreprocessor could leave behind temporary files (#50182)

4.4 Updated Compile and Link Information

Linux Compile and Link Information

Compiler version

Red Hat Linux ES 9.1: gcc/g++ 11.3.1
 Red Hat Linux ES 8.7: gcc/g++ 8.5.0
 Red Hat Linux ES 7.9: gcc/g++ 4.8.5
 Ubuntu 20.04 LTS: gcc/g++ 9.4.0
 Ubuntu 22.04 LTS: gcc/g++ 11.3.0

Debugger version

Red Hat Linux ES 9.1: gdb 10.2-10.el9
 Red Hat Linux ES 8.7: gdb 9.1
 Red Hat Linux ES 7.9: gdb 9.1
 Ubuntu 20.04 LTS: gdb 9.1
 Ubuntu 22.04 LTS: gdb 112.1

Compiling a user action or GCI application

```
g++ -fmessage-length=0 -fcheck-new -O3 -ggdb -m64 -pipe
    -D_REENTRANT -D_GNU_SOURCE -pthread -fPIC -fno-strict-aliasing
    -fno-exceptions -I$GEMSTONE/include -x c++ -c userCode.c
    -o userCode.o
```

The following warn flags are recommended for compilation:

```
-Wformat -Wtrigraphs -Wcomment -Wno-aggregate-return -Wswitch
-Wshadow -Wunused-value -Wunused-variable -Wunused-label
-Wno-unused-function -Wchar-subscripts -Wmissing-braces
-Wmissing-declarations -Wmultichar -Wparentheses -Wsign-compare
-Wsign-promo -Wwrite-strings -Wreturn-type
-Wno-format-truncation -Wuninitialized
```

Linking a user action library

```
g++ -shared -Wl,-hlibuserAct.so userCode.o $GEMSTONE/lib/gcualib.o
    -o libuserAct.so -m64 -Wl,-Bdynamic,--no-as-needed -lpthread
    -Wl,--as-needed -lcrypt -ldl -lc -lm -lrt -znoexecstack
```

Linking a GCI application

```
g++ userCode.o $GEMSTONE/lib/gcirtlobj.o -m64
    -Wl,-Bdynamic,--no-as-needed -lpthread -Wl,--as-needed -lcrypt
    -ldl -lc -lm -lrt -z noexecstack -Wl,-traditional -Wl,-z,lazy
    -o userAppl
```


AIX Compile and Link Information

Compiler version

IBM XL C/C++ for AIX, V16.1

Debugger version

dbx

Compiling a user action or GCI application

```
xlC_r -O3 -qstrict -qalias=noansi -q64 -+ -qpic
-qthreaded -qarch=pwr6 -qtune=balanced -D_LARGEFILE64_SOURCE
-DFLG_AIX_VERSION=version -D_REENTRANT -D_THREAD_SAFE
-qminimaltoc -qlist=offset -qmaxmem=-1 -qsuppress=1500-010:1500
-029:1540-1103:1540-2907:1540-0804:1540-1281:1540-1090:1540-2988
-qnoeh -I$GEMSTONE/include -c userCode.c -o userCode.o
```

Depending on your version of AIX, you need to include `-DFLG_AIX_VERSION=71` or `-DFLG_AIX_VERSION=72`.

Also note that there is no space in the `-qsuppress` arguments that are continued on the following line.

Linking a user action library

```
xlC_r -G -Wl,-bdatapsize:64K -Wl,-btextpsize:64K
-Wl,-bstacksize:64K -q64 userCode.o $GEMSTONE/lib/gciualib.o
-o libuserAct.so -e GciUserActionLibraryMain -LcompilerInstallDir/lib
-lpthreads -lc_r -lC_r -lm -ldl -lbsd -lpam -Wl,-berok
```

Linking a GCI application

```
xlC_r -Wl,-bdatapsize:64K -Wl,-btextpsize:64K
-Wl,-bstacksize:64K -q64 userCode.o $GEMSTONE/lib/gcirtlobj.o
-Wl,-berok -L/compilerInstallDir/lib -lpthreads -lc_r -lC_r -lm -ldl
-lbsd -lpam -Wl,-brtllib -o userAppl
```

DARWIN Compile and Link Information

Compiler version

Apple clang version 14.0.0 (clang-1400.0.29.102)

Debugger version

lldb-1400.0.30.3

Apple Swift version 5.7

Compiling a user action or GCI application

```
g++ -fmessage-length=0 -O3 -ggdb -m64 -pipe -fPIC
-fno-strict-aliasing -D_LARGEFILE64_SOURCE -D_XOPEN_SOURCE
-D_REENTRANT -D_GNU_SOURCE -I$GEMSTONE/include -x c++
-c userCode.c -o userCode.o
```

The following warn flags are recommended for compilation:

```
-Wno-format -Wtrigraphs -Wcomment -Wsystem-headers -Wtrigraphs
-Wno-aggregate-return -Wswitch -Wshadow -Wunused-value
-Wunused-variable -Wunused-label -Wno-unused-function
-Wchar-subscripts -Wno-conversion -Wmissing-braces -Wmultichar
-Wparentheses -Wsign-compare -Wsign-promo -Wwrite-strings
-Wreturn-type -Wno-nullability-completeness
-Wno-expansion-to-defined
```

Linking a user action library

```
g++ -dynamiclib userCode.o $GEMSTONE/lib/gciualib.o
-o libuserAct.dylib -m64 -lpthread -ldl -lc -lm -lpam -undefined
dynamic_lookup
```

Linking a GCI application

```
g++ userCode.o $GEMSTONE/lib/gcirtlobj.o -undefined dynamic_lookup
-m64 -lpthread -ldl -lc -lm -lpam -o userAppl
```

Windows Compile and Link Information

Compiler/Debugger version

Microsoft Visual Studio Professional 2019 Version 16.11.24

VisualStudio.16.Release/16.11.24+33328.57

Microsoft .NET Framework Version 4.8.09032

Visual C++ 2019 00435-60000-00000-AA889

Compiling a GCI application

```
cl /W3 /Zi /MD /O2 /Oy- -DNDEBUG /TP /nologo /D_LP64 /D_AMD64_
/D_CONSOLE /D_DLL /DWIN32_LEAN_AND_MEAN
/D_CRT_SECURE_NO_WARNINGS /EHsc -volatile:iso
/DNATIVE /I 'VisualStudioInstallPath\include'
/I 'VisualStudioInstallPath\Tools\MSVC\14.29.30133\atlmfc\include'
/I 'VisualStudioInstallPath\Tools\MSVC\14.29.30133\include'
/I '%GEMSTONE%include' /c 'userCode.c' -FouserCode.obj
-FduserCode.pdb
```

Linking a GCI application

```
link /LIBPATH:VisualStudioInstallPath\Tools\MSVC\14.29.30133\lib\x64
/LIBPATH:VisualStudioInstallPath\Tools\MSVC\14.29.30133\atlmfc\lib\x64
/LIBPATH:WindowsKitsInstallPath\10\lib\10.0.18362.0\um\x64
/LIBPATH:WindowsKitsInstallPath\10\lib\10.0.18362.0\ucrt\x64
/DEBUG /RELEASE /OPT:REF /INCREMENTAL:NO /MAP /nologo
/MANIFEST /MANIFESTFILE:userAppl.exe.manifest
/MANIFESTUAC:level='asInvoker' /FORCE:MULTIPLE userCode.obj
%GEMSTONE%lib\gcirtlobj.obj %GEMSTONE%lib\gcirpc.lib ws2_32.lib
netapi32.lib advapi32.lib comdlg32.lib user32.lib gdi32.lib
kernel32.lib winspool.lib Secur32.lib Dbghelp.lib
/NODEFAULTLIB:libcmt.lib /out:userAppl.exe
```

The following bugs were present in v3.6.5 and are fixed in this version.

Checkpoint when tranlogs full can result in Stone shutdown

When the transaction logs are full, the Stone pauses to wait for space to become available. If a checkpoint starts while the Stone is waiting, it may have resulted in the Stone shutting down. (#49704)

Shrpcmon crash during crashed client frame lock recovery

The shared page cache monitor frame lock recovery may crash under some circumstances with ongoing page reads, resulting in the message "Freezing shared page cache" (#50082)

Cache Warming Issues

Cache warming could have caused commit record backlog

Cache warming runs in transaction, and when performing a large amount of warming over a slower connection, it could result in a commit record backlog. Now, while still running in transaction, cache warming monitors the commit record backlog and aborts if necessary. (#48419)

Cache warming could hang

There is a race condition between the multiple threads of the cache warmer, which could result in a state flag being set incorrectly, such that all threads were waiting and no progress was made. (#50258)

Race condition results in Gem SEGV in copyFrom:to:

The primitive that supports `Array` and `OrderedCollection` >> `copyFrom:to:` (prim 817), contains an unsafe object allocation. There is a race condition, if a scavenge is triggered by an object faulting into memory after the new object is allocated; the new object could be initialized to zero rather than `OOP_NIL`, causing a Gem SEGV #(50277)

Login related issues

A stuck login with UserSecurityData locked may block further logins

During login, the session briefly locks the UserSecurityData for update. If the login process has problems and does not complete, so the UserSecurityData remains locked, further logins for that user may fail. (#49866)

Slow timeout failure on login when network connection table is full

If the network connection table was full, including with sessions that were zombies or in login or logout status, new sessions attempting to log in could fail login with the error (depending on version), "the maximum number of users is already logged in". Now, the Stone is more aggressive about processing zombies when the table is full. (#49927)

Improved reporting on client disconnect errors

The errors reported when a Gem is unexpectedly disconnected has been improved to make debugging easier: now the specific code and the affected socket are included in the error messages. (#50242)

Gems in login during stone shutdown may fail to exit

A timing condition when the Stone shuts down during a window in the Gem login process may leave the Gems alive after the Stone is shut down. (#50204)

Login log did not record failed logins

The login log, enabled by the configuration parameter `STN_LOGIN_LOG_ENABLED`, was intended to record failed logins as well as logins that succeeded. However, failed logins were not being recorded. (#51069)

When commits suspended, symbol creation causes SymbolGem failure

When commits are suspended (as by `Repository>>suspendCommitsForFailover`), the SymbolGem cannot commit. If a symbol is created (e.g. by executing code), the SymbolGem errored and shutdown, and the Gem hung. (#50185)

SecurityError on objectSecurityPolicy: for a class with no instance variables

If a class has no instance variables, `instVarNames` is an empty Array, which is canonicalized to the object with oop 233217, which is in `SystemObjectSecurityPolicy`. If this class is sent `objectSecurityPolicy`, the code was incorrectly attempting to reassign the object security policy of this empty Array, which reported a `SecurityError`. (#49921)

Bugs in Configuration Parameters

Improved distribution over extents for DBF_ALLOCATION_MODE

Previously, allocation to extents was by free pages, not by extent size, which meant that in certain configurations the extents would not grow per the weights in DBF_ALLOCATION_MODE. (#49992)

Computation of SHR_PAGE_CACHE_NUM_PROCS too small

When SHR_PAGE_CACHE_NUM_PROCS is default (-1), it computes a value based on STN_MAX_SESSIONS and other configuration parameters. This computation was incorrect, resulting in a value that was somewhat smaller than necessary, in the case where every slot was required. (#49458)

Backup and Restore issues

Backup and restore failed to check OOP upper bound

It was possible for programmatic backup to produce a backup file that included OOPs that were higher than the OOP high water for the backup file (which is stored in the backup when the backup is initiated). Restoring this backup produced errors and potentially crashed or corrupted the Gem and/or Stone. (#50081)

Backup to a file with a nonexistent directory produced unclear error message

When making a programmatic backup and including a directory with the pathname, the error message reported was not helpful. (#50239)

Commit as restoreFromBackup starts could cause Stone shutdown

When `restoreFromBackup:` is initiated, another session such as the `SymbolGem` or `ReclaimGem` could commit. This would cause the Stone to shutdown. The restore code now checks for this condition and errors. (#50184)

Finding references failed to report objects in large IdentityBags

The operations `allReferences:`, `findAllReferences`, and `findReferences` do not include references for an objects in a large `IdentityBag` (more than 2K elements). (#40163)

GsFile primitives not interrupted by SIGTERM

When a Gem is executing `GsFile` read primitives, and the Stone was shutdown, the Gem was not interrupted, which could leave the shared memory allocated. (#50186)

Upgrade fails with nonstandard decimalPoint Locale with GsPackagePolicy

If a repository has `GsPackagePolicy` enabled (generally a `Seaside` or `GsDevKit` application), and the `Locale` does not specify a period (US-style) `decimalPoint`, `previousVersion` testing failed, and `upgradeImage` reported an error. (#50063)

Gem crash recovery may cause Stone to SEGV

When a Gem crashes, the Shrpcmon performs recovery to clean up spin locks. In versions before v3.6.5, the Shrpcmon process could SEGV if an address into the process table was out of range (bug #49988, fixed in v3.6.5). Additional codepaths existing outside of the Shpcmon, such as the Stone, which could also SEGV on an attempt to access in invalid offset; these are fixed in this version. (#50005)

After nbLogin, GsTsExternalSession wait methods SIGSEGVed or errored

After invoking `GsTsExternalSession >> nbLogin`, either `waitForReadReady`, `waitForReadReadyTimeout:`, `waitForResult`, or `waitForResultForSeconds:`, should be called, to detect when the login has completed.

`waitForResult` resulted in a SIGSEGV.

`waitForReadReady`, `waitForReadReadyTimeout:`, `waitForResultForSeconds:` errored with call not in progress, but did not crash. (#50203)

Memory leak in GciTsNbLogin

The function `GciTsNbLogin`, which is invoked when using `GsTsExternalSession`, contains a small memory leak. (#50215)

TransactionBacklog signaled incorrectly

When the Gem configuration parameter `#GemAutoServiceSigAbort` is set to false, sessions could still receive signal `#3007/TransactionBacklog`. (#50142)

Object >> subclassResponsibility error was not reported correctly

The messages composed by `subclassResponsibility` and `subclassResponsibility:` did reported the receiver, rather than the class of the receiver; in addition `subclassResponsibility:` incorrectly parsed the selector for class methods. (#50072)

stoned exited with code 3 on clean shutdown

On a clean stone shutdown using `stopstone`, `stopstone` exited with code 0. However, the `stoned` process itself exited with code 3. Now, `stoned` will exit with code 0 on clean `stopstone`. (#50222)

Gems may be slow to shutdown due to memory checks

With very large temporary object memory, a Gem make take considerable time to verify memory before shutting down. The verification is not critical, and is now skipped in the standard (fast) environment. (#50314)

Hot standby related issues

Starting hotstandby continuous restore with obsolete tranlogs may crash

When the STN_TRAN_LOG_DIRECTORIES directory of the slave system in a hotstandby system contains obsolete tranlogs, starting continuous restore may crash when attempting to read logs. In addition, now if a previously running continuous restore (hotstandby) has stopped due to a failure, the restoreStatus includes "continuous restore failed". (#49863).

Hot standby logreceiver errors if there are other logreceivers for the same logsender

A master stone's logsender did not correctly service multiple slave systems logreceivers. If multiple logreceivers are connected to a single logsender, the data sent to the logreceivers was out of sync, and the logreceivers reported validation errors. (#50064)

Problems after circular failOverToSlave

After a failOverToSlave from NodeA to NodeB and a subsequent second failOverToSlave from NodeB back to NodeA, there could be issues on the new master NodeA.

NodeA's transient free oop list could have contained the oops of committed objects, which could result in corruption; this issue was cleared by a stone restart. (#50147)

The master stone NodeA was also left with commit suspended, which required executing resumeCommits. (#50155)

Checking for vote state incorrect

The methods `Repository >> reclaimAllWait:` and `waitForVoteStateIdleSecs:` check if voting is occurring. This was checking the wrong voteState result. This bypassed reporting on specific sessions that were holding up voting. (#50114).

\$GEMSTONE paths containing symlinks that change

The NetLDI and gemnetobject and related scripts did not properly handle the case where the \$GEMSTONE path included a symbolic link, and the target of the symbolic link changed while the NetLDI was running. The gemnetobject (and related scripts) used the current value of \$GEMSTONE, rather than the argument from the NetLDI, which created an unclear state. Now, if the resolved \$GEMSTONE in the environment in which login is initiated does not match the \$GEMSTONE path resolved when the NetLDI was started, login will fail with the message "expanded value of GEMSTONE environment variable changed" (#50173)

SymbolGem printed excessive messages to its log on GC

The SymbolGem's log file could excessively bloat with log messages, particularly commit messages related to symbol garbage collection. (#50112)

Unicode Compares causes strings to not be returned in ExportedDirtyList

The ExportedDirtyList is used in GBS to correctly handle the state of objects replicated to the client, independent of having references on the client. However, the process of performing an ICU-based unicode comparison (regardless of whether the server is in Unicode Comparison Mode), caused a bit to be unset and the object was not returned in the ExportedDirtyList. (#50266)

FFI structures may fault out of memory and lose C data

There are cases in which CByteArrays and CPointers that used by FFI can be committed, faulted out, and lose C data when faulted in. (#49769).

Recent releases include workarounds to avoid the subsequent Gem crash due to bug 49765; the underlying condition is fixed in v3.7.

performOnServer: failed for configurations with client s-bit set

When the client executable (linked topaz or Gem) has the s-bit set, and is started by a user other than the owner of the executable, `System class >> performOnServer: failed` due to permission errors on temporary files. (#50191)

Unnecessary commit conflicts

With Indexing operations

BTreePlus indexes may encounter commit conflicts from internal structures; that is, operations that update an indexed collection and should succeed fail with a commit conflict. (#50146)

On RcidentitySet/Bag

If internal leaf balancing code executes during an add operation, the RcRead sets will not be correct and could result in commit conflicts. (#50220)

Numeric and Time related issues

A Time did not compare equal to equivalent SmallTime

A SmallTime and a Time that represent equivalent points in time did not compare as equal using `=` (equal sign). (#50118)

SmallDateAndTime passivate activate failed for years 2035-2072

Activating a passivated SmallDateAndTime in the approximate year range 2035 to 2072 errored; the activation expected a SmallScaledDecimal value in the internal storage, but the stored value was not in the SmallScaledDecimal range. (#50172)

Number class >> parseLiterals:exponent: handling nil character argument

If the `parseLiterals:` argument to this method was nil, rather than a valid character, it resulted in an error. Now, the method completes without effect as documented. (#49817)

Division by SmallInteger minimumValue threw error

Dividing an Integer by `SmallInteger minimumValue` resulted in an `InternalError, StoreSmallInt out of range`. (#50111)

ScaledDecimal, Fraction >> asFloat could return incorrect results

For some input values, the method `ScaledDecimal >> asFloat` and `ScaledDecimal asFloat` could return slightly incorrect values. These methods have been reimplemented, and return the closest floating point number, rounding to nearest even in case of a tie. (#50071, #50170)

CharacterCollection >> isDigits incorrect for empty string

For an empty String, this method incorrectly returned true. (#50145)

Float >> asDecimalFloat results not as precise as possible

For many values, converting a Float to a DecimalFloat using `asDecimalFloat` returned a result that was not the closest possible DecimalFloat to the given Float. (#50212)

Protocol error on incomplete LGC_PAD_N read over socket

Gems may encounter a GCI protocol error, due to socket read returning bytes that contain an incomplete `LGC_PAD_N` packet (#50045)

topaz set cachename did not work for linked topaz

Setting the name recorded in statmonitor data using the topaz command `set cachename`, in linked topaz, did not actually update the name of the process as recorded in statmonitor. (#50053)

Class subclassed from nil did not have default doesNotUnderstand handling

If a class was subclassed from nil, and `doesNotUnderstand:` was not implemented for that subclass, sending a message selector to an instance that was not understood resulted in infinite recursion. (#49768)

Bitmap written by #saveWriteSetUnionToFile not readable

The Admin GcGem option `#saveWriteSetUnionToFile`, which is false by default, triggers the AdminGem to write a bitmap file containing the write set union. This file was not readable by hidden set nor by GsBitmap protocol. (#49730)

GsTestCase assert:isEquivalentTo: failed to catch errors

This method has been removed, to avoid a false impression of providing test coverage. This method supported returning true for two numeric arguments that are somewhat close, but not exactly equal. This resulted in failure to detect actual incorrect results. (#50159)

Transaction log debug level change

The configuration parameter `STN_TRAN_LOG_DEBUG_LEVEL`, when set to values greater than 1, cause additional information to be written to the transaction logs; this should only be done as instructed by GemTalk Engineering. In v3.7, the output when `STN_TRAN_LOG_DEBUG_LEVEL` is set to level 1 includes page allocation information that previously required a higher level. The additional transaction log space requirement at level 1 is expected to be on the order of 1% of additional space over the requirement for level 0. (#50241)

PositionableStreamLegacy positionW, positionW: did not signal Warning

These methods are intended to signal a Warning to verify code when transitioning an application from 1-based to 0-based legacy streams (that is, not `PositionableStreamPortable`). Previously, these wrote a stack to the Stone now; now, a Warning is signaled.

PPParser >> child included a send to #assert:

The `PetiteParser` class `PPParser` method `#child` included a `send to #assert:`, although `#assert:` is not implemented within that hierarchy. (#50193)