
GemStone®

GemStone/S 64 Bit Release Notes

Version 2.3

October 2008

GEMSTONE[™] S 64

INTELLECTUAL PROPERTY OWNERSHIP

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

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

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 GemStone Systems, Inc. under separate license agreement. This software may only be used pursuant to the terms and conditions of such license agreement. Any other use may be a violation of law.

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

This software is provided by GemStone Systems, Inc. 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 GemStone Systems, Inc. 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 © 1986-2008 GemStone Systems, Inc. All rights reserved by GemStone Systems, Inc.

PATENTS

GemStone is covered by U.S. Patent Number 6,256,637 "Transactional virtual machine architecture", Patent Number 6,360,219 "Object queues with concurrent updating", and Patent Number 6,567,905 "Generational Garbage Collector". GemStone may also be covered by one or more pending United States patent applications.

TRADEMARKS

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

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

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

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

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

Microsoft, **MS**, **Windows**, **Windows 2000** and **Windows XP** 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.

AIX and **POWER4** are trademarks or registered trademarks of International Business Machines Corporation.

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. All terms mentioned in this documentation that are known to be trademarks or service marks have been appropriately capitalized to the best of our knowledge; however, GemStone 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.

GemStone Systems, Inc.

1260 NW Waterhouse Avenue, Suite 200
Beaverton, OR 97006

Preface

About This Documentation

These release notes describe changes in the GemStone/S 64 Bit version 2.3 release. We recommend that everyone migrating to this version read these release notes before beginning installation, testing or development.

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

These documents are also available on the GemStone customer website, as described below.

Terminology Conventions

This document uses the following terminology:

The term “GemStone” is used to refer both to the product, GemStone/S 64 Bit, or previous GemStone/S server products; and to the company, GemStone Systems, Inc.

Technical Support

GemStone provides several sources for product information and support. The product-specific manuals and online help provide extensive documentation, and should always be your first source of information. GemStone Technical Support engineers will refer you to these documents when applicable.

GemStone Web Site: <http://support.gemstone.com>

GemStone’s Technical Support website provides a variety of resources to help you use GemStone products. Use of this site requires an account, but registration is free of charge. To get an account, just complete the Registration Form, found in the same location. You’ll be able to access the site as soon as you submit the web form.

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

Help Request allows designated support contacts to submit new requests for technical assistance and to review or update previous requests.

Documentation for GemStone/S 64 Bit is provided in PDF format. This is the same documentation that is included with your GemStone/S 64 Bit product.

Release Notes and **Install Guides** for your product software are provided in PDF format in the Documentation section.

Downloads and **Patches** provide code fixes and enhancements that have been developed after product release. Most code fixes and enhancements listed on the GemStone Web site are available for direct downloading.

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

TechTips, also in the Learning Center section, provide information and instructions for topics that usually relate to more effective or efficient use of GemStone products. Some Tips may contain code that can be downloaded for use at your site.

Community provides customer forums for discussion of GemStone product issues.

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

You may need to contact Technical Support directly for the following reasons:

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

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

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

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

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

For non-emergency requests, the support website is the preferred way to contact Technical Support. Only designated support contacts may submit help requests via the support website. If you are a designated support contact for your company, or the designated contacts have changed, please contact us to update the appropriate user accounts.

Email: support@gemstone.com

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

Requests for technical assistance may also be submitted by email or by telephone. We recommend you use telephone contact only for more serious requests that require immediate evaluation, such as a production system that is non-operational. In these cases, please also submit your request via the web or email, including pertinent details such as error messages and relevant log files.

If you are reporting an emergency by telephone, select the option to transfer your call to the technical support administrator, who will take down your customer information and immediately contact an engineer.

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

24x7 Emergency Technical Support

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

Training and Consulting

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

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

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

Chapter 1. GemStone/S 64 Bit 2.3 Release Notes

Overview 1

Changes and New Features 1

 New platforms 1

 Improved Performance 2

 Pinned page cleanup 2

 NSC performance 2

 Commit record disposal performance enhancements 2

 Added support for direct I/O 2

 Mark/sweep performance 2

 Extended Character support enhancements 2

 QuadByteString support 2

 UTF encoding 3

 Persistent Shared Counters 4

 GsFile support for GZIP compressed files 4

 Ability to disable commits for sessions 5

 Ability to turn off handling of terminating OS signals 5

 Stack traces on Linux and Mac 6

 Added Stream methods 6

 #, now legal selector 6

 System gems not affected by GEM_TEMPOBJ_CACHE_SIZE 6

 Improved printing of very long Strings 6

 ProfMonitor enhancements 6

 objectAudit now requires reclaim 7

 New GcGem parameter 7

 Utilities with new parameters 8

 startcachewarmer 8

 topaz 8

 Changes in Exceptions/Errors 8

Legacy Exception Framework now provides error message for ANSI exceptions.	8
Added Errors.	8
Configuration Parameter Changes	8
Statmonitor and statistics changes.	10
Programmatic access to system statistics	12
Removed Cache Statistics.	12
Renamed Cache Statistic	12
Added Cache Statistics	12
Reclaim GcGem session stats.	14
Bugs Fixed	15
Signal handler timestamps have slight risk of deadlock	15
AIO page server failures may introduce corruption.	15
Stone did not respond to SIGUSR1	15
Conversion issues	15
Segment-related issues	15
IndexManager autoCommit not initialized	15
Hashed lookup failed on String/DoubleByteStrings containing y	16
Commit during login could have failed, causing later errors	16
Extent pregrow may have left defunct pgsvr processes.	16
Errors in Statically linked custom gems	16
Loading second user action failed on Linux	16
Application load could cause uncatchable code memory space overflow	16
System cacheStatsForGemWithName: sometimes incorrectly returned nil.	17
Risk of retry commit failure in RcEqualityIndexes	17
numArgs incorrectly returns 0 when selector has embedded \$_.	17
Misleading error message on stopSession:	17
Empty error on checkpoint suspend in partial logging mode	17
Authorization error modifying own SymbolList	17
Error executing usersInGroup:.	17
GsSocket connect: may succeed but return false.	17
Fraction, ScaledDecimal raisedTo: returned Float.	17
Object audit re-scan reported twice as many errors.	17
copydbf -i appeared to hang on active tranlog.	18
Delayed response to terminateSession:timeout:	18
Risk of Netldi SEGV.	18
Lint warnings on Windows	18
renameAssociationFrom:to: deleted origin on duplicate target	18
Install on Linux may result in incorrect MANPATH	18

GemStone/S 64 Bit 2.3 Release Notes

Overview

GemStone/S 64 Bit 2.3 is a new version of the GemStone/S 64 Bit object server. This release provides many new features and performance improvements, and fixes a number of bugs; we recommend everyone using GemStone/S 64 Bit upgrade to this new version. The details of these changes are provided in this document.

These release notes provide changes between the previous version of GemStone/S 64 Bit, version 2.2.5.4, and version 2.3. If you are upgrading from a version prior to 2.2.5.4, please also review the release notes for each intermediate release to see the full set of changes.

For details about installing GemStone/S 64 Bit 2.3 or upgrading from earlier versions of GemStone/S 64 Bit or other GemStone server products, see the *GemStone/S 64 Bit Installation Guide* for version 2.3.

This release includes updated documentation:

- ▶ *System Administration Guide*
- ▶ *Gembuilder for C*
- ▶ *Topaz Programming Environment*

Changes and New Features

New platforms

This release adds support for the following platforms:

- ▶ HP-UX 11.23 and 11.31 on the Itanium processor
- ▶ Solaris 10 on the Intel x86 platform
- ▶ Linux Red Hat ES 5.0

In addition, full installations for Mac OSX 10.5.5 on the Intel x86 platform are available. These are not fully QAed nor supported, but may be used for development only.

Improved Performance

A number of optimizations and improvements in algorithms have been made in this release.

Pinned page cleanup

On session termination, there is now faster cleanup of object table pages pinned by that session in the shared page cache.

NSC performance

There have been improvements in performance of certain cases of difference and intersect operations of UnorderedCollections (NSCs).

Commit record disposal performance enhancements

Changes have been made in internal structures and algorithms to improve the performance of commit record disposal by about 64%. Page disposal and processing of requests from the Page Manager should also be somewhat faster.

Commit records are now cached, to avoid the cost of locating commit records for disposal. The size of this cache is set by the new configuration parameter `STN_COMMIT_RECORD_QUEUE_SIZE`. We recommend leaving this at the default, provided `STN_SIGNAL_ABORT_CR_BACKLOG` is set appropriately for the application.

Increasing the setting for `STN_COMMIT_RECORD_QUEUE_SIZE` will improve commit record disposal performance somewhat, but the stone will consume much more heap space.

New cache statistics have been added for monitoring and tuning commit record and page disposal, see “Utilities with new parameters” on page 8.

Added support for direct I/O

Support has been added for direct I/O, and concurrent I/O on AIX. See the added configuration parameters “`STN_EXTENT_IO_FLAGS`” and “`STN_LOG_IO_FLAGS`”, starting on page 9.

Mark/sweep performance

Algorithm changes in the mark sweep garbage collection structures and code improve performance for `Epoch`, `markForCollection`, `findDisconnectedObjects` and `findObjsConnectedTo`.

Extended Character support enhancements

QuadByteString support

Historically, instances of `Character` were limited to values of 65535. Instances of `String` could contain one-byte `Character`s (with values up to 256). If any `Character`s with values between 256 and 65535 were used, the `String` transparently became a `DoubleByteString`, which allowed two bytes per `Character`. `Symbol` and `DoubleByteSymbol` behaved similarly.

However, many character sets, such as Chinese, cannot be expressed in two byte `Character`s. To allow strings with all character sets, GemStone has introduced

QuadByteStrings, which contain Characters that are represented with four bytes. Similar to the way DoubleByteStrings are handled, operations on CharacterCollections transparently manage the relationships between Strings, DoubleByteStrings, and QuadByteStrings.

The range for Character has been correspondingly increased from 65535 (16rFFFF) to 2147483647 (16r7FFFFFFF).

Class MultiByteString has been added as a new abstract subclass of CharacterCollection. Class DoubleByteString, previously a subclass of CharacterCollection, is now a subclass of MultiByteString. Class QuadByteString has been added as a subclass of MultiByteString.

Class QuadByteSymbol has been added as a subclass of QuadByteString, but is not fully implemented. Customers should not use this class until the implementation is completed in a future release. Attempts to generate instances of QuadByteSymbol will generate an error.

The compiler does not currently accept instances of QuadByteString. Attempting to compile a QuadByteString will generate an error.

Added Methods

CharacterCollection >> _convertToQuadByte

Takes an existing String/DoubleByteString and converts it into an equivalent QuadByteString.

CharacterCollection >> asQuadByteString

Generates a new QuadByteString that is equivalent to the receiver.

charSize

All CharacterCollection subclasses and instances of CharacterCollection subclasses now understand the method charSize, which returns the number of bytes that make up a character in instances of this class.

In addition, many methods have been moved within the hierarchy, including a number of String methods which have been moved to CharacterCollection.

Unicode Representation

The Unicode tables provided in \$GEMSTONE/examples CharTableUnicode.dat have been updated to reflect version 5.1 (March 2008) of the Unicode standard, and now include the full tables (prior tables were truncated due to the 0xFFFF character range limit).

UTF encoding

The methods

```
String >> encodeAsUTF8
String >> decodeFromUTF8
String >> decodeIntoDoubleByteStringFromUTF8
MultiByteString >> decodeFromUTF8
MultiByteString >> encodeAsUTF8
```

have been added, implemented in primitives for speed. These methods respectively encode a String, DoubleByteString, or QuadByteString into the equivalent in UTF8

encoded, and decode a String in UTF8 encoding and return the equivalent String, DoubleByteString, or QuadByteString.

Persistent Shared Counters

Support has been added for persistent shared counters. Persistent shared counters provide a means for multiple sessions to share common integer values. There are 128 persistent shared counters, numbered from 1 to 128 (the index of the first counter is 1).

Each update to a persistent shared counter causes a round-trip to the stone process. However reading the value of a counter is handled by the gem (and its page server, if any) and does not cause a round-trip to the stone.

Persistent shared counters are globally visible to all sessions on all shared page caches.

Persistent shared counters hold 64 bit values and may be set to any signed 64 bit integer value. No limit checks are done when incrementing or decrementing a counter. Attempts to increment/decrement the counter above/below the minimum/maximum value of a signed 64-bit integer will cause the counter to 'roll over'.

Persistent shared counters are independent of database transactions. Updates to counters are visible immediately and aborts have no effect on them.

The values of all persistent shared counters are written to the primary database extent at checkpoint time. Updates between checkpoints are written to the transaction log by the stone. Therefore the state of the persistent shared counters is recoverable after a crash, restore from backup, and restore from transaction logs.

Added methods:

```
System class >> persistentCounterAt: index incrementBy: amount
System class >> persistentCounterAt: index decrementBy: amount
System class >> persistentCounterAt: index put: value
System class >> persistentCounterAt: index
```

GsFile support for GZIP compressed files

GsFile has added support for reading and writing files in gzip file format. New methods have been added, similar to existing methods, to work with compressed files.

The following methods have been added:

```
GsFile >> gzOpen: aPathName mode: openMode
GsFile >> gzOpen
GsFile >> isCompressed
GsFile >> contents
GsFile class >> validFileModes
GsFile class >> validFileModesForCompression
GsFile class >> open: aPathName mode: openMode onClient: clientBool
    withCompression: compBool
GsFile class >> openCompressed: aPathName mode: openMode
GsFile class >> openCompressed: aPathName mode: openMode
    onClient: clientBool
GsFile class >> openOnServerCompressed: aPathName mode: openMode
GsFile class >> openReadCompressed: aPathName
GsFile class >> openWriteCompressed: aPathName
```

```
GsFile class >> openReadOnServerCompressed: aPathName  
GsFile class >> openWriteOnServerCompressed: aPathName
```

For details on these methods, see the documentation in the image comments.

For compressed files, the set of openModes in arguments to open:mode:..., etc., are somewhat different. The set of supported modes is 'r', 'w', 'a', 'rb', 'wb', 'ab' (see the man page for fopen() for the meanings of these modes). Update mode (modes with +) are not supported for compressed files. In addition, with compressed files, any valid write mode argument (containing 'a' or 'w') may have a '1' or '9' appended. Appending a '1' means compress using the best possible speed while appending a '9' means compress to attain the best possible compression.

The method GsFile class >> validFileModesForCompression returns a list of all valid file modes for compressed files.

NOTE

As a result of these changes, shared libraries from previous versions of GemStone/S 64 Bit will not work with v2.3.

Ability to disable commits for sessions

GemStone/S 64 Bit 2.3 provides additional control over the ability of individual users and sessions to commit to the repository. Commits can be disabled for a particular user or for a specific session. A reason can be provided for disabling commits. It is also possible to disable commits until a session performs an abort.

The following methods have been added:

```
System class >> commitsDisabledUntilLogout  
System class >> commitsDisabledUntilAbort  
System class >> disableCommitsWithReason: aString  
System class >> disableCommitsUntilAbortWithReason: aString  
System class >> sessionCanCommit  
UserProfile >> disableCommits  
UserProfile >> enableCommits  
UserProfile >> isReadOnly
```

OtherPassword privilege is required to disable commits. Commits cannot be disabled for system accounts (SystemUser, DataCurator, SymbolUser, GcUser, and Nameless).

Ability to turn off handling of terminating OS signals

A new environment variable has been added, GS_DISABLE_SIGNAL_HANDLERS. When this environment variable is enabled (by setting it to any value; the actual setting does not matter) in the environment for a gem process, this gem sessions will not attempt to handle certain types of signals, but will shut down immediately. This avoids side effects with user action or client Smalltalk code.

When the environment variable GS_DISABLE_SIGNAL_HANDLERS is set, and the session receives a SIGSEGV, SIGBUS or SIGILL signal:

- ▶ Process will immediately exit
- ▶ Core files will not be generated
- ▶ C call stacks will not be printed

The applies only to gem sessions; other GemStone processes behave as before.

If this environment variable is enabled, a notice is printed in the header of the gem log.

Stack traces on Linux and Mac

To avoid problems with the installed location of gstack on Linux distributions, GemStone now includes gstack in the bin directory. This provides more reliable stack traces on Linux.

The gstack script is also provided for Mac.

Added Stream methods

The following methods have been added:

```
PositionableStream >> nextLine:
PositionableStream >> peekFor:
```

#, now legal selector

Per the ANSI specification, the compiler will now accept ,, (comma comma) as a method selector.

System gems not affected by GEM_TEMPOBJ_CACHE_SIZE

In cases where application gems need increased temporary object memory, the large amount of memory was also unnecessarily allocated to system gems, such as the SymbolGem and GcGems. Now, the SymbolGem uses a GEM_TEMPOBJ_CACHE_SIZE of 1000, and page manager and GcGems use 5000.

Improved printing of very long Strings

Printing of very long strings is truncated to avoid overflow problems. This truncation has been improved to avoid the results when there are embedded String demarcation characters.

The following methods have been added:

```
PrintStream >> maxSize
PrintStream >> isFull
PrintStream >> overflow
PrintStream class >> printingOn:maxSize:
```

```
CharacterCollection >> printString now calls PrintStream class >>
printingOn:maxSize:, and will truncate the contents at 100000 (100K) bytes. To
override this, use the method CharacterCollection >>
printStringWithMaxSize:.
```

ProfMonitor enhancements

The profiling output from running ProfMonitor has been improved and expanded.

The new output is similar to the gprof flat profile and the call graph:

```
http://www.cs.utah.edu/dept/old/texinfo/as/gprof.html#SEC5
```

The first section of the ProfMonitor report, titled STATISTICAL SAMPLING RESULTS, is unchanged. A new second section, STATISTICAL STACK SAMPLING RESULTS, has

been added. This section has similar information, but the total and percent are cumulative; each line time information includes the time for all message sends from that method.

The third section, titled STATISTICAL METHOD SENDERS RESULTS, equivalent to the gprof call graph, is entirely different than in previous versions. For each method, the method's parent and children methods are included as a block. The columns reported for the method and its parents and children are:

% self Time	% total Time	total ms	local %	Parent Method Child
The percent of time spent in the method itself.	The percent of time spent in the method and its children.	The total time spend in the method and its children for each method - parent, self and children	The percentage breakdown of the time across the parents and children. Always 0.0 for self.	The parent and child method names are indented in the display

In addition, if the currently unsupported option traceObjectCreation was turned on, profiling used excessive temporary object memory. The internal code has been disabled, avoiding this problem (#38807).

objectAudit now requires reclaim

The methods `Repository >> objectAuditNoReclaim` and `Repository >> auditWithLimit:reclaimAll:` are no longer supported and have been removed from the image. Audit results that do not immediately follow reclaim are not reliable; the audit may detect problems with objects that are dead, but not yet reclaimed. If the audit encountered such objects, race conditions could result in session errors.

New GcGem parameter

To allow throttling of dead object reclaim, a new parameter has been added for the GcGem, `#deadObjsReclaimedCommitThreshold`. This parameter determines the maximum number of dead objects to reclaim in a single transaction. This has a default value of 20000 objects.

Now, the reclaim gem will commit as soon as any one of the following conditions is met:

- 1) Number of live objects moved exceeds `#objsMovedPerCommitThreshold`.
- 2) Duration of the transaction exceeds `#maxTransactionDuration`.
- 3) Number of dead objects reclaimed exceeds `#deadObjsReclaimedCommitThreshold`.

Utilities with new parameters

startcachewarmer

The startcachewarmer utility now includes the -L flag, which specifies the location to put the cache warmer log files.

topaz

The topaz executable has added -q and -I switches.

-q starts topaz in quiet mode, suppressing printout of the banner and other information.

-I <topaziniFile> allows specification of a complete path and filename for a .topazini file, overriding the default topazini file location.

Changes in Exceptions/Errors

Legacy Exception Framework now provides error message for ANSI exceptions

If an ANSI exception is passed to the legacy exception framework, the error information that would have been provided by the ANSI handlers is now included.

Added Errors

The following errors have been added in this release:

2501	RT_ERR_GET_GC_LOCK_FAILED	An operation that need to acquire the gcLock was not able to get the lock within a reasonable amount of time.
4300	GS_ERR_MAX_REMOTE_CACHE_LIMIT	Login attempt failed - the maximum number of remote caches already exist.

Configuration Parameter Changes

The following configuration options have been removed:

GEM_SEND_STN_MSGS_VIA_PGSRV

The following configuration options have been added:

GEM_TEMPOBJ_POMGEN_PRUNE_ON_VOTE

The percent of the pom generation area to be thrown away when voting on possible dead objects. Only subspaces of pom generation older than 5 minutes are thrown away. The most recently used subspace is never thrown away.

If this value is not specified, or the specified value is out of range, the default is used.

Default: 50

Minimum: 0

Maximum: 90

STN_COMMIT_RECORD_QUEUE_SIZE

Determines the size of the stone's internal commit record cache. The stone will keep copies of up to this many commit records in heap memory. Stone is able to dispose commit records more quickly when a copy of the commit record is found in this cache.

When the default value of -1 is specified, stone sets this value to be two times the value of the STN_SIGNAL_ABORT_CR_BACKLOG option.

Default: -1

Min: 16

Max: 1000000

Units: Commit Record Pages

STN_EXTENT_IO_FLAGS

Specifies what (if any) special I/O flags will be used to open the database extents. Two kinds of special I/O are supported: direct I/O and concurrent I/O.

Direct I/O tells the operating system avoid caching extent data in the file system cache. Enabling direct I/O tells the operating system to treat the database extents as if they were on raw partitions. Direct I/O may greatly improve database performance in some cases. Concurrent I/O is only available to extents running on the Enhanced JFS file system (aka JFS2) on AIX. Setting this flag has no effect on other operating systems.

If the requested I/O mode is not available, the stone will fail to start and an error message will be printed in the stone log. If this happens, change this option back to zero and restart the stone.

STN_EXTENT_IO_FLAGS has no effect on extents which reside on raw partitions.

Once the stone starts, all processes which open the database extents (gems and page servers) will open the extents using the same I/O flags. This behavior is required by some operating systems.

STN_EXTENT_IO_FLAGS has the following possible values:

0 - no special I/O flags. This is the default.

1 - enable Direct I/O on all extents on file systems.

2 - enable concurrent I/O (AIX only)

Default: 0

Minimum: 0

Maximum: 2 (AIX only), 1 (All Others)

STN_LOG_IO_FLAGS

Specifies what (if any) special I/O flags will be used to open the database transaction logs. Two kinds of special I/O are supported: direct I/O and concurrent I/O.

Direct I/O tells the operating system avoid caching extent data in the file system cache. Enabling direct I/O tells the operating system to treat the database tranlogs as if they were on raw partitions. Direct I/O may greatly improve database performance in some cases.

Concurrent I/O is only available to tranlogs running on the Enhanced JFS file system (aka JFS2) on AIX. Setting this flag has no effect on other operating systems.

If the requested I/O mode is not available, the stone will fail to start and an error message will be printed in the stone log. If this happens, change this option back to zero and restart the stone.

STN_LOG_IO_FLAGS has no effect on tranlogs which reside on raw partitions.

STN_LOG_IO_FLAGS has the following possible values:

- 0 - no special I/O flags. This is the default.
- 1 - enable Direct I/O on all tranlogs on file systems.
- 2 - enable concurrent I/O (AIX only)

Default: 0

Minimum: 0

Maximum: 2 (AIX only), 1 (All Others)

STN_MAX_REMOTE_CACHES

The maximum number of remote shared page caches that the system may have.

Default: 16

Minimum: 0

Maximum: 65535

Statmonitor and statistics changes

The statmonitor interface has been enhanced to provide more configurable control, especially over the collection of system statistics. Not all of the additional capabilities are available for all platforms.

Statmonitor now has the following new features:

- ▶ When collecting system statistics, you can now record statistics per individual CPU rather than seeing the average of all CPUs; see the `-C` option.
- ▶ You can now collect statistics for a list of Process IDs; see the `-Q` option
- ▶ You can now collect statistics for all processes owned by one or more UNIX user ids; see the `-U` option.
- ▶ You can collect statistics for the host system only, without attaching to the shared cache or collecting GemStone statistics; see the `-X` option.

This is the full list of statmonitor options. Not all of the additional capabilities are available on all platforms. The Solaris platform supports all these functions, see below for other OS specific limitations.

- A Collect all available system statistics.
- C Collect system statistics for each individual CPU.
- D Collect system statistics for all disks and partitions
- f *<fileName>* The output file name (default is statmon<pid>.out). Use 'stdout' to send output to stdout instead of a file.
- i *<interval>* The interval in seconds (default is 20).
- I *<intervalMs>* The interval in milliseconds (default is 20000).
- h *<hours>* The number of hours (default is forever).
- m *<stoneHostName>* The stone host name (default is localHost).
- n *<numAppStats>* The number of application statistics (default is 0).
- N Collect system statistics for all network interfaces.

- p <sessionId> A GemStone sessionId to monitor. May be repeated. The default is to monitor all sessions.
- P Sample the stone, shared cache monitor and all AIO page servers only.
- q Quiet mode. Only print messages if an error occurs.
- Qpid1,pid2,... Record statistics for a list of process IDs.
- r Restart a new output file when the current one completes. Each file will be given a unique name. The -h or -t switches control when a restart is done.
- s0 (deprecated) Same as -Y
- s1 (ignored)
- s2 (deprecated) Same as -D
- s3 (deprecated) same as -D -N
- s4 (deprecated) same as -A
- S Sample only the stone and shared cache monitor.
- t <times> The number of samples (default is forever).
- u <seconds> The maximum number of seconds to wait before flushing the cached information to the output file (default is 60). If 0 then the flush will be done every interval.
- W Collect system statistics for system memory pages.
- U uid1,uid2,... Record statistics for all processes owned 1 or more UNIX user IDs.
- X Run in host-only mode. Sample host system statistics only and do not attach to a shared page cache. May be combined with other flags EXCEPT: -m, -n, -p, -P, -S, or -Y.
- Y Disable collection of all system stats, including per-process data.
- z Write the output in compressed gzip format.

Host system stats are available on the following platforms:

	CPU -C	Network -N	Disks -D	System Pages -W
HP-UX on PA-RISC	Yes	-	Yes	-
HP-UX on ia64 (Itanium)	Yes	-	Yes	-
Sun Solaris on SPARC	Yes	Yes	Yes	Yes
Sun Solaris on x86	Yes	Yes	Yes	Yes
Linux on x86	Yes	Yes	Yes	-
Apple Mac	-	-	-	-
AIX	Yes	Yes	Yes	-

-s flags are deprecated and should not be used.

To collect all available system flags, use the -A flag. This using this flag will not cause an error even on systems that have no system stats, such as Apple.

Disk stats include stats for both disks and partitions, if supported.

Similarly, network stats include simple networks, complex networks, and TCP, if supported.

CPU stats collect stats for each individual processor (but not core) in the system.

Programmatic access to system statistics

On most platforms, you can now access OS-level system statistics from GemStone Smalltalk. This is the same information collected by statmonitor.

System class>> fetchSystemStatNames

Return an array of Strings which describe the statistics returned by the #fetchSystemStats method. The length of the result array is host technology dependent. Returns nil if the host system does not support system statistics.

System class >> fetchSystemStats

Return an array of Numbers with the statistics specified in the #fetchSystemStatNames method. The length of the result array is host technology dependent. While most elements in the result array will be SmallIntegers, the result may also contain other types of Numbers such as SmallDoubles, Floats, LargePositiveIntegers, etc. Returns nil if the host system does not support system statistics.

Removed Cache Statistics

The following statistic has been removed:

CommitRecordsDisposable

Renamed Cache Statistic

The following statistic has been renamed:

stnAioWriteFailures is now **StnAioWriteFailures**

Change in capitalization only.

Added Cache Statistics

The following statistics have been added:

AioWriteFailures (Pgsvr)

Total number of write errors detected by this page server, including write errors that succeeded on retry.

CrPageLocateForDisposeCount (Stone)

Number of times a commit record page was not found in the stone's commit record cache at commit record disposal time.

PagesWaitingForRemovalDeferred (Stone)

Number of deferred persistent pages waiting to be processed by the Page Manager gem.

PagesWaitingForRemovalPersist (Stone)

Number of persistent pages waiting to be processed by the Page Manager gem.

PagesWaitingForRemovalTemp (Stone)

Number of temporary pages waiting to be processed by the Page Manager gem.

RemoteSharedPageCacheMax (Stone)

Current setting of the STN_MAX_REMOTE_CACHES configuration parameter

TotalBmPageReads (SPC Monitor)

Total number of bitmap pages read by all processes currently attached to the shared page cache.

TotalDataPageReads (SPC Monitor)

Total number of data pages read by all processes currently attached to the shared page cache.

TotalFramesFromFindFree (SPC Monitor)

Total number of frames found by scanning the cache for all processes currently attached to the shared page cache.

TotalKFramesAddedToFreeList (SPC Monitor)

Total number of frames (expressed in thousands) added to the free list by all processes currently attached to the shared page cache.

TotalKFramesFromFreeList (SPC Monitor)

Total number of frames (expressed in thousands) taken from the free list by all processes currently attached to the shared page cache.

TotalLocalPageCacheKHits (SPC Monitor)

Total number of local page cache hits (expressed in thousands) for all processes currently attached to the shared page cache.

TotalLocalPageCacheKMisses (SPC Monitor)

Total number of local page cache misses (expressed in thousands) for all processes currently attached to the shared page cache.

TotalMiscPageReads (SPC Monitor)

Total number of miscellaneous pages read by all processes currently attached to the shared page cache.

TotalOtPageReads (SPC Monitor)

Total number of object table pages read by all processes currently attached to the shared page cache.

TotalPageReads (SPC Monitor)

Total number of pages read by all processes currently attached to the shared page cache.

TotalPageWrites (SPC Monitor)

Total number of pages written by all processes currently attached to the shared page cache.

TotalWaitsForOtherReader (SPC Monitor)

Total number of PageRead operations avoided by all processes currently attached to the shared page cache.

TranlogKBytesWritten (Gem)

Total number of kilobytes written to the tranlog by stone on behalf of this session.

TranlogRecordKind (Gem)

The kind of tranlog record last written to the tranlog by stone on behalf of this session.

TranlogRecordsWritten (Gem)

Total number of physical tranlog records written to the tranlog by stone by this session.

TteCrPageFreeCount (Stone)

Total number of free commit record page entries in the stone's internal commit record cache.

TteCrPageFreePoolSize (Stone)

Total number of commit record page entries allocated in the stone's internal commit record cache.

TteFreeCount (Stone)

Total number of free TTEs (transaction table entries) in stone.

TteFreePoolSize (Stone)

Total number of TTEs (transaction table entries) allocated by stone.

Reclaim GcGem session stats

The session statistics recorded by the Reclaim GcGem have been increased in this release, to improve reclaim tuning. As of GemStone/S 64 Bit 2.3, the following Reclaim GcGem session statistics are recorded:

0. Write log record time (ms)
Time in calls for stone to write the initial tranlog record for the commit
1. Sort deltas time (ms)
Time to sort the object table deltas
2. Flush data pages time (ms)
Time to flush new data pages from the private cache to the shared cache
3. Get token time (ms)
Call stone to get the commit token (aka TimeWaitingForCommit)
4. Empty deltas prepare time (ms)
5. Empty deltas time (ms)
Time to update the object table
6. Build persist pages used time (ms)
Time to build a list of pages that this commit will make persistent
7. StnCallTranCommit time (ms)
Time to call stone to complete the commit
8. Post commit time (ms)
9. Page reads during empty deltas (stat 5)
10. Total Live oops moved
11. Total Dead objs reclaimed
12. Total OT updates (10 + 11)
13. Total Pages reclaimed
14. Total new data pages added to SPC
15. Live oops moved this transaction (new in v2.3)
16. Dead objs reclaimed this transaction (new in v2.3)
17. OT Updates this transaction (new in v2.3)
18. Pages scavenged this transaction (new in v2.3)
19. New data pages flushed to cache this transaction (new in v2.3)
20. (unused)
21. (unused)
22. value of #deadObjsReclaimedCommitThreshold reclaim gem setting (new in v2.3)
23. value of #objsMovedPerCommitThreshold reclaim gem setting

Bugs Fixed

The following bugs in GemStone/S 64 Bit 2.2.5.4 have been fixed in GemStone/S 64 Bit 2.3.

Signal handler timestamps have slight risk of deadlock

If a SIGTERM or SIGUSR1 was received by a process that is executing a C level `printf()`, `fprintf()`, or `fflush()`, there was a slight risk that the call to get the timestamp would result in a deadlock. (#37748)

AIO page server failures may introduce corruption

Under very rare circumstances, it is possible for the stone to not immediately detect and shut down in response to an aio page server write failure. This allows inconsistent data to be written to the extents and subsequent page cache faults. (#39359, #39364)

Now, aio page servers will retry following write failures up to 5 times. Each failure is reported to the log, even if the write passes on retry. The new cache statistic `AioWriteFailure` tracks all such write failures. If all retries fail, the stone will shut down, since dirty pages cannot be reliably written to the extents.

AIO page servers no longer respond to kill -TERM. Normally you should not need to terminate AIO page servers. If it is necessary, use kill -9.

Stone did not respond to SIGUSR1

Sending kill -USR1 to a GemStone process should print the C stack to the process log. This was not occurring for the stone process. Other processes, such as gems, did not have this problem, and did print stacks on SIGUSR1. (#38928)

Conversion issues

Segment-related issues

When a repository that originated in a version of GemStone that supported Segments, was upgraded from GemStone/S 64 Bit 1.x, the Segment update code created several new Segments rather than using the existing Segments. (#39195)

Repositories that originated in very early versions of GemStone, which had unusual Segment layouts, could also encounter problems during conversion. (#39273)

In addition, repositories that contained exactly 16 Segments, or where upgrade code ran and created Segments to reach this number, ran into an incorrect limit check during upgrade. (#39195)

Other minor Segment conversion issues were also fixed. (#39292)

IndexManager autoCommit not initialized

The IndexManager's variable `autoCommit` was not initialized during conversion, resulting in errors. (#38279)

Hashed lookup failed on String/DoubleByteStrings containing Ÿ

There was an error in the hashing algorithm for DoubleByteStrings containing the Character with ASCII value 255 - the character Ÿ. As a result, a String containing Ÿ and an otherwise equal DoubleByteString would not hash to the same value. This means that in hashed collections (such as Bag and KeyValueDictionary), elements that are DoubleByteStrings containing Ÿ could not be looked up in the collection using the equivalent String, and vice versa. Indexes do not use hashing and are not affected.

The bug in the hashing algorithm has been fixed, and new insertions into hashed collections will be correct. However, existing hashed collections in repositories that are upgraded will continue to have DoubleByteStrings containing Ÿ in the incorrect location. We recommend scanning hashed collections in your repository, and if any contain a DoubleByteString containing the Character Ÿ, that the hashed collection be rebuilt. (#39476)

Commit during login could have failed, causing later errors

During the login process, a new session updates the account's last login time, write-locks the security data information, and commits. In scenarios with many logins occurring simultaneously, the write lock may fail or be dirty, or the commit may fail. The code did not handle these scenarios and several subsequent problems could occur:

- ▶ The first user commit after login by the new session may fail, since no abort was done following the failed commit.
- ▶ The write lock may remain on the security data information.
- ▶ User information, such as login count, may not get updated. This affects results such as number of logins permitted before a password change is required.

(#39478)

Extent pregrow may have left defunct pgsvr processes

Extent pregrow starts up page serve processes to perform its operation. In some cases, these processes may have been left defunct, not shut down, after the pregrow was complete. (#39477)

Errors in Statically linked custom gems

AIX and Linux only

Statically linked custom gems encountered runtime errors on Linux and AIX (#35226, #35227)

Loading second user action failed on Linux

Linux only

Attempting to load a second user action on Linux failed. (#39275)

Application load could cause uncatchable code memory space overflow

With many large methods, the application load could cause out of memory errors that do not trigger regular memory limit handling. This is related to the way code is loaded into the code_gen area of VM memory. (#39226)

System cacheStatsForGemWithName: sometimes incorrectly returned nil

The method `System class >> cacheStatsForGemWithName:` could return nil if any session was found that had been killed by lostOT and not yet removed from the internal tables. (#39371)

Risk of retry commit failure in RcEqualityIndexes

Due to holes in the retry logic, there was a risk of errors when using `RcEqualityIndexes` with heavy load and large numbers of failures. (#36664)

numArgs incorrectly returns 0 when selector has embedded \$ _

If a method selector contains `$ _`, other than as the first character of a keyword, the `numArgs` method will return 0. (#38633)

Misleading error message on stopSession:

When a session is shutdown forcibly using `stopSession:`, the message in the gem log refers to "stopGcSession". (#39257)

Empty error on checkpoint suspend in partial logging mode

In partial logging mode, it is not permitted to suspend checkpoints. Sending the message `suspendCheckpointsForMinutes:`, however, returned error 0, with no specific error message. (#38769)

Authorization error modifying own SymbolList

Modifications to your own `SymbolList` incorrectly required `#OtherPassword` as well as `#CodeModification`. While modifications to another user's `SymbolList` requires both privileges, only `#CodeModification` is required to modify your own `SymbolList`. (#38317)

Error executing usersInGroup:

Executing this method resulted in an error. (#38721)

GsSocket connect: may succeed but return false

A successful `GsSocket` connect on an IP address was checking the peer name to verify success. On network configurations that do not support reverse name lookup, this caused the method to return false, in spite of a valid connection. (#38894)

Fraction, ScaledDecimal raisedTo: returned Float

This operation returned a `Float`. The ANSI specification states that for an `Integer` argument, `raisedTo:` should return the same type as the receiver. (#39102)

Object audit re-scan reported twice as many errors

After the initial scan that `objectAudit` performs, if errors are found a detailed re-scan is performed. This did not reset the error count to 0, so the final reported error count was twice as large. (#39361)

Also, the object audit report has been enhanced to print if it is run in single-user mode and all checks would be performed, or in multi-user mode when only partial checking is done.

copydbf -i appeared to hang on active tranlog

Executing copydbf -i on the currently active tranlog opened and read the file as a compressed file, which is very slow. As a result, copydbf -i on the active tranlog appeared to hang. (#39446)

Delayed response to terminateSession:timeout:

Sending the message `System class >> terminateSession:timeout:`, on some platforms, did not immediately terminate the OS process. (#36806)

Risk of Netldi SEGV

There was a potential risk that during netldi processing (which includes RPC logins and netldi shutdown, for example), a string compare operation could receive a null argument, resulting in a SEGV. It is unclear how this could occur in practice. (#38850)

Lint warnings on Windows

When compiling GCI code or user actions on Windows, depending on the application the compiler returned warnings regarding conversion between numeric types with possible loss of data. (#39305)

renameAssociationFrom:to: deleted origin on duplicate target

When renaming a Symbol using `SymbolDictionary >> renameAssociationFrom:to:`, if an association with the target name already existed, the association that was being renamed was deleted. (#39426)

Install on Linux may result in incorrect MANPATH

Linux only

On Linux, the system default man path is used when MANPATH is not set, or when MANPATH starts or ends with ':' or contains '::' (i.e. contains an empty substring). GemStone modifies the MANPATH during installation, to add \$GEMSTONE/doc. However, if there was previously no MANPATH, GemStone was setting a default MANPATH that did not include ':' or '::', resulting in the system not being able to find man information. (#38143)