
GemStone®

GemStone/S™ Release Notes

Version 6.6.5

July 2014



INTELLECTUAL PROPERTY OWNERSHIP

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

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

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

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

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

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

COPYRIGHTS

This software product, its documentation, and its user interface © 1986-2014 GemTalk Systems, LLC. All rights reserved by GemTalk Systems.

PATENTS

GemStone software is covered by U.S. Patent Number 6,256,637 "Transactional virtual machine architecture", Patent Number 6,360,219 "Object queues with concurrent updating", Patent Number 6,567,905 "Generational garbage collector with persistent object cache", and Patent Number 6,681,226 "Selective pessimistic locking for a concurrently updateable database". GemStone software may also be covered by one or more pending United States patent applications.

TRADEMARKS

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

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

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

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

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

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

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

Linux is a registered trademark of Linus Torvalds and others.

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

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

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

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

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

GemTalk Systems

15220 NW Greenbrier Parkway
Suite 240
Beaverton, OR 97006

Preface

About This Documentation

These release notes describe changes in the GemStone/S™ version 6.6.5 release. Read these release notes carefully before you begin installation, conversion testing, or development with this release.

No separate Installation Guide is provided with this release. For instructions on installing GemStone/S version 6.6.5, or upgrading or converting from previous products or versions, see the Installation Guide for version 6.6.2.

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

Terminology Conventions

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

Technical Support

GemTalk Website

<http://gemtalksystems.com/techsupport>

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

- ▶ **Documentation** for released versions of all GemTalk products, in PDF form.
- ▶ **Downloads**, including current and recent versions of GemTalk products.

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

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

Help Requests

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

Requests for technical assistance may be submitted online, by email, or by telephone. We recommend you use telephone contact only for more serious requests that require immediate evaluation, such as a production system down. The support website is the preferred way to contact Technical Support.

Website: <http://techsupport.gemtalksystems.com>

Email: techsupport@gemtalksystems.com

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

When submitting a request, please include the following information:

- ▶ Your name and company name.
- ▶ The versions of GemStone/S and of all related GemTalk products, and of any other related products, such as client Smalltalk products.
- ▶ The operating system and version you are using.
- ▶ A description of the problem or request.
- ▶ Exact error message(s) received, if any, including log files if appropriate.

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

24x7 Emergency Technical Support

GemTalk offers, at an additional charge, 24x7 emergency technical support. This support entitles customers to 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.

Chapter 1. GemStone/S 6.6.5 Release Notes

Overview 7

Supported Platforms and GBS Versions 7

 Platforms 7

 GBS version summary 8

 GemConnect 8

Changes in this Release 9

 Login User Id 9

 Process log headers 9

 Login Log file 9

 Programmatic access to all User Ids 9

 Ability to search tranlogs by UNIX user name 10

 Enhancements to login log feature 10

 Login log now includes user information 10

 Login log file now has limited permissions 10

 New feature to specify directory for login log file 10

 Risk of crash if OOP map tracking disabled 10

 Risk of crash when clustering large byte objects 10

 Improved session handling 11

 Remote sessions considered disconnected if socket send timed out 11

 The page server may not be informed of lost OT root 11

 Process exit should be unclean on stuck spin lock 11

 Potential for stuck spin locks from EINTR 11

 Page servers on mid-cache not blocked from cache access 11

 Cleanup of zombie sessions 11

 Remote sessions can hang in partly logged in state 12

Sessions in the process of logging in may delay voting 12
Added method to get age of oldest CR 12
Log message when commit record released without logout 13
commitRestore issues 13
 commitRestore may fail under login race condition. 13
 Incorrect handling of error during commitRestore 13
printlogs.sh and searchlogs.sh used a temporary directory not always present . 13
printlogs.sh filter for session id not working. 13
Excessive Linked login warnings removed. 13
Missing stack trace information on Linux 13
Missing timestamps in some log files. 13
[Debug] log line prefix changed to [Info]. 13
Changes in cache statistics 14
 Recently-added cache statistics incorrect 14
 Statmonitor host system statistics on AIX 14
 Added Cache Statistics 14

GemStone/S 6.6.5

Release Notes

Overview

GemStone/S 6.6.5 is a new version of the GemStone/S object server. This release includes feature enhancements and fixes for a number of bugs.

These release notes provide changes between the previous version of GemStone/S, version 6.6.4, and version 6.6.5. If you are upgrading from a version prior to 6.6.4, review the release notes for each intermediate release to see the full set of changes.

Customers on 32-bit GemStone/S are encouraged to consider migrating to GemStone/S 64 Bit, for additional scalability and reliability as well as many important new and enhanced features.

No separate Installation Guide is provided with this release. For installation instructions, use the Installation Guide for version 6.6.2.

Supported Platforms and GBS Versions

Platforms

GemStone/S version 6.6.5 is supported on the following platforms:

- ▶ Solaris 8, 9 and 10 on SPARC
- ▶ AIX 5.3, 6.1, and 7.1
- ▶ Red Hat Linux ES 5.0, 5.5, and 6.1
- ▶ Windows 7 and Windows 2008 R2

For more information and detailed requirements for each supported platforms, please refer to the GemStone/S v6.6.2 Installation Guide for that platform.

GBS version summary

The following versions of GBS are supported with GemStone/S version 6.6.5, with the following client Smalltalk and platforms versions.

GBS version 7.6.1

VW 7.10.1	VW 7.9.1
<ul style="list-style-type: none"> ▶ Windows 7 and Windows 2008 R2 ▶ Solaris 9 and 10 on SPARC ▶ RedHat Linux ES 5.5 and 6.1 	<ul style="list-style-type: none"> ▶ Windows 7 and Windows 2008 R2 ▶ Solaris 9 and 10 on SPARC ▶ RedHat Linux ES 5.5 and 6.1

GBS version 5.2.7

VA Smalltalk 8.6	VA Smalltalk 8.5.2	VA Smalltalk 8.0.2
<ul style="list-style-type: none"> ▶ Windows 7 and Windows 2008 R2 	<ul style="list-style-type: none"> ▶ Windows XP, Windows 7, and Windows 2008 R2 	<ul style="list-style-type: none"> ▶ Windows XP, Windows 7, and Windows 2008 R2

GemConnect

Released GemConnect shared libraries will not load into GemStone/S version 6.6.5 on Solaris and AIX; shared libraries previously supplied for 6.6.4 will load. On these platforms, if you have not obtained new shared GemConnect libraries, please contact GemTalk Technical Support for recompiled libraries. (#43362, #43363)

Other platforms are not affected.

Changes in this Release

Login User Id

In previous releases, the gem log headers and other tracking reported the Real User Id and the Effective User Id. These distinguish between the Gem process user when the gem was spawned by a netldi running as root.

For audit purposes, it is also useful to track the login user id, which is the original id of the user that logged in. This allows the original user to be determined even when sudo or su are used.

If the system is unable to determine a login user id, then -1 is reported; on Windows, 0 is always reported.

To get the login user id, on linux, `/proc/pid/loginuid` is checked, then `getlogin_r()`. The login process also now sets an environment variable `GS_LOGIN_NAME`, which is used if the above options fail.

Process log headers

The login user id is now included in the log file headers. For example:

```
| USER IDS: REAL=gsadmin (300) EFFECTIVE=gsadmin (300) LOGIN=gsadmin (300)
```

Login Log file

The login user name is now also included in the login log, as the last space-delimited field on each line.

Programmatic access to all User Ids

The following methods have been added to get the real, effective, and login user id for the currently gem process.

System >> `realUserId`

Returns a `SmallInteger` indicating the real user ID of the gem process.

System >> `realUserIdName`

Returns a `String` indicating the real user ID of the gem process.

System class >> `effectiveUserId`

Returns a `SmallInteger` indicating the effective user ID of the gem process.

System class >> `effectiveUserIdName`

Returns a `String` indicating the effective user ID of the gem process.

System class >> `loginUserId`

Returns a `SmallInteger` indicating the login user ID of the gem process. The login user ID is the user ID the owner of the session used to log in to the host system.

Returns -1 if the login user ID could not be determined or if the host system does not support this feature.

System class >> `loginUserIdName`

Returns a `String` indicating the login user ID of the gem process. The login user ID is the user ID the owner of the session used to log in to the host system. Returns 'unknown user' if the login user ID could not be determined or if the host system does not support this feature.

Ability to search tranlogs by UNIX user name

searchlogs.sh and printlogs.sh now accept additional arguments:

```
luid      <integer> - filter by gem's login user ID
luidstr   <string>  - filter by gem's login user name
```

Enhancements to login log feature

The ability to log logins and logouts to a separate text file was introduced in v6.6.4. There are a number of enhancements to this feature in version 6.6.5.

Login log now includes user information

As noted under “Login User Id”, the login log now includes an additional entry for the login user id.

Login log file now has limited permissions

The permission for the login log file are now 600, to reflect its nature as a security audit feature. (#44392)

New feature to specify directory for login log file

The directory for the login log file, by default, is placed in the same directory as the stone log. Now another location can be specified using a new configuration parameter.

```
STN_LOGIN_LOG_DIR
```

Specifies a directory where the login log is located when the STN_LOGIN_LOG_ENABLED option is set to TRUE. Otherwise this option is ignored. If STN_LOGIN_LOG_DIR is not specified and STN_LOGIN_LOG_ENABLED is set to TRUE, then the login log will be placed in the same directory as the stone log.

It is a fatal Stone error if the directory specified is not writable by the Stone process.

Default: NONE Min: 0 entries Max: 1 entry

Risk of crash if OOP map tracking disabled

The environment variable GS_DISABLE_OOP_REF_TRACKING disables OOP map tracking. If this is set, there is a risk of a crash when the bitmaps are initialized or cleared. (#44390)

Risk of crash when clustering large byte objects

The clustering code that handled large objects (objects too large to fit on a single page) was not consistent in the way size was computed when making the allowance for tag fields. In worst case large objects, this could result in a SEGV. (#43722)

Improved session handling

There were a number of issues with handling remote sessions and sessions using mid-level caches, particularly under load. The consequences of encountering these bugs included hung and zombie sessions, commit record issues, and delays in voting.

The following bugs have been fixed that relate to handling remote sessions, and new features added to help in diagnosing and repairing after any remaining issues. Version 6.6.5 also includes enhanced printing for errors and signals, more internal error checking, and code cleanup and improvements.

In addition to these changes, a number of cache statistics have been added to improve tracking of session status; see "Added Cache Statistics" on page 14.

Remote sessions considered disconnected if socket send timed out

On a heavily loaded network, the Stone's send to a remote gem on the OOB socket may time out. The Stone misinterpreted this as a socket disconnect and session termination. Now, after such timeouts the send() is retried; the session is only considered disconnected if the poll() on the socket fails.

Each session has a queue of 64 OOB messages pending due to socket right failures. This limit is not expected to be reachable in the absence of another problem. (#44427)

The page server may not be informed of lost OT root.

When a process gets a lostOTRoot or stopSession, status in the system process are set accordingly to prevent the process for taking further action. This bug allowed the page server to get incorrect status, which could have allowed pages to be read or returned to the cache. This would potentially cause page cache faults or other errors. (#44422)

Process exit should be unclean on stuck spin lock

When a process encounters a stuck spin lock, or an error is detected while waiting on a semaphore, the process should shut down uncleanly to ensure cleanup takes place. (#44431)

Potential for stuck spin locks from EINTR

EINTR is now ignored and the operation retried if it occurs in HostShrSemClear(), to avoid the risk that this may cause a stuck spin lock. (#44432)

Page servers on mid-cache not blocked from cache access

When a remote gem is being terminated or having it's view taken away, the page server's process table slot should be, but was not, being "poisoned" to prevent it from accessing the cache. (#44434)

Cleanup of zombie sessions

Zombie sessions are sessions in which an OS process does not exist, but resources within GemStone are still held, such that the system is impacted. While the occurrence of session in this state is a bug, there still may exist unknown pathways to enter this state. Methods have been added to allow systems to clean up such zombie sessions without restarting the stone.

The following methods have been added:

System class >> allZombieSessions

Returns a list of session IDs that are partially logged out and are in the process of being disposed by the stone. These zombie sessions may or may not reference a commit record.

System class >> allZombieSessionsRefOldestCr

Returns a list of session IDs are partially logged out and are in the process of being disposed by the stone which also reference the oldest commit record.

System class >> stopZombieSession: aSessionId

Forces the stone to finish logout processing for the given session ID. Takes no action if the given session has not begun the logout process or if the session no longer exists. When logout processing is completed, references to a commit record by that session are released and resources consumed by the session are recycled.

Returns true if the action was successful and the stone acted on the specified session. A message will be written to the stone log in this case. Returns false if no action was taken. No message is added to the stone log when false is returned.

WARNING: This method bypasses certain internal mechanisms which prevent the immediate termination of active sessions, and therefore should only be used if the oldest CR session(s) could not be stopped by any other mechanisms.

Remote sessions can hang in partly logged in state

When there are issues with startup of a remote or mid-level cache, it is possible for gem sessions to get partially logged in, but fail to complete login. (#44405)

Timeout have been added to the Stone and Page Manager. The Page Manager's timeout is STN_REMOTE_CACHE_PGSRV_TIMEOUT, on the remote NetLDI to start the cache. On timeout, the Gem will get an error, and a message is printed to the Page Manager log file.

The Stone's timeout is on the waiting sessions and is the greater of STN_REMOTE_CACHE_PGSRV_TIMEOUT or 15 seconds. After this timeout, it is a fatal error for Gem waiting on a remote cache, who are not yet fully logged in, but not for Gems waiting on a mid-level cache, who can retry the operation. In both cases, messages are printed to the Stone log.

Sessions in the process of logging in may delay voting

Sessions that are in the process of logging in, but had not completed login, could have held up voting on possible dead objects and delay garbage collection. This was most likely to be visible for sessions that were waiting for the startup of a remote cache to complete, before they finished the login process. Now, sessions that have not completed login are immune from voting. (#44406)

Added method to get age of oldest CR

The following method has been added:

System class >> oldestCommitRecordAge

Return a SmallInteger representing the age of the oldest commit record in seconds. This value represents the elapsed time since the commit that created the oldest commit record occurred.

Log message when commit record released without logout

When the stone releases a commit record referenced by a gem when that gem has not performed an explicit logout, a message is now printed to the stone log.

commitRestore issues

commitRestore may fail under login race condition

commitRestore requires that the session be the only session logged in, and the method checks for other sessions, and suspends logins, before performing the commitRestore. However, there is a race condition if other sessions are in the process of logging in, particularly in the case of logins that take some time to complete. In this case, the commitRestore checks can be missed, and commitRestore will fail. (#43463)

Incorrect handling of error during commitRestore

There was an code path in which some kinds of error occurring during the commitRestore would not have been treated as a fatal error; the code performed an internal abort. This code path was noted as the result of problems reported when calling `restoreFromBackup*.newSystemUserPassword.`, which appeared to succeed, but did not reset the SystemUser password. (#42354)

printlogs.sh and searchlogs.sh used a temporary directory not always present

The printlogs.sh and searchlogs.sh scripts used `/usr/tmp`, which is not installed by default on some OS distributions. The scripts now use `/tmp`. (#44408)

printlogs.sh filter for session id not working

The printlogs.sh script includes a number of filters, including sessionid. This filter was not functional; any specification for filter on sessionid was ignored. (#44461)

Excessive Linked login warnings removed

The stone log previously logged a warning after each 200 linked GBS logins. This warning is now printed at stone startup, but not otherwise. (#44416)

Missing stack trace information on Linux

Linux C level stack traces did not include important information, including source code line numbers/offsets. GemStone compilation has been adjusted so that this information is available; the executables will be larger as a result. (#44430)

Missing timestamps in some log files

Some entries in the NetLDI, GcGem, and Stone logs did not have timestamps. (#44366)

[Debug] log line prefix changed to [Info]

Statements printed during login previously had the prefix `[Debug]`; this has been changed to `[Info]`. (#44410)

Changes in cache statistics

Recently-added cache statistics incorrect

The following cache statistics have been added in previous releases, since v6.6.2:

```
FreeTempOopCount
FreePersistentOopCount
CommitTokenTimeout
CheckpointDeferTimeout
CheckpointDeferState
RemoteCachePgsvrTimeout
```

These statistics were not ordered correctly when writing the statistics information to the statmonitor file, so that the values shown in VSD do not correspond with the labels. They are reported correctly when using System methods to access cache statistics. (#43536).

Statmonitor host system statistics on AIX

Previously, host statistics were not collected on AIX. Now, host stats are collected; however, note that CPU, TCP, and disks statistics are not available.

Added Cache Statistics

The following cache statistics have been added:

CommitRecordReleases (Stone)

Number of times the stone released a session's reference to a commit record. Does not include releases due to a commit, abort or logout requested by the session. See the LastCommitRecordReleaseReasonCode statistic to determine the cause of the last commit record release.

LastCommitRecordReleaseReasonCode (Stone)

A code indicating the reason of the last commit record release by Stone as follows:

- 1 = released due to System stopZombieSession: method.
- 2 = released due to System terminateAllSessionsReferencingOldestCr method.
- 3 = released while waiting for logout processing to finish.
- 4 = released due to timeout after SigLostOtRoot.
- 5 = released due to session disposal at the completion of logout processing.

LastCommitRecordReleaseSessionId (Stone)

Session ID of the last session for which the Stone released the session's commit record reference. See the LastCommitRecordReleaseReasonCode statistic to determine the cause of the last commit record release.

LastSigLostOtRootSessionId (Stone)

Session ID of the last session to be sent a SigLostOtRoot.

LastSigStopSessionSessionId (Stone)

Session ID of the last session to be sent a SigStopSession.

LastSigTermGemSessionId (Stone)

Session ID of the last gem to be sent a SIGTERM.

LastSigTermPageServerSessionId (Stone)

Session ID of the last page server to be sent a SIGTERM. The session ID of a page server matches the session ID of the page servers client gem.

NumInOobWriteQueue (Stone)

Number of sessions with pending writes on their out of band sockets.

TotalSigAbortsSent (Stone)

Total SigAbort messages sent to all sessions.

TotalSigLostOtRootsSent (Stone)

Total SigLostOtRoots messages sent to all sessions.

TotalSigStopSessionsSent (Stone)

Total SigStopSession messages sent to all sessions.

TotalSigTermsSentToGems (Stone)

Total number of times a SIGTERM signal was sent to a local or remote Gem. SIGTERM signals are sent to local gems by the Stone. SIGTERM signals are sent to remote gems by the cache page server on the remote gem's host.

TotalSigTermsSentToPageServers (Stone)

Total number of times the Stone sent a SIGTERM signal to a page server.