

---

GemStone®

# *GemStone/S* *Release Notes*

Version 6.5.7.1

July 2010

GEMSTONE <sup>TM</sup>

---

## 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 © 1986-2010 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.

**Intel** and **Pentium** are registered trademarks of Intel 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.

**Microsoft**, **MS**, **Windows**, **Windows XP**, **Windows 2000**, **Windows 2003**, **Windows Vista** and **Windows 7** are registered trademarks of Microsoft Corporation in the United States and other countries.

**VisualWorks** is trademark of Cincom Systems, Inc.

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 version 6.5.7.1 release. We recommend that everyone migrating to this version read these release notes before beginning installation, testing or development.

No separate Installation Guide is provided with this release. For instructions on installing GemStone/S version 6.5.7.1, or upgrading or converting from previous products or versions, see the Installation Guide for version 6.5.7.

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

### **Technical Support**

GemStone provides several sources for product information and support. The product-specific manuals provide extensive documentation, and should be your first source of information.

**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:

**Documentation** for GemStone/S is provided in PDF format. This is the same documentation that is included with your GemStone/S 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, and past and current versions of GemBuilder for Smalltalk.

**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.

**Community Links** provide 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.

## Help Requests

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.

**Website:** <http://techsupport.gemstone.com>

**Email:** [support@gemstone.com](mailto:support@gemstone.com)

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

Requests for technical assistance may be submitted online, or 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 6.5.7.1 Release Notes**

Overview . . . . . 9

    Installation . . . . . 9

Supported Platforms . . . . . 10

Changes and New Features . . . . . 10

    Remote cache optimizations and tuning enhancements . . . . . 10

        On Solaris, page manager now uses /dev/poll . . . . . 10

        Compression for lists of pages sent to remote caches for removal . . . 10

        Increased batch size for pages to be removed by page manager . . . . 10

        Warnings on slow responses from remote caches. . . . . 11

        Recycle buffers sent to remote caches . . . . . 11

    Page audit now audits data pages by default . . . . . 11

    Tranlog debug level changeable at runtime . . . . . 12

    Temporary page information written to transaction logs . . . . . 12

        Explicit write via new method. . . . . 12

        Automatic write on page cache fault . . . . . 12

    GcGem log files not deleted by default . . . . . 12

    Additional timestamps for messages in page manager log . . . . . 12

    Added Cache Statistics. . . . . 13

        Additional process statistics on Solaris. . . . . 13

Bugs Fixed . . . . . 13

    Performance problems in page manager poll() with remote caches . . . . . 13

    Race condition in page manager handling of dead remote caches . . . . . 14

    Remote cache timeouts counted sequentially, resulting in delayed timeout. . . 14

    Mid-level cache issues . . . . . 14

        Authorization failures with mid-level caches . . . . . 14

        With mid-level cache, cannot detect problem with leaf cache . . . . . 14

        Deadlocks during shutdown with mid-level cache . . . . . 14

Code changes to reduce risk of corruption . . . . . 15  
Stone shutdown if page manager was not session 1. . . . . 15  
Possible problem on tranlog full conditions . . . . . 15  
Potential memory leak in stone . . . . . 15  
VSD on Linux display problem for Y axis . . . . . 15



# *GemStone/S 6.5.7.1*

## *Release Notes*

### **Overview**

GemStone/S 6.5.7.1 is a new version of the GemStone Smalltalk object server. This special build release provides new features, improved performance, and fixes a number of bugs, including issues relating to remote and mid-level caches.

These release notes provide changes between the previous version of GemStone/S, version 6.5.7, and version 6.5.7.1. If you are upgrading from a version prior to 6.5.7, please also review the release notes for each intermediate release to see the full set of changes.

### **Installation**

No separate Installation Guide is provided with this release. For installation instructions, use the Installation Guide for version 6.5.7.

While we recommend running the `upgradeImage` step of the upgrade process, if you wish to omit this step, you should manually rebuild the `ConfigurationParametersDictionary`. This rebuild is done normally by `upgradeImage`. The new runtime configuration parameter options in this release will not be accessible without this rebuild.

To manually rebuild the structure, execute the following as `SystemUser`:

```
run
  Globals at: #ConfigurationParameterDict put:
    System _configFileParameterDict.
  ^System commitTransaction
%
```

In addition, the method `System Class >> writePageDebugDataToTranlog` will not be available if `upgradeImage` is not run.

## Supported Platforms

GemStone/S version 6.5.7.1 is supported on the following platforms:

- ▶ Solaris 8, 9 and 10 on SPARC
- ▶ Red Hat Linux ES 5.0

For more information and detailed requirements for each supported platforms, please refer to the GemStone/S v6.5.7 Installation Guide for that platform.

## Changes and New Features

### Remote cache optimizations and tuning enhancements

#### On Solaris, page manager now uses /dev/poll

/dev/poll is an optimization in Solaris over poll(), providing much improved performance as the number of connections gets larger. The page manager has been modified to use /dev/poll rather than poll() for polling sockets to remote cache page servers.

#### Compression for lists of pages sent to remote caches for removal

The page manager sends lists of pages to be removed to each remote shared page cache. A new configuration parameter has been added to allow this list of pages to be compressed, reducing network traffic.

##### **STN\_PAGE\_MGR\_COMPRESSION\_ENABLED**

Determines if the page manager will compress the list of pages it sends to remote shared page caches for removal. If set to TRUE, all lists of pages larger than 50 will be compressed before transmission. The same compressed list is used to send to all remote shared page caches; i.e., the compression operation is performed no more than once for each list of pages to be sent. Has no effect on systems which do not use remote shared page caches.

Runtime equivalent: #StnPageMgrCompressionEnabled

Default: FALSE

#### Increased batch size for pages to be removed by page manager

The page manager requests pages from the stone to be processed in batches of between 1 and 16384 pages. The existing configuration parameter STN\_PAGE\_REMOVAL\_THRESHOLD controls the minimum batch size; the page manager waits until this threshold is reached before requesting pages. As this value gets larger, the page manager will make less frequent requests for pages from the stone, and process a larger number of pages when it does.

In version 6.5.7.1, the maximum limit for STN\_PAGE\_REMOVAL\_THRESHOLD has been increased from 1792 to 16384.

## Warnings on slow responses from remote caches

When the page sends a request to a remote cache, after the time period defined by `STN_REMOTE_CACHE_PGSRV_TIMEOUT` is exceeded, the remote cache is considered dead and all gems on that remote cache are terminated.

A new configuration parameter has been added, to log warnings when the responses from the remote cache are taking a significant amount of time, but not as long as the remote cache timeout. This allows a long remote cache timeout to be configured, but still monitor the response times of remote caches.

### **STN\_PAGE\_MGR\_PRINT\_TIMEOUT\_THRESHOLD**

The threshold in real seconds used by the page manager to determine if a slow response from a remote shared page cache should be printed to the page manager log file. If a remote cache takes longer than this number of seconds to respond to the page manager, the page manager will print a message to the log file. If a remote cache takes less than this number of seconds to respond, no message is printed.

Note that this value controls the writing of log messages only. The connection to the remote cache will not be terminated by page manager unless `STN_REMOTE_CACHE_PGSRV_TIMEOUT` is exceeded.

Runtime equivalent: `#StnPageMgrPrintTimeoutThreshold`

Default: 5

Minimum: 0

Maximum: 3600

## Recycle buffers sent to remote caches

Rather than build a new, similar buffer to be sent to each remote caches, now the buffer is updated before sending to each cache. This is especially significant if compression is enabled (`GEM_PGSRV_COMPRESS_PAGE_TRANSFERS`).

## Page audit now audits data pages by default

The utility `pageaudit` previously only audited non-data pages - object table pages, bitmap pages, etc. Data pages were audited separately by `objectAudit`.

Now, `pageaudit` also checks data pages, reading them into memory and confirming that they are data pages. Data page audit is the new default. To disable audit of data pages and revert to the `pageaudit` behavior of previous releases, use the new `pageaudit -d` option.

```
Usage: pageaudit [-d] [-h] [-f] [-e execfg] [-z syscfg] [name]
  -d    disables auditing of data pages.
  -e    specifies executable specific configuration file.
  -f    keeps running beyond first error if possible
  -h    prints usage information and exits.
  -z    specifies system configuration file.
  name  the name of the stone (default is gemserver65-audit).
```

## Tranlog debug level changeable at runtime

The configuration parameter #StnTranLogDebugLevel has been added, to allow STN\_TRAN\_LOG\_DEBUG\_LEVEL to be modified at runtime. Previously, the tranlog debug level could only be modified when restarting the Stone.

With a higher STN\_TRAN\_LOG\_DEBUG\_LEVEL, more detailed debug and analysis information is written to the transaction logs. This parameter should only be used under the direction of GemStone Technical Support. Note that the transaction logs may be much larger with higher debug levels.

The value of #StnTranLogDebugLevel can only be changed by SystemUser. It takes effect for each gem on its transaction boundary (commit or abort). The change in value is logged in the Stone log and in the transaction logs.

To set to a higher value, so that much more data is written to the transaction logs, execute:

```
System configurationAt: #StnTranLogDebugLevel put: N
```

where N is 1 or above; most commonly 3. To reset to normal, execute:

```
System configurationAt: #StnTranLogDebugLevel put: 0
```

## Temporary page information written to transaction logs

### Explicit write via new method

A method has been added to write temporary page information to the transaction log:

```
System >> writePageDebugDataToTranlog
```

This method can only be run by SystemUser. The list of temporary pages allocated to each gem on the system is written to the transaction log. On a busy system, execution of this method may cause the Stone to become unresponsive for several seconds.

This method should only be run if requested by GemStone Technical Support.

### Automatic write on page cache fault

When a page cache fault error is encountered, the list of all temporary pages assigned to each gem is written to the transaction logs. To avoid performance impact when a series of page cache faults occurs, this is done on for the first page cache fault; subsequent page cache faults do not trigger this.

## GcGem log files not deleted by default

Admin and Reclaim GcGem logs previously were deleted by default when the process exited. Now, these logs are not deleted.

## Additional timestamps for messages in page manager log

Some messages were written to the Page Manager log without timestamps; this has been corrected.

## Added Cache Statistics

The following cache statics have been added:

**NumProcsSleepingForLock** (SPC Monitor)

Number of processes on this shared page cache which are currently sleeping while waiting to acquire a spin lock.

**PageMgrPrintTimeoutThreshold** (Stone)

Current value of the STN\_PAGE\_MGR\_PRINT\_TIMEOUT\_THRESHOLD configuration parameter.

**PageMgrCompressionEnabled** (Stone)

Current value of the STN\_PAGE\_MGR\_COMPRESSION\_ENABLED configuration parameter.

## Additional process statistics on Solaris

These per process statistics are now collected, on Solaris only, for each process monitored by statmonitor:

<b>AllOtherSleepTime</b>	<b>OutputBlocks</b>
<b>CharacterIO</b>	<b>PercentCpuUsed</b>
<b>DataFaultSleepTime</b>	<b>PercentMemoryUsed</b>
<b>ImageKBytes</b>	<b>RSSKBytes</b>
<b>InputBlocks</b>	<b>SignalsReceived</b>
<b>IVolCSW</b>	<b>StoppedTime</b>
<b>KernelFaultSleepTime</b>	<b>SwapCount</b>
<b>LockWaitSleepTime</b>	<b>SystemCalls</b>
<b>LwpCurCount</b>	<b>SysTime</b>
<b>LwpTotalCount</b>	<b>TextFaultSleepTime</b>
<b>MajFlt</b>	<b>TrapTime</b>
<b>MinorFaults</b>	<b>UserTime</b>
<b>MsgRecv</b>	<b>VolCSW</b>
<b>MsgSent</b>	<b>WaitCpuTime</b>

## Bugs Fixed

### Performance problems in page manager poll() with remote caches

When the page manager read responses from remote cache page servers, it performed non-blocking I/O, and therefore called code to wait and poll for response. On a heavily loaded system with many remote caches, this resulted in very high CPU use by the page manager, and high values for the cache statistic TimeInNetPgsvrReads.

Since the page manager has already been informed that a response is available, there is no need to wait or poll. Now, the page manager uses blocking I/O to read the response.

In addition, the page manager no longer calls poll() when conversing with the Stone.

(#40733)

## Race condition in page manager handling of dead remote caches

The sequence of recycling pages and the detection of a dead remote cache was not correct, allowing the pages returned by gems on a remote cache that had timed out to potentially be reused before the gems were terminated. This resulted in cache coherency errors in gems on the dead remote cache. (#40745)

## Remote cache timeouts counted sequentially, resulting in delayed timeout

The `STN_REMOTE_CACHE_PGSRV_TIMEOUT` configuration parameter controls how long the page manager waits before determining that a remote cache is unresponsive. These timeouts were not handled correctly with multiple remote caches, resulting in potentially delayed handling of unresponsive remote caches.

When the page manager was waiting for responses from multiple page servers, instead of counting the timeout from the time the message was sent to a specific remote pgsvr, the timeout was started at the time a response was received from any remote pgsvr. This resulted in potentially nearly cumulative timeouts.

In addition to simultaneous timeouts, when the remote cache response is received after more than 50% of the remote cache timeout period has elapsed, a warning is printed to the page manager log file. For example,

```
--- 06/21/10 13:15:43.345 PDT :[Warning]: Slow response from page server on host  
bebop.gemstone.com with process ID 12975: 13849 ms.
```

(#40726)

## Mid-level cache issues

### Authorization failures with mid-level caches

When the system is configured with `netldi` running as root, which requires authentication to start processes, attempts to start a mid-level cache failed authentication. (#40718)

### With mid-level cache, cannot detect problem with leaf cache

When a mid-level cache was used, the page manager sent page removals to the mid-level cache, rather than to the remote leaf cache (the one directly used by the gem). The mid-level cache pgsvr then handled communication to the leaf cache. If the leaf cache timed out or became unresponsive, this case was not handled correctly, resulting in the stone treating the mid-level cache as unresponsive, or cache coherency errors in the mid-level cache.

Now, the page manager sends page removals directly to the leaf caches as well as to the mid-level caches, so regular remote cache timeout handling logic is used in all cases.

(#40742)

### Deadlocks during shutdown with mid-level cache

When mid-level caches were used, there were cases in which a page server could deadlock during shutdown.

## Code changes to reduce risk of corruption

Several holes in the code have been fixed, which could have caused or been a factor in causing repository corruption. (#40686)

- ♦ Dirty pages may have been written out by a session other than the owner, if the session id was reused quickly.
- ♦ Page writes by remote gems done asynchronously, with possible windows when page was not in either remote or stone's cache.
- ♦ Globally dirty pages were written out by AIO page server if free frame limit below 1%.

## Stone shutdown if page manager was not session 1

Under some circumstances, the page manager may not be session 1. This resulted in any further logins causing the stone to shutdown with a UTL\_GUARANTEE error. (#40753)

## Possible problem on tranlog full conditions

When the tranlog run out of disk space, there is a narrow timing window in which a gem's transaction state may become incorrect. (#40741)

## Potential memory leak in stone

Code existed for increasing the size of an empty buffer that could have resulted in a memory leak. While it is not expected that the specific circumstances that would cause a leak can occur, it has been fixed. (#40731).

## VSD on Linux display problem for Y axis

On the VSD provided as part of the 6.5.7 release on Linux, the y-axis of graphs was not displayed correctly. This was due to compiler optimizations in Linux on RH 5.0. (#40752)