
GemStone®

GemStone/S 64 Bit™
Release Notes

Limited Distribution Special Release

Version 2.2.6.1

vmware®

GEMSTONE STM 64

INTELLECTUAL PROPERTY OWNERSHIP

This documentation is furnished for informational use only and is subject to change without notice. VMware, 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 VMware, 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 VMware, 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 VMware, Inc.

This software is provided by VMware, 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 VMware, 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 © 1986-2012 VMware, Inc., and GemStone Systems, Inc. All rights reserved by VMware, Inc.

PATENTS

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

TRADEMARKS

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

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

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

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

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

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

Microsoft, **MS**, **Windows**, **Windows XP**, **Windows 2003**, **Windows 7** and **Windows Vista** 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. VMware 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.

VMware, Inc.
15220 NW Greenbrier Parkway
Suite 150
Beaverton, OR 97006

About This Documentation

These release notes describe changes in the GemStone/S 64 Bit™ version 2.2.6.1 limited distribution special 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 64 Bit version 2.2.6.1, or upgrading or converting from previous products or versions, see the Installation Guide for version 2.2.5.2.

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; the GemStone Smalltalk programming language; and may also be used to refer to the company, previously GemStone Systems, Inc., now a division of VMware, Inc.

Technical Support

GemStone Website

<http://support.gemstone.com>

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

- ▶ **Documentation** for released versions of all GemStone products, in PDF form.
- ▶ **Downloads** and **Patches**, including past and current versions of GemBuilder for Smalltalk.

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

Email: techsupport@gemstone.com

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

When submitting a request, please include the following information:

- ▶ Your name, company name, and GemStone server license number.
- ▶ The versions of all related GemStone 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 GemStone holidays.

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, for issues impacting a production system. For more details, contact your GemStone account manager.

Training and Consulting

Consulting is available to help you succeed with GemStone products. Training for GemStone software is available at your location, and training courses are offered periodically at our offices in Beaverton, Oregon. Contact your GemStone account representative for more details or to obtain consulting services.

Chapter 1. GemStone/S 64 Bit 2.2.6.1 Release Notes (Limited Distribution Special Release)

<i>Overview</i>	7
<i>Supported Platforms</i>	7
Platforms	7
<i>Changes in this release</i>	8
Added support for Solaris on x86	8
Internal Performance Improvements.	8
Changes to improve commit record disposal performance	8
Do not check private page cache when disposing persistent pages.	8
Modify use of STN_COMMIT_QUEUE_THRESHOLD to consider size of other queues.	8
Stone defers service of page manager	8
Support for 64K memory pages.	9
Support for 16M Memory Pages for SPC on AIX	9
Determining the required number of large memory pages	9
Configure AIX to use the required number of Large Memory Pages.	10
Enable Unix user to use Large Memory Pages	10
Enabling Large Memory Pages in GemStone/64	10
GemStone Shared Memory Debugging	11
Distribution Changes	11
Distributing static GCI RPC library.	11
checksum, woman, and snefru removed from distribution	11
Changes in cache statistics.	11
Added Cache Statistics	11
Changed Cache Statistic	11

Configuration Parameter Changes 12
 Added Configuration Parameter 12
 Changed Configuration Parameter 12

Bugs Fixed **13**

Backup with Reclaim can result in corruption 13
Unsafe aborts by Reclaim gems 13
Configuration file setting for STN_COMMIT_QUEUE_THRESHOLD not used 13
Performance degradation when number of extents increases 13
Stone preempted by commit record backlog disposal 13
Gem could have core dumped on page server death 13

GemStone/S 64 Bit

2.2.6.1 Release Notes

(Limited Distribution Special Release)

Overview

GemStone/S 64 Bit 2.2.6.1 is a new limited distribution special release version of the GemStone/S 64 Bit object server. This release contains a number of internal performance improvements, adds new features and support for Solaris on x86, and fixes a number of bugs.

These release notes provide changes between the previous version of GemStone/S 64 Bit, version 2.2.6, and version 2.2.6.1. If you are upgrading from a version prior to 2.2.6, please also review the release notes for each intermediate release to see the full set of changes.

No separate Installation Guide is provided with this release. For installation instructions, use the Installation Guide for version 2.2.5.2.

Supported Platforms

Platforms

GemStone/S 64 Bit version 2.2.6.1 is supported on the following platforms:

- ▶ AIX 5.3, AIX 6.1, and AIX 7.1
- ▶ Solaris 10 on x86

For more information and detailed requirements for each supported platform, please refer to the GemStone/S 64 Bit v2.2.5.2 Installation Guide for that platform.

Changes in this release

Added support for Solaris on x86

Support for Solaris on x86 has been added. Previous 2.2.x releases supported Solaris on SPARC, but not Solaris on x86.

The Installation Guide v2.2.5.2 for Solaris provides OS configuration information that applies to Solaris on both SPARC and x86. Specific information for compiling and debugging varies between Solaris on SPARC and on x86; for details for Solaris on x86, please contact GemStone Technical Support.

Internal Performance Improvements

Internal changes have improved performance in a number of areas. This includes a new page ID hash algorithm; see “Performance degradation when number of extents increases” on page 13. There have also been improvements in the page cache entry (PCE) search and copy operations.

Changes to improve commit record disposal performance

With large commit record disposal load, the system-wide performance can degrade significantly; see “Performance degradation when number of extents increases” on page 13. A number of changes have been made to improve this.

Do not check private page cache when disposing persistent pages

When disposing a page, the private page cache only needs to be checked for temporary pages.

Modify use of STN_COMMIT_QUEUE_THRESHOLD to consider size of other queues

Previously, the STN_COMMIT_QUEUE_THRESHOLD setting would defer commit record disposal based on the number of sessions in the commit queue. Now, it additionally considers sessions in the run queue and in the SMC queue, and will defer commit record disposal while there are a greater number of sessions on any of these queues. For the updated description of this configuration parameter, see “STN_COMMIT_QUEUE_THRESHOLD” on page 12.

Stone defers service of page manager

Previously, the page manager serviced tasks from the page manager in the same queue as other sessions, so application tasks could be deferred while the stone was working on page disposal for the page manager.

Now, the stone will not service the page manager while the stone has work to do, if a session is holding the commit token, or if there are sessions on the run queue or SMC queue.

If free space becomes low, however, the stone will immediately service the page manager and dispose pages.

To avoid a busy stone not attending to the page manager for an extended period, a timeout has been added, which can be set using the added configuration parameter

STN_PAGE_MGR_MAX_WAIT_TIME. This timeout specifies a maximum time after the page manager has requested service that the stone will service the page manager. For details, see “STN_PAGE_MGR_MAX_WAIT_TIME”, below.

Support for 64K memory pages

GemStone/S 64 Bit now uses 64K page size for temporary object memory on all platforms.

The default size for memory pages on AIX is 4K (small). Current releases of AIX also support 64K (medium), 16MB (large) and 16GB (very large) memory pages. AIX automatically adjusts between 4K and 64K pages, but using the large memory page sizes require special configuration. In addition to using 64K page sizes for temporary object memory, this release provides support for configuring GemStone to use 16MB pages for the Shared Page Cache only, as described in the next section.

Support for 16M Memory Pages for SPC on AIX

GemStone now supports using Large Memory Pages on AIX for the shared page cache. Large memory pages are 16MB. While default 4KB pages and medium 64KB page sizes are handled automatically, AIX requires special configuration to use large pages.

To use large memory pages on AIX for the shared page cache, you must do the following:

1. Determine how much space you require for the SPC and configure AIX to allocate at least this many large memory pages. Before AIX can allocate large memory pages, the AIX kernel must be rebuilt.
2. The UNIX user running the shared page cache monitor must be configured with permission to use large memory pages.
3. You must enable GemStone/S 64 Bit to request large memory pages by setting an environment variable prior to startup.

Determining the required number of large memory pages

You will first need to determine how much space is needed for the shared page cache. This is a two-step process.

(1) The easiest way to get the required values is to start up GemStone with the desired configuration, and inspect the shared page cache monitor log. These example numbers are for a system with a 1TB (1000GB) cache. These are the values that are needed:

```
Number of pages           62500000
Number of processes       3066
TargetFreeFrameCount      5000
Number of shared counters 1900
```

(2) Once you have the GemStone parameters, start the shrpcmonitor directly and enter the parameters to compute the target size.

```
unix>$GEMSTONE/sys/shrpcmonitor
<prints header information>
SHRPCMON>62500000 setnumpages
SHRPCMON>3066 setnumprocs
SHRPCMON>1900 setnumsharedctrs
SHRPCMON>5000 settargetfflimit
```

```
SHRPCMON>getrequiredsize
For 62500000 pages, 3066 processes and 1900 shared counters,
required cache size is 1076836515840 bytes.
```

In this example, the result of the computation is that the shared page cache will require 1076836515840 bytes.

Large memory pages are 16 MB, so divide 1076836515840 by the page size (16777216), rounding up. This gives a minimum of 64185 large memory pages that are required for a SPC with this configuration.

Configure AIX to use the required number of Large Memory Pages

AIX by default is not configured to allocate large memory pages. Before large memory pages can be used, AIX must be configured to use a specific number of large memory pages. These commands must be run as root.

the following example command configures AIX with 64185 large memory pages:

```
#vmo -r -o lpgg_regions=64185 -o lpgg_size=16777216
#vmo -p -o v_pinshm=1
```

After this has been executed, the bosboot command must be run to build a new kernel image, and the system must be rebooted

Enable Unix user to use Large Memory Pages

The UNIX user running the shared cache monitor must also be given permission to use large memory pages.

```
#chuser capabilities=CAP_BYPASS_RAC_VMM,CAP_PROPAGATE <user id>
```

Alternatively, you can run the SPC monitor process with an effective user ID of root:

```
#chown root $GEMSTONE/sys/shrpcmonitor
$GEMSTONE/sys/startshrpcmon
#chmod u+s $GEMSTONE/sys/shrpcmonitor
$GEMSTONE/sys/startshrpcmon
```

you can confirm large memory pages are available for use using `vmstat -l`

Enabling Large Memory Pages in GemStone/64

GemStone will attempt to allocate the shared page cache using large memory pages based on the value of the added environment variable `GS_SPC_USE_LARGE_PAGES`.

- ▶ If this is not set, large pages are not used.
- ▶ if this is set to 1, it is advisory mode: Large memory pages are requested but the cache still starts if the request is not granted by the system.
- ▶ If this is set to 2, it is mandatory mode. Large memory pages are requested and the cache fails to start if the request is not granted by the system.

When you start GemStone, the shared page cache monitor log will use a line of the form:

```
[Info]: Shared page cache was successfully created using large
memory pages.
```

GemStone Shared Memory Debugging

The SPC monitor process will print extra debugging information to its log file if the following environment variable is set:

```
export GS_DEBUG_SHARED_MEM=1
```

Distribution Changes

Distributing static GCI RPC library

The product distribution now includes a static library suitable for linking with custom GCI applications. This file is: `$GEMSTONE/lib/libgcirpc64-226.a`.

For the Solaris on x86 distribution only, a version of the RPC library that was compiled with gcc v4.3.4 is also included in the distribution. This file is: `$GEMSTONE/lib/libgcirpc64-226-gcc.a`.

checksum, woman, and snefru removed from distribution

The checksum, woman, and snefru executables are no longer included in the distribution.

Changes in cache statistics

Added Cache Statistics

HighWaterPageExtentId (Stone)

The extent ID of the current high water page

HighWaterPageRecordId (Stone)

The record ID of the current high water page

PageManagerMaxWaitTimeMs (Stone)

The current value of the `STN_PAGE_MGR_MAX_WAIT_TIME` configuration option.

PageManagerStarvedCount (Stone)

Number of times the page manager session waited more than `STN_PAGE_MGR_MAX_WAIT_TIME` milliseconds for service from the stone.

TimeInGetPagesForPageMgr (Stone)

Total amount of time in milliseconds the stone spent processing requests from page manager for pages to remove from the cache.

TimeInProcessPagesFromPageMgr (Stone)

Total amount of time in milliseconds the stone spent processing pages removed from the cache by the page manager session.

Changed Cache Statistic

The existing cache statistic **TimeInStonePageDisposal** previously included the time spend in all page disposal operations; now, this statistic does not include the times tracked in **TimeInGetPagesForPageMgr** and **TimeInProcessPagesFromPageMgr**

Configuration Parameter Changes

Added Configuration Parameter

STN_PAGE_MGR_MAX_WAIT_TIME

Maximum time the stone will defer servicing the page manager session because the stone is busy with other tasks. Normally the stone services the page manager session whenever it has idle time and no session is performing a commit. If the time the page manager has been waiting for service exceeds this value, the stone will service the page manager unconditionally and increment the cache statistic PageManagerStarvedCount.

Runtime equivalent: #StnPageMgrMaxWaitTime

Units: Milliseconds

Default: 200

Minimum: 1

Maximum: 1000

Changed Configuration Parameter

The configuration parameter STN_COMMIT_QUEUE_THRESHOLD has been changed; now, rather than just considering the commit queue, the run queue and SMC queue are also considered. The new description of this parameter is:

STN_COMMIT_QUEUE_THRESHOLD

Determines if the stone defers disposing commit records based on the number of sessions in the commit queue, the run queue, and the SMC queue. If the size any of these queues is greater than this value, the stone will defer commit record disposal until all queues have sizes less than or equal to the value.

This setting will be ignored if the commit record backlog exceeds the value of STN_CR_BACKLOG_THRESHOLD.

Runtime equivalent: #StnCommitQueueThreshold

Default: -1 (never defer CR disposal)

Min: -1

Max: 1024

Bugs Fixed

Backup with Reclaim can result in corruption

Under certain conditions, the Reclaim Gem may ignore the high water page sent by the stone, or not refreshing the high water page correctly. This can cause problems such as corrupted backups. (#42461, #42465)

Unsafe aborts by Reclaim gems

If the reclaim gem must pause during a transaction, it was aborting but not discarding the list of pages to reclaim, which was not safe. (#42470)

Configuration file setting for STN_COMMIT_QUEUE_THRESHOLD not used

The setting for the configuration parameter STN_COMMIT_QUEUE_THRESHOLD as set in the configuration files is not used by the system, although it was being set correctly using the runtime parameter. (#42444, #40255)

Performance degradation when number of extents increases

The page hash function did not scale well, so performance degraded for the same operation with an increasing number of extents. The page hash algorithm has been improved in this release. (#42475)

Stone preempted by commit record backlog disposal

The stone CPU usage may go to 100% with a very large commit record disposal, which caused excessive latency for other sessions. This is fixed as described by the changes under "Changes to improve commit record disposal performance" on page 8. (#42459)

Gem could have core dumped on page server death

If a page server was killed, a linked gem could have core dumped, rather than died normally. (#39723)