
GemStone®

GemStone/S 64 Bit™ **Release Notes**

Version 3.2.6

April 2015



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 © 1986-2015 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.

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

Microsoft, **MS**, **Windows**, **Windows 7**, **Windows 2008**, and **Windows 8** are registered trademarks of Microsoft Corporation in the United States and other countries.

Linux is a registered trademark of Linus Torvalds and others.

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

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

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

AIX, **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 64 Bit™ version 3.2.6 release. Read these release notes carefully before you begin installation, conversion testing, or development with this release.

For information on installing or upgrading to this version of GemStone/S 64 Bit, please refer to the *GemStone/S 64 Bit Installation Guide* for version 3.2.6.

These documents are available on the GemTalk 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

Support Website

<http://gemtalksystems.com>

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

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



Table of Contents

Chapter 1. GemStone/S 64 Bit 3.2.6 Release Notes

Overview	7
Supported Platforms	7
Platforms for Version 3.2.6	7
GBS Versions	8
VSD Version.	8
Changes and New Features	9
Updated SSL libraries	9
GsProcess isForked method added.	9
startcachewarmer for remote cache may now use mid-level cache.	9
startnetldi added option to configure socket listening backlog	9
Change in handling of Linux out of memory kill protection	10
Change in handling of read authorization errors	10
Added configuration parameter	10
Duplicate cache names now allowed.	11
Added method physicalSizeOnDisk	11
Reflection now includes tagSizeOf:.	11
Added Gci Function	11
Bugs Fixed	12
Risk of repository corruption if large OopNumberHighWaterMark growth during MFC	12
Change in printRecursiveRepresentationOn: to avoid risk of recursion	12
Hang in SSL code during login	12
Mid-level cache related issues	12
Death of mid-cache caused gems to terminate	12
Shared Page Cache Monitor death not handled properly	12
Page server's death not handled properly	12
Page Servers on mid-level cache not using Gem's #log for log location	12
GsSocket getHostNameByAddress: risk of hang	13

Unresolved symbol in non-upgrade recompile of GemStone kernel methods . . .	13
sigAborts may be lost to GBS or GCI applications	13
Risk of SEGV from thread-unsafe code	13
Gem error when commit on tranlog full	13
GCI Client Socket not closed on Gem abnormal shutdown	13
Stone crash with many remote caches	13
Reclaim Issues	14
Attempt to reclaim a page more than once	14
Slow reclaim with idle/unused extents	14
Reclaim Gem log logging of session count off by one	14
Improvements that address slow reclaim performance issues	14
Connection refused errors on NetLDI connect backlog	14
Cache warmer main thread detaches cache unclearly	14
System currentSessions result may have included nils	14
Socket disconnect may result in stuck spin lock	15
Errors on restoreFromBackup from NFS-mounted drive.	15
Unnecessary atomic operations	15
Changes in FFI interface code	15
Statmonitor and Cache Statistics Changes	16
Inaccuracies in page read cache statistics.	16
SPC Monitor total stats were sums	16
statmonitor -J flag did not collect PageManager statistics	16
Incorrect values for network stats on Linux hosts	17
Memory page cache statistics incorrect on AIX	17
ReclaimGem cache statistic for PinnedPagesCount is incorrect	17
Change in GemStone process statistics on Linux	17
Added GemStone process cache statistics	18
Added Linux System cache statistics	19
Removed cache statistics	20

GemStone/S 64 Bit

3.2.6 Release Notes

Overview

GemStone/S 64 Bit 3.2.6 is a new version of the GemStone/S 64 Bit object server. This release adds new features, new cache statistics, and fixes a number of bugs in v3.2.4; we recommend everyone using or planning to use GemStone/S 64 Bit v3.2.x upgrade to this new version.

These release notes provide changes between the previous version of GemStone/S 64 Bit, version 3.2.4, and version 3.2.6. Versions 3.2.4.1, 3.2.4.2, 3.2.4.3, and 3.2.5 were limited distribution special releases. All changes in these releases are included in these release notes. If you are upgrading from a version prior to 3.2.4, review the release notes for each intermediate release to see the full set of changes.

For details about installing GemStone/S 64 Bit 3.2.6 or upgrading from earlier versions of GemStone/S 64 Bit, see the *GemStone/S 64 Bit Installation Guide* for v3.2.6 for your platform.

Supported Platforms

Platforms for Version 3.2.6

GemStone/S 64 Bit version 3.2.6 is supported on the following platforms:

- ▶ Solaris 10 and 11 on SPARC
- ▶ Solaris 10 on x86
- ▶ AIX 6.1, TL1, SP1, and AIX 7.1
- ▶ Red Hat Linux ES 6.5; Ubuntu 12.04; and SUSE Linux Enterprise 11 Service Pack 3, on x86
- ▶ Mac OSX 10.6.8 (Snow Leopard), with Darwin 10.8.0 kernel, on x86

Note that on Linux, GemStone/S v3.2.6 has been compiled on a later kernel; Red Hat 6.1 and 6.4 are not supported with this version.

For more information and detailed requirements for each supported platforms, please refer to the GemStone/S 64 Bit v3.2.6 Installation Guide for that platform.

GBS Versions

The following versions of GBS are supported with GemStone/S 64 Bit version 3.2.6. You must use GBS version 7.6.1 or later for VisualWorks, or 5.4.2 or later for VA Smalltalk, with GemStone/S 64 Bit v3.2.6.

GBS version 7.6.1

VisualWorks 7.10.1 32-bit	VisualWorks 7.10.1 64-bit	VisualWorks 7.9.1 32-bit
<ul style="list-style-type: none"> ▶ Windows 8, Windows 2008 R2 and Windows 7 ▶ Solaris 10 on SPARC ▶ Ubuntu 12.04, RedHat Linux ES 6.5, and SUSE Linux ES 11 SP3 	<ul style="list-style-type: none"> ▶ Windows 8, Windows 2008 R2 and Windows 7 ▶ Solaris 10 on SPARC ▶ Ubuntu 12.04, RedHat Linux ES 6.5, and SUSE Linux ES 11 SP3 	<ul style="list-style-type: none"> ▶ Windows 2008 R2 and Windows 7 ▶ Solaris 10 on SPARC ▶ SUSE Linux ES 11 SP3

GBS version 5.4.2

VA Smalltalk 8.6	VA Smalltalk 8.5.2
<ul style="list-style-type: none"> ▶ Windows 8, Professional or above ▶ Windows 2008 R2 ▶ Windows 7, Professional or above 	<ul style="list-style-type: none"> ▶ Windows 2008 R2 ▶ Windows 7

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

VSD Version

The GemStone/S 64 Bit v3.2.6 distribution includes VSD version 4.0.2. The previous version of GemStone/S 64 Bit, v3.2.4, included VSD v4.0.

Changes between v4.0 and v4.0.2 include:

- ▶ Ability to select a background color to differentiate executions
- ▶ Ability to copy a statistics value from a graph
- ▶ F12 is now hotkey
- ▶ Per-second calculation of statistics values were incorrectly 1000x high
- ▶ Double-click load in VSD X-emulated truncated contents
- ▶ Corrected handling of statistics values that exceed 2^{32}

For more details, see the *Release Notes* for VSD v4.0.1 and v4.0.2

Changes and New Features

Updated SSL libraries

The version of OpenSSL used by GemStone/S 64 Bit v3.2.6 has been updated to 1.0.2a.

GsProcess isForked method added

The instance method `GsProcess>>isForked` has been added.

startcachewarmer for remote cache may now use mid-level cache

When the startcachewarmer script is used to warm a remote cache, it can now be configured to use or create a mid-level cache. With this configured, the cachewarmer will either load pages into the remote cache using the mid-level cache, or warm the mid-level cache as it loads pages in the remote cache.

The following options have been added to startcachewarmer:

- M host name or IP address where the mid-level cache is running or will be created. The -H option (specifying the host name or IP of the Stone's host) must also be specified with this option.
- C size of the mid-level cache in KB. If omitted, a value of 75000 is used. Only applies if the -M option is also specified and the mid-level cache does not exist.
- N The maximum number of processes that can use the mid-level cache. If omitted, a value of 50 is used. Only applies if the -M option is also specified and mid-level cache does not exist.

If a mid-level cache host name or IP address is specified (via -M), the mid-level cache will be created if it does not already exist. The -C and -N options will be used to specify the size and number of processes that can attach the mid-cache respectively. If the mid-cache already exists, the -C and -N options are ignored.

startnetldi added option to configure socket listening backlog

The default socket backlog for NetLDI, SPC Monitor, and Stone has been increased to 64.

To allow netldi to queue a configurable number of login requests, an option has been added to startnetldi. This addresses bug #45008, "Connection refused errors on NetLDI connect backlog" on page 14.

The netldi has a new option, -b, to specify the maximum backlog on the listening socket.

Usage: startnetldi [-b backlog] [-h] [-d] [-g|-s] [-n]
[-a account] [-l logFile] [-t seconds] [-P portNumber]
[-A address] [name]

Note that if a value passed in with the -b argument is larger than the OS configuration allows, as on Linux per `/proc/sys/net/core/somaxconn`, it will be truncated to that limit.

Change in handling of Linux out of memory kill protection

When a Linux system runs low in memory, each process's `oom_score_adj` setting is used to determine which processes are killed first. Applications can adjust this value; if the Unix user has the `CAP_SYS_RESOURCE` capacity, the `oom_score_adj` can be set to a lower value, otherwise it can only be increased.

The previous behavior had shortcomings:

- ▶ For users without `CAP_SYS_RESOURCE`, the `oom_score_adj` was increased to +250 for Gem and Topaz; however, it did not adjust page servers. Page servers for remote sessions and Free Frame Page servers are now treated the same as Gem and Topaz sessions.
- ▶ For users with `CAP_SYS_RESOURCE`, the `oom_score_adj` was decreased to -500 for most processes, including user Gems and Topaz. This is inappropriate, since killing Gem sessions would be preferably to killing critical non-GemStone processes. Now, Gems, Topaz, and remote and Free Frame Page servers are set to +250.

Change in handling of read authorization errors

In GemStone/S 64 Bit, read authorization checks occur when the object is faulted into the VM, rather than when the actual read occurs. This changes the timing of the error relative to 32-bit GemStone/S, and creates conditions under which you could get a `SecurityError` even when the operations you were performing should not trigger that error. (#45040)

A variation occurred under some conditions of updating an `RcKeyValueDictionary`. (#45054).

This release includes changes and a new configuration parameter to change the way unauthorized objects are handled. By default, there is no change in behavior from that in previous GemStone/S 64 versions. A new configuration parameter, `GEM_READ_AUTH_ERR_STUBS`, has been added. This defaults to `FALSE`; when this is set to `TRUE`, instead of triggering a `SecurityError`, when a read authorization error is encountered, an instance of the new class `UnauthorizedObjectStub` is created.

These changes require corresponding changes in GBS, and should not be used for GBS sessions, unless directed otherwise by GBS Engineering when running with certain versions of GBS.

Added configuration parameter

The following configuration parameter has been added:

`GEM_READ_AUTH_ERR_STUBS`

If `TRUE`, an in-memory instance of `UnauthorizedObjectStub` is constructed for an object fault instead of signalling a `SecurityError` for read authorization denied.

Runtime equivalent: `#GemReadAuthErrStubs`

Default: `FALSE`

This should remain `FALSE` for GBS sessions, unless directed otherwise by GBS Engineering when running with certain versions of GBS.

Duplicate cache names now allowed

The name of a gem in cache statistics can be set using the method `cacheName:`, and this name is visible in VSD and when using the programmatic interface to cache statistics. In version 3.0, it was disallowed to assign a name to a Gem when that name was already in use, and attempting to do so would raise an error.

While it is more unambiguous to avoid duplicate cache names, this error was inconvenient in practice, and duplicate names are now allowed.

Note that accessing statistics via `System >> cacheStatsForGemWithName:` will return statistics for the first gem with a given name. If your cache naming does not create unique names, use other cache statistics lookup methods that rely on PID or `sessionId`, to ensure that you get predictable results.

Added method `physicalSizeOnDisk`

When determining the amount of physical space required to hold objects on disk, the method `physicalSize`, which returns the space required for the in-memory representation, may overstate the requirement for the object on disk. For more accurate calculation, the following method has been added:

```
Object >> physicalSizeOnDisk
```

Returns the number of bytes required to represent the receiver on disk. If the receiver is in special format (which implies that its representation is the same as its OOP), returns zero.

Reflection now includes `tagSizeOf:`

The method `Reflection>>tagSizeOf:` has been added.

Added Gci Function

The following GCI function has been added:

```
(int64) GciFetchTagSize(  
    OopType obj  
);
```

Returns the number of oops of dynamic `instVars` that are allocated in the object. Returns 0 if `obj` is a special object.

Bugs Fixed

The following bugs in v3.2.4 have been fixed in v3.2.6:

Risk of repository corruption if large `OopNumberHighWaterMark` growth during MFC

When the `OopNumberHighWaterMark` grows by a very large amount (more than 50M objects) during MFC, internal structure resizing was not done correctly. This could result in a SEGV, or depending on the specifics of the occurrence, OOPs may be handled incorrectly and introduce corruption. (#45106)

Change in `printRecursiveRepresentationOn`: to avoid risk of recursion

The method `Object>>printRecursiveRepresentationOn`: called `self asString`, which in some cases could cause unwanted recursion. (#44825)

Hang in SSL code during login

Version 1.0.2 of OpenSSL has better retry logic handling than that in 1.0.1x. In addition, the retry logic in GemStone's SSL login (`GciLogin`) did not correctly handle retries in all cases. (#44927)

Mid-level cache related issues

A number of fixes and changes have been made in v3.2.6 in handling of problems in mid-level caches. Systems using mid-level caches are now more tolerant of the loss of connection to a mid-level cache host, and of problems with individual processes supporting a mid-level cache.

Death of mid-cache caused gems to terminate

If a mid-level cache terminated or the mid-level cache machine became unavailable, all remote sessions using that mid-level cache would encounter a fatal error. Now, these sessions will continue running without the mid-level cache; a message is printed to the Gem's stdout. (#44993)

Shared Page Cache Monitor death not handled properly

Under some conditions, the death of the Shared page cache Monitor on a mid-level cache host may not be handled properly, resulting in errors attempting to connect to the mid-level cache. (#45005).

Page server's death not handled properly

When a session's page server on the mid-level cache dies, the session may encounter an error as the remote session attempts to recover the connection. (#44992).

Page Servers on mid-level cache not using Gem's `#log` for log location

A Gem process's `GEMSTONE_NRS_ALL` may include a setting for `#log`, indicating a directory to which the associated logs should be written. If the Gem uses a mid-level cache, it has a page server process on that machine, and the log for this page server should

be but was not using the Gem's #log setting. Instead, this log was using a #log setting from the mid-level cache machines NetLDI environment, or to the unix user's home directory. (#45004)

GsSocket getHostNameByAddress: risk of hang

If the DNS server cannot resolve an address, the execution of `GsSocket class >> getHostNameByAddress`: may hang, depending on your configuration. Now, it will try five times before reporting an error (#45077)

Unresolved symbol in non-upgrade recompile of GemStone kernel methods

Some obsolete classes, such as `ObsoleteSymbol`, were moved to the "ObsoleteClasses" dictionary within `Globals` as part of the 3.0 upgrade; by leaving the class, existing references would continue to be functional. However, recompile of `GemStone` kernel classes (other than by upgrade) resulted in unresolved symbols. While this is not generally needed or recommended, tools such as `STORE` may have caused recompile. (#44990)

sigAborts may be lost to GBS or GCI applications

Changes in handling of asynchronous events to avoid recursive handling were made in 3.2. There was a case in which this code would suppress a second `sigAbort` to a GCI or GBS applications performing certain GCI executions. (#45067)

Risk of SEGV from thread-unsafe code

It was possible for internal repository I/O time tracking code to execute thread-unsafe, with a risk of a null pointer and SEGV. (#44511)

Gem error when commit on tranlog full

When the transaction logs are full, the stone performs special handling of commits and other operations, pausing until space becomes available. If a Gem was performing a commit, some state was cleared by the code that handled the tranlog full condition. In this case, the Gem would terminate with a invalid stone command error. (#44894)

GCI Client Socket not closed on Gem abnormal shutdown

When the Gem terminated due to loss of its page server, the error was not handled correctly. The timing of the close of the socket to the client was incorrect, which for some cases in GBS logins, could cause the GBS client image to hang. (#45065)

Stone crash with many remote caches

An internal value was computed incorrectly, producing a number usually in the vicinity of 250. If the number of remote caches attached to a system grew larger than this, the Stone crashed with a `UTL_GUARANTEE` error. (#45064)

Reclaim Issues

Attempt to reclaim a page more than once

A operation during reclaim is not thread-safe, so under some conditions, one thread may remove a page from the reclaim pages list, while another thread puts the page back on the list. This results in an attempt to reclaim the same page more than once. The conditions for this bug appear to be rare; however, while this bug may not trigger an immediate error, page errors may occur subsequently. (#45041)

Slow reclaim with idle/unused extents

When a reclaim thread finds no pages in an extent that need reclaim, it sleeps for one second before continuing, even if other extents have a large amount of reclaim. This may cause slow performance in systems configured with sequential allocation mode and that have a large amount of free space. In such system, the data (and therefore pages needing reclaim) may be entirely in the first extents, leaving later extents in the sequence empty. (#45015)

Reclaim Gem log logging of session count off by one

When the number of Reclaim Gem sessions is changed, the new number of Reclaim Gem sessions is logged in both the Stone log and the Reclaim Gem log. The number in the Stone log is correct, but the number reported in the Reclaim Gem log is one lower than the actual number. (#45013).

Improvements that address slow reclaim performance issues

This release includes code changes that may improve handling for certain cases of slow reclaim performance.

Connection refused errors on NetLDI connect backlog

On login, gems connect to the netldi on its listening socket. The backlog for this socket is set at 20, and if the number of login requests is much higher than the netldi can process and the backlog exceeds 20, the login will error with "Connection refused". (#45008)

Now, the default socket backlog for NetLDI, SPC Monitor, and Stone has been increased to 64. The netldi has a new option, `-b`, to specify the maximum backlog on the listening socket. see "startnetldi added option to configure socket listening backlog" on page 9.

Cache warmer main thread detaches cache uncleanly

When the cache warmer completed and detached from the shared cache, the main cache warmer thread's disconnect was not clean, resulting in the need for slot recovery. (#45003)

System currentSessions result may have included nils

When the number of sessions was greater than 2034, the results returned by System currentSessions may have included nils. (#45012)

Socket disconnect may result in stuck spin lock

When a socket to a remote cache disconnects while the page server was holding the free PCE spin lock, it was possible for the page server to exit without releasing the spin lock, leaving it stuck. (#45009)

Errors on restoreFromBackup from NFS-mounted drive

Attempting to restore from a backup that was located on an NFS-mounted drive could error. (#45019)

Unnecessary atomic operations

Some atomic operations in threads were being performed while holding a mutex which precluded other threads from modifying the data structure. This has been cleaned up to improve performance. (#45037)

Changes in FFI interface code

Some adjustments have been made in header parsing code that will allow support for later versions of Mac OS. (#45034)

Statmonitor and Cache Statistics Changes

This release includes a number of bug fixes related to cache statistics, and a number of new statistics.

Inaccuracies in page read cache statistics

The following cache statistics were not updated for reads done by a mid-level cache on behalf of a remote gem:

- BitmapPageReads**
- DataPageReads**
- ObjectTablePageReads**
- OtherPageReads**
- PageIoCount**
- PageIoTimeOverallAvg**
- PageIoTime10SampleAvg**
- PageIoTime100SampleAvg**
- PageReads**

The stone, logsender, and logreceiver processes incremented these statistics by the number of read operations, rather than by the number of pages read. (#45006)

SPC Monitor total stats were sums

Previously, some SPC Monitor statistics were calculated as the sum of the corresponding processes for all active processes in the cache. When a session logged out, the SPC Monitor statistics value could drop, counter-intuitively.

Now, the following SPC Monitor statistics are cumulative for the life of the cache, and will not decrease:

- TotalLocalPageCacheHits**
- TotalLocalPageCacheMisses**
- TotalWaitsForOtherReader**
- TotalPageReadsTotalPageWrites**
- TotalFramesFromFreeList**
- TotalFramesFromFindFree**
- TotalFramesAddedToFreeList**
- TotalOtPageReads**
- TotalDataPageReads**
- TotalBmPageReads**
- TotalMiscPageReads**
- TotalPcesRemovedFromFreeList**
- TotalPcesAddedToFreeList**

statmonitor -J flag did not collect PageManager statistics

The -J flag to statmonitor specifies to collect statistics for the Stone, Shared page cache, and Page Manager only. However, the Page Manager statistics were not being collected. (#45081)

Incorrect values for network stats on Linux hosts

The statistics related to network performance were incorrect for Linux hosts; some values were unreasonably high, others were zero. (#45051)

Memory page cache statistics incorrect on AIX

The following cache statistics were expressed on AIX as the number of 4KB pages, rather than KB, and so were understated 4x in the statmonitor data. This has been corrected, and these statistics are now correctly recorded in KB. (#45058)

DataRSS
TextRSS
DataVmSize

ReclaimGem cache statistic for PinnedPagesCount is incorrect

Reclaim Gem cache statistics for PinnedPagesCount is incorrect. (#45020)

Change in GemStone process statistics on Linux

On Linux, cache statistics are collected from `/proc/pid/status` rather than from `/proc/pid/statm`. This provides some additional statistics, which are available for all GemStone processes.

MaxImageSize (All on Linux)

The maximum (high water) size of the process's image in kilobytes.

MaxRSS (All on Linux)

The high water mark of the processes resident set size. Note that this counter is always 0 on Solaris.

RSSStack (All on Linux)

The stack resident set size.

PageTablesMemoryKB (All on Linux)

The amount of memory dedicated to low-level page tables.

ThreadCount (All on Linux)

Number of threads currently active in this process. An instruction is the basic unit of execution in a processor, and a thread is the object that executes instructions. Every running process has at least one thread.

VolCSW (All on Linux)

The number of voluntary context switches done by the process. Note that this counter is always 0 on HP-UX.

IVolCSW (All on Linux)

The number of times the process was forced to do a context switch. Note that this counter is always 0 on HP-UX.

As a result of this change, **SharedKBytes** and **RSSDirty** are no longer collected on Linux.

Added GemStone process cache statistics

The following GemStone process cache statistics have been added:

CommitRecordPageReads (All)

The number of commit record pages read by the process since it was started.

PagesAddedToCacheFromDisk (All)

Number of pages added to the shared cache by this process which were read from disk.

PagesAddedToCacheFromMidCache (All)

Number of pages added to the shared cache by this process which were copied from a mid-level shared page cache.

PagesAddedToCacheFromPrimaryCache (All)

Number of pages added to the shared cache by this process which were copied from the primary shared page cache.

PagesAddedToCacheNewlyCreated (All)

Number of pages added to the shared cache by this process which were newly created. For gems and the stone, the pages were created by the process. For page servers, the pages were created by gem connected to the page server.

PagesInCacheCreatedInLeafCache (ShrPcMonitor)

Number of pages present in the shared cache which were created in a remote shared page cache.

PagesInCacheCreatedInPrimaryCache (ShrPcMonitor)

Number of pages present in the shared cache which were created in the primary shared page cache.

PagesInCacheFromDisk (ShrPcMonitor)

Number of pages present in the shared cache which were read from disk.

PagesInCacheFromMidCache (ShrPcMonitor)

Number of pages present in the shared cache which were copied from a mid-level shared page cache.

PagesInCacheFromPrimaryCache (ShrPcMonitor)

Number of pages present in the shared cache which were copied from the primary shared page cache.

TotalCommitRecordPageReads (ShrPcMonitor)

Total number of commit record pages read into the shared page cache by all processes since the cache was created.

TotalPagesAddedToCacheFromDisk (ShrPcMonitor)

Total number of pages which were read from disk and added to the shared page cache by all processes since the cache was created.

TotalPagesAddedToCacheFromMidCache (ShrPcMonitor)

Total number of pages which were copied from a mid-level shared cache and added to the shared page cache by all processes since the cache was created.

TotalPagesAddedToCacheFromPrimaryCache (ShrPcMonitor)

Total number of pages which were copied from the primary shared cache and added to the shared page cache by all processes since the cache was created.

TotalPagesAddedToCacheNewlyCreated (ShrPcMonitor)

Total number of pages which were newly created and added to the shared page cache by all processes since the cache was created.

Added Linux System cache statistics

The following system stats may now be collected on Linux.

ActiveAnonMemoryKB

The amount of non-file backed memory that has been used more recently.

ActiveFileMemoryKB

The amount of memory used for buffering files that has been used recently.

ActiveMemoryKB

The amount of memory that has been used more recently and usually not reclaimed unless absolutely necessary.

AnonHugePagesKB

The amount of non-file back memory backed by huge memory pages.

AnonymousMemoryKB

The amount of non-file backed memory mapped into userspace page tables.

BounceMemoryKB

The amount of memory used for bounce buffers for block devices.

CachedMemoryKB

The amount of memory used as cache memory.

CachedSwapKB

The amount of swap used as cache memory.

CommitLimitKB

The total amount of memory currently available to be allocated on the system.

CommittedAsKB

The amount of memory presently allocated on the system, including memory allocated by processes that has not yet been used.

FileBufferSizeKB

The amount of memory used in file buffers.

HardwareCorrupted

A boolean indicating if the system has detected a memory failure.

HugePagesFreeKB

The amount of memory in the huge pages pool that has not yet been allocated.

HugePagesRsvdKB

The amount of memory in the huge pages pool for which a commitment to allocate from the pool has been made, but no allocation has yet been made.

HugePageSize

The size of a huge memory page in bytes.

HugePagesSurpKB

The amount of memory in the huge pages pool above the value in `/proc/sys/vm/nr_hugepages`.

HugePagesTotalKB

The total amount of memory in the huge pages pool.

InactiveAnonMemoryKB

The amount of non-file backed memory that has not been used recently.

InactiveFileMemoryKB

The amount of memory used for buffering files that has not been used recently.

InactiveMemoryKB

The amount of memory which has been less recently used. It is more eligible to be reclaimed for other purposes.

KernelDataMemoryKB

The amount of memory used by the kernel for caching data structures.

KernelDataReclaimableMemoryKB

The amount of memory used by the kernel for caching data structures that may be reclaimed.

KernelDataUnreclaimableMemoryKB

The amount of memory used by the kernel for caching data structures that cannot be reclaimed.

KernelStackMemoryKB

The amount of memory used by the kernel stack.

LockedMemoryKB

The amount of memory that has been locked using `mlock(2)` or similar calls. Locked memory cannot be swapped.

MappedMemoryKB

The amount of memory which has been mapped to files.

NfsUnstableMemoryKB

The amount of memory used by NFS pages sent to the server, but not yet committed to stable storage.

PageTablesMemoryKB

The amount of memory dedicated to low-level page tables.

SharedMemoryKB

The amount of memory enabled for sharing between multiple processes via `shmat(2)` and `mmap(2)` with the `MAP_SHARED` attribute set

UnevictableMemoryKB

The amount of memory that cannot be swapped.

WritebackMemoryKB

The amount of memory which is actively being written back to disk.

WritebackTmpMemoryKB

Amount of memory used by FUSE (Filesystem in Userspace) filesystems.

Removed cache statistics

The following statistic has been removed:

MilliSecPerIoSample

Also note **SharedKBytes** and **RSSDirty** are no longer collected on Linux; see "Change in GemStone process statistics on Linux" on page 17