

---

*GemStone*®



# Visual Statistics Display (VSD) User's Guide

**Version 5.6**

April 2023



## 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.

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 © 1986-2023 GemTalk Systems LLC. All rights reserved by GemTalk Systems.

## PATENTS

GemStone software has been 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.

**UNIX** is a registered trademark of The Open Group.

**Intel** is a registered trademarks of Intel Corporation.

**Microsoft**, **Windows**, and **Windows Server** are registered trademarks of Microsoft Corporation.

**Linux** is a registered trademark of Linus Torvalds and others.

**Red Hat**, **Red Hat Enterprise Linux**, **RHEL**, and **CentOS** are trademarks or registered trademarks of Red Hat, Inc..

**AlmaLinux** is a trademark or registered trademark of AlmaLinux OS Foundation.

**Rocky Linux** is a trademark or registered trademark of Rocky Enterprise Software Foundation.

**Ubuntu** is a registered trademark of Canonical Ltd., Inc.

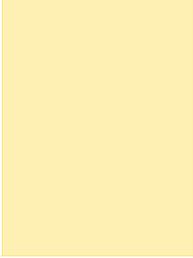
**AIX**, **Power**, **POWER**, **Power8**, **Power9**, and **VisualAge** are trademarks or registered trademarks of International Business Machines Corporation.

**Apple**, **Mac**, **macOS**, and **Macintosh** are trademarks of Apple 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. 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

This manual describes how to use the graphical tool, Visual Statistics Display (VSD), to examine the statistics generated by a GemStone/S 64 Bit or 32-bit GemStone/S repository, or by GemBuilder for Smalltalk (GBS). This information can be used for understanding the performance of your system and diagnosing problems, and for tuning your GemStone configuration. This manual does not describe the underlying GemStone processes, nor how to tune your system; that information is specific to the GemStone product and version, and described in *System Administration Guide* for your server product and version.

## 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 LLC, previously GemStone Systems, Inc. and a division of VMware, Inc.

## Technical Support

### Website

#### **[gemtalksystems.com](http://gemtalksystems.com)**

GemTalk’s website provides a variety of resources to help you use GemStone 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:** [support.gemtalksystems.com](http://support.gemtalksystems.com)

**Email:** [techsupport@gemtalksystems.com](mailto:techsupport@gemtalksystems.com)

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

When submitting a request, please include the following information:

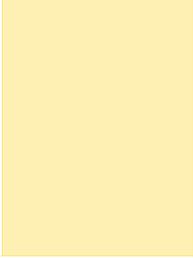
- ▶ Your name and company name.
- ▶ The versions of all related GemTalk products, and of any other related products, such as client Smalltalk products.
- ▶ The operating system, version, and operating system platform you are using.
- ▶ A description of the problem or request.
- ▶ Exact error message(s)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

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

## Training and Consulting

GemTalk Professional Services provide consulting to help you succeed with GemStone products. Training for GemTalk products 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

---

<b><i>Chapter 1. VSD Overview</i></b>	<b>9</b>
1.1 Statmonitor and VSD . . . . .	9
Statmonitor in the GemStone server . . . . .	10
Statmonitor statistics data files . . . . .	11
File Size . . . . .	11
GBS Stat Monitor in GBS . . . . .	11
GBS statistics data files . . . . .	12
1.2 Installing VSD . . . . .	12
Distribution . . . . .	12
Installation on all platforms . . . . .	12
<b><i>Chapter 2. VSD's Graphical Interface</i></b>	<b>13</b>
2.1 Starting VSD . . . . .	13
2.2 Main VSD Window . . . . .	14
Loading data for viewing . . . . .	15
Loading an existing statmonitor file . . . . .	15
Updating the view if the data file is updated . . . . .	15
Loading Multiple Files . . . . .	15
Subsecond sample times . . . . .	15
Monitoring statistics on a running GemStone server . . . . .	16
Process List . . . . .	18
Sorting . . . . .	19
Selecting Process . . . . .	20
Searching for Processes . . . . .	20
Selecting Sets of Processes . . . . .	20
Statistics List . . . . .	20
Searching for Statistics . . . . .	21
Controlling visibility of statistics . . . . .	21

Hiding statistics with only values of zero. . . . .	21
Meaning of statistics. . . . .	21
Selecting and viewing statistics . . . . .	22
2.3 VSD Statistics Chart . . . . .	23
Chart Window. . . . .	23
Chart Options . . . . .	23
Zoom . . . . .	23
Time Format . . . . .	24
Selecting a Line . . . . .	24
Copying value from a line . . . . .	24
Saving a graphic of the chart. . . . .	25
Changing Line Options. . . . .	25
Filters . . . . .	25
Scaling . . . . .	25
Other Transformations . . . . .	25
Summary of Information on a line. . . . .	26
Summary of Information on a section of line. . . . .	26
Other ways to control the view . . . . .	26
Combining statistics values for multiple processes . . . . .	26
Graphing against scales on both right and left side Y-axis . . . . .	27
Rolling Left Trim . . . . .	27
Process names in GemStone . . . . .	27
Changing Hours Offset. . . . .	28
Commas in Numbers . . . . .	28
2.4 Using and Creating Templates . . . . .	28
Gem and Pageserver Names in Templates . . . . .	29
2.5 Working with Multiple Files. . . . .	29
Loading multiple data files. . . . .	30
Loading multiple Individual data sets . . . . .	30
Loading multiple files as a single data set . . . . .	30
Current File and the file mode. . . . .	30
Chart with Statistics from Multiple Files . . . . .	31
Comparing statistics for multiple sequential runs. . . . .	32
2.6 Monitoring. . . . .	32
Auto update . . . . .	32
See only the most recent statistics . . . . .	32
Auto Append of newly created files . . . . .	33
Font colors indicating live processes . . . . .	33
2.7 Aliases for statistic names . . . . .	33
Global aliases . . . . .	34
Process-specific alias names . . . . .	34
Managing aliases in VSD support files . . . . .	35
Alias templates for system operations statistics . . . . .	35
Applying internal operation alias templates to application Gems . . . . .	36
Removing Aliases . . . . .	36
2.8 Setting Background Color . . . . .	37

2.9 Menu Options . . . . .	39
Main Window Main Menu . . . . .	39
Main Window File Menu . . . . .	40
Main Window Chart Menu . . . . .	41
Main Window Template Menu . . . . .	42
Main Window Help Menu . . . . .	42
Process List Pane Pop-up Menu . . . . .	43
Chart Window Chart Menu . . . . .	43
Chart Window Line Menu . . . . .	45
Chart Window All Lines Menu . . . . .	46

## ***Chapter 3. VSD Template and configuration files*** **47**

3.1 VSD Configuration Files. . . . .	47
\$VSDHOME / %VSDHOME% . . . . .	47
.vsdconfig . . . . .	48
.vsdrc. . . . .	48
Configuration file format . . . . .	48
Main Window options . . . . .	49
Chart Window Display Options . . . . .	50
Monitor Window Options . . . . .	51
Window locations and display . . . . .	51
3.2 VSD Chart Templates . . . . .	51
VSD Templates File . . . . .	52
Template Syntax . . . . .	52
3.3 Data in VSD TCL files . . . . .	53
Effect of statistic type on combined lines . . . . .	54

## ***Appendix A. VSD Command-Line Syntax*** **55**

Order of Loading Multiple Files . . . . .	55
---	----



# VSD Overview

---

VSD (Visual Statistics Display) is a graphical tool that allows you to view and interpret the statistical information that can be generated by a running GemStone server or by GemBuilder for Smalltalk. GemStone utilities write this statistical data to a formatted file, and the VSD utility reads these files and displays the data graphically.

## 1.1 Statmonitor and VSD

Every commercial aircraft contains a flight data recorder, commonly called a black box. Black boxes are used to help determine what went wrong if the aircraft ever crashes. Without it, crash investigators would have very little to go on and the cause of many crashes would never be known.

In a production GemStone application, the *statmonitor* program serves the role of the black box. It runs in the background, logging the values of all cache statistics to a disk file. Should a problem with the repository occur, you can review these statistics to determine the cause, allowing you to identify problems and tune GemStone applications for better performance.

Data alone, however, is not enough; you need tools to display this data to allow analysis. VSD allows you to view and interpret the statistical information gathered by GemStone tools such as *statmonitor*. VSD reads the *statmonitor* files and displays the values in a graph. This graphical user interface gives you a meaningful way to see how your GemStone system changes over time, based on periodic samples of the system's state.

Many questions about repository performance are impossible to answer after the fact unless statistics data files are available. For example:

- ▶ Why did `markForCollection` take longer than normal?
- ▶ Why did the backup run slower than before?
- ▶ Why was there more repository growth than normal?
- ▶ Which Gem session was consuming the most CPU or I/O during a time of poor performance?
- ▶ Which Gem had the most number of commit failures?

If you're trying to answer questions such as these, GemStone statistics are your best source of information.

While VSD is a standalone executable, it requires statmonitor or GBS Stat Monitor to generate data:

- ▶ *statmonitor* is a stand-alone GemStone server executable with a command-line interface. This is provided as a utility with your GemStone server product and version, and is required to be the same product and version as the shared page cache it will be attached to. It does not log into GemStone nor use a GemStone session. Statmonitor writes data files to disk, in a platform- and version- independent format, in a text format.
- ▶ *GBS Stat Monitor* is a GBS/VW process that can be enabled and disabled within a VW image. Like statmonitor, this writes data files to disk, in a platform- and version- independent format, as a binary .gfs file
- ▶ *VSD* provides a graphical user interface for viewing the files that statmonitor or GBS produces. The VSD executable is provided for a variety of platforms.

Any version of VSD running on any platform can load statistical data files generated from any other platform and version of the GemStone server or GBS, although older versions of VSD may not be able to read statmonitor files produced by more recent GemStone server or GBS versions.

## Statmonitor in the GemStone server

The statmonitor executable is provided as part of the GemStone/S 64 bit or 32-bit GemStone/S Server, and is described in the *System Administration Guide* for that version of the server.

Statmonitor requires a shared page cache, either the Stone's shared page cache, if the stone is running on that machine, or a remote shared page cache; or, in some versions of GemStone, the -X option can be used to collect system statistics.

For example, to start statmonitor, at the command line, enter a command similar to:

```
statmonitor -f outputFile -i sampleInterval stoneOrCacheName
```

Statmonitor provides a large number of command line options to control monitoring, including the number and range of processes to monitor and how long to run. These options are provided in the *System Administration Guide* for your platform, or by executing `statmonitor -h` at the command line.

Statmonitor can be terminated by CTRL-C, or when a limit is reached as specified on the command line, and will shut down when the repository is shut down.

GemStone recommends running statmonitor at all times for production systems. If the GemStone system encounters an error, it may be able to report the nature of the error, but diagnosis of what led to that error often requires understanding previous activity and the actions of other processes, which can only be understood from statmonitor data. And some important questions, such as why you ran out of memory or why the repository is out of disk space, can only be analyzed using data from before the problem occurred.

## Statmonitor statistics data files

Statmonitor writes text files to disk, either as plain text or in gzip compressed form.

These files contain formatted information, platform- and version- independent, that is read and displayed in VSD. VSD does not directly connect to processes in the GemStone server, but it can startup the statmonitor utility via the command line interface, and read the output files that are generated by statmonitor.

Statmonitor data files may be generated as plain text or compressed using gzip; VSD can read plain text files and files zipped generated or zipped afterwards using gzip or lz4 compression.

The statmonitor data file format changes very rarely. The latest version of VSD can always read statmonitor files generated by older as well as recent versions of the GemStone server.

These files normally have the extension `.out`, `.out.gz`, or `.out.lz4`, although you may specify any file name.

## File Size

For each sample, statmonitor writes a line for each process being monitored to a data file. This data file continues to grow as long as statmonitor is running.

With small intervals and a large number of processes, over time this file may become very large. A longer sampling interval will keep file sizes more reasonable, but may not be fine enough granularity to catch certain types of problems.

Writing the file in compressed format will save space.

Statmonitor includes options to regularly terminate the statmonitor file and begin a new one, either by time or by the number of samples collected. For example, you may wish to start a new file every 24 hours, and delete old files after a few days.

When you are diagnosing specific behaviors, a 1-second sample is the most convenient, although sub-second intervals may be needed. Monitoring running systems will usually require a longer interval to keep the file sizes manageable.

## GBS Stat Monitor in GBS

GBS/VW can be configured to collect statistics describing the performance of its internal operations. These statistics are written to a file named `GBSStatArchive-<dateAndTime>.gfs` in the current VisualWorks working directory.

GBS main statistic types are for `GbsSession` and `GbsSessionManager`. GBS can also be configured to collect per-class cache inventories, in which each process line represents a Class that has instances in the classMap.

to start collecting both main and cache inventory statistics, in GBS execute:

```
GBSM statsEnabled: true
```

to start collecting only the main statistics, in GBS execute:

```
GBSM mainStatsEnabled: true
```

to start collecting only the cache inventory statistics, in GBS execute:

```
GBSM statMonitorRunning: true
```

To stop collecting statistics, use the same method with a false argument.

For more information, refer to the *GBS User's Guide* chapter on profiling.

### GBS statistics data files

Statistical data files written by GemBuilder for Smalltalk (GBS) are written in a binary form, with the extension `.gfs`, for GemTalk Flexible Statistics file. `.gfs` files encode the statistic names and values; and, unlike server-generated statmonitor data files, `.gfs` files also the definition and other information about each statistic.

## 1.2 Installing VSD

There is no need to run VSD on the same machine or on the same platform as the one on which the GemStone server or GBS is running. VSD on any machine and platform can read statistics generated on any other machine and platform. If you will be running VSD on a machine other than the one in which the GemStone server software is installed, you can download the appropriate platform of VSD and install it on that machine.

On UNIX platforms, VSD requires X-Windows.

### Distribution

VSD is provided as part of the server distribution for all versions of GemStone/S and GemStone/S 64 Bit on all platforms.

You may also download VSD separately for any platform. VSD can read data files that are produced on any supported platform by any GemStone product, so you can analyze statistics on any VSD platform regardless of where the server is running.

VSD downloads for all platforms are provided on the GemTalk website, at <http://gemtalksystems.com/vsd/>

### Installation on all platforms

If you are using VSD as part of your GemStone server distribution, VSD is already installed and ready to use, and is on your machine search path.

Installation of standalone VSD is simple. The following instructions use `/` as a path separator, although the process is the same on Windows. On the mac, VSD is provided as an Apple-notarized `.dmg`.

1. Download the VSD distribution, and select a VSD installation directory, *vsdInstallDir*. Unzip the `.zip` file into *vsdInstallDir*, or double-click the `.dmg` file to open, and drag the installation tree to *vsdInstallDir*.

This will create two subdirectories in *vsdInstallDir*; *vsdInstallDir/bin*, which contains the VSD executable files, and *vsdInstallDir/lib*, which contains several subdirectory trees containing the VSD library files.

2. It is convenient, but not required, to add *vsdInstallDir/bin* to the machine search path.

If the `bin` directory is not added to the machine search path, you may execute `vsd` directly from *vsdInstallDir*. VSD will search for the library files under the `/lib` subdirectory of the executable's parent directory, so you do not need to add this directory to the path.

# VSD's Graphical Interface

---

VSD provides a graphical way to view the data that is recorded by statmonitor. When you open VSD on this data, you will see a very large number of statistics that were recorded for potentially a very large number of GemStone processes. Not all of these are going to be useful; your task is to locate the interesting and useful information.

## 2.1 Starting VSD

You can start VSD from your windowing system or from the command line.

- ▶ Double click on `vsd` in a graphical file browser or `vsd.exe` in Windows Explorer.
- ▶ Execute `vsd` at the command line.  
You can either provide the full path to the `vsd` executable, navigate to the directory containing `vsd`, or, if the directory containing the `vsd` executable is on your machine search path, just enter '`vsd`'.

You may start VSD and automatically open a specific data file by:

- ▶ execute `vsd`, specifying the statistics file name on the command line.

For example:

```
prompt> vsd statmonDataFileName  
vsdinstalldir/bin/vsd statmonDataFileName
```

- ▶ On windows, you drag a statmonitor data file and drop it on `vsd.exe` in Windows Explorer

The VSD Main window will appear.

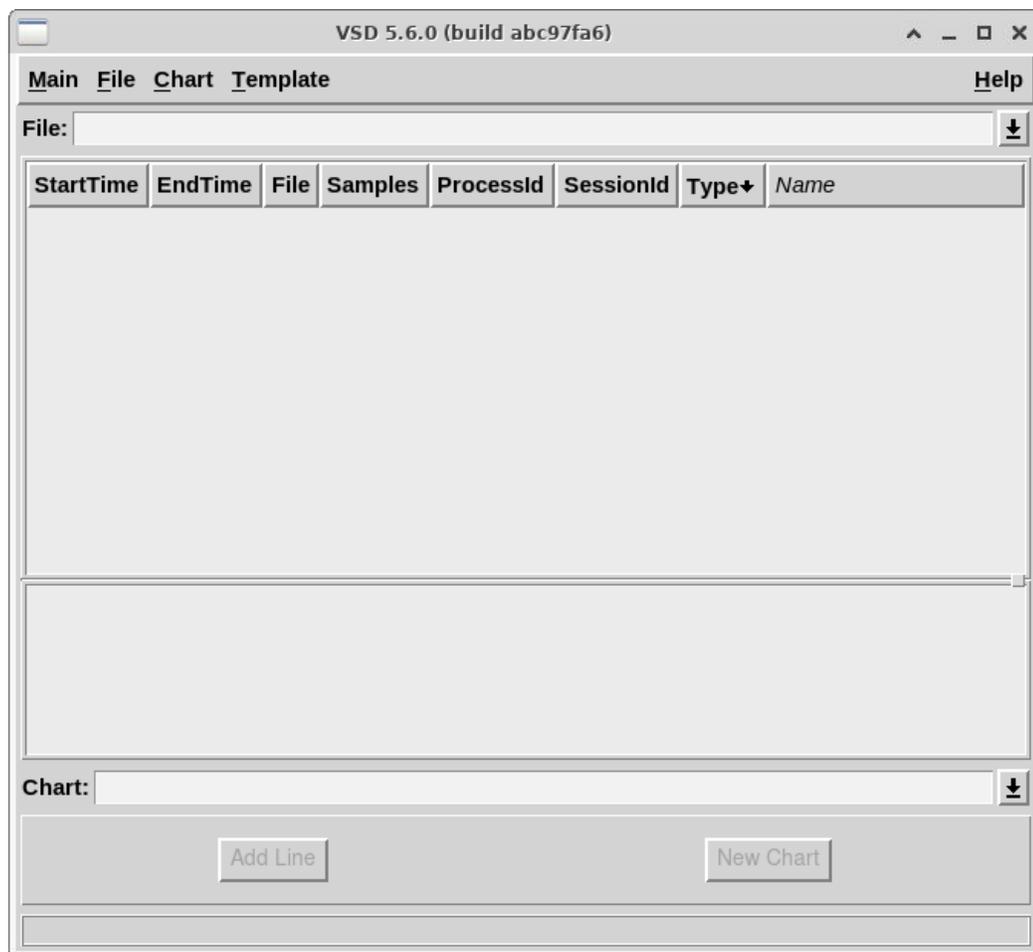
## 2.2 Main VSD Window

When VSD is started, it displays the **Main window**. This window has a single instance and remains open while you are running VSD. When you are within VSD, the F12 key will bring the Main window to the top.

The Main window displays the list of processes and other information about the statmonitor file; here, you will locate and select processes and statistics for which you want to view charts.

This window also provides menus that allow you to set configuration options that control how your charts are displayed. The details for the various menu options are described in section 2.9 on page 39.

Most of these display options are written to a .vsdrc file when VSD exists, which is read on startup. The next time you open VSD you will get the same options. Details of .vsdrc are on page 48.



## Loading data for viewing

If you opened VSD with a particular data file or files as an argument, the main window will display the processes listed from that file. Otherwise, you need to load a data file, or start monitoring a running server.

### Loading an existing statmonitor file

To load a statmonitor file into VSD, use the menu item **Main > Load Data File(s)...**

This brings up a file dialog that allows you to browse for an existing statmonitor output file.

If you know the name of the file you would like to load, you can also type the full path in the File entry box.

### Updating the view if the data file is updated

You may load a statmonitor data file that is still being updated by the live system. As new data is added to this file, you have the option to have vsd update its view.

If you select **File > Auto Update**, VSD automatically updates your display, and any associated charts, any time the data file changes. Otherwise, if auto update is disabled, you can choose **File > Update** periodically to update the displays when you choose.

### Loading Multiple Files

You can load multiple statmonitor data files, either keeping the statistics separate by using the **Main > Load Data File(s)...** menu option, or appending them as if they were a single data set by using the **Main > Append Data File(s)...** menu option.

For more details on working with multiple files, see “Working with Multiple Files” on page 29.

### Subsecond sample times

The statmonitor utility that is part of the GemStone server allows you to specify a sampling time in milliseconds rather than seconds, by using the **-I** option. Since versions of VSD earlier than v4.0 did not handle samples of less than one second, statmonitor historically did not write exact timestamps for subsecond sample rates; it treated the timestamp a counter that was incremented with each sample. The result was a series of timestamps that looked like 1-second intervals but represented statistic values at the requested sample increments.

GemStone/S 64 Bit version 3.3 and later, and GBS, write subsecond timestamps, adding milliseconds to the timestamp if needed. Earlier versions of GemStone, however, write relative timestamps for subsecond samples.

For statistics from older GemStone servers, there is no way to distinguish a true 1-second sample rate from a subsecond sample rate with nominal 1-second timestamps. When analyzing data with subsecond samples from earlier versions of the GemStone server, you must be aware of the sample rate when you load the file, and ensure you use relative timestamps and a time format of seconds.

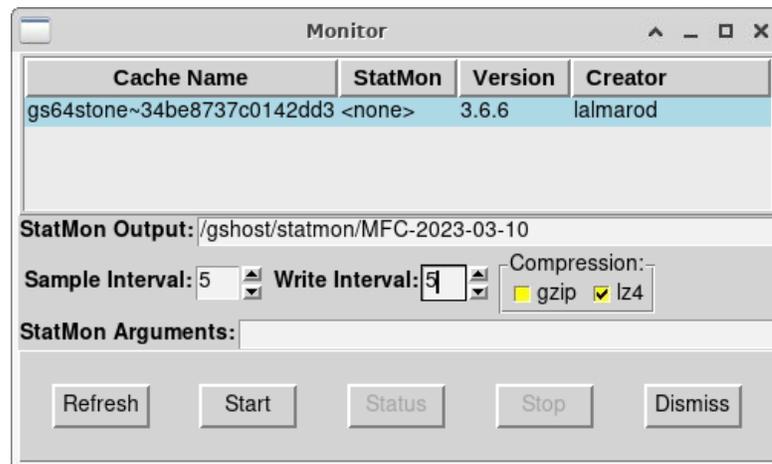
## Monitoring statistics on a running GemStone server

You can use VSD to start statmonitor and automatically read directly from the data file this statmonitor is writing to. This options is not available on Windows.

In order to monitor, you must have \$GEMSTONE set to a GemStone server installation, and the machine search path set to include \$GEMSTONE/bin, to allow VSD to provide a list of running GemStone caches to monitor. However, you can monitor any GemStone version or product, regardless of the version of GemStone specified in the environment of VSD. VSD uses the current GemStone environment only to list available caches. It reads information on the processes, and starts the correct version of statmonitor for the selected stone.

To monitor a live system:

**Step 1.** Bring up the Monitor window using the menu item **Main > Monitor...**



The Monitor window will display a list of all shared caches on this machine.

**Step 2.** Select the cache you wish to monitor, and the monitoring options.

These options are a subset of those offered by the server statmonitor utility. Options are:

- ▶ The statmonitor output pathname.
- ▶ The sample interval (in seconds) – that is, how frequently to statmonitor will read the cache to collect a sample.
- ▶ The write interval – the maximum number of seconds to wait before flushing the samples to the output file. This determines the rate your monitoring view is updated.
- ▶ Whether to compress the output file.
- ▶ Any other statmonitor arguments you may wish to use; see `statmonitor -h`, or the *System Administration Guide* for the version of the GemStone server you are monitoring, for supported arguments.

**Step 3.** Click the **Start** button to start monitoring.

This will start statmonitor, and create a statmonitor data file in the specified location, or in the current working directory. If statmonitor startup fails, the Monitor status window will appear which shows the statmonitor command line with the error message.

The Monitor window will remain open. You can use the **Dismiss** button to close the window; this does not affect monitoring.

VSD loads the data file as soon as it has some samples in it (unless you have explicitly turned off Auto Update).

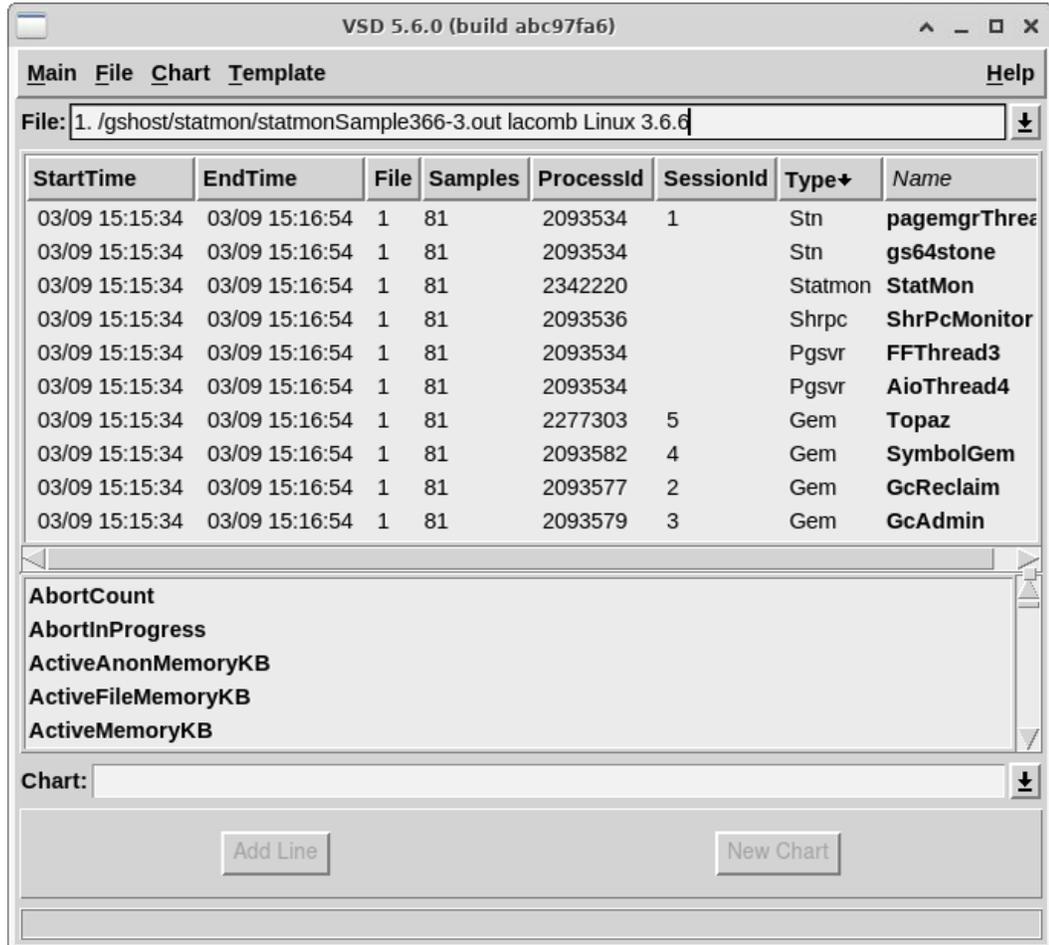
The statmonitor data file that is created by monitoring will grow as long as you are monitoring. If you have a large system and collect samples frequently, this file may become very large.

**Step 4.** Click the **Stop** button to stop monitoring.

If you have dismissed the Monitor window, you can re-open it using the menu item **Main > Monitor**. Select the cache you have been monitoring to enable the **Stop** button.

Note that the statistics data that is created will remain, and can be viewed later. It must be explicitly deleted if you no longer need it.

After you've loaded the data file or started monitoring, the VSD window looks something like this:



## Process List

The upper list pane contains a list of all the entities – processes, thread-based sessions within processes, and other entities such as the disks – for which data is recorded in the statmonitor data. Depending on the options used to collect the data, this may include all GemStone processes or a small subset, and may include a number of non-GemStone entities including the operating system, disks, etc. While these are not all strictly “processes”, in the following discussion this term is used to include all entities that appear in the process list.

For each process, session, or other entity, the display includes:

- ▶ The **StartTime** - the time that process first recorded a sample, adjusted to local time on the machine executing VSD.
- ▶ The **EndTime** - the time that process last recorded a sample, also adjusted to local time.
- ▶ The **File** id for the file in which it was recorded. This is of interest if you have more than one file loaded; the file id corresponds to the number in the file select drop-down.

- ▶ The total number of **Samples** for that process recorded in the data file.  
If you are monitoring, processes that are currently running and for which samples are being taken are shown in green, and these number will change as new data samples are taken.
- ▶ the OS **ProcessID** for the process. Multi-threaded processes in which the threads are GemStone sessions will have individual entries for each thread, and list the same process Id for each.
- ▶ the GemStone **SessionId** for that processes. This only applies for GemStone sessions.
- ▶ the process **Type**.

GemStone processes fall into a number of types, which govern which statistics apply. The common process types are:

**Stn** – Stone (and Page Manager in cases where this is a thread in the Stone )

**Shrpc** – Shared page cache monitor

**Pgsvr** – AIO or free frame page server, or a page server for a remote session

**Gem** – Gem; includes system Gems, RPC Gems, and linked sessions such as linked Topaz

**Statmon** – statmonitor

Additional types will appear if you are also monitoring OS information, including platform-specific statistics.

- ▶ the process **Name**.  
This is assigned by statmonitor or GBS's Stat Monitor, and displayed in VSD. It may be assigned for Gem processes within the GemStone server; see "Process names in GemStone" on page 27.

To reduce the number of processes that are visible in the list, the processes of a given type can be hidden, using the menu item **Hide Instance By Type...**

## Sorting

By selecting a column header in the main windows, you can sort all the processes by that column's attribute. The current sort column is indicated by an up or down arrow next to the column label.

Within primary sort order, the results are further sorted by a secondary sort attribute. This is indicated by the column label in italic font. To specify the secondary sort attribute, right click on the column label.

Note that the secondary sort key direction (ascending or descending) is always the same as the direction of the primary sort.

For example, if you want to find the Stone, you can sort by Type, descending, which sorts the type Stn first. Or, if you know the session ID of the Gem, you can sort by SessionId and scroll until you find that Gem.

## Selecting Process

When you select a process, the list of statistics relevant for that process is displayed in the statistics pane. Multiple selection is allowed.

## Searching for Processes

You may search for a process by name or type, using the Process List pane popup menu item **Search...**, control-S, or control-F when focus is on the Process List pane. This opens a search line just above the Process List pane.

## Selecting Sets of Processes

The menu option **Main > Select Instances By...**, also available on the process list popup menu, allows you to select all processes matching a particular criteria. The options are:

**All** – select all processes.

**Clear** – deselect all processes.

**Name...** – open a dialog allowing you to enter a string, select processes matching that name. The string supports wild cards \* and ?, and allows you to search case-sensitive or case-insensitive.

**Process ID...** – open a dialog allowing you to enter a process Id, and select processes with that process Id.

**Session ID...** – open a dialog allowing you to enter a session Id, and select processes with that Session Id.

**Statistic** – select all processes for which the currently selected statistic or statistics apply. Does nothing if no statistic is selected.

**Type** – select all processes of the same type as the currently selected process or processes. Does nothing if no process is selected.

## Statistics List

The lower pane displays a list of statistics.

Each process type has a specific set of statistics that are relevant for that process. Many of these are unique for that process type; for example, `CommitCount` applies to Gem processes, while `TotalCommits` is meaningful for the Stone. Others, such as `UserTime`, apply to all processes.

When no process is selected, all statistics in the system are shown. When a particular process is selected, only statistics appropriate for that process type are listed. When multiple processes are selected, only statistics that are common to all the selected process types will be listed.

You may scroll down the list to locate the statistic you want. Typing in the first character will select the first statistic with a name beginning with that letter.

Notice that some statistics are listed in bold font, while others are in regular font. Statistics with the font in bold have non-zero values recorded.

## Searching for Statistics

You may search for a statistic by entering the first characters in the Statistic name, or by using the Statistic list popup menu item **Search for Counter...**, control-S, or control-F when focus is on the Statistics List Pane. These open a search line just above the Statistics List Pane.

## Controlling visibility of statistics

### Hiding statistics with only values of zero

You may hide statistics for which the value of that statistic is 0 for the entire sample set. These are the statistics that are listed in non-bold font.

To hide statistics where there are no non-zero values, check **No Flatlines** in the **Main** menu or in the right mouse button popup menu. Note that "No Flatlines" means "No Flatlines at zero", and does not apply to statistics whose values present a flat line at a non-zero value.

### Meaning of statistics

While many statistics have intuitive names, you will often need to look up details. You should not expect to understand every statistic recorded for your system. Some of these statistics require an intimate knowledge of GemStone server internals, some are experimental, and some may only apply under specific server conditions. GemStone documentation, Techtips, and analysis from GemStone Technical Support will include specific statistics that are particularly helpful.

The information in the current version of VSD, including more detail, is displayed in the Statistics Information Window. To open the Statistic Information, use the menu item **Main > Show Statistics Info**.



This singleton window can be left open; it will display the details for any selected statistics as it is selected. You can also lookup statistics using the drop down.

## Selecting and viewing statistics

After you have selected one or more processes and one or more statistics, you can display these results in a new **Chart Window** by clicking on the **New Chart** button. You can create a named chart by entering a name in the **Chart** name entry field.

If you have a chart open, you can add the selected statistics to an existing chart. Select the chart by name in the drop-down with the **Chart** name entry field, then click on the **Add Line** button. The Chart name is set to the most recently selected chart, so if you do not explicitly select a chart, the statistics will be displayed in the most recently selected chart.

The following shows the main window with one process selected, the Stone, and two statistics selected, TotalAborts and TotalCommits.

VSD 5.6.0 (build abc97fa6)

Main File Chart Template Help

File: 1. /gshost/statmon/statmonSample366-3.out lacomb Linux 3.6.6

StartTime	EndTime	File	Samples	ProcessId	SessionId	Type	Name
03/09 15:15:34	03/09 15:16:54	1	81	2093534	1	Stn	pagemgrThrea
03/09 15:15:34	03/09 15:16:54	1	81	2093534		Stn	gs64stone
03/09 15:15:34	03/09 15:16:54	1	81	2342220		Statmon	StatMon
03/09 15:15:34	03/09 15:16:54	1	81	2093536		Shrpc	ShrPcMonitor
03/09 15:15:34	03/09 15:16:54	1	81	2093534		Pgsvr	FFThread3
03/09 15:15:34	03/09 15:16:54	1	81	2093534		Pgsvr	AioThread4
03/09 15:15:34	03/09 15:16:54	1	81	2277303	5	Gem	Topaz
03/09 15:15:34	03/09 15:16:54	1	81	2093582	4	Gem	SymbolGem
03/09 15:15:34	03/09 15:16:54	1	81	2093577	2	Gem	GcReclaim
03/09 15:15:34	03/09 15:16:54	1	81	2093579	3	Gem	GcAdmin

TimeWaitingForIo  
 TotalAborts  
 TotalCommits  
 TotalGemFatalErrors  
 TotalLostOtsSent

Chart: [ ]

Add Line New Chart

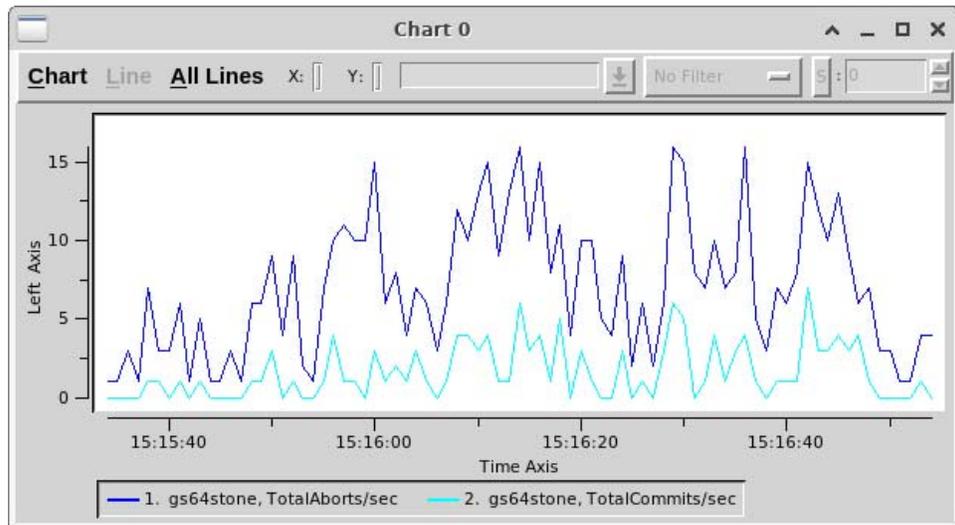
## 2.3 VSD Statistics Chart

To view data for particular statistics, you will open a **Chart Window**. These VSD statistics charts allow you to view and analyze the data for the statistics you have selected. You may have many Chart Windows open.

There are a variety of ways to control the display and to perform operations on selected lines. The details for the various menu options are described under section 2.9 on page 39.

### Chart Window

When you first open a chart, all the values for all statistics are displayed with its default settings. Each line has a unique color and/or pattern.



At the bottom is the Legend - the list of process statistics that are viewed in this chart. Each is described by Process name and statistics name.

The axis of the graph are labeled; the horizontal axis is always time.

The vertical axis is labeled with the number of the units for all the statistics; this may or may not be meaningful, if the statistics used different units of measurement. For example, comparing the counts of the number of commits and the time waiting for I/O, by units, does not directly make sense. However, it can be very useful to see how each varies relative to the other over time.

### Chart Options

You can customize and manipulate a VSD chart in many ways to make your analysis easier. A few of the more useful are described here.

#### Zoom

To zoom in on a subsection of the chart, use the menu item **Chart > Zoom In**. This prompts for a rectangle to zoom in to, allowing you to view with more detail the section of the chart that you are interested in. You can zoom in multiple times to narrow your view further.

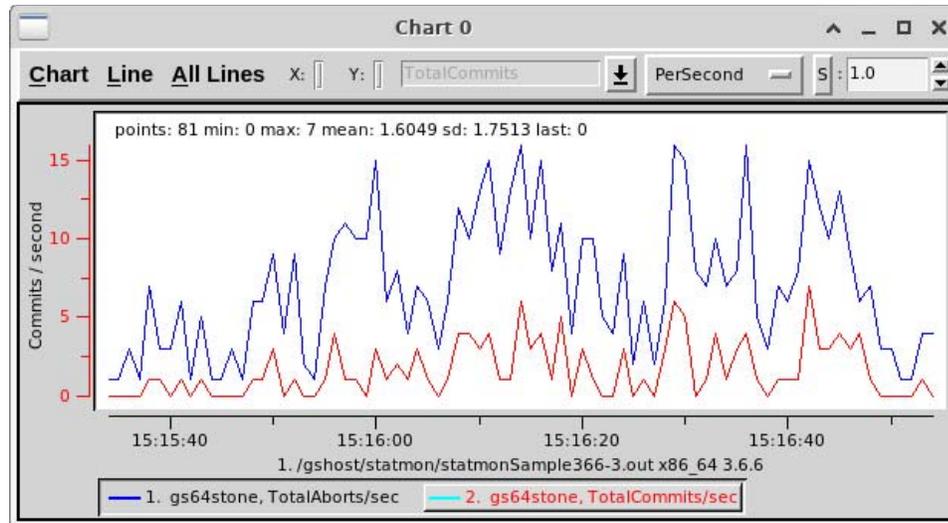
**Zoom Out** returns you to the previous level of zoom.

## Time Format

The times marked on the horizontal axis, by default are in HH:MM:SS format (with hours in 24-hour format). If your sample spans multiple days, you can set the time format to include the month and day using the **Chart > Time Format** menu item.

## Selecting a Line

You can select an individual line from the Legend at the bottom, or by clicking on that line in the graph. Selecting a line colors it red and displays information about that line.



When a line is selected, the following information is available:

- ▶ On the graph itself, a brief **summary statistics** for the line. This includes the number of data points for that statistic, the minimum, maximum, and standard deviation are displayed. These units stay the same regardless of scale or transformation.
- ▶ When the mouse is over a data point in the graph, the **X** (time) and **Y** (units) of that point are displayed. These units stay the same regardless of scale or transformation.
- ▶ the **name** of the statistic.
- ▶ the current **filter**. The filter allows you to distinguish absolute values from rates of change. Each statistic has a default filter based on its type, which can be changed as needed.
- ▶ The scale or other transformation. The scale allows you to compare lines with widely different ranges of values, by scaling one or the other to bring the ranges closer together.

## Copying value from a line

When you have a selected line, you can copy the value of a point on the line using the Chart windows pop-up menu item **Copy Y value**. This menu item is only displayed when you right click while the cursor is over a point on the active line.

The copied value can be pasted into another application using control-V. Middle-button paste depends on the target application copy buffer and may or may not work.

## Saving a graphic of the chart

Using the Chart Window menu item **Chart > Snapshot**, you can save a .gif image of the graphics on the Chart Window to a disk file for later reference.

## Changing Line Options

### Filters

The filter allows you to view either absolute values, rates of change per second or sample. Filter options are:

**No Filter** – Displays the values for the statistic exactly as they are expressed in the statmonitor text file.

**PerSample** – Displays the difference between two consecutive samples of the statistic.

**PerSecond** – Displays the difference between two consecutive samples of the statistic, divided by the number of elapsed seconds between the two samples.

**Aggregate** – Displays a running total for the per-sample deltas of the statistic.

These filters can be adjusted for an individual line using the filter drop-down on the menu bar, , or for all lines on a chart using the menu items under the menu **All Lines**.

### Scaling

When you have multiple lines with widely different ranges of values, you can increase the visibility of a line relative to the other line by scaling it.

Select the line, and in the upper right, use the arrow indicators to scale up or down the line.

Using the **Line > Compute Scale** or the **Chart > Compute Scale All** menu options will calculate and apply a scale based on the chart data values.

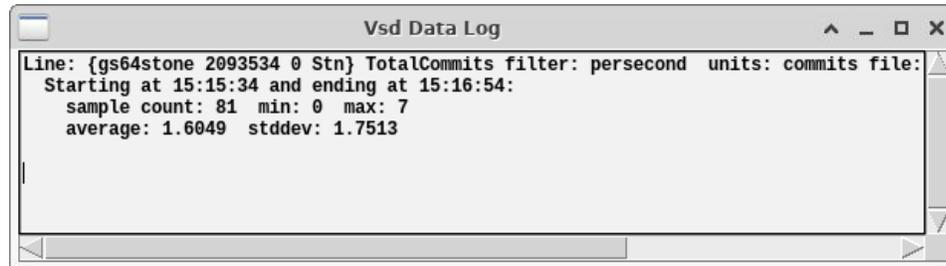
### Other Transformations

By default, **S** is selected for the transformation setting, between the filter list and the scaling value entry field. When **S** is selected, the numeric scrolling field is used to scale the selected line.

You may also select other options, rather than **S**. This allows you to transform the line by adding, subtracting, multiplying, or by dividing the value in the numeric scrolling field.

## Summary of Information on a line

The menu item **Line > Log Info** will open (or append to an already open log window), summary information about a single line. For example:



```

Vsd Data Log
Line: {gs64stone 2093534 0 Stn} TotalCommits filter: persecond units: commits file:
Starting at 15:15:34 and ending at 15:16:54:
sample count: 81 min: 0 max: 7
average: 1.6049 stddev: 1.7513

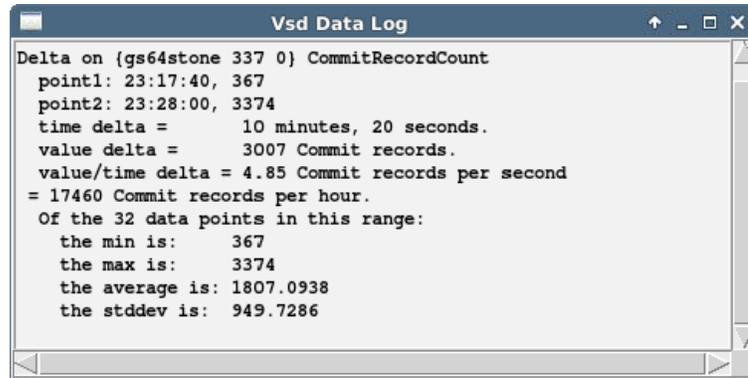
```

## Summary of Information on a section of line

A very useful tool is collecting statistics on a section of line. For example, you may want to know the commit rate over a single period of intense activity, rather than for the entire period of sampling, or the rate of growth of a CommitRecordCount.

The menu item **Line > Log Delta** will allow you to select two point on the selected line. It then opens (or appends to an already open log window), summary information about that section of line.

The useful information will be different for counter statistics vs. values. For CommitCount, which is a counter type and therefore already a delta, the average will be of interest; for values is as CommitRecordCount that provide unit values, the value delta will be useful. For more on statistics types, see the descriptions on page 53.



```

Vsd Data Log
Delta on {gs64stone 337 0} CommitRecordCount
point1: 23:17:40, 367
point2: 23:28:00, 3374
time delta = 10 minutes, 20 seconds.
value delta = 3007 Commit records.
value/time delta = 4.85 Commit records per second
= 17460 Commit records per hour.
Of the 32 data points in this range:
the min is: 367
the max is: 3374
the average is: 1807.0938
the stddev is: 949.7286

```

## Other ways to control the view

There are many ways to control what you are viewing to make analysis of your data easier. Here are a few that are particularly useful.

### Combining statistics values for multiple processes

Normally, when opening a chart or adding lines to an existing chart, one line is added for each selected process for each selected statistic. If you have selected a large number of processes — for example, all Gems — this may be cumbersome to see in the resulting chart.

In cases like this, you can combine values for a single statistic for all the selected processes. To combine values, in the Main VSD Window, check the value **Combine** in the **Main** menu or in the right mouse button popup menu.

### Graphing against scales on both right and left side Y-axis

By default, all lines in a graph are plotted according to Y-axis units scaled on the left-hand side of the chart. You can instead plot some of the lines according to units scaled on the right-hand side of the chart, which will allow them to use a different scale. This is useful when the lines use different units of measurement from those plotted on the left.

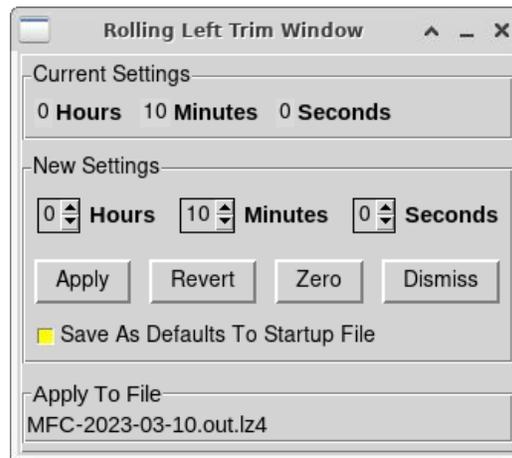
To change the side on which a line is selected, with the line selected in a Chart Window, uncheck the value **Graph on left axis** on the Line menu.

### Rolling Left Trim

When monitoring a running GemStone server, over a long period of time there may be many samples recorded, so that the most recent samples are in a very narrow section of the graph along the right side.

Since in many cases the most recent statistics are of particular interest, you can configure the display to only show the samples for a fixed period of time before the latest loaded sample; as new samples are loaded, the older ones are no longer displayed.

The **File** menu item **Set Rolling Time Window...**, or Cntrl-shift-R, brings up the following dialog:



Rolling left trim, like auto update, applies to the current file, which is shown at the top of the Main window, in the field labeled **File**. The name of the current file is displayed at the bottom of the Rolling Left Trim Window.

### Process names in GemStone

When statmonitor records statistics for Gem processes, it records them with the name "Gem", or for linked topaz logins, "Topaz". This can be modified to an application-meaningful name in the GemStone server, so this name is recorded in the statmonitor data. This can be particularly useful for regular maintenance processes, such as your backup or markForCollection, that are run from scripts.

The options to set the process name vary by the GemStone server version.

- ▶ The topaz command line `-u` argument name is used to set the process cache name, in GemStone/S 64 Bit v3.4 and later.
- ▶ Before logging in, the command:

```
topaz> set cachename MyProcessName
```

can be executed to set the cache name. This is available for linked sessions in v2.4 and later, and additionally for RPC sessions in v3.4 and later.

- ▶ After login, in the session you want to rename, execute:

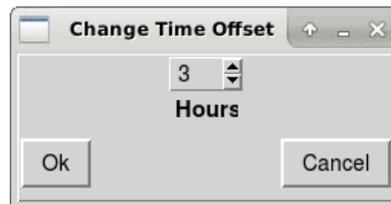
```
System cacheName: 'myName'
```

Depending on how soon you can execute this after login and the statmonitor sample rate, you may have one or more samples showing the original default name.

## Changing Hours Offset

Statmonitor data files include a header file that specifies the timezone in which statmonitor recorded this file. VSD automatically adjust the timestamps displayed according to this timezone, rather than the timezone of the node actually running VSD. This allows you to correlate VSD with log file messages, regardless of the timezone in which VSD is running.

Depending on your configuration, this may not always be what you want to see. You can manually adjust the display offset using the main menu option **Change Time Offset...**, and select the number of hours plus or minus to adjust VSD's display:



## Commas in Numbers

To make it easier to read very large numbers, by default VSD displays comma separators, e.g. 6,000,000. VSD does not have localization; if the separators are distracting, you can disable them using the **Main** menu item **Preferences > Show Comma In Numbers**.

## 2.4 Using and Creating Templates

VSD templates let you quickly create a chart with a commonly used set of statistics. Templates are helpful if you find yourself performing the same task frequently in VSD – for example, monitoring the same five or six statistics. By creating a template for the statistics that you want to monitor most frequently, you can automate the task of building charts.

VSD is shipped with a set of predefined templates. You can open charts on these templates using **Template > New Template Chart**, and choosing any of the predefined templates. If you create your own custom templates, they will also appear in this list.

You can create a custom template based on a chart you have created in VSD, or by editing the formatted text file that contains the templates. See "VSD Configuration Files" on page 47 for more on the text file and how it is formatted, and editing this file outside of the GUI.

To create a custom template from VSD:

**Step 1.** set up a chart with the processes, statistics, scaling, and so on, that you want in your template.

**Step 2.** Chose the **Chart** menu item **Save Template**. This will prompt you to enter a name for the custom template. Hit the enter key when you are done.

You can now use your new template, from the **Template > New Template Chart** menu item.

if you want to remove your template, edit the .vsdtemplates file, and use the menu item **Chart > Reload Template File** to reload the templates file.

## Gem and Pageserver Names in Templates

When you create a template that includes processes that are Gems or PageServers, the name of the process, as well as the type, is included in the template. This helps avoid an excessive number of lines appearing when the chart is opened.

The name that is included does not include any serial numbers that are displayed with the name in VSD. So for example, if the chart on which you are creating a template includes a Gem named Topaz19, the template will use Topaz\* as the name. When you later use the template to create a chart, lines will be added for all Gems with the name Topaz.

Since you may have many Topaz processes, and cannot determine which one is performing a particular task such as backup or MFC, including Gem and Page server statistics in templates requires more consideration. Providing the Gem with a name, as described under "Process names in GemStone" on page 27, is one solution.

## 2.5 Working with Multiple Files

VSD provides a number of ways to work with multiple statmonitor data files. While analysis is often done using a single statmonitor data file or monitoring session, there are a number of cases in which you need to work with multiple files. For example:

- ▶ Multiple Files that represent a single logical data set

Production systems should normally run statmonitor at all times, and archive the results for some period of time, to ensure that this information is available for diagnosis if a problem occurs. This usually involves automatically restarting statmonitor periodically, for example, starting a new data file every day at midnight.

When studying the progress of a operation such as a markForCollection, the statistics of interest may span more than one data file.

- ▶ multiple files that interleave

If you have statistics for remote caches, you may wish to display these along with the statistics for the Stone's cache to track the interactions.

- ▶ Entirely different data sets that you want to compare against each other

For example, if you are studying the effect of a configuration or code change on performance, you may have data sets for two runs of a particular operation, and want to compare them to understand the differences.

## Loading multiple data files

### Loading multiple Individual data sets

You can load multiple statmonitor files using the menu item **Main > Load Data File(s)....**, or by specifying multiple files on the command line.

### Loading multiple files as a single data set

If your multiple data files are logically a single data set (for example, you start statmonitor on the same repository, shut statmonitor down, and start statmonitor again), then you can append the second and later data files to the first file, allowing you to see all the processes and samples at once.

To open multiple files so they are appended, use the menu item **Main > Append Data File(s)**, or the **-a** command line argument. The File list displays the name of the first file, with **[Appended]** following.

You must ensure that the files to be appended are specified in the correct order. The append operation does not detect and prevent loading data sets that are in incorrect order or overlap, and loading invalid sequences results in unpredictable behavior and possible errors.

## Current File and the file mode

VSD always has a current file, which is displayed in the **File:** list. Operations on the **File** menu apply to the current file. If you have multiple files loaded, you can use the **File:** drop down to change the current file.

Each loaded file has an sequence number, in the order in which they are loaded; this number is displayed in the **File:** field and drop-down list and in the **File** column in the process list.

In Single File mode, the current file is the only file that is displayed. In Multiple File mode, all files are displayed. This is toggled using the **Main** menu item **Single File Mode**.

In multiple file mode, you can see which file is associated with a given process in the **File** column. The following example shows three sequential sample sets on the same stone and topaz process.

The screenshot shows the VSD 5.5.0 (build 32a7ba25) interface. The main window displays a table of process statistics. The table has columns for StartTime, EndTime, File, Samples, ProcessId, SessionId, Type, and Name. The data is grouped by process and session. The 'File' column shows the file associated with each process. The 'Type' column shows the process type, and the 'Name' column shows the process name.

StartTime	EndTime	File	Samples	ProcessId	SessionId	Type	Name
08/06 18:02:14	08/06 18:03:34	3	81	6328		Pgsvr	AioPgsvr3
08/06 17:58:50	08/06 17:59:47	2	58	6328		Pgsvr	AioPgsvr3
08/06 17:56:05	08/06 17:56:29	1	25	6328		Pgsvr	AioPgsvr3
08/06 18:02:14	08/06 18:03:34	3	81	9199	5	Gem	Topaz
08/06 17:58:50	08/06 17:59:47	2	58	9199	5	Gem	Topaz
08/06 17:56:05	08/06 17:56:29	1	25	9199	5	Gem	Topaz
08/06 18:02:14	08/06 18:03:34	3	81	6347	4	Gem	SymbolGem7
08/06 17:58:50	08/06 17:59:47	2	58	6347	4	Gem	SymbolGem7
08/06 17:56:05	08/06 17:56:29	1	25	6347	4	Gem	SymbolGem7

Below the table, a list of metrics is shown, with 'CommitCount' selected. The metrics are:

- CodeGenGcCount
- CommitCount**
- CommitRecordPageReads
- CommitRecordsReadAbortCount
- CommitRecordsReadCommitBeforeCommitQueueCount

The 'Chart' field is empty. There are 'Add Line' and 'New Chart' buttons at the bottom.

## Chart with Statistics from Multiple Files

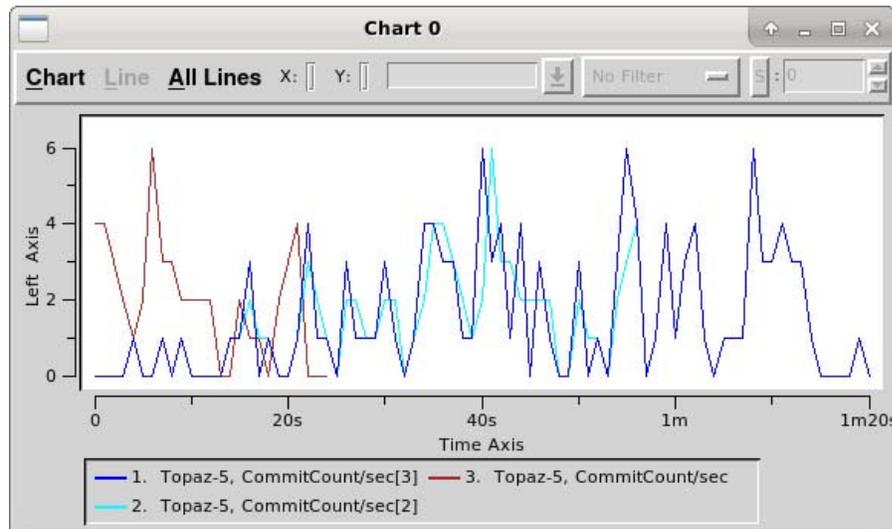
When you have multiple files loaded separately, the legend in the Chart Windows indicates which file the statistics values originated with by appending a [N] to the statistic name for each file, except File 1. The N corresponds to the File list index sequence number of that file, no number indicates it is from file 1.

## Comparing statistics for multiple sequential runs

If you are testing the impact of changes in code or configuration values, you may have multiple sets of statmonitor data that represent similar operations, that are offset in time. You can view these data sets overlaid by using relative time stamps, so each data set is charted with the same starting point.

The following chart includes the CommitCount for the Topaz process in each of the three files. Since Absolute Time stamps is unchecked, they all are displayed as if the starting time was the same.

Note the index following the process name in the legend that indicates which file each line is based on.



## 2.6 Monitoring

VSD can be used for ongoing monitoring of a active production system, allowing an administrator to see potential problems and correct them before they affect performance.

Monitoring a system, using **Main > Monitor...**, is described under “Monitoring statistics on a running GemStone server” on page 16. However, since production systems should be running statmonitor in any case, it is more efficient to use VSD to read the production statmonitor data files as they are being written.

### Auto update

Using the menu item **File > Auto Update** when or after loading a file will tell VSD to load new records in the statmonitor data file as they are written.

### See only the most recent statistics

Monitoring is not concerned with a view of system behavior over time; the most recent statistics are all that are relevant. Using the Rolling Left Trim feature, only the most recent statistics are displayed. For more on this features, see “Rolling Left Trim” on page 27.

## Auto Append of newly created files

To make statmonitor data files easier to manage, most production systems will automatically start writing to new statmonitor files periodically. There are a number of **statmonitor** features that support this, such as **-r**, **-R**, and **-k**.

The **File > Auto Append Next File** will open any new statmonitor files that appear in the same directory as the original statmonitor file, and append these statistics to the current view. This allows monitoring to continue seamlessly when statmonitor data files roll over.

Note that this feature cannot detect if the files are from the same sequence of statmonitor data files, or from an unrelated system or series. It is not recommended to mix statmonitor files from different sequences in the same directory.

## Font colors indicating live processes

While a statmonitor data file is being updated on disk, and VSD is automatically updating its view based on the disk file, the font of the "Samples" number for that process will be shown in green font.

When that process is no longer having samples recorded, either because the file is no longer being updated or that particular process exited, after a timeout the Samples count will be shown in regular black font. This allows you to detect if a process has exited unexpectedly or if something has gone wrong with the monitoring itself.

The timeout by default is 90 seconds. If the statmonitor interval is similar or larger than this, the process may be marked as inactive although it is functioning as expected. You can adjust the timeout using **Main > Settings > Change Inactive File Timeout**.

## 2.7 Aliases for statistic names

When GemStone software generates statmonitor data, each statistic is recorded with a meaningful name, such as "CommitCount". GemStone also provides application-defined statistics, which have non-meaningful names: for Gems, SessionStat00... SessionStat47, and for the Stone, GlobalStat00..GlobalStat47. These can be programmatically updated by customer applications to track application-specific information, and may be used differently by different application processes.

Some GemStone processes, such as markForCollection, reclaim, and backup, use Session statistics to track internal progress. While this is designed for GemStone Engineering analysis, this may be helpful in understanding your system.

Since analysis is easier with meaningful names, there are several kinds ways to manage aliases in VSD:

- ▶ Global statistics rename. This is useful if you find a particular name problematic, but is not appropriate for Session Statistics, which have different meanings in different Gem processes.
- ▶ Custom aliases can be defined that are automatically applied to Gem or Stone processes, both SessionStat\* statistics and to other statistics. These names appear in parenthesis following the statistics name.
- ▶ Automatic system aliases. These are automatically applied, for example, to the ReclaimGem and AdminGem, and provide aliases for the SessionStat\* statistics that

are used by that process. These names appear in parenthesis following the statistics name.

- ▶ Manually applied alias templates for processes running GemStone operations. You can manually apply templates to Gem processes that are running certain GemStone operations, such as markForCollection and backup, to alias the SessionStat\* statistics that are in use by that operation. These names appear in parenthesis following the statistics name.

## Global aliases

Statistics can be globally renamed in VSD, by including a directive:

```
sl_stat -alias existingStatName newName
```

in `.vsdconfig` or `.vsdrc`. This causes all statistics named *existingStatName* to be replaced by *newName* wherever it is used. *newName* should not be an existing statistic name.

## Process-specific alias names

VSD custom aliases allow you to associate a name, description, and other details that VSD uses for display, with an existing statistic name. This association is set by the name of the process (and, optionally, by GemStone version).

This is designed in particular to define meanings for SessionStat00.. SessionStat47. However, aliases can be applied to any Gem statistic.

The aliases are defined via commands that are included in a VSD support file, either `.vsdconfig` or `.vsdrc`. To update an alias, you must shut down VSD, edit the file, and restart.

The `setCtrAlias` command has the following arguments:

```
setCtrAlias CacheNamePattern StatName Alias Type Filter Units Description
[versionPatterns [writeToStartupFile]]
```

where the arguments are:

*CacheNamePattern* - cache name of processes to match, case sensitive

*StatName* - name of an existing cache statistic

*Alias* - name of the alias

*Type* - one of counter, counter64, uvalue, uvalue64, svalue, svalue64.

*Filter* - one of none, persecond or persample.

*Units* - Unit of measure for the statistic.

*Description* - A brief description of the statistic.

*versionPatterns* (optional) - \* for all GemStone versions, or a list of GemStone version numbers that the alias applies to. This is primarily intended for the system aliases described starting on page 35.

*writeToStartupFile* (optional) - Boolean (1 or 0) indicating if the alias will be written to `.vsdrc`. 1 means write the alias definition to `.vsdrc`, 0 means do not write to `.vsdrc`.

For example, if you have a process that acts as a watchdog to kill misbehaving sessions, your code can programmatically update session stat 0 with the number of sessions killed. To view the statistics in VSD, you can add the alias 'NumSessionsKilled' to SessionStat00.

Defining the following alias in `.vsdrc`:

```
setCtrAlias WatchDog SessionStat00 NumSessionsKilled counter64
none {The number of sessions killed} * 1
```

In VSD, when you load the `statmonitor` file or `monitor`, it will display the statistic name for the `WatchDog` process as:

```
SessionStat00 (NumSessionsKilled)
```

The Statistics Information window will provide the description.

Process name matching can include a `*`, which will match any portion of a string or nothing. However, use the `*` with caution, since process names with or without a `*` are treated as unique patterns, and only the one that appears first will be used. Once a process name is matched to `WatchDog`, `WatchDog*`, or `*`, other potential matches to the process name are ignored, and the statistics will not be applied.

## Managing aliases in VSD support files

You may set aliases in either `.vsdconfig` or `.vsdrc`. If you may run older versions of VSD on this node that will access the same `.vsdconfig` and `.vsdrc`, you should ensure that the definitions are conditionalized:

```
if {[llength [namespace which setCtrAlias]]} {
    definitions of aliases
}
```

If the alias is defined in `.vsdconfig` with `writeToStartupFile` as 0 (false) or if **Save Settings On Exit** is false, then the alias definition will not be written to `.vsdrc`, but will be read from `.vsdconfig`.

If you wish to add aliases to `.vsdrc`, either `writeToStartupFile` must be true, or **Save Settings On Exit** must be false, otherwise the definition will be lost next time VSD is shut down.

## Alias templates for system operations statistics

The code that implements important GemStone operations such as `reclaim` and `markForCollection` collect a number of session statistics, to provide additional information for GemTalk Engineering analysis. These are undocumented, and many of them are obscure in the absence of source code. However, there are some that may be useful or suggestive when analyzing problems or performance issues.

The VSD distribution includes `vsd.internal.stat.aliases.tcl`, which contains aliases for the session statistics internally produced by GemStone operations. These are defined per GemStone version, and the statistics definition for each process in specific versions are provided, in addition to the alias names.

Note that since internal session statistics are not a supported or documented feature, there may be inaccuracies or version-specific differences that VSD does not detect.

Version 3.2 and later:

- ▶ Reclaim and Admin Gem
- ▶ `markForCollection`
- ▶ `fullBackup` and `restoreFromBackup`
- ▶ page and object Audit
- ▶ `cacheWarming`

- ▶ logsender and logreceiver

Version 2.4.x:

- ▶ PageManager
- ▶ Reclaim and Admin Gem
- ▶ markForCollection and Fast FDC
- ▶ cacheWarming

32-Bit GemStone/S:

- ▶ PageManager
- ▶ GcGems, Reclaim and Epoch
- ▶ markForCollection

## Applying internal operation alias templates to application Gems

Aliases work by matching the pattern of the Gem name in the statmonitor data file. For example, the ReclaimGem has a name like GcReclaim6. Internal session statistics aliases are automatically applied to named sessions.

Processes such as markForCollection and restoreFromBackup are executed by ordinary Gems, which by default have the cache name GemN or Topaz. If you know the session that ran the given operation, you can apply the alias by name.

To do this, in the Main window, select the session, and use the right menu button pop up menu. **Apply Alias Template For Type >** provides a list of templates that you can select from to apply to the selected Gem process.

You may also set the cache name to exactly match the template definition. For example, for an MFC session, execute:

```
System cacheName: 'MarkForCollection'.
```

And the MarkForCollection template will be automatically applied when the statmonitor file is loaded.

## Removing Aliases

The Main Window process pane pop up menu item Remove Aliases is enabled if the process has an alias, either a custom alias from .vsdconfig or .vsdrc, or an applied system alias. Using the Remove Aliases removes either type of alias.

Removing custom aliases does not affect writing these alias to .vsdrc on shutdown; removes in this way are transient.

A system alias can be reapplied using the **Apply Alias Template For Type** menu. You cannot reapply custom alias; you will need to restart VSD to retrieve these aliases.

## 2.8 Setting Background Color

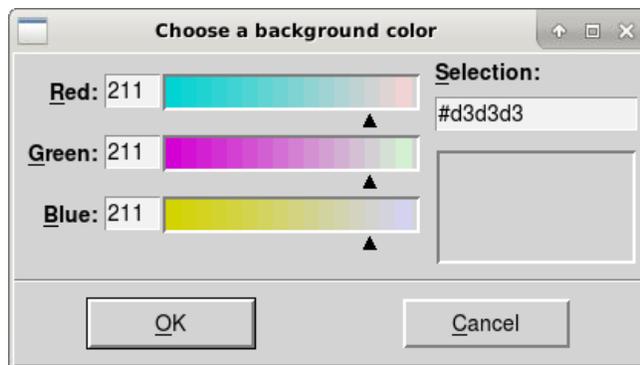
When running multiple instances of VSD, it can be helpful to be able to distinguish the VSD charts belonging to different instances. To support this, you may set the background color of the VSD instance. Once selected, all VSD windows for that VSD instance will display this color in all windows and charts.

Setting a master background color changes the window background, which ordinarily is a light gray, and also changes the background color for process lists, fields, and graphs which ordinarily have a white background.

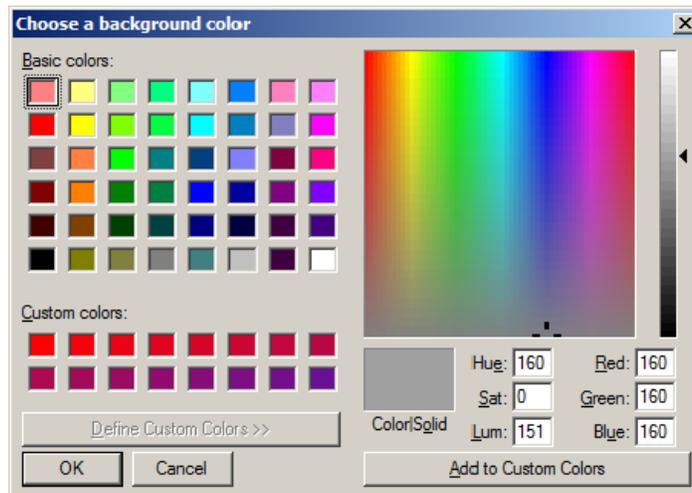
The color may be set in a number of ways:

- ▶ The **Choose a background color** dialog, which can be opened using either by <control>b or by using the main window menu item **Main >> Set Master Background Color...**

On UNIX platforms:



On Windows:



- ▶ Using the **-b** argument to the vsd executable. In addition to accepting a RGB hex color value, you may also use a TK symbolic color name.

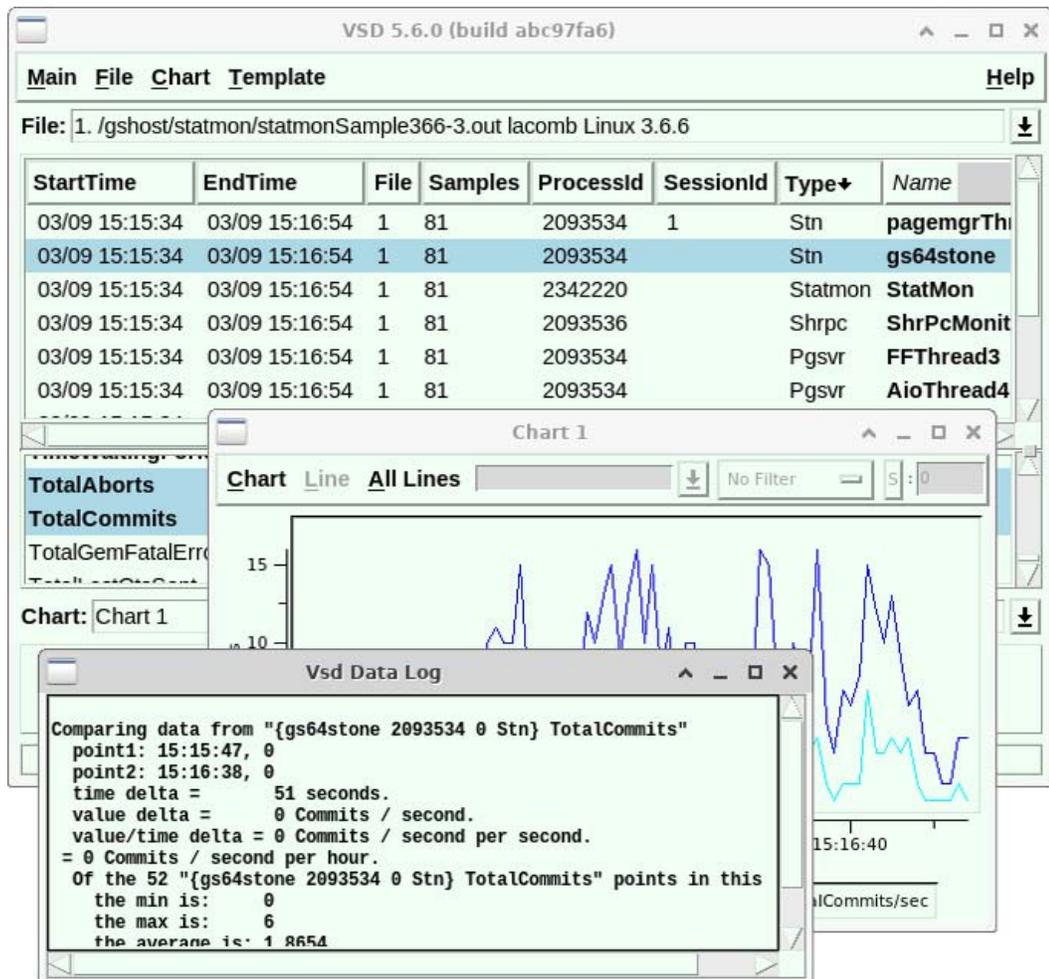
TK color names and RGB values are provided here:

<http://www.tcl.tk/man/tcl8.6/TkCmd/colors.htm>

TK provides over 750 color names describing over 500 unique colors.

- ▶ Setting the environment variable VSD\_MASTER\_BG\_COLOR to a hex RGB value or to a symbolic color name.

After setting a background color, for example:



Note that the background color is **not** saved to the .vsdrc file, since it is intended as a way to distinguish different instances of VSD on the same machine, rather than as a configured color for VSD in general.

## 2.9 Menu Options

This section describes the menu items available on the Main and Chart Windows in VSD.

### Main Window Main Menu

<b>Load Data File(s)...</b>	Load one or more data files. Each file is loaded independently.
<b>Append Data File(s)...</b>	Load one or more data files. If more than one file is selected, merge the data in all files into the current file if a file is already loaded, or into the first file in the list. For correct operation, the files must represent sequences in the same repository and be loaded in sequence.
<b>Monitor...</b>	Open a dialog to display the process and statistics for a running stone. This invokes statmonitor with the selected options, and reads the resulting file with automatic updates. Not available on Windows.
<b>Change Directory...</b>	Change the current VSD directory. This is useful when monitoring, to set the directory in which the statistics file will be written.
<b>Change Time Offset...</b>	Set the offset between the current system time zone and the time zone of the data file.
<b>Set Master Background Color...</b>	Open a dialog to set the background window color.
<b>Copy Selection</b>	Copy the selected processes and statistics to the VSD clipboard.
<b>Show Statistics Info</b>	Open the Statistics Information Window
<b>Select Instances By...</b>	Allows you to: <ul style="list-style-type: none"> <li>▶ <b>All</b> selects all processes.</li> <li>▶ <b>Clear</b> unselects all selected processes</li> <li>▶ <b>Name...</b> opens a dialog to let you enter a string, and selects all process with names that match. Supports wild cards * and ?, and case-sensitive or insensitive.</li> <li>▶ <b>Process ID...</b> opens a dialog to let you enter a process id, and selects all process with that process Id.</li> <li>▶ <b>Session ID...</b> opens a dialog to let you enter a session id, and selects all process with that session Id.</li> <li>▶ <b>Statistic</b> selects all processes that define the currently selected statistic or statistics.</li> <li>▶ <b>Type</b> selects all processes that are the same type as the currently selected process or processes.</li> </ul>
<b>Hide Instances by Type...</b>	Allows you to hide specific process types in the Process list.
<b>Combine</b>	When selecting multiple processes, combine the statistics data for all selected process into a single line.

<b>No Flatlines</b>	Exclude statistics that do not change.
<b>Single File Mode</b>	When this is checked, the display includes only the currently selected file. When not checked, the statistics for all open files are displayed together.
<b>Absolute Timestamps</b>	If checked, timestamps are in real time; if not checked, timestamps are relative to the beginning of the data file.
<b>Copy Referenced Lines</b>	In a chart window, when performing line arithmetic, make copies of the lines rather than using the lines values directly.
<b>Settings</b>	Allows you to set or unset the display options: <ul style="list-style-type: none"><li>▶ <b>Change Inactive File Timeout</b> sets the period of time that VSD waits before determining that no new samples will be added to a given statistic and changing the font color to inactive.</li><li>▶ <b>Choose Fonts...</b> contains submenu items that open a Font Dialog, allowing you to select fonts for the <b>Main Window</b>, <b>Charts</b>, <b>Chart Data Logs</b>, and <b>Help Text</b>.</li><li>▶ <b>Confirm Exit</b> determines if VSD will prompt you before exiting.</li><li>▶ <b>Save Settings On Exit</b> determines if VSD will write the current settings to .vsdrc on exit.</li><li>▶ <b>Show Commas In Numbers</b> allows you to see numeric values over 9999 with a comma.</li></ul>
<b>Save Settings</b>	Save the current configuration to .vsdrc. This is useful when <b>Save Settings On Exit</b> is false.
<b>Exit</b>	Exit VSD.

## Main Window File Menu

<b>Info...</b>	Open a dialog providing information about the current file.
<b>Update</b>	Check for additional data samples in the current disk file, and update the display with changes.
<b>Update All</b>	Check for additional data samples in the disk files for all open statistics file, and update the display with changes.
<b>Auto Update</b>	When this is checked, the system will automatically check for additional data samples in the current disk file, and update the display with added entries.
<b>Auto Append Next File</b>	When this is checked, the system will automatically check for new files in the same directory as the original file (regardless of the name, and automatically open and append these samples to the existing view. Checking this option also checks the <b>Auto Update</b> item. This option is not available on Windows.

<b>Enabled</b>	Used when in multiple file mode (that is, <b>Main &gt; Single File Mode</b> is not checked); the statistics for this file will not be displayed.
<b>Untrim Left</b>	Remove left trims in all charts.
<b>Untrim Right</b>	Remove right trims in all charts.
<b>Set Rolling Time Window...</b>	Brings up a dialog to allow you to set a time period. When monitoring a live system, a chart will only show the values collected up to the give time period before the current moment.

## Main Window Chart Menu

<b>Add to Chart</b>	Add the currently selected process statistics to the most recently selected chart.
<b>New Chart...</b>	Open a new chart with the currently selected process statistics.
<b>Show Legend</b>	Set the default for displaying the legend with the process and statistic name for the various lines at the bottom of the chart window.
<b>Time Format</b>	Set the default time format. Options are Seconds, Hour:Minute:Second (the default), or Month/Day Hour:Minute:Second. Note that if <b>Main &gt; Absolute Timestamps</b> is not checked, all formats are in relative to the first sample in the file.
<b>Default Line Style</b>	Set the line style; options are linear (the default), step, natural, and quadratic.
<b>Show Time Axis Title</b>	Display the label on the horizontal axis, generally the name of the data file
<b>Show Left Axis Title</b>	For statistics graphed on the left axis, display the label on the left side vertical axis, the name of the units for the selected statistic.
<b>Show Right Axis Title</b>	For statistics graphed on the right axis, display the label on the right side vertical axis, the name of the units for the selected statistic.
<b>Show Current Values</b>	Include, on the menu bar, display areas for the X and Y values of the point under the cursor on the selected statistic.
<b>Show Min and Max</b>	Include, on the menu bar, display areas for the maximum and minimum values of the selected statistic
<b>Show Line Stats</b>	Display a line indicating the number of points, minimum, maximum, mean, and standard deviation for the selected statistic, in the upper left of the graph area.
<b>Show CrossHairs</b>	display cross hairs centered on the cursor location.

<b>Show Grid Lines</b>	display grid lines.
<b>Selected Line Color...</b>	Set the color to be used for the currently selected line in all charts.
<b>Close All Charts</b>	Close all open chart windows.

## Main Window Template Menu

<b>Reload Template File</b>	Reread the template file from disk.
<b>New Template Chart &gt;</b>	Select from the templates and open a new chart using that template
<b>Templates Use Selection</b>	If selected, for templates that have patterns using *, open templates using only the selected processes.

## Main Window Help Menu

<b>How to...</b>	Open Help with list of how-to topics.
<b>Main Window...</b>	Open Help for the Main window.
<b>Chart Window...</b>	Open Help for the Chart window.
<b>All Topics...</b>	Open Help with list of all topics.
<b>All Help Text...</b>	Open Help with all help in a single window.
<b>About VSD...</b>	Open About VSD dialog.

## Process List Pane Pop-up Menu

<b>Search...</b>	Search for a process using name or part of a name.
<b>Apply Alias Template For Type</b>	Allows you to apply a predefined set of aliases to a process that is performing a specific GemStone function, such as markForCollection, in order to view session stats associated with that function.
<b>Remove Alias</b>	Removes an applied alias template from the selected process.
<b>Hide Instances by Type...</b>	Allows you to hide specific process types in the Process list.
<b>Select Instances By:</b>	Allows you to: <ul style="list-style-type: none"><li>▶ <b>All</b> selects all processes.</li><li>▶ <b>Clear</b> unselects all selected processes</li><li>▶ <b>Name...</b> opens a dialog to let you enter a string, and selects all process with names that match. Supports wild cards * and ?, and case-sensitive or insensitive.</li><li>▶ <b>Process ID...</b> opens a dialog to let you enter a process id, and selects all process with that process Id.</li><li>▶ <b>Session ID...</b> opens a dialog to let you enter a session id, and selects all process with that session Id.</li><li>▶ <b>Statistic</b> selects all processes that define the currently selected statistic or statistics.</li><li>▶ <b>Type</b> selects all processes that are the same type as the currently selected process or processes.</li></ul>
<b>Combine</b>	When selecting multiple processes, combine the statistics data for all selected process into a single line.
<b>No Flatlines</b>	Exclude statistics that do not change.
<b>Single File Mode</b>	When this is checked, the display includes only the currently selected file. When not checked, the statistics for all open files are displayed together.
<b>Absolute Timestamps</b>	If checked, timestamps are in real time; if not checked, timestamps are relative to the beginning of the data file.

## Chart Window Chart Menu

<b>Add from Template</b>	Open the list of templates and add the process statistics specified by selected template to the current window.
<b>Save Template</b>	Create a new template from the current set of processes and statistics. The prompt for the name is in the menu bar area. Note that the resulting template may need further editing if Gem processes are selected.

<b>Paste</b>	Paste a line that was previously put into the VSD clipboard, and clear the clipboard.
<b>Snapshot</b>	Make a .gif of the chart, and allow you to select a directory and name for the resulting file.
<b>Help...</b>	Open Help for the chart window
<b>Zoom In</b>	Select a rectangle and make this the new graph display area.
<b>Zoom Out</b>	Undo the most recent zoom.
<b>Compute Scale All</b>	Compute a scale for every line in the chart.
<b>Unscale All</b>	Remove any scaling done to any lines.
<b>Show Legend</b>	Display the legend showing the process and statistic names for the various lines at the bottom of the chart window.
<b>Time Format</b>	<p>Set the time format. Options are Seconds, Hour:Minute:Second (the default), or Month/Day Hour:Minute:Second.</p> <p>If <b>Main &gt; Absolute Timestamps</b> is checked, Seconds provides the GemStone timestamp value; other values provide the sample's time, adjusted by a value set by the Main windows' <b>Change Time Offset</b>.</p> <p>If <b>Main &gt; Absolute Timestamps</b> is not checked, the display is either in seconds since the beginning of the file, or <i>mmhmmms</i>, or <i>mmddmmhmmms</i> since the beginning of the file.</p>
<b>Show Time Axis Title</b>	Display the label on the horizontal axis, generally the name of the data file
<b>Show Left Axis Title</b>	For statistics graphed on the left axis, display the label on the left side vertical axis, the name of the units for the selected statistic.
<b>Show Right Axis Title</b>	For statistics graphed on the right axis, display the label on the right side vertical axis, the name of the units for the selected statistic.
<b>Show Current Values</b>	Include, on the menu bar, display areas for the X and Y values of the point under the cursor on the selected statistic.
<b>Show Min and Max</b>	Include, on the menu bar, display areas for the maximum and minimum values of the selected statistic
<b>Show Line Stats</b>	Display a line indicating the number of points, minimum, maximum, mean, and standard deviation for the selected statistic, in the upper left of the graph area.
<b>Show CrossHairs</b>	display cross hairs centered on the cursor location.
<b>Show Grid Lines</b>	display grid lines.
<b>Close</b>	Close this chart window.

## Chart Window Line Menu

<b>Log Info</b>	Add information on the currently selected statistic or statistics to the log window.
<b>Log Delta</b>	Allows you to select a start and end point on any lines, and adds information on the delta between these two points to the log window.
<b>Compute Scale</b>	Compute a scale for this line.
<b>Unscale</b>	Remove any scale on this line.
<b>Graph on Left Axis</b>	When selected, the line is graphed against the axis on the left side of the chart; if unselected, graph against the axis on the right side.
<b>Symbol</b>	Set the symbol to use to indicate points on the line; options are none (the default), scross, splus, square, triangle, circle, diamond, plus, or cross.
<b>Style</b>	Set the line style; options are linear (the default), step, natural, and quadratic.
<b>Update</b>	Reread the data file associated with this line for new entries.
<b>Add Lines</b>	Add the next line selected to the selected line.
<b>Diff Lines</b>	Subtract the next line selected to the selected line.
<b>Divide Lines</b>	Divide the next line selected by the selected line.
<b>Change Integer Type</b>	Allows you to convert the interpretation of a 32-bit statistic to 64-bit.
<b>Normalize</b>	If this item is checked then the line will be normalized. Cannot be used with a line offset.
<b>Trim Left</b>	Allows you to select a point on the selected line, everything to the left of this point (earlier) will be hidden.
<b>Trim Right</b>	Allows you to select a point on the selected line, everything to the right of this point (later) will be hidden.
<b>Untrim Left</b>	Undo any left trims done to this chart.
<b>Untrim Right</b>	Undo any right trims done to this chart.
<b>Copy</b>	Copy the selected line (process and statistic) to the VSD clipboard.
<b>Cut</b>	Remove the selected line, and add it to the VSD clipboard.
<b>Delete</b>	Remove the selected line.

## Chart Window All Lines Menu

<b>No Filter</b>	Remove any filter from all lines on the chart; display the values exactly as they are in the statmonitor text file.
<b>PerSample</b>	Apply the per-sample filter for all lines on the chart; displays the difference between two consecutive samples of the statistic.
<b>PerSecond</b>	Apply the per-second filter for all lines on the chart; displays the difference between two consecutive samples of the statistic, divided by the number of elapsed seconds between the samples.
<b>Aggregate</b>	Apply the aggregate filter for all lines on the chart; displays a running total for the per-sample deltas of the statistic.

# VSD Template and configuration files

---

The first time you use VSD, the executable creates default template files and configuration files in the VSD home directory. These files are used on subsequent executions; this allows it to remember window sizes and locations, aliases, and similar details.

You can delete these file at any point, and they will be recreated from the default the next time you use VSD.

The VSD home directory is, by default, `/home/username/` on UNIX and `c:\users\username\` on Windows. If you set the environment variable `$VSDHOME` or `%VSDHOME%` in the environment before starting the executable, this directory is used instead of the user's home directory.

## 3.1 VSD Configuration Files

VSD provides a number of display options, which are set by menu items in `vsd`. These options are recorded in platform-independent text files that are located the VSD home directory.

There are two files that together manage the parameters that control what you see on startup: `.vsdconfig` and `.vsdrc`.

### **`$VSDHOME / %VSDHOME%`**

The VSD home directory is, by default, the home directory of the user starting VSD. You may define an environment variable, `VSDHOME`, to specify a directory. This will override the default location for the configuration files.

To determine the exact locations of the files that were read by running VSD, see the About VSD window, opened from Main Window **Help > About VSD...** menu item. The **Files:** section includes the full path to all files associated with the currently running VSD.

## .vsdconfig

This file, by default, contains a small number of system definitions. This file is created by VSD if it does not exist; VSD never writes updates to this file.

You may use a text editor to add settings. Any settings in this file override the same definition in `.vsdrc`. This provides another level of customization for the various default settings. For example, if you set

```
set vsd(wv:statmonargs) {-d $GEMSTONE/statmons}
```

It ensures that the monitor window will always default to write data files in a specific directory, regardless of what is in `.vsdrc` or what was done during in previous VSD monitoring.

If you delete the `.vsdconfig` file, then the next time VSD is started, the file will be recreated with default values.

## .vsdrc

This file is used by VSD to record modifications to the display options, including menu selections and window locations and sizes. VSD reads this file when it starts up, and by default, writes the file when it exits, so the next time you open VSD it will have the same settings.

If you delete the `.vsdrc` file, VSD will start with defaults from `.vsdconfig` or defaults from the VSD executable. When VSD subsequently exits, the `.vsdrc` file will be recreated with the current settings from the executable.

Note that VSD reads and applies settings specified in both `.vsdconfig` and `.vsdrc`. The setting in `.vsdconfig` overrides the same setting read from `.vsdrc`. If VSD writes `.vsdrc` on exit, it will update the values written to `.vsdrc` with the values that were read from `.vsdconfig`. If you configure VSD to not write the `.vsdrc` on exit, you may configure VSD at both levels, `.vsdconfig` and `.vsdrc`, without being affected by transient changes when VSD is executed. This provides additional flexibility in managing general settings such as aliases as well as individual settings such as preferred window layouts.

If you would like VSD to always start with a fixed set of display options, rather than continuing with the settings from the previous VSD execution, you can disable saving changes on exit. This is done by unchecking the Main Window's menu item **Main > Save Options On Exit**. To make this setting persistent, you must also execute the **Main > Save Options**.

## Configuration file format

In most cases, you can configure your default display by taking advantage of the natural way VSD saves your settings to `.vsdrc` and reads them on restart. By doing this, you get a consistent display without any extra work.

However, if you want to set a configuration that is always used, regardless of what you do during any particular VSD session, you will need to edit the VSD configuration file `.vsdconfig`. The configuration file format is the same in `.vsdconfig` and `.vsdrc`, so you may copy lines between files.

The following sections document many of the relevant configuration parameter names and meanings.

Configuration options are specified in a format that is specific to VSD/TCL, and are not intended for general use by customers.

To include a comment line, prefix the line with #.

Incorrect configuration names and settings encountered in reading the files are ignored; default values will be used in this case.

Configuration options have the format:

```
set vsd (configurationOption) value
```

for example:

```
set vsd(confirmonexit) 1
```

Acceptable *value* arguments depend on the parameter. Many are Integers, but some parameters are strings. Lists are indicated by {}. Boolean statistics are specified as 0 or 1.

## Main Window options

The following statistics relate to the main VSD window.

Parameter name	Type	Default	Description
confirmonexit	Boolean (0 or 1)	1	Whether to confirm before exiting VSD, <b>Main &gt; Confirm Exit</b>
def:sortbytype	Column label	Type	On file load, the column to use to sort processes.
def:sortdecreasing	Boolean (0 or 1)	1	On file load, the direction of sort for the process list; 0 is decreasing, 1 is increasing.
combineFiles	Boolean (0 or 1)	0	Whether to combine values for all selected processes in all open files when displaying a statistic; <b>Main &gt; Combine Across Files</b>
combineInstances	Boolean (0 or 1)	0	Whether to combine values for all selected processes when displaying a statistic; <b>Main &gt; Combine</b>
filelist	list of strings	empty list	List of filenames displayed in the Main Window File: dropdown.
noFlatlines	Boolean (0 or 1)	0	Whether to show statistics for a process if all values for that statistic are zero; <b>Main &gt; No Flatlines</b>
showCtrInfo	Boolean (0 or 1)	0	Whether to have the Statistics Information window open. <b>Main &gt; Show Statistics Info</b>

Parameter name	Type	Default	Description
singleFileMode	Boolean (0 or 1)	1	Whether working with multiple data files is allowed. <b>Main &gt; Single File Mode</b>
templatesUseSelection	Boolean (0 or 1)	0	Whether the template applies only to the selected processes.

## Chart Window Display Options

The following parameters describe the defaults for including and omitting information from chart windows.

Parameter name	Type	Default	Description
absoluteTSMode	Boolean (0 or 1)	1	Whether to use absolute timestamps for charts. <b>Main &gt; Absolute Timestamps</b>
activecolor	string	red	The color of the selected line on a chart. <b>Chart &gt; Selected Line Color...</b>
def:linestyle	linear   step   natural   quadratic	linear	How the line is rendered on charts; <b>Chart &gt; Default Line Style</b>
def:showcrosshairs	Boolean (0 or 1)	0	Whether to show cursor crosshairs on charts, <b>Chart &gt; Show CrossHairs</b>
def:showcurxy	Boolean (0 or 1)	1	Whether to show current X and Y values on a chart; <b>Chart &gt; Show Current Values</b>
def:showgridlines	Boolean (0 or 1)	0	Whether to show gridlines on charts, <b>Chart &gt; Show Grid Lines</b>
def:showlegend	Boolean (0 or 1)	1	Whether to show the legend (the process and statistic name and line color for the chart lines) on charts; <b>Chart &gt; Show Legend</b>
def:showlinestats	Boolean (0 or 1)	1	Whether to show text with points, min, max, mean, and standard deviation on graph; <b>Chart &gt; Show Line Stats</b>
def:showminmax	Boolean (0 or 1)	0	Whether to show max and min values on chart menu bar; <b>Chart &gt; Show Max and Min</b>
def:xaxis:showtitle	Boolean (0 or 1)	1	Whether to display label on x axis, <b>Chart &gt; Show Time Axis Title</b>
def:xformat	date   time   elapsed	time	How to display the time on the x-axis of a chart; <b>Chart &gt; Time Format</b>

Parameter name	Type	Default	Description
def:yaxis:showtitle	Boolean (0 or 1)	1	Whether to display y axis label on left side, <b>Chart &gt; Show Left Axis Title</b>
def:y2axis:showtitle	Boolean (0 or 1)	1	Whether to display y axis label on right side, <b>Chart &gt; Show Right Axis Title</b>

### Monitor Window Options

The following parameters apply to the Monitor window, describing options for how to start statmonitor.

Parameter name	Type	Default	Description
wv:statmoninterval	integer	20	Monitor window, default statmonitor sample interval
wv:statmonflush	integer	0	Monitor window, number, default statmonitor flush interval
wv:statmonargs	list of options	empty list	Monitor window, the list of statmonitor arguments.
wv:statmoncompress	Boolean (0 or 1)	0	Monitor window, default for if statmonitor should write compressed file

### Window locations and display

The location, size, and proportion of the Main Window, Chart window, Statics info, and other dialogs are recorded.

`def:geometry` sets the size and location of the main window.

`instancePaneSize` sets the height of the Process pane in the main window.

`statisticPaneSize` sets the height of the Statistics pane in the main window.

`geometry:.datawin` sets the size and location of the chart window.

`geometry:.ctrinfo` sets the location of the statistics information window.

There are a number of other dialog that have similar settings.

## 3.2 VSD Chart Templates

VSD templates specify a set of processes and statistics that together represent a useful picture of one behavior of the system. VSD is shipped with a set of predefined templates. As described under "Using and Creating Templates" on page 28, you can open charts on these templates, and create your own templates based on charts you have constructed.

To further customize the templates you use, you may want to edit the templates directly. The templates are stored in a text file using specific syntax. This section provides information you need if you wish to edit the template files directly.

## VSD Templates File

The default templates are in the file `.vsdtemplate`, which is written to the VSD home directory if it does not already exist in that location.

If you delete the `.vsdtemplates` file, then the next time VSD is started, the file will be recreated with default values.

## Template Syntax

The template files are written in a formatted form that is understood by the TCL code that supports VSD.

Templates in the `.vsdtemplate` have the following form:

```
set vsdtemplates(TemplateName) {
  {Type {ProcessName} StatName scale filter axis}
  {Type {ProcessName} StatName scale filter axis}
  ...
}
```

For example:

```
set vsdtemplates(Garbage) {
  {Stn {*} PagesNeedReclaimSize 1 none y2}
  {Stn {*} ReclaimedPagesCount 1 none y2}
  {Stn {*} EpochGcCount 1 persecond y}
  {Stn {*} ReclaimCount 1 persecond y}
  {Gem+ {*Gc*} CommitCount 1 persecond y}
}
set vsdtemplates(CacheTooSmall) {
  {Shrpc ShrPcMonitor FreeFrameCount 1 none y2}
  {+ {*} FramesFromFreeList 1 persecond y}
  {+ {*} FramesFromFindFree 1 persecond y}
}
```

The data for each line in each template includes:

**type** – One of the following:

```
Stn
Shrpc
Gem
Pgsvr
```

'+' can be used, to mean that all processes that match this line spec will be combined into a single line. + can be used either standalone, meaning all process types, or added to a type name. For details on types, see page 19.

**name** – Identifies the specific process by process name.

'\*' can be used to match multiple process names.

'!' can be used to exclude specific processes by name. For example,

```
{Stn {!pagemgrThread} ...
```

matches all Stone processes whose name does not include 'pagemgrThread'.

For more on process names, see page 19.

**stat** – The name of a specific existing statistic.

**scale** – The number of the scale for this particular statistic. It will usually be 1. For more details on scales, see “Scaling” on page 25.

**filter** – The name of the filter, which will be one of the following:

- none
- persecond
- persample
- aggregate

For details on these filters, see “Filters” on page 25

**axis** – Either y, the left axis, or y2, the right access. Units for this statistic line are labeled on the specified side y-axis.

Once you have created or edited your template file, save the `.vsdtemplates` file, and use the menu item **Chart > Reload Template File** to load this file. If you have made syntax errors, reload will report an error and the file will not load.

For more about VSD templates and the syntax and options available, see the VSD help under the topic “Template Syntax”.

### 3.3 Data in VSD TCL files

The VSD distribution includes a file containing information about all statistics that VSD expects to encounter. This includes statistics names and information that impacts how the values are handled in VSD.

For each statistic, there is a line of the form:

```
{STATNAME} {TYPE LEVEL UNITS IS_FROM_OS}
```

with the following information:

STATNAME is the name of the statistic.

TYPE is one of the following:

**counter**

32-bit counter. Counters are expected to start from zero or an arbitrary value and always increase. For counter types, the useful information is the changes between values over time, not the values themselves; for example, the number of times garbage collection has executed. Counter types have a default filter of per second. Also, for counter types "No flatlines" means the per-second value is zero; the unfiltered statistic value itself may be non-zero.

**counter64**

64-bit counter. Similar to counter, but used when the values can become extremely large.

**svalue**

Signed 32-bit value. If data values read from a file are larger than 32 bits, it is automatically promoted to a 64 bit.

**svalue64**

Signed 64-bit value.

uvalue

Unsigned 32-bit value. If data values read from a file are larger than 32 bits, it is automatically promoted to a 64 bit.

uvalue64

Unsigned 64-bit value.

float

Floating point value.

LEVEL is one of the following:

common

advanced

wizard

While the LEVEL field remains in the records, it has no effect in VSD v5.5.3 and later versions.

UNITS is a string that describes what the statistic measures (e.g., pages, operations, etc.). This is used to label the vertical axis.

IS\_FROM\_OS is "true" if the statistic comes from the operating system and is "false" if not.

## Effect of statistic type on combined lines

VSD allows you to perform arithmetic operations between statistic lines. For example, it is useful to select a number of Gems processes and show the combined values of a statistic, such as the CommitCount for all the ReclaimGem processes.

For non-counter types, adding or combining is always a simple sum, in which the values for each process that includes a sample value at the given time are added together.

When one of the processes ends before the time of the final sample for all the processes, however, this is handled differently for statistics of counter type that are designed to be viewed per-second.

For statistics of counter type, on a chart that shows combined values, for processes that ended before the end of the chart display, the combined line continues to include the value of the final sample for any process that ends. This keeps the per-second graphs correct. If you examine the unfiltered combined data, however, this will appear incorrect.

If this adjustment was not made for counters, it would produce invalid decreases for the total after a process ended. For example, the per-second time spent waiting for IO by a set of processes is not actually decreased when one of the processes ends.

# VSD Command-Line Syntax

This section presents the formal command syntax and the available command-line options.

```
vsd [ -b color ] [ -A | -a | -u ] [ dir ] [ file+ ]  
vsd -h | -v
```

- |                        |   |
|------------------------|---|
| <b>-a</b>              | Open and append (as if they were one file) one or more files listed on the command line.  |
| <b>-A</b>              | Open and append (as if they were one file) one or more files listed on the command line. In addition, enable continuous appending with auto-append mode.  |
| <b>-b</b> <i>color</i> | Set the master background color. Color may be a hex RGB value (e.g.: #d3d3ff) or a valid color name (e.g.: white). The list of recognized color names may be found at:<br><a href="http://www.tcl.tk/man/tcl8.6/TkCmd/colors.htm">http://www.tcl.tk/man/tcl8.6/TkCmd/colors.htm</a> |
| <b>-h</b>              | Displays a usage line and exits.  |
| <b>-u</b>              | Open one or more files as separate files. In addition, enable continuous appending with auto-append mode  |
| <b>-v</b>              | Print version information and exit.   |

*dir* is the name of a directory; VSD will set this as the working directory for the executable, and after startup, opens the Open Statmonitor File dialog on this directory.

*file* is the name of a statmonitor file to be loaded as part of startup. If the **-a** or **-A** argument is specified, multiple files will be appended, otherwise, they will be loaded separately.

## Order of Loading Multiple Files

Note that if you are loading multiple files using a wild card \*, these files will be loaded in the operating system sort order.

By escaping the command line \*, VSD can sort the files from oldest to newest, and load in that order.

For example, to load and append a set of files and allow auto-append of new files:

```
vsd -A $GEMSTONE/statmons/\*.gz
```

On windows, you can escape the path separator using the caret, for example:

```
vsd -a C:\statmons\stat^*.gz
```