

[Introduction and History](#)

[Goals](#)

[Issues with Current VSD](#)

[Desired features in JVSD v1.0](#)

[Future features for post 1.0 versions of JVSD](#)

[Documentation](#)

[Build and Usage Requirements](#)

[Statistics File Format](#)

[JVSD Features From VSD](#)

[Classic VSD Features](#)

[JVSD Functional Differences](#)

[Appendix](#)

1. Introduction and History

The Geode framework software creates statistical archive files logging collected metrics from various subsystems such as the number of puts for a region. The primary application to read the statistics files for over the last 10 years has been the Visual Statistics Display (VSD) tool which can read the sampled statistics data and can produce time-series/value graphs. The original design and usage of VSD was to address a need by consulting and engineering to do performance analysis of the commercial GemStone/S Smalltalk server and later the GemFire product line. The VSD tool has since grown out of its initial role as a consulting tool and is now used by support, users and engineering as a critical post-mortem tool to identify performance issue such as GC and along with other kinds of failures within a distributed system.

2. Goals

Replace the current TCL based VSD with a new Java based version (JVSD). The initial release will focus on features of VSD used for post-mortem analysis of performance issues and will be essentially a clone of the functionality in the current classic VSD.

3. Issues with Current VSD

- 3.1. 32 bit only binaries and difficult to setup (install issues)
- 3.2. Difficulty managing large files
- 3.3. Difficulty managing multiple files
- 3.4. Issues running on different platforms (ex. must use Windows XP compatibility mode and additional X11 application is needed on OSX)
- 3.5. Not using a modern UI design (ex. Copy/Paste on different menus)
- 3.6. UI issues (ex. selecting two points for log delta, selecting two lines to diff)
- 3.7. Non-native look and feel (Motif L&F on all platforms, not visually appealing)
- 3.8. Requires shipping a TCL interpreter
- 3.9. Possibly licensing issues in making VSD open source
- 3.10. Issues with finding developer talent with TCL skills

4. Desired features in JVSD v1.0

- 4.1. Meet critical features of legacy VSD
- 4.2. Minimize differences in initial usage and GUI from classic VSD
- 4.3. Open source and community developed

- 4.4. Split out library for users to independently process Statistics files for use by third party software or for monitoring applications
- 4.5. Provide test cases and data to validate functionality
- 4.6. Make new architecture amenable to extension and future features
- 4.7. Manage timezones correctly
- 4.8. Use human readable scales (nanoseconds ->milliseconds/seconds, bytes->megabytes)
 - 4.8.1. Needs to be selectable preference
- 4.9. Identify default or common statistics/metrics
 - 4.9.1. Filter stats list displayed by common/default, advanced, no flatlines

5. Future features for post 1.0 versions of JVSD

- 5.1. Enhance Performance
- 5.2. Best practices usage guide
- 5.3. Snap to sample
- 5.4. Increase abilities to manage a large number of files concurrently
- 5.5. Increase abilities to manage large files with potentially millions of samples
- 5.6. Provide a robust template system to enhance chart reuse
- 5.7. Simplify sharing of statistics and charts (email, web uploads...)
- 5.8. Annotate charts with notes
- 5.9. Have the possibility to merge two charts instead of having to create a third and add the metrics from the two charts
- 5.10. Change timezone without having to change the environment timezone to easier sync stats with logs
- 5.11. Provide an overview of the loaded gfs files with the gfs file header information and the start and end time of the stats file.
- 5.12. Save charts and state of application for fast restarts
- 5.13. Provide plug-ins to process statistics data for third party applications
- 5.14. Provide export to alternative file formats (csv, excel...)
- 5.15. Provide import for alternative file formats (csv, excel...)
- 5.16. Provide intelligence/rules to identify and highlight potential issues
 - 5.16.1. Consider Javascript or similar for rules language to make it easy to extend for other users, ideally an interface that permits usage of user preferred rules engine or scripting interface (JRules, Dools...)
- 5.17. Interface to live system via JMX
- 5.18. Log file analysis (match log file exceptions to times in statistics for events)
- 5.19. Add new chart type - timeline, showing when a process is up relative to others
- 5.20. Encourage outside development and customization (community building)
- 5.21. Mobile version (iOS, Android) stretch goal to ease system management
- 5.22. Web based version (possibly integrate with Pulse)
- 5.23. Support for compressed stats. (gz, zip, etc.)
- 5.24. Easy uploads to Pivotal support or other third parties.
- 5.25. Separate out the statarchiver for use without Geode or a distributed system.

6. Documentation

Classic VSD public documentation	http://www.gemstone.com/docs/tools/VSD/index.html
----------------------------------	---

Current VSD public documentation	http://gemfire.docs.pivotal.io/latest/userguide/index.html#gemfire_tools/vsd/chapter_overview.html
VSD analysis tips and tricks	http://gemfirexd.docs.pivotal.io/1.0.0/userguide/manage_guide/Topics/system_performance.html
Analysis advice	http://blogs.vmware.com/vfabric/2012/10/using-visual-statistics-display-to-analyze-gemfire-runtime-configuration-resources-and-performance.html
JavaFX (Interface)	https://docs.oracle.com/javase/8/javase-clienttechnologies.htm
Web based plotting tools (example/ideas from third parties)	https://plot.ly/feed/ https://zeppelin.incubator.apache.org/
Vernier tools and plotting (example/ideas from third parties)	http://www.vernier.com/products/software/ga/

7. Build and Usage Requirements

The build environment should follow common open source practices and hence use tools that support those principles. A secondary goal will be to minimize the usage of third party libraries outside of the JVM to minimize the requirements for users during installation and to reduce the knowledge requirements for new developers to begin contributing to the project.

- 7.1. JDK 8 - will target and support only JDK 8 given JDK 7 has been EOL'ed
- 7.2. JavaFX as primary UI library - currently ships with JDK and replaces Swing (no additional libraries)
- 7.3. Minimal third party libraries including not requiring either GemFire or Geode classes and where required will use the common and well known libraries
 - 7.3.1. JUnit will be used for basic testing
- 7.4. Maven - build and dependency management
 - 7.4.1. The Geode Gradle should be updated to drive the maven build used by JVSD, the maven build will remain separate to allow JVSD to be forked into a possible future independent project.
- 7.5. The development branch will be jvsdfx-mm and other branches will be retired or deleted. The goal will be to move this prototype forward into a production ready release.
 - 7.5.1. The specifics may change once pushed to the Apache hosted repository under Geode but this internal branch and prototype will be the seed for the new JVSD
- 7.6. Current internal prototype hosted in Apache Geode repository
 - 7.6.1. Currently in branch feature/GEODE-78
- 7.7. No classic VSD source code will be reused in the new JVSD, individuals with licenses for GemFire can use and compare VSD to address functional requirements for JVSD.

8. Statistics File Format

- 8.1. No file or format changes of the statistics will be required or expected for the 1.0 release

- 8.1.1. Dependency on Geode and tool will need to track any statistics file format changes within it
- 8.1.2. Will continue to use default gfs file extension for file filters to identify statistics files
- 8.1.3. Support for gzip files needs to be reviewed and determined if this is critical
- 8.2. Provide ability to export data in CSV format for third party usage

9. JVSD Features From VSD

Fundamentally the initial JVSD release will be a clone of the classic VSD in regards to basic functionality and features. The basic interface should remain similar for most users with a default table or data selection window which will be the primary interface to selecting which data to chart. This data selection window will continue in the new version with only minor updates in form or function to accommodate the change in interface libraries or modern GUI design style. This window will also remain a singleton interface element and home base for user to access other elements such as the charts created.

The chart window is still the primary method for visualizing the statistics data. As with the table window, there are no significant changes expected in how users will interact with the chart window and will continue to only draw time-series line graphs. All GUI enhancements in this release will only to simplify usage such as selecting a line or saving a snapshot.

9.1. GUI

- 9.1.1. Table Window (Statistics selection) - See Figure 5 prototype.
- 9.1.2. Chart Window - See Figure 6 prototype.
- 9.1.3. Log Info Window - a simple scrolling text window where data produced by "Log Delta" operation is displayed - No prototype yet.

9.2. Analysis tools - There are currently no analysis tools to aid in the processing or understanding of events within statistical files. The first release will not change this basic usage or workflow of correlating events via adding one or more data sets to a time series line graph/chart. But hooks to plug in processing will need to be added for future options.

10. Classic VSD Features

The following captures in text the basic functional features (menus and actions) of the classic VSD tool and provides a baseline of expected functionality for any replacement tool to meet. As JVSD is initially intended to be a clone of the classic VSD functionality, developers will use and reference the old version for comparison of a feature and to insure it works in a similar and familiar manner.

Features marked with a (*) will be targeted for inclusion in the first JVSD release.

- **Basic Features of Table Window:** See Figure 1.
 - Load Data File * - open file
 - Monitor - Smalltalk usage to connect to a statsmonitor process
 - Append Data File * - Add new data to currently active data set based on PID, concatenate file to currently open files
 - Statistic Info * - Help screen which defines statistics data types, can set specifics about stat such as default sample style/filter (per Second, no filter...) and the statistics level (common, advanced). See Figure 4.
 - New Chart *

- Close All Charts *
- Close Chart *
- Print Chart
- Add to Chart / Add Line *
- Sort table data via column (PID, Statistic Type...)*
- Help *
 - How To *
 - Main Window *
 - Chart Window *
 - All Topics
 - All Help text
 - Choose Text Font
 - About VSD *
- **Chart Options in Table Window:**
- Settings for default creation of new charts or adding data/lines to an existing chart.
 - Show Legend *
 - Time Format *
 - Elapsed Time Seconds *
 - Hour:Minute:Second *
 - Month/Day Hour:Minute:Second *
 - Default Line Style
 - linear *
 - step
 - natural, quadratic
 - Show Time Axis Title *
 - Show Left Axis Title *
 - Show Right Axis Title *
 - Show Current Values *
 - Show Axis on Left or Right (Graph on Left Axis) *
 - Show Min/Max *
 - Show Line Stats *
 - Show Cross Hairs *
 - Select Line Color (dialog)
 - Choose Chart Font (dialog)
 - Sample Point Symbol
 - None <default> *
 - Cross
 - Plus
 - Square
 - Triangle
 - Circle *
 - Diamond
 - Plus
- **Table Selection Options Context Menu:**

Use right mouse button(context menu) to select a group of statistics for inclusion in a chart. Search for name of statistic type in table.

 - Search *

- Select *
 - clear *
 - by Statistic *
 - by Type *
 - All
 - Chart Select
 - **File:**
 - Browse Open Files
 - Select Single File
 - Select Recent File
 - **General Options:**
 - Statistics Level - Intended to provide a filter for statistics that are useful and commonly used and the more specialized ones that may additional understanding of the product to be useful.
 - Common
 - Advanced
 - All
 - Absolute Time Stamps
 - Copy Reference Lines
 - Single File Mode *
 - No Flatlines *
 - Combine Across Files *
 - Combine *
 - Change Directory - changes working directory for vsd, effects saving of snapshots or writing of postscripts data used for printing (temp data)
 - Update - reload data from one file
 - Update All - reload data from all files
 - Auto Update - periodically update one one
 - Templates
 - Reload Template File
 - New Template Chart
 - Templates Use Selection
 - **Chart Display Options:** See Figure 2.
Information about a selected line or chart including display of values.
 - X value (time) *
 - Y value *
 - Min *
 - Max *
 - Standard Deviation *
 - Sample Style
 - No Filter * - Displays the raw values for the statistic with no filtering at all.
 - Per Second * - Displays the difference between two consecutive samples of the statistic, divided by the number of elapsed seconds between the two samples.

- Per Sample * - Displays the difference between two consecutive samples of the statistic.
 - Aggregate * - Displays a running total of per-sample deltas for the statistic. Resets to zero when the delta is zero or changes direction.
- Scaler (*,/,+,-) with value
- Legend for each statistic *
- Delete statistic via legend * - using middle mouse button
- Select legend to highlight line in chart *
- Select line to highlight by clicking on a point on the line *
- **Chart Functions:**
 - Adjusts the visual presentation of the charted data
 - Update Chart
 - Zoom/unzoom using middle/right mouse buttons *
 - Zoom In *
 - Zoom Out *
 - Compare Two Points - similar to Log Delta
 - Compute Scale
 - Compute Scale All
 - Unscale
 - Unscale All
 - Snapshot - generate a gif of the chart for sharing *
 - Add From Template
 - Save Template
 - Copy Line * - Copy line from chart to clipboard
 - Paste Line * - Add line from clipboard to chart
 - Cut Line * - remove line from chart and store on clipboard
 - Delete Line * - remove line from chart
 - Trim/untrim Left * - Ignore older data to the left of trim line
 - Trim/untrim Right * - Ignore later data to the right of trim line
 - Normalize
 - Add Lines * - Add the values of two lines and put result on current chart
 - Diff Lines * - Subtract second selected line from first and put result on current chart
 - Divide Lines * - Divide second selected line from first and put result on current chart
 - Log Delta * - Computes mean, min, max, stddev between two points on a line. Results displayed in Log Info window. See Figure 3.
 - Log Info * - results from Log delta and Compare two points operations
 - Select/highlight line/data via Legend *

11. JVSD Functional Differences

11.1. Changes

- 11.1.1. Menus will be significantly rearranged, in particular will follow the common GUI standards for “File”, “Edit”, “View” and “Help” menus capturing specific context related functionality.
 - 11.1.2. Scaling of data/charts will only have limited support in first release, this feature is not critical for many users and multiple independent vertical axis will be supported rather than only two as in classic VSD. For a few users this is a critical feature for certain kinds of performance analysis and aids in comparing two sample runs with different configurations.
 - 11.1.3. Chart display is based on sampled data with refinements when zooming
 - 11.1.3.1. Future option to allow non-sampled or high resolution view with associated degrading of performance when drawing chart
 - 11.1.4. Multiple vertical axis will be support with one for each data type (bytes, entries...) being the default rather than only two in VSD (a single left and right)
 - 11.1.5. Many of the charting options should be part of a preferences dialog rather than as menu items of the table window as the defaults once set for a user are rarely changed
 - 11.1.6. Statistics Info window will dropped and its display info added to the help system
 - 11.1.7. Displaying of metadata such as stat file path and machine name that generated file when selecting a legend item needs to addressed to simplify analysis tasks
 - 11.1.8. Chart legend needs to be scrollable or limit the number of items that can on an individual chart (legend can not consume the whole chart window)
 - 11.1.9. Will not support classic VSD preferences stored in the .vsdconfig, .vsdrc and .vsdtempates under the user home directory
 - 11.1.10. Gif output of snapshots will be dropped in favor of PNG output
 - 11.1.11. Legacy features specific to older GemStone products such as GemStone/S will be dropped (help descriptions, monitor...)
 - 11.2. Support library for reading statistics will be provided with examples on usage
-

Appendix

1. Figure 1 Classic VSD Table Window
2. Figure 2 Classic VSD Chart Window
3. Figure 3 Classic VSD Log Info Window
4. Figure 4 Classic VSD Statistics Info Window
5. Figure 5 JVSD Table Window
6. Figure 6 JVSD Chart Window

Figure 1. Classic VSD Table Window

StartTime	File	Samples	Pid	Type	Name
09/02 18:17:48	1	1005	3412	VMStats	vmStats
09/02 18:17:48	1	1005	3412	VMMemoryUsageStats	vmNonHeapMemoryStats
09/02 18:17:48	1	1005	3412	VMMemoryUsageStats	vmHeapMemoryStats
09/02 18:17:48	1	1005	3412	VMMemoryPoolStats	PSSurvivorSpace-Heapmemory
09/02 18:17:48	1	1005	3412	VMMemoryPoolStats	PSPermGen-Non-heapmemory
09/02 18:17:48	1	1005	3412	VMMemoryPoolStats	PSOldGen-Heapmemory
09/02 18:17:48	1	1005	3412	VMMemoryPoolStats	PSEdenSpace-Heapmemory
09/02 18:17:48	1	1005	3412	VMMemoryPoolStats	CodeCache-Non-heapmemory
09/02 18:17:48	1	1005	3412	VMGCStats	PSScavenge
09/02 18:17:48	1	1005	3412	VMGCStats	PSMarkSweep
09/02 18:17:48	1	1005	7002	StatSampler	statSampler
09/02 18:17:49	1	1003	3412	ResourceManagerStats	ResourceManagerStats
09/02 18:17:49	1	1004	3412	PartitionedRegionStats	/test
09/02 18:17:48	1	1005	1980419617	LinuxSystemStats	ketchum.gemstone.com
09/02 18:17:48	1	1005	3412	LinuxProcessStats	CqServer-proc
09/02 18:17:49	1	1004	3412	IndexStats	sampleIndex-2
09/02 18:17:49	1	1004	3412	DLockStats	dlockStats
09/02 18:17:48	1	1005	3412	DistributionStats	distributionStats
09/02 18:17:50	1	1002	3412	CacheServerStats	ketchum.gemstone.com-0.0.0.0.0.0-5522
09/02 18:17:49	1	1004	3412	CachePerfStats	RegionStats-partitionMetaData
09/02 18:17:49	1	1004	3412	CachePerfStats	RegionStats-partition-test
09/02 18:17:48	1	1005	3412	CachePerfStats	cachePerfStats
09/02 18:17:50	1	1002	3412	CacheClientNotifierStatistics	cacheClientNotifierStats

numKeys
numUpdates
numUses
numValues
updatesInProgress
updateTime-IndexStats
usesInProgress
useTime

Chart:

Add Line New Chart

Figure 2. Classic VSD Chart Window

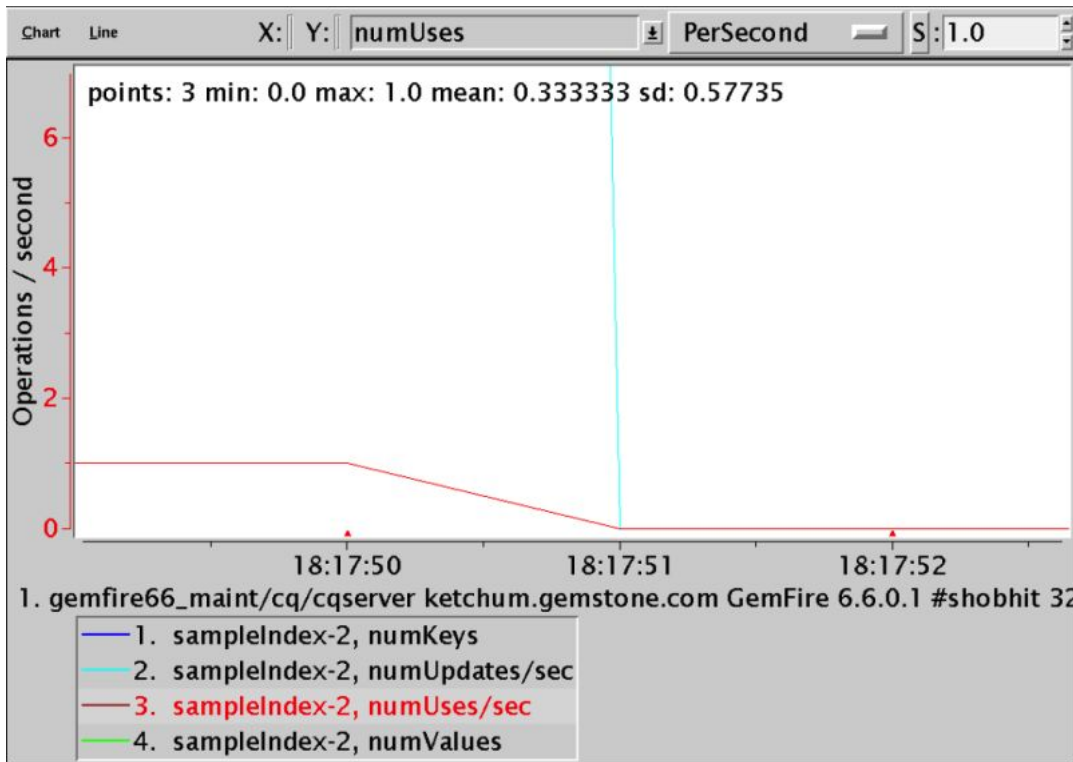


Figure 3. Log Info Window

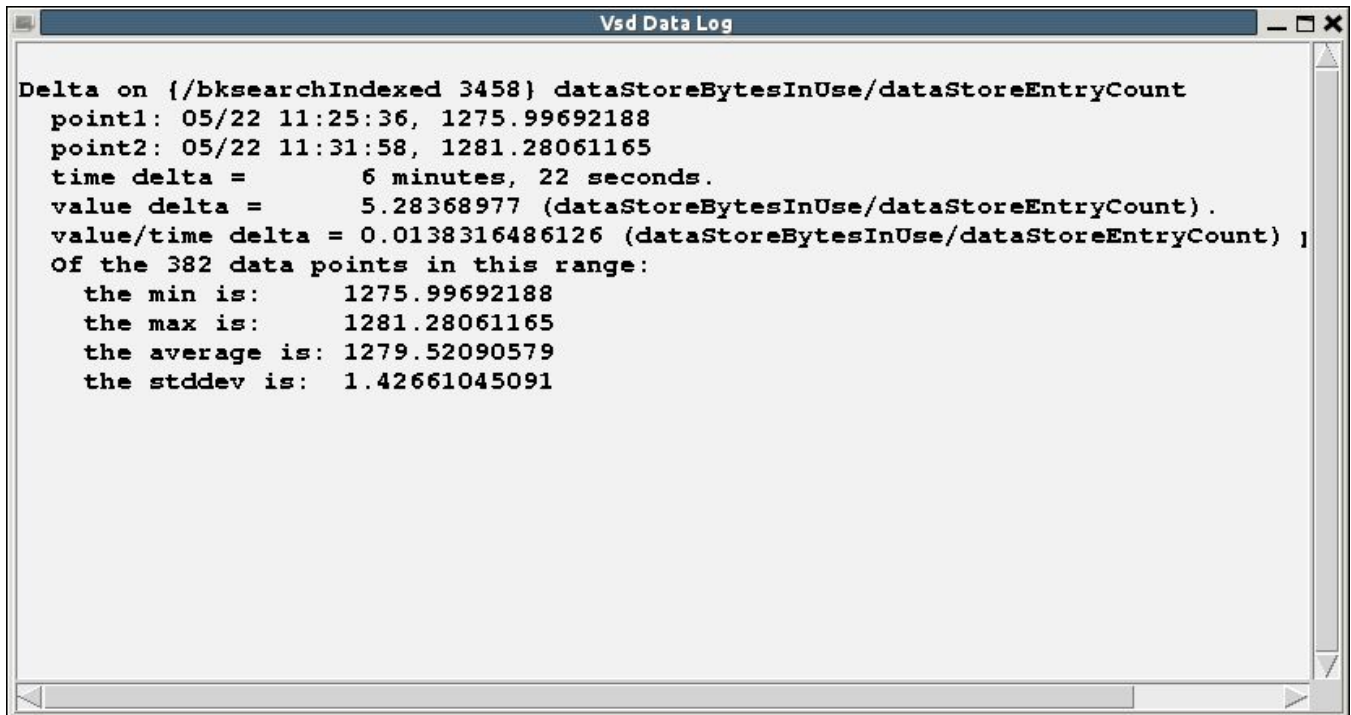


Figure 4. Statistics Information Window

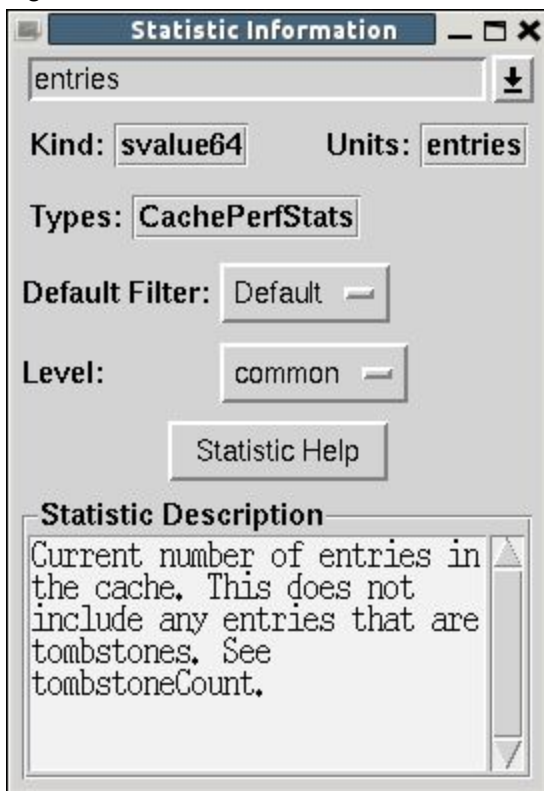


Figure 5. Prototype JVSD Table Window

The screenshot shows the jVSD application window with a menu bar (File, Edit, Help) and a table of system statistics. Below the table is a list of metrics and a chart control area.

Row	Start Time	File	Samples	PID	Type	Name
0	Sat Jan 25 16:30:40 PST 2014		47955		DistributionStats	distributionStats
1	Sat Jan 25 16:30:40 PST 2014		47955		StatSampler	statSampler
2	Sat Jan 25 16:30:40 PST 2014		47955		VMStats	vmStats
3	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryUsageStats	vmHeapMemoryStats
4	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryUsageStats	vmNonHeapMemoryStats
5	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryPoolStats	Code Cache-Non-heap memory
6	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryPoolStats	Par Eden Space-Heap memory
7	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryPoolStats	Par Survivor Space-Heap memory
8	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryPoolStats	CMS Old Gen-Heap memory
9	Sat Jan 25 16:30:40 PST 2014		47955		VMMemoryPoolStats	CMS Perm Gen-Non-heap memory
10	Sat Jan 25 16:30:40 PST 2014		47955		VMGCStats	ParNew
11	Sat Jan 25 16:30:40 PST 2014		47955		VMGCStats	ConcurrentMarkSweep
12	Sat Jan 25 16:30:40 PST 2014		47955		LinuxSystemStats	cslps15.nam.nsroot.net
13	Sat Jan 25 16:30:40 PST 2014		47955		LinuxProcessStats	AMCNDSGemfireStorage1_cslps15-proc
14	Sat Jan 25 16:30:40 PST 2014		47955		CachePerfStats	cachePerfStats
15	Sat Jan 25 16:30:41 PST 2014		47953		ResourceManagerStats	ResourceManagerStats

Below the table, the following metrics are listed:

- sentMessages
- sentMessagesTime
- sentMessagesMaxTime
- receivedMessages
- receivedBytes
- sentBytes
- processedMessages
- processedMessagesTime

At the bottom, there are three buttons: "New Chart", "Add to Chart", and a "Chart" dropdown menu.

Figure 6. Prototype Chart Window

