

Instrumentation Viewer User's Guide

Contents

Overview	4
Introduction	4
Viewer Documents	5
Viewers	6
Setting up a Viewer Document.....	11
The Viewer Document Window	11
Layout Mode	12
Choosing and Placing Viewers	12
Attaching Data to a Viewer	13
Setting Viewer Options	14
Browsing	15
Browse Mode	15
Viewing the Data.....	15
Setting the Time Resolution	16
Using Viewers	17
Trace Time Line Viewer	17
Trace Histogram Viewer.....	17
Trace Call Profile Viewer.....	18
Trace Text Viewer	18
Magnitude Graph Viewer, Growth Graph Viewer	18
Statistics Text Viewer	19
Other Viewers.....	19
Utilities.....	20
Comparing Two Sets of Data	20
Creating Viewer Summary Reports	20
Creating Trace Ranking Reports.....	21
Exporting Data	22
Saving Trace Subsets	22
Printing	23
Running "Live"	23
Viewing Data Live.....	23
Controlling Instrumentation	24
Reference	24
Viewer Reference.....	24
 Trace Time Line Viewer.....	24
Options.....	25
 Trace Histogram Viewer - Text Mode.....	26
Options.....	27
 Trace Histogram Viewer - Graphical Mode.....	27
 Trace Call Profile Viewer	28
Options.....	28
 Trace Text Viewer	28
Options.....	29
 Magnitude Graph Viewer	29

	Options.....	30
	Growth Graph Viewer	30
	Options.....	31
	Statistics Text Viewer.....	32
	Options.....	32
	Histogram Graph Viewer	32
	Options.....	33
	Histogram Text Viewer	33
	Tally Viewer.....	34
	Options.....	34
	Static Text Viewer	34

Overview

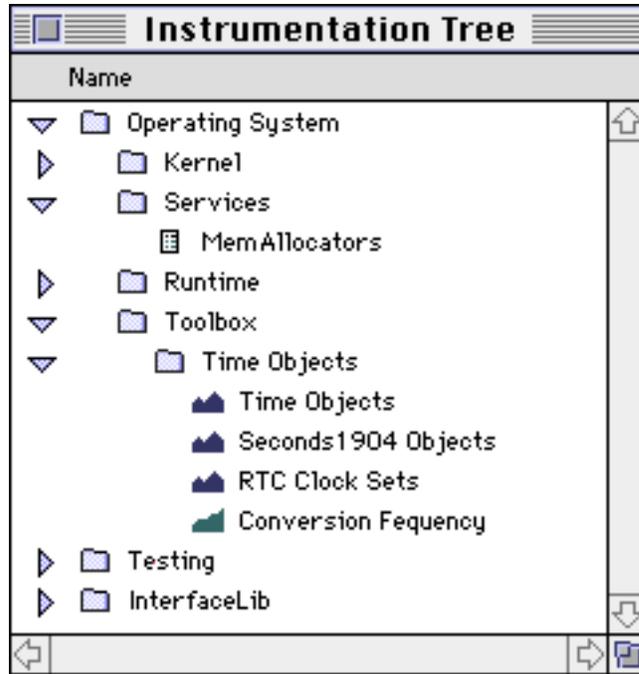
Introduction

The Instrumentation Viewer is a visualization tool that helps you understand the behavior of running software and investigate its performance. It uses plug-in viewers to display the data produced by the MacOS Instrumentation System. You can configure a set of viewers to show particular instrumentation data and save the configuration as a document. A document may be reused to display subsequent runs of the same data.

The MacOS Instrumentation System provides two services to its clients. The first allows a client to define a named event. Whenever the client determines that the event has occurred, it informs the instrumentation system. The instrumentation system creates an event record, including the event identifier and its time stamp, and writes it out to permanent storage – usually to disk. These event records are called “traces.”

The second service allows a client to create a unique, named container, into which it can drop numbers of interest. There are a variety of different “container types,” each of which is suitable for recording a different type of number, such as a count, a transient value, a value within a range, and so on. The instrumentation system will periodically sample these containers and write their contents out to permanent storage. This data is referred to collectively as “statistics.”

Inside the viewer, all of the different trace event and statistics containers appear within a tree. Because of this, they are often referred to as, “Instrumentation Nodes,” or just “nodes.” (For different reasons, they are also sometimes called “Instrumentation Classes.”)



The data files produced by the Instrumentation System appear as time-ordered sets of traces and statistics. For trace nodes there is one record for every event, or in some cases, one record for the beginning of the event and one for the end. For a statistics node, there is typically one data record present for each sampling interval, which is usually one second. (The sampling interval controls how often the current values of the statistics will be written out.) These time stamped trace and statistics records are referred to generically as “data points”.

The Instrumentation Viewer allows you to create customized views which can be used to display these data points over time.

For more details on the Instrumentation Library, the services it provides, and how to use them, see the “Instrumentation Programmer’s Guide.”

Viewer Documents

Documents in the Instrumentation Viewer have two different modes: Layout mode allows you to define how the data will be displayed, and Browse mode allows you to view the data.

The viewer document model is similar to OpenDoc’s. A new document is created initially empty. Viewers can be dragged from the floating Viewer Palette and dropped into the document in Layout mode, where they can be moved, resized and configured. Once a

viewer is in place you can drag instrumentation nodes to it from the instrumentation tree, to specify the data that the viewer will display. In Browse mode, you can set the display time; the data corresponding to the this time appears within the viewers.

The source of the information displayed can be "live" data or data recorded during a previous run. "Live" data can be shown by the viewer when it's running in the presence of software that is generating instrumentation data, or when connected to a such a system that is spooling its data to a serial port. Data from a previous run is packaged in an instrumentation data file, which can be viewed at any time.

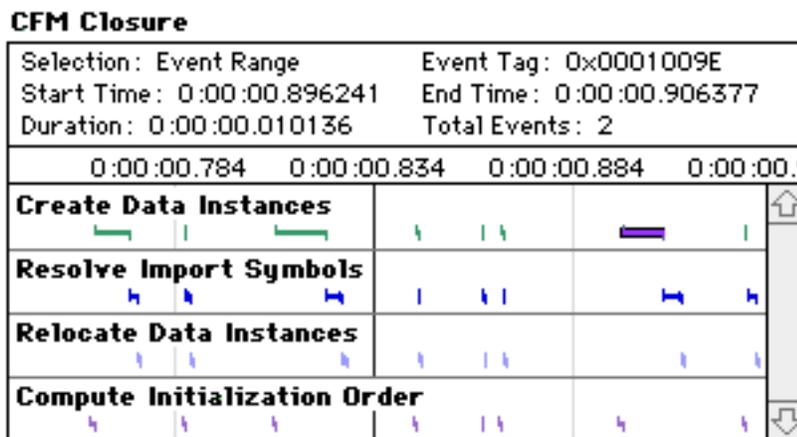
Once the viewer documents have been set up, a typical session with the viewer consists of opening existing documents in Browse mode, selecting a set of instrumentation data, and using the viewers to look at it.

Viewers

The Instrumentation Viewer supplies one or more viewers for each data type that can be recorded by the Instrumentation System.

Trace data is used to record events such as function entry and exit. Clients of the Instrumentation Library produce trace data by calling `InstLogTraceEvent()`, optionally specifying that the call is being made at the beginning of the event, at the end, or somewhere in the middle. The Instrumentation Library records the nanosecond time stamp..

The **Trace Time Line Viewer** shows begin/end trace pairs as ranges on a time line, and individual traces as points. They are color-coded by instrumentation node. This allows you to identify clustering of various events, and to pick out relationships between one event and another.



The Trace Time Line Viewer

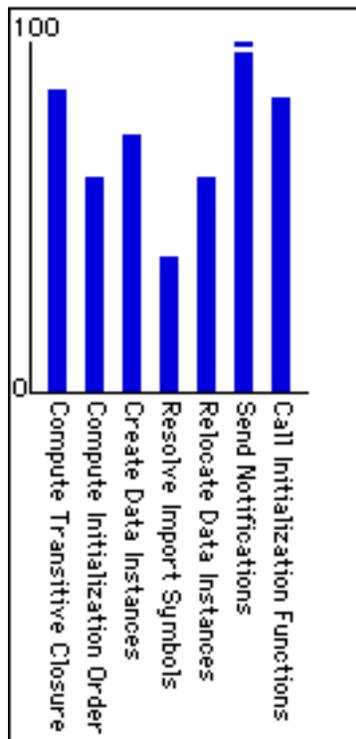
The **Trace Histogram Viewer** displays the traces logged to a set of trace nodes as a histogram. Using this viewer, you can classify each trace node by the number of times it was called. If the trace data consists of start/end pairs, you can also configure the viewer to show time spent inside each pair, or the time spent inside each pair exclusive of the time spent inside another pair in the same task. This last option effectively shows you the time spent inside a routine minus the time spent inside other instrumented routines that it calls, directly or indirectly.

CFM Closure

<u>@Trace</u>	<u>Count</u>	<u>Total (µs)</u>	<u>Minimum</u>	<u>Average</u>	<u>Maximum</u>	
Create Data Instances	13	450697	608	34669	378596	↑
Resolve Import Symbols	13	2407213	148	185170	1569681	
Relocate Data Instances	11	1116582	111	101507	986702	
Send Notifications	71	205554	29	2895	81452	↓

The Trace Histogram Viewer (Text Mode)

CFM Closure



The Trace Histogram Viewer (Graphical mode)

The **Trace Call Profile Viewer** displays the distribution of start/end event durations for a given trace node. You specify a bucket size in microseconds and it shows you the number of calls within each bucket.

FrontWindow() call time

<u>Range (µs)</u>	<u>Count</u>	
30..34	1	↑
40..44	6	
45..49	23	
50..54	26	
55..59	2	
60..64	1	
105..109	1	↓

The Trace Call Profile Viewer

Finally, the **Trace Text Viewer** shows the trace details in text form. It also shows any additional data supplied by a client to the InstLogTraceEventWithData() call.

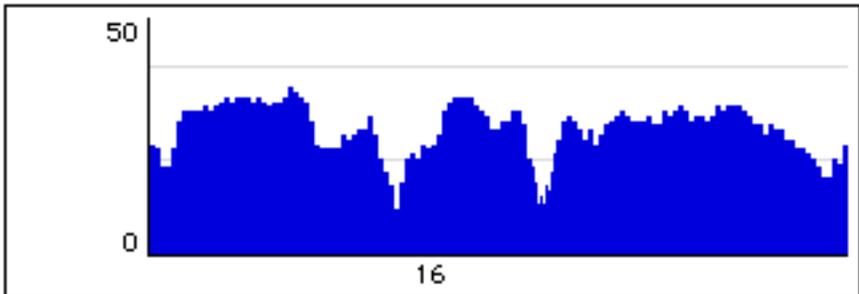
Check Event - detail view

<u>Time Stamp</u>	<u>Event Tag</u>	<u>Event Order</u>	<u>Trace String</u>	
0:08:42.049455	0x0001088F	End		↑
0:08:42.063789	0x0001088F	Middle		
0:08:42.099035	0x0001088F	End		
0:08:42.101080	0x0001088F	Middle		↓

The Trace Text Viewer

Magnitude data can be any 32-bit number that a client can measure which increases or decreases over time. Each data point consists of a value, the minimum value up to this point, the maximum value, the total of all values, and the number of times the value was updated. The last two taken together can be used to calculate the average value. All but the total are signed 32-bit numbers; the total is a signed 64-bit number. Magnitudes are typically used to track varying values such as request sizes and operation durations. It can be displayed graphically using the **Magnitude Graph Viewer**, and as text with the **Statistics Text Viewer**.

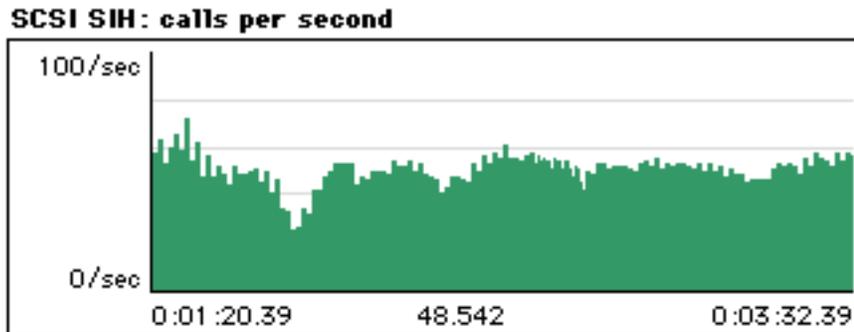
Free Pages in RAM



The Magnitude Graph Viewer

Growth data is a running total of some client-defined count, such as number of page faults or cache misses. Clients update growth nodes by calling the instrumentation library with unsigned 32-bit increments. Each data point consists of a value, which is the total of all increments passed prior to taking the sample, the minimum increment that was passed, the maximum increment, and the number of times the total was updated. All but the total value are 32-bit numbers; the total is 64-bit.

The growth rate can be displayed using the **Growth Graph Viewer**. Like magnitude, growth information may be displayed using the **Statistics Text Viewer**.



The Growth Graph Viewer

VM Monitor

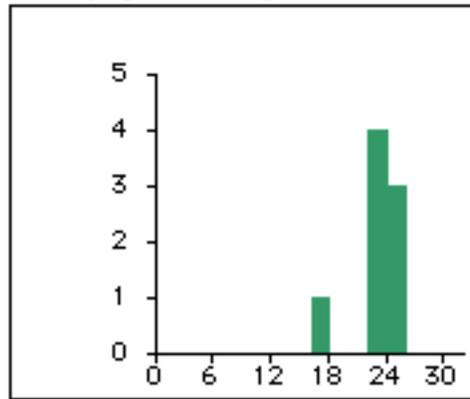
<u>Name</u>	<u>Value</u>	<u>Minimum</u>	<u>Average</u>	<u>Maximum</u>
BackingRead	561	359	10853	480821
Faults	1315	1	-	1

The Statistics Text Viewer

Histogram data produced by the Instrumentation System consists of a number of evenly-spaced “buckets”; each bucket has a count. When the client creates the histogram node, it specifies the width of the buckets and what the total range should be. In order to generate the data, it periodically calls the Instrumentation Library with a signed 32-bit value and a count. If the value falls within the range of a particular bucket, the supplied count is added to the bucket’s count. If it does not, then the count is added to a special bucket called the Overflow bucket instead. Histograms are useful for such tasks as creating profiles of input parameters.

The **Histogram Graph Viewer** displays a single histogram in graphical form, as a series of vertical bars. The **Histogram Text Viewer** displays a single histogram as several lines of text, where each line contains a bucket range and its count. Either one (or both) may be suitable, depending upon the nature of the data within the histogram and how many buckets are defined by it.

Free pages survey



The Histogram Graph Viewer

Tally data is similar to histogram data, in that it is defined by a number of buckets and a count of the “hits” in each bucket. However, where the buckets in the histogram are defined statically when the histogram node is created, the buckets in a tally node are defined dynamically. When the tally is created, the client specifies how many buckets the tally is to keep track of. In order to generate the data, the client periodically updates the tally node with a bucket identifier and a count. If the identifier corresponds to a known bucket, the count is added to the bucket’s count. If not, and there are still buckets undefined, the supplied identifier is given a new bucket whose count is initialized with the supplied count. If there are no buckets left undefined, the count is added to the Overflow bucket.

The **Tally Viewer**, in text mode, shows a list of the buckets; each line shows one bucket identifier and its count.

Cache Line Hits

Bucket Name	Count	
00000002	29,887	↑
00000001	17	
00000008	2,549	
00000006	2,139	
00000007	31	
00000009	1	↓

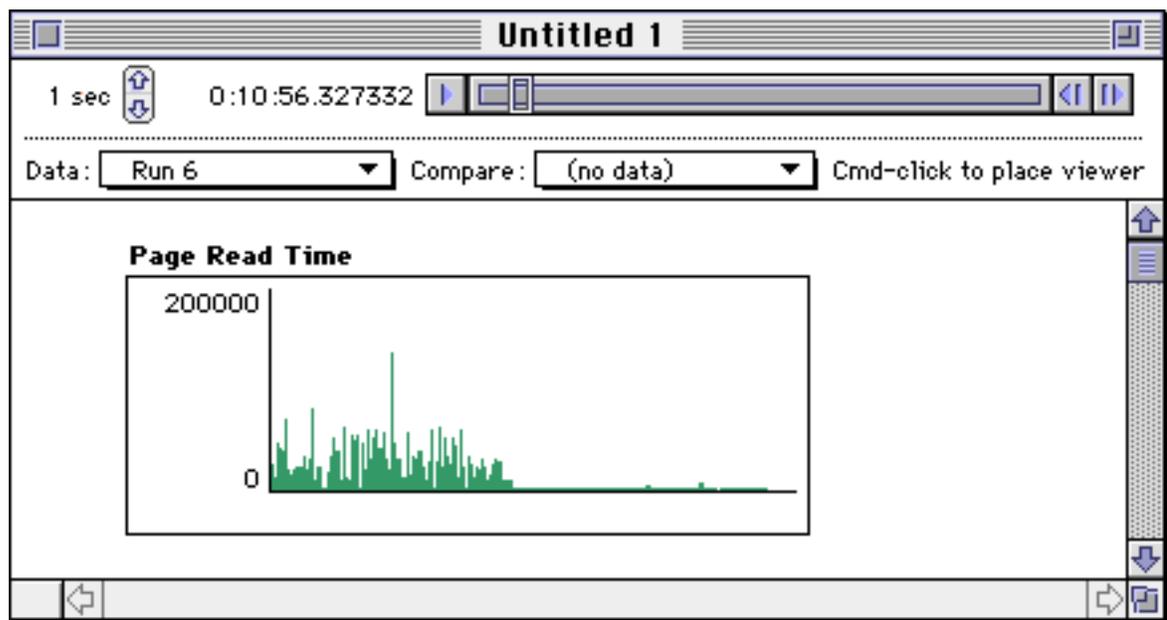
The Tally Viewer

Each of these viewers is covered in more detail in the section, “Viewer Reference.”

Setting up a Viewer Document

The Viewer Document Window

When you choose “New” from the file menu, or open an existing viewer document, the Viewer Document window appears. The viewer document can be in one of two modes: Layout mode, and Browse mode. Layout mode is used to place the appropriate viewers into the window and to configure them with the instrumentation nodes that they are to display. Once they have been set up, you can enter Browse mode to display the contents of an instrumentation data file.



The top section of the window contains controls for setting the current view time and view resolution; it is not active in layout mode. The space just below the time controls contains popups for specifying the data source. The space below that is used for displaying viewers.

Layout Mode

When a new document is created, it is initially in Layout mode. In addition to the Instrumentation Tree floating window, two other floating windows may appear: the Viewer Palette, and the Viewer Options window. If they are not visible, you can show them from the Windows menu.

Layout mode is used to choose the viewers you want to include in your document, set their size and position correctly, attach data to them, and configure their options.

Choosing and Placing Viewers

The first step in defining a new viewer document is to decide which viewers to use to display your data. Some data types may be displayed by several different viewers. In the list below, each data type is shown with the icon that the Instrumentation Tree window uses to represent it.

	<u>Data Type</u>	<u>Available Viewers</u>
	Trace	Trace Time Line Viewer Trace Histogram Viewer Trace Call Profile Viewer Trace Text Viewer
	Magnitude	Magnitude Graph Viewer Statistics Text Viewer
	Growth	Growth Graph Viewer Statistics Text Viewer
	Histogram	Histogram Graph Viewer Histogram Text Viewer
	Split Histogram	Histogram Text Viewer
	Tally	Tally Viewer

Each of these viewers is present in the Viewer Palette floating window. You can place a new instance of any one of them into a viewer document by dragging it from the Viewer Palette and dropping it into the viewer area of the document window.

You can also place a viewer by holding down the Command key and clicking the mouse in the document window where you want the new viewer to go. A popup menu will appear with the list of available viewers. Choosing one will place an instance of that viewer at the mouse position.

Once it is in place, it is represented by a rectangular frame and a title, which is “Untitled” by default. Clicking on it will select it; the frame shows “marching ants”, and resize handles appear at the corners. A viewer can be resized from any corner. Clicking on a viewer without releasing the mouse button allows you to “pick up” the viewer and drag it to another position within the window.

Holding down the shift key while clicking will select multiple viewers. They can be copied by selecting “Duplicate” from the edit menu, or removed by selecting “Clear.”

Selecting a single viewer also allows you to change its options; see the “Setting Viewer Options” section, below.

Attaching Data to a Viewer

In order to for a viewer to display data, it must be given a reference to an instrumentation node of the appropriate type. It uses the full pathname to look up the class within an instrumentation data file, which it then uses to find the data.

You supply a viewer with a reference to a node by dragging it from the Instrumentation Tree and dropping it onto the viewer in the viewer document window. If the viewer can accept the node, the full pathname will appear as text within the viewer. Some viewers can accept an unlimited number of nodes; others can accept only one or two. (See the “Viewer Reference” section for more details.) If you try to drag a node into a viewer which already contains as many nodes as it can display, the viewer will reject the drag.

Note that in order to see nodes within the Instrumentation Tree, you must have a data file open. You can open a file by choosing “Open...” from the File menu. Once open, all of the nodes within the file are displayed inside the Instrumentation Tree floating window. If more than one data file is open, a separate Instrumentation Data window will be created for each file.

(You can attach a node to a viewer even if the node is not present inside the instrumentation data file that the viewer document currently has open. In this case, the viewer will display the “Specified Class Not Found” message while in Browse mode.)

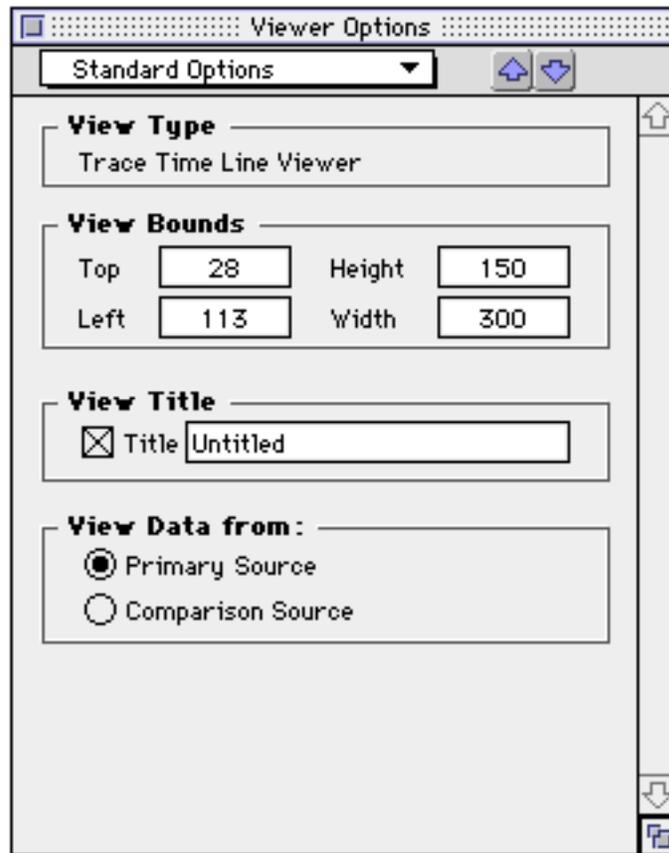
In order to remove a node from a viewer, click on the viewer so that it activates, then click on the pathname to select the node within the viewer. Shift-click to select more than one node. Choose the “Clear” command from the menu or hit the “Delete” key to remove the node from the viewer.

Note that the data node must be compatible with the viewer that you attach it to. (See section “Choosing and Placing Viewers”, above, for a list of which data types the

standard viewers can display.) If you attempt to drag an incompatible node into a viewer, the viewer will reject the drag.

Setting Viewer Options

Viewers can be configured to present their data in various ways. This is done through the Viewer Options floating window, which can be displayed in Layout mode by choosing “Show Viewer Options” from the Window menu, or by double-clicking a viewer. The Viewer Options window has a popup menu which allows you to navigate through the various option sets.



In order to set the options of a particular viewer, click on it in the viewer document window to select it. It will display “marching ants” to confirm that it is selected. Its current settings will appear inside the Viewer Options window.

Each viewer supports a common set of options, called “Standard Options” in the Viewer Options popup menu. They control the view bounds, view title, and data source. Many

viewers support additional options. These are documented for each particular viewer in the “Viewer Reference” section.

Browsing

Browse Mode

Once the viewers have been set up in Layout mode, you can enter Browse mode to display some data. Select “Switch to Browse Mode” from the View menu and open an instrumentation data file using "Open..." from the File menu, if you have not already done so.

A document will “remember” an instrumentation data file that is opened within it, and will attempt to find and open that file again the next time the document itself is opened. If the data file is deleted and replaced with another of the same name, the document will open the new file instead. You can change the active data file by opening another one via the File menu and selecting the name of the newly opened file with the "Data:" popup menu.

(If you are suitably equipped, you can also view data “live”. See the “Viewing Data Live” section in the chapter, “Running Live.”)

In Browse mode, the time and time resolution controls become active. Both the Viewer Palette and the Viewer Options floating windows are hidden. The time slider works just like the standard QuickTime control that it resembles. Instrumentation data consists of a series of points logged over a range of time; the time slider allows you to set where within the time range you’d like to look.

Viewing the Data

Whenever the time slider is advanced by one position, the viewers display another “snapshot” of the data at the new time, rather like the frames of a movie. In response, the various viewers update themselves in different ways. The Trace Time Line viewers, Magnitude Graph viewers, and Growth Graph viewers will each move their data over one slot to the left, displaying the data of the new time in the right-most slot. The Histogram viewers and Statistics Text viewers may update their values, assuming that

new data is present in the new interval. The Trace Text viewers will display only the traces which are present in the new interval.

The beginning of the interval displayed – the “current time” – is shown in the top-left corner of the viewer document, beside the time slider control. It corresponds to the amount of time that had passed since the Instrumentation Library was initialized. The “current time interval” is the period of time as long as the current time resolution (see below) that begins at the current time. Each viewer displays the data which is most closely associated with this time interval; for more information on each viewer, see the “Viewer Reference” section.

In order to advance the data by one frame, click on the right-arrow to the extreme right of the time slider control. In order to go back one frame, click on the left-arrow beside it. Clicking on the “play” button on the left side of the time slider causes the data to advance continuously. The Instrumentation Viewer attempts to show the data in real time, but it does not skip frames, so smaller time steps in complicated viewer documents may cause the data to advance more slowly. Clicking anywhere in the slider body will cause the current time to skip to that point.

Some viewers permit interaction: clicking on a data point in the Trace Time Line viewer will display a summary of it at the top of the viewer; moving the mouse over a bar in the Trace Histogram viewer in graphical mode will show the bar’s value as text.

Selecting “Copy” from the Edit menu will place a picture of the document on the clipboard.

It is possible to use a viewer document to look at an instrumentation data file which contains no data for the node that the viewer was originally set up with. This can occur for one of two reasons: the system that produced the file never defined the instrumentation class node, or the node was defined, but no data was produced for it in the interval displayed. In the first case, the viewer will display the message, “No Class Specified.” In the second, it will display, “No Data for Specified Class.”

Setting the Time Resolution

The time resolution defines the duration of each snapshot. If the time resolution is 500ms, then each snapshot represents 500ms in time; each “frame” is 500ms of the movie. Large time resolutions can be useful for seeing a lot of data in each snapshot; small time resolutions allow you to “zoom in” for a close-up look at a period of particular interest.

The time resolution control is in the top-left corner of the viewer document window, to the left of the time slider control. To increase the duration of each snapshot, click the up-

arrow; to decrease the duration click the down-arrow. The minimum duration is one microsecond; the maximum is 30 minutes

Using Viewers

Note that each of these viewers is covered in more detail in the section, “Viewer Reference.”

Trace Time Line Viewer

The Trace Time Line Viewer is used to display the traces of up to eight trace nodes as a time line. Time is represented in the horizontal direction; each trace is plotted as a vertical line with its horizontal position representing the time it was logged. You can use it as a quick guide to identifying periods of high activity in particular routines, and to correlate things happening in one area of your software to events occurring elsewhere.

You can also click on a trace, or on a Start/End trace pair, to see more information displayed about that particular event.

The viewer’s area is divided into horizontal and vertical bands. The horizontal bands represent each trace node; the vertical bands represent intervals of time equal to the current time resolution. The dark vertical line represents the current time; as time advances, the traces move to the left.

Trace Histogram Viewer

The Trace Histogram Viewer calculates how many trace events were logged up to the current time interval for each of a set of trace nodes. It can present the result in a table or as a graphical histogram. It allows you to get an idea of which of the trace nodes in the set are seeing the most activity.

If you log Start trace points at the beginning of a procedure and End trace points at the end of it, you can also use the Trace Histogram viewer to show the total time spent inside each procedure. This allows you to examine a set of procedures to determine which is using the most (or the least) time.

You can move the mouse over each histogram bar in graphical mode to show a numerical summary of the data for each node. This is the same data that appears in the table when in text mode.

Trace Call Profile Viewer

The Trace Call Profile Viewer is designed to show the distribution of call times for a particular trace node. It assumes that the trace data consists of start/end event pairs. It divides the possible call time into a number of discrete ranges using a user-defined bucket size. It then scans through the trace data looking for start/end pairs. Every time it finds one, it subtracts the start time stamp from the end time stamp to get the call time. It then determines the range that the call time falls into and increments the count of the corresponding bucket. It ignores middle events and all other traces.

It displays the results in a table. Each row shows a time range in microseconds and the number of start/end pairs whose duration fell into that time range. Empty ranges are not displayed.

Trace Text Viewer

In order to see the “custom data” which is logged along with the trace by a `InstLogTraceEventWithData()` call, you can use the Trace Text Viewer. It shows every trace which is present in the current time interval for a single trace node. The last column of is the “Trace String,” which the result of the Instrumentation Library taking the custom data and formatting it into a display string. (You can also retrieve this information by performing a Text Export on the trace node. See the section, “Exporting Data,” in the “Utilities” chapter.)

Magnitude Graph Viewer, Growth Graph Viewer

You can use the Magnitude Graph Viewer to produce a two-dimensional graph of a single-valued measurement that varies in time, such as the amount of free space in a heap. You can use the Growth Graph Viewer to produce a two-dimensional graph of a rate of growth, such as total page faults per second, which is calculated from the number of new page faults that appear in each sampling interval.

Either viewer can have up to two nodes attached to it; when there are two, one gets overlaid on top of the other. Since the second node gets drawn last, it often works best if

the value of the second node is less than that of the first. You can use the Viewer Options floating window to set the color of each node's line.

You can use the auto-ranging capability to get an idea of what vertical range will work best for the data. After that, you can switch to a static range to prevent the display from rescaling whenever it detects a new maximum.

As the data appears, it starts at the right-hand side of the graph and moves left. On the x-axis, one horizontal pixel represents a time interval equal to the current time resolution; the value for the current time is shown by the right-most line.

Statistics Text Viewer

You can use the Statistics Text Viewer to get a compact summary of one or more magnitude and growth statistics. It is most useful when you are only interested in seeing the current value of a statistic, and do not need to see the entire history.

Normally, four columns of information are presented – the value, minimum, maximum, and average. You can configure the viewer to hide some of these, so that only the ones of interest are displayed.

Other Viewers

You can use the Histogram Graph Viewer or the Histogram Text Viewer to examine histogram nodes. These statistics typically aim at constructing a usage profile of some operation whose main parameter varies over a linear range, such as a memory allocator whose main parameter is a requested block size. For histograms with a small number of buckets, the Histogram Text Viewer provides concise totals for each one. It also supports Split Histograms, where the bucket widths are larger in one part of the range than in another. For histograms with a large number of uniform buckets, or where complicated statistical patterns might emerge, the Histogram Graph Viewer shows the data as a standard histogram graph.

The Tally Viewer provides a simple view of tally statistics, showing the bucket identifier followed and its "hit count." Tallies tend to be used for specialized purposes, such as providing totals for the number of files processed on a per-file type basis.

Finally, the Static Text Viewer is mainly an adornment tool that allows you to add auxiliary labels to the viewers that actually display the data.

Utilities

Comparing Two Sets of Data

If you are interested in comparing two sets of data - perhaps two instrumentation runs recorded under slightly varying circumstances - you can use the Instrumentation Viewer to display them both in the same document, synchronized to a single time slider control.

In order to make comparison data available in a viewer document, you must first open a data file using the File menu's "Open..." command or a live data connection, using the File menu's "Open Live Data" command. Once you select this file or live connection with the "Data:" popup menu, it becomes the primary data set. Next, select another data set using the File menu's "Open..." or "Open Live Data" command. Identify this data set as the "comparison" data set by selecting its name with the "Compare:" popup menu.

Every viewer has a "View Data From:" setting available in the Viewer Options floating window. It specifies whether a viewer is to draw its data from the primary data set or the comparison data set. By default, a viewer is set up to display the data from the primary data set. To have a particular viewer show data from the comparison data set, select it in Layout mode and choose "Comparison Source" from the "View Data From:" section.

Upon return to Browse mode, viewers configured in this way will display data from the comparison data source. The other viewers will display data from the primary data source, as before. Specifying a comparison data source causes the range of the time slider control to be reset to include the time ranges of both the primary and comparison data. Moving the time slider control will advance the data from both data sets simultaneously.

Often, a second instrumentation run will begin at a slightly different time than the first. To compensate for this, you can use the View menu's "Align Comparison Data..." command to specify the times in the primary and comparison data sources which are to be aligned. This allows you to line up the interesting data in the second set with that of the first.

Creating Viewer Summary Reports

While the Instrumentation Viewer allows you to examine the data from a single run of your software, it is often interesting to keep a record of various results over a longer

term. You can create text summaries of the data currently displayed, called viewer summary reports, to assist this process.

Each viewer produces a summary of the data that it displays. Thus, the form that the report takes depends upon how the viewers are set up within the document.

In order to create a viewer summary report, open a viewer document, load in an instrumentation data file using the “Open...” menu item of the File menu, and position the time slider so that it displays the data that you are interested in. Then, choose “Create Viewer Summary” from the Report menu. A standard file dialog will appear: name the report and choose “Save.” A text file will be created and each viewer will add its own summary to it.

Each report contains the name of the document used to generate the report and the date. Also included are the number of traces that were lost due to trace buffer overrun, and an estimate of how long traces were being lost for.

Not all viewers will generate report data. Currently, the Trace Histogram viewer writes out the node name and number of calls for each trace node. If the trace data is set up as start/end pairs and the trace histogram viewer is configured to display times, then it will also write out the total time spent, minimum call time, maximum call time, and average call time. All times are expressed in nanoseconds.

The Statistics Text Viewer also writes out a summary. For each node, it writes out the current value, the number of times the node was updated, the minimum value, the maximum value, and if the node is a magnitude statistic, the average value. All values are expressed in the same units that the client used to update the node.

The Trace Call Profile, Histogram Text and Tally Viewers will all write out the tables that they display.

Creating Trace Ranking Reports

When analyzing data for the first time, it is often useful to get a summary of all the traces in the file. This is particularly true when there are a large number of different trace nodes.

Using the Trace Ranking feature of the Instrumentation Viewer, you can produce a tab-delimited text file that shows how many traces were logged to each trace node, and how much time was spent between Start and End events – both cumulative, and exclusive of instrumented sub-calls.

In order to create a trace ranking report, open a viewer document and load in an instrumentation data file using the “Open...” menu item of the File menu. Then, choose

“Generate Trace Ranking” from the Report menu. A standard file dialog will appear: name the report and choose “Save.”

Exporting Data

Using the Instrumentation Viewer, you can export some or all of the data points in an instrumentation data file to a tab-delimited text file. This allows you to run scripts on the data to further filter or process it, and to import it into a variety of spreadsheet, data-analysis, and presentation packages for further investigation.

In order to transfer data from an instrumentation data file to a text file, you must have the instrumentation data file open within a viewer document. It is not necessary that the viewer document contain viewers associated with the data; an empty document will do. Select the instrumentation node(s) which contain the data in the Instrumentation Tree floating window and choose “Export Data” from the File menu. A Standard File dialog will appear; give the text file a name and choose “Save.” The Instrumentation Viewer will convert all the data points associated with the selected nodes to text and write them to the file.

In the text file, each node will be prefixed by the full colon-delimited path of its node name, and a line of column headers. Below the column headers will be the data itself; one data point to a line. Where time stamps are included, they are expressed as nanoseconds.

Normally, data points are exported in time stamp order. There is one exception to this rule: when trace data is exported, trace events marked as Start, Middle or End events and linked together by a common trace tag will be grouped together. This simplifies any analysis which wishes to treat these traces as a single event.

Saving Trace Subsets

Although it may be valuable to keep an individual record of every trace logged, the disadvantage of doing so is that the trace files can become very large. Browsing a 25MB file can take a long time, simply because of the volume of data that must be read.

The Instrumentation Viewer allows you to specify a subset of the trace events in a file and save them into a second file. To do so, select the trace classes that contain the interesting traces in the Instrumentation Tree window. Then activate the viewer document window that is displaying the data file and select the time range containing the interesting data. This is done by holding down the Shift key when you move the time slider control.

Finally, select “Save Traces” from the File menu. A Standard File dialog will be presented; name the file and choose “Save.” The Instrumentation Viewer will produce a second instrumentation data file that contains only the trace events in the specified time belonging to the specified trace classes.

If you do not choose a time range, every trace event belonging to the selected trace classes is saved into the file. If no trace classes are selected, every trace event in the specified time range will be saved.

Printing

The Instrumentation Viewer can print the contents of its viewer document windows. To do so, open the viewer document that you are interested in printing, and load its instrumentation data file. Position the time slider so that the viewer document displays the data that you wish to print. Select “Page Setup” from the File menu to ensure that the correct printer has been chosen and that its options have been set up correctly. Then select “Print” from the File menu.

Selecting “Print One” from the file menu automates the Page Setup and Print steps. It will print a single copy of the document with the current Page Setup options. If no Page Setup options have been specified, it will use the printer’s default settings.

Running “Live”

Viewing Data Live

You can watch the performance data as it is being generated by running the Instrumentation Viewer on a system with active Instrumentation Spoolers. To do so, launch the viewer and select the "Open Live Data" menu item from the File menu. Active instrumentation nodes will appear within the "Live Data" Instrumentation Tree floating window.

In Browse mode, clicking the “play” arrow of the time control will continually advance the time as new instrumentation data comes in. Whenever the slider is up against the right side of the time control, the most recent time interval is being displayed. The slider can be

dragged backward and forward without disrupting the data flow. It is also possible to enter Layout mode to reconfigure the viewer document and return to Browse mode without losing data.

Controlling Instrumentation

You can use the Instrumentation Viewer to turn data collection on and off for individual instrumentation points, or for every active point in the system.

If you are using the Instrumentation Viewer to view live data, the Instrumentation Tree window will display two additional columns labeled N and C. When you click on the dot in the N column of an active instrumentation node, the node will be enabled and the dot will turn green. Clicking the dot again will disable the node. The spoolers only collect data from enabled nodes.

The “C” column is not yet operational.

A viewer document window that is displaying live data contains a “Collect data” check box. Unchecking it will disable every instrumentation node. Checking it will re-enable all nodes which were previously enabled.

Reference

Viewer Reference

The following sections provide a detailed reference to the operation and options of each of the viewers included in the standard release of the Instrumentation Viewer. The icons to the left indicate the instrumentation data type(s) that the viewer can display.

Trace Time Line Viewer

The Trace Time Line Viewer displays trace data for up to eight trace instrumentation nodes on a time line. Its display is divided into a number of vertical segments. Each segment represents a period of time equal to the current time resolution, shown in the top

left-hand corner of the viewer document window. The dark vertical line indicates the position of the current time, which is also shown next to the time slider control.

Traces are represented within these segments by colored vertical lines. By default, each trace node is given a different color. If there are multiple trace points for the same node in the same interval, they are staggered vertically so that they do not overlap. If the trace data includes Start/End trace pairs, then a horizontal line connects each start/end pair. The viewer matches up the start and end events by looking for a matching event tag within the trace records.

If a trace is clicked on, it is highlighted and a summary of the point is presented at the top of the viewer. It shows the trace type (Start, Middle, or End), the event tag, and the time the event was logged.

If the line joining a start/end pair is clicked on, the summary includes the times of the start and end events, the time between the start and end events, and how many traces were recorded for that event. (This includes middle events with the same event tag.)

Hitting the Tab key will cause the selection to move to the next event of the selected node. Shift-tab will move to the previous event of the selected node.

Option-clicking inside a row will search forward for the next trace belonging to that node.

You can zoom in to a particular set of points by clicking and dragging out a rectangle. The viewer will zoom in and adjust the current time to best display the points within the rectangle, to the closest degree allowed by the fixed set of view time resolutions.

Options

Each node can be configured separately with a number of options. Choosing “Trace ClassN Info” from the popup menu in the Viewer Options floating window will show the options for the Nth class displayed. For each class, the color shown in the popup menu is the one that is used to draw the trace events. Selecting “Default Name” will cause the traces to be labeled with the name of the instrumentation node; selecting “Custom Name” allows a different name to be used. Checking “Match Paired Events” will cause the line to be drawn between matching Start and End events; checking “Show Middle Events” will cause the events between them to be drawn.

Additionally, there are two “Time Line Options” which apply to the viewer as a whole: Allow User Interaction, and Display Class Labels. The first creates the area at the top of the viewer for the point summary and fills it in when a point is clicked. The second causes the names of the classes to be drawn above their traces. Both are on by default, but may be turned off to reduce the screen space that the viewer requires.

Trace Histogram Viewer - Text Mode

The Trace Histogram Viewer counts traces from the beginning of the file to the end of the current time interval. It can display the results either as text in a table, or as a histogram. This section describes the text mode and common viewer operations; the next section, “Trace Histogram Viewer: Graph Mode,” describes the graphical histogram mode. For a description of the Instrumentation System’s trace data, see the “Viewers” section in the “Overview” chapter.

By default, the Trace Histogram Viewer is configured to display the number of times an event was logged against each trace node. If the data was recorded as a series of Start/End traces pairs, then the viewer can also display the total amount of time that was spent inside the Start/End pairs, which normally corresponds to the amount of time spent inside a particular routine. The viewer can also show the total time spent inside the Start/End pairs minus the time spent inside other Start/End pairs in the same time interval. This normally corresponds to the time spent inside a particular routine, minus the time spent inside all other instrumented routines that it called.

(Be careful, however, as it is not required that all routines log their start/end pairs within the scope of a single routine. For example, consider the case of an asynchronous PBWrite call with a completion routine. The start trace could be logged just before calling PBWrite, and the end trace could be logged inside the completion routine. If another Start trace is logged after the PBWrite call returns but before the completion routine is run, the instrumentation system will interpret that as a sub-call of the PBWrite trace.)

When counting trace events, the Trace Histogram Viewer counts all traces except Middle and End trace points. This is to ensure that multi-part events are correctly accounted for.

When calculating Time Spent, the viewer counts all Start/End pairs where both the Start and End events appear either before or within the current time interval. Thus, although a Start event for a node may appear within the current time interval, it is not included in the trace histogram totals until the current time is advanced to include the corresponding End event as well. Events that are not part of a Start/End interval do not add time to the total.

You can place as many trace nodes as necessary into this viewer; each node is identified by its class name.

When in text mode, the results are displayed in a table. Each row represents a trace node; the first column is the class name of the node. The second column shows the trace count – the number of trace events present up to the current time interval. If the viewer is configured to display Time Spent, the third column displays the total time spent between start/end pairs up to the current time interval, in microseconds. The fourth, fifth and sixth columns display the minimum start/end time, the average, and the maximum time, respectively. (For a more detailed display of the call time distribution, see the “Trace Call Profile Viewer” section.)

There is a “zero” button in the upper left-hand corner of the viewer. Pressing this button will set the “zero point” of the histogram to the current time. All totals are set to zero, and counting begins from the zero point. You can set the zero point repeatedly, to any point in the time range. Closing and reopening the document will also reset the zero point.

Clicking on a column header will sort the data by that field. The data is resorted as it is updated.

Options

The Trace Histogram Viewer supports a single set of options, accessed via the “Trace Histogram Options” item of the popup menu in the Viewer Options floating window. It controls the counting mode of the histogram, and how the results are displayed. By choosing the appropriate radio button, the viewer can be configured to count the number of calls*, the time spent between Start/End trace pairs, or the time spent exclusive of calls to instrumented subroutines.

Choosing “Text Histogram” will cause the viewer to show the trace counts in a table, as described in this section. Choosing “Histogram Graph” will cause the viewer to display the results as a graphical histogram, as described in the next section.

The vertical range of the histogram is set in the “Range” text box. If the viewer is counting calls, the range should be specified as a number of calls. If it is calculating time spent, the range should be given in microseconds. The range has no significance when the viewer is in text mode.

Trace Histogram Viewer - Graphical Mode

Most of the information in the previous section also applies to a Trace Histogram Viewer in graphical mode. This section describes features specific to the graphical mode.

When in graphical mode, the viewer shows the display limit of the histogram in the top left-hand corner, above the vertical axis. This is a count if the viewer is set to display the number of calls; if the viewer is set to display times then the number will be in microseconds. If the value of particular bar exceeds the display limit, then the bar is shown “broken” at the top.

The bars are labeled along the horizontal axis. Each bar corresponding to a particular trace node is labeled with that node’s class name.

When the mouse is moved over a bar in the histogram, the viewer shows numerical values for it at the top of the display. If the viewer is displaying Number of Calls, this is

* Which corresponds to the number of non-middle, non-end trace points encountered.

just the count. If the viewer is displaying times, it shows the total time spent within the Start/End pairs - which corresponds to the bar height - the minimum Start/End interval logged, the average, and the maximum. The format is Total Min/Avg/Max.

If there is only one trace node present in the histogram, the text summary is always shown.

In graphical mode, the zero button is placed at the origin of the histogram. It operates in the same way, regardless of mode.

Trace Call Profile Viewer

The Trace Call Profile Viewer shows the distribution of call times for a single trace node. It can show you how long the majority of the calls to a particular routine took, how much variance of call time there was, and how many outliers turned up at the extreme edges of the range. This can help you to decide whether worst-case optimization is justified.

The display consists of a series of rows. The first column of each row shows a time range in microseconds. The second column shows how many start/end pairs were encountered whose end time stamp minus the start time stamp fell into that time range. The rows are sorted by time range; rows which do not contain any events are not shown.

Clicking on a column header will sort the data by that field. The data is resorted as it is updated.

Calls on the edge of a bucket are assigned on the basis of where they fall with respect to the lower limit of the bucket. For instance, a call whose time was 49.85 μ s would be counted in the 45..49 μ s bucket, not the 50..54 μ s bucket.

Options

There is a single Trace Call Profile Viewer option: the time range bucket width. By changing this value, you can control the granularity with which the viewer reports its findings. A small bucket width will produce a larger number of ranges with a smaller total in each range; a larger bucket width will produce fewer ranges containing more events.

Trace Text Viewer

The Trace Text Viewer displays all the trace data points in the current time interval for a single node. The traces are displayed as lines of text showing the time stamp, Event Tag, and event order (Start, Middle, or End).

There is also a column for “Trace String.” In addition to the trace data described in the “Viewers” section in the “Overview” chapter, the client of the Instrumentation Library can define and record private data at runtime through the use of the `InstLogTraceEventWithData()` call. This data is formatted by the Instrumentation Library as text using printf-style descriptors, and displayed as the Trace String by the viewer.

Options

The Trace Text Viewer options consist of a series of check boxes that control the visibility of the various columns. The trace time stamp, event tag, event order and trace string can all be turned on or off. There are also check boxes for showing the trace’s kernel process name, task name, and task priority; these options are currently unsupported by the Instrumentation Library.

Magnitude Graph Viewer

The Magnitude Graph Viewer displays instrumentation magnitude data. The data point displayed is the last sample taken prior to the current time. Since this is a statistic, there is typically one data point present for each sampling interval during which the node was enabled to record data.

The data is presented on a two-dimensional graph, with time along the horizontal axis and magnitude value along the vertical axis. The vertical axis is labeled with the magnitude range at top and bottom.

The data appears as vertical bars. The bar that appears on the right-most edge of the graph corresponds to the data point at the current time - the time which is shown beside the time slider control. Bars to the left of it represent successively earlier times, with each pixel representing the current time division. For example, if the current time division is 2 seconds, then four adjacent bars cover a time span of 8 seconds.

The bar height represents the value of the data point; see the section below on options for more information. When the mouse is moved over data on the graph, the value corresponding to the point under the mouse is displayed at the bottom of the graph.

It is possible to present two magnitude nodes at once on the same graph; in this case their bars are drawn in different colors. This viewer cannot show more than two nodes.

Sometimes the viewer will be drawn disabled with the message, “No Data for Specified Class.” This means that the Instrumentation Spooler did not sample the node and write out the results during the current time interval.

Options

The Magnitude Graph Viewer supports three sets of options. The first controls scaling and the position of the legend. If scaling is set to “Auto scale”, then the limits of the vertical axis will be determined by the range of the data being viewed. In order to avoid having to scan through the data, the range is set dynamically. If a data point is encountered whose value is larger than the current range, the range will be reset to accommodate the new value. (This can cause the graph to “jump around.”) If “Manual scale” is chosen, then the range of the vertical axis is determined by the numbers typed into the “Min value” and “Max value” boxes. The units are the same as those of the data being displayed.

The viewer can also display a legend, which shows which color belongs to which node.

The other options control the presentation of the data of the first and second nodes. Each node may have a different color and line style associated with it. In addition, you can control how the data magnitude (the height of the line) is calculated.

Choosing “Instantaneous” will display the last magnitude value set prior to taking the sample. Choosing “Sliding Average” will calculate the result by averaging all data points present in the interval between the current time and the current time minus the time specified in the “Range” box.

The “Exponential Decay” calculation is similar to the “Sliding Average”, except that the weight of the data in the average decreases as it occurs further away from the current time.

“Running Average” is currently unsupported.

Growth Graph Viewer

The Growth Graph Viewer displays instrumentation growth data. Since this is a statistic, there is typically one data point present for each sampling interval during which the node was enabled to record data.

The data is presented on a two-dimensional graph, with time along the horizontal axis and growth rate along the vertical axis. The maximum rate displayable at the current scale is shown at the top of the vertical axis.

The data appears as vertical bars. The bar that appears on the right-most edge of the graph corresponds to the data point at the current time - the time which is shown beside the time slider control. Bars to the left of it represent successively earlier times, with each pixel representing the current time division. For example, if the current time division is 2 seconds, then four adjacent bars cover a time span of 8 seconds.

The bar height represents the growth rate calculated from the data point for that particular time; see the “Options” section, below, for details on how the growth is calculated.

It is possible to present two growth rates at once on the same graph, by dragging two nodes into the viewer in Layout mode. In this case, their bars are drawn in different colors. This viewer cannot show more than two nodes.

Sometimes the viewer will be drawn disabled with the message, “No Data for Specified Class.” This means that the Instrumentation Spooler did not sample the node and write out the results during the current time interval.

Options

The Growth Graph Viewer supports three sets of options. The first – “General Settings” – controls scaling, the legend, and user interaction. The growth rate can be expressed as a number of increments per microsecond, millisecond, second, or minute. If scaling is set to “Auto scale”, then the limits of the vertical axis will be determined by the range of the data being viewed. In order to avoid having to scan through the data, the range is set dynamically. If a data point is encountered whose growth rate is larger than the current range, the range will be reset to accommodate the new value. (This can cause the graph to “jump around”.) If “Manual scale” is chosen, then the range of the vertical axis is determined by the growth rates typed into the “Min value” and “Max value” boxes.

The viewer can also display a legend, which shows which color belongs to which node.

User Interaction is not currently supported.

The other options control the presentation of the data of the first and second nodes. Each node may have a different color and line style associated with it. In addition, you can control how the growth rate (the height of the line) is calculated.

Choosing “Instantaneous” will calculate the growth rate by dividing the amount that the running total has increased in the most recent sample by the sample time. Choosing “Sliding Average” will calculate an average of all of the instantaneous growths present in the interval between the current time and the current time minus the time specified in the “Range” box.

The “Exponential Decay” calculation is similar to the “Sliding Average”, except that the weight of the data in the average decreases as it occurs further away from the current time.

 **Statistics Text Viewer**

The Statistics Text Viewer allows you to display both Magnitude and Growth instrumentation points in a compact text-based display. (For descriptions of Magnitude and Growth data, see the “Viewers” section in the “Overview” chapter.)

The data is displayed as lines of text within the viewer, one line to an instrumentation node. In the default configuration, each magnitude line contains the name of the node, the value, the minimum value, the average, and the maximum. Each growth node is shown by the name, the total value, the minimum increment and the maximum increment. Note that growth statistics are displayed here as totals, not as rates.

See the section on “Options,” below, for more information on configuring the display.

For each node, the data point displayed is the one whose time stamp is closest to (but not later than) the current time, which is displayed beside the time slider control. If no data point exists anywhere between the current time and four sampling periods before the current time, the values are shown as zeros to reflect the fact that the data was no longer being recorded.

Clicking on a column header will sort the data by that field. The data is resorted as it is updated.

Options

By default, the scale is 1:1; choosing a different number will cause all the values in the display to be scaled by the new factor. (This option is not currently supported.)

Each of the Value, Minimum, Average and Maximum columns may be shown by checking the appropriate box, and hidden by unchecking it.

 **Histogram Graph Viewer**

The Histogram Graph Viewer displays instrumentation histogram data. Since this is a statistic, there is typically one data point present for each sampling interval during which the node was enabled to record data. Each data point consists of a count of the number of “hits” in each histogram bucket, including the Overflow bucket.

The data is displayed as a two-dimensional vertical histogram, with the bucket ranges along the x-axis and the counts along the y-axis. Collectively, all the counts represent a single data point; the one displayed is the one closest to (but not greater than) the current time.

It is possible to present two histogram nodes at once on the same graph, by dragging two nodes into the viewer in Layout mode. In this case, their bars are drawn over one another in different colors. This viewer cannot show more than two nodes.

Sometimes the viewer will be drawn disabled with the message, “No Data for Specified Class.” This means that the Instrumentation Spooler did not sample the node and write out the results during the current time interval.

Options

By selecting “General Settings” in the Viewer Options popup menu, you can set the horizontal and vertical ranges of the histograms manually, or specify that they be calculated dynamically from the values of the histogram. (This can cause the bar heights to “jump around” as new maximums are encountered.) You can also specify whether the viewer will draw a legend showing which node is drawn in which color. The user interaction option is not currently supported.

The main and the secondary histogram - if present - may be configured separately by choosing the “Main Histogram” or “Secondary Histogram” items from the Viewer Options popup menu. You can set the color that the bars get drawn in, or specify that each bar be drawn in a different color from its neighbor. You can also set the name that is drawn in the legend representing the histogram node, and specify whether the bars be filled in (the default) or as empty frames (better for displaying two histograms at once).

Histogram Text Viewer

The Histogram Text Viewer displays instrumentation histogram data for a single node. Its function is essentially the same as the Histogram Graph Viewer (see above); the major difference is that it displays the histogram as a list of bucket ranges in one column, with the counts for each bucket in the other. It’s advantage over the graphical version is that it allows you to see exact counts, which is sometimes important.

The Histogram Text Viewer can also display split histograms.

Clicking on a column header will sort the data by that field. The data is resorted as it is updated.

Currently, it has no viewer-specific options.

Tally Viewer

The Tally Viewer displays instrumentation tally data for a single node. Since this is a statistic, there is typically one data point present for each sampling interval during which the node was enabled to record data.

Collectively, all the buckets and all the counts represent a single data point; the one displayed is the one closest to (but not greater than) the current time. In text mode, the viewer presents each bucket as a line of text; the first column shows the bucket identifier; the second shows the count. Note that you can sort by either column by clicking on the column header text. In graph mode, each bar represents the tally count for the bucket identifier specified along the horizontal axis.

Sometimes the viewer will be drawn disabled with the message, "No Data for Specified Class." This means that the Instrumentation Spooler did not sample the node and write out the result at any time from the beginning of the data sample to the current time.

Options

The Tally Display Options enable you to set the display mode to text or graph. The range of the vertical axis shown in graph mode can be set by typing the appropriate value in the "Range" text box. Finally, the viewer can be configured to present the bucket identifiers in decimal, hexadecimal, or text OSType format.

Static Text Viewer

The static text viewer allows you to place static text into your viewer document for the purpose of labeling other views or for adding notes. You can choose to have the text appear in a large or a small typeface. Choosing "Draw Border Box" from the Viewer Options window causes the text to be drawn within a box.