

Performance Tuning using Sampler

The Sampler is an application that will help you discover performance problems with your application. To do this Sampler will sample the call stack of your program over an period of time, and then present the results in a browser.

To start the sampling process click the Start button. If a pid (Process ID) has been provided, the sampling begins immediately. If not, a Process Open Panel appears, and you can choose a process to be sampled. Once started, the Start button becomes a Stop button, which you click when you've finished sampling your program.

The Interval: field is the time in milliseconds between each sample. The default time is 20ms.

Browsing the collected samples

Once the sample is collected, you can now browse the data in two modes: tree and flat. The tree mode start you at the `main()` or `start()` function, and if you click on this entry in the browser, the children of this function are displayed in the next column.

`browser.tiff` →

The numeric field in each of the browser items is the number of times that procedure was found during the sample operation.

Clicking the flat radio button puts you into flat mode. In this mode, the entire tree is flatten, where the number of time a procedure occurs in the sampled stacks is counted and then sorted.

`flat.tiff` →

Clicking on a browser item in flat mode will show you the flattened subtree

under the browser item you clicked on.

Filtering the collected samples

You can reduce the amount of sampled data you are viewing by using the filtering mechanism. Clicking the Filter check box will bring up the Find Panel.

find.tiff ↵

Type in the text you want to filter with, and then click the Filter button. To disable the filtering, turn off the Filter check box in the main window.