

# BestPerf

## Performance & Best Practices

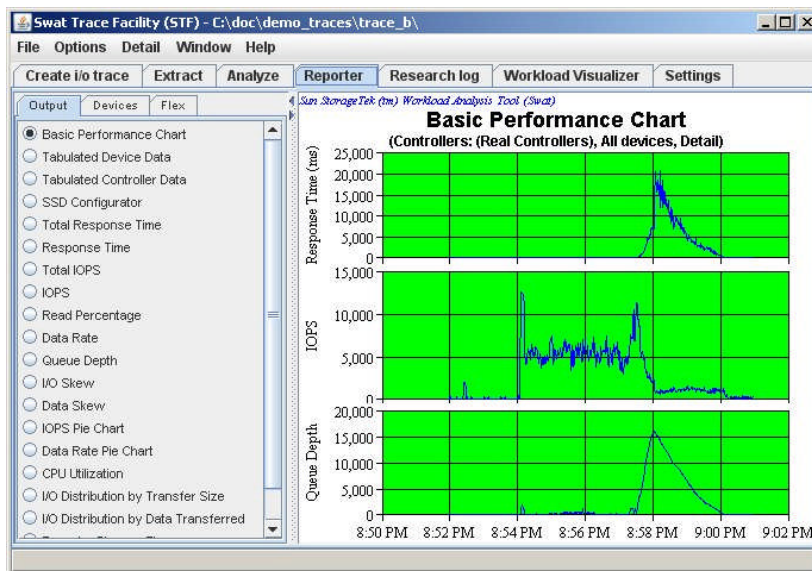
Tuesday Jul 14, 2009

## Storage performance and workload analysis using Swat.

Swat (*Sun StorageTek Workload Analysis Tool*) is a host-based, storage-centric Java application that thoroughly captures, summarizes, and analyzes storage workloads for both Solaris and Windows environments.

This tool was written to help Sun's engineering, sales and service organizations and Sun's customers understand storage I/O workloads.

Sample screenshot:



Swat can be used for among many other reasons:

- Problem analysis
- Configuration sizing (just buying x GB of storage just won't do anymore)
- Trend analysis: is my workload growing, and can I identify/resolve problems before they happen?

Swat is storage agnostic, so it does not matter what type or brand of

Currently Viewing

You are viewing the entry *Storage performance and workload analysis using Swat*.

« [World Record TPC-H@3...](#) | [Main](#) | [Vdbench: Sun Storage...](#) »

About This Blog

BestPerf contains many different performance results and the best practices learned from doing this work on the broad range of Sun's products. The authors of this blog are engineers from Sun's Strategic Applications Engineering group.

Tag Cloud

[\\$/perf 6048](#) [7410 abaqus](#) [amd](#) [ansys](#)  
[application bandwidth](#)  
[benchmark blade cfd](#) [cluster cmt](#)  
[consolidation container cpu2006 current](#)  
[database](#) [ddr dss dtrace exadata f5100](#)  
[flash fujitsu fusion-middleware gcc](#)  
[glassfish hp hpc ib ibm infiniband](#)  
[intel java linux mail mcae mpi](#)  
[mysql nastran oltp oow opensolaris](#)  
[openstorage oracle peoplesoft](#)  
[per-core performance](#) [qdr ras](#)  
[redhat sap-sd scalability seg seismic](#)  
[solaris sparc64 spc-2 spec](#)  
[speccpu specfp specint](#)  
[specjappserver specjbb specmail](#)  
[specomp specpower specweb](#) [ssd](#)  
[storage sun sun-studio sunstudio](#)  
[suse swat t5220 t5240 t5440 t6320](#)  
[tpc-c tpc-h vdbench virtualization](#)  
[vmware watt/perf watts weather web](#)  
[web2.0 windows world-record x2200](#)  
[x2270 x6270 x6275 x64 x86 zeus zfs](#)

Search

Latest Entries

[BestPerf Index 2 August 2010](#)

[BestPerf Index 1 July 2010](#)

[Sun Fire X2270 M2 Achieves Leading Single Node Results on ANSYS](#)

storage you are trying to report on. Swat reports the host's view of the storage performance and workload, using the same Kstat (Solaris) data that iostat uses.

Swat consists of several different major functions:

- [Swat Performance Monitor \(SPM\)](#)
- [Swat Trace Facility \(STF\)](#)
- [Swat Trace Monitor \(STM\)](#)
- [Swat Real Time Monitor](#)
- [Swat Local Real Time Monitor](#)
- [Swat Reporter](#)

## Swat Performance Monitor (SPM):

Works on Solaris and Windows. An attempt has been made in the current Swat 3.02 to also collect data on AIX and Linux. Swat 3.02 also reports Network Adapter statistics on Solaris, Windows, and Linux. A Swat Data Collector (agent) runs on some or all of your servers/hosts, collecting I/O performance statistics every 5, 10, or 15 minutes and writes the data to a disk file, one new file every day, automatically switched at midnight.

The data then can be analyzed using the Swat Reporter.

## Swat Trace Facility (STF):

For Solaris and Windows. STF collects detailed I/O trace information. This data then goes through a data Extraction and Analysis phase that generates hundreds or thousands of second-by-second statistics counters. That data then can be analyzed using the Swat Reporter. You create this trace for between 30 and 60 minutes for instance at a time when you know you will have a performance problem.

A disk I/O workload traced using Swat can be replayed on any test system to any type of storage using Vdbench (an other of my tools, available at <http://vdbench.org>). This allows you to trace a production I/O workload, bring the trace data to your lab, and then replay that I/O workload on whatever storage you want. Want to see how the storage performs when the I/O rate doubles or triples? Vdbench Replay will show you. With this you can test your production workload without the hassle of having to get your data base software and licenses, your application software and licenses, or even your production data.

*Note: STF is currently limited to the collection of about 20,000 IOPS. Some development effort is required to handle the current increase in IOPS made possible by Solid State Devices (SSDs).*

### [FLUENT Benchmark](#)

[Sun Fire X2270 M2 Demonstrates Outstanding Single Node Performance on MSC.Nastran Benchmarks](#)

[Sun Fire X2270 M2 Achieves Leading Single Node Results on ABAQUS Benchmark](#)

[Sun Fire X2270 M2 Sets World Record on SPEC OMP2001 Benchmark](#)

[Sun Fire X4170 M2 Sets World Record on SPEC CPU2006 Benchmark](#)

[Sun Blade X6270 M2 Sets World Record on SPECjbb2005 Benchmark](#)

[Sun Fire X4270 M2 Sets World Record on SPECjbb2005 Benchmark](#)

[Sun Fire X4470 Sets World Records on SPEC OMP2001 Benchmarks](#)

### Feeds

[All](#)

[/Benchmark](#)

[/Best Practices](#)

[/Other topics](#)

[/Sun](#)

[Comments](#)

### Categories

[All](#)

[Benchmark](#)

[Best Practices](#)

[Other topics](#)

[Sun](#)

### Other Blog Features

[blogs.sun.com](#)

[Weblog](#)

[Login](#)

*Note: STF, while collecting the trace data is the only Swat function that requires root access. This functionality is all handled by one single KSH script which can be run independently. (Script uses TNF and ADB).*

## Swat Trace Monitor (STM):

With STF you need to know when the performance problem will occur so that you can schedule the trace data to be collected. Not every performance problem however is predictable. STM will run an in-memory trace and then monitors the overall storage performance. Once a certain threshold is reach, for instance response time greater than 100 milliseconds, the in-memory trace buffer is dumped to disk and the trace then continues collecting trace data for an amount of seconds before terminating.

## Swat Real Time Monitor:

When a Data Collector is active on your current or any network-connected host, Swat Real Time Monitor will open a Java socket connection with that host, allowing you to actively monitor the current storage performance either from your local or any of your remote hosts.

## Swat Local Real Time Monitor:

Local Real Time Monitor is the quickest way to start using Swat. Just enter './swat -l' and Swat will start a private Data Collector for your local system and then will show you exactly what is happening to your current storage workload. No more fiddling trying to get some useful data out of a pile of iostat output.

## Swat Reporter:

The Swat Reporter ties everything together. All data collected by the above Swat functions can be displayed using this powerful GUI reporting and charting function. You can generate hundreds of different performance charts or tabulated reports giving you intimate understanding of your storage workload and performance. Swat will even create JPG files for you that then can be included in documents and/or presentations. There is even a batch utility (Swat Batch Reporter) that will automate the JPG generation for you. If you want, Swat will even create a script for this batch utility for you.

Some of the many available charts:

- Response time per controller or device
- I/O rate per controller or device
- Read percentage
- Data transfer size
- Queue depth

### Bookmarks

---

[Best Practices Wiki](#)

---

[blogs.sun.com](#)

---

[glennf \(SAE Eng\)](#)

---

[henk \(SAE Eng\)](#)

---

[vince \(SAE Eng\)](#)

---

[MPITips \(SAE Eng\)](#)

---

[mheckel \(SAE Eng\)](#)

### Important Sun Links

---

[java.com](#)

---

[java.net](#)

---

[opensolaris.org](#)

- Random vs. sequential (STF only)
- CPU usage
- Device skew
- Etc. etc.

Swat has been written in Java. This means, that once your data has been collected on its originating system, the data can be displayed and analyzed using the Swat Reporter on ANY Java enabled system, including any type of laptop.

For more detailed information go to [\(long URL\)](#) where you can download the latest release, Swat 3.02.

You can find continuing updates about Swat and Vdbench on my blog:  
<http://blogs.sun.com/henk/>

Henk Vandenberg

Posted on: Jul 14, 2009 ::: Posted by: Henk Vandenberg ::: Category: Best Practices ::: Tags: [performance](#) [storage](#) [swat](#) [vdbench](#) :::  
Permanent link to this entry :::

## Comments:

## Post a Comment:

Comments are closed for this entry.

This blog copyright 2010 by John Henning