

Sun StorEdge™ T3 Array for the Enterprise DAS Performance -- RAID Levels 0, 1+0, and 5

This paper will show you T3 direct access storage (DAS) test performance results based on the PeopleSoft GL 8.14 batch load benchmark for customers running Oracle 8i.

Detailed performance tests results include RAID-0, RAID-1+0 and RAID-5 on three different array configurations: 2.6 TB, 5.2 TB and 10.4 TB of available RAW DRIVE storage. These results prove that:

1. **Performance scales with available storage.** Each test shows a clear increase in array performance as available storage is increased, regardless of RAID level. *As you increase available storage, the T3 arrays become even faster!*
2. **T3 DAS RAID 5 performance approaches that of RAID level 0, and then surpasses it as available storage is increased.** *Why debate RAID-0 speed vs. RAID-5 safety -- get the best of both with a T3!*
3. **Optimum performance right out-of-the-box!** No tweaking or fine-tuning is needed because the T3 storage solution is fully optimized for customers running Oracle. *Simply install, power-up, and configure the T3 with a volume manager and go to work!*

You Can Easily Validate These Results Yourself

This performance test suite is designed to help customers to quickly and easily validate these test results inside their own data centers; on their own systems; using out-of-the-box Sun StorEdge™ T3 arrays.

The results show in this paper can be easily duplicated because it uses an industry standard benchmark, which is based on tests conducted in Broomfield, CO on July 2001 by PeopleSoft and Sun. The results of those tests were later published by PeopleSoft and titled "Peoplesoft 8 General Ledger (With Combination Editing) Using Oracle8i On A Sun Microsystems' Sun Fire™ 6800 Server -- Revision 1.0" (C/N 0416-0801).



STORAGE TERMS USED IN THIS PAPER

Direct Access Storage (DAS): Direct connection to server: Rigid configuration with the fastest possible throughput.

RAID-0: Raid Level-0 offers data striping only. It is typically used for high-speed, short-term applications such as digital video editing.

RAID-1+0: Two RAID levels (1+0) are combined to provide RAID-1 disk mirroring plus RAID-0 data striping. Data remains safe even if half of the drives suffer a catastrophic failure.

RAID-5: This RAID level offers full data striping of both data and error correction bytes across multiple drives. RAID-5 is faster than RAID-1+0, but Raid 1+0 is much more reliable.

Data Striping: A single drive storage system can only use one drive for everything. Arrays can use data striping, and distribute file blocks across multiple drives. Because data striping combines the performance of multiple drives, arrays can offer significantly faster read and write performance than single drive systems.

Disk Mirroring: As the name implies, a copy of each data block is written to two or more drives. *Also known as RAID-1.*

Wide Thin Disk Striping Methodology: As a general rule, 80% of the storage array accesses are confined to 20% of the data, which can result in inefficient hot spots. Sun's implementation of the Wide Thin Disk Striping methodology eliminates hot spots because it builds a wide thin strip across all available spindles. Consequently, each major database object (large index or table) is striped over so many disks that hot spotting is impossible. This is because a part of every table and part of every index is on every disk. No matter which table gets hot, all of the disks are used, thereby eliminating I/O bottlenecks.

Duplicating The Test Results In Your Own Data Center

The system under test (SUT) used to obtain these results was a Sun Fire 6800 with the following configuration:

- Sun Fire 6800: 24 x 750Mhz UltraSparc III CPUs. Typical configuration for customers using the full PeopleSoft application suite.
- Installed Memory: 48 GB. In order to ensure an I/O intensive test, the default Oracle System Global Area (SGA) was reduced to 78 MB.
- A total of 16 Sun StorEdge™ T3 Arrays with nine 73GB drives each, for a total of 10.4 TB of available storage. The arrays were tested out-of-the-box and the default block size of 64Kb. Cache mirroring was disabled.

The T3 SUT arrays were configured as RAW DRIVES to ensure consistent test results from one system to the next. This is because array drives that have been formatted with the UNIX File System (UFS) or the VERITAS UNIX File System (VxFS) yield different results.

SUT Default Configuration Software Stack

All of the software and applications used in these performance tests were installed using the default out-of-the-box configurations. No tweaking or fine-tuning is required.

- PeopleSoft 8.14 GL batch benchmark
- Oracle 8i, version 8.1.6
- Solaris8 update 7/01, with the default off setting for the Direct I/O option
- VERITAS Volume Manager, Version 3.2

Configuring Your Own Test System

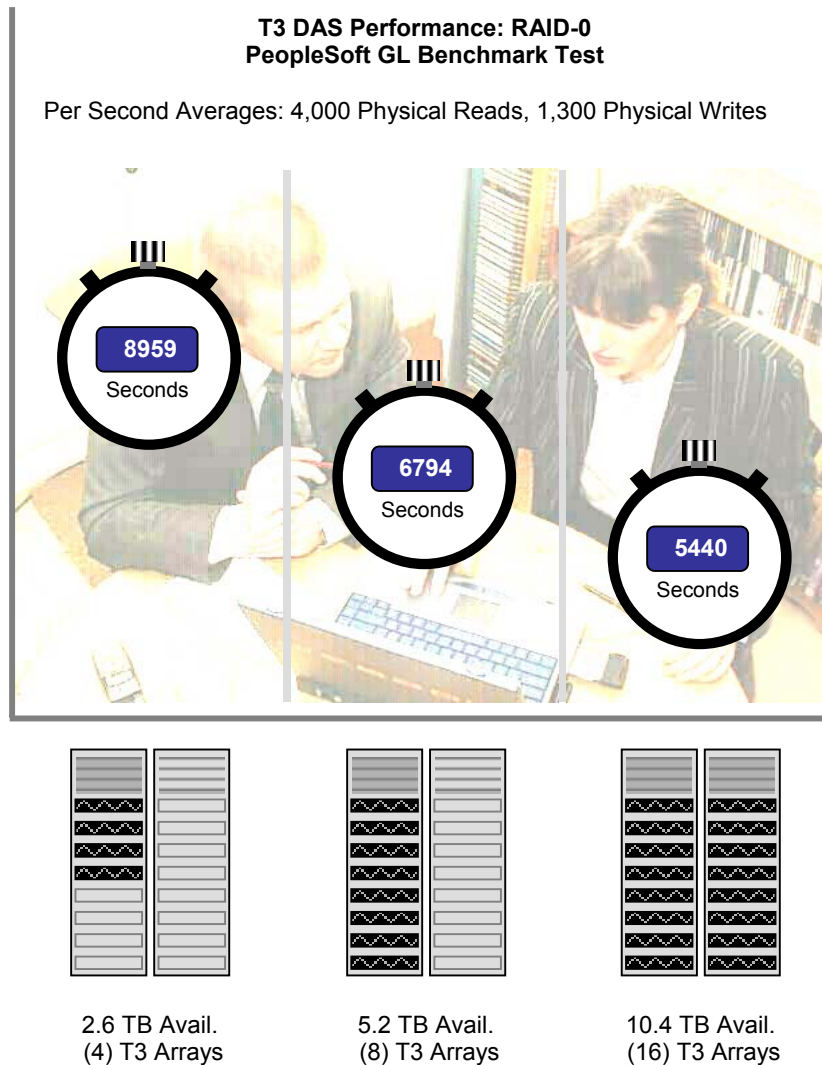
After installing the T3 arrays as RAW DRIVES, install Solaris version 8.0 or greater, Oracle 8.0 or greater, Veritas Volume Manager 3.2 or greater and the PeopleSoft the 8.14 batch loader program, you must then configure your performance test system as follows:

- `db_file_multiblock_read_count = 64`
- `db_block_buffers = 10,000`
- `db_block_size = 8,192`
- `log_buffer = 1048576`
- `timed_statistics = true`
- `shared_pool_size = 34,299,084`
- `sort_area_size = 5,120,000`
- `sort_area_retained_size = 65536`

To ensure consistent results, please note that all of the performance test results shown in this paper were obtained from a dedicated SUT, and that each test was run for 10 minutes.

RAID-0 T3 DAS Performance

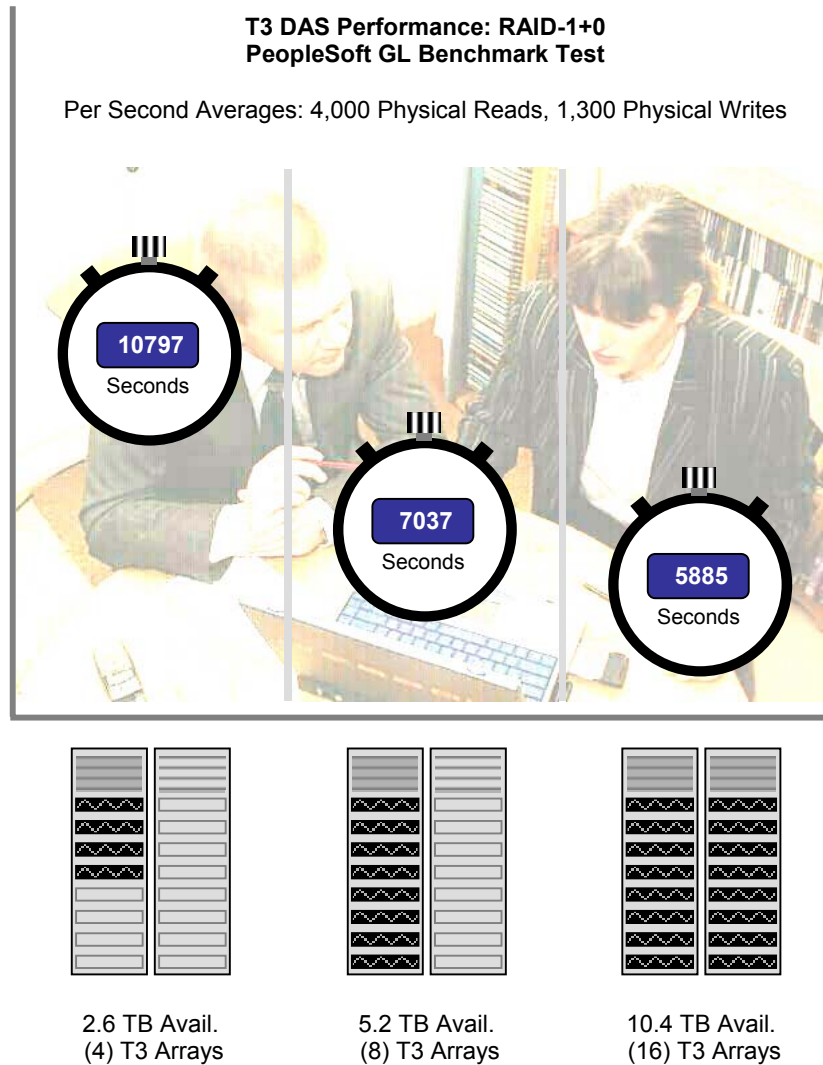
The fastest RAID level of all is RAID-0, because it offers no mirroring backup or error correction. In certain applications where speed is paramount such as digital video editing, RAID-0 on RAW DRIVES offers the fastest performance of any array configuration. For this reason, RAID-0 serves as an excellent baseline when comparing the results of other RAID levels.



RAID-0 on a single drive storage system offers little advantage because all of the data blocks are written to a single drive. However, as drives are added to the array, RAID-0 data striping allows the file blocks to be divided up between multiple drives. The more drives added to the array, the faster the performance, which also holds true for the other RAID levels as well.

RAID-1+0 DAS Performance

In those applications where speed plus the ability to recover from a catastrophic failure, such as fully redundant web servers, SUN has implemented the Wide Thin Disk Striping methodology on the T3, which is referred to as RAID-1+0. This methodology eliminates performance-draining I/O bottlenecks because it builds a wide thin strip across all available spindles.

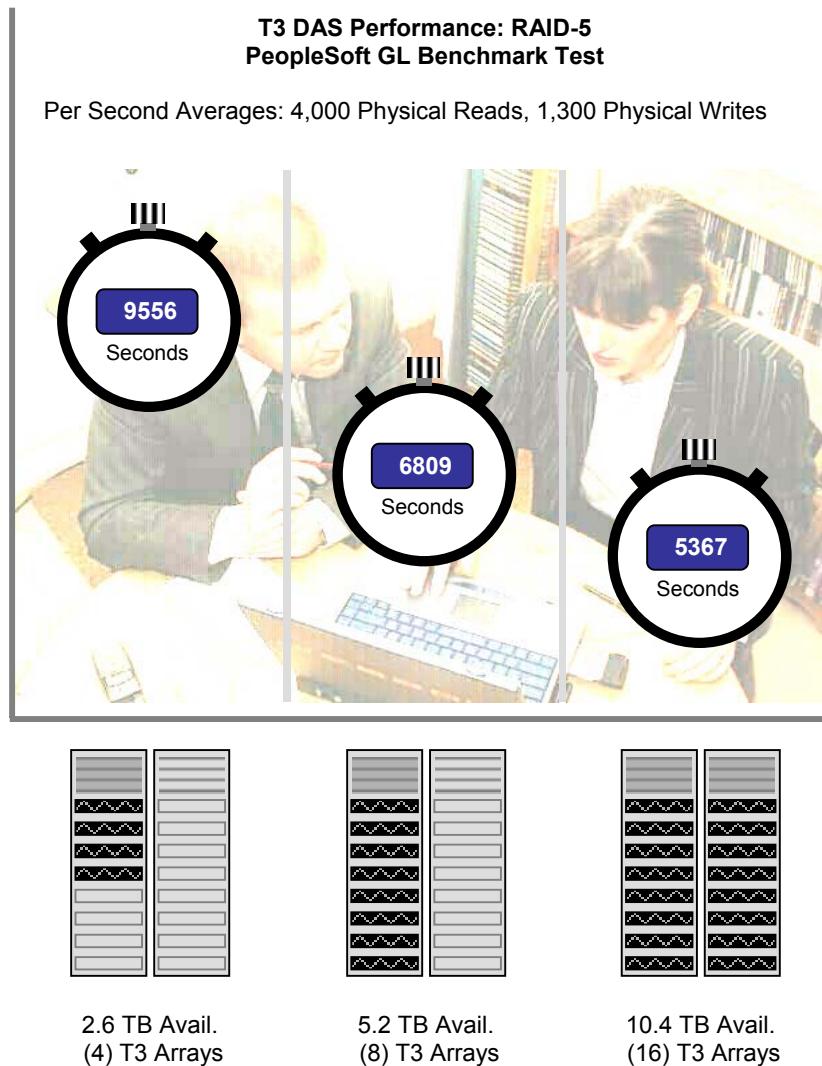


This cache-optimizing methodology is based on Pareto's principle that 80% of production volume usually comes from 20% of the producers. Vilfredo Pareto was a famous 19th century Italian economist and sociologist, known for his application of mathematics to economic analysis.

In terms of storage arrays, Pareto's principal has translated into a Wide Thin Disk Striping methodology on the T3 that automates disk layouts to same degree of efficiency as that of a hand-tuned disk layout.

RAID-5 DAS Performance


For years, enterprise storage users often have had to debate between RAID-1 and RAID-5 in terms of speed and safety. RAID-1 offered redundancy through mirroring but at the expense of having to install twice as many hard drives to store as much data as RAID-5, which offers error correction for maintaining data integrity.



However, things have changed since SUN introduced the T3, because it offers RAID-5 safety with near-RAID-0 performance. This makes RAID-5 on the T3 the open standard choice for the best possible combination of speed and safety.

Summary of Performance Test Findings

DAS T3 arrays in the 2.6 TB - 5.2 TB range perform in a typical manner, where RAID-0 performance is best. However, once the available storage of the T3 reaches the 10.4 TB level, RAID-5 performance is actually superior to RAID-0, right out-of-the-box.



Taking The Next Step

If you need a powerful combination of speed and mirroring safety, SUN's StorEdge™ T3 Array for the Enterprise offers RAID-5 peace of mind with RAID-0 performance, which increases as you add available storage.

To see if the T3 makes sense for your enterprise storage needs, contact your SUN Sales engineer to request a detailed, no obligation analysis of your specific RAID storage requirements.

