

Platypus Technology, Inc.

**79 East Wilder Road
West Lebanon, NH 03784
(603) 298-7455
www.platypus.net**

Platypus Technology: Accelerating Key Application Performance

Preface

It is well known that the amount of electronically stored digital information that enterprises must keep on hand is growing to unprecedented scale. However, media hype about storage growth overlooks a critical fact: only a portion of that stored information demands reliable, near-immediate access. Almost no information that relates types of stored data to application performance has heretofore been available.

In nearly every mission- or enterprise-critical application, the application refers to database information. Maintaining application-database performance in the face of ever increasing user loads can present enterprise IT managers with myriad problems. If the manager uses conventional disk technology, the cost-performance curve tends to be nonlinear (i.e., cost typically increases at a rate much faster than performance as the amount of storage increases), and the IT buyer gets less and less bang for the buck as the application grows. However, most unconventional technology may involve intolerable disruption to operations and — worst of all — risk. Clearly, what is needed is a safe, cost-effective way to accelerate database application performance without a risky bet on new technology.

This *Profile* examines how Platypus Technology presents enterprises with a cost-effective, high-performance complement or alternative to disk storage. The Platypus solution includes time-tested, no-single-point-of-failure solid-state memory with built-in data protection, which can be applied to solve database I/O performance bottlenecks, as well as to replace disk-to-disk performance in myriad other non-database applications.

Executive Summary

Enterprises depend on applications that refer to a database to substantially run their businesses. However, internal and external customer-facing operations, or both, are often plagued by database performance (i.e., response time) problems.

Storage used for critical application data has a cost/performance hierarchy. At the base of the pyramid is the least expensive, least robust type of disk array, JBOD (just a bunch of disks) storage. Network attached storage (NAS) is a level above. A storage area network (SAN) fits in higher up the spectrum. A notch above the SAN is the cache storage used in a SAN, and solid-state storage is at the pyramid's apex.

Every enterprise is bound to have an application that justifies the use of a certain amount of each type of storage. In particular, applications that must scale to large numbers of users and maintain high performance, and which characteristically transfer small blocks of randomly located data to/from storage, are particularly well suited to solid-state storage.

Most mission-critical applications involve databases that are less than 64 gigabytes (GB) in size. Oracle, for example, has stated that 75% of its users' database applications are under the 64-GB mark. Additionally, by segmenting databases or using tuning techniques, database administrators can trim the active portions of larger databases to operate in less than 64 GB of storage.

Many applications serving multiple users simultaneously are bound by disk I/O performance and could benefit by using Platypus Technology. Database applications — which can become I/O bound — are a good fit for solid-state storage. Therefore, examining the typical application access profile and what really happens when users pile on to make I/O hit the wall is worthwhile.

There's Something About a Database ...

Databases typically reside on conventional rotating disk storage. Usually, that disk storage is direct-attached storage, perhaps with hardware equipped with redundant array of independent disk (RAID) protection (e.g., mirroring, striping with parity, clustering, etc.). The general hierarchy of data content goes from less dense (unstructured) to dense (semi-structured) to more dense (structured).

Videos and music are good examples of unstructured files — they tend to be large, and searching and sorting within them is impossible.

Office productivity files, such as Word documents, typify semi-structured files. They tend to be somewhat bulky because of stored attributes (e.g., presentation, font color, formatting, etc.), and they usually allow searching, but not sorting, within a file.

Databases are highly structured files that contain information in compact form. They are designed to support search, sort, and query operations against the data in a file. Usually, the only time that database files are accessed contiguously is when reports or queries are generated. Most access to database files is driven by applications that query for small items located randomly within the database.

...That Fits Well with Solid-State Storage

Any time an entire database — or its most active portions — can be completely contained within solid-state storage, all of its data can be accessed uniformly at electronic speed — microseconds or hundreds of nanoseconds.

In contrast, mechanical disks locate and deliver data in tens of milliseconds. Various schemes use multiple disks with data laid out in “stripes” on teamed disks to create disk subsystems that can present data to a requesting host computer in fewer milliseconds.

Several milliseconds equal several thousand microseconds and are thousands of times longer than hundreds of nanoseconds. In an effort to reduce the orders-of-magnitude chasm between electromechanical and electronic speeds, high-performance disk subsystems use cache memory to enable data to be stored contiguously to the original disk storage data. This data can then be accessed at electronic speed. In general, enterprise-class disk subsystems work using cache memory in this way:

1. Make the initial access at electromechanical speed — milliseconds, at best.
2. Automatically load the cache unit with disk data stored contiguously.
3. Hope for the best — if the requested data has been loaded into the cache, it can be sent to the user at electronic speed. If not, repeat steps one and two.
4. Somehow, automatically manage cache resources in relation to its size ratio, compared with the amount of disk storage, because all users and applications compete for scarce cache space. That is, the disk subsystem probably does not know anything about the application and file types in use, and it may assign cache resources unwisely.

Algorithms Do Not Deal Effectively with Randomness

Cache management algorithms simply fail to yield good results when a business-critical application typically reads or writes small portions of highly random file data. Instead, the disk fetches of contiguous data into a cache are futile; subsequent data accesses will tend to perform at electromechanical rates because they are forced to be — in fact — initial accesses.

This scenario raises a question: *Why use disk and cache at all?* With solid-state storage, all the data is immediately available at nearly main-memory speed. There may be no need for disk drives or cache management logic. Instead, databases, key small files, or indices of large databases for the most enterprise-critical applications can be kept on solid-state storage, providing maximum performance.

Enterprises Should Not Be Forced to Purchase Unnecessary Equipment

Besides their dependence on cache memory, large disk storage subsystems also require that IT managers add disk storage to gain additional ports for more simultaneous users. More capacity, in turn, increases access randomness, which lowers the performance of cache and the entire disk storage system, and makes management more difficult.

In fact, adding solid-state storage can enable an existing storage system based on hard disk drive (HDD) technology to realize its actual capacity. With solid-state storage taking over the most heavily accessed files, RAID arrays no longer need to limit their capacity to just a small percentage of each disk in order to squeeze out performance.

By targeting the true data I/O bottleneck, users of solid-state storage can reduce the quantity of servers needed to host their applications, thereby creating substantial savings not only in hardware, administration, and real estate costs, but also processor-based software license fees.

Holding the Line on Operating Expense

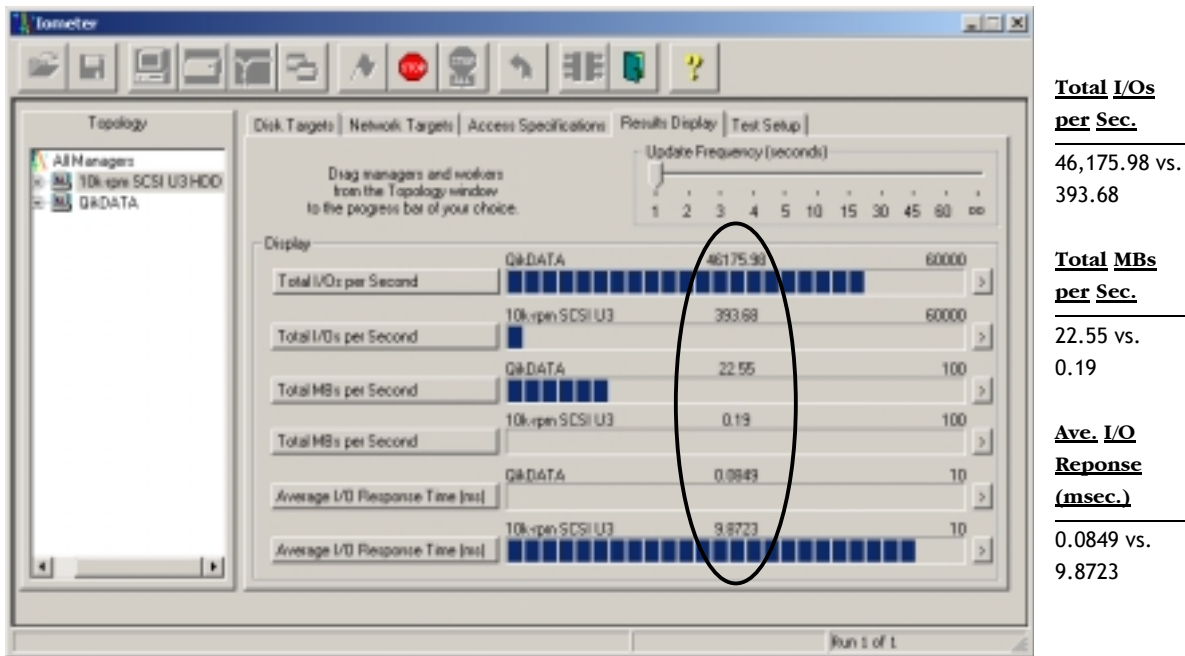
The true cost of owning a storage solution is a combination of capital expense (CAPEX) and operating expense (OPEX). Large HDD-based storage systems can have a lifetime OPEX several fold greater than its CAPEX.

Platypus Technology offers small, easily housed database storage that can access and deliver data *substantially faster than disk storage with CAPEX of less than \$50,000 and negligible OPEX*. In terms of capital outlay, solid-state storage costs more than simple direct-attach disk, but it performs tremendously faster; if database performance counts, the solid-state ROI (return on investment) might be substantially better than HDD-based storage.

Enterprise Application Success Stories

The performance information that a standard benchmark, such as Intel's "Iometer" (Figure 1), may provide is a useful indicator, but real-life results typically provide solid pointers.

Figure 1: *Qik*DATA Versus HDD Iometer Comparisons (Small Data Blocks)



Source: Platypus Technology, August 2002

Monster.com

At Monster.com, which claims to be the world’s largest career search Web site, the key application used is a Siebel Systems customer relationship management (CRM) suite. The CRM suite is used by more than 800 of its salespeople and sales managers to access a 28-GB Microsoft SQL (Structured Query Language) database for corporate interactions and billing. The system also reconciles more than 18,000 accounts each month to produce various reports, plan the next month’s activities, and calculate sales commissions.

Monster.com moved its most heavily accessed “hot” files from HDD RAID storage to a pair of 16-GB Platypus *Qik*DATA units. The results included a 30% reduction in query time, a 50% improvement in database indexing time, the ability of existing servers to handle 25% more simultaneous users, and protection from downtime in a

\$2.3 million-per-day revenue environment. Average customer response time in the Monster.com call center dropped from 10 seconds to about two seconds.

Monster.com’s e-mail distribution was another mission-critical application. The three redundant pairs of servers were delivering 120 e-mails per second during peak periods. Company growth demanded this throughput to double. Rather than

adding another six servers, Monster.com transferred the application's message queue to 8 GB of *QikDRIVE* solid-state storage installed within only one redundant server pair. The resulting increase to 1,200 e-mails per second meant a considerable cost saving and dramatic performance improvement.

Computer Associates eTrust Directory

Many enterprises depend upon the CA eTrust Directory to provide secure, simple sign-on functionality reliably, with unlimited size directories and reliability using an RDBMS (relational database management system).

A large financial institution's Web-based services were severely limited by the number of simultaneous users able to log on, but replacing HDD storage with Platypus *QikDATA* solid-state storage boosted simultaneous user log-on capability from 3 to 20 per second with fewer servers — reducing software license costs.

Platypus provides *QikSPY* diagnostic software gratis, upon request, to help companies discover and analyze application server underperformance, as well as identify files that could be moved to *QikDATA* storage in order to remedy application performance problems.

HyperFeed Technologies, Inc.

HyperFeed provides real-time securities quotes, news, business information, and other financial information to professional clients and individual investors. It uses HyperServer rack servers to receive, store, compress, and encapsulate more than 600,000 equity, option, commodity and OTC (over the counter) issues from exchanges in North America.

Each HyperFeed server's directly attached HDD RAID storage was creating performance bottlenecks for HyperFeed's key applications — HyperFeed, HyperView, Orbit, and Mercenary — especially when user demand spiked. Of course, demand for information spikes precisely when important or unsettling news reaches investors, and HyperFeed's users were impatient. Adding Platypus *QikDRIVE8* solid-state cards in the servers eliminated the performance bottleneck. HyperFeed claims that the installation of a *QikDRIVE8* into each HyperServer has increased data flow by one thousand times.

Hyperion Solutions

Hyperion Solutions' core Essbase database engine technology provides mission-critical business intelligence analytics to a number of enterprises. Hyperion placed an entire 32-GB Essbase 6.5 database on a Platypus *QikDATA* and tested the database load, calculate, and restore phases.

The greatest improvement came in the calculation stage, which involved the most complex processing against small amounts of random database information. In the test, the calculation phase took 98% of the overall processing time. With no tuning changes to the database, *QikDATA* immediately improved performance by up to five times. Even after database tuning, *QikDATA* reduced processing time by 40%. Also, with *QikDATA*, Essbase reporting took half the time it had with HDD storage, and the load phase ran 20% faster.

Energis

Energis is a telecommunications and e-business solutions provider listed among the top 100 in the *Financial Times*' FTSE classification system. Energis had 70% growth for the past six years, with 2001 revenue of £840 million.

Already handling 12 million e-mail messages per day on its Exim mail-server spooling directories that run under Red Hat Linux, Energis faced pressure to greatly increase its ability to handle simultaneous e-mail users without increasing server farm floor space. Using *QikDATA* units to replace HDD units storing the "hottest" files, Energis realized an 8- to 10-fold throughput improvement, which allowed it to re-deploy its Internet data centers.

CAD/CAM

Electronic Data Systems (EDS) produces Solid Edge, a leading CAD/CAM (computer-aided design/computer-aided manufacturing) application system. EDS and Platypus tested *QikDRIVE* units in various configurations, and obtained performance improvements ranging from 31% to 124%.

Operational Ease

For operational purposes, Platypus solid-state storage is a disk to both the operating system and the people administering it. In fact, other than by examining its properties, there is no way for the OS or administrator to discern that the storage is not a mechanical disk. No special software is required; any necessary operating system drivers come with the storage on a CD-ROM.

Platypus storage, either directly attached or via a SAN, can be set up as one or more disk volumes that contain databases and/or files and folders, and all application and administrative software can be used without configuration changes. Storage mirroring is RAID-1 (exact duplicate) mirroring. Internal power-fail detect mirroring is automatic.

Software from vendors such as Veritas, BMC Software, Legato, or Computer Associates can be used for additional external (e.g., tape) point-in-time copying or remote mirroring. Database (and certain other) applications embed error recovery and transaction rollback capabilities that can be used as if a disk were in use. Redundant data paths can also be configured.

Aberdeen Conclusions

Platypus Technology has designed a very special architecture, using solid-state storage that is more than simply a level of storage between disk and main computer memory. It is also a replacement for HDD-based systems, operating on the system I/O bus and fully rigged for mission-critical use.

For the future, an Infiniband interface could be added, when the PCI bus or Fibre Channel offerings no longer appear in high-end servers. But for now, Platypus storage is a low-risk, performance-oriented solution.

Early Platypus Technology adopters have come to appreciate that the true cost of storage is the cost of application delivery. For applications that require high-performance delivery, Platypus is an attractive complement or alternative to HDD storage. Accordingly, Aberdeen recommends that businesses with mission-critical application performance challenges seriously consider deploying Platypus solid-state storage.

To provide us with your feedback on this research, please go to www.aberdeen.com/feedback.

*Aberdeen Group, Inc.
One Boston Place
Boston, Massachusetts
02108
USA*

*Telephone: 617 723 7890
Fax: 617 723 7897
www.aberdeen.com*

*© 2002 Aberdeen Group, Inc.
All rights reserved
August 2002*

Aberdeen Group is a computer and communications research and consulting organization closely monitoring enterprise-user needs, technological changes and market developments.

Based on a comprehensive analytical framework, Aberdeen provides fresh insights into the future of computing and networking and the implications for users and the industry.

Aberdeen Group performs specific projects for a select group of domestic and international clients requiring strategic and tactical advice and hard answers on how to manage computer and communications technology. This document is the result of research performed by Aberdeen Group. It was underwritten by Platypus Technology. Aberdeen Group believes its findings are objective and represent the best analysis available at the time of publication.