

Silicon-Based Accelerators:

A Storage Alternative for Improved Application Performance

By Ken Cameron and Toni Sutherland

While processor speeds have raced ahead in the last 20 years, the relatively slow speed of storage systems, most notably hard drives, has bottlenecked the overall data flow in many mission-critical IT infrastructures. Solid-state storage using silicon-based accelerators can circumvent this bottleneck and improve application efficiency. This article details how the e-business Monster.com achieved dramatic performance improvements for both its customer relationship management (CRM) and e-mail applications by adding silicon-based accelerators.

The hard disk drive (HDD) has served IT departments incredibly well for nearly 50 years, but the increasing demands of the new information economy have stretched this technology to its limits. The inability of the HDD to match the speed of other system elements results in unused capacity throughout the network, leaving IT departments with systems that cannot perform to their full potential.

When the speed of disk I/O limits system performance, adding large numbers of spindles to increase performance can be expensive, management-intensive, and ineffective. The latest 15,000 rpm HDDs are capable of streaming 50 MB/sec, but writing data takes twice as long as reading because the controller must verify that each write has been accepted.

Silicon-based application accelerators, also known as solid-state storage, provide dramatic speed and reliability improvements. Silicon-based accelerators can stream data at theoretical rates exceeding 1000 MB/sec and can perform reads and writes at the same speed. Silicon also has an average access time of just 20 microseconds, more than 250 times faster than an HDD, which is significant when reports require billions of accesses. As memory

prices continue to fall dramatically, silicon-based application accelerators become an appealing alternative to HDDs.

Figure 1 shows how silicon-based application accelerators work within a system infrastructure.

Monster.com: A case study of silicon-based application accelerators

Whenever a business-critical application is not currently maximizing its CPU utilization and needs additional capacity and performance, silicon-based application accelerators can help. Applications may face such a need during daily critical peak periods, when batch jobs consume too much time, or when general company growth surpasses available resources.

Continued company growth caused Monster.com, the world's largest online careers site, to experience significant pressure on its existing customer relationship management (CRM) and e-mail applications. By integrating silicon-based application accelerators, Monster.com improved the performance of these applications and avoided the financial and systematic burdens of adding more HDDs.

Silicon-based accelerators
can stream data at
theoretical rates
exceeding 1000 MB/sec.

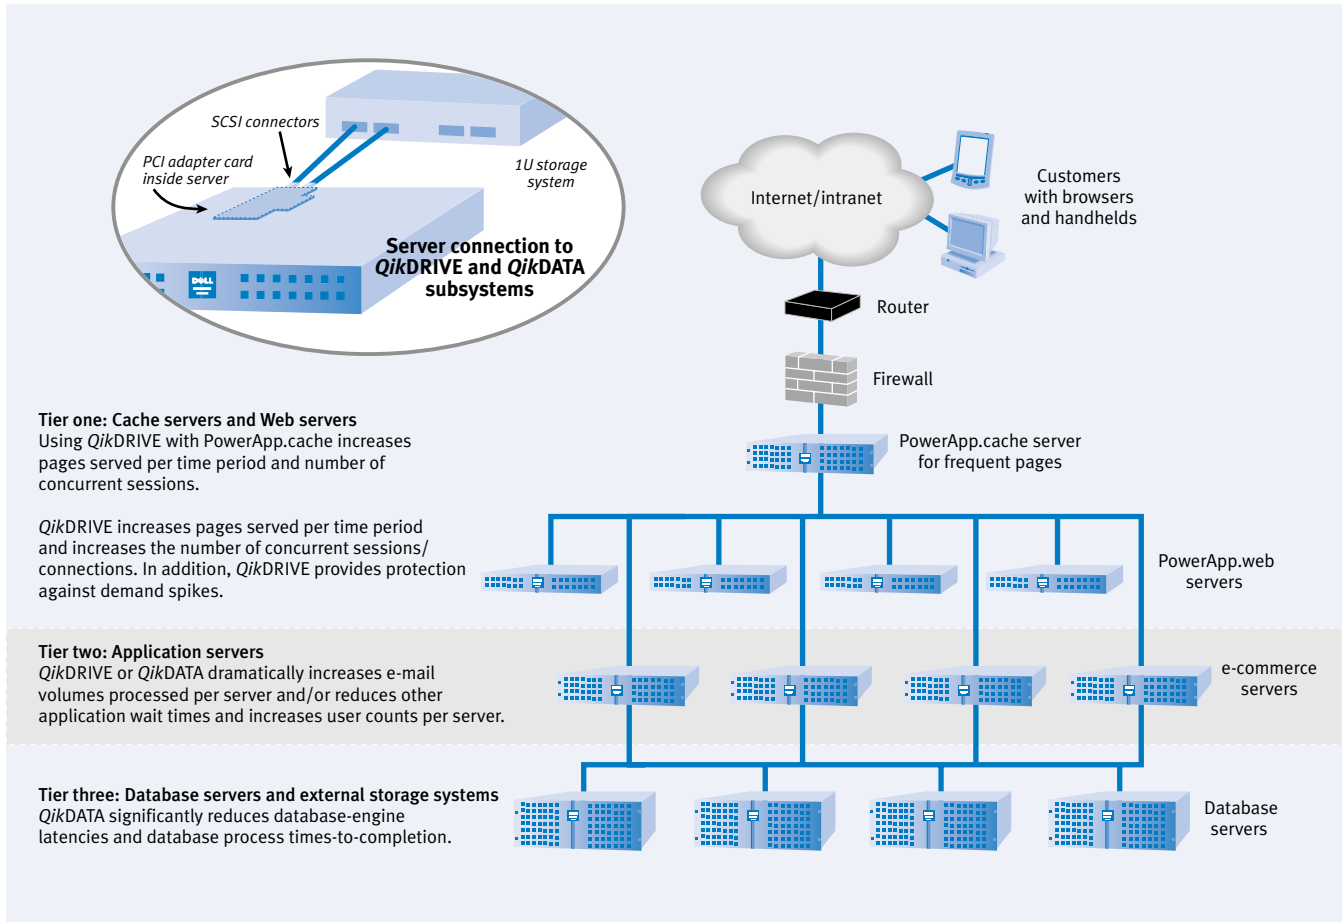


Figure 1. Integrating silicon-based application accelerators into multitiered infrastructures

Performance issues with the Monster.com CRM application

More than 800 Monster.com salespeople and sales managers use the Siebel® Systems CRM software application to manage customer data, track sales opportunities, monitor sales performance, allocate territories, and forecast business.

The SQL database supporting the CRM system is approximately 28 GB. At the close of each month, the CRM system reconciles more than 18,000 sales accounts, generates sales performance reports, calculates commissions, reassigns territories, and updates sales forecasts. Information gathered from this monthly process then forms the foundation for the following month's sales plan and account activity.

After extensive system analysis, Monster.com located significant data I/O bottlenecks in its CRM application. Within these bottlenecked areas, CPU cycles were lost during I/O wait states, resulting in low CPU usage. To remedy this situation, Monster.com implemented two directly attached 16 GB Platypus Technology QikDATA™ application accelerators to off-load the most highly accessed, or "hot," files from existing HDD-based storage. This implementation sought to accomplish several important performance goals: improved end-user experience, reduced end-of-month

batch process time, increased server capacity, and guaranteed first-of-month sales readiness.

Effects of implementing silicon-based accelerators

The QikDATA application accelerators increased system performance with their 350 MB per second and 40,000 I/Os per second capabilities, surpassing the performance improvement goals. See Figure 2.

Several long-term benefits also resulted from this implementation. Faster transaction processing propelled more system uptime, which raised the number of selling hours in a month and ultimately increased sales productivity. Higher performing applications also improved technology-based business strategies.

Monster.com CRM system performance objectives	Impacts of QikDATA silicon-based accelerator
Improve end-user experience	30% reduction in query time
Reduce end-of-month batch process time	50% improvement in database indexing time
Increase server capacity	25% more simultaneous users
Guarantee first-of-month sales readiness	Protection from \$2.3 million/day cost of downtime

Figure 2. Impact of silicon-based accelerator on CRM application

Monster.com business planners could now rethink major service models and capital planning issues, such as how many data centers to deploy and where to place support networks. In addition, end users experienced faster query response time when accessing the Siebel Systems CRM database.

Capacity issues with the Monster.com e-mail application

Every time a Monster.com user participates in a job search, numerous e-mails pass between the job seeker, hiring companies, and Monster.com. The Web site handles three to five million complex e-mail messages each day, which inflicts demanding I/O processing on its storage systems.

Initially comprised of three Dell® PowerEdge® 6450 servers, the Monster.com e-mail system generated a message throughput of 40 messages per second per server—120 messages per second total. Monster.com deployed three additional servers for redundancy.

Due to rapidly growing e-mail volumes, Monster.com needed to double this message throughput without using significant data center floor space or increasing management costs for the e-mail server farm. The performance of the qmail e-mail application was bound by HDD performance, not that of the CPUs, so simply adding more processors to the existing servers would not achieve the goals of Monster.com.

To address the I/O bottleneck, Monster.com implemented silicon-based application accelerators. The most heavily accessed qmail files—the message queues—were easily transferred onto two 8 GB Platypus Technology QikDRIVE™ accelerators.

Effects of implementing silicon-based accelerators

With the QikDRIVE application accelerators in place, the Monster.com e-mail system transferred 1,200 e-mail messages per second per server, far exceeding the performance objectives (see Figure 3). Monster.com could handle 30 times its previous e-mail capacity without increasing floor space requirements or system complexity.

Performance improvements with silicon-based application accelerators

A wide range of businesses and e-businesses can benefit from the additional system capacity that solid-state storage offers. Removing

Monster.com e-mail system architecture objectives	Impacts of QikDRIVE silicon-based accelerator
Raise message throughput from 40 to 80 messages/second	Throughput raised from 40 to 1,200 messages/second
Use no additional floor space	Server footprint reduced from six to two servers
Minimize server farm management costs	66% fewer servers requiring support

Figure 3. Impact of silicon-based accelerator on e-mail application

PLATYPUS SILICON-BASED APPLICATION ACCELERATORS

The Platypus QikDATA application accelerator:

- ▶ Occupies 1U of rack space
- ▶ Supplies up to 128 GB per 64-bit peripheral component interconnect (PCI) interface
- ▶ Addresses redundancy issues with internal uninterruptible power supply (UPS) and recoverable disk system
- ▶ Supports Microsoft® Windows NT® 4.0, Windows® 2000, and Windows XP™; Linux® kernels 2.2.x and 2.4.x (such as Red Hat® Linux 7.0, 7.1, and 7.2); HP-UX™; Compaq® Tru64 4.0x and 5.x; and Sun® Solaris® 2.6, Solaris 7, and Solaris 8 (Certified Solaris Ready)

The Platypus QikDRIVE application accelerator:

- ▶ Supplies up to 8 GB per 32-bit PCI slot card
- ▶ Plugs directly into any Dell PowerEdge 6450 PCI bus
- ▶ Supports Microsoft Windows NT 4.0, Windows 2000, and Windows XP; Linux kernels 2.2.x and 2.4.x (such as Red Hat Linux 7.0, 7.1, and 7.2); FreeBSD; Compaq Tru64 4.0x and 5.x; and Sun Solaris 2.6, Solaris 7, and Solaris 8 (Certified Solaris Ready)

data bottlenecks created by mechanical HDDs allows an existing infrastructure to perform to its full potential.

Silicon-based application accelerators from vendors such as Platypus Technology are designed to bear the I/O load of mission-critical systems. Application accelerators complement previous processor investments and provide additional capacity without increasing floor space requirements, system complexity, processor-based software license fees, and administration costs. ☞

Ken Cameron (*kenc@platypus.net*) is vice president of Technical Services at Platypus Technology, Inc. He works closely with customers by analyzing system performance bottlenecks and implementing solid-state storage solutions. Prior to Platypus, Ken worked with Giantloop Networks and Tessera Enterprise Systems in deploying storage sub-systems, databases, and optical networks.

Toni Sutherland (*tonis@platypus.net*) is vice president of Marketing Communications at Platypus Technology, Inc. Toni has spent almost a decade working with silicon-based storage in U.S., Asia-Pacific, and U.K. markets.

FOR MORE INFORMATION

For more information about Platypus Technology products and services, visit <http://www.platypus.net> or call 1-877-718-0900.