



Answers.com[®]
The world's leading Q&A site

HP Storageworks IO Accelerator is
the Answer for Answers.com

FUSiON-iO



HP Storageworks IO Accelerator is the Answer for Answers.com

Leading Q&A site uses Fusion-io powered devices to recognize full and immediate ROI, improve user experience, and slash IT overhead by 75%.

The Challenge

Answers.com ("Answers") is the leading Q&A web site. It is ranked among the top 20 web sites in the U.S. with over 50 million unique users. Their data centers faced two challenges:

1. Keeping end-users' query response times low. This challenge grew as data volume and site traffic increased.
2. Keeping replication performance high. Users are constantly updating data and slow replication from the Master to a read server could result in responses containing stale (outdated) data.

According to Dan Marriott, Director of Production Operations, these challenges created a constant battle for backend resources to tune queries for performance.

To overcome these challenges, Answers maintained an innovative and scalable MySQL database system that allowed them to drop in servers to horizontally scale performance.

A typical Answers database cluster consisted of a master database server, which handled all writes and distributed the load of incoming user requests, plus around 20 read (or slave) database servers that replicated in real-time from the master.

This design ensured their customers got best-in-class performance, but it saddled them with increased operating costs for power and cooling, rack space, and server administration, and was reaching the limits of the design's performance capabilities. However, at the time the system was implemented, this solution was their only option.



SUMMARY OF BENEFITS

- 30x faster database replication, ensuring responses include most up to date data
- 9x greater query processing throughput
- 30% faster query response times, improving user experience
- Over 8x faster to disaster recovery time
- 75% reduction in server footprint, power costs, and IT overhead
- HP qualified and supported product
- Nearly instant deployment time using the HP BladeSystem architecture

"The value of the cards is tremendous. By reducing the servers I needed from around twenty to five, I estimate my capex per cluster improved by the cost of about ten servers, plus I have significant headroom for future growth. Also important, I reduced my database layer's operating costs by 75%."

Dan Marriott – Director of Production Operations



The Solution


On a data storage specialist's recommendation, Derek Balling, Answers' Data Center Manager, visited the Fusion-io booth at a 2008 conference. After just a few minutes discussion he realized "Fusion-io was perfect for what we did in our databases."

At that time, Fusion-io did not have a product compatible with Answers's blade servers, but as soon as HP began producing the HP Storageworks IO Accelerator, Answers immediately bought two cards to test.

AMPING UP REPLICATION

Derek got his first taste of what the IO Accelerator could do as he was configuring two test servers, one with a card and one without. The wait for the master to update the read server without the card took over six and a half hours. The wait for the server with the card was just twelve and a half minutes.

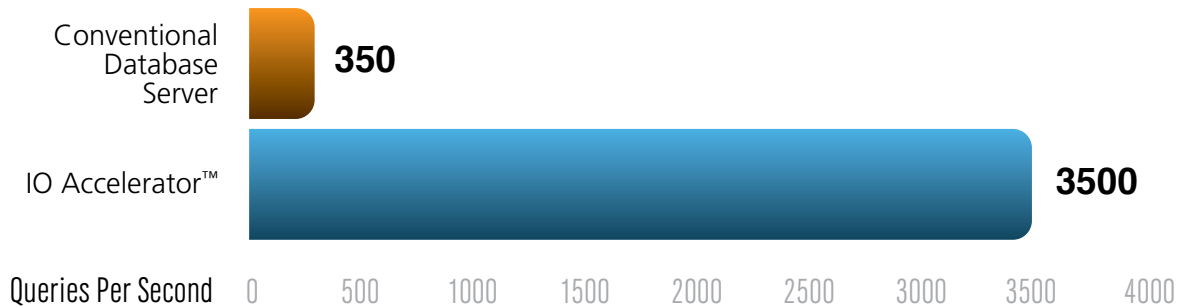
"We ended up doing it a couple times because we thought, 'That's just not possible,'" Derek said. "It was what we call the "Aha!" moment," Dan added.



Replication Update Improvement

EMPOWERING THE DATABASE

To test the systems, they ran complex queries against the test databases. A server without the card reached around 350 queries per second, at which point the database began locking up. Comparatively, a server equipped with the IO Accelerator could run 3,500 queries per second, with no problems whatsoever. Not only could the databases process nearly nine times more queries per second, it processed them 30% faster, a benefit particularly valuable to their end-users.



Query Performance Improvement

SLASHING DISASTER RECOVERY TIME

Answers discovered an additional benefit as they were setting up a redundant data center to its New Jersey site, located in Utah. Installing a 320GB IO Accelerator into one of their key master database servers in Utah reduced the time it took to copy and uncompress the other data center's backup from one hour and forty minutes to under 15 minutes. Combined with the greatly reduced "catch up" times of the read servers, this reduced the worst-case data center recovery time from eight hours to less than sixty minutes.

SCALING FOR SERVER CONSOLIDATION

While the benchmarks showed Answers just how much the cards could improve performance, they learned just how scalable their performance was when they deployed them in production. Dan initially deployed five IO Accelerator-equipped servers in a cluster that previously had 19 slave database servers. This met Answers's comfort level for high availability and growth requirements, and would accommodate server failures and taking servers offline for maintenance.

Derek and Dan soon saw they had more than ample application performance and, as a test, reduced the number of active servers in the cluster to two, “with no appreciable change in user experience. Though we did not try it, I’m fairly confident just one server could have handled the traffic of the site at that time,” Derek said. “It was literally astonishing to see. The cards completely resolved our database replication and query response issues,” Dan stated.

ENHANCING RELIABILITY

The server consolidation improved more than the performance of its database layer, it also enhanced its reliability. With 75% less hardware, its database layer had far fewer mechanical failure points. Additionally, the IO Accelerators added no embedded processors or memory to processing, reducing the risk of soft errors and coding errors. Finally, Fusion-io’s patent-pending self-healing “Flashback Protection,” with N+1 redundancy ensured they would never lose data due to NAND Flash failures.

IMMEDIATE AND CONTINUING ROI

By repurposing 75% of their read servers, per cluster, Answers was able to cut capital expenditures in other areas of its data center, delivering an immediate ROI that surpassed their purchase price. On top of the capital cost returns, the server consolidation reduced operating costs for power, cooling, rack space, and server administration by 75%.

Dan declared, “The value of the cards is tremendous. By reducing the servers I needed from around twenty to five, I estimate my capex per cluster improved by the cost of about ten servers, plus I have significant headroom for future growth. Also important, I reduced my database layer’s operating costs by 75%.”

Above the capital and operating cost returns, Answers gained hours of valuable back-end database team time that they could devote to actively improving applications rather than tuning the database.

EASY MIGRATION

The migration path to the IO Accelerator cards was simple. The cards drop into Answers’s existing blade system, allowing Derek to leverage the benefits of their horizontally scalable architecture.

To update a server, Derek takes it out of rotation, installs the new card, moves the data onto the card, and brings the server back up. The database comes current within minutes of the time it boots. The entire process takes around eight minutes per server.

HP SUPPORTED PRODUCT

Another advantage is that, unlike other SSD products, the IO Accelerator is an HP qualified and supported product, giving Answers the assurance of HP’s stamp of quality and access to their world class support.

Typical Answers Database Read (Slave) Cluster

SYSTEM BEFORE



(20) HP BL460c, 32-48GB RAM

- OS: CentOS 5 - 64bit
- Applications: MySQL 5.0.x
- Disks: (2) 146GB Hard Drives (SAS/10k)

CHANGES TO SYSTEM

Repurposed 15 BL460c servers

Added one 320GB HP IO Accelerator to each server



Summary

Implementing Fusion-io gave Answers the following benefits:

- 30x faster database replication, ensuring responses include most up to date data
- 9x greater query processing throughput
- 30% faster query response times, improving user experience
- Over 8x faster to disaster recovery time
- 75% reduction in server footprint, power costs, and IT overhead
- HP qualified and supported product
- Nearly instant deployment time using the HP BladeSystem architecture

Needless to say both Dan and Derek are thrilled, and are looking at deploying additional cards in a number of Answers' other high I/O database environments.

Dan said, "Prior to purchasing the IO Accelerators, we were adding database servers to this cluster every two months or so. I estimate that two of the five IO Accelerator-equipped servers can handle the same load that required around twenty of the non-equipped servers. The three other servers give us enough headroom that our traffic can double before we need to expand this cluster."

About the Company

Answers Corporation (NASDAQ: ANSW) owns and operates Answers.com, the leading Q&A site, which includes WikiAnswers and ReferenceAnswers. The site supports English, French, Italian, German, Spanish, and Tagalog (Filipino). WikiAnswers is a community-generated social knowledge Q&A platform, leveraging wiki-based technologies. Through the contributions of its large and growing community, answers are improved and updated over time. The award-winning ReferenceAnswers includes content on millions of topics from over 250 licensed dictionaries and encyclopedias from leading publishers, including Houghton Mifflin, Barron's and Encyclopedia Britannica.



FUSION-io

Phone: 801.424.5500 | www.fusionio.com | info@fusionio.com