



Smart analytics
deliver **wisdom**.
Foresee change.

Challenge

After LinkedIn's launch in May 2003, the number of users grew to 8 million in 2006 then exploded to more than 200 million in 2012. This exceptional increase demanded continuous platform development and the rollout of new features.

Solution

LinkedIn deployed HGST PCIe SSDs in a "near pervasive" strategy that delivered a 200% improvement in IOPS, a 3x reduction in latency and significant power savings for 200 million users and growing.



LinkedIn is the world's fast-growing professional networking site that allows members to create business contacts, search for jobs and find potential clients. LinkedIn forms a vital link for more than 200 million professionals in more than 200 countries and territories who count on the LinkedIn network and consider it vital to their day-to-day professional life. They use it to stay connected with business contacts, expand their professional networks, develop closer relationships and build personal brand equity. LinkedIn has a diversified business model with revenues from member subscriptions, advertising sales and talent solutions.

The company has been the most successful social-media stock on the market since its initial public offering in May 2011. LinkedIn surpassed expectations by recording its first quarter with more than \$300 million in revenue in February 2013.

Benefits

- Moving to HGST PCIe cards resulted in positive benefits for LinkedIn.
- Overall performance has improved, application latency has decreased and servers are better utilized.
- Power and cooling requirements are significantly lower. Flash brings persistence of data, handles peaks and valleys more effectively, consumes one-tenth the power of conventional storage, has no moving parts, dissipates less heat and reduces data center power and cooling costs.
- Total cost of ownership has decreased.
- User experience with mail, search and other data-intensive applications has improved across internal and end user surveys.



HGST PCIe SSD

The Challenge

LinkedIn infrastructure started off with RISC UNIX systems, traditional storage-area networks and expensive databases. After experiencing unprecedented growth in just five years, LinkedIn felt it was becoming increasingly difficult to manage and grow its IT infrastructure to meet the needs of the user base.

As with other Web 2.0 industries, the workload of LinkedIn's infrastructure was very unpredictable because even a small increase in the number of users caused a significant increase in stored data, such as blogs, videos and photos. All of this data had to be processed in a very short time. In addition, user experience and application availability were of prime importance. The company's services had to be available 24 x 7 x 365 and these benefits had to be achieved at the lowest possible cost.

It was clear to LinkedIn that its existing data center infrastructure could not scale to meet its performance needs without an architecture overhaul. The company needed to move to an open-systems architecture, allowing it the flexibility to scale, to deploy applications based on business needs and to deliver a quick response for millions of concurrent users. In short, the company needed to build an agile and flexible infrastructure to manage the large and continually increasing user base.

The Solution

LinkedIn crunches 120 billion relationships per day and blends large-scale data computation with high volume. Meeting these needs required a storage architecture that could deliver high IOPS, low latency and the ability to maximize the compute capabilities of the server by pairing it with superior storage. LinkedIn considered moving to the new Flash technology and chose HGST PCIe cards for a variety of reasons, including its ability to deliver higher performance through higher server utilization. HGST's PCIe SSDs were also operationally efficient, providing high-performance storage within the server and close to the CPU, with no moving parts. Plus, the SSD delivered the IOPS and low latency required to sustain business applications. LinkedIn sought to use the most reliable Flash technology available at that time. The initial deployment comprised Flash SLC (single-level cell) chips that offered significant write performance for critical applications.

End-to-End Solution Support

Beginning in 2010, the HGST team worked with LinkedIn to help the company understand the performance, latency and TCO benefits of Flash; explain the use cases; and go over the deployment details. The team provided technical guidance about system bottlenecks within the infrastructure to achieve best performance. The deep technical partnership in early engagement, plus the uncompromising performance, enhanced reliability and data availability of Flash-aware RAID, were key to LinkedIn's selection of HGST as its choice of partner. LinkedIn began the relationship with the purchase of ten cards to test the PCIe SSDs' capabilities in their particular environment.

After one year, LinkedIn was comfortable that HGST's Flash technology was capable of enterprise-grade deployments. Since LinkedIn's IPO in 2011, the number of users had climbed to more than 130 million. LinkedIn evaluated the newer MLC (multi-level cell) Flash technology available and felt that it met their growing needs. MLC offers a good combination of performance, endurance and reliability at a relatively lower cost.

“When we implemented the newer MLC Flash technology, our goal was to take advantage of the performance, latency and TCO benefits over our existing infrastructure. We were immediately impressed with HGST’s customer support as it ramped up our team on the new technology. Once fully implemented, we experienced the high IOPS, low latency and the ability to maximize the compute capabilities of the server, proving we made the right decision.”

Sonu Nayyar
Senior Director of Production Operations at LinkedIn.

200%

IMPROVEMENT IN IOPS

3x

REDUCTION IN LATENCY

200M

USERS AND GROWING

Steps to Widespread Deployment

LinkedIn first deployed HGST PCIe cards in its back-office applications. Once convinced of the product’s performance and results, LinkedIn decided to implement HGST Flash in one of its business-critical applications running on the Voldemort distributed database system. Now HGST storage was on the front lines, supporting applications that had direct impact on online user experience. For example, HGST Flash PCIe cards enable features such as “Who’s Viewed My Profile,” which produces high write loads. Other applications at LinkedIn were similarly challenging on the scaling front, such as the “Find Similar Profiles” feature. While the set of all user profiles is very large, even a modest subset of all user-profile pairs is huge. Other applications that needed to handle hundreds of millions of reads and writes per day were moved to the HGST storage solution. Simultaneously, LinkedIn continued to increase the amount of data it stored.

Project Voldemort was a key growth area for the company. HGST PCIe SSDs were used to derive the highest read/writes from the databases and to achieve server consolidation by providing more than one TB of Flash per server. Hundreds of HGST PCIe cards were deployed across all Voldemort servers, a process that took more than six months. Continued deeper technical engagement and knowledge-sharing was critical to the success of this project.

With the success of Voldemort Flash-based storage, LinkedIn launched a major initiative for rebuilding infrastructure. Today, all of LinkedIn’s data engineering efforts are focused on building services that can work together easily. Flash is being deployed across servers in a “near pervasive” strategy.

As part of this initiative, one of the most important things that LinkedIn is building is a new in-house database system originally designed to provide a usability boost for LinkedIn’s InMail messaging service. These moves are all part of a mission to create an innovative data environment at LinkedIn, with thousands of HGST PCIe cards now deployed in pre-production and production environments.

“Flash is transforming the data center and HGST is leading the Flash platform transformation,” said Ulrich Hansen, Vice President of SSD Product Marketing at HGST. “LinkedIn has embraced the transformation and is a testament to the fact that the future data center will benefit from a server-side platform. LinkedIn and the Web 2.0 community are only the beginning and as IT becomes more familiar and comfortable with the technology we’ll begin to see enterprises mimicking innovators such as LinkedIn and deployed server-side Flash platforms.”

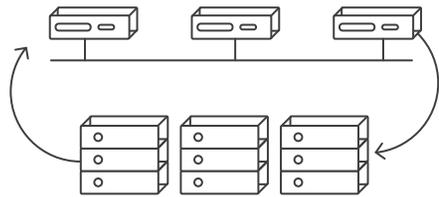
“The success of Flash deployment at LinkedIn reflects a partnership built on collaboration, technology and value milestones, that evolved from a product centric, technology engagement to a trusted business partner who provided a platform and architecture that allowed LinkedIn to scale its infrastructure to support its hyperscale growth,” said Jaggi Subramaniam, HGST Sales Director. “This was a mutually beneficial learning experience in building a partnership with an early innovator experiencing hyperscale business growth that ultimately delivered a successful template to the industry for leveraging Flash as platform to deliver technology and business value.”

The increasing number of users, the company’s stock market success, and its growing revenue are all indicators of a successful Web 2.0 company. LinkedIn’s forward momentum is not expected to slow anytime soon. To continue its growth and success LinkedIn depends on the performance and reliability of HGST PCIe SSDs.

Transformation of the LinkedIn Infrastructure Brings Business Value and Savings

BEFORE

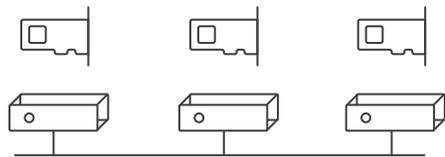
SC Servers on SAN



Monolithic SAN Arrays

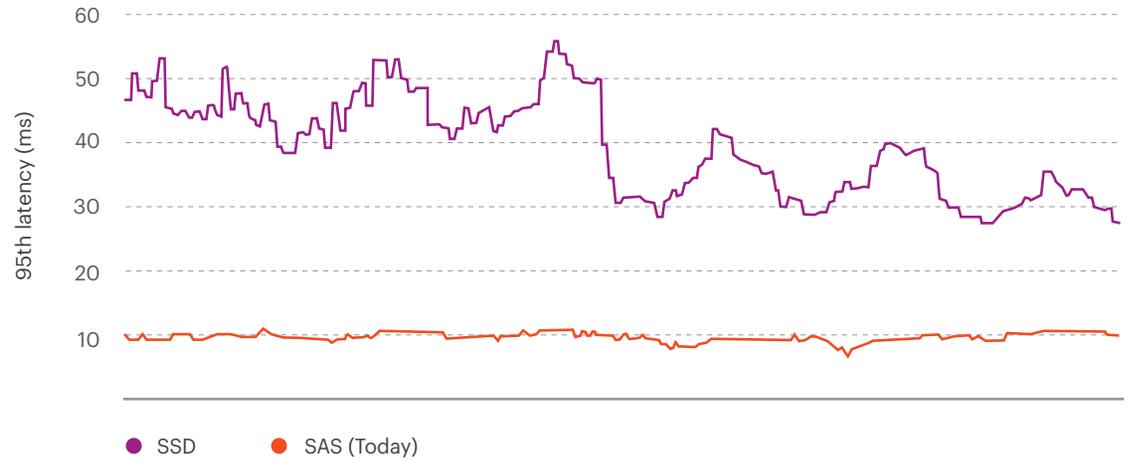
AFTER

Commodity Servers with HGST PCIe Cards



Near Pervasive Flash

Latency Comparison Test Confirms SSD Superiority Over SAS



For more information on HGST PCIe SSDs, visit HGST at <http://www.hgst.com/products/solid-state-drives>