

slide[®]

Building Scalable Storage for a Web 2.0 Site
with 144 Million Unique Visitors per Month
Using DataDirect Networks S2A NAS Solution

www.datadirectnet.com

Challenge

Driven by rapid user growth and burdened by the complexities and cost of managing dozens of white-box storage servers, Slide needed a storage solution that would allow them to scale performance and capacity and simplify their environment efficiently and cost effectively.

Application

Slide enables users to personalize and share photos and videos within the most popular Web 2.0 communities. With a vast amount of user content and over one million widgets created daily, Slide has to consistently deliver capacity and performance to its 45 million user installed base and the 134 million monthly viewers of these personalized images and videos worldwide.

Solution

Slide chose DataDirect Network's S2A Network Attached Storage (NAS) Solution. Easy to deploy, this enterprise-class, turnkey solution is fully assembled and delivered with storage capacity, clustered NAS servers, networking, file system and management software and advanced data protection features. The S2A NAS Solution utilizes cost-effective SATA drives and supports up to 960TB in just two standard data center racks.



Building Scalable Storage for a Web 2.0 Site with 144 Million Unique Visitors per Month Using DataDirect Networks S2A NAS Solution

Slide's Background

Slide is the largest personal media company in the world, reaching more than 144 million unique global viewers each month and 30 percent of the U.S. and 15% of the global Internet audience. Slide helps people express themselves and tell emotionally engaging stories through personalized media created on Slide.com which can be viewed anywhere on the web or desktop.

Slide utilizes Flash technology-based widgets to provide personalization tools, including: Slideshows, Guestbooks, SkinFlix and FunPix. Posting images and videos that contain graphics and/or animation or wrapping them in personalized players and frames are a popular trend on top social networking and blog platforms, including MySpace, Facebook, Bebo, Hi5, Friendster, Tagged, Piczo and Blogger. Slide is also the leading application developer on Facebook with more than 50 million applications installed. Slide's "Top Friends" has more than 3 million active daily users, making it the most active application by more than four times that of any other 3rd party developer.

“White-box Storage Quickly Became Too Difficult to Scale and Complex to Manage”

— Ken Brownfield,
Director of IT, Slide

Like many early stage internet companies, Slide found that scaling capacity with Linux-based white-box storage can be an effective and economical approach up to a point. However, having reached critical mass, Slide's 144 million unique monthly viewers are driving capacity growth at a rate exceeding 3TB/week and accelerating.

In the beginning, Slide's business model was simple: hosting and online photo sharing. Storage allowed users to upload and serve out images to invited friends and family members. Slide began with all their user images on a single Linux white-box server, which reached its maximum capacity in just one month. As their user community and data stores grew, they would add another server with internal SATA drives. Each time capacity would reach the server's threshold they would add another.

As their business evolved, Slide found a unique niche and opportunity to add value by providing personalization tools and extending that capability to reside on online communities, blogs and sites. These enhanced features differentiated their offering and allowed them to address a wider audience that included the general public that happened to be on a blog or online community page. As their user base continued to grow rapidly, additional users and tools not only increased capacity, but challenged performance - so they began to add servers in pairs in order to address bandwidth demands. They also sought to maintain reliability for their customers in case one of the white-box servers went down or a SATA drive inside them failed.

The Challenge

Their environment quickly mushroomed from 8, 12, and 15 . . . all the way up to 20 servers. Adding pairs of Linux boxes, each time additional capacity or performance was required, was no longer a sustainable strategy. Driven by explosive user traffic and data growth, IT quickly found itself faced with managing a farm of white-box servers to facilitate storage - which became a complex and costly proposition. In the San Francisco Bay Area, power to data centers is limited and floor space is quite expensive. Slide needed to centralize, ease management, improve reliability and performance and create more headroom to sustain growth.

Slide's Storage Requirements and Considerations

Ken Brownfield, Slide's Director of IT, sought to centralize and simplify his storage environment, but was also looking for a solution platform that would meet Slide's aggressive data growth and bandwidth requirements. Slide's data characteristics included both small files for images and larger files for movies.

Ken explained, "The top three factors when searching for a NAS Solution were:

1. Robust System Management Capabilities - we wanted to be able to partition without a lot of additional work.
2. Density - we knew that capacity would continue to grow and we wanted to conserve floor space, power and cooling.
3. High Performance with Small Files - it's relatively easy to stream large files, but managing

slide®



FAST FACTS

Launched
August 2005

Headquarters
San Francisco, California

Founder
Max Levchin, co-founder of
PayPal

- Reach**
- 200+ countries
 - 144 million unique viewers/month (December 2007)
 - 1 million new Flash widgets created every day
 - 63+ million apps installed on Facebook
 - 7.5+ million daily active users on Facebook

Major Networks
Facebook, MySpace, Hi5, Orkut,
Bebo, Friendster & more

Case Study

small files with performance but without a lot of wasted space is difficult."

Ken selected DataDirect Networks because of the S2A NAS Solution's ability to excel on all three fronts.

“We evaluated solutions from the usual NAS suspects, but only DataDirect Networks was able to resolve my challenges.”

DataDirect Networks S2A NAS Solution

The S2A Network Attached Storage Solution is a scalable, high-performance data service designed to provide highly available, standards-based access across the enterprise. Unlike traditional NAS approaches that limit flexibility and namespace scalability, the S2A NAS Solution scales performance through the use of clustered NAS servers and S2A Appliance-based storage systems. Easy to deploy, use and manage, the open platform S2A NAS Solution is routinely deployed for a variety of online, near-line and backup applications.

The S2A NAS Solution scales NFS, CIFS and HTTP file services across as many as 16 S2A clustered NAS servers - all presenting a common data set. S2A NAS delivers high performance (up to 8X greater than traditional filers) and maintains quality of service without any performance penalties even under failure conditions.

Slide's configuration consists of: one S2A9550 Appliance Couplet, four integrated clustered NAS Servers (or heads), and ten 48-bay drive enclosures, populated with 500GB SATA drives, for a total capacity of 240TB. As Slide's growth continues, Ken can seamlessly expand with up to 12 additional NAS heads and 480 more drives within his existing system.

In contrast to white-box storage, Slide is able to realize administration, floor space, power/cooling savings that the S2A's dense, centralized NAS Solution provides. No longer are they required to manage an entire server farm and all of the associated networking equipment, software licenses and maintenance just to facilitate storage. This affords a significant reduction in TCO.

“Now I have less downtime from Linux boxes crashing, less latency when serving files and our end-users see much less interruption.”


In addition, Slide's quality of service and data integrity has been greatly enhanced. Up to 8X faster than competitive systems, this solution is capable of reading and writing in excess of 2GB/s, even when the system is fully populated. Unlike other systems where RAID 6 means slow performance, the S2A's DirectRAID™ 6 provides blazing reads and writes and valuable double parity protection, which is a best practice for SATA implementations.

Slide also benefits from the S2A's high Quality of Service. While the occasional drive failure is inevitable in any storage system, with the S2A NAS there is no performance impact when a drive fails or during its rebuild. This gives Slide very predictable performance without concern of degradation during peak traffic periods.

Furthermore, the S2A's SATAssure™ technology validates data integrity on every read to ensure that there is no data corruption, without penalizing performance or consuming storage capacity. If corruption is detected, data is corrected on-the-fly and delivered to the host without any delay, ensuring that Slide's users always get the data they expect and helping to avoid application crashes.

Slide continues to grow in the marketplace and Ken anticipates his 3TB per week growth rate to increase to 4TB per week by mid 2008. Ken is so pleased with the S2A NAS Solution's performance and reliability, he recently added another 240TB.

S2A NAS Solution



DataDirect Networks is the leading provider of open, scalable storage systems for performance and capacity driven applications. DataDirect Networks' S2A (Silicon Storage Architecture) appliance enables modern applications such as video streaming, content delivery, modeling and simulation, backup and archiving, cluster and supercomputing, and real-time collaborative workflows, that are driving the explosive demand for storage performance and capacity. DataDirect Networks' S2A technology and solutions solve today's most challenging storage requirements, including providing shared, high-speed access to a common pool of data, minimizing data center footprints and storage costs for massive archives, reducing simulation computational times, and capturing and serving massive amounts of digital content.

Major corporations, supercomputing centers and rich media organizations, including AOL, Ascent Media, Boeing, CINECA, CGGVeritas, CNN, Disney, Federal Reserve Board, Ford, Hess, Kodak Gallery, Lawrence Livermore National Laboratories, Mediaset, MTV, NASA Ames, RIOT, Sandia National Laboratories, Sony, Technical University Dresden, Technicolor, Time Warner, Thomson, Trinity College Dublin and Universal, utilize DataDirect high performance, high capacity solutions.

DataDirect
NETWORKS
Performance. Capacity. Innovation.

9351 Deering Avenue . Chatsworth . California 91311
phone +1.800.TERABYTE (837.2298) . fax +1.818.700.7601
sales@datadirectnet.com
www.datadirectnet.com