



SUCCESS STORIES

NVIDIA MAXIMIZES ENGINEERING PRODUCTIVITY WITH A GRID ARCHITECTURE

“Two and a half years ago we had 20 NetApp systems and one storage administrator. Today, we have 42 systems with 440TB of capacity, and everything is still managed by a single person.”

KELLY ALEXANDER Manager of Engineering Services, NVIDIA



NVIDIA

KEY HIGHLIGHTS

Location Santa Clara, California

Industry Semiconductor manufacturing

Key Business Challenges

- 100% uptime required
- Extreme processing/storage needs
- DAS limits available horsepower
- Continuously evolving demands

Key Solution Components

- Over 40 NetApp FAS and NearStore® systems
- NetApp SnapMirror® software
- NetApp DataFabric® Manager

Key Business Benefits

- 24x7 availability
- Improved designer productivity
- IT performance eliminates delays
- Single admin manages 440TB

THE CUSTOMER

A global leader in advanced graphics processing technology, NVIDIA Corporation (www.nvidia.com) has received more graphics awards from the PC industry than any other company. Companies worldwide choose NVIDIA graphics processing units to enhance the digital media experience on desktops, workstations, notebooks, handhelds, and other devices. Headquartered in Santa Clara, California, NVIDIA has over 2,000 employees worldwide with revenue approaching \$2 billion annually.

THE CHALLENGE:

Provide 100% Availability and Cost-Effective Scalability for a Busy EDA Environment

The primary challenge faced by NVIDIA's Engineering Services group is to ensure the productivity of NVIDIA's more than 800 engineers. Maintaining a technological edge is critical to continued success in a very competitive market. An extremely active chip design team uses the latest engineering design automation (EDA) tools to create NVIDIA's state-of-the-art graphics chips. These efforts produce an extreme demand for both processing and storage. According to Ed Yee, vice president of Technology, “Several years ago we found that using local storage on our big Sun™ enterprise servers was taking away horsepower needed to process massive engineering jobs. We fixed that by deploying NetApp storage to offload NFS file operations and allow the servers to use their CPU cycles for processing jobs.”

Yet as demands for processing capability continue to rise, NVIDIA must continuously evolve its infrastructure.

THE SOLUTION:

A Scalable Grid Architecture

NVIDIA upgraded to a grid infrastructure to meet growing demands for storage. This architecture consists of 3,000 Linux compute nodes for the majority of its EDA compute jobs, 13 Sun 6800s for large memory compute jobs, and a total of 42 NetApp primary and secondary storage systems with about 440TB of storage capacity. Engineers submit jobs for processing on the NVIDIA grid using LSF (Load Sharing Facility) software from Platform Computing. With LSF, jobs are transparently allocated the resources they need based on predefined rules without engineer or administrator intervention.

NVIDIA utilizes a multitiered storage architecture. The first tier consists of high-end NetApp fabric-attached storage (FAS) systems configured with high-performance, low-capacity drives to ensure a high spindle count. These systems support 1,000 to 1,500 simultaneous clients and provide the large volumes of scratch space that EDA applications require.

Midrange NetApp primary storage systems configured with high-capacity drives store project data, output data files, and enable group collaboration. Data that doesn't require high performance and fast access—including previous-generation chip-design files—is stored on economical NetApp NearStore near-line storage systems.

A separate high-performance FAS cluster is dedicated to storing the NVIDIA crown jewels: engineering data—including all final chip designs and software source trees—and other critical information managed by Perforce SCM software.

To ensure that no data is lost in the event of an unexpected site failure, NVIDIA uses NetApp SnapMirror software to mirror all critical data to a disaster recovery site located about 100 miles away in Sacramento, California. SnapMirror also makes it simple to replicate engineering data between sites for use by engineers in remote locations.

BUSINESS BENEFITS:

High Availability and On-the-Fly Scalability Keep Engineers Productive and Products on Schedule

By delivering 99.99+% uptime, rapid scalability, and streamlined management, the Network Appliance infrastructure is helping NVIDIA succeed in a highly competitive market. “The NetApp systems help us get product to market faster by ensuring that

jobs are processed reliably and efficiently,” explains Kelly Alexander, manager of Engineering Services. “If our design and manufacturing schedule slips, we can miss being first to market with a new product or miss getting a product out for a key selling season like Christmas. It’s critical that our IT infrastructure performs for our engineers.”

“In all the years we’ve relied on NetApp, NVIDIA has never lost data stored on NetApp systems,” points out NVIDIA storage administrator Rod Hernandez. “NetApp provides exceptional availability and requires almost no administration,” adds Alexander. “NetApp technology does what it is designed to do without content maintenance. We know our storage systems will be up 24x7, so we don’t have to worry about them.”

In recent years NVIDIA has more than doubled the number of NetApp systems supporting its engineering infrastructure without the need to add administrative headcount. Hernandez provides the perspective of someone who works with

NetApp hardware and software every day. “NetApp has made my life a lot easier and eliminated a lot of sleepless nights. If hardware maintenance is needed, for example, the systems are designed so that parts can be easily hot swapped without special tools in a matter of minutes.”

In addition, NetApp Snapshot™ technology, which enables administrators to perform online backups using the Snapshot feature with minimal disk space, is a significant time saver. “We create Snapshot copies every four hours and save them for up to two days,” Alexander comments. “Snapshot directories are visible to the engineers, so when they need to retrieve an earlier version of a file, they can do it themselves. This reduces the number of requests we have to address while helping keep our designers productive.

“A great differentiator of NetApp technology is its consistency,” Alexander adds. “Once someone learns how to operate a NetApp system, that person can manage any of our primary storage or NearStore systems because they all use the same interface, and that interface doesn’t change every time there’s a software revision.” Hernandez agrees, “Since I was already familiar with NetApp technology, managing a NetApp Fibre Channel SAN environment required only a few hours of training.”

“Having NetApp on the back end has enabled us to take storage out of the equation,” concludes Alexander. “Our grid architecture allows NVIDIA engineers to focus on chip design without worrying about compute resources or storage. That makes them more productive and ultimately improves our time-to-market.”

ABOUT NETWORK APPLIANCE

Network Appliance is a world leader in unified storage solutions for today’s data-intensive enterprise. Since its inception in 1992, Network Appliance has delivered technology, product, and partner firsts that simplify data management. Information about Network Appliance™ solutions and services is available at www.netapp.com.

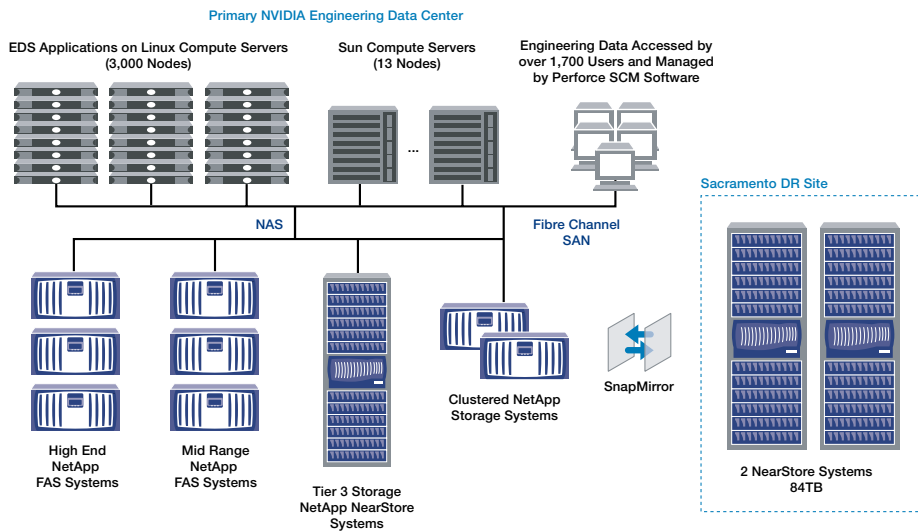


Figure 1) NVIDIA Global Engineering Infrastructure.

NVIDIA utilizes a multitiered storage architecture. High-end NetApp FAS systems support compute jobs, while mid-range systems store project data, output data files, and enable group collaboration. Data that doesn’t require high performance and fast access is stored on NetApp NearStore systems, while a separate FAS cluster is dedicated to supporting a Perforce SCM environment. Critical data is mirrored to a remote site.

