CASE STUDY

"The key factor was how the FalconStor software performed, from various perspectives. For example, were the client IPStor-generated disks able to import without needing to reboot? What about its disk mirroring capabilities? How does the failover work? Following detailed testing, it was clear that FalconStor achieves the better results. At long last we can look at our storage pool, determine precisely who should be given what capacity, and allocate everything at the click of a mouse."

– Dr. Günther Kroβ, Dept. of Computer & Media Services, Humboldt University

Humboldt University, Germany

Using IPStor to allocate, manage, and back up 40TB of data across two sites and get the most out of a pre-existing SAN

About Humboldt University

Founded in Berlin in 1810, Humboldt University (HU) revels in its reputation as the "mother of all modern universities," which stems from the vision of its founder, Wilhelm von Humboldt. Von Humboldt's concept was to provide students with a well-rounded humanist education as well as standard teaching and research, a concept since adopted around the world.

Following the reunification of Germany in 1990, Humboldt University went through an extraordinary process of reorganization and now comprises eleven faculties and two central institutes, including Europe's largest department of medicine.

Business Need

Since 1998, the Mathematics and Science departments of Humboldt University have gradually been relocating from their existing cramped location to new academic buildings on the Science, Business, and Media site at Berlin Adlershof. So far, the institutes of Mathematics, IT, Chemistry, and Physics have moved, as have the Central Science Library and the Central Computer and Media Service.

Early on in the process of relocating, HU decided to redesign its storage infrastructure, partly because of Y2K considerations and partly because it recognized the trend towards centralized management of storage resources— the opposite of the then-current policy of local storage. Additionally, the move meant that storage management would need to span two sites some 30km apart.

With these considerations in mind, HU developed a SAN. Planning began in 1999 and the base installation, comprising two IBM file servers and RSS storage systems, was put into use a year later. The SAN installation at each of the Berlin Mitte and Adlershof sites is based on two 64-port Brocade SilkWorm 12000 core fabric switches. A total of eight switches are deployed between the institutes and locations.

Humboldt University's SAN uses IDE hard drives in the SAN pool. While these don't offer the same performance as SCSI disks, they are only about a quarter of the cost. A number of arrays are each fitted with eight 120 GB disks in RAID5 configuration, and provide an external SCSI interface. A Chaparral SAN-Head controller links the arrays with the SAN via a 2Gb/sec Fibre Channel interface. This compensates for the performance deficit of the IDE drives without compromising availability or security, while at the same time generating savings that enable the procurement of additional disk space.

The storage pool operated by the Computer and Media Service currently provides the institute with more than 40TB of storage, with 16TB located at Berlin Mitte and 24 TB in Adlershof.

FalconStor Solution

By 2003, IT managers at the university discovered what many other organizations also often discover after implementing a SAN: that many of the storage resources in the SAN remained unused. It was against this backdrop that the university's IT managers began to consider virtualization as a means of improving resource usage. Once they had decided that virtualization was the correct approach, the next step was to find an

Industry

Education

Applications

File and print serving, databases, email

IT Challenges

- 40TB of data currently and storage requirement *tripling* annually, with routine additional allocation of storage required
- Labor-intensive management of a complex SAN infrastructure across two sites
- Existing environments required significant administrative attention
- Continuity, integrity and high availability of data required at both locations
- Lack of solid disaster recovery plan
- Unreliable backup and restore

FalconStor Solution

Humboldt University had already installed a SAN infrastructure, but found many of the storage resources remain unused.

After detailed testing and comparison with other potential solutions, the University chose FalconStor IPStor software to provide virtualization and associated storage management functions.

IPStor Software Deployed

4 IPStor appliances
IPStor SAN clients for FC
IPStor SAN clients for IP
IPStor Active-Active Failover
IPStor Synchronous Mirroring
IPStor Local Replication
IPStor TimeMark®
IPStor DynaPath®

Benefits

- On-demand allocation of storage from the virtualized storage pool at the click of a mouse at a central console
- Lower total cost of ownership through:
  - Lower storage administration overhead
  - Reduced cost of acquiring advanced SAN technologies
- Greater return on investment through:
  - Elimination of downtime
  - Reliable data availability
  - Seamless scalability

appropriate solution that would provide both virtualization *and* the additional services that render it so useful.

IT managers at the university were supported in this process by systems firm Controlware, with which the university had already worked on other projects. Several potential solutions, including FalconStor IPStor software, were tested. The key factor was how the software performed from various angles: for example, were the client IPStor-generated disks able to import without needing to reboot? What about disk mirroring capabilities? How did the failover work? These latter two points were the subject of particularly rigorous testing, since the winning solution would need to manage a complex SAN infrastructure over multiple sites.

Consolidated Storage Capacity and Centralized Storage Management

Following all of the detailed testing, it was clear that FalconStor's IPStor solution achieved the best results. It enabled the university's IT managers to pool the storage capacity from all of their disparate disks, centrally view the storage pool, determine precisely who should be given what capacity, and allocate everything at the click of a mouse. This was previously impossible to achieve across the 700+ disks in the SAN.

The first virtualized storage volumes were allocated to central database and mail servers, as well as to a catalogue server in the university's library. Dynamic increases in disk size using IPStor have proceeded smoothly.

High Availability of Data

Maximizing the University's productivity, additional IPStor features have provided HU with advanced data availability benefits, preventing downtime due to hardware or path malfunction, data corruption, user error, viruses, etc.

The ***IPStor Active-Active Failover*** option allows one of the university's IPStor appliances in a pair to immediately and transparently take over the processing tasks of a second IPStor appliance in the pair should the first IPStor appliance fail for any reason or the network path to it be disrupted. Once the primary IPStor appliance is back online, the primary appliance takes over its workload again automatically and transparently, via a process called *failback*. There is no disruption to the university's users, and neither processing nor activities-in-progress are interrupted. This is part of IPStor's end-to-end redundant architecture, which protects against single points of failure such as network connectivity failure, storage connectivity failure, storage device failure, or failures of the appliance running the IPStor software.

The ***IPStor Synchronous Mirroring Option*** provides the university's IT administrators with the capability to define a synchronous mirror for any disk managed by IPStor. The process of creating the mirror does not cause any downtime to active applications. A mirror provides an active, redundant data set to protect against the consequences of device/cabinet/frame level failure, and is also a useful tool enabling the migration of data from older disks to newer disks without any downtime. A mirror can be defined on space from any available disk in the pool, even across cabinets; and this disk is not required to be identical to the primary disk in terms of vendor-brand, type, or even interface.

The ***IPStor TimeMark Option*** has enabled administrators to journal data changes via continuous/scheduled or on-demand point-in-time delta-snapshot copies of data volumes. This provides them with the ability to instantly recover a single file or an entire volume back to a known-good point in time in the event of an accidental deletion, file corruption, or virus attack. The TimeMark option also integrates with the IPStor application-aware Snapshot Agents resident on the application servers to enable online, incremental backups of databases and message stores with transactional integrity and point-in-time consistency.

The ***IPStor DynaPath Agent*** is an application-server-resident software utility that maximizes data availability for IPStor storage systems and provides peak performance across the SAN. The DynaPath Agent performs Fibre Channel and iSCSI HBA load-balancing, transparent failover, and fail-back services to maximize availability and performance of application servers. It takes full advantage of the bandwidth and multiple-server connectivity potential of Fibre Channel and/or iSCSI protocols to optimize resiliency, availability, and performance capabilities at the storage-system level. In the case of the Humboldt University, the DynaPath Agent is used to ensure constant data availability across the SAN by creating parallel active storage paths that transparently reroute application server traffic to a redundant storage path without interruption in the event of a storage network problem. Load balancing enhances peak performance of the SAN by automatically distributing server traffic among the server's multiple storage paths for higher throughput and to eliminate bottlenecks.

Deployment Details

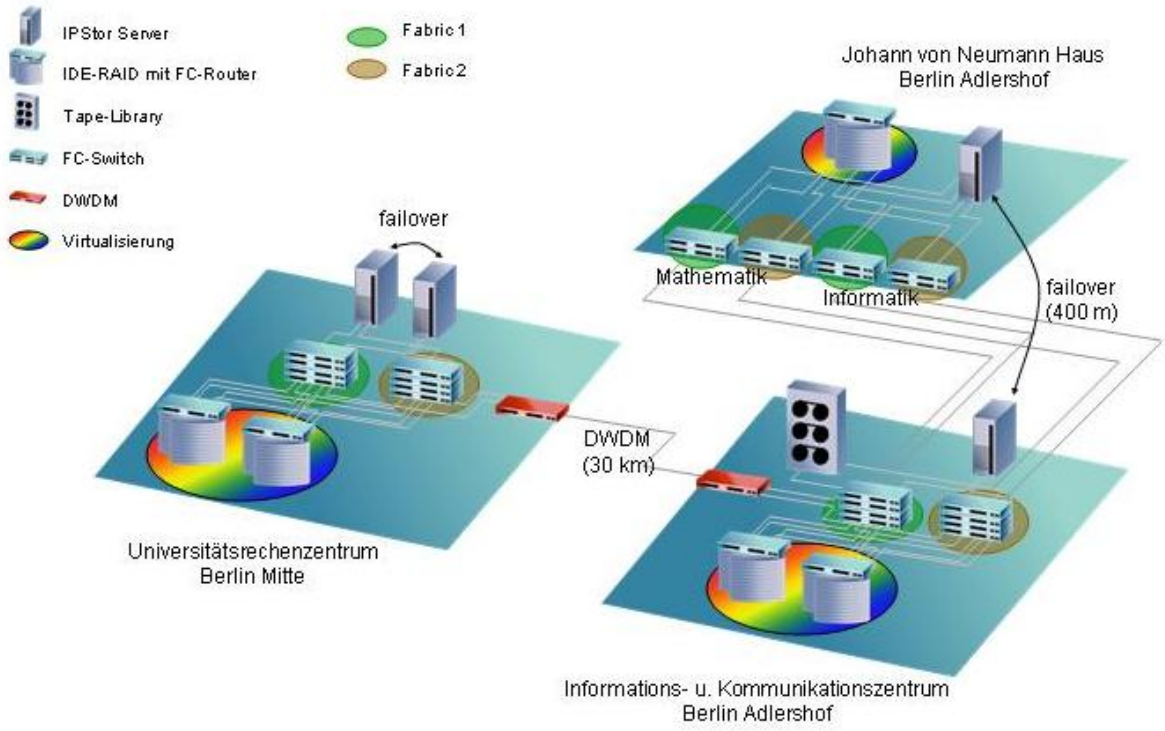

Hardware Configuration	
Primary location 2 x IPStor Appliances Dell 4600 servers RedHat 7.3 1 GB RAM 7 PCI slots Broadcom NICs QLogic 2302 HBAs Switches Brocade Silkworm 12000 and 3800	30 Application Servers IBM (AIX O/S) IBM Emulex HBAs Applications: file services
Secondary location 1 x IPStor Appliance Dell 4600 server RedHat 7.3 1 GB RAM 7 PCI slots Broadcom NICs QLogic 2302 HBAs Switches Brocade Silkworm 12000 and 3800	10 Application Servers Sun (Solaris O/S) JNI HBAs Applications: file services
Third location 1 x IPStor Appliance Dell 4600 server RedHat 7.3 1 GB RAM 7 PCI slots Broadcom NICs QLogic 2302 HBAs Switches Brocade Silkworm 12000 and 3800	10 Application Servers MS (Windows 2000) and RedHat Linux Emulex HBAs Applications: file services

Summary

Humboldt University had found itself in the company of organizations who discover that a SAN is not, in and of itself, the panacea for all storage problems. After in-depth research, IT administrators recognized that to truly derive full benefit from their SAN installation, it was imperative to see the entire gamut of their SAN resources as a single pool of virtualized storage and, moreover, to be able to layer advanced storage services on top of this virtualized pool in order manage and protect storage assets most effectively.

The FalconStor IPStor solution has enabled the university's IT administrators to centrally manage their annually tripling storage resources (currently 40TB of data spread across 700 disks ranging between 36 and 120GB each) at the easy-to-use, go-anywhere IPStor Console. This single point of management significantly reduces total operational costs by allowing fewer IT administrators to manage far greater amounts of storage in far less time, freeing up personnel to attend to other pressing IT tasks. The more efficient utilization of available disk capacity also means that the university is spending dramatically less on raw storage capacity.

The ongoing, gradual migration of the university's data between sites could have been a recipe for data access problems, downtime, or even loss of data. However, despite all of this migration, IPStor has ensured that data is continuously available and the threat of downtime is eliminated, all of which maximizes the return on the university's investment in storage and in IPStor.

FalconStor Software
2 Huntington Quadrangle
Melville, NY 11747
631.773.5859
www.falconstor.com