

# The Evolution of Storage Service Providers: Techniques and Challenges to Outsourcing Storage

Ragib Hasan  
rhasan@ncsa.uiuc.edu

William Yurcik  
byurcik@ncsa.uiuc.edu

Suvda Myagmar  
myagmar@ncsa.uiuc.edu

National Center for Supercomputing Applications (NCSA)  
University of Illinois at Urbana-Champaign (UIUC)

## ABSTRACT

As enterprise storage needs grow, it is challenging to manage storage systems. The costs of locally managing, supporting, and maintaining resilience in storage systems has skyrocketed. Also, companies must comply with a growing number of federal and state legislations mandating secure handling of electronic information.

In this context, outsourcing of storage to utility-model based service providers has emerged as a popular and often cost-effective option. However, this raises issues related to data safety and storage techniques. In this paper, we discuss the business model and evolution of service-oriented companies known as Storage Service Providers and examine the challenges organizations should consider when outsourcing their storage management. To our knowledge, this is the first work to study the SSP model from both technical and business viewpoints. Lastly, we present two case studies, one of a failed SSP and the other of a successful market leader.

## Categories and Subject Descriptors

D.4.2 [Software]: Operating Systems—*Storage Management*

## General Terms

Management, Measurement, Performance, Design, Economics, Reliability, Security, Standardization, Legal Aspects

## Keywords

Storage Service Provider, Data Protection, Outsourcing

## 1. INTRODUCTION

In recent years, information generation, analysis, and archival has reached an all time high while storage of information has also increased in organizational importance. The expectation by most organizations is that storage needs will increase and expand. Managing storage is inherently difficult, and

with the proliferation of multi-terabyte storage system usage, it is becoming more challenging.

Traditionally an organization would buy storage systems for each of its sites or acquire equipment for its data center locations. This is a major cost for companies that are growing fast or are geographically distributed. Not only is the cost of purchasing dedicated storage systems a dominant factor, but the expense of locally managing, supporting, and maintaining resilience in storage systems is costly. In fact, storage management is one of the largest expenses in an organization. According to a Gartner group study [23], the cost of managing data protection and storage is 5 to 7 times the cost of hardware, and 74% of the total storage related costs. The amount of regulation around the legal use

Table 1: Federal/State Legislation

Legislation	Description
Health Insurance Portability and Accountability Act (HIPAA) [3]	Mandates confidentiality, integrity and availability of patient and other medical records and requires healthcare providers to secure the stored records.
SEC 17a-3 and 17a-4 [1]	Requires financial entities to retain client correspondence and all electronic records preserving integrity for auditing.
Gramm-Leach-Bliley Act [2]	Requires financial entities to disclose policies for protecting confidential customer information and ensure integrity and confidentiality while preventing unauthorized access.
Sarbanes-Oxley Act [9]	Ensures accuracy and reliability of corporate disclosures by requiring validation of integrity and accuracy of financial records.
California State Law SB 1386 [8]	Requires all state agencies and businesses that store client information to promptly disclose security breaches.

and retention of electronic information is increasing and so too are the penalties for non-compliance. Table 1 summarizes laws and regulations related to information usage and storage. Sarbanes-Oxley Act [9], HIPAA [3], Gramm-Leach-Bliley Act [2] and others mandate that financial information, patient records, and other client-related information must be handled in a secure manner. Penalties include corporate, civil, and criminal sanctions, with individual accountability to the CEO level. To meet these responsibilities, good in-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

StorageSS'05, November 11, 2005, Fairfax, Virginia, USA.  
Copyright 2005 ACM 1-59593-223-X/05/0011 ...\$5.00.

formation and storage management is vital. SSPs also offer a cost effective hosted environment that can grow to accommodate increased demands and can be readily used to store less critical, longer term data, thereby reducing the burden on a business' main storage facilities.

To reduce storage-related costs while complying with regulations, many companies choose to outsource their data to third party service providers, called Storage Service Providers (SSPs). SSP is a company that provides information storage space and the related management services [26]. SSPs manage the storage, archival, backup, recovery, and security of information. On-demand or Utility storage is another terminology used for such a service [40]. Several enterprises share storage infrastructure and typically only need network connections to the storage provider's data center. This enables companies to cut costs in terms of outlay on equipment and less management, although informed supervision of the SSP is still required. Moreover, as part of a huge infrastructure, often the more storage that is used, the cheaper per unit it becomes.

Storage may be better managed and cost-effective when handled as an outsourced service, which has contractual properties, instead of capital-intensive locally managed storage [34]. This is significant in three ways: (1) handling large storage requires expertise. It is not practical for many organizations to have their own storage specialists; (2) by outsourcing and having a binding contract and service-level agreement on the safety and availability of data, the economic losses and legal liabilities can be transferred to the service provider; (3) in a storage utility model, each organization only pays for as much storage as it needs. An organization can start with a small amount of storage, and gradually increase the storage size, or can have the additional data stored with different providers, possibly based on importance of data.

While storage outsourcing can offer customers a system tailored to their specific data requirements and business model, it is important that businesses evaluate storage providers based on multiple criteria including:

- What type of storage provisioning is appropriate for the company?
- What type of management services does the company need?
- Is company's data protected from threats?

In this paper, we provide analysis that will help to answer these questions. We trace the evolution of the SSP model through its ups-and-downs to a recent revival. Then we present classifications of SSPs based on various criteria such as management style, storage provisioning, and available services. Organizations may decide which type of SSP model is right for them based on this analysis. Our discussion of advantages and disadvantages of SSPs gives insight into what storage-related issues are solved and unsolved by outsourcing. Finally, we present case studies of two SSPs. Although there have been work researching technical aspects of storage systems and information about the business model of SSPs, to our knowledge, this is the first study of SSPs from both the technical and business viewpoints.

The remainder of the paper is organized as follows. Section 2 discusses the evolution of the Storage Service Providers, their business models, and the classification of their services.

Section 3 presents the main advantages of using SSPs for storage. In Section 4, the main challenges in the outsourcing of storage are explored. In Section 5 we present a case study of a successful and an unsuccessful SSP. We conclude with a summary in Section 6.

## 2. THE EVOLUTION AND CLASSIFICATION OF SSP

SSPs propose to help organizations deal with the problem of deploying and managing storage. To fully understand the current trend of SSPs, we first look at the evolutionary path of the SSP market. Then we dissect the SSP business model along multiple dimensions. With more companies demanding efficiency and flexibility from their storage system, utility storage seems to be an attractive SSP model [40]. However, the true utility model barely exists in reality, most companies fit somewhere in between the traditional storage model and utility storage.

### 2.1 Evolution of the SSP Market

The SSP industry has lived through an entire boom-to-bust lifecycle in just about 10 years. Figure 1 shows how the SSP market landscape shifted: initially dominated by small startup companies, and then later taken over by traditional storage vendors.

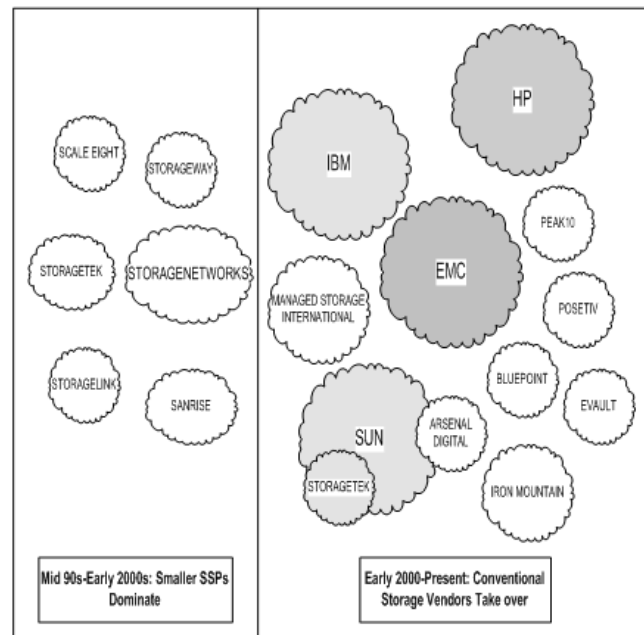


Figure 1: Evolution of the SSP market.

#### 2.1.1 The Boom

The SSP market was initiated in the mid-1990s, during the dot-com boom years, when many startups began as storage service providers. SSPs emerged with much fanfare promising to transform storage into a utility. Some of the early players in storage service business include StorageNetworks, ScaleEight, Storability, StorageWay, and Sanrise. One estimate from IDC [14] predicted that storage services would achieve 139 percent growth per year and will grow up to \$11 billion by 2005. By mid-2001, there were more than 20 storage service providers competing for the market [31].

Most of the startups operated in the utility storage model and focused on providing hosted storage service.

Many of the early adopters of SSP services were new e-business ventures, most of them with very few IT staff on site to support storage. The SSP model suited them well because they were accustomed to purchasing services— application, network, and then storage— on a commodity-type basis. Traditional storage vendors such as IBM, EMC, EDS, Compaq, HP, and Sun, said they would be SSP enablers— providing the storage, servers, consulting and integration on the back-end while continuing to push Storage Area Network (SAN) hardware and software directly to their corporate customers. These vendors had no immediate plans to provide managed storage services themselves at the time.

### 2.1.2 The Bust

The SSP market experienced a bust in early 2000s. Several factors can be attributed to this phenomenon. First, the economic downturn forced many dot-coms to shut down and SSPs lost large portions of their revenue from these customers. Second, many enterprises did not have enough trust on SSPs to outsource primary disk operations [35]. Thus, it was difficult for SSPs to resell the disk capacity they had purchased. Another reason was requirements for small capacity storage and availability of cheap systems to meet the storage needs locally [20]. When large data hosting company Exodus went bankrupt in 2001 [31], it affected many SSPs because they depended on Exodus to deliver services [13]. Also, many SSPs lost big clients when they moved to traditional vendors like IBM or HP.

Many SSPs quickly shifted gears and changed their business model. Some of them assumed the role of managing existing storage environments for large enterprises. Others converted themselves into storage management software vendors. The traditional vendors still operated as “enablers”. These vendors believed they could leverage their storage expertise and familiarity with their own platforms while maintaining their offerings in storage-network consulting, design, and integration as part of total storage solutions.

### 2.1.3 The Revival

The revival of the SSP market started when traditional storage vendors transitioned into managed-storage market. IBM, HP, EDS, and Sun entered the market with strong reputations and expertise in enterprise solutions. Also, traditional document archiving vendors like Iron Mountain extended their business to include electronic storage services.

There are other reasons why the storage landscape has changed and SSPs may now succeed. The economic downturn seems to have ended. The demand for storage capacity continues to grow while IT budgets have grown only slightly [25]. Although price of storage capacity has gone down significantly, cost and overhead of storage management have increased at much faster rate. This means data storage must be managed more efficiently. Regulatory requirements are having a great impact on storage. Addressing legal requirements on top of dealing with capacity and management demands when budgets are tight appear to be the catalysts for considering SSP services in 2005.

## 2.2 The Business Model of SSP

In the SSP model, SSPs sell storage and its management as a service to their customers. Because the SSP model im-

plicitly carries with it the ability for an organization to effectively buy or return storage capacity to suit its expanding or shrinking needs, the SSP model actually treats storage as a commodity. The SSP model can be classified by several criteria: Management Style, Storage Provisioning, Hosting Site, and Service Types.

Management style examines whether the SSP manages the storage for its customers or simply leases out its storage infrastructure and media. Storage provisioning refers to the method by which SSP allocates storage capacity to its customers. Storage services can be provided for storage residing either at client’s site or SSP’s site. There are many types of services that SSPs provide as part of their total storage solutions. Storage backup is considered as the most sought after storage service.

**Management Style.** SSPs differ by their storage management style [30]. The first type of SSPs can be termed as *Pure-Play* SSPs. The business model followed by these SSPs is closer to the utility model. These SSPs offer storage on a per-megabyte basis on systems owned and managed by them. In most cases the storage service is provided over the Internet rather than through proprietary networks. The second group of SSPs are *Traditional Storage Vendors* which provide managed storage on their own proprietary equipment, often owned by and residing at the customer’s facilities. In this case, the storage vendors take the responsibility of managing the storage system.

**Storage Provisioning.** SSPs can be classified according to their storage provisioning. Some of the SSPs provide *capacity-on-demand* service [27], which means that clients pay in advance a flat fee for a fixed maximum capacity they intend to use and are allowed to use up to that capacity. The other type of provisioning is *utility computing* where there is no fixed capacity; clients are billed according to their usage, similar to electricity and other utilities. The rest of provisioning methods lie somewhere in between these two models.

**Hosting Site.** Finally, SSPs can also be classified as providing *on-site* or *off-site* services [6]. In the former case, the SSP provides management and other services with the storage residing at the client’s site, while in the latter case, the SSP hosts storage on its own site.

## 2.3 Classification of Services

The services provided by SSPs can be classified into five categories, as shown in Table 2. The desirability of these technical services varies with innovations and new cost structures. **Managed Storage.** An SSP offers managed stor-

**Table 2: Classification of Services.**

Service	Description
Managed Storage	SSP provides pay-per-use service, leasing disk space to enterprises.
Remote Backup/Recovery	SSP handles backup tasks to prepare for unforeseen disasters.
Data Protection	SSP provides confidentiality and integrity of data.
Replication	Mirror the data for better performance and availability.
Archiving	SSP handles archival of data to comply with regulations.

age services, which means it provides disk storage space that

it leases to companies. It is also termed as pay-per-use service. Managed storage can be provided either on-site, or off-site. Management of storage is a large portion of total cost for information storage systems. By taking over the management of storage from companies, SSPs offer a cost-effective way of expanding storage without greater capital investment. SSPs manage terabytes of storage and provides all backup, fault tolerance, high availability, security, fire protection, and 24-hour service for organizations that decide not to manage these services on their own.

The economics of storage is measured by *Total Cost of Ownership (TCO)*, which includes software, hardware, training, management, operational, and data unavailability costs [24]. Handling storage locally can cost up to 70% of total IT spending [19]. A 2001 study by Goldman-Sachs estimates that in-house storage management costs \$140 per gigabyte per month, while outsourcing storage to SSPs costs \$40 per gigabyte per month [16]. Outsourcing storage to SSPs also saves the client from investing a large capital into storage hardware, thus increasing the TCO.

With the rise in volume, storage management has become a significant overhead. There is a shortage of personnel with expertise in storage management tasks. Curran *et al.* [17] state that 85% of all IT departments experience a storage staffing problem. According to [22], storage management costs about \$80 billion a year. Outsourcing management tasks to an SSP can reduce the management overhead and the staffing problem to a great extent.

**Remote Backup/Recovery.** Backup is a promising area of opportunity for SSPs. Some enterprises only complete 50-60% of the backup jobs they attempt [17]. This number can be significantly improved if outsourced to SSPs. Improving backup success rates makes enterprises better positioned to respond to unforeseen disasters. Backup outsourcing also frees up internal IT staff from a time-consuming task and allows them to focus on other IT priorities. A Gartner/Dataquest survey shows that 58% of IT managers focus on outsourcing backup and recovery to SSPs [16].

**Data Protection.** According to IDC [15], 60% of corporate data is unprotected, and vulnerable to unauthorized access and tampering. However, ensuring the security of data while maintaining compliance with regulations is a not cost-effective for smaller businesses. Therefore, many SSPs provide data protection services where the SSP manages protection of data using cryptography (encryption, decryption, and key management).

**Replication.** Some SSPs provide replication services for storage where the data is effectively replicated logically or geographically to provide better performance and availability. Replication also provides security and integrity of storage, and may move data closer to customers.

**Archiving.** Many SSPs provide archival storage of information. In some cases, the use of SSPs provide better sharing of archived information, such as patient medical images/records, while in other cases, it is mandated by law. Many companies need to store financial communication, such as email, instant messages etc, for up to 7 years [9].

### 3. SSP ECONOMIC ADVANTAGES

Reducing storage management costs is the primary benefit provided by the SSP model. However, SSPs offer other advantages over the enterprise data center model:

**Efficient Data Sharing.** By consolidating storage into a common location, companies that share a common goal or partnership can achieve efficiency. For example, health-care organizations can share the patient records with insurance agencies and doctors. With proper access control, information sharing can reduce the redundancy to a great extent.

**High Availability.** Availability is a significant issue to the service oriented businesses. According to [32], the cost of downtime in an organization varies from \$1000 to \$100,000 per hour; it may even be as high as \$100,000 for real-time transactions. A study by Meta group calculates the average cost of downtime at \$1 million per hour [38]. By handing over the task to SSPs, a company can ensure the availability. SSPs can enhance their availability by using redundant systems and replication, and technologies like SAN which have only 1 hour of downtime per year (on average).

**Quick Provisioning of Storage.** Storage outsourcing allows organizations to buy extra or less storage during times of fluctuation in their marketplace. A retailer may need a significant increase in storage provision over the Christmas period but a sharp decline in storage in January [40]. It is simply not cost-effective for such companies to buy large amounts of storage for temporary storage spikes. By provisioning storage on demand, SSPs can cater to the requirements of clients needing storage for limited times in a cost effective way.

**Better Storage Utilization.** Many companies cannot accurately predict how their business will grow or how their storage needs will evolve. Storage utilization is typically very low during the initial period. A 2002 Study by IT-CENTRIX [39] show that only about 40% of direct attached storage is utilized, while utilization can be increased to more than 60% through consolidating storage. By outsourcing the storage to SSPs, a company can arrange to only pay for as much storage as it needs. On the other hand, SSPs cater to many organizations at a time, so the storage systems they use are better utilized.

## 4. SSP CHALLENGES

In spite of these advantages, there are many open issues related to storage services. In this section, the SSP challenges are discussed.

### 4.1 Scalability

Most SSPs today handle many petabytes of storage. With dropping hardware costs, an SSP needs to organize storage out of many components. The challenge is to service many different client, simultaneously, with many different storage system components.

### 4.2 Providing QoS and SLA

SSPs need to maintain Quality of Service (QoS) according to the service level agreement (SLA) made with its clients. The SLAs include guarantees about performance, capacity, cost, availability, reliability and security. The SLA also negotiates the penalties and monitoring rules for achieving the QoS objectives. To monitor whether SSP provides service as promised in SLA, monitoring the usage and performance is important. In general, such monitoring is provided by a third party. However, it is still interesting to see if such monitoring could be provided without relying on a third party. The challenge is to negotiate many unique SLAs and then manage the individual performance of each when QoS is of-

ten largely determined by factors beyond SSP control (e.g. traffic congestion, outages, attacks, etc).

### 4.3 Bandwidth Cost

SSPs are focused on storing data, so transmission of storage contents to and from clients requires high bandwidth. However, many of the SSPs use Fibre-Channel for high speed connection. While Fibre-Channel is fast, it has distance limitations, so storage outside the local area cannot be accessed via these high-bandwidth connections [7]. The challenge is that most telecommunication providers do not often have dedicated optical fiber available from a customer to an SSP and the cost for installation would be prohibitive.

### 4.4 Latency and Connectivity

One of the most important performance issues related to an SSP is the latency faced by clients in accessing outsourced storage. Most of the SSPs use Fibre channel-based SANs for internal connectivity. Inter-site connectivity, that is, the connectivity between the client and the co-location facility, is usually done by means of Metropolitan Area Network (MAN) and Wide Area Network (WAN) [4]. Two types of technologies provide high-speed, low-latency connectivity in WAN and MANs: *Coarse/Dense Wavelength Division Multiplexing (C/DWDM)* and *Synchronous Optical Network/Synchronous Digital Hierarchy (SONET/SDH)*. Using these optical technologies, it is possible to achieve a one-way latency of 40 milliseconds between New York and Los Angeles [18]. Micro and Macro-benchmark results from Ng *et al.*[33] show that the latency can be hidden by use of file system and operating system caching. To provide satisfactory access to outsourced data, I/O performance needs to be fast enough not to affect application performance at the client side. [33] also shows that applications accessing outsourced storage are affected by network latencies, but read/write caching and prefetching can lower the latency significantly. As data outsourcing becomes a global phenomenon, SSPs would need to implement different techniques to better match latencies in order to provide acceptable performance to clients. Also, the decision between implementing storage on-site and off-site also depends on the nature of applications. High bandwidth/low-latency applications like audio and vide editing should not be outsourced to off-site storage; such storage would benefit from having an SSP-managed in-site storage facility [11].

### 4.5 Business Longevity

In the early days of SSPs, especially during 1999-2001, many startup SSPs provided storage services in different models. However, the initial attempt to create the SSP business largely failed [31]. Although an economic downturn can be attributed as a reason, the reluctance of companies to hand over data to a third party also contributed. The issue about the longevity of SSPs still is a problem, and many companies trust traditional storage vendors rather than start-up SSPs in outsourcing their storage, since outsourcing storage links their business success with the success of the SSP.

### 4.6 Security and Trust

One of the reasons for the early failure of the SSP model is the lack of trust. While client companies may trust an SSP to provide services such as reliability, availability, fault-

tolerance, and performance, they cannot trust if an SSP is going to use information for other purposes. For example, the design document of a new Intel chip should be much more expensive than a subscription fee Intel pays to an SSP. Because of this, data should be encrypted by the client (writer) and sent to the SSP and later decrypted by the client (reader). By doing this, an SSP does not have to worry about leakage of information by one of its employees and is better able to focus on other services. However, this End-to-end encryption approach has several drawbacks. Now the client has the burden to manage keys and key management is not considered to be an easy problem. Second, search and indexing (which should be provided at the SSP) becomes much harder (if not practically impossible) since data is already encrypted [37]. Lastly, while this approach can solve the read access control problem (i.e. an SSP cannot read data), write access control is still questionable. For example, suppose A writes a file and later the file is modified by B. An SSP may provide the old version of the file written by A to a reader C. Therefore, an integrity mechanism should guarantee that the file C receives is the file most recently updated by B. [29]. Also various kinds of denial-of-service attacks, such as deleting files, are always possible. Since no solutions currently exist, a client will have reduced trust in its SSP if it does not take actions to mitigate such attacks.

### 4.7 Legal Liabilities

The recent wave of security breaches in storage systems has resulted in drafting of legislation to make companies more liable for the loss of data. For example, California State law SB1386 [8] mandates that any company that keeps confidential customer records (including financial, personal, non-publicly available information), must notify the customers whenever a security breach occurs. The notification must be written, electronic, and also in case of large scale breaches, be published in statewide media. Other regulations [1, 9, 2, 3] focus on providing security for financial and health-care data. These legislations pose interesting questions as to who should be held responsible when a security breach occurs. Storage attacks at a single SSP can affect all of its clients.

## 5. CASE STUDIES

In this section, case studies on two SSPs are presented. The first one, StorageNetworks, was one of the early market leaders, but it failed due to its business model. The second case study features IBM, the traditional business-service vendor, which has been successful with its managed storage services even though it entered the market late.

### 5.1 StorageNetworks

StorageNetworks was one of the pioneers in the SSP market. When it was established as a startup in 1998, StorageNetworks took advantage of growing bandwidth and storage to create a new market segment for the management and hosting of storage. During the dot-com boom, StorageNetworks was able to raise \$203 million in venture capital [42]. When it went public with its IPO in 1999, StorageNetworks managed to raise \$243 million, with their share prices rising to \$90.25, more than triple the value, on the very first day. At its peak, the market value of the company rose to \$7.96 billion.

**Business Model and Performance.** StorageNetworks' initial business model was to be a pure-play SSP. It envisioned huge data centers in each city, servicing clients in the same way as other utility companies. By 2000, it had opened 36 data centers in Chicago, New York, Los Angeles and San Francisco, with close to 100,000 Gigabytes of on-line disk storage at each center [5]. It offered several levels of services [21]: the NetPACS service used Network Attached Storage (NAS) servers to connect to a company's servers through an IP network, while the DataPACS service utilized SANs connected to a company's file or application servers via Fibre-Channel links. The company kept growing with the dot-com boom, gaining more than 132 clients by 2001, and in 2001 Forbes ranked it as the second fastest-growing technology company [41].

However, one potential problem was to gain trust of the big businesses. In 2001, smaller clients accounted for 57% of the revenue, while larger companies made up only 43% [42]. This distrust and asymmetric business scenario proved fatal for StorageNetworks when the dot-com bubble burst, and many of the smaller clients went out of business. StorageNetworks faced huge operating losses, and its share value dropped down to 83 cents in February 2003 [28]. Another possible reason for the dramatic fall may be the 2001 bankruptcy of Exodus Communications, which hosted many of the data centers used by SSPs like StorageNetworks [31]. Trying to adapt to the changing market situation, StorageNetworks closed many data centers and switched to being a storage management software vendor. This was actually driven by the major clients of StorageNetworks including Merrill Lynch, Ford, Microsoft, and Sun, which preferred on-site management rather than remote management. The switch of business model incurred write-offs of \$114.4 million adding to a total loss of \$132.9 million in the fourth quarter of 2002 [42]. Massive layoffs and loss of clients continued and in 2003, StorageNetworks finally liquidated its assets and went out of business [28].

The significant lesson from StorageNetworks' failure is that even though its idea of providing utility storage was novel, it lacked the necessary business relationship with its customers, who were wary of trusting a startup with their critical data.

**Technical Overview.** Storage Networks planned to create a nationwide Fibre channel based *Virtual Private Network (VPN)*, that would connect the clients to Storage Networks' data warehouses [36]. These data warehouses, also called *StoragePOPs*, were to be built using EMC Symmetrix arrays, Connectrix connectivity devices and also Hitachi and IBM storage devices. For file systems, it supported NetApp and Sun file server. These servers were to be located in large metropolitan areas and would be equipped with redundant storage servers, and power supplies. Besides providing storage hosting, Storage Networks provided other services like replicating Oracle databases, managed backup/recovery, off-site data mirroring, and data migration facilities.

## 5.2 IBM

IBM, the traditional giant in business-oriented computing, has entered the Storage Service Provider market late. With its experience in enterprise management, IBM has managed to establish a profitable storage service in the managed storage model.

**Business Model and Performance.** As of 2005, IBM leads the storage services market. According to a Gartner group report [12], IBM's total market share was \$22 billion in 2004 and is expected to grow up to \$30 billion by 2008. Also, IBM is projected to remain the market leader in storage services through 2008. IBM's managed storage service [10] provides customers with high-availability, and utility access to storage. IBM's Managed Storage Service provides customers with three options: Basic, Premium and Customized services. Availability in Basic and Premium service is 99.9% [10], while in custom service, the clients can define their desired level of availability. Basic and Premium services need maintenance windows of only about 8 hours/month and 8 hours/quarter respectively. Storage provisioning can be done in less than 48 hours.

**Technical Overview.** The storage is provided from IBM's pool of SAN storage built with IBM, EMC and Network Appliance storage connected through Cisco, McData and Brocade switches. The system internally runs Unix with Tivoli and EMC's storage management software. IBM's Storage pool is monitored from its Storage Operations Center in Boulder, Colorado and a back up center in Raleigh, North Carolina. These two geographically distributed centers help in providing fail-safe operations in case of natural disasters and large-scale power failures. The centralized storage operation center also manages storage located in customer premises, the management being conducted remotely. Storage provisioning and capacity requests are managed via a web portal, with storage available on-demand. Security issues are handled seriously. The storage is implemented on a fibre channel backbone network. The SAN-switched fibre channel fabric connects servers to storage devices.

IBM's success shows that its reputation as a vendor with established business relationship with organizations, and a business model more suited to the existing needs of clients are some of the critical factors in the SSP market. While StorageNetworks initiated the SSP business, it tried to create a new business model, which clients were hesitant to adopt. The deciding factor may very well be its long reputation as a vendor, which allowed clients to entrust their data with IBM. Also, IBM's series of storage-related hardware and software enabled it to have a solid storage infrastructure.

## 6. SUMMARY

Outsourcing storage for cost efficiency and competitive advantage is a new service enabled by the Internet that has become attractive to many organizations. In this paper we outline the SSP market from both business and high-level technical viewpoints highlighting business models, competitive advantages, and case studies. Despite extreme business life cycle dynamics, the SSP market is young and may yet prove to be a stable and long-loved business sector. We hope this work provides a background upon which future work (either business or technical) may build to go further in-depth on some of the issues we highlight.

## 7. ACKNOWLEDGEMENTS

We first acknowledge Yongdae Kim (University of Minnesota) who was instrumental in shepherding this paper for this workshop and contributing expertise on storage security issues. We next thank Prof. Yuanyuan Zhou (University of

Illinois) for sharing her experiences with StorageNetworks. We must mention the indirect contributions of our colleagues (in alphabetical order) at NCSA and UIUC: Larry Brumbaugh, Adam Lee, Paul Stanton, Joseph Tucek and Prof. Marianne Winslett. Lastly, the comments and constructive criticism of the anonymous reviewers on previous drafts were invaluable and were incorporated into this final paper.

## 8. REFERENCES

- [1] *General Rules and Regulations Promulgated under the Securities Exchange Act of 1934*. Federal Trade Commission, <http://www.mfainfo.org/washington/1934.pdf>.
- [2] *Gramm-Leach-Bliley Act of 1999*. Federal Trade Commission, <http://www.ftc.gov/privacy/glbact/>.
- [3] *The Health Insurance Portability and Accountability Act of 1996 (HIPAA)*. Centers for Medicare & Medicaid Services, <http://www.cms.hhs.gov/hipaa/>, 1996.
- [4] Application note: Expanding storage service provider connectivity solutions. *LuxN/San Valley*, 2000.
- [5] Storagenetworks up 234%. *CNN Money*, June 2000.
- [6] The emerging storage security challenge. *Somix Technologies Inc., White Paper*, August 2001.
- [7] SSP application brief. *FalconStor White Paper*, May 2001.
- [8] *California Database Breach Act (SB 1386)*. California Senate, [http://info.sen.ca.gov/pub/01-02/bill/sen/sb\\_1351-1400/sb\\_1386\\_bill\\_20020926\\_chaptered.html](http://info.sen.ca.gov/pub/01-02/bill/sen/sb_1351-1400/sb_1386_bill_20020926_chaptered.html), 2002.
- [9] *Sarbanes-Oxley Act of 2002*. U.S. Securities and Exchange Commission, <http://www.sarbanes-oxley-forum.com/>, 2002.
- [10] Managed storage services: Improving performance without increasing capital expense. *IBM Global Services*, June 2004.
- [11] Storage: Thinking outside the box. *iSCSI Storage*, 2004.
- [12] IBM leads in storage services. *DSstar News*, July 2005.
- [13] D. Allen. Storage service providers: By any other name. *Network Magazine*, May 2002.
- [14] G. Blackwell. Storage service providers play. <http://isp-planet.com/services/2001/ssp.html>.
- [15] F. Brick. Mass exodus: The movement of data off site. *Disaster Recovery Journal*, 17(3), 2004.
- [16] P. Curran and B. Youngren. Making the case for ssps. *Infostor Magazine*, April 2001.
- [17] P. D. Curran and B. D. Youngren. What's ahead for storage service providers? *Infostor*, May 2002.
- [18] R. Dufrene. Data replication over the wan - connectivity. *Computer Technology Review*, October 2003.
- [19] L. Erlanger. Storage as universal data tone: New technologies provide global access to your storage systems and your data. *Internet World*, August 2001.
- [20] C. Forelle. Data-farm flop points to flaws in new tech fad. *The Wall Street Journal*, August 2003.
- [21] L. Freed. Outsourced network storage. *PC Magazine*, April 2001.
- [22] J. Gray. Storage service providers? *Panel talk at the 27th International Conference on Very Large Data Bases (VLDB)*, 2001.
- [23] B. Hoard. What is behind the move to storage networks? *Storage Networking World Online Magazine*, April 2001.
- [24] R. Kerns. The economics of a storage strategy. *ComputerWorld*, October 2001.
- [25] R. Kerns. Storage service provider revival? *Computerworld*, March 2005.
- [26] Y. Kim. Security and accounting for storage service provider. In *Intelligent Storage Workshop, University of Minnesota*, 2005.
- [27] J. F. Kovar. Storage on demand. *Storage Pipeline Magazine*, November 2003.
- [28] J. F. Kovar. Storagenetworks files plans to liquidate. *CRN*, July 2003.
- [29] J. Li, M. Krohn, D. Mazires, and D. Shasha. Secure untrusted data repository (sundr). In *Proceedings of the Sixth Symposium on Operating Systems Design and Implementation (OSDI)*, pages 121–136, San Francisco, CA, USA, December 2004.
- [30] J. Madden. Take my data ... please! storage service providers offer viable alternative for businesses. *DM Review Magazine*, May 2001.
- [31] J. Maitland. SSPs:rip. *Byte and Switch*, January 2002.
- [32] G. Nagle. The cost of owning and storing data. <http://www.storagesearch.com/costdata.html>, July 1999.
- [33] W. T. Ng and B. K. Hillyer. Obtaining high performance for storage outsourcing. In *ACM SIGMETRICS*, June 2001.
- [34] B. Ozden, E. Gabber, B. Hillyer, W. T. Ng, and E. Shriver. Storage service providers: a solution for storage management? In *Proceedings of the 27th International Conference on Very Large Data Bases (VLDB)*, 2001.
- [35] S. Schmelling. This year's storage model. *The Net Economy*, April 2002.
- [36] Z. Shess. Storage outsourcing provider opens shop. *Infostor*, May 1999.
- [37] D. X. Song, D. Wagner, and A. Perrig. Practical techniques for searches on encrypted data. In *SP '00: Proceedings of the 2000 IEEE Symposium on Security and Privacy*, page 44, Washington, DC, USA, 2000. IEEE Computer Society.
- [38] J. W. Toigo. The ins and outs of storage management. *Network Connections*, 2001.
- [39] D. Vellante. In search of storage ROI. *SMS*, 7(4), 2002.
- [40] C. Voo. The case for utility computing. *TechTarget*, February 2005.
- [41] C. Willis. Ramp champs. *Forbes ASAP*, February 2001.
- [42] T. Witkowski. Storagenetworks pays \$114m price for software shift. *BizJournals*, February 2002.

**Appendix A**  
**Table 4: SSP Market Survey**

SSP	Type of Services					Offsite/ Onsite	Comment	Business Model		
	Managed Storage	Remote Backup/ Recovery	Data Protection	Replication	Archiving			C D	P G	M S
AmeriVault <a href="http://www.amerivault.com">www.amerivault.com</a>	No	Yes	Yes	Yes	Yes	N/A		N	Y	
Arsenal Digital Solutions <a href="http://www.arsenaldigital.com">www.arsenaldigital.com</a>	No	Yes	Yes	No	No	Offsite	Has partnership with StorageTek/Sun for providing data protection through its ViaRemote Service			Y
BluePoint Data Storage <a href="http://www.bluepointdata.com">www.bluepointdata.com</a>	Yes	Yes	Yes	Yes	N/A	N/A		Y		
EDS <a href="http://www.eds.com">www.eds.com</a>	Yes	Yes	Yes	Yes	Yes	Both	Manages more than 10 petabytes of storage	N	Y	Y
EMC <a href="http://www.emc.com">www.emc.com</a>	Yes	Yes	Yes	Yes	Yes	Both			Y	
eVault <a href="http://www.evault.com">www.evault.com</a>	Yes	Yes	Yes	No	No	N/A	Mainly focuses on online backup and recovery			
HP <a href="http://www.hp.com">www.hp.com</a>	Yes	Yes	Yes	Yes	Yes	Both	Provides on-demand storage and managed storage services	Y	Y	Y
IBM <a href="http://www.ibm.com">www.ibm.com</a>	Yes	Yes	Yes	Yes	Yes	Both	Provides multi-layered managed storage services	Y		Y
InSiteOne <a href="http://www.insiteone.com">www.insiteone.com</a>	No	Yes	No	No	Yes	N/A	Mainly handles medical image storage offsite and in site	N	Y	N
IronMountain <a href="http://www.ironmountaindigital.com">www.ironmountaindigital.com</a>	No	Yes	Yes	No	Yes	N/A	Also provides storage of non-digital paper records, images etc.			
LiveVault <a href="http://www.livevault.com">www.livevault.com</a>	No	Yes	Yes	No	Yes	N/A	Provides disk based storage solutions	Y		
Managed Storage International <a href="http://www.managestorage.com">www.managestorage.com</a>	Yes	Yes	Yes	No	No	N/A	Provides remote monitoring and managed storage services			Y
Peak10 <a href="http://www.peak10.com">www.peak10.com</a>	Yes	Yes	No	Yes	N/A	N/A				
PoseTiv <a href="http://www.poseativ.com">www.poseativ.com</a>	Yes	N/A	No	No	N/A	Onsite	Provides remote management of on-site storage			
ScaleEight (Out of Business)	Yes	N/A	N/A	N/A	N/A	N/A	Went out of business			Y
Storage Networks (Out of business)	Yes	Yes	N/A	N/A	N/A	N/A	Went out of business in 2002			
StorageLink (Out of business)	N/A	N/A	N/A	N/A	N/A	N/A	Went out of business			
Sun/StorageTek	Yes	Yes	Yes	Yes	Yes	Both	Merged with Storability and StorageTek			Y
X Drive <a href="http://www.xdrive.com">www.xdrive.com</a>	Yes	No	No	No	No	N/A	More oriented towards small business or individual users	Y	N	N

CD = Capacity on Demand/Metered Storage      PG = Pay as you Go      MS = Managed Storage