

Storage Service Providers: a Solution for Storage Management?

Banu Özden, Eran Gabber, Bruce Hillyer, Wee Teck Ng, Elizabeth Shriver

Computing Systems Research Department
Bell Labs, Lucent Technologies
600 Mountain Avenue
Murray Hill, NJ 07974
USA

{ozden,eran,bruce,weeteck,shriver}@research.bell-labs.com

Panelists¹

- David DeWitt (University of Wisconsin)
- Bruce Gordon (StorageNetworks)
- Jim Gray (Microsoft Research)
- Banu Özden (Panel Moderator), (Bell Labs)
- John Wilkes (Hewlett-Packard Labs)

Panel Objectives

From an enterprise point of view, data storage is a huge and growing cost center, whose correct operation is vital to the success of the business. Moreover, the ongoing expense of operating and managing a business data center greatly exceeds the equipment purchase price. Some observers believe that storage is far more attractive when viewed as an outsourced service with contracted properties, rather than as a costly in-house facility. Hence the emergence of the Storage Service Provider (SSP) industry. In addition to new SSP vendors, existing communication, hosting and application service providers have recognized storage as a significant revenue opportunity, and are rapidly building data centers to offer managed storage services. IDC forecasts that storage services will be a \$10 billion market in 2005.

In this panel, we anticipate a lively discussion of the issues of storage management, the viability of the SSP model, and its implications on database systems.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, requires a fee and/or special permission from the Endowment.

**Proceedings of the 27th VLDB Conference,
Roma, Italy, 2001**

Storage Management: Although it is understood that managing large amounts of storage is complex and expensive, the issues in large-scale storage management are not well identified and quantified. We will debate aspects of storage management that are difficult, expensive, and critical, and we will seek an understanding of storage management issues involving heterogeneity, interoperability, and performance that arise as we move from host-attached storage to network-attached storage and storage-area networks. We will also discuss whether advances such as storage over IP and the InfiniBand Architecture will ease storage management.

Storage Service Providers: One way to deal with the problem of storage management is to let someone else do it. Potentially, the difficulties of storage management can be transferred to the SSP, who would handle backups and disaster recovery, capacity planning, resource allocation, and technology evolution. Moreover, the economic losses caused by a storage failure could be transferred to the SSP via contractual obligations under service-level agreements. But remote storage faces performance hurdles, and is vulnerable to external events such as network failures, security breaches, and denial-of-service attacks. It also calls for Quality of Service guarantees to meet the service-level agreements and requires the necessary support for monitoring and billing. The SSP model assumes that the customer is willing to place business-critical data into the care of an external organization. This assumption is not only challenged by the potential functional problems mentioned above, but also by other reasons why companies prefer to store data in-house. Can the customer depend on the SSP to purge sensitive files before the subpoena is served? Can the customer depend on the SSP to stay in business? Can a customer effectively transfer legally mandated

¹Confirmed at the time of printing

records-retention obligations to an SSP?

Database Implications: Current database systems have widely varying requirements with respect to performance, security, and reliability. We wish to understand the relationships among storage management, the SSP model, and the requirements of on-line transaction processing, data warehousing, and web and e-commerce back ends. For example, can locking and commit processing tolerate the latency of remote storage? Can a DBA tune remote storage?