

Storage Trend Top Ten Predictions in 2008

— A Perspective from Hitachi Data Systems Corporation —

Jack Domme

OVERVIEW: Hitachi is working with organizations of all sizes to reduce their IT costs and complexity, manage risk, and increase operational efficiency by delivering Services Oriented Storage Solutions that closely align storage infrastructure with business requirements. These solutions are comprised of hardware, software, and services. We announced this in May 2007, and we are addressing a full range of storage requirements with a single integrated platform. As customers learn the clear business advantages of Services Oriented Storage Solutions, they will begin to demand and expect it, and the architecture will be further adopted. We will constantly innovate around storage solutions to enable organizations and users to optimize storage infrastructure while reducing cost and complexity by integrating the functions needed by customers such as Green Technology, as predicated for storage in 2008.

INTRODUCTION

DATA growth continues its exponential expansion, fueled by increased capabilities and expanding regulatory mandates. Concerns about the impact of increasing data center sizes and the energy they use, coupled with the ongoing uncertainty about the stability of the global economy, are now significant

factors in IT (information technology) purchasing decisions. Those decisions favor vendors with advanced technology that is eco-friendly, scalable to meet increasing demand, and capable of consolidating multi-vendor solutions to deliver a lower total cost of system ownership and reduce complexity. With these factors in mind, here are my top ten predictions for

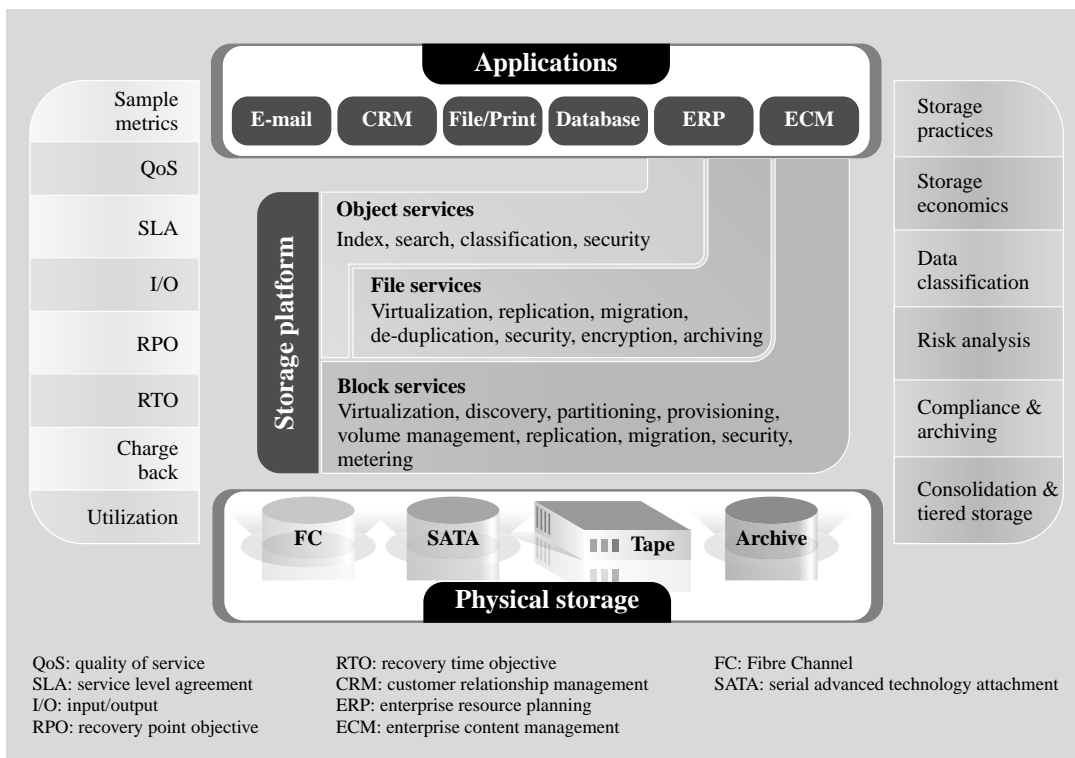


Fig. 1—Hitachi Services Oriented Storage Solutions Architecture. Services Oriented Storage Solutions is a business-centric framework for aligning IT (information technology) storage resources with constantly changing business requirements. It provides a dynamic, flexible platform of integrated storage services enabling organizations and users to optimize storage infrastructure while reducing cost and complexity.

storage in 2008.

FURTHER ADOPTION OF SERVICES ORIENTED STORAGE SOLUTIONS

In May 2007, Hitachi globally announced the industry's first set of Services Oriented Storage Solutions. Services Oriented Storage Solutions is a business-centric framework for aligning IT storage resources with constantly changing business requirements.

As business seek ways of getting more out of their IT infrastructure, data centers will need to become more dynamic in the future. Services Oriented Storage Solutions will become a requisite complement to Services Oriented Architecture in the application space and to Services Oriented Infrastructure. Services Oriented Architecture that enables the cost efficiencies of a dynamic data center, depends on a virtualization layer that allows applications to share information and utilize common services. Services Oriented Infrastructure depends on a virtualization layer provided by products like VMware*¹, which enables operating systems to share the resources of a processor platform. Services Oriented Storage requires a virtualization layer in the storage control unit which enables other storage systems to leverage its services like a high-performance global cache, distance replication, tiered storage, and thin provisioning. Treating storage as a set of services is the foundation of the Services Oriented Storage Solutions strategy, which reduces costs and minimizes management complexity, while optimizing storage infrastructure and improving efficiency.

CONTROLLING CARBON EMISSIONS

With increasing concern about global warming, we will see more governments impose guidelines and legislation around carbon emissions. Major corporations will set targets for the reduction of carbon emissions. A major source of carbon emissions comes from the generation of electricity. The increasing demand for compute power, network bandwidth, and storage capacity will increase the need for data center power and cooling. The US government just completed a study that estimated that the IT sector consumed about 61-billion-kW•h in 2006. That adds up to 1.5% of the total US electricity consumption, at a cost of about US\$4.5 billion, a figure that is expected to double

in the next five years. Some cities like London and New York are running out of electrical capacity, and data centers are forced to relocate to other areas where power is available. This will require facility upgrades and investments in Green technology.

INCREASING USE OF ARCHIVING

Structured data like databases will continue to grow, as they are required to hold more data and hold it longer to meet compliance regulations. Semi-structured data (e-mail, web pages, and document management) is increasing dramatically. Corporate e-mail quotas will increase from less than 200 Mbytes to 2 Gbytes in order to support new knowledge workers and compete with free mail box offerings from Google*², YAHOO!*³, and AOL*⁴. An avalanche of unstructured data will be driven by RFID (radio-frequency identification) tags, smart cards, and sensors that monitor everything from a person's heart to people crossing international borders. New aircraft from Airbus S.A. and The Boeing Company will generate Tbytes of data about each flight. All these pressures will drive the need to archive data in order to reduce the working set of production data. This will call for new types of archiving systems that can scale to Pbytes and provide the ability to search for content across different modalities of data.

CONVERGENCE OF CONTENT, FILE, AND BLOCK BASED STORAGE SERVICES

Instead of separate stove-pipe systems for content (archive) storage, file storage, and block storage, we will see the convergence of these storage types to a common virtualization platform. High availability clusters of content servers and file servers will use a common block virtualization services platform, under one common set of management tools. This will enable content servers or file servers to leverage common block services like distance replication, thin provisioning, or virtualization of heterogeneous storage systems.

AWARENESS OF STORAGE DEFICIENCIES

There will be a growing awareness that the storage of data has become highly inefficient, with low utilization, stranded storage, too many redundant copies, low access speeds, inefficient search and disruptive movement and migration. Continuing to buy

*1 VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions.

*2 Google is a trademark of Google Inc.

*3 YAHOO! is a trademark of Yahoo! Inc.

*4 AOL is a registered trademark of America Online, Inc.

more of the same old storage architectures or increasing capacity on the same 20-year-old architectures will no longer be viable. New storage architectures will be required to meet these changing demands that can scale performance, connectivity, and capacity, nondisruptively to multiple Pbytes.

DATA MOBILITY AS A KEY REQUIREMENT

With the need for continuous application availability, IT will need the ability to move data without disrupting the application. While software data movers have been used in the past, they steal processor cycles from the application and are limited to slow speed IP (Internet protocol) links to move data. As the volume of data increases, this becomes too disruptive. The movement of data will have to be offloaded to a storage system that can move data over high speed Fibre Channel links without the need for the application's processor cycles.

CONTROL UNIT VIRTUALIZATION OF STORAGE

This will be recognized as the only approach to storage virtualization that can add value to existing storage systems. Industry analysts like Dr. Kevin McIsaac of Intelligent Business Research Services in Australia points out, "The idea of being able to layer [network-based] virtualization over existing storage arrays is seriously flawed." This type of virtualization adds another layer of complexity and introduces a performance bottleneck, becoming another potential source of failure and a vendor lock in. A control unit based approach to virtualization is able to leverage all the rich functionality of the control unit to enhance the functionality of lower cost or legacy tiers of storage. A control unit based approach to virtualization will enable less capable storage systems to utilize the value added services in that control unit, like data mobility functions or thin provisioning capabilities.

THIN PROVISIONING

This will provide the biggest benefit in increasing the utilization of storage by eliminating waste from allocated but unused storage capacity. This savings is multiplied many times over by eliminating the need to copy allocated but unused capacity every time a copy is required for backup cycles, replication, data mining, and development test and data distribution. The implementation of thin provisioning should be provided as a service on a storage virtualization platform so that it can benefit existing storage systems

through virtualization. This ability to increase utilization will be embraced by Green advocates and will also be seen as a way to contain costs.

DE-DUPLICATION

De-duplication will be implemented by all the major backup vendors. It is especially effective in eliminating duplicated data in backups. The ability to reduce a stream of data by 20 to 30 times will be extremely valuable in reducing the cost of storing data to the point that it will be feasible to store backup data to disk rather than tape where the operational, availability, and reliability characteristics are better.

CONTINUED ECONOMIC UNCERTAINTY

The collapse of the housing market in the US, high oil prices, and the falling value of the US dollar will create economic uncertainty. Budgets will be tight, and IT will once again have to do more with less. Doing more with less will drive IT organizations to find ways to:

- consolidate IT resources through virtualization;
- increase utilization of resources such as server cycles and storage capacities;
- eliminate redundancies wherever possible through deduplication and single instance store; and
- reduce the working set of production data through the aggressive use of archive products.

CONCLUSIONS

In conclusion, the convergence of application processes such as data mining, integrated data discovery and data searching, and data protection and repurposing will require unprecedented scalability and integration that will be provided in the storage infrastructure. Increased mobile network bandwidth will put greater pressure on core storage and data services to satisfy the consumption of mobile information on demand. All that coupled with the demand for efficient, cost effective and eco-friendly storage will continue to fuel the growth of Hitachi Storage Solutions group.

ABOUT THE AUTHOR



Jack Domme

Joined Hitachi Data Systems Corp. in 2003.

As chief operating officer of Hitachi Data Systems,

Jack Domme leads a worldwide executive team

responsible for sales, corporate and product strategy,

IT management, and logistics planning and operations.