Enterprise Storage

Developing Return on Investment and Business Case Support for Storage Area Networks

A White Paper

by David R. Merrill

Hitachi Data Systems®—Built for fast times™

Executive Summary

Many companies have recognized new efficiencies when upgrading their IT infrastructures to include a Storage Area Network (SAN). It is these varied efficiencies and benefits that should prompt steps toward SAN procurement. SAN deployment should not be undertaken with the heightened expectation that near-term cost savings will always be the result.

Hitachi Data Systems recommends that a return on investment (ROI)/business case analysis be conducted by an organization prior to SAN installation. SANs *do* offer demonstrated areas in which their use can save money for the enterprise, but these savings will vary with differences in topology and the relevance of hard- and soft-cost savings to the organization. This white paper provides a structured approach for calculating SAN ROI. It identifies 29 cases—ranging from Increased Disk Utilization to Reduce/ Eliminate Batch and Backup Windows—in which HDS has observed SAN ROI. This case information with accompanying savings calculations can help our clients build strong and accurate business cases for SAN deployment.

Contents

Developing Return on Investment and Business Case Support for Storage Area Networks

by David R. Merrill

Introduction

Many IT departments are investigating advanced storage options and implementation plans to deploy centrally pooled storage. Storage area network (SAN) technology provides a variety of business and technical advantages over current disk attachment and data storage architectures. Most clients who investigate SANs recognize the technical or business value, but they often need further justification for the required infrastructure investment. This paper outlines a structured approach to determining financial payback and justification for advanced storage or SAN systems.

Since a SAN can radically impact the operational IT infrastructure, and the cost of a SAN infrastructure itself (without considering the cost of storage arrays) is often very high, a solid business case analysis should be performed in the planning phases of SAN design and justification. Without proper planning and impact analysis, the "sticker-shock" of Fibre Channel (FC) SAN technology components may inadvertently drive away some organizations that really should be considering this technology.

At Hitachi Data Systems (HDS), our observation is that the possibility of financial savings alone is an insufficient rationale for implementing Storage Area Networks. SANs and other pooled storage architectures can be expensive, and they do add complexity to IT infrastructures and operations. Often, new processes and technologies are required, as well as training to improve staff skills and the adoption of robust operating procedures. Storage systems attached to SANs are more expensive than the JBOD (just a bunch of disks) or attached disks on low-end servers. At is time, our research concludes that SAN deployment is not necessarily a cost-saving action; rather, it is an investment. As we see second-generation SAN adopters deploying this technology, we observe that more thought is required for the investment rationale, and demands for improved ROI cases are growing. This paper outlines justification and business impact analyses that can be completed to show offsetting terms for the necessary SAN investment.

As HDS continues to provide leadership in the implementation of this technology for worldwide clients, we are noticing clear investment payback and total cost of operation (TCO) improvements that validate the models outlined in this paper. As more empirical data becomes available and is certified by more users, it is plausible that the data may identify the SAN as a major cost-saving technology by mid-2002. This paper comprises theoretical approaches that use industry standards and units of measure (where available) to assist SAN investigators in determining business value and financial impact when implementing storage area networks. These models are being fine-tuned as further empirical data is made available.

HDS has adopted an open architecture—a topology-independent approach to SAN design and implementation—so not all IT departments have to consider complex Fibre Channel switch or director technology to achieve the benefits of central storage pools. As part of the HDS SAN planning methodology, return on investment (ROI) planning models are created to assist in early preparation for cost, benefit, and acquisition approval of new infrastructure technology. HDS has found that ROI has different payback and impact with different topologies. ROI is also different for each IT infrastructure. This knowledge is used routinely to help each client design and implement the SAN solution that best fits operational, technical, and financial requirements of their organization.

HDS has identified several functional areas that demonstrate valid ROI in business case planning for a SAN. Not all cases are equal in terms of measurement, justification, and validation; much of this will depend on the individual client environment. Most of the benefits identified appear for open systems, primarily because these systems tend to have the largest share of data management and scalability problems, and they are usually the prime target for pooled storage or consolidation on the SAN. The 29 SAN ROI cases are listed below, and a portion of these will be outlined later in this paper.

1. Increased disk utilization
2. Deferring disk procurement
3. TB-per-administrator improvement
4. Reduce data center rack/floor space
5. Deferring tape library procurement
6. New disaster recovery (DR) capabilities
7. On-line recoverability options
8. Improved data path availability
9. Reduction of general-purpose UNIX and Windows NT servers
10. Improve LAN/WAN performance, avoid upgrade
11. Reduce/eliminate backup servers
12. Reduce/eliminate batch, backup windows
13. Storage on demand
14. Improved protection of critical data
15. Management costs as a percentage of storage costs
16. Increased I/O performance, bulk data movement
17. Reduced storage maintenance
18. Staff utilization for server management
19. Increased life of current storage
20. Improved disaster recovery capabilities
21. Non-disruptive scalability
22. Avoid data area network growth
23. Impact on new or migrating applications
24. Impact on applications development, testing

25. Extending life of servers
26. Reduce CPU load on servers
27. Support server clustering
28. Secondary security services
29. Vendor consolidation

ROI Theory

Just now, empirical data supporting ROI and financial payback for pooled storage architectures (relative to the current status of SAN technology and the current technical reference model) is emerging. Yet, much SAN ROI analysis may remain theoretical in nature. HDS has found that most SAN savings will be soft-dollar savings—that is, savings that are less tangible and harder to validate by saving actual IT budget. These are often "feel-good" numbers that IT and company management should understand and appreciate, recognizing that normally, no tangible cost savings will be recaptured. However, soft-dollar savings typically represent the larger percentage of SAN ROI and payback dollars, so they should not be underestimated or disqualified from the analysis. Hard-dollar savings also will emerge from the ROI analysis; these are real savings that could be removed from future budgets or operating cost structures.

The combination of hard- and soft-dollar savings will yield the payback picture for the SAN infrastructure that can be used for management and technical decisions around SAN technology. Each IT organization will have to consider their infrastructure and determine which savings are hard, and which are soft. HDS generally has categorized (from observed situations and analysis) where the 29 ROI models fall in terms of hard- or soft-dollar savings. These characterizations are shown graphically in Figure 1.

Estimated savings can be characterized for conservative, aggressive, or nominal states (calculations) and should be considered in the analysis. Conservative states usually will include hard savings and highly likely savings areas (seen or predicted), while an aggressive approach will take hard- and soft-savings examples together for the ROI calculations. Nominal cases will take a middle-of-the-road approach and will be consistent with most acceptable/believable ROI case studies. Aggressive estimates are often used for soft-dollar savings, or when baseline data is not verifiable.

Each of the 29 case options identified in this report have different hard/soft considerations, and tend to be best applied in an approach that is clearly conservative, nominal, or aggressive. The best category for each case option is illustrated in Figure 1.

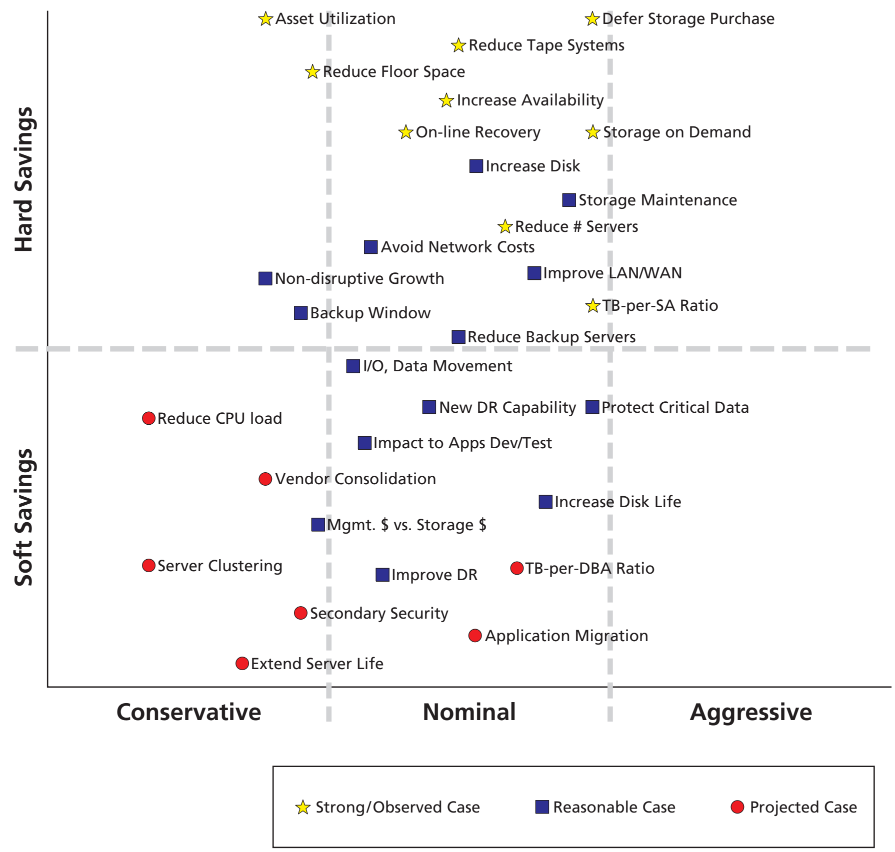

Figure 1: Comparing and Plotting the ROI Cases.

Simple ROI Determination Rules

- To achieve simplicity in ROI, cash flow, internal rate of return (IRR), and depreciation not factored
- Moore's Law used with price stability for all servers and disk price decline
- Lease costs converted to purchase costs
- All hardware costs can be pulled to year 0, at Net Present Value 0 (NPV_0)
- Or, hardware costs can be spread out over the time period in which it is acquired
- Our experience has found that three-to-four-year planning usually is adequate for most SAN payback models

A Structured Approach (Methodology) for Calculating SAN ROI

1. Gather baseline data (see examples in each case model)
2. Define the proposed SAN topology and then determine SAN costs
3. Select ROI cases that best meet the current and target optimum situations (typically, five to eight models are chosen for a business case justification)
4. Run the ROI case models and create summary charts and conclusions
5. Review conclusions and make adjustments to topology type, assumptions, and parametric values of the ROI models

Step 1: Gathering Baseline Data

Before starting a SAN ROI exercise, there are several baseline costs that will be required in order to determine current- or future-year costs. Each of these cost elements could be applied to any of the 29 ROI cases outlined in this paper, and each case will outline the necessary baseline elements needed. A summary of these baseline data points is shown below:

IT Infrastructure Cost Element	Unit of Measure	Time Period
System administrator, database administrator (DBA), storage manager, tape operator	Fully burdened labor rate	Man year
New storage (S/390®, UNIX, Windows NT)	$/MB	Current
Rate of storage price decline	XX %	Per year
Average unit cost of new servers (UNIX, Windows NT)	Purchase/lease cost	At the time of purchase
Average unit cost of Windows NT or UNIX servers with SAN kits (h/w, s/w, middleware, setup)	Purchase/lease cost	At the time of purchase
Expected increase/decrease of servers	+X% or –X%	Per year
Tape library systems	Purchase/lease cost	Current year/future year
Tape media (DAT, DLT, 8mm, other)	Each	Project yearly usage
Off-site tape storage	Per xx tapes or volume ft_	Yearly
LAN operational costs	Per node, NIC	Yearly
LAN degradation impact	% of time above threshold	Yearly
LAN build-out costs	$ per node	Current
Data center floor space	$/square foot	Yearly
Data center environmental costs	Per server or per ft_	Yearly
Application(s) downtime	Hours	MTBF (mean time between failures) or SLA (service level agreement)
Data off-line impact	Hours	Based on SLA
Recovery time for different types of data	Hours	Based on SLA
Revenue potential per person	$/per person/year	Based on revenue + staff
Monitoring cost	Per managed element	Yearly
Maintenance costs for all storage	$ per server or disk array	Monthly or yearly
Maintenance/license fee for storage s/w	$ per server or disk array	Monthly or yearly
Value of improved I/O, performance	Per transaction	Based on SLA
Time to procure/request net new disk	Lost time/hassle factor	Weeks

Table 1: Sample Baseline Costs to Be Used in ROI Calculations.

Step 2: Determining SAN Costs

In order to develop the proper payback models for the business case, the proposed SAN infrastructure has to have completed a logical design that will capture the design parameters and the "grocery list" of components for the baseline costs. This can be accomplished with a variety of services that can provide rough order magnitude (ROM) designs and costs based on information gathered from workshops or an on-site assessment. Once the design is roughed-out and blueprints are available, a preliminary list of SAN components (h/w, s/w, middleware, cabling, integration labor, etc.) can be summarized and priced.

HDS often finds that applying ROI models on more than one topology provides a useful comparative review of design and operational needs. Some topology design options can be very similar in price, but yield different TCO cost streams or ROI payback terms. Having two or three topology options—network attached storage (NAS), point-to-point (PtP), FC switch or director, etc.—will create a balanced financial approach to the analysis and cost/benefit determination.

Some (but not all) of the SAN ROI analysis is completed without considering the detailed cost of storage, since storage growth is assumed to continue regardless of the attachment architecture (SCSI, FC, SAN, ESCON, etc.). Therefore, when selecting and implementing a payback model/option, the cost of additional or replacement storage may need to be separated from any new SAN infrastructure costs. In other words, most ROI cases are prepared from the perspective of justifying the new SAN infrastructure, not necessarily the attached storage.

If the IT department needs to migrate to SAN via PtP and small fabric installations (ie. FC-AL or small switch), then the future price of SAN infrastructure can be normalized to NPV cost or the average year costs. Future year, incremental growth will need to factor price erosion of future SAN components. If an incremental approach to SAN is planned, it is often best to show total costs per usable SAN port as a comparison in the phased implementation. Otherwise, a separate ROI calculation would be needed for each SAN development phase.

Step 3: Selecting SAN ROI Cases

HDS has compiled several cases from which clients can choose when performing a SAN ROI exercise. Not all cases are equal in value and applicability to the target environment. Some cases are more specific to very large or mature IT infrastructures. Others are best suited for small sites, XSP, decentralized sites, etc. Those conducting the analysis should be familiar with the following options in order to choose those that best apply to the target environment. HDS has found that most IT environments choose between five and eight ROI cases in order to build a business plan/ justification. Each case has specific characteristics that may vary, depending on the payback methods needed in the analysis, such as:

- Immediate cost reduction
- Long-term TCO reduction
- Operational efficiencies
- Growth attainment
- Achieving SLA targets

In the interest of brevity, only the most popular or widely used of the 29 models will be outlined in detail in this paper.

Case 1—Increased Disk Utilization

Pooled storage inherently will provide improved storage asset utilization. DAS storage can vary by platform, as shown below:

- Windows NT—20 to 40 percent average disk utilization
- UNIX—30 to 40 percent average disk utilization
- S/390—60 to 70 percent average utilization

The exercise to determine these savings (usually nominal, hard) is to determine the white space or un-utilized disk capacity aggregated across the enterprise, and apply the average $/GB (gigabyte) paid for the disk. This cost becomes the sunk cost or wasted cost of un-utilized disk capacity. By extrapolating the existing trend (low utilization x average $/GB), future year wastage (storage) and expenditures can be determined.

Comparatively, by adopting a pooled storage architecture, the average utilization rates will improve (60 to 90 percent, or whatever the target utilization rate needs to be), and future growth can be assessed showing the new level of utilized storage. The delta savings are highlighted as offsetting part of the SAN investment. Even though enterprise storage arrays are higher priced, the savings are usually very good in medium to large enterprise storage environments.

Baseline data for this ROI condition must include:

- Cost per GB today
- Projected cost per GB for future years
- Storage utilization rate today
- Projected storage growth and utilization for future years

Case 2—Deferring Disk Procurement

As described in Case 1, a pooled storage architecture provides better asset utilization and better efficiencies in procurement. However, regardless of how well the first consolidation effort goes, future storage needs will emerge. With larger, aggregated storage purchases possible (see also, Case 13—Storage on Demand), the future price of any disk bought will be cheaper per GB than storage bought today. Since storage will be cheaper in the future, deferring as much capacity need fulfillment to the future as possible will have positive impact on near-term cash flow. Any improvement that enables the utilization of "white space" capacity that exists in current disk sub-systems will allow more capacity to be deferred to out-years.

Comparisons need to be made regarding storage-on-demand plans and negotiating a long-term disk price contract. Trends in spending can be identified on just-in-time needs vs. having a large reserve on-hand for unexpected growth. Single spindle disk density doubles roughly every 15 months. As the price-per-GB drops, the capacity per frame will continue to increase.

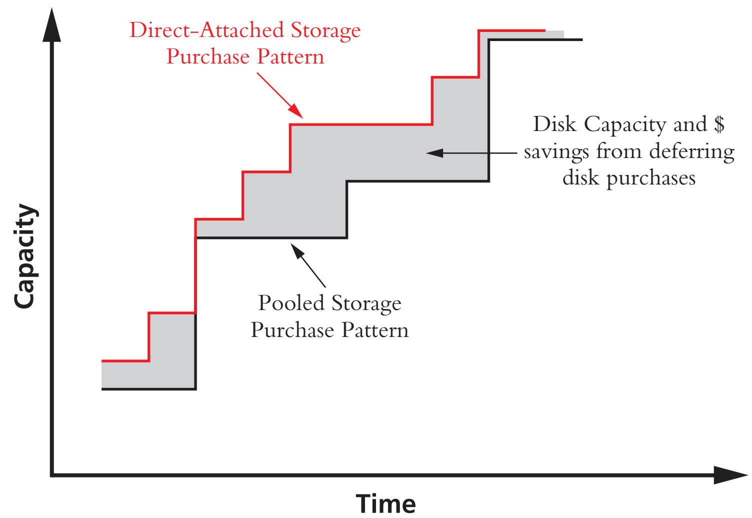

Figure 2: Purchase Deferral and the Effect on Procurement Costs.

Baseline data for this ROI condition must include:

- Rate of price erosion in the storage segment
- Rate of growth for each class of server in the enterprise
- Asset utilization projections before and after the SAN (see Case 1)

Case 3—Terabytes (TB)-per-Administrator Improvement

As the following table shows, pooled storage increases the total capacity of data management that can be assigned to a DBA or storage administrator. Implementation of central, pooled, or SAN strategies can reduce the staff growth requirements in comparison to staff needed to support server-attached, distributed storage.

Topology	GB managed per person
Distributed server-attached	200
Co-located, but server-dependent	400
Use of enterprise storage (PtP)	1,600
1st implemented SAN (medium fabric)	2,000
SAN phase 2 (typically after 6-8 months)	4,000
SAN phase 3 (typically after 8-18 months)	8,000
Target capacity in mature SANs	15-30,000

Table 2: Capacity Managed Per Person.

There are many other benefits of centralized storage, since this architecture also tends to provide uniform methods for change control, version control, and capacity planning, etc. With large enterprise storage pools and a mature management team, follow-the-sun operational coverage (with common processes) is also widespread. Many SAN-enabled enterprises show a good TB-per-person improvement with pooled storage and a SAN. Those organizations that choose not to reduce staff numbers may re-apply staff effort saved by this improvement elsewhere in the enterprise. These actions usually improve employee satisfaction and may lead to additional diversity in an IT career plan.

Baseline data for this ROI condition must include:

- Labor cost for DBA or storage administrator per year (fully burdened labor rate)
- Current storage-capacity-per-person of managed storage
- Rate of storage growth over multiple years

Case 4—Reduce Data Center Rack/Floor Space

Server and disk consolidation represents real savings in data center floor space and environmental costs to the IT department. Through aggressive consolidation and pooling of resources, less space and power is needed to service data storage needs. These costs are true hard-dollar savings, and can start providing payback almost immediately. Many IT shops are running out of space, so the cost avoidance of building a new data center can be factored into the SAN infrastructure investment.

Space savings many not be at once obvious with fewer larger servers and storage arrays, but detailed rack inventory and server/storage inventory will yield real space savings for medium-large data centers.

A base cost of $63 per square foot is an average cost that includes environmental [power and heating, ventilation, and air conditioning (HVAC)] costs to operate a raised floor data center. This may be used for calculating footprint reduction achieved with use of larger storage arrays. Some space (racks) will be needed for FC infrastructure, and this may need to be added back to offset storage savings. Not only should storage array volumes be calculated, but also the resulting drop in number of tape libraries, the reduced number of backup servers, and the fewer number of total servers (due to clustering, etc.) should be included in the calculations. A SAN will provide a radical change to the operating infrastructure, both positive and negative, in terms of data center floor space.

Baseline data for this ROI condition must include:

- Total square footage (and environmentals) for current and projected DAS, server attached storage (SAS), and servers
- Rate of growth in servers, storage, storage arrays, racks, cost of electricity, HVAC, etc.
- Projected space requirements (environmentals) for the SAN infrastructure
- Cost per square foot in the data center

Case 5—Deferring Tape Library Procurement

If the total disk capacity is smaller and the efficiencies of pooled storage work for the IT environment, then the rationale here is that fewer tape devices will be needed to back up the enterprise data. Some local drives that already exist can be used for local backups of application data, OS, etc., but larger enterprise tape libraries will be used to move off-line data from the central storage pool. Many clients have multiple tape stackers, local digital linear tape (DLT), and 8mm drives attached to servers. Many of these can be replaced in favor of a better central library system. General estimates put a 15-to-25 percent tape unit savings through consolidation to larger silos. Efficiencies of 50-to-100 percent also have been observed, depending on the decentralized state of the servers and the quantity of wasted tape space. Fewer tapes will be needed, since better tape utilization will be evident with pooled resources. Less time, handling, and moving of tapes to off-site locations can also be calculated.

Additional savings, although hard to quantify, will show that an enterprise approach to data management, libraries, and information protection reduces the need to maintain multiple tape libraries and catalogs, and thus decreases demand for staff support time. Since most tape library growth is colinear with storage growth, future projections of storage and tape can be correlated. See Figure 3 below.

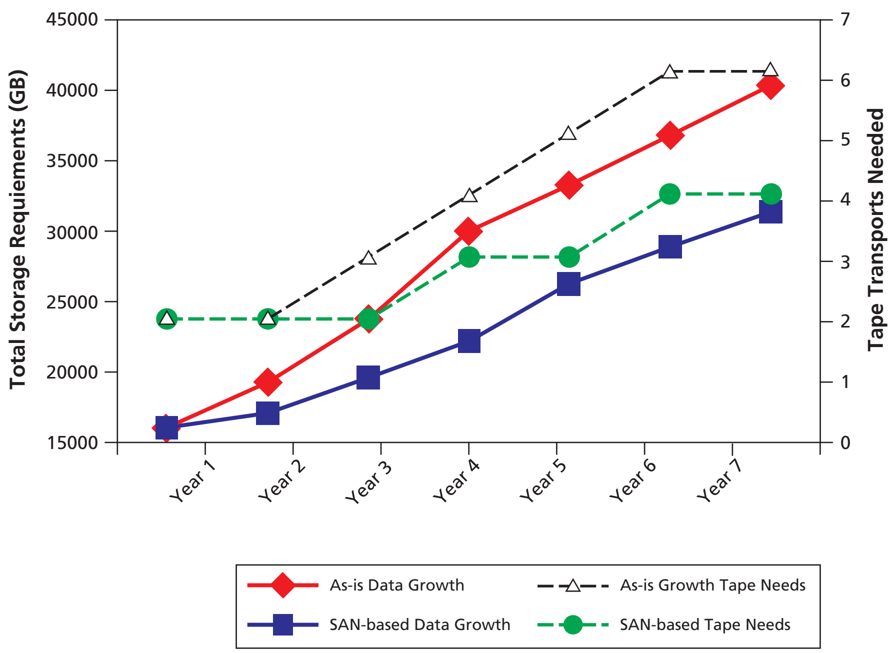

Figure 3: Representative Tape Growth Compared to Storage Requirements.

Baseline data for this ROI condition must include:

- Future year procurement cost of large tape libraries
- Cost of procuring small, single-tape drives, stackers, and small silos
- Cost of tape cartridges
- Tape cartridge transportation costs
- Maintenance costs for multiple tapes
- Staff costs (often hidden) for small server tape handling, management, loading, etc.

Case 6—New Disaster Recovery (DR) Capabilities

The separation of storage from the servers provides an opportunity to manage data and processing separately, and also to plan for recovery with better optimization for data. Data replication can be applied to storage (at the controller level, switch level, or OS level) to create multiple copies of critical data in other parts of the SAN. Since specific datasets and relational database management systems (RDBMS) can be targeted according to the value, appropriate levels of protection can be applied.

Most data center managers realize that catastrophic data loss is less likely to come from natural disasters than from more common events. Having plans in place for the high probability events can be accomplished with target data planning. Here are some additional statistics, which, when applied to Windows NT and UNIX (that tend to have less-than-stellar DR plans in place) provide additional rational for addressing these needs:

- Most businesses experience two hours of downtime per week.
- Most companies set the value of 100 megabytes of data at more than $1 million.
- Forty-three percent of lost or stolen data is valued at $5 million.
- Forty-three percent of companies that have experienced disasters never reopen, and 29 percent close within two years.
- It is estimated that 1 out of 500 data centers will have a severe disaster each year.
- Forty percent of respondents to a computer security survey had detected and verified incidents of computer crime during the previous year.
- Computer crimes cost firms that detect and verify incidents of computer crime between $145 million and $730 million each year.
- Companies that experience a computer outage lasting more than 10 days will never fully recover financially. Fifty percent will be out of business within five years.

It is important to note that SANs do not, within their own architecture, provide DR services and capabilities. However, they can be powerful enablers to plans that already exist. If plans do not currently exist for certain servers or data elements, SANs can provide a starting point to consider remote separation and recovery options as a feature of the SAN design. In performing ROI and payback, some clients may not be performing routine backups for remote Windows NT or UNIX servers. The addition of basic backup and recovery capabilities can use some of the same logic (as outlined in the backup case models), in calculating savings from adding DR capabilities.

Baseline data for this ROI condition must include:

- The IT data center has to be able to determine the business loss or opportunity loss of losing some or part of their data.
- Since the probability of data loss is not consistent across all platforms, an aggregate value needs to be placed on key servers, applications, or datasets that currently are not protected in any DR plan.
- Short-term cost savings, gained by implementing a SAN to achieve fast protection options of critical data (local/remote replication, data movement, etc.), must be determined.

Case 7—On-line Recoverability Options

Gartner Inc. estimates that 70 percent of Fortune 500® companies will implement replication-based backup methods for 10 percent of their data. Moving critical data to near-line or on-line areas for very fast restore will require increased scalability and replication methods that can be found in some SAN topologies.

These calculations are made by determining the downtime impact and associated cost, as well as the reduction of time to restore from on-line. The net time-savings in restoring files or databases can be calculated. Dollarizing the off-line impact or cost and then multiplying by the 99.xxx percent factors (see Table 3) will calculate the yearly off-line costs. Multiplying by the xxx will approximate the incident rate and cost per year. This is the baseline cost. Implementing the SAN may improve the 99.xxx percent availability factor, and by running the same calculations, the delta savings can be determined.

The cost savings for on-line recovery must be balanced with cost of duplicate storage for on-line datasets. Some clients will have multiple instances of replicated data (up to five point-in-time snapshots, for example) to rotate throughout the day. Using snapshots means the client will experience no more than one to two hours of lost data in the event of a major downtime incident. However, the cost of providing five times the disk space should be enough of a deterrent that only key applications will be considered for these techniques.

Baseline data for this ROI condition must include:

- Probability or rate of outages and frequency of outages
- Business cost or opportunity cost for one hour of data loss
- The projected SAN data path rating

Case 8—Improved Data Path Availability

SAN and advanced storage options can increase value of availability to the business as well as IT operations costs. The improvement value of having data more available to the enterprise can be factored into a business case for SANs.

The type of topology selected determines the Quality of Connection (QoC) rating, which in turn determines the data path availability. Different topology selections are shown below.

QoC Level	Fault Tolerance	Architecture	Connection Availability
N/A	N/A	Distributed JBOD	90-96%
N/A	N/A	Centralized, RAID	96-97%
N/A	N/A	Clustered and RAID	99%
1	Failure Sensitive	PtP or Loop	99%
2	Failure Resistant	Loop or Simple Switch	99%
3	Failure Resilient	Switched Fabric	99.9%
4	Failure Tolerant	Multi-stage Fabric or Directors	99.99%
5	Fault Tolerant	Dual Directors	99.999%

Table 3: QoC Measurement by Topology and Percent Connection.

Business availability may be hard to quantify, but in applications like e-commerce and decision support systems (DSS), not having critical data could be catastrophic. Samples of data loss per outage compared to business cost, are shown below.

Description	Cost Per Hour of Outage
Brokerage Operations	$6,450,000
Credit Card Sales Authorization	$2,600,000
ERP	$780,000
Supply Chain Management	$660,000
Electronic Commerce	$600,000
Internet Banking	$420,000
Universal Personal Services	$360,000
Customer Service Center	$222,000
Point of Sales/EFT	$210,000
Pay per View	$150,000
Home Shopping Network	$113,750
Average Cost (Survey of 400 Companies)	$100,000
Home Catalog Sales	$90,000
Airline Reservation System	$89,500
Tele-ticket Sales	$69,000
Messaging	$60,000
Cellular Service Activation	$41,000
Package Shipping Service	$28,250
On-line Network Connection Fees	$25,250
ATM Service Fees	$14,500
Standish Group Survey (Average)	$11,240
Yankee Group Survey (Average)	$1,000

Table 4: Business Cost of an Outage.

The formula for determining the data path outage calculation is:

Value of Business Hour ($) x (Access hours-per-day x 365) x (1-Availabilty Rating)

The difference of the before and after SAN topology decision can be calculated and shown as a year-to-year cost avoidance.

Baseline data for this ROI condition must include:

- Business cost or opportunity cost for one hour of data loss
- Calculations to determine what part or percentage the storage data path (FC, NAS, IP, etc.) plays in the overall outage rating of servers and applications
- The current data path uptime rating
- The projected SAN data path rating

Case 9—Reduction of General-purpose UNIX® and Microsoft® Windows NT® Servers

Much of an IT department's Windows NT growth can be attributed to servers (up to 10 percent) that are installed for disk capacity only, not CPU cycles. This phenomenon is not as prevalent in UNIX environments. As server growth continues, we can expect to see situations in which servers are acquired for the sake of storage only. The assumption for this case is that some percentage of Windows NT servers will not need to be procured thanks to centralization that is enabled with the SAN. This projected number of servers that could be avoided is multiplied by the current or projected cost per server and shown as hard-dollar savings.

There are several hidden costs associated with acquiring and maintaining servers, since the cost of deploying net-new servers is significantly greater than the cost of the server alone. Some of the life cycle costs for servers are categorized below:

Acquisition	Operation	Disposition
Configuration	Electricity	Data Backup
Requisitioning	Cooling	Drive Reformatting
Procurement	Network	Packing
Delivery	Floor Space	Software License Transfer
Cabling	Software Maintenance	
Electrical Wiring	Hardware Maintenance	
Receiving	Asset Tracking	
Installation	Insurance	
Invoice Processing		

Table 5: Hidden Life Cycle Costs of Servers.

Baseline data for this ROI condition must include:

- The number of servers that are used for NFS of CIFS mounted storage only and the rate of growth of these types of servers
- Cost to acquire and TCO cost of these types of servers

Case 10—Improve LAN/WAN Performance, Avoid Upgrades

Network infrastructure will need constant upgrades (about every year) if data movement, backup, and replication is migrated to the IP backbone (assuming incremental data growth). New FC SAN topologies off-load the traffic to the new storage infrastructure, thus relieving the network load and potentially deferring upgrades. Twenty percent of LAN traffic can be attributed to backup processes occurring over the net. From projections on server growth, determine how many new sub-nets or virtual private networks (VPNs) will be required to support server-to-server transfers and network-based backup. The costs of new sub-nets (NIC, hubs, and routers) can be determined from the projected growth, and some portion of these costs can be counted as savings when avoided in favor of FC-based topologies.

This payback is clearly in favor of FC-based topologies. New protocols (iSCSI, CxFS) and gig-Ethernet may drive some options to keep the storage data path on the IP backbone.

Baseline data for this ROI condition must include:

- Itemization of all components that make up the $/node for the client's LAN (routers, hubs, software, management manpower, cabling, etc.). These figures need to include building new sub-nets or VPNs, as well as simple upgrades.

Case 11—Reduce/Eliminate Backup Servers

In addition to, or as part of the analysis from Case 9, there may be a specific instance in which to reduce dedicated backup servers that are servicing small sub-nets or groups of servers. The role of dedicated backup servers is likely to continue within advanced storage architectures, but the need for numerous small servers is diminished. This function will determine how the need for small, dedicated backup servers can be reduced or prevented by using larger pooled storage with tape connections to silos and bypassing the need for host processing (LAN-less backup).

If small backup servers are in use at the current time, then estimations are required on how these will grow (dedicated backup server: TB or server ratios) based on server and storage growth. Included in this category are:

- Cost of server hardware and OS
- Backup software
- Small or local tape systems
- Off-site media management costs
- Local operation costs (manpower)

As part of the SAN deployment, the client may invoke OS/390®-based backup processes to utilize large silos, media management, or off-site tape management. These costs may incrementally increase due to potential growth of SAN-attached storage and data. These costs may be offsetting or determined as part of the TCO for the SAN. A staff element also can be assigned with backup server consolidation, with a sample target of .5 FTE per TB of backed-up data. There may be little or no opportunity to remove existing backup servers. Additionally, their role may be reduced over time and the dedicated tape stackers or silos either added to the SAN infrastructure or written off after being depreciated.

Baseline data for this ROI condition must include:

- Procurement cost and TCO of distributed backup servers and rate of growth of these servers
- Cost of labor for systems backup

Case 12—Reduce/Eliminate Batch, Backup Windows

Backup windows are continuously under pressure to be reduced to save time. Over the next one to two years, most clients will find that the rate of growth with UNIX and Windows NT data will make current windows impossible to meet. Streamlining with new tapes will only provide stopgap measures.

IT departments are exploiting pooled storage benefits with backup improvements. The centralization of data to larger, fewer storage systems enables a more productive backup approach. Snapshots of data can be made, as well as mirroring for critical data sets (see also Case 7—On-line Recovery Options). The ability to break mirrors or back up snapshot copies of data enables servers and applications to be on-line more, and alleviate some shrinking backup window issues. FC-attached disks and tapes can provide better streaming of data to near-line devices. Many organizations are exploiting low cost disks to keep on-line versions of data for restore requirements. Server-less backups and LAN-free backups all contribute to many significant savings in reducing backup errors.

Some of the types of cost savings and areas for improvement that can be captured/modeled include:

- Business or operational cost to take servers off-line for backups
- Business risk if backups are not completed or there is an abend (abnormal end of task)
- Cost of off-line time while data is being restored (see Table 4)

Baseline data for this ROI condition must include:

- Business impact on a failed backup
- Business impact to have data off-line for backup processes

Case 13—Storage on Demand

Many advanced SAN architectures allow for highly scalable growth, both in raw capacity and in connections to the storage/tape pools. This agility in providing capacity on demand often has associated critical business benefits, since a protracted process is not required to acquire, install, configure, and apply management rules to new capacity. (See Figure 4: DAS versus SAN Deployment Time.) The concept of the IT department providing a storage "utility" becomes real, with the opportunity to manage costs (chargeback), service level commitments, and operational effectiveness without having to turn to outside services.

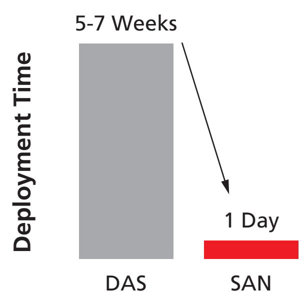

Figure 4: DAS versus SAN Deployment Time.

Baseline data for this ROI condition must include:

- Time and costs involved with storage procurement and any associated opportunity costs due to lack of fast access to storage capacity

The Remaining Case Models

Sufficient space is not available in this paper to outline the remaining SAN ROI cases, but they are listed here for reference purposes.

14. ***Improved Protection of Critical Data***—Some SAN architectures can provide pin-point protection of specific data, either on one server or spread throughout the enterprise. LUN management, RAID, replication, mirroring, snapshots, encryption, and lock management all have varying degrees of security to datasets in the storage pool. Rather than apply large-scale processes to an entire server or disk array, protection and security schemes can be applied at the micro level. This saves in resources and effort, in that applicable levels of protection are applied only where needed.

15. ***Management Costs as a Percentage of Storage Costs***—This is determined by general calculations that demonstrate a lower percentage of the TCO for SAN-attached architectures when compared to direct-attached architectures. As an example, disk management cost is 60 percent of the storage costs for distributed, server-attached storage, while it is estimated at nine percent for SAN.

16. ***Increased I/O Performance and Bulk Data Movement***—An enterprise storage infrastructure with robust data movement software potentially can remove some operation barriers that exist with batch jobs, data loads, or streaming data functions. The overall system improvement may or may not be enhanced by the nature of a SAN. More typically, better performance occurs thanks to improvements in storage subsystems and access methods, etc.

17. ***Reduced Storage Maintenance***—With fewer and larger disk systems being used in the enterprise, there is a potential for a net decrease in hardware and software maintenance costs as the storage is consolidated.

18. ***Staff Utilization for Server Management***—When SANs can reduce the number of servers (infrastructure, storage servers, backup servers, etc.), the resulting impact to server managers can be realized.

19. ***Increased Life of Current Storage***—The useful life of Windows NT servers is three years. Therefore, the useful life of server-attached disk on the Windows NT servers is three years. UNIX servers may last a little longer, but not much. This model calculates the extended useful life of central/pooled storage—up to five or six years—where reasonable.

20. ***Improved Disaster Recovery Capabilities***—Moving critical data to near-line or on-line areas for very fast restore will require the increased scalability and replication methods that can be found in some SAN topologies. The downtime impact must have an associated cost, and thus, reducing the time to restore from on-line, replicated images can be shown to have a positive business value impact.

21. ***Non-disruptive Scalability***—With traditional small server data growth, there are interruptions in service as new servers and disks are added to the network. This interruption has some business impact if it causes the data to be off-line during business hours. Also, to make the transition, manpower is required (usually off-shift), involving DBAs, system managers, technical staff, field service, etc. Some FC topologies provide for non-disruptive maintenance and growth.

22. ***Avoid Data Area Network Growth***—In some NAS and SAN architectures, a new FC SAN may offset the costs necessary to upgrade or expand the IP network that supports file storage or backups. IP network growth costs can be shifted to FC SAN investments.

23. ***Impact on New or Migrating Applications***—HDS has observed application migration benefits through shortened development time, application data coexistence, reduced network impact, and bulk movement of data to/from dissimilar environments.

24. ***Impact on Applications Development, Testing***—Enterprise storage impacts the development process by supporting rapid data movement from a mainframe to a Windows NT environment. This allows developers to test on current data or copies of live data. In many environments, data refresh requires hours of reload from tapes. Testing can improve with direct access to replicated live data provided through the SAN.

25. ***Extending Life of Servers***—With reduced network overhead, CPU overhead, and storage I/O, useful life determinations of some servers may increase, saving organizations the early cost of replacement.

26. ***Reduce CPU Load on Servers***—Some new generation host bus adapter cards can reduce the CPU workload for I/O and significantly improve overall server performance.

27. ***Support Server Clustering***—SAN installation or storage consolidation is usually a prelude to server clustering and application consolidation. The benefits of these types of activities can be attributed partially to the SAN infrastructure investment.

28. ***Secondary Security Services***—SAN fabrics inherently can provide another layer of data path protection and security through hard and soft zoning. Although not fully certified by government security rating agencies, fabric zoning can augment network traffic partitions, authentication/authorization, and ACL systems to provide further data isolation and segregation.

29. ***Vendor Consolidation***—Pooled storage architectures can consolidate the number of vendors involved in providing infrastructure services. This reduced overhead impact can be measured and summarized.

Steps 4 and 5—Summary Charts and Calculations

In any ROI analysis, the results can be very subjective. There are no guarantees to the models, parameters, assumptions, and payback performances outlined in this paper. We remind readers that SAN technology should not be considered on cost savings or strong ROI expectations alone. SAN is a business-enabling technology, which, when planned and deployed correctly, can have a tremendous impact on data management, growth, performance, and investment protection. There are opportunities for SANs to impact IT operations and expenditures financially in the future, but these should be secondary to SANs' capabilities to help companies achieve business and operational goals.

Trade-offs on values and parameters (conservative to aggressive) can be made within the spreadsheets, and the calculations run again, to determine the best ROI results. The overall approach is to summarize savings by year and by category. Separation of soft-dollar savings and hard-dollar savings is recommended.

Hard Dollar Only ROI						
A. Net Investment Required						$ 1,122,100
B. Yearly Cash Flows						$ 1,483,132
			Years			
	2000	2001	2002	2003	2004	
Windows NT Hard Savings	231,068	392,984	399,548	448,167	423,879	
Novell® Hard Savings	13,232	17,327	21,081	22,761	25,950	
UNIX Hard Savings	137,612	72,861	73,412	82,513	82,513	
MVS® Hard Savings	35,000	49,000	76,832	98,345	123,915	
Other Cost Reduction	35,910	38,880	32,400	32,940	31,860	
Recurring Expense	(299,120)	(325,044)	(307,764)	(301,404)	(283,524)	
Net Cash Flow	153,701	246,008	295,509	383,321	404,592	$ 1,483,132
C. Simple ROI determined from B / # of years / A						26.43%

Figure 5: Sample Hard Dollar ROI Summary.

Hard & Soft Dollar ROI						
A. Net Investment Required						$ 1,122,100
B. Yearly Cash Flows						$ 7,584,228
			Years			
	2000	2001	2002	2003	2004	
Windows NT Hard Savings	231,068	392,984	399,548	448,167	423,879	
Novell Hard Savings	13,232	17,327	21,081	22,761	25,950	
UNIX Hard Savings	137,612	72,861	73,412	82,513	82,513	
MVS Hard Savings	35,000	49,000	76,832	98,345	123,915	
Soft Dollar Benefits	35,000	283,309	322,952	369,322	415,433	
Other Cost Reduction	245,533	38,880	32,400	32,940	31,860	
Recurring Expense	(299,120)	(325,044)	(307,764)	(301,404)	(283,524)	
Net Cash Flow	1,061,677	1,344,798	1,512,286	1,773,752	1,891,715	$ 7,584,228
C. Simple ROI determined from B / # of years / A						135.18%

Figure 6: Sample Hard and Soft Dollar ROI Summary.

The formula to calculate simple ROI is as follows:

Total projected savings / number of years / total net investment (NPV_0).

Charts 1-3 below are samples that can be created from the ROI worksheets built for a SAN payback analysis.

	Years				
	Year 0	Year 1	Year 2	Year 3	Year 4
Purchase, Recurring Costs	(1,131,618)	(295,150)	(304,338)	(313,984)	(293,114)
Hard Savings		152,017	273,408	309,939	289,269
Soft Savings		393,279	414,597	444,705	469,107
Net	**$ (1,131,618)**	**250,146**	**383,668**	**440,659**	**465,262**

Chart 1: Cost Savings by Year, Numeric Chart.

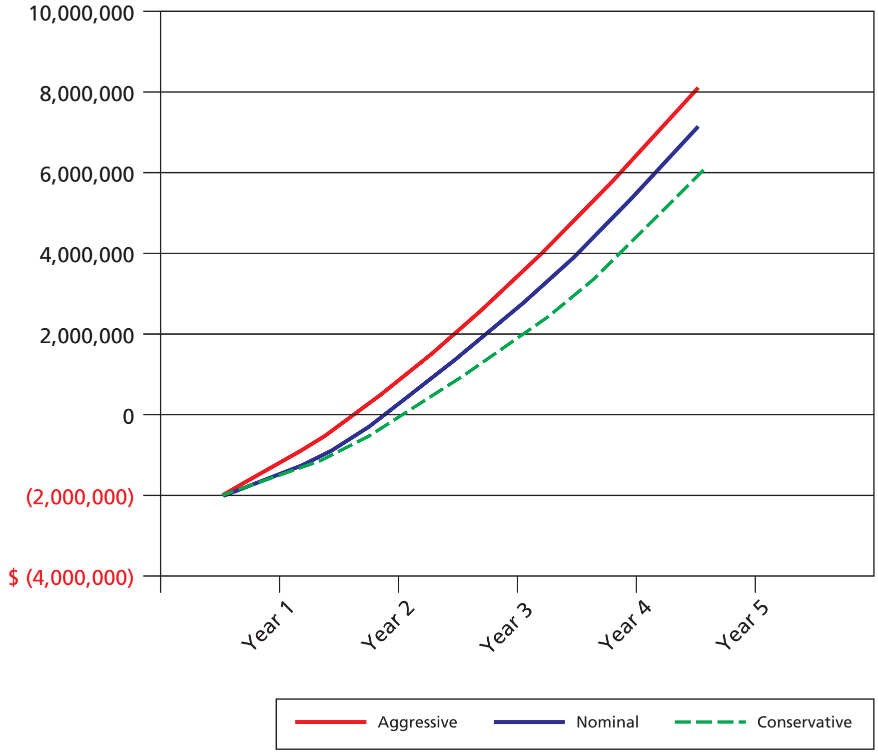

Chart 2: SAN Payback with Aggressive, Nominal, and Conservative Parameters.

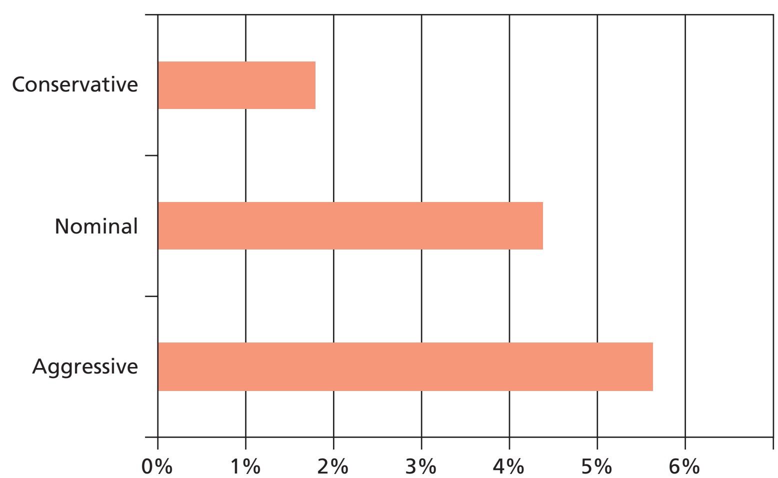

Chart 3: Internal Rate of Return.

ROI Impact on a SAN Topology Decision

Our research into ROI impacting a SAN topology decision, although not exhaustive, concludes that early cost and TCO planning in the design work is a valuable process to employ when choosing and implementing a long-term pooled storage strategy. ROI is not the only determinant in topology and design decisions; there are dozens of factors that need to be included in any logical or physical pooled storage design. Some of the other technical and business considerations may include:

- Geographic considerations of the servers, storage, and staff
- Performance requirements
- Availability needs for the data
- SAN infrastructure serviceability and maintainability
- Scalability and non-disruptive upgrades
- Improvements in data restoration
- Chargeback costs to storage users

SAN ROI and TCO are important elements to consider in parallel with other design qualities. HDS has found that each topology has differing cost and payback characteristics that need to be explored before the design is committed to production. A complete list of topologies compared to each of the 29 SAN ROI payback models is not possible at this time, but a summary of a few popular topologies and payback areas are summarized in Table 6.

Note: Not all SAN topologies and payback considerations are shown. Table 6 is intended to demonstrate the differing characteristics of various topologies to the ROI cases.

ROI Case	NAS	FC PtP	FC Switch Fabric	FC Directors
Increased Disk Utilization	Yes, for file system structures.	Yes, for block system structures.	Yes, for block system structures.	Yes, for block system structures.
New DR Capabilities	Can be achieved depending on the file system, replication s/w. Data can be replicated over long distances.	Yes, disk-based replication techniques, typically within 10km distances.	Yes, disk, switch, host software replication options through different fabrics, on different array locations.	Yes, but may still be some limit in total fabric size and distance.
Staff-per-TB Management Ratio	NAS tends to have 1-2 TB limits per instance. Large pools of centralized storage are not achievable, negating most of the effects of this category.	Yes, with larger enterprise storage arrays.	Yes, with larger enterprise storage arrays.	Yes, with larger enterprise storage arrays.
Improved Availability	Although better than SCSI, dependence on the network availability has to be factored. Overall ratings are in the 90-95% range.	Better than SCSI.	Yes, as long as the right switch topology (Star, Mesh, Cross Connect) is utilized.	Most reliable topology for high availabily in the storage data path. Ranges of 99.9 – 99.999% are achievable with the right configurations.
Non-disruptive Scalability	Yes.	Difficult to achieve when approaching the maximum ports per storage array.	Best case, since the fabric switches can be configured to support non-disruptive upgrades.	Can be achieved, but at a high cost.
Reduce Backup Windows	Not typical, since IP network may be the weak link.	Improvement will be seen through tape spooling.	Yes, if fabric attached tapes are provisioned.	Yes, if fabric-attached tapes are provided.
Improved Performance	Typical IP network throughput is not as good as FC. Some degradation can be expected.	This should be observed since this topology is well-suited for high performance I/O situations.	Possible if the fabric switch mesh complexity is kept to no more than 2-3 hops. Otherwise, new problems will emerge in switch to switch latency.	In a simple 1 or 2 director flat mesh, performance is very good.
Improve LAN/WAN Performance	This will not happen if the NAS traffic is on the production LAN/WAN.	This may be observed.	This may be observed.	This may be observed.

Table 6: ROI Items Charasterisitc in Different Pooled Architecture Topologies.

Glossary of Terms

Acronym	Definition
DAS	Direct attached storage
DB	Database
DBA	Database Administrator
DR	Disaster Recovery
DLT	Digital Linear Tape
FC	Fibre Channel
HBA	Host Bus Adapter
HVAC	Heating, ventilation, and air conditioning
IRR	Internal Rate of Return
LUN	Logical Unit Number
MTBF	Mean Time Between Failures
NAS	Network Attached Storage
NFS	Network File System
NIC	Network Interface Card
NPV	Net Present Value
Point to Point (PtP)	Fibre Channel topology with a simple direct connection between two nodes
QoC	Quality of Connection
RAS	Reliability, Availability, Serviceability
RDBMS	Relational Database Management System
ROI	Return on Investment
ROM	Rough order of magnitude
SAN	Storage Area Network
SAS	Server Attached Storage
SLA	Service Level Agreement
TCO	Total cost of ownership
Windows NT	Windows NT operating system

References

Contingency Planning Research and Strategic Research Corporation

Dataquest/Gartner Research, "The Storage Utility Market: Strap In and Hold Tight!" May 8, 2000

Gartner Group, Data Storage, 1999.

Gartner Group, Network Business Management Research Note, May 1, 1998

Gartner Operations Benchmark Report

HDS Synthetic Storage Architecture Template, 2001

Hitachi Data Systems, Business Continuity Service Description, 1999

Hitachi Data Systems, SAN EIA Final Report, 1999

Hitachi Data Systems, USAG ROI Executive Summary, 1999

Information Economics, Linking Business Performance to Information Technology (Parker & Benson), 1988

International Data Corporation, 1998 Worldwide Disk Systems Market Forecast and Review by Robert C. Gray

ITcentrix, Inc., September 1999

ITcentrix, Inc., Storage Networks: On the Road to Virtual Consolidation, 1999, Page 46

Meta Stepped Consolidation Strategy, 2000

Ontrack Data International

Strategic Research Corp., September 1998

Strategic Research Corp., SAN: Quality of Connection

Hitachi Data Systems

www.hds.com

storage@hds.com

Corporate Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1000
info@hds.com

Asia Headquarters

Suite 3301-6, Shell Tower
Times Square, 1 Matheson Street
Causeway Bay
Hong Kong
2525-2385
infoasia@hds.com

Australia/New Zealand Headquarters

Level 3
82 Waterloo Rd.
North Ryde, NSW 2113
Australia
+61-2-9325-3300
info.australia@hds.com

Canada Headquarters

2550 Victoria Park Avenue
Suite 601
Toronto, Ontario M2J 5A9
(416) 494-4114
www.hds.com
Cananda.Sales@hds.com

Europe Headquarters

Sefton Park
Stoke Poges
Buckinghamshire SL2 4HD
United Kingdom
+44 (0) 1753-618000
info.eu@hds.com
www.eu.hds.com

Latin America Headquarters

750 Central Expressway, MS 3468
Santa Clara, California 95050-2627
U.S.A.
(408) 970-7447
infolatin@hds.com

U.S. Headquarters

750 Central Expressway
Santa Clara, California 95050-2627
U.S.A.
(408) 970-1001
ussalesinfo@hds.com

WHP-103-00
July 2001