June 2003

MANAGEMENT BRIEF

VALUE PROPOSITION FOR ENTERPRISE STORAGE

Cost/Benefit Case for IBM TotalStorage

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net

TABLE OF CONTENTS

EXECUTIVE SUMMARY 1

General Conclusions 1
Vendor Comparisons 4
Industry Trends 5
Infrastructure Opportunity 6

PROFILE SUMMARIES 8

Diversified Bank 8
Business Profile 8
Business Strategy 8
Storage Scenarios 10
Strategic Benefit 12
Efficient Infrastructure 12
Vendor-specific Costs 14
Manufacturing Company 16
Business View 16
New ERP Strategy 17
Storage Scenarios 19
Strategic Benefit 21
Efficient Infrastructure 22
Vendor-specific Costs 23
Retail Company 25
Business View 25
Business Strategy 25
Storage Scenarios 26
Strategic Benefit 29
Efficient Infrastructure 30
Vendor-specific Costs 31

TECHNOLOGY VIEW 34

State of Storage 34
General Picture 34
Specialization Issues 34
Cost/Benefit Quantification 36
Measuring Efficiency 37
Measurement Challenges 37
Performance Variables 38
Scalability 40
Capacity Utilization 40
Service Quality 41
Real-time Implications 42
Storage Management 42
Enterprise Challenges 42
Vendor Strategies 43
Designed-in Capabilities 44
Autonomic Computing 45

DETAILED DATA 48

Vendor Comparisons 48
Diversified Bank 48
Inefficient Scenario 48
Efficient Scenario 52
Manufacturing Company 57
SAP User Base 57

Inefficient Scenario 57
Efficient Scenario 59
Retail Company 65
Inefficient Scenario 65
Efficient Scenario 69
METHODOLOGY 72
Sources of Data 72
General Survey 72
Composite Profiles 72
Basis of Configurations 73
IT Cost Calculations 74
Cost Components 74
Inefficient Scenarios 74
Efficient Scenarios 75
Other Costs 75

List of Figures

1. Summary of Five-year IT Costs for Inefficient and Efficient Scenarios 2
2. Summary of Five-year Business and IT Costs for Inefficient and Efficient Scenarios 3
3. IT Costs for Vendor-specific Options: Efficient Scenarios 4
4. IT Cost Breakdowns: Midrange Component of Efficient Scenarios 5
5. Key Bank Business Statistics 8
6. Bank Scenarios: Summary Data 11
7. Bank Efficient Scenario: Projected Disk Storage Growth 13
8. Bank Efficient Scenario: Overall Vendor Cost Comparisons 14
9. Bank Efficient Scenario: Midrange Vendor Cost Comparisons 15
10. Manufacturing Company: Improvements in Supply Chain Efficiency 16
11. Manufacturing Company: New Strategy for ERP Solutions 18
12. Manufacturing Company Scenarios: Summary Data 20
13. Manufacturing Efficient Scenario: Projected Disk Storage Growth 22
14. Manufacturing Company Efficient Scenario: Overall Vendor Cost Comparisons 23
15. Manufacturing Company Efficient Scenario: Midrange Vendor Cost Comparisons 24
16. Retail Company Scenarios: Summary Data 28
17. Retail Efficient Scenario: Projected Disk Storage Growth 31
18. Retail Company Efficient Scenario: Staffing Detail 31
19. Retail Efficient Scenario: Overall Cost Comparisons 32
20. Retail Company Efficient Scenario: Midrange Vendor Cost Comparisons 33
21. Representative Storage-related Job Titles 35
22. Representative Efficiency Variables 38
23. Representative ERP Workload Breakdowns 39
24. Storage Management Levels 43
25. Levels of Autonomic Capability 45
26. Storage Applications of Autonomic Computing: Current Design Capabilities 46
27. Storage Applications of Autonomic Computing: Future Design Possibilities 47
28. IT Costs for Vendor-specific Options: Efficient Scenarios 48
29. IT Cost Breakdowns: Midrange Component of Efficient Scenarios 48

30. Bank Mainframe Bases 50
31. Bank Inefficient Scenario: Initial Server Base 50
32. Bank Inefficient Scenario: Initial Disk Storage Base 51
33. Bank Inefficient Scenario: Projected Disk Storage Growth 51
34. Bank Inefficient Scenario: IT Costs Detail 52
35. Bank Inefficient Scenario: Staffing Detail 52
36. Bank Efficient Scenario: Initial Server Base 53
37. Bank Efficient Scenario: Vendor Platforms 53
38. Bank Efficient Scenario: Vendor Software Stacks 54
39. Bank Efficient Scenario: Staffing Detail 54
40. Bank Efficient Scenario: EMC Configurations 55
41. Bank Efficient Scenario: HDS Configurations 55
42. Bank Efficient Scenario: IBM Configurations 56
43. Bank Efficient Scenario: Hewlett-Packard Configurations 56
44. Manufacturing Company: Numbers of Users for Principal SAP Systems 57
45. Manufacturing Company Inefficient Scenario: Initial Server and Disk Storage Bases 58
46. Manufacturing Company Inefficient Scenario: Projected Disk Storage Growth 58
47. Manufacturing Company Inefficient Scenario: IT Costs Detail 59
48. Manufacturing Inefficient Scenario: Staffing Detail 59
49. Manufacturing Company Efficient Scenario: Central Data Center Environments 60
50. Manufacturing Company Efficient Scenario: Vendor Platforms 61
51. Manufacturing Company Efficient Scenario: Vendor Software Stacks 61
52. Manufacturing Company Efficient Scenario: Staffing Detail 62
53. Manufacturing Company Efficient Scenario: EMC Configurations 62
54. Manufacturing Company Efficient Scenario: HDS Configurations 63
55. Manufacturing Company Efficient Scenario: IBM Configurations 63
56. Manufacturing Company Efficient Scenario: Hewlett-Packard Configurations 64
57. Retail Company Inefficient Scenario: Initial Server Base 66
58. Retail Company Inefficient Scenario: Initial Disk Storage Base 67
59. Retail Inefficient Scenario: Projected Disk Storage Growth 68
60. Retail Company Inefficient Scenario: IT Costs Detail 68
61. Retail Company Inefficient Scenario: Staffing Detail 68
62. Retail Company Efficient Scenario: Initial Server Base 69
63. Retail Company Efficient Scenario: Vendor Platforms 69
64. Retail Company Efficient Scenario: Vendor Software Stacks 70
65. Retail Company Efficient Scenario: EMC Configurations 70
66. Retail Company Efficient Scenario: HDS Configurations 71
67. Retail Company Efficient Scenario: IBM Configurations 71
68. Retail Company Efficient Scenario: Hewlett-Packard Configurations 71
69. SAN Deployment Rates in Inefficient Scenarios 74
70. Maintenance and Support Values for Efficient Scenarios 75
71. Salary Assumptions 75

Executive Summary

General Conclusions

At the end of 1996, the typical Fortune 500 corporation contained less than seven terabytes of server-based disk storage. By the end of 2002, this had increased to more than 48 terabytes. On current trends, it will increase to more than 230 terabytes by the end of 2007.

This document is about the implications of this trend. It deals with the challenges that are posed, and with the opportunities that are created by the growth and increasing strategic importance of storage resources. Results are based on data from, and analysis of the experiences of 124 large organizations.

Three general conclusions emerge:

1. ***Business impact***. There is a direct correlation between storage and business performance. Storage resources are the backbone infrastructure for an organization's information. They contribute materially to the effectiveness with which this information is used.

   The critical variables of competitiveness in the 21st century are information-dependent. Flexibility and responsiveness, integration of processes, reduction of cycle times, collaboration and cooperation, cost-effectiveness and profitability, and processes of relationship-building with customers and partners are affected by quality, availability and timeliness of information.

   The organizational storage infrastructure magnifies – or impairs – the impact of all of these variables. Its significance increases as organizations move toward real-time strategies. As three composite company profiles presented in this document demonstrate, real-time competitiveness mandates effective cross-organizational storage infrastructures.

2. ***IT productivity***. There is a direct correlation between the organization of storage resources and the productivity of the IT organization as a whole. An effective storage infrastructure improves the efficiency of a wide range of IT functions, including application development, system and database management, network management, operations, and end-user support.

   By increasing across-the-board IT efficiency, an effective storage infrastructure accelerates application delivery processes. It enables personnel, funds and attention to be focused on the delivery of high-impact business solutions, rather than on the minutiae of implementing, administering and maintaining underlying hardware and software platforms. The business contribution of IT may be significantly improved without additional staff or expenditure.

3. ***Efficiency***. Infrastructure-based approaches to storage offer the potential for breakthrough reductions in IT costs. In the company profiles presented in this document, five-year costs for storage, hardware, software, maintenance, personnel and facilities are reduced between 2.4 and 3.9 times by transitioning from inefficient to efficient scenarios for storage utilization.

   These savings are realized through consolidation, technology upgrades, deployment of storage area networks (SANs) and network attached storage (NAS), and adoption of "best practices" for storage management under new corporate-level structures with organization-wide responsibility for storage resources. Results are summarized in figure 1.

   Each of these initiatives provides significant gains. But "the whole is more than the sum of the parts." Their combined impact is cumulative. An organization-wide, infrastructure-based storage strategy will yield economies that far exceed piecemeal adoption of new techniques and technologies.

Figure 1

Summary of Five-year IT Costs for Inefficient and Efficient Scenarios

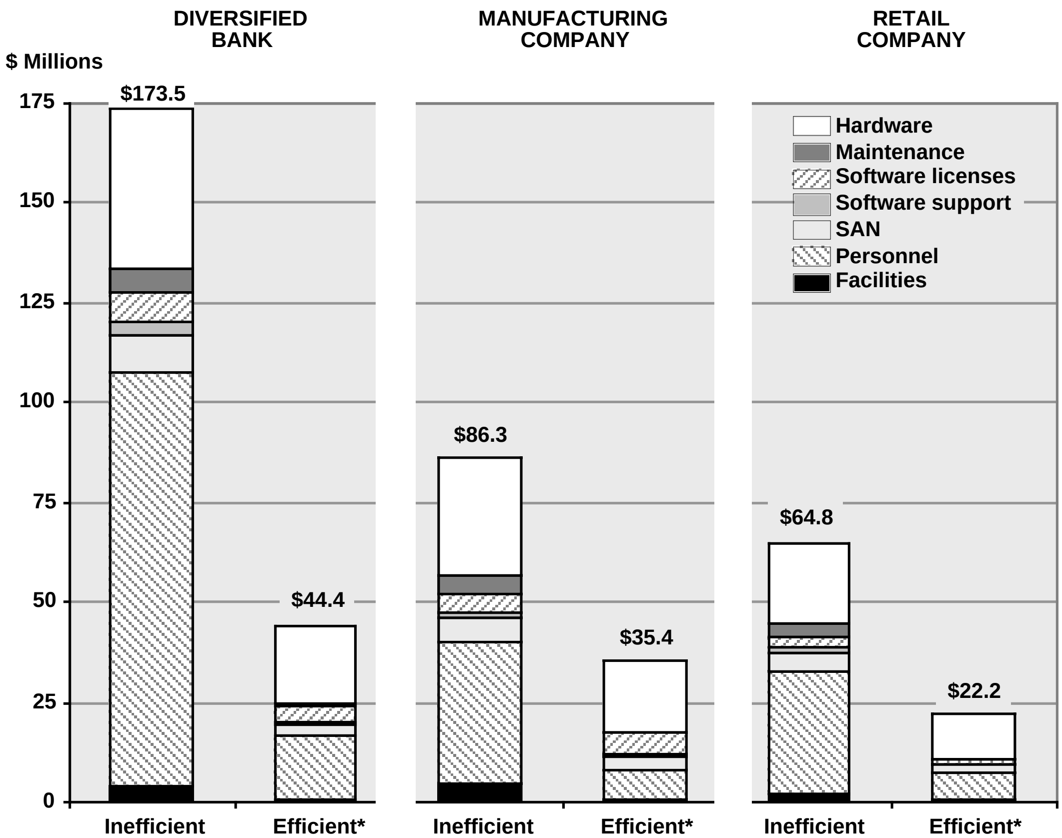

*Lowest cost vendor option (IBM)

At the same time, business costs due to inadequate service quality – including costs of outages, performance shortfalls and limitations in backup and recovery capability – are reduced by wide margins. Figure 2 shows the extent of these for the same three company profiles.

Vulnerability to the effects of severe outages also declines. Recovery times are reduced from days to minutes. For the diversified bank, the bottom-line impact of such outages is reduced from a range of $4.5 to $27.3 million under the inefficient scenario to $1.3 million under its efficient counterpart. For the manufacturing company, the impact is reduced from $2.7 to $66.2 million to less than $300,000. For the retail company, there is a reduction from $27.0 to $155.1 million to less than $600,000.

These are conservative calculations, and are based on storage contributions to overall levels of availability, response time, application-level throughput, and recovery speed and effectiveness. The overall impact of all of these – and corresponding savings in business costs – would be much greater.

The effects of increased overall IT productivity are not included in efficient scenario savings. These would, however, be substantial.

Figure 2

Summary of Five-year Business and IT Costs for Inefficient and Efficient Scenarios

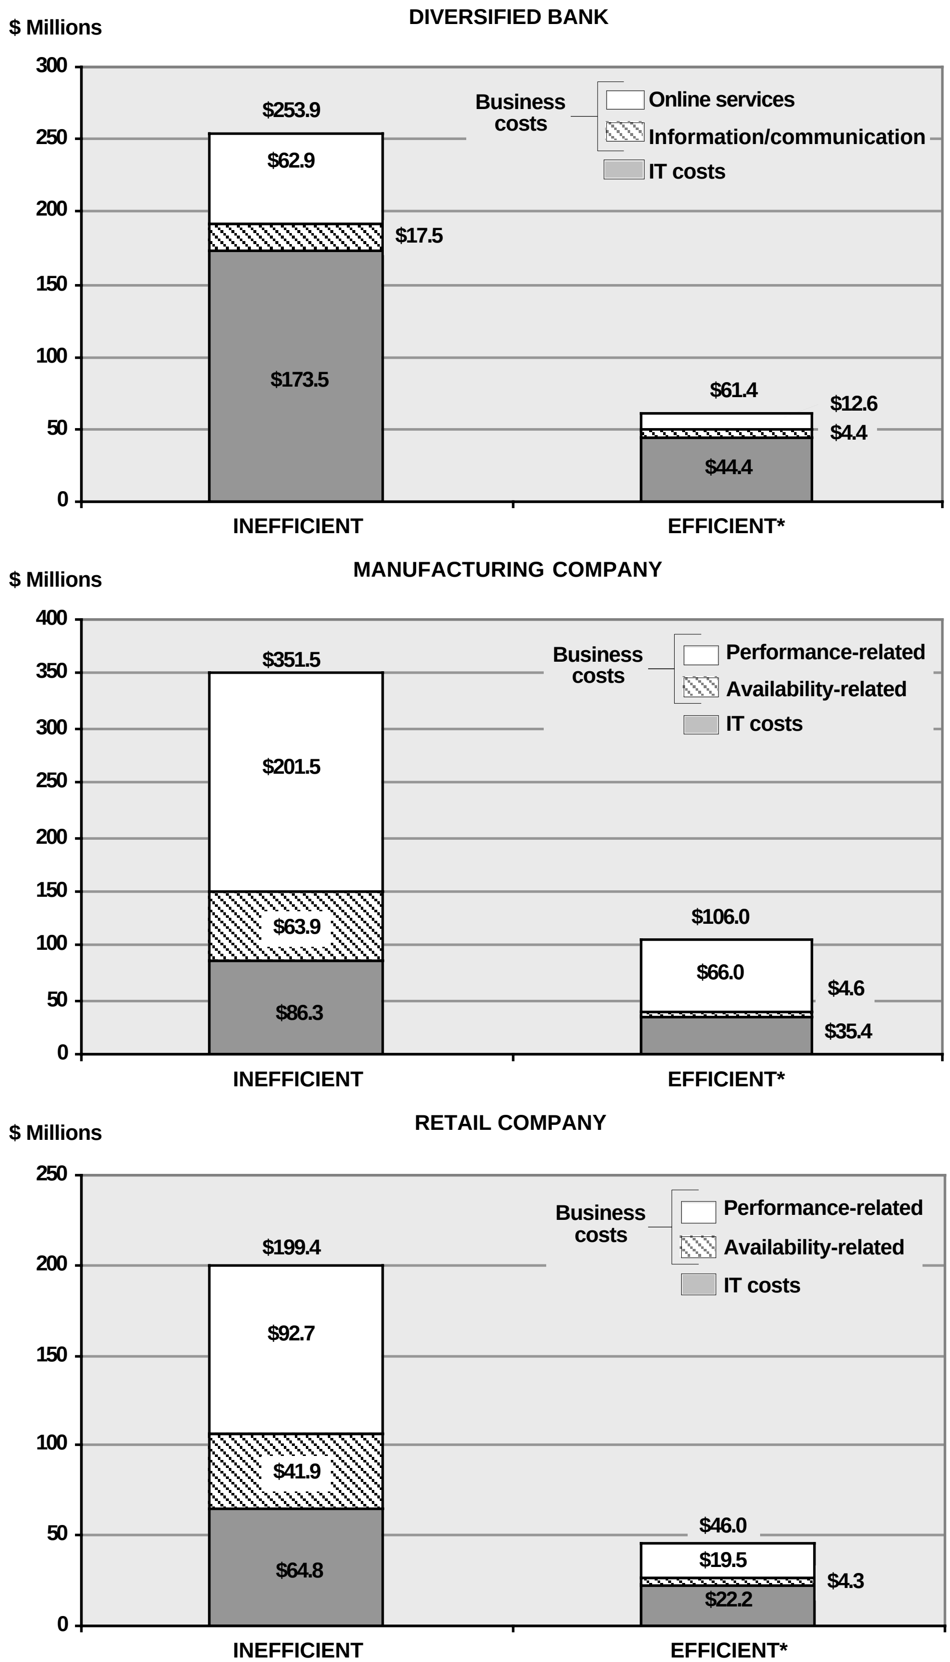

***Lowest cost vendor option (IBM)**

Vendor Comparisons

For efficient scenarios presented above, five-year IT costs were calculated for a number of vendor-specific options. Results for inefficient scenarios are for multivendor storage environments.

These included comparable configurations of latest-generation high-end and midrange storage systems from EMC, Hitachi Data Systems (HDS) and IBM high-end and midrange storage systems, along with equivalent suites of software tools, and SAN and NAS solutions from these vendors.

Hardware, software and SAN costs are based on "street" prices; i.e. prices actually paid by users. The user survey indicated that, in serious competitive bid situations, vendor street prices for hardware did not vary widely. Software licenses and SAN and NAS offerings were also heavily discounted.

Certain vendor systems were said to result in better performance for certain applications, but the differences were not major. The general view was that, for organizational workloads, differences in system-level performance or price/performance were not significant. The same cost values were thus employed for hardware and software licenses for all three vendors.

There was, similarly, little difference in footprints and environmental overheads for comparable systems. The same values are thus employed for facilities costs for all three vendors.

There were, however, a number of differences in the duration of hardware and software warranties, and in annual maintenance and software support costs between vendors that affected overall costs. Personnel costs for IBM systems were also marginally lower, due to higher levels of automation effectiveness. Maintenance, software support and personnel costs thus vary between vendors. Overall results are as summarized in figure 3.

Figure 3

IT Costs for Vendor-specific Options: Efficient Scenarios

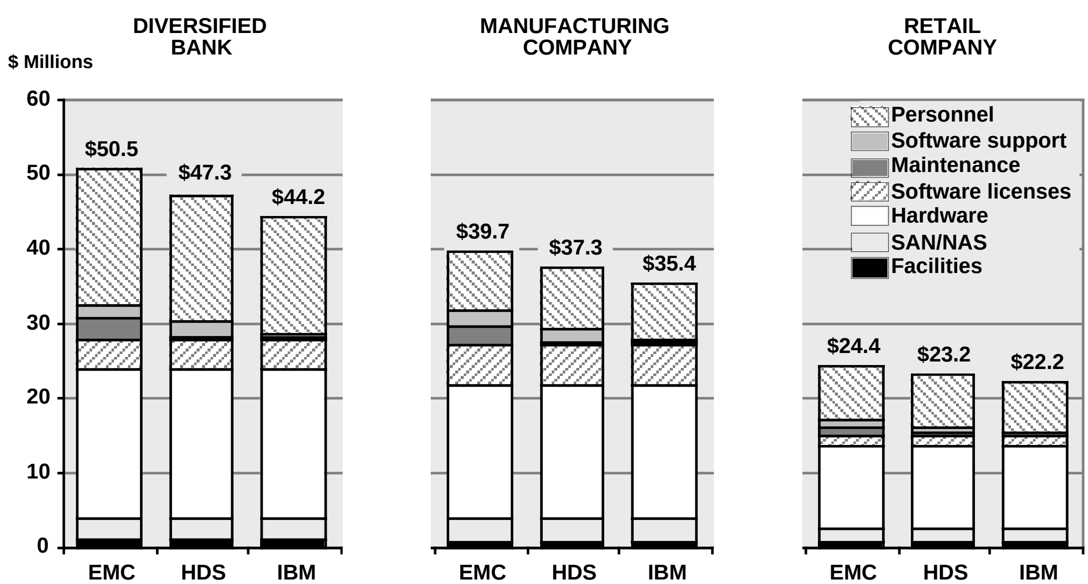

Five-year IT costs for IBM options were 12.3 percent less than for EMC and 6.3 percent less than for HDS for the diversified bank, 11.0 percent and 5.3 percent less respectively for the manufacturing company, and 8.6 percent and 4.3 percent less respectively for the retail company.

Comparable costs for the HDS option were 6.4 percent, 6.0 percent and 4.5 percent less than for EMC for the diversified bank, manufacturing company and retail company respectively.

Separately, hardware, maintenance and software costs were calculated for the midrange components of the same scenarios using storage systems from these vendors and Hewlett-Packard (HP).

Because hardware and software license costs tended again to be similar, the same values for these are employed for all vendors. Variations in overall costs are thus primarily due to differences in warranty duration, and in annual maintenance and software support costs.

Overall results are summarized in figure 4.

Figure 4

IT Cost Breakdowns: Midrange Component of Efficient Scenarios

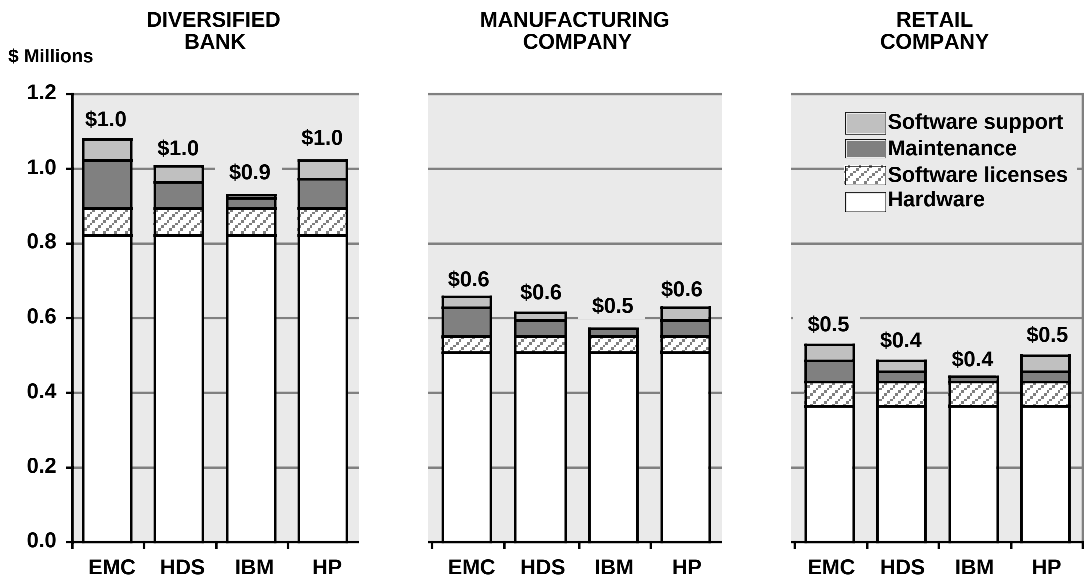

Since midrange systems form a relatively small part of scenarios, this comparison does not include SAN, NAS, personnel or facilities costs.

IBM again emerged as the lowest-cost option, followed by HDS, HP and EMC.

It should be emphasized that values used to calculate these results were based on broad survey averages. Actual vendor pricing varied between user organizations, and it can be expected that such variations will continue to be experienced.

Industry Trends

Among the companies surveyed, there was widespread agreement that storage deployment was moving away from standalone, platform-specific installations. The trend was toward more complex, networked environments in which storage systems supported an increasingly diverse range of applications and platforms.

At the same time, there were pervasive pressures to improve service quality – requirements for 24x7x365 availability were particularly cited – for an increasingly broad range of applications, and to control growth in storage capacity and management costs.

These trends are “raising the ante” for storage vendors. These must supply progressively more advanced management and optimization functionality for large-scale, multiple-platform, multiple-technology environments. They must also provide competence and support that extends across the entire enterprise storage base, and encompasses system-level capabilities as well as SANs, NAS and storage resource management (SRM) solutions.

It can be expected that these trends will become pervasive. Companies evaluating supplier relationships should thus evaluate the impact of changing requirements, and of the varying abilities of vendors to meet these.

Attention should be paid not only to the future evolution of system-level technology, but also to such variables as multivendor interoperability (particularly compliance with emerging Storage Network Industry Association standards), breadth of software and integration capability, and the ability to provide in-depth support for increasingly complex and diverse storage environments.

IBM appears to be moving into a leadership position in all of these areas. The company’s investment – in autonomic (advanced automation) technology represents a particularly significant long-term competitive edge. It can be expected that autonomic capabilities in IBM disk and tape systems, storage management software and SAN Volume Controller (“Lodestone”) will be greatly expanded, and will form a critical component of the company’s forthcoming SAN File System (“Storage Tank”).

Other vendors will probably remain competitive at the level of the storage device or system, and many companies will continue to employ their offerings. But it can be expected that the gap between IBM and its competitors in the ability to supply and support solutions for enterprise-level storage infrastructures will widen over the next five years. The implications are important.

Infrastructure Opportunity

One of the individuals interviewed for this document commented that their organization “didn’t like quantum leaps in technology.” This is a common and perhaps reasonable view. Few organizations wish to face the risks of attempting large-scale technological change on their own.

But a quantum leap has nevertheless become possible. If risks can be managed, and appropriate skills can be applied to the transition process, the benefits of creating an efficient storage infrastructure are potentially massive. Coupled with larger gains in IT productivity, and with enablement of new capabilities for information use and real-time competitiveness, the cost savings and service quality gains from investment in storage infrastructure represent a compelling value proposition.

The overall infrastructure opportunity is much greater. Although IT expenditure and staffing has increased consistently since the 1980s, there is evidence that IT productivity has not kept pace. An increasing proportion of resources have been channeled into managing overly complex, fragmented bases of servers, storage and middleware. The same factors of complexity and fragmentation have undermined processes of application development and deployment.

The creation of efficient infrastructures is emerging as the next phase of breakthrough progress in IT. The potential emerges for transformative gains that extend across all areas of IT implementation and expenditure, and across all of the business activities that IT supports.

The realization of even, say, a 20 to 30 percent increase in overall IT productivity would have a dramatic impact on competitiveness of any organization. The potential to achieve significantly larger gains exists now. That potential will increase rapidly over time.

One question must, however, be resolved. How are the broad skills and resources that will be necessary for infrastructure transformation to be acquired? The answer must be in rethinking the nature of vendor relationships. The traditional model of the vendor supplying systems which customers then build into their existing installed bases cannot, almost by definition, address greater challenges of infrastructure design, deployment and assurance at the enterprise level.

From this perspective, IBM emerges clearly as the leading player in storage solutions. Its ability to deliver lower costs, while significant, is part of a much broader picture. It is the only major vendor whose capabilities cover the full range of components – servers, storage, middleware and networks – and skills – project management, implementation and integration, support, training and financing – which will be required to achieve infrastructure transformation in large organizations.

PROFILE SUMMARIES

Diversified Bank

Business Profile

This company is a full-service retail bank that has expanded aggressively into other areas of financial services. Over the last decade, it has conducted a major series of acquisitions, including retail brokerage, loan, mutual fund, financial advisory and insurance businesses. It has also acquired a number of regional and local bank chains.

There are four main business units: Retail Banking (including small business banking), Commercial Banking, Capital Management (including brokerage, insurance, trust, 401(k), and other wealth and asset management businesses), and Capital Markets (including investment and corporate banking). At the close of its most recent fiscal year, the company reported approximately $250 billion in assets. Figure 5 summarizes other important statistics.

Figure 5

Key Bank Business Statistics

10 million retail banking customers	2,200 branches
4 million credit card accounts	4,000 ATMs
3 million brokerage account	500 financial centers
2 million home mortgages	350 brokerage offices
800,000 small business customers	7,000 brokerage representatives
20,000 commercial customers	3,000 financial advisors
3.5 million online accounts	10 call centers

A wide range of online services has been introduced, and the company has approximately 3.5 million online accounts of all types. There has steady growth in online banking and brokerage, electronic bill payment (EBP) and other Internet services.

There are approximately 70,000 employees, including 12,000 operations and IT, and 5,000 call center personnel, along with a 2,000-person corporate services group.

Business Strategy

Management is targeting aggressive long-term growth in earnings per share (EPS), return on equity (ROE) and shareholder value added (SVA); along with continued expansion of deposits, and of loan and investment portfolios.

The company is moving from a business model based on geographies and lines of business toward one built around customer groups, such as consumers, small businesses, and premium customers. An important objective is to more effectively exploit cross-sell opportunities for a full range of products and services across all customer segments. This shift involves major changes in reporting structures, business processes, compensation plans and incentive programs.

Two other themes in the company's business strategy are particularly significant from a storage perspective:

1. ***Information***. A central objective of IT strategy is to use information more effectively. One priority, according to management, is to "leverage customer information to retain and acquire customers, and deepen and enhance relationships through tailored services." Another is to "use new tools to anticipate customer behavior" in order to maximize sales opportunities for a wide range of product and service offerings.

   To achieve informational goals, a series of IT initiatives has been launched. The existing corporate data warehouse and a range of specialized data marts will be further expanded in size and functionality, and data from recent acquisitions will be integrated into these. Major investments will also be made in new analytical solutions and in database infrastructures to support customer relationship management (CRM) capability.

   There is a strong focus on creating mechanisms to deliver customer information more rapidly, in more useable form, to marketing functions and to front-line sales and service personnel. Management sees the creation of "real-time" information availability as a critical contribution to competitiveness. Other areas of informational investment include company-wide intranets.

2. ***Service quality***. A competitive goal of "being Number One in customer service" has been set. In achieving this goal, management intends to provide integrated, responsive service to all customers, across all touch points. A major objective is to provide 24x7x365 availability for Internet and call center channels.

   Online service quality is seen as particularly important. Online customers tend to be wealthier, more profitable and more loyal than the customer base as a whole, and to represent greater cross-sell growth opportunities. The fact that Internet costs are lower for transactions and customer interaction compared to other channels also represents long-term savings potential.

   Market research has shown that outages, poor response time, and delays in processing online transactions and queries have a significant impact on overall customer satisfaction. Service quality thus materially affects attrition levels, which have a major – and quantifiable – bottom-line impact.

   As the company moves toward real-time information, availability and performance become increasingly important for a wide range of systems. Messaging networks, along with intranets, query and analytical tools also become "business-critical." As such, they become as sensitive to the effects of disruption and delay as conventional systems.

   The importance of security and business continuity has increased significantly over the last few years. Responding to customer concerns, the bank has made leadership in online security a major marketing theme. It is thus particularly important to avoid negative publicity that would be generated by security breaches or privacy violations.

   Business continuity has become an increasing concern for business, legal and regulatory reasons. Effective backup and recovery coverage must be put in place not only for conventional transaction-processing systems and data, but also for a wide range of other applications and platforms.

Cost control and improved risk management across all lines of business remain important management priorities.

Storage Scenarios

Two scenarios are presented:

1. ***Inefficient***. This scenario involves five data centers operated by the principal bank and by two recent acquisitions: a regional bank, and a finance company offering brokerage, mutual fund, insurance and other services. As of yearend 2002, a diverse base of 13 mainframes and more than 2,400 RISC/UNIX and Intel-based servers was employed.

   Server and storage bases have expanded rapidly since the mid-1990s. Online services added more than 500 RISC/UNIX and Intel-based servers, and high rates of growth have also been experienced for other applications. These trends are expected to continue. Server and storage strategies are not coordinated at the corporate level.

   Under this scenario, server and storage bases will continue to expand over the next five years in the same manner. It is expected that at least 1,000 servers will be added, while the disk storage base is projected to increase from slightly more than 300 terabytes at yearend 2002 to close to 1,200 terabytes by yearend 2007.

2. ***Efficient***. In this scenario, the number of data centers has been reduced to two, and major server and storage consolidation initiatives have been undertaken. A “Chief Storage Officer” post has been created to coordinate procurement and deployment strategies, and to implement company-wide policies and procedures for storage management.

   The original 808 RISC/UNIX servers have been replaced by 236 IBM pSeries models, and 1,625 Intel-based servers have been replaced by 103 IBM xSeries 440 models, including new eight- and 16-way platforms. Consolidation of small Intel-based servers has been pursued in a particularly aggressive and successful manner.

   Disk storage resources are now organized on a cross-platform basis, with varying levels of RAID 10 (mirrored) and RAID 5 capability depending on business criticality and performance requirements. Improved backup and recovery mechanisms have been implemented. High-speed, fault-tolerant SAN infrastructures connect all servers, disk storage and tape subsystems. NAS gateways are located at both data centers.

   Under this scenario, the company’s base of disk storage is projected to increase from 229 terabytes at yearend 2002 to 673 terabytes by yearend 2007.

These scenarios, and their five-year staffing and cost implications are summarized, in figure 6. Efficient scenario values are for the lowest-cost vendor option, IBM.

In the efficient scenario, capacity utilization is significantly improved and the rate of storage growth slows. FTE staffing is significantly reduced and growth in numbers of personnel is stabilized. Major savings in all categories of storage-related IT expenditure are realized.

In moving from the inefficient to the efficient scenario, availability levels improve from 99.7 percent to 99.95 percent for critical processes of information access and communication through e-mail networks, strategic intranets and decision support query, and from 99.75 percent to 99.95 percent for customer-facing online services. Major improvements in response time and application-level throughput also occur for these and other systems.

Figure 6

Bank Scenarios: Summary Data

	SCENARIO	
	INEFFICIENT	**EFFICIENT**
Primary Data Centers	5	2
Other Server Locations	163	21
Servers		
Yearend 2002	13 mainframes 808 RISC/UNIX 1,625 Intel-based	8 zSeries 236 pSeries 103 xSeries
Disk Capacity		
Yearend 2002	301 TB	229 TB
Yearend 2007	1,179 TB	673 TB
Storage FTEs		
2003	181	34
2007	235	34
5-year IT Costs ($000)		
Hardware	39,723	19,939
Maintenance	5,836	534
Software	10,940	4,207
SAN/NAS	9,315	2,931
Personnel	103,350	15,815
Facilities	4,304	926
Total ($000)	**173,468**	**44,352**
5-year Business Costs ($000)		
Information/communication	17,534	4,417
Online services	62,861	12,572
Total ($000)	**80,395**	**16,989**
Severe Outage		
Duration	7 hours	15 minutes
Cost ($ million)	4.5 to 27.3	1.3

These result in net five-year business cost savings of $63.4 million through improved internal productivity and reduced attrition among online customers. Calculation of business costs for online services includes lost customer value over a seven-year period caused by inadequate availability and performance of systems accessed through the Internet.

Security and recoverability improve not only for major systems, but also for the many small servers that the bank continues to employ. The potential impact of severe outages is minimized. The bank is able to meet regulatory data protection and privacy requirements more easily. Risk exposure is materially reduced.

Scenario comparisons do not include servers and storage at bank branches and local offices. In these, only 3 to 10 gigabytes of data are maintained locally. Server-attached storage devices would be employed under both scenarios.

Strategic Benefit

In addition to reducing business and IT costs, the creation of an efficient storage infrastructure realizes broader benefits for the bank. By reducing the amount of work that must be performed for relatively low value-added storage administration tasks, an efficient storage infrastructure significantly improves the productivity of the entire IT organization. Improved IT productivity, along with reduced data management complexities, contribute to faster delivery of new applications.

The efficient scenario also directly supports the realization of central goals of business strategy. This is particularly the case in two areas:

1. ***Customer service***. Customer satisfaction is increased by improving availability and performance of Internet services, and by providing more timely access to information for customer service representatives (CSRs) and online self-service tools. Greater customer satisfaction results in higher retention rates, particularly among high-value customer groups, with positive effects on profitability.

2. ***Information***. The bottom-line impact of more effective use of information is potentially even greater. Better-quality, more current information increases the effectiveness of analytical tools employed to maximize profitability and reduce risk exposure.

   The ability to consolidate and share information more comprehensively and rapidly across business units has important implications. Corporate goals of improved cross-sell leverage, more granular customer segmentation, and higher-yield targeted marketing are materially facilitated.

   An efficient storage infrastructure enables the creation of new mechanisms for collection, distribution and exploitation of information across the entire organization. It thus materially facilitates the bank's transition from a product to a customer focus.

   More broadly, the bank is equipped to exploit information in real time for a wide range of high value-added applications. The bottom-line impact of real-time capability, while difficult to quantify, would clearly be substantial.

Compact, centralized storage resources will also assist long-term expansion. Development of new lines of business, and integration of future acquisitions, will be a faster, more reliable and more cost-effective process than under the distributed scenario.

Efficient Infrastructure

In the efficient scenario, a storage infrastructure based on two main tiers is implemented in each data center. These are:

1. ***Critical systems*** are production systems most sensitive to downtime. They include mainframe-based core banking systems; automated teller machine (ATM), point of sale (POS), and branch, call center and Internet banking systems; major line of business and divisional systems; e-mail and intranet backbones; and others.

   These are supported by high-end RAID 10 (mirrored) configurations equipped with tools for remote concurrent copy, point-in-time copy and core management functions. Parallel access volume (PAV) capability is also provided for certain mainframe applications.

2. ***Secondary systems*** include the remaining production business systems; corporate finance and human resources (HR) systems; the bank's large-scale data warehouse, along with data marts and analytical applications; departmental systems; end-user file and print services; Citrix services; non-critical intranet and Internet applications; and development, test and support systems for these.

   These are supported by high-end (secondary systems 1) or – in the case of certain RISC/UNIX- and Intel-based applications – midrange (secondary systems 2) disk storage configurations equipped with point-in-time copy and core management tools.

There are also NAS configurations in both data centers, and midrange RAID 5 systems are located at three small regional facilities. Initial (yearend 2002) installed bases and projected five-year growth for this scenario are summarized in figure 7. Storage volumes are raw capacity for IBM configurations. There are minor differences in EMC and HDS configurations.

Figure 7

Bank Efficient Scenario: Projected Disk Storage Growth

	2002	2003	2004	2005	2006	2007
DATA CENTER 1						
Critical systems*	32	40	48	63	80	100
Secondary systems 1	82	94	111	130	158	186
Secondary systems 2	10	12	14	18	22	28
NAS	1	2	2	2	2	3
Total (TB)	**125**	**148**	**175**	**213**	**262**	**317**
DATA CENTER 2						
Critical systems*	39	50	64	83	110	140
Secondary systems 1	40	51	65	84	110	144
Secondary systems 2	18	24	30	36	48	56
NAS	2	2	3	4	5	6
Total (TB)	**99**	**127**	**162**	**207**	**273**	**346**
REMOTE						
Regional centers (TB)	**5**	**5**	**10**	**10**	**10**	**10**

*Mirrored

Between yearend 2002 and yearend 2007, the disk storage base of will more than triple. Storage-related staffing will, however, remain stable, with 34 FTEs for IBM systems, and 37 and 39 FTEs for HDS and EMC systems respectively. In contrast, under the inefficient scenario, the number of storage FTEs increases from 118 to 235 FTEs over the same period. These numbers include all categories of IT personnel affected by storage efficiency.

Vendor-specific Costs

Five-year IT costs for high-end and midrange disk storage hardware, maintenance and software, along with costs for SAN and NAS, personnel and facilities for the efficient scenario are presented for EMC, HDS and IBM in figure 8.

Figure 8

Bank Efficient Scenario: Overall Vendor Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	11,398	1,915	1,756	1,666	1,775	1,429	**19,939**
Maintenance	–	–	–	798	932	1,055	**2,785**
Software licenses	2,048	399	63	282	522	669	**3,983**
Software support	–	239	364	388	424	496	**1,911**
SAN/NAS acquisition	2,175	65	92	89	87	91	**2,599**
SAN/NAS support	–	–	79	81	84	88	**332**
Personnel	–	3,614	3,614	3,614	3,614	3,614	**18,070**
Facilities	–	115	142	176	220	273	**926**
Total ($000)	**15,621**	**6,347**	**6,110**	**7,094**	**7,658**	**7,715**	**50,545**
HDS							
Hardware	11,398	1,915	1,756	1,666	1,775	1,429	**19,939**
Maintenance	–	–	–	–	270	311	**581**
Software licenses	2,048	399	63	282	522	669	**3,983**
Software support	–	230	352	374	408	477	**1,841**
SAN/NAS acquisition	2,175	65	92	89	87	91	**2,599**
SAN/NAS support	–	–	79	81	84	88	**332**
Personnel	–	3,423	3,423	3,423	3,423	3,423	**17,115**
Facilities	–	115	142	176	220	273	**926**
Total ($000)	**15,621**	**6,147**	**5,907**	**6,091**	**6,789**	**6,761**	**47,316**
IBM							
Hardware	11,398	1,915	1,756	1,666	1,775	1,429	**19,939**
Maintenance	–	–	–	–	247	287	**534**
Software licenses	2,048	399	63	282	522	669	**3,983**
Software support	–	–	–	–	102	122	**224**
SAN/NAS acquisition	2,175	65	92	89	87	91	**2,599**
SAN/NAS support	–	–	79	81	84	88	**332**
Personnel	–	3,163	3,163	3,163	3,163	3,163	**15,815**
Facilities	–	115	142	176	220	273	**926**
Total ($000)	**15,621**	**5,657**	**5,295**	**5,457**	**6,191**	**6,113**	**44,352**

For reasons outlined earlier, certain of the overall costs presented – hardware, software licenses, SAN and NAS, and facilities – are the same for all vendors, while maintenance, software support and personnel costs are vendor-specific.

Costs for the midrange component of the scenario are presented for EMC, HDS, IBM and HP in figure 9.

Figure 9

Bank Efficient Scenario: Midrange Vendor Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	542	65	76	44	61	32	**820**
Maintenance	–	–	–	38	43	48	**129**
Software licenses	76	–	–	–	–	–	**76**
Software support	–	9	11	11	11	11	**53**
Total ($000)	**618**	**74**	**87**	**93**	**115**	**91**	**1,078**
HDS							
Hardware	542	65	76	44	61	32	**820**
Maintenance	–	–	–	–	31	35	**66**
Software licenses	76	–	–	–	–	–	**76**
Software support	–	7	9	9	9	9	**43**
Total ($000)	**618**	**72**	**85**	**53**	**101**	**76**	**1,005**
IBM							
Hardware	542	65	76	44	61	32	**820**
Maintenance	–	–	–	–	13	15	**28**
Software licenses	76	–	–	–	–	–	**76**
Software support	–	–	–	–	4	4	**8**
Total ($000)	**618**	**65**	**76**	**44**	**78**	**51**	**932**
HP							
Hardware	542	65	76	44	61	32	**820**
Maintenance	–	–	–	–	35	39	**74**
Software licenses	76	–	–	–	–	–	**76**
Software support	–	9	11	11	11	11	**53**
Total ($000)	**618**	**74**	**87**	**55**	**107**	**82**	**1,023**

Midrange hardware and software license costs are again the same for all vendors, while maintenance and software support costs are vendor-specific. SAN, NAS, personnel and facilities costs are not broken out separately, as they would be determined by overall installations of high-end as well as midrange systems. In practice, they would be similar for all four vendors.

Manufacturing Company

Business View

This is a global manufacturer of branded consumer packaged goods (CPG) with approximately $12 billion in revenues and 40,000 employees. It currently operates approximately 120 manufacturing plants and 200 distribution centers, along with 15 research, development and engineering centers and more than 400 administrative and sales facilities. The profile focuses on the use of SAP AG systems.

During the 1990s, the company invested heavily in SAP R/3 ERP systems and in reengineering of related business processes. R/3 was deployed across 18 regional and national business units accounting for over 98 percent of worldwide sales. Major gains were achieved in competitiveness and profitability. By 2000, the company's gross margin was over 50 percent, compared to approximately 40 percent a decade before.

Improvements in supply chain efficiency, illustrated in figure 10, were particularly striking.

Figure 10

Manufacturing Company: Improvements in Supply Chain Efficiency

BUSINESS METRIC	BEFORE SAP R/3	AFTER SAP R/3
Order acquisition & processing	1 – 7 days	4 – 6 hours
Purchase order to delivery cycle	12 – 15 days	4 – 5 days
On-time delivery	92%	98%
Order fulfillment accuracy	97%	99%
Inventory reduction	–	30%
Delivery cost reduction	–	10%
Working capital reduction	–	15%
Receivables outstanding	31 days	22 days

Like many other organizations that deployed ERP systems in the 1990s, however, the management found that such gains were not, in themselves, sufficient to guarantee long-term competitiveness. By the end of the decade, major new challenges were emerging.

By 2000, the CPG share of overall consumer spending was declining in relatively mature markets such as North America and Western Europe. The company continued, moreover, to experience severe competitive pressures.

These were reinforced by concentration trends among retail customers. Fewer, larger chains were able to exert more control over terms of business with their suppliers. In addition to demanding lower prices, retailers increasingly required revenue and profit guarantees, customized products, and improved quality and timeliness of service. Penalties and charges (e.g. for late delivery or mislabeling) were escalating.

As retailers expanded internationally, they began demanding multinational supply arrangements and pricing, along with supply support for global and regional promotions. This shift obliged the adoption of a more global business model. Industry-wide shifts to vendor managed inventory (VMI) also required further improvements in supply chain efficiency.

Management anticipated that VMI operations would soon account for more than 40 percent of North American sales, and that other geographies would move in the same direction. All of these pressures were magnified by retailer moves to expand store brand sales.

It was also necessary to improve product innovation. By 2000, approximately 40 percent of sales were of products less than five years old, but it would be necessary to increase this proportion significantly. Greater sophistication in marketing and promotional activities would be a competitive necessity, obliging the company to move beyond traditional forms of advertising toward coordinated, multiple-channel marketing and promotional programs.

These challenges led to major shifts in ERP strategy, which are described below.

New ERP Strategy

The company's ERP plans focus on three new themes:

1. ***Consolidation***. Initial deployment of R/3 systems reflected a decentralized business structure. Marketing and supply chain operations had been organized on a country-by-country basis, with country-specific business units responsible for planning, purchasing, production and distribution to local customer bases.

   The company plans to move to a business model based on four regional structures: North America; Latin America; Europe, Middle East and Africa (EMEA); and Asia/Pacific. Customer relationships, as well as procurement, manufacturing and logistics will be managed at the regional level.

   Global accounts structures are also being put in place, and a global sourcing program is being launched to coordinate regional procurements. Finance, payments processing, HR and other functions are being consolidated to new regional shared services centers to reduce costs.

   To facilitate innovation, however, responsibility for new product development will be decentralized from corporate to regional level. Regional organizations will have more flexibility to develop and commercialize products geared to their markets.

2. ***New solutions***. Original R/3 deployment focused on transaction-processing solutions. The ERP emphasis moves to new SAP solutions for supply chain optimization, business intelligence, e-procurement, online sales, intranet self-service tools and other applications.

   Deployment of these will build upon the core R/3 environment. New solutions will exploit data generated by R/3 operational systems, and will extend core business process through the Internet and intranets. The company's overall approach is illustrated in figure 11.

   With this transition, the primary source of competitive advantage is seen as shifting from transactional efficiency to more effective use of information and better exploitation of online opportunities. It is planned to eventually migrate the underlying R/3 environment to more flexible and technologically sophisticated mySAP solutions.

Figure 11

Manufacturing Company: New Strategy for ERP Solutions

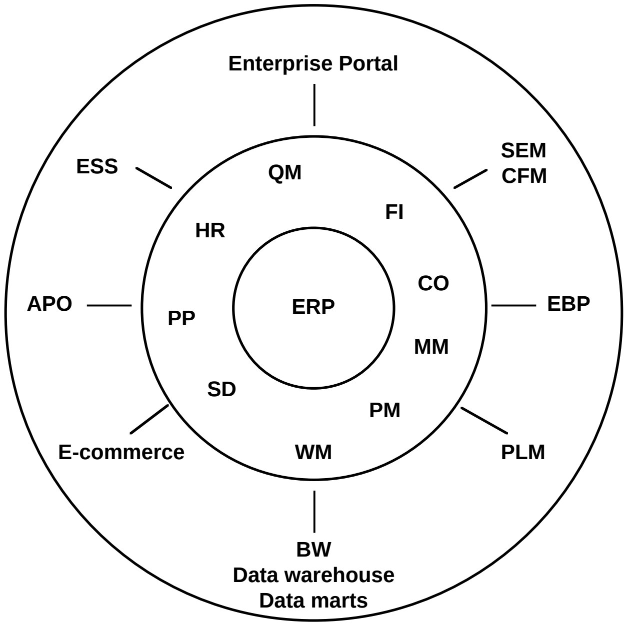

Initial Solutions

CO: Controlling
FI: Financial Accounting
HR: Human Resources
MM: Materials Management
PM: Plant Maintenance
PP: Production Planning & Control
QM: Quality Management
SD: Sales & Distribution
WM: Warehouse Management

New Solutions

APO: Advanced Planning & Optimization
BW: Business Information Warehouse
CFM: Corporate Finance Management
CRM: Customer Relationship Management
EBP: Enterprise Buyer Professional
ESS: Employee Self Service
PLM: Product Lifecycle Management
SEM: Strategic Enterprise Management

3. ***Real-time model***. The company must become more responsive to fast changing, increasingly exacting customer requirements. It must also further reduce cycle times, inventory levels and costs throughout the supply chain; continue to improve performance in on-time delivery and delivery accuracy; and more effectively project short- as well as long-term demand.

   The ability to assemble information from all stages of the supply chain, and to interpret and exploit it in real time, becomes increasingly central to the company's competitive model.

   Sources of information are expanding to include direct interfaces to customer systems through VMI relationships. Volume and complexity of information are magnified as the company moves toward global and regional contracts with larger customers, and toward coordination of purchasing and distribution across national and regional boundaries.

   This new breadth of information must be exploited not only to increase operational efficiency, but also to accelerate product development, marketing and promotional cycles. The challenges of real-time competitiveness extend beyond the supply chain to include all business activities.

The extent to which new challenges are met will depend, in large part, on the IT resources that underpin the expanded and globalized ERP environment. The company's storage infrastructure will, in turn, play a critical role in determining the effectiveness of these. The speed, effectiveness and reliability with which storage enables the collection, distribution and exploitation of information will be particularly important in achieving real-time competitiveness.

The choice of a storage strategy will be a critical decision.

Storage Scenarios

Two scenarios are presented:

1. ***Inefficient***. In this scenario, the new ERP strategy is realized using the decentralized IT structures created in the 1990s. Separate R/3 systems continue to be employed by 18 major subsidiaries, each with their own servers and storage. There are 27 data centers worldwide.

   New SAP applications such as Advanced Planning and Optimization (APO), Customer Relationship Management (CRM), Enterprise Buyer Professional (EBP) and Product Lifecycle Management (PLM) are deployed by each subsidiary. Four large data warehouses and various smaller data marts are employed by regional and in some cases national units.

   Under this scenario, a diverse range of 642 RISC/UNIX- and Intel-based servers, along with a variety of storage devices are initially – as of yearend 2002 – employed to support SAP applications. Over the next five years, the server base is expected to more than double. Driven by growth in data warehouses as well as operational systems, worldwide disk storage is projected to increase from around 273 to more than 1,000 terabytes.

   Storage administration, backup and recovery arrangements, and use of SANs and NAS remain the responsibility of subsidiary IT organizations. Tools and practices continue to vary between these.

2. ***Efficient scenario***. Under this scenario, the number of R/3 instances has been reduced from 18 to 4 – one for each region – and most new SAP applications are deployed as regional or worldwide systems. A global data warehouse has been put in place.

   The company's 27 data centers have been consolidated to two. Major investments have made in wide area network capability to support global operations while maintaining adequate response time and application-level throughput for local users. The two data centers, both of which are located in North America, are fully duplexed for disaster recovery purposes.

   Extensive server and storage consolidation initiatives have taken place, and technology bases have been upgraded. The central data centers initially contains 65 pSeries and 156 xSeries servers, along with 143 terabytes of high-end and midrange disk storage.

   All servers and storage are attached to fault-tolerant SANs. Over the next five years, servers and disk storage are expected to expand to around 300 and 450 terabytes respectively.

   Smaller data centers with pSeries and xSeries servers are maintained in Europe, Latin America and Asia to support regional CRM systems. Initially these employ 11 pSeries and 24 xSeries servers with 12 terabytes of disk storage, projected to increase 48 servers and 39 terabytes by yearend 2007.

   Although IT operations have been consolidated, regional business units retain responsibility for applications software. Certain application development responsibilities, particularly in such areas as research and development, product management, decision support, marketing and CRM have been decentralized to regions.

Disk storage volumes cited above are for production instances of SAP systems, as well as for non-production instances for development, testing, sandbox, quality assurance, staging, integration, training and other support functions.

These scenarios, and their five-year staffing and cost implications, are summarized in figure 12. Efficient scenario values are for the lowest-cost vendor option, IBM.

Figure 12

Manufacturing Company Scenarios: Summary Data

	SCENARIO	
	INEFFICIENT	**EFFICIENT**
ERP Instances	18	4
Data Warehouses	4	1
Data Centers	27	2*
Servers		
Yearend 2002	351 RISC/UNIX 291 Intel-based	76 pSeries 180 xSeries
Disk Capacity		
Yearend 2002	273 TB	155 TB
Yearend 2007	1,022 TB	489 TB
Storage FTEs		
2003	74	17
2007	109	17
5-year IT Costs ($000)		
Hardware	29,587	17,755
Maintenance	4,587	487
Software	6,198	5,627
SAN/NAS	5,536	3,222
Personnel	36,021	7,580
Facilities	4,356	695
Total ($000)	**86,285**	**35,366**
5-year Business Costs ($000)		
Availability-related	63,855	4,562
Performance-related	201,490	65,947
Total ($000)	**265,345**	**70,509**
Severe Outage		
Duration	72 hours	15 minutes
Cost ($ million)	2.7 to 66.2	0.3

*Primary data centers. Three regional centers continue to be employed for CRM systems.

Under the efficient scenario, capacity utilization is significantly improved and the rate of disk storage growth slows. FTE staffing is significantly reduced, and growth in numbers of personnel is stabilized. Major savings in all categories of storage-related IT expenditure are realized.

Availability levels improve from 99.3 percent to more than 99.95 percent for R/3 systems, as well as for new SAP systems such as APO, EBP and PLM, along with business-critical intranets, and online sales, order processing and customer self-service applications. Five-year business costs due to outages are reduced from $63.9 million for the inefficient scenario to $4.6 million.

The business impact of improved performance is even more significant. Response times improve by a average of more than 30 percent. Batch throughput for critical applications is increased by an average of more than 3.4 times. Significant improvements in online throughput also occur. Operational processes, as well as delivery of information for decision support, marketing, customer service and other applications are accelerated.

Exposure to severe outage risks is greatly reduced. Under the inefficient scenario, subsidiary-level recovery times in the event that a data center was disabled were typically three or more days. Depending on which subsidiary is affected, potential losses range from $2.7 million to $66.2 million. Under the efficient scenario, recovery time is reduced to 15 minutes and losses to less than $300,000.

Business costs given here represent lost gross or operating profit, profit net of cost of goods sold, but before deduction of selling, general and administrative (SG&A) and other expenses. This metric is used in preference to lost revenues for quantifying the cost impact of outages.

Lost revenue calculations are normally based on the assumption that a given volume of revenues is lost during an outage. This approach tends, however, to understate actual bottom-line impacts. In a tightly integrated ERP environment characterized by just in time (JIT) practices, real-time operations and minimal inventory levels, outages cause “ripple effects” that impact not only internal processes, but also the supply chain as a whole. The effects last long after service has been restored.

Lost gross profit is used in the same manner to quantify the effects of delays and disruptions caused by shortfalls in response time and application-level throughput.

Strategic Benefit

Transition to the efficient scenario reduces IT and business costs. Application deployment cycles are accelerated, and overall IT productivity improved in the same way as for the diversified bank.

Creation of an efficient, centralized kernel of storage resources significantly reduces the costs, difficulties and time required for the company’s consolidation program. It will deliver similar benefits for future system upgrades and version changes – including migration from R/3 to R/3 Enterprise and eventually to a fully integrated mySAP environment – and for addition of new applications to the core ERP complex.

The efficient scenario also materially improves the resilience of the company’s highly integrated global systems to the effect of outages and disruptions. Consequently, on-time delivery rates are improved. Sales shortfalls, customer dissatisfaction and inventory build-up, as well as the penalties for late delivery may be better avoided.

A further benefit is that spikes and surges in demand can be more easily met, and storage resources more rapidly and efficiently reallocated in a centralized environment than was the case for subsidiary-level operations. This enables the company to manage seasonal variations in a more cost-effective manner. Its ability to deal with the greater volatility in demand caused by industry-wide moves toward promotion-driven marketing also improves.

Most importantly, the efficient storage infrastructure enables real-time competitiveness. Increasingly large amounts of information must be pulled from transaction-processing systems, on a continuous basis, with minimal delays in extraction, distribution and exploitation.

It would be difficult, if not impossible, to realize real-time capability across the original base of distributed servers and storage. Complexities of data integration, as well as the potential for data movement bottlenecks, would be massively greater than in a centralized, efficiently managed environment. Backup and recovery challenges would be equally daunting.

An important conclusion emerges. Efficient storage scenarios do not merely facilitate new ERP strategies driven by information transparency and real-time operations. They are an essential prerequisite for these. This company – and others confronted with similar challenges – could not achieve transformative gains in competitiveness without an effective enterprise-level storage infrastructure.

Efficient Infrastructure

Under the efficient scenario, a storage infrastructure based on two main tiers is implemented in each data center. These are:

1. ***RAID 10 systems*** are high-end configurations equipped with remote concurrent copy, point-in-time copy and core management tools. They support critical systems for which 24x7x365 availability must be maintained.

   Applications in data center 1 include production instances of the R/3 suite and of APO, EBP, PLM, e-commerce systems, and complementary applications including electronic data interchange (EDI), forms, fax, and R/F device and scanner support. Mirrored configurations in data center 2 are used for non-production instances, and are available for failover purposes.

2. ***RAID 5 systems*** are high-end configurations equipped with point-in-time copy and core management tools. In data center 1, RAID 5 systems are used primarily for business intelligence applications. They support the company's data warehouse; specialized data marts for financial, human resources and market research analysis; and SAP Corporate Finance Management (CFM) and Strategic Enterprise Management (SEM) systems.

   In data center 2, RAID 5 systems support the company's North American CRM system, along with SAP Enterprise portal and intranet applications. They are also used for support instances for these and other SAP systems, and for a number of third-party applications.

Midrange storage systems equipped with point-in-time copy and core management tools are employed in EMEA, Latin America and Asia/Pacific data centers to support regional CRM systems. There are also small NAS installations in data centers 1 and 2. Initial disk storage bases and projected growth for these are summarized in figure 13. Volumes are raw capacity for IBM systems.

Figure 13

Manufacturing Efficient Scenario: Projected Disk Storage Growth

YEAREND	2002	2003	2004	2005	2006	2007
DATA CENTER 1						
RAID 10	40	48	59	70	90	110
RAID 5	30	40	54	71	93	129
NAS	2	3	4	5	6	8
Total (TB)	**72**	**91**	**117**	**146**	**189**	**247**
DATA CENTER 2						
RAID 10	40	48	59	70	90	110
RAID 5	30	38	48	60	71	88
NAS	1	2	2	3	4	5
Total (TB)	**71**	**88**	**109**	**133**	**165**	**203**
REGIONAL CENTERS						
Latin America	4	5	6	8	10	12
EMEA	5	6	8	10	12	15
Asia/Pacific	3	4	5	8	10	12
Total (TB)	**12**	**15**	**19**	**26**	**32**	**39**

Although the storage base increases from 155 to 489 terabytes between yearend 2002 and yearend 2007, storage-related staffing will remain stable, with 17 FTEs for IBM systems and 18 FTEs for HDS and EMC systems. Under the distributed scenario, the number of storage FTEs increases from 74 to 109 FTEs over the same period. These numbers include all categories of IT personnel affected by storage efficiency.

Vendor-specific Costs

Five-year overall costs for disk storage hardware, maintenance and software, along with costs for SAN and NAS, personnel and facilities for the entire installation are presented for EMC, HDS and IBM in figure 14.

Figure 14
Manufacturing Company Efficient Scenario: Overall Vendor Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	10,338	1,477	1,621	1,422	1,511	1,386	**17,755**
Maintenance	–	–	–	724	827	941	**2,492**
Software licenses	2,029	452	496	507	948	982	**5,414**
Software support	–	228	356	428	504	630	**2,146**
SAN/NAS acquisition	2,247	102	78	141	145	162	**2,875**
SAN/NAS support	–	–	81	85	88	93	**347**
Personnel	–	1,601	1,601	1,601	1,601	1,601	**8,005**
Facilities	–	82	104	131	167	211	**695**
Total ($000)	**14,614**	**3,942**	**4,337**	**5,039**	**5,791**	**6,006**	**39,729**
HDS							
Hardware	10,338	1,477	1,621	1,422	1,511	1,386	**17,755**
Maintenance	–	–	–	–	243	275	**518**
Software licenses	2,029	452	496	507	948	982	**5,414**
Software support	–	185	287	346	408	509	**1,735**
SAN/NAS acquisition	2,247	102	78	141	145	162	**2,875**
SAN/NAS support	–	–	81	85	88	93	**347**
Personnel	–	1,601	1,601	1,601	1,601	1,601	**8,005**
Facilities	–	82	104	131	167	211	**695**
Total ($000)	**14,614**	**3,899**	**4,268**	**4,233**	**5,111**	**5,219**	**37,344**
IBM							
Hardware	10,338	1,477	1,621	1,422	1,511	1,386	**17,755**
Maintenance	–	–	–	–	228	259	**487**
Software licenses	2,029	452	496	507	948	982	**5,414**
Software support	–	–	–	–	91	122	**213**
SAN/NAS acquisition	2,247	102	78	141	145	162	**2,875**
SAN/NAS support	–	–	81	85	88	93	**347**
Personnel	–	1,516	1,516	1,516	1,516	1,516	**7,580**
Facilities	–	82	104	131	167	211	**695**
Total ($000)	**14,614**	**3,629**	**3,896**	**3,802**	**4,694**	**4,731**	**35,366**

For reasons outlined earlier, certain of the overall costs presented here – hardware, software licenses, SAN and NAS, and facilities – are the same for all vendors, while maintenance, software support and personnel costs are vendor-specific.

Costs for the midrange component of the scenario are presented for EMC, HDS, IBM and HP in figure 15.

Figure 15

Manufacturing Company Efficient Scenario: Midrange Vendor Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	308	49	37	50	32	31	**507**
Maintenance	–	–	–	22	25	28	**75**
Software licenses	44	–	–	–	–	–	**44**
Software support	–	5	7	7	7	7	**33**
Total ($000)	**352**	**54**	**44**	**79**	**64**	**66**	**659**
HDS							
Hardware	308	49	37	50	32	31	**507**
Maintenance	–	–	–	–	18	21	**39**
Software licenses	44	–	–	–	–	–	**44**
Software support	–	4	5	5	5	5	**24**
Total ($000)	**352**	**53**	**42**	**55**	**55**	**57**	**614**
IBM							
Hardware	308	49	37	50	32	31	**507**
Maintenance	–	–	–	–	8	9	**17**
Software licenses	44	–	–	–	–	–	**44**
Software support	–	–	–	–	2	2	**4**
Total ($000)	**352**	**49**	**37**	**50**	**42**	**42**	**572**
HP							
Hardware	308	49	37	50	32	31	**507**
Maintenance	–	–	–	–	20	23	**43**
Software licenses	44	–	–	–	–	–	**44**
Software support	–	5	7	7	7	7	**33**
Total ($000)	**352**	**54**	**44**	**57**	**59**	**61**	**627**

Midrange hardware and software license costs are again the same for all vendors, while maintenance and software support costs are vendor-specific. SAN, NAS, personnel and facilities costs are not broken out separately, as they would be determined by overall installations of high-end as well as midrange storage systems. In practice, they would be similar for all vendors.

Retail Company

Business View

This company is a grocery and drug retailer with approximately $35 billion in revenues, 180,000 employees, and 1,700 stores throughout North America. These include 1,200 supermarkets, of which 850 contain in-store pharmacies, along with 200 “superstores” and 300 standalone drug stores. Outlets average more than 35,000 stock keeping units (SKUs) each, and carry more than 1,200 major vendor and 10 store brands.

The company grew rapidly during the 1990s through a major series of mergers and acquisitions. More recently, however, it has focused on acquiring smaller regional chains and groups of stores from competitors. Since 2000, a steady pace of such acquisitions has continued, typically in increments of 5 to 40 stores. Some stores have also been closed and disposed of to rationalize the overall network.

The logistics organization operates 25 distribution centers, along with approximately 1,000 tractors and 4,000 trailers, and moves more than 450 million cases of products per year. It employs approximately 6,500 people.

In 2001, the company acquired a loss-making online grocery business. By leveraging economies of scale in purchasing, and using its existing distribution infrastructure to cut logistics costs, this has become operation profitable. Internet sales are expected to exceed $150 million during 2003, and are increasing by a rate of more than 20 percent per year. Significant investments continue to be made in this operation to exploit its long-term growth potential.

The company has been an industry leader in other areas of retail IT. It has employed data warehousing for a range of applications since the early 1990s. More recently, aggressive deployment of self-checkout systems, and of R/F technology for logistics and stocking has begun.

Business Strategy

The company’s business strategy is built around five main themes: (1) aggressive cost control and process streamlining (2) maximizing return on capital (3) maintenance of growth (4) company-wide focus on technology, and (5) continued emphasis on employee management and motivation.

In a business characterized by low profit margins, management has historically been cost-sensitive. Weak economic conditions and increasing pressure on profitability has, however, recently led to greater emphasis on cost control and process streamlining. In practice, this is the main focus of business strategy.

A large-scale restructuring exercise has begun. At least 200 underperforming outlets will be closed or sold. Distribution operations will be rationalized. A series of divisional mergers and reorganizations is underway, which will cut approximately 1,500 non-store management and administrative positions. Management is also targeting cost reductions through greater centralization of purchasing.

Organization-wide process reengineering and best practices initiatives have been launched in these and other areas. The overall goal is to realizing savings of at least $250 million per year over a three-year period while maintaining conservative revenue growth.

Management expects IT to play a major role in realizing these objectives. More effective enterprise-level procurement and management of IT resources, rationalization of applications, and data center and server consolidation are seen to offer cost savings potential. This is particularly the case in that the company continues to employ diverse systems inherited from various acquisitions in marketing, procurement, distribution and other areas. Outsourcing opportunities are also being investigated.

IT is also expected to play a more proactive role in increasing organizational efficiency. R/F and self-checkout technologies are expected to yield new store-level savings. Business intelligence capability will remain a major area of investment to identify new opportunities for growth and cost reduction. There is a new focus on supply chain integration and optimization, and on realization of online procurement economies.

More broadly, it is recognized that the company must transition to a real-time competitive model. Faster procurement and logistics cycles, reduced inventory and logistics costs, and improvements in the speed and accuracy of forecasting, planning and replenishment processes are targeted. IT is seen as high-impact enabler of "real-time, fact-based decisions" across the entire supply chain.

Storage Scenarios

Two scenarios are presented:

1. ***Inefficient***. Under this scenario, the company continues to operate with application portfolios, IT structures and installed bases inherited from the original organization and from two recent acquisitions. There are three sets of data warehouses. Core retail, procurement, distribution, sales and marketing, finance, HR and other types of application also vary between chains.

   Servers and storage are located at three primary data centers, and at 48 distribution centers, divisional headquarters and regional sites. Intel-based servers supporting e-mail, intranet and end-user applications are located at these and 116 other sites, including individual stores.

   Eight mainframe systems with 19.9 terabytes of disk storage are divided between the three primary data centers. There are 141 RISC/UNIX servers with 91.2 terabytes of disk storage supporting business intelligence, transaction processing, online and divisional applications. Three data warehouses account for 41.9 terabytes of this total.

   In addition, there are 13 AS/400s with 2.9 terabytes of disk storage employed at a number of distribution centers, and 696 Intel-based servers with 12.7 terabytes of disk storage. Intel-based servers support a mix of Lotus Notes and Microsoft Exchange e-mail and groupware networks, along with intranet, Internet, departmental, file/print serving and other applications.

   If growth continues at the same rate as during 2001 and 2002, under this scenario the company's server base will expand from 858 to close to 1,500 over the next five years. Its disk storage base is projected to increase from 126.7 terabytes at yearend 2002 to more than 685 terabytes by yearend 2007.

   Storage administration, backup and recovery arrangements, and use of SANs and NAS vary between the three chains. Although purchases of storage hardware and software are centralized, this is not the case for management tools and practices. There is no overall corporate storage strategy.

2. ***Efficient scenario***. Under this scenario, the company has standardized key systems. A single set of legacy mainframe-based core retail applications will be employed company-wide, although these will eventually be replaced with more functional packaged solutions. Other corporate standards include PeopleSoft Financials and HRMS, Lotus Domino, i2 Technologies online procurement, and Computer Associates EDI software.

   The company also deploys Retek Merchandising Optimization and Planning, Inventory Optimization and Planning, and Supply Chain Management suites. These replace a diverse collection of packages on RISC/UNIX, AS/400 and Windows platforms, and play an important role in achieving the company's objectives for supply chain integration and optimization. The number of distribution centers has been reduced from 36 to 25.

   Aggressive consolidation initiatives have been undertaken, and server and storage technology bases have been upgraded. The company's diverse base of RISC/UNIX and Intel-based servers have been replaced by 78 pSeries and 121 xSeries servers, including high-end eight- and 16-way xSeries 440 models. Its eight mainframes have been reduced to three. It is expected that the company will employ around 300 servers by yearend 2007.

   This facility contains all of the company's mainframes and most of its RISC/UNIX and Intel-based servers. The only exceptions are small RISC/UNIX and Intel-based servers located at distribution centers and three regional sites.

   Three primary data centers have been replaced by a single new, state-of-the-art facility backed up by a third-party disaster recovery arrangement. Within the facility, disk storage has been reorganized on a cross-platform basis, using a combination of RAID 10 and RAID 5 systems, and NAS. All servers and storage are attached to high-speed duplexed SANs.

   There are initially around 95 terabytes of disk storage, projected to increase to 453 terabytes over the next five years. The company's principal data warehouse will account for most of this growth.

   A new Enterprise Storage organization has been made responsible for storage strategy, support and operations. Fully integrated storage resource management (SRM) capabilities have been put in place at the new central data center, and data on the few remaining remote servers is also backed up this. A "best practices" program for storage, backup and recovery management has been put in place.

Under the efficient scenario, capacity utilization is significantly improved and the rate of growth is slower than for the inefficient environment. FTE staffing is reduced, and growth in numbers of personnel is stabilized. Major savings in all categories of storage-related IT expenditure are realized.

These scenarios, and their five-year staffing and cost implications, are summarized in figure 16. Efficient scenario values are for the lowest-cost vendor option, IBM.

In moving from the inefficient to the efficient scenario, availability levels for core retail and supply chain systems (including Internet procurement and supplier self-service) improve from 99.9 percent and 99.7 percent respectively to 99.99 percent, significantly reducing the impact of outages. Average response times are halved, and batch and online throughput times are reduced by wide margins. Batch order-processing windows for replenishment are, in particular, reduced almost 2.7 times.

Figure 16

Retail Company Scenarios: Summary Data

	SCENARIO	
	INEFFICIENT	**EFFICIENT**
Data Centers	3	1
Servers		
Yearend 2002	8 mainframes 13 AS/400s 141 RISC/UNIX 696 Intel-based	3 zSeries 78 pSeries 121 xSeries
Disk Capacity		
Yearend 2002	127 TB	95 TB
Yearend 2007	685 TB	453 TB
Storage FTEs		
2003	74	15
2007	54	15
5-year IT Costs ($000)		
Hardware	20,073	11,173
Maintenance	3,278	227
Software	3,938	1,484
SAN/NAS	4,426	1,899
Personnel	30,801	6,835
Facilities	2,248	568
Total ($000)	**64,764**	**22,186**
5-year Business Costs ($000)		
Availability-related	41,928	4,339
Performance-related	92,660	19,549
Total ($000)	**134,588**	**23,888**
Severe Outage		
Duration	120 hours	20 minutes
Cost ($ million)	27.0 to 155.1	0.6

Five-year availability-related business costs for these systems are reduced from $39.5 million under the inefficient scenario to $4.2 million, while performance-related business costs are reduced from $84.3 million to $19.0 million for net savings of $100.6 million. Costs are for lost gross profit, and are calculated in the same way as for the manufacturing company. Performance-related costs and savings are again greater than for their availability-related counterparts.

For the Internet sales operation, availability is increased from 99.8 percent to 99.95 percent, and response and order processing times are significantly improved. Five-year business costs due to immediate lost sales and customer defections are reduced from $8.3 million under the inefficient scenario to $403,000. Costs of customer defections are calculated based on seven-year CLV.

The business cost impact of improved availability and performance for Internet sales is magnified by the relatively large size of online orders. Management estimates that online orders average approximately $120, or more than four times the average of purchases by in-store customers.

For the company’s 40,000-user e-mail network and business-critical intranets, availability increases from 99.4 percent under the inefficient scenario to 99.99 percent, while response times are halved and application-level throughput times are increased by an average of close to 2.9 times. Five-year business costs due to lost productivity decline from $2.5 million under the inefficient scenario to $276,000.

Combined five-year business costs for all of these systems decline from $134.6 under the inefficient scenario to $23.9 million under the efficient scenario.

Exposure to severe outage risks is also greatly reduced. Under the inefficient scenario, recovery times if a data center was disabled are up to five days. Depending on which data center was affected, business costs range from $27.0 million to $155.1 million. Under the efficient scenario, recovery time is reduced to 20 minutes and business costs to $586,000. Business costs are for lost gross profit due to supply chain disruption and lost sales for high-turnover items.

Comparisons do not include servers and storage at individual stores. Server-attached storage devices would be employed under both scenarios.

Strategic Benefit

Transition to the efficient scenario reduces IT costs, along with performance- and availability-related business costs by wide margins. Creation of a compact, centralized kernel of storage resources also reduces the costs and difficulties of, and accelerates completion of data center consolidation, server consolidation and application rationalization initiatives.

Key gains are also realized in the following areas:

- ***IT productivity***. Like the bank and manufacturing company, the retail company realizes significant improvements in IT productivity. Numbers of storage-related FTEs have been dramatically reduced. By minimizing underlying data complexities, particularly significant gains are made in database administration and application development productivity.

  More can thus be achieved by the company's application organization with the same number of personnel. A principal benefit will be to facilitate plans to use the core data warehouse as an aggressive tool to expand sales, cut costs and increase profitability.

  Demand for applications that exploit the data in this system has increased rapidly over the last few years, and there is already a large backlog of applications requested by around 250 headquarters and divisional users. It is expected that demand will escalate as access is extended to more than 2,000 store managers and specialists.

  An efficient storage infrastructure will thus not only result in faster processing and availability of information for users, but will enable their application needs to be met more rapidly. The bottom-line impact of these gains will clearly be substantial.

- ***Growth***. By minimizing complexities and increasing IT productivity, the storage infrastructure will assist not only in reducing costs, but also in maintaining growth while this occurs. Centralized storage resources enable the company to more easily integrate store acquisitions, and to more rapidly exploit economies of scale in purchasing and logistics for these than would be the case in an inefficient environment.

  The efficient infrastructure will also enable faster, more cost-effective expansion when economic conditions improve. It will, in particular, make the assimilation of major new acquisitions – including supply chain as well as store operations – a faster, more cost-effective process. The company's short- to medium-term infrastructure investments will place it in a stronger long-term position than many of its competitors.

- ***Real-time information***. The efficient storage infrastructure will be a critical enabler of the company's real-time plans. Processes of data collection, organization, and exploitation will be accelerated across all retail and supply chain operations.

  More than 100 applications – ranging from complex data mining processes to simple, ad hoc queries by managers and front-line personnel in stores and distribution centers – will be faster, more accurate and more effective because of the investment in storage infrastructures. Beneficiaries will include not only senior executives and staff specialists but also line managers and even sales associates, stock clerks and truck drivers.

  Achievement of corporate targets for cost control, profitability and growth is materially facilitated. The ability to deliver information on demand to large segments of its employee base also advances another strategic goal – that of "energizing" the company's workforce. There are few forms of empowerment more powerful than supplying information that enables individuals to make better, faster decisions.

As for the bank and manufacturing company, an important conclusion emerges. Not only would it be a more difficult, expensive and time-consuming process to achieve strategic goals using an inefficient, fragmented base of servers and storage – it would be impossible to do so.

Efficient Infrastructure

Under this scenario, five tiers of disk storage are implemented in the principal data center facility. These are as follows:

1. ***Critical systems 1*** are production systems most sensitive to downtime. They include mainframe-based core retail and merchandising systems, supply chain systems other than data marts and analytical tools, segments of the e-mail network, and strategic Internet and intranet applications, including the online grocery system.

   These are supported by high-end RAID 10 (mirrored) configurations equipped with tools for remote concurrent copy, point-in-time copy and core management functions.

2. ***Critical systems 2*** include corporate finance and HR systems, and less critical mainframe systems and intranets, along with development, test and other support systems for these and critical systems 1. These are supported by high-end RAID 5 configurations equipped with point-in-time copy and core management tools.

3. ***Business intelligence*** includes the data warehouse, loyalty program database, along with data marts and applications for financial, supply chain and marketing analysis. These are supported by high-end RAID 5 configurations equipped with core management tools. This is the largest tier, projected to increase from 43 terabytes at yearend 2002 to 318 terabytes by yearend 2007.

4. ***Secondary systems*** include sales and marketing, end-user support, departmental and other applications that run primarily on small RISC/UNIX and Intel-based servers. These are supported by midrange configurations equipped with point-in-time copy and core management tools.

5. ***NAS*** applications are supported by midrange configurations equipped with core management tools.

Initial (yearend 2002) disk storage bases and five-year growth projections for this scenario are summarized in figure 17. Volumes are raw capacity for IBM configurations.

Figure 17

Retail Efficient Scenario: Projected Disk Storage Growth

YEAREND	2002	2003	2004	2005	2006	2007
Critical systems 1*	22	27	32	37	46	55
Critical systems 2	18	22	25	29	33	39
Business intelligence	43	68	106	168	236	318
Secondary systems	8	12	18	22	26	30
NAS	4	5	7	8	9	11
Total (TB)	**95**	**134**	**188**	**264**	**350**	**453**

*Mirrored

Although the base will expand from 155 to 489 terabytes over a five-year period, storage-related staffing will remain stable, with 15 FTEs for IBM and 16 each for HDS and EMC systems. These include all categories of IT personnel affected by storage efficiency, broken out as shown in figure 18.

Figure 18

Retail Company Efficient Scenario: Staffing Detail

IT FUNCTION	EMC	HDS	IBM
Management/administration	1	1	1
Technical staff	–	–	–
Applications	2	2	2
Database administration	1	1	1
System/storage management	10	10	9
Network management	–	–	–
Operations & support	2	2	2
End-user computing & help desk	–	–	–
TOTAL	**16**	**16**	**15**

Under the inefficient scenario, equivalent staffing increases from 74 in 2003 to 81 in 2004, before dropping to 54 in 2007. The decline reflects personnel savings through aggressive SAN deployment. By 2007, all servers and storage systems in the data centers are SAN-attached. It is unlikely that FTE staffing would decline further after this point.

Vendor-specific Costs

Five-year costs for high-end and midrange disk storage hardware, maintenance and software, along with costs for SAN and NAS, personnel and facilities for this infrastructure are presented for EMC, HDS and IBM in figure 19.

For reasons outlined earlier, certain of the overall costs presented here – hardware, software licenses, SAN, NAS and facilities – are the same for all vendors, while maintenance, software support and personnel costs are vendor-specific.

Figure 19

Retail Efficient Scenario: Overall Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	4,385	823	1,148	2,824	1,041	952	**11,173**
Maintenance	–	–	–	338	396	476	**1,210**
Software licenses	983	72	–	87	145	101	**1,388**
Software support	–	111	156	158	168	188	**781**
SAN/NAS acquisition	1,045	43	94	87	283	181	**1,733**
SAN/NAS support	–	–	38	39	43	46	**166**
Personnel	–	1,452	1,452	1,452	1,452	1,452	**7,260**
Facilities	–	55	77	108	143	185	**568**
Total ($000)	**6,413**	**2,556**	**2,965**	**5,093**	**3,671**	**3,581**	**24,279**
HDS							
Hardware	4,385	823	1,148	2,824	1,041	952	**11,173**
Maintenance	–	–	–	–	114	131	**245**
Software licenses	983	72	–	87	145	101	**1,388**
Software support	–	91	128	131	139	155	**644**
SAN/NAS acquisition	1,045	43	94	87	283	181	**1,733**
SAN/NAS support	–	–	38	39	43	46	**166**
Personnel	–	1,452	1,452	1,452	1,452	1,452	**7,260**
Facilities	–	55	77	108	143	185	**568**
Total ($000)	**6,413**	**2,536**	**2,937**	**4,728**	**3,360**	**3,203**	**23,177**
IBM							
Hardware	4,385	823	1,148	2,824	1,041	952	**11,173**
Maintenance	–	–	–	–	105	122	**227**
Software licenses	983	72	–	87	145	101	**1,388**
Software support	–	–	–	–	46	50	**96**
SAN/NAS acquisition	1,045	43	94	87	283	181	**1,733**
SAN/NAS support	–	–	38	39	43	46	**166**
Personnel	–	1,367	1,367	1,367	1,367	1,367	**6,835**
Facilities	–	55	77	108	143	185	**568**
Total ($000)	**6,413**	**2,360**	**2,724**	**4,512**	**3,173**	**3,004**	**22,186**

Costs for the midrange component of the scenario are presented for EMC, HDS, IBM and HP in figure 20.

Figure 20

Retail Company Efficient Scenario: Midrange Vendor Cost Comparisons

YEAREND	2002	2003	2004	2005	2006	2007	TOTAL ($000)
EMC							
Hardware	232	44	32	23	18	18	**367**
Maintenance	–	–	–	16	19	21	**56**
Software licenses	60	–	–	–	–	–	**60**
Software support	–	7	9	9	9	9	**43**
Total ($000)	**292**	**51**	**41**	**48**	**46**	**48**	**526**
HDS							
Hardware	232	44	32	23	18	18	**367**
Maintenance	–	–	–	–	13	16	**29**
Software licenses	60	–	–	–	–	–	**60**
Software support	–	5	7	7	7	7	**33**
Total ($000)	**292**	**49**	**39**	**30**	**38**	**41**	**489**
IBM							
Hardware	232	44	32	23	18	18	**367**
Maintenance	–	–	–	–	6	7	**13**
Software licenses	60	–	–	–	–	–	**60**
Software support	–	–	–	–	3	3	**6**
Total ($000)	**292**	**44**	**32**	**23**	**27**	**28**	**446**
HP							
Hardware	232	44	32	23	18	18	**367**
Maintenance	–	–	–	–	15	18	**33**
Software licenses	60	–	–	–	–	–	**60**
Software support	–	7	9	9	9	9	**43**
Total ($000)	**292**	**51**	**41**	**32**	**42**	**45**	**503**

Midrange hardware and software license costs are again the same for all vendors, while maintenance and software support costs are vendor-specific. SAN, NAS, personnel and facilities costs are not broken out separately, as they would be determined by overall installations of high-end as well as midrange systems. They would in practice be similar for all four vendors.

TECHNOLOGY VIEW

State of Storage

General Picture

In preparing this document, storage and IT management functions in 124 large companies in North America were surveyed. Among the survey objectives was to document enterprise-level experiences with storage in general, and storage consolidation and rationalization initiatives in particular.

Most of these companies had begun or planned to implement SANs and SRM software. More than 43 percent had undertaken, or planned to undertake storage consolidation exercises. Benefit experiences and expectations were generally similar. Most initiatives were, however, clearly at an early stage. Respondents cited lack of funds and difficulties in quantifying benefits as factors limiting more aggressive adoption.

One benefit of upgrading storage infrastructures was reported to include improved capacity utilization, leading to reduced purchases of capacity. Another, lower costs for administrative personnel, was cited – particularly in that it was possible to support future storage growth without adding staff. SANs had also reduced, or were expected to reduce tape capacity and costs.

Reduction of total cost of ownership (TCO) and realization of return on investment (ROI) were frequently cited as benefits. It was, however, clear that performance, availability, backup and recovery capability, and other service quality benefits were equally or more important factors in decisions to implement SANs, SRM deployment, consolidation or combinations of these.

A number of other commonalities in responses should be mentioned. High-end and midrange storage systems coexisted in most organizations. High-end systems were typically employed for major data center applications. Midrange systems were the norm for smaller, less business-critical applications, particularly when servers were located at remote sites or regional centers.

A continuing role for direct-attached storage (DAS) devices was also recognized. These were seen as appropriate for bank branches, retail outlets, local sales and administrative offices, and other sites with only a few servers and relatively small volumes (less than 100 gigabytes).

The issue of SANs "versus" NAS appears to have largely disappeared. These were typically seen as complementary rather than competitive. The focus in most companies was on SANs, with a secondary role seen for NAS in file-serving applications.

Survey results confirmed general industry perceptions about storage trends in large organizations. But a number of issues emerged which suggest that the picture is more nuanced than is often recognized. These included the following.

Specialization Issues

In most IT organizations, storage administration tasks are undertaken by individuals with titles other than storage administrator (e.g. system administrator, operations, technical support, platform specialist), as part of a broader job mix. The primary job responsibilities and skill sets of these individuals are not storage-specific. This is particularly the case for less business-critical applications running on small RISC/UNIX and Intel-based servers.

Survey responses, coupled with searches of IT recruitment sites, identified more than 60 job titles – listed in figure 21 – which involve dealing with storage resources in one way or another.

Figure 21

Representative Storage-related Job Titles

Application administrator	End-user support	SAN administrator
Application developer	Hardware planner	SAN architect
Archiving specialist	Information security analyst	SAN engineer
Asset management specialist	Information security specialist	SAN IP network specialist
Availability management	LAN management	Security architect
Backup/recovery specialist	LAN specialist	Server administration
Business continuity specialist	Media management	Service-level management
Business systems analyst	Network administrator	Software engineer
Capacity manager	Network analyst	Storage administrator
Capacity analyst/planner	Network engineer	Storage engineer
Change management specialist	Network management	Storage manager
Computer operator	Network operations	System administrator
Computer specialist	Operations analyst	Systems architect
Configuration management	Operations management	Systems analyst
Database administrator	Performance analyst	Systems developer
Database analyst	Problem management	Systems engineer
Database performance and tuning	Production control	Systems programmer
Data management specialist	Production support	Systems support specialist
Data security specialist	Programmer	Technical support
Data warehouse administrator	Programmer/analyst	UNIX specialist
Disaster recovery specialist	Quality assurance	Windows specialist

Unsurprisingly, it is difficult to quantify savings in personnel costs only in terms of storage administrators. The real impact of storage efficiency – or inefficiency – is clearly a great deal broader.

Lack of specialization extends to such functions as architecture, planning, financial management, service-level management and quality assurance. Only eight of the companies surveyed had functions within IT organizations that were responsible for storage strategy and practices at the enterprise level. In most cases, only procurement was coordinated on a company-wide basis.

Little attention is, similarly, paid to the development and implementation of storage-specific “best practices.” Storage management is often not seen as a significant discipline, or is treated as a subset of such specializations as system management, network management, operations management or data center management.

Lack of specialization limits the sophistication applied to the development of storage strategy. One effect is that, in at least some organizations, there is tendency to adopt product-centric approaches – such as deployment of SRM software, or installation of SANs – on the assumption that these will, in themselves, deliver far-reaching benefits.

Experience with such approaches has been mixed. In several companies, SRM deployment had failed to yield expected savings. High software implementation overhead, extensive training requirements and protracted “learning curves” for administrators had led to disenchantment with the SRM concept.

Others cited disappointingly low returns, and unexpectedly high installation and support costs for SANs. In one case, deployment of a SAN in a heterogeneous, poorly optimized environment characterized by large numbers of servers and storage systems had increased complexities, and caused higher rather than lower costs. Many other such examples can be cited.

In these and other companies, attempts had been made to deploy specific technologies without adequately addressing the broader environment in which they would operate.

Significant conclusions emerge. The highest-yield strategies address multiple rather than single parameters of efficiency – including consolidation, network integration, and automation – with the objective of maximizing the performance, service quality and cost-effectiveness of the entire infrastructure. Such strategies must, moreover, address specialization issues – including skills development, management practices, and cultural change – as well as technical challenges.

Cost/Benefit Quantification

Many of the companies surveyed had clearly had trouble in quantifying the benefits of storage initiatives. This is the case even where apparently successful deployments of SANs and SRM, and consolidation initiatives had been completed.

The reasons for this situation include:

- ***Focus on personnel savings***. A narrow focus on reducing storage administrator head counts inevitably restricted savings projections. One respondent commented, for example, that even if all posts with the title "storage administrator" in their organization were eliminated, savings would still not be sufficient to justify a major initiative at a time of budgetary constraints.

  Justifications based on personnel cutbacks were opposed for other reasons. In one case, it was noted that there was a critical need for better-qualified, more skilled storage administration specialists. It made little sense to target removal of precisely this category of personnel.

  In another company, management was concerned to retain skilled and experienced staff familiar with the company's operations. It was decided that cost-justification had to be based on more substantive criteria of business benefit.

- ***Service quality limitations***. Many companies had trouble determining the value of gains in service quality – including such variables as availability, response time, application-level throughput time, security and backup/recovery capability.

  The impact of outages on business-critical transaction-processing systems might, for example, be quantified in terms of "lost revenues," but the number of systems to which such quantification could be easily applied was limited. There was a great deal of uncertainty about how to quantify the effects of outages, data loss or performance shortfalls for others.

  Quantification techniques were often inadequate even for obviously business-critical systems. One manufacturing company had, for example, decided not to invest in real-time replication and failover capability because the effects of "a couple of days" of downtime for its global ERP system were not seen as serious. This decision was made because the company's IT organization had failed to produce an adequate ROI justification.

  This is a startling point of view. There is a large body of evidence that in tightly integrated, low-latency supply chains, even short outages have major repercussions. These may include lost sales, late delivery penalties, increased inventory costs, idle capacity costs and a range of other bottom-line impacts.

  In this company – and clearly in many others – there was a fundamental disconnect between definitions of value employed for IT decision-making, and those which formed the basis of investment strategies elsewhere within the business.

- ***Workload complexity***. Another obstacle to effective cost/benefit quantification was that storage environments were – even after rationalization – relatively complex. Even where numbers of systems and interconnection complexity were reduced, it was still necessary to support a wide range of applications with different workload characteristics, and varying performance and service quality requirements.

  This was particularly the case in large data centers supporting diverse application portfolios and server bases. In one company, for example, it was estimated that storage resources had to support more than 20,000 separate processing jobs in the course of an average 24-hour day.

  It was noted, similarly, that the practice of quantifying administration overhead in terms of storage capacity (e.g. one FTE per four terabytes) could be highly misleading. It was necessary to manage storage resources, and maintain appropriate performance and service levels on an application-by-application basis. The application portfolio to be supported, as well as overall capacity, had a major impact on administrative costs and difficulties.

  Equally, the proposition that storage administration costs are significantly higher than hardware costs (most estimates are in the range of five to eight times higher) is true in general terms. But in data center environments, there are wide variations in cost structures on an application-by-application, system-by-system basis. Any attempt to reduce overall administration costs must, again, involve a great deal of granularity.

Three conclusions emerge. First, cost/benefit justifications for storage initiatives will be most compelling – and most accurate – if they are addressed at the infrastructure level, dealing with aggregated store resources supporting overall application portfolios. Second, justifications must be based on substantive criteria of business benefit rather than simple IT cost savings.

Finally, no single technology or technique guarantees success. Creation of an effective storage infrastructure – and realization of the full benefits that it can deliver – requires a sophisticated, enterprise-level approach that combines high levels of granularity in meeting performance and service quality requirements with high levels of system-wide integration and optimization.

Measuring Efficiency

Measurement Challenges

A variety of techniques are available to measure the performance, effectiveness and cost of storage systems. Most, however, possess inherent limitations. Performance metrics, for reasons discussed in more detail below, tend be highly application- and workload-specific, and offer imprecise guides to how well any given system will meet the requirements of any particular user.

Simple price metrics such as cost per megabyte ($/MB) or gigabyte ($/GB) offer at best approximate indicators of overall costs. System-level costs do not necessarily track to costs of disk capacity, particularly when software, maintenance, personnel and other categories are taken into account. It is generally recognized that hardware represents only a small component of overall storage costs.

TCO metrics may be useful in general terms, but do not allow for workload and service quality variations, or at least do so only in a rudimentary manner. TCO calculation is only valuable if configurations, service levels, staffing and other variables have been reliably determined – and realistically projected over, say, a three- or five-year period – based on organization-specific application profiles, workload profiles and service-level requirements.

The major challenge is not to measure prices or costs – but to measure the value an IT resource brings to the business. In selecting storage solutions, the challenge is to determine the relative value of different options. This means employing composite metrics that deal not only with technology costs and functionality, but also with the business implications of achieving – or failing to achieve – required service quality levels.

The term "efficiency index" is employed here to describe this set of metrics, which includes the variables summarized in figure 22.

Figure 22
Representative Efficiency Variables

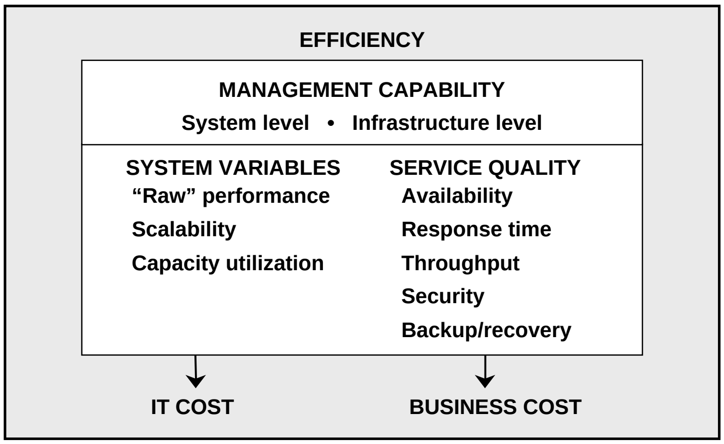

Broader effects such as the degree to which storage resources facilitate speed and quality of application delivery, increase the effectiveness with information is used for business advantage, or enable real-time competitive models may also be included in the definition of "efficiency."

Performance Variables

Performance should obviously be a factor in the selection of storage systems, in the sense organizations should verify that systems are capable of meeting their needs. But comparative measurements of the performance of different vendors' systems are proving to be difficult.

One reason for this is that storage workloads are highly diverse. There are, at the most basic level, differences in metrics between online (response time) and batch (elapse time) processes. The composition of workloads for any specific application or application suite may vary widely.

Figure 23 shows examples for major ERP systems.

Figure 23

Representative ERP Workload Breakdowns

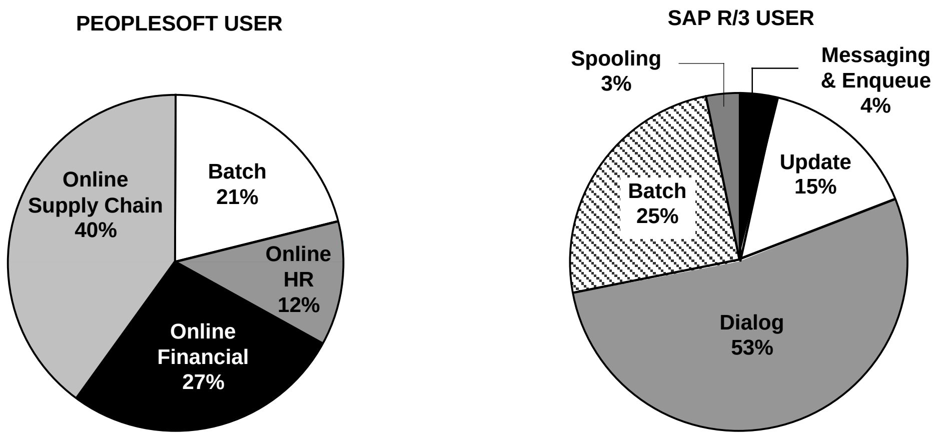

With these and other types of system, workload mixes tend to change substantially over time due normal time-of-day or periodic fluctuations, or because of workload spikes. Unpredictable peaks are particularly common for online systems, but may also be experienced by conventional systems because of marketing campaigns, promotions and a wide range of other events.

There is, moreover, the greater challenge of measuring end-to-end performance for a particular unit of work, such as completing a transaction, responding to a query, or forwarding and delivering a message. This may involve processes occurring on multiple front-end and back-end servers, which may be supported by different storage systems.

Metrics such as application-level throughput (also referred to service time) may be useful in measuring end-to-end performance. But they are highly specific, and it is difficult to generalize from them to predict performance with other types of workload, or in mixed-workload environments. This is, for example, the case for generalized benchmarks such as those developed by the Storage Performance Council.

Attempts to compare performance based on single component-level metrics (e.g. HDS' emphasis on internal bandwidth) or techniques (e.g. EMC's emphasis on direct matrix architecture as an alternative to crossbar switching) also suffer from inherent limitations. A wide range of factors determines the overall performance of a storage system. These include disk-level performance, cache and RAID efficiency, internal and I/O bandwidth, intrasystem routing mechanisms and others.

Overall system-level performance is, moreover, materially influenced by the effectiveness with which component-level capabilities are managed and optimized by microcode, operating systems and software-based tools. Allowance must, for example, be made for the speed at which systems copy for concurrent and point-in-time applications, particularly when these involve large volumes of data.

System-level performance leadership would in any case be moot. The trend in most large organizations is toward storage networking and distribution of workloads across multiple systems. Variables such as SAN bandwidth, switching efficiency and fabric management capability enter the picture. Cross-system mechanisms for such functions as load balancing and workload management will also play an essential role in overall performance.

Performance is a significant variable of storage efficiency, but it is only one of a number of such variables. System-level performance measurements may assist in developing vendor shortlists. But they are unlikely to provide sufficient precision to decide between comparable offerings from leading vendors such as EMC, HDS, IBM and (for midrange systems) Hewlett-Packard.

To distinguish between these, the ante must be raised. Decisions should be based on the ability of vendor offerings to meet the full range of organizational storage requirements in a networked environment over time.

Scalability

Scalability may be defined in terms of the maximum workload size that may be realistically handled by a single storage system (“vertical” scalability) or networked complex of systems (“horizontal” scalability). “Realistically” in this context means without unacceptable degradation in performance or service quality.

Workload size, it should be noted, does not necessarily equate to storage volume. From a scalability point of view, the workload impact of, say, an I/O-intensive business-critical transaction-processing system with 800 gigabytes of data may be greater than that of a multiple-terabyte data warehouse.

Most large organizations will require high levels of horizontal as well as vertical scalability. Although single disk storage systems with capacities of hundreds of terabytes will be technically feasible within the next few years, the potential for bottlenecks in systems of this size is still unclear. Most organizations will, moreover, tend to avoid over-dependence on single systems in case these should fail.

The optimum mix of vertical and horizontal scalability will vary between applications and organizations. But in large data center environments, the most important differentiator will be the extent to which growth in workloads and storage volumes can be distributed and managed across a networked complex. The overall infrastructure, rather than individual storage systems, will be the key determinant of scalability.

Capacity Utilization

Capacity utilization may be defined as the proportion of system processing power, switching capability and storage capacity that may be realistically employed in production operations (“realistically” means the same as above).

In most organizations, only a fraction of capacity is used. This is, in many cases, because storage resources are dedicated to individual servers and are configured to handle peak loading for that server only. Overall utilization may thus be improved by implementing SANs.

Other factors should, however, also be taken into account. Data may be unnecessarily duplicated across applications and servers (in which case consolidation will reduce overall capacity requirements), or may be old and infrequently accessed (in which case tools and procedures which move inactive data to tape or optical media may be used to up additional disk capacity).

Similarly, administrators may overconfigure systems to minimize risks of bottlenecks. Where this occurs, significant increases in utilization will require highly granular, highly effective management of resources occurring in real time.

As for performance and scalability, the most important determinant of capacity utilization will thus tend to be the efficiency of the overall storage complex. The largest gains will come from the efficiency with which workloads are distributed, and resources allocated and managed across the infrastructure as a whole.

Service Quality

Service quality may be considered to include the following :

- ***Availability*** means the role of storage resources in delivering IT service over time so that users experience no outages. This means not only minimizing risks of hardware or software failure, operator error and other disruptive events, but also ensuring that failover processes occur in a manner that ensures that no disruption of service is visible to users.

  Avoidance or (if an application does not require 24x7x365 user access) minimization of planned outages for hardware or software upgrades, scheduled maintenance, data replication (e.g. for batch backups, reporting and data warehouse loads) and other such functions also forms an important part of the availability equation.

- ***Response time*** means the role of storage resources in ensuring that a system responds to user actions fast enough to assure that the user's train of thought is not broken, productivity is not impaired, and business is not lost because internal users, partners or customers become frustrated and log off.

- ***Application-level throughput*** means the role of storage resources in determining the end-to-end time required to complete a unit of work, such as processing a transaction, responding to a customer request, fulfilling an order, or delivering an electronic message. It impacts a wide range of applications, and the contributions these make to business competitiveness and costs.

- ***Security*** means the role of storage resources in preventing unauthorized access to system resources and data, and – if such access does occur – blocking intrusion, containing damage, and identifying sources and methods of penetration.

- ***Backup/recovery*** capability means the contribution of storage resources maintaining continuity of business operations without loss or corruption of records, or disruption in access to these in case of an outage. In practice, it may be considered to equate to disaster recovery (i.e. the ability to maintain continuity in case of an exceptionally severe outage disabling major systems or entire data centers).

All of these variables map to one or more "hard" measures of business costs.

Avoidance of outages, and rapid response and application-level throughput in customer-facing systems, for example, translates into less lost revenue and fewer lost customers. The bottom-line impact of the latter may be particularly significant if customer lifetime value (CLV) or equivalent metrics are employed. Examples of these effects are given in the diversified bank profile, and in business cost calculations for online grocery systems in the retail company profile.

The internal effects of outages and performance shortfalls may be quantified using a variety of industry and process-specific metrics, or broader measures such as the impact of delays and disruptions on supply chain efficiency. Examples of these are given in the manufacturing company and retail company profiles.

The value of backup/recovery is in preventing catastrophic shutdowns in business operations for protracted periods. The effects of these may, again, be quantified using industry norms for each type of incident. Examples of such calculations are given in all three profiles.

The value of security includes avoidance of electronic theft and fraud, costs of sabotage and virus damage, potential legal and human resources costs, and effects of negative publicity. These are not discussed in detail in this document.

In a complex IT environment, it may not always be possible to directly quantify storage-specific business costs. Where this is the case, costs may be estimated based on the contribution that storage resources make to overall service quality levels (e.g. storage accounts for 15 percent of outages, servers account for 20 percent, applications account for 30 percent). This approach is employed in the profiles presented in this document. Overall business costs due to service quality shortfalls would be significantly greater.

Real-time Implications

The quantification techniques outlined here tend to understate rather than overstate both the IT cost and business cost implications of storage efficiency or inefficiency.

The bottom-line effects of being able to make better decisions faster have, for example, not been addressed in detail. Yet, the cumulative impact of increased information timeliness and quality on the hundreds to thousands of decisions that are made every day, at all organizational levels, would clearly be massive in most companies.

It would, in particular, be difficult to quantify the bottom-line contribution that an efficient storage infrastructure can make to enabling real-time competitiveness. Its business value can, however, be simply stated. Without an efficient storage infrastructure, it would be virtually impossible to achieve real-time competitiveness at the enterprise level.

Storage Management

Enterprise Challenges

Management capability is emerging as an increasingly important factor in the selection of storage solutions. Recent industry debates have tended to focus on such themes as SRM software, virtualization and SAN management. The challenges of enterprise storage management are, however, substantially broader than all of these.

Effective management of organizational storage resources will require capabilities implemented at multiple levels, including middleware solutions; cross-system infrastructure capabilities such as SAN file systems, virtualization and device resource management; and functions built into SAN hardware and software, and underlying systems and devices. Figure 24 illustrates these.

Figure 24

Storage Management Levels

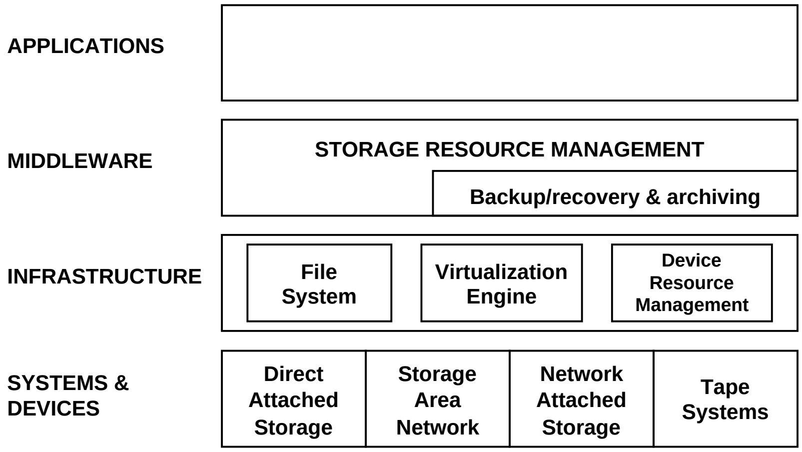

There is widespread agreement about the general objectives of storage management. They include the ability to maximize service quality, minimize technical complexities, and increase cost-effectiveness by optimizing capacity utilization and reducing administrative overheads.

An important goal is to reduce the amount of time administrators must spend in relatively low value-added manual tasks. Over time, basic functions such as monitoring, identifying and resolving problems, distributing workloads, allocating and re-allocating capacity, and maintaining service level compliance would be progressively automated. This would also, increasingly, be the case for more sophisticated analytical, predictive, capacity planning and scheduling functions.

At present, however, the processes of defining, selecting solutions for and implementing storage management are somewhat less clear-cut that those for SANs or storage systems. Multi-level approaches, in particular, are comparatively new, and there is not a great deal of industry experience with implementation and operation of these.

Vendor Strategies

All of the major disk storage vendors have defined strategies that address all of these levels.

EMC's AutoIS strategy is built around its CommandCenter framework and software tools, and Symmetrix DMX, Clariion, Connectrix and Celerra platforms. The company has acquired a series of specialist storage management software companies since 1999 to build up its offerings, and plans to continue this process.

HDS' equivalents are the TrueNorth strategy, HiCommand framework and software tools, and Freedom platforms. The company expects to incorporate IBM block-based virtualization technology into future storage management capabilities.

IBM strategy focuses around its Tivoli-branded family of SRM solutions, into which software technology from recently acquired TrelliSoft will be incorporated, its ESS and FAStT disk platforms, and its Virtual Tape Server (VTS) and tapes systems. The company began delivery of a virtualization engine in early 2003, and its Storage Tank SAN file system is expected to follow later in the year.

Hewlett-Packard’s strategy includes plans to integrate its Continuous Access Storage Appliance (CASA) offerings, which originated with its 2001 acquisition of StorageApps with former Compaq VersaStor virtualization technology. SRM solutions will be offered as part of the company’s OpenView management suite.

A number points of commonality between vendor offerings are emerging. These will include support for the principal Storage Networking Industry Association (SNIA), and interoperability with de facto standards such as Brocade SANs and major third-party tools for SRM, backup/recovery and other applications. It can be expected that most user environments will, at least in large organizations, require significant multivendor interoperability.

It would be outside the scope of this document to attempt to compare the competitive functionality these strategies will deliver. It would also be somewhat premature, in that they are far from completion. But certain conclusions can be drawn about the capabilities these vendors have implemented, and can be expected to implement in underlying systems and at the infrastructure level.

From this perspective, there are clear indications that IBM is developing a lead in two areas of technological differentiation: designed-in capabilities, and autonomic computing.

Designed-in Capabilities

The purpose of storage management is not simply to reduce administrative costs. An effectively managed infrastructure will maximize performance, scalability and capacity utilization. It will also contribute materially to all of the principal variables of organization-wide service quality.

The extent to which a storage system delivers multiple-parameter efficiency will depend largely on how well core system designs and hardware and software components are optimized to achieve key goals. If capabilities are not “designed in,” it will be necessary to implement them through software overlays, hardware add-ons, or both. Embedded inefficiencies will be created which will tend to reduce performance, undermine service quality optimization and increase costs.

IBM has established an industry-track record in building such functions as system management, workload balancing and security, along with reliability, availability and serviceability (RAS) features into underlying hardware, microcode and systems software across its eServer product line. Key benefits have included the company’s ability to draw upon mainframe manageability and service quality strengths and the capabilities of its vertically integrated microelectronics organization.

Comparable capabilities have been delivered, and will clearly continue to be delivered in IBM storage platforms and software. EMC, HDS and Hewlett-Packard have implemented many of the same functions. The breadth of IBM implementation, and the technical sophistication of the company’s approach are, however, significantly greater.

IBM possesses a number of significant advantages in this area. One is the extent of its resources not only in storage hardware and software, but also in servers, middleware and microelectronics. The company is able to realize high levels of technology sharing across its different platforms and lines of business, as well as major economies of scale in development and engineering.

In comparison, EMC is a specialist vendor of disk storage solutions. The company retains a strong market position, and continues to invest heavily in technology, particularly in software development (which, according to management, accounts for around 75 percent of its R&D spending) and acquisition of software companies. But its microelectronics capability is limited, and the company’s focus on developing its software business suggests that hardware-based functionality is a lower priority.

Hewlett-Packard and Hitachi are server vendors, and Hitachi is a major semiconductor manufacturer. IBM has, however, since the mid-1990s realized significantly higher levels of development and engineering integration between its storage, server, software and microelectronics businesses.

In comparison, Hitachi's decentralized overall structure is less conducive to cross-divisional technology transfers. It is likely that Hewlett-Packard will for, for at least the next few years, remain preoccupied with integrating former Compaq platforms. It can be expected that IBM will maintain and probably improve its leadership position in designed-in capabilities.

Autonomic Computing

A second – and in the long term, more important – area of IBM technological differentiation is the company's development of autonomic computing capabilities.

The IBM autonomic computing program extends across all of its major lines of business. It was formally launched in early 2001 as Project eLiza, although some development activities were initiated earlier. Since that time, autonomic computing has been the subject of more than $1 billion in IBM R&D investment. More than 100 significant autonomic features have been introduced into the company's server, storage and software product lines.

Autonomic computing, in the IBM definition, involves applying advanced artificial intelligence technologies to IT resource management and service quality optimization tasks. These are being embedded both into underlying hardware and systems software platforms, and into higher-level middleware such as the company's DB2 database and Tivoli system, storage, application and network management tools.

Four areas of capability are targeted: self-optimization (systems will adjust parameters to achieve optimum performance), self-healing (identification and resolution of problems without human intervention), self-protection (including security as well as availability management, backup and recovery functions) and self-configuration (automated start-up and ongoing change management).

Autonomic technologies will be progressively phased into IBM storage as well as server offerings to achieve five progressive levels of capability, which are summarized in figure 25.

Figure 25

Levels of Autonomic Capability

LEVEL 1 BASIC	LEVEL 2 MANAGED	LEVEL 3 PREDICTIVE	LEVEL 4 ADAPTIVE	LEVEL 5 AUTONOMIC
Characteristics				
Multiple sources of system-generated data	Consolidation of data & actions through management tools	System monitors, correlates & recommends actions	System monitors, correlates & takes action	Integrated components dynamically managed by business rules/ policies
IT Staff Role				
Performs wide range of tasks	Analyzes & takes actions	Approves & initiates actions	Manages performance against service level agreements (SLAs)	Focuses on enabling business needs

Application of autonomic computing capabilities to the IBM ESS, FAStT and VTS platforms has already begun. Current design capabilities, as described by IBM, are summarized in figure 26.

Figure 26

Storage Applications of Autonomic Computing: Current Design Capabilities

SELF-CONFIGURATION	SELF-HEALING	SELF-OPTIMIZATION	SELF-PROTECTION
Platform: Enterprise Storage Server			
Automatic reconfiguration when adding capacity Capacity upgrade on demand CIM enablement for remote monitoring, configuration & control	Hardware-based features RAID auto rebuilds Predictive failure analysis with automatic correction Failover/failback Battery-backed cache Non-volatile store (NVS) write cache Multiple RAID levels	z/OS Parallel Access Volumes z/OS I/O priority Tag command queuing Clustered SMP Multi-path load balancing	LUN masking & partitioning Zoning Configuration version control GDPS support
Platform: FAStT Systems			
Dynamic cache configuration Dynamic volume expansion	Dual hot-swap RAID Controllers Battery-backed cache Multiple RAID levels Call home	Dynamic RAID level migration & segment size changes Performance monitor	LUN masking & partitioning Zoning
Platform: Virtual Tape Server			
Capacity upgrade on demand Automatic file access selection	Hot-swappable disks Automatic self-assessment/ corrective action Media management Automated media recovery Call home	Cache manager dynamic priority assignment Advanced policy-based management Peer-to-peer workload balancing	Security & privacy service Automatic switch password updating Configuration version control

CIM: Common Information Model
GDPS: Geographically Dispersed Parallel Sysplex
LUN: Logical Unit Number

Potential future applications are shown in figure 27.

Figure 27

Storage Applications of Autonomic Computing: Future Design Possibilities

SELF-CONFIGURATION	SELF-HEALING	SELF-OPTIMIZATION	SELF-PROTECTION
Storage Tank			
CIM enablement Policy-based management provisioning & placement Non-disruptive LUN addition	Automated failover on node outage Automated spillover on pool full	Policy-based file group placement Balanced pool-by-pool LUN usage Transparent data movement between volumes & pools	Standard UNIX & Windows file system security Standard authentication schema
Storage Virtualization			
CIM enablement Non-disruptive LUN changes Central point of control	Automated failover on node outage Integrated multi-path selection Automated readmission to cluster of failed controllers	Load balancing across multiple paths	Asynchronous concurrent copy
Device Resource Management			
Aggregation of configuration changes for multiple devices LUN provisioning analysis	Central monitoring & logging Cross-device heterogeneous problem determination Detection & repair of SAN component incompatibility	Central performance monitoring Bottleneck analysis & LUN allocation recommendations	Replication services policies Event log audit trail Automated create replica on addition of new LUN

CIM: Common Information Model
LUN: Logical Unit Number

Although the overall IBM schedule for incorporating autonomic capabilities into its storage offerings is unclear, it can be expected that most if not all of the company's goals will be achieved within the next five years. These capabilities will thus affect the manageability and cost structure of IBM storage offerings relative to competitors within the planning timeframe of most large companies.

Existing autonomic capabilities are largely responsible for the lower IBM FTE staffing levels relative to EMC and HDS systems in the efficient scenarios of all three company profiles presented in this document. Five-year staffing for IBM systems for these scenarios is based on 2003 levels of autonomic implementation in the ESS and FAStT product lines. In practice, it could be expected that disparities between IBM and competitive staffing would increase significantly during this period.

DETAILED DATA

Vendor Comparisons

Detailed breakdowns of vendor-specific IT costs for efficient scenarios are shown in figures 28 and 29. These correspond to summary data presented earlier in figures 3 and 4 respectively.

Figure 28

IT Costs for Vendor-specific Options: Efficient Scenarios

CATEGORY	DIVERSIFIED BANK			MANUFACTURING COMPANY			RETAIL COMPANY		
	EMC	HDS	IBM	EMC	HDS	IBM	EMC	HDS	IBM
Hardware	19,939	19,939	19,939	17,755	17,755	17,755	11,173	11,173	11,173
Software licenses	3,983	3,983	3,983	5,414	5,414	5,414	1,388	1,388	1,388
SAN/NAS	2,931	2,931	2,931	3,222	3,222	3,222	1,899	1,899	1,899
Facilities	926	926	926	695	695	695	568	568	568
Total ($000)	**27,779**	**27,779**	**27,779**	**27,086**	**27,086**	**27,086**	**15,028**	**15,028**	**15,028**
Maintenance	2,785	581	534	2,492	518	487	1,210	245	227
Software support	1,911	1,841	224	2,146	1,735	213	781	644	96
Personnel	18,070	17,115	15,815	8,005	8,005	7,580	7,260	7,260	6,835
Total ($000)	**22,766**	**19,537**	**16,573**	**12,643**	**10,258**	**8,280**	**9,251**	**8,149**	**7,158**
TOTAL ($000)	**50,545**	**47,316**	**44,352**	**39,729**	**37,344**	**35,366**	**24,279**	**23,177**	**22,186**

Figure 29

IT Cost Breakdowns: Midrange Component of Efficient Scenarios

CATEGORY	DIVERSIFIED BANK				MANUFACTURING COMPANY				RETAIL COMPANY			
	EMC	HDS	IBM	HP	EMC	HDS	IBM	HP	EMC	HDS	IBM	HP
Hardware	820	820	820	820	507	507	507	507	367	367	367	367
Software licenses	76	76	76	76	44	44	44	44	60	60	60	60
Total ($000)	**896**	**896**	**896**	**896**	**551**	**551**	**551**	**551**	**427**	**427**	**427**	**427**
Maintenance	129	66	28	74	75	39	17	43	56	29	13	33
Software support	53	43	8	53	33	24	4	33	43	33	6	43
Total ($000)	**182**	**109**	**36**	**127**	**108**	**63**	**21**	**76**	**99**	**62**	**19**	**76**
TOTAL ($000)	**1,078**	**1,005**	**932**	**1,023**	**659**	**614**	**572**	**627**	**526**	**489**	**446**	**503**

Diversified Bank

Inefficient Scenario

This scenario is based on the following groups of applications and platforms:

- ***Mainframe systems***. These include core transaction-processing systems for banking, brokerage, insurance and other lines of business. Branch banking, ATM/POS and electronic funds transfer (EFT) systems are also mainframe-based. The company initially – as of yearend 2002 – operates 13 IBM- and IBM-compatible mainframe processors with a total of 14,972 MIPS and 72.2 terabytes of disk storage.

- ***Corporate data warehouse***. This is a company-wide system used for customer profiling, market segmentation, data mining and other analytical applications. This initially runs on multiple IBM SP configurations with 15 frames, 240 nodes and 56.6 terabytes of disk storage.
- ***Major data marts and analytical systems***. Data marts are employed for channels analysis, call detail analysis, commercial banking, credit card, credit scoring, market research, private client services and other applications. These initially run on 21 high-end and midrange HP, IBM, Sun and Silicon Graphics RISC/UNIX servers with 19.9 terabytes of disk storage.

  A PeopleSoft Enterprise Performance Management (EPM) system is also employed for company-wide financial analysis. The system initially runs on two IBM SP frames with 13 terabytes of disk storage.
- ***RISC/UNIX business systems***. These include 39 systems running on 230 high-end and midrange HP, IBM, Sun and Compaq RISC/UNIX servers with 81.2 terabytes of disk storage.

  Applications include retail, commercial, small business, commercial and corporate banking; consumer loans, money store and mortgage operations; brokerage, insurance and asset management; capital markets; online banking and brokerage; and HR and financial systems.
- ***Small RISC/UNIX servers***. Initially approximately 540 small RISC/UNIX servers from a range of vendors with a total of 19.4 terabytes of disk storage are employed. These are employed in application serving, departmental, Internet and intranet infrastructure, networking, call center, image and document management, workstation serving (primarily in the Capital Markets division) and other roles.
- ***Intel-based servers***. At the beginning of the period, the company has at least 1,625 Intel-based servers running Microsoft Windows and NT, NetWare, Linux and other operating systems with at least 38.6 terabytes of disk storage.

  Intel-based servers are variously employed for e-mail (primarily Microsoft Exchange, but other systems are operated by some smaller business units), departmental applications, end-user file and print services, Web and intranet serving, Citrix serving (for around 15,000 remote branch users), test and development, and other applications.

All mainframe systems and storage, along with most RISC/UNIX servers, are located in the company’s five primary data centers. There are 94 midrange and small RISC/UNIX servers at 34 remote sites employed for call center, document serving, and local and departmental applications.

At least 717 Intel-based servers are located at 163 remote sites. The company’s e-mail network is, in particular, highly dispersed, and end-user file and print servers are located at numerous regional and local offices.

Data centers 1 and 2 support the bank’s principal operations. Data centers 3 and 4 are operated by a recently acquired regional bank. In both organizations, data centers are duplexed for disaster recovery purposes. A recently acquired financial services company operates data center 5. Its business-critical systems are mirrored to data center 1.

The company employs a range of different storage systems. Mainframe disk storage at data centers 1 and 2 is a mix of EMC Symmetrix and HDS equipment, while data centers 3 and 4 employ IBM ESS and data center 5 employs HDS systems. Mainframe storage is typically dedicated, and attached with a mix of ESCON and FICON connections.

RISC/UNIX servers employ a variety of disk subsystems from HP, IBM, Sun, Compaq, EMC (Symmetrix and Clariion), HDS and other vendors, and the storage base for Intel-based servers is even more diverse. Overall, the company employs disk storage from more than 20 vendors.

Additional detail for this scenario is presented on initial mainframe (figure 30), server (figure 31) and disk storage bases (figure 32), projected disk storage growth (figure 33), IT costs (figure 34), and staffing (figure 35).

Figure 30

Bank Mainframe Bases

	INEFFICIENT SCENARIO					EFFICIENT SCENARIO	
	1	2	3	4	5	1	2
Systems	1C8 Pilot Z97 Skyline	ZZ7 XY7 Millennium RY5	IC7	XX7 R85	RX6 X57	2 x 2C8 210 Z97	2 x 210 1C8 1C7
MIPS	5,438	4,850	1,416	1,626	1,642	7,458	7,463
Disk	26.8 TB	19.1 TB	6.5 TB	7.8 TB	12.0 TB	–	–

Figure 31

Bank Inefficient Scenario: Initial Server Base

SERVER TYPE	DATA CENTER						TOTAL
	1	2	3	4	5	REMOTE	
MAINFRAME							
Total	4	4	1	2	2	–	**13**
LARGE RISC/UNIX							
Corporate data warehouse	15	–	–	–	–	–	**15**
Data marts/analytical	6	4	–	4	–	–	**14**
Business systems	12	10	3	–	7	–	**32**
MIDRANGE RISC/UNIX							
Data marts/analytical	7	2	–	–	–	–	**9**
Business systems	31	47	18	6	74	22	**198**
SMALL RISC/UNIX							
All applications	105	221	38	22	82	72	**540**
Total	**176**	**284**	**59**	**32**	**163**	**94**	**808**
INTEL-BASED							
End user	95	90	12	10	25	192	**424**
E-mail	23	18	5	4	13	339	**402**
Web/related	127	33	16	5	93	16	**290**
Departmental	56	45	8	10	35	31	**185**
Citrix	50	20	–	–	–	61	**131**
Other	50	25	15	2	23	78	**193**
Total	**401**	**231**	**56**	**31**	**189**	**717**	**1,625**

Figure 32

Bank Inefficient Scenario: Initial Disk Storage Base

INITIAL INSTALLED BASE (TB)	DATA CENTER 1	2	3	4	5	REMOTE	TOTAL (TB)
MAINFRAME							
Total (TB)	26.8	19.1	6.5	7.8	12.0	–	**72.2**
LARGE RISC/UNIX							
Corporate data warehouse	56.6	–	–	–	–	–	**56.6**
Data marts/analytical	17.0	8.2	–	3.0	–	–	**28.2**
Business systems	10.6	12.8	4.7	–	7.7	–	**35.8**
MIDRANGE RISC/UNIX							
Data marts/analytical	2.7	2.0	–	–	–	–	**4.7**
Business systems	9.0	12.9	5.4	3.1	11.2	3.8	**45.4**
SMALL RISC/UNIX							
All applications	4.2	5.4	1.6	1.3	4.4	2.5	**19.4**
Total (TB)	**100.1**	**41.3**	**11.7**	**7.4**	**23.3**	**6.3**	**190.1**
INTEL-BASED							
End user	1.1	0.7	0.1	0.2	0.5	1.8	**4.4**
E-mail	2.2	1.3	0.3	0.3	0.8	2.0	**6.9**
Web/related	4.7	1.5	0.7	0.3	2.0	0.5	**9.7**
Departmental	3.5	2.7	0.5	0.6	2.0	1.9	**11.2**
Citrix	0.6	0.3	–	–	–	0.7	**1.6**
Other	2.2	1.0	0.2	0.3	1.1	–	**4.8**
Total (TB)	**14.3**	**7.5**	**1.8**	**1.7**	**6.4**	**6.9**	**38.6**

Figure 33

Bank Inefficient Scenario: Projected Disk Storage Growth

PROJECTED GROWTH (TB)	YE 02	YE 03	YE 04	YE 05	YE 06	YE 07
MAINFRAME						
Total (TB)	**72.2**	**89.5**	**111.3**	**138.4**	**172.2**	**214.4**
RISC/UNIX SERVERS						
Corporate data warehouse	56.6	67.9	82.3	99.3	119.5	143.6
Data marts/analytical	32.9	43.3	57.7	76.9	102.7	137.6
Business systems	81.2	109.4	149.7	205.2	281.8	388.7
Small servers	19.4	26.3	36.3	49.9	68.8	95.3
Total (TB)	**190.1**	**246.9**	**326.0**	**431.3**	**572.8**	**765.2**
INTEL-BASED SERVERS						
End user	4.4	5.2	6.4	7.8	9.8	12.2
E-mail	6.9	9.5	13.3	18.6	26.0	36.2
Web/related	9.7	13.5	19.7	29.2	44.3	67.9
Departmental	11.2	15.4	21.6	30.1	41.8	58.4
Other	6.4	8.0	10.4	13.7	18.2	24.5
Total (TB)	**38.6**	**51.6**	**71.4**	**99.4**	**140.1**	**199.2**
TOTAL (TB)	**300.9**	**388.0**	**508.7**	**669.1**	**885.1**	**1,178.8**

Figure 34

Bank Inefficient Scenario: IT Costs Detail

IT COST CATEGORIES	2003	2004	2005	2006	2007	TOTAL ($000)
Hardware acquisition	8,346	8,103	7,899	7,741	7,634	**39,723**
Hardware maintenance	865	1,021	1,173	1,321	1,456	**5,836**
Software licenses	1,112	1,304	1,482	1,665	1,838	**7,401**
Software support	483	572	693	818	973	**3,539**
SAN/NAS hardware/software	1,323	1,632	1,803	1,893	1,851	**8,502**
SAN/NAS maintenance/support	49	97	156	221	290	**813**
Personnel	15,206	18,363	21,448	24,578	23,755	**103,350**
Facilities	576	676	810	994	1,248	**4,304**
TOTAL ($000)	**27,960**	**31,768**	**35,464**	**39,231**	**39,045**	**173,468**

Figure 35

Bank Inefficient Scenario: Staffing Detail

IT FUNCTION	ALL STAFF	STORAGE FTEs				
	2002	2003	2004	2005	2006	2007
Management/administration	**99**	8	9	11	12	10
Technical staff	**101**	4	5	6	6	5
Applications	**765**	14	16	19	22	19
Database administration	**83**	11	13	14	16	14
System/storage management	**265**	110	135	156	169	148
Network management	**226**	6	7	7	8	7
Operations & support	**284**	16	17	18	21	18
End-user computing & help desk	**374**	12	13	15	17	14
TOTAL	**2,197**	**181**	**215**	**246**	**271**	**235**

Efficient Scenario

In this scenario, a number of systems have been standardized across the bank as a whole, based on “best of breed” decisions about which to retain.

These include mainframe-based core banking, branch automation, ATM/POS and EFT applications; along with treasury and accounting systems, and a number of RISC/UNIX- and Intel-based data marts and business systems. Common Web interfaces and customer service applications have also been merged across the principal bank and its two recent acquisitions.

PeopleSoft Financials, Human Resources Management System (HRMS) and EPM have been adopted as corporate standards. There are approximately 1,000, 800 and 200 users respectively of these systems. PeopleSoft HRMS is used for HR administration, benefits and payroll for all employees.

In the efficient scenario, the vast majority of servers and storage are located at the bank’s two primary data centers. Apart from branches and local offices that employ server-attached storage devices, the only remote platforms are 15 Intel-based Citrix, e-mail and end-user servers located at five regional centers; and 44 small and midrange RISC/UNIX servers employed for call center and remote document distribution applications that employ direct-attached storage.

Additional detail on the initial server base for this scenario is presented in figure 36. Vendor-specific calculations are based on the platforms, software stacks and staffing levels shown in figures 37, 38 and 39 respectively. Additional detail on vendor configurations is presented in figures 40 through 43.

Figure 36

Bank Efficient Scenario: Initial Server Base

SERVER TYPE	DATA CENTER 1			DATA CENTER 2			REMOTE			TOTAL SERVERS
	Large	Midrange	Small	Large	Midrange	Small	Large	Midrange	Small	
RISC/UNIX										
Corporate data warehouse	11	–	–	–	–	–	–	–	–	**11**
Data marts/analytical	6	7	–	5	8	–	–	–	–	**26**
Business systems	15	5	–	19	70	–	–	–	–	**109**
Small servers (all applications)	–	–	8	–	–	38	–	4	40	**90**
Total	**32**	**12**	**8**	**24**	**78**	**38**	–	**4**	**40**	**236**
INTEL-BASED										
End user	2	–	–	2	–	–	–	5	–	**9**
E-mail	2	1	–	2	–	–	–	5	–	**10**
Web/related	5	1	7	8	2	11	–	–	–	**34**
Departmental	3	1	3	4	3	7	–	–	–	**21**
Citrix	2	2	–	2	–	–	–	5	–	**11**
Other	–	5	5	1	2	5	–	–	–	**18**
Total	**14**	**10**	**15**	**19**	**7**	**23**	–	**15**	–	**103**

Figure 37

Bank Efficient Scenario: Vendor Platforms

	EMC	HDS	IBM	HP
DATA CENTER 1				
Critical systems	Symmetrix DMX	Lightning 9980V	ESS 800	–
Secondary systems 1	Symmetrix DMX	Lightning 9980V	ESS 800	–
Secondary systems 2	Clariion CX600	Thunder 9570V	FAStT 900	EVA 2C12D
NAS	Celerra CNS Clariion CX200	Enterprise NAS Gateway Thunder 9570V	NAS Gateway 300 FAStT 200	NAS e7000 EVA 2C2D
SAN	Connectrix	Brocade 12000	IBM SAN Switch	–
DATA CENTER 2				
Critical systems	Symmetrix DMX	Lightning 9980V	ESS 800	–
Secondary systems 1	Symmetrix DMX	Lightning 9980V	ESS 800	–
Secondary systems 2	Clariion CX600	Thunder 9570V	FAStT 900	EVA 2C12D
NAS	Celerra CNS Clariion CX400	Enterprise NAS Gateway Thunder 9570V	NAS Gateway 300 FAStT 500	NAS e7000 EVA 2C6D
SAN	Connectrix	Brocade 12000	IBM SAN Switch	–
REMOTE				
Regional centers	Clariion CX200	Thunder 9570V	FAStT 200	EVA 2C2D

Figure 38

Bank Efficient Scenario: Vendor Software Stacks

SYSTEMS	EMC	HDS	IBM	HP
Critical systems	CommandCenter Package ESN Manager PowerPath Enterprise Symmetrix Remote Data Facility (SRDF) TimeFinder PAV	HiCommand Resource Manager Complete TrueCopy ShadowImage HPAV	ESS Specialist ESS Expert Peer-to-Peer Remote Copy (PPRC) FlashCopy PAV	–
Secondary systems 1	CommandCenter Package ESN Manager PowerPath Enterprise TimeFinder PAV	HiCommand Resource Manager Complete ShadowImage HPAV	ESS Specialist ESS Expert FlashCopy PAV	–
Secondary systems 2	CommandCenter Package ESN Manager PowerPath Enterprise	HiCommand Resource Manager Complete	ESS Specialist ESS Expert	–
NAS	Access Logix Navisphere Manager SnapView	HiCommand Resource Manager Open ShadowImage	Base package Host access FlashCopy	Virtual Controller Software (VCS) Snapshot
Regional centers	Access Logix Navisphere Manager	HiCommand Resource Manager Open	Base package Host access	Virtual Controller Software (VCS)

Figure 39

Bank Efficient Scenario: Staffing Detail

IT FUNCTION	EMC	HDS	IBM
Management/administration	2	2	2
Technical staff	1	1	1
Applications	4	4	3
Database administration	4	3	3
System/storage management	21	20	18
Network management	1	1	1
Operations & support	4	4	4
End-user computing & help desk	2	2	2
TOTAL	**39**	**37**	**34**

Figure 40

Bank Efficient Scenario: EMC Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
Critical systems*	Symmetrix DMX	2 x 16 TB	2 x 16 TB 1 x 8 TB	3 x 16 TB	3 x 16 TB 1 x 15 TB	4 x 16 TB 1 x 15 TB	6 x 16 TB
Secondary systems 1	Symmetrix DMX	2 x 32 TB 1 x 16 TB	2 x 32 TB 1 x 20 TB 1 x 12 TB	2 x 32 TB 1 x 26 TB 1 x 24 TB	3 x 32 TB 1 x 28 TB 1 x 8 TB	3 x 32 TB 1 x 28 TB 1 x 27 TB 1 x 6 TB	4 x 32 TB 1 x 28 TB 1 x 18 TB 1 x 15 TB
Secondary systems 2	Clariion CX600	1 x 7 TB 1 x 4 TB	2 x 7 TB	1 x 9 TB 1 x 7 TB	1 x 11 TB 1 x 9 TB	1 x 13 TB 1 x 11 TB	2 x 15 TB
NAS	Clariion CX200	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB
DATA CENTER 2							
Critical systems	Symmetrix DMX	3 x 13 TB	3 x 15 TB 1 x 6 TB	4 x 15 TB 1 x 6 TB	5 x 15 TB 1 x 11 TB	6 x 16 TB 1 x 14 TB	8 x 16 TB 1 x 15 TB
Secondary systems 1	Symmetrix DMX	1 x 16 TB 1 x 13 TB 1 x 12 TB	1 x 21 TB 2 x 15 TB	1 x 27 TB 1 x 20 TB 1 x 19 TB	1 x 27 TB 2 x 25 TB 1 x 9 TB	1 x 32 TB 1 x 30 TB 1 x 27 TB 1 x 21 TB	4 x 32 TB 1 x 18 TB
Secondary systems 2	Clariion CX600	3 x 7 TB	3 x 9 TB	3 x 11 TB	3 x 13 TB	4 x 13 TB	4 x 15 TB
NAS	Clariion CX400	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 7 TB
REMOTE							
Regional centers	Clariion CX200	5 x 1 TB	5 x 1 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB

*Mirrored

Figure 41

Bank Efficient Scenario: HDS Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
Critical systems*	9980V	2 x 16 TB	2 x 16 TB 1 x 8 TB	3 x 16 TB	3 x 16 TB 1 x 13 TB	1 x 20 TB 3 x 19 TB	4 x 20 TB 1 x 19 TB
Secondary systems 1	9980V	2 x 32 TB 1 x 16 TB	2 x 38 TB 1 x 20 TB	1 x 44 TB 1 x 43 TB 1 x 23 TB	2 x 51 TB 1 x 28 TB	2 x 51 TB 1 x 35 TB 1 x 19 TB	2 x 51 TB 1 x 43 TB 1 x 41 TB
Secondary systems 2	9570V	1 x 7 TB 1 x 4 TB	2 x 7 TB	1 x 8 TB 1 x 7 TB	1 x 11 TB 1 x 9 TB	1 x 13 TB 1 x 11 TB	2 x 15 TB
NAS	9570V	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB
DATA CENTER 2							
Critical systems*	9980V	3 x 13 TB	2 x 17 TB 1 x 16 TB	2 x 22 TB 1 x 21 TB	2 x 28 TB 1 x 27 TB	3 x 28 TB 1 x 25 TB	2 x 29 TB 3 x 28 TB
Secondary systems 1	9980V	1 x 16 TB 2 x 12 TB	1 x 21 TB 2 x 15 TB	1 x 27 TB 1 x 20 TB 1 x 19 TB	1 x 36 TB 2 x 25 TB	1 x 36 TB 1 x 32 TB 1 x 30 TB 1 x 12 TB	1 x 41 TB 1 x 40 TB 1 x 36 TB 1 x 27 TB
Secondary systems 2	9570V	3 x 7 TB	3 x 9 TB	3 x 11 TB	3 x 13 TB	4 x 13 TB	4 x 15 TB
NAS	9570V	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 6 TB
REMOTE							
Regional centers	9570V	5 x 1 TB	5 x 1 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB

*Mirrored

Figure 42

Bank Efficient Scenario: IBM Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
Critical systems*	ESS 800	2 x 16 TB	2 x 16 TB 1 x 8 TB	3 x 16 TB	3 x 16 TB 1 x 15 TB	5 x 16 TB	5 x 20 TB
Secondary systems 1	ESS 800	2 x 33 TB 1 x 16 TB	2 x 37 TB 1 x 20 TB	2 x 37 TB 1 x 24 TB 1 x 13 TB	2 x 37 TB 2 x 28 TB	3 x 41 TB 1 x 35 TB	1 x 42 TB 3 x 41 TB 1 x 21 TB
Secondary systems 2	FAStT 900	1 x 6 TB 1 x 4 TB	2 x 6 TB	1 x 8 TB 1 x 6 TB	1 x 10 TB 1 x 8 TB	1 x 12 TB 1 x 10 TB	2 x 14 TB
NAS	FAStT 600	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB
DATA CENTER 2							
Critical systems*	ESS 800	3 x 13 TB	3 x 13 TB 1 x 11 TB	4 x 16 TB	4 x 16 TB 1 x 19 TB	5 x 20 TB 1 x 10 TB	7 x 20 TB
Secondary systems 1	ESS 800	1 x 16 TB 1 x 13 TB 1 x 11 TB	1 x 21 TB 2 x 15 TB	1 x 27 TB 1 x 20 TB 1 x 18 TB	1 x 36 TB 2 x 24 TB	1 x 36 TB 2 x 31 TB 1 x 12 TB	1 x 41 TB 1 x 40 TB 1 x 36 TB 1 x 27 TB
Secondary systems 2	FAStT 900	3 x 6 TB	3 x 8 TB	3 x 10 TB	3 x 12 TB	4 x 12 TB	4 x 14 TB
NAS	FAStT 500	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 6 TB
REMOTE							
Regional centers	FAStT 600	5 x 1 TB	5 x 1 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB

*Mirrored

Figure 43

Bank Efficient Scenario: Hewlett-Packard Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
Secondary systems 2	EVA 2C12D	1 x 7 TB 1 x 4 TB	2 x 7 TB	1 x 9 TB 1 x 7 TB	1 x 11 TB 1 x 9 TB	1 x 13 TB 1 x 11 TB	2 x 15 TB
NAS	EVA 2C2D	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB
DATA CENTER 2							
Secondary systems 2	EVA 2C12D	3 x 7 TB	3 x 9 TB	3 x 11 TB	3 x 13 TB	4 x 13 TB	4 x 15 TB
NAS	EVA 2C6D	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 7 TB
REMOTE							
Regional centers	EVA 2C2D	5 x 1 TB	5 x 1 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB	5 x 2 TB

Manufacturing Company

SAP User Base

Numbers of named and concurrent users for the company's principal SAP systems are shown in figure 44. Calculations for both the inefficient and efficient scenarios employ the same numbers of users.

Figure 44
Manufacturing Company: Numbers of Users for Principal SAP Systems

APPLICATION/MODULE	USERS	
	Named	Concurrent (Max.)
R/3 Suite	21,000	3,000
Advanced Planning & Optimization (APO)	250	50
Business Information Warehouse (BW)		
Corporate data warehouse	1,350	250
Data marts	650	100
Corporate Finance Management (CFM)	2,300	600
Customer Relationship Management (CRM)	3,350	700
Enterprise Buyer Professional (EBP)	3,000	500
Product Lifecycle Management (PLM)	2,550	650
Strategic Enterprise Management (SEM)	750	150
Enterprise Portal	23,000	3,000
Employee Self Service (ESS)	23,000	3,500

R/3 configurations are based on Version 4.6c of this system.

Inefficient Scenario

Additional detail for this scenario is presented on initial server and disk storage bases (figure 45), projected disk storage growth (figure 46), IT costs (figure 47), and staffing (figure 48)

Figure 45

Manufacturing Company Inefficient Scenario: Initial Server and Disk Storage Bases

INSTALLED BASE	SERVERS		DISK STORAGE (TB)	
	RISC/UNIX	Intel-based	R/3	Other
NORTH AMERICA				
U.S. & Canada	53	32	39.0	26.0
LATIN AMERICA				
North/central	25	–	15.2	1.6
Brazil	21	9	16.1	1.2
Southern	12	16	18.8	1.3
EUROPE				
EMEA	16	18	–	29.0
Germany	48	17	12.1	3.2
United Kingdom	18	25	6.4	2.0
France	23	–	6.0	1.9
Benelux	13	16	2.8	1.9
Italy/Southern	26	–	6.6	1.2
Iberia	6	25	3.8	1.3
Nordic/CIS	–	24	2.3	0.6
Eastern	12	19.	5.0	1.3
MIDDLE EAST/AFRICA				
South Africa	14	3	3.0	0.8
Middle East/Africa	5	17	1.6	0.6
ASIA/PACIFIC				
Australia	10	13	3.2	1.6
Singapore	29	9	16.0	3.7
Hong Kong	–	31	5.0	1.2
Japan	20	17	16.0	13.5
TOTAL	**351**	**291**	**178.9**	**93.9**

Figure 46

Manufacturing Company Inefficient Scenario: Projected Disk Storage Growth

PROJECTED GROWTH (TB)	YE 02	YE 03	YE 04	YE 05	YE 06	YE 07
EMEA	93.4	119.9	154.5	215.6	274.7	352.2
North America	65.0	80.6	100.1	124.5	155.1	193.6
Latin America	54.2	72.4	96.8	129.4	173.1	231.7
Asia/Pacific	60.2	78.8	103.6	137.2	182.6	244.9
TOTAL (TB)	**272.8**	**351.7**	**455.0**	**606.7**	**785.5**	**1,022.4**

Figure 47

Manufacturing Company Inefficient Scenario: IT Costs Detail

IT COST CATAGORIES	2003	2004	2005	2006	2007	TOTAL ($000)
Hardware acquisition	6,229	6,052	5,897	5,762	5,647	**29,587**
Hardware maintenance	691	814	929	1,026	1,127	**4,587**
Software licenses	586	735	971	1,048	1,161	**4,501**
Software support	235	280	357	388	437	**1,697**
SAN/NAS hardware/software	848	962	1,022	1,086	1,154	**5,072**
SAN/NAS maintenance/support	25	55	90	127	167	**464**
Personnel	5,636	6,390	7,443	8,117	8,435	**36,021**
Facilities	595	692	837	1,005	1,227	**4,356**
TOTAL ($000)	**14,845**	**15,980**	**17,546**	**18,559**	**19,355**	**86,285**

Figure 48

Manufacturing Inefficient Scenario: Staffing Detail

IT FUNCTION	ALL STAFF	STORAGE FTEs				
	2002	2003	2004	2005	2006	2007
Management/administration	**34**	3	4	5	6	6
Technical staff	**28**	3	3	4	4	4
Applications	**267**	8	9	10	11	12
Database administration	**40**	5	6	7	7	7
System/storage management	**89**	33	37	43	47	49
Network management	**102**	2	2	2	2	2
Operations & support	**119**	17	19	22	24	25
End-user computing & help desk	**160**	3	3	3	4	4
TOTAL	**839**	74	83	96	105	109

Efficient Scenario

Additional detail on centralized data center environments is shown in figure 49. Vendor-specific calculations are based on the platforms, software stacks and staffing levels shown in figures 50, 51 and 52 respectively. Additional detail on vendor configurations is presented in figures 53 through 56.

Figure 49

Manufacturing Company Efficient Scenario: Central Data Center Environments

DATA CENTER 1	DATA CENTER 2
APPLICATIONS	
Core Systems Production R/3 Applications Advanced Planning & Optimization (APO) Product Lifecycle Management (PLM) Enterprise Buyer Professional (EBP) Complementary applications EDI server E-commerce	**Core Systems Support** R/3 Applications Advanced Planning & Optimization (APO) Product Lifecycle Management (PLM) Enterprise Buyer Professional (EBP) Complementary applications EDI server E-commerce
Business Intelligence Corporate data warehouse Data marts • Finance (profitability analysis) • Human resources • Market research Corporate Finance Management (CFM) Strategic Enterprise Management (CEM)	**Business Systems** Customer Relationship Management (CRM) – North America Enterprise Portal Employee Self Service Various
NAS Applications	**NAS Applications**
SERVER BASE	
Core Systems Production 4 x pSeries 690 3 x pSeries 670 10 x pSeries 650 45 x xSeries	**Core Systems Support** 5 x pSeries 690 5 x pSeries 670 8 x pSeries 650 72 x xSeries
Business Intelligence 6 x pSeries 690 5 x pSeries 670 8 x pSeries 650 12 x xSeries	**Other Systems** 2 x pSeries 670 2 x pSeries 650 27 x xSeries
NAS Applications 3 x pSeries 650	**NAS Applications** 4 x pSeries 650

Figure 50

Manufacturing Company Efficient Scenario: Vendor Platforms

	EMC	HDS	IBM	HP
DATA CENTER 1				
RAID 10	Symmetrix DMX	Lightning 9980V	ESS 800	–
RAID 5	Symmetrix DMX	Lightning 9980V	ESS 800	–
NAS	Celerra CNS Clariion CX400	Enterprise NAS Gateway Thunder 9570V	NAS Gateway 3000 FAStT 500	NAS e7000 EVA 2C6D
SAN	Connectrix	Brocade 12000	IBM SAN Switch	–
DATA CENTER 2				
RAID 10	Symmetrix DMX	Lightning 9980V	ESS 800	–
RAID 5	Symmetrix DMX	Lightning 9980V	ESS 800	–
NAS	Celerra CNS Clariion CX200	Enterprise NAS Gateway Thunder 9570V	NAS Gateway 300 FAStT 600	NAS e7000 EVA 2C6D
SAN	Connectrix	Brocade 12000	IBM SAN Switch	–
REMOTE				
Regional centers	Clariion CX600	Thunder 9570V	FAStT 900	EVA 2C12D

Figure 51

Manufacturing Company Efficient Scenario: Vendor Software Stacks

SYSTEMS	EMC	HDS	IBM	HP
RAID 10	CommandCenter Package ESN Manager PowerPath Enterprise Symmetrix Remote Data Facility (SRDF) TimeFinder	HiCommand Resource Manager Open TrueCopy ShadowImage	ESS Specialist ESS Expert Peer-to-Peer Remote Copy (PPRC) FlashCopy	–
RAID 5	CommandCenter Package ESN Manager PowerPath Enterprise TimeFinder	HiCommand Resource Manager Open ShadowImage	ESS Specialist ESS Expert FlashCopy	–
NAS	Access Logix Navisphere Manager	HiCommand Resource Manager Open	Base package Host access	Virtual Controller Software (VCS)
Remote	Access Logix Navisphere Manager SnapView	HiCommand Resource Manager Open ShadowImage	Base package Host access FlashCopy	Virtual Controller Software (VCS) Snapshot

Figure 52

Manufacturing Company Efficient Scenario: Staffing Detail

IT FUNCTION	EMC	HDS	IBM
Management/administration	1	1	1
Technical staff	1	1	1
Applications	2	2	2
Database administration	2	2	2
System/storage management	9	9	8
Network management	1	1	1
Operations & support	2	2	2
End-user computing & help desk	–	–	–
TOTAL	**18**	**18**	**17**

Figure 52

Manufacturing Company Efficient Scenario: EMC Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
RAID 10	Symmetrix DMX	3 x 10 TB 1 x 11 TB	4 x 11 TB 1 x 8 TB	5 x 11 TB 1 x 9 TB	7 x 11 TB	9 x 11 TB	11 x 11 TB
RAID 5	Symmetrix DMX	2 x 15 TB	2 x 20 TB	2 x 27 TB	2 x 27 TB 1 x 17 TB	3 x 27 TB 1 x 13 TB	4 x 27 TB 1 x 23 TB
NAS	Clariion CX400	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB
DATA CENTER 2							
RAID 10	Symmetrix DMX	3 x 10 TB 1 x 11 TB	4 x 11 TB 1 x 8 TB	5 x 11 TB 1 x 9 TB	7 x 11 TB	9 x 11 TB	11 x 11 TB
RAID 5	Symmetrix DMX	2 x 15 TB	2 x 19 TB	2 x 19 TB 1 x 10 TB	3 x 20 TB	2 x 25 TB 1 x 23 TB	3 x 25 TB 1 x 16 TB
NAS	Clariion CX200	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB
REMOTE							
Latin America	Clariion CX600	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB
EMEA	Clariion CX600	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB	1 x 16 TB
Asia/Pacific	Clariion CX600	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB

Figure 54

Manufacturing Company Efficient Scenario: HDS Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
RAID 10	9980V	4 x 11 TB	4 x 11 TB 1 x 8 TB	5 x 11 TB 1 x 9 TB	7 x 11 TB	9 x 11 TB	11 x 11 TB
RAID 5	9980V	2 x 15 TB	2 x 20 TB	2 x 27 TB	2 x 27 TB 1 x 18 TB	3 x 32 TB	3 x 32 TB 1 x 36 TB
NAS	9570V	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB
DATA CENTER 2							
RAID 10	9980V	4 x 11 TB	4 x 11 TB 1 x 8 TB	5 x 11 TB 1 x 9 TB	7 x 11 TB	9 x 10 TB	11 x 11 TB
RAID 5	9980V	2 x 15 TB	2 x 19 TB	2 x 19 TB 1 x 11 TB	3 x 20 TB	2 x 25 TB 1 x 23 TB	3 x 25 TB 1 x 16 TB
NAS	9570V	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB
REMOTE							
Latin America	9570V	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB
EMEA	9570V	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB	1 x 16 TB
Asia/Pacific	9570V	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB

Figure 54

Manufacturing Company Efficient Scenario: IBM Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
RAID 10	ESS 800	4 x 10 TB	4x 10 TB 1 x 8 TB	5 x 10 TB 1 x 9 TB	7 x 10 TB	9 x 10 TB	11 x 10 TB
RAID 5	ESS 800	2 x 15 TB	2 x 20 TB	2 x 27 TB	2 x 27 TB 1 x 17 TB	3 x 31 TB	3 x 31 TB 1 x 36 TB
NAS	FAStT 500	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 6 TB	1 x 8 TB
DATA CENTER 2							
RAID 10	ESS 800	4 x 10 TB	4 x 10 TB 1 x 8 TB	5 x 10 TB 1 x 9 TB	7 x 10 TB	9 x 10 TB	11 x 10 TB
RAID 5	ESS 800	2 x 15 TB	2 x 19 TB	2 x 19 TB 1 x 10 TB	3 x 20 TB	2 x 24 TB 1 x 23 TB	3 x 24 TB 1 x 16 TB
NAS	FAStT 600	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB
REMOTE							
Latin America	FAStT 900	1 x 4 TB	1 x 5 TB	1 x 6 TB	1 x 8 TB	1 x 10 TB	1 x 12 TB
EMEA	FAStT 900	1 x 5 TB	1 x 6 TB	1 x 8 TB	1 x 10 TB	1 x 12 TB	1 x 15 TB
Asia/Pacific	FAStT 900	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 8 TB	1 x 10 TB	1 x 12 TB

Figure 56

Manufacturing Company Efficient Scenario: Hewlett-Packard Configurations

	Platform	2002	2003	2004	2005	2006	2007
DATA CENTER 1							
NAS	EVA 2C6D	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB
DATA CENTER 2							
NAS	EVA 2C2D	1 x 1 TB	1 x 2 TB	1 x 2 TB	1 x 3 TB	1 x 4 TB	1 x 5 TB
REMOTE							
Latin America	EVA 2C12D	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB
EMEA	EVA 2C12D	1 x 5 TB	1 x 7 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB	1 x 16 TB
Asia/Pacific	EVA 2C12D	1 x 3 TB	1 x 4 TB	1 x 5 TB	1 x 9 TB	1 x 11 TB	1 x 13 TB

Retail Company

Inefficient Scenario

This scenario is based on the following groups of applications and platforms:

- ***Mainframe systems***. These include core merchandising and sales consolidation systems operated by the original company and its two recent acquisitions. The company initially – as of yearend 2002 – operates eight IBM- and IBM-compatible mainframe processors with a total of 1,725 MIPS and 19.9 terabytes of disk storage.

- ***Data warehouses***. At the beginning of the five-year projection period, the main company data warehouse runs on five IBM SP frames with a total of 80 nodes and 28.5 terabytes of disk storage. Data warehouses continue to be employed by the two retail chains recently acquired by the company. These run on Sun Microsystems E10000 and E6500 servers with 10.1 and 3.3 terabytes of disk storage respectively.

- ***Corporate systems***. Separate finance and HR systems are operated by the three company entities. The original company employs PeopleSoft Financials and HRMS running on high-end IBM RS/6000 servers. The other two employ mainframe-based applications; and J.D. Edwards Financials and HRMS running on Sun Microsystems servers.

  The RS/6000- and Sun-based systems employ a combined total of 14.1 terabytes of disk storage. Capacity for mainframe-based systems is included in the mainframe total above.

- ***Supply chain systems***. These include systems for procurement, warehouse management, inventory and other supply chain applications, along with optimization and analytical tools operated by each of the three company entities. There are 17 systems running on a mix of 104 Hewlett-Packard and IBM RISC/UNIX servers, IBM AS/400s and Intel-based Windows servers with a combined total of 31.8 terabytes of disk storage.

- ***Sales and marketing systems***. These include systems for category and space management, marketing and promotions, loyalty program management, sales audit, analytical and other applications. There are 18 systems running on a mix of 52 Hewlett-Packard, IBM and Sun Microsystems RISC/UNIX servers, and Intel-based servers from these and other companies with a combined total of 7.3 terabytes of disk storage.

- ***Other systems***. The three company entities employ a mix of Lotus Notes and Microsoft Exchange e-mail and groupware networks; Microsoft and Novell NetWare LANs; IBM, Microsoft, Sun Microsystems, and open source (primarily Apache) web and intranet software; Windows-based departmental applications; test and development systems and others which run primarily on Intel-based servers.

  The company also employs a Computer Output on Laser Disk (COLD) system running on Sun Microsystems servers. Other systems are supported by 16 RISC/UNIX and 657 Intel-based servers with 11.7 terabytes of disk storage.

Mainframe disk storage includes EMC (10.5 terabytes) and IBM (9.4 terabytes) systems. RISC/UNIX disk storage includes a variety of EMC, HDS, HP, IBM, Sun Microsystems and compatible systems, while Intel-based servers employ Dell Computer, HP, IBM and other vendors’ storage.

Additional detail for this scenario is presented on initial server (Figure 57) and disk storage (figure 58) bases, projected disk storage growth (figure 59), IT costs (figure 60), and staffing (figure 61).

Figure 57

Retail Company Inefficient Scenario: Initial Server Base

	DATA CENTER			TOTAL
	1	2	3	
MAINFRAME				
Total	3	2	3	**8**
DATA WAREHOUSE				
Large RISC/UNIX	5	2	1	**8**
Midrange RISC/UNIX	–	–	1	**1**
Small RISC/UNIX	–	–	–	**–**
CORPORATE SYSTEMS				
Large RISC/UNIX	2	–	–	**2**
Midrange RISC/UNIX	2	2	1	**5**
Small RISC/UNIX	3	–	2	**5**
SUPPLY CHAIN				
Large RISC/UNIX	4	2	–	**6**
Midrange RISC/UNIX	18	4	10	**32**
AS/400	–	13	–	**13**
Small RISC/UNIX	21	7	9	**37**
Intel-based	–	2	14	**16**
SALES & MARKETING				
Large RISC/UNIX	1	–	–	**1**
Midrange RISC/UNIX	5	3	1	**9**
Small RISC/UNIX	11	6	2	**19**
Intel-based	6	7	10	**23**
E-MAIL & GROUPWARE				
Midrange RISC/UNIX	–	1	–	**1**
Small RISC/UNIX	–	3	–	**3**
Intel-based	174	106	66	**346**
END USER				
Intel-based	73	56	22	**151**
WEB/RELATED				
Intel-based	44	19	16	**79**
OTHER				
Midrange RISC/UNIX	3	–	–	**3**
Small RISC/UNIX	4	5	–	**9**
Intel-based	41	25	15	**81**
TOTAL	**420**	**265**	**173**	**858**

Figure 58

Retail Company Inefficient Scenario: Initial Disk Storage Base

INITIAL STORAGE BASE (TB)	DATA CENTER 1	2	3	TOTAL (TB)
MAINFRAME				
All systems	10.5	5.8	3.6	**19.9**
DATA WAREHOUSE				
Large RISC/UNIX	28.5	10.1	2.5	**41.1**
Midrange RISC/UNIX	–	–	0.8	**0.8**
Small RISC/UNIX	–	–	–	**–**
CORPORATE SYSTEMS				
Large RISC/UNIX	8.6	–	–	**8.6**
Midrange RISC/UNIX	2.0	0.7	1.8	**4.5**
Small RISC/UNIX	0.4	–	0.6	**1.0**
SUPPLY CHAIN				
Large RISC/UNIX	13.8	3.3	–	**17.1**
Midrange RISC/UNIX	3.0	1.6	3.6	**8.2**
AS/400	–	2.9	–	**2.9**
Small RISC/UNIX	1.1	0.8	0.3	**2.2**
Intel-based	–	0.2	1.2	**1.4**
SALES & MARKETING				
Large RISC/UNIX	0.7	–	–	**0.7**
Midrange RISC/UNIX	1.5	0.7	0.4	**2.6**
Small RISC/UNIX	0.8	0.7	0.2	**1.7**
Intel-based	0.4	1.0	0.9	**2.3**
E-MAIL & GROUPWARE				
Midrange RISC/UNIX	–	0.3	–	**0.3**
Small RISC/UNIX	–	0.6	–	**0.6**
Intel-based	1.5	0.4	0.4	**2.3**
END USER				
Intel-based	1.0	0.5	0.3	**1.8**
WEB/RELATED				
Intel-based	0.9	0.5	0.4	**1.8**
OTHER				
Midrange RISC/UNIX	1.2	–	–	**1.2**
Small RISC/UNIX	0.2	0.4	–	**0.6**
Intel-based	1.6	1.1	0.4	**3.1**
TOTAL (TB)	**77.7**	**31.6**	**17.4**	**126.7**

Figure 59

Retail Inefficient Scenario: Projected Disk Storage Growth

CATEGORY	YEAREND INSTALLED BASE (TB)					
	YE 02	YE 03	YE 04	YE 05	YE 06	YE 07
Mainframe	19.9	22.5	25.5	28.9	32.7	37.0
Data warehouse	41.9	70.8	119.6	202.2	289.3	393.7
Corporate systems	14.1	16.9	20.2	24.3	29.3	35.4
Supply chain	31.8	42.0	55.5	73.4	97.2	128.9
Sales & marketing	7.3	10.1	13.8	19.0	26.2	36.0
Email & groupware	3.2	4.6	6.5	9.2	13.1	18.7
End user	1.8	2.2	2.7	3.3	4.0	4.8
Web/related	1.8	2.5	3.4	4.8	6.6	9.1
Other	4.9	6.6	8.9	11.9	16.1	21.7
TOTAL (TB)	**126.7**	**178.2**	**256.1**	**377.0**	**514.5**	**685.3**

Figure 60

Retail Company Inefficient Scenario: IT Costs Detail

IT COST CATEGORIES	2003	2004	2005	2006	2007	TOTAL ($000)
Hardware acquisition	3,703	4,056	4,681	3,981	3,652	**20,073**
Hardware maintenance	483	573	659	743	820	**3,278**
Software licenses	338	478	655	682	685	**2,838**
Software support	135	174	224	260	307	**1,100**
SAN/NAS hardware/software	666	780	822	848	866	**3,982**
SAN/NAS maintenance/support	30	58	88	118	150	**444**
Personnel	6,257	6,937	6,968	5,671	4,968	**30,801**
Facilities	266	331	428	541	682	**2,248**
TOTAL ($000)	**11,878**	**13,387**	**14,525**	**12,844**	**12,130**	**64,764**

Figure 61

Retail Company Inefficient Scenario: Staffing Detail

IT FUNCTION	ALL STAFF	STORAGE FTEs				
	2002	2003	2004	2005	2006	2007
Management/administration	**36**	2	2	2	1	1
Technical staff	**28**	1	1	1	1	–
Applications	**359**	9	10	10	8	7
Database administration	**44**	6	7	7	6	5
System/storage management	**113**	40	44	43	34	30
Network management	**123**	2	2	2	2	1
Operations & support	**131**	12	13	13	10	9
End-user computing & help desk	**119**	2	2	2	1	1
TOTAL	**953**	**74**	**81**	**80**	**64**	**54**

Efficient Scenario

In this scenario, PeopleSoft Financials and HRMS systems (including HR administration, benefits, and payroll modules) have been deployed on a company-wide basis. There are more than 2,000 HRMS and 700 Financials users.

Additional detail on the initial server base for this scenario is presented in figure 62. Vendor-specific calculations are based on the platforms and software stacks shown in figures 63, 64 and 65 respectively. Additional detail on vendor configurations is presented in figures 66 through 68.

Figure 62

Retail Company Efficient Scenario: Initial Server Base

MAINFRAME SYSTEMS 2 x z900 1C3 S/390 R74 **DATA WAREHOUSE** 5 x pSeries 690	**SALES & MARKETING** 1 x pSeries 690 3 x pSeries 670 8 x pSeries 650 9 x pSeries 630 16 x xSeries 440
CORPORATE SYSTEMS 2 x pSeries 690 5 x pSeries 650 **SUPPLY CHAIN** 3 x pSeries 690 3 x pSeries 670 5 x pSeries 650 (central) 25 x pSeries 650 (remote) 4 x xSeries 440	**E-MAIL/GROUPWARE** 25 x xSeries 440 **END USER** 10 x xSeries 440 **WEB/RELATED** 45 x xSeries 440 **OTHER SYSTEMS** 2 x pSeries 670 2 x pSeries 650 5 x pSeries 630 21 x xSeries 440

Figure 63

Retail Company Efficient Scenario: Vendor Platforms

SYSTEMS	EMC	HDS	IBM	HP
Critical Systems 1	Symmetrix DMX	Lightning 9980V	ESS 800	–
Critical Systems 2	Symmetrix DMX	Lightning 9980V	ESS 800	–
Business Intelligence	Symmetrix DMX	Lightning 9980V	ESS 800	–
Secondary Systems	Clariion CX600	Thunder 9570V	FAStT 900	–
NAS	Celerra CNS Clariion CX600	Enterprise NAS Gateway Thunder 9570V	NAS Gateway 300 FAStT 700	NAS e7000 EVA 2C12D
SAN	Connectrix	Brocade 12000	IBM SAN Switch	–

Figure 64

Retail Company Efficient Scenario: Vendor Software Stacks

SYSTEMS	EMC	HDS	IBM	HP
Critical systems 1	CommandCenter Package ESN Manager PowerPath Enterprise Symmetrix Remote Data Facility (SRDF) TimeFinder	HiCommand Resource Manager Complete TrueCopy ShadowImage	ESS Specialist ESS Expert Peer-to-Peer Remote Copy (PPRC) FlashCopy	–
Critical systems 2	CommandCenter Package ESN Manager PowerPath Enterprise TimeFinder	HiCommand Resource Manager Open ShadowImage	ESS Specialist ESS Expert FlashCopy	–
Business intelligence	CommandCenter Package ESN Manager PowerPath Enterprise	HiCommand Resource Manager Open	ESS Specialist ESS Expert	–
Secondary systems	Access Logix Navisphere Manager SnapView	HiCommand Resource Manager Open ShadowImage	Base package Host access FlashCopy	Virtual Controller Software (VCS) Snapshot
NAS	Access Logix Navisphere Manager	HiCommand Resource Manager Open	Base package Host access	Virtual Controller Software (VCS)

Figure 65

Retail Company Efficient Scenario: EMC Configurations

DATA CENTER	Platform	2002	2003	2004	2005	2006	2007
Critical systems 1*	Symmetrix DMX	2 x 8 TB 1 x 6 TB	3 x 9 TB	2 x 10 TB 1 x 12 TB	2 x 12 TB 1 x 13 TB	2 x 15 TB 1 x 16 TB	3 x 16 TB 1 x 8 TB
Critical systems 2	Symmetrix DMX	2 x 9 TB	1 x 10 TB 1 x 12 TB	1 x 12 TB 1 x 13 TB	1 x 14 TB 1 x 15 TB	1 x 16 TB 1 x 17 TB	1 x 19 TB 1 x 20 TB
Business intelligence	Symmetrix DMX	1 x 21 TB 1 x 22 TB	2 x 22 TB 1 x 23 TB	3 x 33 TB 1 x 8 TB	5 x 34 TB	6 x 34 TB 1 x 33 TB	9 x 34 TB 1 x 13 TB
Secondary systems	Clariion CX600	2 x 4 TB	2 x 7 TB	2 x 9 TB	2 x 11 TB	2 x 13 TB	2 x 15 TB
NAS	Clariion CX600	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 8 TB	1 x 9 TB	1 x 11 TB

*Mirrored

Figure 66

Retail Company Efficient Scenario: HDS Configurations

DATA CENTER	Platform	2002	2003	2004	2005	2006	2007
Critical systems 1*	9980V	2 x 8 TB 1 x 6 TB	3 x 9 TB	2 x 10 TB 1 x 12 TB	2 x 12 TB 1 x 13 TB	2 x 15 TB 1 x 16 TB	2 x 18 TB 1 x 19 TB
Critical systems 2	9980V	2 x 9 TB	1 x 11 TB 1 x 12 TB	1 x 12 TB 1 x 13 TB	1 x 14 TB 1 x 15 TB	1 x 16 TB 1 x 18 TB	1 x 19 TB 1 x 20 TB
Business intelligence	9980V	1 x 21 TB 1 x 22 TB	2 x 34 TB	1 x 52 TB 1 x 54 TB	3 x 56 TB	4 x 56 TB 1 x 12 TB	5 x 56 TB 1 x 36 TB
Secondary systems	9570V	2 x 4 TB	2 x 7 TB	2 x 9 TB	2 x 11 TB	2 x 13 TB	2 x 15 TB
NAS	9570V	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 8 TB	1 x 9 TB	1 x 11 TB

*Mirrored

Figure 67

Retail Company Efficient Scenario: IBM Configurations

DATA CENTER	Platform	2002	2003	2004	2005	2006	2007
Critical systems 1*	ESS 800	2 x 8 TB 1 x 6 TB	3 x 9 TB	2 x 10 TB 1 x 12 TB	2 x 12 TB 1 x 13 TB	2 x 15 TB 1 x 16 TB	2 x 18 TB 1 x 19 TB
Critical systems 2	ESS 800	2 x 9 TB	1 x 10 TB 1 x 12 TB	1 x 12 TB 1 x 13 TB	1 x 14 TB 1 x 15 TB	1 x 16 TB 1 x 17 TB	1 x 19 TB 1 x 20 TB
Business intelligence	ESS 800	1 x 21 TB 1 x 22 TB	2 x 34 TB	2 x 35 TB 1 x 36 TB	4 x 42 TB	4 x 42 TB 2 x 34 TB	6 x 42 TB 2 x 33 TB
Secondary systems	FAStT 900	2 x 4 TB	2 x 6 TB	2 x 9 TB	2 x 11 TB	2 x 13 TB	2 x 15 TB
NAS	FAStT 700	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 8 TB	1 x 9 TB	1 x 11 TB

*Mirrored

Figure 68

Retail Company Efficient Scenario: Hewlett-Packard Configurations

DATA CENTER	Platform	2002	2003	2004	2005	2006	2007
Critical systems 1	–	–	–	–	–	–	–
Critical systems 2	–	–	–	–	–	–	–
Business intelligence	–	–	–	–	–	–	–
Secondary systems	EVA 2C12D	2 x 4 TB	2 x 7 TB	2 x 9 TB	2 x 11 TB	2 x 13 TB	2 x 15 TB
NAS	EVA 2C12D	1 x 4 TB	1 x 5 TB	1 x 7 TB	1 x 8 TB	1 x 9 TB	1 x 11 TB

METHODOLOGY

Sources of Data

General Survey

A survey was conducted of 124 companies in the three industries for which profiles are presented to establish patterns of and plans for storage deployment, organization and staffing for IT as a whole, and storage in particular; installed bases and growth rates for servers and storage; technologies, platforms and software employed; and actual vendor prices, terms and conditions.

Survey respondents employed a variety of centralized and distributed deployment models, and included a representative mix of customers of all of the vendors covered in this document. Most had already begun significant deployment of SANs.

Companies ranged from $1 billion to more than $50 billion in revenues or – in the case of banks and financial services companies – $12 billion to more than $600 billion in assets during their most recent fiscal year. Companies were headquartered or had a significant base of operations in North America.

Results indicated that, in competitive bids, vendor hardware pricing tended to be similar for equivalent configurations. Software license fees, and SAN and NAS offerings were also heavily discounted. There was, however, more variation in maintenance and software support pricing.

It was decided, on the basis of these results, to adopt the approach to efficient scenario comparisons described earlier (i.e. to employ the same street price hardware and software license cost values for all vendors, while varying maintenance, software support and personnel costs). There was little obvious variation in facilities costs between vendors, and the same values were again employed for equivalent configurations for all vendors.

General survey inputs were also employed to calculate values for inefficient scenarios.

Composite Profiles

Profiles are based on actual company business models, strategies and experiences, organization structures, application portfolios, server and storage bases, workloads, service quality levels and IT staffing. The profiles are composites drawn from a number of different companies, which were the subject of more in-depth survey.

The diversified bank profile includes data from nine North American banks with yearend fiscal 2002 assets of $100 to more than $600 billion. The manufacturing company profile includes data from 10 global consumer products manufacturers with $2 to $40 billion in 2002 revenues employing SAP R/3 and mySAP applications. The retail company profile contains data from five North American and two European grocery and drug retailers with $20 to $50 billion in 2002 revenues.

Basis of Configurations

For the inefficient as well as efficient scenarios, server and storage configurations as of yearend 2002, along with projected growth rates for these, are based on values reported by users for the categories, groups of system, and individual systems detailed for each profile.

Efficient scenario configurations for IBM pSeries, xSeries and zSeries servers are based on server numbers and models reported by users. These included HDS, IBM and Fujitsu (former Amdahl) mainframes; Fujitsu Siemens, Hewlett-Packard (including former Compaq), IBM, Silicon Graphics, Sun Microsystems, and other RISC/UNIX servers; and a wide range of Intel-based platforms.

Mainframe configurations were updated to zSeries models based on MIPS capacity, while RISC/UNIX and Intel-based servers were translated into pSeries models based on vendor relative performance data and industry norms. Degrees of consolidation were determined using best practice norms for types of system and server in each profile.

Disk storage systems configurations for efficient scenarios were based on raw capacity in terabytes reported by users. Systems were configured on a cross-platform basis according to availability and performance requirements of individual systems and groups of systems (e.g. applications most sensitive to downtime were supported by mirrored systems with concurrent replication as well as point-in-time copy capabilities). Others were supported by various RAID 5 configurations.

Configurations reflect practices reported by users. Respondents commonly, for example, preferred not to configure systems above a particular size – usually in the 10 to 20 terabytes range. This was to avoid performance constraints, vulnerability to system failures (there was a tendency to try to avoid what one respondent described as “putting too many eggs in the same basket”) or both. These concerns were less apparent for systems supporting less downtime-sensitive applications.

EMC, HDS, IBM and Hewlett-Packard systems in efficient scenarios are equipped with 146 GB drives. There are some differences in capacity between equivalent vendor configurations. These are, however, relatively minor and do not affect cost comparisons. Reflecting user practices, configurations do not exceed 80 percent of total raw capacity of any system.

There is considerable evidence that levels of availability delivered by latest-generation high-end RAID 5 systems are now equivalent to those of their RAID 10 counterparts. Companies surveyed tended, however, to continue to employ mirrored storage for their most critical applications. This situation, rather than inherent superiority of RAID 10 to RAID 5 technology, is reflected in efficient scenario configurations.

In all three profiles, efficient scenarios include consolidation of data centers, servers and storage, and rationalization and standardization of systems. These do not, however, have a significant impact on initial bases of and projected growth rates for disk storage compared to inefficient scenarios.

Calculations assume that overall capacity requirements for each type of system remain consistent between scenarios. Lower initial bases and growth rates in efficient scenarios are due primarily to improved utilization rather than reductions in numbers of data centers, servers and systems. This tends to understate the actual savings in capacity and costs that would be realized under efficient scenarios.

In inefficient scenarios, SAN deployment occurs at the rates shown in figure 69. Percentages refer to the overall volume of disk storage in data centers accounted for by SAN-attached systems.

Figure 69

SAN Deployment Rates in Inefficient Scenarios

COMPANY	2002	2003	2004	2005	2006	2007
Diversified bank	30%	50%	75%	85%	95%	100%
Manufacturing	15%	35%	55%	80%	90%	95%
Retail	20%	35%	60%	85%	95%	100%

In efficient scenarios, data centers are SAN-enabled, and all servers and storage are SAN-attached as of yearend 2002. Subsequent acquisition costs are for new attachments and infrastructure upgrades.

IT Cost Calculations

Cost Components

IT costs include hardware acquisition, maintenance, license and support costs for storage-specific software (including concurrent copy, point-in-time copy, core storage management and other tools as indicated for each profile); FTE personnel engaged in storage-related tasks; and facilities, including real estate, environmental and other data center costs for disk storage systems, SANs and NAS.

SAN and NAS costs are divided into two categories: acquisition (including hardware and software license fees) and support (including hardware maintenance and software support). This reflects the practice of some vendors of bundling hardware and software components for pricing purposes. Cost figures do not include servers, software other than as indicated above, networks other than SANs or NAS, costs of other IT personnel, or installation or other services costs.

Inefficient Scenarios

For inefficient scenarios, hardware costs are calculated based on $ per gigabyte ($/GB) "street price" values for the disk storage systems and capacities in each profile. Maintenance costs were calculated based on vendor warranties and percentages of "street price" acquisition cost for each year (e.g. 5.5 percent of hardware acquisition cost per year starting after a two-year warranty period). Software license and support costs were calculated in a similar manner.

Costs were calculated for various mixes of 28 mainframe and RISC/UNIX disk storage platforms from EMC, Hewlett-Packard (including former Compaq platforms), IBM, Silicon Graphics and Sun Microsystems platforms, along with software and SAN offerings from these and other vendors. Generalized cost assumptions were employed for Intel-based disk storage platforms and software. Costs include server-based software for direct-attached storage configurations.

Costs do not include hardware or software licenses acquired before 2003, but maintenance and software support calculations allow for systems installed up to five years before 2003.

Hardware costs, and maintenance costs calculated as a percentage of these, assume a decline of 25 percent per year in loaded costs per unit of capacity for disk storage systems, and 15 percent per year for SANs and NAS.

Efficient Scenarios

Hardware and software license acquisition costs were calculated using "street price" values as described earlier. Hardware maintenance and software support costs were calculated using the values shown in figure 70.

Figure 70

Maintenance and Support Values for Efficient Scenarios

VENDOR	HIGH-END HARDWARE		MIDRANGE HARDWARE		SOFTWARE	
	Warranty Duration	Maintenance per Year	Warranty Duration	Maintenance per Year	Warranty Duration	Maintenance per Year
EMC	2 years	7%	2 years	7%	3 months	15%
HDS	3 years	2.25%	3 years	5.75%	3 months	12.15%
IBM	3 years	2.175%	3 years	2.45%	3 years	5%
HP	–	–	3 years	6.45%	3 months	15%

Maintenance and support costs for SANs and NAS are based on two-year warranties and annual costs equivalent to 3.62 percent of acquisition. Costs include hardware and license acquisitions conducted during 2002, which constitute initial storage bases. Maintenance and software support calculations assume that warranty coverage for hardware and software licenses begins on January 1, 2003.

Other Costs

Personnel costs for inefficient and efficient scenarios are based on the annual salary averages shown in figure 71.

Figure 71

Salary Assumptions

FUNCTION	SALARY
Management/administration	$ 54,646
Technical staff	79,519
Applications	70,240
Database administration	83,824
System/storage management	66,106
SAN administration	96,352
Network management	59,212
Operations & support	51,058
End-user computing & help desk	45,456

Salaries are increased by 28.2 percent to allow for benefits, bonuses and other non-salary compensation, along with training, development and related items. It is assumed that all personnel costs per FTE remain stable over projection periods.

Facilities costs were calculated based on square footage occupied by disk storage systems, SAN and NAS devices, and related equipment in raised floor environments. A norm of $50 per square foot per year was employed. It is again assumed that this remains consistent over projection periods.

ABOUT THE INTERNATIONAL TECHNOLOGY GROUP

ITG sharpens your awareness of what's happening and your competitive edge . . . this could affect your future growth and profit prospects

The International Technology Group (ITG), established in 1983, is an independent information services and management consulting firm specializing in timely analyses of current developments and future trends in information systems technologies and markets.

Particular attention is given to analyzing trends in emerging technologies and markets. ITG's staff and supporting consortium of international analysts are well experienced and skilled in creative assessment techniques and research methodologies.

Services provided to clients are designed to provide factual data and reliable documentation to assist in the decision-making process. The information provided establishes the basis for developing tactical and strategic plans. Important developments are analyzed and practical guidance is offered on the most profitable ways to respond to the changes influencing and shaping complex IT deployment.

A broad range of services is offered, furnishing clients with the information necessary to complement their internal capabilities and resources. Customized programs are designed to provide a continuous source of technology, product and market information involving delivery of various combinations of the following services:

Status Reports	In-depth studies analyzing important topics and future industry trends
Flash Reports	Timely documents covering latest developments and significant market activity
Management Briefings	Periodic interactive planning meetings with management to address tactical issues
Telephone Consultation	Immediate response to clarify conclusions and review current informational needs
Executive Presentations	Scheduled interactive strategic meetings with decision-makers to examine future directions

The firm's **Enterprise Information Systems** service is particularly noteworthy. This service covers various architectures, technologies and techniques that may be effectively deployed in order to achieve the goals of enterprise-wide information systems. Scenarios being proposed by leading system suppliers and software vendors are objectively evaluated. Clients include systems suppliers, software vendors, service firms, financial institutions, government agencies, trade associations and end users.

International Technology Group

4546 El Camino Real, Suite 230
Los Altos, California 94022-1069
Telephone: (650) 949-8410
Facsimile: (650) 949-8415
Email: info-itg@pacbell.net