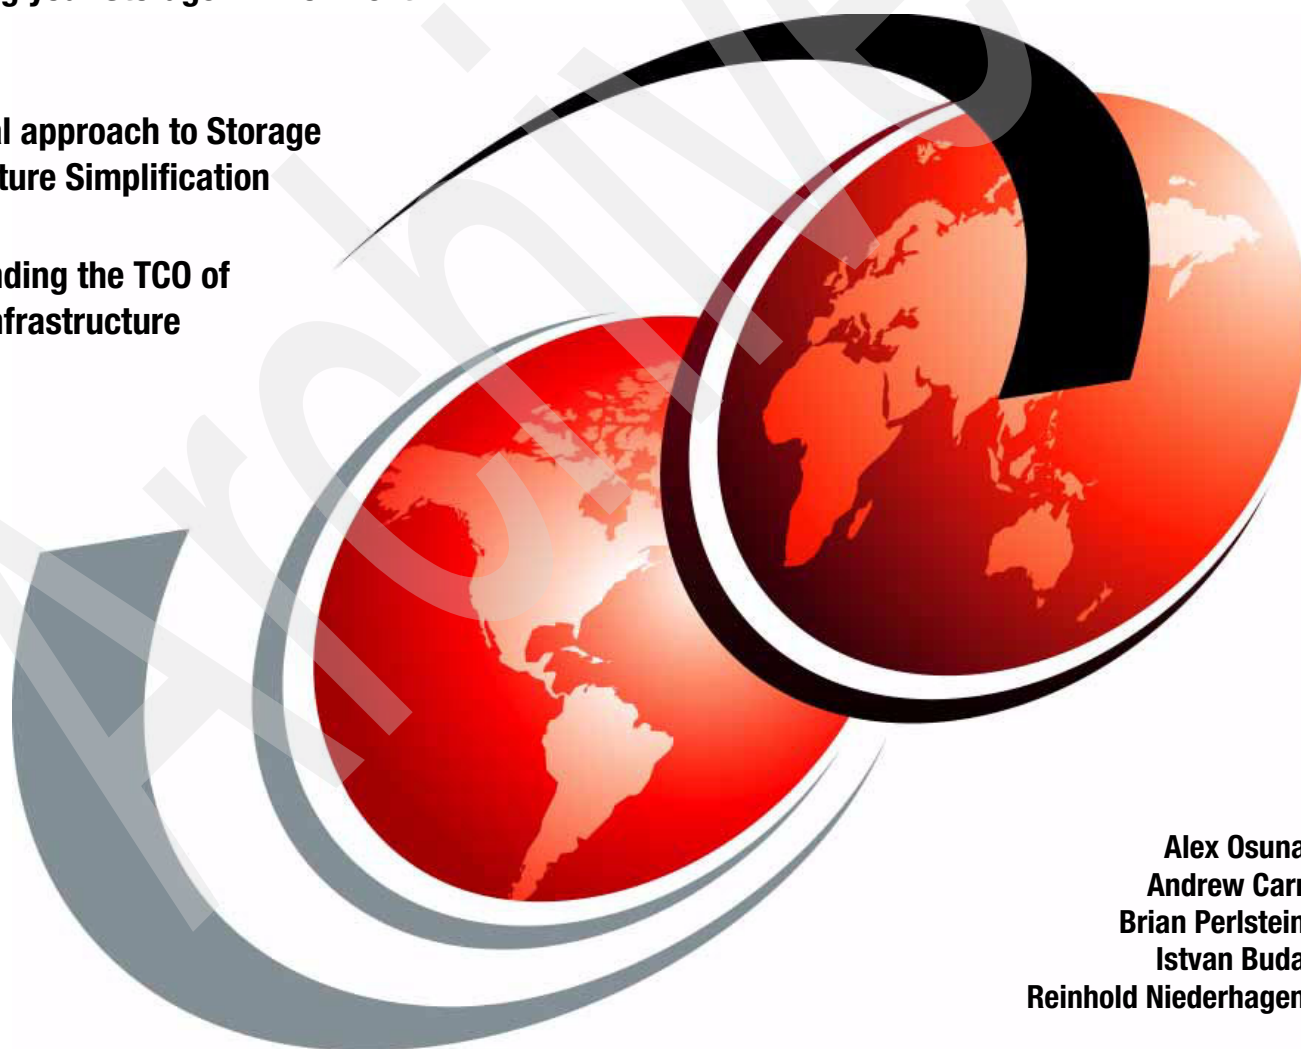


Introduction to Storage Infrastructure Simplification

Simplifying your Storage Environment

A practical approach to Storage Infrastructure Simplification

Understanding the TCO of Storage Infrastructure



**Alex Osuna
Andrew Carr
Brian Perlstein
Istvan Buda
Reinhold Niederhagen**

Redbooks



International Technical Support Organization

Introduction to Storage Infrastructure Simplification

January 2006

Archived

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

First Edition (January 2006)

This edition applies to version one of the current release of IBM TotalStorage and System Storage products as of August 2005.

© Copyright International Business Machines Corporation 2006. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures	vii
Notices	xi
Trademarks	xii
Preface	xiii
Become a published author	xiv
Comments welcome	xiv
Part 1. Infrastructure Simplification principles	1
Chapter 1. Introduction to Storage Infrastructure Simplification	3
1.1 Storage Infrastructure Simplification introduction	4
1.2 Infrastructure challenges	7
1.2.1 Mission data: Tough and getting tougher	7
1.3 Key Storage Administrator tasks	12
1.3.1 Manage Disk Performance	13
1.3.2 Storage Replication	14
1.3.3 Manage Storage Network Fabric	14
1.3.4 Provisioning and storage management task automation	15
1.3.5 Protect Data and Optimize Data Storage	16
1.4 Summary	16
Chapter 2. Total Cost of Ownership of storage infrastructure	17
2.1 What is TCO?	18
2.2 What a complex infrastructure costs	22
2.3 TCO and skilled staff	23
2.4 Storage administration costs	24
2.5 Costs and return on investment	25
2.6 Measuring total costs	27
2.6.1 Steps to calculate TCO	28
2.7 Costs of problems or faults	29
2.8 TCO with IBM TotalStorage	30
2.8.1 Infrastructure simplification	31
2.8.2 About the TCO examples	31
2.8.3 The cost advantages of DS6000	32
2.8.4 DS8000: A new standard for lowering TCO	32
2.8.5 Cost advantages of the SAN Volume Controller	32
2.8.6 DS6000 operational costs in detail	32
2.8.7 Operational costs in detail	33
2.8.8 DS8000 additional potential consolidation benefits	35
2.8.9 Analysis background information	36
2.8.10 How we calculate environmental costs	36
Part 2. Virtualization of IT resources	39
Chapter 3. Simplifying your fiber Storage Network	41
3.1 From where does the complexity come?	42
3.2 Interoperability	43
3.3 Simple Storage Networks	44

3.3.1 zSeries storage network development	45
3.3.2 How did these components simplify their infrastructure	46
3.3.3 Comparing a S/390 Storage Network with an open SAN	48
3.3.4 Simplifying disk storage	53
3.4 Tape backup infrastructure	57
3.4.1 Principles of backup solutions.	57
3.4.2 Simplifying tape infrastructure.	59
3.4.3 The next simplification step was to consolidate	60
Chapter 4. Case study in Infrastructure Simplification	63
4.1 The solution.	64
4.1.1 The benefits	64
4.2 Demonstrating Infrastructure Simplification	64
4.2.1 A strategic alliance with one focus: Enable the highest quality healthcare.	64
Part 3. Automated storage management	69
Chapter 5. Storage virtualization	71
5.1 Introduction to storage virtualization	72
5.1.1 Storage virtualization concepts	73
5.1.2 IBM and storage virtualization technology	75
5.2 IT storage architectural directions	80
5.2.1 Automated Management.	81
5.3 Types of storage virtualization	81
5.3.1 The IBM TotalStorage vision.	83
5.3.2 Storage virtualization models	83
5.3.3 SAN Volume Controller in-band storage virtualization benefits	86
5.4 IBM TotalStorage SAN Volume Controller	87
5.4.1 Block virtualization	88
5.4.2 SAN Volume Controller characteristics.	90
5.4.3 SAN Volume Controller interoperability.	94
5.5 IBM TotalStorage SAN File System	97
5.5.1 File virtualization	98
5.5.2 SAN File System characteristics.	99
5.5.3 SAN File System architecture	100
5.6 IBM System Storage N Series	103
5.6.1 Windows environments.	104
5.6.2 Consolidation for all Platforms	105
5.6.3 High Availability.	106
5.7 Management and productivity	107
5.7.1 Storage Management Initiative	108
5.7.2 Open storage management with CIM	108
5.7.3 IBM storage management of the virtualized SAN	108
Chapter 6. Software to simplify managing your storage systems	111
6.1 IBM TotalStorage Productivity Center.	113
6.1.1 IBM TotalStorage Productivity Center for Data.	114
6.1.2 IBM TotalStorage Productivity Center for Disk	118
6.1.3 IBM TotalStorage Productivity Center for Fabric.	120
6.1.4 IBM TotalStorage Productivity Center for Replication.	121
6.1.5 Benefits.	122
6.2 IBM Tivoli Provisioning Manager.	123
6.3 IBM Tivoli Storage Manager	126
6.4 Customer C: Productivity Center in real life	127

Chapter 7. Infrastructure Simplification using the IBM System Storage N series	131
7.1 Challenge: Expanding Linux workload: Wild server crashes.	132
7.2 Unified solution: N Series System Storage	132
7.3 Business benefits: Fast sequencing	133
7.4 Secure Storage and Backup of Regulated Data	134
7.5 Simplicity and minimal IT overhead	134
Chapter 8. A storage provisioning solution for SAN File System	137
8.1 An on demand storage environment	138
8.2 The SAN File System solution	139
8.2.1 SAN File System physical environment	139
8.2.2 Storage usage.	140
8.3 Storage Provisioning to simplify volume management	142
8.3.1 SAN File System automation tasks.	143
8.3.2 Modelling of SAN File System in TPM	144
Part 4. Advantages of using IBM TotalStorage for Infrastructure Simplification	145
Chapter 9. Competitive advantages to Infrastructure Simplification using IBM TotalStorage	147
9.1 Virtualized versus array replication services	148
9.2 Non-disruptive data migration	149
9.3 Storage pooling	149
9.4 Virtual volumes	150
9.5 Big Performance and Capacity, Smaller Package	151
9.6 Based on the best	152
9.7 SAN File System Architecture.	152
Part 5. Storage environment consolidation	157
Chapter 10. Consolidation principles	159
10.1 Storage consolidation	160
10.2 Benefits of consolidation	163
10.3 Key IBM storage consolidation products.	163
10.4 Storage consolidation examples	163
10.4.1 DS6000 storage consolidation	163
10.4.2 DS8100 storage consolidation	164
10.4.3 SVC promotes consolidation.	166
Chapter 11. Infrastructure Simplification in action	169
11.1 Introduction	170
11.2 History	170
11.3 Phase One research and design (Don't rush it)	171
11.4 Product acquisition	173
11.4.1 Configuring the switch.	173
11.4.2 Lessons learned troubleshooting cable issues	175
11.4.3 The core-edge debate.	175
11.4.4 Spending money is spending money! Finding the ROI in building a SAN	176
11.4.5 Backup and recovery	179
11.4.6 Environments: Have you seen my electric bill lately?	179
11.4.7 Post implementation: Was the environment simplified?	180
11.5 Phase two: More consolidation	180
11.5.1 Could it have been too good to be true?	181

11.5.2 Lessons learned: Educate and translate	182
11.5.3 Next steps	183
11.5.4 If I could do it again with more money, and technology was available	183
Abbreviations and acronyms	187
Related publications	189
IBM Redbooks	189
Other publications	189
Online resources	190
How to get IBM Redbooks	190
Help from IBM	190

Figures

1-1	Heterogeneous IT environment	5
1-2	Benefits of Infrastructure Simplification	6
2-1	TCO indirect and direct	18
2-2	Direct costs of cars	20
2-3	Hidden costs of the IT departmental car	21
2-4	The cost of managing a complex environment	24
2-5	Simple environment	25
2-6	TCO analysis	28
2-7	Costs of down time	30
2-8	DS6000 storage consolidation	33
2-9	DS8100 storage consolidation	34
2-10	DS8100 storage consolidation compared with ESS F-20 and OEM	34
2-11	DS8300 storage consolidation	35
3-1	SAN sprawl	42
3-2	Sample Interoperability matrix for an IBM DS8000	43
3-3	Interoperability: Everything has to fit	44
3-4	Development of zSeries storage network	45
3-5	Port view of 9032 model 3	48
3-6	Meshed fabric	51
3-7	High availability meshed fabric	52
3-8	HA fabric with high port count switches	53
3-9	Managing a complex heterogeneous SAN	54
3-10	Common disk system	55
3-11	SVC with heterogeneous storage	57
3-12	Basic backup solution	58
3-13	Consolidated solution with shared drives	59
3-14	LTO technology	59
3-15	Tape backup block or module	60
3-16	Overview of a tape backup module in a 3584	61
3-17	WAN-connected clients	62
4-1	Customer M's simplified Storage Infrastructure	66
4-2	Lan free backups	66
4-3	Storage reservoir	67
5-1	IBM vision of the On Demand storage environment	72
5-2	Basic concept of storage virtualization	74
5-3	The IBM SAN Volume Controller 2145-8F2	75
5-4	Cisco MDS 9000 Caching Services Module	76
5-5	N Series versus SAN	78
5-6	N Series connectivity	79
5-7	IBM N3700	79
5-8	N5200	79
5-9	Trends	80
5-10	In-band virtualization model	84
5-11	Out-of-band virtualization model	85
5-12	Scalability: 256/1024 Host Support	87
5-13	SNIA block aggregation model	89
5-14	IBM plan for block virtualization	90
5-15	SAN today	91

5-16	Block level virtualization	91
5-17	IBM SAN Volume Controller	92
5-18	Infrastructure simplification with SVC	93
5-19	Capacity utilization	94
5-20	IBM TotalStorage SAN Volume Controller supported environments	95
5-21	Cluster creation	96
5-22	SAN File System	98
5-23	SNIA file aggregation model	99
5-24	IBM plan for file virtualization	100
5-25	IBM TotalStorage SAN File System	101
5-26	Supported operating systems and storage devices	102
5-27	Consolidation Architecture	105
5-28	Consolidation	106
5-29	High availability with the N Series	107
6-1	Productivity Center Launchpad	113
6-2	Productivity Center for Data: General overview	115
6-3	Productivity Center for Data: Asset information	116
6-4	Productivity Center for Data: Total free space	117
6-5	Productivity Center for Data: Wasted space	117
6-6	Productivity Center for Data: Project future needs	118
6-7	Productivity Center for Disk: storage devices panel	119
6-8	Productivity Center for Fabric console	120
6-9	IBM Tivoli Support Web page	121
6-10	Replication Manager's tasks	122
6-11	Storage provisioning: Manually	124
6-12	Storage provisioning supported by Tivoli Provisioning Manager	125
6-13	TPM workflow: Growing a file system on a Windows host	125
6-14	IBM Tivoli Storage Manager server	127
6-15	Customer C data in the Productivity Center for Data: Enterprise-wide Summary	128
6-16	Customer C: Trending utilization of disk space	128
6-17	Customer C: Largest Files by filesystem Report: Shows possible duplicate files	129
6-18	Customer C: Report on large file systems running full and containing stale files	129
7-1	Before and after consolidation	133
8-1	On demand environment	138
8-2	Physical SAN File system environment	140
8-3	Shared storage environment	141
8-4	Orchestrated Web serving using SAN File System	141
8-5	Provisioning tasks to introduce a new SAN File System client	142
8-6	SAN File System user pool managed by TPM	144
9-1	Traditional SAN versus SVC	148
9-2	Non-disruptive data migration	149
9-3	Storage pooling	150
9-4	Virtual volumes	151
9-5	SAN File System architecture	152
9-6	Single name space	153
9-7	Single point of management	154
9-8	Heterogeneous file sharing	155
9-9	Policy-based file placement	155
9-10	Storage provisioning	156
10-1	Traditional SAN with distinct storage islands	160
10-2	Consolidated SAN environment	161
10-3	Consolidated storage environment	162
10-4	DS6000 storage consolidation	164

10-5	DS8100 storage consolidation	165
10-6	DS8100 storage consolidation versus ESS F-20	165
10-7	DS8300 storage consolidation	166
10-8	Array pooling	167
11-1	SCSI attachment	170
11-2	First SAN	171
11-3	Full scale SAN implementation	173
11-4	Cable strategy	175
11-5	Large servers	177
11-6	Small servers	178
11-7	IBM BladeCenters	181
11-8	SVC inclusion	182
11-9	If I had to do it over again	186

Archived

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law. INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server®
@server®
Redbooks (logo) ™
eServer™
iSeries™
pSeries®
xSeries®
z/OS®

zSeries®
AIX®
BladeCenter®
Enterprise Storage Server®
ESCON®
FlashCopy®
IBM®
POWER5™

Redbooks™
S/390®
System Storage™
System/390®
Tivoli®
TotalStorage®
Virtualization Engine™

The following terms are trademarks of other companies:

Solaris, Sun, and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows server, Windows, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

NetApp, the Network Appliance logo, the bolt design, DataFabric, FAServer, FilerView, MultiStore, NearStore, NetCache, SecureShare, SnapManager, SnapMirror, SnapMover, SnapRestore, SnapVault, SyncMirror, and WAFL are registered trademarks and Network Appliance, ApplianceWatch, BareMetal, Camera-to-Viewer, Center-to-Edge, ContentDirector, ContentFabric, Data ONTAP, EdgeFiler, HyperSAN, InfoFabric, NetApp Availability Assurance, NetApp ProTech Expert, NOW, NOW NetApp on the Web, RoboCache, RoboFiler, SecureAdmin, Serving Data by Design, SharedStorage, Smart SAN, SnapCache, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapMigrator, Snapshot, SnapSuite, SohoCache, SohoFiler, The evolution of storage, Vfiler, VFM, Virtual File Manager, and Web Filer are trademarks of Network Appliance, Inc. in the U.S. and other countries. All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

Preface

This IBM® Redbook introduces Infrastructure Simplification. Infrastructure Simplification is the methodology of analyzing the complete enterprise: business processes, workflow environment end-to-end, and IT for simplification. This analysis yields opportunities to save you time and money and eliminate unnecessary complexity that impedes the flow of information. This IBM Redbook discusses Storage Infrastructure Simplification and demonstrates multiple ways that IBM TotalStorage® and Infrastructure Simplification can help you reduce complexity, save time and money, and release the flow of information in your business.

The team that wrote this book

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization, San Jose Center.



Figure 1 Alex, Istvan, Reinhold, and Andy

Alex Osuna is a Project Leader in the San Jose International Technical Support Organization. He has more than 20 years of experience with IBM Storage in areas of planning, development, support, performance, product education, product introduction, and systems engineering.

Andrew Carr (Andy) is a storage specialist in the UK. He has 16 years of experience in the storage field. His areas of expertise include tape, SAN, and disk storage systems. He has written previously about SAN and pSeries®. His previous experience includes implementing storage systems with global services, pre-sales support, product planning, and ITS hardware support.

Brian Perlstein (not pictured) is a technical architect with CareTech Solutions, an information technology outsourcing provider that services Oakwood Healthcare System's IT department. During his nine-year career with Oakwood, Brian has specialized in the implementation of hardware across the healthcare system. Currently, Brian is concentrating on enhancing the storage infrastructure throughout Oakwood's Information Technology environment.

Istvan Buda is a Senior Support Engineer and has been working for IBM Business Partner Distributor Trade Ltd. in Budapest, Hungary, since 2002. His areas of expertise include IBM pSeries and IBM TotalStorage products. He has over 10 years of IT experience. Istvan holds the following certifications: IBM Certified Specialist High-End disk Solutions, Open Systems Storage Solutions, and pSeries Enterprise Technical Support.

Reinhold Niederhagen is working as a System Engineer in IBM Global Services in Germany. He has more than 25 years of field experience in the IT industry, almost 20 years within IBM. His areas of expertise include application and system programming on several

platforms. He has experience as a project leader in the mainframe area. In his current position, he consults with several larger German customers in the mainframe and workstation area.

Thanks to the following people for their contributions to this project:

International Technical Support Organization, San Jose California

Mary Lovelace
Emma Jacobs
Bert Dufrasne
Leslie Parham
Jon Tate

IBM Storage Marketing

Rajesh Krishnan
Jim Tuckwell
Tony Pearson
Roger Wofford

Steve Strutt
IBM Consulting IT Specialist IBM TotalStorage

Jerry L. Bellomy
Sr. Competitive Specialist-Storage IBM Software Group

Tricia Jiang
IBM TotalStorage Technical Evangelist

Evan Salop
IBM Global Services, Integrated Technology Services

Karin von Kaenel
IBM Research

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks™ to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an e-mail to:

redbook@us.ibm.com

Archived



Part 1

Infrastructure Simplification principles

In this part, we give you a basic introduction to Infrastructure Simplification and discuss key topics related to Infrastructure Simplification. In addition, we identify the Storage Administrator tasks so that you can better understand which areas of Infrastructure Simplification are most beneficial to your business.

Archived



Introduction to Storage Infrastructure Simplification

This chapter discusses the challenges associated with the increasing demands to store more and more information. The amount of data today that a company needs to keep is constantly growing. The necessity to store this amount of data and access it from different places and systems has caused dramatic changes in the IT infrastructure.

This chapter introduces and describes an important concept called Storage Infrastructure Simplification.

1.1 Storage Infrastructure Simplification introduction

Increasing amounts of disks, switches, hubs, disk systems, and tape systems have to be installed to satisfy the growing need of information. Heterogeneous hardware and software environments of multiple vendors are installed and have to be maintained. Various generations of technology have to be compatible with each other. Your IT department has to work with all different types of software and hardware using a multitude of different tools.

Handling this vast aggregation of hardware and software is getting more and more complex and time consuming. Problem determination and management in such a complex environment can become a nightmare.

Complexity has been and continues to be the biggest deterrent to achieving access. Most of you, maybe all of you, did not choose this complexity. You inherited it (or maybe it inherited you). It grew from acquisition: Of competitors, resources, people and processes. Other factors have compounded the problem, the need to retain data for longer periods, sometimes indefinitely. Security concerns also come into play, the need to protect your business, your customer's sensitive information, your competitive advantage, and intellectual property. Information is what drives growth, but information has varying relative value to the business, until now, it has been difficult to align storage costs with this value.

Storage Complexity

Today's storage infrastructure complexity is an inhibitor. Storage infrastructure is inflexible, difficult to manage, and its costs are not aligned with the value of information.

Compounding the problem:

- ▶ Tight budgets and the need for higher return on investment (ROI)
- ▶ Competitive pressures and demands for data availability
- ▶ Regulatory compliance and security concerns

Nowadays, IT people face a broad range of different systems, architectures, and software components with which they have to work. They have to keep their knowledge of hardware and software from different vendors on a current level and still have to work with equipment from a previous generation. Figure 1-1: Heterogeneous IT environment shows a fully grown complex environment.




Figure 1-1 Heterogeneous IT environment

Infrastructure Simplification is all about helping your business meet its goals. Complexity can prevent even the best companies from being able to act nimbly to meet ever-changing market and client demands. Goals that are set and agreed upon can stall when the infrastructure makes it too complex or slow to accomplish these goals. Infrastructure Simplification can help you meet your business goals and:

- ▶ Remove difficulty in accomplishing administrative tasks and deploying new applications
- ▶ Reduce training requirements to move with greater speed
- ▶ Help eliminate risks of human error that can lead to data loss
- ▶ Save time that can be deployed to other, higher-level organizational goals
- ▶ Cut costs

Infrastructure Simplification is part of the journey toward creating an On Demand business. It provides the basis for a stable storage environment that helps, not hinders, business growth and adaptation.

So how do you escape this complexity and return to an infrastructure which can reliably support the rigors of business growth...through consolidation, virtualization, and automated management. These capabilities are the basis for changing from a complex storage environment to a simpler one.

These are the capabilities which you should consider when simplifying your storage environment. In this redbook, we discuss hardware, software, and procedures that help you simplify your environment. Guidelines help you to find the right point to start with simplification and show you what you can do next. Infrastructure Simplification is all about reducing complexity and costs across your IT environment. IBM can help you with consolidation, management centralization, and capabilities for unifying heterogeneous storage islands

through virtualization, resulting in new levels of flexibility and choice. See Figure 1-2: Benefits of Infrastructure Simplification.

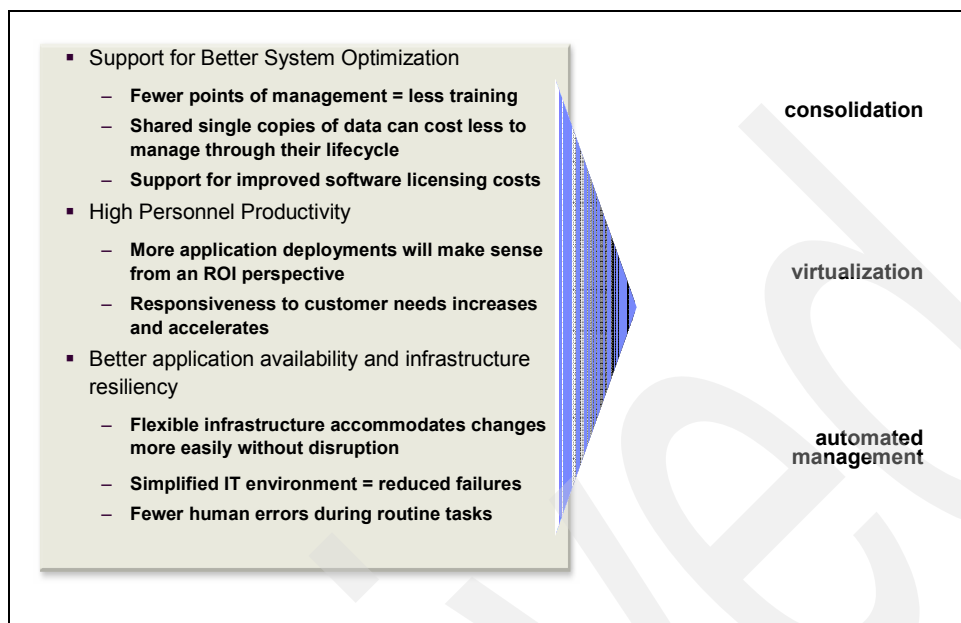


Figure 1-2 Benefits of Infrastructure Simplification

Figure 1-2: Benefits of Infrastructure Simplification represents how you can escape the complexity and return to an infrastructure which can reliably support the rigors of business growth through consolidation, virtualization, and automated management. These capabilities are the basis for changing from a complex storage environment to a simpler one.

Infrastructure Simplification is part of the journey toward becoming an On Demand business. Infrastructure Simplification provides the basis for a stable storage environment that helps, not hinders, business growth and adaptation.

Is it time for you to simplify your infrastructure?

Consider the answers to the following questions:

- ▶ Are the costs of managing your server and storage infrastructure increasing?
- ▶ Do you deal with increasingly complex server and storage infrastructure management?
- ▶ Have you been asked to reduce the budget for your IT organization without reducing service levels?
- ▶ Are you tired of paying maintenance for servers and storage of which you are only using a minimal percentage?
- ▶ Do you need to lower your data center facility cost?
- ▶ Do you have problems with data backup because of your dispersed and inconsistent storage landscape?
- ▶ Do you have resource constraints?
- ▶ Will you get continuing support for your current operating system environment?

1.2 Infrastructure challenges

The following material is paraphrased from “*Take the next step in creating an effective storage infrastructure*” by Marc Farley, which is available for purchase as an IBM publication, order number G224-7269.

1.2.1 Mission data: Tough and getting tougher

We do not know any successful businesses storing less data today than yesterday. Businesses today strive to collect as much information as possible to understand their customers, partners, vendors, and competitors better. As a result, executives ask IT managers to manage increasing amounts of data on all types of systems, from e-mail servers to database systems. The pressure on the IT staff to provide responsible data stewardship is also on the rise. Among the challenges that the IT staff faces are:

- ▶ Lack of interoperability. Theoretically, storage networks can be built with a variety of products from multiple vendors, designed to best fit the needs of the organization. Unfortunately, storage networking technology has suffered from proprietary approaches and a lack of agreement about how to implement fundamental functions, causing confusion and prolonged, expensive certification testing. This has resulted in homogeneous islands of technology.
- ▶ Working with multiple, complex management interfaces. Given the backdrop of homogeneous solutions, storage products have been managed as point products with distinctly different management interfaces and control methods. This fragmented management environment can create training and coverage challenges for IT organizations as well as increase administrator errors.

Businesses today strive to collect as much information as they can, and, as a result, IT managers are being asked to manage increasing amounts of data.

- ▶ Managing multiple access control methods. While storage networks can deliver many benefits, the access paths have to be carefully monitored and managed. Data integrity is typically protected by the combination of switch zoning and logical unit (LUN) masking techniques, which are not managed together under a single coherent connection management system. Therefore, all changes to the network generally have to be checked and double-checked for safe operations.
- ▶ Creating and managing very large storage resources. Data growth is unpredictable, potentially causing availability problems and putting data at risk when storage subsystems reach their maximum sizes. While storage networks have largely been justified for their ability to scale, the process of scaling storage capacity still holds unnecessary risks. In addition, there are many industries with applications that need extremely large file system implementations, avoiding the lost productivity and confusion associated with multiple file systems and mount points.
- ▶ Storage administration costs are too high. The administration costs per gigabyte of data in open-systems storage are much higher than they are for mainframe storage. Current solutions have multiple, overlapping control points, potentially resulting in duplication of effort, and additional, costly audit procedures.
- ▶ Migrating data can be difficult and disruptive. The simple concept of migrating data from one repository to another becomes monumental in large-scale, high-availability environments, which are increasingly common.
- ▶ Storage solutions may not be cost-effective. Applications have various performance requirements and these applications are valued differently by the organization. However, many storage products can have somewhat rigid configurations with narrow pricing options to address those differences.

- ▶ Premature obsolescence of storage products. Storage subsystems eventually reach their full capacity and are then replaced by larger subsystems. The historical lack of interoperability and management standards has made it difficult to continue using these still-valuable resources. This is far from optimal for a responsible, cost-conscious storage resource plan.
- ▶ Providing appropriate, cost-effective business continuity support. Business continuity has been achieved and proven, but many current products are too costly for many applications. A range of solutions to match application priorities is needed to create a comprehensive business continuity strategy.
- ▶ Storage resource constraints on business continuity. Business continuity solutions implemented in storage subsystems are often able to work only for data stored on those subsystems. This means an application can be limited to storing its data on a single subsystem, which can lead to scalability and capacity problems and can ignore the possibility that an application, such as a database, can have different storage requirements for its multiple types of data.
- ▶ Implementing capable backup protection for all applications. Backup is chronically problematic due to the large amounts of data and the disappearing backup window for performing the job. Considered by many to be the fundamental application for responsible data stewardship, backup is not working correctly for many IT professionals, and it continues to bewilder and frustrate diligent efforts to solve it.
- ▶ Responding successfully to emerging governmental regulations. A growing stream of governmental regulations is being passed that has direct impact on data and storage management operations. In the United States, various new federal, security, and health information laws contain mandates for data retention, privacy, and availability. The writers of these regulations make their intent clear, but sometimes leave the actual measurable objectives open for speculation. Responding to regulations that lack metrics is frustrating, especially when they include opposing objectives, such as ensuring that data is available on demand to auditors while simultaneously guaranteeing it is also held secure.

While there are many symptoms to address, such as those listed, it is important to diagnose the source of the problem correctly. The enormous amount of data and its rapid growth rates are pushing storage technologies into uncharted territory on a regular basis. Heritage direct access storage (DAS) products are severely underpowered to deal with today's data and storage management problems. That is, of course, why the market has turned to storage networking technology for solutions.

What is infrastructure and what does it do?

If we are going to think about network storage infrastructure, we may want to step back from the details and make a few observations about other infrastructures and what makes them effective.

The role and importance of infrastructure

We create infrastructure to facilitate the use of essential resources. Centuries ago, Roman aqueducts ran for hundreds of miles throughout Europe, providing a stable and reliable source for water. A similar type of infrastructure exists today in the form of electric power distribution. Transportation infrastructures, such as highway, rail, shipping, and air travel systems combine to provide distribution of essential resources throughout the world. In general, infrastructures that provide a broad range of services are likely to be productive over an extended period of time.

A more recent example of a global infrastructure is the fiber-based communication networks that were funded and built during the "dot com" boom in the nineties.

This infrastructure is the base for:

- Call centers in India
- Software Development houses in China
- Departments where team members work on every continent

These new developments are based on inexpensive access to large bandwidth standard networks that circle the globe. A standard high bandwidth storage network in your organization can encourage similar innovation and benefits.

Another example is locating restaurants for easy access by car. The fast-food industry took full advantage of the transportation infrastructure.

Infrastructures provide nonlinear economic benefits by enabling unforeseen business opportunities. Anyone could have predicted the growth of the petroleum and automobile industries with the construction of the U.S. highway system, but far fewer people envisioned the explosive growth of the fast-food restaurant industry. Successful restaurateurs quickly learned the key to success was leveraging people's driving habits. By locating restaurants for easy access by car and building drive-through ports, the fast-food industry took full advantage of the transportation infrastructure.

You can argue that the Internet is the only true, global infrastructure in existence, with its ability to instantly locate resources anywhere in the world. Speculations of its ability to spawn new commercial opportunities are everywhere, but it is far from clear what the eventual impact will be. The effects are profound and the businesses that can leverage this new infrastructure to the fullest are the businesses most likely to reap the greatest rewards. Sound network storage infrastructures will provide optimal data availability for Internet applications.

Goals of infrastructures

Infrastructures are built to be used broadly and heavily. They need to be accessible to individuals and organizations that want to use them and they need to support the demands of many users. The two primary characteristics of the most successful infrastructures are:

- ▶ Flexibility—to facilitate usage and availability of resources
- ▶ Stability—for a reliable, safe environment that supports businesses and lifestyles

Balancing flexibility and stability

The twin goals of flexibility and stability often work at cross purposes. Stability is preserved by limiting change where flexibility is created by enabling change. Creators, planners, and managers of infrastructures often struggle with these conflicting goals.

It is difficult and expensive to create stability in an infrastructure that was designed to be inherently flexible.

Tip: Following Murphy's Law, the need for increased stability often is not recognized until the infrastructure is already heavily used.

As we saw during the massive power blackout that struck the Eastern U.S. and Canada in the summer of 2003, one of the most obvious weaknesses in the power grid infrastructure is the presence of too many control points without a single system-wide control center. You can argue that the failure occurred because the lack of a control structure rendered the grid unmanageable.

The opposite approach is to make a stable infrastructure more flexible by adding access points or removing existing barriers. This has certainly happened in the information systems infrastructure, where 30 years ago, most of the systems were mainframes, and you had to be

a systems programmer to access data. Of course, today, almost anybody can walk into a library and access the Internet on a Personal Computer. A lot of technology needed to be invented along the way, which resulted in the birth of very large industry segments. There are many lessons and valuable perspectives, not to mention technologies, that could be used in the development of storage network infrastructures.

A third approach, and one that seems to be the most effective, is to create an infrastructure from its inception with the understanding that flexibility and stability need to be balanced. The telephone switching infrastructure developed by AT&T in the United States during the past century is a terrific example of a stable and flexible infrastructure. It was fast, easy to use, and worked even when power disruptions occurred. It also convincingly proved the power of automation as a stabilizing infrastructure force.

Flexibility and cost-efficiency of infrastructures

An infrastructure will not last if it is not flexible enough to provide economic incentives throughout the range of services it offers. Even relatively expensive infrastructure elements should have justifiable costs for both service providers and their customers. The service options and pricing for freight services illustrate how a complex infrastructure can provide a wide range of service levels with competitive prices. The cost of using an infrastructure and the value it delivers determine its cost-efficiency. If an infrastructure is not cost-efficient, it will be replaced by other technologies and services, regardless of the investment made in the infrastructure.

Selecting partners to build a storage infrastructure

Creating a storage network infrastructure is a large process that should involve many different perspectives and inputs. Companies undertaking this work should identify partners who can help with the design of the infrastructure, the selection of products, and the operating plan, including maintenance.

Design partners

The initial design of a storage network infrastructure does not have to encompass everything that is eventually desired, but it needs to provide for the flexibility to add future functions. Care should be taken to avoid proprietary capabilities that can become obsolete.

The selection of partners, that will help with the design, needs to be made based on the partners' experience and perspectives on storage, networking, and facilities management. Vendors with professional services organizations that can implement third-party products can be excellent candidates, because they know how to build cross-vendor solutions, but they also have access to specialized resources from within their product groups.

Vendor partners

If a chain is only as strong as its weakest link, it is imperative to have strong, lasting vendor partnerships. The selection of vendors for an infrastructure can include new vendors with new technology, but it should obviously include stable vendors that have long-term strategic goals for their network storage products and that believe in forging strategic partnerships with their customers. The number of companies that can do this is not very large, but these companies also tend to have solid, trustworthy technology that helps them maintain their strong customer relationships.

Operations partners

An infrastructure is a sizable, dynamic facility that needs to be monitored and controlled. There are good reasons for having partners that can help operate the environment and assist with troubleshooting. Maintenance and service partners are likely to include vendor partners, but they can also include experienced, third-party companies for services, such as business

continuity and audit preparations. Obviously, the products in an network storage infrastructure should be covered by service and maintenance agreements, that include options for specifying minimum response times and maintaining on-site spare parts.

Moving forward

Data stewardship is not a noble task; it is a necessity. The ability to respond to its challenges depends on the capabilities of the network storage infrastructure that is built.

Developing an infrastructure is a serious undertaking. While there are certainly challenges, the process can be managed effectively by setting clear goals and attainable objectives at each step along the way. It is also important to have solid metrics in place early in the process to help drive decisions, gain consensus, and minimize confusion. Here are a few thoughts on how to do this.

Forming a team

It is important to get a team in place if you do not already have one. The team should represent different groups within IT to ensure that their requirements are considered as well as making sure that important information about the project is communicated throughout the organization. The team depends heavily on individuals who have in-depth technical expertise of the systems and architectures already in place. The team also depends on disciplined processes and patience. One of the primary objectives in forming a team is to get individuals who can strike a balance between technology exploration and thorough planning.

Establish goals, objectives, and metrics

Keeping an eye on high-level strategic goals makes the process much easier for everybody involved. For any infrastructure development, the two primary goals are to create flexibility and stability. Storage-specific goals need to include scalability, availability, and manageability. Objectives are short-term targets for incremental, tactical projects. Examples of objectives include such things as researching, testing, comparing, selecting, and implementing designs and products. Metrics provide the means to measure products as well as the work that has been done. You need to include the guiding principles of virtualization, redundancy, and integrated management in the metrics for any storage network products.

Building experience

Storage networking technology has a definite learning curve for which you must account, not just by businesses, but also by storage networking vendors. It is simply not realistic to assume that products will be able to deliver on all their future promises, nor is it realistic to think that every decision made by the team will be optimal, despite the team's best efforts. It is essential for the team to get real-world experience with products and designs. The most effective infrastructures will include detailed insights based on designing, implementing, and troubleshooting product implementations.

Establishing a budget

An infrastructure cannot be built without spending money. The budget needs to include human resources, products, and services. You need adequate funding from the inception to explore products and designs. A good strategy is to select products with the flexibility to adapt to architectural changes. For instance, a product that is initially considered as the central part of an implementation may need to make the transition into a role where it functions as part of a distributed function in the infrastructure. Products that lack standards-based management can become prematurely obsolete.

At some point, the budget needs to include a build-out of strategic components. It is difficult, if not impossible, to attain stability without adequately funding the critical mass of strategic products that creates the foundation for stability. This build-out can begin as soon as the

strategic elements of the infrastructure have been selected. IT management will need to synchronize its budget processes with milestones in the infrastructure project.

CONCLUSION

The “perfect time” to start working on a storage network infrastructure is now. There is a good chance you will be working on it for several years, and the only way to finish the major work sooner is to start as soon as possible. The element of experience should not be underestimated, and, although, I write books and articles on the topic, there is only so much you can learn from books, most of the things you need to learn come from working closely with vendors, trial and error, and from comparing notes with team members and partners.

1.3 Key Storage Administrator tasks

Businesses are finding that complexity is hurting their ability to act and respond in the marketplace, due to:

- ▶ High costs both for equipment acquisition and ongoing operation
- ▶ Large training requirements and unpredictable staffing needs
- ▶ Unwanted risk of human error and failure that comes with complexity
- ▶ Excess time spent on IT management to deal with day-to-day requests
- ▶ Difficulties in responding rapidly to needs because of the inflexibility of older systems without virtualization

The last three bullets have the greatest effect on the Storage Administrator. The following are examples of Key Storage Administrator tasks and items to consider when developing an approach to Infrastructure Simplification.

Non-Database File system Management

One of the key challenges a Storage Administrator today has is managing the Non-Database File systems. These File systems can be located on any number of devices:

- ▶ Network-Attached Storage (NAS)
- ▶ SAN-Attached Storage (SAN)
- ▶ Locally Attached Storage

In addition, the number of File system options you have is increasing:

- ▶ UFS
- ▶ QFS
- ▶ VXFS
- ▶ JFS
- ▶ JFS2
- ▶ FAT
- ▶ NTFS
- ▶ EXT3
- ▶ EXT2
- ▶ Reiser

The above list is just a snapshot of file system combinations available today.

All of these options make the tasks of a Storage Administrator complex. Some of the tasks involved with management are (Table 1-1):

- ▶ Capacity monitoring
- ▶ Heterogeneous access

- ▶ Periodic migration
- ▶ Backup
- ▶ Space management
- ▶ Non-business use of space
- ▶ Performance
- ▶ Data location

Table 1-1 TotalStorage simplification tools

Tasks	TotalStorage simplification
Capacity monitoring	Productivity Center for Data
Heterogeneous access	SAN File System, SAN Volume Controller
Backup	Tivoli® Storage Manager
Space management	Tivoli Storage Manager, Productivity Center for Data
Non-business use of space	Productivity Center for Data
Performance	Productivity Center for Disk
Data location	Productivity Center for Disk

Database File system Management

The management of database file systems is much less complex, but often more critical because data, such as payroll, design, employee records, and customer transactions are key to day to day operations. The database file systems have the same management characteristics that are mentioned for file systems, but are more acute in the performance and capacity planning area.

1.3.1 Manage Disk Performance

For the Storage Administrator, Disk Performance is a daily task, if not sometimes a hourly one.

Configure and Optimize Disk Performance

Among the things driving performance configuration and optimization are:

- ▶ Application performance
- ▶ User response time
- ▶ File access response time

Among the factors that influence how the Storage Administrator addresses the configuration and optimization requirements are (See Table 1-2):

- ▶ Reviewing subsystem performance characteristics
- ▶ Analyzing device performance characteristics
- ▶ Validating the SAN configuration
- ▶ Determining the read to write ratio of applications

Table 1-2 TotalStorage Automated Management tools

Tasks	Automated Management
Subsystem Performance Characteristics	Productivity Center for Disk, Productivity Center for Data

Device Characteristics	Productivity Center for Data, Productivity Center for Disk
SAN configuration	Productivity Center for Fabric
Read to Write ratio	Productivity Center for Disk

1.3.2 Storage Replication

In a Storage Administrator's 12 hour day, replication requests are numerous if not constant. These requests are the result of:

- ▶ Data relocation for purposes of hardware upgrade
- ▶ Data replication for purposes of availability and disaster recovery
- ▶ Data duplication for purposes of testing
- ▶ Data replication for purposes of backup
- ▶ Data relocation for performance reasons

The requests can differ in the resources it takes to satisfy them (Table 1-3):

- ▶ Volume to volume copying
- ▶ Subsystem to subsystem replication
- ▶ LUN to LUN replication
- ▶ Virtualized volume migration
- ▶ Location to Location replication
- ▶ Virtual Storage Gateway to another Virtual Storage Gateway

Table 1-3 TotalStorage products address these resource requests

Resource Requirement	Product
Volume to volume copying	Flashcopy
Subsystem to subsystem replication	Flashcopy, Copy Services
LUN to LUN replication	Copy Services
Virtualized volume migration	SVC vdisks
Location replication	Copy Services
Virtual Storage Gateway to Virtual Storage Gateway	Copy services, SVC

1.3.3 Manage Storage Network Fabric

It may seem that the SAN administrator needs to be a separate job from the Storage Administrator, but it is the Storage Administrator who can make the most intelligent decisions on backend fabric management as it applies to storage devices. Zoning and access play key parts in the Storage environment. Key tasks are (See Table 1-4):

- ▶ Zoning - Making storage available to the right servers and protecting data from unauthorized access are key tasks associated with zoning. *Zoning* is the means of creating virtual storage environments within the fabric network.
- ▶ Topology Rendering - You can check with any LAN administrator and the key to managing and operating their network is the ability to display the total network or networks to verify connectivity and status. The topology map must not only give a diagram at a glance, but the ability to drill down on each endpoint for identification, type, HBA, version, and WWN to simplify fabric management.

- Fabric Alerts - Since the Storage Administrator cannot be constantly looking at the topology display for highlighted change of status, the Fabric Manager must provide an alerting facility or SNMP alerts that can be forwarded to a centralized alert and monitoring application, therefore, reducing the Storage Administrator workload.

Table 1-4 TotalStorage and Fabric Management tools

Task	TotalStorage solution
Zoning	Productivity Center for fabric
Topology Rendering	Productivity Center for fabric
Fabric Alerts	Productivity Center for Fabric, SAN Volume Controller, DS4000, DS6000, and DS8000

1.3.4 Provisioning and storage management task automation

Daily activities for the Storage Administrator include storage provisioning and storage management, such as special requests of departments, applications, users, and special projects, and so on. There are times though that the manual process of provisioning will not meet the business needs of these groups and a more automated and efficient process is needed. Automated provisioning or automation of storage management tasks are needed for those critical applications or file systems where the time of day or the time taken to provision or execute have to be around the clock and immediate. The automation has to take the manual tasks normally done below and automate them in a workflow fashion.

- Add a volume (Storage Subsystem).
- Select storage subsystem.
- Select or create a new volume.
- Select host HBA ports (WWNs).
- Select subsystem controller ports (WWNs).
- Map volume to controller ports.
- Map volume to host HBA ports.
- Set Paths (SAN fabric switches).
- Determine if multiple paths required.
- Create or update zones.
- Get active zone set.
- Add zone to zone set.
- Activate Zone set.
- Set up replication (if necessary).
- Map the HBA LUNs to the operating system and file system.
- Update volume group and file system (Host server).
- Add physical volume to volume group.
- Add physical volume to logical volume.
- Create or extend file system.
- Extend application to use additional space.
- Reconfigure backup.

Tivoli Provisioning Manager and the Tivoli Productivity Center Suite of products automate the provisioning storage administrator tasks listed above.

1.3.5 Protect Data and Optimize Data Storage

Since the Storage Administrator's duties include managing the placement of data, the performance of data access, the availability of data access including uptime and connectivity, often times the protection of data is just one more tasks that is added to his responsibilities. Data Protection is often an afterthought in it's assignment but in times of disaster it becomes the number one priority and most focused on task. The following are some Storage Administrator tasks associated with Data Protection.

- ▶ Government regulated retention of information
- ▶ Controlling access to backups
- ▶ Setting up procedures to recover for local and remote recovery
- ▶ Setting up data protection policies by user, department, type of data, OS, retention, versioning.
- ▶ Monitoring daily backups and reacting to growth rates by adding more resources
- ▶ Setting up backup schedules based on criticality of data
- ▶ Identifying Service Level Agreements for Restore times and structuring resources to meet these agreements.
- ▶ Insuring success of backups
- ▶ Identify inactive files and space manage them or archive and delete them automatically

The features of IBM Tivoli Storage manager and IBM Tivoli Productivity Center for Data will address all of the tasks identified above

1.4 Summary

Important: It is worth mentioning that one of the real benefits of Infrastructure Simplification is to help make the typical storage administrator's day vastly shorter than 12 hours.

This redbook is divided into three parts. Part 1, "Infrastructure Simplification principles" on page 1 covers the basic principles: What leads to high total cost of ownership (TCO) and first steps in Infrastructure Simplification. Part 2, "Virtualization of IT resources" on page 39 contains chapters that discuss customer examples and topics such as simple storage networks, simplifying tape backup, and storage virtualization. We also discuss SAN volume controller and SAN file system in this part. Part 3, "Automated storage management" on page 69 contains chapters about Infrastructure Simplification in action and describes a storage provisioning solution as well as software that can be used to simplify your Infrastructure. Part 4, "Advantages of using IBM TotalStorage for Infrastructure Simplification" on page 145 gives you competitive advantages of using IBM TotalStorage products for Infrastructure Simplification. Part 5, "Storage environment consolidation" on page 157 talks about how you can use IBM TotalStorage products to simplify your environment and shows you a customer example of consolidation.



Total Cost of Ownership of storage infrastructure

This chapter discusses the Total Cost of Ownership (TCO) of both storage and storage infrastructure with details regarding:

- ▶ What is TCO?
- ▶ What a complex infrastructure costs
- ▶ TCO and skilled staff
- ▶ Measuring total costs
- ▶ Costs of problems or faults

2.1 What is TCO?

Outside of the Information Technology industry, TCO is a simple accounting measurement used as part of any purchase or investment plan. A widget manufacturer needs a new widget machine. They look at various machines and compare:

- ▶ Initial investment
- ▶ Depreciation over five years
- ▶ Widget output
- ▶ Maintenance costs
- ▶ Cost per hour of the operator, including training

A recommendation about the widget machine with the best TCO can be made from this investigation.

Note: In information technology, the problem is that many of the costs are indirect or hidden. Interaction in a complex environment means that a product with low TCO can drive up the TCO for the total environment due to interoperability issues, networking, consolidated storage, support, and operational problems, absorbing a disproportionate or unplanned part of shared resources.

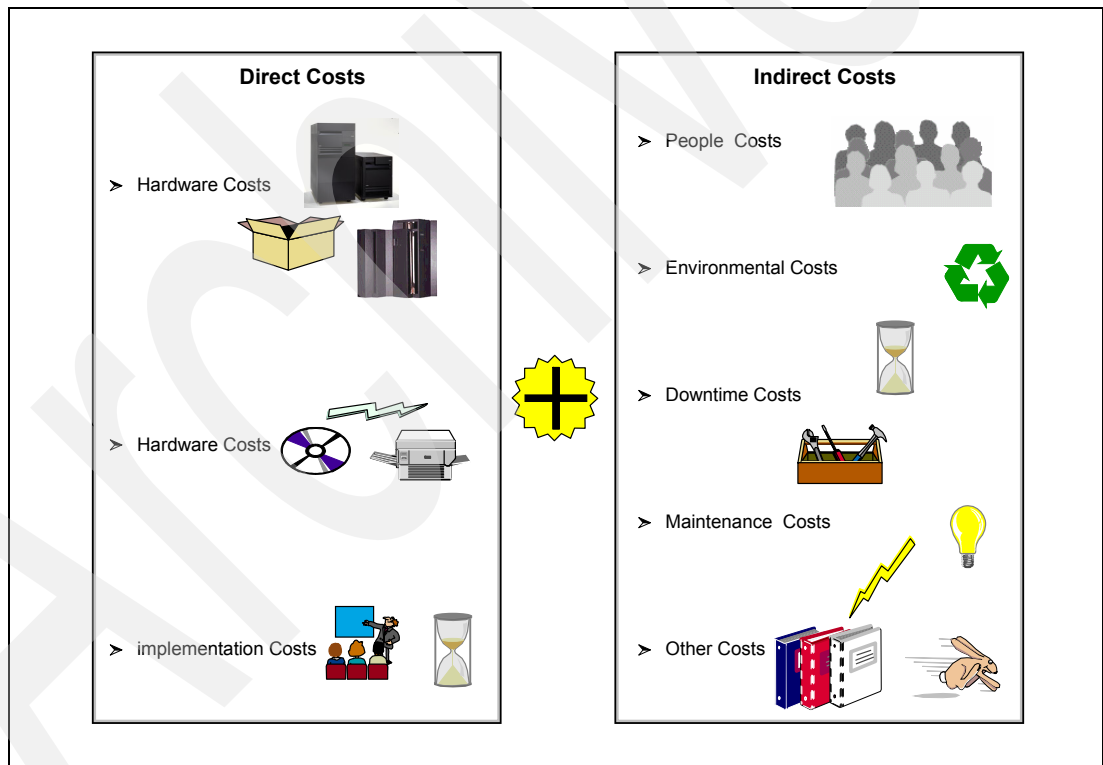


Figure 2-1 TCO indirect and direct

TCO, as part of IT purchasing, became widely used in the 1980s. The initial example of TCO was the cost of owning, deploying, managing, and maintaining workstations or personal computers (PCs). IT shops found in the 1980s that it costs nearly \$10,000 per year to own and maintain a PC, compared to the initial purchase price of about \$2,000. This shocked users and vendors of IT products and made TCO one of the key decision factors in purchasing new technology or products.

The important factors about TCO are:

- ▶ **Awareness of the real Total Cost of Ownership.**
 - Not just the product price, but what the solution costs over its life span, and the effect on the total cost of the entire infrastructure.
- ▶ **TCO is only one measurement out of several measurements to evaluate.**
 - A laptop computer has a high TCO when compared to a pencil and a sheet of paper. Punch cards had a very low TCO when compared to RAM memory.
- ▶ **TCO strips emotion and descriptions away and leaves facts.**
 - Rather than looking at:
 - Biggest
 - Fastest
 - Largest
 - Features
 - Name
 - Look at facts:
 - Initial price
 - Maintenance costs for the expected useful life
 - Training costs
 - Additional requirements, network, storage, and support
 - Outage costs
 - How will this affect my existing environment?
- ▶ **Documents why a decision was made.**
 - IT is a dynamic industry with high turnover. Decision makers may not be around in a year to justify why a product or solution was purchased. A documented TCO study can clarify the decision, and show that the solution met the business needs that prevailed at that time.

Buying a car as an example

Look at buying something that has some of the similar emotional and personal attachments and opinions that people bring to purchasing IT solutions, such as a car.

A company needs a vehicle for the IT department. Operators need to travel to the unstaffed data center to:

- ▶ Transport backup tapes to the disaster recovery site
- ▶ Assist with installations
- ▶ Accompany engineers working in the data center
- ▶ Upgrade software

The company already has a fleet of company cars, but needs one more car. Since this car is specifically for the IT department, an IT Manager has been asked to recommend a suitable car. The IT department will pick up most of the costs.

The operators need to carry approximately one hundred cartridges in a carrying case and drive twenty to thirty miles a day on a mixture of busy highways and quiet roads.

The IT Manager's experience is that Company Z-manufactured cars are reliable, so the IT Manager asks for quotes from three local dealers for mid-sized station wagon type cars that can safely transport 100 cartridges in the appropriate carrying cases.

The IT Manager does not look at maximum speed, acceleration, or luxury features, just the costs to run the vehicle for three years. In Figure 2-2, look at the direct costs.

Direct costs

	Car A	Car B	Car C
Price	21,150	22,730	21,505
Urban m.p.g.	26.4	37.7	35.8
Extra urban m.p.g.	42.8	58.9	55.4
Fuel tank (gallons)	13.2	14.3	13.6
C02(g/km)	193	153	164
Insurance group	9	12	8
<i>Table 2-1 Typical insurance quote</i>	372	440	358
% of value retained (3 years/36,000mls)	48	54	40
cost per mile in dollars	.54	50	55
Servicing costs 3 years	661	819	866
Miles between service intervals	20k	12.5k	12k
Fuel type	Unleaded	Diesel	Diesel
Running costs over 3 years average miles per year = 10k	16320	14910	16410

Figure 2-2 Direct costs of cars

The decision

After studying the direct cost, the recommendation is to buy car A. Although car B has low running costs, the experience is car A is a good vehicle. The 20,000 mile service interval is the deciding factor because availability is critical.

Indirect or hidden costs

Three years later, car A was a good choice and worked well for three years. The costs of owning this vehicle seemed to be lot higher than expected, but no one really knows why. The car performed as expected but there were a lot of issues. See Figure 2-3.

- The insurance company insists that anyone driving a car above insurance group 8 needs an advanced one day driving course. Any operator in the department could be required to drive this car so everyone had to be trained.

- \$200 per person for the course, twenty person days to arrange and attend training.
- ▶ Maintenance of pool cars is done by company mechanics. Since car A was the first car from this manufacturer they had seen, all the mechanics had to be trained. We needed spare parts sourced and stocked, and a support agreement with the local dealer was negotiated for any serious problems.
 - After this training, 50% of the mechanics left to work for the dealer that supplied car A. Replacements were difficult and expensive to recruit.
- ▶ All the vehicles owned by the company are fueled by diesel supplied by company pumps next to the maintenance facility. Car A runs on unleaded gas. This meant that the company had to install and maintain a special tank and pump.
 - No one knows what the tank and pump cost to install and maintain.
- ▶ The 20,000 mile service interval meant that every service was a major outage and took one or two days.
 - The IT department had to use and pay for a rental car while car A was unavailable.
- ▶ The IT car was perceived by other departments to be faster and better than their cars. It was borrowed for lots of special or critical jobs.
 - The mileage was much higher than expected and the IT department ended up negotiating for time on a shared resource.




Figure 2-3 Hidden costs of the IT departmental car

Conclusion of the TCO of the car example

Car A was a good choice, or not?

It performed as expected and performed well through its life span. Additional costs were incurred because the impact of this new car on the existing infrastructure was not understood.

Tip: Talking to the people who administer, maintain, and fuel the existing pool cars, combined with the purchaser's own experience and information from the vendors, would have given the purchaser a better understanding of the indirect costs incurred by this purchase, and a more accurate estimate of the total cost of owning this car.

2.2 What a complex infrastructure costs

If a purchase is simple, it is easy to measure cost-effectiveness and identify problems. If a purchase is complex, it is difficult to measure true cost because it is hard to break down costs and identify issues.

In Chapter 5, "Storage virtualization" on page 71, we discuss how complex SANs are created.

Limited budgets and too few skilled people have pushed IT organizations into looking for short-term solutions. When a new application or project appears, the easy and inexpensive option is to add another low cost server. While this "server sprawl" or proliferation of UNIX® and Windows® Intel® servers is an attractive short-term solution, the infrastructure costs to support and maintain these inexpensive servers often exceed the purchase price of the server.

Now add storage to the sprawl. Every server has two or four fiber Host Bus Adapters (HBAs) and a share of the consolidated storage. As we add more servers, we run out of SAN ports, so we add another switch, and then another, and then another. Now we have SAN sprawl with a complex interlinked fabric that is difficult to maintain and change.

To make things more difficult, we buy different storage systems, which look less costly at the time, each with unique management and monitoring tools and configuration characteristics. The decisions are made on purchase cost, suitability to a specific application, or simple personal preferences. These storage systems increase costs by requiring:

- ▶ Specialist training and skills required to operate, monitor, and manage them
- ▶ Specific software and licensing
- ▶ Multipathing software
- ▶ Additional interoperability issues that have to be checked and administered
- ▶ Unique advanced functions, licensed to a specific system
- ▶ An isolated logical or physical environment
- ▶ Outages to update or upgrade multiple systems

There are numerous possible tiers of storage. One example is:

Tier 1: Enterprise class storage for production

Tier 2: Midrange disk for clients that do not need high availability, development, or quality assurance

Tier 3: Advanced Technology Attachment for Archiving, backup, backups of the backup, low access files, or documents

Tier 4: Network Attached Storage (NAS)

Managing this structure requires a team with the training and skills to manage, monitor, and operate four storage systems, each one with its own tools and characteristics. Then administrators have issues with:

- ▶ Interoperability
- ▶ Migrating between tiers
- ▶ Different and/or disparate advanced functions

- ▶ Different licensing
- ▶ Maintenance/warranty
- ▶ Different administration
- ▶ Change control risk management

Tiering is a good way of breaking down costs and showing your clients the difference in direct costs between a terabyte of tier one storage and a terabyte of tier three storage.

However, consider the indirect costs of adding complexity to your environment. The tier three ATA (Advanced Technology Attachment) drives are inexpensive to purchase. A TCO study analyzing the total cost to the environment of adding this system may show that it is less expensive to upgrade the tier one or tier two systems to meet the new requirements than it is to add the complexity of tier three drives.

There are numerous ways to mask TCO costs, and IT purchasers need to be aware of this:

- ▶ Prepaid licenses: Be sure that prepaid software does not lock the IT department into a high TCO. Often, prepaid multipathing software can force higher costs overall for many years. IBM does not currently charge for its multipathing software.
- ▶ Future year maintenance: Future prices are not set; a low cost today can increase tomorrow.
- ▶ Power and cooling: These costs can be significant.
- ▶ Software licenses

2.3 TCO and skilled staff

The driving force in business today is to do more with less. Companies push to make their employees more productive, or move to a location where employees are just as productive, but cost less.

Just for simplicity, consider a storage specialist or administrator costs the company \$100,000. With the combination of salary, benefits, training, and support, it does not take long to get to \$100,000.

In a complex environment with heterogeneous servers, the operating system, SAN, and storage systems, a storage administrator is working very hard to effectively manage 20 TBs of data.

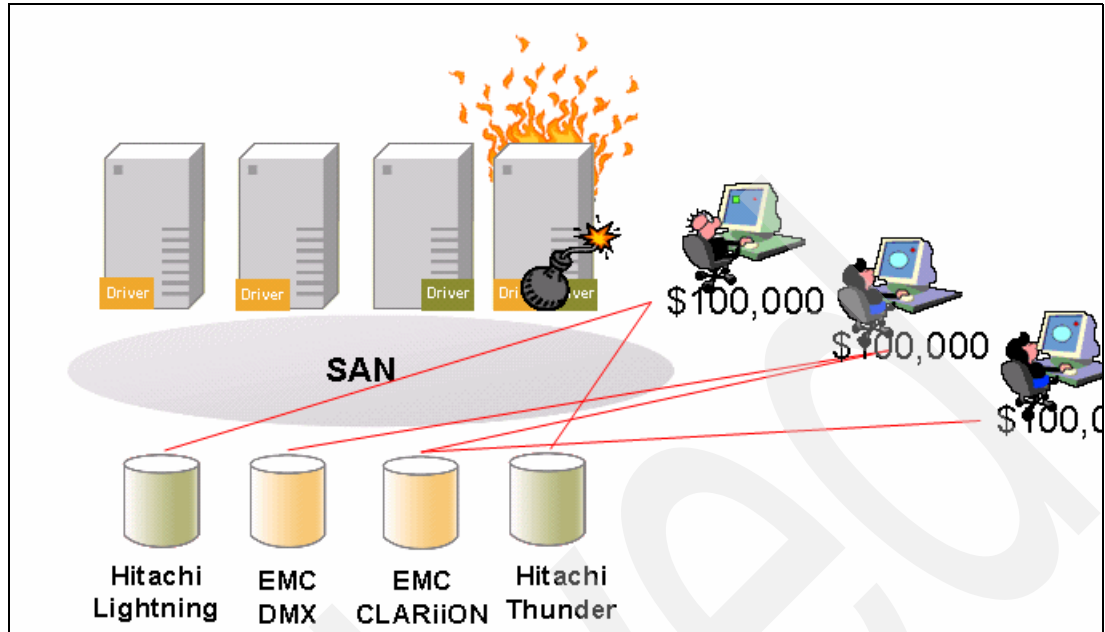


Figure 2-4 The cost of managing a complex environment

In a simple homogeneous environment where the servers and operating systems are the same, the SAN is a simple two of four switch fabric, the storage is homogeneous or has a standard virtual interface.

2.4 Storage administration costs

A storage administrator/specialist can manage 100 TBs or more. A simple calculation of \$100,000 over 100 TBs = \$1,000 per year per 1 TB.

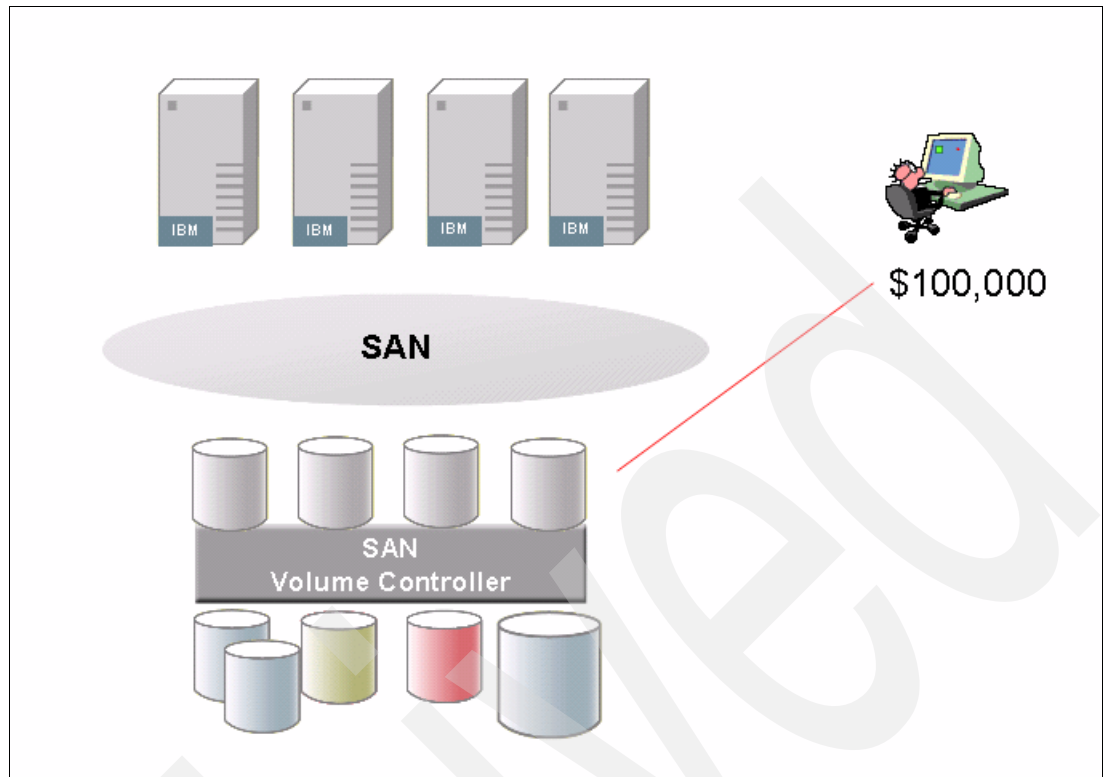


Figure 2-5 Simple environment

Every successful business is storing more data today than yesterday. Businesses today collect as much information as they can about their customers, partners, vendors, and competitors. As a result, IT managers are asked to manage increasing amounts of data on all types of systems, from e-mail servers to database systems. There is pressure to mine these mountains of data for nuggets of information. Requirements to protect, to secure, and, especially, to provide wide, but responsible access to information are increasing all the time.

To meet these demands, we need simple infrastructures and the tools to standardize and effectively manage and profit from this tide of information.

2.5 Costs and return on investment

The driving force to simplify your environment is to expose and eliminate hidden, indirect costs.

- ▶ Simplify
- ▶ Document
- ▶ Standardize

Until you know exactly how much a terabyte of storage costs to:

- ▶ Purchase
- ▶ Install
- ▶ Configure and implement into your existing environment
- ▶ Manage
- ▶ Utilize to the maximum extent
- ▶ Back up
- ▶ Restore effectively

You have no idea whether you are getting a good return on your investment (ROI) or not.

Simplify

If you simplify your SAN to two high port count, high availability switches, or two pairs of high port count high availability switches, you know:

- ▶ Purchase price
- ▶ Maintenance cost
- ▶ You can concurrently upgrade, switch code and hardware
- ▶ You can minimize the cost of down time
- ▶ A standard management interface for your SAN
- ▶ You can minimize training costs
- ▶ The cost per port and per blade of any internal upgrade
- ▶ The incremental cost to add another two switches if all ports are utilized

Document

Document your TCO study, especially the business constraints that affected the decision at the time, and the requirements of the solution or project. This documentation:

- ▶ Clarifies your decision for later review
- ▶ Gives an unbiased basis for making your decision
- ▶ Can be used as a discussion tool with vendors to explain why a certain product was chosen
- ▶ Includes aspects of the complete environment that make up the TCO
- ▶ Contributes as a deciding factor in future upgrades and new projects
- ▶ Can be a differentiator between the original plan and final implementation

Standardize

Simple, standard solutions are good, easier to cost, easier to troubleshoot, easier to fix, easier to grow, and easier to manage.

Your CIO, CFO, or other C-level executives with cross-functional responsibilities (hardware, software, and networking) should be able to understand where the information your company depends on is:

- ▶ How access is controlled.
- ▶ How resilience has been added.
- ▶ How the data is protected.
- ▶ How the data is backed up, and, more importantly, how it is restored in a timely manner.
- ▶ Why you have standardized on a single type of storage, or tiered different storage systems.
- ▶ The tools and skills needed to manage this critical resource.

If you cannot explain in broad terms to a non-technical person how you manage your storage, then it is too complex. Everyone from the CEO to accounting to the guys who install the power and lay the cables should be able to understand why you have at least two SAN switches in a redundant configuration, how storage is allocated, and how it is backed up and restored.

If you do not understand something, how can you measure its value?

2.6 Measuring total costs

When TCO analysis was developed, it was a recommendation about how to calculate costs. This methodology is used in a number of different areas, not just in IT.

Today, a number of companies have developed their “own” analysis of TCO, based on the original recommendations and their own practices. To create a complete calculation, the purchaser must collect and categorize all the areas where the money can go. Without hands on experience, the purchaser can only calculate a rough estimate of the total costs. Every vendor shows flashy charts about how the purchase prices are going down in IT, but user costs are still increasing. What is driving this increasing expense?

When calculating the true costs of ownership, the purchaser's hardest challenge is to collect all the information necessary. Some of these costs are obvious costs. These are called *direct costs*. Direct or obvious costs are:

- ▶ Initial hardware, software, and license purchase costs
- ▶ Deployment
- ▶ Planned obsolescence
- ▶ Maintenance contracts, including spare parts, support, and consumable supplies
- ▶ Operations, including labor costs for operators, administrators, and internal support
- ▶ Power, cooling, and floor space
- ▶ Administration, procurement, financial administration, and training

Over half of the costs are still invisible, because you tend to forget items that are not directly related to actual purchased items, but the purchaser deals with them sooner or later. Some costs are not easily quantified in money, but in this case, you must estimate or calculate the cost of missed opportunities.

These are called *indirect costs*. Indirect, less, or not so obvious costs are:

- ▶ System down time in terms of lost opportunities, sales, and productivity
- ▶ Being “locked” into a technology, and therefore unable to take advantage of new technologies
- ▶ Maintaining nonstandard configurations
- ▶ User training and Help desk support
- ▶ Self-supporting times, changing parameters, and resolving difficulties
- ▶ Impact on existing environment, interoperability, and complexity
- ▶ Down time interruptions to user operations during deployment, failures, and upgrades
- ▶ Custom application development

The concept behind TCO is that you cannot control an environment that you do not know.

Note: If you do not know where the money goes, you cannot decrease the costs without disrupting the quality of services.

Comparing solutions provided by different vendors is difficult. Each vendor uses different calculation methodologies to measure TCO. There are too many dynamic elements to expect a completely accurate financial forecast to predict the costs. The result of TCO is not a forecast, but just a factor which helps eliminate subjective elements and emotions.

TCO analysis

one-time costs outweighed by recurring savings

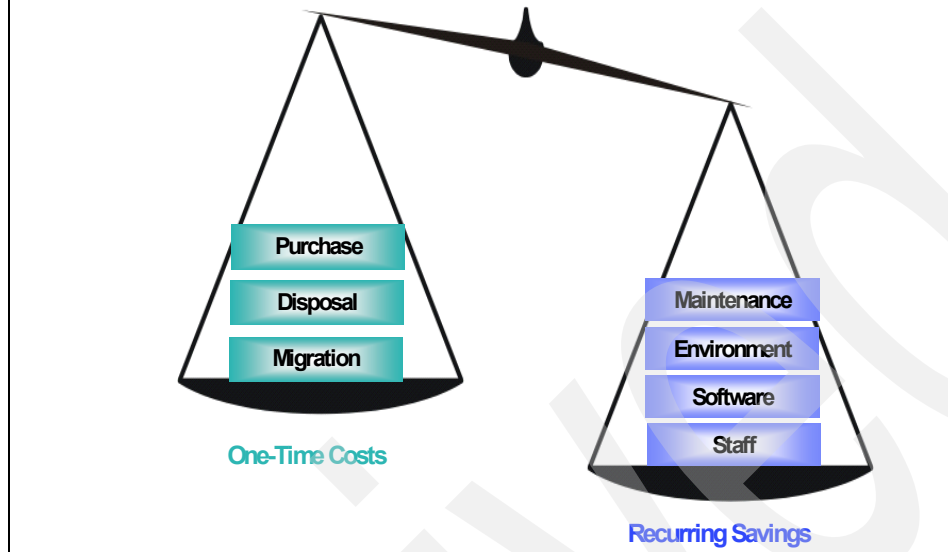


Figure 2-6 TCO analysis

2.6.1 Steps to calculate TCO

There are as many different TCO calculations you can make as there are questions that you can ask. The basic steps to calculate TCO are:

1. Decide the primary goal of the calculation and express the clear expectations that you have against the final result.
2. Determine the time period to consider for the calculations.
3. Collect all related financial data.
4. Generate reports to calculate the actual TCO using these numbers.
5. Compare results to industry averages.
6. Generate more reports to highlight potential problems and suggest action items.
7. Do a "What if?" analysis, plan acquisitions, and alter existing plans, if necessary.

By collecting all financial data, the analysts ask about questions, such as:

- ▶ What is your actual cost per minute of storage-related down time?
- ▶ What is your estimated lost productivity and lost revenue cost per minute of storage-related down time?
- ▶ What is your productivity factor?

Sometimes you cannot answer these questions easily. Since down time can be a major hidden cost in any storage environment, it is important to calculate the number of planned and unplanned outages per year. In the IT world, since one employee often has multiple job responsibilities, the TCO calculations introduced a term, *full-time equivalents* (FTEs), to calculate human costs.

2.7 Costs of problems or faults

The down time measurement has many faces.

Go back to the car analogy. If a driver wants to buy a gas for a car, the driver goes to a gas station. The gas station attendant tells the driver that the gas pump is down, to please wait or come back later. Depending on the driver's actual situation, this problem can have many outcomes.

- ▶ The driver just goes to another gas station. In this case, the first gas station has lost one customer.
 - The direct loss to the gas station is approximately \$30 of gas.
 - Indirect loss is the customer is inconvenienced for a few minutes.
- ▶ The driver does not have enough gas to get to the next gas station; therefore, the driver has to wait one and a half hours for the gas pump system to come back online.
 - Here, there is no direct loss to the gas station, the driver still has to buy \$30 of gas.
 - The driver is a business person who has missed an important meeting. Now the driver is upset, angry, and embarrassed. The driver promises to never use the brand of gas again, and tells everyone how bad the gas stations are. The indirect loss is possibly thousands of dollars to the gas stations.
 - The customer is also CEO of a large distribution company with hundreds of trucks that all refuel at those gas stations using company fuel cards.
 - There is still no direct cost to the gas station.
 - The customer eventually gets back to his office, cancels the contract with the gas station company and moves to a competitor. The indirect costs are millions of dollars.

In most TCO calculations, the cost of down time is either ignored or treated as a simple calculation. We make \$3,65ML per year. If the system is down for a day, it costs us \$10,000.

If the day that you lose is your financial year-end processing or the day that you launch a new product, the cost can be millions of dollars lost for the whole company.

If you want shocking figures, only six percent of companies that suffer a catastrophic data loss survive. See Figure 2-7.

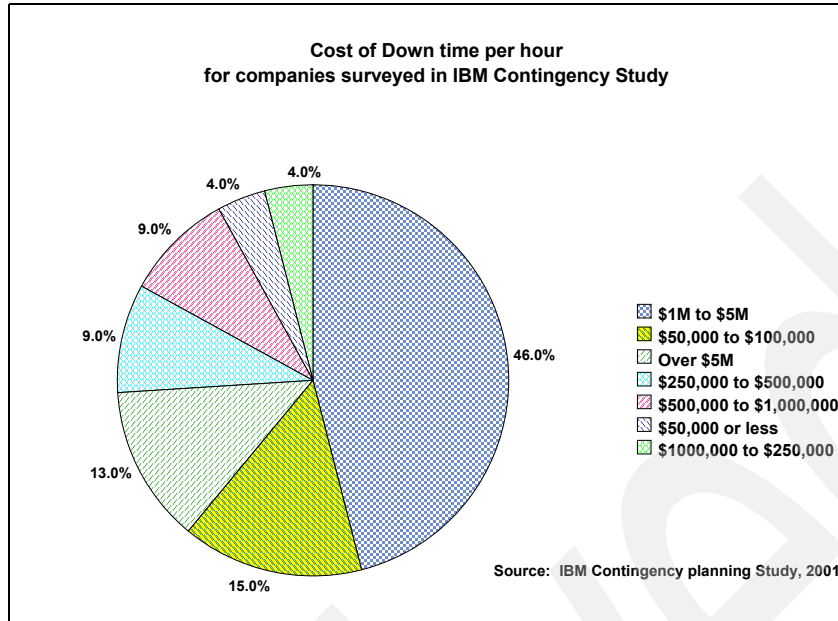


Figure 2-7 Costs of down time

When you calculate the cost of down time, do not just take the easy average (Revenue/days or hours). Spend time looking at critical scenarios: Year-end, month-end, or a critical period in your environment, and document the average cost and the cost of a failure at a critical time in your business cycle.

The lesson to take from TCO is document everything. Look at the purchase price last. You or your procurement people can put pressure on the sales people any time. Look at what it is going to cost to:

- ▶ Implement
- ▶ Integrate
- ▶ Manage
- ▶ Back up and restore
- ▶ And what it will cost if it fails.
 - That unique, temporary, nonstandard, undocumented, and inexpensive solution will fail at a critical point.
 - Year-end
 - Quarter-end
 - Month-end
 - The day the CEO wants to see it.
 - The day you think you are going on vacation.

2.8 TCO with IBM TotalStorage

In the following sections, we discuss the TCO advantages of IBM TotalStorage.

2.8.1 Infrastructure simplification

Infrastructure simplification is all about reducing complexity and costs across your IT environment. IBM can help you with consolidation, management centralization, and capabilities for unifying heterogeneous storage islands, resulting in new levels of flexibility and choice, lower TCO, and improved ROI by:

- ▶ Consolidating dispersed storage resources
- ▶ Providing a unified, strategic view of your data
- ▶ Breaking through traditional storage complexity with advanced management capabilities
- ▶ Innovating to unify and simplify heterogeneous storage environments

There are advantages in overall long term cost of ownership which can be realized by consolidating data stored on older generation storage arrays such as IBM TotalStorage Enterprise Storage Server® (ESS) Model F20 or competitive other equipment manufacturers' (OEM) systems, and replacing these older arrays with new IBM TotalStorage DS6000 or DS8000.

We designed the following examples to show you the current operational costs of these older arrays, so you can compare these costs over a four year horizon to costs associated with the new IBM TotalStorage DS6000 and DS8000 over the same period.

Note, we only quantify operating expenses. In our study, we did not quantify additional potential savings in the following areas, but you can also realize:

- ▶ Personnel time through management efficiencies offered by managing one storage array, instead of two, four, or eight
- ▶ Personnel time through improved efficiencies available due to enhancements in management capabilities (DS Storage Manager) and the Command-Line Interface
- ▶ Savings in floor space
- ▶ Improved point-in-time copy function, performance, and flexibility
- ▶ Improved mirroring function, performance, and flexibility
- ▶ Improved flexibility through enhancements and increases in addressing, a larger number of subsystems (LSS), larger volumes, and larger LUN sizes
- ▶ Cost savings potential offered by more flexible advanced function licensing
- ▶ Increased management flexibility and efficiency offered by DS8000 storage system LPARs
- ▶ Superior scalability

2.8.2 About the TCO examples

The following examples compare the maintenance and environmental costs of the installed older generation storage. These examples assume the assets being replaced have been fully depreciated (or are past lease expiration), so that there are no further costs associated with acquisition. Note that ESS Model F20 has hardware maintenance charges only, which include advanced function support.

Based on IDC data, there are costs associated with hardware and advanced function software for the OEM storage arrays. The ongoing operational costs for the older systems are compared to the costs associated with acquisition and operation of the cited DS6000 or DS8000 hardware, advanced functions, operational costs, as well as estimated migration costs over the same four year period. For the purpose of this analysis, we assume the time value of money is 6%.

Note: A number of considerations beyond the scope of this redbook can influence actual cost savings and cost avoidance. You can have cost considerations that we do not have in these examples.

We created the following TCO examples using previous IBM storage arrays and competitors' similar arrays and models that we refer to as OEM.

2.8.3 The cost advantages of DS6000

- ▶ Up to 50% reduction in price
- ▶ Support for both Open Systems and mainframe systems
- ▶ Higher density storage
- ▶ Reduced installation and service complexity

2.8.4 DS8000: A new standard for lowering TCO

- ▶ Four year warranty
- ▶ Linear scalability
- ▶ Dual-clustered server POWER5™ technology
- ▶ Increased capacity (3 times the ESS Model 800)
- ▶ Increased performance: Up to 6 times the ESS Model 800
- ▶ Flexibility: Dramatic addressing enhancements
- ▶ More flexible feature licensing
- ▶ Larger capacity volumes supported
- ▶ Increased opportunities for consolidation

2.8.5 Cost advantages of the SAN Volume Controller

With a virtualized environment, you can pool all of the available capacity behind the SAN Volume Controller. To the host systems, this appears as single, very large disk controller. Your administrators can more easily reassign unused capacity simply by remapping physical capacity to virtual disks that need more space. As a result, the overall utilization of your physical disks can be increased, and future purchases deferred. Think about the savings involved. Last year, we conducted research with 200 customers. On average, they thought this technology would improve their overall utilization by approximately 20%. Assume you have 10 TBs of storage for which you paid an average of \$0.05 per MB. A 20% improvement in utilization represents a savings of \$100,000.

2.8.6 DS6000 operational costs in detail

Figure 2-8 on page 33 represents a combined example of comparing the operational costs of two ESS F-20s and similar class OEM subsystems compared to acquisition and operational costs for DS6000 over a four year period. As noted, OEM customers with two OEM arrays spend over 50% more in operational costs than the acquisition and operational cost for DS6000 over four years. This example includes the “upgrade” warranty cost to provide for same day IBM maintenance during the four year warranty period. Note the significance of how much of the acquisition and operational costs over four years for a new DS6000 is covered by the current operational cost expenditures for the older IBM ESS or OEM storage arrays. This offers you an excellent opportunity to begin discussions with ESS F-20 and OEM customers about potential operational cost benefits as well as additional potential benefits and savings based on their specific environment, application development plans, workload projections, and so on. Highlights of these comparisons are in Figure 2-8.

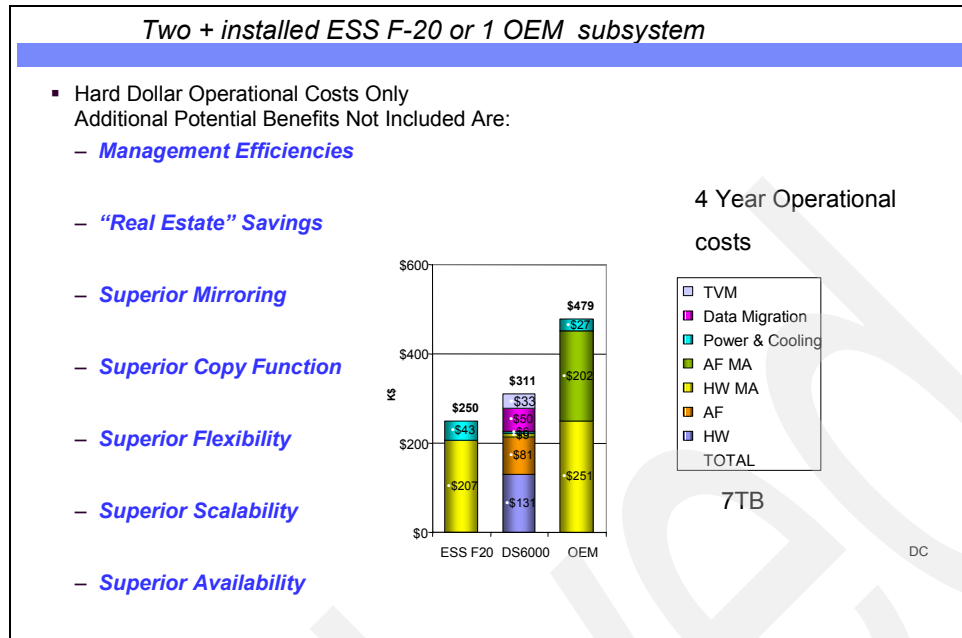


Figure 2-8 DS6000 storage consolidation

2.8.7 Operational costs in detail

Figure 2-9 represents a combined example of the operational costs for two ESS F-20s or OEM arrays versus the acquisition and operational costs for DS8100 over a four year period. As noted, OEM customers with two OEM similar class arrays could potentially spend roughly the same in operational costs as the acquisition and operational costs over four years for DS8100. Note the significance of how much of the acquisition and operational costs over four years for a new DS8100 is covered compared to the current operational cost expenditures for the older IBM ESS or OEM storage arrays. This offers you an excellent opportunity to begin discussions with customers about additional, potential benefits and savings based on their specific environment, application development plans, workload projections, and so on.

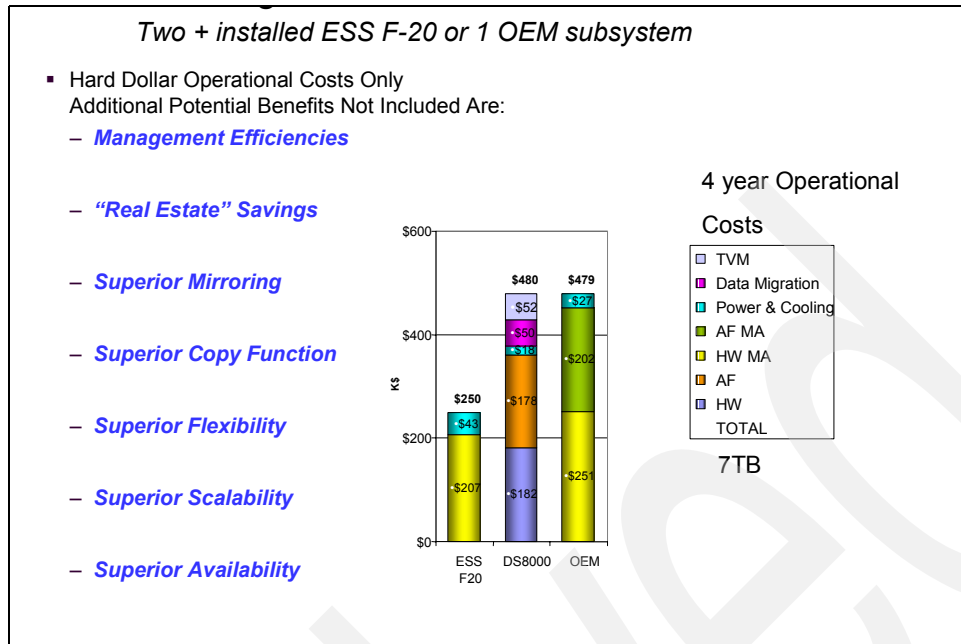


Figure 2-9 DS8100 storage consolidation

Figure 2-10 is a combined example of the projected operational costs of four ESS F-20s or similar class OEM models compared to projected acquisition and operational costs for DS8100 over a four year period. As noted, OEM array customers could spend more in operational costs than the acquisition and operational cost over four years for DS8100. Note the significance of how much of the acquisition and operational cost over four years for a new DS8100 is covered compared to the current operational cost expenditures for the older IBM ESS or OEM arrays.

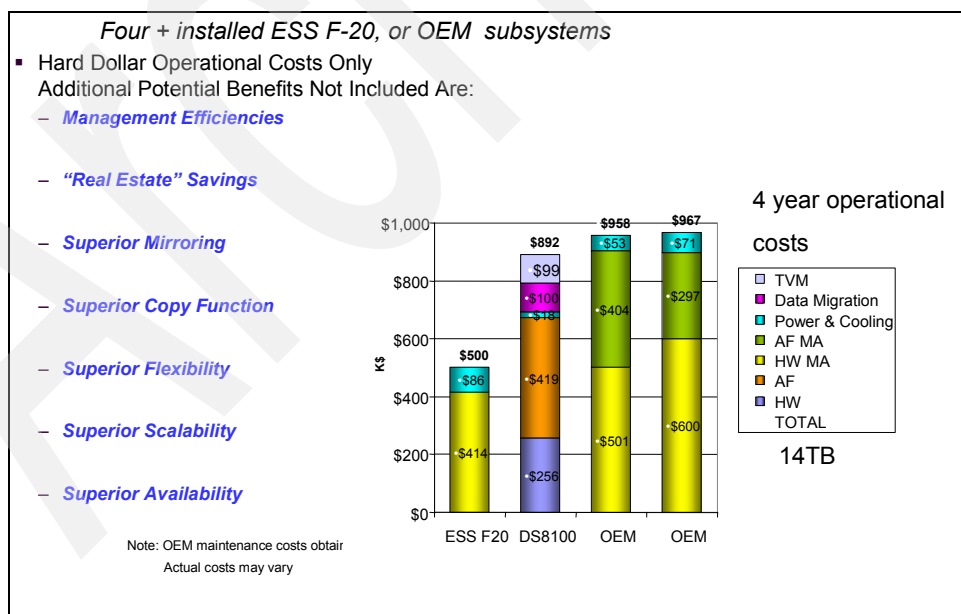


Figure 2-10 DS8100 storage consolidation compared with ESS F-20 and OEM

Figure 2-11 on page 35 is a combined example of the projected operational costs for ESS F-20 or similar class OEM over a four year period. As noted, OEM customer arrays could

expect to spend significantly more in operational costs than the acquisition and operational cost over four years for DS8300. Note the significance of how much of the acquisition and operational costs over four years for a new DS8300 is covered compared to the current operational cost expenditures for the older IBM ESS or OEM storage arrays. This offers you an excellent opportunity to begin discussions with customers about additional, potential benefits and savings based on their specific environment, application development plans, workload projections, and so on.

The potential benefits offered by storage system LPARs can also provide advantages in this environment and should be explored with each of these customers.

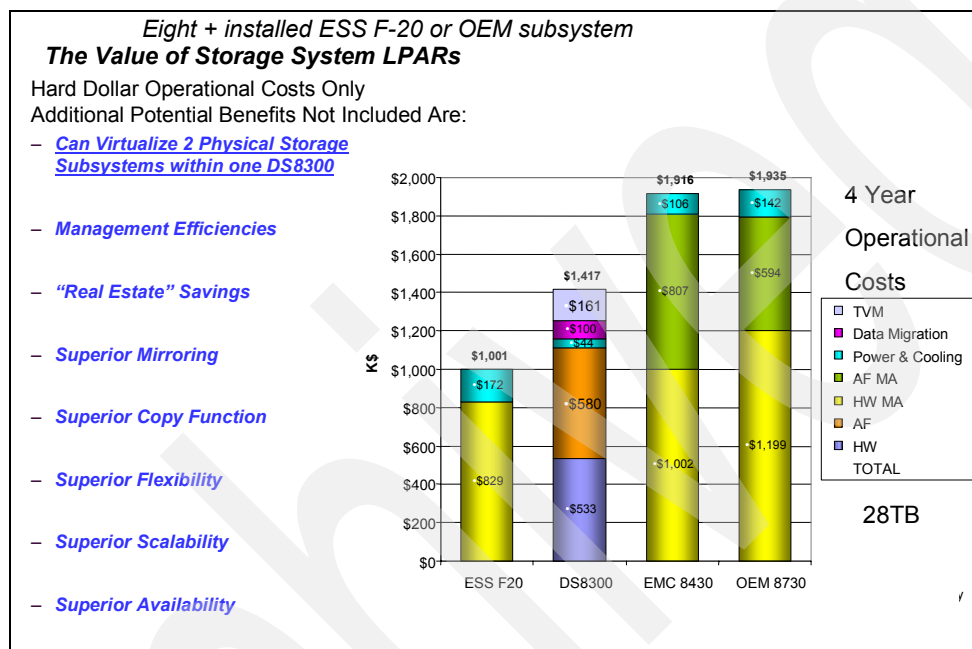


Figure 2-11 DS8300 storage consolidation

2.8.8 DS8000 additional potential consolidation benefits

- Management efficiencies
 - Fewer management consoles
 - ESS Hardware Management Console (HMC) supports up to eight ESSs, but most customers use one or two
 - Easier server to disk assignment
 - Easier server/disk allocation
 - Improved storage utilization
 - Potential to defer capacity upgrades
 - Reduced tuning
 - Tuning only one array instead of many
 - Larger volumes
 - More efficient management
 - One copy environment
 - Improved copy management

- Fewer data paths to manage
 - Improved efficiency
- Reduced “other” per system management actions
 - Improved efficiency
 - Other “per subsystem” management tasks that take up storage administrator time
- Reduced tuning/improved performance
- ▶ Infrastructure
 - Fewer SAN ports
 - Frees up ports, potentially delaying future purchases of ports and/or directors
 - Fewer channels
 - 2 Gb FC/FICON on the DS6000/DS8000
 - 2 Gb is not available on ESS F Models or older OEM
- ▶ Data center “real estate”
 - “Mileage will vary” depending on current customer data center space utilization
 - Potential significant hard dollar justification based on greatly reduced floor space requirements
 - Utilization of existing available 19” rack space with DS6000
 - Reduced table space needed for management consoles
- ▶ Capacity utilization
 - Improved point-in-time copy target utilization by having one target pool on one subsystem instead of multiple target pools on multiple subsystems
 - Reduced number of disk spares

2.8.9 Analysis background information

- ▶ 6% assumed interest rate
- ▶ Standard field discounts (no special bid pricing)
- ▶ 10K rpm/146 GB drives used in all DS6000/DS8000 configurations
- ▶ \$5K -- \$7K per TB migration costs
- ▶ DS6000 “IBM Maintained-24/7” Maintenance Agreement upgrade included in DS6000 example

2.8.10 How we calculate environmental costs

(14 TB DS8100 example)

- ▶ Power
 - $\text{KVA} \times \text{cost/kilowatt hr.} \times \text{hr./day} \times 365$
 - $5.5 \times .075^* \times 24 \times 365$
 - \$3,614/yr. power cost

- ▶ Cooling (12,000 BTUs per ton of air conditioning)
 - DS8100 requires 1.5 tons of air conditioning for cooling (18,772 BTU)
 - A/C x cost/kilowatt hr. x hr./day x 365
 - $1.5 \times .075^* \times 24 \times 365$
 - \$987/yr. cooling cost
- ▶ \$4,600 annual DS8100 environmental cost

Archived



Part 2

Virtualization of IT resources

In this part, we discuss consolidation and virtualization: key methods to improve management of IT environments. In addition we give two customer examples of using IBM TotalStorage and System Storage™ products to simplify your Infrastructure.

Archived



Simplifying your fiber Storage Network

In this chapter, we look at why we need to simplify our storage infrastructure. The following reasons drive the need to simplify storage networks:

- ▶ Massive vertical, capacity, and horizontal granular growth.
 - Huge growth with no additional head count. We need to do more with less.
- ▶ Heterogeneity is a fact. Consistency is a requirement.
 - Interoperability in Storage Area Networks is still a problem.
 - No single solution from any vendor meets all requirements.
 - Lots of time and money are spent managing infrastructure.
- ▶ The desire for quality of service options. Matching resources by their importance.
 - Maximize ROI by matching resources to their importance to the core business.
 - To optimize across different workloads or isolate specific workloads or applications.
- ▶ Difficult and complex disaster recovery processes.
 - Problems with managing data.
 - Getting data to a business continuity location in a reliable cost-effective solution.
 - The need to support multiple complex replication solutions.

We look at actions to achieve this simplification.

- ▶ Options to unify and simplify heterogeneous storage environments:
 - We briefly look at how SANs became complex.
 - Interoperability problems, especially in heterogeneous SANs.
 - We then cover options how to simplify your disk subsystems.
 - Simplify your storage network and your backup and tape infrastructure.

Note: IBM has been building fiber-based storage networks for twenty years. We use the mature zSeries® storage network as an example for comparison in this section.

3.1 From where does the complexity come?

A Storage Area Network (SAN) is a simple thing, a path from a server to a common storage resource. So, from where does the complexity come?

Limited budgets and too few skilled people have pushed IT organizations into looking for short term solutions. When a new application or project appears, the easy, inexpensive option is to add another low cost server. While this “server sprawl” or proliferation of UNIX and Windows Intel servers is an attractive short term solution, the infrastructure costs to support these inexpensive servers often exceeds the purchase price of the server.

Now add storage to the sprawl. Every server has two or four fiber Host Bus Adapters (HBAs) and a share of the consolidated storage. As we add more servers, we run out of SAN ports, so we add another switch, and then another, and then another. Now we have SAN sprawl with a complex interlinked fabric that is difficult to maintain or change. See Figure 3-1: SAN sprawl.



Figure 3-1 SAN sprawl

To make things more difficult, the servers were probably purchased from multiple vendors, with decisions made on cost, suitability to a specific application, or merely someone's personal preference. Different vendor's servers were tested on very specific SAN configurations. Every server vendor has their own interoperability matrix or list of SAN configurations that the vendor has tested, and that particular vendor supports. These lists have the:

- ▶ Operating system level
- ▶ Specific HBAs and the drivers for these HBAs
- ▶ Multipathing software
- ▶ Types and models of SAN switches
- ▶ Code level for these switches
- ▶ Storage subsystems and the microcode level that was tested

Because the SAN and storage are common to all servers, therefore, the interoperability matrix for every server then has to be checked out and compared before any change is made to the storage network, or even to individual servers.

Important: A reasonably complex SAN could have servers from Sun™, IBM, HP, and Dell and storage subsystems from the same or other vendors. There is no guarantee that SUN, Dell, and IBM support the same level of SAN or disk microcode. Changes, upgrades, or new implementations can be delayed for months, while you negotiate with multiple vendors to support a common code stream or driver.

3.2 Interoperability

Note one: The rule for interoperability in SANs is that it is the manufacturer of the storage device that defines interoperability. They do the testing that produces a matrix of supported configurations, from the storage device to the server.

Note two: There is an additional complication in backup solutions. Where the vendor of the backup software also produces a matrix of tested, supported hardware.

If you had Sun servers and you are planning to connect them to an IBM ESS disk storage system you would go to the IBM ESS support matrix and confirm.

- ▶ What SAN Switches and micro code are supported
- ▶ What fibre channel HBAs have been tested and supported, and the driver level
- ▶ The server Models and Types
- ▶ The Level of the Operating System
- ▶ What multipathing software is recommended and the tested level
- ▶ Any advanced functions that may be supported
 - Flash or point-in-time copy
 - Remote mirroring
 - Boot from SAN
 - Server clustering

Servers	Operating Systems	Host Adapters	Fabric Support
Enterprise <ul style="list-style-type: none"> • 220R, 250, 420R, 450 • 3000, 4000, 5000, 6000 • 3500, 4500, 5500, 6500 • 10000 	Solaris <ul style="list-style-type: none"> • 2.6 <ul style="list-style-type: none"> • VERITAS Volume Manager with DMP • 3.1 and 3.2 • VERITAS Cluster Server • 1.1.2, 1.3, and 2.0³ • 7 <ul style="list-style-type: none"> • VERITAS Volume Manager with DMP • 3.1 and 3.2 • VERITAS Cluster Server • 1.1.2, 1.3 and 2.0³ • 8 <ul style="list-style-type: none"> • Sun Cluster 3.0¹⁰ • MPxIO 4.1, 4.2 • Sun Cluster 3.1¹³ • MPxIO 4.2, 4.3 • VERITAS Volume Manager with DMP • 3.1.1, 3.2, 3.5, and 4.0¹⁴ • VERITAS Cluster Server • 1.3 and 2.0 • 3.5 and 4.0¹⁴ • PPRC/VCS Agent¹⁴ • 9^{3,7} <ul style="list-style-type: none"> • Sun Cluster 3.0¹⁰ • MPxIO 4.1, 4.2 • Sun Cluster 3.1¹³ • MPxIO 4.2, 4.3 • VERITAS Volume Manager with DMP • 3.5 and 4.0¹⁴ • VERITAS Cluster Server • 3.5 and 4.0¹⁴ • PPRC/VCS Agent¹⁴ 	AMCC / JNI ¹² <ul style="list-style-type: none"> • FC64-1063-N • FCC-6562 / FCC2-6562³ • FCX-6562 / FCX2-6562³ • FCI-1063-N • FCE-1473-N¹ • FCE-6460-N¹ Emulex <ul style="list-style-type: none"> • LP8000 • LP8000S⁴ • LP9002C³ • LP9002L / LP9002DC¹ • LP9002S¹ • LP9402DC¹ • LP9802¹ QLogic <ul style="list-style-type: none"> • QLA2200F • QLA2310F/QLA2310FL¹ • QLA2340/L / QLA2342/L¹ • QCP2340⁹ Sun ¹⁰ <ul style="list-style-type: none"> • 6727A • 6757A • 6767A (SG-XPCI1FC-QF2) • 6768A (SG-XPCI2FC-QF2) • 6799A 	Cisco <ul style="list-style-type: none"> • MDS 9120 and 9140¹¹ • MDS 9216,A,I¹ • MDS 9506⁹ and 9509¹ CNT (Inrange) <ul style="list-style-type: none"> • FC/9000-64⁶, -128⁶, -256⁸ • 2042-N16¹⁴ IBM <ul style="list-style-type: none"> • 2005 Model B32¹⁴ • 2005 Models H08 & H16¹⁴ • 2026 Model E12, 224¹⁴ • 2027 Models 140 & 232¹⁴ • 2027 Models R16 and 256 • 2045 Model N16¹⁴ • 2103 Model H07⁵ • 2109 Models F16 and F32¹ • 2109 Model M12¹, M14¹⁴ • 2109 Models S08 and S16 • 3534 Model F08¹ McDATA <ul style="list-style-type: none"> • ED 5000, • ES-3016 and ES-3032 • Intrepid <ul style="list-style-type: none"> • 6064 and 6140¹ • Sphereon <ul style="list-style-type: none"> • 3216 and 3232¹ • 4300¹¹ and 4500¹
Netra <ul style="list-style-type: none"> • 1125 			
SPARCcenter / SPARCserver <ul style="list-style-type: none"> • 1000, 1000E • 2000, 2000E 			
Sun Fire <ul style="list-style-type: none"> • V120¹ • V210¹¹ • V240¹¹ • V250¹¹ • V280R¹ • V440¹¹ • V480, V880¹ • V490, V890¹⁴ • V1280¹¹ • 3800¹ • 4800¹, 4810¹, 6800¹ • 120000², 15000² 			
Ultra <ul style="list-style-type: none"> • 1, 2, 5, 10, 20, 30, 60, 80 			

Figure 3-2 Sample Interoperability matrix for an IBM DS8000

The key to managing interoperability in a SAN is to check and question everything. If you have equipment from three storage vendors, you must look, check, and compare the details of all three interoperability matrixes.

Important: Never make assumptions if you read anything that is unclear. Ask the vendor for clarification and get their answers documented.

Figure 3-3: Interoperability: Everything has to fit shows what pieces must fit for interoperability.

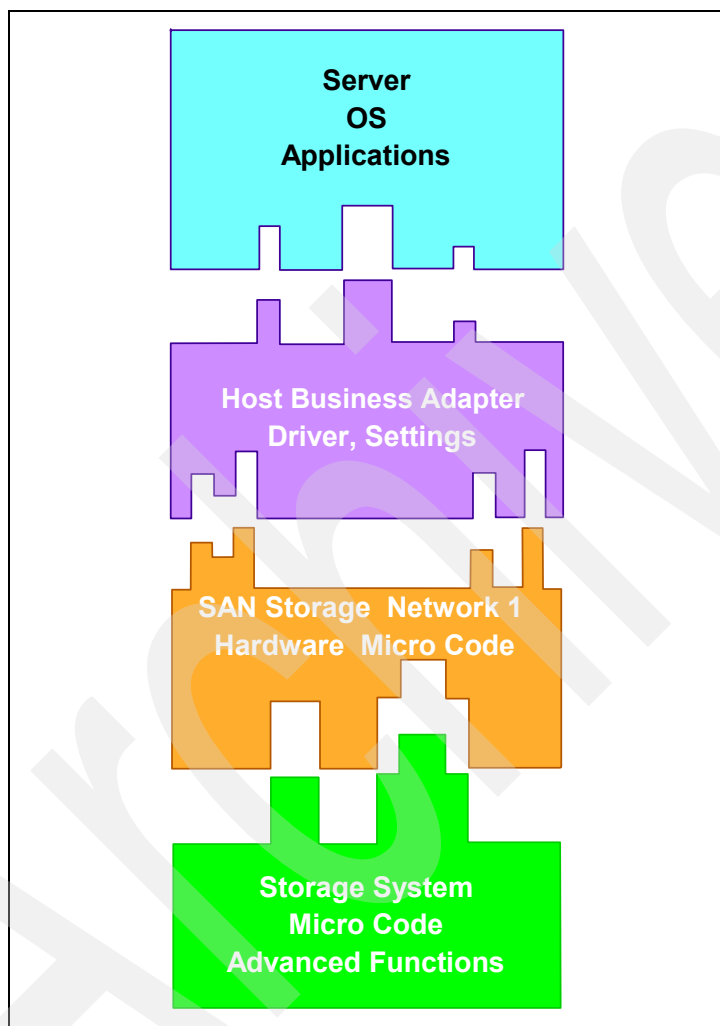


Figure 3-3 Interoperability: Everything has to fit

3.3 Simple Storage Networks

IBM has been building fiber storage networks for twenty years. The oldest fiber storage networks and the most stable are those used in the zSeries or mainframe environment. This redbook is not intended to endorse Open Systems or zSeries, but it is intended to look at the zSeries storage network as a mature, highly available, efficient example of a SAN and for you pick up some lessons. Why invent the wheel when there is already one there to look at and improve?

3.3.1 zSeries storage network development

zSeries went through the same development stages as open systems.

- ▶ From direct attached decentralized storage
- ▶ To centralized storage islands
- ▶ To common infrastructure
- ▶ to a form of virtualized resources with automated storage management

Figure 3-4: Development of zSeries storage network shows how disk attachment evolved in zSeries.

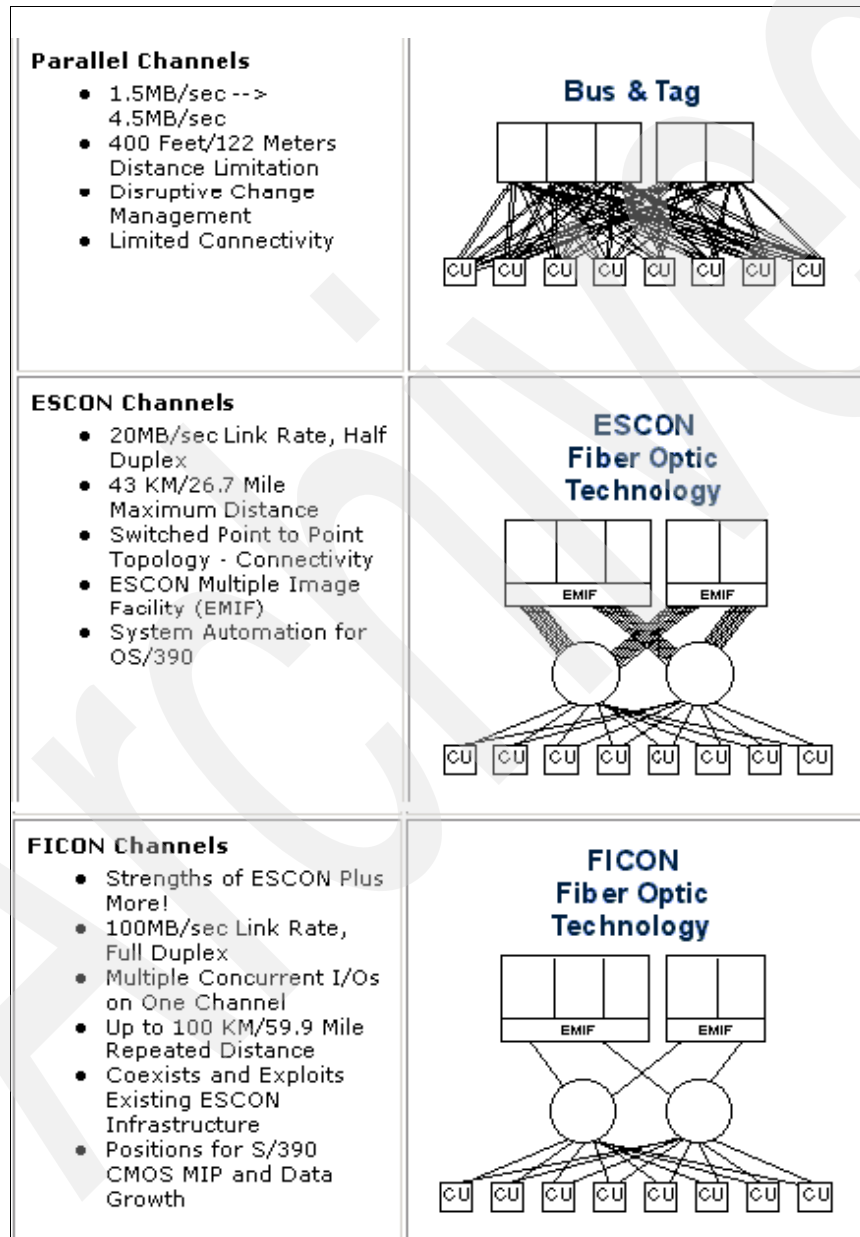


Figure 3-4 Development of zSeries storage network

The first implementation of Fibre ESCON® channels was a replacement for large, heavy copper Bus and Tag cables. The new channels were still used for direct connections to storage systems because they were faster and easier to deploy.

The step forward into what we now call storage networks came in the early 1990s with the introduction of high availability, high port count switches called ESCON Directors (Model IBM 9032).

The other big step was the introduction of ESCON Multiple Image Facility (EMIF) on zSeries. This enabled multiple logical systems to access consolidated storage subsystems through a common infrastructure.

Important: This is similar to functions now available on high end Open Systems, such as the IBM pSeries where you have an I/O LPAR that allows multiple LPARs or systems to share fiber and LAN adapters, therefore, reducing the number of adapters and simplifying the infrastructure.

3.3.2 How did these components simplify their infrastructure

In this section, we look, at a high level, at the components of a zSeries storage network that contribute to its high availability and efficiency.

- ▶ Homogeneous system and operating system
- ▶ Standard volume, LUN sizes, and definitions
- ▶ Simple fabrics, small number of high port count, and high availability switches
- ▶ Management software
- ▶ Virtualized resources and automation

Homogenous systems and operating systems

All the systems in a S/390® storage network are the same. Since IBM makes both the servers and the operating system, interoperability is very simple.

The only question is: Does the storage system support zSeries?

Note: Adding a virtualization layer to an open heterogeneous environment gives some of the advantages of having homogeneous systems or storage, because all of the I/O goes through a common interface (the virtual layer).

Standard volume or LUN sizes and definitions

The volumes in a zSeries system are called 3390 or 3380 volumes. The 3390 was the standard IBM disk in the early 1990s when IBM introduced ESCON. The 3390s were physically large disks approximately 19 inches in diameter mounted in subsystems with up to 16 drives per system. These systems were then linked to dedicated 3390 disk controllers.

The size of the modern virtual 3390 volumes are still defined by the characteristics of the original physical disk (number of tracks, cylinders x heads) in a specific model of 3390.

Tip: These volume sizes are small by modern standards, so functionality was added to allow users to customize some volumes by adding more tracks.

Because these volumes are standard, all disk systems from all vendors look the same to the system. A zSeries storage network can have disks from various manufacturers such as HDS, IBM, and EMC. Once they have been installed and configured, the storage administrator does not need to know or care from where the disks come. It is just a pool of storage accessed through common interfaces through a common infrastructure.

Some advanced functions like point-in-time copy and mirroring are tied to the subsystem and manufacturer.

Table 3-1 shows standard volume sizes in zSeries.

Table 3-1 Standard Zseries volumes and capacities

Logical device types	Cylinders	Bytes per cylinder	Logical device capacity (GB)	Physical capacity used (GB)	534 byte sectors per cylinder
3320-2	2,228	840,080	1.802	1.982	18,680
3390-3 custom (1)	3,330		2.838	2.943	
3390-0	1 - 3,330		0.00085 - 2.838	0.00178 - 2.048	
3390-0 custom (1)	10,017 (2)		8.514	8.828	
3390-2 (3380)	340 - 32,780 (2)		2.839 - 27.845	2.044 - 28.888	
3390-2 (3380)	2,228	712,140	1.585	1.821	1,580
3390-2 (3390) custom (1)	3,330		2.377	2.781	
3320-2	1 - 3,330		0.00071 - 2.377	0.00168 - 2.791	
Notes:1. A CKD volume that has a capacity different from that of a standard 3390 device type is called a custom volume. 2. In an interleaved partition, the number of cylinders is 10,017. In a non-interleaved partition, the number of cylinders can be from 3,340 to 32,760.					

Simple fabrics, high port count, and high availability switches

This is the description of the ESCON Director SAN. From the beginning, ESCON switches or “directors” were high availability (HA) high port count switches. Figure 3-5 on page 48, Port view of 9032 model 3 is from approximately 1995 - 1996. This is a one hundred and twenty port switch with the high availability N+1 features, additional power, controller card, and cooling that you see in modern HA core or director class switches.




Figure 3-5 Port view of 9032 model 3

Install these switches in pairs in order to give alternate consolidated paths from the host to the storage. In a modern SAN, this simple configuration is called a *collapsed core design*. We discuss this design in greater detail later in this chapter.

3.3.3 Comparing a S/390 Storage Network with an open SAN

One of the strengths of an Open Systems SAN is that you can have different systems from different manufacturers sharing the same infrastructure and resources.

Homogeneous systems

You can select different systems based on which vendor offers the best price, best performance, or particular systems specified as part of a new project or application. This ability to share resources is one of the important selling points for SANs, but this flexibility comes at a cost.

A small team of people can manage a large number of systems, a lot of storage, and a large infrastructure if it is all the same, or looks the same. Every time you add a different storage system, server type, or operating system to your infrastructure, you are increasing complexity because of the differences in the costs, number of people, and specialized skills required, and the increased difficulty managing your SAN and infrastructure.

Note: Adding a virtualization layer between heterogeneous servers and storage gives you many of the advantages of homogeneous servers and standard storage devices.

Standard LUN sizes

The question we need to ask is, “Does having standard LUNs bring any benefits to an Open Systems environment?”

The easy answer is it is too difficult, there is no ideal LUN size, because every application, server, and storage system are different. Another benefit is that disk can be carved up into LUNs once at install, then those standard sizes can be aggregated easily, using technology such as virtualization. Often, operating systems have difficulty expanding or shrinking existing LUNs. By forcing the LUN carving to happen only once at install, there can be greater simplicity in allocating them later.

To come to a decision, you need to look at your environment, then make an informed choice about whether or standardization will benefit your environment. The questions we need to ask are:

- ▶ Talk to the application people. Ask them how big their databases are. How much storage they allocate to each system or server. Do they see performance problems or bottlenecks? What is their average annual growth? Are there specific recommendations for SAP or Oracle?
- ▶ Do the operating systems offer volume or storage management functions to improve performance and make management easier? If so, what are the limits and recommendations for the number and size of volumes?
- ▶ Look at the architecture of the disk storage system. What are the optimal LUN sizes to get the most efficient utilization of the physical disks or arrays?

In Example 3-1: Deciding on a Logical Unit Number (LUN) size standard, we look at a database application on a pSeries server running AIX® 5.2 with an IBM ESS disk system with 73 GB 15K RPM disk drives.

Example 3-1 Deciding on a Logical Unit Number (LUN) size standard

With ESS, the minimum configurable LUN size is 0.1 GB and the maximum is the total effective capacity of the RAID array that the LUN is defined on. In most cases, the choice of LUN size has minimal effect on performance. However, in an effort to simplify Storage Administration tasks, customers may wish to adopt a LUN size standard. This allows LUNs to be allocated and subsequently de-allocated and re-allocated in an orderly fashion, without wasting space. A consistent LUN size is also a key component of the “striping and spreading” technique. Since the usable capacity of any ESS rank is some multiple (6, 7, 3 or 4) of the disk size, using the physical disk size (roughly) as the standard LUN size allows for efficient allocation of the available ESS disk capacity, even when multiple array configurations are used. This would equate to a LUN size of 35.0 GB (for 36.2 GB drives), 70.1 GB (for 72.8 GB drives) or 140.2 GB (for 145.6 GB drives).

When there is a mix of physical disk sizes in your environment, consider basing LUN size on the size of the smallest disk. Or, for environments where ESS storage is shared across a large number of relatively small servers and smaller allocation units are desirable, consider some fraction of the disk size – such as 8.7, 17.5, or 35.0 GB (when using 72.8 GB drives). It is also possible to define more than one LUN size standard for an enterprise (e.g. 35.0 GB for large environments and 4.3 GB for small environments). Having multiple standard LUN sizes somewhat increases the complexity of the storage management task, but may provide for somewhat more efficient storage allocation if properly managed.

The AIX Logical Volume Manager (LVM) has limits on Physical Partition size and the number of physical partitions per physical volume. Therefore, from an AIX LVM point of view, the most “optimum” LUN sizes are determined by the formula $2^n \times 1016 \times 1 \text{ MB}$ (where $n = 0$ through 7).

Since ESS LUN sizes need to be in 1 GB increments, this would equate to possible LUN sizes of .9, 1.9, 3.9, 17.9, 15.8, 31.7, 63.5, or 127 GB. These are somewhat different from the LUN sizes suggested earlier. The AIX “optimum” LUN sizes may result in slightly better performance (through smaller Physical Partition sizes) and allow for Volume Groups with slightly larger capacity. However, since the ESS Physical Disk sizes are not evenly divisible by the optimum AIX LUN sizes, their use will typically reduce the effective ESS storage capacity by 10-15%. (For example, for an 6+P array of 36.2 GB drives, it is possible to get six 35.0 GB LUNs for a total capacity of 210 GB, or six 31.7 GB LUNs for a total capacity of 190.2 GB – a 9.4% difference.)

In this case it was decided that there were advantages to having standard LUN sizes. The discussion then centered whether to size these LUNs for optimal utilization of the available storage, or tuning for optimal performance with AIX / Logical Volume Manager

The decision was to go for 17.5 Gb which was a compromise between the 17.9 recommended for AIX Logical Volume Manager and an efficient fraction of the physical disk to optimize utilization

Note: Even if you decide against having standard LUNs, it is valuable to gather and document information on your primary servers, applications, and storage systems.

Simple fabrics, high port count, high availability switches

In open storage networks, the first and most commonly adopted switches were small eight or sixteen port switches. These switches were not high availability (HA) with a single power supply, fan, control, or processor and had to be rebooted as part of a microcode upgrade. Since these switches were a single point of failure, you needed to have at least two of them, probably connected by Inter Switch Links (ISLs). The limited port count on these switches and the demands of server and storage sprawl discussed in “From where does the complexity come?” on page 42 mean that you soon had to buy another pair and another pair. Then you had fabric sprawl.

There were lots of “buzz” words (nicknames) for these designs. Most of them were misleading:

- ▶ Meshed fabric
- ▶ Skinny tree
- ▶ Leaf
- ▶ Core edge
- ▶ Edge core edge
- ▶ Expanded core

See Figure 3-6: Meshed fabric.




Figure 3-6 Meshed fabric

There were problems with these fabrics, such as complexity and the number of ports wasted as ISLs (Inter Switch Links). The problems affect the management of the SAN.

- ▶ Finding or identify faults in these fabrics is very hard. The component switches are basic and have very little error logging or diagnostic functions.
- ▶ Identifying performance bottlenecks and hotspots in the fabric is challenging. Since the design is complex, it is easy to get into a situation where high performance servers or storage must make multiple hops between switches to get to data.
- ▶ Management is a difficult task. The switches talk to each other all the time giving updates about themselves and the status of the fabric. Any change, such as adding a device to a port, removing a device, changing zoning, or adding a switch, causes a huge amount of communication between the switches. If you make any change, you must wait for this commutation to finish before making another change. There are various recommendations regarding how long to wait, varying from five minutes to 24 hours, depending on how big and complex your fabric is and how cautious you are.

▶ **Note:** We recommend you wait at least two minutes for every switch in the fabric. So in a ten-switch fabric, wait 20 - 30 minutes between changes.

- ▶ Updating microcode or firmware is another problem. Sometimes a complete outage is required necessitating negotiating with the users to take down all the servers and storage. Remember to check interoperability on all hosts and storage devices for the new code, or you update one switch at a time. With each switch, you need to check the configuration of the switch as documented, and that everything you expect to be online is online. Update the code which can reset the switch and initiate a huge amount of inter-switch traffic. When the reset has completed, check the configuration of the switch to make sure all the attached devices have come back and are as expected and documented. It can take at least an hour to check, update, and check each switch. A ten-switch fabric can take two person days to update. While this upgrading is occurring, you have two streams of code running in the same fabric, which is an uncomfortable situation.

- This complexity adds to the risk of zoning or multipathing errors that cause problems later and are difficult to isolate.

The switches themselves are not High Availability, so any hardware failure can disable the entire switch either due to the fault or as part of the repair action.

A solution to these challenges is to build a High Availability (HA) fabric, such as two meshed fabrics making up two isolated paths for the host to the storage. You still have the same problems listed for a single meshed fabric, but any changes only affect half of the switches.

Important: It is critical in a configuration such as this one that any changes made on one side of the fabric are not replicated on the other side for at least 24 hours.

Figure 3-7: High availability meshed fabric is an example of a complex HA fabric.

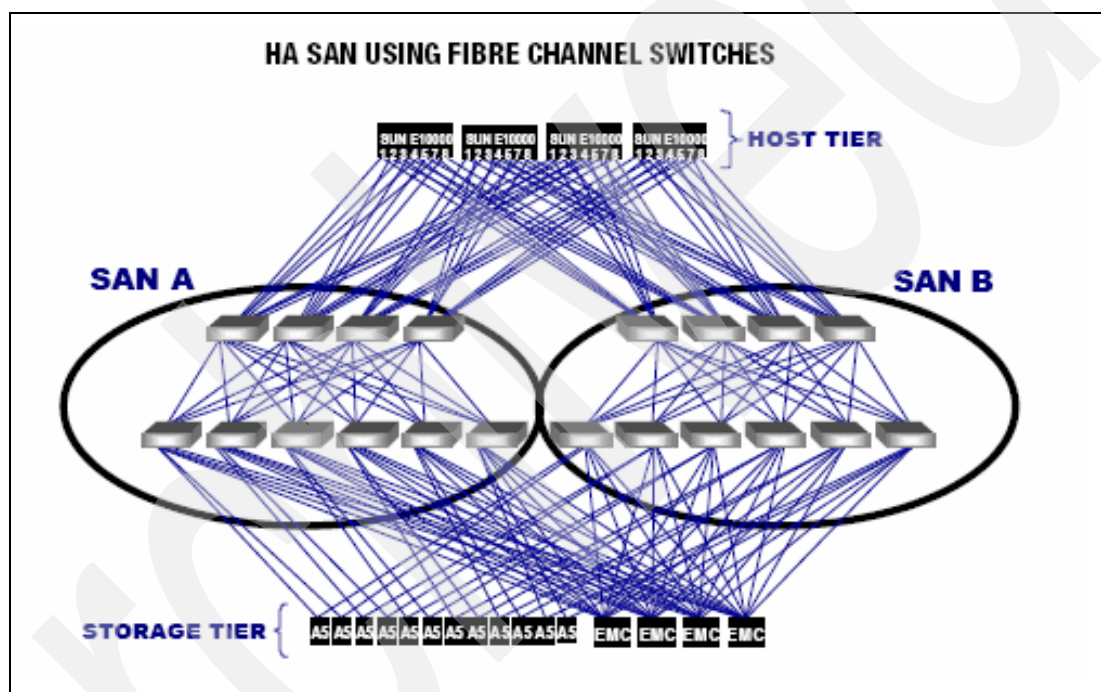


Figure 3-7 High availability meshed fabric

To get back to our comparison between zSeries and open storage networks. The zSeries networks started with high port count, high availability switches, so they never had the complex fabrics that grew and grew in the large open accounts. The first Open Systems Fibre Channel high port count HA switches came from companies such as McData that had their roots in the early ESCON networks.

A Director class or HA switch eliminates many complex fabric problems

A Director class or HA switch has excellent error logging, notification, and diagnostic tools. These switches have 99.999% availability with additional power, cooling units, no single point of failure, and the ability to hot replace most components.

A Director class or HA switch provide good performance monitoring because there is a single switch with multiple cards on blades on a common backplane. There are no problems with bottlenecks or hot spots.

A Director class or HA switch allows you to do code updates “hot” without affecting access from the servers to the storage. We still recommend waiting 24 hours before replicating any changes made on one side of the fabric to the other.

The configuration of the fabric is simple, typically two switches, so it is difficult to make zoning or multipath errors.

The simplest Storage Area Network has two switches

We need high port count switches to meet the requirements of large open organizations. These switches have the availability characteristics to build a 99.999% available SAN, and the management tools to make running a large port count network as easy as possible.

Figure 3-8: HA fabric with high port count switches is a HA fabric with high port count HA switches.




Figure 3-8 HA fabric with high port count switches

3.3.4 Simplifying disk storage

The simplest disk storage for your servers is internal or direct attached disk. The problems with this solution are:

- ▶ Difficult to scale, grow, or upgrade
- ▶ Cannot be shared as a common resource
- ▶ Expensive and inefficient to have storage pools limited to a server
- ▶ Expensive and difficult to manage in a large organization

The solution to these problems is to move to consolidated disk systems where dozens or hundreds of servers can share a common storage resource. This is:

- ▶ Comparatively easy to grow or upgrade
- ▶ The resource can be shared
- ▶ Relatively inexpensive storage
- ▶ Easy for one person or a small team to manage many terabytes of storage

The last point is important and needs to be discussed in more detail. The concept that the ratio of the number of storage administrators to terabytes of storage is proportional to the simplicity or complexity of the storage pool, in other words, the greater the complexity, the fewer terabytes per storage administrator managed. The simplest, consolidated disk storage pool is a single, high availability, high performance system like the IBM DS8000. It is easy to check the interoperability of your SAN and servers and easy to grow and manage. These advantages continue to be true even if you have many DS class systems.

Figure 3-9: Managing a complex heterogeneous SAN shows a complex heterogeneous environment which you have to manage.

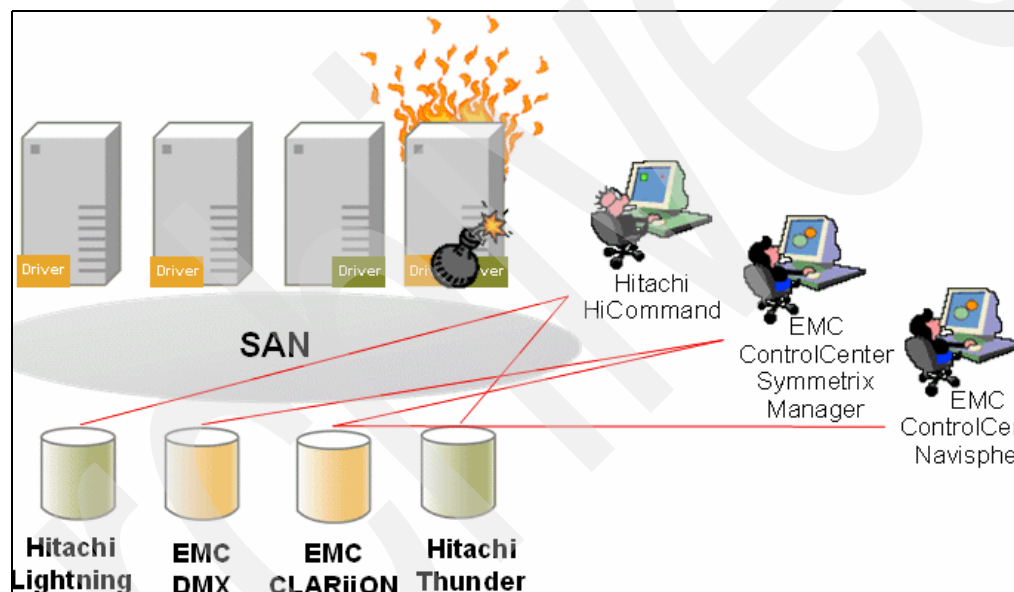


Figure 3-9 Managing a complex heterogeneous SAN

Complexity appears when you have multiple disk systems from different vendors making up your consolidated disk storage pool. Every time you add a different disk system, you are multiplying cost and difficulty in managing this common resource.

- ▶ Interoperability problems among different storage systems, the SAN, and servers
- ▶ Different management tools to configure and manage the storage
- ▶ Different characteristics or “sweet spots” in performance and configuration
- ▶ Disk systems can have different availability or performance standards
- ▶ Limitations on the number or size of volumes or LUNs can be different
- ▶ High-end functions such as point-in-time copy and remote mirroring are incompatible and require different management and monitoring tools

For both the reasons listed and many others, the ideal pool of consolidated disk storage is made up from disk systems with common:

- ▶ Management tools to reconfigure and monitor
- ▶ Interoperability
- ▶ Configuration and performance characteristics
- ▶ High-end functions

Consolidating all your storage to a single high-end product gives you all of these advantages and allows very efficient and cost-effective management of large amounts (hundreds of terabytes) of disk storage.

Figure 3-10: Common disk system shows an environment with heterogeneous servers using a common disk system.




Figure 3-10 Common disk system

Unfortunately, for most large businesses, it is too difficult to migrate all of their different storage pools onto one high-end platform. The reasons include:

- ▶ It is too expensive to buy new technology and throw away old technology.
- ▶ Migration is too difficult and risky.
- ▶ Financial reasons where business policy requires they buy from multiple vendors to make sure they get the best prices on any new purchases.

- Unwillingness to be dependent on a single vendor.

In a situation where we have heterogeneous disk systems and it is not possible to migrate to a standard platform, the solution is to install a virtualization layer between the storage devices and the servers. This virtualization layer can be in-band within the storage network or out-band in a separate server. Virtualization intercepts all I/O communication between the servers and storage. The servers see the LUNs or volumes they expect with no knowledge or requirement to know on which storage system the data is actually written. The storage talks to the virtualization device so the storage only sees one server type. This means we have:

- A common tool to reconfigure and monitor the storage.
- Simpler interoperability since servers and storage are talking to a single device.
- Standard high-end functions such as point-in-time copy and remote mirroring across all virtualized disk systems. This gives you the option to copy from expensive high-end storage to less expensive low-end systems for backup and disaster recovery.
- Back-end storage can be purchased from any vendor at the best price possible without adding to the complexity and cost of managing this storage.
- Once the disk systems have been virtualized, it is simple and transparent to the host systems to move data between disk systems as part of a plan to reduce maintenance costs or realistically tier your storage between high-end high performance, mid-range, and low-end low performing disk systems.

In a traditional SAN, capacity on physical disks is assigned to individual host systems. Unless you are really good at forecasting, you are going to end up with some disk systems that are really highly utilized and others that are not. Because of the physical limitations of dealing with physical storage, it is difficult, and often impossible, to load balance the environment. You end up underutilizing storage.

Figure 3-11 on page 57, SVC with heterogeneous storage, shows an environment with a virtualization layer, giving a common interface to different disk subsystems.

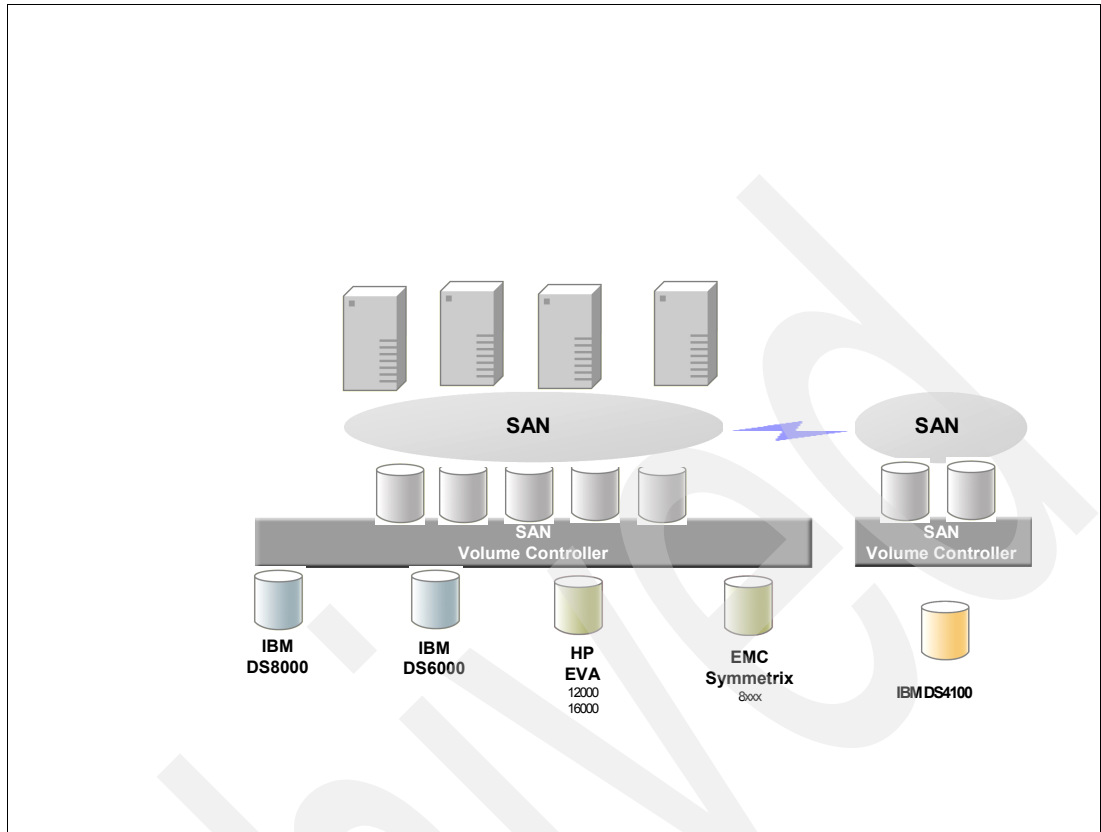


Figure 3-11 SVC with herterogenous storage

3.4 Tape backup infrastructure

In this section, we look at simplifying your backup infrastructure. The examples we use are tape backup, but the principles are true for solutions that use low cost disk as an interim backup.

3.4.1 Principles of backup solutions

The simplest backup solution is a tape drive mounted in or directly connected to a server to back up and restore the data or applications for that server. In an environment where you have hundreds or thousands of servers, this simple solution becomes too difficult because it is too expensive in people and time. It is unmanageable because it is impossible to manage that many tapes, confirm that backups have completed, and all tapes are able to be restored.

IMPORTANT: Remember that in the backup and restore equation **RESTORE** is the only reason to do backups. You do backups so that you can restore the data effectively if required.

Most backup solutions are designed, priced, and measured on their ability to backup data as fast as possible. The ability to recover the data fast is seldom considered. When you have to

restore a database or systems, the recovery time can be the most important thing in the world.

The solution to the problem of backing up hundreds or thousands of systems was to develop dedicated servers running backup software such as IBM Tivoli Storage Manager, Veritas Netbackup, and Legato. Initially, this software backed up data from other servers, normally over a TCP/IP LAN, and wrote multiple backups to direct attached high performance, high capacity drives.

Figure 3-12: Basic backup solution shows a simple way to do backups.

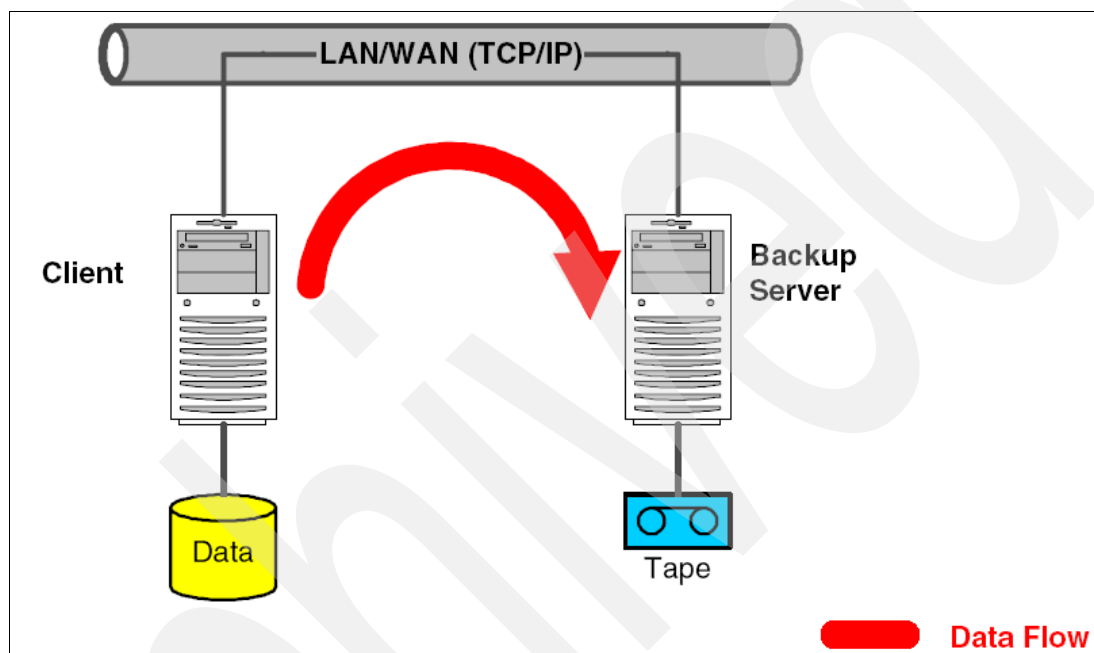


Figure 3-12 Basic backup solution

The next step in this consolidation process was to install high performance drives into robotic tape libraries or *silos* where simple robotics perform the mechanical functions of loading new cartridges, removing and storing full cartridges, and reloading the cartridges if a restore was required.

These drives were SCSI-attached, seriously limiting the distance between the server and the tape drives. The number of drives accessible by the server was also limited. In the 1990s fiber interfaces were added to the end drives to allow tape drives to be shared across the storage network. This enabled backup servers to share these tape drives, reducing costs by increasing the ability to share tape drives.

Figure 3-13: Consolidated solution with shared drives shows an environment with shared drives.

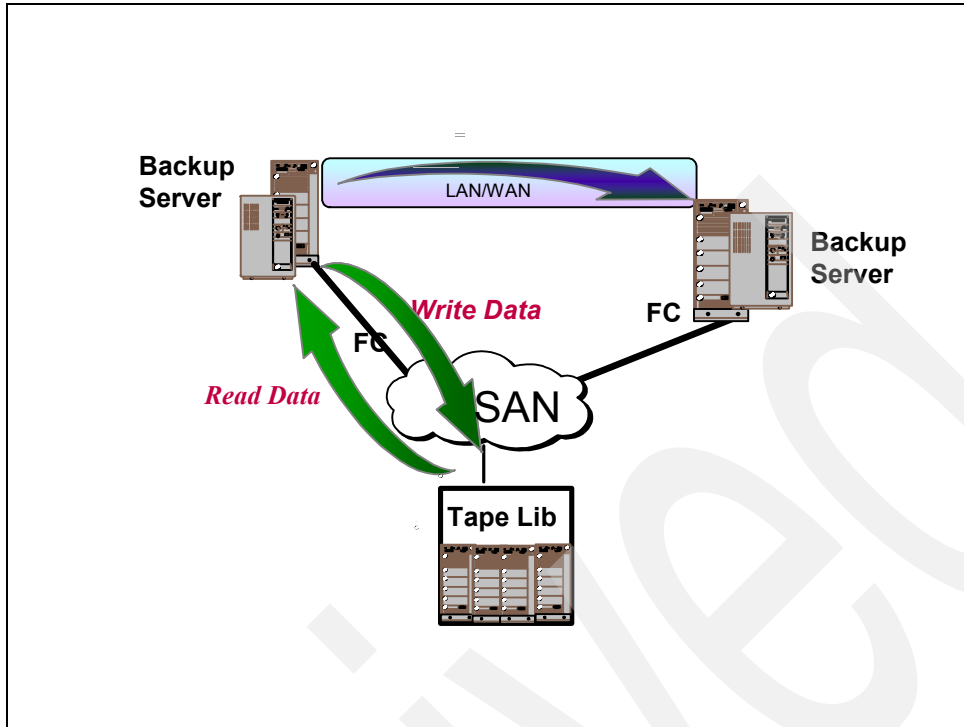


Figure 3-13 Consolidated solution with shared drives

This solution of sharing expensive SAN-attached drives mounted in automated libraries is still the standard for large organizations that need to back up huge numbers of servers and data.

3.4.2 Simplifying tape infrastructure

The key to simplifying your tape infrastructure was the introduction three years ago of the Linear Tape Open (LTO) tape drive. LTO was developed to be a new standard in tape by a coalition of IBM, HP, and Seagate.

When you are streaming large amounts of data to tape, LTO performance, throughput, and capacity can be adequate for many workloads when compared to enterprise tape. See Figure 3-14: LTO technology.

LTO Ultrium	Generation 1	Generation 2	Generation 3	Generation 4
Capacity				
– Native	100GB	200GB	400GB	800GB
– Compressed	200GB	400GB	800GB	1600GB
Transfer Rate	10-20 MB's	20-40 MB's	40-80 MB's	80-160 MB's
(Native)				

Figure 3-14 LTO technology

Note: The reason this is relevant to simplification is that for the first time we have high performance tape drives at a cost point where the infrastructure to share these drives (SAN ports) costs as much as the drives themselves.

3.4.3 The next simplification step was to consolidate

Consolidation of tapes and libraries connected to powerful backup servers was next, and, then they connected them to the SAN. This maximizes the ability to share the tape resources.

The next step in tape simplification is to treat these backup servers and their dedicated drives as building blocks or backup modules.

For example, we use a commodity server running Tivoli Storage Manager with two fiber HBAs, two Gigabyte Ethernet adapters, and four LTO2 drives. If we assume a compression ratio at the drive of 2:1, each drive can write 70 MB per second or 250 GB per hour. We have four drives, so 1 TB per hour. Figure 3-15: Tape backup block or module shows this configuration.




Figure 3-15 Tape backup block or module

Divide this throughput by the average amount of data on your servers (100 GB) and you have a rough sizing model of your backup window.

Note: With incremental backup, you are only backing up file data that changed, possibly only 5% to 10%. This changes the sizing model. Also, as a result, the number of servers backed up by each module dramatically increases. IBM Tivoli Storage Manager has a database to keep track of these incremental backups and facilitate a fast restore. This is not true with all backup vendors.

Each building block can back up ten servers per hour. In an eight hour backup window, you can back up eighty servers per building block.

The drives associated with a backup module are installed in a scalable tape library such as the IBM 3584. The 3584 is scalable to 16 frames with 192 drives and capacity for 6000 cartridges. Every drive can have a control path to the library robotics. Any one or group of drives can be defined as a virtual library within the physical frames with its own drives, cartridge slots, and media.

For example, you can define a backup module with four drives as a virtual library within the IBM 3584 with a pool of 200 slots and cartridges. This small library is easy to manage and makes the growth and cost of your backups more transparent.

Figure 3-16: Overview of a tape backup module in a 3584 shows this configuration.

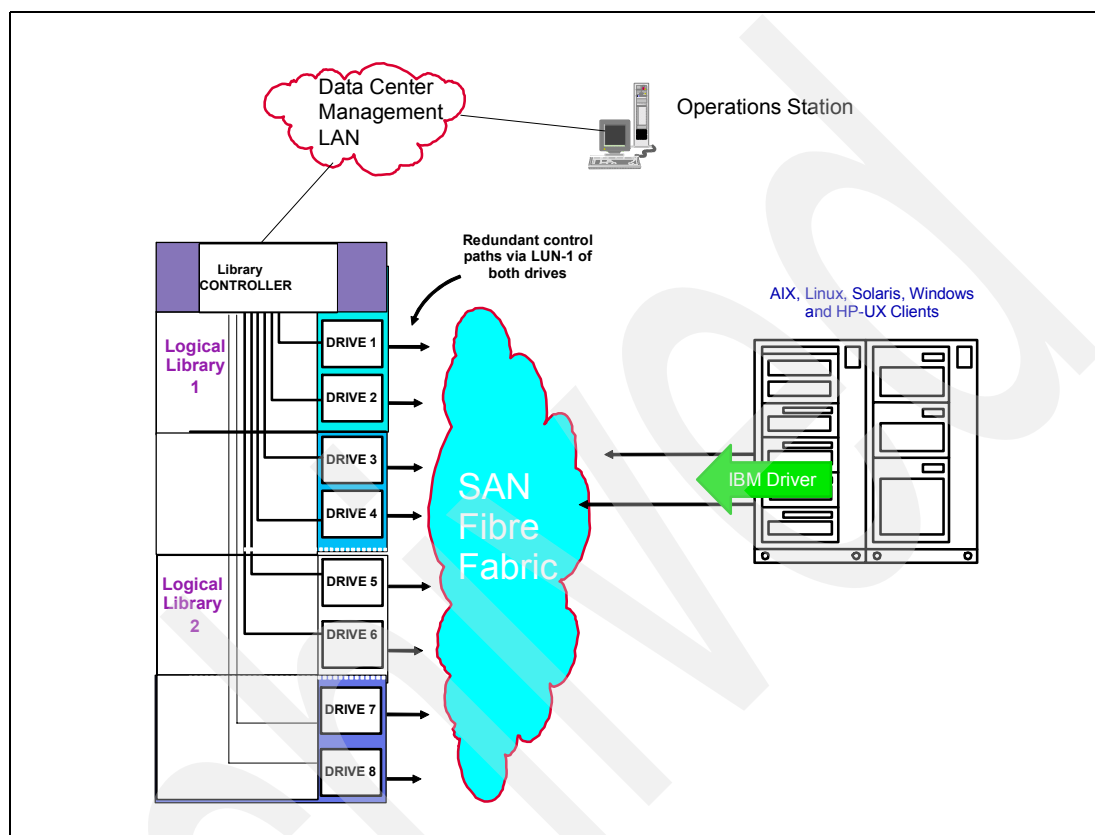


Figure 3-16 Overview of a tape backup module in a 3584

Note: This is a simplification of the sizing process to illustrate how to simplify backups. We break the solution down into easy to understand blocks. In reality, you must conduct a detailed examination of the data, the LAN, and the realistic throughput of the servers in order to correctly size this solution.

For remote clients, the backup data is transmitted over a wide area network (WAN) to the backup module and onto tape.

Figure 3-17: WAN-connected clients shows this type of configuration.

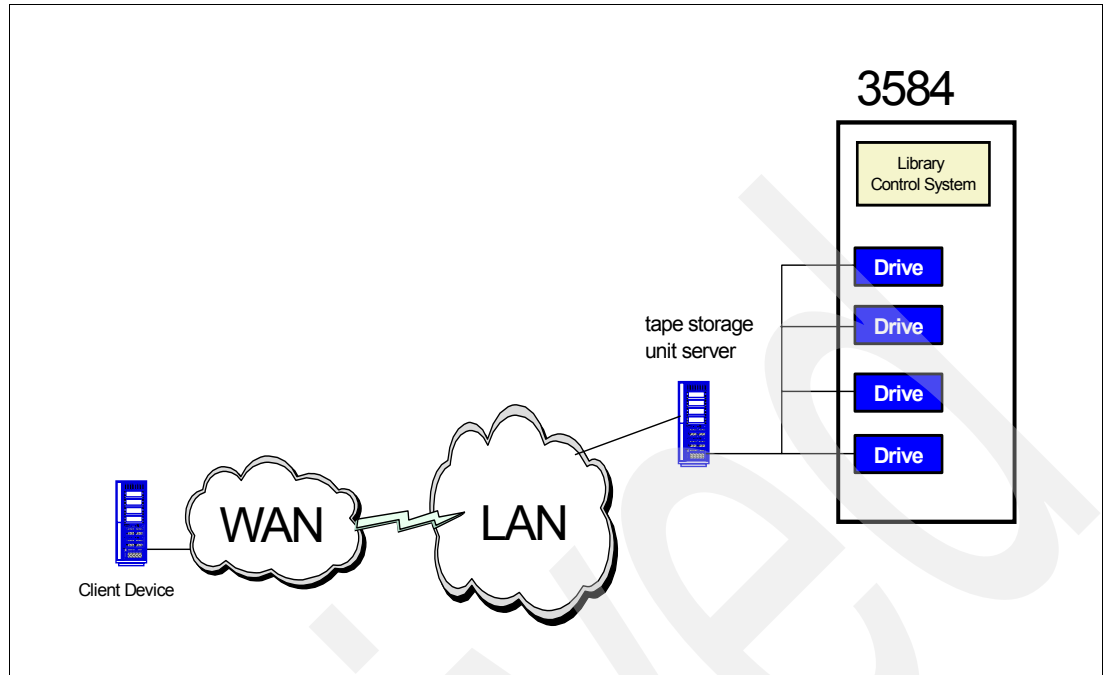


Figure 3-17 WAN-connected clients

For critical servers or servers with more than 500 GBs of data to back up, traditional LAN backup may not be suitable. If we follow the same argument on cost and simplification, then giving these servers dedicated drives makes sense.

- ▶ It is cost-effective.
- ▶ It is simple.
- ▶ It gives these important servers the resources they need to back up, and, more importantly, to restore data in the required time.

Note: An interesting combination of server and backup simplification is a high-end pSeries server (595), for example, where one LPAR with dedicated tape drives is used as a backup server for production LPARs in the same system.

Disk backup or tapeless backup

A hot topic at the moment is disk backup. The hype is “Disk is inexpensive. Why bother with tape?” Some vendors who only make disk products may say your data is totally secure, and that they make multiple copies of the data, on-site, off-site, anywhere you want. There are places where tapeless backup to disk is still important, and IBM offers these capabilities via IBM disk and software (IBM Tivoli Storage Manager).

There is little information about the total cost of having multiple copies of your data spinning on low-cost, low reliability disks. Raw disk priced per terabyte is declining year to year. The reliability characteristics of these disks mean RAIDing is essential. RAID5 has a capacity overhead of at least 20%. RAID10 or mirroring has an overhead of 50%. To get ten terabytes of usable disk backup, you need to buy twenty terabytes of raw disk.



Case study in Infrastructure Simplification

This chapter describes the challenge of a healthcare facility we refer to as Customer M to improve the quality, speed, and accuracy of patient care. Customer M centralized patient information and achieved its success using IBM TotalStorage Products.

4.1 The solution

This chapter goes into more detail about the solution we used to meet the challenge we describe. The solution itself was a virtualized storage management suite that enables higher availability, scalability, and control over data for streamlined management, improved data recovery, and ease of migration.

4.1.1 The benefits

Customer M demonstrated that it was one of the finer academic healthcare providers by establishing a flexible, reliable IT infrastructure through centralized patient data on demand, more proactive diagnostics, and medical collaboration.

4.2 Demonstrating Infrastructure Simplification

Customer M knew that continuing to deliver high-quality, personalized care required improving the accuracy of patient information while maximizing resources and enhancing research capabilities. This also meant taking advantage of new medical technologies to improve the use of information to deliver care. Customer M needed a means for accessing and sharing timely, accurate patient data throughout the hospital. To raise the standard of patient care through information-based medicine, Customer M identified four immediate goals:

- ▶ Improve clinical outcomes for patients
- ▶ Increase patient safety and reduce the risk of medical error
- ▶ Improve access to health services across the entire continuum of care
- ▶ Reduce the cost of delivering care

Customer M knew that accomplishing these goals would require a major technology commitment. To begin with, its storage was isolated on individual systems, resulting in poor utilization, low administrative productivity, and higher risk of down time. Redesigning Customer M's underlying IT infrastructure was necessary to enable electronic patient records and to provide uninterrupted data access on demand.

4.2.1 A strategic alliance with one focus: Enable the highest quality healthcare

Customer M launched a multiyear project divided into five phases designed to enhance patient care, hospital operations, and improved work life for the physicians and staff. The IT staff at Customer M knew that it needed the help of technology partners to achieve the hospital's goals. The hospital selected IBM as its primary information infrastructure provider. The IBM proven track record in healthcare, deep domain expertise, and vision for information-based medicine were key decision making factors for Customer M.

Supporting Customer M in its extensive multiyear commitment, IBM dedicated a full-time specialist for the project's duration. This on-site consultant helped the hospital identify and implement the newest IBM technologies that best support its ongoing needs and long-range objectives.

Capitalizing on IBM dedication to information-based medicine, Customer M laid the groundwork for the transformation of its care delivery by building a solid electronic foundation of reliable servers and storage.

Phase One

Phase One of the project began with an assessment of the new data center's requirements, during which IBM helped develop a virtual storage strategy for its SAN components. To manage the demands of patient and other clinical data, Customer M needed a reliable, scalable storage infrastructure that could handle large data volumes and throughput while eliminating down time. IBM designed a consolidated storage environment to help Customer M access information on demand, providing rapid information delivery, responsiveness, and resilience around the clock.

By establishing a scalable foundation for information-based medicine, Customer M is taking a bold step that will further its reputation as a leading provider of academic healthcare. To help Customer M achieve its goals, IBM outlined the following strategic steps:

- ▶ Create a centralized patient record that you can access in near realtime, enabling safer, more proactive patient care.
- ▶ Define an IT infrastructure that eliminates the need for future retooling.
- ▶ Build a scalable storage infrastructure with demonstrated reliability for disaster tolerance.
- ▶ Enable medical images to be stored, retrieved, and viewed electronically by multiple users by transitioning its imaging processes from film to a state-of-the-art picture archive communication system.
- ▶ Enhance research by enabling advanced collaboration among clinicians across the enterprise.

Multilayered data management built on a solid storage foundation

IBM consolidated Customer M's disparate storage infrastructure by connecting Customer M's IBM eServer™ pSeries and iSeries™ servers to an 8 TB IBM Enterprise Storage Server 800 and two IBM TotalStorage DS4500 storage devices, enabling fast access across a secure picture archive communication systems environment. See Figure 4-1. To centralize Customer M's tape backup processes, IBM installed an IBM 3584 Ultrium Ultrascable Tape Library. IBM Tivoli Storage Manager software automates backup of server data and applications and enables the SAN-connected servers and Storage Manager client computers to make maximum use of their direct network connection to storage.




Figure 4-1 Customer M's simplified Storage Infrastructure

LAN Free Backup describes a method of transferring data to be backed up with out use of local area networks except for transfer of meta data. Figure 4-2 depicts a Lan free backup environment.




Figure 4-2 Lan free backups

Reducing complexity and costs with virtualized storage management

At the heart of the IBM solution is the TotalStorage SAN Volume Controller, which drives down the cost and complexity of storage management while providing Customer M with greater flexibility. TotalStorage SAN Volume Controller (SVC) supports centralized data management and improves overall access to data by allowing Customer M to pool storage from disparate storage controllers into a single managed resource. See Figure 4-3.

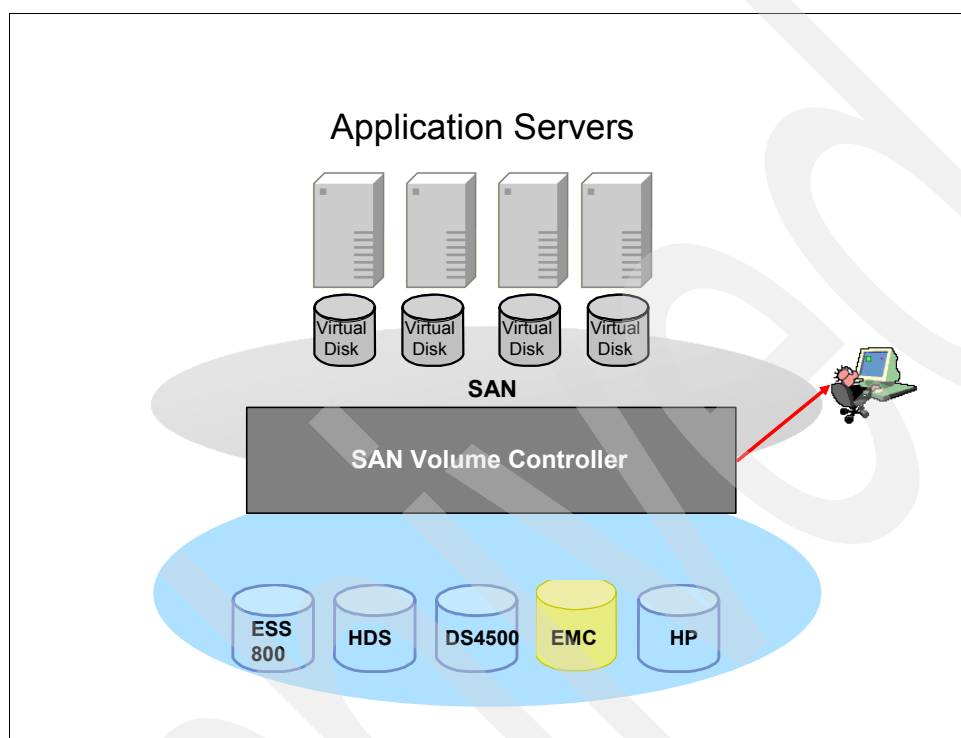


Figure 4-3 Storage reservoir

The virtualization of storage means the hospital can quickly and easily reallocate resources for optimal data backup and recovery within a secure environment. This gives Customer M a more resilient, highly available storage infrastructure that increases uptime to 99.99%. Customer M can also readily adapt its infrastructure to accommodate new users, applications, and storage technologies over time. In fact, the hospital has already made plans to capitalize on its technology-ready advantage.

IBM Tivoli Storage Manager delivers scalability, intelligent data technology, disaster preparation, and support for Customer M's information platform and medical applications. To protect patient data from hardware failures and other errors, IBM Tivoli Storage Manager software provides backup and archiving in offline storage. For increased recovery management, IBM Tivoli Storage Manager ensures high availability for critical data and applications in the event of hardware, software, or network failures. IBM Tivoli Storage Manager's intelligent, automated, and policy-based data movement and storage techniques decrease disaster-recovery time and increase service levels, while reducing administrative costs.

The foundation for information-based medicine

With this robust and flexible infrastructure in place, Customer M is positioned to enable an electronic medical record system that includes a new clinical data repository with an improved

viewing application. These enhancements will yield immediate benefits for Customer M's patients, physicians, and administrators, such as:

- ▶ A single, centralized source for patient information
- ▶ Enterprise-wide, integrated clinical applications and databases
- ▶ Faster access to radiology reports and images
- ▶ Distributed, simultaneous access to medical images

Enabling the future of quality care

From the project's launch, the IT and medical staff at Customer M approached the project collaboratively with a focus on detail, right down to education and training for staff, physicians, residents, and students.

Subsequent phases of the project will include enabling electronic access to advanced clinical applications with clinical documentation, physician order entry, a pharmacy system and medication safeguards, patient scheduling, and systems that include patient tracking and monitoring. The end result will be an IBM On Demand medical organization where critical patient information is digitally available at the point of care, enabling physicians, clinicians, and researchers to ensure that patients receive the highest quality care.



Part 3

Automated storage management

In this part, we discuss In-Band and Out-of-Band virtualization, TotalStorage software, and automated storage management. In addition, we provide an example of using Tivoli Provisioning Manager and SAN File System.

Archived

Storage virtualization

This chapter gives you an overview of storage virtualization and how IBM approaches different types of storage virtualization. This chapter also provides an overview of the IBM TotalStorage contributions to storage virtualization.

Deploying a storage network requires many choices. Not only are there Storage Attached Networks (SANs) and Network Attached Storage (NAS) to consider, but also other technologies, such as iSCSI. The choice of when to deploy a SAN or use NAS continues to be debated. CIOs and IT professionals must plan to ensure that all of the components from multiple storage vendors work together in a storage virtualization environment to enhance their existing storage infrastructures, and/or build new infrastructures, while keeping a sharp focus on business efficiency and business continuance.

This chapter covers:

- ▶ Types of storage virtualization
- ▶ IBM TotalStorage SAN Volume Controller
- ▶ IBM TotalStorage SAN File System
- ▶ IBM TotalStorage DS Family
- ▶ IBM N Series System Storage
- ▶ Management and productivity

5.1 Introduction to storage virtualization

Successful businesses require real-time responsiveness to change, whether it is due to new customer needs, changes in the supply chain, unexpected competitive moves, external threats, or changes in the economic climate. Rapid response to change requires an IT infrastructure that can turn information into a competitive advantage; the IT infrastructure must provide maximum benefit at an affordable cost, and must have the flexibility to support changes in business processes. An on demand operating environment provides a cost-effective and flexible IT environment. With information at the heart of competitiveness, storage becomes an ever more critical component of an On Demand operating environment.

Success in the on demand world will depend on the ability to leverage information technology. A greater dependence on information means a greater dependence on storage infrastructure. What differentiates an On Demand Business is the ability to quickly sense and rapidly respond to a dynamic marketplace. To do this, there are challenges that an On Demand Business must overcome.

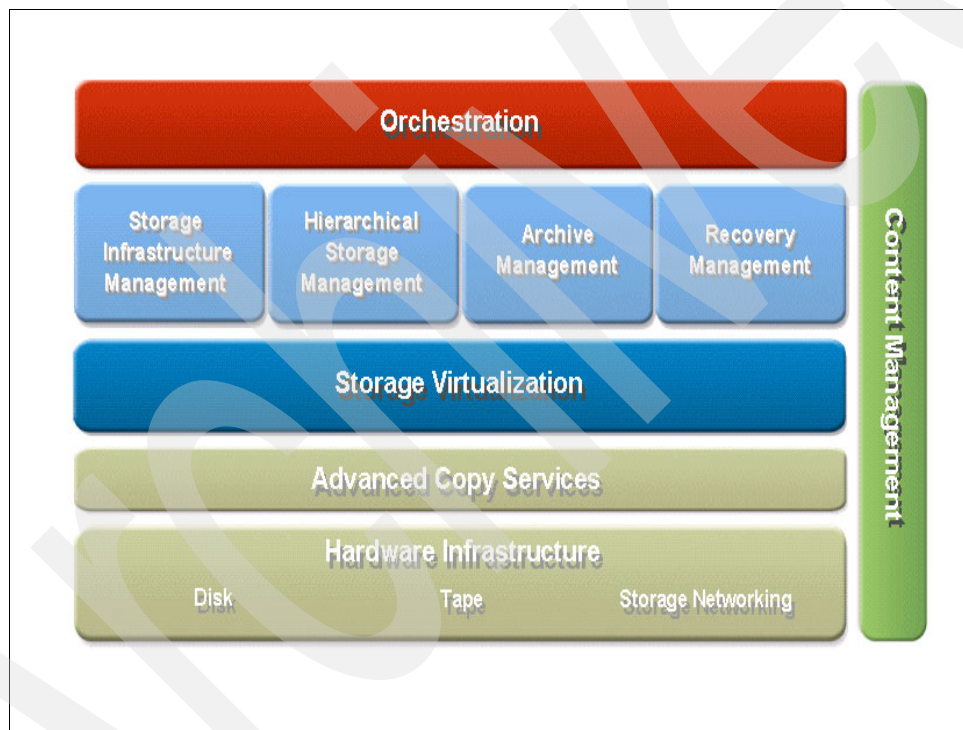


Figure 5-1 IBM vision of the On Demand storage environment

At the business level, customers face three major challenges:

- ▶ **Managing infrastructure growth:** Storage needs continue to grow at over 50% per year. Managing storage infrastructure becomes more complex every year, since we now have to deal with multiple server platforms and different operating systems, which may be connected to a storage area network (SAN) with multiple and diverse storage platforms.
- ▶ **Increasing complexity:** Although the declining cost of storage per megabyte makes it attractive to add additional disks, the increasing complexity of managing this storage results in overutilized staff and underutilized IT resources. Combining this with the shortage of skilled storage administrators, it is possible to add significant cost and introduce risk to storage management.

- **Maintaining availability:** The added complexity of 24x7 environments significantly reduces, for example, the efficiency of conducting routine maintenance, scheduling backups, data migration, and introducing new software and hardware. This problem is compounded by the fact that as availability increases, so does the cost inherent with making it so.

These challenges still exist, although large SANs do offer desirable and tangible benefits, for example, better connectivity, improved performance, distance flexibility, and scalability. Yet even these benefits may be outweighed by the added complexity that they introduce.

As an example, large enterprise SANs often contain different types of storage devices. These differences can be in the size and types of disk deployed, their level of performance, or the RAID configurations available. Often customers have different vendor storage devices as the result of mergers or consolidations or as a means of creating competition among vendors. The result, however, is that storage infrastructure and SAN administrators need to configure storage to servers, and then keep track of which servers own or have access to, that storage. The storage administrative tasks can become daunting as the SAN grows and as the storage administrators manually attempt to manage the SAN.

Furthermore, the complexity of different file systems in the same SAN requires that storage administrators know how to administer each client operating system (OS) platform. The management interfaces for each can differ, and each storage device requires its own unique multipathing drivers. Lastly, since the file systems are tied to each of the servers, storage management functions potentially have to be run on hundreds of servers. It is easy to see why manageability and interoperability are the top areas for concern, especially in a SAN where the number of possible storage and OS platform permutations are considerable.

These challenges are at odds with the commonly held belief that storage cost per megabyte is decreasing. It is clear that the cost of managing storage is greater than the initial purchase price. You need to simplify the storage infrastructure to address storage manageability, while at the same time addressing the need for interoperability. The solutions which are described in this chapter are designed to help simplify your storage infrastructure, optimize your storage utilization, and enable your business to adapt quickly and dynamically to variable environments. These solutions represents the next stage in the evolution of storage infrastructure.

Storage virtualization is one of the current buzzwords in the industry, especially with the increased acceptance of Storage Networks. But, besides all the hype, there is a lot of confusion, too. Companies are using the term *storage virtualization* and its characteristics in various and different forms. This chapter describes the reasons and benefits of storage virtualization in a technical and neutral way. The audience will understand the various terms and will receive a clear picture of the different storage virtualization types, levels, and approaches.

5.1.1 Storage virtualization concepts

Most customers are growing their number of installed TBs at a very rapid rate – 50%, 100% per year. The initial purchase price of all this disk is significant. With software and maintenance, the four year TCO can easily double. The financial problem is that you are installing new TBs at a faster pace than per TB prices are coming down. One of the strategies that IT managers and CIOs have used to combat this problem is to implement multi-vendor storage environments (either by choice or through necessity such as acquisition or data center consolidation) to generate competition between the vendors in order to get the best price. More recently, CIOs have started implementing multi-tier storage environments, some enterprise class storage and some midrange storage, the intent is to implement Information Lifecycle Management. The problem is that multi-vendor and even single vendor/multi-tier

environments come with technical challenges. Storage virtualization is a solution for improving storage utilization. By virtualizing the storage environment, you can easily solve problems such as:

- ▶ Difficult to maximize the utilization of the physical assets
- ▶ Difficult to keep up-to-date several multipathing drivers
- ▶ Difficult to manage different storage boxes

By virtualizing, you can leverage the following benefits:

- ▶ Improve application availability
- ▶ Optimize storage resource utilization
- ▶ Enhance storage personnel productivity

The Storage Networking Industry Association (SNIA) defines storage virtualization as:

“The act of integrating one or more (back-end) services or functions with additional (front-end) functionality for the purpose of providing useful abstractions. Typically virtualization hides some of the back-end complexity, or adds or integrates new functionality with existing back-end services. Examples of virtualization are the aggregation of multiple instances of a service into one virtualized service, or to add security to an otherwise insecure service. Virtualization can be nested or applied to multiple layers of a system.”

Note: The SNIA uses the term “aggregation” instead of “virtualization”.

Or, to put it in more practical terms, storage virtualization is the pooling of physical storage from multiple network storage devices into what appears to be a single storage device that is managed from a central console or a single reservoir of storage. The goal of virtualization is to logically simplify and generalize the physical infrastructure and its management. In Figure 5-2, you can see the basic concept of storage virtualization.

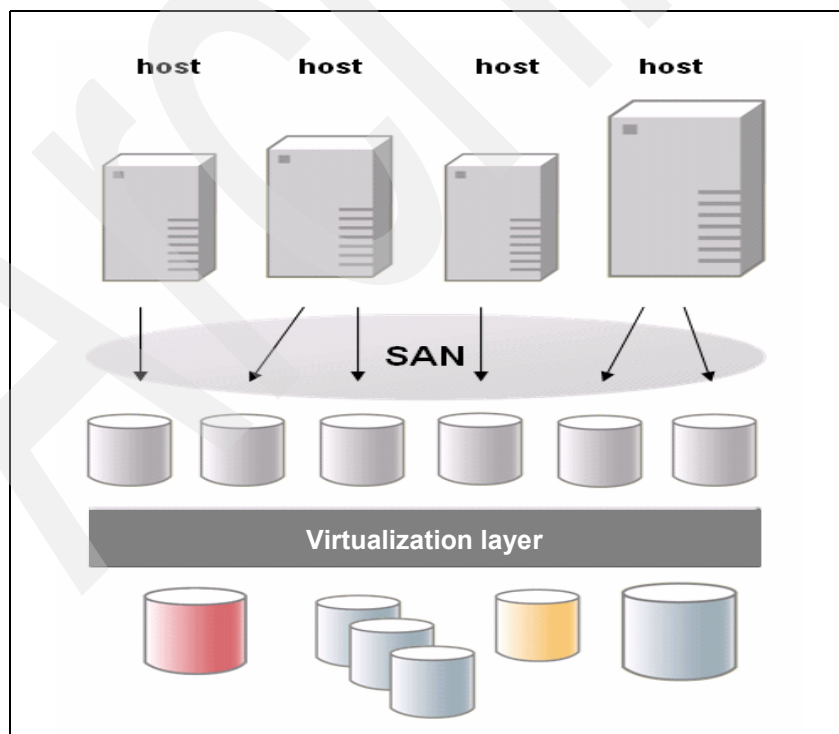


Figure 5-2 Basic concept of storage virtualization

Storage virtualization techniques are becoming increasingly more prevalent in the IT industry today. Storage virtualization forms one of several layers of virtualization in a storage network, and can be described as *the abstraction from physical volumes of data storage to a logical view of data storage*.

This abstraction can be made on several levels of the components of storage networks and is not limited to the disk subsystem. Storage virtualization separates the representation of storage to the operating system (and its users) from the actual physical components. Storage virtualization has been represented, and taken for granted, in the mainframe environment for many years.

The SAN is making it easier for customers to spread their IT systems out geographically, but even in networks, different types of servers that use different operating systems do not get the full benefit of sharing storage. Instead, the storage is partitioned to each different type of server, which creates complex management and inefficient use of storage. When storage must be added, applications are often disrupted. At the same time, the reduced cost of storage and the technology of storage networks, with faster data transfer rates, have enabled customers to use increasingly sophisticated applications, such as digital media. This has caused even greater complexity and difficulty of management because the amount of storage required grows at unprecedented rates. The IBM TotalStorage vision introduces ways to eliminate these problems.

5.1.2 IBM and storage virtualization technology

In this section, we describe the IBM TotalStorage technology and initiatives that form the IBM approach to storage virtualization, such as:

- ▶ IBM TotalStorage SAN Volume Controller (SVC)
- ▶ IBM TotalStorage SAN File System (SFS)
- ▶ IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000
- ▶ IBM N Series System Storage

SAN Volume Controller

Figure 5-3 shows the IBM TotalStorage SAN Volume Controller (SVC) hardware.



Figure 5-3 The IBM SAN Volume Controller 2145-8F2

SAN Volume Controller is a virtualization appliance solution that maps virtualized volumes visible to hosts and applications to physical volumes on storage devices. All servers that are served by the SAN can be connected to the SVC. This can include all servers on the SAN or only a subset. This enables the system administrators to view, access, and control a common pool of storage on a SAN, so they can use storage resources more efficiently. The SAN Volume Controller provides centralized management through a single interface to support easier storage allocation and address application demands. This flexibility provides the benefit of better storage utilization by reducing or eliminating the problem of unused storage found in direct attached storage implementations, and reducing required administrative time and resources. Each server within the SAN has its own set of virtual storage addresses which are mapped to a physical address. If the physical addresses change, the server continues running using the same virtual addresses it had before. This means that volumes or storage can be added or moved while the server is still running. The IBM virtualization technology

improves management of information at the block level in a network, enabling applications and servers to share storage devices on a network.

IBM TotalStorage SAN File System

A SAN-wide file system is for accessing data on storage networks across multiple application OS platforms and heterogeneous storage devices. The SAN File System also provides centralized, policy-based management of the data in the SAN. With the SAN File System, all of the files owned by the servers are made visible to clients of the SAN File System as a single file system in the SAN. This means that you can make all files in the SAN accessible to all of the servers, if you choose. This eliminates the need to maintain copies of the same file for use by multiple servers. Because there is a central catalog of all files in the SAN, policies can be established for each file in areas such as file placement, security, and service level requirements. With the SAN File System, the storage administrators do not need to assign storage volumes to individual servers. By not having to partition the storage across the application servers, you can have more efficient use of the storage and less storage administration required. For those familiar with DFSMS in the z/OS® environment, think of the SAN File System as the Open Systems version of DFSMS.

SAN Volume Controller Storage Software for Cisco MDS 9000

The Caching Services Module is shown in Figure 5-4.




Figure 5-4 Cisco MDS 9000 Caching Services Module

The SAN Volume Controller Storage Software for Cisco MDS 9000 is a joint project between IBM and Cisco, in which IBM provides the virtualization software (SAN Volume Controller), and Cisco provides the hardware platform (MDS 9000) and the Cisco MDS 9000 Caching Services Module (CSM). The SAN Volume Controller Storage Software for Cisco MDS 9000 is a storage virtualization solution that creates a pool of managed disks from the attached storage subsystems, which are then mapped to a set of virtual disks for use by various attached host computer systems. The system administrators can view and access a common pool of storage on the SAN, which allows them to use storage resources more efficiently, and provides a common base for advanced functions similar to those provided by the SAN Volume Controller. IBM has been working in the virtualization field for decades. Across its broad product line, IBM has achieved various storage virtualization milestones including:

- ▶ First in the industry to virtualize a “mainstream” operating system with OS/VS1 in 1972.
- ▶ First in the industry to introduce virtual tape storage solutions in 1997 with the IBM Virtual Tape Server.
- ▶ First in the industry to prototype the concept of a SAN-wide file system with policy-based storage management in 1997, resulting in the IBM TotalStorage SAN File System product, which debuted in 2003.
- ▶ First in the industry to virtualize two separate Virtual Tape Storage images into a single Peer-to-Peer Virtual Tape Storage image, in 2000.
- ▶ First in the industry to demonstrate and provide heterogeneous storage file virtualization using storage pools of different vendors, in 2001.
- ▶ First in the industry to have storage virtualization software achieve industry standards and pass SNIA's Conformance Testing Program (SNIA-CTP) for the Storage Management Initiative Specification (SMI-S).
- ▶ First in the industry to offer support for multiple disk systems from multiple different vendors with the IBM TotalStorage virtualization software.
- ▶ First in the industry to develop and ship storage disk solutions that enable clients to run multiple storage workloads, allocate resources and create “virtual” storage images using the IBM Virtualization Engine™ technology with the IBM TotalStorage DS8000.

IBM N Series System Storage

IBM N Series System Storage is designed to offer you fast data access with extremely low maintenance requirements for a highly capable data storage solution when attached to an IP network. The N Series storage system integrates storage and storage processing into a single unit, facilitating affordable network deployments. These advanced storage systems leverage a proven storage architecture and offer standard IBM N Series elements, including integrated I/O, high availability via clustering, and Fibre Channel disk drives. N Series models are designed to integrate easily into existing IT environments to deliver unified storage for organizations with NAS, iSCSI, FCP or combined environments, making enterprise-level storage a realistic goal for company sites regardless of size or staffing. Small to medium business customers who want pooled storage but do not have the time, money, or skills to implement a Fibre Channel SAN are one of the segments that benefit from the N Series.

N Series model

Though the N Series supports multiple protocols and access methods including TCPIP and NFS or CIFS which are the more widely used protocol and access methods. See Figure 5-5 on page 78.

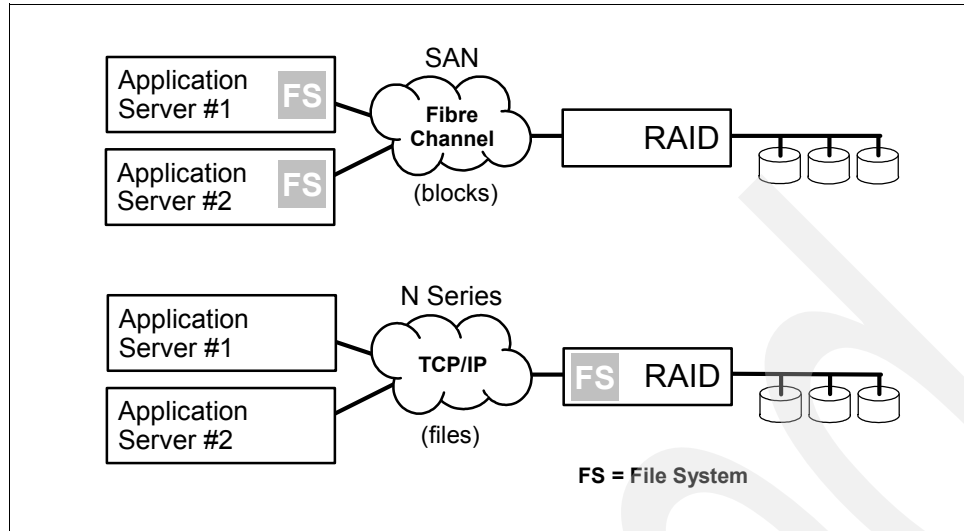


Figure 5-5 N Series versus SAN

The need to move to a single Storage Infrastructure often drives the decision to move to the N Series. If you are a customer still running applications on Direct Attached Storage, then you are faced with a decision to move to SAN, Network, or iSCSI-attached storage and determining which one is the best for your environment.

Key decision making points for you are:

- ▶ Current Network Infrastructure
 - Cost of adding, changing, or adding a new network
- ▶ Training costs associated with new infrastructure
- ▶ Implementation costs
- ▶ Management costs, including tools

The N Series is often the best answer for this initial move to a single Storage Infrastructure, providing you with a combination of SAN, iSCSI, and Network-attached Storage connectivity. See Figure 5-6 on page 79. The benefits to you are:

- ▶ Utilization of existing Network
 - Reduces training costs
 - Reduces amount of management and/or tools required
 - Makes use of existing skills
 - Ease of Implementation
- ▶ Entry level SAN storage
 - This can help with testing, training, and implementation experience until you need to implement products like the SVC, DS6000, or DS8000.
- ▶ iSCSI connectivity

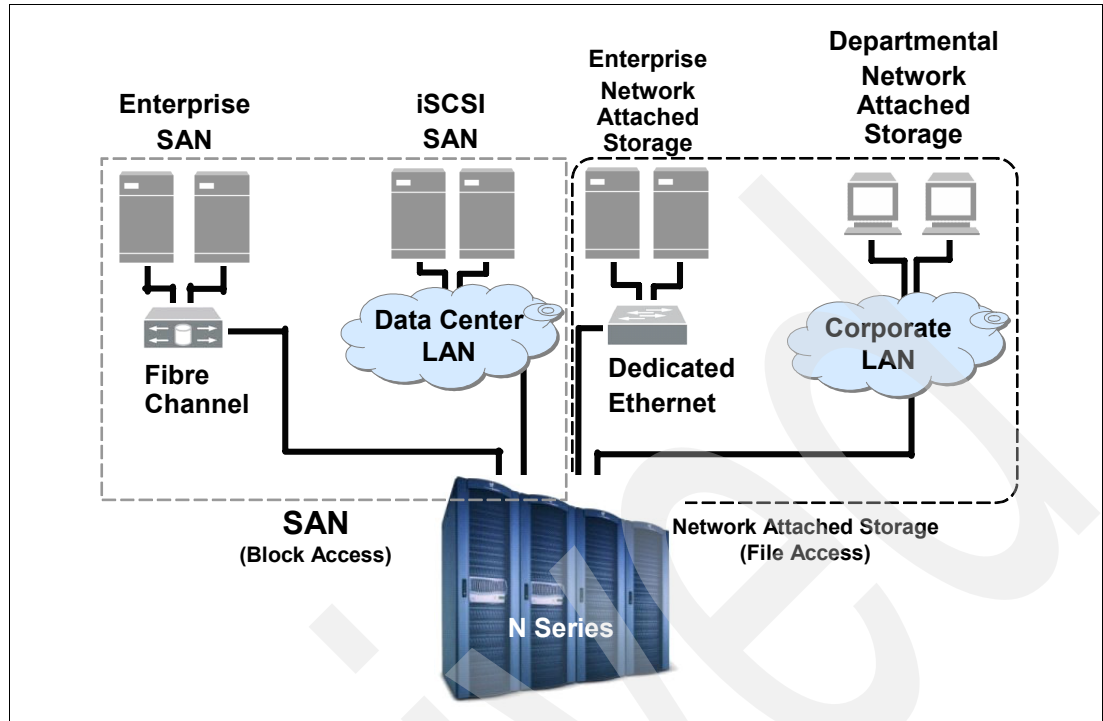


Figure 5-6 N Series connectivity

Figure 5-7 shows you the model N3700 of the N Series.



Figure 5-7 IBM N3700

Figure 5-8 shows you the model N5200 of the N Series.



Figure 5-8 N5200

5.2 IT storage architectural directions

Next, we describe the architectural influences and standards that drive the IBM TotalStorage Software vision.

Many contemporary authors state that storage has become a commodity. People want to be able to simply use storage without limitations or worries, to completely disregard its whereabouts and management, yet they always want to be sure of its abundance and availability. At the same time, storage costs have been steadily decreasing, and people have been implementing new ways of connecting storage devices. The volume of data storage required in daily life and business has exploded. Each year capacity is growing by over 50% and hardware cost is decreasing by approximately 30%, but availability requirements are approaching 100%. See Figure 5-9. Users are mobile, access patterns are unpredictable, and the content of data is more interactive.




Figure 5-9 Trends

Storage itself may well be treated as a commodity, however, the management of it is certainly not. It has been estimated that the cost of managing storage can be up to 11 times the cost of the storage itself.

Infrastructure simplification addresses the increasing complexity of managing storage and reduces the associated costs dramatically. The primary purpose of infrastructure simplification is the full exploitation of the benefits promised by a SAN. Storage virtualization can become an enabler for sharing data, ensuring higher availability, providing disaster tolerance, and improving performance. Storage virtualization allows for storage consolidation of resources, provides policy-based automation, and enables several other benefits, which do not automatically result from the implementation of SAN hardware components.

5.2.1 Automated Management

IBM TotalStorage Productivity Center is an open storage infrastructure management solution designed to help reduce the effort of managing complex storage infrastructures, to help improve storage capacity utilization, and to help improve administrative efficiency. It is designed to enable an agile storage infrastructure that can respond to on demand storage needs.

- ▶ IBM TotalStorage Productivity Center for Data¹
- ▶ IBM TotalStorage Productivity Center for Fabric²
- ▶ IBM TotalStorage Productivity Center for Disk³
- ▶ IBM TotalStorage Productivity Center for Replication⁴
- ▶ IBM Tivoli Provisioning Manager
- ▶ IBM Tivoli Storage Manager

IBM TotalStorage Productivity Center (Productivity Center) is a software package that has been designed to enable administrators to manage SANs and storage from a single console. This software solution is designed specifically for managing networked storage components based on the SMI-S, such as:

- ▶ IBM TotalStorage SAN Volume Controller
- ▶ IBM TotalStorage SAN File System
- ▶ IBM TotalStorage DS Family
- ▶ IBM Enterprise Storage Server (ESS)
- ▶ Other vendors' subsystems, including EMC, Hitachi, and Hewlett-Packard (HP)

These storage virtualization products are part of the IBM commitment to the open standards adopted by the Storage Networking Industry Association (SNIA). They implement standard CIM-based Application Program Interfaces (APIs) to allow management applications from IBM and other vendors to administer and monitor their activities.

We describe these products in more detail in Chapter 2, "Total Cost of Ownership of storage infrastructure" on page 17.

5.3 Types of storage virtualization

In the 1980s as the trend started to move away from centralized mainframes to distributed computing, the demand for personal computers increased dramatically. Hundreds of companies wanted hundreds of personal computers (PCs) for all sorts of different applications and all sorts of unique purposes. When IBM entered the market, the IBM PC was embraced as the standard PC and was eventually copied by many manufacturers.

Because the PC was built from off the shelf parts, it was called *open architecture*, and it became a standard on which to build and use as a base for development. It was good for the vendors, because they easily developed "interface cards". It was good for the customer, because they can use adapter cards from any vendor.

In storage today, there is a large demand to manage more and more data, and to share critical information. The solution to this problem is seen as today as storage virtualization. The question is which type of storage virtualization will be adopted as the standard for the future. No longer can vendors simply build the biggest and/or the fastest storage, but the

¹ formerly IBM Tivoli Storage Resource Manager

² formerly IBM Tivoli SAN Manager

³ formerly IBM TotalStorage Multiple Device Manager Performance Manager

⁴ formerly IBM TotalStorage Multiple Device Manager Replication Manager

simplest design and the least expensive storage to operate (lowest TCO) will be used as a base for future development.

Depending on where and how storage virtualization is implemented, you can differentiate solutions such as the following:

What is created:

- ▶ Disk drive virtualization
- ▶ Storage system partitioning (DS8000 series)
- ▶ Block virtualization (SVC)
- ▶ File system virtualization (SFS)
- ▶ File system virtualization (N Series)
- ▶ Tape, tape drive, and tape library virtualization (VTS)

Where it is done:

- ▶ Host or server-based virtualization (AIX LVM)
- ▶ SAN or network-based virtualization (SVC, SFS, and N Series)
- ▶ Storage device or storage subsystem virtualization (DS family, VTS)

How it is implemented:

- ▶ In-band virtualization (SVC, N Series)
- ▶ Out-of-band virtualization (SFS, N Series)

The IBM TotalStorage Open Software virtualization solution is SAN-based, which helps allow for a more open virtualization implementation. Locating virtualization in the SAN, and therefore, in the path of input/output (I/O) activity, helps to provide a solid basis for policy-based management. The focus of IBM on open standards means its storage virtualization solution supports freedom of choice in storage-device vendor selection.

Disk drive virtualization

The disk firmware abstracts the physical disk parameters such as cylinders, heads, and sectors into a single logical block address. Storage subsystem virtualization can take this to the next level by enabling the presentation of different disk drives.

Storage subsystem partitioning

Disk storage systems can provide some level of virtualization already by subdividing disks into smaller virtual drives. Conversely, more storage devices can be consolidated to form one large virtual drive. RAID subsystems are an example of virtualization at the storage level. Storage virtualization can take this to the next level by enabling the presentation, and the management, of disparate storage systems.

Block virtualization

Block virtualization can enable the independence of storage pools from heterogeneous servers into a single storage resource. The SAN fabric is zoned to allow the storage virtualization appliances to see the storage subsystems, and for the servers to see the storage virtualization appliances. Servers are not able to directly see or operate on the storage subsystems.

File system virtualization

The virtualization software abstracts multiple individual file systems into a single shared file system. The file system virtualization provides the highest level of virtual storage. It can also provide the highest benefit, because it is data that is shared, allocated, and protected; not volumes.

Tape virtualization

Virtualization software abstracts tape drives to provide immediate response to replication seeking access. Typically, this involves the use of disks as well as tapes.

5.3.1 The IBM TotalStorage vision

You can implement storage virtualization at the host, network, and storage level. The IBM vision is to move the storage device management intelligence out of the server, reducing the dependency of having to implement specialized software, like Logical Volume Managers (LVM), at the server level. By putting virtualization function in the network where it can be accessed by all connected servers and storage, organizations can avoid duplication of storage function, simplifying the environment.

By implementing at a fabric level, storage control is moved into the network, which gives you the opportunity to use all storage for storage virtualization, and at the same time reduces complexity by providing a single view of storage. The storage network can be used to leverage all kinds of services across multiple storage devices, including storage virtualization.

By implementing at a file system level, file details are effectively stored on the storage network instead of in individual servers and storage devices. This design means the file system intelligence is available to all application servers. Doing so provides immediate benefits: a single namespace and a single point of management. This eliminates the need to manage files on a server by server basis.

5.3.2 Storage virtualization models

For storage virtualization, two models can be drawn. The two models are *in-band* and *out-of-band*. These models are not mutually exclusive. In many environments, a combination of both models may be desired. Both models have their strengths and applications.

In-band virtualization

In-band virtualization is sometimes called *symmetric or synchronous virtualization*, otherwise known as *block aggregation*. The in-band virtualization model is shown in Figure 5-10.




Figure 5-10 In-band virtualization model

In a conventional SAN, the logical unit numbers (LUNs) that are defined within the storage subsystem are directly presented to the host or hosts. When we implement an in-band virtual storage network, both data and control flow over the same path. This means having a virtualization engine sit between the hosts and their storage devices in the data path that can take physical storage from one or more storage subsystems and offer it to hosts in the form of a virtual disk (VDisk). This engine manages storage virtualization by receiving the hosts' I/O commands and redirecting the data to the assigned storage device. The engines range from switch-based appliances to devices such as routers. Levels of abstraction exist in the data path, and storage can be pooled under the control of a domain manager. In general, in-band solutions are perceived to be simpler to implement, especially since they do not require special software installed in servers (other than conventional multipathing software). In-band solutions can also provide caching and advanced functions within the storage network. This can help improve the performance of existing disk systems, extend their useful life, and reduce the cost of new storage capacity by enabling the use of lower function, lower cost disk systems without the loss of performance. The IBM plan for block virtualization is shown in Figure 5-14 on page 90.

Other advantages include:

- ▶ Offloading function from the host
- ▶ Providing storage management for the SAN
- ▶ Performing performance optimizations in the data path
- ▶ Supporting host systems that are not in a cluster
- ▶ Supporting multiple heterogeneous hosts
- ▶ Integrating well with storage management software
- ▶ Releasing the customer from a particular vendor's storage
- ▶ Integrating with storage to create a better management picture
- ▶ Offering excellent scalability

Out-of-band virtualization

Out-of-band virtualization is sometimes called *asymmetric* or *asynchronous* virtualization. IBM uses out-of-band virtualization in its SAN File System. The out-of-band virtualization model is shown in Figure 5-11.




Figure 5-11 Out-of-band virtualization model

In an out-of-band implementation, the data and meta-data (data about the data) are separated into different places. Separating the flow of control and data in this manner allows the I/O to use the full bandwidth that a SAN provides, while control could go over a separate network, or routes in the SAN that are isolated for this purpose. This means the virtualization appliance is not in the data path. In an out-of-band solution, the servers request authorization to data from the meta-data controller, which grants it, handles locking, and so on. Once they are authorized, servers access the data directly without any meta-data controller intervention. Once a client has obtained access to a file, all I/O goes directly over the SAN to the storage devices. For many operations, the meta-data controller does not even intervene. This results in performance that is nearly equal to local file system performance with all of the benefits and added functionality that comes with an out-of-band implementation. Out-of-band virtualization involves moving all mapping and locking tables to a separate server (the meta-data controller) that contains the meta-data of the files. Typically, out-of-band virtualization is more targeted toward file sharing across the SAN than in-band virtualization. Out-of-band virtualization normally involves a single file system in a single name space.

File virtualization is a similar technique to block virtualization. However, rather than dealing with blocks of data, file virtualization addresses the needs of accessing and sharing files in a storage network. In the SNIA model, hosts get file meta-data from file system or Network Attached Storage (NAS) controllers, and then access the data directly. File virtualization can be used in conjunction with or independent of block virtualization. The IBM plan for file virtualization is shown in Figure 5-24 on page 100.

Other advantages include:

- ▶ Releasing the customer from a particular vendor's storage
- ▶ Providing storage management for the SAN
- ▶ Offering excellent scalability
- ▶ Offloading host processing
- ▶ Supporting storage management from multiple vendors
- ▶ Integrating well with storage management software
- ▶ Supporting multiple heterogeneous hosts
- ▶ Relatively low overhead in the data path

SAN File System is an example of an out-of-band virtualization implementation.

5.3.3 SAN Volume Controller in-band storage virtualization benefits

When we implement an in-band virtual storage network, both data and control flow over the same path. Levels of abstraction exist in the data path, and storage can be pooled. The SVC In-band solution also provides caching and advanced functions within the storage network. This can help to improve the performance of existing disk systems and can extend their useful life, and reduce the cost of new storage capacity by enabling the use of lower function and lower cost disk systems without the loss of performance.

Additional advantages include:

Table 5-1 Benefits

Function	Benefit
Ability to offload function from the host	Reduce contention for resources
Virtualization location	Between hosts and storage
Storage pooling	Over 1,024 hosts using Cisco or McData, 256 hosts per I/O group. See Figure 5-12 on page 87.
Host impact	No client/host software required.
Multipath software	Yes, on some OSs, free of charge.
Caching	4 GB of memory.
Performance impact	Performance neutral or, in some cases, a performance improvement.
Heterogeneous subsystem attachment	SVC supports the major subsystem vendors.
Adaptive performance	The SVC can intermix high performance and low performance subsystems as required by application needs. This capability extends the life of older technology by virtualizing it to less critical applications.
Centralized resource management	Once subsystem LUNs are mapped to the SVC resource, management is all done centrally from the SVC.

Function	Benefit
Increased Scalability	Accommodates growth.
WWPNs	2048.
Fabrics supported	Up to 4.
Metro Mirror consistency groups	Up to 256.

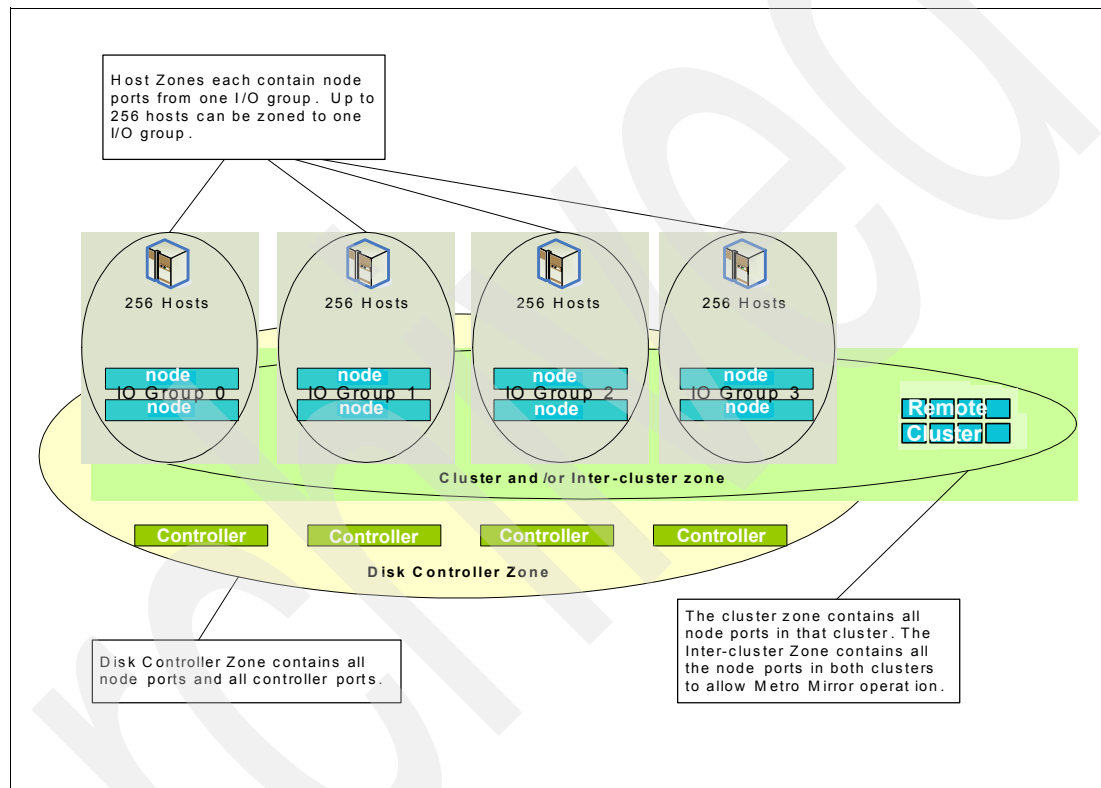


Figure 5-12 Scalability: 256/1024 Host Support

5.4 IBM TotalStorage SAN Volume Controller

The IBM TotalStorage SAN Volume Controller provides block virtualization and volume management for disk storage within the SAN. In simpler terms, this means that the SAN Volume Controller manages a number of back-end disk subsystem controllers and maps the physical storage within those controllers to logical disk images that can be seen by application servers and workstations in the SAN. The SAN must be zoned in such a way that the application servers cannot see the same back-end LUNs seen by the SAN Volume Controller, preventing any possible conflict between the SAN Volume Controller and the application servers that are both trying to manage the same back-end LUNs.

As described earlier, when an application server performs I/O to a VDisk assigned to it by the SAN Volume Controller, it can access that VDisk via either of the nodes in the I/O group. Each node can only be in one I/O group and since each I/O group only has two nodes, the

distributed redundant cache design in the SAN Volume Controller only needs to be two-way. The SAN Volume Controller I/O groups are connected to the SAN in such a way that all back-end storage and all application servers are visible to all of the I/O groups. The SAN Volume Controller I/O groups see the storage presented to the SAN by the back-end controllers as a number of disks, known as *managed disks or MDisks*. Because the SAN Volume Controller does not attempt to provide recovery from physical disk failures within the back-end controllers, we recommend, but do not necessarily require that MDisks are a RAID array. The application servers should not see the MDisks at all. Instead, they should see a number of logical disks, known as *virtual disks or VDisks*, which the SAN Volume Controller presents to the SAN.

MDisks are collected into groups, known as *managed disk groups (MDGs)*. The MDisks that are used in the creation of a particular VDisk must all come from the same MDG. Each MDisk is divided into a number of extents. The default, and minimum, extent size is 16 MB, and the maximum extent size is 512 MB, based on the definition of its MDG.

The virtualization function in the SAN Volume Controller maps the VDisks seen by the application servers to the MDisks presented by the back-end controllers. I/O traffic for a particular VDisk is, at any one time, handled exclusively by the nodes in a single I/O group. Although a cluster can have many nodes within it, the nodes handle I/O in independent pairs. This means that the I/O capability of the SAN Volume Controller scales well (almost linearly), since you can obtain additional throughput by simply adding additional I/O groups.

5.4.1 Block virtualization

Block level virtualization provides servers with a logical view of physical storage. The IBM TotalStorage SAN Volume Controller product provides advanced block virtualization capabilities.

Block virtualization manages multiple storage devices and volumes as groups. These groups are managed independently of the physical layout of the storage. Because of this independence, you can add new disk systems to a storage network, and you can migrate data to them without causing disruption to applications that use the storage.

Since the storage is no longer controlled by individual servers, any server can use it as needed. You can add or remove capacity on demand without affecting the application servers. Storage virtualization simplifies storage management and reduces the cost of managing the SAN environment. In Figure 5-13, we show the SNIA block aggregation model.

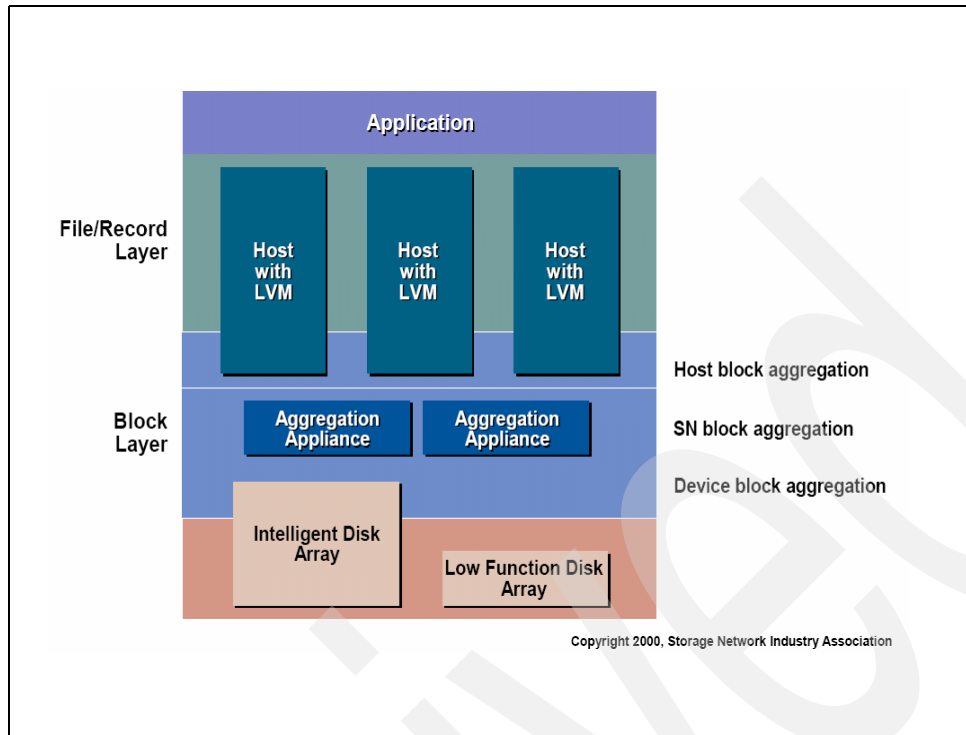


Figure 5-13 SNIA block aggregation model

Block aggregation provides the following significant benefits to customers:

- ▶ **Increased storage administrator productivity:** Administrators can manage, add, and migrate physical disks transparently. You accomplish this by providing insulation between the server's view of the logical disks and the actual physical disks. You improve productivity by reducing planned downtime and allowing administrators to perform management functions when convenient rather than waiting for ever-decreasing downtime windows.
- ▶ **Advanced functions provided by a common platform:** By providing a logical view of physical storage, you can perform advanced functions at a single point in the SAN in a common way, regardless of the underlying physical storage. You can also perform FlashCopy®, peer-to-peer data copy, and data migration in a common way. This common platform will be used to provide other advanced functions over time such as advanced security and quality of service capabilities.
- ▶ **Improved capacity utilization:** You can reallocate spare capacity on underlying physical disks without impact on servers, irrespective of the server operating system or platform type. You can create logical disks from any of the physical disks that the virtualization device manages.

In Figure 5-14, we show the IBM block virtualization plan as:

- ▶ In the network
- ▶ In the data path
- ▶ Move intelligence of controller into network
- ▶ Enterprise reliability

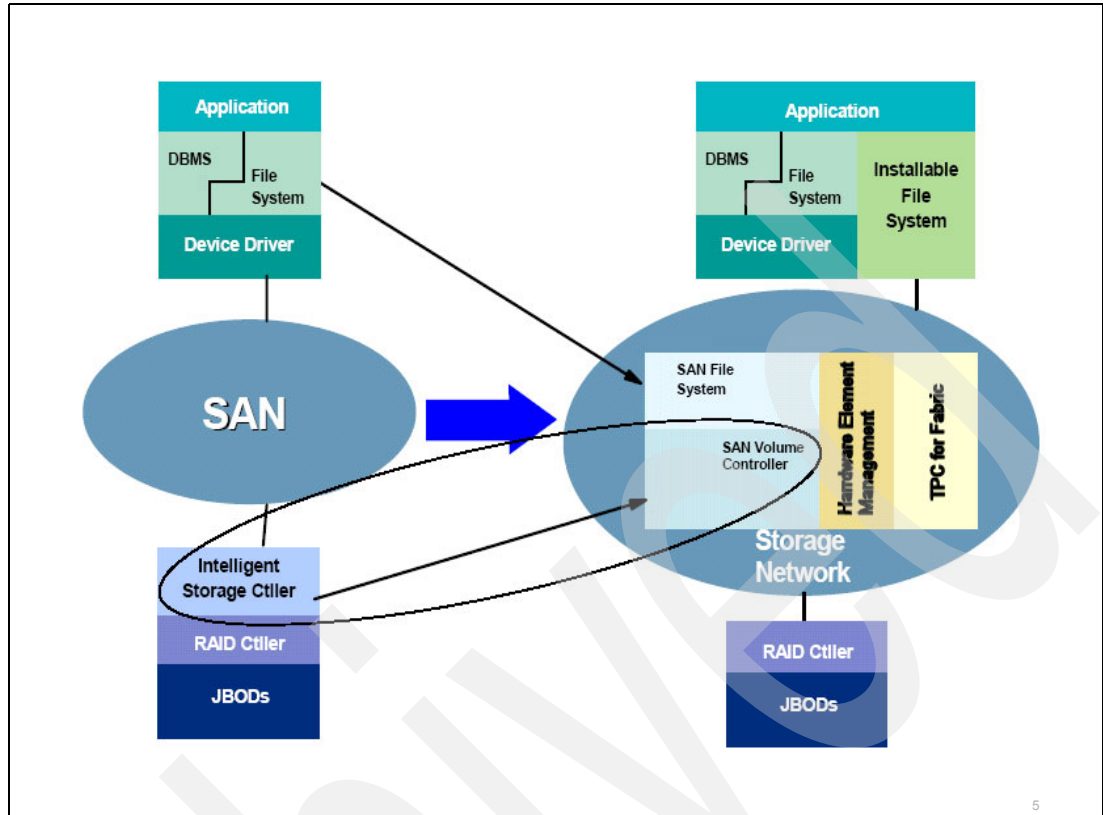


Figure 5-14 IBM plan for block virtualization

We have chosen to develop our block virtualization product in the storage network using an in-band approach. We think that this approach will provide a superior solution for customers needing the benefits of block virtualization. Our solutions are designed to be modular, redundant, and scalable. We have based our solutions on clustered IBM @server® xSeries® servers, which support high availability and performance that is horizontally scalable. The system allows for the addition of nodes (engines) non-disruptively to provide enterprise-class scalability. Our long history of storage controller development has enabled us to develop systems where, in the rare case that a component failure occurs, the storage virtualization device can continue to operate without disruption. The SAN Volume Controller and the SAN Volume Controller Storage Software for Cisco MDS 9000 are the IBM TotalStorage solutions for block virtualization.

5.4.2 SAN Volume Controller characteristics

The IBM TotalStorage SAN Volume Controller is an in-band implementation that minimizes the dependency on unique hardware and software, decoupling the storage functions expected in a SAN environment from the storage subsystems and managing storage resources.

In SANs today, shown in Figure 5-15, we show servers mapped to specific devices, called *physical mapping*.

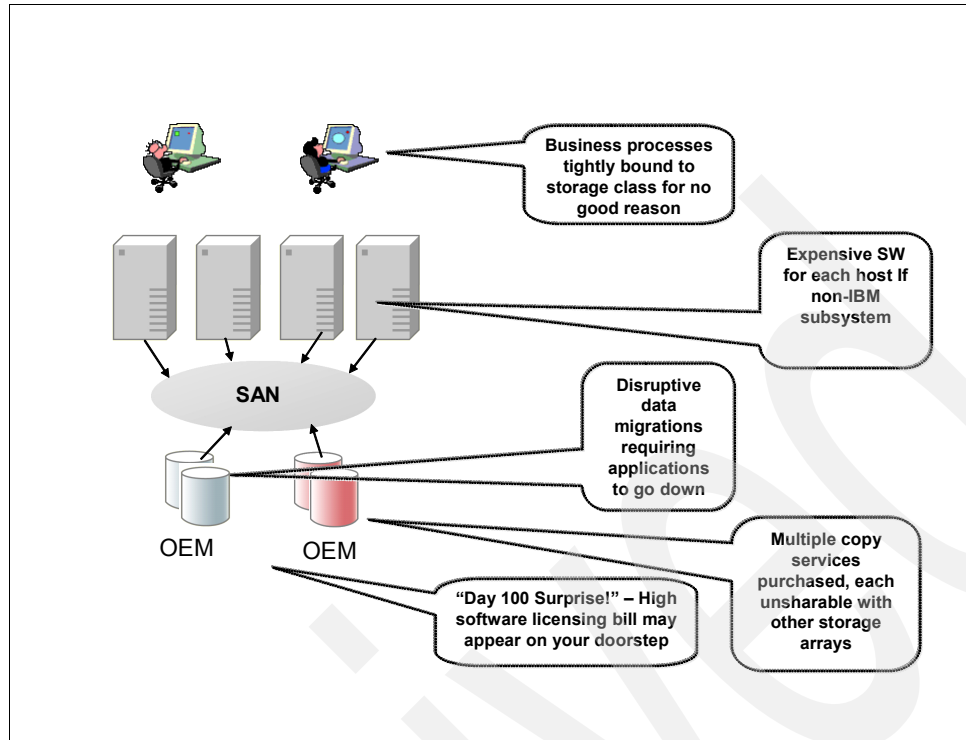


Figure 5-15 SAN today

With the SAN Volume Controller, shown in Figure 5-16, servers are mapped to virtual disks, therefore, creating a virtualization layer called *logical mapping*.

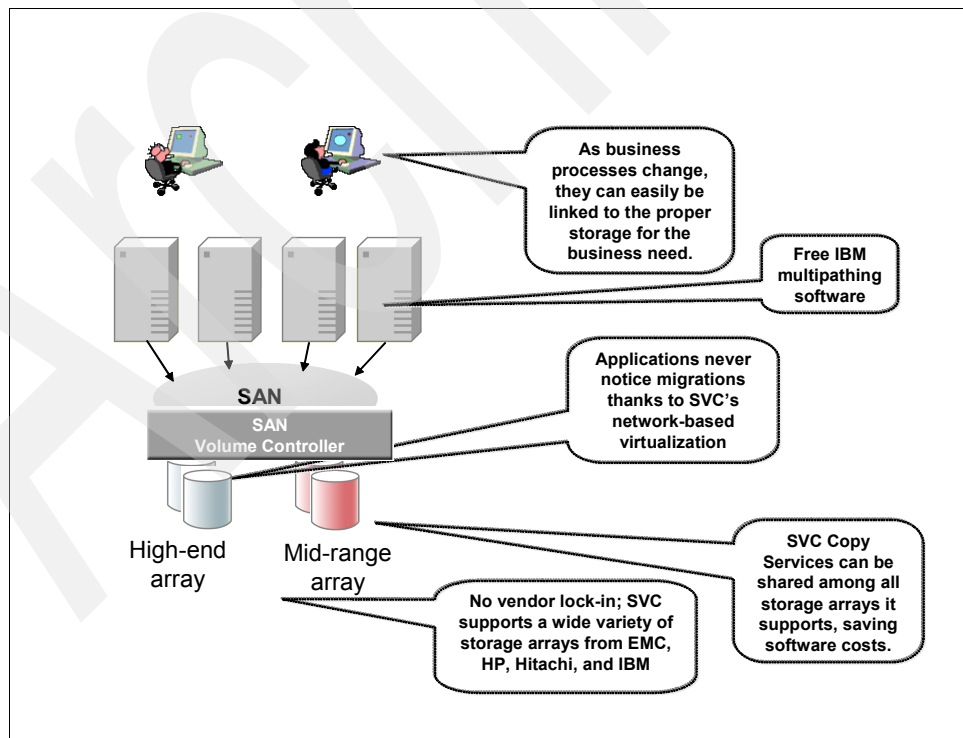


Figure 5-16 Block level virtualization

The SAN Volume Controller implementation creates two zones, the host zone and the disk zone, to ensure disk storage devices are protected from being accessed by the application servers, and to ensure the SAN Volume Controller handles I/O management.

The IBM SAN Volume Controller is designed to provide a redundant, modular, scalable, complete solution, as shown in Figure 5-17.

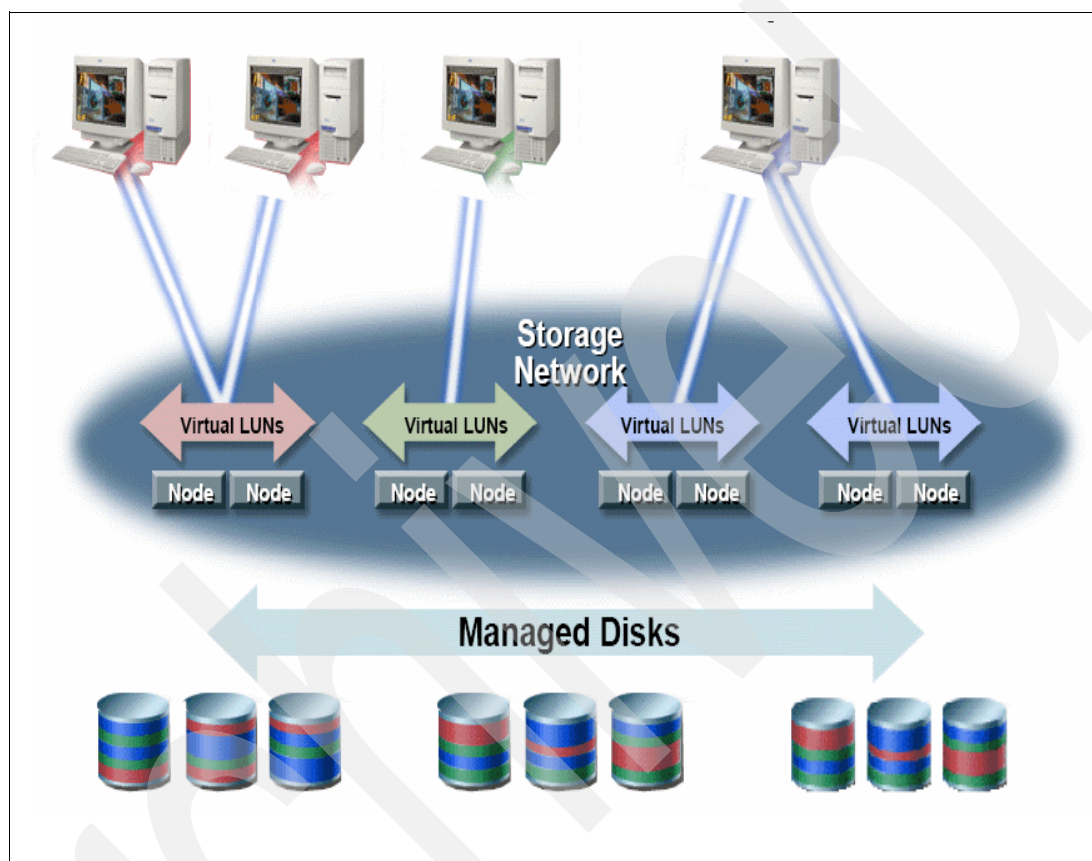


Figure 5-17 IBM SAN Volume Controller

Each SAN Volume Controller consists of one or more pairs of engines, where each pair operates as a single controller with failover redundancy. Each node is an xSeries eServer with a large read/write cache mirrored across the pair. Virtual volumes are shared between a pair of nodes. The pool of managed disks is controlled by a cluster of paired nodes.

The SAN Volume Controller is designed to provide complete copy services (see figure 5-16) for data migration and business continuity. Since these copy services operate on the virtual volumes, dramatically simpler replication configurations can be created using the SAN Volume Controller, rather than replicating each physical volume in the managed storage pool. These copy services include point-in-time FlashCopy, and Metro Mirror (formerly Synchronous Peer-to-Peer Remote Copy PPRC). This support includes:

- ▶ FlashCopy from one box to another box, even across different vendor devices (by contrast, ESS and DS4000 require FlashCopy source and destination to be in the same disk subsystem. Most other disk subsystems have similar limitations.).
- ▶ Metro Mirror from one SVC cluster to another, or in-house on a single SVC cluster. Source and destination can be different vendor devices.
- ▶ Copy services are not required to be licensed on the underlying managed devices (ESS, DS4000, and so on).

- Destination volumes can be lower-cost, saving the customer money in deploying a two-site business continuity solution.

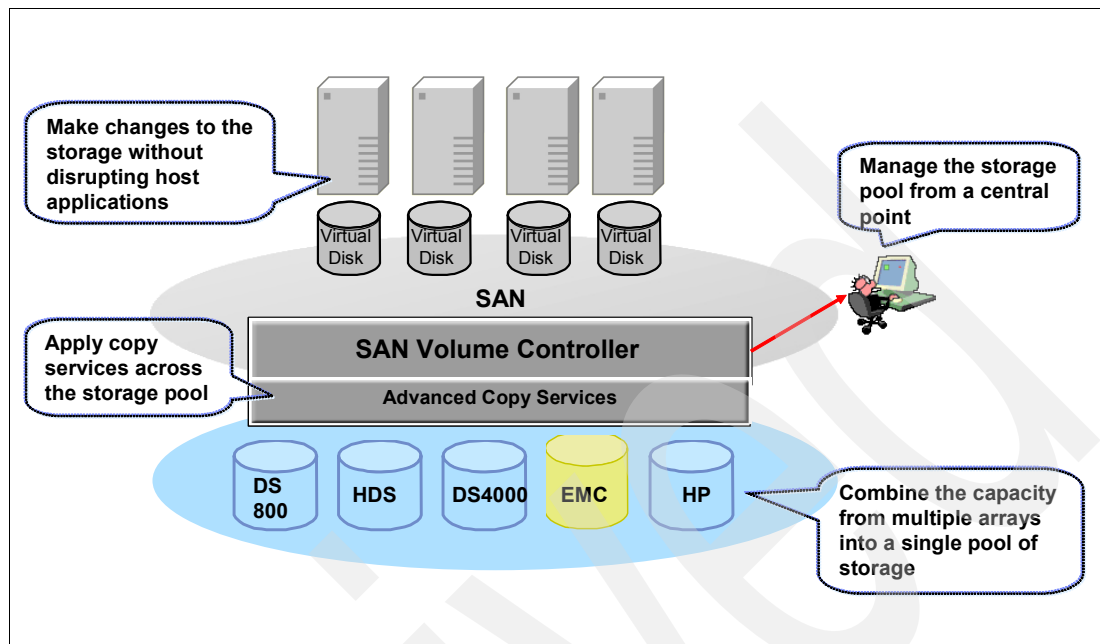


Figure 5-18 Infrastructure simplification with SVC

The IBM SAN Volume Controller improves storage administrator productivity, provides a common base for advanced functions, and provides more efficient use of storage.

The SAN Volume Controller simplifies the infrastructure by helping you:

- Reduce IT administration costs by improving IT administration productivity with a single point of control, administration, planning, and security.
- Increase the managed data to IT administrator ratio.
- Improve flexibility by providing system availability and data protection.
- Improve overall capacity utilization across all application servers, regardless of operating system or platform type.
- Reutilize older storage enclosures.
- Provide a single consistent set of high value functions across all OS or platforms usable on all SAN-attached storage (for example, copy services, mirroring, remote mirroring, and backup/restore).
- Provide scalable performance with the addition of relatively low cost components (HBAs, Intel Server memory).

The SAN Volume Controller consists of software and hardware components delivered as a packaged appliance solution in a variety of form factors. The IBM SAN Volume Controller solution can be preconfigured to the customer's specifications, and will be installed by an IBM customer engineer. The architecture of the SAN Volume Controller is designed to bring enterprise class reliability and performance to open-systems environments. It features hardware redundancy and elements of the IBM advanced autonomic computing technologies. The intent is to help minimize downtime and improve availability while performing remote mirroring; point-in-time copies; backup and restore; maintenance functions; and performance, capacity, and connectivity upgrades. IBM has designed and tested the SAN Volume Controller for easy integration into existing environments, including

heterogeneous hardware and operating systems. It is interoperable with a wide range of servers running Linux®, UNIX, and Microsoft® Windows operating systems, whether from IBM or other vendors.

The SAN Volume Controller provides centralized management through a single interface to support easier storage allocation and address application demands. This flexibility helps provide the benefit of better storage utilization by reducing or eliminating the problem of unused storage found in direct attached storage (DAS) implementations, and to reduce required administrative time and resources. These are key factors in realizing a lower total cost of ownership (TCO).

Customer data can be migrated from existing storage environments into an IBM SAN Volume Controller environment, and thereafter grown into a SAN Volume Controller-managed environment, providing protection of your investment and lower storage TCO.

In summary, the value that the IBM TotalStorage SAN Volume Controller solution provides is increased system availability, greater storage capacity utilization, improved protection capability, and enhanced scalability. See Figure 5-19.

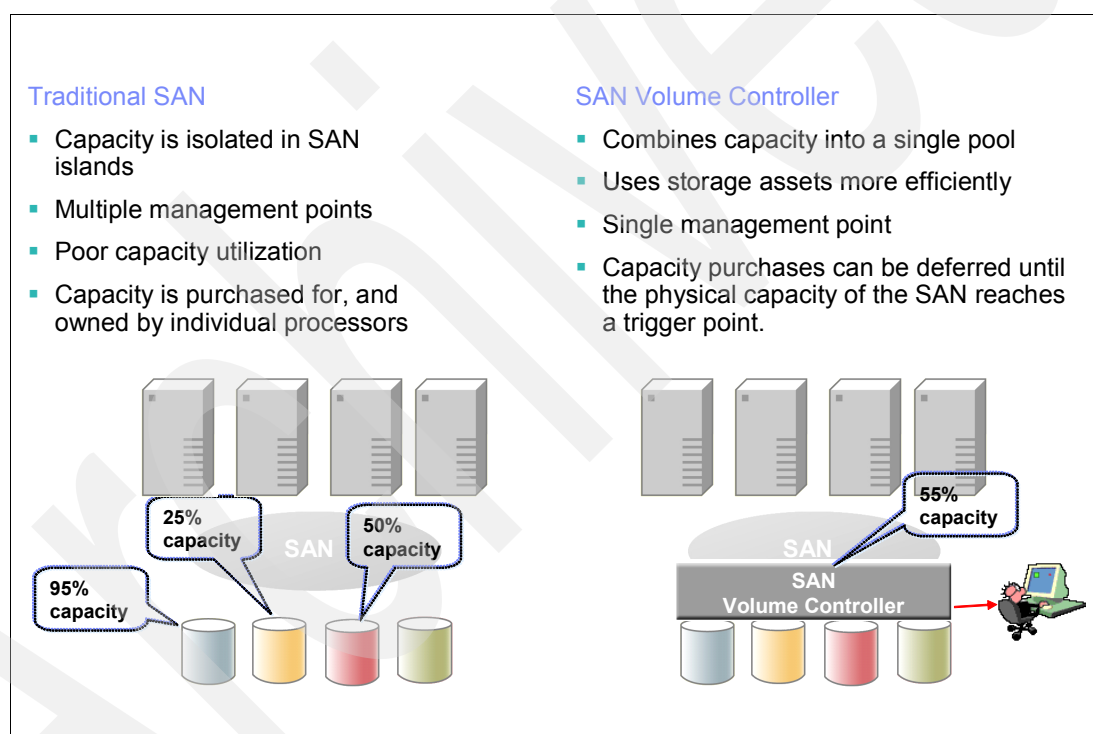


Figure 5-19 Capacity utilization

The SAN Volume Controller supports all the major operating systems, SAN switches, and storage subsystems, including IBM DS4000 (FASTT), IBM DS6000, IBM DS8000, IBM ESS, and other vendors' storage.

5.4.3 SAN Volume Controller interoperability

The IBM TotalStorage SAN Volume Controller Software V3.1 supported operating systems and storage devices are show in Figure 5-20 on page 95.

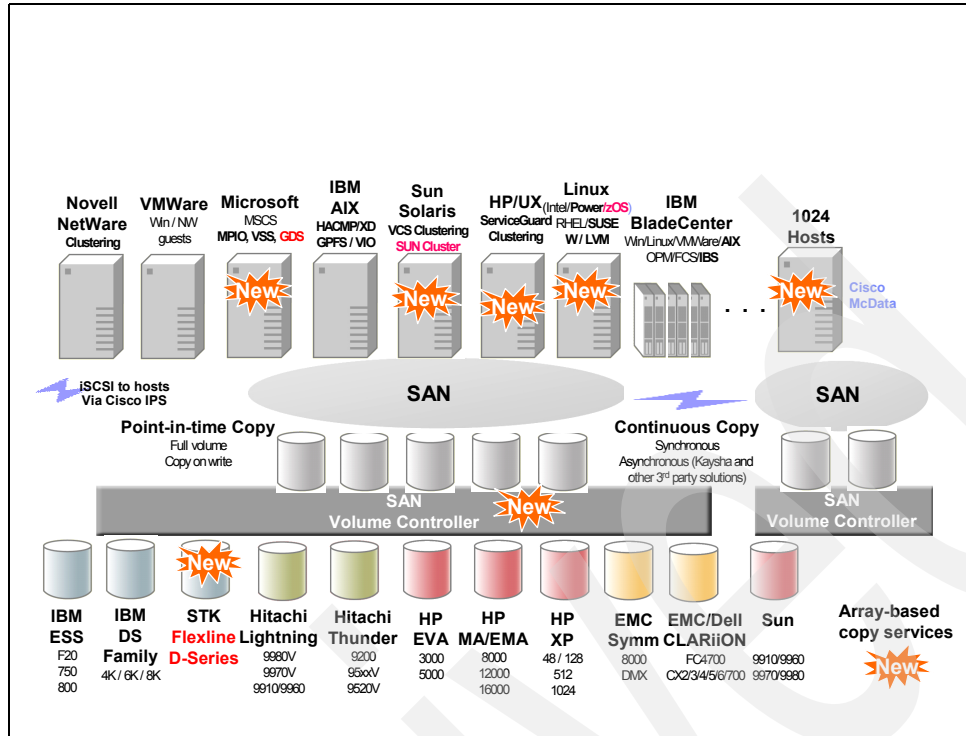


Figure 5-20 IBM TotalStorage SAN Volume Controller supported environments

For a complete and recent list, go to the following Web site:

<http://www-1.ibm.com/servers/storage/software/virtualization/svc/interop.html>

SVC Software for Cisco MDS 9000 characteristics

The IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 uses in-band virtualization. Single storage virtualization engines, which are known as nodes, are combined to create clusters.

A cluster contains four nodes. A node is a single engine. Each Caching Services Module (CSM) supports two engines, or nodes. Nodes within the cluster are grouped in pairs known as an I/O group. Nodes within an I/O group back up one another. Data written to the nodes is duplicated across caches in both nodes. Virtual disks are shared between nodes in an I/O group. To eliminate any single point of failure, nodes in an I/O group must be on a separate CSM. In Figure 5-21, we show what our cluster will look like when it is created.

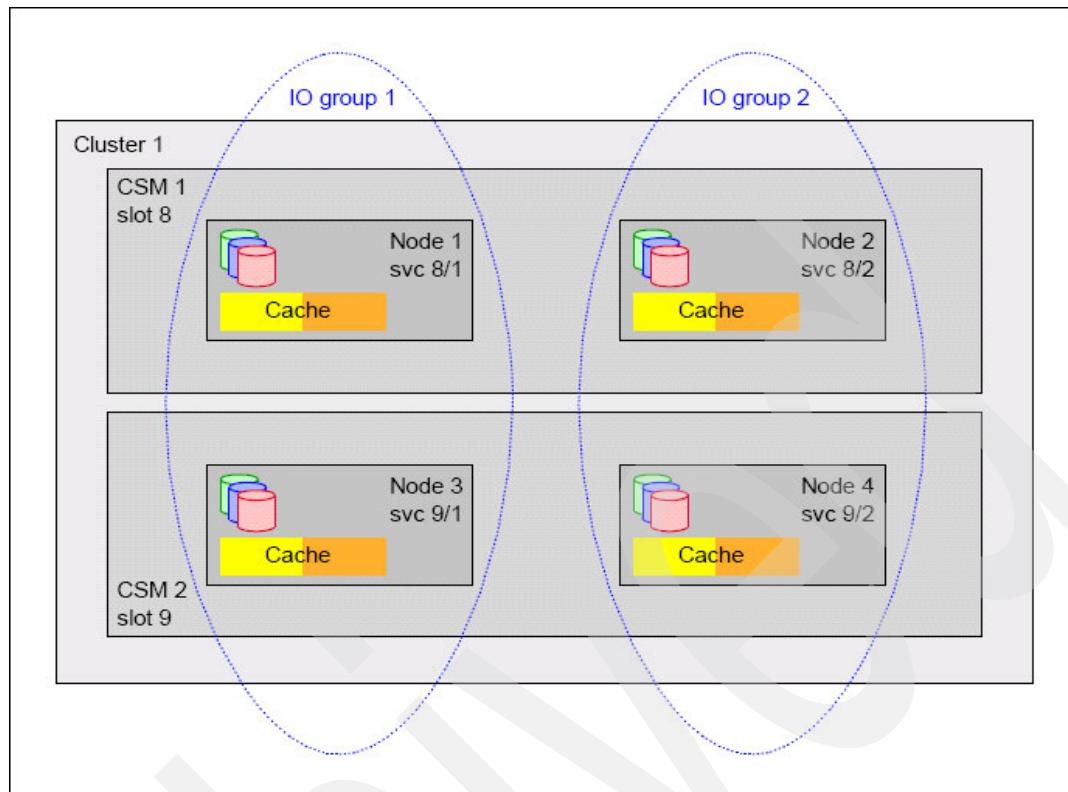


Figure 5-21 Cluster creation

The SAN Volume Controller Storage Software for Cisco MDS 9000 I/O groups see the storage presented to the SAN by the back-end controllers as a number of disks, known as *managed disks*. The application servers do not see these managed disks. This is achieved by zoning or by using Virtual SANs (VSANs). Instead, they see a number of logical disks, known as *virtual disks*, that are presented to the SAN by the SAN Volume Controller Storage Software for Cisco MDS 9000. Each node must only be in one I/O group and provide access to the virtual disks in the I/O group. The SAN Volume Controller Storage Software for Cisco MDS 9000 helps to provide continuous operations and can also optimize the data path to ensure performance levels are maintained.

The solution offers the following benefits and advantages:

- ▶ Reduces complexity
- ▶ Lowers the cost of managing SAN-based storage
- ▶ Creates a single pool of storage from disparate storage devices to increase capacity utilization
- ▶ Implements a cache-based, clustered architecture to provide a highly available solution
- ▶ Provides the scalability and performance required in today's demanding storage environments

In summary, the value that the SAN Volume Controller Storage Software for Cisco MDS 9000 provides is increased system availability, greater storage capacity utilization, improved protection capability, and enhanced scalability similar to that provided by the SAN Volume Controller.

5.5 IBM TotalStorage SAN File System

IBM TotalStorage SAN File System is a multiplatform, robust, scalable, and highly available file system, and is a storage management solution that works with Storage Area Networks (SANs). It uses SAN technology, which allows an enterprise to connect a large number of computers and share a large number of storage devices via a high-performance network.

With SAN File System, clients using different operating systems can share data directly from large, high-performance, high-function storage systems, such as IBM TotalStorage SAN Volume Controller (SVC), IBM TotalStorage Enterprise Storage Server (ESS), IBM TotalStorage DS Family Servers, (DS4000, DS6000, and DS8000), as well as non-IBM storage devices. The SAN File System is built on a Fibre Channel network, and is designed to provide superior I/O performance for data sharing among heterogeneous computers.

SAN File System differs from conventional distributed file systems in that it uses a data-access model that separates file meta-data (information about the files, such as owner, permissions, and the physical file location) from actual file data (contents of the files). The meta-data is provided to clients by meta-data servers, the clients communicate with the meta-data servers only to get the information they need to locate and access the files. Once they have this information, the SAN File System clients access data directly from storage devices via the clients' own direct connection to the SAN. Direct data access eliminates server bottlenecks and provides the performance necessary for data-intensive applications.

SAN File System presents a single, global namespace to clients where they can create and share data, using uniform file names from any client or application. See Figure 5-22. Furthermore, data consistency and integrity are maintained through SAN File System's management of distributed locks and the use of leases. This shared name space enables files to be accessed directly over the SAN at full speed, but also shared between AIX, Linux, Solaris™ and Windows operating system applications. Rather than managing hundreds of individual file systems, or dozens of NAS boxes, you now only have to manage a single name space.

SAN File System also provides automatic file placement through the use of policies (Figure 5-22 on page 98) and rules. Based on rules specified in a centrally-defined and managed policy, SAN File System automatically stores data on devices in storage pools that are specifically created to provide the capabilities and performance appropriate for how the data is accessed and used. The SAN File System is a unique environment that allows files to be managed based on policies on different tiers of storage.

Advantages:

- Single global view of file system
- Metadata Server processes only metadata operations
- Linear scalability of global file system by adding Metadata Server nodes
- Advanced, centralized, file-granular policy-based mgmt.
- Non-disruptive management of physical assets under storage pool

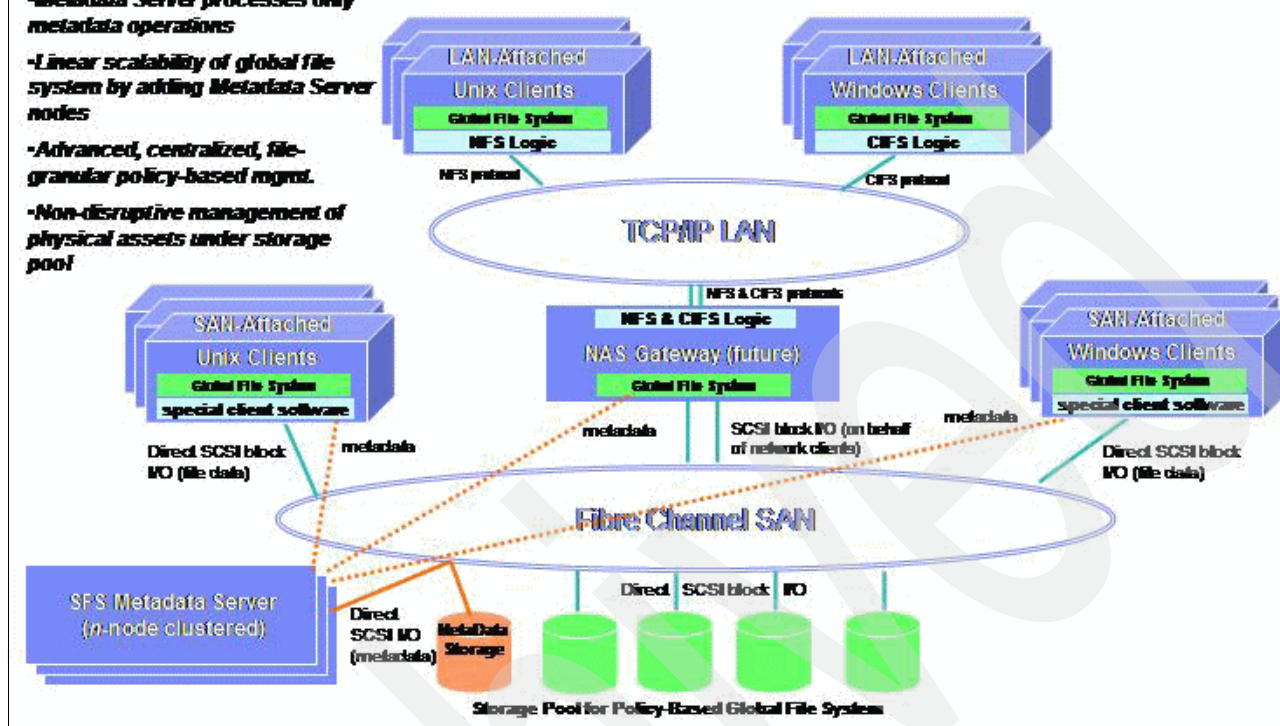


Figure 5-22 SAN File System

5.5.1 File virtualization

While block virtualization provides flexibility when working with blocks of data stored in volumes, file virtualization provides flexibility when accessing and managing data stored in files. The IBM TotalStorage SAN File System product provides advanced file virtualization.

In Figure 5-23, we show the SNIA approach to file aggregation.

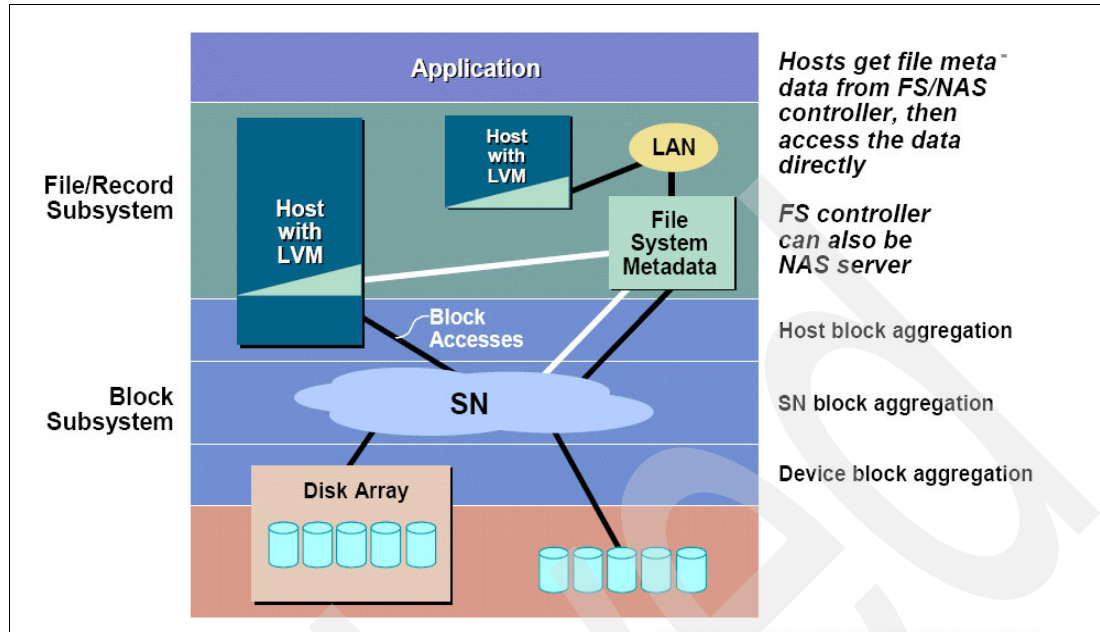


Figure 5-23 SNIA file aggregation model

In the SNIA model, hosts get file meta-data from file systems or NAS controllers, then access the data directly. With NAS, each device is a self-contained file system island. The IBM approach to file virtualization is to provide a meta-data controller in the storage network providing a single global namespace for accessing data on storage devices.

5.5.2 SAN File System characteristics

The IBM TotalStorage SAN File System architecture makes it possible to bring the benefits of the existing mainframe system-managed storage (SMS) to the SAN environment. Features such as policy-based allocation, volume management, and file management have long been available on IBM mainframe systems. However, the infrastructure for such centralized, automated management has been lacking in the open systems world of Linux, Windows, and UNIX. On conventional systems, storage management is platform dependent. The SAN File System provides a single, centralized point of control to better manage files and data, and is platform independent. Centralized file and data management dramatically simplifies storage administration and lowers TCO.

We show the IBM plan for file virtualization, based on the SAN File System architecture, in Figure 5-24, as:

- ▶ Common file system in a SAN or even an enterprise
- ▶ Meta-data controller
- ▶ Direct I/O from application servers to either virtual volumes or real volumes
- ▶ SAN/NAS convergence

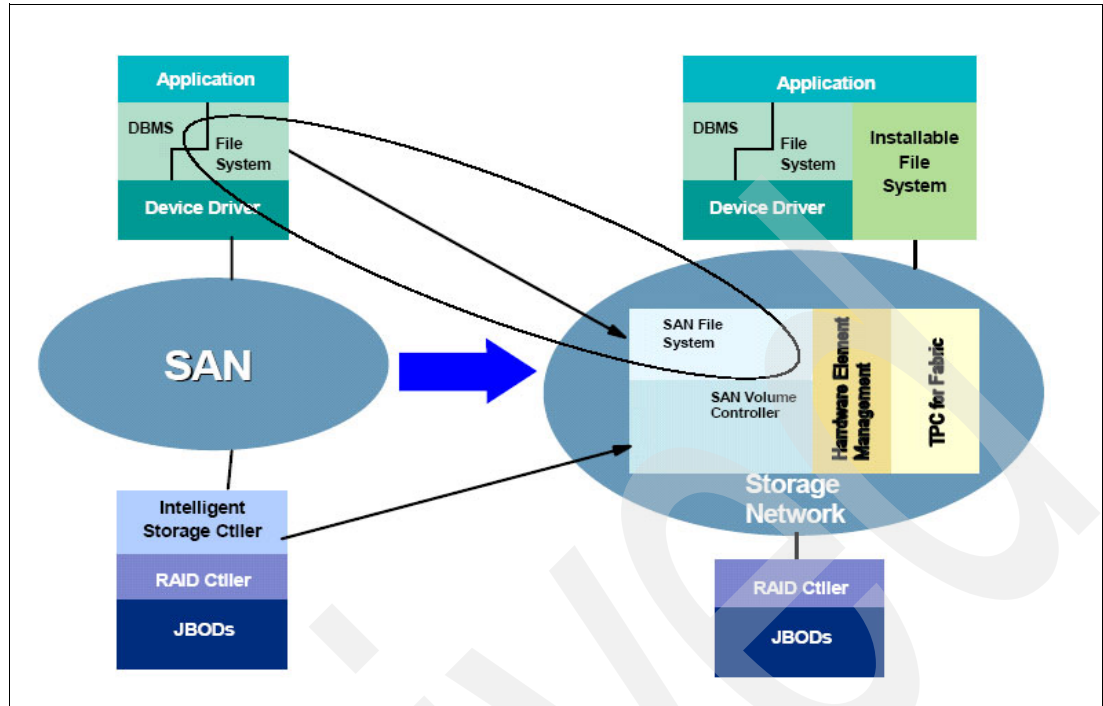


Figure 5-24 IBM plan for file virtualization

The SAN File System is a common file system specifically designed for storage networks. By managing file details (via the meta-data controller) on the storage network instead of in individual servers, the SAN File System design moves the file system intelligence into the storage network where it can be available to all application servers. Doing so provides immediate benefits: a single namespace and a single point of management. This eliminates the need to manage files on a server by server basis.

The SAN File System automates routine and error-prone tasks such as file placement and handles out of space conditions. SAN File System allows true heterogeneous file sharing where the reader and writer of the exact same data can run different operating systems.

5.5.3 SAN File System architecture

SAN File System meta-data controller is designed as a cluster of servers attached to a SAN and a small software addition to the application servers. Other than installing the SAN File System client on the application servers, no changes are required to applications to use the SAN File System.

In Figure 5-25, we show a pictorial representation of the SAN File System environment.

Advantages:

- Single global view of file system
- Metadata Server processes only metadata operations
- Linear scalability of global file system by adding Metadata Server nodes
- Advanced, centralized, file-granular policy-based mgmt.
- Non-disruptive management of physical assets under storage pool




Figure 5-25 IBM TotalStorage SAN File System

Application servers that request a file, obtain information about the file (the meta-data) from the SAN File System meta-data controller that manages file locks and all other file information. SAN File System then provides that information to the application server, which then accesses the blocks comprising that file directly through the SAN. By caching the meta-data in the client and providing direct access from the application server to the underlying storage, the SAN File System provides local file system performance over the SAN.

The SAN File System consists of a small module of enablement code that runs on application servers and a meta-data controller based on clustered IBM xSeries servers for redundancy and fault tolerance.

The features of the SAN File System work together to provide a variety of benefits to customers. One of the major benefits is a single image or global namespace. This function shields the end user from storage network complexity and dramatically reduces administrative workload.

Since the SAN File System is designed to be implemented on a variety of operating systems from Windows to various flavors of Linux and UNIX, it allows all of these operating systems to share files. A file created in Windows is as accessible from a Windows client as it is from Solaris, AIX, or any other supported platform, and vice versa. However, keep in mind that an application is still required to be able to read that file, however accessible it is.

The supported operating systems and storage devices by the IBM TotalStorage SAN File System, are shown in Figure 5-26.

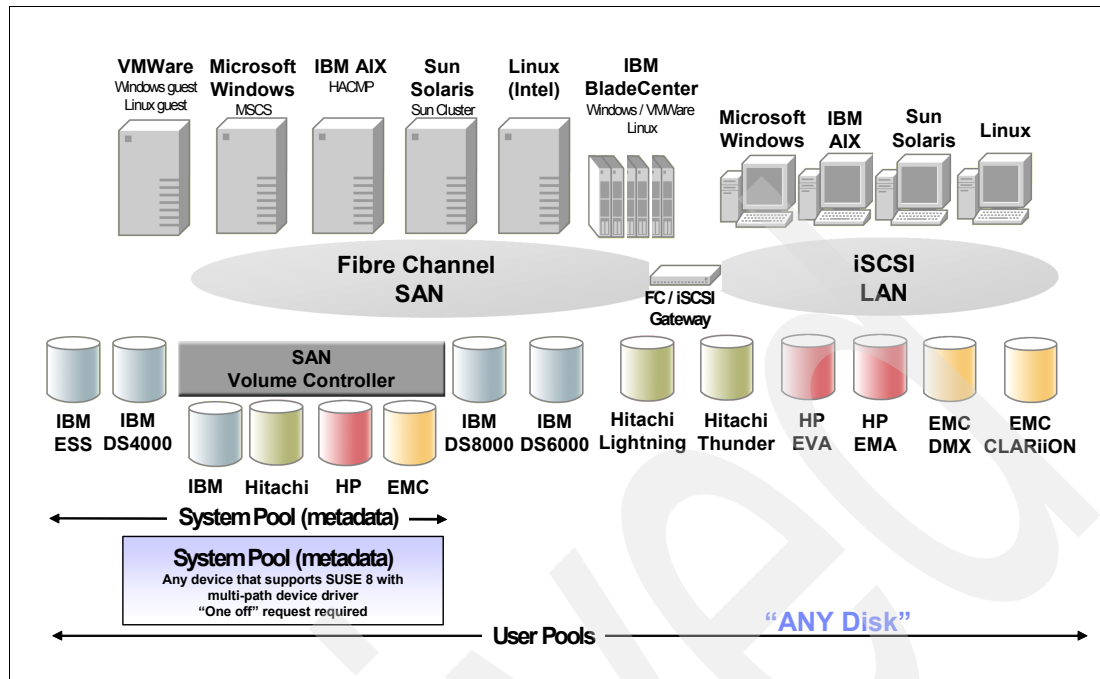


Figure 5-26 Supported operating systems and storage devices

For a complete and recent list of supported systems, go to Web site:

<http://www-1.ibm.com/servers/storage/software/virtualization/sfs/interop.html>

Since the SAN File System has a complete understanding of all files on the SAN, including the essential meta-data to make important decisions, it is a logical point to manage the storage on the network through policy-based controls. For example, the SAN File System can decide where to place each file based on user-defined criteria, such as file type, using policy-based automation. Setting these policies can help administrators gain more time by moving aged files, user files, and other files on a given schedule or based on specific criteria.

SAN File System provides the ability to group storage devices according to their characteristics, such as latency and throughput. These groupings, called *storage pools*, allow administrators to manage data according to the characteristics that matter to them. For example, an administrator can define a storage pool for mission-critical applications using highly reliable storage arrays that are backed up nightly and have full disaster recovery capabilities. The administrator can also define a storage pool for less critical applications based on Jabots with weekly tape backups and minimal disaster recovery capabilities.

Because the SAN File System meta-data is separate from the application data, files can be manipulated while remaining active. For example, files processed by a mission critical application can be non-disruptively moved within or across storage pools without stopping the application. Data migration from one storage system to another can be handled non-disruptively by having the SAN File System move the pools (data) to new physical disks, then disconnecting the old disks, all done without quiescing applications.

The SAN File System approach allows users and administrators to access, save, share, and centrally manage files on storage networks. It can leverage policies to direct files into specific storage pools with different class of service characteristics. For example, these may include mirrored pools for disaster recovery, striped pools for performance, or a pool of slower, low cost drives. Storage can be added to these pools dynamically and be immediately available for use by applications. When files are removed from service, the SAN File System can

automatically reallocate the space without disruption. If a LUN is removed from the SAN File System control, the data on that LUN is automatically moved. SAN File System helps reduce TCO by simplifying the management of files in a storage network. No application changes are required to realize these benefits.

SAN File System offers a logical extension to current NAS and SAN environments. Our approach to NAS for customers with SANs is to add NAS capabilities to the SAN File System, therefore, allowing storage administrators to manage the NAS file data with the same tools they use for their application servers and SAN File System. Data does not have to be duplicated across multiple NAS devices. This approach of SAN/NAS convergence can lower TCO in these environments.

In summary, the SAN File System is a common SAN-wide file system that permits centralization of management and improved storage utilization at the file level. SAN File System is delivered in a highly available configuration based on IBM eServer xSeries with clustering for the meta-data controllers, providing redundancy and fault tolerance. SAN File System is designed to provide policy-based storage automation capabilities for provisioning and data placement, non-disruptive data migration, and a single point of management for files on a storage network. The use of the SAN File System can greatly help simplify the management of files on SANs and result in a significant reduction in TCO.

The IBM TotalStorage SAN File System is designed on industry standards, so it can:

- ▶ Allow data sharing and collaboration across servers over the SAN with high performance and full file locking support, using a single global namespace for the data.
- ▶ Provide more effective storage utilization by reducing the amount of duplicate data and by sharing free and temporary space across servers.
- ▶ Improve productivity and reduce the “pain” for IT storage and server management staff by centralizing and simplifying management through policy-based storage management automation, therefore, significantly lowering the cost of storage management.
- ▶ Facilitate application server and storage consolidation across the enterprise to scale the infrastructure for storage and data on demand.
- ▶ Simplify and lower the cost of data backups through built-in, file-based FlashCopy image function.
- ▶ Eliminate data migration during application server consolidation, and also reduce application downtime and failover costs.

5.6 IBM System Storage N Series

All of the IBM N Series models are powered by a specialized OS – Data ONTAP which is designed to offer the best networked storage performance and data availability. The N Series reduces complexity by using an “appliance” philosophy; keep it simple by doing one thing but do it better than anyone else. Data ONTAP is specifically designed for storage and is simpler to use than general purpose OSs such as Windows or Unix. The flexibility in N Series hardware provides a “future proof” investment for enterprises – no matter which file sharing protocols (CIFS, NFS, FTP, or HTTP) and storage connections (FCP or iSCSI) are used, N Series hardware can be plugged into any of these environments.

Write AnyWhere File Layout

Data ONTAP incorporates the highly optimized WAFL (Write Anywhere File Layout) file system and storage virtualization layer that is designed to minimize disk head movement for efficient reads and writes. With every write to the file system, the N Series stripes the data to

all the disks in a volume simultaneously. This results in very high performance and automatically balances the I/O load across all the available disks in a volume.

Utilize existing software

The N Series allows you to take advantage of software that exists today.

Backup software

The N Series storage system works with all of the top backup software vendors including IBM Tivoli Storage Manager, Veritas, Legato, CA, Commvault, Connected, Atempo, and Bakbone.

Storage management

The N Series storage system has an open access policy for storage and system management. The Data ONTAP suite of APIs are published so that vendors can integrate control and operational functions into their products. Adopters of the Manage ONTAP APIs include leading storage management vendors such as IBM TotalStorage Productivity Center for Data, ApplIQ, BMC Software, Computer Associates, CreekPath, Fujitsu Softek, NuView, NTP, Storability, Tek-Tools, TeraCloud, and Veritas/Precise-Wquinn. The N Series includes monitoring plug-ins that work with widely used management software, such as HP Openview and IBM Tivoli Enterprise.

Anti-Virus software

The N Series storage system also supports Symantec, CA, Trend Micro, Network Associates, and Sophos.

5.6.1 Windows environments

Many customers have deployed Windows file servers throughout their organization for home directories, departmental shares, and Web servers. Over time companies are finding themselves managing 100s or even 1,000s of file servers, most of which are severely underutilized. The N Series provides a Windows Data Consolidation solution that is quick, efficient, and a cost-effective solution to these problems.

Why is the N Series good for consolidating Windows file serving?

- ▶ Radically reduce your Windows file server count.
 - N Series storage system consolidated file serving and home directory storage reduces costs. You get more efficient storage utilization, fewer servers to manage, and fewer software upgrades. And, you can expand storage quickly and easily, when you need it, at a lower incremental cost. The administration load is high for Windows servers with continuous streams of patches and OS updates. See Figure 5-27.
 - Create a shared storage pool and consolidate processing power.
 - Reduce facilities costs by reducing the multiple underutilized servers taking up valuable data center space.
- ▶ Increase your organization's productivity.
 - N Series storage system enables more efficient data sharing throughout the enterprise across any mix of computing platforms and protocols. With a simplified architecture and fewer servers to manage, your administrators can focus on tasks that add value for users.
- ▶ Get data faster.
 - Optimized for serving Windows data, N Series solutions deliver faster response times and greater throughput than the traditional Windows file server infrastructure.

- ▶ Integrate into the Windows environment.
 - Active Directory and Group Policy support.
 - Integrate into the Windows 2000 domain just like a Windows server™, and manage just like any other Windows server. Shares can be defined and deleted, offered up to users, and the administrator can use the same Windows tools they use today with Windows servers.
- ▶ Kerberos and LDAP support.
- ▶ Leverage existing Windows administration tools.
- ▶ Integrated with Volume Shadow Copy Services (VSS).

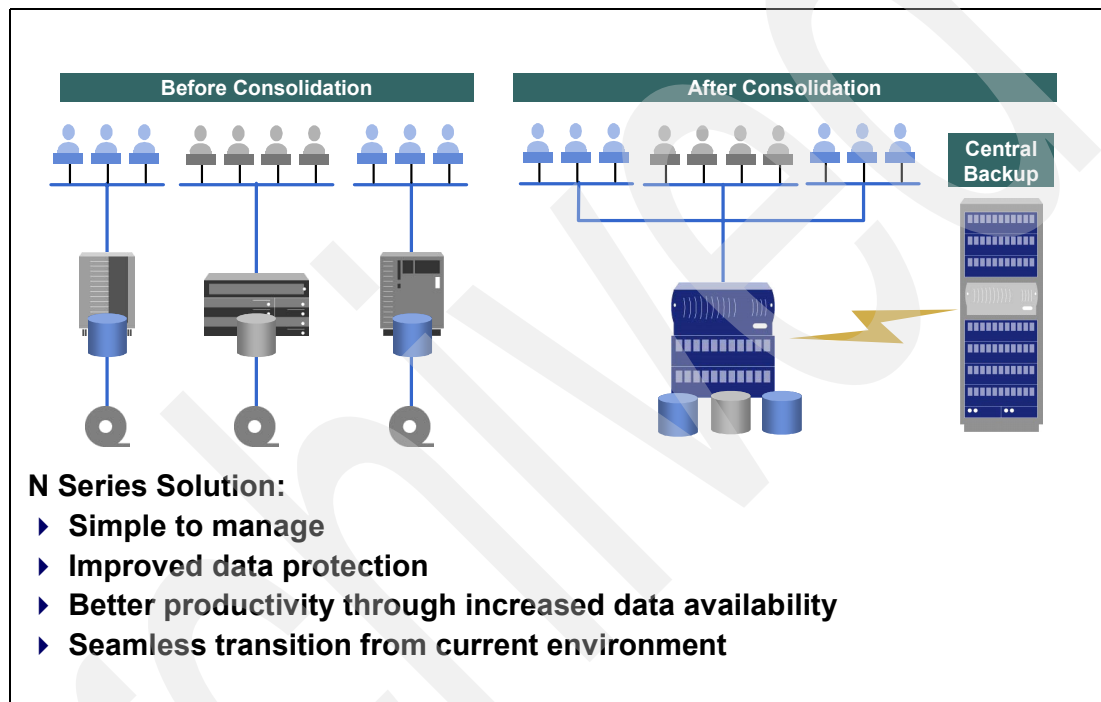


Figure 5-27 Consolidation Architecture

5.6.2 Consolidation for all Platforms

In Windows environments, we discussed how the N Series is a perfect partner for consolidating your Windows environment. This consolidation capability is not limited to Windows, but works just as well in a UNIX or LINUX environment. Customers who have several older servers, such as Windows, Unix, or Linux coming to end-of-life, who need to start replacing these servers, consolidating many small servers into a smaller number of larger capacity servers, find the N Series is a good vehicle to do so.

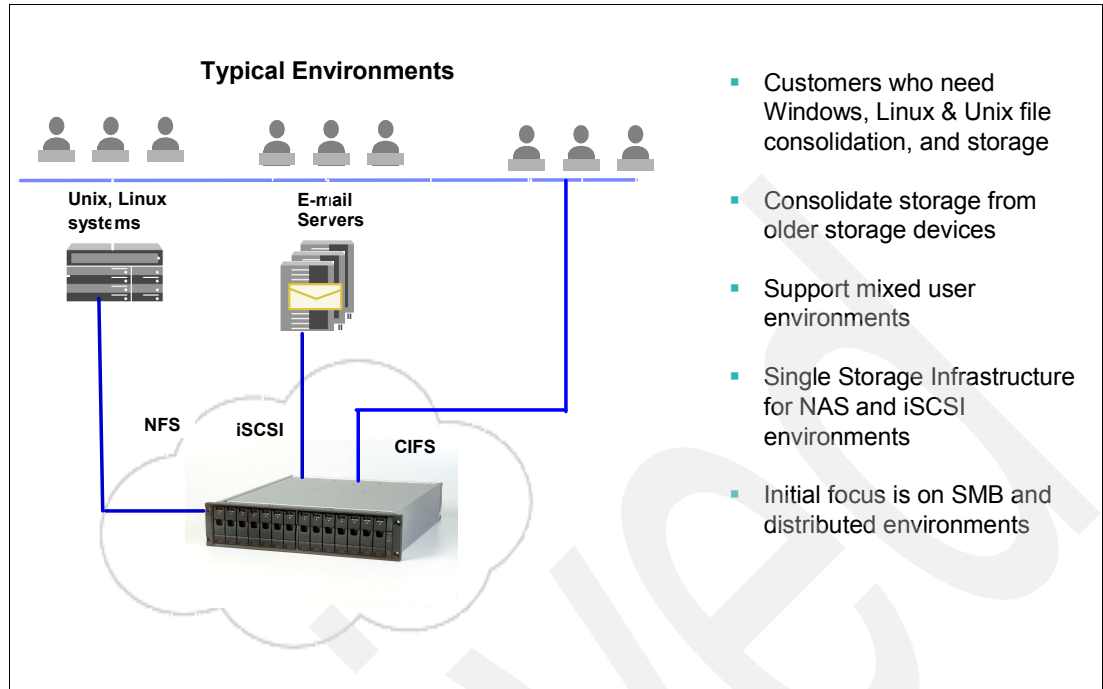


Figure 5-28 Consolidation

5.6.3 High Availability

For those customers unable to upgrade to a SAN-based storage network, the N Series offers copy services-like capabilities such as the SVC, DS6000, and DS8000. In the N Series this out of band copy is called SnapMirror. It provides a continuous mirror of the primary site to a secondary site and N Series System Storage.

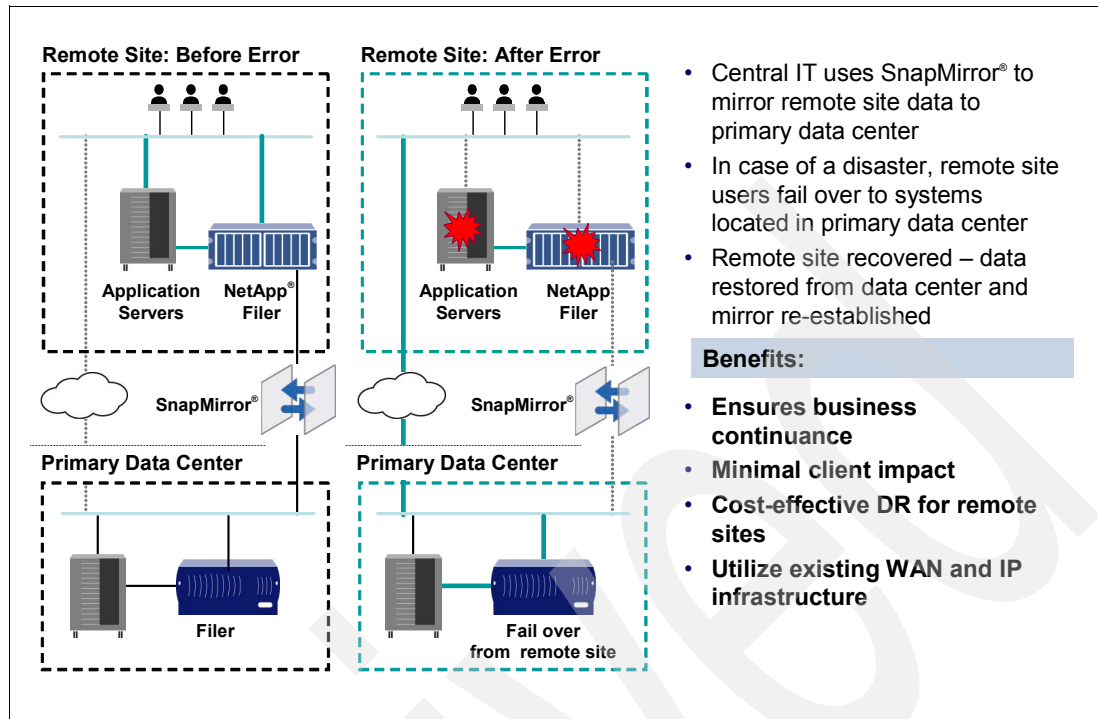


Figure 5-29 High availability with the N Series

5.7 Management and productivity

In 1999, the Storage Networking Industry Association (SNIA) and Distributed Management Task Force (DMTF) introduced open standards for managing storage devices. These standards use a common protocol called the Common Information Model (CIM) to enable interoperability. The Web-based version of CIM (WBEM) uses XML to define CIM objects and process transactions within sessions. This standard proposes a CIM Object Manager (CIMOM) to manage CIM objects and interactions. CIM is used to define objects and their interactions. Management applications then use the CIM object model and XML over HTTP to provide for the management of storage devices. This enables central management through the use of open standards.

IBM is committed to implementing the SNIA standards-based model to allow IBM products, such as IBM TotalStorage Productivity Center for Data, and other vendor management applications, to more easily administer, monitor, and control IBM storage devices. Standard CIM support is included with the IBM TotalStorage SAN Volume Controller, IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000, and the IBM TotalStorage SAN File System products, as well as IBM TotalStorage DS4000 series, IBM TotalStorage DS6000 series, IBM TotalStorage DS8000 series, IBM TotalStorage ESS, and storage from other software vendors.

Following these standards, ensures that you would be able manage any storage device regardless of the manufacturer, from a single management application. You do not have to setup as many management environments no matter how many devices you have. This not just simplifies your infrastructure, but saves purchasing costs, management costs, and administrative costs. For more detailed information about saving costs, go to Chapter 2, "Total Cost of Ownership of storage infrastructure" on page 17. IBM prefers to use open industry standards management for its management software rather than use vendors' proprietary

interfaces, and IBM encourages other storage vendors to release open Storage Management Interface Specification (SMI-S) compliant interfaces as soon as possible.

5.7.1 Storage Management Initiative

SNIA is using its Storage Management Initiative (SMI) to create and promote adoption of a highly functional interoperable management interface for multi-vendor storage networking products. The SNIA strategic imperative is to have all storage managed by the SMI interface by 2005. The adoption of this interface allows the development focus to switch to the development of value-add functionality. IBM is one of the industry vendors promoting the drive toward this vendor-neutral approach to SAN management.

The Storage Management Interface Specification (SMI-S) for SAN-based storage management provides basic device management, support for copy services, and storage virtualization. As defined by the standard, the CIM services are registered in a directory to make them available to device management applications and subsystems.

SNIA uses the xmlCIM protocol to describe storage management objects and their behavior. CIM allows management applications to communicate with devices using object messaging encoded in xmlCIM.

For more information on SMI-S, go to the following Web site:

<http://www.snia.org>

5.7.2 Open storage management with CIM

SAN management involves configuration, provisioning, LUN assignment, zoning, and masking, as well as monitoring and optimizing performance, capacity, and availability. In addition, support for continuous availability and disaster recovery requires that device copy services are available as a viable failover and disaster recovery environment. Traditionally, each device provides a command line interface (CLI) as well as a graphical user interface (GUI) to support these kinds of administrative tasks. Many devices also provide proprietary APIs that allow other programs to access their internal capabilities.

For complex SAN environments, management applications are now available that make it easier to perform these kinds of administrative tasks over a variety of devices.

The CIM interface and SMI-S object model adopted by SNIA provide a standard model for accessing devices, which allows management applications and devices from a variety of vendors to work with each other's products. This means that customers now have more choices available of devices that work with their chosen management application, and more choices of available of management applications they can use with their devices.

IBM has embraced the concept of building open standards-based storage management solutions. IBM management applications are designed to work across multiple vendors' devices, and IBM devices are being CIM-enabled to allow them to be controlled by other vendors' management applications.

5.7.3 IBM storage management of the virtualized SAN

IBM TotalStorage products use the SNIA model for CIM enablement to ensure you can manage IBM products easily in a simple and open way. In addition, each device supports a command line interface (CLI) to allow scripting of repeatable operator tasks. Advanced user interfaces are also provided to simplify the management of each device.

In a SAN environment, it is typical for multiple devices to work together to create a storage solution. The IBM TotalStorage Productivity Center is an open storage infrastructure management solution designed to help reduce the effort of managing complex storage infrastructures, to help improve storage capacity utilization, and to help improve administrative efficiency for interacting SAN devices.

Interacting SAN devices include:

- ▶ IBM TotalStorage SAN Volume Controller
- ▶ IBM TotalStorage SAN File System
- ▶ IBM TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000
- ▶ IBM TotalStorage DS4000 family
- ▶ IBM TotalStorage DS6000 family
- ▶ IBM TotalStorage DS8000 family
- ▶ IBM TotalStorage ESS devices
- ▶ Other vendors' products

The IBM TotalStorage Productivity Center is designed to enable an agile storage infrastructure that can respond to on demand storage needs.

Archived

Software to simplify managing your storage systems

This chapter provides basic information about software that simplifies managing your storage subsystems. These software tools can help you become more productive in your storage environment.

We describe in this chapter how IBM TotalStorage and IBM Tivoli products can help you get answers to questions you may have concerning your storage environment. Questions such as:

- ▶ What is our average storage utilization?
- ▶ When and where will we need additional storage resources?
- ▶ Where do we have potential problems in our storage environment which may become severe problems soon?
- ▶ What kind of data is filling up our disks?
- ▶ Which files stored on the disks are still in use?
- ▶ Which files are old and can be moved to less expensive storage?

The answers to these questions can help you to:

- ▶ Project storage growth
- ▶ Identify growth hotspots
- ▶ Identify data patterns that are key to business processes
- ▶ Identify data patterns that are counterproductive to the business goals or data that can be stored on less costly storage
- ▶ Help you to reduce or control your storage budget

Concerning backup and disaster recovery, these questions may be on your mind:

- ▶ Which data is not being backed up?
- ▶ How much total storage does this application need online and for backup?

- ▶ Is this group or line of business' data protected by the appropriate backup and/or replication method?
- ▶ How can I ensure that my backups do not fail because of lack of space?

The answers to these questions can help you make sure you are protecting vital business data and/or meeting government regulations for security and retention.

Additional questions that you may have are:

- ▶ How much money did we spending managing all of our storage last year?
- ▶ What percentage of each work week do our systems administrators spend “managing” all of our storage?
- ▶ How much new storage will our organization need next year?

The software we discuss in this chapter makes it easier to you to answer these questions and can help you on your way toward infrastructure simplification. It can help you increase productivity, effectiveness, and, because of automation, require you to do fewer manual interactions in your storage environment.

6.1 IBM TotalStorage Productivity Center

The IBM TotalStorage Productivity Center is an open storage infrastructure management solution. It has been designed to reduce the effort of managing complex storage environments with an integrated set of tools to assist system administrators to get the most out of their storage.

IBM TotalStorage Productivity Center is comprised of the following four products:

- ▶ *IBM TotalStorage Productivity Center for Data* (Productivity Center for Data) concentrates on data in the open systems environment. Productivity Center for Data can provide you 300 enterprise wide reports in the heterogeneous environment. These reports give you the capability to monitor your storage on a file-based level.
- ▶ *IBM TotalStorage Productivity Center for Disk* (Productivity Center for Disk) provides you information about your storage subsystems on a higher, more physical level than Productivity Center for Data does. IBM Productivity Center for Disk enables you to configure your SAN-attached devices from a single console. It also manages the performance for the ESS and SAN Volume Controller (SVC).
- ▶ *IBM TotalStorage Productivity Center for Fabric* (Productivity Center for Fabric) helps you to monitor and customize your fabric from a single console. Automatic disk discovery, changing the topology of your SAN, error detection, zone control, real-time monitoring, alerts, and event management for heterogeneous SAN environments are provided by Productivity Center for Fabric.
- ▶ *IBM TotalStorage Productivity Center for Replication* (Productivity Center for Replication) provides copy services management for the ESS. IBM Productivity Center for Replication helps you to configure and manage the Point-in-Time copy and Metro Mirror capabilities of the ESS in supported configurations.

Figure 6-1, shows where you can choose with which component you want to work.

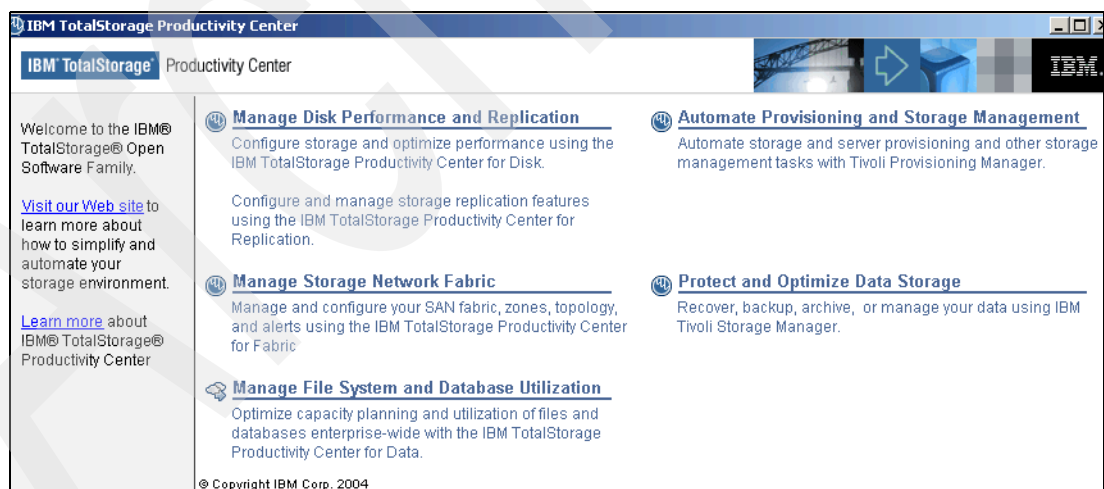


Figure 6-1 Productivity Center Launchpad

For information about how to install and customize these components and for more details, see the IBM Redbooks:

- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *Managing Disk Subsystems using IBM TotalStorage Productivity Center*, SG24-7097

6.1.1 IBM TotalStorage Productivity Center for Data

The job of TotalStorage Productivity Center for Data (based on Tivoli Storage Resource Manager) is to help you get a clear picture of the business information you are storing on your storage infrastructure. It lets you look underneath the hardware to understand how the data is being stored and used.

TotalStorage Productivity Center for Data starts by automatically identifying the databases, file systems, and files in your environment and providing analysis about the types of data being stored. This analysis is based on best practice categories shipped with the product and can be customized to your particular environment.

Based on this categorization, you can use TotalStorage Productivity Center for Data to automate actions. For example, suppose you find that you have a large quantity of files that need to be deleted. TotalStorage Productivity Center for Data can delete them automatically and routinely. Or you may find that you have data that needs to be archived. TotalStorage Productivity Center for Data can drive your archive product to archive the data, automatically.

With the environment now automatically managed and cleaned up, TotalStorage Productivity Center for Data can give you current and historical capacity utilization metrics about each category of data so you can do a better job of capacity planning.

IBM TotalStorage Productivity Center for Data gives you the ability to look at your data from several points of view:

- ▶ Host perspective
- ▶ Application perspective
- ▶ Accounting perspective
- ▶ File perspective

The TotalStorage Productivity Center for Data collects information about the storage environment (via a unified agent) and stores it in a database repository. Storing the information in a database gives you the capability to access historical data too. The stored information can be displayed from a native GUI client or browser interface anywhere in the network. The GUI or browser interface gives access to the other functions of TotalStorage Productivity Center for Data, including creating and customizing a large number of different types of reports and setting up agents.

The data, having been collected, may be used to discover, monitor, and create enterprise policies for your disks, storage volumes, file systems, files, and databases. You will be able to identify potential areas of exposure, evaluate the data residing on your servers, set up control mechanisms for autonomic management, and start the capacity planning process by predicting growth.

Note: As of release 2.3, TotalStorage Productivity Center for Data now supports DS6000 and DS8000.

Figure 6-2 shows the summary panel that displays, if you enter the TotalStorage Productivity Center for Data product. Several statistics with enterprise wide summary information display on this panel.

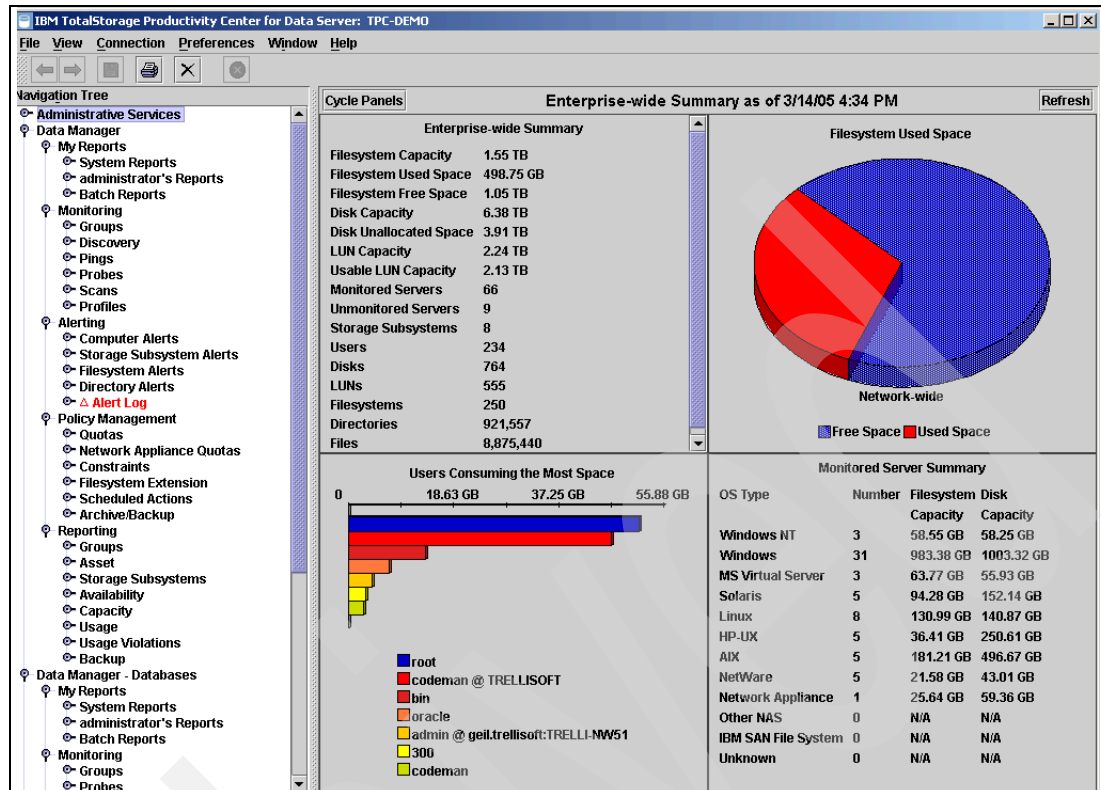


Figure 6-2 Productivity Center for Data: General overview

You can display this summary panel on any workstation. This panel allows the administrator or management a quick view of what is going on in storage environment.

You have access to over 300 standardized reports, and you can customize reports to your needs, if necessary. These reports give you detailed information about:

- ▶ Assets
- ▶ Availability
- ▶ Capacity
- ▶ Usage
- ▶ Usage violation
- ▶ Backup

Figure 6-3 gives you quick and simple access to basic information about the systems in your network. You get:

- ▶ Network address
- ▶ IP address
- ▶ Type of operating system with its version
- ▶ Time zone in which the system is running
- ▶ Manufacturer
- ▶ Model
- ▶ Serial number
- ▶ Processor type
- ▶ Processor speed
- ▶ Processor count
- ▶ Installed RAM size
- ▶ Swap space
- ▶ Disk capacity

- Unallocated disk space
- File system free space
- Last boot time
- CPU architecture

Computer	Network Address	IP Address	OS Type	OS Vers
TOTAL =>				
Cluster Group.SHAKENBAKECS.TRELLISOFT	SHAKENBAKECS.TRELLISOFT	9.48.179.18	MSCS Virtual Server	N/A
Group 0.SHAKENBAKECS.TRELLISOFT	TRELLISOFT	N/A	MSCS Virtual Server	N/A
Group 1.SHAKENBAKECS.TRELLISOFT	TRELLISOFT	N/A	MSCS Virtual Server	N/A
NW51-DEV1.geil.trellisoft:TRELLISOFTNDS	NW51-DEV1.geil.trellisoft	0.0.0.0	NetWare	N/A
NW6-DEV1.geil.trellisoft:TRELLI-NW6	NW6-DEV1.geil.trellisoft	0.0.0.0	NetWare	5.60.0
PATRICK.geil.trellisoft:TRELLI-NW51	PATRICK.geil.trellisoft	0.0.0.0	NetWare	5.0.10
SPONGEBOB.geil.trellisoft:TRELLI-NW51	SPONGEBOB.geil.trellisoft	0.0.0.0	NetWare	5.0.10
SQUIDWARD.geil.trellisoft:TRELLI-NW51	SQUIDWARD.geil.trellisoft	0.0.0.0	NetWare	5.0.10
adrian	adrian.tivlab.austin.ibm.com	9.48.197.2	AIX	5.2.0.0
aix433	aix433.trellisoft.com	9.48.179.212	AIX	4.3.3.75
ausaps5a	ausapp5a.tivlab.austin.ibm.com	9.48.197.83	Windows	5.0:Serv
bake	bake.TRELLISOFT	9.48.179.166	Windows	5.2
blacklab	blacklab.trellisoft.com	9.48.179.129	AIX	5.2.0.10
blog.trellisoft.com	blog.trellisoft.com	9.48.179.22	Linux	2.4.20-3
browser380	browser380.trellisoft.com	172.20.1.11	Linux	2.4.7-10
chodacki	chodacki.trellisoft.com	9.48.179.165	Windows	0:Servic

Figure 6-3 Productivity Center for Data: Asset information

Customers often use this information to track OS level and versions for compliance and security considerations. In addition, the asset information provides an accurate processor count which is often used in today's software pricing. Also, by displaying a hardware profile, machines eligible for upgrade are easily identified.

Figure 6-4, and Figure 6-5, show sample reports for total free space and wasted space. This information helps you to locate systems that have too much space assigned or are going to run out of space soon. Therefore, you can prepare to take the appropriate action before the system stops working because it runs out of space. The wasted space chart shows you which file systems are larger than they need to be. You can reduce these file systems to assign the freed up space to other file systems.

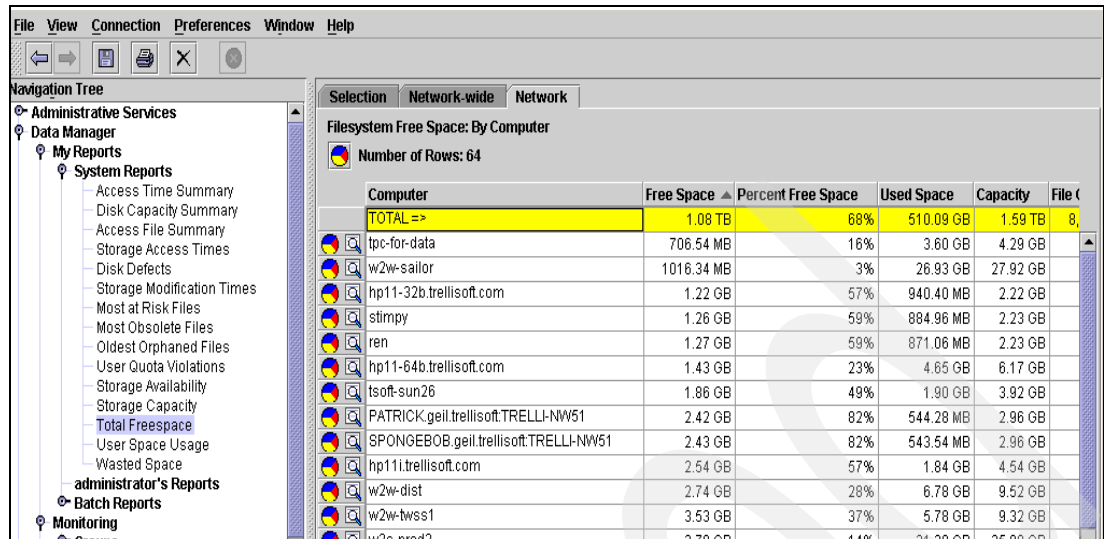


Figure 6-4 Productivity Center for Data: Total free space

Also, the free space identifies those areas where you can consider IBM Tivoli Provisioning Manager for automated action on a low or out of free space condition. The wasted space analysis (see Figure 6-5) can be used to determine where virtualization, such as the SVC, can be used to balance storage utilization across the enterprise.

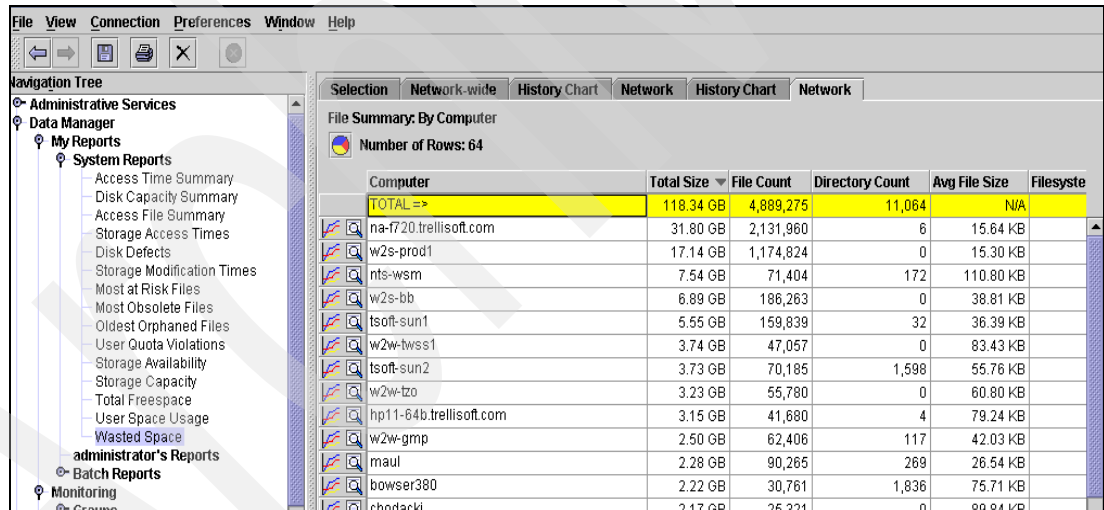


Figure 6-5 Productivity Center for Data: Wasted space

Figure 6-6 shows a sample of a projection of future needs of disk space.




Figure 6-6 Productivity Center for Data: Project future needs

You can use the information available from TotalStorage Productivity Center for Data to:

- ▶ Discover and monitor storage enterprise-wide.
- ▶ Create enterprise-wide reports on assets, files and file-systems, databases, users, and applications.
- ▶ Provide alerts on issues such as capacity problems or policy violations.
- ▶ Support chargebacks by usage or capacity.

For more details about the product, see the IBM Redbooks:

- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *IBM Tivoli Storage Resource Manager: A Practical Introduction*, SG24-6886
- ▶ *Managing Disk Subsystems using IBM TotalStorage Productivity Center*, SG24-7097

6.1.2 IBM TotalStorage Productivity Center for Disk

IBM TotalStorage Productivity Center for Disk is designed to centralize management of network storage devices that implement the SMI-S standard established by the Storage Networking Industry Association (SNIA). This includes the IBM TotalStorage Enterprise Server (ESS), IBM TotalStorage DS4000 series family, and IBM TotalStorage SAN Volume Controller (SVC), and any device they manage. In Productivity Center V2.3, support for the DS8000 and DS6000 is provided. Please check the following Web site for supported storage systems:

<http://www.ibm.com/support/docview.wss?uid=ssg1S1002592>

The IBM TotalStorage Productivity Center for Disk is designed to:

- ▶ Help reduce storage management complexity and costs while improving data availability
- ▶ Centralize management of different storage devices through open standards (SMI-S)
- ▶ Enhance storage administrator productivity
- ▶ Improve storage resource utilization
- ▶ Offer proactive management of storage devices

IBM TotalStorage Productivity Center for Disk provides access to single-device and cross-device configuration functionality. It allows you to view important information about the storage devices that are discovered by the IBM TotalStorage Productivity Center for Disk, examine the relationship between those devices, or change their configurations.

Performance information and alerts from the discovered devices can be selected from just one single application. You do not have to start several instances of a software product or access many different disks via a separate Web browser window. Just one window gives you access to a lot of important information about any of your disks. Therefore, you can be much more productive in managing your disk environment and you can easily improve data availability and storage utilization. Productivity Center for Disk gives you enough information and data to act proactively in your storage environment.

Figure 6-7 shows a panel from TotalStorage Productivity Center for Disk. You can access all of your subsystems from here.

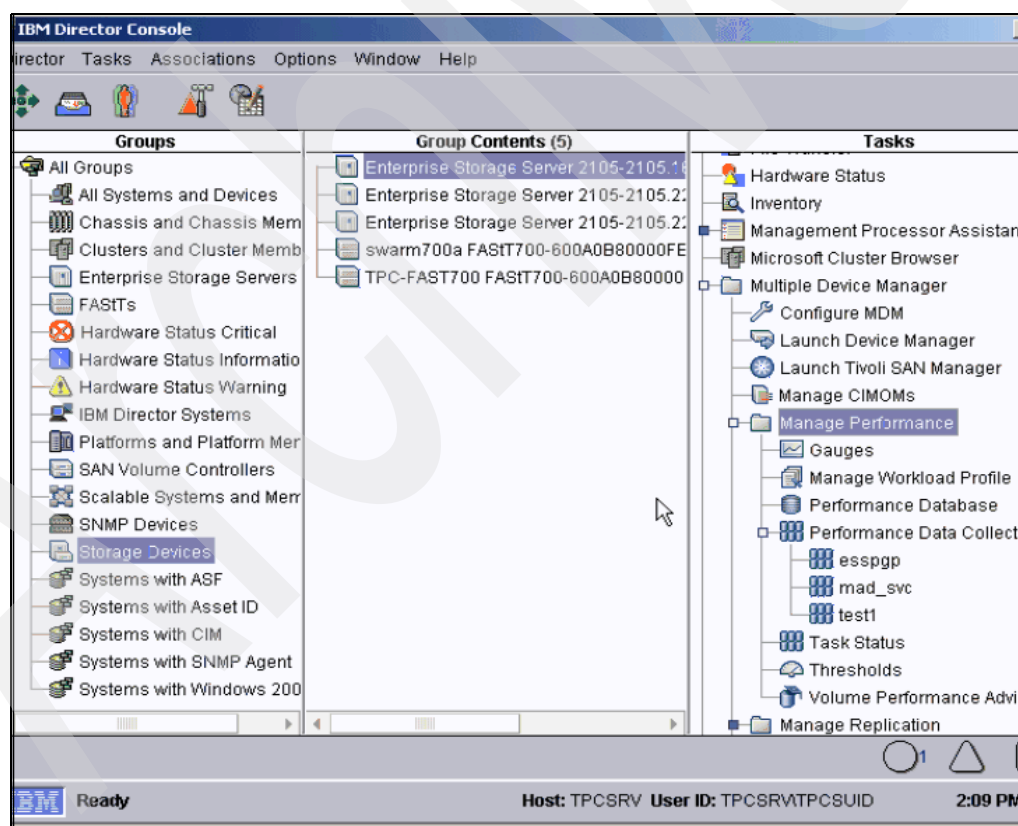


Figure 6-7 Productivity Center for Disk: storage devices panel

For more details about Productivity Center for Disk, see IBM Redbooks:

- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *Managing Disk Subsystems using IBM TotalStorage Productivity Center*, SG24-7097

6.1.3 IBM TotalStorage Productivity Center for Fabric

IBM TotalStorage Productivity Center for Fabric is a SAN Fabric management tool focused on discovering and monitoring the health of SAN islands that exist within an organization. A SAN island is a group of SAN switches, storage devices, and hosts that are connected together to form one network. TotalStorage Productivity Center for Fabric is architected to common industry standards and therefore allows you to choose best-of-breed hardware products for your storage infrastructure.

TotalStorage Productivity Center for Fabric uses a combination of agents installed on SAN connected servers and SNMP calls directly to SAN switch hardware to discover and monitor SAN health and operations. TotalStorage Productivity Center for Fabric performs a topology discovery and renders the components and storage resources. This enables you to validate the intended connections between systems and storage devices. Therefore, it is quite easy to check to see whether or not the real configuration meets the plans. Using TotalStorage Productivity Center for Fabric helps you to be more effective.

Figure 6-8 shows the TotalStorage Productivity Center for Fabric console.

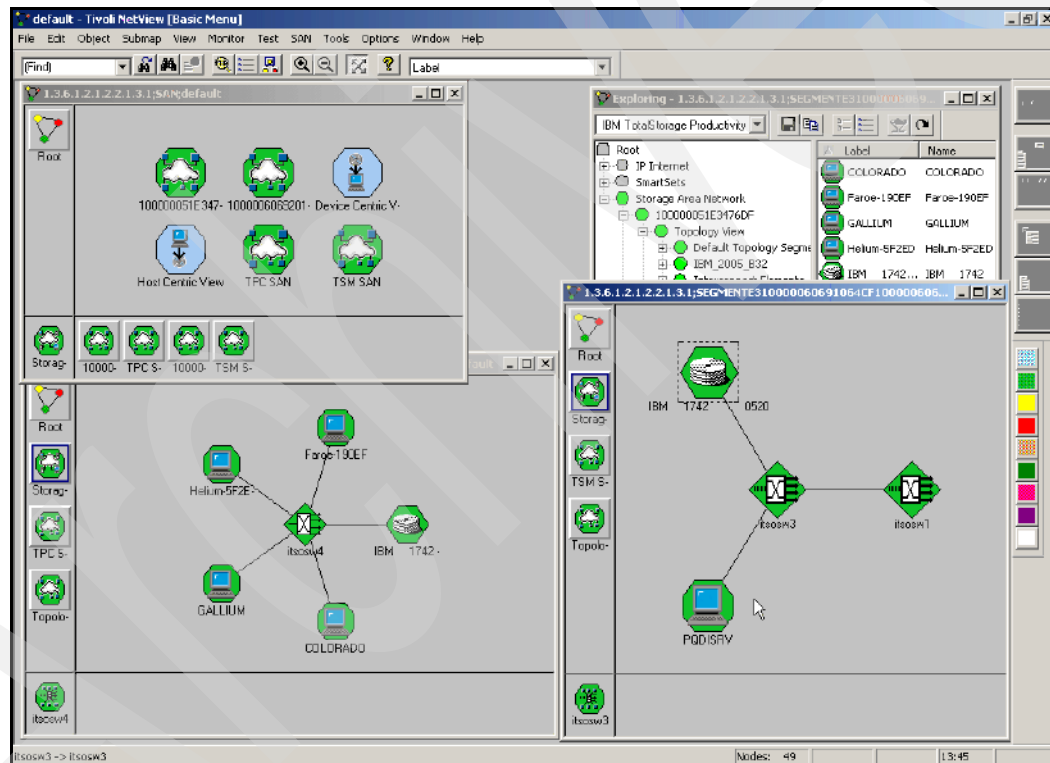


Figure 6-8 Productivity Center for Fabric console

From this console, you retrieve detailed information about your SAN elements. You display zoning information and host to device relationships. TotalStorage Productivity Center for Fabric has access to any element in your fabric that meets the standards. You do not need to use a different tool for different devices, which makes it much easier for you to change something in the topology of your fabric.

TotalStorage Productivity Center for Fabric also has the ability to view, change, and create SAN fabric zones on supported switches. Zoning functions are available as a GUI-based user interface for direct manipulation of fabric zoning or as API functions for external software to call.

Please check for supported hardware at:

http://www.ibm.com/software/sysmgmt/products/support/IBM_TSANM_Device_Compatibility.html

Figure 6-9 shows the Web page mentioned above. On this Web site, you can select the kind of hardware for which you need support information:

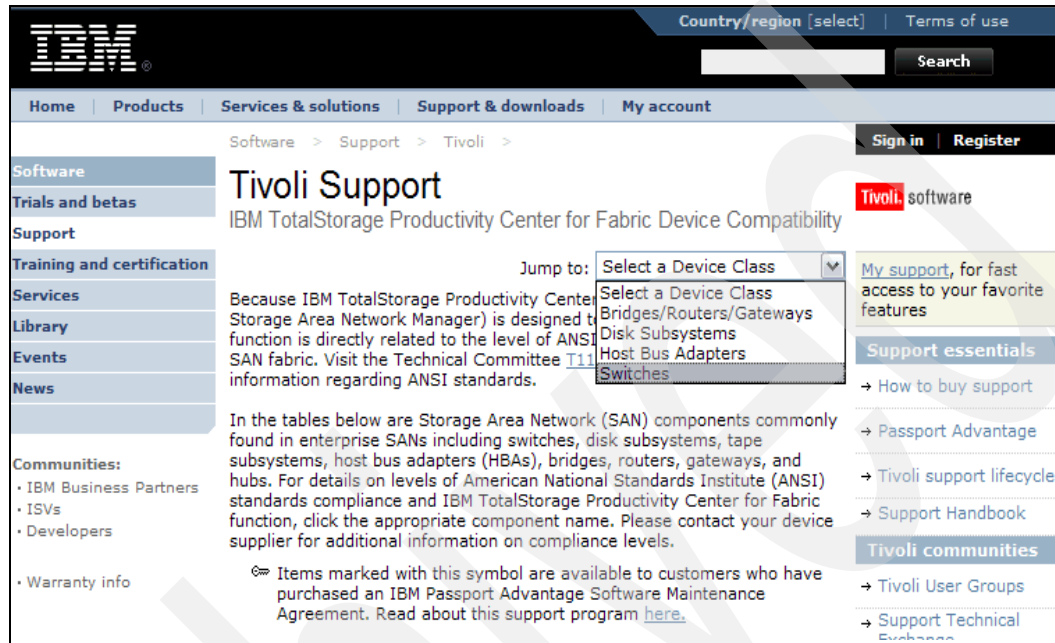


Figure 6-9 IBM Tivoli Support Web page

For more details about the product, see the IBM Redbooks:

- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *Managing Disk Subsystems using IBM TotalStorage Productivity Center*, SG24-7097

6.1.4 IBM TotalStorage Productivity Center for Replication

TotalStorage Productivity Center for Replication helps you administer and configure copy services functions from SMI-S compliant storage devices. From this tool, you are able to monitor the replications you initiated. Multiple pairs are handled as a consistent unit. Freeze-and-Go functions can be performed when errors in mirroring occur.

TotalStorage Productivity Center for Replication is designed to control and monitor the copy services in large scale client environments. You can use Productivity Center for Replication to perform the following tasks:

- ▶ Create replication groups
- ▶ Set up a group for replication
- ▶ Create, save, and name a replication task
- ▶ Create a replication session
- ▶ Manage a replication session

TotalStorage Productivity Center for Replication provides integrated administration, optimization, and replication features for interacting SAN devices. This provides an integrated view of the system, so you can perform complex configuration tasks and be more productive in managing your SAN infrastructure.

Figure 6-10 shows what you actually see.

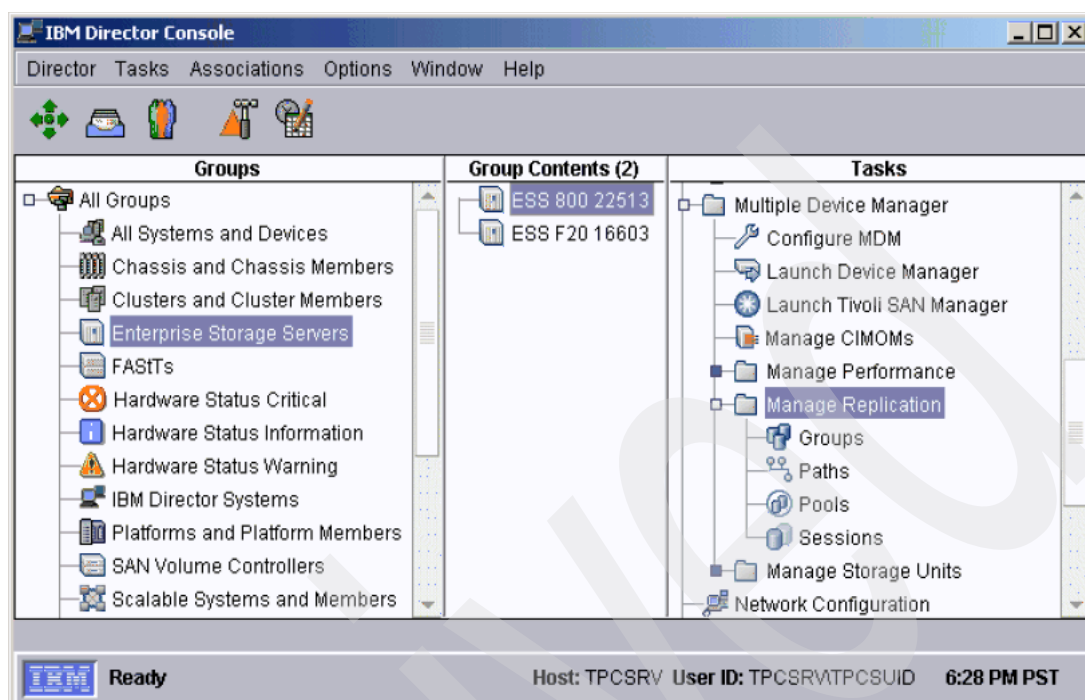


Figure 6-10 Replication Manager's tasks

6.1.5 Benefits

- ▶ Operates using device-level copy services
 - Operates "out-of-band"
 - Supports z/OS and open systems
- ▶ Single point of control for Point-in-Time and Remote volume replication services
- ▶ Automated source-target matching
 - Administrator defines target volume candidate pools
- ▶ Cross-device consistency groups
 - Single commands that execute across groups of volumes (script elimination)
 - Start/Stop Point-in-Time Copy
 - Start/Stop/Suspend/Resume Continuous Copy
 - Status Query
 - Resynchronize remote volumes
 - Used to synchronize databases or applications across multiple storage devices

For more details about the product, see the IBM Redbooks:

- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *Managing Disk Subsystems using IBM TotalStorage Productivity Center*, SG24-7097

6.2 IBM Tivoli Provisioning Manager

IBM TotalStorage Productivity Center with Advanced Provisioning is an integrated storage capacity provisioning solution designed to simplify and automate complex cross-discipline tasks for provisioning storage capacity in the enterprise environment and is designed to move you from just-in-case provisioning to intelligent, automated on demand provisioning. Simply put, it allows software to execute the tasks chosen by administrators rather than asking the administrator to perform each individual task manually.

With increasing demand for their services and no increase in personnel to handle the demands, IT staffs are often forced to overprovision IT resources, including storage capacity, based on worst-case, peak demands. The introduction of storage area networks has added to this challenge, potentially making storage capacity provisioning a very labor-intensive process requiring as many as 50 individual steps and up to several days of an expert's time. The result, the IT environment can become inflexible, expensive, under-utilized, and difficult to manage.

The TotalStorage Productivity Center with Advanced Provisioning is designed to help you:

- ▶ Reduce management costs through workflow automation.
- ▶ Improve availability by reducing the human error factor and automating your best practices.
- ▶ Provision new servers and applications quickly by facilitating the integrated provisioning of servers and storage through shared workflows and a common tool.

How IBM Tivoli Provisioning Manager provides storage provisioning capabilities by utilizing the Productivity Center storage management suite is covered in detail in Table 6-1.

Table 6-1 Comparison of IBM Tivoli Provisioning Manager and Productivity Center

Attribute	Tivoli Provisioning Manager	TotalStorage Productivity Center
Product description	Data center solution with automated storage provisioning capabilities.	Complete storage management solution with additional automated storage provisioning capabilities (through IBM Tivoli Provisioning Manager).
Positioning	Release management	Availability management
User	Server administrator deploying a new server	Data center storage administrator
Device management	Basic: IBM Tivoli Provisioning Manager automation packages (workflows) comprise a toolkit that can be customized to meet customer requirements	Advanced: Complete, comprehensive out of the box solution for storage management
Device interfaces	Interfaces directly to supported storage (Cisco, Brocade, and McData) using CLIs	Exploits multiple standards-based interfaces for comprehensive management (reporting and control) across many vendors (via SMI-S standard)
Platform support	Limited to platforms supported with IBM Tivoli Provisioning Manager automation packages	Heterogeneous: Supports multiple devices and vendors

The next two figures show how IBM Tivoli Provisioning Manager makes it easier to provide storage. Figure 6-11 shows the workflow that is necessary to provide a new disk to a server without using IBM Tivoli Provisioning Manager.

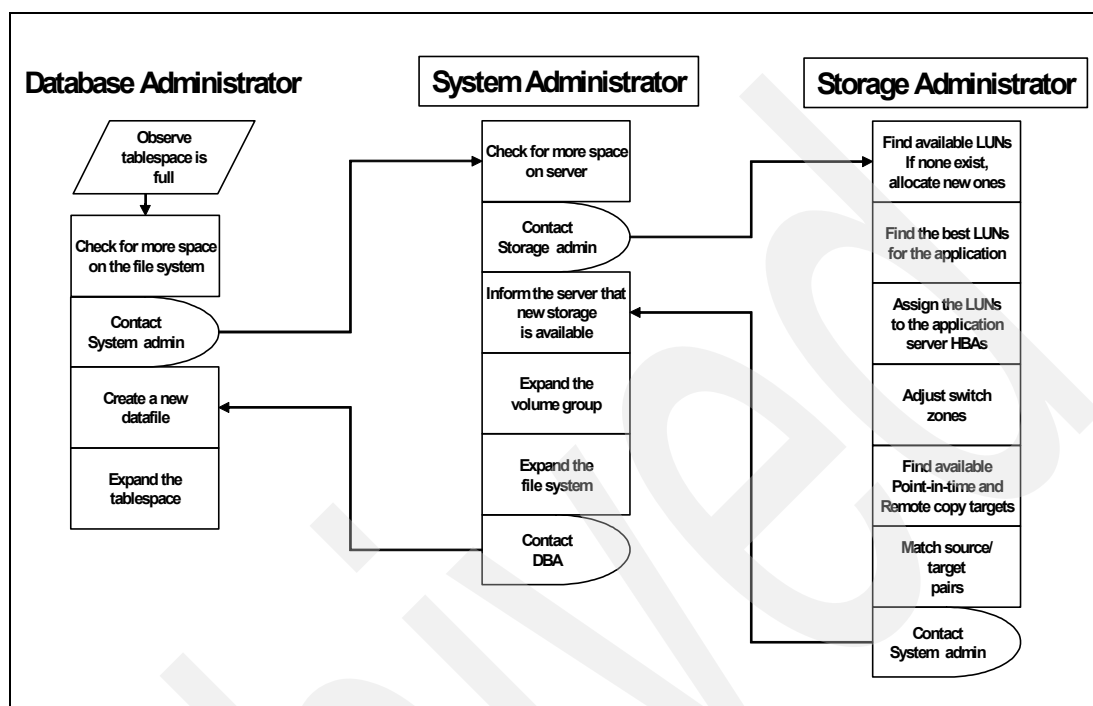


Figure 6-11 Storage provisioning: Manually

These are the steps:

1. The database administrator (DBA) will find out there is a need for more space on the file system because a tablespace is going to be full. The DBA will have to contact the system administrator who will have to contact the storage administrator, if there is no more space available on the server.
2. The storage administrator will have to allocate a new LUN, if none exists, assign the LUN, adjust the zoning, prepare copy targets, and match source and target pairs since the database requires mirrored data. Finally, the storage administrator will have to contact the system administrator again when the LUN is ready.
3. The system administrator will have to add the new LUN to the server, expand the volume group, and expand the file system. Finally, the system administrator will have to contact the database administrator.
4. The database administrator now can create a new data file and expand the tablespace.

Figure 6-12 shows the workflow to do the same thing if you have Tivoli Provisioning Manager supporting you in this task. After defining the tasks within Tivoli Provisioning Manager once, the storage administrator just has to initiate the process.

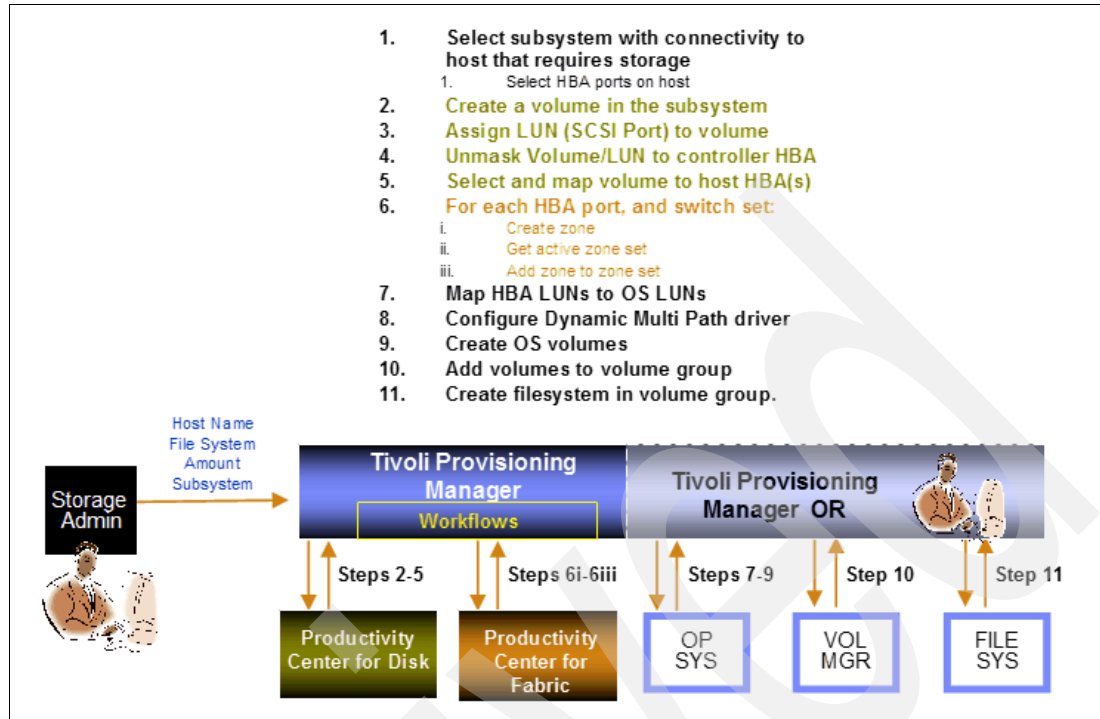


Figure 6-12 Storage provisioning supported by Tivoli Provisioning Manager

A workflow to increase the size of a windows file system with Tivoli Provisioning Manager can look like the workflow shown in Figure 6-13.

Workflow Execution			Simple view	Detail view
Workflow: Grow a file system on Windows Host			Logical Operation: none	
Description: Grow a file system on Windows Host				
User: tioappadmin				
Parent Workflow: None				
Execution Time	Transition Execution	Status		
2-26 16:37:04.03	Get current file sys size from SRM			
2-26 16:37:15.236	Check Size			
2-26 16:39:09.841	Expand SAN Volume Controller disk			
2-26 16:38:23.705	Get filesystem to vdisk mapping from SRM			
2-26 16:38:38.406	Get SVC host name from SRM			
2-26 16:39:09.45	Expand SVC Disk			
2-26 16:39:09.06	Expand SVC Disk			
2-26 16:39:55.316	Rescan Windows disk			
2-26 16:39:55.176	Rescan Disks			
2-26 16:43:10.397	Expand Windows disk			
2-26 16:40:41.493	Get filesystem to windisk mapping from SRM			
2-26 16:42:00.486	Create script to expand disk			
2-26 16:41:22.993	Get Drive letter from filesystem			
2-26 16:41:41.489	Substitute drive letter			
2-26 16:42:00.176	Substitute Disk number			

Figure 6-13 TPM workflow: Growing a file system on a Windows host

For more detailed information, see the IBM Redbooks:

- ▶ *Exploring Storage Management Efficiencies and Provisioning - Understanding IBM TotalStorage Productivity Center and IBM TotalStorage Productivity Center with Advanced Provisioning*, SG24-6373
- ▶ *Provisioning On Demand Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888

6.3 IBM Tivoli Storage Manager

IBM Tivoli Storage Manager (IBM Tivoli Storage Manager) protects an organization's data against hardware failures and other errors by storing backup and archive copies of data in offline storage. It can scale to protect hundreds of computers from laptops to mainframes, running a dozen different operating systems, connected via the Internet, WANs, or LANs.

IBM Tivoli Storage Manager is a full-function storage software product that addresses the challenges of complex storage management across distributed environments. It protects and manages a broad range of data, from the workstation to the corporate server environment. More than 39 operating platforms are supported; all include a consistent graphical user interface. IBM Tivoli Storage Manager provides:

- ▶ Centralized administration for data and storage management
- ▶ Efficient management of information growth
- ▶ High-speed automated server recovery
- ▶ Full compatibility with hundreds of storage devices, as well as LAN, WAN, and SAN infrastructures
- ▶ Customized backup solutions for major groupware, Enterprise Resource Planning (ERP) applications, and database products

IBM Tivoli Storage Manager is a premier choice for complete storage management in mixed platform environments. It is used by more than 80 of the Fortune 100 companies, and it protects more than a million systems around the world. Today, IBM Tivoli Storage Manager features are more powerful, modular, flexible, and easy to use than ever before.

Figure 6-14 gives you an idea of how IBM Tivoli Storage Manager server works.

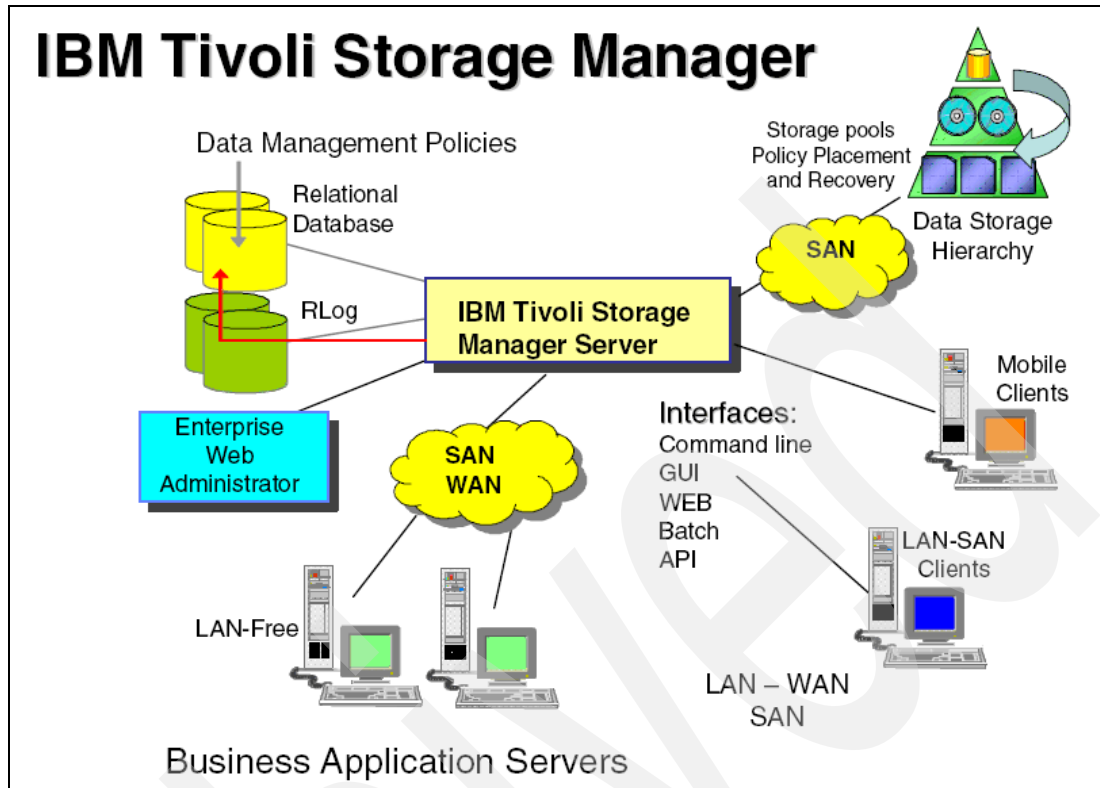


Figure 6-14 IBM Tivoli Storage Manager server

The incremental forever or progressive incremental technology works to decrease tape, tape slot, tape drive, and network costs to help provide you a better ROI.

You can find more information on the following Web site:

<http://www.ibm.com/software/tivoli/products/storage-mgr/>

The following IBM Redbooks cover IBM Tivoli Storage Manager topics:

- ▶ *IBM Tivoli Storage Management Concepts*, SG24-4877
- ▶ *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638

6.4 Customer C: Productivity Center in real life

This section describes how an IBM customer, called Customer C, used Productivity Center for Data to help simplify their storage infrastructure. Customer C installed Productivity Center for Data in the environment and started to collect data. After one week of data collection, Customer C reviewed the first results and made several changes to the TotalStorage Productivity Center configuration to look more closely into business directories. Another week of collecting data passed and the findings were presented again to refine the next phase.

The next figures show the charts created during the analysis of the installation at Customer C. See Figure 6-15.

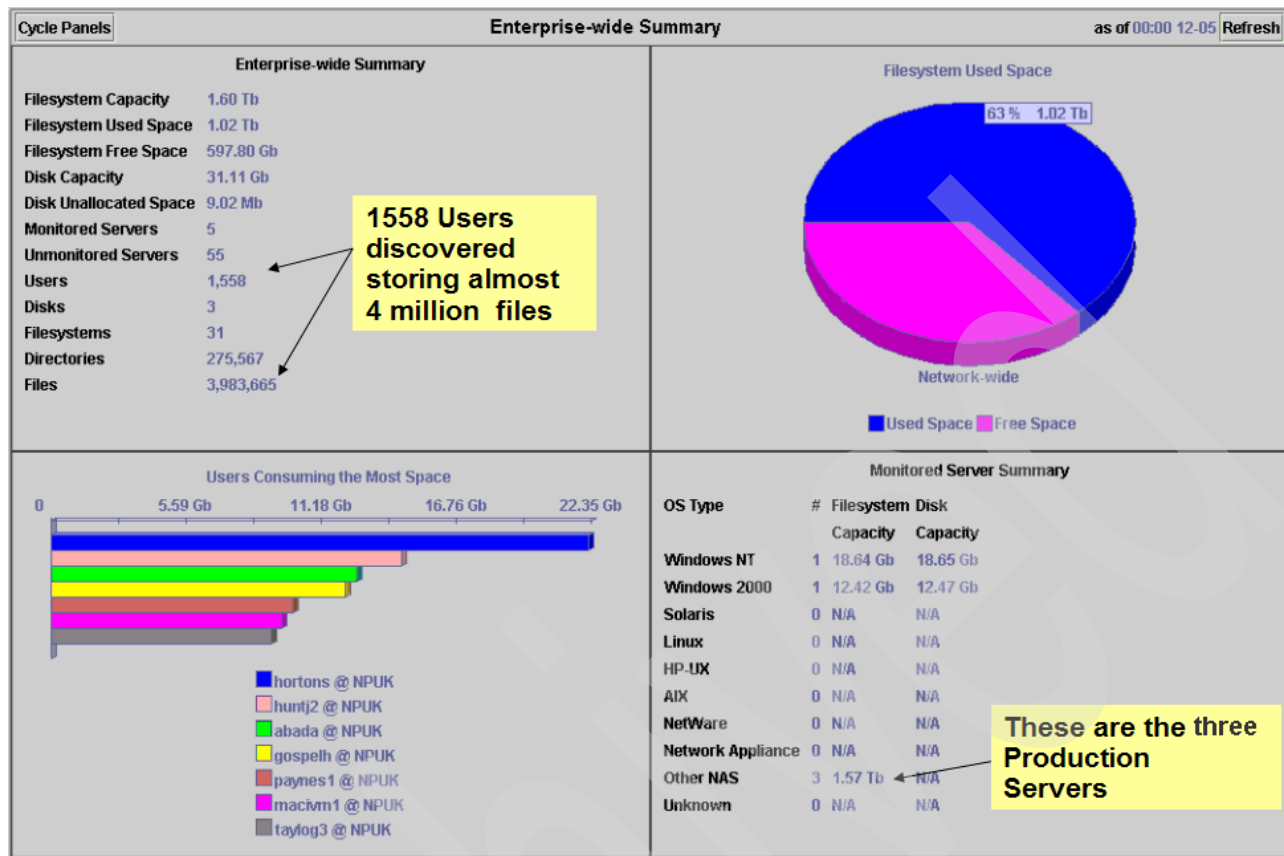


Figure 6-15 Customer C data in the Productivity Center for Data: Enterprise-wide Summary

See Figure 6-16 for several systems.




Figure 6-16 Customer C: Trending utilization of disk space

Look at Figure 6-17 and files you need to watch!

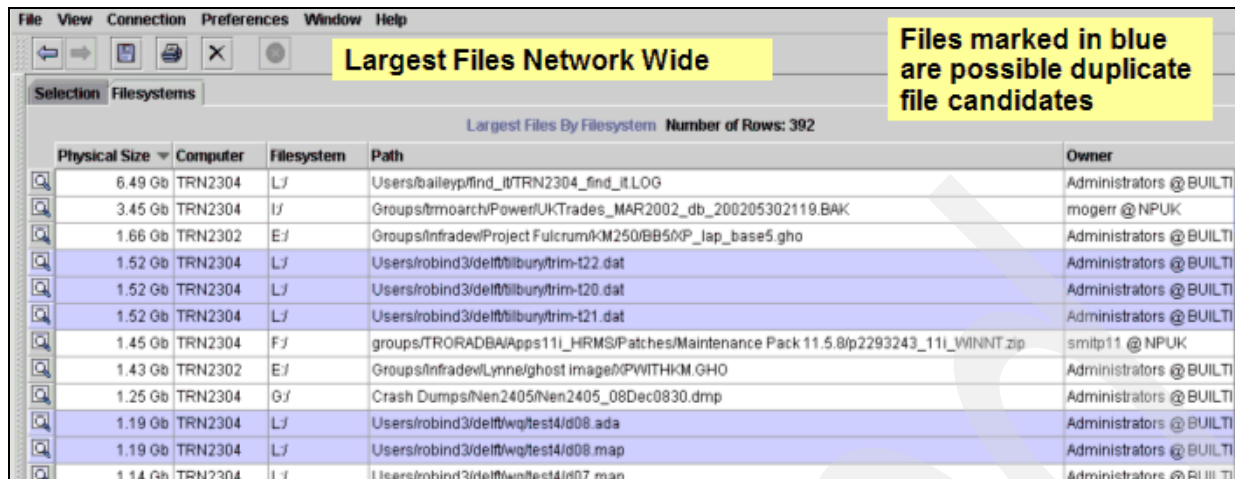


Figure 6-17 Customer C: Largest Files by filesystem Report: Shows possible duplicate files

Figure 6-18 shows two file systems which are at risk of running out of space. These file systems have a lot of *stale files*, or files that have not been accessed in more than six months.




Figure 6-18 Customer C: Report on large file systems running full and containing stale files

There are many additional reports that we did not discuss here. Reports such as:

- ▶ Largest network wide users
- ▶ File type summary (duplicate music files included in the backup)
- ▶ File sizing summary
- ▶ Trending data growth network wide
- ▶ Production file systems utilization

All these reports give you greater, more detailed insight into your installation. Here are sample results from the reports created after the analysis:

- ▶ Customer C exhibits at least 250% growth in creation of data over the last two months.
- ▶ Customer C has no policies to take action to groom the disks of stale, dump, temp, non-business, or duplicate data.
- ▶ If Customer C implements policies to automate the grooming of disks every day, at least 35% of the disk space can be reclaimed on an ongoing basis (this may take place over a period of months).
- ▶ The total cost of ownership of adding unnecessary capacity should be considered such as cost of the disk, tape, upgrade of the controllers, employee-time to perform the upgrades, backup team time, and slower data recovery times.
- ▶ The size of the IBM Tivoli Storage Manager backup database can also be better controlled, if only mission critical data is stored online and is included in the backup window.

A one sentence summary of the analysis at Customer C:

Customer C realized that if they take action on these findings on stale and redundant data cleanup, 400 GB is potentially reclaimable. According to Gartner Group, this can add up to a savings of more than \$300,000 per year.



Infrastructure Simplification using the IBM System Storage N series

This chapter discusses a customer example and explains how the N series storage system helped simplify a customer's storage infrastructure, including addressing their data growth and decreasing their storage management burden. This customer is a leader in the medical industry in the field of medical (genome) research. For the sake of their privacy, we refer to the customer as Customer N.

7.1 Challenge: Expanding Linux workload: Wild server crashes

Over a three year period, lab data volumes grew from under 1 TB to 18 TB and the commodity hardware Linux cluster topped the 200-node mark. Customer N was bursting at the seams in terms of Linux file server capacity.

“Sequencing production processes generate tens of millions of files, in fact, a single volume can contain nearly 30 million files. They had absolutely reached the performance and manageability limits of the existing file server structure. A number of disruptive file server crashes further emphasized the urgency of finding a more robust solution. System reliability is a principal concern in sequencing processes, one production stop can jeopardize several hundred thousands of dollars’ worth of data moving through the lab.”

Solution cost and complexity are important considerations for a lab funded largely through grant awards.

“As is the case at most research organizations, it’s easier to find budget for science than budget to run an IT group.”

With limited headcount to support our infrastructure, we always strive to deploy highly flexible and easily manageable technologies.

7.2 Unified solution: N Series System Storage

Today the bulk of Customer N’s data resides on an N Series fabric-attached System Storage with 14 TB capacity. The types of data include:

- ▶ Production pipelines for genome mapping and DNA sequencing
- ▶ Scientists’ project data
- ▶ Computational analyses
- ▶ Encrypted patient data
- ▶ Lab information management system data

The N Series System Storage is part of a unified storage solution that provides network-attached storage (NAS) for all lab applications. Regulated patient data is encrypted at the server and stored on an isolated iSCSI LUN on the N Series storage system.

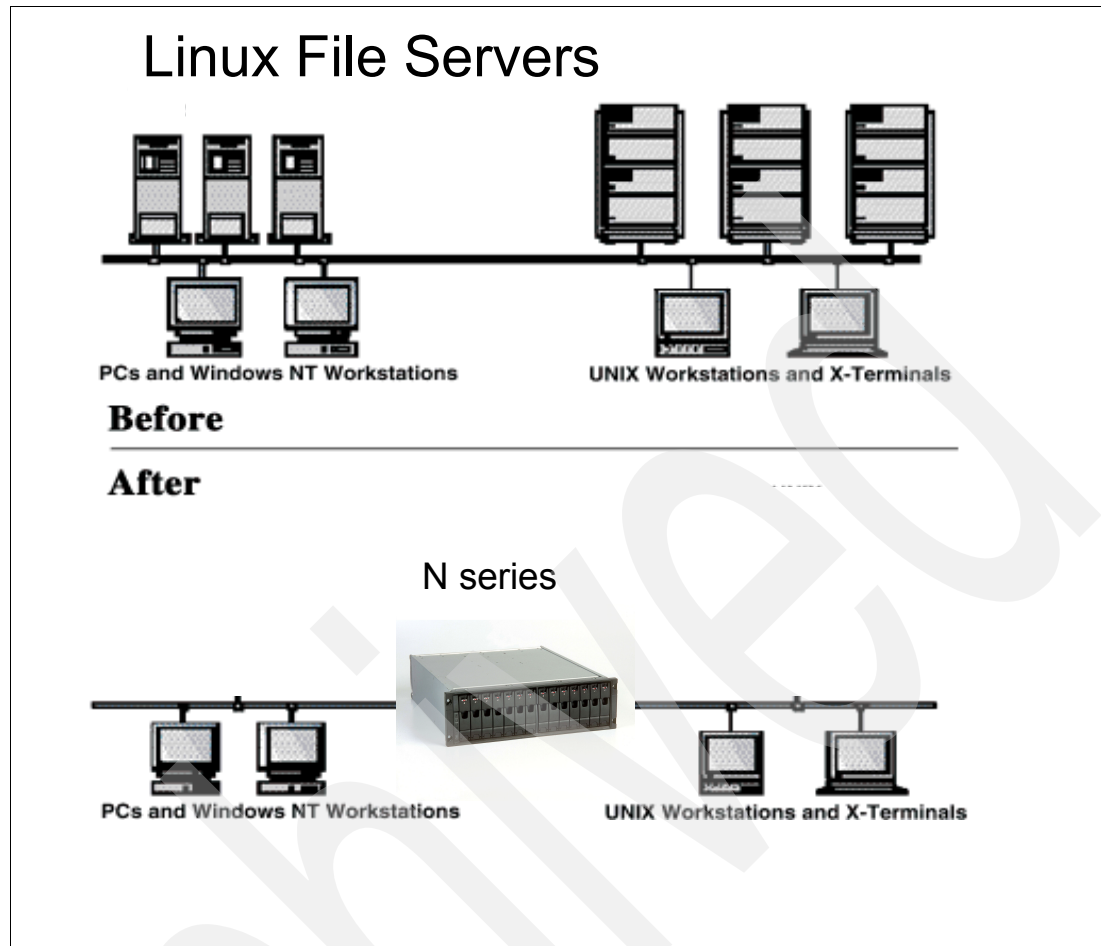


Figure 7-1 Before and after consolidation

Lab data is archived to another N Series System Storage, which provides cost-effective nearline storage. DataFabric Manager software automates storage structure monitoring and trend reporting to facilitate growth rate prediction.

The N Series solution was a natural choice for Customer N. The bulk of their data is held in file systems, an arena in which the N Series excels. Customer N knew that the N Series systems could also handle the heavy I/O loads in their lab environment. It made sense to complement their Linux cluster with a storage architecture equally straightforward and scalable.

7.3 Business benefits: Fast sequencing

The N Series solution keeps production pipelines operating interruption-free, while it also supports more than 200 users and research collaborators running an array of Linux desktop and Windows applications.

“Performance has been excellent, even when the system was being highly utilized.”

Customer N moves a lot of data around, especially during heavy analysis on the Linux cluster. The fact that they can now push close to 400 MB per second to the cluster, yet still have responsive NFS mounts on the Linux desktops is essential. They do not have to worry about error-prone data staging just to isolate cluster traffic, nor do we have to duplicate data

all over the place to try and squeeze more performance out of a file server. Additionally, because they hold all their raw sequence and mapping data on the secondary N Series System Storage, scientists have extremely fast access to archive data at a comparatively low cost per gigabyte.

Since the deployment of the N Series solution, there have been no process stopping failures, a benefit that people estimate may have already saved the company hundreds of thousands of dollars.

“When a failure occurred in the old file server structure, there was a risk that they could have to restart a weeks-long process. Science productivity was negatively impacted, and the potential existed to lose large quantities of costly chemical and reactive agents used in the sequencing processes. The N Series solution is inherently more reliable with the added benefit of Snapshot technology for near instantaneous data recovery.”

On demand storage scalability ensures that production does not stop for expansion. Analysis of lab data can require significant computational time. It is important that they do not run out of capacity in the middle of a large job. Recently, near the end of a two-week process on the Linux cluster, they had to add a shelf to expand a volume that DataFabric Manager had identified as being close to capacity. It took only 20 minutes to bring an additional 2 TB online, and they added the shelf while the N Series System Storage was live and pushing over 30,000 NFS ops/second. If they would have had to stop the process to add capacity, they would have lost two weeks of computation.

“The ability to meet very dynamic per project and per-team storage requirements without downtime gives them greater flexibility to accept new research challenges. As just one example, their processing resources, combined with unique applications, enabled them to complete the base assembly of a coronavirus in just five days.”

7.4 Secure Storage and Backup of Regulated Data

N Series plays an important role in the protection of sensitive patient data, helping the business comply with government information privacy laws. Government regulations specify the types of personal data that can be collected, how it is stored, precautions that must be taken to protect the data, and more.

“Storing Windows XP encrypted patient information on the N Series SCSI SAN allows them to easily segregate and automatically back up encrypted data. Tape backups are completely secure, if someone attempted an unauthorized playback of one of the patient data tapes, all they would see is a large encrypted file. The iSCSI LUN, which is easy to set up, expand, or shrink, simplifies the process and keeps our security high, even on backups.”

7.5 Simplicity and minimal IT overhead

N Series technology is simple to understand and manage.

“With an IP-based storage structure, it is very easy to manipulate and prioritize traffic and processes, share files, or isolate data, such as patient information stored on the iSCSI SAN. Customer N used to spend a lot of time dealing with file servers and NFS issues and creating islands of storage for various purposes. Now they have a simple, unified structure from which they can seamlessly provision capacity. The N Series solution eliminates the need to deal with complicated hierarchical storage management (HSM) schemes. They spend almost no

resources administering the N Series System Storage and have saved at least a full headcount. The structure is equally simple from a user perspective. Scientists do not want to waste time dealing with complex data access patterns. They want to do science, and that is what the N Series storage system solution enables.”

The knowledgeable support staff has been another benefit of the N Series solution.

Archived



A storage provisioning solution for SAN File System

This chapter is the Redpaper, *An Introduction to Storage Provisioning with Tivoli Provisioning*, REDP-3900, by Steve Strutt, reproduced here with permission of the author.

This chapter looks at an example of a TPM-based storage provisioning solution written to manage IBM TotalStorage SAN File System.

8.1 An on demand storage environment

To enable an IT Server Provider to provide on demand storage services, with pay by use charging and low administrative overhead, a solution based on IBM TotalStorage SAN File System and Tivoli Provisioning Manager was developed and implemented. Key requirements specified for the storage environment by the provider were:

- ▶ Low administrative overhead and skill requirements
- ▶ Automated management
- ▶ Pooling and sharing of capacity between users
- ▶ Sharing of data between hosts
- ▶ Support for Web server orchestration
- ▶ Heterogeneous platform support

The need was for a single storage infrastructure which would be managed simply, but also offer great flexibility and low costs. Flexibility and pooling of storage were delivered by deploying SAN File System and IBM TotalStorage SAN Volume Controller (SVC). This provides a common infrastructure with a single set of management tools for a heterogeneous environment. Storage provisioning using IBM Tivoli Provisioning Manager (TPM) was implemented to automate and simplify management of SAN File System's SAN infrastructure and support an orchestrated Web server environment. The use of provisioning reduced the need for key, scarce storage skills and delivering low management costs. Figure 8-1 shows a logical view of the solution, with applications sharing virtualized storage under control of the IBM orchestration and provisioning tools.

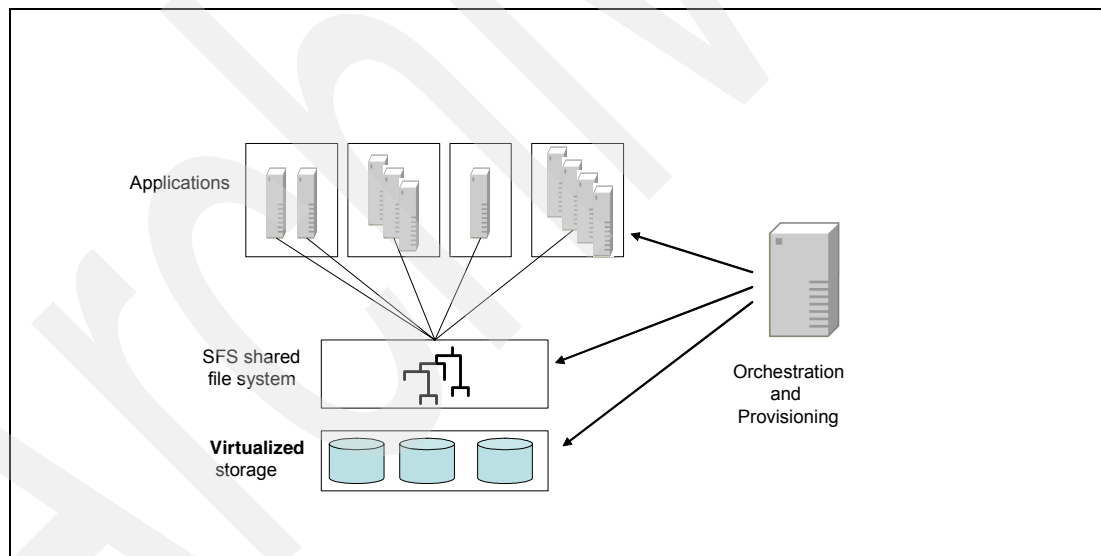


Figure 8-1 On demand environment

The focus here is primarily on the storage provisioning solution developed for this environment. The SAN File System architecture provisioning approach described here is also applicable to other SAN File System environments, since the management tasks described are common to any environment. This provisioning solution is a specific example of the power of TPM to create a fully automated storage environment using storage provisioning.

8.2 The SAN File System solution

The SAN File System environment and the associated provisioning solution support two user communities:

- ▶ An on demand storage environment offering usage-based pricing and capacity on demand for traditional applications and servers.
- ▶ A fully orchestrated application environment. Server and application orchestration provides dynamic provisioning of IBM @server BladeCenter® blade servers as application loads demand. This additionally exploits SAN File System's data sharing capabilities for high-performance, scalable, and secure sharing of application and Web site file storage between servers in the application clusters.

The use of SAN File System enables the service provider to deliver competitively priced storage service offerings to all its customers though the use of pooled disk storage and low file system administration overheads. A key requirement was low management cost, and since SAN technologies were new to the service provider, management of the SAN infrastructure needed to be task-oriented by administrators not skilled in SAN or storage technologies. The storage provisioning solution based on TPM provided the ability to manage the environment simply by using a number of defined provisioning tasks, significantly reducing the need for skilled administrators.

8.2.1 SAN File System physical environment

Figure 8-2 shows the environment as deployed with Fibre Channel-attached SAN File System clients sharing virtualized storage presented via an IBM TotalStorage SAN Volume Controller (SVC). The SAN fabric is CISCO MDS 9509 directors with 9140 edge switches. Disk storage is over 40 TB of DS4300 with FC and SATA disk. Hosts are Intel HS20 blades in IBM BladeCenter racks running Microsoft Windows and RedHat Linux, also standalone HP Intel servers and Sun Solaris servers.

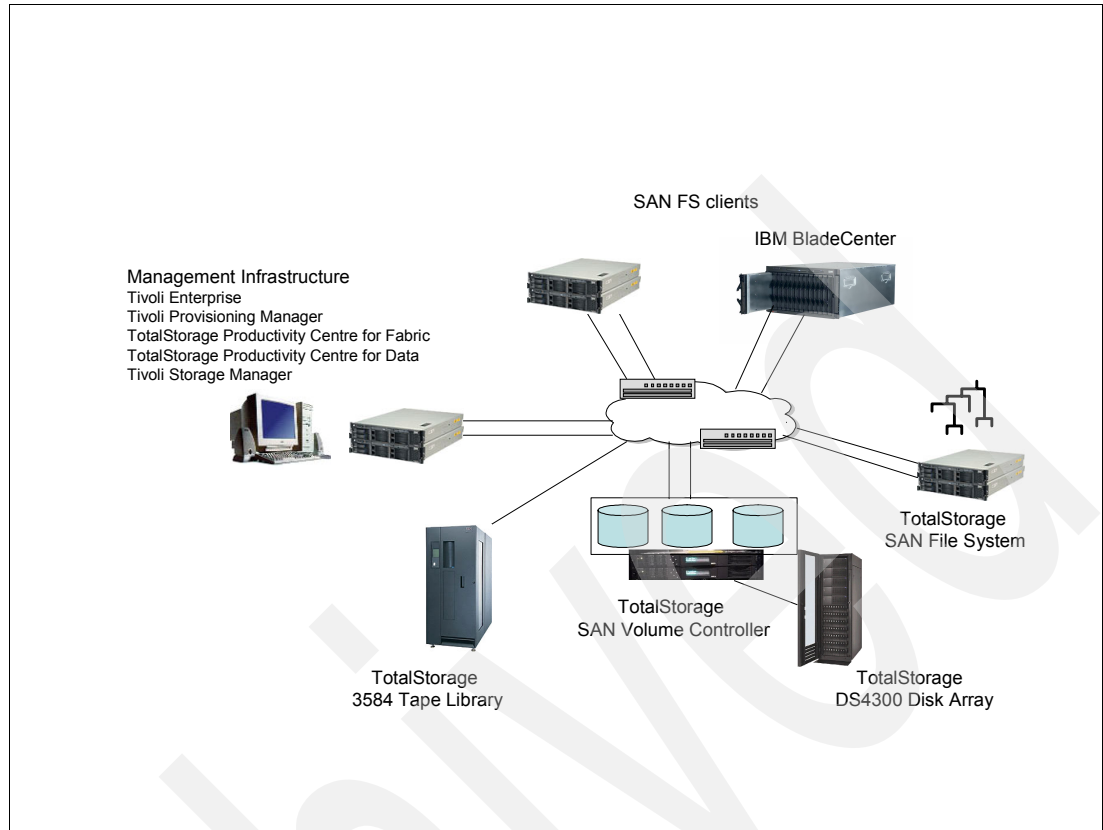


Figure 8-2 Physical SAN File system environment

8.2.2 Storage usage

The two user communities share the SAN File System environment. All users exploit the storage pooling and sharing of capacity managed through TPM.

Storage on demand environment

The on demand storage environment offers usage-based pricing and capacity on demand for traditional applications and servers. The different application user groups exist in different branches of the SAN File System file system structure, but share the same pools of SAN File System disk storage as shown in Figure 8-3.

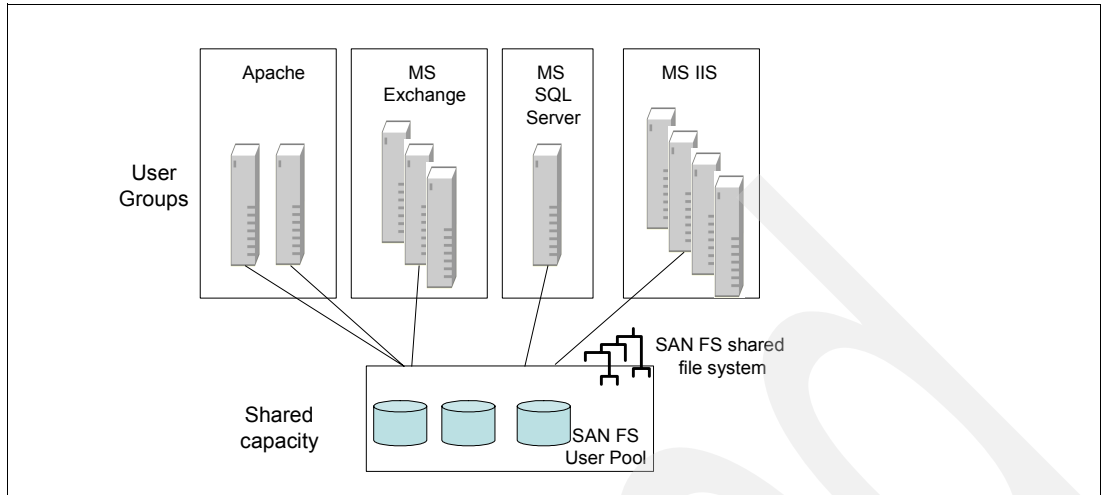


Figure 8-3 Shared storage environment

The different user groups share the same SFS Pools, sharing capacity. Free space is aggregated across all the users of the pool and is available on demand for capacity growth. Users are only charged for their usage and not for unused space. The cost of providing the free space buffer is spread across all users by the cost model.

Orchestration environment

The orchestrated application environment offers dynamic provisioning of IBM eServer BladeCenter blade servers as application loads demand using orchestration and provisioning. This is shown in Figure 8-4 on page 141.

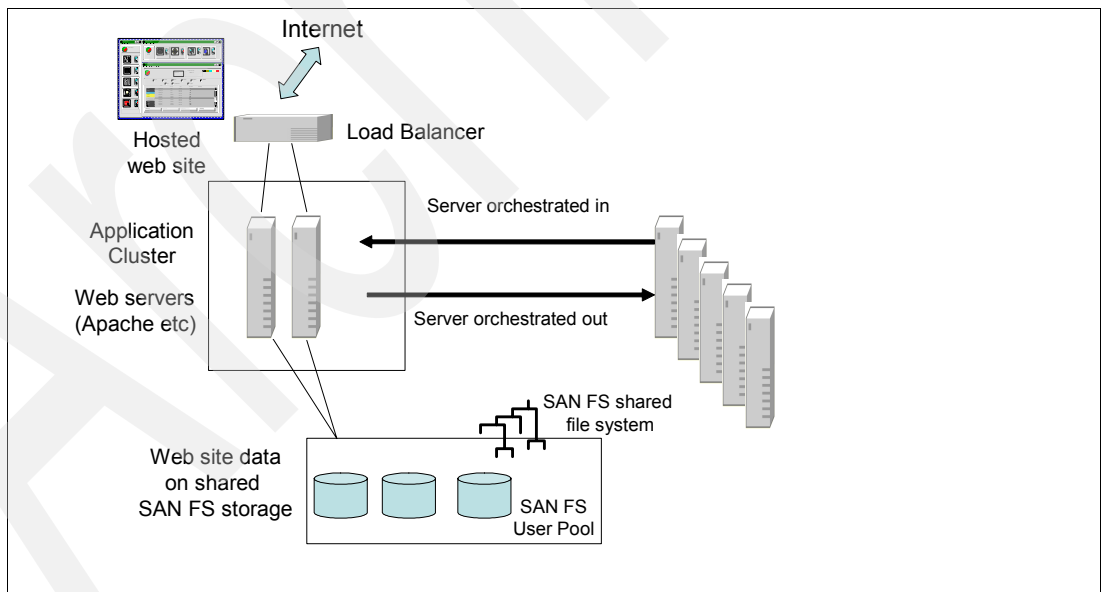


Figure 8-4 Orchestrated Web serving using SAN File System

This exploits SAN File System's data sharing capabilities, to share a single copy of Web site data between multiple Web servers. As load increases on the Web site due to increased user activity, additional servers are provisioned into the application cluster, spreading the workload evenly across the new larger pool of servers. SAN File System provides for

high-performance, scalable, and secure sharing of the application and Web site file storage between servers in the application clusters.

The ability of the storage provisioning solution to automatically add new SAN File System clients on demand without manual intervention is essential to the ability to rapidly orchestrate in new Web servers as utilization levels change. The new servers connect into and share the existing single copy of the Web site data in the SAN File System environment.

8.3 Storage Provisioning to simplify volume management

A SAN File System solution reduces and simplifies volume management tasks compared to traditional file systems because there are typically far fewer LUNs to manage. Provisioning can further simplify administration by automating the remaining volume management tasks required to add new clients or capacity. Another benefit is that administrators do not require substantial SAN administration skills, which are scarce.

Volume management tasks for SAN File System are only needed when new client servers are added or additional storage capacity is required. The volume management tasks associated with configuring SAN-attached storage in a SAN File System environment are similar to those for server-centric file systems, including the Linux extended file system V3 (EXT3), windows NTFS, and the AIX journaled file system (JFS). Compared to these server-centric file systems, these tasks are performed much less frequently, typically only when all free space in a pool has been exhausted. Tasks include:

- ▶ Loading and configuring the host bus adapter (HBA) driver
- ▶ Loading and configuring multipathing software, such as IBM Subsystem Device Driver
- ▶ Zoning the SAN
- ▶ Masking the storage subsystem LUNs

Figure 8-5 illustrates the actions automated by the storage provisioning solution.

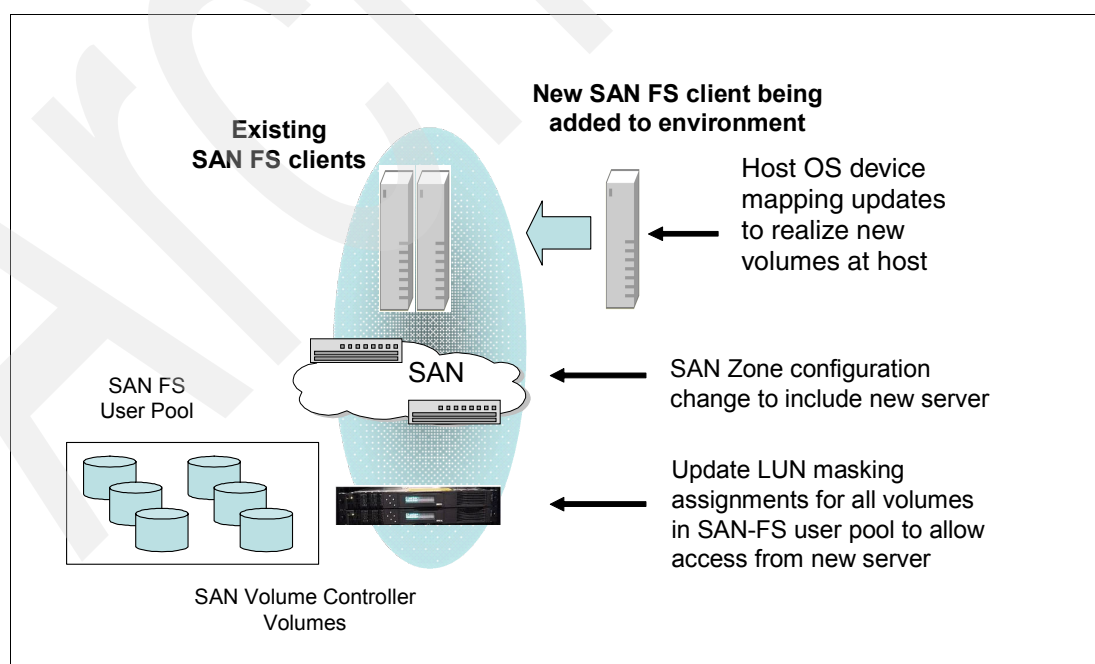


Figure 8-5 Provisioning tasks to introduce a new SAN File System client

When an administrator wants to add a new SAN File System client, the SAN zoning must be updated to enable the new host to see the storage subsystem, by adding it to existing zones or creating new zones. The subsystem LUN masking assignments for all the volumes in the SFS user pool must be updated to enable the new client to see the subsystem volumes. Finally, the HBA driver and multipathing software at the host must be reconfigured to map the LUNs presented from the storage subsystem to the operating system and the SAN File System client as physical disks.

While relatively simple for an administrator to perform, these tasks have to be carried out for each volume in a SAN File System User Pool as new clients are added or for each server using the pool when a new volume is added to the pool. The tasks also require coordination to ensure successful execution and to guarantee that prerequisite tasks have been completed. These SAN File System management tasks are ideal candidates for automation using provisioning workflows, reducing administrative effort for day-to-day operation of a SAN environment.

8.3.1 SAN File System automation tasks

These are the principal automated tasks that were created to configure and manage the SAN File System environment:

- ▶ Addition and removal of SAN File System clients
- ▶ Creation and deletion of SAN File System storage pools
- ▶ Creation and deletion of SAN File System filesets
- ▶ Addition and removal of client access to storage pools
- ▶ Addition and removal of space in SAN File System storage pools

Performing these tasks by hand for a large environment could require a significant amount of effort and high potential for human error. Instead, SAN File System combined with provisioning software enables the management tasks to be broken into repeatable steps, automated through workflows. This helps reduce human error and simplifies the tasks of administrators so that they do not need specialized skills to perform complex SAN administrative tasks.

SAN File System workflows

Several workflows are used to manage the SAN File System environment and simplify administration. These coordinate the configuration of the several infrastructure components. The six basic SFS management workflows are:

- ▶ Add and remove hosts
 - SFS_AddHostToSFSPool
 - SFS_RemoveHostFromSFSPool
- ▶ Add and remove pools
 - SFS_CreateSFSPool
 - Create a new SFS User pool with no storage
 - SFS_DeleteSFSPool
 - Delete an SFS User pool, removing all data
- ▶ Add and remove storage space
 - SFS_AddSpaceToSFSPool
 - SFS_RemoveVolumeFromPool

Some of the other workflows which are involved by the basic management workflows are:

- ▶ SFS_CheckClientStatus

- Check to see if SAN File System client is active and talking to MDS cluster
- ▶ SFS_RediscoverLuns
 - Performs SFS LUN rediscover to find new volumes after pool changes
- ▶ SFS_CheckVolumeStatus
 - Performs SFS query to check status of SFS volume

8.3.2 Modelling of SAN File System in TPM

Central to enabling easy management of the SAN File System environment is the ability to model the SAN File System configuration in the TPM DCM. Figure 8-6 on page 144 shows a SAN File System user pool managed by TPM.



Figure 8-6 SAN File System user pool managed by TPM

Within TPM, the SAN File System user pool is represented by a TPM Volume Container. Servers, which have been zoned and masked to be able to see the volumes in the pool, are shown as “Access Servers”. The SAN File System volumes in the user pool are represented as TPM Logical Volumes. To simplify management and enable SAN File System resources to be easily tracked and identified, the volumes created on the SVC use the pool name as a prefix, for example, Webpool_1

Using the TPM DCM to model SAN File System in this fashion simplifies the task of management, since there is a ready list of servers that use and have access to the pool and the list of volumes in the pool. The server list is used as input to the AddSpaceToSFSPool workflow to determine which servers require the newly created volume to be masked to them. The volume list is used as input to the AddHostToSFSPool workflow, to determine the storage volumes which must be masked to the new SFS client so it can access the pool.

Advantages of using IBM TotalStorage for Infrastructure Simplification

In this part of the book, we review the competitive advantages of using IBM TotalStorage and System Storage products to simplify your storage infrastructure.

Archived



Competitive advantages to Infrastructure Simplification using IBM TotalStorage

In this chapter, we discuss the competitive advantages of the architectural design of the IBM TotalStorage products. This is not an exhaustive list and covers the most important advantages.

9.1 Virtualized versus array replication services

IBM TotalStorage portfolio offers numerous capabilities for simplification using standard methods such as consolidation up to more powerful options such as automated management and virtualization. Like all major networked systems before it, including the telephone system and the Internet, virtualization is a natural part of the evolution of IT systems to maintain function and eliminate complexity. There are numerous ways virtualization, particularly block and file virtualization, add to the simplicity of the overall storage environment.

For block virtualization, IBM offers the SAN Volume Controller. Part of its features include the use and management of replication services, which can dramatically improve a traditional SAN. See Figure 9-1 on page 148. Let's paint the picture. Your lines-of-business have important applications. You want to use Point-in-time or Continuous copy services to protect those applications, but you are faced with two challenges.

First, most vendors offer their copy services on their high-end storage servers and they require that both the source data and the target copy are on the same type of high-end storage device. That means you have to buy twice as much high-end storage, driving up cost. Second, in order to build processes for, and integrate applications with these copy services, you have to use the vendor's APIs and each vendor has their own. The problem is that this approach locks you in to a particular vendor, or makes it very difficult to maintain a multiple supplier strategy to help you keep capital costs down. Analysts speculate that these rigid limitations have prevented many enterprises from taking advantage of the tremendous value of copy services for improving application availability.

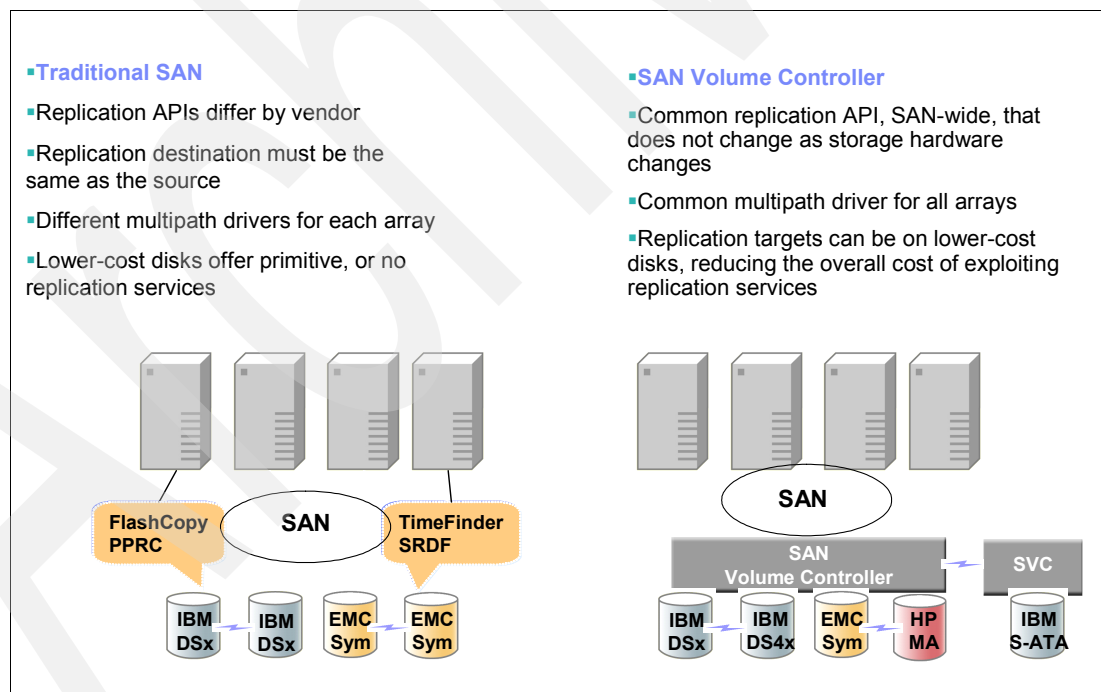


Figure 9-1 Traditional SAN versus SVC

Virtualization reduces the cost and improves the flexibility of replication services. First, it gives you a common API that works across all physical storage devices, meaning you can integrate your processes and applications without the dark cloud of future change hanging over your head. Second, it allows you to use high-end storage for your primary data and lower-cost storage for your target copies, driving down the cost of implementing replication. SAN

Volume Controller even allows some use of replication software in underlying arrays, allowing the software license to expire so that SVC can take over replication responsibilities from the underlying array.

9.2 Non-disruptive data migration

In a traditional SAN, there is a static connection between the host system and the physical disk. See Figure 9-2. In order to move the data that belongs to the application, the application must be stopped to allow the reestablishment of a new set of static connections between the host and the new disk system. Only then can the application be restarted. Lease expiration also can mean making a change in the physical infrastructure, disrupting applications again.

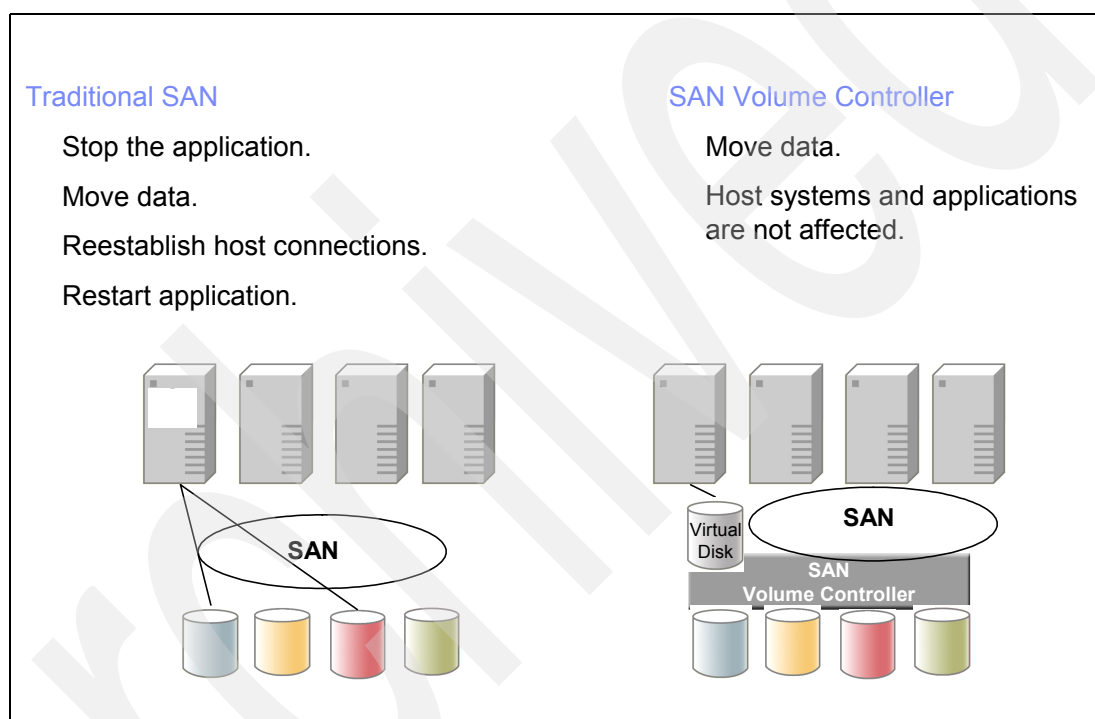


Figure 9-2 Non-disruptive data migration

With a virtualized environment, however, things are different. The host system is dealing with a virtual disk, not the physical disk that went off lease. So, all the administrator has to do is to tell the SAN Volume Controller to remap the data on the virtual disk to another physical disk. The host system and the business application (that is running on the host system) cannot tell that a change was made in the physical infrastructure.

9.3 Storage pooling

SAN Volume Controller can help increase the storage capacity that is available to host applications. Pooling (Figure 9-3) the entire storage network's capacity for DS8000 and other storage devices enables host applications to access capacity beyond their islands of SAN or direct-attach storage capacity. Pooling can also help improve administrator productivity by enabling management at the cluster level, resulting in a single point of control over the entire network.




Figure 9-3 Storage pooling

By pooling storage into a single reservoir, SVC helps insulate host applications from physical changes to the storage pool. This allows applications and, in turn, organizations to continue running without disruption. SVC also includes a dynamic data-migration function that can help administrators migrate storage from one device to another without going offline. This allows administrators to reallocate and scale storage capacity without disrupting applications.

9.4 Virtual volumes

In 2003, IBM introduced and delivered the SAN Volume Controller. With this solution, virtual volumes (Figure 9-4) are presented to the host, and the physical blocks are scattered among different devices from different hardware vendors. Virtual volumes can be shrunken or expanded as needed, and the physical blocks can be moved from one physical device to another with little or no disruption to applications.




Figure 9-4 Virtual volumes

9.5 Big Performance and Capacity, Smaller Package¹

Where performance and capacity have been enhanced in the new TotalStorage DS products, floor space, density and complexity have been significantly reduced. Consolidation scenarios benefit strongly from having dense storage, particularly where floor tiles command premium rent, such as in popular Telco hosting sites and financial centers. Packing multiple enterprise storage features into a tiny package, the DS6000 series is designed to achieve up to nearly twice the scalability of its refrigerator-sized EMC DMX800. The contrasts are marked and impressive: the DS6000 requires as little as 4 percent of the volume at 5 TB that the EMC DMX800 requires and is one-tenth of the weight. The DS6000 fits as much as 1.9 times more storage per square foot than EMC's DMX3000 and as much as 3.0 times than that of the DMX800. Interoperability with the DS8000 for copy services offers the combined benefits of these two storage systems.

Comparison

- ▶ 4% of the space for comparable enterprise offerings at 5TB
- ▶ 1/10 the weight compared to IBM biggest competitor
- ▶ 1/4 the power consumption compared to IBM biggest competitor 5TB
- ▶ 1/15 the entry point capacity compared to our biggest competitor (as measured in minimum number of disks)
- ▶ Almost 2X the scalability of our biggest competition
- ▶ Up to 1.5X the performance compared to our competitors

¹ TotalStorage Magazine February/March 2005

9.6 Based on the best²

The rack-mountable DS6000 has a physical storage capacity of up to 65.4 TB in its compact 5.25- by 19-inch, 125-pound frame. “An EMC box like the DMX800 has the same capacity, but it’s 75 by 24 inches and weighs nearly 1,600 pounds.”

The DS8000 consumes 30-percent less floor space than the ESS and extends the relationship between the TotalStorage line and POWER* technology. POWER5 technology brings partitioning capabilities to the DS8000, which is the first POWER5 technology-based storage server. The optional Virtualization Engine features available with the 4-way model are designed to allow customers to run two separate “virtual” storage subsystems within a single DS8000 server.

The rack-mountable DS6000 has a physical storage capacity of up to 65.4 TB in its compact 5.25- by 19-inch, 125-pound frame

9.7 SAN File System Architecture

The following are the key Architectural advantages of the IBM SAN File System compared to our competitors (Figure 9-5).

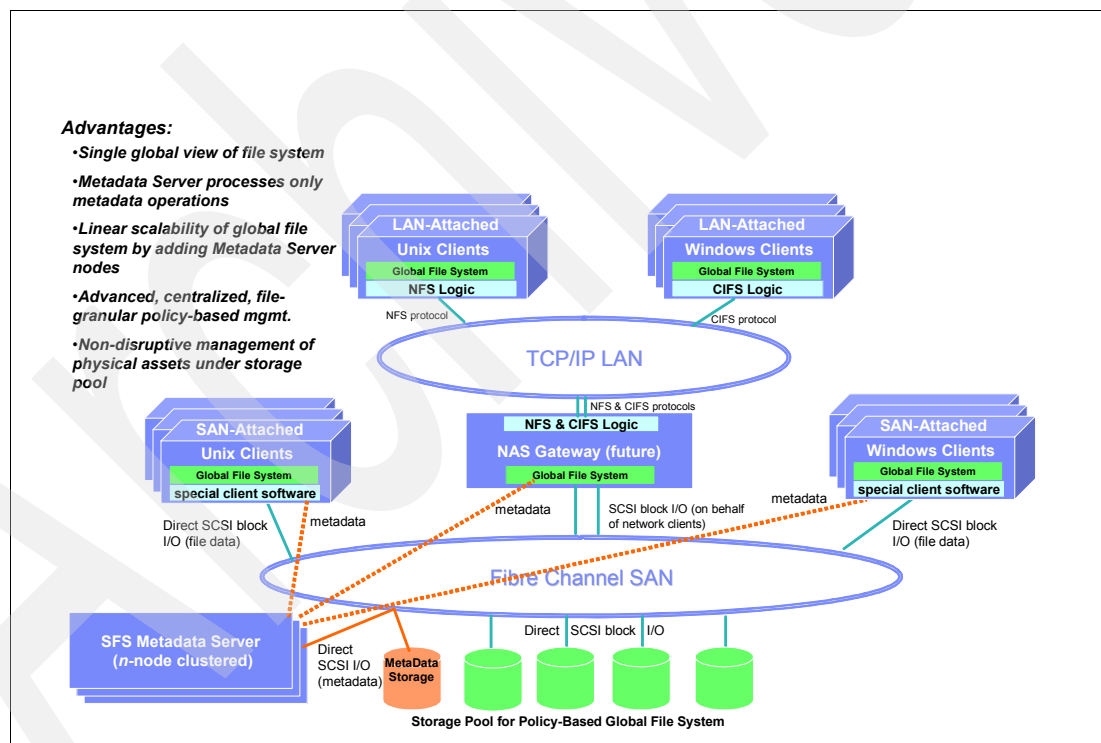


Figure 9-5 SAN File System architecture

² TotalStorage Magazine November/December 2004 Neil Tardy

Advantages

Single file system with one name space

SFS provides common name space, a single file structure which all participating host systems and applications can plug into (Figure 9-6) which enables a single point of management Figure 9-7 on page 154.

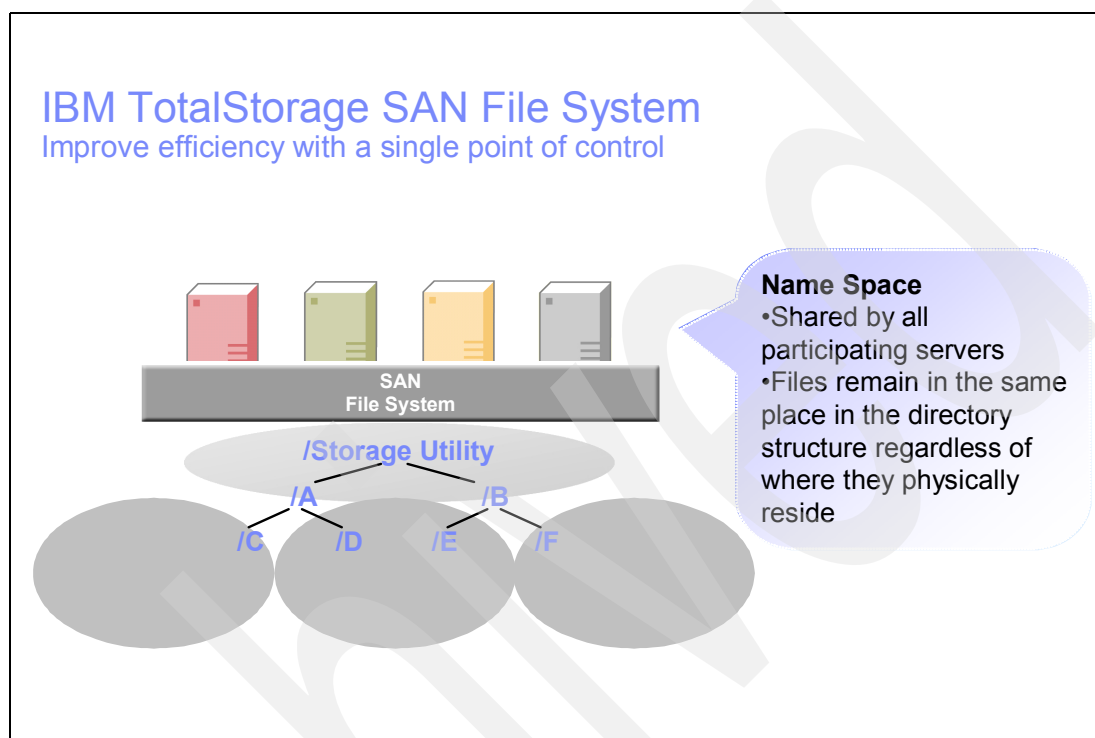


Figure 9-6 Single name space

IBM TotalStorage SAN File System

Multi-Network, Multi-Tier, Centralized Data Management

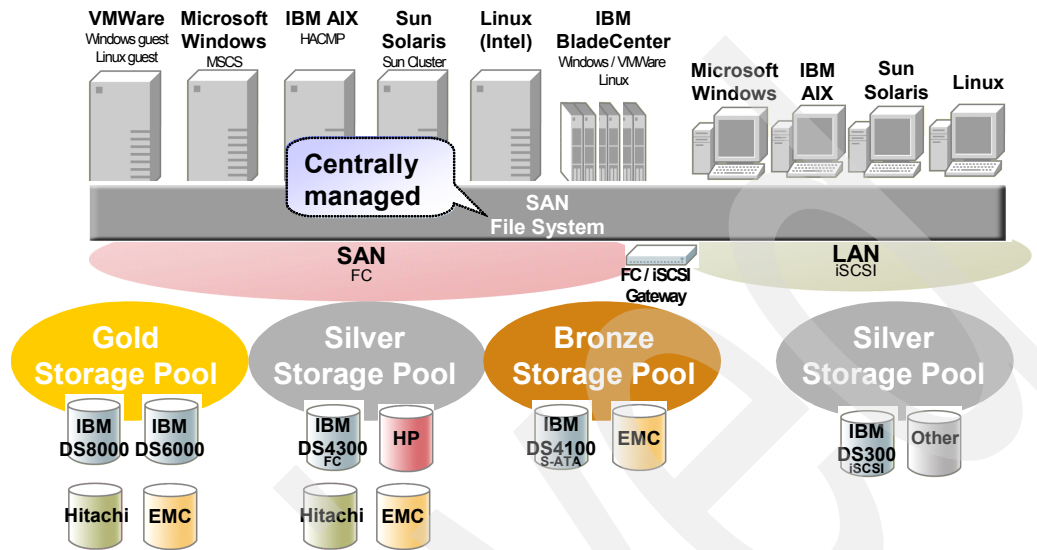



Figure 9-7 Single point of management

Heterogeneous file sharing

SAN File System allows flexible, secure sharing of files between different platforms (Figure 9-8). The user map identifies a Windows user and a UNIX user that are to be treated as equivalent for the purposes of checking file access and permissions. You can share a file among Windows and UNIX users and be very specific about which users you want to be able to access the file.

IBM TotalStorage SAN File System

Secure sharing of files between specific Windows and Unix users



Heterogeneous File Sharing

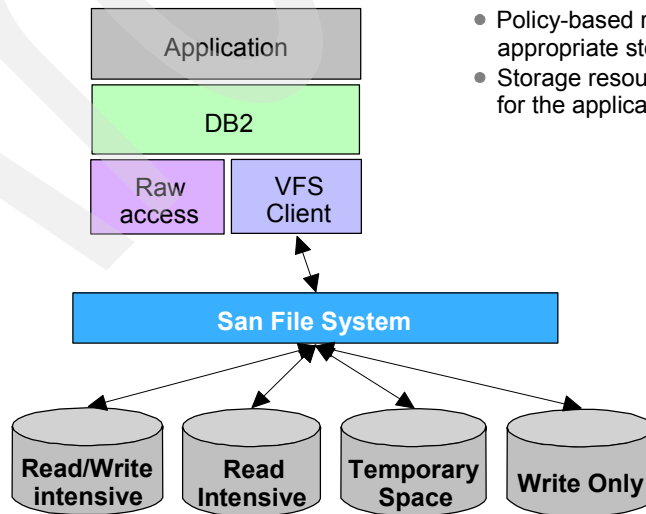
- Reduce storage needs by eliminating need to maintain multiple copies of files for data sharing
- Streamlined turnaround times

Figure 9-8 Heterogeneous file sharing

Policy-based placement of files

With the categories of data understood (and as a result the different values of information) and storage resources pooled, you can now implement automated policies. Automated policies help ensure the right categories of data are stored on the right pools of capacity. You are also automating the process of matching the value of information to the right cost of storage (Figure 9-9).

Policy-based File Placement



- Policy-based rules map data into appropriate storage pools
- Storage resources can be optimized for the application's requirements

Figure 9-9 Policy-based file placement

Non-disruptive provisioning of physical storage

The SAN File System allows a storage administrator to allocate all storage to the SAN File System. When a SAN File System host needs additional space, the SFS administrator issues a simple command to increase the amount of storage the SAN File System host can use (you can also decrease the amount of storage a SFS host can use as well). This can be done concurrently without rebooting the SFS host (Figure 9-10).

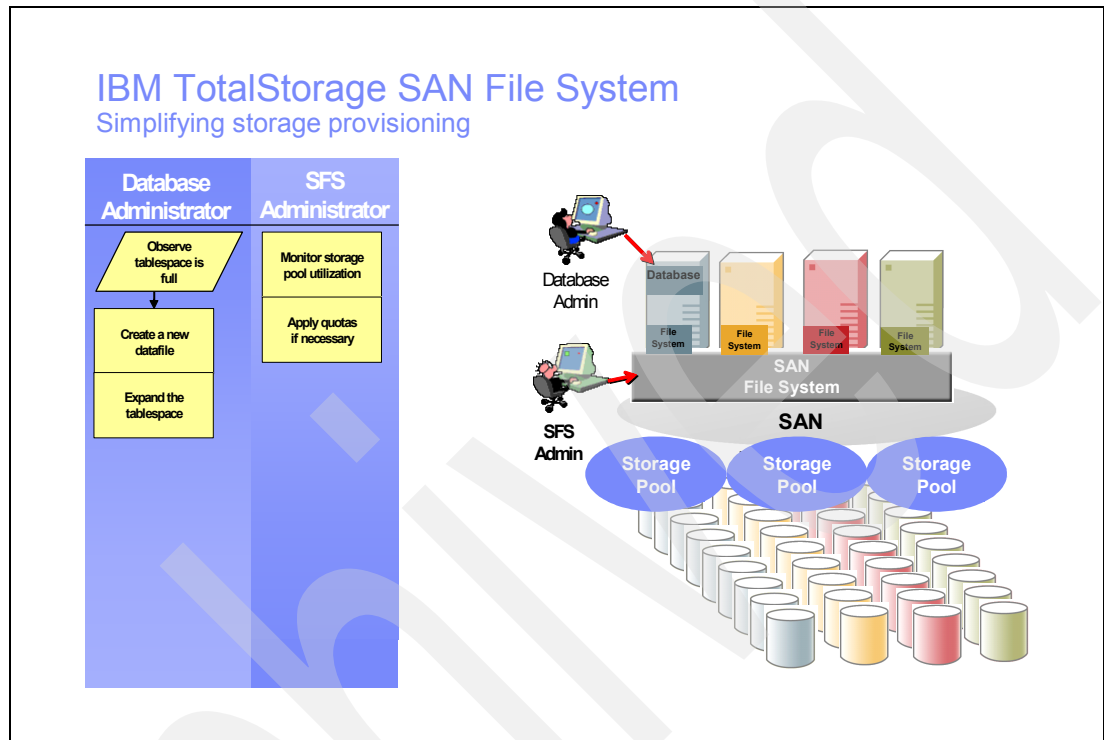


Figure 9-10 Storage provisioning



Part 5

Storage environment consolidation

In this part, we discuss how IBM TotalStorage Products can help consolidate your storage environment. We also provide an example of a real life solution to consolidation using IBM TotalStorage Products.

Archived



Consolidation principles

In this chapter, we discuss how consolidation of your storage environment can lead to more efficient management and utilization of resources.

10.1 Storage consolidation

In the early 1980s, data sharing emerged as a critical necessity for improving efficiency within enterprises. This term is used somewhat loosely. It is sometimes interpreted to mean the replication of files or databases to enable two or more users, or applications, to concurrently use separate copies of the data. The applications concerned may operate on different host platforms.

Data sharing can also be used to describe multiple users accessing a single copy of a file. This is called *true data sharing*. In a homogeneous server environment where all servers run the same OS, with the appropriate application software controls (for example, a clustered file system), multiple servers can access a single copy of data stored on a consolidated storage subsystem. If attached servers are heterogeneous platforms (for example, a mix of UNIX and Windows), sharing data between such unlike operating system environments is complex. Storage consolidation and management of information are important requirements to effectively share information.

In the late 1990s, storage networking emerged in the form of SANs, Network-Attached Storage (NAS), and Internet Small Computer System Interface (iSCSI). These technologies were aimed at reducing the TCO of storage by managing islands of information among heterogeneous environments with disparate operating systems, data formats, and user interfaces in a more efficient way.

SANs enable you to consolidate storage and share resources, because storage capacity can be connected to multiple servers, and at a greater distance than smaller, direct-attached disks. By separating storage resource management from individual hosts, a SAN enables disk storage capacity to be consolidated. The results can be lower overall costs through better utilization of the storage, lower management costs, increased flexibility, and increased control. This storage consolidation can be achieved physically or logically.

Physical storage consolidation

In this approach, you take two or more pieces of hardware and consolidate them into fewer and/or bigger hardware. In Figure 10-1, you can see a traditional SAN environment.




Figure 10-1 Traditional SAN with distinct storage islands

Data from disparate storage subsystems can be combined within shared disk arrays, which may be located at some distance from the servers. The capacity of these disk arrays can be shared by multiple servers, and the array capacity may be partitioned and zoned, so that each server has access to an appropriate portion of the available volumes. In Figure 10-2, you can see a consolidated SAN environment.




Figure 10-2 Consolidated SAN environment

Users may also benefit from the advanced functions typically offered with such subsystems. For example, RAID capabilities, remote mirroring, and instantaneous data replication functions may not be available with smaller, direct-attached disks. In Figure 10-3, you can see a consolidated storage environment.




Figure 10-3 Consolidated storage environment

Available capacity can be dynamically allocated to any server requiring additional space. Capacity not required by a server application can be reallocated to other servers. Dynamic capacity allocation avoids the inefficiency associated with free disk capacity that is attached to one server and not being used by other servers. Extra capacity can be added, in a non-disruptive manner.

Physical storage consolidation does not mean that all wasted space concerns are addressed, but it is a significant step forward in available capacity management.

For example, one DS6000 can deliver the performance and capacity of two ESS F20 (see Figure 10-4 on page 164, or one DS8100 can deliver the performance and capacity of four ESS F20s, or one DS8300 the performance and capacity of eight ESS F20s (see Figure 10-5 on page 165).

You can save hardware and maintenance costs by physical storage consolidation. Storage consolidation allows you to leverage the benefits of clustering technology and increasing application availability. The administrative tasks become easier as system managers have fewer pieces of hardware. You can save floor space, power, and cooling capacity. But you have to deal with the security risks introduced by storage consolidation, and the resulting multiple access points with wider accessibility to sensitive stored data.

Logical storage consolidation

It is possible to achieve shared resource benefits from the SAN without moving existing equipment. A SAN relationship can be established between a client and a group of storage devices that are not physically co-located (excluding devices that are internally attached to servers). A logical view of the combined disk resources may allow available capacity to be allocated and reallocated between different applications running on distributed servers, achieving better utilization.

In this approach, you are differentiating the physical resource and the logical representation of that resource. In this environment, the upper layer no longer deals with the underlying physical implementation. The upper layer deals simply with a generic, logical layer underneath. The physical storage consolidation can help you simplify your storage infrastructure. But, by the growth of the data being stored and the growth of the requirements and services you may need, your enterprise could continually be evaluating and undertaking logical storage consolidation.

10.2 Benefits of consolidation

Consolidation is the leveraging of advances in storage density to consolidate many older devices into fewer or one newer device. Some benefits of consolidation are:

- ▶ Cost savings
- ▶ Scalability with large capacity products or scalability out with rack and stack products
- ▶ Fewer elements to manage

10.3 Key IBM storage consolidation products

Clearly physical storage consolidation can help you simplify your storage infrastructure. But in some ways, with the growth in data being stored and the advancements in storage density, consolidation is an activity that you can continually evaluate and undertake. IBM has the products and devices to assist with this ongoing task:

- ▶ SAN products
- ▶ IBM Network-Attached Storage gateways
- ▶ IBM N Series systems
- ▶ Tape libraries
- ▶ Disk storage servers
- ▶ IBM TotalStorage SAN Volume Controller software
- ▶ IBM TotalStorage SAN File System software
- ▶ Migration and Installation Services

10.4 Storage consolidation examples

The following examples compare the physical footprint of the depicted storage systems. What you should also consider are maintenance and environmental costs of the installed older generation storage. In determining your storage consolidation savings, also consider whether the assets being replaced have been fully depreciated (or are past lease expiration) such that there are no further costs associated with acquisition. It should be noted that ESS Model F20 has hardware maintenance charges only which include advanced function support. With reference to Other Equipment Manufacturer (OEM) hardware, we are comparing against our leading competitor of storage arrays. Note there are costs associated with hardware and advanced function software for these OEM storage arrays.

10.4.1 DS6000 storage consolidation

The IBM TotalStorage DS6000 is the first storage system to offer enterprise-class functionality at a price point required by many midsize businesses. With its rack mountable package and modular, pay as you grow scalability, the DS6000 delivers enterprise functionality at a fraction of the cost of other enterprise-class storage systems. The DS6000

can replace multiple ESS E and F models or older OEM subsystems, saving you money and space. See Figure 10-4.




Figure 10-4 DS6000 storage consolidation

10.4.2 DS8100 storage consolidation

IBM TotalStorage DS8000 is the new IBM standard in vertical scalability and represents an ideal storage consolidation platform for large enterprises. One of the innovative things IBM has done in designing the DS8000 is converge many of scalable technologies that it has pioneered in mainframe and POWER-based computer systems with its storage systems. The DS8000 is built around the IBM POWER5 processors and introduces Logical Partitioning (LPAR) technology to the storage industry. The DS8000 can replace multiple ESS E and F models or older OEM subsystems, saving you money and space. The following example shows a comparison of one DS8100 versus two installed ESS F-20s or older OEM storage arrays. See Figure 10-5.




Figure 10-5 DS8100 storage consolidation

The DS8000 is an ideal platform for consolidating existing storage assets into a single, smaller footprint, simplifying the overall storage infrastructure. The DS8000 can replace multiple ESS E and F models or older OEM models. The following example shows a comparison of a DS8100 versus four installed ESS F-20s or older OEM storage arrays. See Figure 10-6.




Figure 10-6 DS8100 storage consolidation versus ESS F-20

The DS8000 is also ideal for consolidating workloads from both distributed systems and mainframe systems. Historically, many enterprises have maintained physical separation between these environments. With DS8000 storage system LPARs, these workloads can be physically consolidated while still maintaining logical separation.

The DS8000 can replace multiple ESS E and F models or older OEM models. The following example shows a comparison of a DS8300 with storage system LPARs versus eight installed ESS F-20s or older OEM storage arrays. See Figure 10-7.

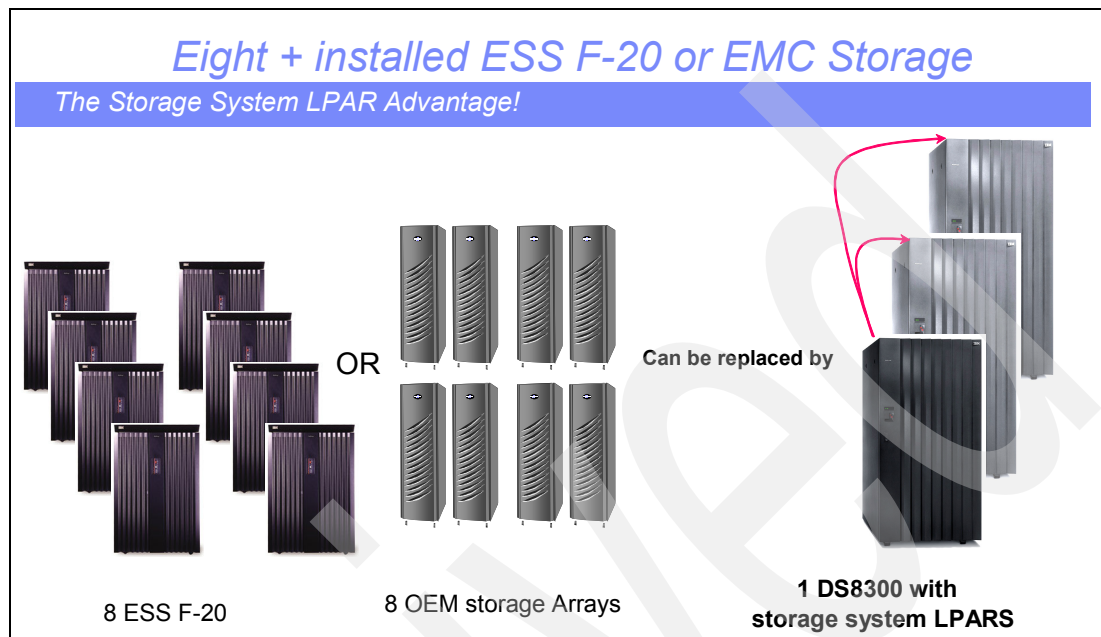


Figure 10-7 DS8300 storage consolidation

10.4.3 SVC promotes consolidation

With a virtualized environment, all the available capacity can be pooled behind the SAN Volume Controller. The host systems think they are dealing with a single, very large disk controller. Your administrators can much more easily reassign unused capacity simply by remapping physical capacity to virtual disks that need more space. As a result, the overall utilization of your physical disks can be increased and future purchases deferred. See Figure 10-8.

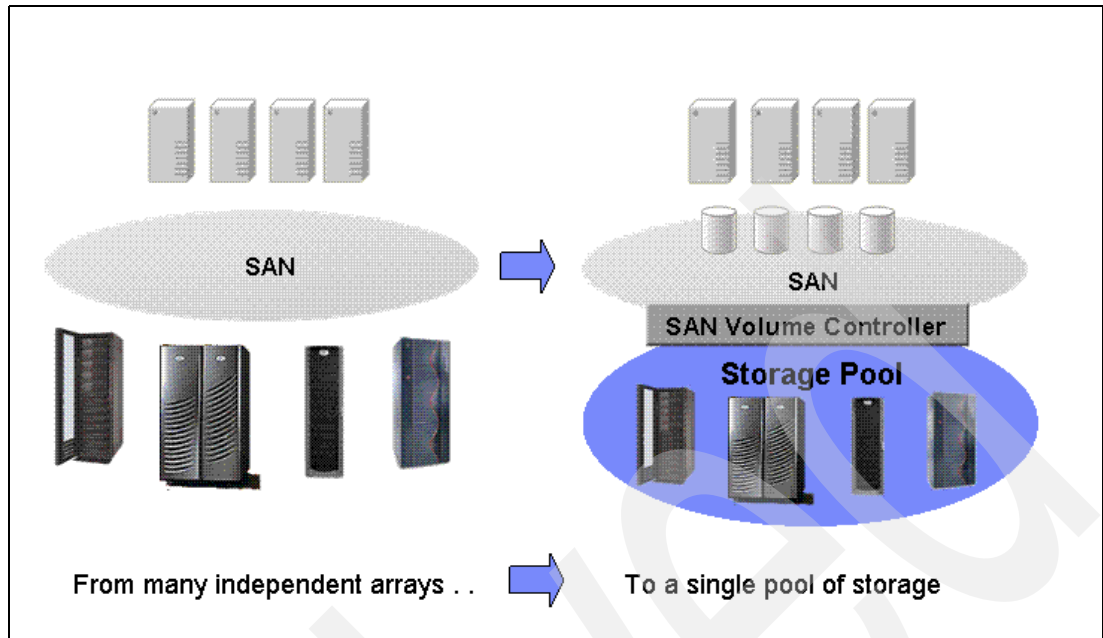


Figure 10-8 Array pooling

Archived



Infrastructure Simplification in action

This chapter, given from a customer perspective, describes a TotalStorage Infrastructure Simplification project that was done in a major health care center. The healthcare center is called Customer H in this example. This chapter discusses the environment, process, and results. Actually, the steps to success in this consolidation and simplification effort by Customer H reflect many of the principles outlined in “Infrastructure challenges” on page 7.

11.1 Introduction

As with many business environments, Customer H's use of information technology is growing by leaps and bounds. An area where growth is most dramatic is data storage. Customer H's data storage needs are growing by 100 to 200 % each year, a rate that if left unmanaged could lead to serious issues in the areas of Disaster Recovery and migrations. With this growth came a massive increase in the need for improved IT infrastructure. For this reason, Customer H embarked on a project to implement a storage area network infrastructure.

11.2 History

Prior to implementing a storage area network infrastructure, most of the storage was directly attached to the server that utilized it. Previously, Customer H initiated storage consolidation with the purchase of an IBM ESS that included 400 GB of storage and eight SCSI adapters (Figure 11-1). Their purpose was to simplify UNIX system storage by migrating all of the UNIX systems off of multiple direct-attached SSA and SCSI disk cabinets into a single managed platform. This consolidation project was a success and a positive introduction to management staff to the idea of storage consolidation.

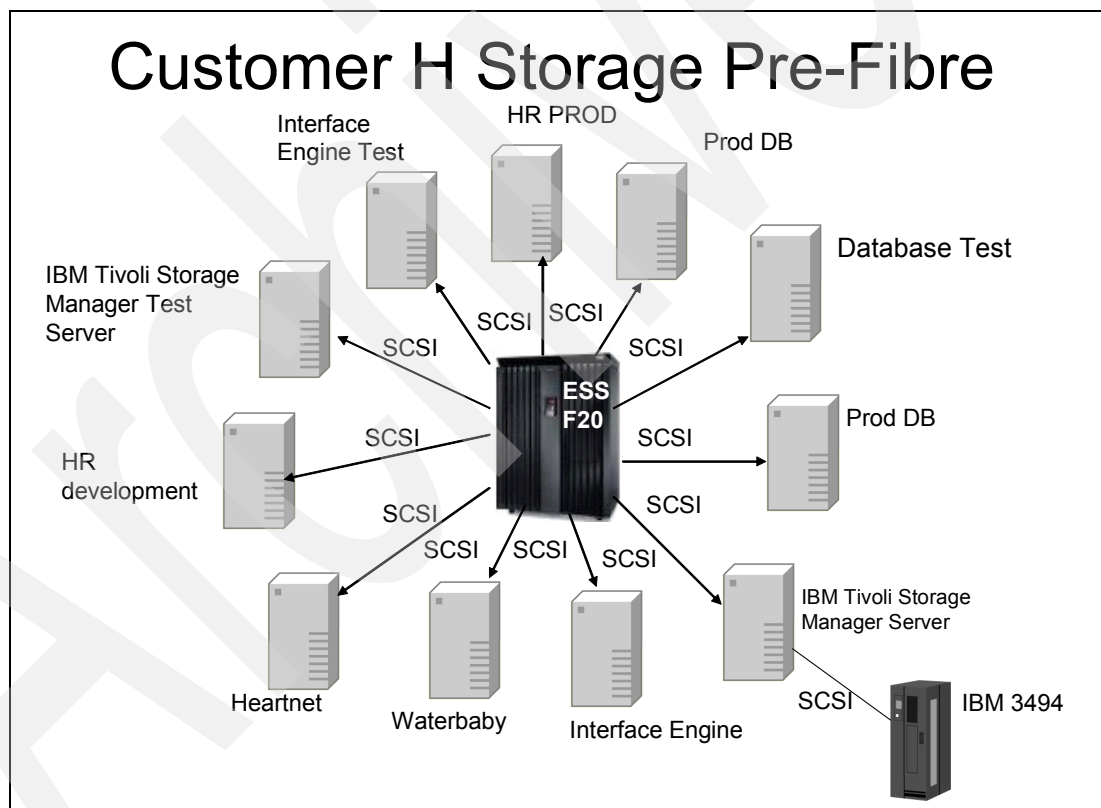


Figure 11-1 SCSI attachment

By 2000, the IT department had successfully placed 11 UNIX servers onto the ESS, but they were also beginning to notice issues with I/O performance on their financial DB server. Along with this, there was a major upgrade of the financial system that would increase the number of users 10-fold as well as the amount of data being recorded. After analyzing the situation, they determined that the SCSI connectivity was the bottleneck and that they needed improved performance to handle the impending upgrade and user growth. This led to the

company's first test of Fibre Channel SAN in an open environment and the purchase of two IBM 2109 8-port SAN switches, which offered 1 Gb performance (Figure 11-2).

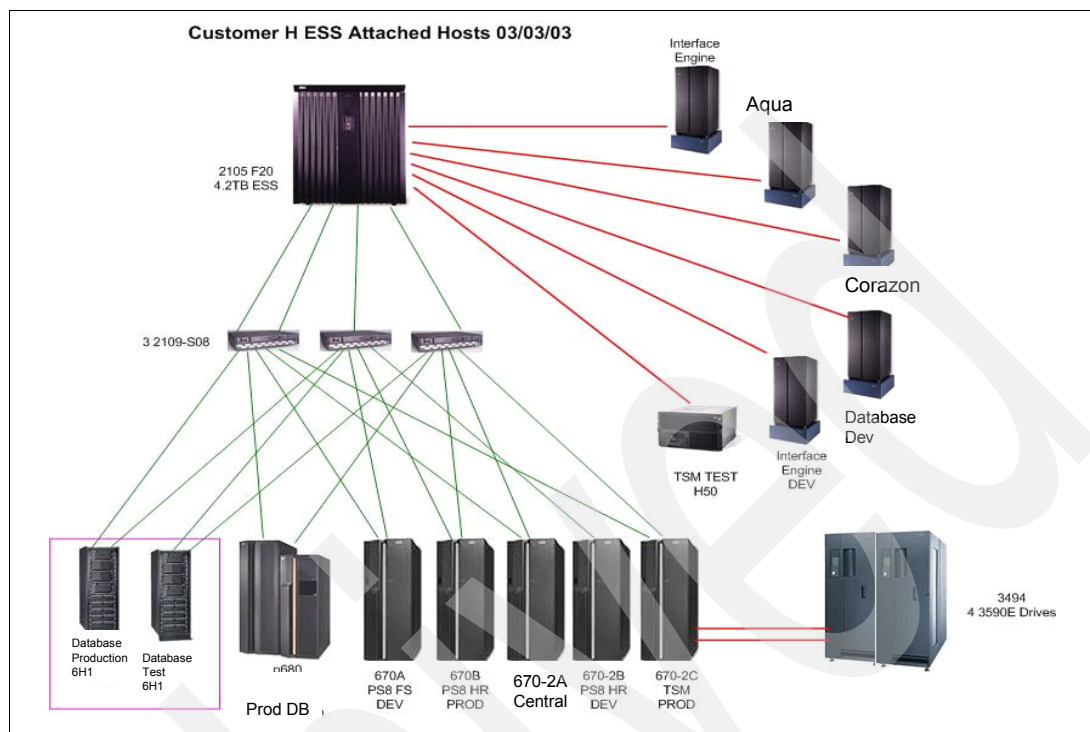


Figure 11-2 First SAN

By implementing Fibre Channel, the throughput capabilities were increased from 60 MB/s to 100 MB/s, a 67% increase. Overall, the I/O performance was improved on the server. This again showed management that implementing a SAN and storage consolidation had benefits to the organization. This success also increased the demand for connectivity to the SAN and the ESS. However, Company H needed to grow the SAN and ensure better availability before the company would risk attaching additional systems.

With the success of the small storage area network, many of the operations staff wanted to take the system to the next level, implementing in their environment a full scale SAN to which any system could attach. To do this, the storage administrators needed to come up with a plan to prove the value of the capital investment to the organization. Based on the concept and previous successes, management gave the approval to move forward. As part of this agreement, it was decided to phase the project in order to prevent unnecessary issues and not overwhelm themselves.

11.3 Phase One research and design (Don't rush it)

Phase one of the organization's plan was to determine what the current needs from a storage perspective were. To make this easier, the project was broken down by server platform, and a consulting group that specialized in storage area networks was brought in to assist in the design. The overall goal was to find a cost-effective solution that would meet the growing needs and concerns of the organization.

Prior to having the consulting group involved, the storage team compiled a list of goals and requirements in order to develop a base roadmap for the project. This was done so that they

could maintain control of the plan, and to keep the consultants in line with what needed to be accomplished. Included in the list were:

- ▶ Availability and resiliency
- ▶ Minimal components in order to reduce the complexity
- ▶ Scalability and upgrade ability
- ▶ Ease of use
- ▶ High performance
- ▶ Reduced TCO
- ▶ Improve personnel productivity
- ▶ Improve backups and restores

The consulting group was able to bring useful information to the table. Together, there was a good understanding of storage area networks, and the technology that was involved. The consultants also had vendor contacts and could get better information when requested. The storage team and consultants spent approximately 40 hours together on-site going over the current infrastructure, expected growth, and current marketplace. Two weeks after the meeting, the consultants returned to present a plan to the project team.

A best practice that all management and storage administrators must follow closely is that they should always perform research on their own, and never fully rely on consultants or vendors.

Consultants and vendors tend to have comfort zones. They guide you in that direction versus looking at newer technology in the marketplace. As part of the project, the storage teams performed their own research and were contacting vendors to validate information and technology. In doing this, the teams were able to find a director class SAN switch that was new to the marketplace, but was a better fit based on the predetermined requirements.

The staff was initially focusing on three major switch vendors. With two of the vendors, there were concerns with their licensing, a need for additional software (creating costs and complexity), and what were determined to be single points of failure. The third vendor, Cisco, offered them a director class device that met core requirements, meeting a key requirement of zero downtime for upgrades. However, this issue was that the device was new and that very few customers had it in a production mode. To help quell concerns, Cisco offered them a demonstration device so that testing could be done to see if it met their standards. Because the customer was limited by time and had a concern about availability, a basic test plan was created and used to test the resiliency of the Cisco MDS 9509 director. The test included:

1. Test for resiliency, remove parts
2. Perform software upgrades
3. Data transfers
4. Product ease of use
5. Vendor support

During the Proof Of Concept (POC) testing, they made efforts to run through all possible failure/performance scenarios for the SAN director. Testing included removing and reinserting parts, cables, and power supplies. They also tested upgrading the OS on the switches while performing transactions. No system issues were detected, the management software was very easy to use, and support from Cisco was excellent.

As with any product, the product is only as good as the support behind it. An excellent product with poor support does not merit more investment versus a decent product with excellent support. From the start, IBM and Cisco made sure all the customer's needs were met, from both the design perspective to the implementation. This increased the customer's comfort level with the new product, and gave management the satisfaction that they had made a good

investment. Over two years later, Customer H still says they are extremely satisfied with Cisco's support for their SAN devices.

11.4 Product acquisition

Initially, the first purchase was a single 9509 with 64 ports to add into the existing SAN. The plan was to move half the ports connected to the existing SAN to the 9509, and then purchase the second director at the beginning of the next fiscal year. In reality, this was not the best plan, because it created redundancy issues along with a point of failure in the cabling, which caused a delay in completing the migration and purchase of a second director. Once the second director was purchased, the implementation of the original design was accomplished. See Figure 11-3.

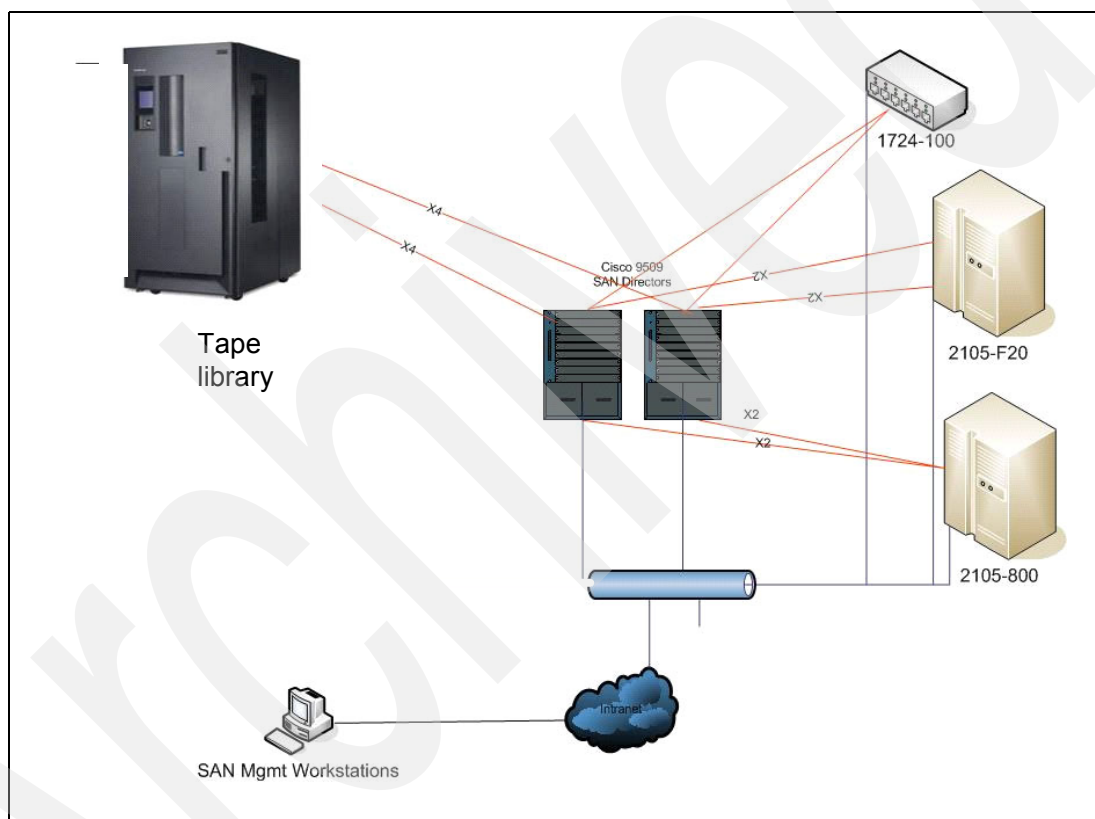


Figure 11-3 Full scale SAN implementation

11.4.1 Configuring the switch

One of the more difficult tasks in implementing a SAN environment is deciding how to configure the switch, for example, zones and VSANs, and determining the proper level of security. Due to the small scale of the SAN, the customer chose to stick with a single VSAN and to zone by World Wide Port Name (WWN) versus port. This gave the flexibility to move cables if necessary without having to modify the zone configuration if there were a port issue. Of course, there are pros and cons to this approach. See Table 11-1.

Table 11-1 Pros and Cons

Pros	Cons
Ease of setup	Less secure (Spoofing).
Ease of port migration	If an adapter on a host fails, then you need to add the host to the zone and have it log into the fabric (may require a reboot of the host).
Easier to maintain	Persistent ID issues.

Another decision that was made in the configuration discussion was to keep the two switches as separate devices, meaning no Inter-Switch Links. By selecting this design, it protected them from fabric corruption or other issues spreading from one switch to another. It also gave them the freedom to upgrade the switch code in a more controlled measure, because the switches would not be required to be at the same exact level. Simplification principles normally dictate reducing differences in the network; however, in this case, the management preference required this configuration.

When designing a fabric, setting up internal policies as to what and how things are added to the SAN is very important. If done correctly, a SAN is a very cost-effective, hearty environment, but like anything else, if there are no controls put in place, not only can performance be affected, but the TCO/ROI for it can be easily decimated.

Some of the policies had to do with which types of servers can connect to the SAN. Servers that have persistent data, or that are non-clustered with storage and need less than 100 GB, may not need to be on the SAN. Policies on cable management also were put in place to avoid “spaghetti” cabling and mishandled cables. SAN redesign is an opportune time for reconsidering and simplifying cabling issues as well.

In the past, cables were run either under the raised floor or in overhead cable trays from the server directly to the SAN switches. This was an issue since cables were more fragile than copper, and if a cable needed to be reclaimed, not only could it be difficult to pull a long length, it could lead to an outage caused by pulling the wrong cable, or catching another cable. To reduce this risk, Customer H strategically placed Fibre Channel patch panel throughout the main data center, limiting the longest cable pull to 24 feet (including the rack height and the necessary slack for cable bends). With the use of patch panels, you need to be aware that one side needs to be uncrossed, in effect, A-A, B-B. With this in mind, Customer H implemented a policy to custom order all cables to eliminate confusion, as well as excess slack that could get caught and tangled. A cable labeling strategy that made it easy to identify a server's cable was also implemented. Figure 11-4 below is an example.

Example:

- ▶ SAN Switch A LC1 Port 5 SAN switch, line card, and port
- ▶ Patch Panel 1 Port 12 Patch panel and port
- ▶ Exchange FC1 Server name and Fibre Channel card



Figure 11-4 Cable strategy

11.4.2 Lessons learned troubleshooting cable issues

When Customer H first began implementing the SAN, sometimes the host would not connect to the port on the switch. Initially, the team would spend hours troubleshooting the issue by rebooting the server and replacing cards, until finally, trying a new cable, which would work. After the fourth time of experiencing cable issues, the team came up with an idea to begin to use a regular flashlight to check the cables for flaws. This worked satisfactorily, but white light is hard to see. An \$11 LED light solved that problem. The red LED light is easy to see, and today, all cables are checked prior to pulling them, eliminating hours of troubleshooting.

Things to look for when checking a cable are:

1. Does the cable have any cuts or breaks?
2. Are the ends polished correctly and cut straight?
3. Are the cables configured for send receive, for example, A-B, B-A?
4. If you are using a patch panel, is one set of cables A-A B-B? Otherwise, the send-receive will be wrong.

Attention: Never stare directly into the cable or the SFP. Use a small piece of white paper or a mirror.

11.4.3 The core-edge debate

Astute readers may wonder why not ISL the two switches together. The answer is easy. First, to have redundancy and high availability, you need to have multiple paths between each switch, which means higher port counts. Secondly, interoperability and change management are processes to simplify as well, such as changing the code on one switch type. Third, if one switch has an issue, then it could broadcast it via the ISL link to the other switches, amplifying

the issue. Is there a place for the core-edge design? Yes, in the case where you need to manage a SAN where departments want separate physical devices, this design is useful, although not always as affordable.

Issues with the core-edge design:

- ▶ Requires additional ports and cabling
- ▶ Not as resilient
- ▶ May require interoperability
- ▶ Issues can arise with device code upgrades
- ▶ Makes troubleshooting more difficult
- ▶ Just not simple!

From the organizational view, the ability to zone along with the ability to create virtual SANs (VSANs), eliminated the need to go core-edge, as well as eliminated the need to have the SAN switches ISL'ed together.

11.4.4 Spending money is spending money! Finding the ROI in building a SAN

As part of the research by the organization, they found that the average utilization of the storage on servers was less than 30 percent. Worse yet was the knowledge that none of the storage could be reallocated to other servers that may need the storage space. Also, as servers needed to be replaced, the storage often could not be reused because newer equipment no longer supported that technology. This lack of flexibility became a large expense as the environment grew and storage needs increased.

In a small environment, having servers with direct-attached storage (DAS) may be justified, but as the environment grows, managing hundreds of servers with Direct Access Storage is a management nightmare. In the Customer H environment, this could add hours of work onto a single system administrator by forcing them to manage the following:

- ▶ Drive type
- ▶ Licensed internal code levels
- ▶ Driver level
- ▶ Controller licensed internal code levels
- ▶ RAID management
- ▶ Storage performance
- ▶ Storage growth
- ▶ System availability
- ▶ Backups and restores

As part of the initial project, a basic cost analysis was performed for storage in the environment. Based on the utilization average and an average price of \$1,000 per disk drive, the organization was wasting \$700 per drive, with an average of five disk drivers per server that added up to \$3,500 per server. Multiply this by 100 Intel servers alone in the year 2000 and there was at least \$350,000 in storage that was not and could not be utilized. Measuring by Full Time Equivalent (FTE) which is a way to measure a worker's productivity and/or involvement in a project. The additional hours that each FTE each employee put in managing just the storage component of the servers and the cost of managing Direct Access Storage (DAS), \$5,000 per server per year, and the costs begin to approach the one million dollar mark. See Table 11-2.

Table 11-2 Costs of storage

Component	Costs
Disk drive	\$1,000

Component	Costs
Utilization on a dollar basis (30%)	\$300
Unused space in dollars (70%)	\$700
Average drives per server	5
Unused space per server in dollars	\$3,500
Number of servers year 2000	100
Total unused space in dollars	\$350,000
Storage management costs per server	\$5,000
Total storage management costs year 2000 100 servers X \$5,000	\$500,000

The environmental aspects are another is advantage of DAS. In order to accommodate DAS, many servers were oversized, meaning that only four servers could be placed in a rack. Over time, this created a space shortage in the primary data center, as well as power issues. These larger servers were also often under-utilized, which meant that the company was spending more money than necessary on server equipment to meet the storage needs (Figure 11-5). Conversely, use of small servers (see Figure 11-6) would save expenses.



Figure 11-5 Large servers




Figure 11-6 Small servers

Direct-attached storage's architecture and storage typically is not always as resilient as that of SANs. In a business environment, the availability of data is critical. To protect data and add levels of availability to it, many companies implement some level of RAID to the storage arrays, typically RAID 1, 5, or 10. This protects the data against a disk failure, but does not protect against controller failures or other non-storage hardware failures. As the data stored in the various IT applications became more critical, Customer H needed to look at how to protect their data better and make it more available. They did this by looking at how data was stored and where points of failure were. The findings helped justify the change in how they implemented their IT infrastructure in the future.

Some key findings of this study were:

- ▶ The applications were not highly available.
- ▶ RAID controllers for DAS in the servers were a single point of failure.
- ▶ Clustering was not implemented due to a poor ROI on external arrays for individual projects.
- ▶ A lot of downtime was required for storage upgrades due to the lack of redundancy.

With most SAN arrays, redundant components come standard as part of their design. Due to this design, outages both scheduled and unscheduled are infrequent as most tasks can be done online. Clustering of servers can also be done very easily as storage arrays are designed to allow the attachment of multiple hosts and LUN zoning, therefore, improving the TCO of the storage array and improving the ROI for a clustered application.

Performance is another area where Customer H improved with the implementation of the SAN, enough so that they have actually provided a definite return on investment for applications. Although adding a SAN into the mix adds "hops" and protocol conversions into the mix, performance using SAN-attached disk has been much better than using DAS. The reason this occurs is two-fold. First, most disk arrays have RAM-based caching which allows the writes to be committed quickly to the array, and similarly so for reads. Secondly, by using SAN-based arrays, data can be spread much easier, making sure that no disk runs too hot. With the migration of backups to the SAN, they can not only reduce the load on the IP infrastructure, but also reduce the time the backups run on the servers via the higher I/O throughput. Also, if they utilize snapshots for their backups, they can reduce the I/O overhead as well as the CPU overhead that is normally utilized when the backup client runs to virtually nothing.

By simplifying the storage environment, the staff needs for managing it also reduce. Yes, the cost of storage itself is decreasing, but left unmanaged, the cost of managing storage equipment is increasing. IDC stated that in 1997 the average storage manager handled about

750 GBs of storage. In 2001, another IDC study found that the average employee was handling about 1.3 TBs of storage. By 2004, the study, which you can read at the following Web site:

http://searchcio.techtarget.com/tip/1,289483,sid5_gci769891,00.html

indicated the average employee would be handling about 5.3 TBs of storage. If this were left to individual system administrators managing direct-attached storage, not only would it be overwhelming from a scale perspective, but the management of the storage alone would eat up a large portion of the person's time. If you were to consider the number of different storage drivers in a DAS environment, as well as the number of different storage devices, it would be very enlightening and understandable why a simplified SAN is more cost-effective. In a SAN, much of these savings come from the ability to consolidate support and centralize management. And with today's virtualization technologies, SAN-based storage is even more cost-effective from a management perspective. Customer H's new environment, the equivalent of one FTE's time is spent managing over 20 TBs of storage in multiple locations.

11.4.5 Backup and recovery

Customer H has been an IBM Tivoli Storage Manager shop since the mid 1990s. Initially, when the number of servers was limited, the implementation was accomplished using multiple IBM DAT drives. As the IT infrastructure grew, and the need for better enterprise backups expanded, an IBM 3494 tape library with six SCSI-attached drives was purchased. This dramatically improved the backups for the environment, yet, by next year, it was beginning to become a bottleneck due to its I/O limitations, both at the drive and the connectivity level.

During the SAN design sessions, Customer H worked with consultants to come up with a solution to this, and two concepts were offered. The first was to put a black box on the back of each tape drive that would convert the SCSI drives to fiber connectivity. The second was to replace the system with a new fiber-based LTO system. Evaluating the performance specifications of both, storage capacity of both, and most importantly, the cost and ROI, it was determined that migrating to an IBM 3584 library with eight LTO2 drives would provide the best return on investment for the price. One of the key benefits of implementing a fiber-based tape unit is that LAN-free backups could be tested, as well as improving restore times based on the higher bandwidth of the tape library. The availability in the industry of Storage logical partitioning may improve the situation further.

Recovery of data in the environment was another issue that Customer H had to review. Prior to the SAN, some systems were backed up to a IBM Tivoli Storage Manager server, while others had system specific backup processes in place. Restoration of data was also problematic, especially with platform upgrades. When the company performed its last file server upgrade, it took on average three days to bring each server online with multiple staff people involved for varying degrees. It also required the system to be down so that no changes were made during the migration. With the SAN in place, server migrations are much simpler, and require fewer resources. When a server that has SAN-connected storage needs to be replaced, often it requires just a move of Fibre Channel cables and an OS-based disk discovery to bring the data back online. This can take a data migration from days, to just a few hours, or sometimes, just minutes. This is a true advantage of simplifying the storage environment.

11.4.6 Environments: Have you seen my electric bill lately?

Another disadvantage of DAS is the environmental aspect. In order to accommodate DAS, many servers were over 6-8U, meaning that only four servers could be placed in a rack. Over time, this created a space shortage in the primary data center, as well as power issues, since

each server required at least two separate power feeds. With the implementation of the SAN, Customer H has been able not only to reduce the scale of the servers, but also the number via the ability to allocate storage on a large scale. They have also been able to reduce their power and cooling needs via the server consolidation (LPARs and blades), which has been extremely beneficial in opening up floorspace and power in the primary data center. Even with additional growth of the storage environment, they do not see any environmental issues since storage arrays are becoming smaller and their need for power and cooling is decreasing.

11.4.7 Post implementation: Was the environment simplified?

With phase one complete, a quick look at the storage environment quickly showed all the areas that had been improved via simplification. First, Customer H only had to manage two switches, versus multiple switches that would have been present in a business-as-usual scenario, or even with a cord-edge design. Secondly, we had 15 servers sharing a single storage device, with server-based storage utilization above 70%, a increase of over 40% from the original scenario. Third, only the equivalent of one Full Time Equivalent (FTE) (two physical FTEs using around 50% of their time) is required to manage the equipment and software. All the attached systems were also easier to manage, since they all had the same storage driver, could have storage allocated dynamically, and they did not experience outages as other systems had. Server-based clustering was also much easier, and added another level of availability to the environment. Finally, with both a new server and a fiber-attached LTO2 tape library, the enterprise-based backup via IBM Tivoli Storage Manager were occurring much quicker, which decreased the backup window for the organization.

11.5 Phase two: More consolidation

With the SAN in place and the design structure stable, it was now time to tackle the complexity around the key platform in the environment, Microsoft Windows. The storage in the Microsoft Windows environment was the most under-utilized at the company. Often storage utilization was less than 30%. The servers were oversized to accommodate the increased storage needs of the application. Another need was to make the Windows-based servers highly available via clustering. With the SAN, clustering became easier and it also allowed allocation of a single LUN to multiple devices without the expense of purchasing storage arrays for specific applications. Using the SAN-based storage also consolidated the storage drivers for the servers, simplifying the administration. When IBM began offering pSeries servers with the ability to LPAR, Customer H capitalized on the opportunity to implement it.

Customer H knew the advantages that consolidating storage brought, and knew that doing the same on the server side would only have positive advantages to the IT shop. With logical partitioning, the team was able to increase the utilization of hardware components, such as CPU, memory, and even power and cooling. Well over 20 IBM pSeries servers have been consolidated onto seven LPARable boxes. The team also implemented three IBM BladeCenter servers (see Figure 11-7) (14 servers capable each) into half of a 42 U rack, and are looking to implement VMWare to further virtualize the server environment.



Figure 11-7 IBM BladeCenters

11.5.1 Could it have been too good to be true?

As SAN usage grew, and as applications were retired or replaced, storage allocation issues emerged. One storage array, an IBM Enterprise Storage Server (ESS) did not then allow removal of a LUN without reformatting the entire LSS, a subset of the storage array. Often, other systems' LUNs are placed on the same LSS, and, therefore, reformatting is not an option. With this issue, the storage team found itself over-allocating storage again to systems just to get the most use out of the purchased disk. If that was not possible, Customer H would migrate the data in use by creating another LUN and mirroring on the server prior to removing it (CPU and I/O overhead). This was time-intensive and not simple. They needed a way to simplify their storage allocations and deallocations.

Customer H brought this issue to the attention of their IBM storage representative who introduced them to the IBM TotalStorage SAN Volume Controller (SVC) concept. This is a block-level virtualization product that allows better management of storage through more granular control. It would give them the ability to move data to different areas on heterogeneous devices without the server knowing, expand LUN allocations, or shrink them. It also allowed for a single storage driver to be used on the servers even with the heterogeneous storage behind it. Another advantage of the product was that it centralized the ability to FlashCopy onto the SVC, eliminating the need to maintain it on individual storage arrays, as well as allowing Flashcopy across any storage device behind the SVC, versus being limited to the individual storage array. Another way SVC simplified things was the method of storage allocation to the client devices from the SVC was now unified (reducing training requirements) after the initial configuration. See Figure 11-8.




Figure 11-8 SVC inclusion

The next phase of simplifying the storage environment began with the purchase of the SVC and implementing tiered storage with the addition of an IBM DS4000 SATA array. With this purchase, Customer H was now able to become even more granular with their storage allocations, as well as improving the ability to reclaim unused storage. This also gave them the ability to spread data across different tiers of storage based on the application's need. Customer H said the SVC installation was easy.

The storage team began an in-depth search into storage administration tools to simplify management of the environment.

Since Customer H has a heterogeneous storage environment, they sought a utility that could centralize the management of all the devices into a single location with minimal impact on clients. The key product to catch their attention was IBM TotalStorage Productivity Center. This suite of integrated tools can manage the data, storage, and fabric from a central Web-based dashboard that allow them to access it from anywhere with Web access. They also liked the product due to the fact that client agents are combined into a single install, so that they do not have to manage multiple agents on individual clients.

11.5.2 Lessons learned: Educate and translate

The key to getting the most value out of a storage environment, as well as maintaining the ability to keep it simple, is to manage the storage allocations and process. Manage what application owners need, not what they want. Also, when talking about utilization, analyze what is actually used on the servers, not the overall array. A key example that was helpful for Customer H was in a project that called for 4 TBs of storage over three years based on the proposal by the vendor. When the equipment arrived and it was time to allocate the storage to the project (just from the development prospective), Customer H contacted the

vendor and asked how much they would need. Their initial response was 4 TBs. Out of concern and curiosity, the storage team questioned the vendor about why they needed it. Their response was that it was easier to allocate all the storage up front versus requiring downtime later to add additional storage. Knowing that this person had either limited or no SAN experience, the storage team asked to have a call set up with the vendor's technical team to determine how they utilized the storage with their product. Based on this call, and education about how storage pools are allocated in a SAN, Customer H determined that the storage environment could grow without downtime or any effect on the application. In the end, a more affordable 1 TB of storage was allocated to the application, a 75% reduction in their request, with future purchases made as storage prices decline.

Another group that can be challenging to a Storage Administrator are the Database Administrators (DBAs). DBAs like to maintain full control of their storage, and like to have lots of it on hand so that their tables can grow easily (often automatically). To improve storage resource utilization, Customer H's storage team needed to show the DBAs that they did not need lots of storage preallocated to their system, and that the Storage Administrator could easily and quickly allocate storage to them within minutes. The SAN is flexible, especially with the integration of the SVC. The Storage team sat down with the DBAs, listened to their needs, and then educated them on how the SAN infrastructure works. As a result, they were able to reclaim 30% of the storage previously allocated to database environments.

11.5.3 Next steps

As with any infrastructure project, there is always more to do or that can be done to improve performance or to make the system more resilient. For their next steps, Customer H is planning the following:

- ▶ Reanalyze snapshots for quick application recovery.
- ▶ Implement Storage Resource Management tools to simplify the SAN management.
- ▶ Research file level virtualization.
- ▶ Implement additional tier levels of storage.
- ▶ Implement LAN free backups on a more standard basis.
- ▶ Expedite disaster recovery.
- ▶ Research business continuance capabilities via SAN.
- ▶ Expand SAN to other sites.
- ▶ Business continuance.
- ▶ Server consolidation.
- ▶ Storage policies.

11.5.4 If I could do it again with more money, and technology was available

As any person who has designed a system would acknowledge, the chance to start over from scratch after gaining the knowledge and experience would be a dream. In Customer H's environment, both the IBM team and the customer team learned many lessons, and because of this, have spent many hours making changes to the storage environment. Many of these changes were not major, but had more to do with further simplifying the management of the storage environment, cost reduction, or the benefit of hindsight. If given the chance to start over, they would have implemented the following:

- ▶ Identify a mission statement for the group.
- ▶ Identify the goals to support the mission statement.
- ▶ Define what service levels they plan to provide.
- ▶ Create a multiple year plan up front to assist in infrastructure expenditures planning.
- ▶ Create policies for the storage environment and data availability.

- ▶ Implement tier levels of storage and connectivity along with block level virtualization to improve allocations and TCO.
- ▶ Use a Storage Resource Management tool to manage storage from fabric, device, and file level storage.

It is very helpful to identify the mission of any group in an organization. This mission defines what the group is about and where its focus should be. Recently, Customer H's storage team sat down to review their actions and plan for the future, they did exactly this. What they came up with was that the

"The Enterprise Storage Management Team is charged with providing highly available, cost-effective storage solutions that maintain data integrity and business continuity for the company and its customers. Not only do they provide their customers with a storage environment, but they are focused on data integrity and availability, as well as acknowledging that we need to focus on TCO and ROI."

With a mission statement in place, they then created goals that met the mission of the group, as well as the goal of the organization. The goals that they defined were:

- ▶ To make all data residing on SAN storage highly available (24X7/365 days per year).
- ▶ Recover data with minimal data loss.
- ▶ Provide content management within the organization's data retention policies.
- ▶ Minimize unnecessary data duplication.

After defining goals, it was important to define the Service Level Agreements (SLAs) that would be provided to the organization. They decided to make their SLAs around availability, based on how quickly they could make data available in the event of a data outage. Good references to get some ideas on how to define SLAs as well as provide business continuance can be found in the IBM Redbook, *IBM TotalStorage Business Continuity Solutions Overview*, SG24-6684. For their planning purposes, they created three levels for SLAs:

- ▶ **Critical (0-1 hour data loss)**
 - **Maximum 1 hour recovery**
 - Scheduled snapshots of the data or mirroring for minimal data and fast recovery.
 - Disk-based backups and recovery.
 - Remote application mirroring.
- ▶ **Medium (1-24 hour data loss)**
 - **1-24 hour recovery**
 - Disk-based backup and recovery.
 - Snapshots of data.
 - Electronic vaulting of system backups.
- ▶ **Low (24 hour data loss)**
 - **1-5 day recovery**
 - Legacy data recovery via off-site tape backups.
 - If corruption required recovery from tape, there would be a maximum of five days data downtime.
 - In the event of a site disaster, the goal would be to have minimal data loss via electronic vaulting of system backups and tape-based recovery.

This helped them create a longer-term plan for storage. Below is a sample three year plan based on the above SLAs.

Three year plan:

Year One:

- ▶ Goals:
 - Research and implement extended on-line archiving of system backups/disk-based recovery.
 - Test data snapshot technologies and the ability to recover applications using this technology.
 - Develop data snapshot policies.
 - Research and test data vaulting.
 - Work with other groups to define data retention policies and data archiving needs.
 - Sunguard data recovery test.
 - Implement Storage Resource Management (SRM) tools to reduce storage complexity.
- ▶ Requirements:
 - Additional storage.
 - FTE to perform snapshot testing.
 - Participation from other IT departments.
 - Adjustment to three year capital plan.

Year Two:

- ▶ Goals:
 - Implement data snapshot policies.
 - Implement data snapshot technology for critical applications.
 - Evaluate existing data retention and modify to meet the business needs of the organization.
 - Develop data recovery testing policy for both new and current applications.
 - Perform disaster recovery testing on all new applications entering the Customer H environment.
 - Provide concise storage utilization reports via the Storage Resource Management tools.
- ▶ Requirements:
 - Policy approval for data snapshot.
 - Additional storage.
 - Additional licensing for snapshot utility.
 - FTE availability.
 - Participation from other IT departments.
 - Addition of SLAs and policies to project development and plan.
 - Identify critical applications.

Year 3:

- ▶ Goals:

- Implement remaining critical SLA recovery solutions.
- Implement data recovery testing policy for all systems.

Requirements:

- Policy approval for data recovery testing.
- Secondary site for mirroring.
- Additional bandwidth to site specifically for data mirroring and storage at the secondary site.
- Participation from other IT departments.
- FTE availability.

Policies are the key method for making simplification stick in an environment. As one can see, enacting a plan for the group is not a solitary concept. The involvement of other Information Technology departments, as well as the users may be required. Policies are also key in implementing, managing, and simplifying a storage environment. Policies are the true guiding principles for how things are managed in the organization. By getting policies approved at the organizational level, you can more easily control the data in the environment. Planned future Customer H policies are (See Figure 11-9):

- ▶ Application backups using snapshot technologies.
- ▶ Data recovery tests.
- ▶ IBM Tivoli Storage Manager data retention policies.




Figure 11-9 If I had to do it over again

Abbreviations and acronyms

UFS

AMD	Advanced Micro Devices	IBM	International Business Machines Corporation
API	Application Program Interface		
ASIC	Application-specific Integrated Circuit	ILM	Information Lifecycle Management
		ISL	Inter-Switch Link
AT&T	American Telephone and Telegraph	IT	Information Technology
		ITSO	International Technical Support Organization
ATA	Advanced Technology Attachment		
BTU	British Thermal Unit	JBOD	Just a Bunch of Disks
CA	Computer Associates	JFS	Journaled File System
CEO	Chief Executive Officer	LAN	Local Area Network
CIFS	Common Internet File System	LDAP	Lightweight Directory Access Protocol
CIM	Common Information Model		
CIO	Chief Information Officer	LED	Light Emitting Diode
CLI	Command Line Interface	LPAR	Logical Partition
CPU	Central Processing Unit	LSS	Large Sub System
CSM	Caching Services Module	LTO	Linear Tape Open
DAS	Direct Attached Storage	LUN	Logical Unit Number
DB	Data Base	LVM	Logical Volume Manager
DBA	Data Base Administrator	MA	Maintenance Agreement
DFSMS	Data Facility System Managed Storage	MB	Megabyte
		NAS	Network Attached Storage
DMTF	Distributed Management Task Force	NFS	Network File System
		NTFS	NT File System
DR	Disaster Recovery	OEM	Other Equipment Manufacturer
EMIF	ESCON Multiple Image Facility	OS	Operating System
ERP	Error Recovery Program	PB	Petabyte
ESCON	Enterprise Systems Connection (architecture, IBM System/390®)	PC	Personal Computer
		POC	Proof of Concept
ESS	Enterprise Systems Storage	PPRC	Peer to Peer Remote Copy
FAT	File Allocation Table	RAID	Redundant Array of Independent Disks
FC	Fibre Channel		
FTE	Full Time Equivalent	RAM	Random Access Memory
FTP	File Transfer Protocol	ROI	Return On Investment
GB	Gigabyte	RPM	Revolution Per Minute
GUI	Graphical User Interface	SAN	Storage Area Network
HA	High Availability	SATA	Serial ATA
HBA	Host Bus Adapter	SCSI	Small computer System Interconnect
HDS	Hitachi Data Systems		
HP	Hewlett Packard	SFS	SAN File System
HSM	Hierarchical Storage Manager	SLA	Service Level Agreement
I/O	Input/Output	SMI-S	Storage Management Initiative Specification
I/T	Information/Technology		

SMS	System Managed Storage
SNIA	Storage Network Industry Association
SNMP	Small Network Message Protocol
SPC	Storage Performance Council
SRM	Storage Resource Manager
SSA	Serial Storage Architecture
STK	Storage Tek
SVC	SAN Volume Controller
TB	Terabyte
TCO	Total Cost of Ownership
TCP/IP	Transfer Control Protocol/Internet Protocol
TPC	TotalStorage Productivity Center
TPM	Tivoli Provisioning Manager
UFS	Uniform File System
UK	United Kingdom
URL	Unicode Record Locator
USD	United Record Locator
USP	Universal Storage Platform
VSAN	Virtual Storage Area Network
VSS	Virtual Shadow Copy Services
VTs	Virtual Tape System
WAFL	Write Anywhere File Layout
WAN	Wide Area Network
WBEM	Web-based Version of CIM
WWN	World Wide Port Name

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information about ordering these publications, see “How to get IBM Redbooks” on page 190. Note that some of the documents referenced here may be available in softcopy only.

- ▶ *An Introduction to Storage Provisioning with Tivoli Provisioning*, REDP-3900
- ▶ *IBM TotalStorage Productivity Center: Getting Started*, SG24-6490
- ▶ *Managing Disk Subsystems Using IBM TotalStorage*, SG24-7097
- ▶ *IBM TotalStorage Business Continuity Solutions Overview*, SG24-6684
- ▶ *Exploring Storage Management Efficiencies and Provisioning*, SG24-6373
- ▶ *Provisioning On Demand Introducing IBM Tivoli Intelligent ThinkDynamic Orchestrator*, SG24-8888
- ▶ *IBM Tivoli Storage Management Concepts*, SG24-4877
- ▶ *IBM Tivoli Storage Manager Version 5.3 Technical Guide*, SG24-6638
- ▶ *Virtualization in a SAN*, REDP-3633
- ▶ *Virtualization and the On Demand Business*, REDP-9115
- ▶ *Introduction to Storage Area Networks*, SG24-5470
- ▶ *The IBM TotalStorage Solutions Handbook*, SG24-5250
- ▶ *IBM TotalStorage SAN Volume Controller*, SG24-6423
- ▶ *IBM TotalStorage: Introducing the SAN File System*, SG24-7057
- ▶ *Implementing the IBM TotalStorage SAN Volume Controller Storage Software on the Cisco MDS 9000*, SG24-7059
- ▶ *IBM TotalStorage: Integration of the SAN Volume Controller, SAN Integration Server and the SAN File System*, SG24-6097
- ▶ *The IBM TotalStorage DS8000 Series: Concepts and Architecture*, SG24-6452
- ▶ *The IBM TotalStorage DS6000 Series: Concepts and Architecture*, SG24-6471
- ▶ *Volume Migration Using SAN Volume Controller*, TIPS0400

Other publications

These publications are also relevant as further information sources:

- ▶ *Take the next step in creating an effective storage infrastructure* by Marc Farley, G224-7269
- ▶ *IBM TotalStorage Infrastructure Simplification sales kit*
- ▶ *TotalStorage Infrastructure Simplification solution pack sales kit*

Online resources

These Web sites and URLs are also relevant as further information sources:

- ▶ IBM TotalStorage Virtualization Family
<http://www.storage.ibm.com/software/virtualization>
- ▶ Server consolidation
<http://www-1.ibm.com/servers/solutions/serverconsolidation>
- ▶ Total cost of ownership
<http://www-1.ibm.com/servers/solutions/serverconsolidation/tco/>

How to get IBM Redbooks

You can search for, view, or download IBM Redbooks, Redpapers, Hints and Tips, draft publications and Additional materials, as well as order hardcopy Redbooks or CD-ROMs, at this Web site:

ibm.com/redbooks

Help from IBM

IBM Support and downloads

ibm.com/support

IBM Global Services

ibm.com/services

Numerics

24x7 73
3380 46
3390 46
3584 60–61
3584 library 179
3990 46
8100 34
8300 35
9032 46–47

A

abstract 74–75
access 75–76
Access Servers 144
accessing 85
accomplished 172
Accounting perspective 114
acquisition 163
adapt 67
adapter 46
administration 23, 93, 126
administrator 24, 72, 102, 123–124
administrators 123
adopted 81
advanced function support 163
advanced functions 84, 89, 161
Advanced Provisioning 123
Advanced Technology Attachment 23
advantage 67
affordable 176
aggregated 141
aggregation 74
AIX 50, 82, 101, 142
alerts 118
allocate storage 180
analysis 114
API 81, 108, 120, 148
API's 148
application 8, 42, 49, 118
application availability 148
application performance 13
Application perspective 114
applications 178
approach 75
appropriate action 116
architectural design 147
architecture 49, 77, 178
archive 65, 133
archive copies 126
archive product 114
archived 114
Archiving 22
arena 133
array 178
assembly 134
assessment 65
assets 115
ATA 23

attachment 178
automate 138
automated 67
Automated source-target matching 122
automated storage management 45
automated storage provisioning 123
automating 142
automation 46, 80, 112
autonomic 93
avail 80
Availability 172
availability 7–8, 47, 73–74, 80, 115, 119, 123, 172, 178
Availability management 123
available 178

B

back-end 56, 74, 87
backplane 52
backup 8, 13–14, 22, 30, 57, 60, 62, 93, 112, 115, 126
backup software 43
backup window 8, 60
bandwidth 179
basic principles 16
benefit 82
black box 179
blade 52
blade servers 139
BladeCenter 139
Bladecenter 180
BladeCenter blade servers 141
block 82
block level 76
block virtualization 82, 84–85, 87–88, 90, 148
boot from SAN 43
bottleneck 49, 97
boxes 74
Brocade 123
budget 11, 42
building block 60
Bus and Tag 45
business continuity 8
business continuity solution 93
business directories 127
business policy 55
business processes 111
buying a car 19

C

c 86
cable 174
cable issues 175
cables 174–175
cache-based 96
Caching Services Module 95
capacity 13, 80, 93, 96, 115, 132, 134, 142, 162
capacity monitoring 12–13
capacity problems 118
capacity provisioning 123
capacity utilization 81, 89

- capital investment 171
- cartridge 60
- categories of data 155
- categorization 114
- centralized 65, 75, 81
- centralized storage 45
- challenges 72
- charge backs 118
- chemical 134
- CIM 81, 107–108
- CIM Object Manager 107
- CIMOM 107
- CIO 73
- Cisco 76, 172
- CLI 108
- clinical 64, 67
- CLIs 123
- cluster 84, 133
- cluster level 149
- Clustering 178
- clustering 77, 162, 180
- code level 42
- collaborate 68
- collapsed core design 48
- collected 114
- co-located 162
- command line interface 108
- common 76
- Common Information Model 107
- common platform 89
- common storage 54
- compete 72
- competition 73
- competitive advantages 147
- competitiveness 72
- complement 133
- complex 72, 134
- complex management 75
- complex storage 113
- complexity 56, 67, 72–73, 80, 96, 119, 132, 172
- compliance 116
- component 72
- computation 134
- computational 132
- computers 126
- concept 171
- configuration 14, 51–52, 54–55, 61
- confusion 73, 174
- consolidate 179
- consolidated 65
- consolidated disk storage 55
- consolidated disk system 54
- consolidated storage 42, 160
- consolidated storage environment 65
- Consolidating 31
- consolidating 180
- Consolidation 151
- consolidation 55, 58, 80, 148, 159, 163, 169–170
- consolidation project 170
- consolidations 73

- consultants 172
- consulting group 171–172
- contemporary 80
- Continuous copy services 148
- control 93, 160
- controller failures 178
- cooling capacity 162
- Copy Services 14
- copy services 14, 93, 148
- cord-edge 180
- core edge 50
- core-edge 176
- corporate server environment 126
- cost 22, 42, 48, 72, 96
- cost analysis 176
- cost effective 62
- cost of storage 178
- cost-effective 7, 72
- costs 7, 64, 80, 93, 119, 160
- CPU 178, 180
- CPU architecture 116
- Cross-device consistency groups 122
- cross-vendor 10
- CSM 95
- csm 76
- customizing 114
- cylinder 46

D

- DAS 8, 94, 177, 179
- data 80, 84
- data center 19
- Data center storage administrator 123
- data collection 127
- data duplication 14
- data growth 7
- data location 13
- data migration 89, 179
- data path 84, 86
- data protection 93
- data recovery 64
- data relocation 14
- data replication 14
- data retention 8
- Data sharing 160
- data sharing 97, 160
- data storage 170
- Data transfers 172
- database 7, 49, 118, 124
- database administrator 124
- database filesystem 13
- database repository 114
- databases 68, 160
- DataFabric™ Manager 133
- DBA 124
- decentralized storage 45
- decimated 174
- Dell 42
- demand 81
- demands 65

- dense storage 151
- deploy 73, 77
- depreciated 163
- depreciation 18
- design 10, 172–173
- device 51
- devices 73
- DFSMS 76
- digital 68
- Direct Access Storage 176
- direct access storage 8
- direct attached storage 45, 94
- direct-attach storage 149
- direct-attached disks 160–161
- directly attached storage 179
- director 47, 173
- director class device 172
- disadvantage 179
- disaster 14, 67
- disaster recovery 56, 102, 111
- discovered devices 119
- disk 78
- disk array 161
- disk arrays 178
- disk capacity 115
- disk drive virtualization 82
- disk space 117
- disk storage 53, 55
- disk zone 92
- disks 76, 119
- disparate 67, 160
- disparate storage systems 82
- disrupted 75
- disruption 150
- disruptive 132
- distributed environments 126
- Distributed Management Task Force 107
- DMTF 107
- DMX800 151
- documentation 68
- domain 64
- domain manager 84
- down 89
- downtime 89, 178
- Drive type 176
- Driver level 176
- DS family 54, 82
- DS4000 94, 97, 107, 109, 118
- ds6 32
- DS6000 31, 94, 97, 107, 109, 114, 118, 151–152, 163
- DS6800 162
- DS8000 54, 77–78, 82, 94, 97, 107, 109, 118, 149, 151–152, 164–166
- DS8000 server 152
- DS8100 162
- DS8300 162
- DSxxxx 81
- dynamic 72

E

- economic 72
- edge core edge 50
- electronic 64, 68
- electronically 65
- element 120
- elements 163
- EMC 46, 81, 151–152
- EMC's 151
- EMIF 46
- encrypt 134
- encrypted 132
- enhancements 68
- enterprise 65, 73, 163
- enterprise class storage 22
- enterprise resource planning 126
- Enterprise Storage Server 81
- enterprise wide summary information 114
- enterprise-class scalability 90
- enterprises 148, 160
- environment 72, 114, 148, 169, 179
- environment. 114
- environmental issues 180
- environments 126
- ERP 126
- ESCON 45–47, 52
- ESCON Director 46
- eServer 92, 103
- ESS 31, 43, 49–50, 81, 94, 97, 107, 109, 113, 118, 164–166, 170–171
- ESS F20 162
- establishing 65
- estimates 134
- Ethernet 60
- event 67
- expanded core 50
- expenses 177
- experience 11
- EXT2 12
- EXT3 12

F

- f 84, 103
- fabric 14, 42, 46–47, 50–53, 82–83, 120, 132, 174
- facilitating 123
- FAT 12
- fault 51
- Fibre 46
- fibre 41, 179–180
- fibre cables 179
- fibre channel 171
- Fibre Host Bus Adapter 22
- fibre storage 16
- fibre storage network 44
- file 73, 83
- file access response time 13
- File perspective 114
- file storage 142
- file system virtualization 82

- file systems 116
- file virtualization 85, 98
- financial system 170
- firmware 51, 82
- flash copy 43
- FlashCopy 89
- Flashcopy 14
- flex 75
- flexibility 9, 98, 160
- flexible 67
- floorspace 180
- footprint 165
- free 82
- free space 116–117
- Freeze-and-Go functions 121
- FTE 176, 180
- FTE's 179
- functions 73
- funding 11

G

- Gartner Group 130
- gear 85
- genome 132
- gigabyte 134
- goal 171
- goals 11, 64
- government 134
- grant 85
- grow 53
- growth 72, 111, 130, 170
- GUI 108, 114, 120

H

- hardware 4, 52, 160, 162
- hardware profile 116
- HBA 14, 22, 42–43, 60, 93, 142–143
- HDS 46, 81
- head 46, 135
- health 64
- health care 169
- healthcare 65
- Heterogeneous 138
- heterogeneous 23, 138
- Heterogeneous access 12–13
- heterogeneous computers 97
- heterogeneous environment 46
- heterogeneous hosts 86
- heterogeneous servers 49
- high 64, 67
- high availability 46
- high availability fabric 52
- high availability switch 50
- High performance 172
- high-performance, 139
- Hitachi 81
- homogeneous 160
- hop 51
- hospital 64, 67

- host 48, 83
- Host Bus Adapter 22, 42
- hotspot 51
- hotspots 111
- HP 42, 59, 81
- HSM 134
- HTTP 107

I

- i 86
- I/O performance 170–171
- I/T applications 178
- I/T infrastructure 178
- I/T shop 180
- IBM 42, 46, 59, 180
- IBM 9032 46
- IBM SAN File System 152
- IBM Tivoli Provisioning Manager 123
- IBM Tivoli Storage Manager 126, 180
- IBM TotalStorage SAN File System 77, 138
- IBM TotalStorage SAN Volume Controller 139
- IBM TotalStorage Software 80
- IBM TotalStorage vision 83
- IDC 178–179
- imaging 65
- in-band 83–84, 90
- in-band virtualization 82–83
- increases 67
- incremental backup 60
- industry 73
- industry standards 120
- info 72
- information 64
- information technology 170
- infrastructure 3, 8, 10, 45–46, 48, 57, 59–60, 72, 172, 178
- infrastructure cost 42
- infrastructure simplification 112
- input/output 82
- InRange 52
- Inter Switch Link 51
- interface 46, 56, 73
- interoperability 7–8, 22, 44, 51, 54–56, 73, 175–176
- inter-operability matrix 42
- interoperability matrix 42
- interruption 133
- Inter-Switch Links 174
- inter-switch traffic 51
- investment 18
- ip 77
- IP address 115
- is 78
- iSCSI LUN 134
- ISL 50–51, 175
- isolate 134
- IT 123, 138
- IT environment 123
- IT infrastructure 170

J

JBOD 102
jeopardize 132
JFS 12
JFS2 12
just-in-case 123

K

key benefits 179

L

lab 132
LAN 14, 46, 61–62, 126
LAN-free 179
large disk 166
last boot time 116
lay 75
leaf 50
Lease expiration 149
LED 175
level 82
lever 83
library 65
library robotics 60
license 23
Licensed internal code levels 176
lifecycle 73
limit 80
Linear Tape Open 59
Linux 94, 99, 101, 132
linux 133–134
Locally Attached Storage 12
location replication 14
logical mapping 91
logical unit number 84
logical view of data storage. 75
Logical Volume Manager 50, 83
LPAR 46, 62
LPAR's 180
LTO 59–60
LUN 7, 14, 46, 49–50, 54, 56, 84, 87, 103, 108, 124, 180
LUNs 143
LUNs 142
LVM 50, 82–83

M

m 114
main 75
mainframe 7, 44, 99
maintenance 10, 18, 23, 162–163
maintenance costs 18
manage 72
manageability 73
managed disk 88
managed disk group 88
management 67, 119, 170, 172, 174, 176
management tool 54–55
manipulate 134

manual 73
manufacturer 115
map 75
mapping 85
matrix 43
matrix of supported configurations 43
max 65
McData 52
MDG 88
MDisk 88
MDS 9000 76
mechanisms 114
medical 64, 68, 131
medicine 64
megabyte 73
meshed fabric 50, 52
meta 85
meta-data 85, 97, 99–103
meta-data controller 99
meta-data server 97
metrics 8, 11
microcode 42, 50–51
microcode level 42
Microsoft 180
Microsoft Windows 94
mid-range disk 22
migrate storage 150
migration 14, 55, 173
million dollar mark 176
mirrored pool 102
mirroring 47, 62, 93
model 43
modular 90, 163
money 77
monitor 56, 118, 121
monitoring tool 54
multi 23, 73
multi pathing 43
multi pathing software 42
Multilayered 65
multi-pathing 52
multipathing 142
multi-pathing driver 74
multi-pathing software 22, 84
multiple copies 62
multiple network storage 74
multiple switches 180
multiple vendors 42, 55
Multiply 176
multi-supplier 148
multi-tier 73
multiyear 64
Murphy's Law 9

N

name 83
name space 153
NAS 22, 77, 85, 99, 103
nested 74
net 76, 78

- network 8, 83
- network address 115
- Network Attached Storage 12, 22, 85, 160
- networked systems 148
- new zones 143
- nfs 133–134
- nodes 95
- non business use of space 13
- non-disruptive 162
- NTFS 12

O

- objectives 11
- OEM 163
- OEM storage arrays 164–165
- OEM subsystems 164
- offline 67
- off-load 84
- off-loading 86
- ON 64
- on demand 8, 139–140
- On Demand Business 72
- online 134
- ontap 104
- open 104
- open architecture 81
- Open Software 82
- open standards 119
- open storage 52
- open system 46, 48–49
- open systems 45
- open-system 7
- operating plan 10
- operating platforms 126
- Operating System 43
- operating system 24, 42, 46, 48–49, 72–73, 94, 101, 115
- operating system environments 160
- operating systems 126
- operations 64
- optimization 121
- optimize 73
- Oracle 49
- orchestrate 142
- orchestrated application 139, 141
- orchestration 138
- Orchestration environment 141
- organization 68, 171, 176
- organization's 171
- organizational 176
- OS 116, 160, 172
- OS/VS1 77
- Other Equipment Manufacture 163
- outage 51
- outages 180
- Outlined 65
- out-of-band 83, 85–86
- out-of-band virtualization 82, 85
- overhead 62
- oversized 177

P

- p5 595 62
- partition 76
- partitioning 179–180
- partner 10
- partners 64
- patch panel 175
- patient 64, 68
- PC 81
- peak demands 123
- Peer-to-Peer 77
- peer-to-peer data copy 89
- people 48
- Performance 133, 178
- performance 7, 13–14, 48–49, 51–52, 54–56, 58, 60, 73, 84, 93, 96–97, 102, 132, 134, 162, 174
- periodic migration 13
- permissions 154
- permutations 73
- Persistent ID 174
- personal computers 81
- personnel productivity 172
- phases 64
- physical asset 74
- physical blocks 150
- physical disk 102
- physical disks 143
- physical footprint 163
- physical infrastructure 74, 149
- physical partition 50
- physical storage 74, 84, 88
- physical storage consolidation 162
- physical volume 50, 75
- physicians 68
- picture 73
- planning 93
- platform 67
- playback 134
- point in time copy 43, 47, 54, 56
- Point-in-time 122, 148
- Point-in-Time Copy 122
- points of failure 178
- Policies 174
- policies 76, 174
- policy 174
- policy violations 118
- policy-based 80, 102
- policy-based automation 80
- polished 175
- pool 67
- pooled 155
- pooling 150
- pools of capacity 155
- port 47, 51, 53, 174
- Port Migration 174
- possible duplicates 128
- potential problems 111
- power and cooling 180
- POWER-based 164
- precaution 134

- predetermined 172
- prerequisite tasks 143
- pressure 7
- prevent 171
- price 73
- primary data center 179
- privacy 8
- problematic 179
- process 134
- processor count 115
- processor speed 115
- processor type 115
- production mode 172
- productive 119
- productivity 7, 64, 74, 89, 93, 119
- products 10
- project team 172
- proliferation of UNIX 42
- Proof Of Concept 172
- protect 67
- protection 134
- provider 64
- provisioned 141
- provisioning 138, 141
- provisioning of servers 123
- provisioning software 143
- provisioning tasks 139
- provisioning workflows 143
- pSeries 62, 180
- purchase 73

Q

- QFS 12
- quality 64, 68
- quantity 114

R

- rack 174, 180
- radiology 68
- raditional file systems 142
- RAID 73, 82, 88, 161, 176, 178
- RAID10 62
- RAID5 62
- RAM 19
- ram 19
- RAM size 115
- rapid 65
- read to write ratio 14
- real volume 99
- realistic 77
- reallocate 150
- reallocated 176
- rebooting 156
- reconfigure 56
- recorded 170
- recovery 14, 58, 126
- Redbooks Web site 190
 - Contact us xv
- reduce 75

- redundancy 173, 175
- redundant 26, 90
- regulated 132
- regulations 8
- Reiser 12
- relationship 119
- reliability 65, 89
- reliable 134
- remapping 166
- remote client 61
- remote mirroring 43, 54, 56, 93
- remote volumes 122
- replication 14, 121
- replication feature 121
- replication group 121
- replication services 148
- replication session 121
- replication task 121
- report 118
- reports 115
- requirements 134, 171
- research 65, 132–133, 172, 176
- reservoir of storage 74
- resiliency 172
- resilient 176
- resources 8, 72, 111, 134
- Restoration 179
- restore 30, 57, 62, 93
- Resume Continuous Copy 122
- retention 112
- retooling 65
- return on investment 178
- revenue 30
- risk 7
- robotic tape library 58
- robust 132
- ROI 127

S

- S/390 46, 48
- safety 64
- SAN 13–14, 22, 41–44, 47–48, 51, 53–54, 59–60, 72–73, 75–77, 80–82, 84–90, 93, 96–97, 99, 101–103, 108–109, 113, 120–121, 126, 139, 143, 148–149, 160–163, 171–174, 176, 178–180
- san 26, 73, 76, 78, 82, 85
- SAN administration skills 142
- SAN arrays 178
- SAN attached disk 178
- San Attached Storage 12
- SAN based arrays 178
- SAN configuration 42
- SAN connected storage 179
- SAN design 179
- SAN elements 120
- SAN environment 173
- SAN File System 13, 75–76, 86, 97–103, 107, 109, 137–143, 156
- SAN File System client 143
- SAN File System clients 143

- SAN File System configuration 144
- SAN File System environment 142, 144
- SAN File System filesets 143
- SAN File System host 156
- SAN File System management tasks 143
- SAN File System storage pools 143
- SAN File System User Pool 143
- SAN File System user pool 144
- SAN File System's 139, 141
- SAN infrastructure 139
- SAN management 108
- SAN ports 42
- SAN sprawl 22, 42
- SAN switch 42
- SAN switches 120, 171
- SAN Volume Controller 75–76, 87–88, 90–94, 96–97, 109, 113, 148–150
- SAN Volume Controller Storage Software for Cisco MDS 9000 75–76, 90, 95–96, 107, 109
- SAN-attached storage 142
- SAN-based 180
- SAN-based storage 96
- san-connect 65
- SANs 160, 178
- SAN-wide 76
- SAP 49
- SATA 139
- Scalability 163, 172
- scalability 84, 86, 96, 134
- scalable 65, 90, 133
- scalable tape library 60
- scale 53, 179
- scale of the servers 180
- scheduling 68
- science 132, 134–135
- scientists 134
- script elimination 122
- SCSI 170
- SCSI drives 179
- Seagate 59
- seamlessly 134
- secure 67
- security 93, 112, 116, 134
- segregate 134
- separate power feeds 180
- sequence 134
- serial number 115
- server 24, 42–43, 48–49, 51, 53–54, 56–57, 62, 134, 161
- server clustering 43
- server equipment 177
- server online 179
- server sprawl 42
- server-based virtualization 82
- server-centric 142
- servers 73
- services 64, 74
- SFP 175
- SFS 75, 81–82, 86, 99, 109, 143, 153–154
- SFS host 156
- SFS management workflows 143
- SFS Pools 141
- shared drive 58
- shared resources 18
- significant 134
- simple 62
- simplification 60, 62, 180
- Simplification principles 174
- simplify 175
- simplifying cabling 174
- simplifying the storage environment 178–179
- simultaneous 68
- single points of failure 172
- skill 48
- skilled administrators 139
- skills 78
- skinny tree 50
- small environment 176
- SMI 108
- SMI-S 77, 81, 108, 118–119
- SMS 99
- snap 106
- snapshot 134
- SNIA 74, 77, 81, 85, 88, 98–99, 107–108, 118
- SNIA-CTP 77
- SNMP 120
- soft 82
- softek 104
- software 4, 16, 172
- software pricing 116
- software product 119
- software upgrades 172
- Solaris 101
- solution 67, 133, 135
- solutions 73
- source and target pairs 124
- source data 148
- space 116, 177
- space management 13
- spaghetti 174
- specialist 22, 64
- Spoofing 174
- sprawl 22
- stability 9
- staff 135
- stale file 129
- standard LUN 49
- standardization 49
- standardized reports 115
- standards 8, 172
- Status Query 122
- storage 48–49, 53, 64, 71–72, 155, 176, 179
- storage administrator 24, 89, 124
- storage administrators 171
- Storage Area Network 42
- storage area network 171
- storage area network infrastructure 170
- storage area networks 172
- storage arrays 163
- storage budget 111
- storage capacity 123, 142, 179

- storage consolidation 160, 162–164
- storage density 163
- storage device 43, 51
- storage devices 120, 149
- storage driver 180
- storage environment 115, 119, 180
- Storage Infrastructure 78
- storage infrastructure 81, 114, 127, 163
- storage management 49, 84, 119, 123, 177
- Storage Management Initiative 108
- Storage Management Interface Specification 108
- storage management software 84
- storage network 42, 46, 103
- Storage Networking Industry Association 74, 81, 107, 118
- storage on servers 176
- Storage performance 176
- storage perspective 171
- storage pool 53–54, 102, 150
- storage pooling 140
- storage pools 55, 77
- Storage provisioning 138
- storage provisioning 139, 142
- storage sprawl 50
- storage subsystem 42
- storage subsystem partitioning 82
- storage system 48–49
- storage system partitioning 82
- storage utilization 74–75, 117, 180
- Storage virtualization 75
- storage virtualization 82–83
- storage virtualization model 83
- storage volumes 114
- strategically 174
- strategy 174
- striped pool 102
- subsystem 47
- subsystem to subsystem replication 14
- suitability 42
- summary panel 114–115
- Sun 42–43
- Sun Solaris 139
- support 64, 173
- support matrix 43
- SVC 14, 67, 75–76, 81–82, 87, 94, 97, 109, 113, 117–118, 144, 149–150
- swap space 115
- sweet spots 54
- switch 22, 42, 46–48, 50–53, 120
- switch code 174
- switch vendors 172
- switches 47, 174
- symmetric virtualization 83
- synchronize databases 122
- synchronous virtualization 83
- system 77
- system administrator 124, 176
- system administrators 113
- System availability 176
- system availability 93

- system managers 162
- system-managed storage 99

T

- tablespace 124
- tape 57, 59, 61, 65, 83, 134
- tape backup 57
- tape drive 57
- tape infrastructure 59
- tape simplification 60
- tape virtualization 83
- tapes 83
- tasks for provisioning 123
- TCO 16–25, 29–30, 73, 82, 94, 99, 103, 160, 172, 178
- technologies 132
- technology 10, 55, 64, 134, 172, 176
- Telco 151
- terms 73
- throughput 61
- time zone 115
- Tivoli Provisioning Manager 81, 124–125
- Tivoli Storage Manager 13, 60, 81, 126
- tools 78
- topology 14, 120
- topology rendering 14
- Total 17
- Total Cost of Ownership 17
- total storage management 177
- TotalStorage 65, 111, 152
- totalstorage 63
- TotalStorage DS Family 71, 81
- TotalStorage DS Family Servers 97
- TotalStorage DS products 151
- TotalStorage DS4000 family 109
- TotalStorage DS4000 series 107
- TotalStorage DS6000 family 109
- TotalStorage DS6000 series 107
- TotalStorage DS8000 family 109
- TotalStorage DS8000 series 107
- TotalStorage Enterprise Server 118
- TotalStorage Enterprise Storage Server 97
- TotalStorage ESS 107, 109
- TotalStorage Infrastructure 169
- TotalStorage portfolio 148
- TotalStorage Productivity Center 81, 109, 113, 127
- TotalStorage Productivity Center for Data 81, 107, 113–114, 118
- TotalStorage Productivity Center for Disk 81, 113, 118–119
- TotalStorage Productivity Center for Fabric 81, 113, 120
- TotalStorage Productivity Center for Fabric console 120
- TotalStorage Productivity Center for Replication 81, 113, 121
- TotalStorage Productivity Center with Advanced Provisioning 123
- TotalStorage SAN File System 71, 81, 97, 109
- TotalStorage SAN Volume Controller 71, 81, 87, 109, 118
- TotalStorage SAN Volume Controller Storage Software for Cisco MDS 9000 109

- TPC 81, 109, 118, 123
- TPC for Data 13, 107, 113, 127
- TPC for Disk 13–14, 113, 119
- TPC for Fabric 14, 113, 120
- TPC for Replication 113, 121
- TPM 123, 137–138, 144
- TPM DCM 144
- TPM Logical Volumes 144
- track 46
- traditional applications 139
- training 78
- transactions 172
- transparent 61
- trending utilization 128
- troubleshooting 175–176
- TSM 126, 130
- type 43

U

- UFS 12
- unallocated disk space 116
- under utilized 177
- unified 77, 132
- UNIX 22, 42, 94, 99, 101, 154, 170
- unmanaged 170
- unscheduled 178
- unused capacity 166
- unused space 177
- upgrade 53–54, 170
- upgrades 172, 176
- usage 115
- usage violation 115
- user response time 13
- utilization 74, 93, 103, 111, 119, 130, 159, 162, 176, 180
- utilized 133

V

- variable 73
- various 76
- varying degrees 179
- VDisk 84, 87–88
- vdisk 14
- vendor 42–44
- vendor/multi-tier 73
- Veritas 104
- virtual 65, 71, 73, 81
- virtual disk 88, 91, 149
- virtual library 60–61
- Virtual SAN 96
- Virtual Storage Gateway 14
- virtual tape 77
- Virtual Tape Server 77
- Virtual Tape Storage 77
- virtual volume 99
- Virtual volumes 150
- Virtualization 148
- virtualization 11, 46, 56, 74–76, 148
- virtualization appliance 75
- virtualization appliances 82

- Virtualization Engine 77, 152
- virtualized 64, 74, 149
- virtualized environment 166
- virtualized resources 45
- virtualized SAN 108
- virtualized volume migration 14
- virtualization 16
- vision 75, 83
- VMWare 180
- volume 46–47, 132
- volume candidate pools 122
- volume management 87, 142
- volume to volume copying 14
- VSAN 96, 173
- VSANS 176
- VSANs 173
- VSS 105
- VTs 82
- VXFS 12

W

- WAN 61, 126
- warranty 23
- wasted space 116–117
- WBEM 107
- web browser 119
- web servers 141
- white light 175
- wide area network 61
- Windows 94, 99, 101
- Windows-based 180
- Wintel 22
- Wintel server 42
- workflow 124
- workflow automation 123
- workflows 143
- WWN 14, 173

X

- XML 107
- xmlCIM 108
- xSeries 92, 101, 103

Z

- z/OS 76, 122
- zone 82, 120, 161, 173
- zones 143
- Zoning 142
- zoning 14, 51–53, 120, 124
- zoning information 120
- zSeries 42, 44–47, 52

Introduction to Storage Infrastructure Simplification

(0.2"spine)
0.17"<->0.473"
90<->249 pages



Introduction to Storage Infrastructure Simplification



**Simplifying your
Storage environment**

**A practical approach
to Storage
Infrastructure
Simplification**

**Understanding the
TCO of Storage
Infrastructure**

This IBM Redbook introduces Infrastructure Simplification. Infrastructure Simplification is the methodology of analyzing the complete enterprise: business processes, workflow environment end-to-end, and IT for simplification. This analysis yields opportunities to save you time and money and eliminate unnecessary complexity that impedes the flow of information. This IBM Redbook discusses Storage Infrastructure Simplification and demonstrates multiple ways that IBM TotalStorage and Infrastructure Simplification can help you reduce complexity, save time and money, and release the flow of information in your business.

**INTERNATIONAL
TECHNICAL
SUPPORT
ORGANIZATION**

**BUILDING TECHNICAL
INFORMATION BASED ON
PRACTICAL EXPERIENCE**

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

**For more information:
ibm.com/redbooks**

SG24-7114-00

ISBN 0738494429