

Storage Area Network Essentials

**A Complete Guide to Understanding
and Implementing SANs**

Richard Barker
Paul Massiglia

Wiley Computer Publishing



John Wiley & Sons, Inc.

NEW YORK • CHICHESTER • WEINHEIM • BRISBANE • SINGAPORE • TORONTO

Storage Area Network Essentials

Storage Networking Essentials is both useful and readable, a rarity for a treatise on a subject as broad and dynamic as storage networking. Paul and Richard add just the right amount of historical perspective to show why the Storage and Networking communities were destined to meet. Going beyond succinct discussions of technology and implementation, the focus on compelling business motivations make this book an essential introduction approachable by a much wider audience than technical professionals.

**Mike Dutch, Director–Strategic Business Development
TROIKA Networks, Inc.**

Storage Area Network Essentials

**A Complete Guide to Understanding
and Implementing SANs**

Richard Barker
Paul Massiglia

Wiley Computer Publishing



John Wiley & Sons, Inc.

NEW YORK • CHICHESTER • WEINHEIM • BRISBANE • SINGAPORE • TORONTO

Publisher: Robert Ipsen
Editor: Carol A. Long
Assistant Editor: Adaobi Obi
Managing Editor: Micheline Frederick
Text Design & Composition: North Market Street Graphics

Designations used by companies to distinguish their products are often claimed as trademarks. In all instances where John Wiley & Sons, Inc., is aware of a claim, the product names appear in initial capital or ALL CAPITAL LETTERS. Readers, however, should contact the appropriate companies for more complete information regarding trademarks and registration.

This book is printed on acid-free paper. ∞

Copyright © 2002 by Richard Barker, Paul Massiglia. All rights reserved.

Published by John Wiley & Sons, Inc., New York

Published simultaneously in Canada.

No part of this publication may be reproduced, stored in a retrieval system or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning or otherwise, except as permitted under Sections 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4744. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 605 Third Avenue, New York, NY 10158-0012, (212) 850-6011, fax (212) 850-6008, E-Mail: PERMREQ@WILEY.COM.

This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold with the understanding that the publisher is not engaged in professional services. If professional advice or other expert assistance is required, the services of a competent professional person should be sought.

Library of Congress Cataloging-in-Publication Data:

Barker, Richard, 1946–

Storage area network essentials : a complete guide to understanding and implementing SANS / Richard Barker, Paul Massiglia.

p. cm.

ISBN 0-471-03445-2 (cloth : alk. paper)

1. Computer networks.
2. Information storage and retrieval systems.
3. Computer storage devices. I. Massiglia, Paul.

TK5105.5 .B347 2001

004.6—dc21

2001044410

Printed in the United States of America.

10 9 8 7 6 5 4 3 2 1

The authors wish to dedicate this work to their respective families, always the unsung heroes of all-consuming projects like this. Not only did they support us throughout, but Barbara Barker spent countless hours in the early stages of the project editing Richard's prose and trying to make the manuscript *look* like a manuscript. For her part, Judy Massiglia spent hours driving across boundless tracts of the American southwest one summer so her husband could write while they rode. Now *that's* dedication.

In a very real sense, it wouldn't have happened without you both.

Thanks.

About the Authors

Richard Barker is a Senior Vice President of VERITAS Software Corporation, where he takes a particular interest in storage area networks and network-attached storage devices, especially as they relate to the deployment of highly scalable and usable storage solutions in the global enterprise market. He is a well-known author on CASE topics having written three best selling books in the Oracle CASE (computer-aided systems engineering) Method series. Before joining VERITAS, he served as Senior Vice President of the Consulting and Technical Services Division of OpenVision International, Ltd. and later as Senior Vice President, Product Division. From 1984 through 1994, he was a Senior Vice President of Oracle, responsible for worldwide development and marketing of Oracle's Development Methodology and CASE product set. Prior to that, he managed the development center for the IDMSX mainframe database product for International Computers Limited (ICL) and developed a distributed database system for the U.K. health service. Mr. Barker is also a former board member of the Storage Networking Industry Association (SNIA), the premiere international organization for development and promotion of storage networking technology.

Paul Massiglia is a Technical Director at VERITAS Software Corporation. He represents VERITAS at several standards and industry marketing associations, and writes and lectures extensively. His 21 years in the data storage industry segment have included marketing, engineering, and engineering management positions with the former Digital Equipment Corporation, Adaptec Inc., and Quantum Corporation.

Mr. Massiglia is the former Vice Chairman of the RAID Advisory Board and is currently a member of the Board of Directors of the SNIA.

Acknowledgments	xix
Foreword	xxi
Introduction	xxiii
Part One: Understanding Storage Networking	1
<hr/>	
Chapter 1 What Storage Networking Is and What It Can Mean to You	3
What Is a Storage Area Network?	3
What Makes a SAN Different?	4
What Makes a Good SAN?	6
What Makes a Great SAN?	7
Why Connect Storage to a Network?	7
The Secret to SANs' Success: Software	9
SAN Software Capability	9
Summarizing SANs and Software	12
The Best Is Yet to Come: Radical Changes in Information Storage and Processing	13
Best-in-Class Computer Systems	13
Smarter Storage and Appliances	14
Heterogeneous Computer Systems	14
Data Storage as a Service	15
Back to Earth. . . .	16
A Couple of Closing Clarifications	16
Is a SAN Storage or a Network?	16
What Is This Area Thing, Anyway?	16
Is NAS Just SAN Spelled Backward?	17
Summary	19
Chapter 2 Benefits: What to Expect from SANs	21
The SAN Paradigm Shift	21
Properties of SANs	22
So What's the Paradigm Shift?	22
Conventional Information Processing	22

	The Primary SAN Paradigm Shift	24
	New Ways of Doing Things	25
	The Secondary SAN Paradigm Shift	25
	More New Ways of Doing Things	27
	A Model for Enterprise Information Processing	27
	Ten Ways the SAN Paradigm Shift Changes Information Processing for the Better	29
	Change 1: More Storage for Less Money	29
	Change 2: More I/O for Less Money	32
	Change 3: Better Data Protection for Less Money	35
	Change 4: Global Access to Global Data	36
	Change 5: Less Duplication, More Consistency	37
	Change 6: Taking It to the Streets: Really Global Access to Really Global Data	38
	Change 7: Clusters Come Out of the Closet	40
	Change 8: Disaster Recovery—Keeping Information Processing Going and Going . . .	42
	Change 9: Cutting Down—More Is Not Always Better	43
	Change 10: Global Computing for Global Enterprises	44
	Summary	47
Chapter 3	Leading Up to SANs: One View of Data Center Evolution	49
	A Bit of History	49
	Mainframe Generation I: The Rise of the Data Center	50
	Mainframe Generation II: Timesharing	52
	Online Transaction Processing: Instant Gratification for Users	54
	Minicomputers and Beyond: Diseconomies of Scale	55
	Local Area Networks	57
	The Standardization of Software	58
	A Pause to Take a Breath	59
	The PC and Client/Server Computing	59
	Specialized Servers	61
	Limitations of Specialized Servers	62
	Global Enterprises and the Need for Global Data	63
	External Storage Subsystems	65
	Summary	66
Chapter 4	Killer Apps for SANs	67
	Killer App 1: Backup—The Application Everyone Loves to Hate	67
	Shrinking Backup Windows	68
	Why a Backup Window Anyway?	68

Consistent Backups	69
Backup Schizophrenia	70
Different Uses, Different Backups 1	70
Different Uses, Different Backups 2	72
Backup on a Single Computer	73
A Little Relief: Multiple I/O Buses	75
Backup in Client/Server Environments	76
Improvement 1: LAN-free Backup	78
Improvement 2: Tape Drive Sharing	80
Improvement 3: The Designated Backup Server	82
Improvement 4: Serverless Backup	83
Improvement 5: Off-Host NAS Backup	85
Killer App 2: Highly Available Data	88
Mirroring 101	89
Drilling Down into Mirroring: Who's Doing What to Whom	90
RAID Controllers: A Lesson in Virtual Reality	92
Volume Managers: The Other Disk Virtualizer	93
The SAN Difference	94
SANs and Available Data	96
Killer App 3: Disaster Recoverability	98
The Third Mirror Theme	99
Split Mirrors for Backup	99
Other Ways to Use Split Mirrors	100
Killer App 4: Clusters—Continuous Computing	102
Requirements for Highly Available Computing	102
Positioning Clusters: Degrees of Availability	105
Tiers of Clusters	106
Detecting and Reacting to Failures	109
Failover Trade-offs in Clusters	109
Clusters and SANs	110
Heterogeneous Clusters	113
Killer App 5: Data Replication	113
Replication for Disaster Recoverability	114
Replication for Data Publication and Consolidation	115
Replication for Data Movement	117
The Mechanics of Data Replication	118
The Bottom Line: Continuous Global Access to Timely Information	119
Summary	120
Part Two: The Parts: What's in a SAN	123
Chapter 5 Storage Networking Architecture	127
The Path from Data to Application	127
From Bits to Records	128
Data Transformation Paths	130

Network Storage Systems	133
Storage Network System Configurations	134
Basic SAN Model	135
NAS Appliance	135
Enterprise NAS Device	138
In-Band SAN Appliance	139
Out-of-Band SAN Appliance	141
Cluster File System with Central Metadata	142
Symmetric Cluster File System	144
RAID Subsystem-Based Volume Replication	144
Server-Based Volume Replication	146
File-Based Data Replication	147
Summary	148
Chapter 6 The Storage in Storage Networking	149
Challenges for Network Storage	149
The Cost of Online Storage	150
Disk Economics	150
Enclosure Cost	151
The Cost of <i>Usable</i> Data Center Storage	151
The Bottom Line: Paying for Online Storage	152
Making SAN Storage Perform	153
Disk Aggregation and Virtualization	153
Disk Aggregation by Striping	154
Server-Based and Controller-Based Striping	157
Keeping SAN Storage Up and Working	157
RAID: Protection Against Disk Failures	158
Mirroring versus RAID versus Data Striping	160
Interconnect Failures	161
RAID Controller Failures	162
Transparent and Nontransparent Failure Tolerance	163
Atomic Operations and Data Integrity	164
Really Big RAID Controllers	165
Embedded RAID Controller Failure Tolerance	166
Volume Manager Failure Tolerance	166
Choosing among Storage Options	167
Summary	169
Chapter 7 The Network in Storage Networking	171
Fibre Channel: The 800-Pound Gorilla of SANs	171
Fibre Channel: The Standards	172
Fibre Channel: The Chips	172
Fibre Channel: The Devices	173
Fibre Channel: The Infrastructure	173
Fibre Channel Variations: Coping with Complexity	174

Fibre Channel Transmission Media	175
Fibre Channel Protocols: One Size Does Not Fit All	177
Fibre Channel Protocol Layers	178
Topology: The Shape of a Fibre Channel SAN	179
Fibre Channel Fabric Topology	182
Connecting Disks to a Fibre Channel Fabric	185
Emerging SAN Interconnect Technologies	186
New Wires for Storage Networks	187
Other Storage Networking Developments	189
Repartitioning Storage Functions	190
Summary	191
Chapter 8 Basic Software for Storage Networking	193
Software for SANs	193
The Data Center I/O Stack	194
Special Considerations for SAN Software	195
Discovering Devices	196
Controlling Access to Storage Devices	197
Controlling Access to Data Objects	198
Shared Access Data Managers	199
Lock Managers	200
How Lock Managers Work	201
Lock Management in Distributed Systems	202
Computer System I/O Performance	202
Cache and I/O Performance	202
I/O Load Balancing	203
Volumes: Resilience, Performance, and Flexibility	204
Technical Properties of Volumes	205
Failure-Tolerant Volumes: What Kind of Protection?	205
High-Performance Volumes: Balancing I/O Load	206
Mirroring, RAID, and Failures	208
System Crash Safeguards	208
I/O Performance with Volumes	210
Striped Volumes	211
Management Flexibility	213
File Systems and Application Performance	213
File System Space Allocation	213
Extent-Based File Systems	214
Preallocation and Alignment	216
Back Doors for Database I/O	217
Using Large Memory Effectively	218
Fast Recovery from System Crashes	219
Online File System Administration	221
File System Defragmentation	221
Moving Active Files	223

Online File System Expansion	223
Backup and Continuous Data Access	224
Frozen Images	224
Split Mirror Frozen Images	225
Copy-on-Write Frozen Images	226
Properties of Consistency	227
Other Types of Checkpoints	229
Summary	229
Chapter 9 Advanced Software for Storage Networking	231
Data Replication	231
Types of Data Replication	232
The Nature of Data Replication	233
Data Replication versus Mirroring	234
Different Types of Data Replication	235
Database Replication	235
File System Replication	236
Storage Device Replication	237
Replication Policies	237
Synchronous and Asynchronous Replication	238
Synchronous Replication	238
Asynchronous Replication	239
Using Data Replication	240
Hybrid Data Recovery	241
Disasters and Failures	241
Replication and Data Consistency	242
Reciprocal Disaster Recovery	242
Clusters: The Processing in Data Processing	244
Data Center Clusters	244
Data Center Clusters in Perspective	245
Why Clusters?	246
Applications and Clusters	247
Kinds of Resources	248
Starting and Stopping Resources	248
Dependencies among Services	249
Failover Services	249
Parallel Application Service Groups	250
Managing Cluster Resources	253
Cluster Interconnects	254
Client and Storage Interconnects for Clusters	255
Storage Devices for Clusters	255
Cluster Data Models	255
Distributed Lock Management	256
Server-Based Volume Management and Clusters	257
Volume Management in Clusters	257

Cluster File Systems	260
The Significance of Shared Data Clusters	261
Disaster Recovery and Global Clusters	262
Global Computing and Disaster Recovery	265
“Following the Sun”	266
Clusters and Storage Area Networks	266
Summary	267
Chapter 10 Enterprise Backup Software for Storage Area Networks	271
Backup Management for SANs	271
Enterprise Data Protection	272
Backup	272
Enterprise Backup Architecture	273
Backup Control and Data Flow	275
Scaling Backup Operations	275
Enterprise Backup Policies	277
Which Data Objects to Back Up	277
When to Back Up	278
Where to Back Up Data	278
Backup Automation	279
Minimizing the Impact of Backup	279
Full and Incremental Backup	279
The Impact of Incremental Backup	280
Different Types of Incremental Backup	281
Combining Cumulative, Differential, and Full Backups	281
Synthetic Full Backups	282
Copy-on-Write Technology and Backup	284
Backup and Databases	286
Changed Block Incremental Backup	286
Multistream Backup: Multiplexed and Parallel Backup Streams	288
Factors Affecting Backup Speed and Cost	288
Multiplexed Backups	289
Parallel Backup Streams	290
Parallelization of Large Backup Jobs	290
Sharing Tapes	291
LAN-free and Serverless Backup	291
Summary	293
Part Three: SAN Implementation Strategies	295
Chapter 11 Adopting Storage Networking	299
Why Adopt Storage Networking?	299
Benefits and the Mechanisms for Delivering Them	300
From Business to Storage	300

Technology and Business Benefits of SANs	300
SAN Benefits in a Larger Information Processing Context	302
How Much Downtime Is Too Much?	303
Who Needs SANs?	306
Developing a SAN Deployment Strategy	306
Nine Steps to Successful SAN Deployment	308
Step 1: Determining Business Information Requirements	308
Step 2: Learning Technology's Capabilities	309
Step 3: Matching Requirements with Technology	310
Step 4: Real Planning	311
Step 5: Auditing the Current Situation	312
Step 6: Merging IT Plans with Enterprise Plans	313
Step 7: Design and Vendor Selection	314
Step 8: Rollout	317
Step 9: Iteration, Review, and Learning	317
Critical Success Factors	318
SAN Adoption Alternatives	319
LAN-free Backup: Easing One's Way In	320
Storage Pooling: Administrative Flexibility	321
Frozen Image Backup: Eliminating the Window	321
Data Sharing: The NAS Way	323
Application Availability: Clustering for Failover	323
Shared Data and Parallel Clustering: Application Scaling	324
Data Replication I: Publication and Decision Support	325
Data Replication II: Disaster Recoverability	326
Going Global: Wide Area Clustering	327
The Potential for Disaster: One Egg, One Basket (Also Known as the SAN Two-Edged Sword)	328
The Other Edge	328
Summary	330
Chapter 12 Managing SANs	331
SAN Management Defined	331
Network Management Tasks	332
Storage Management Tasks	332
What Gets Managed in a SAN?	333
The Basics: What's Special about Managing a SAN	334
Discovery: "Hello, Who's Out There?"	334
But First, "Hello, I'm Here"	335
Next, "Who Else Is There?"	336
What's In a Name?	336
The New Meaning of <i>Host</i>	337
Zoning: Who Gets to Do What to Whom?	338
Routing: Accept What You Cannot Change	340
An Ideal SAN Management Environment	342

SAN Management and Online Storage Cost	343
Virtualization: Allocating Capacity Rather Than Disks	344
SAN Management and Quality of Online Storage Service	346
Storage Capacity	346
Data Availability	347
I/O Performance	348
New Types of Transaction Processing	352
SAN Management and Backup Cost	353
SAN Management and Backup Impact	355
SAN Management and Application Availability	357
SAN Management and Asset Utilization	358
SAN Management Tools	359
Who's in Charge, Anyway? Organizing to Administer a Storage Network	361
A Hidden Cost of SAN Deployment?	362
The State of SAN Management	363
SAN Management and Standards	364
Industry Association SAN Management Activities	365
Affinity Group SAN Management Activities	365
SAN Management Challenges	366
SAN Management Tomorrow: The Race to the Storage Dial Tone	368
Summary	370
Afterword Speculations on the Future of SAN	373
The Enterprise Data Wish List	373
New Interconnects, New Dynamics	374
Low-Latency NAS	374
A Single Enterprise Network Technology	375
Different Models, Different Capabilities	377
Automating Management: Virtualization and Quality of Storage Service	378
Availability: Abolishing Backup	378
Data over Distance: Files across the Water	379
New Partitionings: Using the Intelligence in Storage Devices	379
In Closing . . .	380
Appendix Storage Networking Organizations	381
Industry Associations	381
The Storage Networking Industry Association (SNIA)	381
The Fibre Channel Industry Association (FCIA)	383
Fibre Alliance (FA)	383

Infiniband Trade Association (IBTA)	383
Jiro	383
Standards Organizations	384
The American National Standards Institute (ANSI)	384
The Distributed Management Task Force (DMTF)	386
Storage Systems Standards Working Group (SSSWG)	387
The Internet Engineering Task Force (IETF)	
(http://www.ietf.org)	388
NDMP	388
NSIC	388
W3C	388
Dictionary of Storage Networking Terminology	389
Index	485

ACKNOWLEDGMENTS

My thanks to the board members of the Storage Networking Industry Association, who have contributed ideas and comments and have also reviewed early drafts.

My thanks also to Bill North, Guy Bunker, Paul Scammell, and Fred van den Bosch of VERITAS for their input and support.

Several suppliers of SAN software and hardware products offered assistance in the form of images to help illustrate this book. Their inclusion does not suggest that the particular supplier is any better or worse than any other—simply that the picture fit very well. We thank them all for their assistance.

—Richard Barker

I joined this project, which was originated by my manager and coauthor Richard Barker, in June 2000. Richard often jokes (at least, *I hope* he's joking) that I infiltrated his short project to produce a 100-page overview of SANs and turned it into a 500-page tome. Well, right he is about that! As I began to write, I discovered that there was a lot to be said about SANs and particularly about the systems and capabilities they enable that I felt hadn't been adequately covered in the existing literature. So the more we wrote, the more I felt there was to write. Even as we go to press, I can think of additional topics I'd like to cover, such as the role of a putative storage administrator, as well as more I'd like to say on several topics that are covered, like the emerging SAN technologies, whose descriptions are likely to be obsolete before this book reaches the shelves.

But I couldn't have come close to doing my part of this book alone. I wish to take this opportunity to thank the countless VERITAS colleagues and friends who contributed, sometimes unknowingly, to the constantly changing kaleidoscope through which I perceive SAN technology and the enterprise computer systems based on it. Some people who contributed in ways clearly above and beyond the call deserve to be singled out for special thanks:

To Andrea Westerinen, former Technical Director of the Storage Networking Industry Association, who provided the material for Appendix A on standards organizations and industry associations.

To the entire SNIA Technical Council for allowing me to participate in the discussions that led to Chapter 5 on storage network system architectures and particularly to David Black of EMC who kept my interpretation at least somewhat aligned with the

SNIA model. Also thanks to Bruce Naegel of VERITAS for several helpful suggestions on how to describe storage network architectures.

To Steve Macfarlane, Peter Grimmond, and Alan Randolph of the VERITAS Software Corporation consulting organization for providing real-world expertise on which Chapter 11 on SAN adoption processes is based. Particular thanks to Alan Randolph for his exhaustive review of the chapter and for enlightening me about the current state of SAN component interoperability as seen through the eyes of an implementer.

To Brett Cooper of the VERITAS SANPoint Control team for suggesting how to deal with the sprawling topic of SAN management and for reviewing and redirecting my interpretation of the topic.

I thank all these people for their contributions to this project and hasten to add that any remaining errors and omissions are to be laid solely at my door.

Last, I reiterate my thanks to my coauthor, Richard Barker, who managed to seamlessly merge me into a going project and who remained unruffled each time he discovered that his words had been morphed into something else. Richard proves over and over again that less management is more.

–Paul Massiglia

S*Storage Networking Essentials: A Systems Approach*, by Richard Barker and Paul Masiglia, is one of the most complete works addressing storage networking available today. Richard and Paul's thoroughly refreshing approach addresses storage networking from the data center. This is the first work I have had the opportunity to review that focuses on the complete solution from the enterprise application viewpoint. This unique approach to the highly complex topic of storage networking addresses the relevant topics required by systems administrators, implementers, and users of storage networking.

In order to understand my enthusiasm, I would like to take you on a short history of recent trends in storage networking.

Storage networking has been maturing since the introduction of Fibre Channel in the early 1990s. At first, Fibre Channel interfaces to storage devices were basically server-attached storage (SAS)—a faster, fatter pipe that would allow storage to be located at a greater distance from the server, but it was still a point-to-point connection. All of the storage was dedicated to a single server, and the data was locked to the server. If there was a failure of the system, the data could not be recovered until the system was operational again.

The introduction of Fibre Channel Arbitrated Loop (FCAL) for the first time allowed multiple servers and storage devices to exist on the same connection. FCAL marked the early demonstrations of storage networking, but it was not a very robust server connection. As a bus, it had shared bandwidth connection, it was cumbersome to implement, and it was wrought with early product failures and numerous compatibility issues. FCAL enjoyed limited deployment as a server to storage interconnection. Today it appears mainly as back-end disks to intelligent controller interconnect.

The full, modern-day SAN potential was not realized until the introduction of the first cost-effective Fibre Channel switch, resulting in a switched fabric (FCSW). For the first time, multiple servers as well as data storage devices could exist on a single storage network. Multiple servers could now effectively share large storage devices, allowing open systems the advantage of options previously available only to data center configurations. For the first time data protection, high availability, disaster recovery, and data center wide backup are now available to open system environments.

Bringing us back to the present. Today, a physical interconnection has been defined. What does the future hold? Two discernable directions are emerging.

The first is the quest for a better interconnection scheme. Advancements continue in the Fibre Channel arena—technology that is faster, cheaper, smaller, and cooler, with more intelligence, will be emerging. New interconnection technologies also continue

to emerge. SANs over IP and Infiniband are just two of these emerging technologies that will be competing with Fibre Channel in the future.

These new technologies will enable new functionality within the SAN. Interprocessor communications for data sharing and mobility, increased performance that allows more data to move more freely in the networks, longer and more reliable lines of communications, and much more.

The second development is the exciting new area of storage networking software, embracing the promise of technological breakthroughs. Storage networking has successfully demonstrated that a fundamental, interoperable interconnection infrastructure is alive and healthy, and is well on its way down the deployment path. We can now focus our attention on the real benefits of storage networking, software services, management, and automated policy-based operations. These areas will become the dominant focus of storage networking, propelling storage networking beyond mere interconnection strategies and into the realm of intelligent storage networking. . . . But I am getting ahead of myself.

Richard Barker and Paul Massiglia have assembled the essentials of the state of the art in storage networking with a focus on systems and software. Abundant high-quality illustrations, explanations, and definitions make this a must-read for anyone contemplating storage networking. “How This Book Is Organized” will give you some insight into the overall approach and recommendations for which section is best suited to your background.

I would also highly recommend this book for even the most battle-scarred veteran of storage networking.

It’s in my library.

–Larry Krantz
Chairman, Storage Networking Industry Association (www.snia.org)

Welcome to *Storage Area Network Essentials: A Complete Guide to Understanding and Implementing SANs*, the book that cuts through jargon, simplifies storage networking concepts, and helps you determine how your organization's information services department can exploit this exciting new technology to meet evolving information processing needs.

Why is this book necessary? High-speed interconnect technology has made *storage area networks*¹ (SANs) based on both Fibre Channel and Gigabit Ethernet a reality. Every major storage and server vendor offers *SAN-enabled* products that use these advanced interconnects. We predict that essentially every enterprise will eventually employ storage networking, simply because it will be the only alternative. The question for information services departments is whether they will be able to manage SAN deployment or whether it will just happen to them. The successful information services administrators and executives will be those who understand storage networking and become its master.

Storage area networks enable new ways of moving, storing, and accessing data that promote more reliable and cost-effective enterprisewide information processing. SANs simplify data processing and management by off-loading many resource-hungry routine functions to intelligent devices and networks dedicated to and optimized for storage I/O traffic.

It's a safe bet that your competitors are considering storage networking if they are not already using it. As with most compelling new technologies, you either adapt and adopt or you're left in the competitive dust.

About This Book

This book is designed to help the reader understand storage area networks—the properties, benefits, architectural concepts, technologies, and, most important, the pitfalls of storage networking. You won't find detailed descriptions of Fibre Channel protocols or RAID subsystems; other books cover those and, in the authors' opinion, they aren't of primary interest to our intended audience. What you will find that we believe is unique is a description of how the new capabilities enabled by storage networking are implemented in products you can buy. How do they work? Which ones are easy?

¹ If it's *storage networking*, why isn't it just *storage networks*? Where did the *area* come from? We give our view of how storage area networks got their name in Chapter 1.

Which ones are hard? Which ones should be implemented early and which should be delayed until you have some SAN experience under your belt? These are the questions that we think other books on storage networking *don't* cover.

We've spent relatively little time on the emerging interconnects iSCSI, FCIP, iFCP, and Infiniband, which are undergoing development and standardization as this book goes to press. We have concentrated instead on things you can buy today, because those are the products that create the issues you will have to grapple with for the next couple of years.

Who Are You?

So who is our intended audience? In writing this book, we have made some assumptions about our readership. First and foremost, you are up-to-date enterprise information processing technology professionals and are familiar with distributed client/server computing. Your concerns are not so much with the details of any one technology as with acquiring an overall grasp of what technologies can do for you and what they're going to cost you in new skills and ways of doing things. In addition . . .

Networking

You use networks and understand how they have revolutionized enterprise information usage. Your daily professional life includes electronic mail and other client/server applications. The more adventurous among you may operate home networks to share resources among your personal computers. We've spent little time selling the benefits of interconnected enterprises.

Storage

Likewise, we anticipate that you are familiar with the principles of enterprise data storage and management. You know about formatting and partitioning hard disks. On a larger scale, you are familiar with *enterprise RAID subsystems* that provide storage for several servers and employ advanced techniques such as mirroring to improve data reliability and availability. You understand why backups and archives of important data are business necessities and you're familiar with techniques for creating and managing them and what a pain in the neck they are for the people who have to keep applications running.

Strategy

We assume that you are either undecided about introducing storage networking into your organization or that your organization has made the decision to introduce a storage network and *now what?* We believe you want to learn how a SAN can improve your organization's ability to exploit its information assets before you make significant capital and human investments in the technology.

And Last . . .

We believe that you're a professional who fashions the constantly changing mix of computer and network system products on the market into solutions to your organization's information handling problems. Our goal with this book is to make your job easier by giving you a comprehensive introduction to a radically new enterprise data processing paradigm that's sweeping the information services community: *storage networking*.

How This Book Is Organized

There are three parts to this book.

Part One: Understanding Storage Networking

In this part, we present an overview of the SAN concept. We describe what SANs are, how they evolved, what they do, and why enterprises should be using them. We expect this part to be of interest to executives and managers whose responsibilities include enterprise budgeting and planning. The SAN overview and benefits description contained in this part should give budgeters some ideas about the value of SANs to their organizations.

Part Two: The Parts: What's in a SAN

In this part, we drill down into the architectural, storage, networking, and software components that constitute SAN technology, with particular emphasis on the software components. Because other books have covered storage and networking concepts extensively, we've concentrated more on what we perceived to be the white space—architecture and software. We expect this part to be of interest to SAN-based application developers and managers. The technology descriptions should be of assistance when designing application topologies and data center operational procedures.

Part Three: SAN Implementation Strategies

In this part, we describe strategies for moving an enterprise from *server-centric* computing with local storage to a *storage-centric* information processing environment in which the central resource is universally accessible data. We have also included an extensive chapter (Chapter 12) on taming the beast—managing a storage network once you have one. We expect this part to be of interest to IT planners and technical strategists whose responsibilities include planning and leading SAN adoption for their organizations, as well as to IT administrators who will have to keep storage networks up, running, and servicing the organization's information needs.

Other Stuff

There's a lot of churning going on in the data storage and networking industries as we write this book. New interconnects, new management paradigms, and components with new functional splits are starting to appear. We've included an afterword in which we speculate a bit about the future and what it might bring in the way of changes to enterprise data processing and management.

The appendixes contain handy reference information, including a list of standards organizations and other industry groups that concern themselves with storage networking and a glossary of the terminology likely to be encountered during further study of storage networking.

Some Conventions: Terminology to Start With

We define terms that we believe may be unfamiliar to some readers in place as they are introduced. Our usage corresponds to the Storage Networking Industry Association's *Dictionary of Storage Networking Terminology*,² which is reproduced in part in Appendix 2. In addition, here are a few terms to get you started.

In this book, we mention *enterprises* frequently. We have chosen this as a blanket term to represent corporations, academic and research institutions, government organizations, and other institutions that use computers to support their business. The focus of this book is on larger enterprises that have multiple servers.

We use the term *data center* to refer to a room or set of adjacent or nearby rooms whose primary purpose is to house computing, networking, and storage equipment. Larger enterprises typically operate several data centers.

It is remarkable that in an industry as mature as data storage, consistent terminology for very fundamental concepts has not been widely adopted. Through its online dictionary, the Storage Networking Industry Association is attempting to standardize storage networking terminology. In general, we have tried to be consistent with the SNIA's usage. There are a few key terms whose definitions will be especially helpful. These include:

- The term *disk drive* is used to represent the familiar rotating magnetic media random access nonvolatile data storage device.
- The term *disk* is used to denote either a disk drive (as above) or a virtual disk (sometimes called a logical unit number, or LUN) presented by a RAID controller. This book does not deal with removable disk media such as diskettes or other cartridge devices.
- The term *tape* is used to denote the linear magnetic recording media on which data is stored and accessed sequentially.
- The term *tape drive* is used to denote the transport mechanism in which tapes (as above) may be mounted in order to record and read data. In some cases, the term

² The Storage Networking Industry Association's *Dictionary of Storage Networking Terminology* can be found at www.snia.org/dictionary.

tape drive denotes both the mechanism and the media installed in it (e.g., *data is written to the tape drive . . .*).

- The term *media* is used to refer to objects on whose surfaces data can be permanently stored. Although the term *media* describes the circular platters on which a disk drive stores data, in practice, it is encountered more often when referring to storage objects like tapes and optical disks that can be separated from the mechanisms used to record data on them.
- The terms *library*, *media library*, and *robotic media library* are all used interchangeably to denote a machine that includes removable media storage, one or more tape drives, and a robotic mechanism to move individual tapes between storage and tape drives.
- The term *RAID subsystem* is used to denote a collection of one or more intelligent disk controllers that coordinate the operation of a common set of disk drives. Where the distinction is significant, the terms *embedded RAID controller* and *external RAID controller* are used to distinguish controllers mounted in server enclosures from those packaged separately from the servers they support. Large-scale RAID subsystems, which typically interconnect between dozens of disk drives and two or more computer access ports, are referred to as *enterprise RAID subsystems*.
- The term *storage device* is used to denote either a disk drive, a tape drive, a RAID subsystem, or a tape library containing both tape drives and a robotic media handler.

