

Making Everything Easier!™

2nd Edition

# Storage Area Networks

FOR  
**DUMMIES**®

## Learn to:

- Implement new technologies such as data de-duplication, iSCSI, and Continuous Data Protection
- Design storage area networks that meet specific needs
- Maintain and troubleshoot SANs
- Develop SANs that will aid your company's disaster-recovery plan

**Christopher Poelker**  
**Alex Nikitin**





# *Storage Area Networks*

FOR

# DUMMIES®

2ND EDITION

**by Christopher Poelker and Alex Nikitin**



WILEY

Wiley Publishing, Inc.

## Storage Area Networks For Dummies®, 2nd Edition

Published by  
Wiley Publishing, Inc.  
111 River Street  
Hoboken, NJ 07030-5774

[www.wiley.com](http://www.wiley.com)

Copyright © 2009 by Wiley Publishing, Inc., Indianapolis, Indiana

Published by Wiley Publishing, Inc., Indianapolis, Indiana

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) 646-8600. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permissions>.

**Trademarks:** Wiley, the Wiley Publishing logo, For Dummies, the Dummies Man logo, A Reference for the Rest of Us!, The Dummies Way, Dummies Daily, The Fun and Easy Way, Dummies.com, Making Everything Easier, and related trade dress are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries, and may not be used without written permission. All other trademarks are the property of their respective owners. Wiley Publishing, Inc., is not associated with any product or vendor mentioned in this book.

**LIMIT OF LIABILITY/DISCLAIMER OF WARRANTY: THE PUBLISHER AND THE AUTHOR MAKE NO REPRESENTATIONS OR WARRANTIES WITH RESPECT TO THE ACCURACY OR COMPLETENESS OF THE CONTENTS OF THIS WORK AND SPECIFICALLY DISCLAIM ALL WARRANTIES, INCLUDING WITHOUT LIMITATION WARRANTIES OF FITNESS FOR A PARTICULAR PURPOSE. NO WARRANTY MAY BE CREATED OR EXTENDED BY SALES OR PROMOTIONAL MATERIALS. THE ADVICE AND STRATEGIES CONTAINED HEREIN MAY NOT BE SUITABLE FOR EVERY SITUATION. THIS WORK IS SOLD WITH THE UNDERSTANDING THAT THE PUBLISHER IS NOT ENGAGED IN RENDERING LEGAL, ACCOUNTING, OR OTHER PROFESSIONAL SERVICES. IF PROFESSIONAL ASSISTANCE IS REQUIRED, THE SERVICES OF A COMPETENT PROFESSIONAL PERSON SHOULD BE SOUGHT. NEITHER THE PUBLISHER NOR THE AUTHOR SHALL BE LIABLE FOR DAMAGES ARISING HEREFROM. THE FACT THAT AN ORGANIZATION OR WEBSITE IS REFERRED TO IN THIS WORK AS A CITATION AND/OR A POTENTIAL SOURCE OF FURTHER INFORMATION DOES NOT MEAN THAT THE AUTHOR OR THE PUBLISHER ENDORSES THE INFORMATION THE ORGANIZATION OR WEBSITE MAY PROVIDE OR RECOMMENDATIONS IT MAY MAKE. FURTHER, READERS SHOULD BE AWARE THAT INTERNET WEBSITES LISTED IN THIS WORK MAY HAVE CHANGED OR DISAPPEARED BETWEEN WHEN THIS WORK WAS WRITTEN AND WHEN IT IS READ.**

For general information on our other products and services, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

For technical support, please visit [www.wiley.com/techsupport](http://www.wiley.com/techsupport).

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Library of Congress Control Number: 2008942264

ISBN: 978-0-470-38513-5

Manufactured in the United States of America

10 9 8 7 6 5 4 3 2 1



# *About the Authors*

**Christopher Poelker** has been in the field of computer technology since 1974. Chris was an electronics engineer in the U.S. Army, and tried to stay out of trouble by hiding in tanks while installing laser range finders and computer-aided ballistic trajectory systems. After leaving the service, Chris went to school in New York City at good old Control Data Institute and was hired as a field engineer by Digital Equipment Corporation. In his spare time, Chris started his own software company, developed databases, and became a Microsoft MCSE and instructor. Chris worked for Digital for 18 years until it was bought by Compaq, where he stayed on as a StorageWorks systems engineer until joining Sun Microsystems in 2000. Chris left Sun to become a consulting storage architect for Hitachi Data Systems and became the district storage manager for HDS in New York City. In 2006, Chris left HDS for FalconStor software, where he now works as the Vice President of Enterprise Solutions. Chris has designed and implemented storage networks for many of the Fortune 100 companies in the U.S. and around the world. In his spare time, Chris sometimes speaks at industry forums, writes magazine articles, and has acted as the SAN expert at SearchStorage.com.

**Alex Nikitin**, currently a systems expert at HBO, has logged in 15 years in the Information Technology industry. Alex has worn many hats in this industry, ranging from application programmer and network administrator to the ultimate responsibility over large server farms, and storage and backup solutions for some of the world's top financial and pharmaceutical companies. Alex and Chris worked together at Hitachi Data Systems in New York, where Alex was the "go to" guy for difficult storage designs and implementations. Prior to joining HDS, he also spent time growing the install base of Storage Area Networks as a Professional Services Consultant for EMC Corporation, implementing SAN solutions for various companies, large and small, on Windows NT/2000, Solaris, HP/UX, AIX, and Linux platforms. Seemingly always cast in a storage-centric capacity, his career has focused on caring for and reliably delivering vast amounts of storage to his user community.

# *Dedication*

**Christopher Poelker:** To my sister Nancy, whose love and friendship meant the world to me.



# *Author's Acknowledgments*

**Christopher Poelker:** I would have never been given the chance to write this book if it weren't for my friends at TechTarget, who used to run the SearchStorage.com Web site: Michelle Hope and Maryann Tripp. These two wonderful women were the reason I was introduced to Melody Layne and Susan Christophersen of Wiley Publishing, who made the first book possible, along with Teresa Artman, who spent many a long night copy-editing the manuscript and making up for my horrible writing skills. (I should have paid closer attention during eighth-grade English!) For this second edition, I would like to thank Kyle Looper and Kim Darosett for their patience during deadlines when my day job was getting in the way of keeping this edition on track.

I'd also like to thank my partner in crime and co-author, Alex Nikitin, who again saved my marriage and my duty as a father to my children by taking over some of the load and helping me crank out some of these chapters. Alex is one of the best storage guys I have ever had the privilege of working with.

Thanks also need to go this time to FalconStor software, for letting me proceed with this project and letting me play hooky now and then to crank out a chapter or two. Thanks to ReiJane Huai, Wendy Petty, Wayne Lam, Alan Chen, Tom Strumpf, Bruce Sasson, Joanne Ferrara, and everyone else at FalconStor who filled in or helped me out to give me time during the writing of this book.

Special thanks to all the folks who taught me most of the things I know about storage: my brothers, Lenny Poelker and Greg Poelker, who are a heck of a lot smarter than I! Also to my mentors: Wayne Lam, Wai Lam, Stanley Qin, Irving Moy, David Shyu, Cartic Vengkatraman, Gene Chesser, Steve Sicola, Paul Kruschwitz, Jimmy Wu, Raymond Tong, Paul Mitchell, Mike Mendola, Jo McCausland, John Lallier, Nick Sinish, Brian Rice, Steve O'Rielly, Catherine Brown, Frank Cizin, Leonard Hayward, Charlie Mulrooney, Paul Poon, Don Thatcher, Tony Merschdoff, Jeff Sinisgalli, John Fonseca, Marty Citron, Al Catalano, Tom Lindemann, Roland Song, Pierre Dansereau, Mike Pierro, and most of all, Kevin Shumacker, whose help over the years I could never repay. Thanks to Ken Garnau, Nancy Berliner, and Charlie Santana for help with the mainframe stuff, and everyone I ever met from Brocade/McData, and Emulex.

Tom Clark and Robert Stout, Nancy Jennings, and John Dorl from the original Nishan systems were a big help with helping me figure out SAN extensions, (especially Tom's books!), and Mark Farley's book was the first I read on SAN and still is one of the best. (I guess I learned things the hard way, by just doing it.)

Finally, thanks for the support of my family: Deborah, my wife, for being a single mother again while I was writing the book, and of course, my children Cole, Chris, and Rachel for all their support and being the wonderful people they turned out to be.

## **Publisher's Acknowledgments**

We're proud of this book; please send us your comments through our online registration form located at <http://dummies.custhelp.com>. For other comments, please contact our Customer Care Department within the U.S. at 877-762-2974, outside the U.S. at 317-572-3993, or fax 317-572-4002.

Some of the people who helped bring this book to market include the following:

### ***Acquisitions and Editorial***

**Project Editor:** Kim Darosett

**Executive Editor:** Steven Hayes

**Copy Editors:** Barry Childs-Helton, Susan Pink,  
Kathy Simpson

**Technical Editor:** Michael Vannette

**Editorial Manager:** Leah Cameron

**Editorial Assistant:** Amanda Foxworth

**Sr. Editorial Assistant:** Cherie Case

**Cartoons:** Rich Tennant  
([www.the5thwave.com](http://www.the5thwave.com))

### ***Composition Services***

**Project Coordinator:** Patrick Redmond

**Layout and Graphics:** Samantha Allen,  
Reuben W. Davis, Melissa K. Jester,  
Christine Williams

**Proofreaders:** Broccoli Information  
Management, Amanda Graham

**Indexer:** Broccoli Information Management

---

### ***Publishing and Editorial for Technology Dummies***

**Richard Swadley**, Vice President and Executive Group Publisher

**Andy Cummings**, Vice President and Publisher

**Mary Bednarek**, Executive Acquisitions Director

**Mary C. Corder**, Editorial Director

### ***Publishing for Consumer Dummies***

**Diane Graves Steele**, Vice President and Publisher

### ***Composition Services***

**Gerry Fahey**, Vice President of Production Services

**Debbie Stailey**, Director of Composition Services



# Contents at a Glance

---

<b><i>Introduction</i></b> .....	<b>1</b>
<b><i>Part I: SAN 101</i></b> .....	<b>5</b>
Chapter 1: The Storage Area Network.....	7
Chapter 2: SAN Building Blocks .....	23
Chapter 3: What Makes a SAN Go .....	63
Chapter 4: What Makes a SAN Stop .....	75
<b><i>Part II: Designing and Building a SAN</i></b> .....	<b>91</b>
Chapter 5: Designing the SAN.....	93
Chapter 6: SANs and Disaster Recovery .....	139
Chapter 7: Putting It All Together .....	157
<b><i>Part III: Using Advanced SAN Features</i></b> .....	<b>193</b>
Chapter 8: Networking SANs .....	195
Chapter 9: SAN-Based Backup.....	225
Chapter 10: Mirror, Mirror: Point-in-Time Copies .....	247
<b><i>Part IV: SAN Management and Troubleshooting</i></b> .....	<b>271</b>
Chapter 11: Approaches to SAN Management .....	273
Chapter 12: Troubleshooting SANs .....	305
<b><i>Part V: Understanding the Cool Stuff</i></b> .....	<b>333</b>
Chapter 13: Using Data De-Duplication to Lighten the Load.....	335
Chapter 14: Continuous Data Protection .....	357
Chapter 15: Everything You Ever Wanted to Know about Virtualization.....	379
<b><i>Part VI: The Part of Tens</i></b> .....	<b>405</b>
Chapter 16: Ten Reasons to Use a SAN.....	407
Chapter 17: Ten Reasons NOT to Use a SAN.....	413
<b><i>Index</i></b> .....	<b>419</b>



# Table of Contents

---

<b><i>Introduction</i></b> .....	<b>1</b>
About This Book .....	1
Foolish Assumptions .....	2
Conventions Used in This Book .....	2
How This Book Is Organized .....	3
Part I: SAN 101 .....	3
Part II: Designing and Building a SAN .....	3
Part III: Using Advanced SAN Features .....	3
Part IV: SAN Management and Troubleshooting .....	3
Part V: Understanding the Cool Stuff .....	3
Part VI: The Part of Tens .....	3
Icons Used in This Book .....	4
<b><i>Part I: SAN 101</i></b> .....	<b>5</b>
<b>Chapter 1: The Storage Area Network</b> .....	<b>7</b>
Defining a SAN .....	7
Fiber versus Fibre .....	8
How a SAN Makes Computing Different .....	9
Understanding the Benefits of a SAN .....	10
Finding Out Whether a SAN Is Right for You .....	12
Who should use a SAN? .....	12
Who should not use a SAN? .....	13
Dissecting a SAN (The Four Ps) .....	14
The Parts of a SAN .....	14
The host layer .....	14
The fabric layer .....	15
The storage layer .....	16
Storage arrays .....	16
The SAN Protocols .....	18
The SAN Players .....	20
The SAN Platforms .....	20
Applications that benefit from a SAN .....	21
Applications that require a SAN .....	22
<b>Chapter 2: SAN Building Blocks</b> .....	<b>23</b>
SAN Components and How They're Used .....	23
The Host Layer .....	26
Host bus adapters .....	26
Gigabit gadgets: GBICs and GLMs .....	27



The Fabric Layer .....	29
Understanding storage fabrics .....	29
SAN hubs .....	30
SAN switches .....	33
Data routers .....	38
Cables .....	40
Cable-connector types .....	43
SAN ports and port naming .....	43
Basic SAN port modes of operation .....	44
Protocols used in a Fibre Channel SAN .....	46
The Storage Layer .....	48
Storage arrays: Storing your data .....	48
Explaining Redundant Array of Inexpensive Disks (RAID) .....	49
RAID benefits .....	49
RAID types .....	50
Logical Unit Numbers (LUNs) .....	53
Understanding storage-array classification .....	55
Modular versus monolithic .....	56
Monolithic (Enterprise) .....	56
Modular (Departmental) .....	57
Why Cache memory makes a difference .....	58
<b>Chapter 3: What Makes a SAN Go .....</b>	<b>63</b>
Networking Basics .....	63
Moving Data at the Speed of Light .....	64
Bandwidth .....	66
Fibre Channel Protocols .....	67
The arbitrated loop .....	68
Loop addressing .....	69
The Switched Fabric .....	70
The fabric protocol .....	71
Fabric addressing .....	72
<b>Chapter 4: What Makes a SAN Stop .....</b>	<b>75</b>
Discovering What Causes SAN Problems .....	75
Preventing Poor SAN Design .....	76
Bandwidth .....	77
Too much distance between components .....	78
Excess latency .....	79
Congestion .....	80
Over-subscription .....	81
Using the Right Cables in the Right Way .....	82
Avoiding connection issues .....	83
Macro- and micro-bends, and the patch panel pain .....	83
Cable labeling .....	85
Choosing the right host bus adapter for your computer .....	85

Going with a single vendor ..... 86  
 Mixing switch vendors ..... 89

***Part II: Designing and Building a SAN..... 91***

**Chapter 5: Designing the SAN ..... 93**

Basic SAN Designs: Understanding the Layers ..... 93  
 Point-to-Point Topology ..... 96  
 Arbitrated Loop Topology..... 98  
     Cascading hubs ..... 99  
     Loop of hubs..... 101  
     Creating resilient hub networks..... 102  
     Fault-tolerant loops ..... 104  
 Switched Fabric Topology..... 105  
     Types of SAN switches ..... 106  
     Choosing which switches to use..... 107  
     Using the right bandwidth for the job..... 108  
     Trunking and what it's used for ..... 108  
 Basic Fabric Topologies..... 110  
     Dual switches, the SAN fabric building block ..... 110  
     Loop-of-switches topology ..... 113  
     Meshed fabric topology ..... 114  
     Star topology ..... 117  
     Core-edge topology..... 119  
 Understanding Zoning..... 123  
     The parts of a zone ..... 125  
     Types of zoning ..... 126  
     Zone alias names ..... 127  
 Initial Switch Setup..... 130  
     Setting up a Brocade switch ..... 131  
     Setting up an original McData (Brocade) director switch..... 135  
 Best Practices — Tips from the Trenches ..... 135  
     When to choose a director-class switch ..... 136  
     Standardize on a single vendor's switches..... 136  
     Standardize your firmware versions ..... 136  
     Standardize your HBA drivers..... 137  
     Use unique zone alias names ..... 137  
     Using storage from multiple vendors..... 137  
     Always use two fabrics..... 137

**Chapter 6: SANs and Disaster Recovery ..... 139**

How Much Downtime Can You Afford?..... 140  
 Gathering the data for a disaster-recovery plan..... 141  
 Create a detailed plan that meets your requirements ..... 143

Recognizing the Importance of Distance, Bandwidth, and Latency .....	145
Distance.....	145
Bandwidth.....	145
Latency .....	147
Choosing the Recovery Site .....	148
Existing facility .....	149
Co-location facility .....	150
Choosing Where to Run the Data Replication Process.....	151
Host-based data replication solutions .....	151
Appliance-based data replication solutions.....	152
Array-based data replication solution .....	154
Shipping tapes as a solution.....	155
The Importance of Testing.....	156

## **Chapter 7: Putting It All Together . . . . . 157**

Building a SAN by Hand .....	158
The SAN Plan.....	159
Fabric Zoning 101.....	160
LUN security .....	161
Setting Up the SAN .....	162
Keeping good notes .....	162
Setting up the switches.....	162
Preparing the Servers .....	170
Loading the driver .....	171
Customizing the HBA card's configuration .....	171
Planning the HBA connections.....	171
Configuring the Array.....	172
The hardware .....	172
RAID setup .....	174
Plugging Things In .....	177
Configuring the Zones.....	180
Mapping the zones first.....	180
I'm zoning out . . . ..	181
Back to the Servers: Did It Work?.....	183
Unix servers.....	184
Windows system .....	187
iSCSI, You SCSI, We All SCSI .....	188
Initiators and targets.....	188
IQN: iSCSI qualified name.....	189
SCSI Name Service .....	189
Data domains.....	190
Getting started with iSCSI .....	190
Getting serious with iSCSI.....	191
Data Migration .....	191
Network migration.....	192
Backup/restore migration.....	192
Disk-to-SAN migration .....	192

***Part III: Using Advanced SAN Features ..... 193***

**Chapter 8: Networking SANs .....195**

- Defining a SAN Island ..... 196
- Connecting SAN Islands ..... 198
  - Disk/data sharing ..... 198
  - Data copying ..... 199
- The Storage WAN, MAN, and SWAN ..... 199
  - Using the network only for storage management ..... 200
  - The storage SWAN ..... 202
- Choosing and Using SAN Extenders ..... 202
- Choosing the Correct Link for the Job ..... 204
  - IP connections ..... 205
  - OC-type connections ..... 206
- Reducing Costs with Compression, Data De-duplication  
and WAN Tuners ..... 208
  - Compression ..... 208
  - De-duplication ..... 210
  - WAN tuners ..... 211
- SAN Connection Protocols ..... 211
  - FCIP: The SAN tunnel ..... 211
  - iFCP: The SAN gateway ..... 213
- Stretching the SAN (The Rubber-Band Approach) ..... 216
- Using Connected SAN Islands (The Two-Rubber-Bands Approach) ..... 219
- Using a SAN as Network Attached Storage ..... 220
- iSCSI: An Alternative Method ..... 223

**Chapter 9: SAN-Based Backup .....225**

- Understanding Backup ..... 225
- Understanding SAN Backup ..... 227
  - The backup window ..... 227
  - Tape drives ..... 228
  - Tape libraries ..... 229
  - Backup policy ..... 230
- Choosing a Backup Solution ..... 230
  - Integrated tape drive in each server ..... 232
  - Backup over a corporate LAN to a tape drive  
connected to an independent backup server ..... 233
  - Backup over corporate network to robotic  
tape library connected to an independent backup server ..... 234
  - LAN-less backup to shared tape library over SAN ..... 235
  - Serverless backup to shared tape library through SAN ..... 236
  - Disk-to-disk backup ..... 237
  - Image copy and snapshots in the SAN ..... 239
  - SAN data replication/remote backup ..... 242

Determining How Long a Backup Will Take .....	242
Determining backup speeds .....	244
The formula for backup .....	245
<b>Chapter 10: Mirror, Mirror: Point-in-Time Copies .....</b>	<b>247</b>
The Uses of Point-in-Time Technology .....	248
Make backups .....	248
Make corruption-recovery images .....	250
Save space .....	251
The possibilities are endless .....	252
Complete versus Metadata Copies .....	254
Which PiT Type Should You Use? .....	256
Creating a PiT Copy .....	257
Managing Your Point-in-Time Copies .....	259
Pair up your volumes .....	259
Create the pairs .....	260
Splitting the mirror, snapping a copy .....	261
Doing a resync: Refresh and restore .....	262
I need my data now! .....	263
Using a PiT copy .....	264
The Finer Points of PiT .....	264
Guideline #1: Understand when to snap a copy .....	265
Guideline #2: Keep your PiT pairs separated .....	267
Guideline #3: Use the right disk storage for PiT copies .....	268
Questions to ask your SAN vendors .....	270
<b>Part IV: SAN Management and Troubleshooting .....</b>	<b>271</b>
<b>Chapter 11: Approaches to SAN Management .....</b>	<b>273</b>
Management: From Simple Networking to SANs .....	274
SAN Management from the Ground Up .....	275
Start small; think big .....	275
Documentation is key .....	275
Cable Management: Spaghetti, Anyone? .....	276
Physical cable management .....	277
Logical cable management .....	279
Labeling Your Cables .....	279
Data center coordinate system .....	280
Standard naming convention .....	282
Documenting the cable arrangements .....	283
Using a SAN Management Framework .....	283
Working with SAN management software .....	284
Communicating in a common language .....	284
Speaking the language yourself .....	285
Putting everything together .....	286



What SAN Management Gives You.....	286
A bird’s-eye view of your network.....	287
Agent-based management.....	288
Health monitoring.....	289
Records of events .....	291
Change management .....	291
Predictions of problems.....	292
Streamlining SAN Administration.....	292
Step-by-step administration .....	292
Using a framework .....	293
Automating Your System: “SAN? Do You Read Me, SAN?”.....	294
Backing up .....	295
Managing database storage.....	296
Providing a Service Level Agreement .....	296
Simple versus complex SLAs.....	297
Setting service levels.....	298
Building a Storage Management Team .....	298
SAN architects.....	299
SAN engineers .....	300
Monitoring team.....	301
SLA and performance specialists.....	301
Provisioning staff .....	302
Planning for the future .....	302
Common responsibilities .....	303

**Chapter 12: Troubleshooting SANs . . . . . 305**

The Best Method: Prevention .....	306
Troubleshooting Methodology.....	306
Go with what you know.....	306
Elementary, my dear Watson .....	307
Didn’t read the manual, did you?.....	309
Build a golden configuration .....	309
Typical Problem Types.....	310
Obvious problems .....	310
Phantom problems .....	316
Continuous problems.....	318
Catastrophic problems .....	322
Example Scenarios .....	322
Scenario #1 .....	323
Scenario #2 .....	326
Scenario #3 .....	328

**Part V: Understanding the Cool Stuff..... 333****Chapter 13: Using Data De-Duplication to Lighten the Load . . . . . 335**

Understanding Data De-Duplication.....	336
Benefits of data de-duplication.....	336
How de-duplication works.....	337
Data De-Duplication in the Datacenter.....	338
The de-duplication vendors.....	339
How data gets de-duplicated.....	339
In-band versus out-of-band data de-duplication.....	346
Using Data De-Duplication in a SAN.....	347
Files: Honey, I shrunk the files.....	347
Blocks: Been there, stored that.....	347
What about hash collisions?.....	350
Why Data De-Duplication Is Important.....	352
When to Use Data De-Dupe (And When Not To).....	354
Applications for which data de-duplication makes sense.....	354
Applications for which data de-duplication doesn't make sense.....	355
De-Duplication in Action.....	356

**Chapter 14: Continuous Data Protection . . . . . 357**

Understanding What Continuous Data Protection Is.....	357
How CDP Makes Storage Work Like a Database.....	358
CDP data journaling.....	359
Splitting the writes.....	360
CDP makes everything different.....	363
Sizing a CDP journal.....	365
Best Practices for Storage When Configuring CDP Solutions.....	367
The Truth about Near CDP and True CDP Solutions.....	368
Example 1: Recovering a database with traditional backup.....	369
Example 2: Recovering a database from a snapshot.....	370
Example 3: Recovering a database using near CDP.....	370
Example 4: Recovering a database using true CDP.....	370
CDP versus Snapshots.....	371
Using CDP to Eliminate Backups.....	372
Using CDP to Simplify Recovery and Reduce Costs.....	374
Knowing Your CDP Vendor.....	376

**Chapter 15: Everything You Ever Wanted to Know  
about Virtualization . . . . . 379**

Understanding What Virtualization Is.....	379
Exploring the Types of Virtualization.....	381

Implementing Virtualization in a Datacenter ..... 383  
     Server virtualization ..... 386  
     Storage virtualization ..... 389  
 In-Band versus Out-of-Band Virtualization ..... 399  
     In-band virtualization ..... 399  
     Out-of-band virtualization..... 400  
 The Virtualization Vendors and Where They Play ..... 401  
     Host-based storage virtualization vendors ..... 401  
     Fabric-based virtualization storage vendors..... 402  
     Storage-array-based virtualization ..... 402

***Part VI: The Part of Tens..... 405***

**Chapter 16: Ten Reasons to Use a SAN .....407**

You Want Better Disk Utilization..... 407  
 You Need a Good Disaster Recovery Solution  
     for Multiple Applications ..... 408  
 You Need Better Availability for Your Applications ..... 408  
 You Need More Storage Room..... 409  
 Backup Is Taking Too Long ..... 409  
 You’re Focusing on Server and Storage Consolidation ..... 409  
 You’ve Been Tasked to Save Your Company Money ..... 410  
 You Need to Manage Storage for Many Locations from a Central Site ... 410  
 You Need to Decrease IT Management Costs ..... 411  
 You Need Better Performance for Your Applications ..... 412

**Chapter 17: Ten Reasons NOT to Use a SAN.....413**

You Need Larger File Servers..... 413  
 You Only Have a Few Inexpensive Servers ..... 414  
 You Want to Save Your Company Money This Year..... 414  
 You Want to Use the Latest and Greatest Solutions Available..... 415  
 You Need a Disaster-Recovery Solution for a Single Application ..... 415  
 You Want a SAN but Don’t Have the Budget..... 415  
 You Use Gigabit Ethernet on Your LAN..... 416  
 Everything Already Runs Fine..... 416  
 You Need to Back Up Multiple Remote Offices over Slow Links ..... 417  
 You Need to Replicate Your Data for Disaster  
     Recovery but Can’t Afford Fast WAN Connections ..... 418

***Index..... 419***



# Introduction

---

**W**elcome to *Storage Area Networks For Dummies*, 2nd Edition! The first edition was the book we wish was around when *we* were trying to learn about this stuff! We tried to take a fairly mundane topic and make it a fun read so you could get up to speed on storage networking as quickly and painlessly as possible. This second edition is written in the same spirit as the first.

When the first edition was written in 2003, very few books on storage area networks (SANs) were available. The books that were available were very narrow and extremely technical in focus. They were about as fun to read as the directions for setting up a DVD player. Although more books about storage networking are available now, most are still very technical and about as fun to read as the U.S. tax code (nothing against accountants here!).

Over the years, this book has become the standard bearer in keeping the subject concise, simple, and fun, and has now been updated to include all the new cool stuff and changes that have occurred since the original printing. Let's face it — most folks typically look for a fast and easy way to get information; you want the information you *need* to know, and not everything there *is* to know, so this second edition uses that same point of view common to all of us poor slobs who need to make a SAN work with no budget, no training, and no time. So don't worry, be happy — and just have some *fun!*

## *About This Book*

The mission of this book is to help you find your way around while exploring the topic of storage area networking. The book is dedicated to individuals who, for better or for worse, have been tasked with designing, implementing, configuring, or troubleshooting a storage area network. We hope that the information here will enable both the beginner and the moderately expert storage professional sort through the ins and outs of a SAN. We use familiar language to demystify the technology and translate the jargon as necessary. You'll discover how to choose the right hardware for the job, design a SAN by using the best practices in the industry, connect everything to make it work, and troubleshoot the SAN to fix problems when they occur.

You also get to delve into the hard stuff — the stuff that many companies pay expensive consultants for (who have usually just read a book like this just before they were hired!), so be brave, and we'll do our best to make this painless. If you want to understand what a SAN is and what it does, you've come to the right place. Everything you should need (or want) to know about storage area networks is here in one location.

## *Foolish Assumptions*

We have one or two foolish assumptions about you, the reader of this book:

- ✓ You are responsible for or have worked with computer storage before.
- ✓ You might want to continue working in that field of endeavor.
- ✓ You want to find an easy way out of your networking storage problems so you can get home and play video games, watch TV, and drink beer with your friends.

We also assume that you don't want to become an instant expert. You want to know just enough to be able to speak intelligently on the subject matter but also know when it's time to call in an expert.

This book tries to get you past all the little details that get in the way of understanding a subject by using a real-world approach. We didn't create the technology; we just know how to make it work because we've been in the field for more years than we care to say. This book tries to impart just enough good info to help you make your stuff work — and to understand enough to know when someone is misleading you or trying to rip you off. The best defense is a good offense.

## *Conventions Used in This Book*

We want you to understand all the instructions in this book, and in that spirit, we've adopted a few conventions.

When you hit a chapter in which we ask you to do something, you will be prompted by a numbered list. The numbers in the list are the order of the steps that you need to take to accomplish the task at hand. Just follow the steps listed, and everything should be just fine. If you need to enter something on the keyboard, we ask you to *type* it. If you need to use your mouse, we ask you to *click* it. That's all there is to it.

## How This Book Is Organized

This book is designed as a reference, so you don't have to read the book cover to cover. Just look in the Table of Contents to find the topic you're interested in and start reading. If you're unsure of some of the acronyms you've been hearing out there, check out the Cheat Sheet at the front of this book for easy reference.

The book is organized in seven parts.

- ✔ **Part I: SAN 101:** This part covers the basics of storage area networks, including what you need to know if you're going to buy a SAN, build one yourself, or have someone build it for you.
- ✔ **Part II: Designing and Building a SAN:** This part of the book handles all the fun stuff, such as how to add more storage to your servers and how to connect everything and get it running. It also covers what you need to know if you already have a SAN in place or need to know how to use it or set it up properly.
- ✔ **Part III: Using Advanced SAN Features:** If you've always wanted to know what the heck a snapshot copy was, this is the place to find out. This part also covers advanced topics such as backing up your data, which will help you get more bang for your buck out of your SAN.
- ✔ **Part IV: SAN Management and Troubleshooting:** Every now and then, something goes bump in the night. This part shows you how to manage and troubleshoot problems when they occur as well as how to avoid having to face those problems in the first place.
- ✔ **Part V: Understanding the Cool Stuff:** This part covers the cool new advances in storage area networking since 2003, such as storage virtualization, data de-duplication, and advances in data protection and replication such as Continuous Data Protection (CDP). You find out how to use these technologies to help you save money and become more productive and better prepared when trouble happens.
- ✔ **Part VI: The Part of Tens:** This part includes ten reasons to use a SAN and ten reasons *not* to.



We've also provided a bonus chapter titled "Outsourcing SAN Solutions" that you can download from the book's companion Web site at [www.dummies.com/go/sanfd2e](http://www.dummies.com/go/sanfd2e).

## Icons Used in This Book

To help you get the most out of this book, we've placed icons here and there. Here's what the icons mean:



Next to the Tip icon, you can find shortcuts and tricks of the trade to make you more productive without even realizing it.



Where you see the Warning icon, tread softly and carefully. It means that we've been burned by this already and don't want you to have to learn the hard way, as we did.



Stuff marked with the Remember icon is like jotting a note to yourself in a class. Make an effort to bend the ear of the page so that you don't forget it.



Okay, we probably put too many of these icons in the book. But what the heck . . . sometimes trying to explain this stuff is like writing a book on *Brain Surgery For Dummies*. We need to point out the details at times so you don't end up with a migraine.



# Part I

# SAN 101

## The 5<sup>th</sup> Wave

By Rich Tennant



"They can predict earthquakes and seizures,  
why not server failures?"

## *In this part . . .*

**T**he computer industry is funny. As soon as you get comfortable with the latest technology and become the resident expert, something new comes out, and the whole learning process starts all over again. Some people enjoy the challenge of learning about something new, and some think it just makes life more difficult. This first part of the book tries to make things easier by introducing you to storage area networks, or SANs. Just like in high school shop class, when you were introduced to the drill press, the table saw, and the first aid stations, we introduce you to the various tools that you use to build a SAN. We tell you what each tool does, why it's necessary, and how you can use it in your new SAN project.

## Chapter 1

# The Storage Area Network

---

### *In This Chapter*

- ▶ Understanding storage area networks (SANs)
  - ▶ Determining whether a SAN is right for you
  - ▶ Looking at SAN layers and protocols
  - ▶ Figuring out which operating systems benefit from SANs
  - ▶ Discovering which applications can use or require SANs
- 

**T**his chapter is dedicated to helping you get a handle on what a storage area network (SAN) is, the basics of how one works, and whether one is right for your needs. You'll discover all the parts that make up a SAN, the things that make one run, and who actually makes all the different parts that you can buy. Putting a SAN together is somewhat like putting together one of those high-end stereo systems; you have many components and many different manufacturers to choose from. This chapter helps you choose the ones that suit your needs and create something that you can be proud of.

These days, becoming proficient with SANs can mean a major boost to your career. Perhaps you're bored to death in your current position and would like a change of pace. SAN administration is one of the highest-paying jobs in Information Technology (IT) today. If you add storage area networking to your résumé, you may find your phone ringing off the hook as headhunters vie to offer you a six-figure income (hey, might as well dream big).

## *Defining a SAN*

First, the basics. In today's terms, the technical description of a SAN (Storage Area Network) is a collection of computers and storage devices, connected over a high-speed optical network and dedicated to the task of storing and protecting data.

In a nutshell, you use a SAN to store and protect data. A SAN uses the SCSI (Small Computer Storage Interconnect) and FC (Fibre Channel) protocols to move data over a network and store it directly to disk drives in block format. Today, that high-speed network usually consists of fiber-optic cables and switches that use light waves to transmit data with a connection protocol known as Fibre Channel. (A *protocol* is a set of rules used by the computer devices to define a common communication language.) More and more, regular Internet protocol (IP)–based corporate networks, and even the Internet, are being used as the network part of a SAN. IP networks that are already in place can be used by other storage connection protocols such as iSCSI (internet Small Computer Storage Interconnect) to move and store data.

Using a network to create a shared pool of storage devices is what makes a SAN different. A SAN moves data among various storage devices, allows sharing data between different servers, and provides a fast connection medium for backing up, restoring, archiving, and retrieving data. SAN devices are usually bunched closely in a single room, but they can also be connected over long distances, making a SAN very useful to large companies.

Many of today’s SAN components are pretty much plug-and-play. To create a simple SAN, you just connect all the devices together with cables, and off you go. Creating larger SANs with many storage switches can become complex, though, and that’s the reason for this book: to give you a handle on what you need to know about large, complex SANs.

## *Fiber versus Fibre*

No, it isn’t just a snooty way of spelling *fiber*. (Well, okay, not *only* that.) Networking geeks use the *fibre* spelling (reversing the *er* to *re*) to refer specifically to fiber-optic cables used in a SAN. The idea is to differentiate SAN cables from the optical cables used in other networks (such as TCP/IP Networks). That’s because SAN devices use a different language to communicate with each other than do the devices in other networks. This is why the main protocol used in a SAN (snooty or not) is called *Fibre Channel*.

All network protocols are divided into layers, like a layer cake. All the layers in the cake are logically tied together into a *stack*. Each layer of the stack provides different functionality, and each device in the network uses the stack like a language to communicate with other devices in the network. The bottommost layer of the stack is hardware-based (as opposed to software-based), and thus is referred to as the *physical layer*.

The physical layer consists of tangible hardware stuff such as cables, switches, and connectors. This is where the fiber-optic cables are. On top of the physical layer are the software layers that make up the *protocol stack*. In a Fibre Channel SAN, those layers make up the *Fibre Channel* protocol.

Each type of network uses a different protocol to handle data. The Internet, for example, uses a protocol stack called the Transmission Control Protocol/Internet Protocol (TCP/IP). The physical layers of both Internet and SAN can transmit data as light pulses over fiber-optic cables — which (as you might expect) makes the data move nearly as fast as light. The only difference between regular fiber-optic computer networks such as the Internet and a fiber-optic SAN is the protocol and the switches used by the devices to talk to each other over the network. SANs use the *Fibre Channel* protocol and Fibre Channel switches, and the Internet uses the *TCP/IP* protocol and Ethernet switches. Fibre Channel was developed to move data really fast between computers and disk drives; TCP/IP (or “Internet Protocol”) was developed to move files over long distances between computers.

## How a SAN Makes Computing Different

Using a SAN can really change how you think about computing. In the past, there was the *mainframe*, which was a gigantic computer that could run all the programs in a large business. All the computer stuff was gathered in one place called a *data center*. All the storage that the mainframe needed was directly connected to it. Everything was located and managed as a single, large entity.

The PC revolution changed a lot of things. Everything started to spread out. Data was moved off the mainframe and stored in server computers. The servers were then dispersed throughout the enterprise to bring computing power closer to the actual users. The servers became connected by a network, called a *local area network*, or LAN. This was cool because now the computing power was spread out and made more available to end users. Eventually, LANs were connected to create the Internet.

Networks enabled people who used computers in far-flung places to communicate and share information with each other. In business, problems arose when inter-networking finally took off. A great deal of data was now being stored with no effective way to manage it all. Managing all the scattered data dispersed throughout the network became a nightmare.

Because all data storage was located inside each individual server, you had no effective way to efficiently allocate storage space between all the servers. Sure, users could share files over a LAN, but you still needed a way to share access to physical disks, rather than using dedicated disks inside every server. Hence the advent of the SAN.



Since the original TCP/IP network protocols used in a LAN (Local Area Network) were built to move and share files, they had no built-in way to directly access disk drives. As a result, very high-performance applications needed direct access to *block-based* disk drives to move and store data very fast. (Data is stored as blocks on a disk drive.)

Disk drives in a SAN are stored in a dedicated storage device called a *disk array*. All the servers connect to the storage device over a high-speed network using the Fibre Channel protocol, which enables very fast access to disks over a network. Using a SAN gives businesses shared and consolidated access to data storage — available to any server connected to the SAN.

Putting a SAN in place makes individual server computers less important and more peripheral to the data stored in the SAN. After all, the data is what is important to your business. If you lose a server, you can buy a new one. If you lose your data, it's "Adiós, amigo" for your business.

## *Understanding the Benefits of a SAN*

The typical benefits of using a SAN are a very high return on investment (ROI), a reduction in the total cost of ownership (TCO) of computing capabilities, and a pay-back period (PBP) of months rather than years. Here are some specific ways you can expect a SAN to be beneficial:

- ✔ **Removes the distance limits of SCSI-connected disks:** The maximum length of a SCSI bus is around 25 meters. Fibre Channel SANs allow you to connect your disks to your servers over much greater distances.
- ✔ **Greater performance:** Current Fibre Channel SANs allow connection to disks at hundreds of megabytes per second; the near future will see speeds in multiple gigabytes to terabytes per second.
- ✔ **Increased disk utilization:** SANs enable more than one server to access the same physical disk, which lets you allocate the free space on those disks more effectively.
- ✔ **Higher availability to storage by use of multiple access paths:** A SAN allows for multiple physical connections to disks from a single or multiple servers.
- ✔ **Deferred disk procurement:** That's business-speak for not having to buy disks as often as you used to before getting a SAN. Because you can use disk space more effectively, no space goes to waste.
- ✔ **Reduced data center rack/floor space:** Because you don't need to buy big servers with room for lots of disks, you can buy fewer, smaller servers — an arrangement that takes up less room.
- ✔ **New disaster-recovery capabilities:** This is a major benefit. SAN devices can mirror the data on the disks to another location. This thorough backup capability can make your data safe if a disaster occurs.
- ✔ **Online recovery:** By using online mirrors of your data in a SAN device, or new continuous data protection solutions, you can instantly recover your data if it becomes lost, damaged, or corrupted.