

Storage Networks: The Complete Reference

About the Author

Robert Spalding, a technology consultant, leveraged his 30 years of advanced technology experiences to develop *Storage Networks: The Complete Reference* for McGraw-Hill/Osborne. Being at the computer technology forefront, he was one of the first to use dyadic processing within the mainframe environment, setting the foundations for high-end symmetrical processing. His open systems management research and study led the way in developing the first management tools for massive parallel processing configurations that are still used today. As a research analyst covering the storage industry, he identified and published the early research that evolved into complex storage network technologies of today.

Robert is currently working on advanced storage strategies for imaging technologies and distribution as well as archival technologies for unstructured data products such as audio and video. *Storage Networks: The Complete Reference* is the latest in a string of more than 50 publications that include research papers, technical articles, and music compositions. He can be reached through his independent consultancy at RobertSpalding@kentrow.com

About the Contributing Writers

Ian Spalding has worked as a technical writer and editor for several start-up companies. He is currently at work on his first novel.

Himanshu Dwivedi is a Managing Security Architect at @stake, Inc., where he leads the Storage Center of Excellence (CoE), which focuses research and training around storage technology, including Network Attached Storage (NAS) and Storage Area Networks (SAN). At @stake, Himanshu forms part of the San Francisco-based Professional Services Organization (PSO) providing clients with network architecture assessments, attack and penetration services, secure network design, and secure server analysis. He is also part of the @stake Academy, where he is a lead instructor in several security classes.

Himanshu is considered an industry expert in the area of SAN security, specifically Fibre Channel Security. He has given numerous presentations and workshops regarding the security in SANs, including the SNIA Security Summit, Storage Networking World, TechTarget, the Fibre Channel Conference, SAN-West, and SAN-East. He holds a wide spectrum of well-diversified security experiences, specializing in the telecommunications industry. Additionally, Himanshu has worked with major organizations in the United States, Europe, South America, and Asia, including some of the major software, manufacturing, and financial-based entities.

About the Technical Editor

Giota Vavasis Then (Patty) started her professional career in 1982 as an applications programmer. In 1986, Patty began using her programming skills in the data center to write utilities for data-center automation as an MVS Systems Programmer while supporting 24/7 Enterprise Data Center operations. While in this capacity, Patty developed expertise in areas of Data Center Security, Storage Management, Capacity Planning, Performance Measurement and Modeling, Business Continuance, and Disaster Recovery, to name a few. In 1992, Patty moved from working in the data center to working as a vendor. Past companies she has represented include Hitachi Data Systems, EMC Corp, StorageTek, and, most recently, XIOtech/Seagate Technologies.

Patty has held various positions over the last 21 years in the IT industry, such as Systems Support Engineer, Sales and Consulting Engineer, National Accounts Manager, Customer Service District Manager, Product Marketing Manager, and Sales and Education Consultant.

Awarded numerous honors over the years for her outstanding contributions, Patty currently heads the "Consortium Technologies Group," a consulting and education practice based in the Tampa Bay Area, that brings together a team of world-renowned experts with real-world customer-focused experience.

Storage Networks: The Complete Reference

Robert Spalding



Tata McGraw-Hill Education Private Limited

NEW DELHI

McGraw-Hill Offices

New Delhi New York St Louis San Francisco Auckland Bogotá Caracas
Kuala Lumpur Lisbon London Madrid Mexico City Milan Montreal
San Juan Santiago Singapore Sydney Tokyo Toronto

S I T Library
Valachil, Mangalore



Accn No: 017352

Srinivas Institute of Technology

Acc. No.:

7352

Call No.:



Tata McGraw-Hill

Storage Networks: The Complete Reference

Copyright © 2003 by The McGraw-Hill Companies, Inc. All rights reserved. No part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of the publisher, with the exception that the program listings may be entered, stored, and executed in a computer system, but they may not be reproduced for publication

Tata McGraw-Hill Edition 2003

Seventh reprint 2009
RLCQCRCFRYLDC

Reprinted in India by arrangement with The McGraw-Hill Companies,
Inc., New York

Sales territories: India, Pakistan, Nepal, Bangladesh, Sri Lanka and Bhutan

ISBN-13: 978-0-07-053292-2

ISBN-10: 0-07-053292-3

Published by Tata McGraw-Hill Education Private Limited,
7 West Patel Nagar, New Delhi 110 008, and printed at
Pashupati Printers Pvt. Ltd., Delhi 110 095

The McGraw-Hill Companies

Contents at a Glance

Concepts of Storage Networking

| | | |
|---|---|----|
| 1 | The Data Storage and Data Access Problem | 3 |
| 2 | The Battle for Size and Access | 17 |
| 3 | Decoupling the Storage Component: Putting Storage on the Network | 35 |
| 4 | Decoupling the Storage Component: Creating a Network for Storage | 49 |

Storage Fundamentals

| | | |
|---|-----------------------------------|-----|
| 5 | Storage Architectures | 65 |
| 6 | Device Overviews | 77 |
| 7 | Connectivity Options | 95 |
| 8 | Data Organizational Methods | 115 |

Network Attached Storage

| | | |
|----|--|-----|
| 9 | Putting Storage on the Network: A Detailed Discussion .. | 131 |
| 10 | NAS Hardware Devices | 143 |

| | | |
|----|--------------------------------|-----|
| 11 | NAS Software Components | 157 |
| 12 | NAS Connectivity Options | 175 |

Part II Storage Area Networks

| | | |
|----|--------------------------------------|-----|
| 13 | Architecture Overview | 199 |
| 14 | Hardware Devices | 219 |
| 15 | Software Components | 235 |
| 16 | Configuration Options for SANs | 253 |

Part III Application—Putting It Together

| | | |
|----|---|-----|
| 17 | Defining the I/O Workload | 275 |
| 18 | Applying the SAN Solution | 293 |
| 19 | Applying the NAS Solution | 315 |
| 20 | Considerations When Integrating SAN and NAS | 335 |

Part IV Management—Keeping It Running

| | | |
|----|------------------------------------|-----|
| 21 | Planning Business Continuity | 353 |
| 22 | Managing Availability | 365 |
| 23 | Maintaining Serviceability | 383 |
| 24 | Capacity Planning | 401 |
| 25 | Security Considerations | 421 |

Part V Appendixes

| | | |
|---|--|-----|
| A | NAS Case Study: The International Image Processing Company | 453 |
| B | SAN Case Study: The Import Auto Industry | 465 |
| C | SAN/NAS Management Case Study: The Southwestern CD Company | 477 |
| D | Glossary | 489 |
| | Index | 501 |

Contents

| | |
|-----------------------|------|
| Foreword..... | xix |
| Acknowledgments | xxix |
| Introduction | xxxi |

Part I

Concepts of Storage Networking

| | | |
|----------|---|----|
| 1 | The Data Storage and Data Access Problem | 3 |
| | The Challenge of Designing Applications | 4 |
| | Non-Linear Performance in Applications | 7 |
| 2 | The Battle for Size and Access | 17 |
| | Why Storage Networking Makes Sense | 21 |
| | The Problem: Size | 22 |
| | The Problem: Access | 22 |
| | Business Applications Defined for Storage | 23 |
| | Maintenance and Support Applications | 24 |
| | Building the Business Application | 25 |
| | The Benefits of Storage Networks on Business Applications | 28 |
| | Scalability of Access to Business Application Data | 30 |
| | Consolidation of Business Application Data | 31 |

| | |
|--|----|
| Faster Access to Data Through Improved Bandwidth and Centralized Configurations | 31 |
| Increased Business Application Availability Through Efficiencies in Storage Management | 31 |
| The Effects of Storage Networks on IT Organizations | 33 |

| | |
|---|----|
| 3 Decoupling the Storage Component: Putting Storage on the Network | 35 |
| The NAS Idea | 36 |
| NAS Evolution and Development | 38 |
| An Operational Overview of NAS | 44 |
| Leveraging NAS for Data Access | 44 |
| Leveraging NAS for Data Size | 45 |
| NAS Caveats | 47 |

| | |
|---|----|
| 4 Decoupling the Storage Component: Creating a Network for Storage | 49 |
| The Data-centric World | 50 |
| Distributing Computer Processing | 51 |
| Symmetric Multiprocessing Systems (SMP) | 51 |
| Massive Parallel Processing Systems (MPP) | 52 |
| Distributing I/O Processing | 53 |
| An Enhanced I/O Protocol—Fibre Channel | 54 |
| A High-Speed Interconnect—Switched Fabric Network .. | 55 |
| “Shared Nothing” Architecture | 55 |
| The SAN Idea | 56 |
| SAN Evolution and Development | 57 |
| An Operational Overview of SAN | 58 |
| SANs for Data Access | 60 |
| SANs for Data Size | 60 |
| The Caveats for SAN | 61 |

Part II

Storage Fundamentals

| | |
|---|----|
| 5 Storage Architectures | 65 |
| Storage Elements | 66 |
| Data Storage Hierarchy | 68 |
| Storage Systems | 70 |
| Typical Storage System Configurations | 71 |

| | |
|---|----|
| 6 Device Overviews | 77 |
| Peripheral Connectivity Components and Concepts | 78 |



| | |
|-----------------------------|----|
| The Host Adapter | 81 |
| The Controller | 83 |
| Magnetic Disk Storage | 84 |
| Disk Systems | 85 |
| Disk Arrays | 86 |
| RAID Storage Arrays | 87 |
| Magnetic Tape Storage | 90 |
| Tape Drives | 90 |
| Formats | 92 |
| Read/Write Mechanisms | 92 |
| Tape Systems | 92 |
| Optical Storage | 93 |







| | |
|--|-----|
| 7 Connectivity Options | 95 |
| Connections: Getting on the Bus | 96 |
| Bandwidth | 97 |
| Addressing | 98 |
| Interrupt Functions | 98 |
| Electro-Mechanical | 98 |
| Bus Evolution | 99 |
| Bus Operations | 99 |
| Parallel vs. Serial | 100 |
| Differences in Bus and Network Architectures | 101 |
| The PCI Bus | 103 |
| The SCSI Bus | 103 |
| Fibre Channel | 105 |
| Fibre Channel Operation | 107 |
| Other I/O Buses | 109 |
| USB and Firewire | 110 |
| Creative Connection Strategies | 111 |



| | |
|--|-----|
| 8 Data Organizational Methods | 115 |
| Organizing Data | 116 |
| Interrelationships of Performance | 117 |
| Virtualization Times Three | 117 |
| Requirements for Enterprise-Class File Systems | 117 |
| File Systems | 119 |
| Allocation | 119 |
| Volume Definition | 120 |
| File Allocation | 120 |
| Management | 120 |
| Types of File Systems | 121 |
| Databases | 125 |
| Looking for Data | 125 |
| Database Systems: Different Animals | 127 |

Part III

Network Attached Storage

| | | | |
|---|-----------|---|-----|
|  | 9 | Putting Storage on the Network: A Detailed Discussion | 131 |
| | | The NAS Architecture | 132 |
| | | The NAS Hardware Architecture | 133 |
| | | The Server Part | 133 |
| | | The Storage Part | 134 |
| | | The Network Part | 135 |
| | | The NAS Software Architecture | 136 |
| | | The Micro-Kernel OS Part | 136 |
| | | NAS File Systems | 137 |
| | | Network Connectivity | 138 |
| | | NAS as a Storage System | 139 |
| | | The Departmental NAS Architecture | 139 |
| | | The Internet NAS Architecture | 139 |
| | | The Enterprise NAS Architecture | 141 |
|  | 10 | NAS Hardware Devices | 143 |
| | | NAS Hardware as the I/O Manager | 144 |
| | | NAS Hardware I/O Manager Components | 146 |
| | | The NAS Processing Models | 152 |
| | | NAS Workloads and NAS Hardware | 154 |
|  | 11 | NAS Software Components | 157 |
| | | A Remote File System with Attached Storage | 158 |
| | | The NAS Micro-Kernel OS | 161 |
| | | The NAS File System | 165 |
| | | The Storage Software | 168 |
| | | NAS Software as an I/O Manager | 172 |
| | | NAS Software and NAS Processing Models | 172 |
| | | I/O Workloads and NAS Software | 173 |
|  | 12 | NAS Connectivity Options | 175 |
| | | Putting Storage on the Network | 176 |
| | | NAS Connectivity Architecture: Hardware | 176 |
| | | NAS and LAN Topologies | 177 |
| | | NAS Connectivity Architecture: Software | 189 |
| | | An Overview of TCP/IP Processing | 189 |
| | | NAS Connectivity in Operation | 193 |
| | | Network Configurations and Workloads | 194 |




Part IV**Storage Area Networks**

| | | |
|-----------|--|-----|
| 13 | Architecture Overview | 199 |
| | Creating a Network for Storage | 201 |
| | The Network Part | 202 |
| | The Software Part | 205 |
| | The Hardware Part | 209 |
| | The Connectivity Part | 211 |
| | SAN Configurations | 213 |
| 14 | Hardware Devices | 219 |
| | The Fibre Channel Switch | 220 |
| | The F_Port | 221 |
| | The E_Port | 223 |
| | The G_Port | 223 |
| | Host Bus Adaptors | 224 |
| | Putting the Storage in Storage Area Networks | 227 |
| | JBOD | 227 |
| | RAID | 228 |
| | Bridges and Routers | 229 |
| | Fabric Operation from a Hardware Perspective | 231 |
| | SAN Hardware Considerations | 234 |
| 15 | Software Components | 235 |
| | The Switch's Operating System | 236 |
| | Device Drivers | 239 |
| | The Supporting Components | 241 |
| | Device Sharing | 242 |
| | Data Sharing | 244 |
| | SAN Management Applications | 246 |
| | ISL Functions | 249 |
| | Considerations for SAN Software | 252 |
| 16 | Configuration Options for SANs | 253 |
| | Connecting into the Data Center | 254 |
| | Implementing Multiswitch Configurations | 256 |
| | Supporting the Server Population | 258 |
| | Supporting the External Networks | 262 |
| | The Evolving Network Connections | 262 |
| | The Evolving Device Connections | 264 |
| | Integrating Tape Solutions | 266 |
| | Looking for an Optical Connection | 267 |


| | |
|------------------------------------|-----|
| Connecting to the Future | 267 |
| SAN Configuration Guidelines | 270 |

Part V

Application—Putting It Together


| | | |
|---|--|-----|
|  17 | Defining the I/O Workload | 275 |
| | Storage Planning and Capacity Planning | 277 |
| | The Definition and Characterization of Workloads | 278 |
| | The Business Application | 281 |
| | Deposit Application—Workload Attributes | 281 |
| | I/O Content and Workloads | 285 |
| | The Data Organizational Model | 285 |
| | User Access | 287 |
| | Data Paths | 289 |
| | Considerations for I/O Workloads in Storage Networking | 290 |
| | SAN or NAS Solutions | 292 |
|  18 | Applying the SAN Solution | 293 |
| | SAN Workload Characterization | 294 |
| | Applying the SAN to OLTP Workloads | 299 |
| | The Data Organizational Model | 300 |
| | User Access | 301 |
| | Data Paths | 302 |
| | The Design and Configuration of OLTP-Based Workloads | 302 |
| | Applying the SAN to Web Transactional Workloads | 304 |
| | The Data Organizational Model | 305 |
| | User Access | 305 |
| | Data Paths | 306 |
| | The Design and Configuration of Web-Based Workloads | 306 |
| | Applying the SAN to Data Warehouse Workloads | 308 |
| | The Data Organizational Model | 309 |
| | User Access | 310 |
| | Data Paths | 310 |
| | The Design and Configuration of Data Warehouse Workloads | 310 |
| | Conclusions on Applying SAN Solutions | 312 |
| | Points to Consider Before Proceeding | 312 |
| | Preparing for the Solution | 312 |
|  19 | Applying the NAS Solution | 315 |
| | NAS Workload Characterization | 316 |
| | What NAS Does Well | 317 |
| | What NAS Can't Do | 318 |




| | |
|--|-----|
| Applying the NAS to Departmental Workloads | 324 |
| The Data Organizational Model | 325 |
| User Access | 325 |
| Data Paths | 325 |
| Design and Configuration for Departmental-Based Workloads | 325 |
| Applying the NAS to Enterprise Web Workloads | 327 |
| The Data Organizational Model | 327 |
| User Access | 327 |
| Data Paths | 328 |
| Design and Configuration for Web-Based Workloads | 329 |
| Applying the NAS to Specialized Workloads | 329 |
| The Data Organizational Model | 329 |
| User Access | 330 |
| Data Paths | 330 |
| Design and Configuration for Specialized Workloads | 330 |
| Conclusions on Applying NAS Solutions | 330 |
| Consider These Points Before Proceeding | 330 |
| Preparing for the Solution | 332 |

| | |
|---|-----|
|  20 Considerations When Integrating SAN and NAS | 335 |
| The Differences | 336 |
| The Similarities | 338 |
| The Need to Integrate | 341 |
| Future Storage Connectivity | 342 |
| iSCSI, LAN/WAN Storage | 342 |
| InfiniBand, the Universal Bus | 344 |
| Inside the Box | 346 |
| Storage Integration | 348 |

Part VI

Management—Keeping It Running

| | |
|--|-----|
|  21 Planning Business Continuity | 353 |
| Defining the Environment (Putting Storage in Context) | 356 |
| Categorizing Interruptions in Service | 357 |
| Work from the Top Down | 359 |
| Participation | 359 |
| The Role of Storage Networking in Business Continuity | 360 |
| How Much Insurance Do You Need? | 361 |
| Level of Redundancy | 361 |
| Level of Requirements—How Much It Will Cost? | 362 |
| Level of Associated Requirement—How Much It Will Really Cost? | 362 |

| | | |
|---|--|-----|
| | Storage Design and Implementation of the Business | |
| | Continuity Plan | 362 |
| | SAN | 362 |
| | NAS | 363 |
| | Remote vs. Local | 363 |
| | Backup vs. Disaster | 363 |
| | Expertise and Configuration Data | 363 |
|  | 22 Managing Availability | 365 |
| | Availability Metrics | 367 |
| | External Service Level for Storage Networks | 367 |
| | Internal Service Levels for Storage Networks | 370 |
| | Implementing the Plan | 371 |
| | On-Demand Availability—NAS and SAN | 373 |
| | Highly Available Systems—SAN maybe NAS | 375 |
| | Fault-Tolerant Systems—SAN only | 375 |
| | Storage Availability—Considering RAID | 376 |
| | Finding the Holes | 377 |
| | Where to Look | 377 |
| | Data Recovery | 379 |
| | Closing the Loop | 382 |
|  | 23 Maintaining Serviceability | 383 |
| | Tracking the Configurations | 385 |
| | Physical Configuration Management for SANs | 386 |
| | Logical/Software Configuration Management for SANs | 387 |
| | Physical Configuration Management for NAS | 388 |
| | Logical/Software Configuration Management for NAS | 390 |
| | Investigating the Changes | 392 |
| | Driving Change Through Configuration Management | 392 |
| | How Problem Management/Resolution Becomes the | |
| | Ying and Yang of Storage Networks | 393 |
| | The Effects of Capacity Upgrades on Change | |
| | Management | 393 |
| | Tracking Changes to Physical and | |
| | Logical Configurations | 394 |
| | Guidelines for Establishing Change Management in | |
| | Storage Networking | 395 |
| | Whose Problem Is It Anyway? | 396 |
| | Closing the Loop on Serviceability | 400 |
|  | 24 Capacity Planning | 401 |
| | Developing a Plan for Storage | 404 |
| | The Storage Analysis | 407 |
| | Leveraging a Storage Upgrade | 407 |

| | | |
|-----------|--|------------|
| | Establishing a New Storage Network | 408 |
| | Tools for Storage Analysis | 409 |
| | Collecting Storage Information | 410 |
| | Modeling Performance and Capacity Requirements | 414 |
| | Implementing the Plan | 417 |
| | Closing the Loop | 419 |
| 25 | Security Considerations | 421 |
| | Overview of Information Security | 424 |
| | Authentication | 424 |
| | Authorization | 425 |
| | Auditing | 425 |
| | Integrity | 425 |
| | Encryption | 425 |
| | Availability | 425 |
| | Security Methods | 426 |
| | Storage Security Technology | 427 |
| | Inline Encryption | 428 |
| | Encryption at Rest | 428 |
| | Key-Based Authentication | 430 |
| | Two-Factor Application Authentication | 431 |
| | Storage Security Challenges | 432 |
| | Products and Solutions | 433 |
| | Encryption | 434 |
| | Security Standards and Specifications | 434 |
| | Bandwidth/Functionality Tradeoffs | 435 |
| | Fibre Channel SAN Security | 435 |
| | Zoning | 435 |
| | World Wide Name Spoofing | 436 |
| | Frame Weaknesses | 437 |
| | LUN Masking | 441 |
| | E_Port Replication | 443 |
| | NAS Security | 443 |
| | Common Internet File System | 443 |
| | Network File System | 444 |
| | Best Practices | 445 |
| | Passwords | 446 |
| | Configuration | 446 |
| | Management | 447 |
| | Remote Connectivity | 448 |
| | Operating System Security | 449 |
| | Storage Appliance Services | 449 |

Part VII

Appendixes

| | | |
|----------|--|-----|
| A | NAS Case Study: The International Image Processing Company | 453 |
| | The Situation Analysis | 455 |
| | IIP Systems Infrastructure | 455 |
| | Problem Detail | 458 |
| | The Search and Evaluation | 459 |
| | IIP Storage Solutions | 461 |
| | IIP Storage Evaluation Results | 462 |
| | IIP Storage Acquisition and Implementation | 462 |
| | The IIP NAS Solution | 463 |
| B | SAN Case Study: The Import Auto Industry | 465 |
| | The Situation Analysis | 467 |
| | The Search and Evaluation | 469 |
| | Available Alternatives | 470 |
| | Storage Evaluation Results | 471 |
| | Solution Acquisition and Implementation | 472 |
| | SAN Operational Issues | 472 |
| | New SAN Software Tools | 473 |
| | The SAN Solution | 474 |
| C | SAN/NAS Management Case Study: The Southwestern CD Company | 477 |
| | The Deferred Responsibility Scenario— or Who’s Doing What to Which Data? | 479 |
| | The Problem | 480 |
| | Business Situation | 481 |
| | Technical Scenario | 481 |
| | Problem Dialogue | 481 |
| | Epilogue | 482 |
| | Evaluation | 484 |
| | Steps to Better Storage Networking Management | 487 |
| D | Glossary | 489 |
| | Index | 501 |

Foreword

Sean Derrington
Senior Program Director
Infrastructure Strategies
META Group, Inc.

IT organizations (ITOs) increasingly view storage as a strategic differentiator for existing and new applications. Many ITOs recognize the critical role information plays in the business, and leveraging that information requires an unassailable storage infrastructure foundation. I believe forward-looking ITOs have to continue to decouple storage and server selection. This aids in ensuring storage solutions are viewed as much more than an afterthought for new applications. For most organizations, I believe this is a best practice, and the trend will continue as the storage market matures in technical capabilities and more IT professionals become well versed in the arts and values of storage. Even though current capital expenditures are under extreme scrutiny (and will continue to be through 2003), leading ITOs are acquiring storage solutions on a tactical basis—yet, they are planning and selecting the vendors on a strategic basis.

Overall, storage is moving to a services model, comparable to that of IP networking and (future) application infrastructure. Storage services will be tiered in function, price, and complexity, and delivered to applications and lines of business. For ITOs to deliver on this direction, storage (all aspects of it—disk, tape, disaster recovery, capacity planning, and so on) must be perceived and treated as an infrastructure component. Infrastructure components are common and reusable across as many consumers (for example, applications and servers) as possible; storage should be no different. Consequently, this requires a storage infrastructure to be an initial

consideration (not an afterthought) to drive business priorities. Moreover, infrastructure design and planning teams must have seats at the business project table and translate technology into business value.

Typically, IT architects create architectural principles and guidelines addressing three- to five-year planning cycles, while infrastructure teams design and plan component technologies (vendor- and product-specific), targeting current through 36-month requirements. The challenge going forward will be how ITOs handle the transition from storage infrastructure planning and design to storage operations. This is a critical success factor for ITOs in general, but especially for storage management, because capabilities are rapidly maturing and vendors are rapidly innovating. I believe that, through 2004, storage infrastructure and storage operations may be organizationally in the same group—but in the longer term (2005/06), I believe these responsibilities will diverge and become distinct groups, mirroring current systems, network, and application infrastructure and operations teams. The objective of infrastructure planning is *“to determine the scope, scale, and design of infrastructure necessary to provide application service levels required by the business in the short, medium, and long term.”* Furthermore, *“the primary design goal for information systems must be to enable rapid change in business processes and in the applications and technical infrastructure that enable them!”*

The implementation of automated networked storage (Storage Area Network [SAN] or Network Attached Storage [NAS]), which rationalizes (reduces the number of variants in) storage management applications and optimally integrates these functions, will need to leverage tiered storage software, hardware, and professional services to be able to measure (both internally and for the business) a successful storage infrastructure. Automated networked storage will provide businesses with agility, dramatically increase productivity, and enable ITOs to manage terabytes and petabytes successfully and cost-effectively. This book will address, in detail, the various architectural considerations, the components that satisfy each storage architecture, and various connectivity options for a networked storage environment.

Storage and Application Service-Level Agreements (SLAs)

Because all applications and servers are not created equal (as determined by business requirements), prioritization for storage services (from premium enterprise-class to internal workgroup storage) must be evaluated. By 2004/05, more than 60 percent of the servers in the data center will be connected to networked storage, and I expect more than 70 percent of storage capacity to be networked in this time frame. Although there can be modest increases in procurement costs to network storage, the resulting benefits in management productivity, increased utilization, agility, infrastructure and operations personnel savings, and scalability are substantial.

Servers sharing external storage resources can also aid in holistic capacity planning and increase overall capacity utilization (such as in GB/TB, Fibre Channel [FC] fabrics, and external subsystems—impacting data-center floor space, and so on). For applications

that are more volatile, storage can be reassigned to the servers as the capacity threshold (60–80 percent utilization) is reached. This is an area where enterprise storage procurement and capacity planning are important (compared to internal storage). I generally recommend a six- to nine-month buying cycle for enterprise storage and a competitive storage environment. Overall, ITOs can plan on an approximate 35 percent per year decrease in hardware prices (8–10 percent per quarter) and, given the state of the current global economy, should proceed cautiously on long-term acquisitions (12+ months).

Storage Management

Enterprise storage management will likely surface as the most significant contributor to ITO operational savings, because human resources remain the highest data-center cost. These costs will be increasingly important as more servers participate in networked storage (the disk and tape SAN). ITOs must continue to remove the dependencies and “stove piped” thinking of operating systems and storage management. Many operational economies of scale and flexibility will be gained by employing such an approach (for example, capacity managed per administrator). Certainly, storage management will still need to interface with the other operational disciplines (such as systems management, network management, database administration), but a center of storage management excellence is the design objective.

The goal is to streamline and automate many of the daily tasks that consume staff resources. For example, during the past 24 months, FC fabric management has moved to a core function under the purview of many enterprise storage vendors. And I believe additional tasks, such as storage provisioning/allocating (reducing the number of individuals involved and the time required from hours or days to minutes), resource management (which users/applications are using particular capacities), and topology mapping (which resources—physical or logical—are allocated to applications) will be among the first tasks within multivendor storage administration.

The key is to examine how and which operational tasks can be automated and leveraged (across server operating systems and multivendor storage hardware), enabling storage capacity to increase while storage infrastructure and storage operations staff remain constant—or, hopefully, decrease over time.

For automation efficiencies, storage operations should identify solutions that are seamlessly integrated—that is, a single application that can incorporate various functions (possibly multivendor) and not a loosely coupled launching pad (or iconic launch) for disassociated islands of automation. Solutions leveraging a central repository (persistent database) serving as the “storage management data warehouse” would enable ITO efficiencies gained through the function of individual “management components.” Moreover, harnessing this repository and “feeding” (analogous to the real-time data warehousing feedback loop) information to other management elements or ISV applications (such as for chargeback/billing) will be of significant value. Without this next level of integration/consolidation, users will duplicate efforts, not achieve the ability to scale (performance, capacity, and so on) their environments, nor control storage staff costs.

I believe storage provisioning/allocation will be one of the initial fully automated areas, yet this will require organizational (at minimum, functional) changes, because currently this task typically transcends individual and possibly departmental responsibilities (for example, server, database, storage). Simplifying storage allocation, which requires a multitude of individual tasks, is not a simple problem to solve. It typically involves the following:

- Assigning/configuring the appropriate LUNs from the physical storage subsystem
- Identifying and allocating the number and location of host interconnects on the array
- Masking the LUNs to the particular server(s)
- Zoning the FC switched fabric
- Incorporating and modifying any necessary FC security policies (if available)
- Ensuring the volume manager/file system and the application and/or database on the host(s) are updated, and incorporating the desired changes
- Ensuring other storage management applications (for instance, backup/recovery, remote replication services) comprehend the changes and the implications of such changes in storage policies

Center-of-excellence (COE) best practices dictate investigating the benefits/returns of consolidation and integration tasks, the most significant of which are those that have the greatest, most immediate impact on the costs that govern operations. Those gains will be seen as an annuity to the business, versus one-time savings. This annuity could be significant not just in dollars, but also in business agility and infrastructure flexibility.

Application Recoverability

Application/DBMS recoverability is one of the most compelling reasons for networked storage (predominately SAN-based storage) participation. Based on the mean time to recovery (the time it takes from the recognition of a failure to the time the application begins accepting transactions again—not simply restoring information from tape to server) discussed with the business, numerous technology options can assist in meeting those requirements. Snapshot or volume-based replication (either local or remote) can dramatically reduce the recovery times to minutes, preventing the need for a tape restore (but not eliminating the need for tape backup/recovery).

Although restoration is a significant and necessary capability, the “state” of the information/data can be much more important than the existence of the information. Recovering data in an inconsistent state—or data from one database that reflects a point in time that is significantly different from another reliant database—is often useless. Consistently managing replication and recovery options (possibly from multiple

vendors) will prove to be operationally beneficial, as will rationalizing backup/recovery software and hardware across as many platforms as possible (often excluding the mainframe). Consistently managing various storage vendors' snapshot/volume replication can provide further operational savings, striving for policy-based automation. Moreover, just as all applications are not created equal, storage replication software is no different. Integration with operating systems, clustering software, applications, and databases, as well as quiescent time (the time required for a consistent view of the application), will often differentiate offerings, given that replication software has become much more prominent during the past 36 months.

Data and Media Center of Excellence

There are three high-level layers of the data and media COE: conceptual (policy), physical (infrastructure), and operations (operational). Each of these three is a superset comprised of multiple functional and subfunctional elements. Each of the functions within each category and associated job titles (suggested) are detailed as follows and serve as a guideline for functional and organizational structure.

It is also important to note that, although many of these functional categories exist in the mainframe environment, simply having mainframe personnel responsible for all systems is not necessarily appropriate. There must be a balance of people, processes, and technology to create an effective COE, and the technologies are often quite different between mainframe and non-mainframe environments. Nonetheless, COEs should be viewed holistically and include the following.

Policy

Policies offer the highest-level view of storage management, dictate the required infrastructure to meet policy mandates, and should be governed by a storage policy director (SPD). These policies, which are interdependent and should be viewed holistically, include the following:

- **Service portfolios and service-level agreements** SPDs should determine what service packages they will offer to business units, as well as the associated cost/service tradeoffs, based on the needs of the business units.
- **Data protection strategies** Not all data is created equal, and the data protection mechanisms should be aligned with the service portfolio.
- **Vendor selection strategies** ITOs should determine, as a matter of policy, whether to pursue a best-of-breed multivendor strategy (albeit rationalized and not a vendor "du jour" strategy), or a "one throat to choke" single-vendor strategy. This will vary from one category to another, but should be done with an eye toward strategic objectives, which are not always mutually exclusive. Either way, rigorous interoperability mandates should be a prerequisite for consideration to enable maximum deployment flexibility.

Infrastructure

After the overall storage policies have been determined, a storage infrastructure architect should be responsible for designing and implementing the components needed to achieve the service goals. This will include device and network management (at the functional level) as well as a multitude of subfunctions, including the following:

- **Physical topology design** In most organizations, it will be necessary to use and integrate SAN, NAS, and direct-attached storage either for optimum performance or for security requirements and data protection. The ability of vendors to not only function/coexist, but also support (via customer and professional services expertise) these advanced heterogeneous environments (multiprotocol, multivendor, multifunction, cross-organizational, and cross-regional) will determine the most appropriate platforms (both hardware and software) on which to base infrastructure for the flexibility to adapt, as subelements of technology evolve over time.
- **Performance modeling and monitoring** The storage infrastructure architect should be responsible for not only designing sufficient resources, but also for ensuring that service levels are met. This includes storage components (for example, disk arrays) as well as storage networks (such as FC fabrics) and working with IP networking personnel, in the case of NAS.
- **Security** Security is an underaddressed area in storage management, due mostly to the current lack of capability, but should be made the responsibility of the storage infrastructure architect working in conjunction with the enterprise security architect.

Other subfunctional tasks may include physical subsystem and tape library design, fault and performance management, virtualization, SAN management, and enterprise planning/design/procurement.

Operations

Operational procedures will not vary greatly from current practices, but should be managed by dedicated storage administrators. Functionally, traditional systems management should remain within the systems management group, and the key is to begin separating the storage operations from systems, database, and application management. To what degree will certainly be dependent on each organization, but the three main functional disciplines are:

- **Performance management** This consists of performance modeling and management (including FC-based networks), asset management, chargeback, and billing by application of business unit.

- **Resource management** This involves resource allocation or provisioning—the actual assignment of storage (for example, LUNs, data zones, volumes) according to architectural/application requirements, as well as quota management, usage management (for instance, by application or individual user), metering, and capacity planning.
- **Availability management** This deals with backup/recovery operations. Backup/recovery should be included with storage management, not systems management. This includes backup/recovery scheduling, media management, and problem resolution. High-availability server clustering management (in conjunction with systems management), hierarchical storage management (where applicable), and replication (for example, remote replication for business continuity) are also of consideration.

Similar to balancing the handoffs with enterprise architects, infrastructure, operations, and lines of business, the data and media COE is no different. As can be seen, some of the functions and subfunctions transcend the policy, infrastructure, and operations categories. Communication among groups is a key element for a successful storage infrastructure that is adaptable to change and is measurable.

Actionable Items

ITOs should undertake the following initiatives:

- **Rationalizing storage hardware and storage software** This should encompass all aspects of storage across the server and application portfolio. Certainly, storage life cycles must be considered, yet a rationalization strategy (reducing variations—both vendor and configuration) will provide significant strategic value, even in tactical times.
- **Creating a storage infrastructure** ITOs should begin networking storage resources (SAN, NAS, and backup/recovery architectures), leverage tiered storage offerings (for example, internal storage, midrange, enterprise) and functional software (say, replication, server cluster integration, and backup/recovery), and look to common components (such as FC switches, host bus adapters, and so on) where possible. They should also seek new elements that adhere to intensive interoperability standards and procedures to ensure maximum configuration flexibility.
- **Optimizing storage operations** This includes rationalizing and consolidating management tools and personnel responsibilities; consistent storage management of multivendor environments is beginning to emerge and is a significant strategic directive.
- **Creating a data and media center of excellence** ITOs should employ a center of excellence using the guidelines outlined in the previous section.

Outsourcing

From 2003 through 2006, as ITOs adopt a storage service delivery model resembling that of traditional outsourcers, accurately measuring storage services and metrics will be a critical success factor in ITOs not only gaining credibility, but also in justifying (and making an informed decision about) whether a storage and storage management sourcing strategy should be employed. ITOs can use three categories when determining measurable metrics. The initial set comprises alternatives to the traditional ROI/total cost of ownership senior management mandate and is primarily an internal benchmark of (current and future) capabilities that quantify how the strategy is providing value. The second and third sets, technical and operational, are more strategic because some of these measurements are not possible with some storage vendors. However, ITOs should keep the measurements simple, limiting the metrics to those that are the most important to the business. Often, using network-based SLAs (for example, web hosting) as a framework and mapping them to storage services provides a good starting point. The three categories available to ITOs are:

- **Internal benchmarking** This includes storage managed per administrator, percentage of storage utilization, days storage in inventory, data life-cycle multiplier, data availability, and mean time to recovery.
- **Technical measures** These include storage availability, storage latency, data migration/exit clause, diagnostic, capacity utilization, performance, resource utilization, and mean time and maximum time to recover/resolve.
- **Operational measures** These include maximum time to notify, moves/adds/changes, project-specific SLAs, and vendor responsiveness.

Conclusions

ITOs are increasingly realizing the importance of storage, particularly among midsize companies that have historically not had requirements for enterprise storage capabilities. Through 2005/06, most organizations will have organized around, and created, a storage infrastructure and operations team (the data and media center of excellence). Consequently, it is imperative that ITOs begin to view storage in regards to the strategic importance it will have—even in difficult economic environments that often, and many times incorrectly, result in solely tactical decisions. ITOs should undertake a strategic evaluation examining how daily functions can be leveraged across (and automated by) multivendor capabilities, and how tasks currently consuming significant resources can be dramatically reduced (for example, recovery, provisioning, and procurement).

Moreover, as storage technologies continue to mature and more servers participate in networked storage (SAN and NAS), ITOs will be forced into measuring delivered capabilities as storage becomes a service that is delivered to the business in business

speak (such as performance, availability, flexibility to change), and associated cost tradeoffs based on their selection of tiered services is clearly understood.

Robert Spalding has been a respected authority on storage and storage management for the past 20 years. Having worked with Robert in the past, he was the first person to begin describing the need, and preparations necessary, for an adaptive storage infrastructure. In doing so in 1995, Bob began speaking about the need for a Chief Storage Officer.

Much of this thought leadership has manifested itself in Robert's published works, and specifically in this book, *Storage Networks: The Complete Reference*. As an industry analyst covering the entire storage and storage management market, it has been my experience that simple yet technical reference books are extremely useful. I'm sure storage professionals will benefit greatly by having this as a staple manual on their desk. Enjoy.

Sean Derrington is a leading authority on storage and storage management infrastructures. He specializes in storage subsystems, Fibre Channel, Storage Area Networks, Network Attached Storage, backup/recovery, and disaster recovery architectures, with a concentration on e-business, transactional, and analytic storage infrastructures. Sean joined META Group in July, 1995. He has a B.S. degree in Material Science and Engineering as well as a B.S. degree in Engineering and Public Policy, both from Carnegie Mellon University. Sean can be reached at sean.derrington@metagroup.com.

Acknowledgments

Writing a technical book is very much like building a house. It springs from inspiration, necessity, aesthetics, multiple talents, and hard work. Nothing would get published if the myriad of small tasks were not completed to meet deadlines from a project plan. My appreciation goes to all of those talented people who worked and contributed to this book.

If there is any dedication of this book it is to Donn Thornton for his inspiration. Donn was instrumental in teaching me the value of understanding, listening, and more importantly thinking, not only outside the box, but outside the universe. This book moved forward with the help of Kevin Ansell, whose experiences and discussions helped me through many difficult times. My thanks and appreciation goes out to Kevin for his sense of humor and encouragement.

Franny Kelly, acquisition editor at McGraw-Hill/Osborne, championed keeping this project alive and on schedule. My thanks to Franny for trusting me with the completion of this large project. Regarding the tremendous task of assembling all the writing into a real-live book, I'd like to convey my sincere thanks to my project editor, Jody McKenzie. Jody (along with her talented copyeditors) was critical to my technical dissertations becoming a coherent set of text and artwork and, subsequently, a completed manuscript.

I admit it; technically, I had help. Patty Then, storage consultant extraordinaire proved a much-needed sounding board in the review and validation of a wide set of technical

concepts and specifications that make up this book. Patty's knowledge of systems, storage infrastructures, and data centers was a key resource in illustrating the technical aspects of the SAN. She is one of those rare individuals versed in performance, capacity planning, and operational experiences that span both mainframe and open systems. I also need to thank Himanshu Dwivedi for his work in Chapter 25 on the increasing interests and concerns regarding security in storage infrastructures.

My thanks and appreciation also go to Ian Spalding, technical and creative writer, for his last-minute assistance in writing some of the key chapters from my encrypted notes. In addition, Ian was enlisted again to assist in proofing the book before final publication. We all wish you the best of success on your first novel.

Finally, all of this work would not have been possible without the support, assistance, inspiration, and help from my partner and wife, Artie. Her help in moving this book toward completion was more than anyone could ask. Thank you.

Robert Spalding
February, 2003