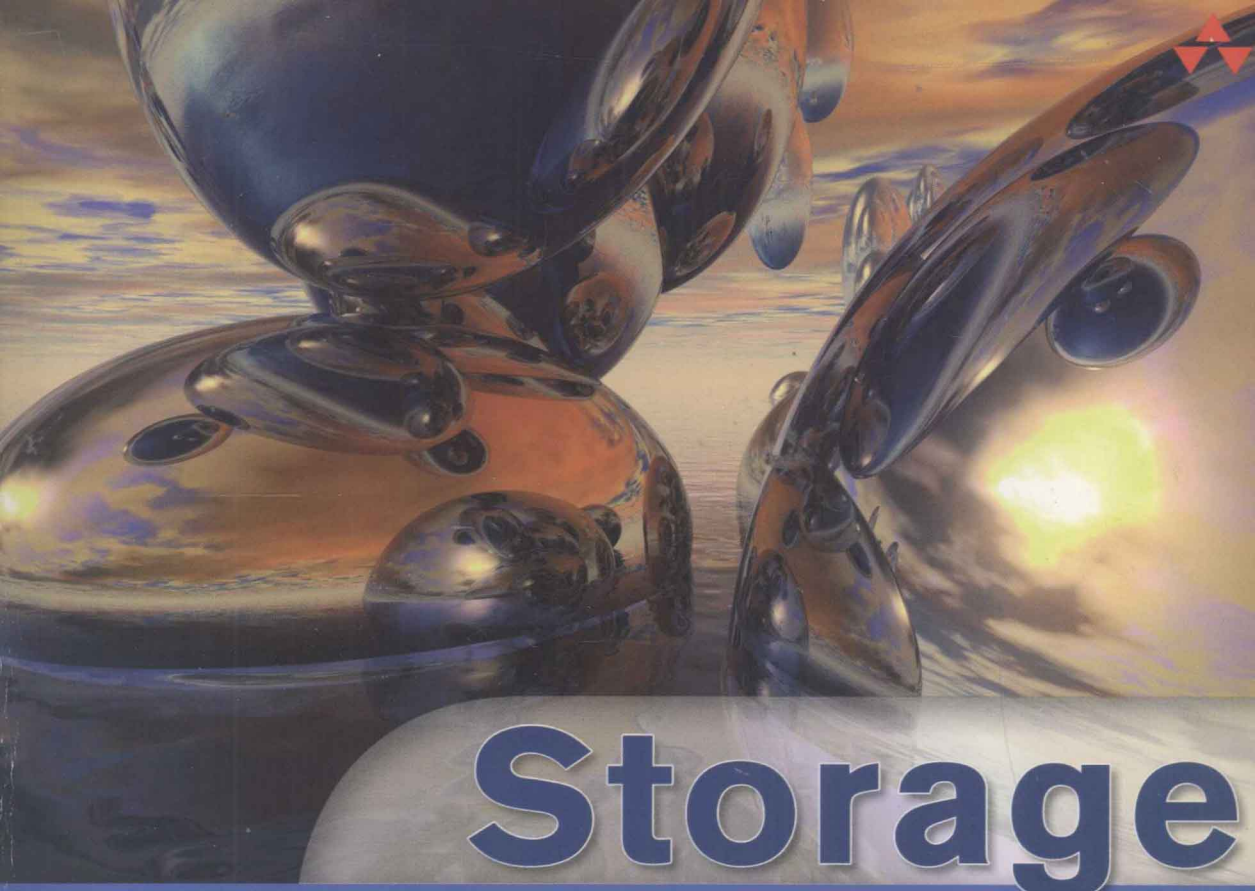

Storage Virtualization

Technologies for Simplifying Data Storage and Management

TOM CLARK

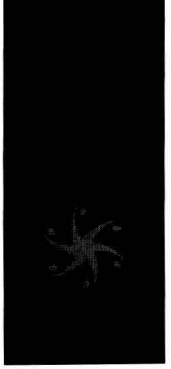

Storage Virtualization

Technologies for Simplifying Data Storage and Management

Tom Clark

Addison-Wesley

Upper Saddle River, NJ • Boston• Indianapolis • San Francisco
New York • Toronto • Montreal • London • Munich • Paris
Madrid • Capetown • Sydney • Tokyo • Singapore • Mexico City

Library of Congress Catalog Number:
2004116051

ISBN 0-32-126251-4
Text printed in the United States on recycled paper at R. R. Donnelley, Crawfordsville, Indiana.
First printing, March 2005

Storage Virtualization

This book is
dedicated to
the anonymous men and women
whose labor, from the
silicon foundries, factories and assembly lines
scattered around the globe,
brings advanced technologies
from the realm of ideas
to practical and productive
implementation

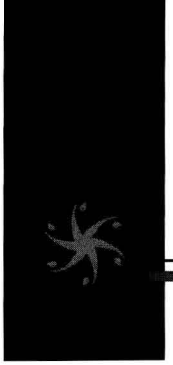

Preface

THE FOLLOWING WORK provides an overview of storage virtualization technology and its myriad manifestations. Like any emerging technical trend, the vague outlines of virtual storage concepts have only begun to sharpen as the technology has matured to productive applications in the real world. Today, there is a diversity of storage virtualization solutions, often tailored to meet specific storage needs. This book attempts to explain the background for that diversity, how different solutions function, and the essential value that is driving storage virtualization towards higher levels of utility.

Chapter 1 (Introduction), discusses the current state of storage virtualization in the market, and provides core concepts for understanding the hierarchy of virtualization-enabled storage operations.

The following three chapters provide foundation knowledge for understanding the relationship between what the user sees and what is done behind the scenes. Chapter 2 (Files and Records) begins with data in its more familiar format, as objects manipulated by applications for persistent storage. Chapter 3 (Data on Disk) provides the link between upper layer file/record structures and lower layer block data storage. Ultimately, all data must reside somewhere, and at some point the content of a file or record will be transformed into data blocks. Chapter 4 (The Storage Interconnect) reviews the connectivity required to link servers with their storage assets, whether by direct-attached SCSI, Fibre Channel or iSCSI.

The next five chapters discuss storage virtualization proper and the various means that have been engineered to support it. Chapter 5 (Abstracting Physical Storage) examines the process of aggregating multiple storage systems into a virtual storage pool. The secret recipe behind this is the mapping of logical block addresses presented by each system to virtual block addresses that are in turn presented to servers. This may be done in various

ways. Chapter 6 (Virtualization at the Host) discusses software virtualization that runs on individual servers. Chapter 7 (Virtualization at the Storage Target) discusses array-based virtualization techniques that are offered in some form by nearly all storage vendors. Chapter 8 (Fabric-based Virtualization) reviews the integration of virtualization technology with fabric switches and initiatives such as the Fabric Application Interface Standard that promise interoperable solutions. Chapter 9 (Virtualization Appliances) examines fabric-attached solutions and the innies/outties dispute between in-band and out-of-band methods. Collectively, these chapters address the "where it is done" category of the SNIA storage virtualization taxonomy.

Chapter 10 (Virtualization Services) discusses the practical application of virtual storage to real problems such as high availability and heterogeneous storage use. Storage virtualization in general is a foundation for higher level storage services such as hierarchical storage management.

Chapter 11 (Virtualized SAN File Systems) reviews another use of virtualization technology to streamline file system management and to enable distributed computing environments.

Chapter 12 (Virtualized Tape) provides an overview of the application of virtualization and system aggregation concepts to classic tape backup operations. New technologies such as RAIT (Redundant Array of Independent Tape devices) are breathing new life into an established fixture of data center environments.

Chapter 13 (Storage Automation and Virtualization) discusses the higher level services that are enabled by storage virtualization. Policy-based storage management, application-sensitive virtualization intelligence and the ability of applications themselves to leverage underlying virtualized services are new areas of development that hold great promise for simplifying storage operations.

In conclusion, Chapter 14 (The Storage Utility) examines the wide range of technical dependencies that have been setting the pace of development of storage virtualization technology. As is typical for a final chapter, Chapter 14 also provides wild speculations on the future of storage virtualization and the benefits it may yet provide as an enabling technology for a storage utility.

I have included a bibliography, which unfortunately seems to be an endangered species in technical works these days. Although there are few works on storage virtualization, there are many fine references on SAN technology that the serious reader can pursue. There is also a glossary of storage virtualization terminology as well as general SAN concepts.

The Appendix section at the end of the book includes additional reference material that hopefully will be of interest. Appendix A (Industry Resources) provides web links to industry and standards organizations. Appendix B (Vendor Resources) provides web links to storage virtualization and storage networking vendors, grouped by product type. Appendix C (Observations & Speculations) is an opportunity for industry observers, analysts, experts and customers to express their opinions on what this technology is and where it might be going.

Intended Audience

The following work should be useful for anyone who wants to understand the higher functions of storage networking. Storage managers, administrators, SAN architects, storage engineers, analysts, vendors, students and anyone involved in data storage technology should appreciate the new opportunities that storage virtualization provides. This book is therefore intended for a fairly diverse audience, from readers who already have experience with SANs to those who are just learning the benefits of shared storage solutions.

It is always difficult, however, to write a technical work with a specific reader in mind. Some readers will want more technical content; others less. Some will appreciate a broader overview, while others will want to get immediately to the point. This book attempts to provide both sufficient technical detail to be meaningful for a technical audience, and sufficient overview to provide an understanding of the subject by a less technical reader. For both types of readers, feel free to fast forward through sections that discuss concepts already well-understood, or simply of less interest.

Throughout the text I have attempted to avoid mention of specific products or vendors. This is done both to preserve objectivity in discussing technical matters as well as to extend the useful shelf life of the work despite the inevitable innovations that will be introduced. Hopefully, the concepts and relationships explained in the following chapters will provide a useful framework for understanding where we are in the space-time continuum of virtualization's maturation into more sophisticated products.

Acknowledgments

TECHNOLOGY IS A COLLECTIVE ENTERPRISE, both within companies and among the companies that comprise an industry. I have had the good fortune of meeting and working with hundreds of technologists, engineers, system architects and administrators in the storage networking industry over the past eight years. In that time, networked storage has evolved into a successful mainstream technology and transformed data storage processes. Although market competition sometimes throws the storage family into dysfunction, the industry as a whole continues to spawn new solutions and new companies while attracting new talent to further technical innovation. The embodiment of the shared storage family is the Storage Networking Industry Association (SNIA), which should be credited for its contributions to standards, interoperability, education and technology advancement, and for bringing order from the chaos of vendor conflict. In the area of storage virtualization specifically, the SNIA has helped to clarify concepts and relationships within a still emerging and often confusing subject area.

Publishing is also a very collective effort. From the initial idea for a book to its appearance on a shelf, a work passes through many hands. For a technical work in particular, peer review is required to ensure technical accuracy and objectivity of content. I would like to thank Steve Blumenau, Mark Carlson, Milan Merhar and David Thiel for subjecting themselves to the review of the manuscript and for their many useful suggestions and observations. I also am indebted to Mary Franz and Catherine Nolan, my editors, and Lori Lyons, Noreen Regina and many others at Addison Wesley for overseeing this project and accommodating my sporadic output due to my heavy travel schedule.

As always, my wife Lou bears the brunt of personal sacrifice that accompanies a writing project. Although she does not share my interest in

logical block address mapping algorithms, she understands that this is something I will probably outgrow at some point. In the meantime, she patiently waits for me to emerge from my study, back into multi-faceted reality to share a stroll through the garden or perhaps to help spread another ton of compose on it.

Due to my position at McDATA Corporation, I've traveled extensively over the past year, meeting customers and technologists throughout the US as well as Europe, Australia, New Zealand, India and Asia. My thanks to the many users of storage networking technology for their profound insight into the challenges of aligning technology to diverse application requirements and for affirming the global value of a technology that still has far more to offer in the future.

Tom Clark
Seattle, Washington

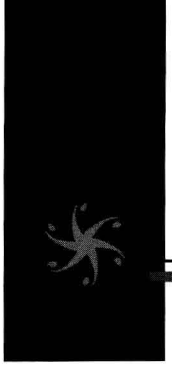

About the Author

TOM CLARK has held director positions at McDATA Corporation and other storage networking companies, conducts SAN seminars and tutorials worldwide, and serves as customer liaison. A noted storage industry author and advocate, he is a former board member of the Storage Networking Industry Association (SNIA) and has held chair positions for SNIA customer initiatives and the SNIA Interoperability Committee. His previous Addison-Wesley books include *Designing Storage Area Networks, Second Edition* (ISBN 0321136500) and *IP SANs: A Guide to iSCSI, iFCP, and FCIP Protocols for Storage Area Networks* (ISBN 0201752778). Tom lives in scenic Duvall, WA.

Also Available from Addison-Wesley and Tom Clark

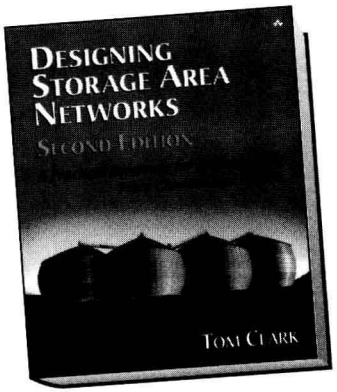

Designing Storage Area Networks, Second Edition

A Practical Reference for Implementing Fibre Channel and IP SANs

BY TOM CLARK

Written for network developers, IT consultants, administrators, and managers, this updated and greatly expanded edition of the best-selling *Designing Storage Area Networks* goes far beyond a straight description of technical specifications and standards. The text offers practical guidelines for using diverse SAN technologies to solve existing networking problems in large-scale corporate networks. With this book you will learn how the technologies work and how to organize their components into an effective, scalable design. In doing so, you will discover today's best methods for managing storage area networks, including new troubleshooting techniques.

ISBN 0-321-13650-0, © 2003, 592 pages

IP SANs

A Guide to iSCSI, iFCP, and FCIP Protocols for Storage Area Networks

BY TOM CLARK

IP storage and networking have traditionally resided in two distinct worlds. Networking professionals from an Internet Protocol (IP) internetworking background are usually not familiar with storage issues, and storage administrators may be unfamiliar with IP internetworking. With IP storage networking, network professionals dealing with storage area networks (SANs) now have an integrated option for improved data storage. IP SANs explains these new IP technologies that enable SANs to keep up with today's networking needs, detailing the various storage solutions that are created when both disciplines are combined.

ISBN 0-201-75277-8, © 2001, 320 pages

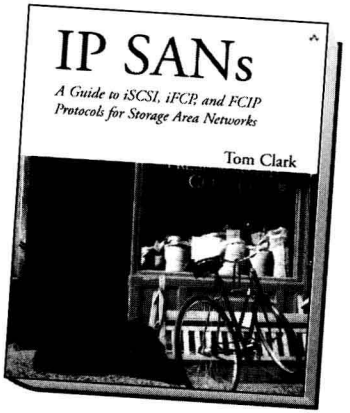

Visit us online for more books and more information, and to read sample chapters:
http://www.awprofessional.com

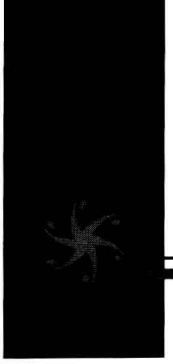

Contents