

Foreword to the Second Edition by Hermann Strass

A book on the subject of storage networks is especially important during these fast-moving times. The technology for storage networking is basically bringing with it new structures and procedures that will remain topical in the foreseeable future regardless of incremental differences and changes in products. This book is based on the experience of its authors in their day-to-day work with the material. It provides system administrators and system planners in particular with the tools they need for an optimal selection and cost-effective implementation of this complex technology, the use and operation of which currently seems indispensable in view of the ever-increasing storage quantities in companies. The technology of networked storage provides demonstrable and important cost savings. Growth therefore continues even in an unfavourable economic climate.

Storage quantities are growing because we are now working much more in colour, in three-dimension and digitally than was the case years ago. Furthermore, legal regulations that exist in the European Union and in other countries make the electronic/digital storage of all business data compulsory. The law no longer allows old business documents to be filed in printed form in archives. Data quantities continue to increase in good times as well as bad. Even lost contracts and the related data must be stored digitally. The legal regulations on their own are thus ensuring that a certain amount of growth in data is inevitable.

In the past, data was stored on disk and tape drives that were connected directly to a server. Storage was operated as a peripheral to the computer. Access rights, virus protection and other functions could thus be performed on the relevant computer (server). For reasons that are explained in detail in this book, this mode of operation is no longer

practical today. Storage has been detached from the servers and combined to form a separate storage network. This has resulted in a fundamentally different approach to dealing with storage. The new procedures required will continue to be developed into the near future. Data storage therefore has a value of its own. It is no longer a matter of attaching another disk drive to a server.

Today stored data and the information it contains are the crown jewels of a company. The computers (servers) needed for processing data can be purchased by the dozen or in larger quantities – individually as server blades or packed into cabinets – at any time, integrated into a LAN or a WAN or exchanged for defective units. However, if stored data is lost, restore of it is very expensive and time-consuming, assuming that all or some of it can even be recovered. As a rule, data must be available ‘around the clock’. Data networks must therefore be designed with redundancy and high availability.

These and related topics are covered in detail in this book. The approach is based upon the current state of technology only to a certain degree. What is more important is the description of the fundamental topics and how they relate to one another. This coverage goes beyond the scope of even lengthy magazine articles and will continue to be topical in the future. This is the only book available in the market today that covers this subject so comprehensively.

The requirements of storage networks are fundamentally different from those of the familiar local networks (LANs). Storage networks have therefore almost exclusively been using Fibre Channel technology, which was specially developed as a connection technology for company-critical applications. Storage networking is not a short-term trend and efforts are therefore currently underway to use other existing (for example, Ethernet-LAN-TCP/IP) network technologies as well as new ones that are coming on the market (for example InfiniBand and FCoE). Under certain circumstances these are totally sensible alternatives. This book highlights which selection criteria play a role here. It is usually not technical details or prejudices that are decisive but rather usage requirements, existing infrastructure and devices, along with a careful assessment of the future development in companies. The aim of this book is to provide valuable help in structural planning and the selection of devices and software.

The importance of networked storage technology has grown substantially since the first edition was printed. For the reasons mentioned in this book and due to regulatory requirements, even medium-sized companies need to manage large quantities of data and make them available for many years. This is why the sections on storage archiving have been considerably expanded in the new edition of this book. In a global economy business continuity is overly important for survival. This second edition devotes extensive coverage to this topic.

Overall this book is an excellent work. It explains the chosen subject comprehensively and in great detail, based on solid technical foundations. It is hoped that it will gain a wide circulation, particularly as it corrects a great many half-truths with its presentation of facts and addresses the usual prejudices.

Hermann Strass

Preface by the Authors

This Preface answers the following main questions:

- What does this book deal with?
- Who should read this book?
- How should this book be read?
- Who has written this book?

WHAT DOES THIS BOOK DEAL WITH?

The technology of storage networks fundamentally changes the architecture of IT systems. In conventional IT systems, storage devices are connected to servers by means of SCSI cables. The idea behind storage networks is that these SCSI cables are replaced by a network, which is installed in addition to the existing LAN. Server and storage devices can exchange data over this new network using the SCSI protocol. Storage networks have long been a known quantity in the world of mainframes. Fibre Channel, iSCSI, FCoE and Network Attached Storage (NAS) are now also taking storage networks into the field of Open Systems (Unix, Windows, OS/400, Novell Netware, MacOS).

Storage networks are a basic technology like databases and LANs. Storage was previously installed in the servers. Now most storage capacity is provided in external devices that are linked to servers over a storage network. As a result, anyone who is involved in the planning or operation of IT systems requires basic knowledge about the fundamentals and the use of storage networks. These networks are almost as widespread as SCSI, SAS and SATA but are more complex than LANs and TCP/IP.

The book is divided into two parts. Part I deals with fundamental technologies relating to storage networks. It guides the reader from the structure and operating method of

storage devices through I/O techniques and I/O protocols to the file systems and storage virtualisation.

The second part of this book presents applications that utilise the new functions of storage networks and intelligent disk subsystems. The emphasis here is on the shared use of resources that are available over a storage network, scalable and adaptable storage architectures, network backup and digital archiving. Another important focus of the book is business continuity with strategies for continuous and loss-free operation as protection against small failures and large catastrophes. Further focal points are the discussions on the management of storage networks and the management of removable media. Last but not least, the SNIA Shared Storage Model provides a reference model to describe storage networks.

At the end of the book we have added a glossary, an index and an annotated bibliography, which in addition to further literature also highlights numerous freely available sources on the Internet.

Section 1.4 sets out in detail the structure of the book and the relationships between the individual chapters. Figure 1.7 illustrates the structure of the book. At this point, it is worth casting a glance at this illustration. Note that the illustration also describes the subjects that we will not be covering.

Long before the second edition was printed, many readers of the first edition wanted to know what the differences are between the two editions. Here we want to express that our approach was successful, we aimed at introducing basic concepts rather than presenting actual products and overly technical details. The chapter on I/O techniques was the only one that required some updating on Fibre Channel and iSCSI. The key distinction of the second edition is the addition of two new chapters covering the topics of digital archiving and business continuity. We have also expanded the coverage on the copy services of intelligent disk subsystems.

WHO SHOULD READ THIS BOOK?

Our approach is, first, to explain the basic techniques behind storage networks and, secondly, to show how these new techniques help to overcome problems in current IT systems. The book is equally suitable for beginners with basic IT knowledge and for old hands. It is more an introduction to the basic concepts and techniques than a technical reference work. The target group thus includes:

- System administrators and system architects
- System consultants
- Decision makers
- Users
- Students

After reading the whole book you will be familiar with the following:

- The concepts of storage networks and their basic techniques
- Usage options for storage networks
- Proposed solutions for the support of business processes with the aid of storage networks
- The advantages of storage networks
- New possibilities opened up by storage networks.

HOW SHOULD THIS BOOK BE READ?

There are two options for reading this book. Those readers who are only interested in the concepts and usage options of storage networks should read Chapter 1 (Introduction) and Part II (Application and Management of Storage Networks); they can use Part I as a reference to look up any basic technical information they might require. Readers who are also interested in the technical background of storage networks should read the book through from the beginning.

WHO HAS WRITTEN THIS BOOK?

Ulf Troppens began work on this book in 2001. Rainer Erkens joined him soon after, providing his contributions on the topics of storage virtualisation, management of storage networks and NDMP for the first edition in 2002. In 2004 Wolfgang Müller-Friedt expanded the English translation – which was presented with the ‘Editor’s Choice Award 2005’ by Linux Journal – with his sound knowledge of magnetic tape, tape libraries and their management. Lastly, the second edition has been expanded considerably through contributions by Nils Hausteine (digital archiving) and Rainer Wolafka (business continuity).

All five authors have different roles at the Storage Competence Center of IBM in Mainz, Germany. Our responsibilities range from the development and testing of new software for storage networks to providing guidance to customers on the procurement of suitable products and the respective underlying concepts as well as on the installation and support of relevant hardware and software for customer environments. We advise customers on how storage networks can help to solve problems in their current IT systems. This experience has made us familiar with the types of questions customers have in respect of storage networks. Our involvement extends to customers with experience in storage networks as well as to those who are novices in this field. The positive feedback we have received from readers of the first edition show that our work has helped us to structure the content of this book and to choose topics in a way that are important to readers of books on storage networks.

Our intention has been to take off our ‘IBM hats’ and to write this book from an unbiased viewpoint. As employees of IBM in the area of storage technology, the experience and opinions that have been formed in our day-to-day work have of course had some influence on this book. In this connection, we have to be very familiar with our own company’s products as well as with those of our competitors and to position these products so that we inevitably have a view that goes beyond the IBM scope. In the end, this book is our personal work and has no connection with IBM apart from our employee relationship. Most importantly, this book does not represent any of the official opinions of IBM.

ACKNOWLEDGEMENTS FOR THE SECOND EDITION

We would like to give special thanks to our technical advisors on the second edition: Dirk Jahn (Archiving), Hans-Peter Kemptner (Business Continuity), Robert Haas (Limitations of RAID 5) and Hermann Strass for the Foreword. Other contributions were made by Jens-Peter Akelbein. We also appreciate the help we received on the publishing side from Rene Wiegand (copy-editing), Ulrich Kilian (LaTeX) and Rene Schoenfeld (editorial), all who helped to make our manuscript ready for printing.

With regard to the second English edition we would like to thank Birgit Gruber, Tiina Ruonamaa, Brett Wells, Liz Benson, Anna Smart, Sarah Tilley, Mary Lawrence and Sarah Hinton (all Wiley & Sons) as well as Deepthi Unni and her team at Laserwords. Last but not the least we thank Hedy Jourdan for the great translation of the new parts from German to English.

ACKNOWLEDGEMENTS FOR THE FIRST EDITION

We would also like to use this preface to thank some of the people who have made a significant contribution to the first edition of this book. From a chronological point of view, we should start by mentioning the editorial department of *iX* magazine and the copy-editing staff of dpunkt.verlag as they set the whole project in motion in March 2001 with the question ‘Could you see yourselves writing a book on the subject of storage in the network?’

Regarding content, our colleagues from the IBM Mainz storage community, especially the former SAN Lab and the current TotalStorage Interoperability Center (meanwhile renamed to Systems Lab Europe), deserve mention: Without the collaboration on storage hardware and software with customers and employees of partner companies, business partners and IBM, and without the associated knowledge exchange, we would lack the experience and knowledge that we have been able to put into this book. The list

of people in question is much too long for us to include it here. The cooperation of one of the authors with the students of the BAITI 2000 course of the Berufsakademie Mannheim (University of Applied Science Mannheim), from whom we have learnt that we have to explain subjects such as ‘RAID’, ‘disk subsystems’, ‘instant copy’, ‘remote mirroring’ and ‘file server’, was also valuable from a didactic point of view.

With regard to quality control, we thank our proofreaders Axel Köster, Bernd Blaudow, Birgit Bäuerlein, Frank Krämer, Gaetano Bisaz, Hermann Strass, Jürgen Deicke, Julia Neumann, Michael Lindner, Michael Riepe, Peter Münch, René Schönfeldt, Steffen Fischer, Susanne Nolte, Thorsten Schäfer, Uwe Harms and Willi Gardt, as well as our helpers at dpunkt.verlag, whose names we do not know.

We should emphasise in particular the many constructive suggestions for improvement by Susanne Nolte, who also contributed a few paragraphs on ‘DAFS’, and the numerous comments from our colleagues Axel Köster and Jürgen Deicke and our manuscript reader René Schönfeldt. In this connection, the efforts of Jürgen Deicke and Tom Clark should also be mentioned regarding the ‘SNIA Recommended Reading’ logo, which is printed on the front cover of the book.

With regard to the first English edition of this book we have to thank even more people: First of all, we would like to thank René Schönfeldt from dpunkt.verlag for convincing Birgit Gruber from Wiley & Sons to invest in the translation. We greatly appreciate Birgit Gruber for taking a risk on the translation project and having so much patience with all our editorial changes. Rachel Waddington did an outstanding job of translating the text and all the figures from German into English. Last but not least, we would like to thank Daniel Gill for leading the production process, including copy-editing and typesetting, and we would like to thank the team at Laserwords for typesetting the whole book.

Closing comments

Finally, the support of our parents, parents-in-law and partners deserves mention. I, Nils Haustein, would like to thank my dear wife Susann who gave me a lot of ‘computer time’ and the opportunity to make a contribution to this book. I, Rainer Wolafka, would like to thank my dear wife Tineke for her support and her constant encouragement and motivation to work on this book and to my son Daniel for understanding why I did not always have the time he deserved during this time. I, Wolfgang Müller-Friedt, would like to thank my dear wife Christel for her patience, her emotional support and for many more reasons than there is room to list in these notes. I, Ulf Troppens, at this point would like to thank my dear wife Silke for her support and for taking many household and family duties off my hands and thus giving me the time I needed to write this book. And I, Rainer Erkens, would like to thank my dear partner Christina, who never lost sight of worldly things and thus enabled me to travel untroubled through the world of storage

networks, for her support. We are pleased that we again have more time for children, our families and friends. May we have many more happy and healthy years together.

Mainz, April 2009

Ulf Troppens
Rainer Erkens
Wolfgang Müller-Friedt
Nils Haustein
Rainer Wolafka