Storage Networks Explained Basics and Application of Fibre Channel SAN, NAS, iSCSI, InfiniBand and FCoE, Second Edition

Ulf Troppens, Wolfgang Müller-Friedt, Rainer Wolafka

IBM Storage Software Development, Mainz, Germany

Rainer Erkens, Nils Haustein

IBM Advanced Technical Sales Europe, Mainz, Germany

Translated by Rachel Waddington, Member of the Institute of Translating and Interpreting, UK

New material for this edition translated from the original German version into English by Hedy Jourdan



Storage Networks Explained

Storage Networks Explained Basics and Application of Fibre Channel SAN, NAS, iSCSI, InfiniBand and FCoE, Second Edition

Ulf Troppens, Wolfgang Müller-Friedt, Rainer Wolafka

IBM Storage Software Development, Mainz, Germany

Rainer Erkens, Nils Haustein

IBM Advanced Technical Sales Europe, Mainz, Germany

Translated by Rachel Waddington, Member of the Institute of Translating and Interpreting, UK

New material for this edition translated from the original German version into English by Hedy Jourdan



First published under the title Speichernetze, Grundlagen und Einsatz von Fibre Channel SAN, NAS, iSCSl und InfiniBand. ISBN: 3-89864-135-X by dpunkt.verlag GmbH Copyright © 2003 by dpunkt.verlag GmbH, Heidelberg, Germany

1st edition of the English translation first published 2004. Translation Copyright © 2004, John Wiley & Sons Ltd.

This edition first published 2009 Translation Copyright © 2009, John Wiley & Sons Ltd.

Registered office

John Wiley & Sons, Ltd, The Atrium, Southern Gate, Chichester, West Sussex PO19 8SQ, United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. 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 or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

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

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data:

Storage networks explained : basics and application of Fibre Channel SAN, NAS, iSCSI, InfiniBand, and

FCoE / Ulf Troppens ... [et al.]. – 2nd ed. p. cm.
Rev. ed. of Storage networks explained / Ulf Troppens. c2004. ISBN 978-0-470-74143-6 (cloth)
1. Storage area networks (Computer networks) 2. Information storage and retrieval systems. I. Troppens, Ulf. II. Troppens, Ulf. Storage networks explained. TK5105.86.T78 2009

004.6-dc22

2009014224

A catalogue record for this book is available from the British Library.

ISBN 978-0-470-74143-6

Typeset in 10/12pt Times by Laserwords Private Limited, Chennai, India. Printed in Singapore by Markono Print Media Pte. *For Silke, Hannah, Nina, and Julia You keep showing me what really matters in life.*

> For Christina, Marie and Tom For your love and support.

For Christel Only your patience and your understanding have made my contribution to this book possible.

> For Susann In Love.

For Tineke, Daniel and Marina For the love, motivation and reassurance you have always given me.

Contents

A	bout t	he Aut	hors	xix
Fo	orewo	rd to th	e Second Edition by Hermann Strass	xxi
Pı	eface	by the	Authors	xxiii
Li	List of Figures and Tables			
1	Intr	oductio	n	1
	1.1	Server	-Centric IT Architecture and its Limitations	1
	1.2	Storag	e-Centric IT Architecture and its Advantages	3
	1.3	Case S	Study: Replacing a Server with Storage Networks	5
	1.4	The St	tructure of the Book	7
PA	ART I	Tech	nologies for Storage Networks	13
2	Inte	lligent l	Disk Subsystems	15
	2.1	Archite	ecture of Intelligent Disk Subsystems	16
	2.2	Hard I	Disks and Internal I/O Channels	18
	2.3	JBOD	Just a Bunch of Disks	21
	2.4	Storag	e Virtualisation Using RAID	22
	2.5	Differe	ent RAID Levels in Detail	24
		2.5.1	RAID 0: block-by-block striping	25
		2.5.2	RAID 1: block-by-block mirroring	26
		2.5.3	RAID 0+1/RAID 10: striping and mirroring combined	26
		2.5.4	RAID 4 and RAID 5: parity instead of mirroring	31

		2.5.5	RAID 6: double parity	35
		2.5.6	RAID 2 and RAID 3	37
		2.5.7	A comparison of the RAID levels	38
	2.6	Cachin	ng: Acceleration of Hard Disk Access	40
		2.6.1	Cache on the hard disk	40
		2.6.2	Write cache in the disk subsystem controller	41
		2.6.3	Read cache in the disk subsystem controller	41
	2.7	Intellig	gent Disk Subsystems	42
		2.7.1	Instant copies	42
		2.7.2	Remote mirroring	45
		2.7.3	Consistency groups	49
		2.7.4	LUN masking	51
	2.8	Availa	bility of Disk Subsystems	55
	2.9	Summ	ary	56
3	1/0	Technic		50
5	3 1	The Pl	μ σ	57 60
	3.1	SCSI	hysical 1/0 1 and from the effortion the Storage System	62
	5.2	321	SCSI basics	62
		322	SCSI and storage networks	65
	33	The Fi	ibre Channel Protocol Stack	66
	5.5	331	Links ports and topologies	69
		332	EC.0: cables plugs and signal encoding	71
		333	FC-1: 8b/10b encoding ordered sets and link control protocol	73
		334	FC_{-2} : data transfer	75
		335	FC-3: common services	82
		336	Link services: login and addressing	82
		337	Fabric services: name server and co	85
		338	FC-4 and ULPs: application protocols	86
	34	Fibre (Channel SAN	88
	011	3.4.1	Point-to-point topology	89
		3.4.2	Fabric topology	89
		343	Arbitrated loop topology	93
		3.4.4	Hardware components for Fibre Channel SAN	96
		3.4.5	InterSANs	100
		3.4.6	Interoperability of Fibre Channel SAN	103
	3.5	IP Sto	rage	105
		3.5.1	IP storage standards: iSCSI, iFCP, mFCP, FCIP and iSNS	105
		3.5.2	TCP/IP and Ethernet as an I/O technology	109
		3.5.3	Migration from Fibre Channel to IP storage	116
	3.6	Infinib	and-based Storage Networks	117
		3.6.1	InfiniBand	118
		3.6.2	Virtual Interface Architecture (VIA)	120

	3.7	3.6.3 SCSI via InfiniBand an Fibre Channel over Ethernet (F	nd RDMA CoE)	122 124
		3.7.1 I/O Consolidation base	ed on Ethernet	125
		3.7.2 FCoE Details		126
		3.7.3 Case studies		131
		3.7.4 Data Center Bridging ((DCB)	132
		3.7.5 Outlook		135
	3.8	Summary		135
4	File	Systems and Network Attache	d Storage (NAS)	137
	4.1	Local File Systems		137
		4.1.1 File systems and datab	ases	138
		4.1.2 Journaling		139
		4.1.3 Snapshots		139
		4.1.4 Volume manager		140
	4.2	Network File Systems and File	Servers	140
		4.2.1 Basic principle		141
		4.2.2 Network Attached Stor	age (NAS)	143
		4.2.3 Performance bottlenecl	s in file servers	145
		4.2.4 Acceleration of networ	k file systems	146
		4.2.5 Case study: The Direct	Access File System (DAFS)	147
	4.3	Shared Disk File Systems		150
		4.3.1 Case study: The Gener	al Parallel File System (GPFS)	152
	4.4	Comparison: Fibre Channel SA	IN, FCOE SAN, ISCSI SAN and NAS	156
	4.5	Summary		159
5	Stor	age Virtualisation		161
	5.1	Once Again: Virtualisation in t	he I/O Path	163
	5.2	Limitations and Requirements		167
		5.2.1 Architecture-related lin	nitations of non-virtualised storage	
		networks		167
		5.2.2 Implementation-related	limitations of storage networks	168
		5.2.3 Requirements of the da		1/1
	5.2	5.2.4 Proposed solution: stor	age virtualisation	172
	5.3 5.4	Definition of Storage Virtualisa	tion	174
	3.4	5.4.1 Dealisation of the virtu	validation antity	174
		5.4.2 Penlacement of storage	ansanon enury	1/4
		5.4.3 Efficient use of recourt	t ut victs	175
		5.4.4 Efficient use of resource	the by data migration	170
		5.4.5 Performance increase	to by data migration	170
		5.4.6 Availability due to the	introduction of redundancy	170
		J.T.O Availability due to the	introduction of requilibrilly	1//

ix

	5.4.7	Backup and archiving	177
	5.4.8	Data sharing	177
	5.4.9	Privacy protection	178
5.5	Storage	Virtualisation on Block or File Level	178
5.6	Storage	Virtualisation on Various Levels of the Storage Network	180
	5.6.1	Storage virtualisation in the server	181
	5.6.2	Storage virtualisation in storage devices	182
	5.6.3	Storage virtualisation in the network	183
5.7	Symme	tric and Asymmetric Storage Virtualisation in the Network	184
	5.7.1	Symmetric storage virtualisation	184
	5.7.2	Asymmetric storage virtualisation	188
5.8	Summa	ry	191

PART II Application and Management of Storage Networks 193

6	Арр	lication	of Storage Networks	195
	6.1	Definit	tion of the Term 'Storage Network'	195
		6.1.1	Layering of the transmission techniques and protocols	196
		6.1.2	Networks in the I/O path	197
		6.1.3	Data networks, voice networks and storage networks	198
	6.2	Storag	e Sharing	199
		6.2.1	Disk storage pooling	199
		6.2.2	Dynamic tape library sharing	201
		6.2.3	Data sharing	203
	6.3	bility of Data	206	
		6.3.1	Failure of an I/O bus	206
		6.3.2	Failure of a server	210
		6.3.3	Failure of a disk subsystem	212
		6.3.4	Failure of virtualisation in the storage network	214
		6.3.5	Failure of a data centre based upon the case study 'protection of	
			an important database'	215
	6.4	Adapta	ability and Scalability of IT Systems	219
		6.4.1	Clustering for load distribution	219
		6.4.2	Web architecture	226
		6.4.3	Web applications based upon the case study 'travel portal'	230
	6.5	Summ	ary	234
7	Nety	work Ba	ackun	237
-	7.1	Gener	al Conditions for Backup	238
	7.2	Netwo	rk Backup Services	238
				-00

7.3 Components of Backup Servers 241

8

	721	Tab askadular	241
	7.3.1	Job Scheduler	241
	7.3.2	Ellor handler Metadeta detabase	241
	7.3.3	Media managar	241
74	7.3.4 Doolar		242
7.4	DacKup	nance Gaine as a Result of Network Rackup	244
7.5	Derfor	nance Oanis as a Result of Network Backup	243
7.0	761	Application specific performance bottlenecks	240
	7.0.1	Parformance bottlanacks due to server centric IT architecture	240
77	7.0.2 Limite	d Opportunities for Increasing Performance	247
1.1	771	Separate I AN for network backup	240
	7.7.1	Multiple backup servers	240
	7.7.2	Reckup server and application server on the same physical	249
	1.1.5	computer	251
78	Next C	eenipuer	251
7.0	7 8 1	Server-free backup	252
	782	LAN-free backup	252
	783	LAN-free backup with shared disk file systems	254
	784	Backup using instant copies	256
	7.8.5	Data protection using remote mirroring	259
	7.8.6	Tape library sharing	259
7.9	Backur	of File Systems	261
1.5	7.9.1	Backup of file servers	261
	7.9.2	Backup of file systems	262
	7.9.3	Backup of NAS servers	263
	7.9.4	The Network Data Management Protocol (NDMP)	265
7.10	Backur	o of Databases	272
	7.10.1	Functioning of database systems	272
	7.10.2	Classical backup of databases	274
	7.10.3	Next generation backup of databases	276
7.11	Organi	sational Aspects of Backup	277
7.12	Summa	ary	278
Are	hiving		281
81	Termin	alagy	281
0.1	811	Differentiating between information and data	201
	812	Archiving	282
	813	Digital archiving	202
	0.1.5		202

	8.1.4	Reference architecture for digital archive systems	283
	8.1.5	Differentiating between archiving and backup	285
	8.1.6	Differentiating between archiving and ILM	288
8.2	Motiva	tion, Conditions and Requirements	291
	8.2.1	Reasons for archiving	291
	8.2.2	Legal requirements	291
	8.2.3	Technical progress	293
	8.2.4	Requirement for stability	294
	8.2.5	Risks from the environment and from society	295
	8.2.6	Requirement for adaptability and scalability	295
	8.2.7	Operational requirements	296
	8.2.8	Cost-related requirements	297
	8.2.9	Conclusion: Archive systems as a strategic investment	297
8.3	Implen	nentation Considerations	298
	8.3.1	WORM storage technologies	298
	8.3.2	Data security	302
	8.3.3	Data integrity	302
	8.3.4	Proof of regulatory compliance	305
	8.3.5	Deletion of data	305
	8.3.6	Continuous operation	307
	8.3.7	Loss-free operation	307
	8.3.8	Data management: storage hierarchy and migration	309
	8.3.9	Component-neutral archiving	310
	8.3.10	Selection of components and vendors	311
8.4	Interfa	ces in Archive Systems	311
	8.4.1	Interface between application and DMS	313
	8.4.2	Java Content Repository (JCR)	314
	8.4.3	Interface between DMS and archive storage	315
	8.4.4	eXtensible Access Method (XAM)	317
	8.4.5	Management interfaces	318
	8.4.6	Interface between DMS systems	318
	8.4.7	Standardised interfaces for archive systems	319
8.5	Archiv	e Solutions	320
	8.5.1	Archiving of emails	321
	8.5.2	Archiving of files	325
	8.5.3	Archiving of ERP systems	332
	8.5.4	Archiving in hospitals	335
	8.5.5	Central archives	338
8.6	Operat	ional and Organisational Aspects	342
8.7	Summa	ary and Outlook	343
Busi	iness Co	ontinuity	345

20 64101		• ••
9.1	General Conditions	346

9

	9.1.1	Terminology	346
	9.1.2	Target audience	347
	9.1.3	Classification of risks	347
	9.1.4	Classification of outages	348
	9.1.5	IT failures in the context of business processes	349
	9.1.6	Resumption of business processes	349
	9.1.7	Business continuity for the web architecture	351
	9.1.8	Cost optimisation for business continuity	351
	9.1.9	Risk analysis and risk management	353
	9.1.10	Creation of a business continuity plan	355
9.2	Strateg	ies of Business Continuity	356
	9.2.1	High availability	357
	9.2.2	Disaster recovery	358
	9.2.3	Continuous business operation	358
9.3	Parame	eters of Business Continuity	359
	9.3.1	Availability	359
	9.3.2	Characterisation of availability (MTBF, MTTR and MTTF)	360
	9.3.3	Calculation of overall availability	361
	9.3.4	Characterisation of failures (RTO and RPO)	362
	9.3.5	Network Recovery Objective (NRO)	364
9.4	Quality	of Service for Business Continuity	364
	9.4.1	Service Level Agreements (SLAs)	364
	9.4.2	High availability versus disaster recovery	365
	9.4.3	The seven-tier model	366
	9.4.4	Tier 0: no data backup	368
	9.4.5	Tier 1: data backup without a backup data centre	368
	9.4.6	Tier 2: data backup with backup data centre	368
	9.4.7	Tier 3: electronic vaulting	369
	9.4.8	Tier 4: instant copies	369
	9.4.9	Tier 5: software mirroring	369
	9.4.10	Tier 6: disk subsystem-based mirroring	369
	9.4.11	Tier 7: fully automated solutions	370
9.5	Busine	ss Continuity Solutions	370
	9.5.1	Basic techniques	370
	9.5.2	Solution segments of the seven-tier model	371
	9.5.3	Backup and restore	372
	9.5.4	Rapid data recovery using copies	372
	9.5.5	Rapid data recovery using mirrors	374
	9.5.6	Continuous availability	377
9.6	Switch	of Operational Location	383
9.7	Organi	sational Aspects	383
9.8	Summa	ary	384

10	Man	agemer	nt of Storage Networks	387
	10.1	Requir	ements	387
		10.1.1	User-related requirements	388
		10.1.2	Component-related requirements	388
		10.1.3	Architectural requirements	389
		10.1.4	One central management system	389
		10.1.5	Five basic services	390
	10.2	Charac	terisation of Management Interfaces	391
		10.2.1	In-band interfaces	391
		10.2.2	Out-band interfaces	392
		10.2.3	Standardised interfaces	393
		10.2.4	Proprietary interfaces	393
		10.2.5	Conclusion	394
	10.3	In-band	d Management	394
		10.3.1	In-band management in Fibre Channel SAN	396
	10.4	Out-ba	nd Management	398
		10.4.1	The Simple Network Management Protocol (SNMP)	399
		10.4.2	CIM and WBEM	403
		10.4.3	Storage Management Initiative Specification (SMI-S)	409
		10.4.4	CMIP and DMI	411
	10.5	Operat	ional Aspects of the Management of Storage Networks	413
	10.6	Summa	ary	414
11	Rem	ovable	Media Management	417
	11.1	The Si	gnificance of Removable Media	417
	11.2	Remov	vable Media	418
		11.2.1	Tapes	419
		11.2.2	CD. DVD and magneto-optical media	419
		11.2.3	Management features of removable media	419
	11.3	Librari	es and Drives	421
		11.3.1	Libraries	421
		11.3.2	Drives	421
		11.3.3	Media changers	422
	11.4	Proble	ms and Requirements in Respect of Removable Media	
		Manag	ement	425
		11.4.1	Efficient use of the available resources	427
		11.4.2	Access control	427
		11.4.3	Access synchronisation	430
		11.4.4	Access prioritisation and mount request queuing	430
		11.4.5	Grouping, pooling	431
		11.4.6	Media tracking and vaulting	433
		11.4.7	Cartridge life cycle management	435
		11.4.8	Monitoring	437

		11.4.9	Reporting	437
	11.5	The IEI	EE 1244 Standard for Removable Media Management	438
		11.5.1	Media management system architecture	439
		11.5.2	Media manager and MMP	440
		11.5.3	Library manager and drive manager	445
	11.6	Summa	ry	446
12	The	SNIA S	hared Storage Model	449
	12.1	The Mo	odel	450
		12.1.1	The functional approach	450
		12.1.2	Graphical representations	451
		12.1.3	An elementary overview	451
		12.1.4	The components	452
		12.1.5	The layers	454
		12.1.6	The file/record layer	454
		12.1.7	The block layer	456
		12.1.8	Combination of the block and file/record layers	458
		12.1.9	Access paths	459
		12.1.10	Caching	460
		12.1.11	Access control	460
		12.1.12	Clustering	462
		12.1.13	Storage, data and information	463
		12.1.14	Resource and data sharing	464
		12.1.15	The service subsystem	464
	12.2	Exampl	es of Disk-Based Storage Architectures	466
		12.2.1	Direct attached block storage	466
		12.2.2	Storage network attached block storage	467
		12.2.3	Block storage aggregation in a storage device: SAN appliance	468
		12.2.4	Network attached block storage with metadata server: asymmetric	
			block services	469
		12.2.5	Multi-site block storage	470
		12.2.6	File server	471
		12.2.7	File server controller: NAS heads	472
		12.2.8	Asymmetric file services: NAS/file server metadata manager	474
	10.0	12.2.9	Object-based storage device (OSD)	475
	12.3	Extensi	on of the SNIA Shared Storage Model to Tape Functions	4/5
		12.3.1	Logical and physical structure of tapes	4//
		12.3.2	Differences between disk and tape	4/8
	12.4	12.3.3 Exempt	Extension of the model	4/8
	12.4		Es of Tape-Based Backup Techniques and Architectures	480
		12.4.1	File system volume healtun	480
		12.4.2	Volume backup	401
		12.4.3	volume backup	482

	12.4.4	File backup to virtual tape	482
	12.4.5	Direct attached tape	484
	12.4.6	LAN attached tape	484
	12.4.7	Shared tape drive	485
	12.4.8	Partitioned tape library	485
	12.4.9	Virtual tape controller	486
	12.4.10	Virtual tape controller with disk cache	488
	12.4.11	Data mover for tape	489
	12.4.12	File server with tape drive	490
	12.4.13	File server with external tape	491
	12.4.14	File server with data mover	492
12.5	Summa	ıry	493
13 Fina	al Note		495
Glossar	y		497
Annota	ted Bibli	iography	527
Append	lix A: Pı	coof of Calculation of the Parity Block of RAID 4 and 5	535
Append	lix B: Cl	hecklist for the Management of Storage Networks	537
B.1	Applica	ations	538
	B.1.1	Monitoring	538
	B.1.2	Availability	538
	B.1.3	Performance	538
	B.1.4	Scalability	538
	B.1.5	Efficient use	538
B.2	Data		539
	B.2.1	Availability	539
	B.2.2	Performance	539
	B.2.3	Data protection	539
	B.2.4	Archiving	539
	B.2.5	Migration	539
	B.2.0	Data sharing	540
D 2	В.2./ Веления	Security/access control	540
B.3	Resour	Ces	540
	B.3.1 D 2 2	Monitoring	540
	D.J.2 D 2 2	Configuration	540
	D.J.J B 2 /	Connguiation Resource use	540 541
	D.J.4 B 2 5	Capacity	541 541
	D.J.J D 2 6	Capacity Efficient resource utilisation	541
	D.3.0	Encient resource utilisation	341

Index			543
	B.4.4	Performance	542
B.4	B.4.3	Availability	542
	B.4.2	Monitoring	542
	B.4.1	Topology	542
	Network		542
	B.3.9	Security	542
	B.3.8	Resource migration	541
	B.3.7	Availability	541

About the Authors



The authors are employed at IBM's storage competence center in Mainz, Germany. They work at the interface between technology and customers. Their duties cover a wide field of responsibilities. They develop and test new software for storage networks. They present the latest hardware and software products in the field of storage networks to customers and

explain their underlying concepts. Last but not least they deploy and support respective hardware and software in customer environments.

Ulf Troppens (centre) studied Computer Science at the University of Karlsruhe. Since 1989 he has been primarily involved in the development and administration of Unix systems, storage systems, data and storage networks and distributed applications.

Rainer Erkens (left) studied Mathematics at the University of Mainz. His experience in the management of computers and distributed applications goes back to 1992. Since 2005 he is a technical support manager in IBM's European Storage Competence Center.

Wolfgang Müller-Friedt (right) studied Computer Science at the FH Darmstadt. He is a software architect focussing on the software development of management applications for storage networks which support open standards such as SMI-S and IEEE 1244.

Nils Haustein (left front) studied Electrical Engineering at the TU Chemnitz. For several years he is with IBM's advanced technical sales support in Europe where he is focussing on digital archiving.

Rainer Wolafka (right front) studied Electrical Engineering at the FH Frankfurt and Software Engineering at the Santa Clara University. Since 1997 he is working in the field of storage networks and the software development of management applications for storage networks.

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 Haustein (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