



Cloud and Virtual Data Storage Networking

Your journey to efficient and effective information services

Greg Schulz



CRC Press
Taylor & Francis Group

AN AUERBACH BOOK

Cloud and Virtual Data Storage Networking

Your journey to efficient and
effective information services

This page intentionally left blank

Cloud and Virtual Data Storage Networking

Your journey to efficient and
effective information services

Greg Schulz



CRC Press

Taylor & Francis Group

Boca Raton London New York

CRC Press is an imprint of the
Taylor & Francis Group, an **informa** business
AN AUERBACH BOOK

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2012 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20110726

International Standard Book Number-13: 978-1-4398-5174-6 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (<http://www.copyright.com/>) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
<http://www.taylorandfrancis.com>

and the CRC Press Web site at
<http://www.crcpress.com>

Contents

Preface	xv
Acknowledgments	xix
About the Author	xxi
Chapter 1 Industry Trends and Perspectives: From Issues and Challenges to Opportunities	1
In This Chapter	1
1.1 Getting Started	2
1.2 The Importance of Data and Storage	2
1.2.1 The Business End of IT Data Storage Impact	3
1.2.2 Addressing Business and IT Issues	5
1.2.3 What Is Driving Data Growth and Information Reliance	5
1.3 Business Issues and IT Challenges	6
1.4 Business and IT Opportunities	7
1.4.1 Traditional Information Services Delivery/Model	7
1.4.2 Information Factories	8
1.5 Opportunity for Cloud, Virtualization, and Data Storage Networking	11
1.5.1 IT Clouds and Virtualization: Not If, Rather When, Where, Why, and How	11
1.5.2 Private Cloud: Coexistence vs. Competing with Legacy IT	14
1.6 Common Cloud, Virtualization, and Storage Networking Questions	14

1.7	Cloud, Virtualization, and Storage Networking: Bringing It Together (for Now)	17
1.8	Chapter Summary	18
Chapter 2	Cloud, Virtualization, and Data Storage Networking Fundamentals	21
	In This Chapter	21
2.1	Getting Started	22
2.2	Server and Storage I/O Fundamentals	22
2.2.1	Server and I/O Architectures	24
2.2.2	Storage Hierarchy	26
2.2.3	From Bits to Bytes	27
2.2.4	Disk Storage Fundamentals	29
2.2.5	Initiators and Targets	29
2.2.6	How Data Is Written to and Read from a Storage Device	30
2.2.7	Storage Sharing vs. Data Sharing	31
2.2.8	Different Types of Storage: Not All Data Storage Is the Same	32
2.3	I/O Connectivity and Networking Fundamentals	38
2.4	IT Clouds	40
2.5	Virtualization: Servers, Storage, and Networking	42
2.6	Virtualization and Storage Services	43
2.7	Data and Storage Access	44
2.7.1	Direct Attached Storage (DAS)	45
2.7.2	Networked Storage: Network Attached Storage (NAS)	45
2.7.3	Networked Storage: Storage Area Network (SAN)	45
2.7.4	Networked Storage: Public and Private Clouds	47
2.8	Common Questions	47
2.9	Chapter Summary	48
Chapter 3	Infrastructure Resource Management	49
	In This Chapter	49
3.1	Managing Data Infrastructures for Cloud and Virtual Environments	49
3.2	Introduction to Infrastructure Resource Management	51
3.3	Understanding IT Resources	54
3.4	Managing IT Resources	58

3.5	Service Offerings, Categories, and Technology Alignment	59
3.6	Gaining Situational Awareness and Control	64
3.7	From SRM to E2E SRA	66
3.8	Search and eDiscovery	67
3.9	Performance and Capacity Planning	68
3.10	Data Movement and Migration	70
3.11	Chapter Summary	71
Chapter 4	Data and Storage Networking Security	73
	In This Chapter	73
4.1	Being Secure Without Being Scared	73
4.2	Eliminating Blind Spots, Gaps in Coverage, or “Dark Territories”	74
4.3	Security Threat Risks and Challenges	76
4.4	Taking Action to Secure Your Resources	78
4.4.1	Physical Security	79
4.4.2	Logical Security	80
4.4.3	Multitenancy	80
4.4.4	Deciphering Encryption	81
4.5	Securing Networks	82
4.6	Securing Storage	84
4.6.1	Removable Media Security	85
4.7	Virtual Servers, Physical Servers, and Desktops	85
4.8	Securing Clouds	86
4.9	Disposing of Digital Assets and Technology	87
4.10	Security Checklist	88
4.11	Common Security-Related Questions	88
4.12	Chapter Summary	89
Chapter 5	Data Protection: Backup/Restore and Business Continuance/Disaster Recovery	91
	In This Chapter	91
5.1	Getting Started	92
5.2	Data Protection Challenges and Opportunities	92
5.3	Protect, Preserve, and Serve Information Services	93
5.3.1	Basic Information Reliability–Availability– Serviceability (RAS)	96
5.3.2	High Availability and Business Continuance	97
5.3.3	Disaster Recovery	98

5.3.4	Data Protection vs. Preservation (Backup vs. Archive)	99
5.4	SLO and SLAs: How Much Availability Do You Need vs. Want	100
5.4.1	RTO and RPO: Balancing Data Availability vs. Time and Budgets	100
5.4.2	Reconciling and Assessing RTO and RPO Requirements	101
5.4.3	Tiered Data Protection	103
5.5	Common-Sense Data Protection	103
5.6	Virtual, Physical, and Cloud Data Protection	104
5.6.1	Tools and Technologies	105
5.6.2	Virtual and Physical Machine Movement	109
5.6.3	Enabling High Availability	110
5.6.4	Snapshots and Continuous Data Protection	112
5.6.5	Backup and Recovery	114
5.6.6	Data Replication (Local, Remote, and Cloud)	120
5.6.7	Data Protection Management	122
5.7	Modernizing Data Protection and Backup	123
5.7.1	Expanding from DR to BC, Shifting from Cost Overhead to Profit Center	124
5.7.2	Using Virtualization and Clouds to Enhance Data Protection	124
5.8	Data Protection Checklist	127
5.9	Common HA-, BC-, and DR-Related Questions	127
5.10	Chapter Summary	128
Chapter 6	Metrics and Measurements for Situational Awareness	129
	In This Chapter	129
6.1	Getting Started	129
6.2	Making Sense of Metrics and Measurements	131
6.3	Different Metrics for Different Audiences	132
6.4	Key Performance Indicators	133
6.4.1	Averages, Rates, and Ratios	134
6.4.2	Compound Metrics	136
6.5	Measuring IT Resources and Services Delivery	137
6.5.1	Performance, Availability, Capacity, Energy, and Economics (PACE)	138

6.6	Where to Get Metrics	145
6.7	Accounting and Chargeback	146
6.8	Benchmarks and Simulation Comparisons	147
6.9	Common Metrics-Related Questions	149
6.10	Chapter Summary	150
Chapter 7	Data Footprint Reduction: Enabling Cost-Effective Data Demand Growth	151
	In This Chapter	151
7.1	Getting Started	152
7.1.1	What Is Driving Expanding Data Footprints	154
7.1.2	Changing Data Access and Lifecycles	155
7.1.3	What Is Your Data Footprint Impact?	155
7.1.4	Business Benefits of Data Footprint Reduction	157
7.2	The Expanding Scope and Focus of Data Footprint Reduction	158
7.2.1	Reducing Your Data Footprint	159
7.2.2	Not All Data or Applications Are the Same	160
7.3	DFR Techniques	161
7.4	Metrics and Measurements	161
7.5	What to Look for in a DFR Technology Solution	162
7.6	Best Practices	163
7.7	Common DFR Questions	163
7.8	Chapter Summary	164
Chapter 8	Enabling Data Footprint Reduction: Storage Capacity Optimization	167
	In This Chapter	167
8.1	DFR Techniques	167
8.2	Archiving	168
8.2.1	Tools and Targets	171
8.3	Compression and Compaction	172
8.3.1	Compression Implementation	173
8.3.2	Real-Time and On-the-Fly Compression	173
8.3.3	Postprocessing and Deferred Compression	176
8.4	Consolidation and Storage Tiering	176
8.5	Data De-duplication	177
8.5.1	Dedupe Fundamentals	178
8.5.2	How and Where Dedupe Is Implemented	181

8.5.3	Dedupe Locations (Hardware, Software, Appliance, Source, and Target)	182
8.5.4	Global vs. Local Dedupe	184
8.6	DFR and RAID Configurations	186
8.7	Space-Saving Snapshots	187
8.8	Thin Provisioning	188
8.9	Common DFR Questions	189
8.10	Chapter Summary	189
Chapter 9	Storage Services and Systems	191
	In This Chapter	191
9.1	Getting Started	191
9.2	Tiered Storage	193
9.3	Storage Reliability, Availability, and Serviceability (RAS)	194
9.4	Aligning Storage Technology and Media to Application Needs	195
9.4.1	Hard Disk Drives	197
9.4.2	Hybrid Hard Disk Drives	198
9.4.3	Removable Hard Disk Drives	199
9.4.4	Solid-State Devices	200
9.4.5	Magnetic Tape	203
9.4.6	Different Storage Media Are Better Together	204
9.5	Storage Services and Functionalities	206
9.5.1	Redundant Components	206
9.5.2	RAID and Data Availability	208
9.5.3	Enhancing Data Availability and Protection	212
9.5.4	Automated Storage Tiering and Data Movement	214
9.5.5	Performance Optimization	215
9.5.6	Unified, Multiprotocol, and Function Storage	216
9.5.7	Intelligent Power Management and "Green" Storage	218
9.5.8	Management Tools	219
9.6	Storage System Architectures	220
9.6.1	Servers as Storage, Storage as Servers	221
9.6.2	Clustered and Grid Storage	222
9.6.3	Cloud Storage	225
9.7	Storage Virtualization and Virtual Storage	226
9.7.1	Volume Mangers and Global Name Spaces	226

9.7.2	Virtualization and Storage Services	226
9.7.3	Storage Virtualization Locations	229
9.8	Common Storage Questions	230
9.9	Chapter Summary	230
Chapter 10	Server Virtualization	233
	In This Chapter	233
10.1	Getting Started	233
10.2	Virtual Servers	234
10.3	Inside Virtual Servers and Virtual Machines	235
10.4	Virtual Desktop Infrastructure	241
10.5	Cloud and Virtual Servers	243
10.6	Can and Should All Servers or Desktops Be Virtualized?	244
10.7	Virtualization Beyond Consolidation: Enabling IT Agility	245
10.8	Common Virtualization Questions	247
10.9	Chapter Summary	248
Chapter 11	Connectivity: Networking with Your Servers and Storage	249
	In This Chapter	249
11.1	Getting Started	249
11.2	Networking Challenges	251
11.3	I/O and Networking Bits and Bytes, Decoding Encoding	252
11.4	I/O and Networking Fundamentals	253
11.5	Server (Physical, Virtual and Cloud) Topics	255
	11.5.1 Peripheral Component Interconnect (PCI)	255
	11.5.2 Adapters, NICs, and CNAs	257
11.6	I/O and Networking Devices	259
11.7	Converged and Unified Networking	260
	11.7.1 PCI-SIG I/O Virtualization	261
	11.7.2 Converged Networks	263
11.8	Local Networking (DAS, SANs, and LANs)	265
	11.8.1 The Role of Networking with Storage	265
	11.8.2 Ethernet (802.1)	267
	11.8.3 Fibre Channel (FC)	268
	11.8.4 Fibre over Ethernet (FCoE)	269
	11.8.5 InfiniBand (IBA)	271
	11.8.6 iSCSI (Internet SCSI)	272

11.8.7	Serial Attached SCSI (SAS)	272
11.8.8	The Best Protocol for Block Storage	274
11.9	Enabling Distance (MANs and WANs)	275
11.9.1	Bandwidth and Protocol Optimization	277
11.10	Cloud, Virtualization, and Management Topics	278
11.10.1	Accessing Cloud Storage and Services	279
11.10.2	Virtual I/O and I/O Virtualization (IOV)	280
11.10.3	N_Port ID Virtualization (NPIV)	281
11.11	Configuring for Reliability, Availability, and Serviceability (RAS)	282
11.11.1	Flat Networks	283
11.11.2	Configuration and Topologies	284
11.11.3	Cabling: Tools and Management	285
11.12	Common Networking Questions	286
11.13	Chapter Summary	287
Chapter 12	Cloud and Solution Packages	289
	In This Chapter	289
12.1	Getting Started	289
12.2	Clarifying Cloud Confusion: What Does Your Cloud Need to Do	291
12.3	IaaS, PaaS, SaaS, and AaaS	291
12.4	Accessing Clouds	293
12.5	Public Cloud Services	295
12.6	Private Clouds	297
12.7	Stacks and Solutions	297
12.8	PODs and Modular Data Center Components	299
12.9	Vendor Lock-in: The Good, the Bad, and the Ugly	300
12.10	Evaluating Cloud Servers and Solutions	302
12.11	Common Cloud Questions	303
12.12	Chapter Summary	304
Chapter 13	Management and Tools	305
	In This Chapter	305
13.1	Getting Started	305
13.2	Software and Management Tools	306
13.3	Management Tool Interfaces	308
13.4	End-to-End Management	310

13.5	Licensing Topics	312
13.6	The Evolving Role of Management Tools	314
13.7	Hard vs. Soft Products	315
13.8	The Other IT Resources: People, Processes, and Policies	316
13.9	Common Management-Related Questions	318
13.10	Chapter Summary	318
Chapter 14	Applying What You Have Learned	321
	In This Chapter	321
14.1	Getting Started	321
14.2	Don't Be Afraid, but Look Before You Leap	322
14.3	Addressing Issues and Challenges While Enabling Opportunities	323
14.4	What's Your Vision, Strategy, and Plan?	324
14.5	What to Consider When Evaluating Technologies, Techniques, and Services	327
14.6	Common Cloud, Virtualization and Data Storage Networking Questions	328
14.7	Chapter Summary	329
Chapter 15	Wrap-up, What's Next, and Book Summary	331
	In This Chapter	331
15.1	Wrapping up This Part of the Journey	331
15.2	Where Have We Been, Where Are We Going: From Efficiency to Effectiveness	332
15.3	What to Do Next	334
15.4	Futures, Trends, Perspectives, and Predictions	336
15.5	Chapter and Book Summary	338
Appendix		341
Glossary		343
Index		357

This page intentionally left blank

Preface

Since I wrote *The Green and Virtual Data Center* (CRC Press, 2009), I have had the opportunity to meet and speak with thousands of information technology (IT) professionals around the world, both in person and virtually via different venues. Even as that last book was being printed, the ideas to be found in this new book were coming together, as I refined my premises and identified new industry trends. One of those trends is the result of hype and FUD (fear, uncertainty, and doubt) about material from various sources. However, there is a gap between the hype and the FUD, addressed by a theme of this book: *Don't be scared of clouds and virtualization, but do look before you leap*. What this means is that you should do your homework, prepare, learn, and get involved with proof of concepts and training to build the momentum and success to continue an ongoing IT journey to enable agility, flexibility, scale, and productivity while reducing costs and complexity.

Another notable trend today is that the amount of data being generated, moved, processed, and stored for longer periods of time shows no signs of decreasing. Even during the recent global economic crisis and recessions, there were few, if any, signs of a data or information slowdown. In fact, there was evidence to the contrary: While financial institutions were being plagued with record-breaking negative economic impacts, they were also generating, moving, and processing record amounts of data that had to be preserved. The only thing that has changed is that we need to do even more with less—or more with what we currently have—to support and sustain business growth. This means leveraging existing and new or emerging technologies and techniques to stretch available resources, including budgets, people, floor space, and power, to support business growth, new applications, and even more data and information.

To sustain business growth while enabling new functionalities or services, providers of information services need to look at various options for becoming more efficient. Becoming more efficient means more than cost avoidance; it also includes boosting productivity while streamlining information services delivery. This means stretching resources (people, processes, budgets, hardware, software, energy, facilities, and services) further while enabling better business agility and productivity. This book picks up where

The Green and Virtual Data Center (Intel recommended reading for developers) ended and considers how IT or other providers of information services can do more with available resources (people, processes, products, and services) while reducing per-unit costs and maintaining or enhancing quality of service and customer satisfaction.

Beyond Buzzwords, Hype, and FUD

There are those who think that cloud technology is all about building from the ground up, or at least establishing new protocols, interfaces, management standards, and reference models. Not surprisingly, these tend to be engineers, technical marketers, investors, entrepreneurs, or just fans of the latest “shiny new toy.” Then there are those who believe that cloud and virtualization techniques and their associated technologies can be used to complement or enhance existing environments.

This book looks at clarifying “cloud confusion” and expanding the discussion of virtualization beyond consolidation for enabling agility, flexibility, and ease of management. For some, this will mean private clouds or traditional IT approaches leveraging some new technologies; for others, it will mean public clouds used completely or in a complementary manner. Some moving to public clouds will use technology that is still emerging, perhaps rebuilding or rip-and-replace, and others will see the move as a green-field or clean-sheet opportunity.

Who Should Read This Book

This book cuts across various IT data technology and resource domains to provide a single source that discusses the interdependencies that need to be supported to enable a virtualized, efficient, effective, and agile information services delivery environment. Do you need or want a cloud? Do you have to have or would you like a virtual environment? Do you feel compelled to have a converged data and storage network, or is there a particular business opportunity or challenge? What is the business case, demand, challenge, or opportunity for addressing or enabling clouds, dynamic infrastructure, and virtual technology? This book looks at these and other questions, providing answers, ideas, and insight to stimulate thinking about where, when, why, and how to deploy cloud, virtualization, and data storage networking resources on a public, private, or legacy IT basis. The book is about convergence in terms of technologies, techniques, and various best practices that pertain to cloud, virtualization, dynamic infrastructure, and traditional environments’ delivery of information services.

Audiences that will benefit from reading this book include IT purchasing, facilities, server, storage, networking, database, and applications analysts, administrators, and architects, as well as CIOs, CTOs, CMOs, and CFOs. Also, manufacturers and solution partners (vendors), value-added resellers, consultants, sales, marketing, support, and engineering specialists, public relations, investment communities, and media professionals associated with IT technologies and services can all find something of interest.

The book looks at the changing role of data and storage networks to support and sustain resilient and flexible, scalable virtual and cloud environments, and how to leverage those techniques into existing environments to achieve great efficiency, boosting service while reducing per-unit costs. If this resonates or if you want to learn more, then this book is a must-read for real-world perspectives and insight to address server, storage, networking, and other infrastructure resource management topics to support current and next-generation public or private virtual data centers that rely on flexible, scalable, and resilient data storage and networks. This could be the starting point on your cloud or virtualization journey, but also a great resource for use in traditional environments. It is short on hype and FUD; instead, it focuses on what you need to determine where various technologies and techniques can be applied.

How This Book Is Organized

This easy-to-navigate book is divided into four parts. Part I, “Why Cloud, Virtualization and Data Storage Networks Are Needed,” includes Chapters 1 and 2 and covers the background and basics of information service delivery and clouds; Part II, “Managing Data and Resources: Protect, Preserve, Secure, and Serve,” includes Chapters 3 through 6 and looks at common management tasks along with metrics for enabling efficient and effective data infrastructure environments. Part III, “Technology, Tools, and Solution Options,” includes Chapters 7 through 13 and explores the various resource technologies (servers, storage, and networking) and techniques. Finally, Part IV, “Putting IT All Together,” comprising Chapters 14 and 15, brings together the previous parts and provides a glimpse into the future of cloud, virtualization, and data storage networking.

Is It a Nonstop Flight or a Journey with Intermediate Stops and Layovers?

A major theme of this book is that IT has been on a journey for several decades to get to a place where more can be done with available resources while maintaining or enhancing quality of service, feature functionality, and cost reduction. Challenges of journeys include departing too soon, before proper preparations have been made, or waiting too long and missing an opportunity. On the other hand, rushing in too quickly may lead to surprises that result in less than pleasant experiences. So, don’t be scared of clouds, dynamic infrastructure, and virtualization, but look before you leap. Learn the benefits as well as the caveats of clouds, and understand where the gaps are, so that you can work around them while leveraging what is available to expand your horizons for the long haul.

As you read the chapters in this book, you will discover a mix of existing and emerging technologies; to some this will be review, while for others it may be new. The main idea is that cloud and virtual environments rely on physical or fundamental resources, processes, and people operating collectively and collaboratively in a more

efficient, effective, and agile manner. Whether you are going all in with clouds and virtualization or are simply looking to expand your awareness while continuing with business as usual for your environments, the technologies, techniques, and best practices laid out in these pages apply equally to cloud, virtual, and physical data and storage networking environments of all sizes.

For some environments, there have been numerous intermediate stops during the journey from mainframe to distributed to client server to Web-based to consolidated, virtualized, and cloud computing paradigms, with various initiatives, including service-oriented architectures (SOAs), information utilities, and other models along the way. Some organizations have transitioned “nonstop” from one era to another, whereas others have had clean-sheet or green-field (starting from scratch) opportunities.

For some, the journey is to the cloud (public or private), while for others the cloud is a platform to enable a transition through an information continuum journey of years or decades. For example, despite being declared dead, the mainframe is very much alive for some organizations, supporting legacy and Linux open systems providing SOAs and private or public cloud capabilities. The trick is finding the right balance of old and new, without clinging to the past or rushing into the future without having an idea of where you are going.

An important part of the journey involves measuring your progress, determining where you are and when you will get to where you are going while staying within your budget and on schedule. Keeping resources safe during transition, for business continuance, disaster recovery, or data protection in general, is also important.

When you go on a trip for business or fun, your vehicle, medium, or platform may be foot, bicycle, automobile, plane, train, or some combination. You decide on the method based on performance or speed, capacity, and comfort; space, reliability, and schedule availability; effectiveness, personal preferences, and economics. Often the decision is made entirely on economics, without factoring in time and productivity or enjoyment. Sometimes, an airplane ride from the U.S. West Coast to the Midwest is more productive because of the time saved, even if the cost is higher than traveling by car.

Having said all of that, it is now time to stow your items, place your seat back and tray table in the upright and locked position, and secure your seat belt as we prepare for takeoff. I hope you enjoy your journey through the following chapters and pages.

Greg Schulz

Acknowledgments

Writing a book is more than putting pen to paper or, in this case, typing on a computer: It involves hundreds of hours working behind the scenes on various activities. Writing a book is similar to a technology development project, whether hardware, software, or a service, in that it includes an initial assessment of the need. Having established the need for this project, I chose to write the book myself. Other behind-the-scenes activities in any project include research, design and validation, actual content generation, edits, reviews or debugging, more edits, working with text as well as graphics, project management, contracts, marketing, and production, among others.

Thanks and appreciation to all of the vendors, value-added resellers, service providers, press and media, freelance writers, venture capitalists, bloggers, consultants, twitter tweeps, and IT professionals around the world I have been fortunate enough to talk with while putting this book together. I would also like to thank others who were directly or indirectly involved with this project, including Andy Fox, Mike Connor, Corey Donovan, Rob Dombrowsky, Jim Dyer, Chris Foote, Rich Lillis, Ed Haletky, Mike Hall, Greg Knieriemen and Marc Farley of Infosmack, Keith Norbie, Bruce Ravid, Drew Robb, Hubert Smith, Eric Siebert, Stanley Stevens, the Nelsons and Schoellers, as well as George Terwey.

Special thanks to Tom Becchetti, Greg Brunton and Kay Wylie, Gert Brouwer, Georgiana and Alex Comsa, Preston Deguise, David Marshall, Dr. “J” Metz, and Sherryl Savage. Thanks to John Wyzalek, my publisher, with whom I formalized this project over a great lunch in New York City while watching the St. Patrick’s Day parade at a restaurant in Grand Central Station (one of the best burgers ever!), Chris Manion and everyone else at CRC/Taylor & Francis/Auerbach, as well as a big thank you to Theron Shreve of Derryfield Publishing Services and his crew, including Lynne Lackenbach and Marje Pollack. Very special thanks to Damaris Larson, who continues to push and challenge me in support of this and countless other writing projects: You are the best!

Finally, thanks to my wife, Karen, and to “Big Babe” and “Little Leo,” for having the patience to support me while I worked on yet another book project!

To all of the above and to you the reader, thank you very much.

Greg Schulz

This page intentionally left blank

About the Author

Greg Schulz is an independent IT industry advisor, author, blogger (<http://storage-ioblog.com>), and consultant. Greg has over 30 years of experience across a variety of server, storage, networking, hardware, software, and services architectures, platforms, and paradigms. He brings the rare perspective of having been a IT professional working in various business sectors (e.g., as an IT customer), a vendor, and a value-added reseller, in addition to having been an analyst and advisory consultant across servers, storage, networks, hardware, software, virtualization, and cloud services.

After spending time as a customer and a vendor, Greg became a Senior Analyst at an IT analysis firm covering virtualization, SAN, NAS, and associated storage management tools, techniques, best practices, and technologies in addition to providing advisory and education services. In 2006, Greg leveraged the experiences of having been on the customer, vendor, and analyst sides of the “IT table” to form the independent IT advisory consultancy firm Server and StorageIO (StorageIO). He has been a member of various storage-related organizations, including the Computer Measurement Group (CMG), the Storage Networking Industry Association (SNIA), and the RAID Advisory Board (RAB), as well as vendor and technology-focused user groups. Greg also speaks frequently at conferences, seminars, and private events around the world.

Greg has received numerous awards and accolades, including being named a VMware vExpert and an EcoTech Warrior by the *Minneapolis-St. Paul Business Journal*, based on his work with virtualization, including his book, *The Green and Virtual Data Center* (CRC Press, 2009). In addition to his thousands of reports, blogs, twitter tweets, columns, articles, tips, pod casts, videos, and webcasts, Greg is also author of the SNIA-endorsed study guide, *Resilient Storage Networks—Designing Flexible Scalable Data Infrastructures* (Elsevier, 2004). Greg is regularly quoted and interviewed as one of the most sought-after independent IT advisors providing perspectives, commentary, and opinion on industry activity. Greg has a B.A. in computer science and a M.Sc. in software engineering from the University of St. Thomas. Learn more at www.storageio.com.

This page intentionally left blank

Chapter 1

Industry Trends and Perspectives: From Issues and Challenges to Opportunities

There is no such thing as a data or information recession.

– Greg Schulz

In This Chapter

- Moving beyond the hype
- Navigating the maze of cloud and virtualization stories
- The business demands of IT and data storage
- IT issues and challenges involving data storage
- The business benefit of cloud and virtual data storage networking
- Opportunities to address data storage issues and challenges
- The role of virtualization, cloud, and storage networking
- Maximizing IT resources without compromising quality of service
- What defines a public and private cloud service, product, solution, or paradigm
- The importance of information access, data consistency, and availability

2 Cloud and Virtual Data Storage Networking

This chapter looks at the big picture of business issues and demand drivers that set up the need for cloud, virtualization, and data storage networking. Key themes and buzzwords covered include cloud computing, cloud storage, public and private clouds, information factories, virtualization, business issues or challenges, barriers to productivity, technology tools and techniques, along with best practices. Additional themes and topics discussed include enabling agility, flexibility, scalability, resiliency, multitenancy, elasticity, managed service providers (MSPs), converged networks, Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and IT optimization.

1.1. Getting Started

You probably didn't wake up this morning thinking, "I need to have someone buy or implement a cloud, virtualization, or storage networking solution." Granted, if you are a vendor or a consultant, that may be your job (assessing, designing, selling, installing, or supporting data storage, networks, virtualization, or clouds). However, if you are not a vendor, consultant, analyst, or journalist, but rather someone responsible for information technology (IT)-related solutions for your organization, typically the need to buy and deploy a new solution is tied to solving some business issue or capitalizing on an opportunity.

A common challenge in many organizations is exploding data growth along with associated management tasks and constraints, including budgets, staffing, time, physical facilities, floor space, and power and cooling. Before going further into why you need or do not need a cloud, virtualization, or a storage network, let's take a step back and look at what is driving data growth and the consequent need to manage it more effectively.

1.2. The Importance of Data and Storage

We live in an information-centric world. As a society, we have a growing reliance on creating and consuming data (Figure 1.1), which must be available when and where it is needed. Data and related information services are enabled or provided via information technology services combining applications, facilities, networks, servers, storage hardware, and software resources.

More data can be stored in the same or smaller physical footprint than in the past, thus requiring less power and cooling per gigabyte (GB), terabyte (TB), petabyte (PB), or exabyte (EB). However, data growth rates necessary to sustain business activity, enhance IT service delivery, and enable new applications are placing continued demands requiring more processing, network, or input/output (I/O) bandwidth and data storage capacity.

As a result of this increasing reliance on information, both for home and personal use along with business and professional needs, more data is being generated, processed, moved, stored, and retained in multiple copies for longer periods of time. The

net result is that IT organizations of all sizes are faced with having to do more with what they have (sometimes with less), including maximizing available IT resources while overcoming common footprint constraints (available power, cooling, floor space, server, storage and networking resources, management, budgets, and IT staffing).




Figure 1.1 IT and data storage demand drivers.

1.2.1. The Business End of IT Data Storage Impact

Just as we live in an information-centric society which extends from home to the office, from the small office/home office (SOHO) to the remote office/branch office (ROBO), small/medium-size business (SMB), small/medium enterprise (SME), to ultra-large organizations or enterprises, there is another common theme, and that is economics. Economics are a constant focus, whether it is costs or expense, profits and margins, return on investment (ROI), total cost of ownership, or some other business specific measurement.

On the one hand, there is a need or reliance on having more information; on the other, there are the constants of economics, cause and effect, and supply and demand. You need or want information, but there is a cost to supporting or managing it. Yet information can also directly or indirectly drive profits, so a balancing act is necessary. Thus, to support or sustain economic (business) growth or manage the data necessary to maintain daily activity, there are associated costs (hardware, software, people, facilities, power, etc.) that need to be managed.

Innovation is doing more with what you have: supporting growth and enhancement of services without negatively impacting service-level objectives (SLOs), including quality of services, while reducing per-unit cost for service delivery (as shown in Figure 1.2). The trick is to find the balance among boosting productivity, reducing costs, and maintaining or enhancing customer service delivery.

Figure 1.2 sums up the balancing act of maximizing use of available IT resources while supporting growing business demands in a cost-effective manner. IT resources include people, processes or best practices, time, budgets, physical facilities, power,

4 Cloud and Virtual Data Storage Networking

cooling, floor space, server, storage and networking hardware, along with software and services. All too often, the approach has been to drive cost down by increasing utilization at the expense of quality of service (QoS) and SLOs. An example is leveraging consolidation or migration to a cloud service based on a lower-cost model that trades QoS and SLO for price.

Another variation is to boost QoS and SLOs along with performance to meet demand at the expense of cost or less effectively utilized resources. In other words, it's relatively easy to improve in one area while causing issues or aggravation in another. Innovation occurs when all three categories shown in Figure 1.2 are positively impacted.

Figure 1.2 identifies constraints or barriers to cost-effective service delivery while maintaining or enhancing the service delivery experience including QoS and SLOs. Cloud, virtualization, and data storage networking are tools and techniques that, combined with best practices, can be used to enable innovation and meet the objectives of Figure 1.2.

Clouds, virtualization, and data storage networks can be used to enable cost reduction and stretching of resources by supporting consolidation initiatives. However, these same tools and techniques can also be used for enabling agility, flexibility, and enhanced services that can improve both top- and bottom-line business metrics. For some environments or applications the focus may be on cost reduction while supporting little to no growth, while for others it may mean working with the same or a slightly increased budget while supporting business demand and SLOs. In some organizations this also means reducing costs or stretching available budgets and resources to do more with what they have.

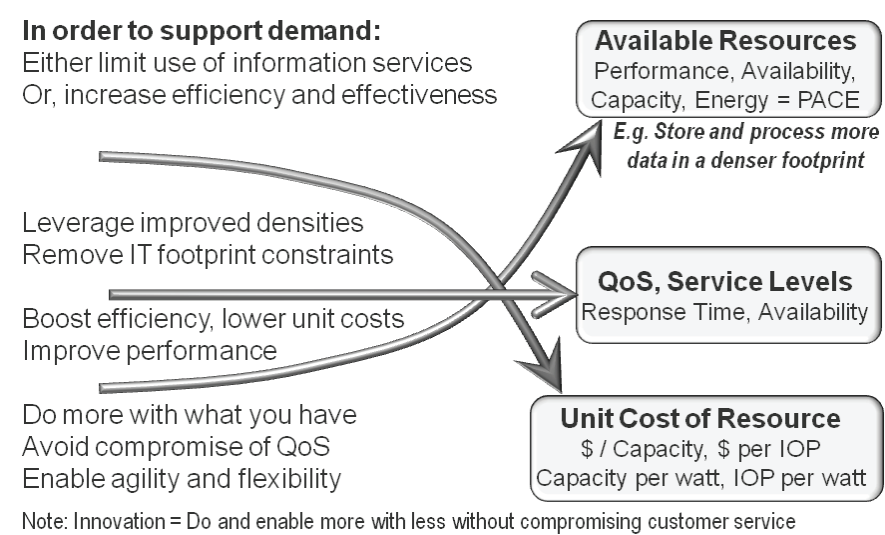


Figure 1.2 Supporting demand, maintaining quality of service (QoS), while reducing costs.

1.2.2. Addressing Business and IT Issues

Clouds, virtualization, and storage networks are tools, techniques, and best practices to help support or sustain growth while reducing per-unit costs, removing complexity, enabling flexibility or agility, and also enhancing customer experiences. Clouds, virtualization, and storage networks are not the objective themselves; rather, they are tools, vehicles, or mechanisms that can be used to help achieve broader business and IT objectives. They can be used for new, start-from-scratch environments; they can also be aligned with existing IT service delivery as well as help with a transition-over-time evolution of IT.

Thus, taking a step back from the technology, tools, and techniques, and keeping the bigger picture in focus, helps to understand what to use when, where, and why, as well as how to go about it in a more effective manner.

1.2.3. What Is Driving Data Growth and Information Reliance

The popularity of rich media and Internet-based applications has resulted in explosive growth of unstructured file data, requiring new and more scalable storage solutions. General examples of unstructured data include spreadsheets, PowerPoint, slide decks, Adobe PDF and Word documents, Web pages, video and audio, JPEG, MP3, and MP4, photos, audio, and video files.

Examples of applications driving continued growth of unstructured data include:

- Gaming, security, and other surveillance video or security
- Unified communications including Voice-over-IP (VoIP)
- Rich media entertainment production and viewing
- Digital archive media management
- Medicine, life science, and health care
- Energy including oil and gas exploration
- Messaging and collaborations (email, IM, texting)
- Internet, Web, social media networking, video and audio
- Finances, marketing, engineering, and customer relations management (CRM)
- Regulatory and compliance requirements

While structured data in the form of databases continues to grow, for most environments and applications the high-growth area and expanding data footprint along with associated performance bottlenecks is centered on semistructured email data and unstructured file data. Unstructured data has varying I/O characteristics that change over time—for example, data starting out with a lot of activity, then going idle for a time, followed by extensive reads, as in the case of a video or audio file becoming known and popular on a media, entertainment, social networking, or a company-sponsored website.

Data footprint is the total data storage needed to support your various business application and information needs. Your data footprint may, in fact, be larger than how much

actual data you have. A general approach to determine your data footprint is to simply add up all of your on-line, near-line and off-line data storage (disk and tape) capacity.

1.3. Business Issues and IT Challenges

I commonly get asked if virtualization and clouds are a passing fad, full of hype, or if they are real and being attacked by fear—uncertainty—doubt (FUD). Granted, and unfortunately, there is a lot of hype along with FUD, leading to confusion about both cloud and virtualization—tending to set them up as popular fads, much as compliance, “green” IT, information lifecycle management (ILM), client server, and storage networking were initially viewed.

Common business issues, challenges, and trends pertaining to IT include:

- Increased reliance on information services being accessible when needed
- Competitive and other market dynamics causing financial constraints and focus
- Regulatory compliance and other industry or corporate mandates
- Stretched resources (staffing levels, skill sets, budgets, facilities)
- The need to reduce costs while increasing services and productivity
- A shift from cost reduction or avoidance to efficiency and effectiveness models

How often do you use data storage or information services? Perhaps you use data storage without realizing it, at home, at work, and elsewhere. Data storage is in play all around us, used for different purposes and in various forms. Some might say that data storage is the most important IT resource compared to servers or computers, networks, desktop, laptops or workstations, or application software tools. On the other hand, some would say that networks are the most important, or servers, or whatever is that individual’s specialty. For the sake of argument I will position data storage as equal to servers, networks, hardware, and software, as all are needed to be effective.

Common IT issues, challenges, problems, and trends include:

- More data to process, move, manage, store, and retain for longer periods of time
- Increased reliance and expectation that information services be available 7×24
- Limited or strained resource constraints causing bottlenecks or barriers
 - People or staffing and applicable skill sets
 - Hardware, software, and networking bandwidth
 - Budgets (capital and operating)
 - Power, cooling, floor space
 - Time for backup or data protection windows
- Regulatory, compliance, and other regulations
- Demand causing performance, availability, capacity, and energy (PACE) impacts
- Software or hardware licensing and maintenance, support as well as service fees
- Aging IT infrastructures along with related interoperability and complexity
- Time involved in aligning IT resources to business or service needs
- Speed and accuracy of IT resource provisioning

When I talk with IT professionals or customers, I ask them if they have a mandate to reduce costs, which is a common industry messaging or rallying theme. Surprisingly, a majority of them indicate that it is not costs per se that they have to reduce (though some do). Rather, they have to do more with what they have with their current budget to support business growth, new applications, and functionality.

1.4. Business and IT Opportunities

Now, back to the question you woke up with this morning: “Do I need to have someone buy or implement a cloud, virtualization, or storage networking solution?”

Or maybe you woke up wondering how you are going to support business growth, demands for more data, flexibility, reduce cost, and enhance service delivery. Or perhaps you need to figure out how to defend your environment or market your environment to the rest of your business as opposed to the business going to external resources.

For some, efficiency and optimization can be avoidance or simply increasing utilization to reduce or spread costs over more work being done. However, another form of efficiency and effectiveness is stretching resources to do more while boosting productivity or removing barriers and constraints.

1.4.1. *Traditional Information Services Delivery/Model*

Delivery of information services continues to evolve. As this evolution of techniques, technologies, best practices, and new products continues, there are decades of legacy applications and data that need to be supported.

The basics of any information services delivery model can be summarized (Figure 1.3) as users accessing business applications running on a server with information stored somewhere, all accessed via some device over a network. The device could be a dumb terminal cabled or networked to the server with dedicated direct attached disk storage or a smart phone via a wireless connection.

Similarly, the server could be virtualized or nonvirtualized with an operating system, database, and other tools to support and manage the applications and storage. From these basic information services delivery models, additional deployment scenarios can be established, including dedicated or shared applications, services, time sharing, or what we call today cloud and managed services, among others. Likewise, there can be different types of tiered servers, storage, and networks, which will be discussed in more detail over the next several chapters.

A common expression is that what is old is new and what is new is old. In the case of cloud and virtualization, for some these are new or perhaps repackaged open variations of what they have seen, heard, or read about previously. For example, IBM mainframes have had propriety virtualization for several decades. Those same platforms—which have been declared dead or dinosaurs—with their tenets of being highly optimized, metered or instrumented with metrics and reporting, scalable, and resilient—are what some cloud attributes seek to emulate from an open perspective. This has led to some

interesting discussions about why not keep things on the mainframe (clinging to the past) versus moving everything to the cloud (racing to the future).

Depending on the types of applications, the answer may be to move some or all to the cloud. On the other hand, leveraging a cloud managed service provider, hosting, or outsourcing can be the answer for other applications to coexist with your environment. For example, if you still have the need for an IBM z/OS class mainframe, but it has become a small part of your environment, then outsourcing it or leaving a hosted or managed service can be an option.

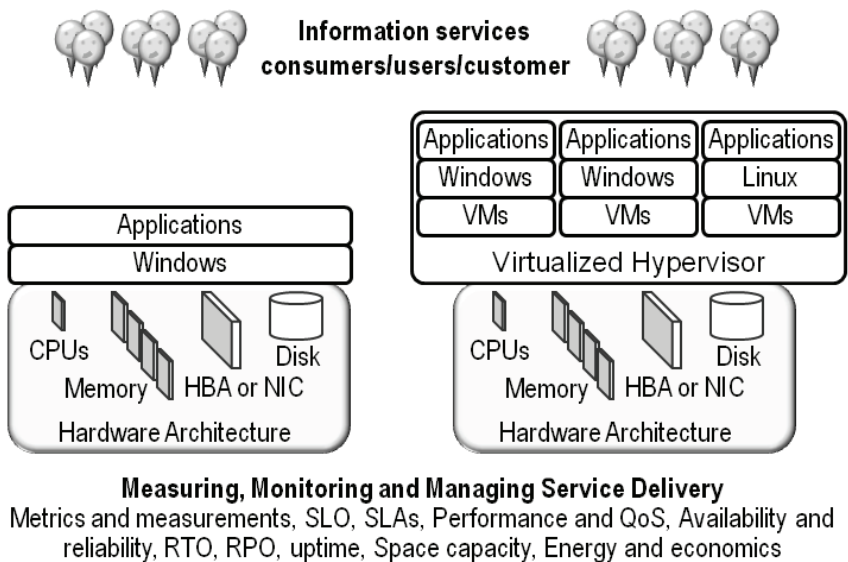


Figure 1.3 Information services delivery basics.

1.4.2. Information Factories

Most IT organizations or infrastructures exist to support the business applications and information needs of an organization. In some cases, the business applications services provided by IT include supporting factories, accounting, marketing, and engineering, among others. However, IT or information providers also often suffer from “shoemaker’s children” syndrome in that they may not have adequate insight or management tools for themselves. For example, an organizations may have accounting and tracking systems supported by IT, but does IT have accounting or metrics on performance, availability, capacity, configuration, energy, and economics for a given service delivery?

Traditional factories (Figure 1.4) leverage different tools, techniques, metrics, measurements, best practices, resources, and people skill sets to build and deliver goods or services to a prescribed service level and price point. Factories can be dedicated or




Figure 1.4 Traditional factory.

private, they can be shared or supplier-based, or they can be hybrid, similar to how IT services can be sourced, used, and delivered. An organization may have its own factory, or its factory could be a virtual or third-party jobbing or other service. Goods/services may be produced by someone under contract. IT services may be delivered via an organization's own factory, via a third party, or virtualized.

Basic characteristics of factories include:

- Reliable, to meet demand, avoid downtime, avoid mistakes and rework
- Scalable, to meet changing workload demands
- Efficient, reduce waste, customer SLOs met in an economical manner
- Work is done quickly, yet reliably, with good quality
- Flexible capacity and ability to retool to meet changing needs
- Factories may be wholly owned, shared, or owned by a third party
- Factories consume materials and resources to create/deliver goods and services
- Those goods and services may in turn be consumed by other factories
- Factories produce product to a blueprint, template, or run book specifications

The notion of the information factory (Figure 1.5) sets up the discussion around cloud, virtualization, and storage networks on either a public, private, or hybrid basis.

For some, the idea of an information factory and cloud may bring *déjà vu* experiences of the information utility model of the late 1980s and early 1990s.

Additional characteristics of factories include:

- Rely on suppliers or secondary and tertiary factories (subs)
- Have bill of materials, metrics and measurements, costing information
- Quality assurances programs to ensure that QoS and SLOs are being met

10 Cloud and Virtual Data Storage Networking

- Focus on reducing defects and waste while boosting productivity to reduce cost
- Build templates for optimized information service delivery
- Best practices, processes, policies and procedures
- Balance of productivity, cost control, waste reduction, utilization, meeting SLOs
- Leverage new technologies that have good payback for enabling goals
- Cost-effective habitats for deploying and using technologies
- Efficiency gained with repeatable processes, and increased workload activity

Information factories can be

- Private
- Public
- Hybrid

Information factories (or clouds) should be or enable:

- Multitenancy, measurability, and accountability
- For service providers, this can include chargeback
- Secure, flexible, dynamic, scalable, and resilient
- Able to relocate services as needed
- Rapid deployment and provisioning of resources
- Efficient, cost-effective resource usage that meets QoS and SLAs
- Automate and guide users or customers to best-fit services selection

The similarities between factories, information factories, clouds, and information services delivery should be clear.

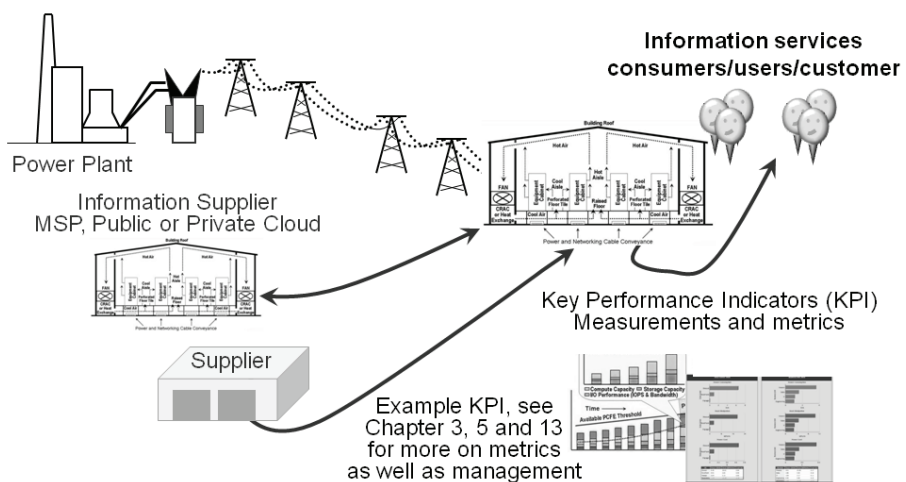


Figure 1.5 Information factory.

1.5. Opportunity for Cloud, Virtualization, and Data Storage Networking

Like a physical factory, some of an information factory's work is done on the premises and some off-site at other locations, including those of subcontractors or suppliers. In the case of information factories, the product being produced is information services, with the machinery being servers, storage, and I/O networking managed with software, processes, procedures, and metrics. Raw materials include data, energy to power and cool the physical facility, and technologies, all operating to deliver the services at a low defect or error rate while meeting or exceeding QoS, performance, availability, and accessibility requirements in a cost-effective manner.

For some cloud compute or storage providers, the value proposition is that they can supply the service at a lower cost than if you use your own capabilities. Similar to service bureaus, out-sourcing, managed service, or hosting facilities of the past, cloud-based services are a means of shifting or avoiding costs by moving work or data elsewhere to be processed or stored.

However, it is a mistake to consider clouds for just for their cost-saving abilities while ignoring performance, availability, data integrity, ease of management, and other factors that can impact service delivery and expenses. Clouds should be looked at not as a replacement or competing technology or technique, but rather as a complementary approach to existing in-house resources.

Cloud computing and storage are simply additional tiers of servers and data repositories that may have different performance, availability, capacity, or economics associated with them to meet specific business and/or application needs. That is, cloud computing and cloud storage coexist and complement what is currently being done, with the objective of boosting quality of service, availability, or customer satisfaction while supporting more data being processed, moved, and stored for longer periods of time at a lower unit cost.

1.5.1. IT Clouds and Virtualization: Not If, Rather When, Where, Why, and How

There are many different types and definitions of clouds, including those of the National Institute of Standards and Technology (NIST) and the Data Management Task Force (DMTF). Cloud computing is a paradigm, and thus its definition is still evolving along with use cases and the underlying technologies, techniques, and best practices.

Some see clouds as the wave of the future, even if they're not sure what that future may be. To others, a cloud is a cloud if, and only if, it is outside of what you currently are doing or have done with IT. Some will argue that a cloud is only a cloud if new hardware or software is involved, while others will assert that a cloud is only a cloud if your applications and data exist outside your location.

Consequently, different people will have different thoughts or perspectives about clouds, depending on their perception or definition. For example, in Figure 1.6, thoughts and opinions based on an ongoing StorageIO blog research poll of a mix of

vendors, IT professionals, and others shows at two extremes; those who see clouds as the solution to everything, and those who see no chance or place for a cloud whatever it happens to be. In the middle are about 81–82% (the poll is ongoing, so results may vary) of the respondents, who vary from seeing a place for clouds depending on the definition or use case to others who are skeptical but want to learn more about what to use when, where, and why.

It's about delivering information services in a cost-effective manner that supports demand while meeting service objectives (Figure 1.2). Figure 1.7 shows various information services delivery models that rely on different techniques, technologies, and best practices that can be competitive or complementary. Cloud metaphor has been used in and around IT for decades as being a means to abstract underlying networking details or an applications architecture.

A key attribute of clouds is that of abstracting or masking underlying complexities while enabling agility, flexibility, efficient, and effective services delivery. This leads to some confusion, which for some creates opportunities to promote new products, protocols, standards, or services while for others it means repackaging old initiatives. For example, some may have a *déjà* moment when looking at a cloud presentation back to the late 1980s during the information utility wave that was appearing with the advent of the x86-based PCs along with client servers. For others, that moment could be time sharing or service bureau, and for others the cloud's Web-based and highly abstracted virtualized environments.

What are IT clouds? Where do they fit? How does tape coexist with clouds? Like many IT professionals, you may already be using or leveraging cloud-based computing or storage techniques, either as a product or as a service, without realizing it.




Figure 1.6 IT cloud confusion and opportunity. (Source: StorageOblog.com.)




Figure 1.7 Various information services delivery and resource models.

Common cloud-based functions or services include:

- Remote or off-site backup, replication, vaulting, or data copy
- Remote or off-site storage on-line, near-line, or off-line
- Email and messaging services including social media networking and Web 2.0
- Archive, fixed content, and other reference or look-up data
- Website, blog, video, audio, photo, and other rich media content hosting
- Application hosting (e.g., salesforce.com, concur expense, social media)
- Virtual server or virtual machine (VM) hosting (Amazon, VCE, etc.)
- General on-line storage or application-specific storage such as Google Docs

Does this mean that if backup, business continuance (BC) or disaster recovery (DR), or archive data is sent off-site to a storage or hosting facility, it has been sent to the cloud? Some say no unless the data were transmitted electronically to on-line disk at a service provider location leveraging programmatic interfaces and other cloud ware (technology, services, or protocols developed, optimized, or packaged for public and private clouds). That might also be a product- or services-based definition. However, in theory, the concept is not that far off, as clouds, in addition to being a product or service, are also a management philosophy or paradigm to do more with what you have without negatively impacting service delivery.

Characteristics of clouds include:

- Ease of service access (self-service)
- Ease of service deployment or provisioning
- Elasticity and multitenancy
- Safety, security, with data integrity

- Flexibility, scalability, and resilience
- Cost effectiveness and measurability
- Abstraction or masking of underlying complexities
- Can move or change the focus and presentation
- Leverage repeatable processes, templates, and best practices
- Efficiency as a result of scale and increased workload or usage

Confusion exists in that there are many different types of clouds, including public and private, products and services, some that use familiar interfaces or protocols with others using different technologies. Clouds can be a service or a product, an architecture, or a management paradigm, similar to previous generations such as the information utility or service-oriented architectures (SOA), client server computing, and others.

What this means is that some of the tenets of cloud storage and computing involve shifting how resources are used and managed, thus enabling the notion of an information factory. They can be external or internal, public and private, housed at a hosting or co-location facility as well as at traditional out-sourcing or managed service providers. Thus a hosting site may or may not be a cloud, and a cloud can leverage hosting services but does not require them. Various information services delivery models are shown in Table 1.1.

1.5.2. Private Cloud: Coexistence vs. Competing with Legacy IT

Clouds and virtualization should be seen for what they really are as opposed to what they are often portrayed to be so that a solution can be sold. In other words, take a step back, look at the business issue, then apply the applicable technology or task at hand to the situation. Instead of clouds being a solution looking for a problem, they become a tool and technique that can be used in different complementary ways. Cloud computing and storage are another tier of traditional computing or servers providing different performance, availability, capacity, economic, and management attributes compared to traditional technology delivery vehicles.

If IT is a core piece of the business, it probably makes more sense to retain tighter control. For example, a manufacturing company may out-source or rely on suppliers for key components, or perhaps even provide contract-under-supervision manufacturing, leveraging proprietary processes and techniques. Than if the related IT functions are also essential, they too would be retained and kept close to the vest while other functions might be out-sourced or sent to the cloud.

1.6. Common Cloud, Virtualization, and Storage Networking Questions

Does cloud storage require special hardware or software, or is it a matter of how those resources are used, deployed, and managed? As with many things, the answer is, “It

Table 1.1 Information and Data Services Deployment Models

Model	Characteristics and When to Use	Examples
Co-location ("colo")	Shared facilities with dedicated space for your equipment. Power, cooling, security, networking, and some monitoring or other optional services provided. Primary or supplemental space for your hardware.	iphouse, Rackspace, Sungard, Timewarner, visi, and many others
Hosting services	Services and or application hosting. These could be email, Web, or virtual machines. In addition, these could be Applications as a Service (AaaS). Many colos provide application hosting services. Instead of renting space for hardware, you rent time and use of software.	VCE, ADP, Amazon, Bluehost, Google, HP, IBM, iphouse, Oracle, Rackspace, Salesforce, and others
Legacy IT	Hardware (servers, storage, and networks) plus software (applications and tools) are bought or leased, operated, and managed by IT staff.	Hardware and software in your existing environment
Managed service provider	Similar if not the same as a hosting service. Provides some service, which might be applications, archiving, backup, storage space, backup, replication, email, Web, blogs, video hosting, business continuance/disaster recovery, among others. Instead of you running or hosting the application, you use a service provided to you that runs on someone else's shared infrastructure. Some may have déjà with service bureaus or time sharing.	Amazon, AT&T, Carbonite, Campaigner, EMC Mozy, GoDaddy, Google, Iron Mountain, Microsoft, Nirvanix, Seagate i365, Sungard, Terremark, Wells Fargo vSafe, among others
Out-sourcing	Could be on- or off-site, where you either move your applications and possibly equipment to a third party who operates and manages to specific service-level objectives (SLOs) and service-level agreements (SLAs).	Dell/Perot, HP/EDS, IBM, Lockheed/Martin, SunGard, Terremark, Tata, Xerox/ACS, and Wipro, among others
Private cloud	Dedicated to an organization need. Could be managed by IT staff or third party on-site or off-site. May use cloud-specific technologies or traditional technologies managed with cloudlike premises or paradigms. May be called in-source or IT 2.0.	Instrumented or metered IT environment for effective service delivery. Many different products available
Public cloud	An IT infrastructure that supports shared computing, storage, and or application services provided by an organization. The services may be available free or for a fee. They can be used to replace or complement existing IT capabilities. Access shared applications such as salesforce, email, backup or archive destination, virtual servers, and storage, among others. Buzzwords include applications as a service (AaaS), infrastructure as a service (IaaS), storage as a service (SaaS), and platform as a service (PaaS), among many other XaaS variations.	AT&T, VCE, Amazon E2C or S3, Google, Iron Mountain, Rackspace, Salesforce, Terremark, HP, and IBM, among many others