

# Data Center Storage

Cost-Effective Strategies,  
Implementation, and Management

Hubbert **Smith**



CRC Press

Taylor & Francis Group

AN AUERBACH BOOK

# **Data Center Storage**

**Cost-Effective Strategies,  
Implementation, and Management**



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

# Data Center Storage

Cost-Effective Strategies,  
Implementation, and Management

Hubbert Smith



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

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

© 2011 by Hubbert Smith  
Auerbach Publications is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works

International Standard Book Number: 978-1-4398-3487-9 (Hardback)

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) (<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 Auerbach Web site at**  
**<http://www.auerbach-publications.com>**

*Dedicated to my wonderful wife Vicki and to my wonderful  
children, Thomas, Hannah, Sasha and Leah.  
Hope you will vote for them when they run for President.*

---

# Contents

---

<b>About the author</b>	<b>xvii</b>
Acknowledgements	xvii
<b>What, Exactly, Will We Accomplish?</b>	<b>xix</b>
The One Reason Every CIO or IT Manager Should Read This	xix
At Risk of Ruining the Ending . . .	xix
Business Data	xx
<b>Part I Building Blocks, Power, and Consolidation</b>	<b>1</b>
<b>Chapter 1 The Disk Drive: The Fundamental Building Block of Enterprise Storage</b>	<b>3</b>
1.1 Using a Metrics-Based Approach to Get Past the Complexity	4
1.2 Metrics for the Basic Building Block: Hard Disk Drives and Solid State Devices	5
1.3 About HDD Capacity and Service Levels	12
1.4 Financial Responsibility and the Tradeoffs between Capacity, Form Factor, and Spin Speed	15
1.5 Demystifying Hard Disk Drive Performance	15
1.6 Hard Disk Drive Quick Guided Tour	17
1.6.1 Hard Drive Spin Speeds and Latency	20
1.6.2 Interface Speeds: SATA Speeds, SAS Speeds, FC Speeds	21
1.6.3 HDD Caching	23
1.6.4 HDD Queuing	24
1.6.5 Drive Performance Summary	25

<b>Chapter 2</b>	<b>Power and AC</b>	<b>27</b>
2.1	Powering Down Inactive Drives, or MAID	28
2.2	Facilities and Power	29
2.2.1	Facilities: Air Conditioning	30
2.3	Establishing a Baseline for Performance, Capacity, and Power	32
2.4	Facilities Cooling Improvements Within Reach	33
<b>Chapter 3</b>	<b>Storage Consolidation</b>	<b>35</b>
3.1	Cost of DAS versus SAN: Overprovisioning	37
3.2	Cost of DAS versus SAN: Backup Efficiency	40
3.3	Hot Data, Cool Data, and Consolidation, a Single-Tiered Systems Approach	47
3.4	Consolidation for Noncritical (Tier 2) Applications	50
3.5	More Silo System Avoidance	51
3.5.1	Corrective Action for Situations Where Silos Are Not Yet Deployed	52
3.5.2	Corrective Action for Situations Where the Silos are Already in Place	53
3.6	Consolidation Conclusion	56
<b>Chapter 4</b>	<b>Service Level Overview</b>	<b>59</b>
4.1	Service Level Agreements (SLAs)	60
4.2	Reviewing SLAs with Business Unit Customers	61
4.3	Delivering on SLA Commitments	63
<b>Chapter 5</b>	<b>Uptime, Reliability, and SLAs</b>	<b>67</b>
5.1	SLA Uptime: RAID Levels	68
5.2	SLA and Point in Time Copy: Avoiding the Backup Window Problem	68
5.3	SLA Uptime: Failover	72
5.4	SLA Cost and Billing	74
5.5	Performance and SLAs	75
<b>Chapter 6</b>	<b>Storage Tiering and SLAs</b>	<b>77</b>
6.1	Comparing Scenarios	78
6.1.1	Scenario 1: Single-Tier Baseline	78
6.1.2	Scenario 2: Two Tiers	79



6.1.3	Scenario 3: Two Tiers, Better Service Levels	82
6.1.4	Scenario 4: Two Tiers, Significantly Better Service	84
6.1.5	From Scenario to Reality	86
<b>Chapter 7</b>	<b>Service Level Agreements and IT Bill-Back</b>	<b>89</b>
7.1	IT Bill-Back Example of SLA	90
<b>Chapter 8</b>	<b>Demonstrating Delivery on SLA Service Levels</b>	<b>93</b>
8.1	Customer Feedback and Service Levels	93
8.2	SLA Conclusion	95
<b>Chapter 9</b>	<b>Planning for Growth and Storage Tiering</b>	<b>97</b>
9.1	Storage Growth, and SLA for Provisioning	98
<b>Chapter 10</b>	<b>Wrap-Up: Projects Within Reach</b>	<b>101</b>
10.1	Improvement 1. Improve Efficiencies with Static Tiering	101
10.2	Improvement 2. Improve Power and Air Conditioning	103
10.3	Improvement 3. Consolidate and Reduce Server-Attached Storage	103
10.4	Improvement 4. Better Backup—Replace Tape	104
10.5	Improvement 5. Establish Tiered Storage and Service Level Agreements	104
10.6	Improvement 6. Migrate Cool data onto to Tier 2, Release Tier 1 Resources for Critical Applications	105
10.7	Improvement 7. IT Bill-Back	106
10.8	Improvement 8. Track and Improve	106
10.8	Conclusion	106
<b>Part II</b>	<b>Managing Aging Data and E-Mail Expenses</b>	<b>109</b>
<b>Chapter 11</b>	<b>Migration and Retiring Aging Systems</b>	<b>111</b>
11.1	File Migration Project Planning	113
11.1.1	File Migration Plan, Discover, Test, QA, Provision, and Cut-Over	114

11.1.2	Aging Data and the Role of Near-Line Archive and Search	116
11.1.3	Block (Database, Email etc.) Application Data Migration	116
11.2	Migration Overview (Virtualized Storage)	118
<b>Chapter 12</b>	<b>Shared Folders and Content Management</b>	<b>119</b>
12.1	Content Management System Cost Analysis, Before and After	121
<b>Chapter 13</b>	<b>Storage Strategies for E-Mail</b>	<b>123</b>
13.1	E-Mail Quotas, PC Folders and PC Backup	123
13.2	E-Mail Hosting Providers	125
13.3	Checklist for E-Mail Project	125
<b>Chapter 14</b>	<b>Spending Wisely on Performance</b>	<b>127</b>
14.1	Establish Clear Performance Criteria (in SLA Terms)	127
14.2	The Yardsticks of Storage Performance	128
14.3	Resolving Performance Issues in Production Systems	129
14.3.1	Paging	129
14.3.2	I/O Size	130
14.3.3	Fragmentation	130
14.3.4	Individual Account Problems	131
14.3.5	Resource Contention	131
14.3.6	Time of Day	131
14.3.7	Storage Array Misbehavior	132
14.3.8	Storage Caching	132
14.3.9	Storage Rearrangement	134
14.3.10	Fibre Channel Zoning	134
14.3.11	Latency	135
14.4	Solid State Drives (SSDs) and Solid State Storage (SSS)	136
14.4.1	Where and How to Use SSDs	136
14.4.2	SSD Endurance, Life Cycle, and Warranty	142
14.4.3	Flash and RAM	143

<b>Chapter 15</b>	<b>Performance and Backup</b>	<b>145</b>
<b>Chapter 16</b>	<b>The Right Tools for Reliability, Uptime, Disaster Recovery, and Archiving</b>	<b>151</b>
16.1	Server Problem (Direct Attached Storage)	151
16.2	Server Problem (External SAN Storage)	152
16.3	Drive Failure	152
16.4	Accidental Data Deletion or Modification, or Virus	153
16.5	Storage Array Controller Failure	153
16.6	Cable Failure	154
16.7	SAN Switch Problem	154
16.8	Data Center Power Problem	155
16.9	Data Center Air Conditioning Problem	156
16.10	Data Center Interior Problem (Sprinklers)	157
16.11	Data Center Act of Nature	157
16.12	Spending Wisely on Reliability	158
<b>Chapter 17</b>	<b>Reliability and Server Failover</b>	<b>161</b>
17.1	How Server Failover Works	161
17.2	Load Balancing	169
<b>Chapter 18</b>	<b>Reliability and Continuous Data Protection, Synchronous Replication</b>	<b>173</b>
<b>Chapter 19</b>	<b>Reliability and Near-Continuous Data Protection, Asynchronous Replication</b>	<b>175</b>
<b>Chapter 20</b>	<b>Reliability and Data Integrity (T10-DIF or T10-PI)</b>	<b>177</b>
<b>Chapter 21</b>	<b>Virtualization Overview: Focus on the Business Benefits</b>	<b>179</b>
21.1	Virtual Machine Sprawl	181
21.2	Aligned Server Virtualization and Storage	181
<b>Chapter 22</b>	<b>Storage Virtualization</b>	<b>183</b>
22.1	Virtualized Storage and Shared Storage for Virtualized Servers	183
22.2	Thin Provisioning	185
22.3	Storage Provisioning Before and After	186

22.3.1	Thin Provisioning and Growth	187
<b>Chapter 23</b>	<b>Virtualization, Storage Tiers, and Manual Data Movement</b>	<b>189</b>
23.1	Information Lifecycle Management (ILM or HSM)	192
<b>Chapter 24</b>	<b>Virtual Desktop Infrastructure (VDI)</b>	<b>195</b>
24.1	Why VDI?	195
24.2	Implementing Virtual Desktops	196
24.2.1	Option 1: The Clone Approach	196
24.2.2	Option 2: The Snapshot Approach	196
24.3	Common Mistakes to Avoid	198
24.4	Data Security	198
<b>Chapter 25</b>	<b>Converged NAS and SAN</b>	<b>201</b>
<b>Chapter 26</b>	<b>Storage for Nontraditional IT Applications</b>	<b>205</b>
26.1	Video Surveillance	206
26.2	Audio (Call Center)	208
26.3	Web Marketing, Web Content, and Converged Communications	208
26.4	Storage Approaches for Converged Applications	210
<b>Chapter 27</b>	<b>Part II Wrap-Up: Projects Within Reach</b>	<b>211</b>
27.1	Improvement 0. Every Proposal to Spend on Expanding Tier 1 Storage Hardware Must Include a Proposal to Migrate/Archive Data off Tier 1.	212
27.2	Improvement 1. Initiate Migration Projects to Retire Old Hardware and Archive Aging Data.	212
27.3	Improvement 2. Establish Performance Profiles for Key Storage Systems, and Establish a Performance SLA.	212
27.4	Improvement 3. Establish a Knowledge-Base to Debug/Resolve Storage Performance Issues Quickly	213
27.5	Improvement 4. Pilot SSDs.	213
27.6	Improvement 5. Use a Storage Array–Based Snapshot/Backup.	213

27.7 Improvement 6. Use Server Failover and Spend Wisely on Storage Reliability.	214
27.8 Improvement 7. Create E-Mail Quotas, PC Mail Folders, and PC Network Backup.	214
27.9 Improvement 8. Replace Shared Folders with a Content Management System.	214
27.10 Improvement 9. Using Manual Storage Tiers, Move Aging Data from Tier 1 to Tier 2.	215
27.11 Improvement 10. Pilot Automated Storage Tiers.	215
27.12 Improvement 11. Use a Virtual Desktop Infrastructure.	216
27.13 Improvement 12. Establish IT Infrastructure Standardization.	216
<b>Part II Conclusions</b>	<b>217</b>
<b>Part III Managed Hosting and Cloud</b>	<b>219</b>
<b>Chapter 28 Managed Hosting and Cloud Computing</b>	<b>221</b>
<b>Chapter 29 The Business Driving Managed Hosting: What It Means to You</b>	<b>227</b>
29.1 Why Managed Hosting?	230
29.2 Data Center Definition(s) According to TIA-942	231
29.3 Working Around the Overfull Data Center	233
29.4 The Dynamic Data Center and Cash Flow	234
29.5 Is Managed Hosting Proven and Viable?	235
<b>Chapter 30 Managed Hosting Vetting Process</b>	<b>237</b>
30.1 Hosting Provider Categories	237
30.1.1 Services-Centric Hosting	237
30.1.2 Dedicated or Shared Hosting	238
30.1.3 Facility-Centric Hosting (Includes Interconnection and Co-Location)	238
30.1.4 RFQ, Selection, SLA and Risk Management	238
30.2 Implementing SLAs	242
30.3 Managing the Risk	243
30.4 Hosted Management Migration Project Planning and Execution	245

30.5 Managed Hosting Simple Cost Analysis: Before and After	246
<b>Chapter 31 Why Cloud is Relevant</b>	<b>249</b>
31.1 Origins and Blueprint of Cloud Computing	250
31.2 Cloud and the CIO	250
31.2.1 Security	251
31.2.2 Availability	252
31.2.3 Performance	252
31.2.4 Cost	253
31.2.5 Standards	253
31.2.6 Refining the Answers	254
31.3 Traditional IT versus Cloud Approaches	254
<b>Chapter 32 Implementing Cloud Storage in Your Operation</b>	<b>257</b>
32.1 Cloud-Based E-Mail	258
32.2 Cloud-Based Content Management and Collaboration	259
<b>Chapter 33 Hybrid Cloud</b>	<b>263</b>
33.1 Steps to Prepare for Cloud Technologies	263
33.1.1 Be Selective on Cloud Migration	264
33.1.2 Have Data Management Retention Policies	264
33.1.3 Have a Migration Process	264
33.2 Examining Key Norms of Traditional IT versus Cloud Approaches	265
33.3 Cloud Performance	267
<b>Chapter 34 Cloud Spectrum of Options</b>	<b>269</b>
34.1 Handling Objections and Resistance to Change	270
34.2 Trends in Public Cloud that Will Materialize in IT and Private Cloud	270

<b>Chapter 35</b>	<b>End Game and Hardware Roadmap to Leverage the Cloud</b>	<b>273</b>
<b>Chapter 36</b>	<b>Strategy and Execution</b>	<b>277</b>
<b>Chapter 37</b>	<b>Constructing a Roadmap</b>	<b>283</b>
	37.1 Managing a Roadmap	288
	37.2 Roadmap Execution, Package Phase Definition and Approvals	289
<b>Chapter 38</b>	<b>Risk Management</b>	<b>295</b>
<b>Chapter 39</b>	<b>Part III Wrap-Up: Projects Within Reach</b>	<b>297</b>
	39.1 Improvement 1. Find the Right Managed Hosting Provider for Your Business	297
	39.2 Improvement 2. Compare Managed Hosting Service Level Agreements and Managed Hosting Risk Management	297
	39.3 Improvement 3. Find the Right Cloud Storage Provider	298
	39.4 Improvement 4. Implement Cloud-Based E-mail	298
	39.5 Improvement 5. Implement Cloud-Based Content Management and Collaboration	298
	39.6 Improvement 6. Implement a Hybrid Cloud Project	298
	39.7 Improvement 7. Create a Storage Strategy and Execution Plan that Includes a ZBB, Roadmap(s), and Supporting Phase Gate Project Management/Approval Process	299
	<b>Part III Conclusion</b>	<b>300</b>
<b>Appendix A</b>	<b>Storage Protocol Basics</b>	<b>301</b>
<b>Appendix B</b>	<b>Project Management</b>	<b>305</b>
<b>Appendix C</b>	<b>People, Process, Technology</b>	<b>309</b>
	C.1 Push Decisions into the Hands of Those Most Capable	309

C.2	Develop a Culture of Peer Reviews	310
C.3	Give Credit Where Credit is Due	310
C.4	Create an Environment Where it's Okay to Ask for Help	310
C.5	Create Status Reports with a Project Management Slant	312
C.6	Encourage Communication and Require Transparency	312
C.7	Insist on Execution; Find and Resolve Dysfunctionality	313
<b>Appendix D Root Cause, Corrective Action Process</b>		<b>315</b>
<b>Appendix E Iometer: Performance Testing in the Lab</b>		<b>317</b>
E.1	Which Benchmark?	317
E.2	The Iometer Fileserver Benchmark	318
E.3	Diagrams	320
E.4	Reporting and Interpreting the Results	324
E.5	Conclusion	324
E.6	References	326
<b>Glossary</b>		<b>327</b>
<b>Index</b>		<b>337</b>



---

## About the author

---

Mr. Smith is an enterprise storage veteran with 25+ years of experience at Kodak, NCR, Intel, WDC, Samsung, and currently with LSI.

He is a published author: Serial ATA Architectures and Applications; and patent holder: USPTO 7,007,142 Network Data Storage-related operations.

Mr. Smith successfully managed 25 tech projects, negotiated and executed 15 significant technology-related business deals, was a significant contributor in establishing the Serial ATA industry standards and in establishing the Enterprise Serial ATA market.

### Acknowledgements

Theron, John, Darice, Tim ... thanks for making this happen; and happen better. You guys are the best, a joy to work with.



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

---

# What, Exactly, Will We Accomplish?

---

## The One Reason Every CIO or IT Manager Should Read This

Today we overspend on data center storage and still often fall very short of business needs for storage. This book takes a balanced “business-meets-technology” approach to delivering data center storage service levels while avoiding overspending.

Let’s expand that statement. Today, we overspend on data center storage. Business data is expanding fast, therefore our storage-related overspending is accelerating. This book reveals where and why we overspend and how to deliver same or better storage service levels, storage capacity, and data protection at a lower cost to your company.

This book is for those who are motivated to apply financial responsibility when spending on data center storage. This book provides alternatives that exceed business needs and are clear, quantifiable, and understandable by both the chief information officer (CIO) and the chief financial officer (CFO), allowing communication, understanding, agreement, and execution.

## At Risk of Ruining the Ending . . .

1. Replace direct attached storage and compartmentalized storage with consolidated storage area network (SAN)-attached storage. Consolidating direct attached storage into shared storage saves money on backup. Consolidation creates the opportunity to create tiers. It creates the opportunity to establish Service Level Agreements (SLAs), which help identify Tier 1 storage and treat it separately from Tier 2 storage. Establish the concept of storage tiering.
2. Use SLAs as a clear and simple way to communicate with your business unit customers. Establish service levels and assign them

to applications. Reserve Service Level 1 for only those applications in which performance impacts user productivity or storage performance impacts the company's ability to earn money. Push all other data to Service Level 2 and Service Level 3 in the interest of avoiding overspending for storage for noncritical applications.

3. Use time and data aging as a tool to work in your favor to save money. Push cooler data to affordable capacity-optimized storage. Implement processes to protect data and to clear data off Service Level 1 storage: Migrate data as it ages from *performance tier* to *capacity tier* to *archive*.
4. Selectively deploy to managed hosting and cloud storage. Steps 1, 2, and 3 change your storage to allow your company to take advantage of managed hosting and the cloud selectively, where it makes sense, is most efficient, and has the least risk.

Put that way, it looks easy. The more aggressively you can apply the storage tiering/service level model, the better off you'll be in your task to deliver storage services that meet and exceed business unit needs and avoid overspending.

The harder questions addressed in this book include overcoming the barriers to change—changing attitudes and behaviors to get these changes implemented in your organization so that you can deliver better quality for a cheaper price. This book includes CIO-relevant analyses to get both the funding and the staffing to implement projects to deliver on service levels, economically.

## Business Data

Business data is the lifeblood of our businesses. In data center storage, we underspend or spend wastefully by continued investment in existing systems. Or, we overspend, we overcompensate by treating all data as mission-critical, Tier 1 data; when in-fact, most data is Tier 2. When we overspend, we waste limited money which should be better applied to building a more competitive business. When we underspend, we are left with IT infrastructure which reduces our ability to execute and compete.

Enterprise storage is shifting, we know it's risky and complex to make changes, we also know standing still is often more painful than the risk and difficulties of plunging ahead. We intend to disassemble big problems

into smaller problems, which are more readily understood, quantifiable, and fixable.

It is no understatement to say IT is under attack (see Table P.1).

Table P.1 Quotes on Storage Growth and IT Expenses

		Source
Storage Growth	"Intel's storage requirements grow 35 percent per year, driven in part by the need to retain data for compliance reasons and to fulfill e-discovery requests"	Diane Bryant, CIO of Intel, at Computerworld, August 17, 2009
	IDC reports that, with a compound annual growth rate of almost 60 percent, the digital universe is projected to hit 1800 exabytes by 2011. That's a tenfold increase within five years."	Machael Friedenberg, President and CEO of CIO Magazine in CIO Magazine, September 1, 2009
	"The urge to collect is primal."	Hubbert
	"All data is NOT hot data."	Hubbert Smith
IT Expenses	"Most IT departments spend at least 50 percent of their budget on salaries, and up to 70 percent of IT staff time is spent on maintenance, according to analysts."	<a href="http://www.cloudemail101.org/should-i-go-to-the-cloud">www.cloudemail101.org/should-i-go-to-the-cloud</a> (accessed September 28, 2010)
	"Taylor Woodrow, a UK-based construction company, has switched its 1,800 employees' e-mail from Microsoft Outlook to enterprise Gmail." "The company's IT director has already seen a cost savings of nearly \$2 million by ditching on-premise e-mail."	<a href="http://www.cio.com/article/429863/Cost_Savings_Found_When_Microsoft_Outlook_Ousted_for_Gmail_at_British_Construction_Firm_">www.cio.com/article/429863/Cost_Savings_Found_When_Microsoft_Outlook_Ousted_for_Gmail_at_British_Construction_Firm_</a> (accessed September 28, 2010).

The treadmill never stops. Data center storage demand is relentless and accelerating: more users, more applications, better uptime, faster provisioning, faster transfer rates, and lower latencies. The pinch continues with lower OpX (spending for operations, mostly people) budgets for

headcount, consultants, and services; and the crunch is worsened with lower CapX (capital for equipment) budgets for equipment and data center infrastructure.

This book tackles all these issues, but its focus is on the biggest: How to deal with business operational mandates for storage infrastructure with expanded capabilities without breaking the bank and outgrowing existing data center infrastructure. At the core of this book, we are asking how to best solve (or at least ease the pain) the business appetite for data center storage, capacity growth, performance demands, uptime SLAs, and low power consumption.

These issues beg the question: What got us where we are today? Why all the pain? Why this disconnect?

Enterprise storage matured several decades ago. I was in data center storage back in the 1980s and 1990s, and the entire problem space was different: databases were small, e-mail users were few, and traffic was low. Only large businesses invested in IT departments. At that time, small businesses had some PCs, but enterprise computing was the rarified domain of big financial institutions. As enterprise storage matured and began to be a commodity, the world was still mostly paper; the Internet was still a DARPA science project. A typical enterprise storage system involved a 10 MB-to-10 GB database, or all-text e-mail, a server, a RAID controller, and a bunch of disk drives, usually 10K rpm.

Compared to enterprise applications today, those old applications were downright tiny, hugely expensive, owned by a few. The bottom line is that enterprise storage got stuck in the 1980s and 1990s, and it more or less just stayed there as the IT world changed.

Fast forward to today. Since those innocent golden years, the world has changed rapidly—normal business people do not even consider running any major component of their business on paper. In most businesses, e-mail is more important than the telephone as a communications tool. I could go on and on; but to the point: Today, enterprise applications are large, and they are many. IT organizations and those who serve IT organizations are tasked to deliver storage infrastructure that is fast enough, big enough, reliable enough, and affordable. It's that last attribute—*affordable*—that gives us pause. Using approaches from the 1980s and 1990s, we can achieve performance, capacity, and reliability if we throw enough money at the problems. But the point of this book is to deliver efficiently and affordably.

Clearly the approaches formed in the 1980s and 1990s are no longer workable or affordable.

The good news: Financial sanity can be found again, through a clarified view of the problem with a clear and simplified view of the variables: price, performance, power, and capacity. There are alternatives to the approaches matured in the 1980s and 1990s, and we're going to cover these alternatives in the next couple hundred pages.

It is my objective to thoroughly engage you, the reader, with tasks along the way that take the book from the theoretical to the real-world. We want the book framework to help you apply intelligence from your own shop into a meaningful analysis and project plan, enabling you to respond to business demand with a new and workable storage system approach with hard data, which will make sense to both CIOs and CFOs. At the end of the day, CIOs and CFOs write the checks. They'll either spend on approaches from the 1990s (increasingly uncompetitive and unsustainable), or they'll spend on the new approaches, which are workable and competitive in the 2010s and beyond.

If you walk a mile in the shoes of the CFO or CIO, you can see that they're looking for alternatives to throwing money at the Cold War—era approaches—and you can help them. To get as close as possible to your IT world, we will create example analysis exercises. I'll fill in the analysis with industry-typical values. The analyses are structured so that you can apply your real-world data within the same structure to illustrate before-and-after proposals. The desired outcomes are improvement plans that will make sense to CFOs and CIOs. They are clear and compelling enough that these executives will authorize staffing and funding for the improvements.

This book is organized into three parts:

- Part I     *Building Blocks, Power, and Consolidation* of many separate server-direct attached storage into shared network storage. Better single-tier storage through consolidation. Service levels.
- Part II    *Tiered Storage, SLAs, Managing Aging Data and E-Mail Expenses* by implementing tiered storage and service levels for individual systems. Cost-effective use of solid state storage in a Tiered+SLA operation.
- Part III   *Managed Hosting and Cloud* using tiered storage and service level approaches.

This book, by design, is organized bottom-up instead of top-down. Although we might have started with big-picture use cases, service levels, sweeping system architectures with big confusing SAN diagrams, and the

like, this is not one of those books that starts at a high level and never gets down to actionable detail. Nor is this a textbook covering information for information's sake. We are interested in getting right to the approaches that yield the quickest return with the lowest risk. Therefore, we start with the basic building block of data center storage: the hard disk drives (HDDs). HDDs are the common denominator in all data center storage systems.

Rest assured that we'll get into the use cases, service levels, systems architectures, and the like. We will definitely cover solid state devices, solid state storage, and cloud storage. But right now, we're starting from the bottom because that's where we can move the needle and get meaningful results.

The approach will hopefully engage you, as a reader, to evolve the content from discussing a theoretical system to discussing *your* system. Hope you'll enjoy the ride.



# Part I

## Building Blocks, Power, and Consolidation

In Part I, we consolidate server direct attach storage, improve on single tier storage, and discuss unrelenting demand growth for performance capacity and the expenses related to hardware price (CapX), power (OpX), management (OpX), data aging (OpX), and data center outgrowth (CapX). We demystify both the technology and the business aspects of storage performance, replication, backup/recovery/archival, and storage virtualization.

Part I establishes a baseline for single-tier storage using critical ratios for performance, capacity, cost, and power as well as data aging (these same ratios will be compared for multitier and cloud storage architectures).



Taylor & Francis

Taylor & Francis Group

<http://taylorandfrancis.com>

# Chapter 1

---

## The Disk Drive: The Fundamental Building Block of Enterprise Storage

---

All enterprise storage is based on the basic building block: the hard disk drive or HDD.

We will start with *key* metrics of the hard disk drive: *price*, *performance*, *power*, and *capacity*. We're going to identify the important stuff and remove the noise. Once we have a command of foundation and key metrics, the systems, processes, and economic delivery on service levels readily falls into place.

HDDs come in various price/performance/power/capacities and we can remove the complexity to arrive at the right tool for the job at hand. In the real world, enterprise storage is *not* about which system offers the cheapest gigabytes for the money. Enterprise storage is *not* about which system has the highest performance benchmark number. In the real world, enterprise storage is about the balance of adequate performance, adequate capacity, adequate reliability, and data protection, all at the best possible price. And this book is about disassembling the problem into manageable but related components to find that balance for the situation at hand.

Most people, including IT people, consider storage performance as the realm of PhDs: complex and difficult to understand. I am here to tell you, it is much simpler than it may appear. We can overcome the perceived complexity of storage performance by deconstructing this messy, complex problem into easy-to-understand components, focusing on the important, pertinent facts, and eliminating the nonrelevant noise.

Our objective is to remove the mystery, remove the mush, remove the fear, uncertainty, and doubt (FUD), and explore the mystery of storage performance and storage reliability. Everywhere possible, we will focus on the quantifiable and push aside the immaterial. HDDs are the fundamental building blocks of enterprise storage, so that's where we will start. We will continue on to a systems view; we'll address operational improvements; and

then we'll review new system architectures with variations (dynamic storage tiering, hybrid cloud, and private cloud, to drop a few names).

## 1.1 Using a Metrics-Based Approach to Get Past the Complexity

Ever run into one of those personalities who, when you ask the time of day, tells you how to build a clock? We're just looking for the time of day—we don't want or need to know how to build a clock. This metaphor applies well to data center storage. Storage vendors and suppliers have been known to claim leadership based on one (and usually only one) dimension of their product.

My favorite example is when a newbie enterprise solid state product marketing person claimed a high performance number (something like 100,000 IOPS, or input-output operations per second) and then crowed that to achieve that same number of IOPS would require some hundreds of 15K rpm drives. While we can relate to the notion that a vendor needs marketing sound bites—such as the world's best storage performance benchmark—we also know performance benchmarks alone are not the whole story. We know IOPS can be large or small; the workload can be random or sequential; the workload can be some mix of reads or writes. The one-dimensional “market-ecture” above, though factually correct, does not remotely resemble anything in the real world. The benchmark above assumed the smallest possible block size (512 bytes), 100 percent random workload, and 100 percent read 0 percent write workload—a situation never encountered in mainstream data storage.

- In the real world, block sizes vary, but the typical block size is 4,000 bytes (not 512 bytes).
- In the real world, the workload is sometimes random and sometimes sequential (not 100 percent random).
- In the real world, there is a mix of reads and writes; the rule of thumb is a 70:30 read:write ratio (not read-only).
- And obviously, the workload (mix of block sizes, read versus write, and random versus sequential) can vary based on the storage task/application, as well as on the time of day, week, month, quarter, or even year.

Our approach is to focus on real-world benchmarks, real-world use cases, and key components. We make a conscious effort to cull the noise, the irrelevant, and the imponderable from the equation.

We'll discover how to establish your own relevant criteria, applicable to your shop, rather than buying into those one-dimensional talking points. To be fair, to counterbalance the self-serving people; the data center storage industry has no shortage of good folks whose first instinct is to make things right with the customer. The techniques and approaches we'll cover will help you clearly identify those good folks in the industry, in contrast to the other kind.

## 1.2 Metrics for the Basic Building Block: Hard Disk Drives and Solid State Devices

For HDD building blocks, our approach is to structure a decision-making process around key metrics: price, performance, power, and capacity. As our objective is to turn raw data into useful information; we can take these four key variables (raw data) and evaluate them using key ratios (useful information), as shown in Table 1.1.

Table 1.1 Hard Disk Drive Key Ratios (Bigger is Better)

Key Ratio for Hard Drives	Ratio	Example
Performance/Cost	IOPS/\$	300 IOPS, \$200 IOPS/\$ = 1.5
Performance/Power	IOPS/watt	300 IOPS, 12 watts IOPS/watt = 25
Capacity/Cost	GB/\$	300G, \$200 GB/\$ = 2
Capacity/Power	GB/watt	300G, 12 watts GB/watt = 33

Notice the *benefit* (performance or capacity) is always in the top of the fraction (numerator), and the *expense* (cost or power) is always in the bottom of the fraction (denominator).

This way, *bigger is always better*.

The *key ratios* chart in Table 1.2 serves to simplify the total storage view. It tells us 10K drives are better in GB/\$, better in GB/watt, and better in

Table 1.2 Hard Disk Drive Key Ratios Raw Data

	10K	15K	7200 2TB	10K 400G		15K 144G		7200 2TB	
performance	300	400	120	IOP/\$	1.5	IOP/\$	1.3	IOP/\$	0.7
capacity	400	144	200	IOP/Watt	25.0	IOP/ Watt	57.1	IOP/ Watt	17.1
power	12	7	7	GB/\$	2.0	GB/\$	0.5	GB/\$	1.1
price	\$200	\$300	\$180	GB/watt	33.3	GB/ watt	20.6	GB/ watt	28.6

Sources: Vendor HDD data sheets, Nextag.com for approximate price, storager-eview.com for approximate Web server performance.

IOPS/\$; but not better in IOPS/watt than 15K rpm drives. It also serves as the underlying data for Figures 1.1 through 1.7.

Storage system engineering is sometimes (but not always) about performance, and it’s also important to see the entire picture including price, power, and capacity (Figure 1.1).




Figure 1.1 Key Metrics for 2.5" Small Form Factor 15K rpm, 146G

Clearly, the strength of this product is IOPS/watt. It’s noticeably anemic in the areas of GB/\$, GB/watt, and IOPS/\$.

The creators of this small form factor 2.5" HDD product were motivated by their IT customers to add more storage performance in over-full

data centers with limited power, limited A/C, and limited floor space (sound familiar?).

In situations where slow storage interferes with end-user productivity (and, as a result, this costs the company money), this class of performance-optimized HDD or SSD is the right tool for the job. But in situations where storage performance has a minimal impact on end-user productivity (e.g., e-mail), there are other, more financially responsible tools for the job.

Let's review the same chart for a typical 10K rpm drive (Figure 1.2).




Figure 1.2 Key Metrics for a 3.5" 10K rpm, 400G (Bigger is Better)

The 10K rpm product diagrammed above shows some balance across IOPS/\$, IOPS/watt, GB/\$, and GB/watt. This characterization is for a 3.5" drive. It consumes more power, but it also has more platter space with better total capacity, better capacity/\$, and good sequential performance. The ratios improve our shared understanding of the merits of a specific HDD (the basic building block of storage).

What does Figure 1.3 tell us? We are looking at the tradeoffs between a 10K rpm 450G drive as compared to a 15K rpm, 144G drive. In this example (no surprise) the 10K rpm drive exceeds the 15K drive in GB/watt and GB/\$. Also (no surprise) the 15K drive exceeded the 10K drive in IOPS/watt. The interesting surprise, however, is that the 10K drive exceeded the 15K drive in IOPS/\$.

So, do we conclude we should use 10K rpm drives throughout your system?




Figure 1.3 Combined Chart for Evaluation of Key Storage Ratios: IOPS/\$, IOPS/watt, GB/\$, GB/watt (Bigger is Better)

This analysis indicates we should use 10K rpm drives as the default. But, when performance is the top criteria, this analysis leads us to apply 15K rpm drives.

This is just a simple example. And the analysis gets even more interesting when we add enterprise solid state devices (SSDs).

We know some systems should be optimized for performance, and other systems should be optimized for capacity, and still other systems should be optimized for a combination of both performance and capacity. With this insight and structure, we'll be able to objectively compare and buy the right tool for the job. Later in the book, our sections on Service Level Agreements (SLAs) will map to this approach.

Sometimes we need a moving van, sometimes a sports car, right? This approach balances technology against power and cost of ownership. This metrics/ratio approach will drive closure on the question of “when is *good enough* really good enough?”

It gets better in Figure 1.4.

We could double the scale of this diagram and the GB/watt (almost 300) and GB/\$ (11) would still be off the scale.

The capacity-optimized disk drive is an incredible tool to achieve economic delivery of service levels. Capacity-optimized drives are not the right tool for situations where storage performance has an impact on user productivity and therefore costs the company money, but in almost every other



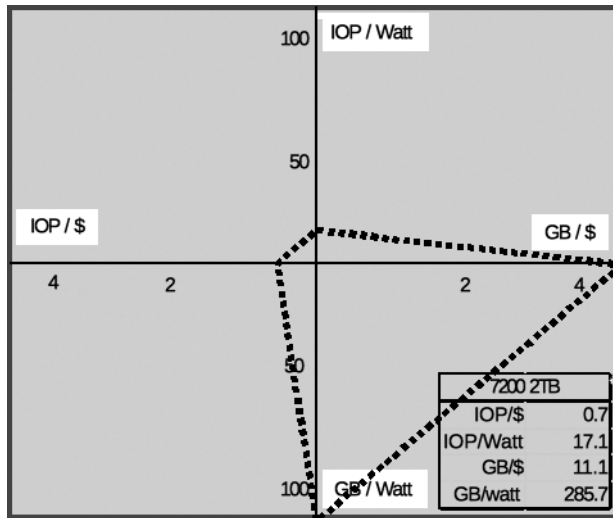


Figure 1.4 Key Metrics for a 5400 rpm, 2000G, 120 est IOPS, est \$180

instance the capacity-optimized drive is a tool that truly can save money and still get the job done.

There are data center professionals who have serious reservations regarding the reliability of high-capacity drives in the enterprise as well as regarding the use of SATA (serial advanced technology attachment) as an enterprise drive interface. It's likely these reservations are based on stale information and assumptions. Highly reliable capacity-optimized drives have been shipping for the better part of a decade. They are available in both SAS interface (for dual controller implementations) and SATA (for single controller and server-direct-attached implementations). These enterprise-class capacity-optimized drives (Raid Edition or NL-Nearline) demonstrate 1.2 million hours mean time to failure, consistent with other 10K and 15K drives.

Although there is much more to the subject than we touch on here (we will cover it in later sections on manual tiering and automated tiering), solid state devices make great sense when used in conjunction with capacity-optimized drives. SSDs make limited sense in general IT applications employing single-tiered approaches. But an approach that uses SSDs plus capacity-optimized HDDs properly, in two-tier applications, offers a significant advantage in IOPS/\$, IOPS/watt, GB/\$, and GB/watt over any single-tier storage system (see Figures 1.5, 1.6, and 1.7).

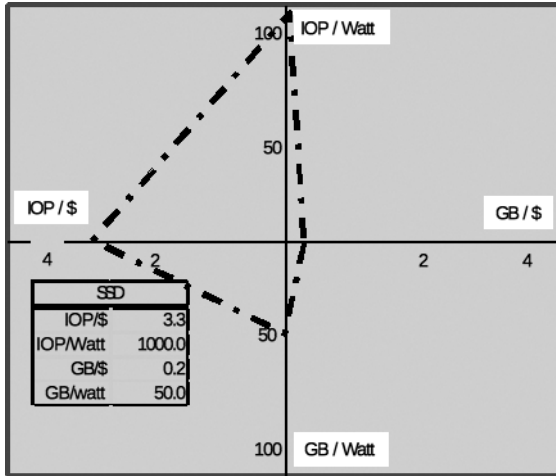


Figure 1.5 Key Metrics Comparison for SSD; Assume 100GB, 2000 IOPS, \$600 (Bigger is Better)

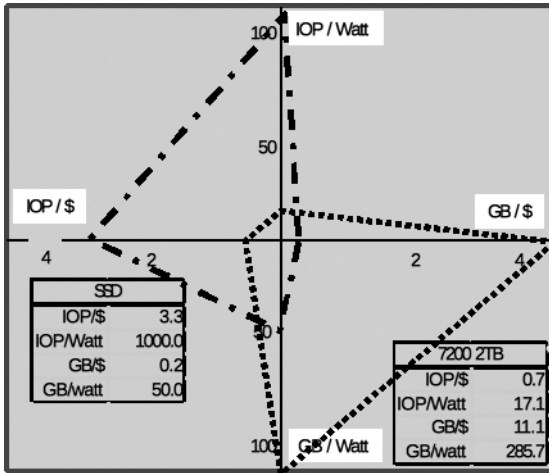


Figure 1.6 Key Metrics Comparison for SSD with Capacity-Optimized 7200 rpm Drives (Bigger is Better)

Notice the storage device classes that are strongest: capacity-optimized and SSD. Everything else is a compromise. For the upcoming sections, we will walk before we run, so I will not mention SSDs for the next few sections, but we'll cover SSDs and where and why they make financial sense later in the book.

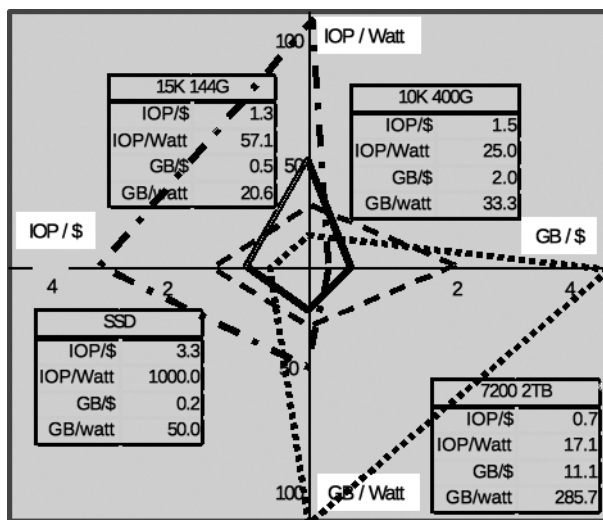


Figure 1.7 Key Metrics Comparison for 15K, 10K, 5400 rpm, and SSD

The bottom line is that these products (high-capacity HDD, 15K rpm HDD, SSDs) align to storage service levels (Tier 2 and Tier 1).

When these technologies, plus people, plus processes, are intelligently applied to deliver on service levels and manage aging data from tier-to-tier-to-tape archival, the operational savings and capital savings are compelling.

Table 1.3 shows the underlying data driving the charts in this chapter.

Table 1.3 Raw Data for Key Metrics Comparison for 15K, 10K, 5400 rpm, and SSD

	10K	15K	7200 2TB	SSD	10K 400G	15K 144G	7200 2TB	SSD
performance	300	400	120	2,000	IOP/\$ 1.5	IOP/\$ 1.3	IOP/\$ 0.7	IOP/\$ 3.3
capacity	400	144	2,000	100	IOP/Watt 25.0	IOP/Watt 57.1	IOP/Watt 17.1	IOP/Watt 1000.0
power	12	7	7	2	GB/\$ 2.0	GB/\$ 0.5	GB/\$ 11.1	GB/\$ 0.2
price	\$200	\$300	\$180	\$600	GB/watt 33.3	GB/watt 20.6	GB/watt 285.7	GB/watt 50.0

Source: Vendor data sheets for typical drives, typical performance data (Web server IOPS) from storagereview.com, typical prices from nextag.com

My point is that considering the ratios of IOPS/\$, IOPS/watt, GB/\$, and GB/watt enables us to avoid getting tangled in information spaghetti. Using key ratios, we trade confusion for clarity as we compare one *class of drives* to another *class of drives*. New HDD products will emerge with improved capacities, improved performance, improved pricing.

I hope that makes sense, and that we can declare “confusion avoided” instead of falling victim to analysis paralysis or stalling our investigation.

A side note on the drives we've examined: You may read this and think, "He said 400 GB, didn't he mean 450 GB?" At the time I put this section together the 400G 10K rpm drive was based on four platters of 100 GB each. Along the way, the platter density changed from 100 GB to 150 GB per platter. Now we see a 450 GB 10K rpm product, not a 400 GB product. That's the nature of the HDD industry.

The 450 GB product is based around three platters of 150 GB each (significantly less expensive to produce; its higher bit density offers higher sequential performance). The raw data will change quickly; it's the ratios that are the main event. Ratios turn data into information.

### 1.3 About HDD Capacity and Service Levels

The march of disk drive capacity increases over the years is a truly amazing story. Every year, the drives increase in capacity (often doubling year after year). Higher capacity drives basically increase the number of bits on the platters. Two good things result: more capacity for the same power, and improved performance, with bits passing under the head at a faster rate. The incentives for data centers to migrate from many smaller drives onto fewer but higher capacity drives include saving on the power bill, saving drive slots, saving data center power system capacity, saving data center A/C system capacity, and reducing the risk of outgrowing existing data centers. (See Table 1.3.)

When you implement, be conscious of the capacity versus performance trade-off: Higher spin-speed, smaller form-factor drives are great in performance-optimized systems.

For situations where performance is the driving concern, it's financially appropriate to use small form-factor 15K rpm drives. The cost of the storage is small as compared to the cost of the user time and productivity. For stock trading and similar Service Level 1 applications, that's a justifiable use of the company's money. But 15K rpm small form-factor drives are not the right tool for every job.

The other side of that coin: It is irresponsible and wasteful to apply 15K rpm small form-factor drives in a situation where storage system performance has low impact on end-user time and productivity (e.g., Service Level 2 or Level 3 applications).

10K drives are popular for this exact reason: For single-tier storage systems that require a compromise between capacity and performance, 3.5" 10K drives have "good enough" performance and better capacity than 15K drives or small form factor HDD. (See Table 1.4.)

Table 1.4 The Tradeoffs Between Form-Factor, Spin Speed, and Capacity

Drive Type	Common Capacities	
15K rpm 2.5" performance	146G, 73G	Note the tradeoff: smaller form factor = much smaller capacity
15K rpm 3.5" performance	600G, 450G, 300G	
10K rpm 2.5" performance	300G, 146G	Note the tradeoff: smaller form factor = much smaller capacity
10K rpm 3.5" performance	400G, 300G	
7200 rpm 3.5" capacity	2TB, 1TB, 750G, 500G	Note the tradeoff: slower spin speed = larger capacity
5400 rpm 3.5" capacity	2TB, 1TB	

Source: HDD data sheets

The idea is to apply the right technology to the right problem:

- 15K rpm small form factor drives should be used only for Service Level 1: high-performance applications
- 3.5" larger capacity drives should be used for Service Level 2: active storage applications
- 3.5" capacity optimized drives should be used for Service Level 3: seldom-accessed data

Data center storage tiering and its associated service levels will receive thorough coverage throughout the book.

Also notable is the fact that drive capacities have grown over the years. There is money to be saved by replacing older, lower-capacity drives with similar up-to-date higher-capacity drives. At one time, a 146 G 15K drive was state-of-the-art; no longer. Now we have 400 G and 600 G 15K drives, and expect larger ones still. That's quadruple the capacity for the same slot and same power footprint. Upgrading drives in an older system is a straightforward and obvious way to save money on power and free up slot capacity in your existing single-tier data center storage.