

ORACLE®

Oracle Press™

Enterprise Grid Computing with Oracle

ORACLE®

Oracle Press™

Enterprise Grid Computing with Oracle

Brajesh Goyal
Shilpa Lawande

McGraw-Hill

New York Chicago San Francisco
Lisbon London Madrid Mexico City Milan
New Delhi San Juan Seoul Singapore Sydney Toronto

McGraw-Hill books are available at special quantity discounts to use as premiums and sales promotions, or for use in corporate training programs. For more information, please write to the Director of Special Sales, Professional Publishing, McGraw-Hill, Two Penn Plaza, New York, NY 10121-2298. Or contact your local bookstore.

Enterprise Grid Computing with Oracle

Copyright © 2006 by The McGraw-Hill Companies, Inc. (Publisher). All rights reserved. Printed in the United States of America. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of Publisher.

Oracle is a registered trademark of Oracle Corporation and/or its affiliates. Screen displays of copyrighted Oracle software programs have been reproduced herein with the permission of Oracle Corporation and/or its affiliates.

1234567890 DOC DOC 019876

ISBN 0-07-226280-X

Sponsoring Editor

Lisa McClain

Editorial Supervisor

Janet Walden

Project Editor

Mark Karmendy

Acquisitions Coordinator

Alex McDonald

Technical Editors

Paul Strong, Robert Mello, Anil Khilani,
and Bob Thome

Copy Editor

Mark Karmendy

Proofreader

Paul Tyler

Indexer

Claire Splan

Production Supervisor

George Anderson

Composition

Apollo Publishing Services

Illustrators

Apollo Publishing Services

Series Design

Jani Beckwith, Peter F. Hancik

Information has been obtained by Publisher from sources believed to be reliable. However, because of the possibility of human or mechanical error by our sources, Publisher, or others, Publisher does not guarantee the accuracy, adequacy, or completeness of any information included in this work and is not responsible for any errors or omissions or the results obtained from the use of such information.

Oracle Corporation does not make any representations or warranties as to the accuracy, adequacy, or completeness of any information contained in this Work, and is not responsible for any errors or omissions.

**To my grandmother,
the late Jhammubai Goyal,
and my grandfather-in-law,
Dr. Moolchand Mittal**

—*Brajesh*

To family and friends

—*Shilpa*

About the Authors

Brajesh Goyal is the senior manager for the application integration program office at Network Appliance, where he is driving Network Appliance's efforts for integrating NetApp technologies into grid efforts of various enterprise vendors. Brajesh has co-authored a short book on enterprise grid computing. Before NetApp, Brajesh was the principal product manager for grid computing at Oracle, where he was among the first few people to explore grid computing and influenced Oracle's strategy and efforts in the grid space. Brajesh has an M.S. in Computer Science from University of Minnesota, Minneapolis, and a B.Tech degree in Computer Science from Indian Institute of Technology, Mumbai.

Shilpa Lawande is an expert in database and data warehousing technologies, and has several years of experience with Oracle products. She is currently with Vertica Systems, developing a high-performance database solution for data warehouses running on grid infrastructure. Previously, at Oracle, she worked on data warehousing and self-managing technologies in the Oracle database. Shilpa has co-authored two books on Oracle—*Oracle 9iR2 Data Warehousing* (Digital Press, 2003) and *Oracle 10g Data Warehousing* (Digital Press, 2004). She has co-authored a short book on enterprise grid computing jointly with Brajesh Goyal. Shilpa has an M.S. in Computer Science from University of Wisconsin-Madison and a B.Tech degree in Computer Science from Indian Institute of Technology, Mumbai.



Contents at a Glance

1 Motivation for Enterprise Grid Computing	1
2 The Evolution of Enterprise Grid Computing	29
3 Grid-enabling the Data Center	59
4 Storage Grid	83
5 Server Grid	127
6 Application Grid	169
7 Information Grid	211
8 Software Provisioning	245
9 Grid Management	269
10 Enterprise Grid Security	305
11 Business Continuity with an Enterprise Grid	329
12 Steps to Grid Adoption	345
Index	371



Contents

FOREWORD	xvii
ACKNOWLEDGMENTS	xxi
INTRODUCTION	xxiii
1 Motivation for Enterprise Grid Computing	1
Business Demands Today	3
Flexibility and Responsiveness to Business Demands	3
Better Value for Money	4
Enterprise IT Architecture Today	5
Hardware Infrastructure	5
Software Infrastructure	7
Business Application and Process Infrastructure	9
Operational Elements	11
Enterprise Grid Computing Defined	12
The Enterprise Grid Computing Model	13
Enterprise Grid Computing at Different Layers of IT	15
Centralized Management	18
Taking Strategic Steps Towards an Enterprise Grid	19
Benefits of Enterprise Grid Computing	20
Lower Costs of IT	20
Flexibility at Every Layer of the IT Stack	22
Predictable Quality of Service	22
What Does Enterprise Grid Computing Mean to You?	23
CEOs	23
Lines of Business	23
IT Executives	24

IT Staff	25
The Oracle Customer	25
Summary	26
References	27
2 The Evolution of Enterprise Grid Computing	29
Grid Classification and Evolution	30
Grid for Scientific and Technical Computing	31
Enterprise Grid	33
Grid Across Partner Organizations	35
Computing as a Utility	36
Grid Definitions and Related Terms	37
What Is Grid Computing?	38
Industry Terms for Grid Computing	40
Terminology Relevant to Enterprise Grid Computing	42
Industry Trends	43
Hardware Trends	43
Software Trends	44
Enterprise Trends	45
Regulatory Trends	46
Role of Standards in Enterprise Grid Computing	46
Enterprise Grid Alliance (EGA)	47
Goals	47
Relevance to Enterprise Grid Computing	47
EGA Working Groups	47
Global Grid Forum (GGF)	48
Goals	48
Relevance to Enterprise Grid Computing	48
Open Grid Services Architecture (OGSA) v1.0	49
Future of GGF and EGA	49
World Wide Web Consortium (W3C)	50
Relevance to Enterprise Grid Computing	50
Organization for the Advancement of Structured Information Standards (OASIS)	51
Relevance to Enterprise Grid Computing	51
Distributed Management Task Force (DMTF)	52
Relevance to Enterprise Grid Computing	53
Storage Networking Industry Association (SNIA)	53
Relevance to Enterprise Grid Computing	53

What Do Standards Mean for Your Enterprise Today?	54
Summary	55
References	55
3 Grid-enabling the Data Center	59
The “IT as a Service” Model	60
The Consumer Perspective on IT	60
The Supplier Perspective on IT	61
Enterprise Grid Alliance (EGA) Reference Model	63
Enterprise Grid Overview	64
Enterprise Grid Components	64
The Grid Management Entity	68
Applying the EGA Reference Model to Oracle Environments	77
Enterprise Applications as Grid Components	77
Storage Grid Components	78
Server Grid Components	79
Oracle Grid Control as a GME Component	79
Application Grid Components	80
Information Grid Components	81
Summary	81
References	82
4 Storage Grid	83
Enterprise Storage Infrastructure	84
Operational Requirements from Storage Infrastructure	85
Challenges with Enterprise Storage Infrastructure Today	89
Islands of Storage	89
High Cost of Storage Infrastructure	93
Inflexible Storage Infrastructure	93
Complexity of Storage Management	95
Storage Evolution Towards Enterprise Grid Computing	97
Storage Arrays	97
Storage Networking Technologies	99
Storage Virtualization Technologies	102
Storage Provisioning Technologies	108
Centralized Storage Management	111
Storage Grid Management and Provisioning Process	114
Understanding the Existing Storage Infrastructure	114
Understanding the Enterprise Storage Needs	115

Service-level Management for Storage Infrastructure	119
Monitoring Utilization and Capacity Planning	120
Storage Standards	121
SMI	121
Data Management Forum	122
Future of Storage Grid	122
Summary	124
References	125
5 Server Grid	127
Enterprise Server Infrastructure	128
Requirements from Server Infrastructure	130
Enterprise Server Infrastructure Today	132
Problems with the Silo Architecture	133
Enterprise Grid Computing Model for Server Infrastructure	136
Advances in Server Hardware	138
Low-cost Modular Servers	139
Commodity Server Clusters	141
Virtualization	144
The Value of Virtualization	144
Virtualization Overview	145
Hardware Virtualization	145
Virtualization of Infrastructure Software	149
Dynamic Server Provisioning and Resource Management	152
Resource Provisioning for the Oracle Database	153
Resource Provisioning for Oracle Application Server	156
Server Grid Management	157
Pooling of Server Resources	158
Exploiting Server Hardware and Software Features	158
Using an Integrated Management Tool	159
Role of Administrators in a Server Grid	160
Service-level Management	161
Server Management Standards	163
CIM	164
ASF	164
SMASH	164
DCML	165
Future Directions	165
New Licensing Models	165
Server Interoperability	165

Multicore Processors	166
Unified Networking for Servers and Storage	166
Summary	166
References	167
6 Application Grid	169
Enterprise Applications and Business Processes Today	170
Business Processes	171
A Closer Look at Enterprise Application Integration	173
Evolution Towards Application Grids	178
Service-oriented Architecture	179
Life Cycle of a Service	183
Core Technology and Standards for Implementing an SOA	186
Benefits of Deploying SOA on Enterprise Grid Infrastructure	194
Implementing SOA in	
Oracle Environments	195
Developing Services	196
Orchestrating Business Processes	199
Monitoring and Optimizing the Business Process	205
Managing Services	205
SOA Challenges	207
Summary	209
References	210
7 Information Grid	211
The Business Need for an Information Grid	213
Operational Information Integration	213
Regulatory Compliance	214
Information as an Enabler for Business	215
Challenges in Information Integration	215
Disparate Sources of Information	216
Information Fragmentation	217
Traditional Approaches to Information Integration	218
What Is Missing?	220
An Information Grid	221
Why Call It a Grid?	221
Evolution Towards an Information Grid	224
Unifying Similar Information Sources	224

Information Integration Spanning	
Databases and Applications	228
Enterprise Search	230
Semantic Information Integration	231
Implementing an Information Grid with Oracle Technologies	232
Data Integration	232
XML DB and XML DB Repository	236
Master Data Management with Oracle Data Hub	238
Enterprise Collaboration Technologies	238
Secure Enterprise Search	239
Information Grid Standards Activities	240
Summary	242
References	243
8 Software Provisioning	245
Need for Enterprise Software	246
Enterprise Software Deployment Today	247
The Software-deployment Life Cycle	247
Challenges in Software Deployment Today	247
Simplifying Software Provisioning with a Grid	251
The Software-deployment Process in a Grid	252
Software Provisioning Using Oracle Grid Control	255
Managing the Software-deployment Life Cycle	256
Configuration Tracking and Change Management	262
Summary	266
References	266
9 Grid Management	269
Role of Enterprise IT Management	270
Typical Enterprise IT Management Functions	271
New Requirements on Enterprise IT Management Today	272
IT Management Today	274
Managing Complexity	275
Reactive Service-level Management	276
No Single Point of Control	277
Evolution Towards Enterprise Grid Management	278
Standardization	278
Consolidation and Centralized Management	280
Service-centric Management	282
Automation	284

Grid Management Using Oracle Grid Control	285
Policy-based Standardization	286
Managing Groups	289
Service-level Measurement and Diagnostics	290
Centralized Management of the Extended Oracle Environment	295
Grid Management Standards	299
DMTF Common Information Model (CIM)	300
Management Interfaces	302
Summary	302
References	303
10 Enterprise Grid Security	305
Overview of Enterprise Security Management	306
System Security	307
Information Security	307
Application Layer Security	308
Identity Management	308
Issues with Enterprise Security Today	309
Physical Isolation	309
Focus on Vulnerabilities Rather than Business Risk	310
Hard-coded Security Policies Within Applications	311
Too Many Identities	311
Security in an Enterprise Grid	312
What Does Grid Bring to Security?	312
IT Infrastructure Security in the EGA Model	314
Information Grid Security	317
Application Grid Security	317
Identity Management	318
Security Standardization Efforts	320
Enterprise Grid Security with Oracle	322
IT Infrastructure Security with Oracle Grid Control	322
Oracle Database Security	323
Oracle Identity Management Solutions	324
Summary	327
References	327
11 Business Continuity with an Enterprise Grid	329
Need for Business Continuity in the Enterprise	330
Protection Against Downtime	330

Protection Against Data Loss	332
Disaster Recovery (DR)	333
Issues with Business Continuity Solutions Today	333
Complexity of Management	333
High Costs of Availability	334
Business Continuity in an Enterprise Grid	334
Benefits of a Grid Architecture	334
Service-centric Business Continuity	337
Business Continuity for an Enterprise Grid	340
Summary	343
References	344
12 Steps to Grid Adoption	345
Transitioning to the Grid	346
Barriers to Grid Adoption	346
Implementing Change	349
Identifying a Sweet Spot	352
Scaling Out	353
Strategic Steps Towards an Enterprise Grid	354
Standardization	354
Consolidation	355
Automation	356
Tactical Steps	357
Define	357
Measure	358
Analyze and Improve	360
Control	363
Financial Quantification of Grid Benefits	
in Oracle Environments	364
Infrastructure Cost Savings	364
Infrastructure Management Cost Savings	367
Revenue Impact and Productivity Gains	368
Reduced Risk Exposure	368
Summary	369
References	370
Index	371

Foreword

An invasion of armies can be resisted, but not an idea whose time has come.
—Victor Hugo, 1852

Grid computing is an idea whose time has come.

It has come not as a dramatic flash of brilliance or as a mind-bending leap, but as part of a progression in computing that appears so natural, so obvious, that in retrospect the innovation in all of it can almost go overlooked.

Innovation in IT has been driven by largely the same goals for the last 50 years—the drive for increased automation, greater efficiency, and the flexibility for systems to accommodate change. Along with other important companies, Oracle has helped drive innovation through past generations of computing and continues to lead as businesses transition to grid computing.

Before diving into the heart of this book to learn the workings of grid computing and the practice of it with Oracle software, let us revisit briefly the innovations that shaped historical generations of computing. For it is out of this context that grid computing emerges as the fifth major era in IT, and it is from within this reality that companies must evolve.

The mainframe generation of computing, led by IBM in the 1960s, automated core operations of the world's most important businesses and set a high bar for the performance, scalability, and reliability expected for enterprise information technology. The innovations in throughput, parallelism, and high availability of the mainframe remain an inspiration today.

One major limitation of mainframe computing was the extremely high cost of big iron. For example, a typical mainframe from the popular IBM 7090/94 series cost \$3,134,500, or about \$18 million in today's dollars. Building applications for the mainframe was also highly specialized and took far too long. This led to an “applications backlog”—unrealized innovations in business automation functionality that needed a new foundation in order to be unleashed.

A new wave of innovations from the late 1970s and into the 1980s defined the minicomputer generation. Dramatically lower-cost hardware, driven by smaller and more powerful CPU and memory chips and new computer architectures, along with simpler, interactive operating systems, introduced a new type of platform for deploying business applications.

At the same time, software innovations such as the C programming language, the Unix operating system, and the relational database management system began to emerge. In 1970, IBM's Ted Codd published the landmark paper, "A Relational Model of Data for Large Shared Data Banks." Oracle's Larry Ellison and his early partners seized the opportunity to offer a commercial RDBMS. The relational database, along with so-called "fourth generation" (4GL) tools, facilitated rapid application development, and enabled, for the first time, end users to query databases and generate reports without application programming at all.

The lower cost of computers and the more flexible and usable tools for application creation also empowered departments to solve their own automation problems, thus creating the dawn of departmental computing. Together, these innovations helped address the pent-up demand for new applications, and expanded access to computing and automation, which spawned further innovation in business applications.

The next era of business computing, the client-server generation, was enabled by the advent of the PC in the early 1980s. As Moore's Law continued to reflect the pace of innovation in processor and memory chips through the 1980's, PCs became accessible to consumers and indispensable to professionals in all business domains. Rich graphical user interfaces became the norm for business applications, and individuals were empowered with tremendous, easy-to-use computing power at their fingertips.

As networking standards, particularly TCP/IP, emerged, it became possible to connect the PCs to server machines, creating the client-server model of application architecture. This architecture differed from previous approaches by separating the data management function, which ran on shared servers, from the business logic and user interface processing, which ran on the PC. Connecting PCs via networks to server machines running relational databases provided for rich user interaction with the application, while ensuring the integrity of the shared data.

While PCs delivered powerful new tools to users, the proliferation of PCs within enterprises led to new problems. Networking limitations made it difficult to achieve adequate performance and scalability. Maintaining all those desktop and laptop computers, installing and upgrading systems and application software, and tracking and protecting thousands of new corporate assets introduced significant variable costs to computing that dwarfed the relatively cheap computers themselves. The cost of labor in managing PCs, servers, and software became recognized as a separate variable in the costs of computing and one that could quickly overwhelm other expenses.

The Internet computing model delivers on some of the unfulfilled hopes and high costs of the client-server era. In this approach, the end user interacts with a web browser, which is responsible only for formatting the display. In contrast to client-server, no business logic runs on the desktop. Dedicated servers, easily managed in the data center, take over the job of processing user interactions and running the application logic. The multitier, network-centric architectures of the Internet computing era reduce the costs of management, while enabling greater adaptability than do client-server architectures.

One barrier that slowed adoption of Internet computing was the need to rewrite applications, many of which had only recently undergone the move to client-server computing. While the Internet and the Web were all the rage for savvy consumers in the mid-1990s, it was not obvious at the time that enterprise applications would also benefit from the new model. Oracle was one of the first software companies to whole-heartedly embrace the Internet age, promoting the notion of thin clients, developing Internet-centric application development tools, and rewriting all its enterprise applications for the Web.

This history of innovation and the migration through four major eras of computing brings us to where we are today, in the midst of a rapid evolution to grid computing. The innovation behind grid computing, including virtualization of every layer of the computing stack, automated provisioning and workload balancing, and centralized management over distributed and diverse grid components, is no less remarkable than the innovation that drove change in past computing generations. With grid computing, however, unlike the other eras, applications do not require rewriting to take advantage of the benefits of grid computing. For the first time in computing history, existing applications from a past generation can continue unchanged into the new generation, automatically reaping many of the rewards.

And the rewards are considerable. In the Internet computing era, a company's assets, including storage and servers, and intellectual assets such as code, were dedicated to a task. These dedicated silos reduced the ability to reuse and redeploy resources and increased the costs of management. With grid computing, in contrast, all types of resources—e.g., storage, processing, development, management, information—can be more fully utilized. All types of resources—e.g., servers, application logic, data elements—have the flexibility to be combined in new ways to solve new problems as needs change. Furthermore, it is possible to achieve mainframe-level performance and reliability by configuring a number of relatively inexpensive and small commodity servers to form a grid. Or, to borrow Oracle's description of the grid, it "runs faster, costs less, and never breaks."

XX Enterprise Grid Computing with Oracle

The new concepts of grid computing have migrated quickly from academia to industry, and businesses are moving rapidly to deploy grids. Whereas nearly a decade elapsed between Ted Codd's seminal paper on relational databases (1970) and Oracle's launch of the first commercial relational database (1979), it was only five years after Ian Foster's first grid blueprint book that Oracle 10g—the first database and application server platform for the grid—was introduced. Less than a year later the industry came together to form the Enterprise Grid Alliance. By 2008, the technologies will be widely deployed.

In the following pages, you will develop a rich conceptual understanding of grid computing, uncover how these grid technologies really work, and learn how to transition to grid computing in your own operations.

Approximately 50 years after the dawn of the mainframe computing era, grid computing now delivers equivalent qualities of service, better experiences for users and developers, and the flexibility to adapt to future changes, and does it at a dramatically lower cost.

Surely, this is an idea whose time has come.

The best thing about the Grid is that it is unstoppable.

—The Economist, June 21, 2001

Acknowledgments



Writing this book has been a long but fruitful process for us; however, we could not have done it without the help and support of many people along the way. We would like to acknowledge the extremely important but often silent contribution of the reviewers who provided us critical feedback and technical guidance for the material. We are extremely grateful to these people for taking valuable time out of their busy schedules and providing thorough and extremely critical review:

- Bob Thome, Senior Manager, Distributed Database Product Management at Oracle.
- Paul Strong, currently a Distinguished Engineer at eBay Inc., previously a systems architect for N1 at Sun Microsystems and chair of both the Technical Steering Committee and the Reference Model Working Group of the Enterprise Grid Alliance.
- Rob Mello, Enterprise Software Solutions Consultant at Palladium Group and responsible for building successful sales engineering teams at technology startups BrainTree Security Software and Relicore.
- Dave Pearson, a Grid expert and Senior Director of Product Marketing for Grid Computing at Oracle Corporation.
- Tom Rarich, Storage industry veteran and Senior Product Manager for Product V-Series at Network Appliance.
- Paul Feresten, Storage industry veteran, creator of the StorageWorks brand as vice president at Digital Equipment Corp, and presently a senior product marketing manager at Network Appliance.
- Erik Peterson, a high availability and scalability solutions expert, currently with Oracle Real Application Clusters development at Oracle Corporation.
- Miranda Nash, a Grid expert and Director of Development for Enterprise Information Integration solutions at Oracle Corporation.

- Anil Khilani, Principal Product Manager for Oracle Grid Control at Oracle Corporation.
- Pavna Jain, Sr. Director for Oracle Fusion Middleware at Oracle Corporation.
- Amit Jasuja, Vice President of Development for Security and Identity Management Products at Oracle Corporation.
- Pankhuri Agrawal, budding anthropologist, architect, and social scientist; a woman of many worlds.

Thanks also to Anik Khilani, Dave Shaffer, Alex Peattie, and Kevin Finn for providing screenshots for the book.

We would like to extend our gratitude to Miranda Nash and Ken Jacobs for writing the foreword to the book.

Our thanks go to Lisa McLain and Alex McDonald for coordinating this project on behalf of McGraw Hill/Oracle Press. We would like to thank Mark Karmendy and the entire production staff at McGraw Hill for the excellent quality and ease of the production process.

Last but not least, we would like to say a big thank you to our spouses Meghna Mittal and Nitin Sonawane and to our families who were a constant source of encouragement and support throughout this project.

Introduction



This book introduces the reader to the concept of enterprise grid computing. This book is written to be a high-level blueprint for evolution towards enterprise grid computing covering the entire IT stack. We have endeavored to make the book easy to read and as self-contained as possible. However, due to the breadth of topics covered and the technology choices available to the end user, we are unable to provide a nuts-and-bolts guide for implementing each solution. We do include a plethora of references and sources for further information, which together with this book, should provide readers with the ammunition they need to tackle the transition to Grid. While this book primarily focuses on Oracle environments, the general concepts and best practices discussed herein apply equally well to all IT environments.

The primary audience for this book would be CIOs and Information Technology (IT) professionals who wish to get a comprehensive overview of this emerging concept, while at the same time, understand its practical benefits and current limitations. CEOs and other executives focused on the enterprise's bottom line may be interested in the high-level overview of enterprise grid computing provided in Chapter 1 and the tactical and strategic implications and the ROI model discussed in Chapter 12. This book may also be used as a textbook on enterprise grid computing.

The structure of the chapters is as follows:

- Chapter 1 provides an introduction to the subject of enterprise grid computing and its benefits to enterprise IT. We also delve into the problems facing enterprises today and how they are solved by enterprise grid computing.
- Chapter 2 provides a brief history of grid computing in general and how it has evolved into enterprise grid computing. We provide an overview of related terminology, industry trends, and standards organizations involved in this effort.
- Chapter 3 introduces the concept of a grid-enabled data center and presents a reference implementation based on the Enterprise Grid Alliance (EGA) Reference model. Here we set the stage for the remaining chapters.

Chapters 4 through 11 each deal with a specific area of the IT stack. In each chapter, we describe what grid computing involves, present technology choices and processes for evolution towards a grid, as well as give an overview of standards activities in that area.

- Chapter 4 deals with Enterprise Storage infrastructure and emerging technologies in Storage Arrays, Storage Networking, Storage Virtualization and Provisioning, and Storage Management space that make it possible to implement a cost-effective storage grid today.
- Chapter 5 delves into the next layer of the IT stack, namely Enterprise Server infrastructure. We discuss technologies such as low-cost modular servers, clustering, virtualization, and server provisioning and resource management, which make a server grid a reality today.
- Chapter 6 takes the reader up into the enterprise applications domain. We show the relationship between Service-Oriented Architectures (SOA) and Grid and how the two can work together to create a very flexible and agile application and business process infrastructure.
- Chapter 7 takes on the still nascent topic of Information Grid, which aims to bring together all enterprise information scattered around various disparate information sources into a common umbrella.
- Chapter 8 talks about software provisioning processes and how they are improved with an enterprise grid.
- Chapter 9 discusses the topic of grid management and how it differs from today's IT management. We discuss how enterprises can evolve from today's IT management of managing complexity to grid management of managing pools of resources.
- Chapter 10 treats enterprise security from a grid perspective, highlighting the benefits and challenges brought about by a grid environment.
- Chapter 11 discusses the important topic of business continuity and how a grid-enabled data center is in fact better prepared for an eventuality that shares multiple redundant components to serve high availability and disaster recovery needs of groups of applications.
- Chapter 12 outlines some tactical and strategic steps enterprises can take to migrate to the grid in an incremental manner, while addressing the pressing needs of an enterprise. We provide examples of pilot projects that enterprises can undertake. We also present a financial ROI model for the grid.