



Linux Handbook

A Guide to IBM Linux Solutions and Resources

Discover the IBM Software solutions for Linux-based servers

Understand how Linux operates on each IBM **@server** platform

Learn about IBM resources for implementing Linux solutions



Nick Harris
Ralph Cooley
Cameron Hunt
Randy Kuseske
Dan Lacine
Tomomi Takada
Bob Waite
Dirk Webbeler
Alexander Zaretsky

ibm.com/redbooks

Redbooks



International Technical Support Organization

**Linux Handbook: A Guide to IBM
Linux Solutions and Resources**

April 2003

Note: Before using this information and the product it supports, read the information in “Notices” on page xi.

First Edition (April 2003)

© Copyright International Business Machines Corporation 2003. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP
Schedule Contract with IBM Corp.

Contents

Notices	xi
Trademarks	xii
Preface	xv
The team that wrote this redbook	xv
Become a published author	xix
Comments welcome	xx
Forward	xxi
Chapter 1. Introduction to Linux	1
1.1 The foundation for Linux	2
1.2 The IBM commitment to Linux	3
1.2.1 IBM's 'penguin' projects	4
1.3 Linux: The operating system	9
1.3.1 Linux BIOS management	11
1.3.2 Where Linux fits in	12
1.3.3 Working with other operating systems	12
1.4 Infrastructure	13
1.5 Data storage	13
1.6 Middleware	14
1.7 Application solutions	16
1.8 Clients	17
1.9 The Linux distributions	18
1.9.1 What is a distribution	19
1.9.2 What is common among the distributions	21
1.9.3 What is open source	21
1.10 Red Hat	23
1.10.1 What's in the distribution?	23
1.11 UnitedLinux	25
1.11.1 What's in the distribution?	25
1.11.2 The Linux Standards Base	27
1.12 Conectiva	28
1.12.1 What's in the distribution?	28
1.13 SCO Open Linux	29
1.13.1 What's in the distribution?	29
1.14 SuSE	37
1.14.1 What's in the distribution?	37
1.15 Turbolinux	39

1.15.1	What's in the distribution?	39
1.16	Other players	41
1.16.1	Yellow Dog Linux	42
1.16.2	Redflag Linux	43
1.16.3	Debian Linux	44
1.17	Going nuts with kernels	44
Chapter 2. Open source software		47
2.1	Creating the code and protecting it	48
2.1.1	Source code	48
2.1.2	Software licenses	49
2.2	The origin of open source	51
2.2.1	The Free Software Foundation	51
2.2.2	The compromise of open source software	52
2.3	Linux and open source	54
2.3.1	The GNU utilities	55
2.3.2	The problem of complexity	55
2.3.3	The benevolent dictator	56
2.4	IBM and open source	57
2.4.1	IBM and open source licenses	57
2.4.2	IBM and open source software	57
Chapter 3. Linux, UNIX, and Windows systems		59
3.1	Features, functions, and technologies	60
3.1.1	Description of the ratings	61
3.1.2	Methodology for the feature, function, and technology rating system	61
3.1.3	Linux, UNIX, and Windows: A side-by-side comparison	62
3.1.4	Linux, UNIX, and Windows: Strengths and weaknesses	64
3.2	Cost of ownership	65
3.2.1	Total cost of ownership	65
3.2.2	Operating environment and product TCO perspective	67
3.2.3	Findings	68
3.2.4	TCO versus ROI	68
Chapter 4. zSeries and Linux		71
4.1	S/390 and zSeries architecture	72
4.1.1	Ways to run Linux on zSeries	74
4.1.2	31-bit and 64-bit options	75
4.2	All about z/VM	75
4.2.1	Integrated Facility for Linux	78
4.2.2	Hardware requirements for z/VM	78
4.2.3	Virtual networking with z/VM	79
4.2.4	VM Guest LAN support	80
4.2.5	z/VM for Linux performance: Scalability	82

4.2.6	Productivity: Development and test workloads	82
4.2.7	Productivity: Production workloads	83
4.2.8	Systems management: Data backup with Tivoli Storage Manager and FlashCopy	85
4.2.9	Consolidation example: Server farm in a box	85
4.2.10	Consolidation example: Data and application servers on a single zSeries server	86
4.2.11	Horizontal growth: Adding another server	87
4.3	zSeries Linux distributions	89
4.4	zSeries customers and Linux	89
4.5	Linux application solutions	90
4.5.1	Mail serving scenario	92
4.5.2	News group serving scenario	95
4.5.3	File/print serving with Samba	99
4.5.4	IBM Software for Linux on zSeries	102
4.6	How IBM customers are using Linux	108
4.6.1	T-Com	109
4.6.2	Boscov's Department Stores	109
4.6.3	zSeries services solutions	111
Chapter 5. xSeries and Linux		115
5.1	Linux and xSeries	116
5.1.1	Linux on xSeries	117
5.1.2	xSeries Linux distributions	120
5.2	Intel-based server customers and Linux	121
5.3	Consolidation	122
5.4	Linux clusters	126
5.4.1	High-availability cluster	126
5.4.2	Load-balancing cluster	127
5.4.3	High-performance computing	128
5.5	Linux in a distributed enterprise	129
5.5.1	Basic implementation strategies	130
5.5.2	Business benefits	131
5.5.3	Customer example: E*TRADE Financial	131
5.6	Linux application solutions	132
5.6.1	IBM Software for Linux on xSeries	132
5.6.2	ISVs and application solutions	133
5.6.3	Customer scenario: Journyx	142
5.6.4	Customer scenario: Flamenco Networks	143
5.7	Infrastructure solutions	144
5.7.1	Customer scenario: weather.com	148
5.8	xSeries architected solutions	149
5.8.1	Blade server	149

5.8.2 Customer scenario: Turning Stone Casino	150
5.9 xSeries services solutions	151
5.9.1 Implementing Linux	152
5.9.2 Support	153
Chapter 6. pSeries and Linux	155
6.1 How Linux runs on the pSeries server	156
6.1.1 Linux-capable pSeries model structure	156
6.1.2 pSeries operating system scenarios	157
6.1.3 AIX 5L and Linux	158
6.1.4 Options to run Linux on the pSeries server	161
6.1.5 Linux and its relationship with AIX	164
6.2 Linux for pSeries hardware enablement	164
6.2.1 Linux for pSeries scalability	165
6.2.2 Linux for pSeries's Reliability, Availability, and Serviceability	166
6.2.3 Other sources of information	166
6.3 Consolidation	167
6.4 Linux for pSeries clusters	168
6.5 Linux for pSeries application solutions	169
6.5.1 IBM Software for Linux on pSeries	169
6.5.2 Other Linux for pSeries software applications	175
6.6 Linux for pSeries infrastructure and architected solutions	176
Chapter 7. iSeries and Linux	177
7.1 How Linux runs on the iSeries server	178
7.1.1 Why Linux on iSeries	179
7.1.2 Linux and LPAR	181
7.1.3 iSeries Linux integration	183
7.1.4 Linux on iSeries models	184
7.1.5 iSeries Linux distributions	188
7.2 Consolidation	189
7.3 Linux clusters	190
7.4 Performance	190
7.5 Linux in a distributed enterprise	191
7.6 Linux application solutions	193
7.6.1 Small office example using OpenOffice	195
7.6.2 Java	195
7.6.3 IBM Software Group products	195
7.7 Infrastructure solutions	196
7.8 iSeries solutions	196
7.9 iSeries service solutions	198
Chapter 8. IBM Software for Linux	199
8.1 Overview of Linux and IBM Software availability	200

8.2 IBM WebSphere brand products on Linux	202
8.2.1 WebSphere product availability on Linux	204
8.2.2 Competitive software.	204
8.2.3 Linux-related WebSphere resources.	206
8.3 IBM Data Management products on Linux	208
8.3.1 DB2 Product Family	208
8.3.2 Informix product family	211
8.3.3 U2 product family	212
8.3.4 Competitive software.	212
8.3.5 IBM Data Management resources	213
8.4 IBM Tivoli products on Linux	215
8.4.1 Tivoli product availability on Linux.	217
8.4.2 Competitive software.	218
8.4.3 Tivoli resources.	219
8.5 Lotus products on Linux	220
8.5.1 Lotus product availability on Linux	220
8.5.2 Competitive software.	221
8.5.3 Lotus resources.	222
8.6 Additional IBM Software	223
8.6.1 Product availability on Linux	223
8.6.2 Additional resources	224
8.7 IBM Integrated Platform for e-business.	224
8.8 IBM Software on Linux deployment examples	225
8.8.1 Server consolidation using WebSphere Application Server on the zSeries server.	225
8.8.2 Computer services using Domino and WebSphere on xSeries and iSeries.	226
8.8.3 DB2 UDB, DB2 Connect, and WebSphere Portal.	227
8.9 IBM and Linux solution customer success	228
8.10 Additional IBM Software for Linux resources	229
8.10.1 Evaluation and demonstration software	229
8.10.2 IBM Software on Linux training	230
8.10.3 IBM Redbooks.	230
8.11 Internal IBM Software for Linux resources	231
8.11.1 Linux sales resources	231
8.11.2 Product information.	233
8.11.3 Technical resources	234
8.11.4 Evaluation and demonstration software	237
8.11.5 IBM Software support organizations.	238
Chapter 9. Linux clusters	243
9.1 An introduction to clusters.	244
9.1.1 Core cluster concepts	244

9.1.2 IBM cluster offerings	252
9.2 Linux cluster components	254
9.2.1 Nodes	254
9.2.2 Networks	257
9.2.3 Support software	261
9.2.4 Applications	262
9.3 Linux cluster history	264
9.4 Linux cluster solutions	268
9.4.1 Cluster 1350	268
9.4.2 ClusterProven®	268
9.4.3 Solution examples	269
9.5 The future of Linux clusters	271
9.5.1 Emerging trends	271
9.5.2 The role of Linux clusters in emerging trends	272
Chapter 10. IBM Services for Linux	275
10.1 IBM Services	276
10.2 IBM Global Services	276
10.2.1 IBM Support Line	276
10.3 Linux consultancy	279
10.4 IBM Services and Software	281
10.4.1 WebSphere	282
10.4.2 DB2 for Linux	283
10.4.3 Tivoli	284
10.4.4 Lotus	287
10.5 Platform services	287
10.5.1 zSeries services	288
10.5.2 iSeries services	289
10.5.3 pSeries services	290
10.5.4 xSeries services	290
10.5.5 Consolidated platform matrix	291
10.6 Examples of Linux services	291
10.6.1 Example #1: A billion dollar financial institution	291
10.6.2 Example #2: A multi-million dollar advertising company	292
10.6.3 Example #3: One of the largest network services providers	292
Chapter 11. Desktop Linux	293
11.1 Linux desktop environments	294
11.1.1 KDE and GNOME	295
11.1.2 StarOffice and Open Office	298
11.2 Windows: Linux hybrids	299
11.2.1 A look at Lindows	299
11.2.2 CrossOver: Running Windows programs on Linux	300

11.2.3 Wine ‘is not an emulator’	301
11.2.4 Trying Linux without installing it	301
11.3 Desktop deployment options	302
11.3.1 Thin client or PDA	303
11.3.2 Fat client: Quality desktop	304
11.3.3 Inexpensive client	304
11.3.4 Packaging and updating Linux	305
Related publications	307
IBM Redbooks	307
Other publications	308
Referenced Web sites	308
How to get IBM Redbooks	309
IBM Redbooks collections	309
Index	311

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:
IBM Director of Licensing, IBM Corporation, North Castle Drive Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrates programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Trademarks

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

@server™
Redbooks (logo) ™
alphaWorks®
ibm.com®
iNotes™
iSeries™
pSeries™
xSeries®
z/Architecture™
z/OS™
z/VM®
zSeries®
AFS®
AIX 5L™
AIX®
APL2®
AS/400®
BladeCenter™
Blue Gene™
BookManager®
Chipkill™
ClusterProven®
CICS®
Domino™
DB2 Connect™
DB2 Universal Database™
DB2®
Enterprise Storage Server®

Everyplace®
ESCON®
FlashCopy®
FICON™
Informix®
Intelligent Miner™
IBM®
IMS™
LearningSpace®
Lotus Notes®
Lotus Workflow™
Lotus®
LANDP®
Multiprise®
MQSeries®
MVS™
Netfinity®
NetView®
NetVista™
Notes®
OS/2®
OS/390®
OS/400®
Parallel Sysplex®
PowerPC®
Predictive Failure Analysis®
PAL®
Red Brick™

Redbooks™
RS/6000®
S/370™
S/390®
Sametime®
SecureWay®
Sequent®
ServerProven®
ServeRAID™
SmartSuite®
SmoothStart™
System/360™
System/370™
System/390®
SANergy™
SP™
Tivoli Enterprise™
Tivoli Enterprise Console®
Tivoli®
TotalStorage™
UniData™
ViaVoice®
VisualAge®
VM/ESA®
VSE/ESA™
WebSphere®
X-Architecture™

The following terms are trademarks of other companies:

ActionMedia, LANDesk, MMX, Pentium and ProShare are trademarks of Intel Corporation in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

C-bus is a trademark of Corollary, Inc. in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

SET, SET Secure Electronic Transaction, and the SET Logo are trademarks owned by SET Secure Electronic Transaction LLC.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Other company, product, and service names may be trademarks or service marks of others.

Preface

IBM® has a strong commitment to Linux as an operating system and Linux-based solutions. This IBM Redbook gives a broad understanding of IBM products and how they relate to the Linux operating system. It tries to encapsulate the IBM Software Solutions that are available for Linux-based servers. It also describes the Linux enablement throughout the IBM @server product line.

This IBM Redbook offers enough information to give you a high-level understanding of Linux. Therefore, it is not meant to be a highly technical document. It also provides pointers to further sources of information.

This redbook is intended for IBM representatives, Business Partners, and Customers who are involved in considering and planning Linux-based software solutions. You will find this redbook particularly helpful if you are planning to test or implement Linux in an enterprise.

The team that wrote this redbook

This redbook was produced by a team of specialists from around the world working at the International Technical Support Organization (ITSO), Rochester Center.



Nick Harris is a Senior Systems Specialist for IBM and has spent the last three years in the ITSO, Rochester Center. He specializes in LPAR, iSeries™ hardware and software, external disk, Integrated xSeries® Server for iSeries, and Linux. He also writes and teaches IBM classes worldwide on areas of AS/400® and iSeries system design and server consolidation. He spent 13 years in the United Kingdom AS/400 Business and has experience in S/36, S/38, AS/400, and iSeries servers. His e-mail address is <mailto:niharris@us.ibm.com>



Dirk Webbeler is an I/T Specialist in Germany. He has five years of experience in iSeries field and more than 10 years in the I/T business. Working in sales and technical support, his areas of expertise include TCP/IP, local area network (LAN) and wide area network (WAN) planning, Windows NT and 2000, server consolidation, implementation of Management Information Systems (MIS), and database design. He has written extensively on MIS, dynamic programming, and Windows NT/2000 Server consolidation. His e-mail address is <mailto:Dirk.Webbeler@dw-solutions.de>



Cameron Hunt is an IT consultant and educator in the United States. He has 10 years of experience in the IT field and has spent the last three years focusing on Linux. His areas of expertise include Linux, Microsoft, and Novell system administration. He has written extensively on computer security.




Randy Kuseske is currently assigned to Server Group Market Intelligence at IBM Rochester. He has 21 years of experience in the IT industry, previously working at Unisys, EMC, and AiC. His areas of focus have included operating system internals, compiler and language technology, globalization, and product competitive analysis. He currently performs product competitive analysis on various vendor environments, products, and technologies for iSeries, IBM @server, and IBM Market Intelligence initiatives. His e-mail address is <mailto:random1@us.ibm.com>



Dan Lacine is an Advisory Education Specialist at the IBM Rochester Support Center. He provided remote support for 21 years ranging from System/390® software to iSeries Communications and Client Access. He currently develops and teaches classes to IBM personnel on iSeries products such as iSeries Navigator and iSeries Access for Web, and PC-based operating systems. He co-authored a redbook titled *Make Your AS/400 Year 2000 Ready*, SG24-5401, and has written articles for *AS/400 Magazine* and the *UK Assist/400* technical newsletter.



Tomomi Takada is an IT Specialist at the Linux Support Center of IBM Japan. He has three years of experience in the Linux field. He has provided technical support for the public sector and been in charge of xSeries. His areas of expertise include Linux, xSeries, security, and systems management for Intel-based servers. His e-mail address is <mailto:e30810@jp.ibm.com>



Alexander Zaretsky is an IT Support Specialist at the IBM Andean pSeries™ Software Support Center in Bogotá, Colombia. He has more than ten years of experience in the UNIX system and C programming fields and five years as an AIX® and Linux System Administrator. His areas of expertise include programming and development in AIX and Linux environments, Tivoli® product support, and network configuration and security. He holds a degree in computer systems engineering from The Colombian School of Engineering. His e-mail address is <mailto:alexzare@co.ibm.com>

Ralph Cooley is a Linux Sales Specialist at IBM and has been an IBM employee since IBM's purchase of Sequent® Computer Systems in 1999. He promotes Linux and the IBM Linux initiative to the public and to IBM prospects, customers, and partners. Ralph has been a Linux user for four years. He has 25 years of UNIX market experience as a system administrator, application developer, kernel developer, instructor, consultant, and salesperson. Ralph holds a Bachelor of Science in ocean engineering and has earned a Masters in Business Administration. His e-mail address is <mailto:cooleyr@us.ibm.com>

Bob Waite is an Advisory Software Engineer. He has worked at IBM Rochester for 24 years writing product documentation for User Technologies. Currently, he is the documentation team leader for Linux and Java on iSeries. Previously, he was the documentation team leader for WebSphere®, San Francisco, and other object technologies. His e-mail address is <mailto:rwaite@us.ibm.com>

Thanks to the following people for their contributions to this project:

Scott Knupp
Worldwide Program Director, Linux Sales

Whei-jen Chen, Data Management
Rufus Credle, xSeries and WebSphere
Greg Geiselhart, zSeries®
Pete Kovari, WebSphere
Edson Manoel, Clusters
Keigo Matsubara, pSeries
David Watts, xSeries
IBM ITSO

Ed Gauthier
Susan Greenlee
zSeries

Joe Banas
Richard C Ferri
Egan Ford
John Hawkins
Norm Snyder
Mark Solomon
Clusters

Jay Ashford
John Bissell
pSeries

Dave Amundson
Craig Johnson
iSeries

Carol Carson
Amy Freeman
Mike Keleher
xSeries

Jakob Carstensen
IBM Software Group

Bob Butler
Ron Houston
Yohichiroh Ishizak
Dr. Harald Neumann
Joseph Rhoden
IBM LDP Alliances

Nick Carr
Red Hat

Greg Anderson
The SCO Group

Takashi Nishioka
Turbolinux

Martina Krahmer
SuSE

Tony Ravinsky
IBM Global Services

We also thank the many people who reviewed this book in its draft stages and provided feedback to IBM Redbooks™. With this first edition, we look forward to more valuable input from you in the future.

Become a published author

Join us for a two- to six-week residency program! Help write an IBM Redbook dealing with specific products or solutions, while getting hands-on experience with leading-edge technologies. You'll team with IBM technical professionals, Business Partners and/or customers.

Your efforts will help increase product acceptance and customer satisfaction. As a bonus, you'll develop a network of contacts in IBM development labs, and increase your productivity and marketability.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Comments welcome

Your comments are important to us!

We want our Redbooks to be as helpful as possible. Send us your comments about this or other Redbooks in one of the following ways:

- ▶ Use the online **Contact us** review redbook form found at:

ibm.com/redbooks

- ▶ Send your comments in an Internet note to:

redbook@us.ibm.com

- ▶ Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. JLU Building 107-2
3605 Highway 52N
Rochester, Minnesota 55901-7829

Forward

IBM embraced the potential of Linux several years ago with an expectation that Linux would provide unprecedented choice, value and flexibility for our customers and partners. Today, this broad investment in Linux across IBM's business is reflected in a vast number of IBM product offerings and solutions available in the marketplace. As this publication goes to print, Linux is the fastest growing server operating system in the industry.

In this handbook, we set out to describe in tangible terms, the breadth and depth of IBM's commitment to Linux. At the same time, we seek to provide a solid background for business and technical management as well as for technical leaders who are seriously considering Linux for implementing a new class of enterprise solutions. The chapters which follow are a result of the significant and enduring investments that IBM has made in Linux through the open source development community, as well as within IBM's own world-wide development labs, services organizations, and partner support programs. The Linux ecosystem is large, diverse, and growing rapidly.

Undoubtedly, the content which follows will continue to change as Linux matures and corporations accelerate their adoption of Linux-based solutions. For that reason, where possible, we included Web sites in an effort to enhance the currency of the material. I sincerely hope that this handbook provides you with both the knowledge and confidence to implement your next projects with an IBM Linux solution!

Dr. D. Scott Knupp

Worldwide Program Director
Linux Sales



Introduction to Linux

This chapter examines the components and features of Linux, its application and middleware support. It reviews Linux capabilities in both the server and client environments. Finally, it look at what comprises a distribution and highlights the major features of the most popular ones.

1.1 The foundation for Linux

Linux is a UNIX-like open-source operating system. It was the original creation of Linus Torvalds from Helsinki, Finland, in 1991. He wrote the first kernel, the underlying program interfacing and running the computer hardware.

Torvalds invited programmers from around the world to comment on and to improve his code. This is one of the key ideas behind the success of Linux. With the world as the laboratory, the number of testers and developers is nearly endless. It is because of this resource that Linux is constantly evolving and improving.

Since the Linux source code is freely available, several companies have developed different *distributions* of Linux. A distribution is a complete system. The key component is the Linux *kernel*. Other utilities, services, and various applications can be included as well, depending on the distribution and the intended use. There is no standard distribution. Each distribution that is available has unique advantages. IBM was early to recognize the value of Linux:

- ▶ Investing in Linux-related product development
- ▶ Forming alliances with key Linux distributors
- ▶ Contributing to the open-source community
- ▶ Aggressively supporting the platform

IBM believes this investment will benefit its customers as they continue to exploit Linux for their IT infrastructures and e-business.

Figure 1-1 shows the Linux ecosystem chart in which IBM is a key component of life in the Linux world.

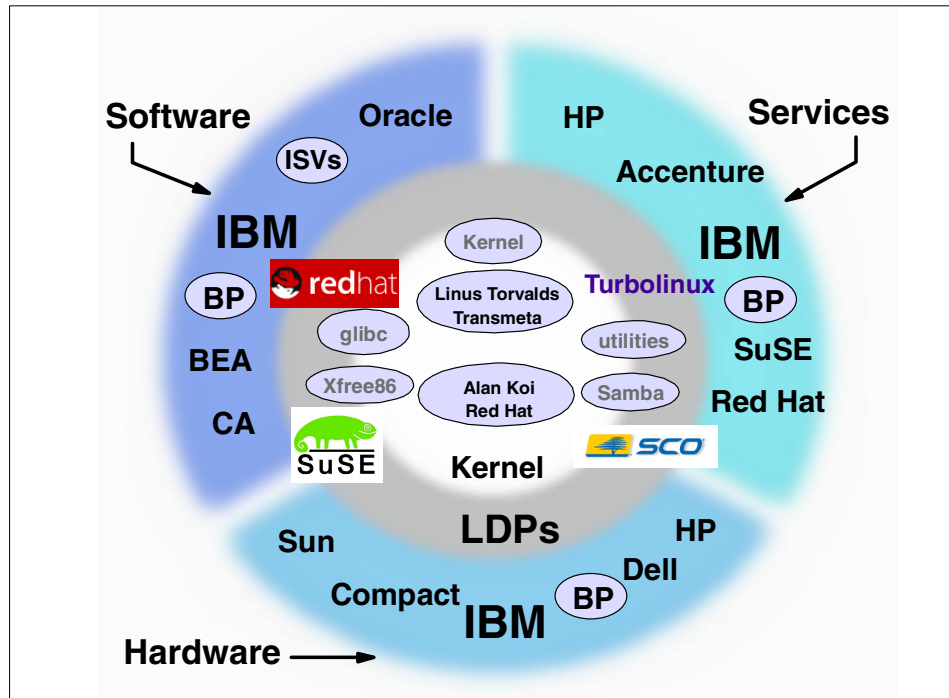


Figure 1-1 The Linux ecosystem

1.2 The IBM commitment to Linux

IBM is fully committed to the open source movement and believes that Linux will emerge as a key platform for e-business. IBM will work with the open-source community. They will leverage relevant technologies and experience to help enhance Linux, to define the standards, and to extend Linux to the enterprise level.

IBM provides continued support and participation throughout the world. You can learn more about this support and IBM's commitment to Linux at the following Web sites:

- ▶ The IBM Linux portal for a general point of entry into IBM and Linux
<http://www.ibm.com/linux>
- ▶ IBM Linux Technology Center (LTC)
<http://www.ibm.com/linux/ltc>

- ▶ IBM Solution Partnership Centers
<http://www.developer.ibm.com/spc/index.html>
- ▶ IBM Linux support line
http://www.ibm.com/services/e-business/linux_8.html
- ▶ The Open Source Development Lab
<http://www.osdl.org/>

As part of this continuing commitment, IBM has teamed with leading commercial Linux distributors, including The SCO Group (previously known as Caldera Systems), Red Hat, SuSE, and Turbolinux. This cooperation allows IBM to port, test, and certify the performance of its offerings running on various Linux distributions. The result is to enable you to exploit the full potential of Linux.

1.2.1 IBM's 'penguin' projects

At IBM, there is a real thrust to adopt Linux technology where it makes sense. The following sections highlight some of the internal IBM projects that use Linux.

Web infrastructure

Development of IBM Web infrastructure is expanding with the use of Linux in the following projects:

- ▶ Linux portal
- ▶ Web content management system
- ▶ Advanced search engine
- ▶ W3 e-workplace development environment
- ▶ Intranet forums
- ▶ W3 e-workplace special events
- ▶ Hone/IBM link Web portal (planned)

Linux portal

The external IBM Linux portal www.ibm.com/linux was deployed on two redundant, load balanced xSeries servers running Apache Server on Red Hat Linux. The software that is running is:

- ▶ IBM HTTP Server
- ▶ TSM Client and Tivoli Monitoring Agent (TMA)
- ▶ Server Resource Management (SRM)
- ▶ Enterprise Security Manager (ESM)

Web content management system

The IBM Linux portal team developed a content management system. This system uses Linux to facilitate the development, delivery, and tracking of the

content for the www.ibm.com/linux site. This includes a backend DB2® server running on Linux on zSeries hardware.

Advanced search engine

The W3 advanced search facility required a more robust and efficient search capability. Inktomi's search engine was selected to run on Intel Linux. The search engine is now hosted on four symmetric multiprocessing (SMP) xSeries servers running:

- ▶ Linux and TSM Client
- ▶ Tivoli Monitoring Agent (TMA)
- ▶ Server Resource Management (SRM)
- ▶ Enterprise Security Manager (ESM)

W3 e-workplace development environment

IBM's own intranet application (W3) development team is developing all of their applications on development and test environments on zSeries VM (virtual machine) guests running Linux.

Intranet forums

This project replaced the VM Forums infrastructure with two Linux logical partitions (LPARs) on each of the two zSeries servers that are accessible from the IBM W3 intranet. IBM used VM Forums for over 20 years.

This collaboration system provides service to over 350,000 IBMers worldwide. It is now a repository of news groups with over 1,000 topics that have around 15,000 new posts every month.

W3 e-workplace special events

Linux hosts *W3 jam sessions*. *Manager Jam* was hosted on zSeries hardware and ran the event on five VM virtual servers running SuSE Linux. These servers ran IBM HTTP Server and WebSphere Application Server 3.5, and one of the images ran DB2 and IBM HTTP Server.

Jamming is the art of strategic brainstorming across a corporation. It's a technique IBM Research began using about a year ago. The term originates from the jazz and blues music concept of a jam session, which is the ultimate gathering of expert musicians to create something great. No one knows what will be produced during a jam session, only that it will be something that's never been heard before.

In a way, this is what IBM jam sessions are about. We gather experts, give them a scenario to ponder or a problem to solve, and then set them loose. Those experts, like a group of great musicians, then create a series of "riffs" or great ideas that may not surface through more traditional problem-solving processes.

Hone/IBM Link Web Portal (planned)

The portal Web site for EMEA's customer order support functions was implemented on a Linux S/390® WebSphere server. This configuration includes Enterprise Access Enabler to VM legacy applications and databases. The functions include inventory retrieval, order submission, and price inquiry. After the successful pilot phase, this project is in development for full production.

Security

Security is always an important project area. The following security projects are Linux based:

- ▶ Security assessments
- ▶ Virus detection
- ▶ Storage Architecture Security Directory
- ▶ E-mail antivirus scanners

Security assessments

IBM Global Services (IGS) Internet Vulnerability Scanning Services are handled by 67 Intel boxes running Linux. Through this configuration, IGS provides internal and commercial Internet and intranet vulnerability scanning services. IBM serves over 100 commercial customers and scans 30,000 internal IP addresses per week through this service. Linux has used these services for over five years.

Virus detection

Linux servers have been running "honey pot" and scanning software for security purposes to detect and track attempted attacks and viruses. One server is up and running in each of the 13 Asia Pacific countries (or regions).

Honey pot: This computer system on the Internet is expressly set up to attract and "trap" people who attempt to penetrate other people's computer systems.

E-mail antivirus scanners

This project is designed to deploy antivirus scanning software into the IBM mail infrastructure. This project involves scanning both incoming and outgoing mail for viruses. Trend Micro Interscan VirusWall is the antivirus software that was selected. It will run on Linux and be deployed on 18 xSeries machines worldwide.

Monitoring

System monitoring projects include:

- ▶ Performance monitoring
- ▶ Asset monitoring
- ▶ Operations for e-hosting and network management

Performance monitoring

Linux supports IBM's End-to-End Probe Platform (EPP), a Java-based monitoring tool primarily built to run on the Windows NT platform but that also runs on Linux. EPP measures the performance and availability of applications from an end user's perspective. The performance and availability metrics are generated by periodically executing and measuring the duration of emulated transactions. Each probe server instance can poll approximately 180 servers per hour using the default timeout values.

There are 17 EPP servers running Linux, deployed in 13 different countries in EMEA. There are six in the U.S.A. and four in Canada. Linux is an ideal operating system for this kind of application because it is reliable and economical. Also its multitasking capabilities provide four times the throughput over the Windows NT solution. This allows IBM to use 75% fewer servers for the same workload.

Asset monitoring

The workstation asset management tools Asset Center, Brio Reporting, and Communication Director track and manage workstation hardware and software assets. These tools run on a Linux image that is running as a guest under z/VM® on S/390 hardware. There is also a DB2 database on the S/390 Linux for Software Inventory.

Operations for e-hosting and network management

Linux is used to monitor over 1,600 commercial URLs hosted by IGS. It is also used for remote console and cluster management. IBM uses Linux to monitor over 2,300 network interfaces.

File and print

The file and print projects include:

- ▶ File serving
- ▶ File and print servers

File serving

IGS is using a network of Linux file servers to provide standard client images worldwide via IBM Standard Client Installer (ISCI). It also provides the IBM Standard Software Installer (ISSI) application and maintenance fix packages to end users. There are currently 156 ISCI/ISSI servers in production in 28 countries (or regions) worldwide.

File and print servers

Australia and New Zealand converted 66 file and print servers previously on OS/2®, Windows 95, and Windows NT to 44 xSeries servers running Linux and one RS/6000® box running AIX 5.1L across 13 sites.

Follow-on projects are probable in other Asia Pacific countries (or regions). This project included a Linux domain controller for each site as well.

IBM manufacturing and development facilities

The projects for manufacturing and development include:

- ▶ Microelectronics 300mm wafer manufacturing and test
- ▶ Electronic design automation
- ▶ VLSI test engineering
- ▶ Manufacturing line kiosks
- ▶ Software development
- ▶ IBM Internal Open Source Bazar

Microelectronics 300mm wafer manufacturing and test

The future 300mm wafer chip manufacturing line tool controllers from IBM's Microelectronics Division run on xSeries machines running Linux. They were slated to be deployed with Windows 2000, but a major reliability concern was detected.

This project demonstrated that Linux was much more reliable in this environment. The chip manufacturing line tool controllers currently run on 174 xSeries boxes running Linux. Two hundred to 300 installations are planned across both the East Fishkill and Burlington, New York, sites. After almost a full year of production support, there have been no Linux-related outages.

Electronic design automation

IBM Microelectronics Division embarked on a multiyear strategy to add Linux as a key element of their design automation platform. This involves a long-term migration of design workstations and development infrastructure to Linux, as well as tools and applications porting.

To date, already 80% of the tools have been ported, with a year-end completion goal. This division is also developing a developer/designer workstation environment. This is based on a vision for a potential for productivity gains, performance benefits, and functional benefits of Linux-based tools.

VLSI test engineering

The microelectronics chip manufacturing and test engineering organization ported their test data generation programs to Linux. This enabled the use of two xSeries servers in the process of generating device test patterns and development of programs for the chip testers in the manufacturing process.

Manufacturing line kiosks

IBM Canada's Bromont manufacturing facility installed 81 kiosks to access Web applications, legacy host applications, and printers using old Pentium 133/166 boxes destined for retirement.

Software development

Over 400 VM guests and 15 LPARs on the S/390 are being used for product development of DB2 and WebSphere Commerce Suite products.

IBM Internal Open Source Bazaar

The IBM Internal Open Source Bazaar is an internally deployed open source project development portal to promote open source style development inside IBM. It is an opportunity for IBM developers to “share” their code with other IBM developers around the world. It is hosted on a Linux-based server. It is a service that offers easy access to the best in CVS, mailing lists, bug tracking, message boards/forums, task management, site hosting, permanent file archival, full backups, and total Web-based administration.

IBM Research facilities

The following projects are examples of where IBM Research facilities are using Linux-based servers.

Blue Gene™ protein folding simulation

Over 110 Linux servers are being used as compute and file servers in this advanced super computing project. This project focuses on the modeling of the life science process of protein folding.

Oceano Web hosting utility

The IBM Research Division is working on a hosting utility that can dynamically adjust capacity in response to workload across multiple commercial customers. They are using a clustered set of Linux servers that can be dynamically reconfigured between customers. There are 75 Intellistations running Linux.

1.3 Linux: The operating system

Over the past few years, the Linux operating system has become a real and viable alternative for PC users as well as corporate servers and users. Linux delivers the power and flexibility of a UNIX server or desktop. It also provides a set of utilities, Internet applications, and a fully functional desktop interface.

The Linux operating system has become a server platform for very powerful Internet applications. The Linux operating system is capable of running corporate Web servers, FTP, Gopher and wide-area information server (WAIS) Web sites.

Linux is a fully functional operating system similar to a UNIX system. It has all the standard features of enterprise UNIX systems. Management of the command structure is enabled through *shells*. We mention two of the many advanced shells that are available for Linux, the Bourne Again Shell (bash) and the tsch shell. Both shells support a complete shell-based programming language that you can use to create shell scripts.

There are four fundamental components of Linux:

- ▶ **Kernel:** Provides program and hardware device management (such as disks and printers)
- ▶ **Shell:** Receives commands from the user and sends those commands to the kernel for the execution through a standard user interface.
- ▶ **File structure:** Files are organized into directories with the disk hardware. Each directory may contain any number of subdirectories each holding files. This is a similar structure to PC operating system file structures.
- ▶ **Utilities**

In some solutions, typically with clusters, the Linux server does not need the traditional PC hardware BIOS. The hardware is directly controlled by Linux. This provides phenomenal boot-up times (three seconds is the current record).

Linux has the same multiuser and multitasking capabilities as large UNIX servers. It provides the same level of system administration that you find on standard UNIX systems. Users can run several programs concurrently. You can create user accounts for different users and define their access rights to the files and system components. Installation of new devices and network connection and control is also provided as standard in the Linux operating system.

As a development environment, Linux has a powerful set of development tools for creating industrial-strength applications. The development toolset includes the *GNU C Compiler*.

The Linux structure provides the ability for programmers to access the hardware of the computer and the networks to which it is connected. This is achieved by the provision of a *hardware abstraction layer* where programs can take advantage of hardware features through a standard applications programming interface (API).

Linux programs can be portable to other versions of UNIX systems. Linux can use ANSI C, combined with one of several portability graphical user interface (GUI) toolkits. These programs can be written for both UNIX systems and Windows servers.

1.3.1 Linux BIOS management

In a cluster environment, the ability to quickly reconfigure hardware is particularly desirable. Linux can help with this task. Linux has the ability to boot without a BIOS. Current PCs used as cluster nodes depend on a vendor-supplied BIOS for booting. The BIOS, in turn, relies on inherently unreliable devices, such as floppy disks and hard drives, to boot the operating system. In addition, current BIOS software is unable to accommodate non-standard hardware, making it difficult to support experimental work. The BIOS is slow, often erroneous, redundant, and most importantly difficult to maintain. Imagine walking around with a keyboard and monitor to each of the 128 nodes in a cluster to change one BIOS setting.

The LinuxBIOS gunzips the Linux kernel from NVRAM and essentially requires no moving parts other than the fan. It performs a minimal amount of hardware initialization before jumping to the kernel start and lets Linux do the rest. As a result, it is much faster (current record three seconds), which has sparked interest in the consumer electronics community as well. Moreover, updates can be performed over the network.

Using a real operating system to boot another operating system provides much greater flexibility than using a simple netboot program or the BIOS. Because Linux is the boot mechanism, it can boot over standard Ethernet or over other interconnects such as Myrinet, Quadrics, or SCI. It can use SSH connections to load the kernel, or it can use the InterMezzo caching file system or traditional network file system (NFS). Cluster nodes can be as simple as necessary, perhaps as simple as a CPU and memory, no disk, no floppy, and no file system. The nodes will be much less autonomous, making them easier to maintain.

The LinuxBIOS Project: LinuxBIOS is an open source project. It aims to replace normal BIOS with hardware initialization and a compressed Linux kernel that can be booted from a cold start. The project began as part of clustering research work in the Cluster Research Lab at the Advanced Computing Laboratory at Los Alamos National Laboratory. The motivation for the project was for the operating system to gain control of a cluster node at power on. Other benefits of LinuxBIOS include needing only two working motors to boot (CPU fan and power supply), fast boot times (current fastest is three seconds), and freedom from proprietary (buggy) BIOS code. Overall, the benefits are numerous and helped gain support from many vendors in the high-performance computing and embedded computing markets.

For more information about LinuxBIOS, refer to the following Web site:

<http://www.linuxbios.org/index.html>

1.3.2 Where Linux fits in

The adoption of Linux has been a rapid and changing process. It can be difficult to determine where Linux fits into the modern IT world. Figure 1-2 shows the modern IT spectrum, from the operating systems to the client interface. Linux plays a role in almost every area and is quickly making inroads where traditional operating systems ruled for years.

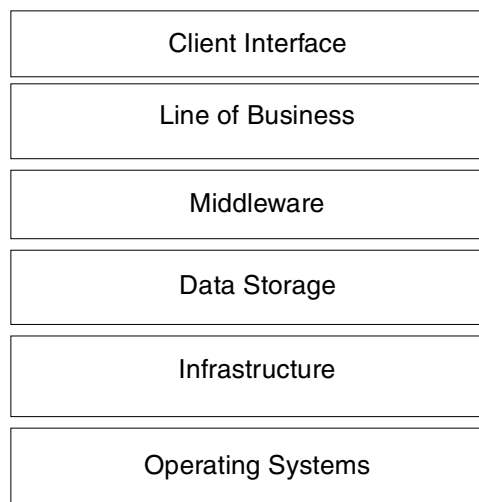


Figure 1-2 The IT spectrum: Logically and physically

1.3.3 Working with other operating systems

Linux is being used not only as the single operating system on a given computer system. It is also frequently used to populate virtual partitions on larger, non-Intel based systems such as the IBM @server iSeries, pSeries, xSeries, and zSeries servers. The virtual or logical partition provide:

- ▶ **Server consolidation:** An organization with existing Linux or UNIX servers can save hardware and personnel costs by consolidating those servers into virtual partitions on an existing IBM @server platform.
- ▶ **High availability:** Multiple virtual partitions can provide a level of availability and application-based load balancing while retaining the management benefits of consolidated hardware.
- ▶ **Infrastructure:** Linux virtual partitions can provide security and infrastructure services to other partitions that share the same virtual network.

1.4 Infrastructure

Linux enters the mainstream markets by providing critical infrastructure services.

Web serving

The combination of Linux and Apache offers an attractive package for customers. It provides a low cost, flexible solution originally for static Web sites, with over 30% of the world's Web sites running this combination. The demand is now moving toward a more dynamic approach with Web sites that users can interact with and that support high transaction rates.

File and print serving

One of the basics for Linux implementation is the provision of inexpensive facilities such as file and print services. Linux offers a rapid return on investment (ROI) in this part of the infrastructure space. The management capabilities and low cost make this an easy solution to justify. Also, this is an important environment, but it does not typically have the operational importance of line-of-business applications. It is a relatively safe place for businesses to test this new technology.

Domain name server (DNS) and DHCP

As a UNIX-clone, Linux is well proven at hosting Berkeley Internet Name Daemon (BIND) name servers and Dynamic Host Configuration Protocol (DHCP) services.

Router

Linux is capable of advanced routing using inexpensive commodity hardware. Also, some router vendors have chosen Linux to be their embedded operating system.

Firewall and Intrusion Detection Services (IDS)

Linux has been a popular provider of firewall and IDS services. Because of the advanced configuration and customization options, along with a small memory footprint, Linux has been an ideal solution for many organizations who want to avoid proprietary solutions.

1.5 Data storage

The explosive growth in data storage needs has driven Linux into a premier position as a low-cost, efficient, and highly customizable platform to handle storage management. In particular, Linux is being used in Internet Small

Computer Systems Interface (iSCSI) environments. iSCSI works by attaching SCSI-based storage resources directly to an Ethernet network and is a low-cost alternative to Fibre Channel solutions. In this environment, Linux operates as a storage node, managing RAID hardware and providing a high level of network performance.

The IBM TotalStorage™ IP Storage 200i is an example of a storage device using Linux as an operating system. You can find more information on the Web at:

<http://ssdweb01.storage.ibm.com/snetwork/iSCSI/index.html>



Figure 1-3 The IBM TotalStorage IP Storage 200i

1.6 Middleware

Linux is currently seeing large growth in the middleware arena. Middleware is called the “glue” that holds IT applications together. It is typically divided into the following four areas.

Database

Linux is quickly growing as a preferred platform for databases because of its low cost, stability, and extensibility. It also offers enterprise-level features such as load balancing and high availability clusters. The two most common database implementations are DB2 and Oracle 9i RAC.

Note: IBM DB2 Universal Database™ provides enterprise-level database features on Linux. For more information, see:

<http://www.ibm.com/software/data/db2/linux/validate/>

Cross system support

Most products with the “middleware” label are really software platforms that provide integration, transaction, development support to tie various IT services together. Linux is the operating system of support for middleware because its

open and flexible nature works well in the heterogeneous environment of modern IT networks. IBM offers two of the strongest examples of middleware:

- ▶ WebSphere
- ▶ MQSeries®

Note: The IBM WebSphere platform provides integration, transaction, and development support services while taking advantage of the open nature of Linux. To learn more about WebSphere, go to:

<http://www.ibm.com/software/info1/websphere/index.jsp>

Knowledge discovery (KD) and collaboration

Increasingly, a major task of middleware is to organize the information flowing through a network, and not to just manipulate and store it. Knowledge discovery refers to the integration of enterprise content management (ECM) and knowledge management (KM). It entails merging multiple technologies that originally (and individually) were responsible for capturing, locating, organizing, and delivering information.

Collaboration includes all technologies that provide integrated messaging, calendaring, and task management. The open heritage of Linux provides a flexible platform in which to integrate the various knowledge discovery and collaboration tools.

Note: IBM provides knowledge collaboration products through its Lotus® line. Follow this link for more information about Linux and Lotus:

<http://www.lotus.com/>

Systems management

Systems management has increasingly become centralized and heterogeneous. With an decreased margin for downtime, middleware must incorporate the ability to monitor all systems at all levels of abstraction. Because Linux is open in nature, it allows a higher level of access to its internal workings than closed or proprietary systems offer.

A management infrastructure built on Linux also takes advantage of its UNIX system heritage and many third-party and open source products that can provide additional functionality.

Note: IBM provides systems management middleware through its Tivoli software family. You can learn more about Tivoli on the Web at:

<http://www.ibm.com/tivoli>

1.7 Application solutions

Application solutions refer to those applications that are critical for day-to-day operation of an organization. They usually encompass Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), Supply Chain Management (SCM), and Business Intelligence (BI) products. Some Linux-capable solution providers are introduced in the following sections. They include SAP, J.D. Edwards, and eOne Group.

Overall, there is a wide range of support for Linux across the entire spectrum of IT system requirements. Table 1-1 summarizes some of the key elements and the applications that fill them.

Table 1-1 *Business areas and solutions*

Infrastructure solutions	Vendor or solution
Database and Application Development Tools	DB2 UDB, Everyplace, Oracle, Informix (Retail/Distribution) Sybase (FSS), Rational
Application Server	WebSphere Application Server, Apache
Network and System Management	Tivoli, VERITAS, BMC, Trustix, Compuware
Security and Firewall	Checkpoint, Symantec, CISCO, Trustix, SuSE Firewall, Tivoli
Clustering (High Availability)	Steeleye, VERITAS, PolyServe, Legato
Server Consolidation	VMware, Samba
Mail and Messaging	Sendmail, Bynari, Domino, Caldera VMS, SuSE eMail Server
e-business solutions	
E-commerce	WebSphere Commerce Suite, eOneGroup
Accounting and CRM	ACCPAC, Sage
Business Intelligence and Data Warehouse	SAP BI, Brio, Crystal Decisions, Dimensional Insight, Hyperion
ERM/CRM/SCM (cross industry)	mySAP.com, QAD, BaaN Invensys
Retail Distribution	Relavis, MarCole, 360 Commerce, Caldera
Finance and Banking	Axiom, Sanchez, Rueters, Intelligent Markets, JD Edwards CRM for Financial Services, SUNGARD
Industrial (EDA)	Synopsis, Cadence, LSCD, Mentor Graphics

SAP

SAP supports Linux versions of their ERP products, but restricts support to specific hardware platforms. For a list of SAP-certified Linux hardware platforms, see:

<http://sap.com/solutions/technology/linux/platforms/index.asp>

J.D. Edwards

J.D. Edwards has announced that they will offer their Financial Services CRM tool with the IBM WebSphere Application Server running on a Linux and xSeries server.

You can learn more about this announcement on the Web at:

<http://www.jdedwards.com/pressreleases/>

eOne Group

eOneCommerce is 100% platform independent. You can integrate it with IBM hardware, software, database, and operating system. It can run in a Linux LPAR or under a native operating system.

1.8 Clients

While desktop versions of Linux traditionally trailed Microsoft products in ease-of-use, that is rapidly starting to change. There is an increasing use of Linux as a client interface in all areas of computing. The following sections provide specific examples.

Thick clients

Thick clients include desktops and laptops that run Linux as their primary operating system. However, it is still common (especially in development groups) to see dual-booting machines to allow some Windows compatibility.

Thick clients are prevalent in the business and technical support organizations.

Thin clients

High maintenance and labor costs have contributed to the trend of replacing thick clients with thin clients.

Linux makes an ideal thin client platform due to the limited memory requirements and the broad embedded processor support. The highly customizable nature of Linux allied with centralized control makes the thin client an attractive business proposition.

Note: The IBM NetVista™ line offers several Linux-based thin client solutions. For more information, see:

<http://www.ibm.link.ibm.com/cgi-bin/master?xh=teg8ZCuCakwz2Q2USenGnN9332&request=salesmanual&parms=H%5F8363%2D%2B03&xhi=salesmanual%5E&xfr=N>

Personal Digital Assistants (PDAs)

Linux has been a popular choice for embedded devices for all types. However, commercial Linux-based PDAs only became available in the last few years.

Application portability is cited as the major benefit, because users can run the exact same applications on their PDAs as they do on their desktops. An excellent example of a Linux-based PDA is the Paron MPC. The Paron was developed primarily for vertical markets that need secure, encrypted access to a Bluetooth-enabled infrastructure like banking, inventory, and government. Paron MPC offers a finger-print recognition chip as part of its security features.

IBM co-developed the Paron MPC with CDL. To learn more about this, see:

<http://www.linuxdevices.com/articles/AT7145548309.html>
http://www.cdlusa.com/products/paron_mpc.shtml

1.9 The Linux distributions

This section introduces you to the Linux distributions. It answers the question: “How does a distribution of Linux differ from the Linux operating system itself?”

How do you know which is the right distribution for your environment? Hundreds of Linux distributions are currently available. The move into enterprise class software drives and separates the distributions. Examples of enterprise class distributions include SuSE Linux Enterprise Server (SLES) and Red Hat Linux Advanced Server (RHAS).

Not all distributions are supported on all IBM @server platforms. In general, IBM supports those distributions that provide the enterprise class of service. These distributions cost more, but include more enterprise class features and multiple levels of support from bug fixes to 24-x-7 support.

The features of enterprise class distributions include:

- ▶ **New offerings:**
 - Optimized release cycles of 12 to 18 months, three to five year version support
 - Independent software vendor (ISV) certification platform

- Increased commitment to maintenance and support
- Enterprise rich functionality such as multithreaded input/output (I/O), improved SMP usage, and more
- ▶ **Customer value:**
 - *Reliability*: Longer integration and testing cycles
 - *Scalability*: Improved SMP; functions such as Piranha from Red Hat for load balancing and SuSE backport of 2.6 scheduler for scalability and performance
 - *Manageability and service*: Improved RAS function, 24-x-7 global enterprise class support offerings
- ▶ **Implementation:**
 - Red Hat Advanced Server 2.1
 - SuSE Linux Enterprise Server 2.1
 - SCO Linux
 - Turbolinux Enterprise Server

This section presents some of the most popular distributions, in particular those that run on IBM *@server* platforms and give you a preview of each. Then it highlights other players in the marketplace. That includes distributions that are not as well known but have an impact in specific geographies or are suited to a different audience.

1.9.1 What is a distribution

What is usually referred to as *Linux* is more accurately called a “distribution” or “distro” for short. Technically, Linux is only the core of the operating system or the kernel. Linux distributions are prepackaged collections of software, generally comprising the kernel and various other packages.

The major distributors include a graphical desktop environment and an installation program to help you get started. The kernel (for which the source code is freely available) is managed by Linus Torvalds and a leadership team. Many developers contribute to updating the kernel with new features as well as device drivers. The packages that can be added to the operating system include compilers, shells, and other applications. Not all of these packages may be available as open source. You can learn about open source in 1.9.3, “What is open source” on page 21.

A Linux distribution includes the Linux kernel plus utilities, programming tools, window managers, and other software that make up a full operating system. Distribution companies, such as SCO, Red Hat, SuSE, and Turbolinux, and non-commercial organizations, such as Debian, download the latest open source

packages from the Internet. Then they add utilities, such as installation programs, and package them on a CD-ROM with a manual. The kernel code in each distribution may be different since distributions pull Linux versions. Linux Standards Base attempts to standardize the aspects of the a distribution to minimize the changes that affect portability.

A *distributor* (such as Red Hat) is a company who packages a Linux distribution for commercial purposes. They technically don't resell the software license, but rather sell a subscription service for support and maintenance. Generally these are one-year subscriptions. Distributors also sell Linux consulting and integration services. They may sell applications or middleware in addition to the Linux operating system.

Application compatibility

In the vast majority of cases, Linux applications are compatible with all distributions of Linux, which accounts for the aphorism "Linux is Linux is Linux". When a new Linux kernel is released, it is uploaded to the main Linux kernel site:

<http://www.kernel.org>

The distribution companies then package the kernel and provide it as an update on their own Web sites. Distribution vendors take the kernel as is, with all changes and fixes that are contributed by members of the development community. The speed of distribution release has slowed considerably in the last few years. At one point, new distributions were released about twice a year. With the emphasis on the enterprise class distribution and wide range of products running on or under Linux, the pace of releases has slowed considerably.

Each distribution categorizes its products uniquely. For example, Red Hat divides its products into enterprise, small/medium business, and home versions. SuSE has two divisions: business and private. Turbolinux showcases its products in a single group as does The SCO Group. Where possible in this chapter, we separate the products into business and workstation markets.

As new hardware is announced and IBM Linux distribution partners release new versions, its a good idea to check the status of IBM support. The following Web links can help you find the latest status for the xSeries server. The IBM development team also writes installation instructions for each system, as it leaves testing, to document any recommendations for running versions of Linux operating systems.

For the version of Linux that is supported by IBM on each model of the xSeries server, see:

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>

For options that are supported by IBM by system and operating system, see:

<http://exist.raleigh.ibm.com/nosmgr/optlist/sprec.asp>

For example, IBM works with SCO, Red Hat, SuSE, and Turbolinux to certify xSeries hardware via the Linux distribution partner certification programs. The results are posted within the Linux distribution partner Web sites. For more information, see:

- ▶ Red Hat:

<http://hardware.redhat.com/hcl/?pagename=hcl>

- ▶ SuSE:

http://hardwaredb.suse.de/index.php?LANG=en_UKd

- ▶ Turbolinux:

<http://www.turbolinux.com/hcl/TTlist.html>

1.9.2 What is common among the distributions

There is much in common among each distribution. For example, each distribution comes with two possible desktop GUI interfaces: the K Desktop Environment (KDE) and GNOME. To the end user, they are functionally equivalent. Choosing one over the other is a matter of personal preference. Both of these Windows desktop GUIs function well as graphical interfaces. For developers, the difference is the set of tools and libraries that each option offers. KDE uses the Qt C++ cross-platform GUI toolkit for KDE development, where GNOME uses the Graphics Tool Kit (GTK) and other tools and utilities.

Each distribution has its unique installation program. Some are GUI programs instead of text-based programs. An office suite is usually included with each distribution such as Star Office, Open Office, or a distribution-created office suite. Most distributors also contain a variety of utilities that include configuration and installation utilities, firewalls, Web browsers, and Web server programs.

The distributions include differences in the versions of the kernel, compilers, libraries, etc., supported languages, and adherence to standards. UnitedLinux seeks to answer some of these issues with standards. Support can also vary from distributor to distributor, for example response times to problems, sales and support coverage by geography, etc.

1.9.3 What is open source

Open source is often thought of as software that is given away or free to be redistributed. In reality, open source is software whose *source code* is available free of charge to the general public. The rationale for this concept is that a larger

group of programmers, who are not concerned with proprietary ownership or financial gain, will produce a more useful and bug-free product for everyone to use.

The Internet is full of open-source software in heavy commercial use. Some examples include:

- ▶ **Apache:** Runs over 70% of the world's Web servers
- ▶ **Perl:** The engine behind most of the “live content” on the World Wide Web
- ▶ **BIND:** The software that provides the domain name service (DNS) for the entire Internet

Note: There are ten items that comprise the Open Source Definition. You can learn more about them on the Open Source Initiative (OSI) Web site at:

<http://www.opensource.org>

You can read more about open source in 2.2.2, “The compromise of open source software” on page 52.

Linux is an example of open source software. Linux programmers, working over the Internet, are the community that improves, adapts, and fixes Linux with amazing speed. This rapid evolutionary process frequently produces better software than the traditional development model.

Linux: An example of open source

Linux is written and distributed under the GNU General Public License (GPL). (GNU stands for “GNU is Not UNIX”.) The GNU General Public License is an example of certified open source software. As is the case for open source, Linux’s source code is freely distributed and available to the general public.

The General Public Licenses are designed to make sure that you:

- ▶ Have the freedom to distribute copies of free software (and charge for this service if you want)
- ▶ Receive source code or can obtain it upon request
- ▶ Can modify the software or use portions of it in new free programs
- ▶ Are aware that you can do these things

You can read the entire script of the GNU General Public License on the Web at:

<http://www.gnu.org/copyleft/gpl.html>

1.10 Red Hat



Founded in 1994, Red Hat provides the most popular distribution in the U.S.A. It is well known for its Red Hat Package Manager (RPM) system for installing and maintaining its components.

Red Hat is a leader in the development, deployment, and management of Linux and open source solutions for Internet infrastructure, ranging from embedded devices to secure Web servers. Red Hat is the basis for several other Linux distributions.

1.10.1 What's in the distribution?

Red Hat offers software solutions for the enterprise, for small to medium businesses, and for the small office/home office. This section describes each of the major offerings and their features.

Red Hat Linux operating systems

Operating systems for the enterprise consist of:

- ▶ Red Hat Linux Advanced Server for large enterprises
- ▶ Red Hat Linux for iSeries, pSeries, zSeries, and S/390.
- ▶ Red Hat Linux for the Itanium Processor for Intel 64-bit Itanium processors

Operating systems for the small/medium business consist of:

- ▶ Red Hat Linux Advanced Server
- ▶ Red Hat Linux 8.0 Professional

Operating systems for the small office/home business consist of:

- ▶ Red Hat Linux 8.0 Professional
- ▶ Red Hat Linux 8.0 Personal

Red Hat Linux business applications

Red Hat also provides applications for enterprise businesses and small to medium businesses. Some of these apply to both environments.

Application solutions for the enterprise consist of:

- ▶ **Stronghold Enterprise:** An Apache-based secure Web server designed to enable secure transmission and storage of mission-critical enterprise data over commerce servers, intranets, and extranets.
- ▶ **Content and Collaboration Management (CCM):** Open source software solutions for managing content and enterprise team collaboration. It was originally developed by ArsDigita.

- ▶ **GNUPro Developer Tools:** A commercial software development suite of tools built around the open source GNU standard. It is delivered by subscription to provide continual access to the latest upgrades and support.
- ▶ **Cygwin:** A UNIX/Linux development and system administration environment for Microsoft Windows. It includes a UNIX/Linux shell environment, Windows ports of popular GNU development tools, and the Cygwin library, which provides the required UNIX system calls and environment.
- ▶ **Embedded Linux Developer Suite:** A collection of development tools and Linux-based embedded runtime technologies that enable the creation, deployment, and testing of target software components for embedded devices.

Application solutions for the small and medium business consist of:

- ▶ **Stronghold Enterprise:** An Apache-based secure Web server designed to enable secure transmission and storage of mission-critical enterprise data over commerce servers, intranets, and extranets.
- ▶ **Content and Collaboration Management (CCM):** Open source software solutions for managing content and enterprise team collaboration. It was originally developed by ArsDigita.

Red Hat Linux Advanced Server contains these features:

- ▶ An extended release cycle for long-term stability
- ▶ An enterprise-tuned kernel
- ▶ Red Hat Cluster Manager, which provides an application failover infrastructure
- ▶ Network Enterprise Service for simplified management of multiple systems

The Advanced Server also includes Bluecurve, which offers a desktop GUI firewall, Web server (Apache), GUI support for Braille, and mobility-limited language interfaces.

Red Hat Linux Professional and Personal include:

- ▶ OpenOffice office suite for creating documents, spreadsheets, and presentations
- ▶ Mozilla open source browser, e-mail client, address book, and HTML author
- ▶ Ximian Evolution e-mail client, contact manager, and calendar
- ▶ Customizable firewall to dynamically adjust security settings

Added features in the Professional version include office and multimedia applications as well as administration tools and extended service.

You can find the Red Hat Application Brief on the Web at:

<http://w3.linux.ibm.com/linux/files/RedHatApplicationBrief.pdf>

For additional information about Red Hat versions, see:

<http://www.redhat.com>

1.11 UnitedLinux

What may be a major event in Linux history started on May 30, 2002. SCO, Conectiva, SuSE, and Turbolinux announced that they would join forces to create a single binary-compatible Linux distribution, thus sharing development costs and future royalties. UnitedLinux was formed as a consortium, and the UnitedLinux product UL 1.0 was ultimately announced in November of 2002. By 4 December 2002, SCO unveiled SCO Linux 4.0, and SuSE unveiled its SuSE Linux Enterprise Server 8. These were the first Linux distributions on the market based on the UnitedLinux 1.0 code. Turbolinux released its Enterprise Server 8 version in Japan.

UnitedLinux is a standards-based Linux operating system. It targets the business user. It seeks to streamline Linux development and certification around a global, uniform distribution of Linux. It is certified to work across the both AMD and Intel 32-bit x86 lines; the 64-bit x86 Athlon, Opteron, and Itanium families; and the complete IBM eServer brand.

While each distributor sells the same core UnitedLinux distribution, they can add additional programs and utilities, provided that application portability is preserved. Each distributor sell its products under its own brand name “Powered by UnitedLinux”.

You can find a UnitedLinux Application Brief on the Web at:

<http://www.ibm.com/linux/UnitedLinux-Application-Brief.pdf>

Read more about UnitedLinux in their technical white paper at:

<http://www.unitedlinux.com/pdfs/whitepaper4.pdf>

Major companies and ISVs can now work with two main Linux versions (the other being Red Hat) instead of five or more major versions. This simplification saves companies from porting applications across multiple platforms. In its first month, over 15,000 downloads of the UnitedLinux beta were initiated.

1.11.1 What’s in the distribution?

UnitedLinux contains a common kernel, APIs, file systems, and installation routines for business Linux. Developers and customers obtain UnitedLinux via one of UnitedLinux’s supporting vendors. That is, you order it from either Conectiva, SuSE, SCO, or Turbolinux. Regardless of the vendor name on the

box, you receive UnitedLinux. This comes in the form of an installation CD or CDs. Each company then adds their own software packages, which come on separate CDs. Figure 1-4 shows the architecture of UnitedLinux.

SCO Components	Conectiva Components	SuSE Components	Turbolinux Components	UnitedLinux Features
Linux Standard Base (LSB)				
Common hardware, database and autoprobe backends		Common packages, libraries, kernel drivers and driver update		Common configuration, file locations, syntax and semantics

Figure 1-4 UnitedLinux architecture

UnitedLinux is a General Public License Linux. The source code is free, but the binaries (executable code) are not. For example, using the UnitedLinux source code, developers can create their own binaries (compiled source code ready for execution), but they won't be allowed to use UnitedLinux branding or UnitedLinux support.

In addition, UnitedLinux supports the Linux Standards Base (LSB) described in 1.11.2, "The Linux Standards Base" on page 27. You can view UnitedLinux as an implementation of the LSB and the Li18nux and GB18030 standards. UnitedLinux also has international support for the English, German, French, Italian, Japanese, Korean, Portuguese, Spanish, Simplified Chinese, and Traditional Chinese languages. For more information, see:

<http://www.unitedlinux.com>

The standards UnitedLinux requires of its members

UnitedLinux adheres to Linux and industry standards, including:

- ▶ **File System Hierarchy Standard (FHS):** A set of requirements guidelines for _le and directory placement under UNIX-like operating systems
- ▶ **Linux Standard Base:** Standards that increase compatibility
- ▶ **LI18NUX:** Internationalization of a core set of APIs and components

- ▶ **Scalability:** To take advantage of larger and more complex systems and handle new categories of applications:
 - A complete set of software and tools to build server farms for workloads that would otherwise be unmanageable for a single machine
 - Scheduler enhancements to improve process scheduling on SMPs and to prevent the scheduler from becoming a bottleneck
 - Asynchronous input/output to minimize waiting on I/O in large, busy systems
- ▶ **High availability:** The ability to avoid downtime, including:
 - POSIX-compliant event logging and notification capability
 - Dynamic probes: Dynamic insertion of breakpoints in code for debugging
 - Non-disruptive and tailored dumping of system data
 - Toolkit to record and trace system events
 - Hotplug PCI support
- ▶ **Security:** Kerberos support, basic firewall support, and a consolidated set of community security enhancements known as *Bastille*
- ▶ **File systems:** Support for the journaling file system (JFS), and the popular community file systems Reiser File System (ReiserFS), XFS, and the ext3 file system
- ▶ **Network/storage/device management:** Supports Logical Volume Manager (LVM) and Enterprise Volume Management System (EVMS) to manage storage
- ▶ **Platform support and interoperability:** Hardware capability across all relevant platforms and architectures, including Intel (32 and 64-bit), AMD, PowerPC® (iSeries and pSeries), and the zSeries mainframe
- ▶ **Development environment:** A development environment for ISVs that includes all the compilers for libraries, sources, text editors, graphical user interface support, and other tools to enable the building of applications for UnitedLinux

For more information, see the UnitedLinux white paper at:

<http://www.unitedlinux.com>

1.11.2 The Linux Standards Base

The Linux Standards Base is an earlier attempt designed to bring the array of Linux distributions into a common core. It is a workgroup of the Free Standards Group.

Supported by the development community and IT industry leaders, the Free Standards Group is an independent, non-profit organization dedicated to accelerating the use and acceptance of open source technologies through the development, application, and promotion of standards. These standards include common behavioral specifications, tools and APIs, to make development easier across Linux distributions.

For more information, see the Linux Standards Base home page at:

<http://www.linuxbase.org>

1.12 Conectiva

Conectiva, a UnitedLinux partner, is a Brazilian Linux distribution with several titles in Spanish and Portuguese. This Linux distribution is created specifically for Latin America. Conectiva Linux is aimed at both desktops and servers. It is RPM-based and APT-enabled. It can be installed using several preconfigured profiles, or it can be completely customized for specific tasks or power users.

1.12.1 What's in the distribution?

Conectiva includes over 1,000 applications including Star Office, Netscape (in native languages), games, image manipulation applications, and other utilities. It also includes a wide array of commercial software, such as Oracle, VariCAD, and ViaVoice®.

The Server Edition comes with software for network administration, e-commerce, and support for RAID and clusters. The current version includes native support for the standard network protocols including IPv6, NIS, and LDAP. Additional features include remote boot, Webmail, load balancing, RAID, FAX, Groupware, virtual file system (VFS), security auditing, and mail and Web servers.

Conectiva contains two different categories of profiles, one for servers and one for work stations. The file systems that are supported include Ext2fs, Reiserfs, NTFS, FAT, and FAT32.

Conectiva ships with Distributed Replicated Block Device (DRBD), a high availability kernel module that mirrors hard drive contents over a network. A high availability utility, the *heartbeat package*, monitors machines within a cluster and notifies the setup when one of the machines within the cluster dies, allowing the box with the mirrored drive to take over.

You can learn more about Conectiva on the Web at:

<http://www.conectiva.com>

1.13 SCO Open Linux



The SCO Group (SCO, formerly Caldera International) is a provider of software solutions for small- to medium-sized businesses (SMBs) and replicated branch offices. SCO solutions include UNIX and Linux platforms; management, messaging, and e-business tools; and services that include technical support, education, consulting, and solution provider support programs.

Based in Lindon, Utah, SCO has a worldwide presence with offices in 18 countries and representation in 82 countries. This infrastructure enables SCO to provide local support and dependable solutions to businesses around the world. In addition, SCO has a channel of more than 16,000 solution providers, a developer network of nearly 10,000, thousands of direct account customers and an installed base of more than two million systems. SCO solutions are divided into three broad areas: operating systems, extended platform and services.

Caldera, Inc. was founded in 1994 by Ransom Love and Bryan Sparks. In 1998, Caldera Systems, Inc. was created to develop Linux-based business solutions. In 2001, Caldera Systems, Inc. acquired the assets of the Server Software Division and Professional Services Division of The Santa Cruz Operation, Inc. (SCO), forming a new company, Caldera International, Inc. In 2002, Caldera changed its name to The SCO Group (NASDAQ: SCOX).

1.13.1 What's in the distribution?

SCO Linux 4.0 powered by UnitedLinux is the next generation enterprise class Linux operating system. It is built on industry standards and contains the reliability and stability characteristics that have long been the hallmark of SCO platforms.

SCO offers four editions of SCO Linux to meet the needs of a range of business customers. All editions include one year of maintenance bundled with the product. To find out which one is right for you and for more information on SCO Linux 4.0 powered by UnitedLinux, go to:

<http://www.sco.com/products/scolinuxserver>

Key features of SCO Linux 4.0 powered by UnitedLinux

The key features of SCO Linux 4.0 include:

- ▶ A powerful and easy to configure LVM. LVM allows system administrators to create arbitrary grouping of physical hard disks. These groups can be modified (enlarged or reduced) based on need, giving unlimited scalability to the server.

- ▶ Better data throughput using asynchronous I/O. Asynchronous I/O provides a more efficient use of system resources by eliminating bottlenecks around data transfers. Instead of the CPU stalling while data is transferred for processing, the system can use that time to process other data that is ready.
- ▶ Support for storage area networks (SAN).
- ▶ Improved scalability due to the O(1) scheduler for systems under heavy workloads.

Support for virtual private networks (VPN)

SCO Linux Server 4.0 includes a complete IPsec VPN solution using FreeSWAN. With VPN technology, administrators can create extension of a private network that encompasses links across shared or public networks.

Webmin System and server administration

SCO Linux Server 4.0 features the Webmin system and server administration tool. This powerful tool simplifies all areas of system administration, from adding and removing users to configuring Web infrastructure servers. Since Webmin is browser based, administrators can manage systems even in remote geographical locations as easily as if they were sitting at the machine.

Improved system security

SCO Linux Server 4.0 includes a broad range of security applications such as the security hardening tool Bastille, SAINT for security audits, Ethereal for network traffic monitoring, and Tripwire for intrusion detection. These applications coupled with SCO's already secure operating system give you the tools you need to ensure system security.

Expanded journaling file system support

Journaling file systems add a higher level of reliability and faster recovery time. JFS, ReiserFS, XFS, and Ext3 journaling file systems are included with SCO Linux Server. Each of these file systems has been tested and optimized for the best performance and stability.

Interoperability with Windows

SCO Linux Server 4.0 includes an enhanced version of Samba 2.2.5 for superior connectivity with networks and users running Microsoft Windows networks. With SCO Linux, users can transparently share files and printers regardless of what operating system they are running on their desktop.

For a full list of SCO Linux 4.0 features, see:

http://www.sco.com/images/pdf/scolinux/SCO_Linux_DS3.qxd.pdf

SCO Linux Server 4.0 for the Itanium Processor Family is an enterprise-class operating system designed for use with Intel Itanium 2 processors. The UnitedLinux core of SCO Linux Server is a standards-based, LSB and Openi18n (formerly Li18nux)-compliant Linux platform that focuses on reliability, availability, stability and security. SCO Linux Server 4.0 provides reliable business services including, but not limited to a secure web server, file and print services (Microsoft Windows, Linux, and UNIX file and print), and network infrastructure services. SCO Linux Server 4.0 also includes twelve months of the SCO Linux Update Service. This service provides notification of availability of updates and an easy to use graphical tool to download and apply updates to your system.

Features and benefits

UnitedLinux Release 1.0 powers SCO Linux Server 4.0 for the Itanium Processor Family. UnitedLinux bears the engineering, testing, and financial resources of four leading Linux vendors: The SCO Group, Conectiva, SuSE, and TurboLinux. Their combined efforts have produced a Linux operating system that is truly ready for the enterprise.

The benefits of UnitedLinux include:

- ▶ **Major OEM and ISV certification:** With UnitedLinux, OEMs and ISVs are no longer forced to choose to which Linux versions they will certify. By certifying their products to a single UnitedLinux core, they can reach customers using any Linux operating system powered by UnitedLinux. This not only is a cost saving to the OEM or ISV. However, it ensures the broadest range of supported hardware and software for SCO Linux users.
- ▶ **Interoperability between Linux distributions:** Since all Linux versions that are powered by UnitedLinux use the same binary base, any application ported to UnitedLinux will run on *four* Linux distributions. This unique interoperability allows users of SCO Linux to take advantage of an ever-broadening range of software.
- ▶ **Improved reliability and stability:** UnitedLinux leverages the expertise of all four UnitedLinux vendors to create a stable, secure, and reliable operating system that has been tested to run in mission-critical environments.

The major features of SCO Linux include:

- ▶ **Support for the Intel Itanium® 2 Processor:** SCO Linux Server 4.0 for the Itanium Processor Family is among the first operating systems for Intel's Itanium 2 processor, Intel's powerful new 64-bit technology. It is designed to take advantage of the dramatic performance levels of these new processors and backed by SCO's legendary stability and performance.
- ▶ **Linux 2.4.19 Kernel:** The core of SCO Linux Server 4.0 is the 2.4.19 Linux kernel. New features include broadened USB support, Logical Volume Manager, improved journaling file system support, POSIX-ACLs, new O(1)

scheduler (improves SMP support), asynchronous I/O, Enterprise Volume Management System (EVMS), PCI hot plug support on supported hardware, and many other performance enhancing capabilities.

- ▶ **Security:** SCO Linux Server includes a broad range of security features. By ensuring that only minimal services are running on boot up, SCO Linux Server allows the user to enable only the services they need and eliminate possible security holes. In addition to this, only processes that must run as root are configured to do so. SCO Linux Server includes a suit of security tools such as SAINT, Bastille, and Portsentry.
- ▶ **Intrusion detection:** Several intrusion detection software packages, including Snort, Strobe, and Tripwire, have been integrated into SCO Linux Server to check for the intrusion of files and ports. They also advise the administrator of these intrusions so that they can take protective action.
- ▶ **Expert package selection:** This installation option allows an expert user to select and deselect individual software packages during installation so that the user can customize the server to the user's needs.
- ▶ **Webmin:** Webmin is a browser-based admin tool that assists the administrator in all areas of system management from adding users to configuring complex servers. Webmin can also be used for secure remote management of servers through a browser on the administrator's local system.
- ▶ **Proactive software management:** Each licensed copy of SCO Linux Server includes one year of automated security and maintenance updates via the Internet from the SCO Linux Update Service.
- ▶ **Software and hardware RAID support:** SCO Linux Server includes tools and drivers for installing and running Linux on RAID systems. The SCO Linux installer allows the user to configure software RAID during the setup process.
- ▶ **Journaling file system:** Journaling file systems add a higher level of reliability and faster recovery time. JFS, ReiserFS, XFS, and Ext3 journaling file systems are included with SCO Linux Server. Each of these file systems has been tested and optimized for the best performance and stability.
- ▶ **Automated installation:** With SCO Linux Server, you no longer have to install multiple servers manually one at a time. The automated installation feature in SCO Linux Server allows the user to create an XML file that the SCO Linux installer can read to perform a server installation with no user interaction.
- ▶ **Support for servers with greater than 4 GB of RAM:** The SCO Server includes support for systems with more than 4 GB of RAM.
- ▶ **Docview:** Docview is an online document viewer that serves up the entire product documentation in Web pages accessible from a Web browser either locally or remotely. Docview also indexes the entire documentation on the

system, including RPM package information, and makes it searchable by the user.

- ▶ **Firewall:** Includes IP-chains and Iptables firewalls.
- ▶ **Virtual private network:** With VPN technology, administrators can create an extension of a private network that encompasses links across shared or public networks. SCO Linux Server uses IPSec (FreeSWAN) to create VPNs.
- ▶ **Choice of Windows managers:** SCO Linux Server features both the KDE 3.0.3 and Gnome 2.0 desktop environments. This allows users to choose their preferred windows manager and have access to both KDE and Gnome applications.
- ▶ **High availability:** The following high availability applications are available for SCO Linux Server:
 - Heartbeat: Simple two-node fail-over for services running across a network such as Apache or Samba.
 - DRBD: Disk-over-LAN mirroring similar to RAID 1 but across a network
 - Linux Virtual Server (LVS): Using LVS, the administrator can build clustered systems for scalability and fault tolerance
 - Mladm: Software RAID administration for disk arrays
 - Multipath I/O on device arrays and logical volumes
- ▶ **Introduction to Linux courseware:** SCO Linux Server includes a sample of the wide range of educational courses available for SCO Linux Server and all SCO products.

SCO Linux Server 4.0 can be configured to run any of the following servers:

- ▶ **Web server:** Using current, secure versions of Apache, PHP, and Tomcat, SCO Linux Server includes everything you need to build advanced web servers.
- ▶ **Web proxy (HTTP/HTTPS/FTP proxy):** SCO Linux Server includes the Squid proxy server for speeding up Internet downloads. It also includes the Squid Web cache redirectory, Squirm.
- ▶ **File and print:** SCO Linux Server includes the tools you need for File and Print servers using Samba for Windows, CUPS (printing) and NFS (files) for UNIX systems, NetAtalk for Mac systems, and Mars_NEW for NetWare 2.x and 3.x systems.
- ▶ **Name servers and DHCP:** Name resolution can be accomplished with either DNS (using Bind 9) or WINS (using Samba). SCO Linux Server can use DHCP in both client and server mode.

- ▶ **FTP servers:** SCO Linux Server includes Very Safe FTP (VSFTP), an FTP server built specifically with security in mind. SCO Linux Server also includes the popular ProFTP and TFTP packages.
- ▶ **Mail and news servers:** SCO Linux Server includes both Postfix (default) and Sendmail mail transfer agents.
- ▶ **SQL database servers:** SCO Linux Server includes both MySQL and PostgreSQL database backend servers with extensions for ODBC, JDBC for heterogeneous OS access support.
- ▶ **Authentication servers:** SCO Linux Server includes a suite of authentication servers such as OpenLDAP, Kerberos 5, Samba (used for Winbind and as a Widows domain controller), NIS, and PAM.
- ▶ **Time server:** The Network Time Protocol (NTP) is used to synchronize the time of a computer client or server to another server or reference time source, such as a radio or satellite receiver or modem.

The core technologies in SCO Linux Server 4.0 include the following versions of some of the most sought-after features in SCO Linux:

- ▶ Linux Kernel 2.4.19
- ▶ Webmin 1.070
- ▶ glibc 2.2.5
- ▶ Apache 1.3.26
- ▶ Xfree86 4.2
- ▶ CUPS Printing System 1.1.15
- ▶ KDE 3.0.3/Gnome 2.0
- ▶ PostgreSQL 7.2.2/MySQL 3.23.52
- ▶ Samba 2.2.5
- ▶ Java2 JRE 1.3.1

OpenLinux Workstation

In OpenLinux® Workstation, SCO® has assembled popular development applications, tools, and utilities to answer the needs of most all Independent Software Vendors (ISVs) and corporate developers. These include:

- ▶ **Compilers/interpreters/assemblers**
 - Shells: bash, tcsh, zsh, korn
 - gcc (C)
 - g++ (C++)
 - PERL
 - Python
 - gawk (text manipulation)
 - flex (lexical manipulation)
 - bison (turns parser grammar into C/C++)
 - as & nasm (assembler)

- ▶ **Development language summary**
 - C/C++
 - Java (JSP & EJB)
 - Python
 - PERL
 - PHP
 - Scripting
- ▶ **Commercial tools**
 - SUN Microsystems Forte
 - Borland JBuilder Foundation
 - Borland Kylix 1.0 Open Edition
- ▶ **Editors/IDEs**
 - kdevelop (various, emphasizes C++/KDE/QT)
 - xemacs (various)
 - QT Designer (C++/GUI)
 - Vim (various)
 - kde advanced editor (various)
 - Quanta (HTML)
 - Gimp (image manipulation)
 - ImageMagick utilities
 - xv & pixie (image display and manipulation)
 - StarOffice™
 - CameleoLight™ (image manipulation)

SCO UnixWare® 7.1.3

This version of SCO UnixWare marks the return of the UnixWare brand name as the premier UNIX operating system for industry standard Intel and AMD processor systems. UnixWare 7 offers a choice for companies who place a high value on the Scalability, Reliability, and Security inherent in UNIX technology-based systems.

Like its predecessor release (Open UNIX 8.0), UnixWare 7.1.3 runs both Linux and native UNIX applications, using the Linux Kernel Personality Technology (LKP). The LKP feature provides a scalable, stable, secure, reliable environment, and has been updated in UnixWare 7.1.3 to support multibyte characters.

UnixWare 7.1.3 has been updated to enable SCO Update Service for UnixWare. This innovative service provides electronic notification and delivery of operating system changes and gives customers control of the upgrade process. SCO Update Services simplifies and streamlines the process of deploying new technology and keeping deployed systems up to date.

UnixWare 7.1.3 consists of a comprehensive family of preconfigured editions and optional products to build and deploy UNIX and Linux applications:

- ▶ **Data Center Edition** is for the highest-end multi-purpose servers demanding continual 24 X 7 availability, supporting hundreds or thousands of end users by supplying access to a wide range of applications from a variety of clients.
- ▶ **Enterprise Edition** is for medium-to-high-end enterprise servers to run large-scale business applications and databases for decision support and online transaction processing.
- ▶ **Departmental Edition** is for departmental servers in medium or large organizations to run applications and reliably share business critical information with a wide range of client devices.
- ▶ **Business Edition** is for small businesses or workgroups requiring file and print services, reliable access to diverse applications, and the ability to expand system capability as the organization grows.
- ▶ **Base Edition** is a solid foundation for building dedicated or specialized server environments, such as telecommunications equipment and other embedded systems. It also excels as a powerful graphical workstation.
- ▶ **Development Kit** supersedes the kit offered with Open UNIX 8 and Feature Supplements 7.1.1b and contains a wide range of tools that enable the rapid development of business critical native UNIX applications. The kit contains C and C++ compilers, extensive Java 2 development tools and popular Open Source tools, so developers can choose the right set of tools for their environment. Compatibility Modules for OpenServer™ are supported and can be used to create binaries for execution on SCO OpenServer 5.0.6 and OpenServer 5.0.7.
- ▶ **Optional Services** are a series of enhancement products that extend and complement the built-in functionality provided in the UnixWare 7.1.3 Edition products. The optional services that are available include:
 - NeTraverse Merge 5.3
 - Advanced File and Print Server 4.0.2
 - Reliant® HA 1.1.3
 - Online Data Manager 3.2
 - Disk Mirroring 3.2
 - UnixWare and OpenServer Development Kit (UDK) 7.1.3
 - SCOoffice Mail Server
 - Samba 2.2.4
 - Samba Multi-byte Version 2.2.5
 - Squid Caching Proxy Server 2.4 STABLE7
 - Cdrtools 1.11a21 record audio/data CDs

For additional information, see:

<http://www.sco.com>

1.14 SuSE SuSE

Founded in 1992, SuSE Linux is based in Nuremberg, Germany, and is the most widely used distribution in Europe. The latest versions have a new graphical installation tool to simplify the installation. The advanced features include hardware acceleration for 3-D graphic chips and a journaling file system. SuSE also develops and supports a series of X Servers for newer, high-end graphic cards. The installation is available in German, English, Italian, and French. SuSE Linux comes with a large group of utilities and applications.

SuSE is available in business versions. These include SuSE Linux Enterprise Server (SLES); servers for zSeries, xSeries, pSeries, and iSeries; and versions that support Alpha, PowerPC, and Intel platforms.

1.14.1 What's in the distribution?

SuSE divides its products into two categories: business customers and private customers. The following sections describe the major offerings and their features.

Solutions for business customers

SuSE Linux Enterprise Server is a server operating system for deployment in IT environments of all sizes. It is available for multiple hardware platforms:

- ▶ SuSE Linux Enterprise Server - SLES 8 for x86
- ▶ SuSE Linux Enterprise Server - SLES 8 for zSeries (64-bit)
- ▶ SuSE Linux Enterprise Server - SLES 8 for S/390 and zSeries
- ▶ SuSE Linux Enterprise Server - SLES 8 for iSeries and pSeries
- ▶ SuSE Linux Enterprise Server - SLES 8 for Itanium Processor Family
- ▶ SuSE Linux Enterprise Server - SuSE Linux 7.3 for SPARC

SuSE Linux Enterprise Server 8 has these features:

- ▶ Automated software installations and updates (AutoYaST)
- ▶ Asynchronous device management
- ▶ Dynamic storage management (Logical Volume Manager)
- ▶ Encrypted and journaling file systems ReiserFS, ext3, JFS
- ▶ Support for databases through raw devices
- ▶ Web servers: Apache and Tomcat
- ▶ File and print support
- ▶ Internet servers: Name server, Mail and News server

- ▶ SQL database server: ODBC and JDBC access
- ▶ Basic firewall

SuSE also supports IBM mainframe capabilities in S/390 and zSeries including channel architecture, virtualization, crypto hardware support, and hipersockets.

Other business solutions include

- ▶ SuSE Linux Openexchange Server: A groupware/e-mail server
- ▶ SuSE Linux Firewall
- ▶ SuSE Linux Database Server: A professional database system
- ▶ SuSE Linux Enterprise Client: A maintained client version for the desktop or laptop

Solutions for private customers

For private customers, SuSE offers these products:

- ▶ SuSE Linux 8.1 Professional
- ▶ SuSE Linux 8.1 Personal
- ▶ SuSE Linux Pro-Office for SuSE Linux 8.0
- ▶ SuSE Linux 7.3 PowerPC Edition
- ▶ SuSE Linux 7.1 for Alpha
- ▶ SuSE Linux Office Server
- ▶ SuSE Linux Office Desktop

SuSE Linux 8.1 Personal includes:

- ▶ An exclusive desktop system
- ▶ Free Microsoft-compatible office package (OpenOffice.org)
- ▶ Secure Internet and e-mail communication
- ▶ Personal fax solution for ISDN and modem with KDEprintFax

SuSE Linux 8.1 Professional contains these additional features:

- ▶ Intelligent software installer
- ▶ USB2 and Firewire hardware detection
- ▶ Integrated databases
- ▶ Configurable firewall
- ▶ Server and network functionality
- ▶ Office network integration
- ▶ Developer tools

SuSE Linux Pro-Office for SuSE Linux 8.0 includes StarOffice 6.0 by Sun and the K Desktop Environment. With their desktop products, SuSE intends to address the growing demand for Linux on the PC for office workers in the small office/home office (SOHO) environment (SuSE Linux Office Desktop) and in large enterprises (SuSE Linux Enterprise server). The latter product (planned for

GA in second quarter of 2003) will be maintained as an SLES. It delivers a Wine-based Windows emulator and a Office Suite, but can be customized on demand ranging from a thick client to a thin client.

For more information, see:

<http://www.suse.com>

1.15 Turbolinux

Turbolinux incorporates a suite of Linux products for both the workstation and server. It is particularly used in the Asia Pacific markets, having both Japanese and Chinese language integrated versions on top of the unified code base that is sold as the international (English) version.

Turbolinux offers Linux solutions geared toward corporate needs and home desktop needs. Turbolinux also offers a unique clustering solution that allows for the construction of highly available and scalable networks based on low-cost commodity components.

Note: Turbolinux, in conjunction with SuSE, its UnitedLinux partner, will supply the SLES8 product family for the zSeries, iSeries, and pSeries servers in first quarter 2003.

1.15.1 What's in the distribution?

The business solutions that Turbolinux provides are:

- ▶ Turbolinux Enterprise Server 8 powered by UnitedLinux (IA32)
- ▶ Turbolinux Enterprise Server 8 powered by UnitedLinux (Itanium Processor Family) in process of release
- ▶ Turbolinux 8 Server (IA32)
- ▶ Turbolinux 8 for AMD X86-64
- ▶ Turbolinux 8 Workstation (IA32)
- ▶ Turbolinux 7 for iSeries, Turbolinux 6.5 for pSeries
- ▶ Turbolinux Enterprise Server 8 for zSeries powered by UnitedLinux in process of release
- ▶ Turbolinux Enterprise Server 8 for iSeries powered by UnitedLinux in process of release

- ▶ Turbolinux Enterprise Server 8 for pSeries powered by UnitedLinux in process of release
- ▶ Turbolinux Cluster Server 7

Turbolinux 8 Server offers enterprise features. These include 64 GB memory support, large file support (LFS), which provides a maximum file size of 4 terabyte (TB) and 16 TB file system support, optimized scheduler, and LVM with disk volume management. It also includes journaling file systems, PXE support, Linux Kernel Crash Dump (LKCD) for problem analysis, and Enterprise Volume Management System (EVMS) support.

Turbolinux Server 6.5 for zSeries and S/390 is the first Turbolinux distribution optimized to run on IBM mainframe hardware platforms. It is a complete Linux server distribution. It includes the Linux operating system, database applications, an Apache Web Server, a Samba network file server, Sendmail mail server, C programming language application development tools, as well as an FTP server, DNS, and Telnet.

Turbolinux 7 for iSeries supports English, Japanese, Korean, and Simplified and Traditional Chinese with the unified code base. The Turbolinux 7 for iSeries distribution includes database applications, an Apache Web Server, a Samba network file server, Sendmail mail server, C programming language application development tools, as well as an FTP server, DNS, and Telnet.

Turbolinux 6.5 for pSeries supports English, Japanese, Korean, and Simplified and Traditional Chinese with the unified code base. Turbolinux 6.5 for pSeries is a complete Linux server distribution including the Linux operating system, database applications, an Apache Web Server, a Samba network file server, Sendmail mail server, C programming language application development tools, as well as an FTP server, DNS, and Telnet.

Turbolinux 7.0 for Itanium extends the Turbolinux Server to Intel's 64-bit architecture. e-business, Internet, print, file, and database configuration options are available in the base package.

Turbolinux Cluster Server 7 delivers availability, redundancy, scalability, security, and flexibility.

EnFuzion clusters all available computing resources on a corporate network to create a powerful “virtual supercomputer” for jobs that demand complex computations and data processing. Traditionally, these jobs were handled by high-end servers.

The Turbolinux server includes journaling file systems, PXE support, Linux Kernel Crash Dump (LKCD) for problem analysis, LVM and Enterprise Volume Management System (EVMS) support, and Linux event logging.

Turbolinux 8 Workstation includes e-mail, Web browsers, office and mail suites, PDA hot-sync capabilities, ADSL setup tool, MP3 and MIDI, and diverse hardware support including USB devices. In addition, Turbolinux 8 Workstation provides an ADSL interface for high-speed connections and extended support for various digital media including digital video data.

To learn more about Turbolinux, visit their Web site at:

<http://www.turbolinux.com>

1.16 Other players

There are hundreds of Linux distributions. You can find between 150 to 300 documented summaries on some Web sites.

Several excellent resources list and summarize distributions. Among these, we recommend:

- ▶ Linux Online:
<http://www.linux.org>
- ▶ LinuxBasis:
<http://www.linuxbasis.com/distributions.html>
- ▶ DistroWatch:
<http://www.distrowatch.com>

Due to the dynamic nature of Linux, we find continual changes as to which distributions are ranked in the top ten. If we don't mention your favorite distributor, it is not an oversight. Rather we highlight those that are currently most visible. While this redbook was written, UnitedLinux Version 1.0 was announced as well as new products by Turbolinux and Red Hat. The Linux world is in constant change, so we encourage you to investigate the resources yourself to find the latest information.

The "other players" we highlight in this section include Conectiva, a UnitedLinux distribution; Redflag, the major distributor in People's Republic of China; Yellow Dog, a Red Hat derivative that focuses on the PowerPC market; and Debian, a non-commercial distribution originally targeted for developers. Table 1-2 compares these "players".

Table 1-2 Basic distributor information

Distribution name	Developed by	Platforms	Server, Workstation, Enterprise	Location on the Web
Conectiva	Conectiva Curitiba, Brazil	Intel	S - W	http://www.conectiva.com
Debian	Ian Murdock	Alpha, Intel, PPC, Sparc, Other	S - W	http://www.debian.org
Redflag	Redflag Software Co., Ltd. Beijing, PRC	Alpha, Intel, PA, RISC, Sparc, S/390	S - W - E	http://www.redflag-linux.com/eindex.html
Red Hat	Red Hat, Inc. Raleigh, NC	Intel x86 Itanium/x86-64 Opteron zSeries, S/390 iSeries, pSeries other	S - W - E	http://www.redhat.com
SCO Group (Caldera)	The SCO Group Lindon, Utah	Intel Itanium	S - W	http://www.sco.com
SuSE	SuSE Linux AG Nuremberg, Germany	Alpha, Intel, Itanium/x86-64, PPC, Sparc, zSeries, S/390 iSeries, pSeries other	S - W - E	http://www.suse.com
Turbolinux	Turbolinux, Inc., Tokyo, Japan	Intel, Itanium, S/390	S - W - E	http://www.turbolinux.com
Yellow Dog	Terra Soft Solutions, Inc. Loveland, CO	PPC	S - W	http://www.yellowdoglinux.com

1.16.1 Yellow Dog Linux

Yellow Dog Linux (YDL) is a Red Hat and RPM-based distribution for the PowerPC architecture. There are two versions:

- ▶ **YDL Champion Server:** The version designed for Internet, intranet, development, and other mission-critical environments that require the most out of their operating system.

- ▶ **YDLGone Home:** The version for the home or small office. It is ideal for ISPs, network servers, and high performance computing.

What's in the distribution?

Yellow Dog has solutions for desktop and clustering.

Black Lab Cluster Management is Terrasoft's cluster product. Black Lab v2.1 is a cluster build and management application for High Performance Computing (HPC) clusters running Yellow Dog Linux. It offers single-click installation and configuration, automated updates through apt-get, and a graphical user interface.

Yellow Dog Linux for PowerPC is the desktop product. Yellow Dog Linux includes a graphical installer; dual-boot using BootX or yaboot; GNOME including Nautilus & Evolution, KDE, the "Liquid" theme default CUPS print system; apt-get for RPM; OpenOffice; Mozilla; Galeon, Evolution; AbiWord; and Netatalk.1. It also includes MAC-on-Linux (MOL), which is software that runs the Mac operating system within Linux via the GNOME or KDE desktop.

You can learn more about Yellow Dog Linux on the Web at:

<http://yellowdoglinux.com/products>

1.16.2 Redflag Linux

Redflag Software Co., Limited first appeared in August 1999, when it was created by the Institute of Software at the Chinese Academy of Sciences. Financial help came from government-owned Shanghai NewMargin Venture Capital. Redflag came out with an integrated desktop office suite called Chinese 2000. It joined the GNOME foundation with the intent to localize GNOME for the Chinese population. It has training and certification programs, such as Redflag Certified Engineer (RCE), and several applications.

Redflag Linux is the government-approved operating system for desktop PCs in the People's Republic of China. It won the Chinese government contract at the start of 2002 and the contract for the Beijing Municipal Government.

What's in the distribution?

Redflag has a variety of server products. These include Red Flag Cluster Server, Webmail Server, Function Server (a Web server), Database Server, High Availability Server, and Linux Server. The Database Server supports the database products Oracle 8i/9i, IBM DB2v7.1/v7.2, Informix® 2000, Sybase v11.9/v12.5, PostgreSQL, and MySQL. This product is targeted for enterprise-level applications.

The desktop version of Redflag includes applications ranging from Internet access, image processing, and multimedia to entertainment and games. Combined with an Office suite, this distribution is capable of processing Chinese content editing and printing.

You can learn more about Redflag on the Web at:

<http://www.redflag-linux.com/eindex.html>

1.16.3 Debian Linux

The Debian distribution is a non-commercial distribution whose goal is to effectively compete in the commercial market. This is one reason it is favored by many developers. Debian comes from the names of the creators of Debian, Ian Murdock, and his wife, Debra.

Debian was started in August 1993 as a new distribution that would be made openly, in the spirit of Linux and GNU. Most of the basic operating system tools come from the GNU project, hence the name GNU/Linux. Debian GNU/Linux provides more than a pure operating system.

It is a free distribution and the largest in terms of packages (has over 9,000 packages). The distribution supports multiple languages including French, German, Italian, Japanese, Portuguese, Spanish, Catalan, and Danish. It also supports several platforms including Intel, alpha, pa-RISC, PowerPC, and sparc64.

What's in the distribution?

As mentioned earlier, there are over 9000 packages available in the Debian distribution. You can install the standard applications for graphics, Web browsers, word processors, desktop managers, and so on. You can search the contents of these packages on the Debian Web site at:

<http://www.debian.org>

1.17 Going nuts with kernels

The kernel is the central module of an operating system. It is the part of the operating system that loads first, and it remains in main memory. The kernel is responsible for memory management, process and task management, I/O, and disk management. It manages files and is the part of the operating system that interfaces with hardware devices. Figure 1-5 shows the Linux structure with the kernel sitting between the shell and the hardware.

The open source portion of Linux that will *always* remain free is the kernel. Hundreds of developers submit their proposed changes to the kernel on an on-going basis. Linus Torvalds and his committee must approve these changes for them to be incorporated into the next kernel version.

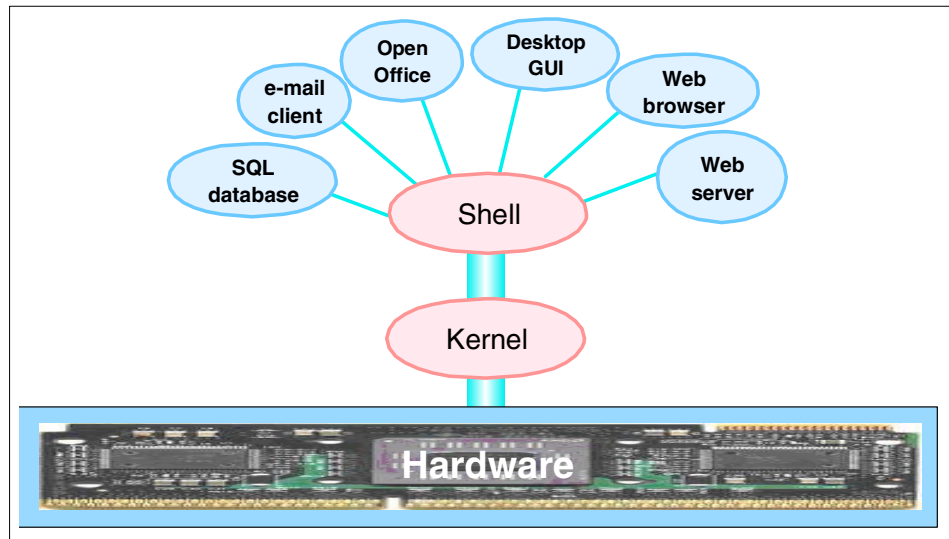


Figure 1-5 Linux structure

Linux kernel versions are divided in two series: developmental (odd series, for example, 2.3.xx or 2.5.x) and production (even series, for example, 2.2.x, 2.4.x, and so on). The developer series are fast moving versions for testing new features, algorithms, device drivers, and so on. Production or stable kernels have a well-defined feature set, a low number of known bugs, and tested drivers. They are released less frequently than the development kernels.

GNU/Linux distributions are usually based on chosen stable kernel versions, but not necessarily the latest production version. A different version number of the kernel appears for various platforms (SPARC, Alpha, PowerPC, and so on.)

You can download the latest Linux kernel from:

<http://www.kernel.org>

Another important area of Linux decision making is for workload performance related to scaling, which is in turn determined by the Linux kernel level. Table 1-3 shows the scaling by kernel.

Table 1-3 Linux scaling

Number of CPUs	Linux kernel	Example distributions
One to four	2.4	RHAS 2.1 Red Hat 7.3 SuSE SLES7
One to eight	2.4.18+	UL 1.0 SuSE SLES8 RHAS 8.0
One to sixteen	2.6	Second quarter 2003

Do not rely on scaling alone to determine increases in application performance. The way in which the workload performs with the whole system is more important. Adding more processors may not gain as much improvement in performance as achieving a balanced system. Enterprise distributions also have additional advanced functions for scaling.

SMP scaling and Linux are important techniques for clusters and multinode scaling. They include network load balancing, grid computing, Beowulf, OpenMosix, etc. The “free” per node software licensing of Linux can lend itself to a more cost-effective multinode solution.

For more information, visit the following sites:

- ▶ Load balancing:
http://www.lcic.org/load_balancing.html
- ▶ Grid computing
<http://www.gridcomputing.com/>
<http://www.ibm.com/grid/>
- ▶ IBM Redbooks grid portal
<http://publib-b.boulder.ibm.com/redbooks.nsf/portals/OnDemand>
- ▶ Beowulf:
<http://www.beowulf.org/>
- ▶ OpenMosix:
<http://openmosix.sourceforge.net/>



Open source software

“Open source” refers to any program whose source code is made available for use or modification as users or other developers see necessary. Historically, the makers of proprietary software generally did not make the source code available. Open source software is usually developed as a public collaboration and made freely available.

This chapter introduces you to some key concepts of open source software. It also helps you to understand the relationship of Linux and IBM with open source.

2.1 Creating the code and protecting it

Any discussion of open source software requires an explanation of how software is taken from a more human-readable form (source code) and converted into execution code (binary or object code). It is also important to define the various types of licenses that serve to protect intellectual capital in the form of computer software.

2.1.1 Source code

Most computer programming is done in one or more languages that are abbreviated, compact, and highly structured versions of English. This language is later converted to a form that a computer can understand in a process called *compiling*. As a result, this language in its original form is called *source code*. Example 2-1 shows source code from the Java programming language.

Example 2-1 Java source code example

```
import javax.swing.*;
import com.sun.java.swing.*;
import java.awt.*;
public class HelloSwingApplet extends JApplet {
public HelloSwingApplet() {
    getRootPane().putClientProperty("defeatSystemEventQueueCheck",
                                    Boolean.TRUE);
}
public void init() {
    JLabel label = new JLabel(
        "You are successfully running a Swing applet!");
    label.setHorizontalAlignment(JLabel.CENTER);
    label.setBorder(BorderFactory.createMatteBorder(1,1,2,2,Color.black));
    getContentPane().add(label, BorderLayout.CENTER);
}
}
```

Compiling the code

The source code, after it is completed, is then submitted to a *compiler*. A compiler is a software program that reads the special instructions that make up the source code and creates a file that can be executed by a specific type of processor.

Because computers read their code using binary numbers, this file is said to be a *binary file*. The process of compiling is similar to a cook following a recipe to create a meal. In compiling, many of the structures and labels in the code are removed, just as in the process of preparing a meal, where many of the original food forms are altered beyond simple recognition.

Binary code

The output of the compiler is in a form made up of binary numbers. Example 2-2 shows a section of a compiled program.

Example 2-2 Binary numbers from a computer program

```
00000101 00011111 00001000 00000000 00000000 00000000 11010010
01100001 01100000 01100001 01001000 01011000 01000000 11010010
01000001 11100000 10111000 01101000 11010010 00000011 01000000
11000000 01001000 01011000 11110000 01000000 00001100 00000101
00000100 11000010 00000000 00000000 01000111 11110000 10111000
01011000 10110000 11000001 00111000 01000111 11110000 10110110
```

Because the compiler removed many labels, instructions, and forms from the source code in creating the binary code, it can be difficult for a human to determine the nature or structures of a program simply by reading the binary output of the compiler. Most software, especially commercial software, is shipped only in a binary format. The source code is heavily protected from publication. This preserves the methods and technologies of the software creator, who typically invests a significant amount of money in their development.

2.1.2 Software licenses

Most software is published with copyright protection, similar to a book. Unlike a book, however, software typically adds an additional layer of intellectual property protection in the form of a license.

Licenses serve to specify, clarify, and restrict the legally permissible use of a software product. Most licenses require the purchase of rights to use the software. Figure 2-1 breaks down the various software licenses into several categories.

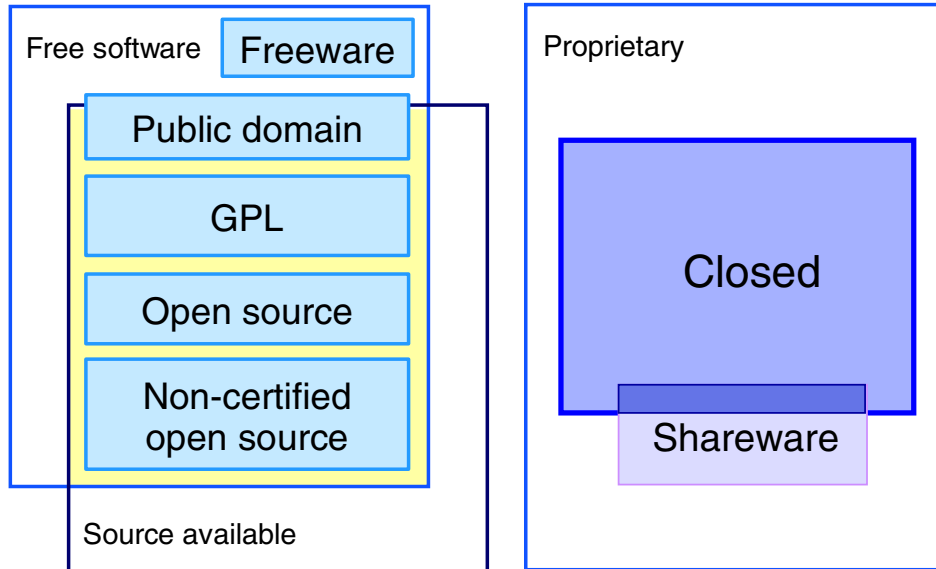


Figure 2-1 Software license categories

Proprietary software includes the following two overlapping subcategories:

- ▶ **Closed software:** This is any software where the source code is not available for review or modifications. However, it can be released as shareware. This is the traditional model of most software.
- ▶ **Shareware software:** This is a common license for many popular programs and applications, including some enterprise-level software. Shareware licenses allow the use of an application, without charge, for a limited time, primarily to allow an end user to evaluate the software. It is common practice to limit the use or functionality of a software program for the duration of the evaluation period. At the end of the evaluation, a user is either prohibited from using the program or is given the option to purchase special code that allows full supported use of the program.

Free software overlaps heavily with open source software, but not freeware. *Freeware* is any software license that allows you to use the software at no charge and with no expectation for payment in the future. Software issued under a freeware license traditionally does not include the source code. It is usually submitted by individuals or small organizations working on small programs and projects. Let's see how open source software fits into this picture.

2.2 The origin of open source

The label “open source” is applied to a philosophy of intellectual property that is the basis of several software licenses (and a general license category formally called open source) and a model of software development. Most proponents think of open source as concept that asserts that the user of a software program should have the right to see and modify the source code that was used to create that program.

The open source “movement” has its origins in the development and licensing of UNIX systems in the 1970s. AT&T, the creator of UNIX, liberally licensed UNIX systems to U.S. research universities to take advantage of any developments made to the operating system by university workers. Because many of these universities were connected to each other through ARPAnet, the precursor to the Internet, individual developers could share improvements and troubleshoot software problems. They could generally participate in a peer-to-peer environment that crossed geographical and organizational lines.

The UNIX system became a commercial-only operating system in 1984. This act ended the liberal code-sharing license that allowed developers to take advantage of each others code improvements. As a result, an attempt was made to continue the tradition of code sharing, most publicly by Richard Stallman, who helped start and still heads the Free Software Foundation (FSF).

2.2.1 The Free Software Foundation

According to the Free Software Foundation, free software allows users to freely run, copy, distribute, study, change, and improve the software. It give users the following types of freedom:

- ▶ **Freedom 0:** Run the program for any purpose.
- ▶ **Freedom 1:** Study how the program works and adapt it to your needs. You must have access to the source code for this.
- ▶ **Freedom 2:** Redistribute copies to help others.
- ▶ **Freedom 3:** Improve the program and release your improvements so the whole community benefits. Again, you must have access to the source code for this.

You can learn more about free software on the Free Software Foundation Web site at:

<http://www.fsf.org/philosophy/free-sw.html>

The Free Software Foundation believes that denying access to the source code of a software program is morally wrong. Because of this, it created the General

Public License (GPL). This license forces the free use (as previously defined) of any software protected by the GPL. A distributor of GPL software can charge to distribute the software code or support services. However, the users of the software can modify or use it in any way without paying the copyright owners for that capability.

GPL-protected software is different than software in the public domain for the simple reason of derivation. Code that is based on public domain software can be copyrighted and closed, while derivatives of GPL software are required (by the GPL) to remain open. The GPL focuses on freedom of use and not freedom from being charged for the software.

2.2.2 The compromise of open source software

GPL was the first license to publicly protect access to software source code. However, most commercial organizations felt that the license did not provide enough protection for their investments in software development. Other organizations were more concerned with controlling the integrity of their applications and limiting the publication of derivative works that might fragment or compete with their control of the original program.

As a result, the 1980s saw the creation of many similar-themed, but less restrictive licenses such as the MIT X license or the BSD license. In 1998, when Netscape announced that they were going to release the source code of their Navigator Web browser to the public, a viable alternative to the GPL was created.

The open source initiative

The phrase “open source” was first proposed during a meeting in February of 1998 of well-known Linux proponents including John “maddog” Hall (from Linux International), Sam Ockman (a member of the Silicon Valley Linux User’s Group), and Eric Raymond (author of *The Cathedral and the Bazaar*). Netscape asked Eric Raymond to help with the process of releasing and managing the code. He wanted to use the opportunity to help push the open source idea and development model into the commercial arena.

As a result of that meeting, the open source initiative was created. This not-for-profit group created a formal description of open source called the Open Source Definition (OSD). It also created a certification process for any license that attempted to bill itself as open source.

Note: Most popular open source software licenses are certified as OSD compliant. You can find a list of them on the Web at:

<http://www.opensource.org/licenses/>

The OSD outlines use of the software based on the following criteria:

- ▶ **Free distribution:** You can sell or give away open source software as part of a greater software package that includes other programs from any number of sources. This eliminates the possibility of the original license imposing redistribution fees or royalties.
- ▶ **Source code:** You must distribute open source software with the source code. Or you must make it easy to obtain the source code without an unreasonable reproduction charge. You must present the code in a common and non-misleading format. If source code is not easily accessible, it is more difficult to improve or evolve the code.
- ▶ **Derived works:** Anyone is allowed to create modifications to the original code and distribute those modifications under the same terms as the original software license dictates. This allows the cycle of improvement and development to continue, while preserving the control of the original developer of the code.
- ▶ **Source code integrity:** While a software license cannot prohibit changes to the source code, it can limit how that modified code is distributed. The most common restriction is to specify that any changes must be distributed as separate patches that can be integrated into the original source code (or ignored) at the discretion of the user. Another method of achieving the same aim is to force modified code to carry a different name or version from the original. Both restrictions serve to establish the responsibility for software changes while preserving the brand or product identity.
- ▶ **Personal discrimination:** Open source licenses are not allowed to restrict any individual or group of individuals from using their software. However, they can warn users about legal restrictions or obligations. The restrictions cannot be built into the software itself.
- ▶ **Use discrimination:** An open source license cannot restrict how the software can be used. This means that any industry or endeavor can take advantage of the software.
- ▶ **License distribution:** All rights granted by an open source license must be without provision. That is, no further agreements (such as a non-disclosure agreement) or licenses can be required for consideration in compliance with the original license.
- ▶ **Non-product specific:** The permissions granted as part of an open source license cannot depend on the code being distributed as part of a package or group of software. Modifications to a specific licensed product can be distributed without including the entire package (or additional products) that made up the original distribution.
- ▶ **Non-restrictive:** An open source license cannot make demands of other software that is distributed with the original program. For example, a license

cannot dictate that all programs distributed with it (as part of a package or media, such as a CD) must also be open source. This protects commercial interest by allowing the bundling of commercial software with open source software.

- ▶ **Technology neutral:** This point aims to remove the requirement that a user performs a specific action (other than installation and use) to accept the open source license. This allows you to distribute the software through multiple platform and media types without requiring a graphical user interface (GUI) window option, for example, to accept the license.

You can learn more about the OSD on the Web at:

<http://opensource.org/docs/definition.html>

These ten points were an attempt to bridge the gap between an apparent philosophical dispute between two diametrically opposed licensing schemes, the forced freedom of the GPL on one side, and the complete control and secrecy of most typical closed source software licenses on the other. The core issue is about licenses. It is also about which software development methodology and process is superior, closed and centralized or open and distributed. The implication is about revenue of a finished product and control of the development process itself.

This is the genius of the open source movement. It is a compromise of both ends of the development and licensing spectrum. It allows the original creator of a program to have some control over how the software develops and is presented to the public. At the same time, individual developers and users retain the ability to manipulate and share improvements of the source code for their own needs.

These theoretical arguments originally were not enough to convince businesses to change established processes and culture and adopt a more decentralized and open developmental model. The success of Linux gave organizations an effective example of how the compromise between distributed development and centralized control could (and did) function successfully.

2.3 Linux and open source

Linus Torvalds began working in 1991 to create a better terminal emulator for his PC. Within a few months of working in his spare time, he created a small but workable clone of the UNIX system kernel. He published this (with the source code) on an Internet news group that was dedicated to operating system kernels. Several news group participants immediately offered ideas (and actual code) as improvements. This began a process that continues today of individuals and

organizations providing ideas, new code, and patches that are evaluated by Linus (and a select group of volunteers) for inclusion into the kernel.

You can find a well-researched overview of the Linux development process on the Web at:

http://firstmonday.org/issues/issue5_3/kuwabara/index.html

2.3.1 The GNU utilities

A crude operating system by itself was not enough to convince more than merely interested to take it seriously. Nor was Linux the first project that relied on an open source development model. That title belongs to the Free Software Foundation. Under Richard Stallman's direction, it created the GPL and released numerous utilities under the license as part of the GNU project. These included the GNU C Compiler Collection (gcc), which currently contains front ends for C, C++, Objective-C, Fortran, Java, and Ada, as well as language libraries (libstdc++, libgjc, etc.). It is a high-performance, multi-platform compiler that remains the standard by which all other compilers are judged.

You can locate the GNU C Compiler Collection on the Web at:

<http://gcc.gnu.org>

Because Torvalds decided to release Linux under the GPL, it was natural to bundle the GNU utilities with it. In fact, the Free Software Foundation argues that because so much of Linux was built on, and distributed with, the GNU utilities, Linux should be called *GNU/Linux*. While little argument exists against the case made by the Free Software Foundation, few have adopted this label, preferring the simpler name of *Linux*.

2.3.2 The problem of complexity

As the development of Linux continued, it was inevitable that significant complexity would emerge. This was especially true as the kernel grew to over one million lines of code. This code was mostly maintained by part-time volunteers (estimated to exceed a total of 40,000 in number) worldwide, who communicated with each other almost exclusively through e-mail.

This complexity is managed by Linux developers using a version control program called the Concurrent Versions System (CVS). CVS works by recording all the changes that occurred in the source code of an application. As a result, bug or error detection is simplified because it is possible to track the exact evolution of events that led to a particular software error.

You can find the CVS on the Web at:

<http://www.cvshome.org>

CVS also works to enable software coordination. Each software developer has their own unique work space. This isolates them from the work of others until their code is ready to be merged, a process that is aided by CVS.

But CVS is only part of the answer to the problems created in distributed development. For example, coordination and communication between developers is still necessary to determine when and how code should merge or branch, whether multiple changes within a file or across multiple files will functionally or logically conflict, and how to resolve those conflicts.

CVS is a tool to aid management and control of the development process, but it does not create or manage that process by itself. How, then, does the chaotic model ascribed to open source software allow for something so complex as an operating system, let alone an operating system that is stable enough to be considered for deployment in critical enterprise applications?

2.3.3 The benevolent dictator

The development of Linux is mostly distributed and decentralized. However, Linus Torvalds himself still exerts authority and ultimately maintains coordination, along with an informal group of trusted, knowledgeable volunteers colloquially called the “Inner Circle”.

Linus tends to focus on developing the beta version of each kernel version. After he considers a version stable enough for general release, he turns over the responsibility of maintaining that stable version to Alan Cox, another well-known programmer. While Linus turns his attention to the next beta kernel, Alan is left with wide discretion about how the current kernel is maintained and updated.

This same pattern is seen with other members of the Inner Circle. That is, an individual is in charge of development and maintenance of an area they originated or took over from another original developer.

Because of the sheer complexity and quantity of the kernel code, its components, and the various patches and feature requests, Linus has encouraged a system where any submitted code is first sent to the individual who oversees the area in question. As a result, Linus’ principles and goals in regard to the development of Linux are still maintained, while significant authority is delegated to the individuals who oversee specific subareas.

In this fashion, Linux follows more of an open source model of development. It combines the innovation that comes from forced free code imposed by the GPL,

yet retains the vision of the kernel's original author through a delegation of authority.

2.4 IBM and open source

IBM is committed to open source as both a license and a development model for several reasons:

- ▶ IBM clients and partners have requested open source software (including Linux) support for all IBM platforms, products, and solutions.
- ▶ Open source software, with its wide distribution and use, typically becomes an industry standard.
- ▶ Innovation within the open source community typically occurs at a higher rate and volume than in closed-source communities.

2.4.1 IBM and open source licenses

IBM's first open source license was the IBM Public License (IPL), introduced in 1999 and certified by the OSI in the same year. As part of the Eclipse release in 2001, IBM updated the IPL and renamed it the Common Public License (CPL). The OSI certified the CPL in May 2001 as an open source license.

For more information about the CPL, see:

<http://www06.ibm.com/developerworks/library/os-cplfaq.html>

2.4.2 IBM and open source software

IBM is a major contributor to the open source community, with over 250 developers worldwide working full time on several initiatives. Such initiatives include:

- ▶ **Linux Technology Center (LTC):** The LTC works to enable the enterprise capabilities of Linux through development and contribution of technology, utilities, tools, and code.

You can learn more about the LTC on the Web at:

<http://oss.software.ibm.com/linux>

- ▶ **Open source software support:** IBM is an active supporter of projects and software including:
 - Open Source Cluster Application Resources (OSCAR)
 - Open Source Developer Lab (OSDL)
 - GNOME

- KDE
- Open Source Initiative (OSI)
- Free Standards Group
- USENIX
- Linux high availability
- OpenLDAP
- USB
- PCI hot plug
- Advanced Power Management (APM)
- PPC-32
- PPC-64
- Stream Control Transmission Protocol
- Free Standards Group (Linux Standard Base)
- Samba

Learn more about IBM's involvement in open source software on the Web at:

<http://www.ibm.com/developerworks/oss/>

- ▶ **Eclipse:** This is an open source software integrated development environment (IDE) sponsored by IBM that acts as a development infrastructure platform. Through extensions (called *plug-ins*) to the base IDE, Eclipse can provide nearly any type of capability, including code versioning functionality and even syntax highlighting. IBM's WebSphere Studio Workbench is a commercial product that is built on Eclipse.

You can find Eclipse on the Web at:

<http://www.eclipse.org>



Linux, UNIX, and Windows systems

The majority of enterprise Linux implementations are found in corporate Internets or in the Internet infrastructure. Other notable uses include file and print server deployment, application development, messaging server deployment, data warehousing, embedded applications, thin client applications, and massively parallel scientific applications. Linux is becoming robust enough to find its way into engineering software development, the healthcare industry, the retail industry, and financial institutions.

Wall Street has adopted Linux in a big way. Major companies have been in the press, sporting their new implementation projects with Linux. These are not just departmental solutions, but large company-wide solutions indicating potential savings of 30 to 1. These results will certainly attract the attention of Chief Financial Officers (CFOs) and Chief Information Officers (CIOs) looking to improve their business with leading edge solutions and major spending reductions.

Linux is also taking hold in the public and government sectors as the IT industry tries reduce expenditures. How does it compare to the traditional UNIX systems vendors or Windows? Does Linux provide a real advantage? This chapter provides a high level overview of some of the key characteristics of the Linux, UNIX, and Windows operating environments.

3.1 Features, functions, and technologies

With all of the features, functions, and technologies available for and within the various operating environments, it is difficult to provide a side-by-side comparison or accurate rating or ranking. The scope, depth, and breadth of the feature, function, or technology varies by vendor or distributor.

We provide a high-level view of some of the major features, functions, and technologies that can be part of a purchasing decision. Each area that is listed can easily be broken down into several subcomponents and analyzed. Add-ons, additional products, and independent software provider (ISV) or third-party solutions can significantly change the perspective of the feature, function, or technology.

We used the following guidelines to create a rating system of the high-level features or functions of Linux, UNIX, and Windows system:

- ▶ The Linux distributions were combined into a single entity, using the most recent standard distribution from the distributors. Only the major Linux distributions were analyzed. The intent was not to provide a side-by-side analysis of Linux distributions. Nor was the intent to provide a best-of-breed Linux distribution, taking the best features of each distribution.
- ▶ The UNIX operating environments were combined into a single entity, using the most recent versions. Only the major players or vendors in the UNIX operating environment, such as Sun Solaris, Hewlett Packard (HP) HP-UX, and IBM AIX, were analyzed. Again, the intent was not to provide a side-by-side analysis of UNIX distributions. Nor was the intent to provide a best-of-breed UNIX distribution, taking the best features of each distribution.
- ▶ The Windows operating environment that was analyzed was Windows 2000. Windows NT and the Windows.NET were not included. Windows is the most prevalent operating system in the world today. Of the Windows versions that are running, approximately 70% are Windows 2000.

3.1.1 Description of the ratings

The rating given to each feature, function or technology is shown in Table 3-1. A three-scale rating system was used to keep the rating as simple as possible.

Table 3-1 Rating system

Rating	Description
+	The function or feature is available, robust, and mature. It may be industry leading.
=	The function or feature may be available, but not robust. The function or feature may require a third-party solution.
-	The function is not available or may not be robust enough to be considered good enough for enterprise deployment.

3.1.2 Methodology for the feature, function, and technology rating system

To explain the rating system, you can use partitioning as the feature technology to be rated. For most enterprise or small and medium business (SMB) environments, partitioning is one of the technologies considered when making a purchasing decision, especially in a server consolidation environment. Partitioning, as a technology, can easily be broken down into several subcategories. These may include dynamic resource allocation, maximum number of partitions supported, processor granularity, memory granularity, or heterogeneity of operating environments supported. However, this level of granularity is not done here.

Various partitioning solutions are available for UNIX operating environments. Sun's UNIX partitioning solution provides hardware partitioning (dynamic system domains) and resource partitioning (containers). HP's UNIX solution provides hardware partitioning (nPars), software or logical partitioning (vPars), and resource partitioning. IBM's UNIX partitioning solution provides software or logical partitioning (LPAR). Although there are UNIX solutions, which vary in scope and depth, the UNIX partitioning solution is integrated within the UNIX environment and can have distinct advantages. Therefore, a "+" rating was given to UNIX systems.

The Microsoft Windows 2000 operating environment has no integrated partitioning solution, although Windows takes advantage of the hyperthreading technology built into the Intel architecture. Third-party solutions, such as VMware's ESX Server and GSX Server, are available. A "-" rating was given to Microsoft Windows.

Linux operating environments have no integrated partitioning solution, although third-party solutions, such VMware's ESX Server and GSX Server and SWSOft's Virtuozzo, are available. A "-" rating was given to Linux operating environments.

3.1.3 Linux, UNIX, and Windows: A side-by-side comparison

Table 3-2 compares some of the key characteristics of the Linux, UNIX, and Windows operating environments. This comparison uses the rating system that was mentioned in 3.1.1, "Description of the ratings" on page 61.

Table 3-2 Comparison of Linux, UNIX, and Windows systems side by side

Feature, function, technology	Linux	UNIX systems	Windows	Description, comments
Uses industry standard processor	+	-	+	IA32, IA64 or equivalent; UNIX proprietary
64-bit or greater support	=	+	=	Supports 64-bit
Device support	=	+	+	Supports most devices; Linux fairly current
Scalability and performance	=	+	=	Scales and performs well
Partitioning	-	+	-	Integrated partitioning; need third-party solution
Clustering	+	+	+	Integrated clustering
Guaranteed high availability	=	+	=	High availability platform, dynamic; good solutions available
Recovery	-	+	=	Dynamic, fault, data; good solutions available
Web applications	+	+	+	Web, proxy, time servers; strength of Linux
Protocols	+	+	+	HTTP, SSL, SMTP, FTP, etc.; Linux is fairly current
Integrated programming language	+	+	+	C, C++, Perl, Java, etc.; IDEs for Linux available
Linux, UNIX, Windows clients	+	+	+	Linux desktop is exponentially improving
Integrated print, file server	+	+	+	Strength of Linux, good solutions available

Feature, function, technology	Linux	UNIX systems	Windows	Description, comments
Mail server	+	+	=	Cost of Microsoft Exchange
Mail client	+	+	+	Not as robust as Microsoft Exchange
Web application server	+	+	+	Strength of Linux, good solutions available
Systems management	=	+	+	CL, GUI, Web, administration; not integrated; lags but better
Heterogeneous systems management	=	=	+	Linux, UNIX, Windows; good solutions available
Storage management	=	+	+	Configuration, SAN, NAS, spares; good solutions available
Software management	=	=	+	Version, package, configuration; good solutions available
Workload management	-	+	-	Resource, cluster, etc.; need more solutions
Security	+	+	=	Authentication, encryption, SSL; good, but several issues

From Table 3-2, you can see that a typical Linux distribution provides a good level of support for many features, functions, and technologies. However, a typical Linux distribution may not be quite as robust as a UNIX system. Linux is tough enough for many enterprise environments. New features, functions, and technologies are literally being added on a daily basis.

3.1.4 Linux, UNIX, and Windows: Strengths and weaknesses

Table 3-3 lists some of the general strengths and weaknesses of the Linux, UNIX, and Windows operating environments.

Table 3-3 *Strengths and weaknesses of Linux, UNIX, and Windows systems*

	Strengths	Weaknesses
Linux	<ul style="list-style-type: none"> ▶ Total cost of ownership (TCO) ▶ Hot technology, industry wide endorsement ▶ Number of environments and applications ▶ Open source ▶ Growing base of Linux knowledge and developers ▶ Security ▶ Flexibility 	<ul style="list-style-type: none"> ▶ System management capability and integration ▶ Lower number of enterprise applications available on non-Intel based hardware ▶ Market perception as a desktop-only environment ▶ Maturity of 64-bit processor support ▶ High-end performance features on 64-bit ▶ Channel structure
UNIX	<ul style="list-style-type: none"> ▶ Disaster recovery, reliability ▶ Web server and Internet capability ▶ Secure networking ▶ System management capability and integration ▶ Partitioning ▶ Integrated clustering ▶ Scalability and performance 	<ul style="list-style-type: none"> ▶ Tied to proprietary processor architectures ▶ Heterogeneous functions and interoperability ▶ Total cost of ownership
Windows	<ul style="list-style-type: none"> ▶ System management capability and integration ▶ Integrated Web server environment ▶ Support for Web services ▶ Large number of enterprise applications available ▶ Familiar environment due to desktop deployment 	<ul style="list-style-type: none"> ▶ Total cost of ownership ▶ Security ▶ Heterogeneous functions and interoperability ▶ Maturity of 64-bit computing ▶ Proprietary technology

3.2 Cost of ownership

This section discusses TCO as it relates to the planning and implementation of Linux systems. It also contrasts TCO with return on investment (ROI).

3.2.1 Total cost of ownership

Vendors, consultants, and market research firms have done numerous total cost of ownership studies. As shown in Figure 3-1, most research seems to agree that there is a TCO advantage for Linux in long-term deployment and operation. Most research illustrates some type of savings with the deployment and operation of a Linux environment, although the amount of savings can differ greatly.

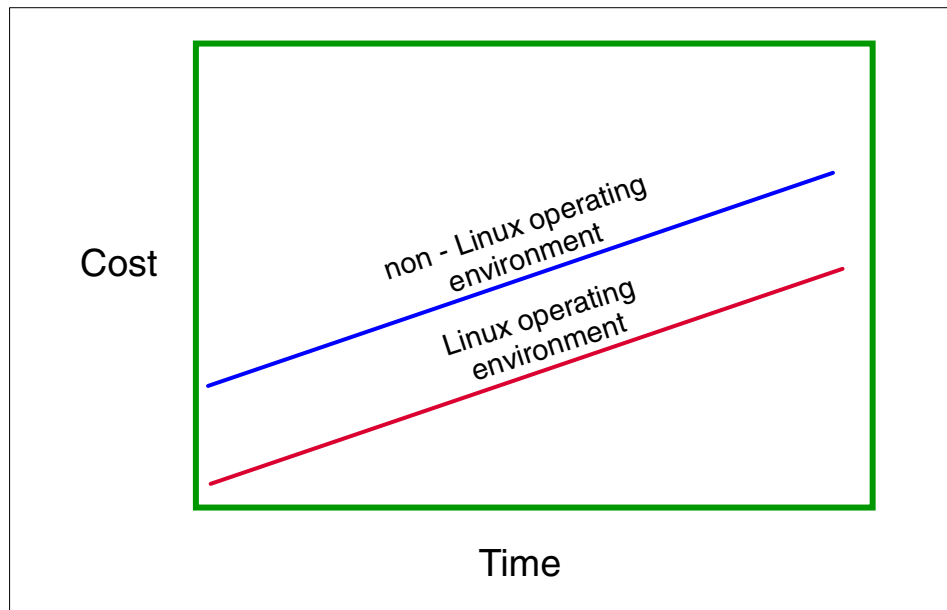


Figure 3-1 Typical Linux TCO study showing cost over time

Many factors that can skew a simple TCO study. Such factors include:

- ▶ A comparison of the Linux and non-Linux operating environment; also which Windows operating environment, UNIX operating environment, and Linux distribution
- ▶ Whether a list price or deeply discounted price was used
- ▶ Whether new hardware was configured into the TCO model
- ▶ Whether a desktop deployment, enterprise deployment, or SMB deployment was used in the TCO model

- ▶ The type of workload used in the TCO model, perhaps including collaborative, Web server, file server, or database server
- ▶ The software and applications used in the configuration
- ▶ Assumptions about administrative capabilities of the chosen configuration and the number of administrators needed
- ▶ The time frame the TCO model covers
- ▶ Whether support and service contracts were configured into the model
- ▶ Whether consideration was given to training a user when transitioning to a different operating environment and the short term loss in productivity
- ▶ The sponsor of the study (to some degree)

The cost savings of Linux may not always be as clear and simple as many open source advocates have proclaimed. You must factor in any hidden costs into the TCO model. Such hidden costs may include the skill level of the IT staff, the complexity of applications, and the cost to migrate from another environment.

The purchase price of hardware and software alone provides only a small indication of the total lifecycle cost for an IT solution. The TCO life cycle includes all staffing costs associated with acquiring, maintaining, and removing an IT asset from an organization.

There are numerous models to categorize the total cost of computing. You can break down the total cost of computing model into the following categories:

- ▶ Hardware
- ▶ Software
- ▶ Staffing
 - Development
 - Support

You can also break down the total cost of computing model into the following categories:

- ▶ **Buy:** Includes making the lease or purchase decision
- ▶ **Use:** Setup, training, use, and technical support
- ▶ **Dispose:** Removing the asset, transfer data, and applications

Most TCO models don't include the cost of external communication, media, power consumption, floor space, etc. since they are somewhat similar. The key to any TCO model is the categories to be measured and the quality of the measurements.

You cannot calculate other key factors normally into a TCO model. Such factors may include complexity, scalability, security, manageability, frequency of software

upgrades, availability of applications, and reliability. Discounting can also play a major role in the TCO of a solution. In many instances when a complete solution is purchased from one vendor, deep discounting is applied.

There is also a productivity paradox with selecting an operating environment that is not immediately measurable in any of the TCO models. For example, will the business become more productive or make a better return on investment if the business uses a Linux, UNIX, or Windows operating environment? Does the decision to move to a Linux, UNIX, or Windows operating environment affect the basic business structure?

3.2.2 Operating environment and product TCO perspective

While it can be difficult to qualitatively analyze the TCO of a total solution that has many intangible aspects, you can analyze some of the tangible costs of the hardware and software. Table 3-4 identifies costs that you can analyze in a Windows versus Linux operating environment solution.

Table 3-4 Cost analysis comparison in a Windows versus Linux environment

Function, product	Windows solution	Cost	Linux solution	Cost
Base operating environment	Windows 2000 Advanced Server	\$\$\$\$\$	Red Hat, SuSE, Turbolinux boxed set	\$
Web server	IIS	0	Apache	0
Antivirus software	Norton Antivirs	\$	Squid, Virulator	0
Database	SQL Server	\$\$\$\$\$	Postgre SQL	0
Mail server	Exchange Server	\$\$\$\$\$	Sendmail, Postfix	0
Office suite	Office	\$\$\$\$\$	StarOffice, OpenOffice	0
Development environment	VisualStudio	\$\$\$	KDE	0
Desktop client	Windows XP Professional	\$\$\$\$\$	Red Hat, SuSE, Turbolinux	0
System management	MOM, SMS, AC	\$\$\$	Webmin	0

3.2.3 Findings

The following list outlines general, key findings with most TCO models that compare and contrast the Linux, UNIX, and Windows operating environments:

- ▶ For many businesses, the Linux operating environment has emerged as a viable alternative for enterprise computing. Using the Linux operating environment has the potential to lower IT costs. This is especially true for businesses that have the right mix of staff, skills, commodity hardware, software, and overall computing environment.
- ▶ To realize the full potential of the Linux operating environment, companies may need to pilot a Linux project and then build a vision or strategy for a longer term deployment.
- ▶ The cost of Linux and Windows operating environments on Intel platforms is highly competitive and low due to the acquisition cost of Intel architecture server hardware.
- ▶ Linux can be deployed with no or low licensing fees, particularly if Linux is deployed without the support services that normally accompany an enterprise-level product. This fact alone can account for a lower TCO.
- ▶ TCO may depend on the configuration and workload. However, with some key workloads, such as Web applications, collaborative applications, firewall, file server, or print server applications, Linux is likely to have a lower TCO.
- ▶ Most companies, whether large or small, operate in at least a dual operating environment.
- ▶ Staffing and skills are a key consideration.
- ▶ Linux can be used as a tool for negotiating the price of proprietary solutions in other operating environments.

The IBM @server offers the possibility for server consolidation. As explained in the following chapters, these servers have the capability of virtualization, from hundreds of virtual Linux servers on zSeries to tens on other IBM @server platforms. This provides additional TCO savings.

3.2.4 TCO versus ROI

While TCO is an important consideration, ROI is equally significant in the customer decision-making process. CFOs take a good look at the ROI they can expect to realize with new servers. While it may be useful to consider the overall total cost of ownership, today customers are interested in saving money in a tough economic environment.

ROI can be more effective than TCO as a measure of business value, especially with new, emerging technologies. For this reason, we expect to see a movement

of the community toward a more comprehensive ROI justification, where TCO is just a portion of the overall analysis.

It's coming to be viewed that ROI analysis is a higher quality approach for making business decisions since it's more inclusive of the risks and returns of making new IT investments. This is what CFOs are looking for and why Linux is becoming more accepted each day.

Numerous case studies have shown Linux to deliver outstanding ROI value. These benefits may derive from much greater deployment flexibility, much greater openness, much higher performance and functional benefits (derived from the prevalent deployment architectures described earlier), and universally accepted better reliability and security. Such emerging technologies as Linux typically excel in ROI analyses, which highlight the benefits of the new technology, one of which is cost savings.

Linux today contributes to business innovation through the following typical high-value deployment methods that provide TCO and ROI advantages:

- ▶ **Using appliance servers** to provide fixed, single-function solutions simplifies use and lowers support, operation, and maintenance costs. Single-function solutions are often deemed preferable for infrastructure workloads over adding software to a general-purpose server.
- ▶ **Consolidation of distributed workloads** onto centralized servers dramatically reduces the overall administration costs and delivers improved resource utilization.
- ▶ **Clustering Linux servers** for high-parallel or repetitive tasks simplifies administration (since tasks can be scripted and repeated across multiple servers). This greatly enhances performance. Customers are deploying Linux widely, often in high availability configurations, which are more suitable, reliable, and secure with Linux.



zSeries and Linux

This chapter provides a brief overview of the Linux solution that is available on the zSeries server. It includes information about the following topics:

- ▶ The various modes that Linux runs on the zSeries server
- ▶ The benefits of running Linux as a z/VM guest
- ▶ Networking options that are available for z/VM Linux guests
- ▶ Consolidation scenarios using Linux on zSeries
- ▶ Software solutions for Linux on zSeries
- ▶ Special IBM offerings for zSeries and Linux

4.1 S/390 and zSeries architecture

The heart of zSeries and S/390 hardware is the multichip module (MCM) that contains up to 20 processing units (PU). On other platforms, these are commonly referred to as central processing units (CPUs) or engines.

The core elements of any computer are the processing units. These contain the z/Architecture™ logic and functions that implement the architectural extensions of S/390 ESA architecture. As implemented in the zSeries 800 and zSeries 900 models, each PU actually has dual internal instruction processors.

Instructions are executed by both internal processors, in parallel, and the results are compared. If the results do not match, an instruction retry process is performed. This is all done automatically, by the PU, and is not visible to the operating system. The normal result of the dual processors in a PU is the execution of a single instruction stream.

We normally refer to a PU as a single processor and ignore the fact that there are really two parallel processors inside each PU. PUs are used in one of five ways:

- ▶ **Central Processor (CP) PU:** Used by the operating system for executing customer work
- ▶ **System Assist Processor (SAP) PU:** A System Assist Processor is a PU that runs the channel subsystem Licensed Internal Code to control I/O operations. One of the SAPs in a configuration is assigned as a Master SAP and is used for communications between the multichip module and the support element.

In LPAR mode, all SAPs perform I/O operations for all logical partitions. The zSeries 900 12-PU MCM models have two SAPs as standard. The 20 PU MCM models have three SAPs as standard configurations.

- ▶ **Spare PU:** This is a PU that is not enabled for any purpose. The system uses this PU to replace a failing processor, if necessary. If all four PUs (plus an SAP) are enabled, then there are no spare PUs. Spare PUs may be used for various upgrade options such as Capacity Backup (CBU) and Customer Upgrade on Demand (CUoD).
- ▶ **Integrated Facility for Linux (IFL) PU:** This method is restricted to running Linux or z/VM.
- ▶ **Integrated Coupling Facility (ICF) PU:** This method is used to run the Coupling Facility function for use in a Parallel Sysplex® environment.

An IFL is a processor reserved for Linux (or Linux under VM). The significance is that it cannot be used to run other operating systems, and its existence is not reflected in the system model number, MIPS rating, or other power ratings. The

system model, MIPS, or other power rating method has significant implications for software costs. Adding an IFL does not affect these costs, permitting the use of Linux without impacting other software costs.

Figure 4-1 shows how IFLs may be allocated to logical partitions (LPARs).




Figure 4-1 The relationship between IFLs and LPARs

Important: IFLs must be dedicated to LPARs running Linux or z/VM. They can be shared by multiple LPARs, provided the LPARs are running Linux or z/VM.

Also, an LPAR cannot contain a mix of IFLs and CPs.

An ICF is used only to run the Coupling Facility licensed code. It cannot be used to run normal operating systems. It is similar to an IFL in that its existence does not change the system model number (or MIPS rating) nor impact software costs for system.

IFLs and ICFs require the use of LPARs. If neither of these are used, a zSeries 800 system can run in basic mode (no LPARs). Note, however, that use of the new z/OS™ operating system package requires the use of LPARs.

This chapter discusses zSeries and Linux. Here are few useful links to additional information:

- ▶ VM and Linux for zSeries Resources
<http://www.vm.ibm.com/linux/>
- ▶ Linux for IBM @server zSeries
<http://www.ibm.com/servers/eserver/zseries/os/linux/>
- ▶ zSeries mainframe servers
<http://www.ibm.com/servers/eserver/zseries/>

Also several service offerings are available. Table 4-1 lists the platform-based information for services that are available for zSeries servers.

Table 4-1 zSeries Linux information matrix

Information	Location on the Web
Linux for zSeries applications	http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxisv.html
Linux for zSeries operation system download	http://www.ibm.com/servers/eserver/zseries/os/linux/dist.html
Linux for zSeries Development tools	http://www.ibm.com/servers/eserver/zseries/os/linux/ldt/
Linux for zSeries Library	http://www.ibm.com/servers/eserver/zseries/os/linux/press.html
Linux for zSeries Education	http://www.ibm.com/servers/eserver/zseries/os/linux/ed.html

4.1.1 Ways to run Linux on zSeries

There are three ways to run Linux on the zSeries server. Each way is explained in the following sections.

Basic mode

Linux can run on the entire machine, without any other operating system. In basic mode, you can use a single Linux image on the zSeries “bare metal”.

Logical partition

You can logically partition the zSeries hardware into a maximum of 15 separate LPARs. A single zSeries, for example, can host z/OS applications in one partition, VM and VSE applications in others, and Linux applications in additional partitions.

z/VM

A customer can run Linux as a virtual machine using z/VM as a hypervisor. z/VM provides virtualization of CPU processors, input/output (I/O) subsystems, and memory. A customer running z/VM can run hundreds of Linux systems on a single zSeries. With z/VM, for instance, a customer can offer a complete Linux server environment to each application developer and host production system all on the same zSeries server. We discuss z/VM and Linux in more detail later in this chapter.

The decision on how to run Linux on zSeries is typically made based on the expected workload:

- ▶ Linux workloads with a large memory footprint and requiring extensive processing power may be best suited for running in an LPAR.
- ▶ Few (if any) workloads require running Linux in basic mode.
- ▶ Running Linux as a z/VM guest offers the maximum flexibility. The number of Linux guests running under z/VM can be adjusted to meet workload requirements while maximizing usage of the zSeries hardware. Most customers choose this option to maximize their investment.

4.1.2 31-bit and 64-bit options

Linux solutions are available for the 31-bit and 64-bit environments. The option availability depends on the zSeries model and the Linux distribution.

Linux for S/390 (31-bit)

Linux for S/390 is a 31-bit version of Linux. It is currently available for S/390 Parallel Enterprise Server G5 and G6 and the Multiprise® 3000 processors. It also runs on zSeries models in 31-bit mode. Because of the 31-bit limitation, the addressable main storage is limited to 2 GB.

Linux on zSeries (64-bit)

Linux on zSeries supports the new 64-bit architecture on zSeries processors. This is all zSeries 800 and 900 models. 64-bit support eliminates the 31-bit storage limitation of 2 GB.

4.2 All about z/VM

This section explains the value, both business and technical, of running tens to hundreds of virtual Linux systems on z/VM.

Note: “VM” means *virtual machine* not virtual memory. The virtual machine concept found in the z/VM product provides the function-rich *virtualization technology* on the zSeries server. In the context of this chapter, the virtual machine concept and virtualization technology are interchangeable. z/VM provides much more functionality than partitioning of physical resources.

IBM has been investing in virtualization technology for over 30 years. This is a proven, function-rich technology base that has benefitted the Linux on zSeries environment since day one. Customers have run business-critical solutions on VM systems just as long.

Virtual machine technology was introduced with the IBM System/360™, the original mainframe, back in the mid-1960s. At that time, VM (known as CP-67) enabled users of S/360 systems to run multiple copies of operating systems on a single mainframe. The ability to run multiple operating systems made it possible to migrate to new releases and versions on the same mainframe system. A user could run a production workload on the old level of software, while performing the work required to migrate the environment to the new operating system.

Internally, IBM used VM to simulate the System/370™ architecture on a System/360 processor. VM would run on a real S/360 processor and create S/370™-mode virtual machines. This let IBM develop its suite of S/370-architected software before a “real” S/370 processor was available for testing.

Note: VM was also used to develop VM/XA, VM/ESA®, and the new 64-bit version of z/VM.

The ability of VM to simulate hardware and architectures today continues to give customers time-to-market advantages for their products and services.

Over the course of 30-plus years, IBM made investments in hardware, architecture, microcode, and the VM product itself to enhance the virtualization technology available with each successive line of mainframe computers. This “total system” investment is why the zSeries line of servers offers the industry's best-of-breed virtualization technology. Today z/VM Version 4.3 is the most current VM release.

The virtualization technology found in z/VM is unique. Some “virtual server” solutions consist of running multiple applications on a single operating system using virtual storage techniques to meet memory demands. z/VM, on the other hand, uses real resources at its disposal (processors, memory, I/O, network, etc.) to create virtual machines. Each machine can run its own independent copy of any operating system that supports the zSeries architecture (which includes Linux, of course).

Hardware architecture and facilities are presented to the “guest” system (which is “virtualized”) by z/VM in such a way that the guest operating system believes it is using a real system, not a virtual one. Figure 4-2 shows an example with multiple virtual machines, including three Linux machines, a z/OS, a VSE, and a CMS guest. The physical hardware is not seen by the guests. Instead, each guest sees the virtual machine presented by z/VM.

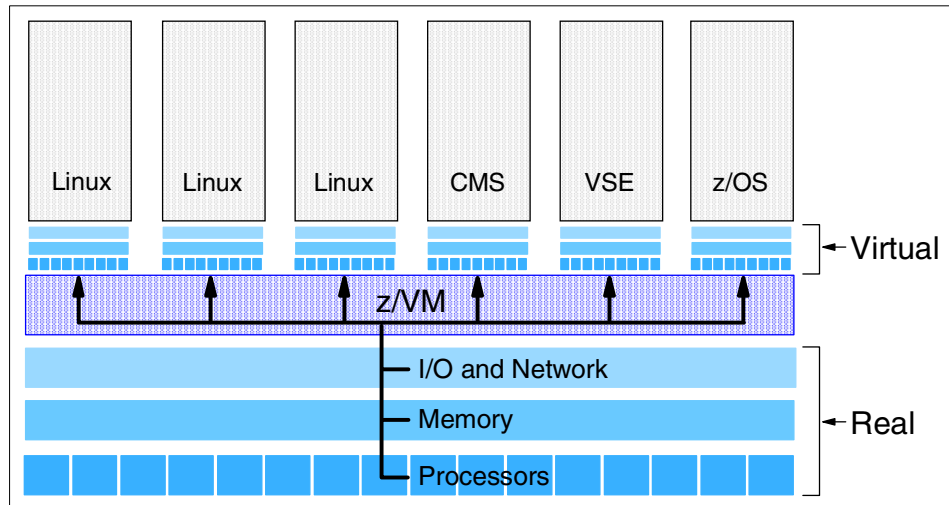


Figure 4-2 z/VM virtualization

A guest operating system can run unchanged on z/VM and not make any effort to detect whether it's running in a virtual machine. Optionally, it is possible to determine whether the “machine” is virtual. If it is, specific VM-only functions are available to guest systems.

z/VM can create a virtual machine environment that is functionally richer than a “real” environment. Examples of this include:

- ▶ Using data-in-memory techniques to transparently enhance guest system performance
- ▶ Adding virtual CPs to z/VM guests, which are seen as real CPs by the guest even when the number of virtual CPs exceeds the number of real CPs
- ▶ Simulating device and networking facilities that don't exist in the real mainframe system
- ▶ Sharing a single copy of a kernel among several guest systems

z/VM V4 requires G5 processor technology or better. This is found in the Multiprise 3000 server as well as the G5, G6, zSeries 900, and recently announced zSeries 800. z/VM V4 (all releases) can run on the IFL and standard processor engines. IFL engines are available on the Multiprise 3000, G5, G6, zSeries 900, and zSeries 800. The Linux-only model of the zSeries 800 uses IFL engines exclusively.

Note: An engine is merely another term for a processor.

4.2.1 Integrated Facility for Linux

Adding capacity to an existing mainframe to run Linux workload or workloads can be done cost-effectively using IFL engines. The addition of IFL engines does not change a server's model designation. Also, the cost of existing software is not affected by the presence of IFL engines.

Figure 4-3 shows a zSeries 900 Model 1C6 configured with three IFL engines. The engines are allocated from the set of available spare CPs on a zSeries server (including the G5/G6). A zSeries 900 Model 1C6 (a six-way, standard engine configuration) has 10 additional processor engines available for use as IFLs, standard processors, or ICFs.

In Figure 4-3, the customer purchased and allocated three IFL engines (IFL0, IFL1, IFL2). The cost of this z/VM environment is one-time charge (based on a three-engine pricing model) in addition to an annual service and support fee.

The presence of the IFL engines does not change a server's model designation. This is still a zSeries 900 Model 1C6. The software fees for the products running on the standard engines (CP0 to CP5) are not affected by the addition of the IFL engines.

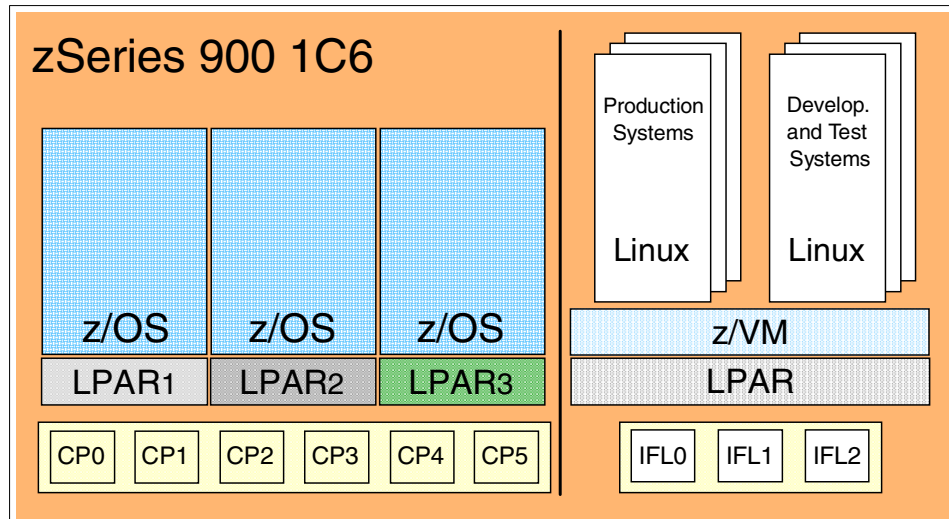


Figure 4-3 zSeries 900 Model 1C6 with three IFL processors and z/VM V4

4.2.2 Hardware requirements for z/VM

z/VM V4 requires at least a Multiprise 3000, G5, G6, zSeries 900, or zSeries 800 processor. All z/VM V4 releases can run on the IFL as well as standard

processor engines. IFL engines are available on the Multiprise 3000, G5, G6, zSeries 900, and zSeries 800. The Linux-only model of the zSeries 800 uses IFL engines exclusively.

4.2.3 Virtual networking with z/VM

Virtual networks are a great way to connect guests running under z/VM. Virtual networks use z/VM virtualization to provide connectivity between guests without needing real hardware. Virtual networks enable fast and cost effective TCP/IP communication between z/VM Linux guests.

Note: Connectivity in a z/VM virtual network is limited to guests running in a single z/VM image. Virtual networks cannot be used for inter-LPAR communication.

Three types of virtual networks are available to Linux guests:

- ▶ **Virtual channel-to-channel (vCTC):** vCTC networks provide point-to-point connectivity between guests without real channel allocation (as required for real CTC connectivity). A real channel-to-channel adapter is used to connect a real mainframe to another using the channel protocol.

z/VM provides the ability to define virtual channel-to-channel adapters so users can connect virtual machines using the CTCA protocol. On z/VM, this is useful for connecting Linux virtual machines to other virtual machines that don't support IUCV, such as VSE/ESA™, OS/390®, and z/OS. Virtual vCTC networks can also connect Linux virtual machines.

- ▶ **Inter-user communication vehicle (IUCV):** Point-to-point TCP/IP connections between Linux guests can be established using IUCV. IUCV is a VM-unique, virtual machine-to-virtual machine communication protocol. The Linux for S/390 and zSeries kernels include an IUCV driver, allowing you to connect two Linux virtual machines.

Linux treats IUCV connections as any other TCP/IP network connection. This results in memory-speed networking between the connected Linux servers. The speed of the IUCV “line” is a factor of processor speed. The faster your zSeries server is, the faster your IUCV network is. When you upgrade to a faster real processor, you automatically increase the speed of your virtual network. IUCV is also supported by z/VM's own TCP/IP stack, so you can use IUCV to connect Linux images to VM's TCP/IP stack.

- ▶ **VM Guest LAN:** This type of network enables local area network (LAN) connectivity between z/VM guests. You can learn more about the details of VM Guest LAN in 4.2.4, “VM Guest LAN support” on page 80.

Figure 4-4 shows a simple virtual network. There are three groups of five Linux servers, each connected to a virtual router using a virtual networking protocol. The virtual routers are virtual machines, each with a real networking interface to the “outside world”. In this scenario, each router is connected to a real Open Systems Adapter (OSA).

This example shows how you can connect a large number of Linux servers to a real network using fewer real network adapters, exploiting z/VM's virtual networking capabilities to essentially share the real network connections among your Linux servers. The value proposition here is one of technology exploitation and simplicity. The virtual networking configuration looks just like something deployed on discrete, real servers. Networking personnel do not need to learn something new when working with virtual networking on z/VM.

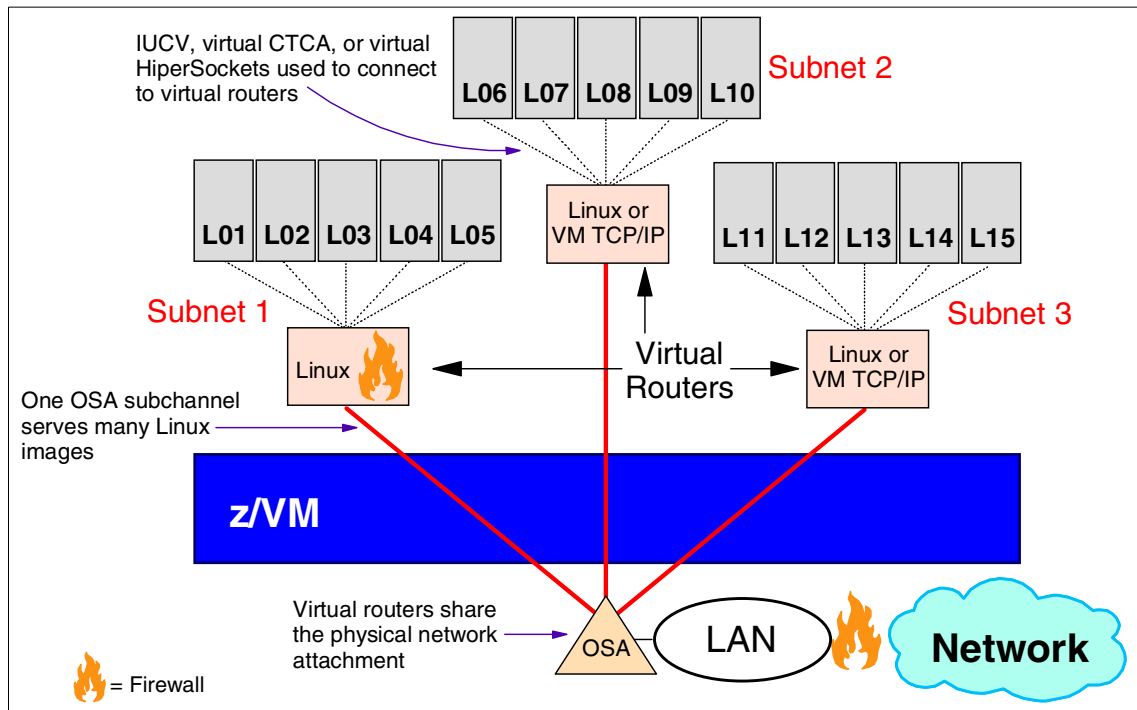


Figure 4-4 z/VM virtual networking

4.2.4 VM Guest LAN support

Introduced in z/VM 4.2 and enhanced in z/VM 4.3, VM Guest LAN provides virtual LAN networking support without relying on real LAN hardware. There are two types of VM Guest LAN:

- ▶ **HiperSockets:** Introduced in z/VM 4.2, the HiperSockets VM Guest LAN emulates real HiperSockets networking in a z/VM image.
- ▶ **OSA:** Introduced in z/VM 4.3, the OSA VM Guest LANs emulates real OSA networking in a z/VM image.

VM Guest LAN support eliminates some point-to-point LAN management challenges when using IUCV and vCTCA. z/VM is designed to enable virtual HiperSockets emulation using VM Guest LANs on processors that don't support real HiperSockets, such as the G5, G6, and Multiprise 3000. This enhances the virtual networking environment on those processors, while allowing customers to prepare for a real HiperSockets environment prior to moving to a zSeries server.

It is possible to define “system” guest LANs and guest LANs that are associated with a specific VM user (for example, virtual machine). System guest LANs exist independently of any active (for example, logged-on) user. Guest LANs are associated with a user exist only while that user is active. For either type of guest LAN, authorized users can link to the LAN to participate in HiperSockets communications. There is no predefined limit on the number of virtual HiperSockets devices that you can link to a guest LAN. Nor is there a limit on the number of guest LANs that can be defined.

The virtual networking environment in Figure 4-5 shows a single Linux image (in the middle) connected to a real network interface. It is serving as a router to the “outside world” for the Linux servers connected to the three guest LANs.

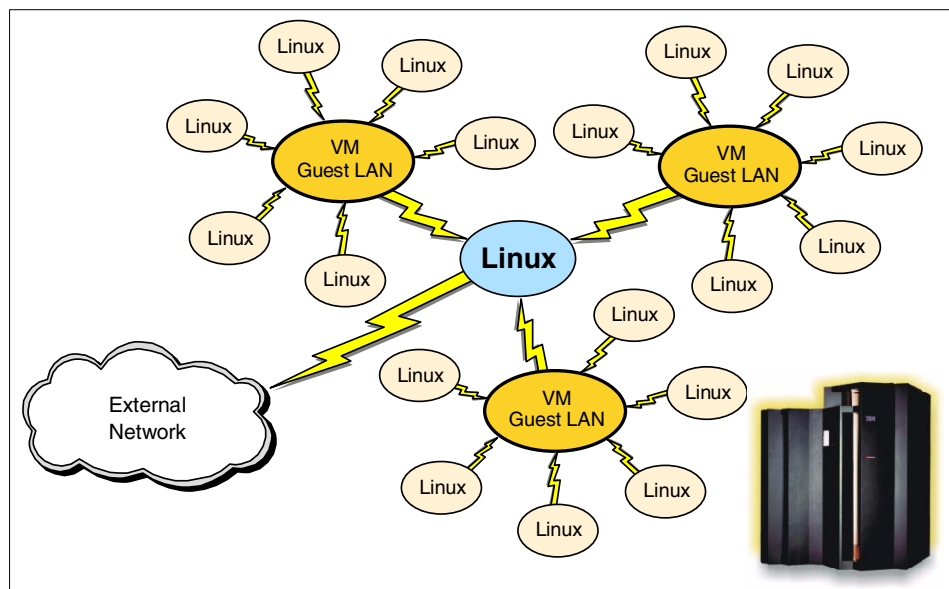


Figure 4-5 Virtual networking example

The Linux router can also transfer data from one guest LAN to another. This allows customers to isolate their Linux images on a guest LAN from Linux images on other guest LANs, using the router to secure traffic between the guest LANs.

4.2.5 z/VM for Linux performance: Scalability

zSeries customers have many options for scaling a Linux workload on z/VM. z/VM supports vertical and horizontal growth. VM exploitation of hardware technologies enhances the performance and scalability of Linux workloads.

With z/VM, you can grow your server workload vertically or horizontally. *Vertical growth* is realized when you add more system resources to an existing virtual machine. You can give a Linux virtual machine more processor capacity, more virtual memory, more I/O devices and bandwidth, plus more networking bandwidth.

Horizontal growth is a typical way to grow UNIX workloads. It is easily accomplished by adding another virtual machine, or two, or three, or more on z/VM. There is an added benefit in the efficient use and sharing of system resources. This makes adding more Linux virtual servers a particularly attractive option when you consider the limited scalability of a single Linux instance.

The added scalability of Linux workloads is realized when you choose to exploit VM's data-in-memory techniques. Virtual disks in storage and Minidisk Cache can boost the performance of Linux servers on z/VM by making disk I/O operations execute at processor memory speeds, avoiding trips to the I/O subsystem, and waiting on spinning disks of storage.

z/VM support for such advanced storage functions as Peer-to-Peer Remote Copy (PPRC) and FlashCopy® benefits Linux virtual servers, although Linux itself doesn't support the function. Exploitation of this technology and other technologies, such as FICON™ support, is transparently passed on to the Linux servers simply because they are running on z/VM.

4.2.6 Productivity: Development and test workloads

z/VM enhances the development and testing of Linux on zSeries. z/VM makes it easier to develop, enhance, and test Linux systems. Anyone can use Linux with z/VM virtual machines. When you run Linux on z/VM, you will realize productivity gains in the areas of development and testing.

Virtual machines are easy to create, add resources to, and decommission when you no longer need them. A z/VM environment is considerably more flexible than logical partitioning. Development and test environments are generally dynamic.

They grow and shrink in size over a short period of time, which is something that is well suited for z/VM.

z/VM includes extensive built-in debug facilities. You can “instruction step” through a Linux virtual machine, interrogating storage locations, registers, and other system facilities after each instruction if you choose. You can set “traps” to stop a Linux virtual machine when a condition occurs, making it easier to target bugs.

Because z/VM virtualizes system resources, you can create a complex test environment using virtual devices, networks, and facilities instead of purchasing real hardware to accomplish the same. With z/VM, you can conduct extreme test cases to see if your server environment can handle the demand.

Virtual machines on z/VM have virtual processors. You can define anywhere from one to 64 virtual processors with z/VM. This is useful if you want to test Linux on a multiprocessor configuration, even if the underlying z/VM system is running on a single processor.

When you have your own Linux virtual machine, or two, or more, you complete your work faster and innovation flourishes. The ability to try something on Linux in a virtual machine environment is a low-cost endeavor. Successful solutions are created by users who had some spare time or the necessary resources to experiment. Linux on z/VM gives you a sandbox to innovate at your best.

4.2.7 Productivity: Production workloads

Production workloads running on z/VM are easy to modify, extend, and adapt to changing business needs. Creating Linux servers on demand can help meet business needs. Time is money, especially when it comes to e-business. If business systems can't handle the demand or aren't flexible enough to react to new opportunities, business suffers.

zSeries servers with z/VM help you deal with unexpected increases in workload. You can bring Linux servers online in a matter of seconds or minutes with z/VM. You can pre-configure Linux servers offline (that is, logged off) and bring them online as soon as you need them. It is possible to provision the server in real-time in a matter of minutes as well.

You can add capacity to an existing server image in real-time (without causing a server outage) using available processor capacity. You can also add real processor capacity with the zSeries CUoD and immediately add it to your existing virtual Linux server environment.

The flexibility of a virtual machine also means that Linux images can stand by as a hot backup for production servers. The attraction of this feature on z/VM is that

the standby images consume minimal resources. In a discrete server world, a standby server typically costs as much as the production server it's backing up. The cost includes acquisition and support expenses.

A key aspect of z/VM's value for Linux is resource sharing. This includes processor capacity, network bandwidth, and disk technology. CMS minidisks, a disk partitioning technology, can be readily shared among Linux images. This enables you to store data on one set of disks to be accessed by multiple Linux server images. A read-only minidisk can be accessed by any number of Linux virtual machines. Such a configuration is ideal for version control and upgrades of application software. With z/VM's minidisk cache support, minidisk data can be cached using processor memory. This enables memory-to-memory data transfer rates for disk I/O to a read-only minidisk.

z/VM and DASD options

Linux on zSeries can use IBM TotalStorage Enterprise Storage Server® (ESS) for disk space allocation. As shown in Figure 4-6, Linux disk space can be allocated from TDISK space, from a single disk volume, or from VM minidisks.

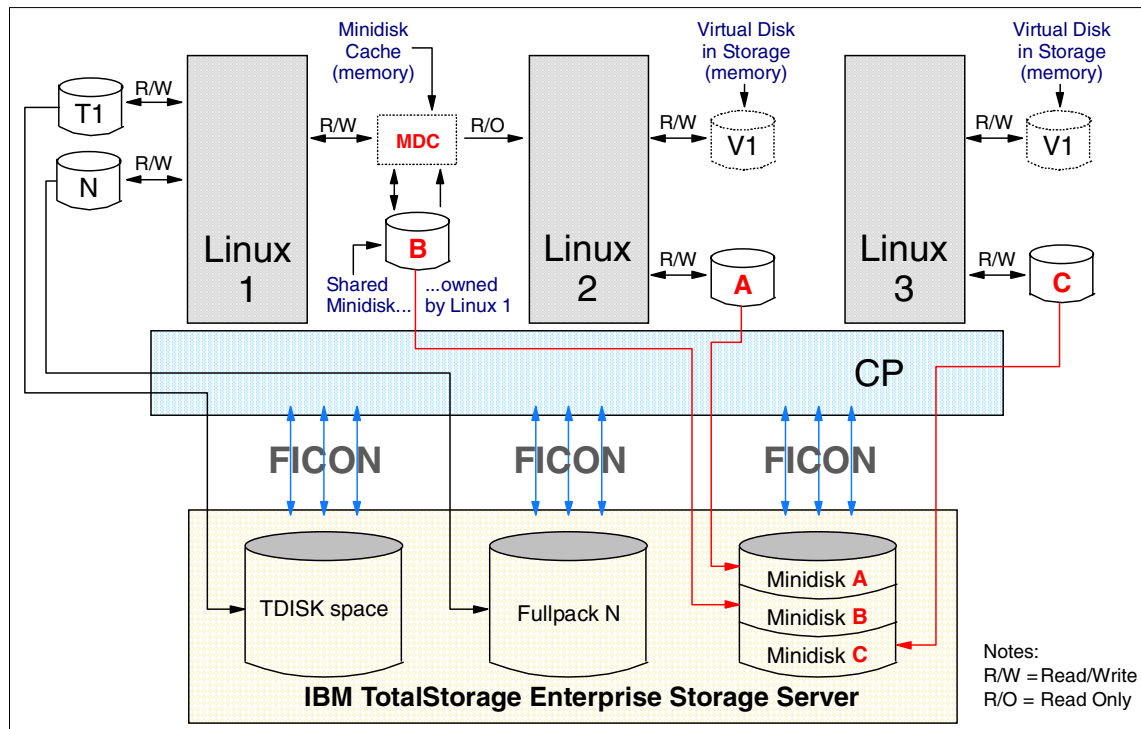


Figure 4-6 z/VM and virtual disks

4.2.8 Systems management: Data backup with Tivoli Storage Manager and FlashCopy

You can exploit z/VM technology for file-level Linux data backups using Tivoli Storage Manager. z/VM technology improves performance of the backup operation. FlashCopy minimizes the outage time required for the backup operation.

z/VM offers a considerable amount of systems management function for Linux servers. A case in point is backup and restore using Tivoli Storage Manager. Figure 4-6 represents a backup scenario that exploits z/VM's support for the FlashCopy function of the IBM TotalStorage Enterprise Storage Server. Running Linux in a z/VM virtual machine means that you can exploit the advanced technology found in FlashCopy even though Linux itself does not support the capability.

The data that is backed up in this example is found on disk address 201. This disk is "owned" by the production application running on the leftmost Linux image. A FlashCopy of this disk is made using the z/VM FlashCopy command. This gives produces the 202 disk shown in the diagram, a duplicate copy of the 201 disk. The FlashCopy operation is a fast operation on the ESS subsystem that minimizes the outage required of the production application running on Linux. After FlashCopy completes, the production application can resume, accessing the 201 disk as required.

The 202 disk is attached to a second Linux virtual machine that was created for the purpose of performing the backup operation. The TSM client code is running in this Linux image. It sends the file-level data to the TSM server that is running in a separate virtual machine, using a high-speed virtual connection between the two virtual machines. The z/VM TSM server, separately licensed but delivered with the z/VM product, sends the backup data to the tape subsystem.

4.2.9 Consolidation example: Server farm in a box

Many expenses (and waste) are associated with real distributed servers. It's typical for a server image to be active only 15 to 20% of the time. When idle, a discrete server is unable to allocate its processor and memory resources to applications on other servers that could use the capacity.

Consolidation of distributed servers on Linux running under z/VM can offer an attractive alternative. Virtual servers on z/VM share total system resources. Processor and memory capacity is allocated to the servers that need it, when they need it.

As shown in Figure 4-7, each discrete server needs its own disk space. A lot of that space contains the same files/data found on other servers, representing a

needless waste of disk space (and money). This figure shows how z/VM lets you share disk space among virtual Linux images. Not only is it an efficient use of your disk investment, but it is also a great way to provide version control for your application software (placing your software on one disk, accessible by all your virtual Linux servers).




Figure 4-7 Server farm in a box

With separate servers, you need to connect real cables to each server. z/VM provides “virtual” cables for your Linux servers. These virtual cables are fast, inexpensive, and provide memory-to-memory data transfer rates.

Command and control operations require a large support staff or expensive software products. They also need corresponding dedicated servers to run the software and client software installed on each server image to communicate with the command and control servers. z/VM offers built-in command and control functions.

If required, compute-intensive server workloads can remain on discrete server images and connect to the rest of the “server farm” running on z/VM.

4.2.10 Consolidation example: Data and application servers on a single zSeries server

Hosting your application servers on zSeries with Linux and z/VM can save you money and enhance the quality of service. Many of the applications connecting to mainframe servers are already available on Linux.

It is relatively simple to move these applications to Linux on zSeries, without requiring changes to the client community, and save money.

The server consolidation scenario in Figure 4-8 is commonly referred to as *application integration*. The cost savings are similar to the *server farm in a box*

environment. The co-residency of data and application servers means that such technologies as HiperSockets can be exploited for fast connectivity between data servers and application servers without the complexities of a real network.

In Figure 4-8, the application servers running on the distributed servers (the mid-tier application servers) are replaced by Linux virtual machines running under z/VM. The real networks are replaced by the networking features offered by zSeries (HiperSockets and VM Guest LANs). Linux on z/VM on zSeries facilitates the three-tier architected environment on two tiers of real hardware. There is no need to redesign the environment. The users in this scenario can continue to connect to an application server. They most likely do so without noticing that their application servers are re-hosted. That is unless they notice that their application server doesn't suffer the outages it used to.

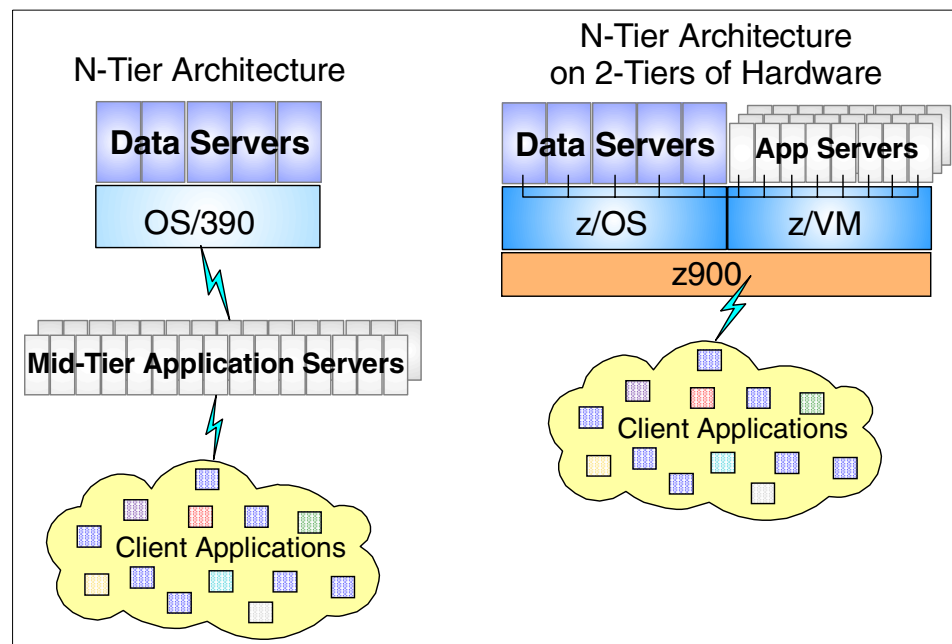


Figure 4-8 Data and application servers moving to a single zSeries

HiperSockets enable memory-to-memory network connectivity between LPARs within the same CEC. In this diagram, an LPAR is running z/OS, and another LPAR is running z/VM with Linux application servers.

4.2.11 Horizontal growth: Adding another server

z/VM is a cost-attractive platform that handles a growing server environment. Virtualization technology is a “must have” for fast-growing server environments.

You can respond to server growth with greater ease and a more efficient exploitation of hardware technologies using z/VM and Linux for zSeries. Figure 4-9 shows an example of how this works.

Sometimes growth plans are underestimated. Growth may begin with a modest number of discrete servers. Soon you find yourself adding more servers to the configuration than you anticipated. Each additional discrete server takes time to setup and bring online, contains resources that can't be shared with other servers (memory, processors, disk space, etc.), and places demand on the people and related software products performing systems management.

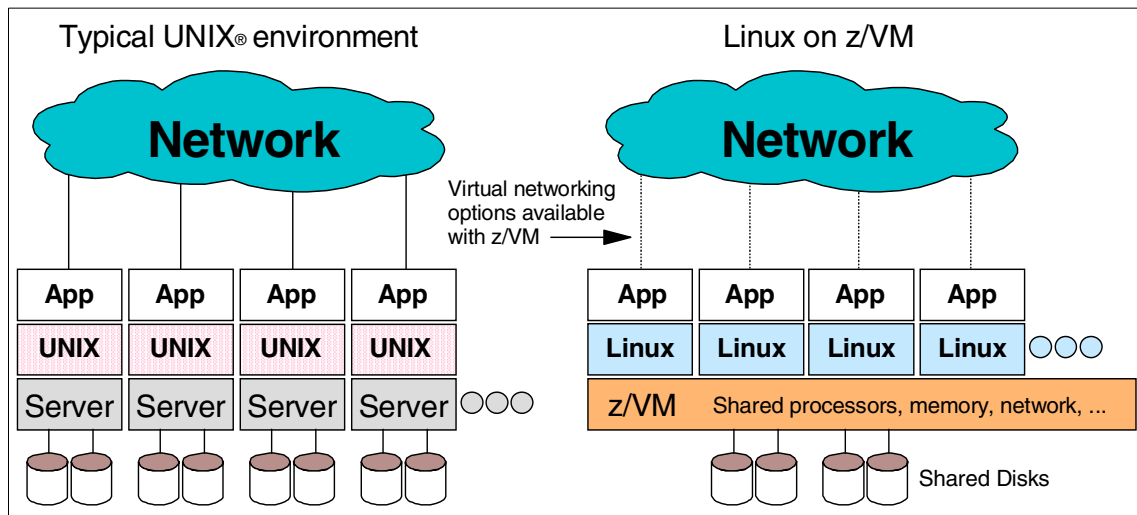


Figure 4-9 Horizontal server growth

Adding another virtual Linux server on z/VM is much less taxing on people and the IT budget. Virtual machines can be brought online in minutes, not hours or days. z/VM efficiently shares the total system resources among all the Linux server images. There's no waste of processor capacity, memory, I/O bandwidth, etc. And z/VM's built-in systems management capabilities quickly and easily adapt to the presence of another server image.

Customers need to carefully consider their server growth requirements when deploying server solutions. A fast-growing server environment, although small to begin with, may be ideally suited for Linux on zSeries with z/VM.

4.3 zSeries Linux distributions

Linux distributions for zSeries are currently available from three major Linux distributors: SuSE, Turbolinux, and Red Hat.

Linux for S/390 has been generally available since early 2000 either from IBM DeveloperWorks or from the Linux distribution partners. All the current distributions (as of November 2002) are based on the Linux 2.4 kernel. Each distribution includes additional middleware and applications. They also include infrastructure services, such as domain name server (DNS), Dynamic Host Configuration Protocol (DHCP), and Network File System (NFS) file servers, and packages, such as the Apache Web server, Squid proxy server, Simple Mail Transfer Protocol (SMTP) mail server, and Samba Windows networking server. These distributions also leverage the HiperSockets technology of the zSeries server to interconnect between different partitions.

The distributions that are available for zSeries are:

- ▶ SuSE
 - SLES 7
 - SLES 8
- ▶ Red Hat
 - Red Hat 7.1
- ▶ Debian 3.0

In the future, a version of UnitedLinux will also be available for the zSeries servers.

4.4 zSeries customers and Linux

The zSeries server is widely recognized as an integrated business server that scales both vertically and horizontally. It is reliable, scalable, and recognized as one of the most flexible, easy to use systems in the industry. And it can run multiple environments and help to quickly deploy applications.

These attributes position zSeries as one of the best platforms to manage the complexity and cost of e-business enablement. Key characteristics of Linux on zSeries, such as a new generation of applications, integration, and consolidation, strongly support the IBM initiatives. They can result in measurable customer benefits for the deployment of e-business solutions.

4.5 Linux application solutions

Linux distributions for zSeries contain such notable open source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail MTA, Squid, and PostgreSQL. Linux applications are sourced from many areas. Open source applications are available for download from the Internet. Applications can also be developed for the zSeries. Compilers are available with the Linux for zSeries distributions.

Solution providers have announced products for Linux on the zSeries server. There are more than 200 applications from more than 120 vendors available for Linux on zSeries. For a detailed product list, see:

<http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxproduct.html>

To assist solution providers with bringing additional solutions to market, a zSeries Test Drive for Linux is available that provides remote access to Linux installations on zSeries. Different types of test drives are available that represent different distributions and amounts of disk space. For more information, refer to either of the following Web sites:

<http://www.developer.ibm.com/welcome/s390/eslinux.html>

<http://etpgw02.dfw.ibm.com/rdp.html>

Infrastructure applications

Examples of infrastructure applications announced for the zSeries Linux solutions are:

- ▶ **extend5:** This is Acucorp's integrated development environment for COBOL, which includes ACUCOBOL-GT (the flagship product). Several enterprise companies and ISVs have written or extended their COBOL code using ACUCOBOL-GT. This integrated development environment provides software developers with the ability to move their existing AUCOBOL-GT applications to Linux for zSeries without recompiling.
- ▶ **GMx Solutions' CM_SAFE Server for zSeries:** This application provides configuration file management capabilities for open system hosts. It also provides server support for zSeries Linux installations and agent support for all major UNIX variations and platforms.
- ▶ **zGuard:** This application from Frank Bernard Informationstechnik, is a Linux-based Internet security solution for S/390 and zSeries. It has firewall, IPsec-VPN, online virus scan of various protocols including mail, HTTP, NNTP, and FTP, and basic IDS functionality.

Line-of-business applications

Examples of line-of-business application solutions available for the zSeries Linux solutions are:

- ▶ **mySAP.com:** This is the successor of SAP's successful enterprise resource planning software, SAP R/3. It is an e-business platform for the entire enterprise. It offers numerous modules from procurement, project management, production, accounting, and controlling.
- ▶ **LOGA 2001:** This is a modern payroll application from Personal & Informatik AG. This application incorporates remuneration in salary and wages along with travel expense accounting. A variety of solution packages can be offered, depending on the branch involved, such as the industrial sector services sector or public administration.

Personal & Informatik AG also offers a human resources management system called LOGA/ERM (Employee Relationship Management). This portal empowers employees to automate certain functions and reduces the cost for administrative work. Using the personalized home page, access is granted for applications, services, and information for daily work and career planning.

- ▶ **OpeN/2 for Linux:** This is a consolidated solution approach from S2 for authorization and payment-based transaction processing. It leverages mainframe-class computing power reduced complexity and implementation time. OpeN/2's affinity with the Linux mainframe environment stems from its Platform Abstraction Layer (PAL [®]) design that isolates platform dependent code from higher-level services and applications.

Web-enablement applications

Examples of Web-enablement application solutions are:

- ▶ **eOneCommerce:** This solution from eOne Group is an e-business software application and Internet publishing tool.
- ▶ **Tamino XML Server:** This application, from Software AG, is a high performance information management platform based on native XML storage and open standard Internet technologies. Tamino XML Server helps to find and manage any type of content across the enterprise and enables rapid implementation of robust, high-performance, mission-critical electronic business applications based on XML standards. The server is built to:
 - Efficiently store XML documents natively, that is in their original format
 - Expose information residing in various external XML or non-XML sources (legacy data) or applications to the outside world in XML format
 - Search effectively on all information available to the Tamino application

- ▶ **CommuniGate Pro:** This application, from Stalker Software, offers a carrier grade messaging solution. It incorporates high performance, speed, reliability, security with an extensive feature set. Its unique dynamic clustering architecture allows you to manage millions of accounts, while providing a 99.999% uptime requirement. It is based on open standards and supports over 24 different operating systems and hardware platforms including all IBM platforms. It provides customizable Web-based e-mail access, Internet Message Access Protocol (IMAP), SMTP/Post Office Protocol (POP), Lightweight Directory Access Protocol (LDAP), a built-in Web server for personal Web pages, a state-of-the-art list server, multilanguage support and straightforward administration tools, and much more.
- ▶ **DI-Atlantis:** This solution from Dimensional Insight, Inc. is a multidimensional software package designed to transform data for intuitive access and analysis. DI-Atlantis allows you to give users the information they need in a format that is easily understood.

4.5.1 Mail serving scenario

Figure 4-10 illustrates a typical electronic mail application that an enterprise may use for in-house mail serving. Typically usage of e-mail applications is sporadic, and demands on the processor are very low. Long periods of idle time are common as servers wait to be contacted by a user to send, receive, or check their mailbox. When the server is active, it usually performs I/O operations since little computation is performed for this application.

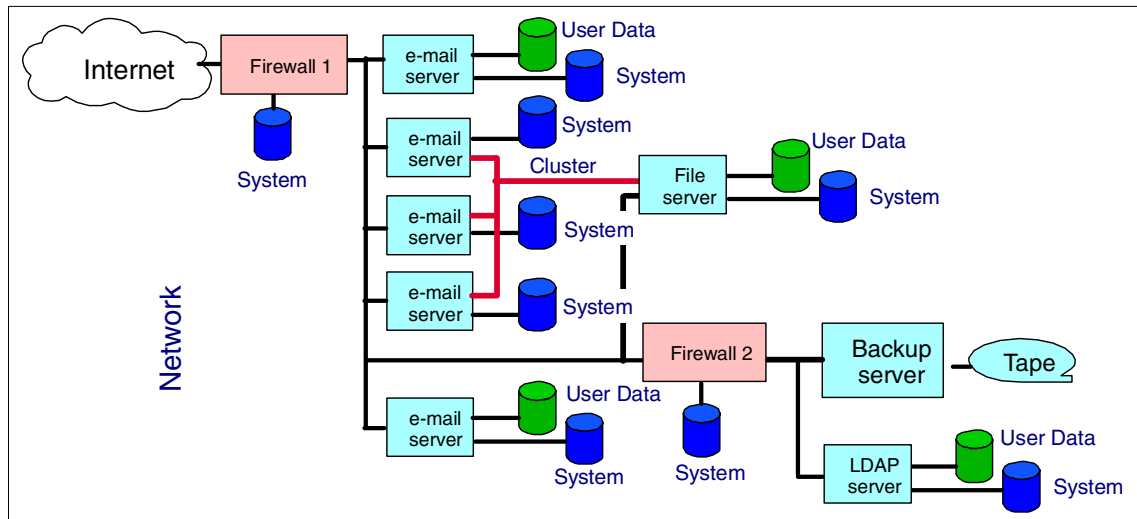


Figure 4-10 Mail serving scenario

The example in Figure 4-10 makes several interesting and valid points. Regardless of the actual solution scenario, the solution-based servers are not the only servers making up the solution. Firewalls, failover protection, directory servers, backup servers, and others are typically included in production environments. Keep in mind that, in most distributed solutions, any of those servers may constitute a single point of failure. Also, each server occupies floor space, consumes power, and must be monitored continuously.

The software that drives this application is uniform and can be used by all of the servers that provide the solution. The service level required for this application is prime shift or 12 hours a day, 5 days per week.

There are two options to this scenario:

- ▶ 30 2-way Intel 1,000 megahertz (MHz) servers
- ▶ One zSeries 900 Integrated Facility for Linux engine

Figure 4-10 illustrates the distributed solution. Notice that the servers are connected to a physical network which requires hardware. Also notice that each server is running a separate copy of the operating system and application code. This is duplicated for each server in the solution, which is 30 times in this case. Each server requires disk space to support this. If additional disk storage is required for a server, the alternatives are to either purchase an additional entire disk volume and install it or to purchase a higher capacity disk, move the required software and data to it, and install it.

Maintenance must be performed on each server individually. Some type of operator automation is required to monitor the 30 separate systems and respond to replies when necessary. Backup is done by a dedicated backup server. Client code on each system communicates with the server to perform backup operations.

Important: In the following scenario, the assumption is that the Intel servers run somewhere in the 1% to 5% utilization range.

Figure 4-11 shows the IBM @server zSeries solution. Notice that the inter-server communication takes place within the machine. This allows the Linux instances to communicate at memory speeds, as opposed to physical networking speeds with the usual latency. In addition to improving inter-server communication, much of the cost of external networking hardware and software is avoided.

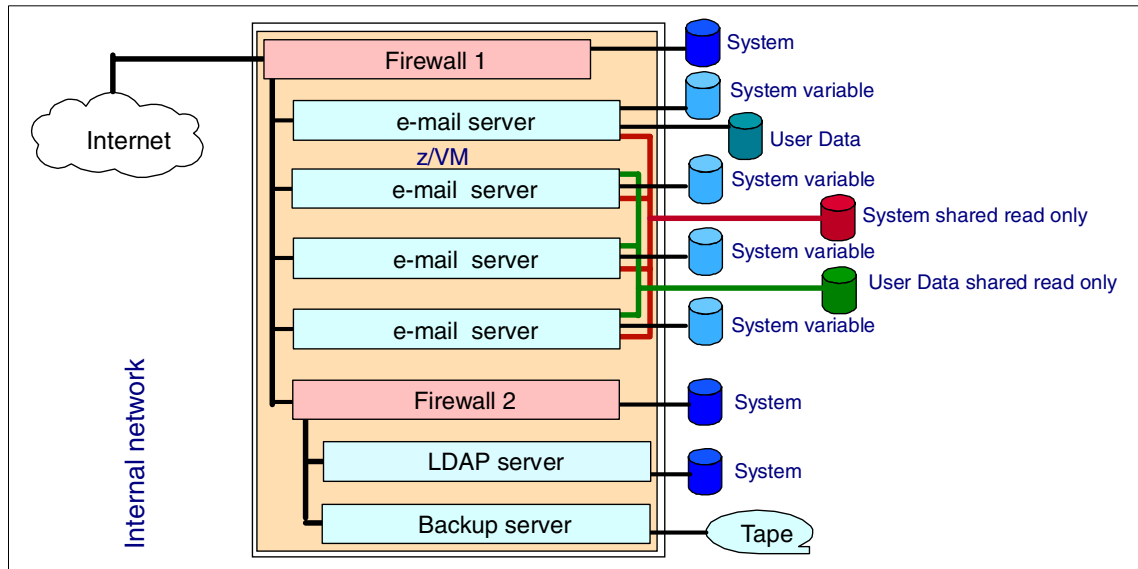


Figure 4-11 zSeries solution

Disk storage for the operating system and application code is shared by all 30 virtual servers while they must be duplicated in the distributed case. Sharing the code offers two advantages, controlling the cost of disk storage, and ease of maintenance. Since only one copy of the code exists, it only has to be serviced in a single place. This results in savings in system management software as well as staffing. Although only one copy of the code is being used, 30 servers are in operation. As in the distributed case, they will issue messages, which must be monitored and, in some cases, responded to. z/VM includes a programmable operator function (called PROP) that may be used for that purpose to avoid the cost of additional software to perform the task. Additional systems management software is required in complex environments.

z/VM also provides the ability to partition the disk space available to the servers. In the distributed model, increasing disk space meant buying new disk. If there was free disk space on other servers, it could not be moved to the server needing increased disk storage. In the IBM solution, z/VM provides the ability to partition disks into minidisks that can be allocated easily to a server that requires additional disk space. Note that in e-mail applications, some mailboxes will grow at greater rates than others and may require expansion.

Backup is done again by a backup client. In the IBM solution, it's done by the Tivoli Systems Manager Backup server, which is pre-installed on z/VM. The Tivoli Systems Manager server communicates with clients on the virtual instances to perform backup operations.

Another factor to consider is that the IBM TotalStorage Enterprise Storage Server disk solution is a unique and far more powerful solution than dedicated disks on stand-alone systems. The reliability characteristics of the Enterprise Storage Server include the advanced functions that are available in the unit, such as FlashCopy. This function allows you to create a copy of a disk volume instantly. This is extremely useful for minimizing downtime for backups or creating test data for new applications.

Cost comparison: Three-year total cost of ownership (TCO)

Figure 4-12 shows an estimate of a three-year cost comparison between the distributed solution and Linux zSeries 900 solution. There are dramatic cost savings in reliability and personnel costs.




Figure 4-12 Cost comparison in the mail serving scenario

4.5.2 News group serving scenario

This case describes servers that are dedicated to news group serving such as those provided by application and Internet service providers. The application is rather simple, and demands on the processor are low. Long periods of idle time are common as servers wait to be contacted by a user that is requesting information. When the server is active, it usually performs I/O operations since little computation is performed for this application.

The software that drives this application is uniform and can be used by all servers that provide the solution. In the case of news group serving, no allowance for outages was made.

As in the previous scenario, this has three different approaches: one Intel based, one S/390 based, and one zSeries based. The alternative solutions in this case are:

- ▶ A distributed solution of 215 Intel-based servers
- ▶ 9672 six-engine Linux G6 X67
- ▶ Four zSeries 900 Integrated Facility for Linux engines

Figure 4-13 shows a logical schematic of the distributed solution. The servers are connected to a physical network that requires hardware. Also each server is running a separate copy of the operating system and application code. This is duplicated for each server in the solution, which is 215 times in this case. Each server requires disk space to support this. If additional disk storage is required for a server, the alternatives are to purchase an additional entire disk volume and install it or to purchase a higher capacity disk, move the required software and data to it, and install it.

Maintenance must be performed on each server individually. Also some type of operator automation is required to monitor the 215 separate systems and respond to replies when necessary. Backup is done by a dedicated backup server. Client code on each system communicates with the server to perform backup operations.

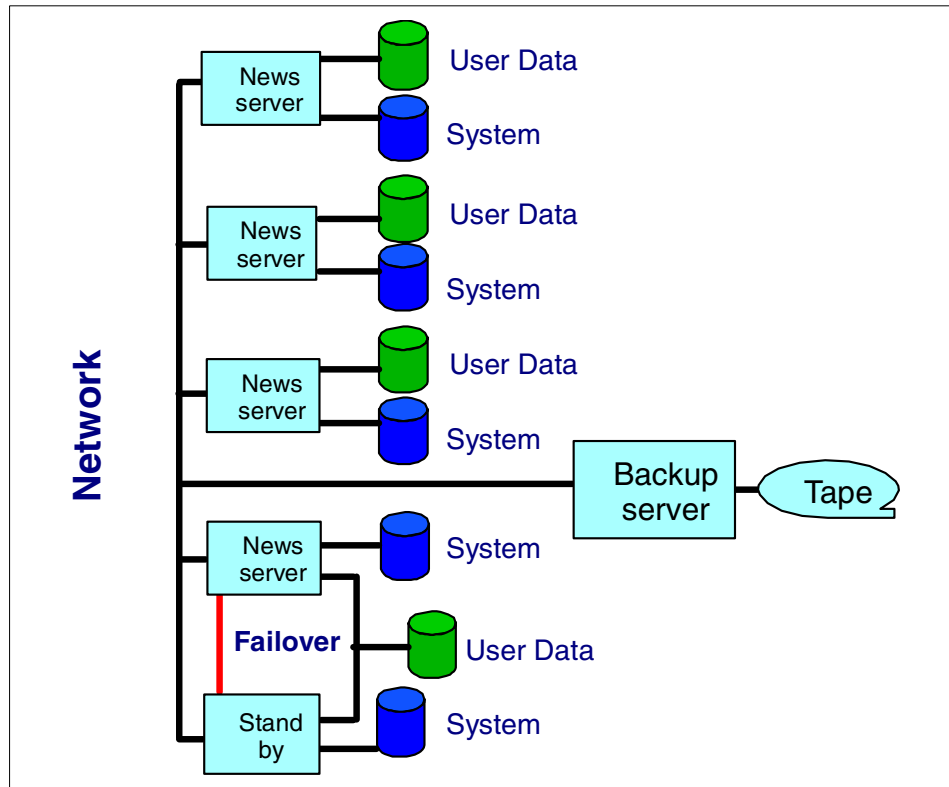


Figure 4-13 Distributed news group server scenario

Figure 4-14 shows the zSeries solution. The inter-server communication takes place within the machine, which allows the Linux instances to communicate at memory speeds, as opposed to physical networking speeds with the usual latency. In addition to improving inter-server communication, much of the cost of external networking hardware and software is avoided.

Disk storage for operating system and application code is shared by all 215 virtual servers, while they must be duplicated in the distributed case. Code sharing offers two advantages:

- ▶ Controlling the cost of disk storage
- ▶ Ease of maintenance

Since only one copy of the code exists, it only must be serviced in a single place, resulting in savings in system management software and staffing. Even though only one copy of the code is being used, 215 servers are in operation. As in the distributed case, they will issue messages that you must monitor and respond to, in some cases. z/VM's PROP function allows you to do that, avoiding the cost of

additional software to perform the task. Additional systems management software is required in complex environments.

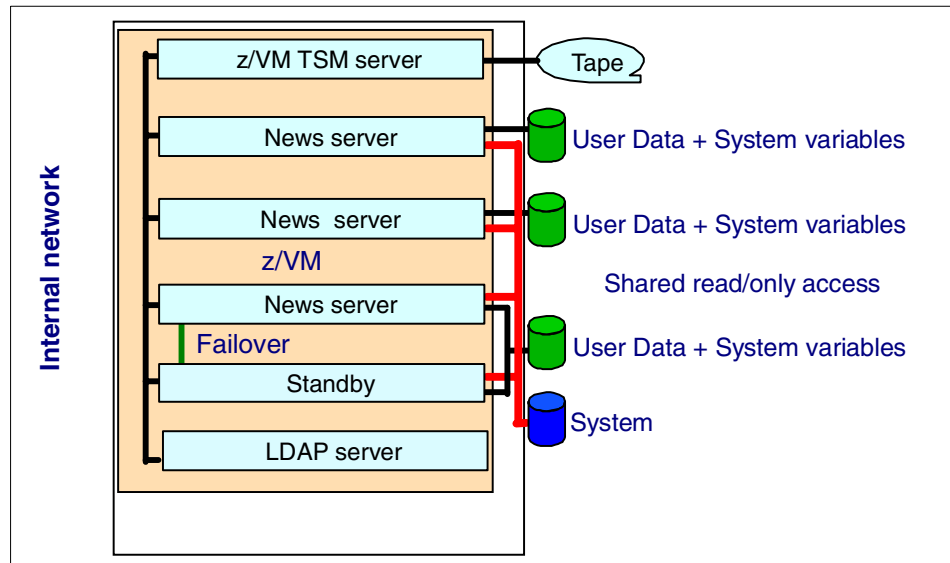


Figure 4-14 zSeries solution

z/VM can help partition the disk space available to the servers. In the distributed model, increasing disk space meant buying new disks. If free disk space was available on other servers, it could not be moved to the server that needed increased disk storage. In the S/390 solution, z/VM can partition disks into minidisks, which can be allocated easily to a server that requires additional disk space.

Backup is done again by a backup client. In the zSeries case, it's done by the Tivoli Systems Manager Backup server which is pre-installed on z/VM. The Tivoli Systems Manager server communicates with clients on the virtual instances to perform backup operations.

Another factor to consider is that the Enterprise Storage Server disk solution is unique and far more powerful than dedicated disks on stand-alone systems. The reliability characteristics of the Enterprise Storage Server include the advanced functions, such as FlashCopy, that are available in the unit. This function allows the user to create a copy of a disk volume instantly. It is extremely useful for minimizing downtime for backups or for creating test data for new applications.

Cost comparison: Three-year TCO

Figure 4-15 shows the cost estimate for each alternative. It includes the values for the IBM solutions on zSeries and S/390. Typically, a zSeries solution requires

fewer engines than a G6 solution since zSeries engines are substantially more powerful. This affects certain costs, such as software and maintenance costs, that are based upon the number of engines required by a solution.

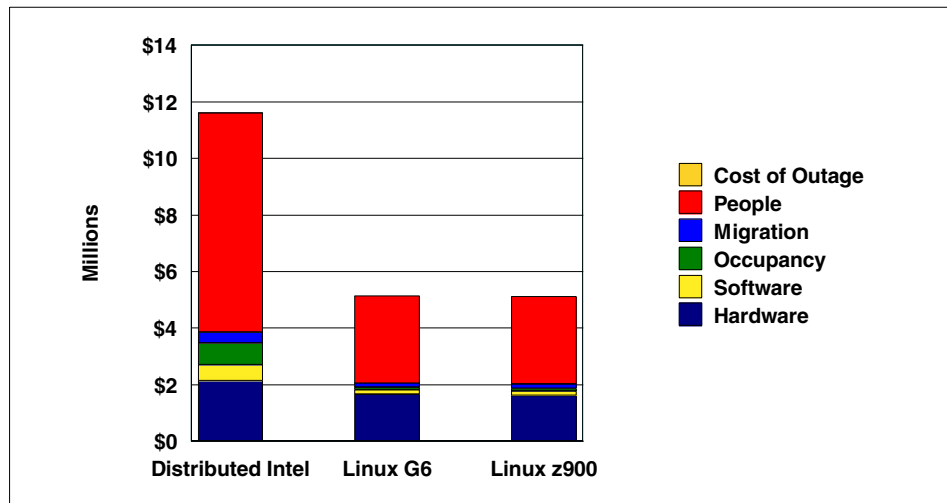


Figure 4-15 News group serving scenario comparison

4.5.3 File/print serving with Samba

This scenario describes providing file/print services using the Samba file print application. Samba is high quality open source code and is integrated in most Linux distributions. File/print serving does not place great demand on the compute resource of the processor. However, it exercises the I/O capabilities of the solution as users upload/download files or route them to printers. Servers spend a significant amount of time waiting for users to request print or file services.

The software that drives this application is uniform and can be used by all servers that provide the solution. The service level required for this application is prime shift, 12 hours per day 5 days per week.

The alternative solutions in this case are:

- ▶ A distributed solution of 150 Intel 2-way 1000 MHz servers
- ▶ 9672 G6 X47

Let's look at the alternative solutions as we did when we looked at the application. Figure 4-16 shows you a logical schematic of the distributed solution. Note that the servers are connected to a physical network that requires hardware. Also notice that each server is running a separate copy of the

operating system and application code. This is duplicated for each server in the solution, which is 150 times in this case. Each server requires disk space to support this. If additional disk storage is required for a server, the alternatives are to purchase an additional entire disk volume and install it, or purchase a higher capacity disk, move the required software and data to it, and install it.

Maintenance must be performed on each server individually. Some type of operator automation is required to monitor the 150 separate systems and respond to replies when necessary. Backup is done by a dedicated backup server. A client code on each system communicates with the server to perform backup operations.

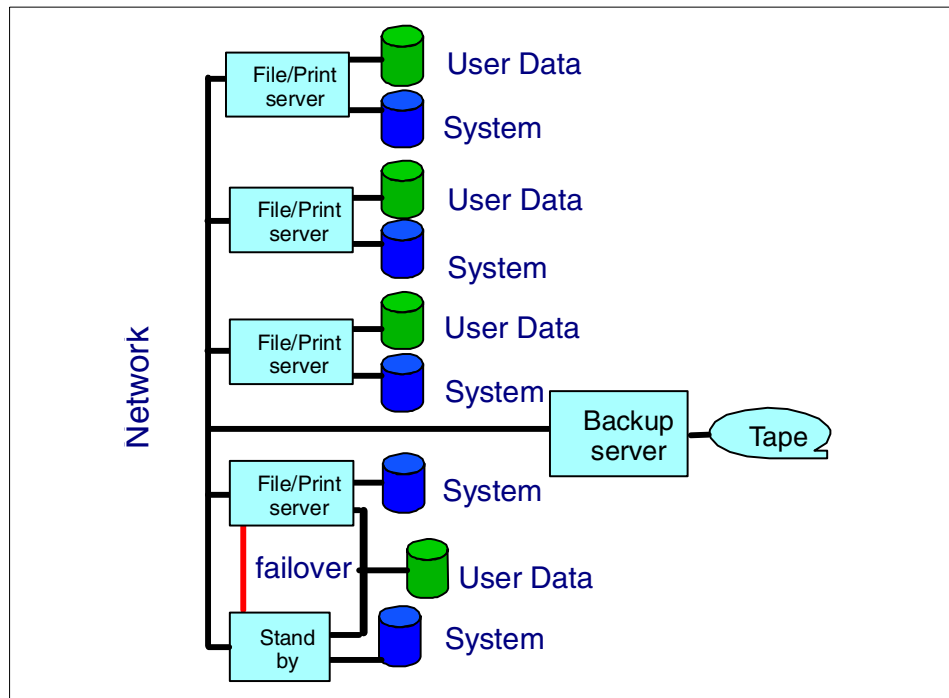


Figure 4-16 Samba file and print scenario

Figure 4-17 shows the zSeries solution. Notice that the inter-server communication takes place within the machine, which allows the Linux instances to communicate at memory speeds, as opposed to physical networking speeds with the usual latency. In addition to improving inter-server communication, much of the cost of external networking hardware and software is avoided.

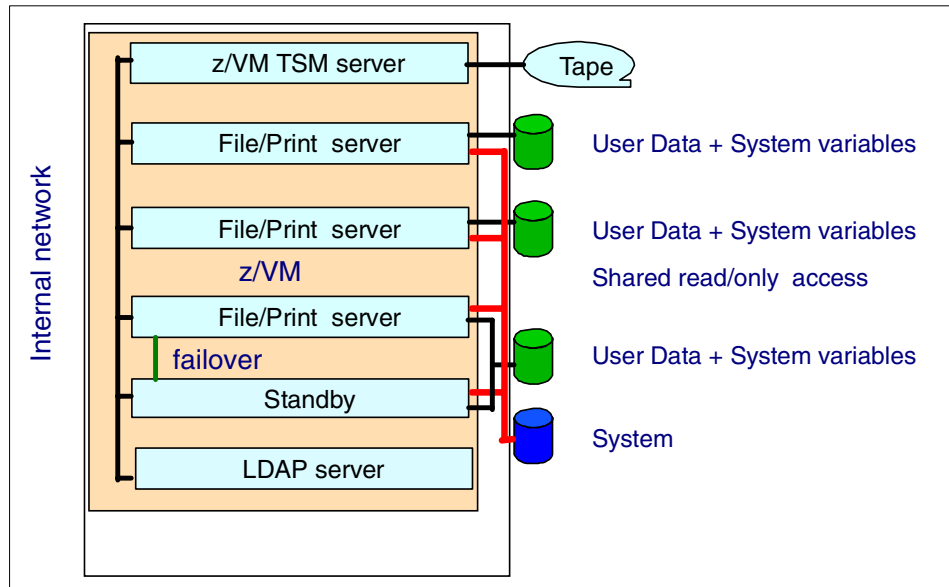


Figure 4-17 Samba file and print on zSeries

Disk storage for the operating system and application code is shared by all 150 virtual servers while they must be duplicated in the distributed case. Sharing the code offers the advantages of controlling the cost of disk storage and ease of maintenance. Since only one copy of the code exists, it must only be serviced in a single place, which results in savings in system management software and staffing. Although only one copy of the code is used, 150 servers are in operation. As in the distributed case, they issue messages, which must be monitored and responded to in some cases. The PROP function in z/VM is designed for that purpose, avoiding the cost of additional software to perform the task. Additional systems management software are required in complex environments.

z/VM can help you to partition the disk space available to the servers. In the distributed model, increasing disk space meant that you needed to buy new disks. If there was free disk space on other servers, it could not be moved to the server that needed the increased disk storage. In the S/390 solution, z/VM allows you to partition disks into minidisks that you can easily allocate to a server that requires additional disk space.

Backup is done again by a backup client. In the S/390 case, it's done by the Tivoli Systems Manager Backup server, which is pre-installed on z/VM. The Tivoli Systems Manager server communicates with clients on the virtual instances to perform backup operations.

Another factor to consider is that the Enterprise Storage Server disk solution is unique and far more powerful than dedicated disks on stand-alone systems. The reliability characteristics of the Enterprise Storage Server include such advanced functions in the unit as FlashCopy so you can create a copy of a disk volume instantly. This is extremely useful for minimizing downtime for backups or for creating test data for new applications.

Cost comparison: Three-year TCO

In this cost comparison, the cost of outage figures highly. Significant savings can be made by reducing these potential costs. People are also a major cost factor that can be reduced.

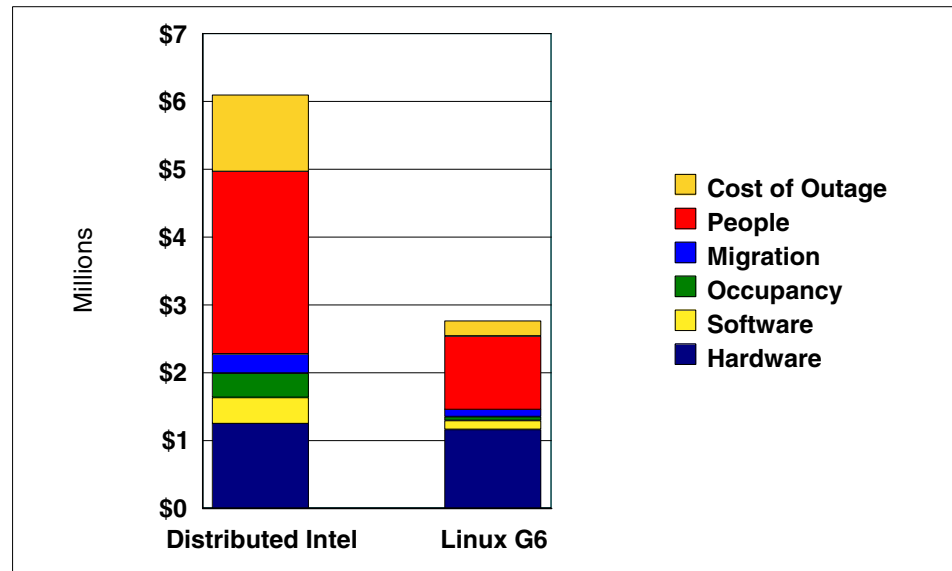


Figure 4-18 File/print serving with Samba cost comparison

4.5.4 IBM Software for Linux on zSeries

Table 4-2 through Table 4-7 on page 108 list IBM middleware that is available to run on Linux for zSeries. Additional solutions are expected to be available in the future.

Table 4-2 zSeries and DB2 software matrix

DB2 Intelligent Miner Visualization	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Net Search Extender	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Net.Data	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Personal Developer's Edition	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Runtime Client	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Spatial Extender	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Universal Developer's Edition	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
DB2 Universal Database Enterprise Server Edition	8.1	zSeries	See DB2 for Linux website for supported platforms and distributions
Product Name	Version/Release	Hardware	Kernel/ Distribution
IBM			
IBM Application Workload Modeler for Linux on zSeries	1.1	zSeries	Red Hat Linux 7.2 for IBM S/390 SLES 7 (64-bit version required for SAP ICL client simulation)
IBM Application Workload Modeler for Linux on zSeries	R1	zSeries	Red Hat Linux 7.2 for IBM S/390 SLES 7 (64-bit version required for SAP ICL client simulation)
IBM CICS Transaction Gateway	5	zSeries	Linux for S/390® SuSE 7.0 (Kernel level 2.2.16 & 2.4) or Turbolinux V6 or Turbolinux V6.5
IBM Developer Kit for Linux , Java 2 Technology Edition Java Technology development environment	1.4	zSeries	Supported Platforms

Table 4-3 IBM IMS™ and Informix

Product Name	Version/Release	Hardware	Kernel/ Distribution
NEW! IBM VSE/ESA	2.7	zSeries	Â
IMS Note: IMS Connect only supports Linux - see note under the Kernel/Distribution column for fur			
Information Management System (IMS) Connect	1.1 1.2	zSeries	Note: IMS Connect running under z/os or os/390 can be used by an application running on Linux s/390 or zSeries to access local applications and data residing on IMS
Informix			
Informix C-ISAM	7.24	zSeries	Kernel 2.2.16 (GLIBC 2.1.3) - Linux for s/390

Table 4-4 IBM Tivoli

Product Name	Version/Release	Hardware	Kernel/ Distribution
Tivoli Access Manager for e-business	4.1	zSeries	Access Manager Base: SLES 7 for S/390® and IBM zSeries (2.4.7 kernel, 31 bit) and SLES 7 for 64 bit zSeries (2.4.17 kernel, 64 bit with a 31-bit compatible mode) Access Manager WebSEAL: SLES 7 for S/390 and zSeries (2.4.7 kernel, 31 bit) and SLES 7 for 64 bit zSeries (2.4.17 kernel, 64 bit with a 31-bit compatible mode)
Tivoli Access Manager for e-business (formerly known as Tivoli Policy Director / Tivoli Policy Director for Application Servers) Access Manager Base Access Manager WebSEAL	3.9	zSeries	SuSE Linux Enterprise Server 7 for s/390 and IBM zSeries Kernel 2.4 Client and Server (32 bit mode)
Tivoli Access Manager for Operating Systems	4.1	zSeries	SuSE Enterprise Level Server 7 for zSeries Red Hat 7.2 for zSeries
Tivoli Configuration Manager	4.2	zSeries	Server and Client: SLES 7 for s/390 and zSeries 2.4 Kernel 32 bit
Tivoli Decision Support for OS/390	1.5.1	zSeries	SuSE 7.0 for s/390 (2.2 Kernel) TurboLinux 6.5 (2.2 Kernel) Server Only
Tivoli Distributed Monitoring	4.1	zSeries	SuSE SLES 7.0 for zSeries Client and Server
Tivoli Enterprise Console	3.8	zSeries	Enterprise Console Server, UI Server, Network Management Engine and Endpoint: SuSE Linux V7.0 for zSeries

Table 4-5 IBM Tivoli continued and UniData™

Product Name	Version/Release	Hardware	Kernel/ Distribution
Tivoli Enterprise Console	3.7.1	zSeries	SuSE SLES 7 on zSeries Client and Server
Tivoli Identity Manager	1.1	zSeries	SuSE SLES 7 Client Only
Tivoli Management Framework	4.1	zSeries	Red Hat 7.2 SLES 7.2
Tivoli Monitoring for Transaction Performance	5.1	zSeries	Quality of Service Endpoints: Red Hat Linux for zSeries 7.1 TurboLinux Server 7 for zSeries SuSe Linux Enterprise Server 7.0 for zSeries
Tivoli Monitoring (formerly known as Tivoli Distributed Monitoring)	5.1.1	zSeries	SuSE 7.0 for zSeries (Client and Server) Red Hat Server 6 for zSeries (Client and Server)
Tivoli Monitoring (formerly known as Tivoli Distributed Monitoring)	5.1	zSeries	SuSE SLES 7.0 for s/390 Client and Server
Tivoli Netview for z/OS MSM Agents (Client Only) NMC Server (Server Only)	5.1	zSeries	Red Hat 7.2 zSeries SuSE SLES 7 for zSeries
Tivoli Remote Control	3.8	zSeries	SuSE Linux 7.0 for S/390 (Server and Gateway) SuSe 7.3
Tivoli Risk Manager	4.1	zSeries	SLES 7 for S/390 and zSeries Client Only
Tivoli Risk Manager	3.8	zSeries	SuSE SLES 7 Client Only
Tivoli Software Distribution	4	zSeries	2.2 Kernel SuSE SLES 7.0 for S/390 Client and Server
Tivoli Storage Manager	5.1.5	zSeries	Back up Archive Client Only: SuSe Linux Enterprise Server 7 for S390 and zSeries
Tivoli Storage Manager	5.1	zSeries	SuSE SLES 7.0 Client Only
Tivoli Storage Manager Client	4.2	zSeries	Linux Kernel 2.2.16 or later SuSE 7.0
Tivoli Switch Analyzer	1.2	zSeries	Â
IBM Tivoli System Automation for Linux	1.1	zSeries	SLES 7
Tivoli User Admin	3.8	zSeries	SuSE SLES 7.0 Client Only
Tivoli Workload Scheduler Tivoli Workload Scheduler Fault Tolerant Scheduling Agent	8.1	zSeries	SuSE Linux Enterprise Server (SLES) 7 for S/390 Red Hat 7.2 zSeries Turbo Linux 7 on zSeries Client Only
IBM UniData	5.2x	zSeries	SuSE for s/390 7.0

Table 4-6 IBM WebSphere

Product Name	Version/Release	Hardware	Kernel/ Distribution
WebSphere			
WebSphere Application Server	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Application Server Advanced Developer Edition for Linux	4	zSeries	SuSE 7.0 - Kernel 2.2.16 Turbolinux Server 6.5 - Kernel 2.2
WebSphere Application Server Advanced Edition for Linux	4	zSeries	SuSE 7.0 (Kernel 2.2.16) Turbolinux Server 6.5 (Kernel level 2.2.) WAS V4.0.3 on zLinux on SuSE 31b 2.4 kernel
WebSphere Application Server Advanced Single Server Edition for Linux	4	zSeries	SuSE Linux 2.4 Kernel
WebSphere Application Server Advanced Edition for Linux	3.5	zSeries	SuSE 7.0 (Kernel 2.2.16) TurboLinux Server for 390 6.0 (Kernel 2.2)
WebSphere Application Server for Developers	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Application Server Enterprise	5	zSeries	SLES with 2.4 Kernel
WebSphere Application Server Enterprise for Developers	5	zSeries	SLES with 2.4 Kernel
WebSphere Application Server Network Deployment	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Commerce Business Edition for Linux	5.4	zSeries	2.4 Kernel SuSE Linux Enterprise Server (SLES) 7.0
WebSphere Host on Demand	7	zSeries	Â
WebSphere Host on Demand	6	zSeries	Server: Linux on zSeries

Table 4-7 WebSphere continued

Product Name	Version/Release	Hardware	Kernel/ Distribution
WebSphere			
WebSphere Application Server	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Application Server Advanced Developer Edition for Linux	4	zSeries	SuSE 7.0 - Kernel 2.2.16 Turbolinux Server 6.5 - Kernel 2.2
WebSphere Application Server Advanced Edition for Linux	4	zSeries	SuSE 7.0 (Kernel 2.2.16) Turbolinux Server 6.5 (Kernel level 2.2.) WAS V4.0.3 on zLinux on SuSE 31b 2.4 kernel
WebSphere Application Server Advanced Single Server Edition for Linux	4	zSeries	SuSE Linux 2.4 Kernel
WebSphere Application Server Advanced Edition for Linux	3.5	zSeries	SuSE 7.0 (Kernel 2.2.16) TurboLinux Server for 390 6.0 (Kernel 2.2)
WebSphere Application Server for Developers	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Application Server Enterprise	5	zSeries	SLES with 2.4 Kernel
WebSphere Application Server Enterprise for Developers	5	zSeries	SLES with 2.4 Kernel
WebSphere Application Server Network Deployment	5	zSeries	SLES (SuSe) 7 or Red Hat 7.2 based on kernel 2.4
WebSphere Commerce Business Edition for Linux	5.4	zSeries	2.4 Kernel SuSE Linux Enterprise Server (SLES) 7.0
WebSphere Host on Demand	7	zSeries	Â
WebSphere Host on Demand	6	zSeries	Server: Linux on zSeries

4.6 How IBM customers are using Linux

The zSeries and S/390 servers are capable of running several operating environments. z/OS, OS/390, and VM can provide industry leading reliability and stability. The system has evolved over 30 years and today can run most integrated applications. In the IT world, zSeries is often called the “Big Iron” to describe its mainframe character and uniqueness in the market.

Today many different applications are available for the zSeries. The following sections describe a few of the Linux-related architected solutions.

4.6.1 T-Com

T-Com, a division of the Deutsche Telekom AG, Germany's former public telephone provider, chose Linux on zSeries to implement a high availability environment for its internal Web platform and mail services. Previously these services ran as UNIX services on 25 Sun servers. T-Com went to Linux on zSeries because of the combination of the open system and mainframe technology ensuring scalability and high availability. Using the Linux on zSeries solution, T-Com significantly decreased operating costs, floor space, and infrastructure expenditures.

By consolidating 25 Sun servers onto a zSeries mainframe running Linux, Deutsche Telekom improved their IT system and increased the availability of software programs. Their overall availability, disaster recovery preparedness, and security has increased because the mirrored mainframe architecture facilitates disaster protection and recovery available in three locations.

Backups can run in real time at different locations. That way, in case of failure of a computing center, another computing center can take over the workload. This is an important criteria for an IT service provider and its customers. The open structure of Linux helps to integrate applications smoothly with low administration costs. Resources are available quickly due to proven good scalability and flexibility. Over 40 percent less utility space and 50 percent less energy consumption are a welcomed advantage to Deutsche Telekom.

An S/390 type 2064-102 with two additional, dedicated Linux processors is used for the production systems. Two dozen virtual Linux partitions run under the hypervisor operating system VM, on which production, backup, and test environments run. Deutsche Telekom uses SuSE Linux for zSeries (S/390). The Perl-based applications, which were consolidated on the mainframe, support the operation of utility servers and mail backup services. In the future, T-Com plans to port their domain name servers onto the mainframe.

For the complete story, see:

http://www.ibm.com/servers/eserver/zseries/news/pressreleases/2002/zlinux_deutschele_07-19-02.html

4.6.2 Boscov's Department Stores

Boscov's is a private, family-owned department store that currently operates 37 stores in Pennsylvania, Delaware, New Jersey, New York, and Maryland in the United States. All of the stores are within an approximate 300 mile radius of the

company's headquarters and the central distribution center in Reading, Pennsylvania. Currently, the company has about 9,000 full- and part-time employees, excluding seasonal help. With an annual revenue of \$1 billion U.S., it has grown by acquiring stores, opening new stores, and catering to a loyal customer base.

Boscov's needed to stop the increasing growth of their Windows NT server farm and support staff while improving availability, scalability, the backup and recovery for many different applications. Boscov's also needed to support the growth in existing CICS®, DB2, and IMS business-critical applications on z/OS v1.1, as well as all new critical application software for Linux on zSeries.

Boscov's relied on mainframes for key business applications for many years and has a significant investment in zSeries including S/390, CICS, DB2, and IMS. The mainframe traditionally offered the characteristics needed for their business-critical computing such as CICS and DB2. When IBM added the ability to host the Linux operating system on their zSeries and S/390 mainframes, it enabled the mainframe's high quality of services to be extended to the many existing distributed applications running within the server farm. The company evaluated the mainframe-based Linux solution capabilities and benefits, which include its unparalleled flexibility, processing power, and ability to host e-commerce applications running under WebSphere Commerce Suite. Next, they decided to consolidate the server farm applications onto the mainframe running Linux for zSeries and S/390 Version 5.1.

Boscov's replaced an S/390 G3 server Model R44 with a zSeries 900 Model 102 in July 2001 to handle existing z/OS business-critical applications and to consolidate the server farm applications onto Linux virtual servers on the mainframe. Boscov's deployed SuSE Linux virtual servers hosted on z/VM in an LPAR powered by one IFL engine in a zSeries 900 to do the job. The zSeries 900 is also used to support new Linux applications. This includes support for planned enhancements to the Boscov's Web site by offering advanced e-commerce applications developed with IBM VisualAge® for Java and WebSphere development tools. They also included support for anticipated growth in the existing business-critical z/OS applications.

A proof of concept for Boscov's solution is:

► **Scalability**

- The applications in 12 of 44 servers in the server farm were ported to Linux virtual servers on the mainframe.
- Four Web servers were replaced by one WebSphere Commerce Suite Pro.

- ▶ **Total cost of ownership**
 - Mainframe application (z/OS) software and environmental costs did not increase as the (IFL powered) Linux application workload increased.
 - The cost to maintain discrete Windows NT servers was reduced or eliminated.
 - The amount of time needed to administer the mainframe was less than that of a server farm.
- ▶ **Flexibility**
 - The zSeries 900 efficiently runs mixed workloads including PeopleSoft, Data Warehouse, IMS, DB2, and Linux applications.
 - The scalability that allows for thousands of Linux for zSeries virtual servers and partitions allows for the consolidation of all workloads to one central machine.
 - Linux for zSeries also offers workload balancing between different Linux LPARs and z/OS DB2 partitions.
- ▶ **Availability**
 - The Parallel Sysplex environment in the single zSeries 900 mainframe provides 24-hour availability for e-commerce-generated credit card authorization using the CICS Point-of-Sale applications.

You can find this story on the Web at:

<http://www.ibm.com/servers/eserver/zseries/os/linux/css/boscovs1.html>

4.6.3 zSeries services solutions

IBM Global Services offers many different services for zSeries and Linux solutions. The following sections provide a glance at the available offerings. For more details, see the IBM Global Services Web site at:

http://www.ibm.com/services/e-business/linux_2.html

Operational support services

Support Line for Linux provides a comprehensive around-the-clock enterprise-level remote usage and defect support for major distributions of the Linux operating system. It also provides such support for all IBM and many non-IBM applications that operate in a Linux environment.

IGS portfolio of services

IGS offers a comprehensive portfolio of Linux services including:

- ▶ **IBM Migration Services:** Consolidates file/print and Web Serving workloads to Linux for zSeries. This service helps customers to consolidate workloads from distributed server farms by providing an assessment of which workloads can be consolidated onto Linux on zSeries. Then it completes the consolidation for them.
- ▶ **IBM Migration Services:** Sendmail Advanced Message Server for Linux. This service migrates a customer's existing Sendmail servers and users (on Windows, UNIX, or Linux servers) to Sendmail Advanced Message Server (SAMS) running on Linux for zSeries.
- ▶ **IBM Migration Services – Bynari Insight Server for Linux:** This service migrates a customer's existing Windows Exchange servers and users to Bynari Insight Server running on Linux for zSeries.

IGS also offers locally delivered services including:

- ▶ Linux Rapid Deployment Solution for installing Linux; installing, setting up, and enabling TCP/IP; and installing and configuring Samba or Apache. This service enables a customer to expedite the deployment of applications on their Linux for zSeries server.
- ▶ Middleware, Database and Connector installation and configuration services, for WebSphere Application Server Advanced Edition, DB2, MQSeries, etc.

Linux Strategy Workshop

This offering includes two days of information gathering at the customer site, a one-day workshop activity, and one day to finalize the conclusions. Based on the customer interviews during the first two days, a tailored Linux strategy is developed with the decision makers and IT management personnel. At the end, the customer receives a high-level, but customized Linux strategy. The decision makers and IT management personnel also receive an overview of what Linux can do for them and the cost efficiencies gained by using Linux.

Linux Solution for e-business

The Linux Solution for e-business service provides rapid deployment of a customer's Linux and zSeries-based e-business environment. This solution is tailored to fit the specific needs of the customers by allowing them to select a wide variety of middleware and applications that will be installed and configured on a base of an IBM-supported Linux operating system.

VM and Linux installation support

Beginning with the LPAR configuration of the zSeries, this offering helps customers to implement Linux in virtual machines running under VM. The service includes installation of VM and Linux as a VM guest.

Linux high availability for zSeries installation support

This specialized offering implements a high availability solution using Linux on zSeries.

Bynari installation support

Bynari Insight Server for Linux services provides a compelling cost and systems management proposition for customers to consolidate messaging workload from their Windows NT server farms to multiple Linux images running on a single IBM @server.

Customers often have Microsoft Exchange as their messaging server platform. IBM Migration Services can provide services to migrate to Bynari Insight Server for Linux. This results in server consolidation of those workloads that can minimize the TCO, enable system management capabilities, and enhance application availability.

Migration to Sendmail Advanced Message Server for Linux

The primary target of this offering is customers who are already using Sendmail on UNIX-like platforms or on multiple servers. During the project, all mail servers are consolidated onto one central Sendmail Advanced Message Server running on Linux on zSeries.

DB2 Connect™ installation support

DB2 Connect enables the connection of applications to the mainframe databases. This service sets up DB2 Connect to make the host data directly available to clients as PCs or network clients.

Workload and server consolidation

Workload and server consolidation is one of the major driving factors in the IT industry. The zSeries is a target platform for consolidation issues. Therefore, IBM Global Services offers several solutions for customers that need to consolidate. Some of the available services are:

- ▶ **Consolidation for file/print services and Web serving workloads**
- ▶ **Migration services for DB2 Universal Database for S/390:** This offering includes a migration to DB2 Version 7 on OS/390 and DB2 for Linux running under a system image (LPAR), a guest system (VM), or stand-alone Linux.

The project includes installation, customization, and verification, plus basic skill training for operations staff.

- ▶ **SmoothStart™ Services for MQSeries:** These standardized projects contain installation planning and on-site project management to ensure a quick and easy installation of MQSeries. These projects also offer a basic skill training for staff.



xSeries and Linux

This chapter provides an overview of the Linux implementation, solutions, and resources available for the xSeries server.

This chapter covers the following topics:

- ▶ How Linux runs on the xSeries server
- ▶ xSeries customers and Linux
- ▶ xSeries and Linux as a server consolidation solution
- ▶ xSeries Linux customer
- ▶ Linux in a distributed enterprise
- ▶ Linux applications for xSeries
- ▶ Infrastructure solutions for xSeries
- ▶ xSeries architected solutions

5.1 Linux and xSeries

In the 1960s, IBM introduced the flexibility of the System/360, and the monolithic mainframe world changed forever. In the 1980s, IBM introduced the IBM PC, which used off-the-shelf parts, and the concept of a personal computer changed forever. Now IBM unveils modular, scalable server technology, and the world of industry-standard servers will never be the same.

IBM @server X-Architecture™ technology is a blueprint for extending the benefits of advanced mainframe technologies to IBM Intel processor-based servers. These benefits are in the areas of availability, scalability, systems management, service, and support. IBM has been delivering on the promise of the X-Architecture model since 1998. It has included such innovative technologies as Active™ PCI, C2T Interconnect™ cabling, Chipkill™ memory, Predictive Failure Analysis®, Light Path Diagnostics™, and IBM Director Software Rejuvenation.

Today, IBM continues to build on the X-Architecture blueprint with Enterprise X-Architecture technologies. These technologies yield revolutionary advances in the input/output (I/O), memory, and performance of xSeries servers. This peerless new server design creates a flexible “pay as you grow” approach to buying high-end 32-bit and 64-bit xSeries servers. The results are systems that can be scaled quickly, easily, and inexpensively.

Enterprise X-Architecture technology enables the following capabilities:

- ▶ XpandOnDemand™ scalability
- ▶ System partitioning
- ▶ PCI-X I/O subsystem: Active™ PCI-X
- ▶ Remote I/O
- ▶ Active Memory™
 - 256-GB memory capacity
 - High-speed (DDR) memory
 - Memory ProteXion™
 - Chipkill™ memory
 - Memory mirroring
 - Hot-add/hot-swap memory
- ▶ 533 MHz front side bus (FSB)
- ▶ Xcel4 Server Accelerator Cache
- ▶ Real Time Diagnostics

If you are looking for all of these abilities in an industry-standard server today, they are available from IBM. These features deliver application flexibility, innovative technology, and new tools for managing e-business. They bring to industry-standard servers the kinds of capabilities that were formerly available

only to users of mainframes and other high-end systems. Combined with existing X-Architecture technologies, these innovations result in unprecedented “economies of scalability,” unmatched flexibility, and new levels of server availability and performance.

In Intel-Architecture (IA) servers, Linux has evolved quickly. When compared with other operating systems, new drivers and functions were incorporated one after another from the great development environment surrounding open source software. Since distributions with kernel-2.4 appeared, Linux for Intel-Architecture servers has been accepted as an operating system with the function and stability that can fully be used as an enterprise operating system. While the market share of other platforms has increase little, the Intel-Architecture server market with Linux has increased dramatically.

5.1.1 Linux on xSeries

The xSeries server is based on X-Architecture, where IBM linked advanced technology from years of enterprise server development and the industry-standard technology represented by the Intel CPU (Figure 5-1).

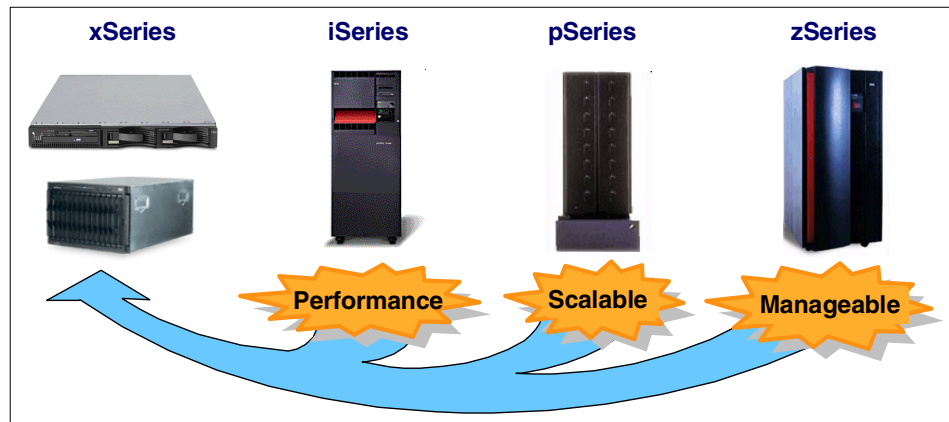


Figure 5-1 The concept of xSeries: The best from all IBM @server product lines

While the xSeries server delivers advanced technology to customers at a low price, many applications and services with industry standards are also usable. Recently, IBM delivered X-Architecture for high-end server technology. In the autumn 2001, IBM announced the Enterprise X-Architecture. This new technology raises the bar on scalability, flexibility, availability, and performance by a dynamic arrangement of server resources. IBM offers Enterprise X-Architecture through core logic developed by IBM and realized through the IBM

XA-chip set corresponding to the IA-32 and IA-64 processors. This technology starts in the xSeries Models 360 and 440.

Operating system support for xSeries changes frequently. Table 5-1 lists several helpful Web sites to which you can refer for platform-based information about the services that are available for the xSeries server. This information is available at the IBM ServerProven® site on the Web at:

<http://www.pc.ibm.com/us/compat/nos/matrix.shtml>

You can find the drivers and documents for xSeries on the Web at:

http://www.pc.ibm.com/support?lang=en_US&page=brand&brand=IBM+PC+Server

Table 5-1 xSeries Linux information matrix

Information	Location on the Web
Linux for the xSeries information	http://www.pc.ibm.com/us/eserver/xseries/index.html
eServer BladeCenter™	http://www.ibm.com/servers/eserver/blades/
Linux clustering on xSeries	http://www.ibm.com/servers/eserver/clusters/
Linux for xSeries independent software vendor (ISV) resources	http://www06.ibm.com/developerworks/offers/linux-speed-start/isv.html

Linux and functions of X-Architecture

X-Architecture technology is an evolving blueprint for xSeries servers that is drawn from the vast enterprise server heritage of IBM. xSeries engineers took the technologies that already revolutionized larger IBM systems and brought them to the Intel-based platform. The xSeries server is the industry-standard server, designed to provide enterprise power, scalability, control, and service at attractive prices. Table 7-1 lists the Linux support for X-Architecture.

Table 5-2 Linux support and X-Architecture

Functions	Supported by Linux
Active PCI (hot-add and hot-swap)	IBM-supplied drivers for Linux
Redundant NIC	IBM-supplied drivers for Linux
Advanced Systems Management processors	IBM-supplied drivers for Linux
IBM Director 4.1	IBM supports native Linux
Storage area network (SAN)	IBM-supplied drivers for Linux

For more information about X-Architecture, see:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/>

You can download drivers from the Web at:

http://www.pc.ibm.com/support?lang=en_US&page=brand&brand=IBM+PC+Server&doctype=Downloadable+files

Enterprise X-Architecture with Linux

With support for both Microsoft and Linux operating systems, the application flexibility of xSeries servers is extended. The hardware is optimized for the latest industry platforms. This enables customers to choose from a broad portfolio of applications to best suit their business needs. To further enhance flexibility, physical partitioning and a common platform design help to ease the transition from 32-bit to 64-bit high-end computing.

New tools make systems management easier than ever before. With self-diagnosing and self-healing technologies, such as Active PCI-X and third-generation Chipkill memory, systems are designed to stay up and running continuously. The xSeries server provides high availability and exceptional performance for systems that need to be scaled quickly, easily, and inexpensively. Table 5-3 shows the Linux support that is available for Enterprise X-Architecture.

Table 5-3 Linux support with Enterprise X-Architecture

Functions	Correspondence by Linux
XpandOnDemand Scalability	Red Hat AS2.1, SuSE8.0
Remote I/O	Kernel-2.4.9-31 or later

XpandOnDemand is part of the X-Architecture. The power of a 16-way server is now available with the xSeries 440 to take charge of advanced enterprise applications and drive a higher level of performance on this flagship server. Powered by Enterprise X-Architecture™ technology, these 4U rack-optimized, industry-standard servers support up to 16-way processing by interconnecting two xSeries 440 chassis as a single 8U configuration. This makes it one of the most rack-dense 16-way servers in the world.

XpandOnDemand offers scalability in the way you buy and grow. This revolutionizes data center servers with a modular, pay-as-you-grow building block design. This design offers low entry price points and upgradability to powerful 16-way SMP and Remote I/O. XpandOnDemand allows you to purchase only the performance and I/O capacity that you need, when you need it, without having to buy costly upfront infrastructure.

For more information about Enterprise X-Architecture, see:

<http://www.pc.ibm.com/us/eserver/xseries/xarchitecture/enterprise/index.html>

5.1.2 xSeries Linux distributions

Linux for the xSeries server is available from leading Linux distributors Red Hat, SuSE, Turbolinux, and The SCO Group (previously known as Caldera). Recently The SCO Group, Conectiva, SuSE, and Turbolinux partnered to create UnitedLinux. This new initiative will streamline Linux development and certification around a global, uniform distribution of Linux, designed for business. IBM will continue to support Red Hat Linux and fully support UnitedLinux, which will make it easier to create a wide variety of Linux solutions.

For a detailed explanation of each distribution, see 1.9, “The Linux distributions” on page 18. You can also refer to the following Web sites:

- ▶ Red Hat Linux:

<http://www.redhat.com>

- ▶ SuSE Linux:

<http://www.suse.com>

- ▶ Turbolinux:

<http://www.turbolinux.com>

- ▶ SCO Group:

<http://www.caldera.com/>

The packages of these distributions are updated frequently. To download fixes and security packages, see the following Web sites:

- ▶ Red Hat:

<http://www.redhat.com/apps/support/errata/>

- ▶ SuSE:

<http://www.suse.com/us/private/download/updates/index.html>

- ▶ Turbolinux:

<http://www.turbolinux.com/security/>

- ▶ SCO Group:

<http://www.caldera.com/support/download.html>

5.2 Intel-based server customers and Linux

Leading analysts consider Intel-based servers to be the most adaptable of all servers, including mainframe, UNIX servers, and business servers. Analysts estimate that the market for Intel-based servers (IA-32 and IA-64) will grow at about 10% compound annual growth rate (CAGR) between now and 2005. They see the 64-bit market for Intel servers showing significant growth. In the expanded market for Intel-based servers, Linux shows the fastest growth of all operating systems.

Most of the current Linux usage comes from small businesses. IBM Market Research shows that most of the deployment of Linux applications is in the infrastructure server area. Within two years, the share of more important applications, such as the database server, will expand significantly. Table 5-4 shows several motivations to move to a Linux platform based on IBM Market Research from February 2001. There is an increasing interest in Linux for line-of-business (LOB) use in large enterprises because of its increasing availability, reliability, and customer successes.

Table 5-4 Worldwide Linux application deployment

Application deployment	Today	Within two years
Web server	51%	76%
Web application server	34%	61%
Firewall server	35%	60%
E-mail server	40%	59%
Network server	39%	59%
Database server	29%	55%
Development system	29%	55%
Workgroup server	18%	40%
Transaction server	11%	22%

The most influential factors in moving to Linux can reduce overall cost. Reliability is proving to be a significant factor. As Linux moves into the mainstream, industry acceptance will increase as indicated by the chart in Figure 5-2.

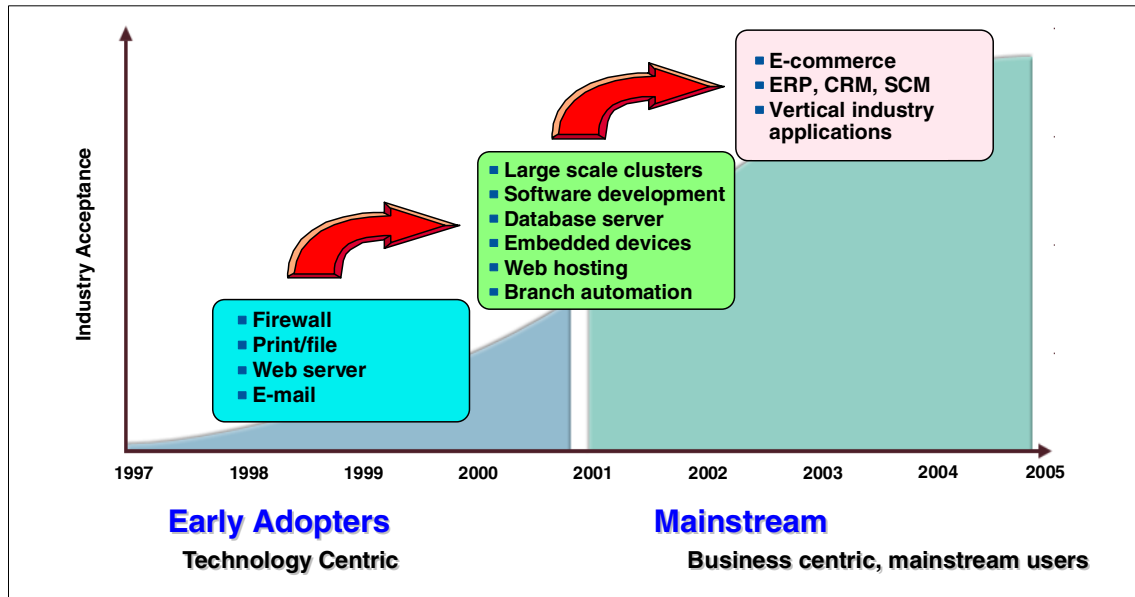


Figure 5-2 Expected growth rate of Linux acceptance

5.3 Consolidation

There are four types of workload consolidation (also referred to as *server consolidation*). Each type offers significant benefits in the following areas:

- ▶ Reduced administrative costs because of central management
- ▶ Better management of system proliferation and more consistent architecture
- ▶ Management of purchasing to achieve volume purchasing discounts
- ▶ Consistent process for security, operating system levels, and updates

The four types of consolidation are:

- ▶ **Consolidating multiple systems previously spread around the enterprise into fewer, more centralized locations**

For this type of consolidation, xSeries and Linux together provide the following benefits:

- Reduced floor space: The xSeries features rack-dense solutions and cable-chaining technology.
- IBM Director: This includes a function called Update Xpress, which allows you to propagate system images and updates to all consolidated systems.

- Workload management solutions: They reschedule work from heavily loaded systems to more lightly loaded systems.
 - Availability via redundancy within systems and across systems: With a consolidated environment, you can configure your solutions for maximum availability. The xSeries server already has redundancy within systems for such things as power supplies. Now, in a consolidated environment, you can set up another system to handle work from a failing system or from a system you may want to take offline to replace.
- **Consolidating many systems into a reduced number of larger servers**
- The xSeries Model 440 offers a strong value proposition with its *Pay as You Grow* feature. This feature allows you to buy only the computer resources you need for the first stage of your consolidation efforts and to add additional computer capacity as your consolidation efforts grow or business expands.
 - The remote I/O feature of the xSeries Model 440 allows you to increase the I/O capacity of your systems without throwing out the base system. You can expand your computer power with Pay as You Grow and customize the I/O needs as your consolidation requirements change.
 - Virtual or physical partitioning: The xSeries server allows you to create physical partitions, each with its own memory and I/O. In addition, VMware offers virtual partitions. With both partitioning configurations, failover between virtual or physical partitions is possible.
 - This environment requires fewer administrators.
 - This environment offers separate resource allocation, billing, and server-level agreements.
 - VMware and virtual partitions: VMware offers several key benefits as you implement server consolidation. These benefits include:
 - The ability to run Windows and Linux partitions side by side. You can run Windows applications in one partition and Linux applications in a second partition. VMware supports both SuSE and Red Hat as Linux operating environments.
 - The ability to change partitions dynamically as computer needs change. You can add a new partition or alter the I/O resources dedicated to it in a few minutes
 - You can set up partitions and customize them to serve specific workload requirements. For example, you can set up two Web-serving partitions, a print-serving partition, and two mail-serving partitions. The mail- and Web-serving partitions may likely have mission-critical requirements, so you can set up the second partition as a backup partition that you can shift workload to if one partition has capacity

problems. The print-serving partitions are likely set up with less I/O capacity dedicated to them than to the mail-serving partition. This makes customization very flexible for your specific workloads.

► **Consolidating data**

The key differentiators and benefits that xSeries and Linux together bring to this area include:

- The ability to create highly available data configurations with system area networks and shared storage environments
- Data mirroring to ensure availability of data if disk failures occur
- Improved security with centralized security policies rather than fragmented data scattered around the enterprise
- Fewer database administrators required

► **Consolidating applications**

The key differentiators and benefits that xSeries and Linux together bring to this area include those mentioned above as well as these:

- A reduction in application licensing costs by consolidating application usage on one system rather than multiple systems across the enterprise
- The ability to better manage new application migration plans with operating system or application-level changes
- A reduction in compute capacity requirements by load balancing multiple application requirements across systems

VMware ESX Server

With the introduction of VMware ESX Server, the xSeries server now offers virtual partition functionality that was previously only found in main-frame environments. VMware ESX Server simplifies server infrastructure by partitioning and isolating server resources in secure and transportable virtual machines. VMware ESX Server enables you to remotely manage, automatically provision, and standardize these server resources on a uniform platform.

Figure 5-3 shows the architecture of VMware ESX Server. Operating systems and applications are isolated in separate virtual machines.

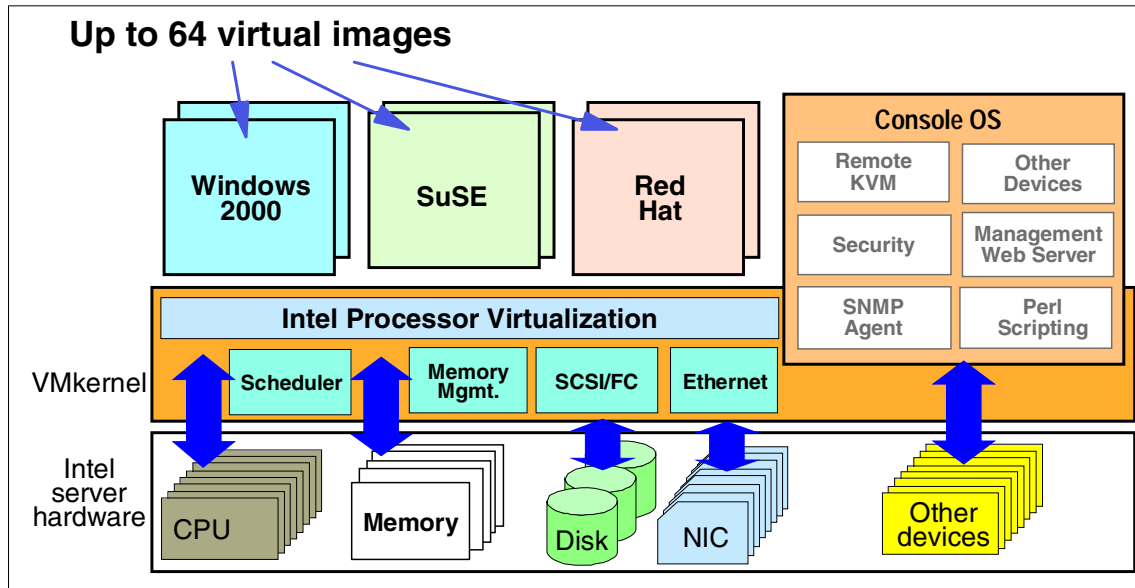


Figure 5-3 The architecture of VMware ESX Server

VMware ESX Server transforms physical systems into a pool of logical multiple virtual machines that reside on a single piece of hardware. System resources are dynamically allocated to any operating system based on immediate need. This gives you mainframe-class capacity utilization and mainframe-class control of the server infrastructure.

VMware ESX Server runs directly on system hardware to provide a secure, uniform platform for easily deploying, managing, and remotely controlling more operating systems. Advanced resource management controls allow you to guarantee service levels of CPU, memory, networking, disk resources, and failover between partitions.

VMware ESX Server offers the following benefits:

- ▶ You can move applications running on dedicated systems into separate virtual machines on a single, more reliable, and scalable system.
- ▶ You can manage servers remotely from any location, simplifying server maintenance.
- ▶ With advanced resource management controls, you can guarantee service levels.
- ▶ You can script common monitoring and management tasks.

For more information about VMware ESX Server, see:
http://www.vmware.com/products/server/esx_features.html

5.4 Linux clusters

A clustered system is basically a system that combines two or more computers into a group to improve reliability or processing performance. The types of clusters can be roughly classified as:

- ▶ High-availability cluster
- ▶ Load-balancing cluster
- ▶ High-performance computing
- ▶ Tightly coupled clusters

For a detailed explanation and references about clusters, see Chapter 9, “Linux clusters” on page 243.

5.4.1 High-availability cluster

As your computer requirements expand, you can choose to expand your systems vertically or horizontally. *Vertical scaling* essentially means implementing larger systems, many times larger than SMP systems. *Horizontal scaling*, or clusters, means implementing multiple smaller systems to handle the workload and leveraging them either as one computer resource or as a large server farm with sophisticated workload balancing.

Cluster implementations begin with high-performance computing environments where customers have hooked many systems or nodes together over a high-speed interconnection. By leveraging software, they split jobs into small chunks to increase speed and efficiency.

Today, clusters are implemented in life sciences, financial, and petroleum environments, where large amounts of data are analyzed and the problems are well suited to be split among multiple machines. Often, clusters are implemented as large server farms, an implementation that has many of the benefits of server consolidation.

IBM is a leader in Intel-processor Linux clusters. The key differentiators include:

- ▶ **Reduced system cost:** With Linux and xSeries solutions, the initial cost is extremely attractive.
- ▶ **System management:** With IBM Director and Cluster System Management, you now have solutions to manage many nodes as one computer resource.

This produces savings in total cost of ownership (TCO) by requiring fewer administrators and improving system availability.

- ▶ **High-availability solutions:** IBM, together with its ISV partners, offers high-availability solutions to allow failover from failing systems within the cluster to nodes that are up and operational.
- ▶ **Cluster integration services:** Clusters are difficult to implement and design. The strategy must include application and workload analysis. It must also include a plan for how all the pieces are put together and how well all the pieces come together at installation time. Cluster integration services provided by an IBM Business Partner or IBM Global Services (IGS) are a key differentiating value proposition that has produced many successful cluster implementations.

High-availability software

The software for high-availability clusters includes the following products. Consult the accompanying Web site for more information:

- ▶ Veritas Cluster Server
<http://www.veritas.com>
- ▶ SteelEye Lifekeeper
<http://www.steeleye.com/products/linux/>
- ▶ Toshiba DNCWARE Cluster Perfect
<http://www.toshiba.co.jp/index.htm>
- ▶ Red Hat Cluster Manager
<http://www.redhat.com/software/rhel/as/cluster/>
- ▶ Linux Virtual Server
<http://www.linuxvirtualserver.org/>
- ▶ Tivoli System Automation
<http://www.ibm.com/software/tivoli/products/sys-auto-linux/>

5.4.2 Load-balancing cluster

Load-balancing clusters raise performance and throughput by distributing many jobs, processes, and requests to two or more servers. Load-balancing cluster software has a flexible setup process, which is a strong point.

Advanced load balancing has an agent that is installed on each nodes. This client performs the health check of the database server, the application server, etc.

Load-balancing software

The major software providers for load-balancing cluster applications include:

- ▶ IBM WebSphere Edge Server
<http://www.ibm.com/software/webservers/edgeserver/index.html>
- ▶ Turbo Linux Cluster Server
<http://www.turbolinux.com/products/tcs/>
- ▶ Ultra Monkey
<http://www.ultramonkey.org/>

5.4.3 High-performance computing

High-performance computing performs calculations that are numerically intense and require high performance and a mass memory. The tasks are carried out by parallel execution of the independent processes.

There are two types of high-performance computing: one that uses a library for message passing and the other that uses a job scheduler. When the library is used for message passing, an independent process is performed in two or more threads and processes.

When the job scheduler is used, one application is simultaneously run in two or more nodes. Each node performs the process with a script or an exclusive tool.

IBM Cluster 1350

Cluster 1350 is the latest generation of affordable integrated Linux cluster offerings from IBM. Cluster 1350 combines the power of the latest xSeries rack-optimized servers with IBM Cluster Systems Management (CSM) for Linux software, IBM storage products, and leading third-party networking components. It offers powerful, flexible solutions for high-performance computing (HPC) and commercial application environments.

Cluster 1350 is available in a wide range of configurations, including single-rack systems of up to 32 cluster nodes and large-scale systems of up to 512 cluster nodes and beyond. Cluster nodes are one- or two-way x335 servers (1U) with Intel Xeon processors running at 2.4, 2.6, or 2.8 GHz. Systems may optionally include up to 32 one- or two-way x345 storage servers (2U), with Intel Xeon processors running at 2.4, 2.6, or 2.8 GHz. In addition, the xSeries Model 345 can be used as a cluster node. Each cluster requires a two-way Model 345 management node.

Cluster 1350 systems are managed by IBM Cluster Systems Management for Linux Version 1.3. This provides resource monitoring, automated operations,

remote hardware control and command execution, configuration file management, and parallel network installation. CSM for Linux helps ease administration and may reduce life-cycle costs by allowing management of an entire Cluster 1350 system from a single management node. As workload demand increases, CSM for Linux allows incremental growth of the Cluster 1350 configuration without necessarily increasing the management complexity.

The highlights of Cluster 1350 include:

- ▶ Hardware
 - xSeries 335 and 345 rack-optimize servers
 - 10/100 Ethernet, Gigabit Ethernet, Myrinet
 - FAStT200 and FAStT700
 - Terminal server and KVM switch
 - 42U rack
- ▶ Software
 - Red Hat 7.2, 7.3, Advanced Server 2.1
 - IBM Cluster Systems Management (CSM) for Linux 1.3
 - IBM General Parallel File System (GPFS) for Linux 1.3
- ▶ Services
 - Base system integration and setup
 - Warranty support
 - Linux Cluster Installation Services
 - Support Line for Linux Cluster

For more information about Cluster 1350, see:

<http://www.ibm.com/servers/eserver/clusters/>

5.5 Linux in a distributed enterprise

Complex enterprises, such as retail stores and banks, often require computers in many sites. The applications run on many geographically dispersed machines and are managed remotely. Because of its low cost and small footprint, Linux on the xSeries is an excellent solution. In distributed enterprises, ease of use and ease of support are critical. Once again, the xSeries and Linux are perfect for this situation because of their high availability and reliability. Linux is also a great solution as a standard application interface that links all applications on the xSeries servers.

For example, Linux provides Burlington Coat Factory with functionality that is the same as or better than Windows NT at about 25% lower TCO, including hardware, software, operating system licensing, and support costs.

Note: At the time Burlington Coat Factory implemented the solution, Sun was charging for the Solaris operating system. However, Sun no longer charges a licensing fee for Solaris for low-end hardware.

Linux provides a standardized interface for the deployment of applications, enabling a reduction in deployment time. With networked terminals, administrators can upgrade applications once per site, rather than once per computer. Linux provides Burlington Coat Factory with a standardized platform for the deployment of new applications, reducing application deployment time by about 50%. With Burlington Coat Factory's Linux solution, application upgrades are about 400% faster with the thin client architecture because programs run remotely and administrators are not required to load as many applications.

Distributed enterprise can reduce support requirements by maintaining one copy of an application at the server level, rather than maintaining multiple copies at the client level. In some cases, support time and personnel required may decrease by up to 90%. Aggregated over many stores or locations, some customers support savings can reach millions of dollars.

5.5.1 Basic implementation strategies

There are three basic implementation strategies for distributed enterprises:

- ▶ **Decentralization**
 - Relocates information technology (IT) operations closer to the user
 - Delivers to the user from a local server that communicates to the user through a thin client
- ▶ **Centralization of corporate data assets**
 - Provides significantly improved business control capabilities and rapid rollout and deployment of enhanced function
- ▶ **Application simplification**
 - Provides identical application functionality to each user via thin client or kiosk
 - Reduces the management and maintenance of servers and applications
 - Reduces TCO

5.5.2 Business benefits

The distributed enterprise solution provides following business benefits:

- ▶ Centralized management via the server and Internet: For a secure, stable desktop environment
- ▶ Accelerated deployment of applications, capability, hardware: For speed and flexibility
- ▶ Access to new and existing applications: That is, seamless access to multiple servers and networks
- ▶ Low TCO:
 - Centralized management; lower training and support cost
 - Rapid application and function deployment
 - Significantly reduced client functionality requirements
 - Reduced software licensing costs

5.5.3 Customer example: E*TRADE Financial

E*TRADE Financial knows what it takes to become a leading financial services brand. With the launch of the www.etrade.com Web site in 1996, the E*TRADE Group, Inc. experienced an immediate explosion in consumer demand for its services. From its beginnings as a discount brokerage firm, the company has grown into a global leader in personal financial services with branded Web sites spanning the globe and total customer assets.

When planning an upgrade to its Web services, the company decided to replace the hardware and software that powers its North American financial services business. After careful analysis, the company replaced its servers with IBM xSeries servers. The company also made a fundamental change in its operating system. It resolved to deploy Linux, because it is an open-source and standards-based platform and holds the promise of cost savings.

With 10,000 concurrent users and four million customer accounts, E*TRADE Financial's success depends on its ability to provide superior customer service. As a fully integrated financial services firm with a significant online presence, providing a high-performance Web site that is up and running whenever customers want to conduct business is the essence of the brand. E*TRADE Financial must also have a very scalable system that can handle the peaks and valleys of the online financial services business.

The company recognized the potential of Linux and needed a complete solution from a vendor who also recognized and embraced Linux technology, and had experience with large Linux deployments. In switching to Linux, E*TRADE Financial was positioned in the vanguard of the open-source movement and as a

financial services technology leader. To make the transition, the company needed a trusted provider. E*TRADE Financial chose IBM.

E*TRADE Financial upgraded its data center operations by deploying 90 xSeries Model 330 servers running Linux. By using Linux, the company sharply reduced software licensing fees and service costs. The company's new deployment is expected to provide the savings and performance levels necessary to meet business goals. In addition, because the solution is based on open standards, E*TRADE Financial has flexibility in the applications the company chooses to deploy.

E*TRADE believes that the combination of the xSeries servers and the IBM commitment to Linux will save the company millions of dollars and enhance the E*TRADE Financial global brand. By using Linux, the company realized sizable software cost savings. By deploying the lower-cost xSeries servers, the company has experienced a rapid return on investment.

For more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/KMOE-5EAMKE?OpenDocument&Site=linuxatibm>

5.6 Linux application solutions

Although the general opinion is that Linux has a limited number of enterprise-level applications available, the Linux applications featured in the IBM Global Solutions Directory has grown to more than 2,900. IBM has 10 porting centers around the world and is continually enhancing this portfolio.

5.6.1 IBM Software for Linux on xSeries

Just about every IBM Software component needed to help customers enable their e-business is available for Linux, including WebSphere, DB2 Universal Database, Tivoli, and Lotus Domino™. The following sections list the readiness of IBM Software for Linux.

Note: The distribution level that IBM Software requires changes frequently. For more information, see Chapter 8, “IBM Software for Linux” on page 199.

WebSphere software platform

WebSphere is infrastructure software for dynamic e-business, delivering a proven, secure, and reliable software portfolio. WebSphere Application Server,

MQSeries, and other product groups and development tools allow the latest Web technology and enable simple extensions of e-business to meet market needs.

For more information about the WebSphere family, see:

<http://www.ibm.com/software/info1/websphere/index.jsp>

DB2 product family

DB2 Universal Database is used as the basis of e-business. It is a powerful database that provides the small and middle-scale user flexibility in response to demand and provides the large-scale user the optimum relational database. DB2 supports industry-standard open-source software such as Linux, Java, and XML.

For more information about the DB2 family, see:

<http://www.ibm.com/software/data/>

Tivoli software product

Broadband provides a common connection for mass communication to the Internet, and e-business expands this connection even further. Many companies are required to provide systems 365 days-a-year, 24 hours-per-day. To provide this service requires an integrated management of the e-business infrastructure.

Tivoli can provide this integrated system management software. The management solution offers flexibility for the needs of e-business, such as reduction of management cost, improvement in reliability of the whole system, and efficient management of the expanding data set. The open architecture, which corresponds to the multiplatform support, and the high scalability, which supports growth and change of business, are key features of the Tivoli product.

For more information about the Tivoli family, see:

<http://www.ibm.com/tivoli>

Lotus Domino

The Domino family provides a multiplatform foundation for collaboration and e-business, driving solutions from corporate messaging to Web-based transactions.

For more information about the Lotus Domino family, see:

<http://www.lotus.com/>

5.6.2 ISVs and application solutions

The list of ISV products that are supported under Linux continues to grow significantly. Many major software vendors have made a commitment to the Linux

platform. The ServerProven program from IBM ensures real business applications and solutions that are tested for the marketplace. To find the right solution, check out the growing list of validated ServerProven applications for Linux on xSeries.

For more information about ServerProven, see:

<http://www.developer.ibm.com/welcome/eserver/eSC.pl?mvcid=Welcme&packageid=000>

Also check the interactive Solution Sizing Tool Web site. It provides the best-fit server solution required to run a particular application environment. Refer to either of the following sites:

<http://myvctest.austin.ibm.com/servlet/ServletEngine?mvcid=Welcme&packageid=000>

<http://www.developer.ibm.com/welcome/myvc.pl>

The sizing tools that are available for Linux solutions include:

- ▶ **Apache HTTP Server for Linux:** The xSeries Apache 1.3 on Linux solution sizing tool is based on recommendations of xSeries servers running Apache 1.3 on Linux and optionally SteelEye's LifeKeeper high availability clustering.
- ▶ **BISON Solution Release A2.9:** The xSeries BISON Solution Release A2.9 solution sizing tool is based on recommendations of xSeries servers running BISON Solution Release A2.9
- ▶ **Changing Worlds ClixSmart Navigator Server:** The xSeries Changing Worlds ClixSmart Navigator Server solution sizing tool is based on performance measurements of xSeries servers running Changing Worlds ClixSmart Navigator Server.
- ▶ **Check Point VPN-1 Pro:** The xSeries Check Point VPN-1 Pro solution sizing tool is based on performance measurements of xSeries servers running Check Point VPN-1 Pro.
- ▶ **DB2 UDB EE 7.2:** The xSeries DB2 UDB EE 7.2 solution sizing tool is based on recommendations of xSeries servers running DB2 UDB EE 7.2.
- ▶ **IBM HTTP Server for Linux:** The xSeries IBM HTTP Server for Linux solution sizing tool is based on recommendations of xSeries servers running IBM HTTP Server for Linux.
- ▶ **Integrated Sensors RTEExpress:** The xSeries and Netfinity® Integrated Sensors RTEExpress solution sizing tool is based on recommendations of xSeries and Netfinity servers running Integrated Sensors RTEExpress.
- ▶ **Lecando Training Server 5:** The xSeries Lecando Training Server 5 solution sizing tool is based on performance measurements of xSeries servers running Lecando Training Server 5.

- ▶ **Linux Gateway (Firewall & Proxy):** IBM xSeries Linux Gateway (Firewall and Proxy) solution sizing tool is based on the Linux Kernel 2.4.5.
- ▶ **mip Webbased Business Solution:** The xSeries mip Webbased Business Solution solution sizing tool is based on performance measurements of xSeries servers running mip Webbased Business Solution.
- ▶ **Net-Tech Linux ASP Solutions:** The xSeries and Netfinity Net-Tech Linux ASP Solutions recommendations are based on performance measurements of xSeries and Netfinity servers running Net-Tech Linux ASP Solutions by Net-Tech Linux Limited.
- ▶ **RealServer 8.01:** The xSeries and Netfinity RealServer 8.01 recommendations are based on performance measurements of xSeries and Netfinity servers provided by RealNetworks.
- ▶ **Samba File and Print Server:** The xSeries Samba File and Print Server solution sizing tool is based on recommendations of xSeries servers running Samba and optionally SteelEye's LifeKeeper high availability clustering.
- ▶ **Sendmail Mail Server:** The IBM xSeries Sendmail Mail Server solution sizing tool is based on recommendations of xSeries servers running Sendmail Mail Server and optionally SteelEye's Lifekeeper high availability clustering.
- ▶ **Shuttlesoft Volanto.Office:** The xSeries and Netfinity Shuttlesoft Volanto.Office recommendations are based on performance measurements of xSeries and Netfinity servers provided by Shuttlesoft.
- ▶ **Srishti PayScribe 3.1 for Linux Quick sizing tool:** The xSeries and Netfinity Srishti PayScribe 3.1 for Linux recommendations are based on performance measurements of xSeries and Netfinity servers provided by Srishti.
- ▶ **Trustix Firewall:** The xSeries Trustix Firewall solution sizing tool is based on recommendations of xSeries servers running Trustix Firewall.
- ▶ **Trustix Mail Server:** The xSeries Trustix Mail Server solution sizing tool recommendations are based on performance measurements of xSeries servers running Trustix Mail Server.
- ▶ **WebSphere Application Server 4.0 for Linux:** The xSeries IBM WebSphere Application Server 4.0 for Linux solution sizing tool is based on recommendations of xSeries servers running IBM WebSphere Application Server 4.0 for Linux.

The following sections list examples of ISV applications that have been ported to Linux.

Non-IBM middleware

Today, Linux is used as an Internet infrastructure server and to build e-business environments. As important middleware, Oracle and Weblogic are mentioned as are IBM products.

- ▶ **Oracle9i Database** is a popular object-oriented relational database. It provides scalability, reliability, and security with rich functions. Oracle9i also has a clustering option named Real Application Clusters (RAC) that provides more scalability, availability, and performance for database applications. For more information, see:

<http://www.oracle.com/ip/deploy/database/oracle9i/>

- ▶ **Oracle9i Application Server (Oracle9iAS)** has an integrated J2EE application server. It provides cost-effective built-in portal function, wireless and voice, Web page caching, powerful business intelligence features, complete integration, and more, integrated in a single product. Oracle9iAS supports major J2EE, Web services, and XML industry standards. For more information, see:

<http://www.oracle.com/ip/deploy/ias/>

- ▶ **BEA Weblogic Server** is a Web application server that offers scalability, pliability, and reliability suitable for being called an enterprise application platform. For more information, see:

<http://www.bea.com/products/weblogic/server/index.shtml>

- ▶ **VMware ESX Server** simplifies server infrastructure by partitioning and isolating server resources in secure and transportable virtual machines. It enables these server resources to be remotely managed, automatically provisioned, and standardized on a uniform platform. For more information, see:

http://www.vmware.com/products/server/esx_features.html

IBM offers outstanding service and support for this product. IBM is the only hardware vendor that offers complete support for Windows, Linux, and key applications all in a VMware environment. VMware provides support for the VMware products only.

The process to obtain this high level of support is:

- a. **Register for Subscription Software Services:** One year of maintenance subscription service is included with your purchase of ESX Server from IBM. This entitles you to any updates or fixes for the VMware ESX Server software released within a year of your purchase, direct from VMware. To receive these updates, you must first register for this service with VMware. You must register for Subscription Software Services within 30 days after purchase.

- b. **Purchase Support for ESX Server:** Support for ESX Server is available through IBM Operational Support Services – Support Line. A single, integrated remote support solution Support Line helps support your complex, distributed IBM and multivendor operating systems and software products. Contact IBM Global Services for more information on the Support Line offering in your area.

Enterprise applications

In recent years, the use of Linux enterprise applications has increased due to the stability of Linux. In the Linux enterprise application arena, companies such as SAP and J.D.Edwards provide solutions. Consider the following solutions:

- ▶ **SAP** is a comprehensive software package that manages the flow of the business processes of a company as represented by R/3. In the ERP software package market, SAP's market share reaches 50% or more and serves as the global standard. In an Intel-Architecture server, since large-scale storage systems or system management software is not perfect, its use in core business is not very prevalent. However, Linux is expected to move into this area in the future. For more information, see:

<http://www.sap.com/linux/>

- ▶ **INSUITE Enterprise** is a Web-type groupware that offers functions for using information quickly and exactly in an all-in-one solution. It can build groupware, EIP, and a knowledge management environment in a short time. INSUITE is usable across the business from day one.
- ▶ **Cybozu Garoon** is a key collaboration application. Workgroup members can check mail, schedules, intranet documents, and various news forums via the Internet by using Cybozu Garoon from any Web browser. For more information, see:

<http://cybozu.com/index.cbml>

Security

The Linux server is used as a gateway server or file server in many cases in an Internet infrastructure. Although there are not many reports of virus infection in Linux environments at present, when a transaction that passes through a Linux gateway contains a virus, and the data resides in a Linux file server, the Windows system machine using the data can suffer damage.

The xSeries server offers a low-cost, secure foundation for your investment in Linux infrastructure and business applications. By providing hardware security features, self-healing, and self-diagnostic services, as well as secure remote administration and monitoring, the xSeries server is the platform for those who need to implement a secure system infrastructure.

IBM works closely with the leading security software providers and the managed security providers to deliver tested and validated security solution configurations and ServerProven security software. IBM security services augment the xSeries features by enabling customers to assess their security risks and develop plans to address the risks. The combination of a secure platform, a full line of solutions, and complete services coverage makes the xSeries server solution an ideal secure server solution.

Server security can be provided at several levels, including hardware, operating system, middleware, and applications. Table 5-5 summarizes how xSeries delivers security at each level.

Table 5-5 Levels of server security

	Security features	Customer value
xSeries hardware	<ul style="list-style-type: none"> ▶ Secure remote monitoring ▶ Secure remote administration ▶ Configuration control ▶ Lights-out support ▶ Self-diagnostics ▶ Hardware security framework 	<ul style="list-style-type: none"> ▶ Low-cost, secure platform for any security application environment ▶ Part of a complete IBM security solution built on IBM and industry-partner offerings
Operating system	<ul style="list-style-type: none"> ▶ Linux distribution ▶ OpenSSL encryption ▶ OpenSSH authentication ▶ PAM authentication ▶ SNORT intrusion detection ▶ Apache secure Web server ▶ Integrated firewalls in many distributions ▶ Tripwire for file integrity ▶ SuSE Linux Enterprise <ul style="list-style-type: none"> – OpenSSH – OpenSL – PAM – Kerberos authentication ▶ Red Hat Advanced Server <ul style="list-style-type: none"> – Firewall – OpenSSH – OpenSSL – Kerberos authentication 	Linux is a low-cost, highly functional, and secure UNIX derivative. It brings the benefits of UNIX at a substantially lower cost. A complete suite of open-source and industry security solutions make Linux an enterprise-class environment. Linux is fully supported by IBM security services.
Middleware	<ul style="list-style-type: none"> ▶ Tivoli Risk Manager ▶ Tivoli Access Manager ▶ Tivoli Identity Manager 	A complete security management environment that can integrate with other IBM and ISV offerings and is fully supported with IBM services.
Applications	Trustix	<ul style="list-style-type: none"> ▶ Mail server with server-side antivirus ▶ File/Print

	Security features	Customer value
Security applications	<ul style="list-style-type: none"> ▶ Checkpoint ▶ Trustix ▶ Trend Micro ▶ Tripwire Security Systems, Inc. ▶ SuSE Linux firewall 	<ul style="list-style-type: none"> ▶ Firewall ▶ Secure Linux operating system ▶ Firewall with VPN ▶ Proxy server ▶ Antivirus software ▶ File-integrity security and policy-compliance applications
Services	<ul style="list-style-type: none"> ▶ IBM Security Assessment Services ▶ IBM Managed Security Services ▶ IBM Detection Services ▶ IBM Security and Privacy Services ▶ IBM Tivoli Consulting Services ▶ IBM Cross-Industry & Industry Solutions 	<p>A comprehensive approach to security from a people and IT perspective. Addresses process, education, training, vulnerability assessment, and response plans. Helps customers identify critical assets and strengthen measures to control their access and use. The customer chooses whether to be security self-sufficient or to use IBM-managed services to protect their corporate and IT assets.</p>

Here are some examples of security applications that run on Linux:

- ▶ **Symantec AntiVirus Scan Engine** provides advanced, high-performance virus scanning and repair services. It is designed specifically to protect traffic served through, or stored on, network infrastructure devices. The solution is easily integrated into server software applications and networking systems. Symantec AntiVirus Scan Engine detects and protects against viruses, worms, and trojan horses in all major file types, including mobile code and compressed file formats. For more information, see:

<http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=173>

- ▶ **Norton AntiVirus for Lotus Domino** offers the most comprehensive, automatic protection against new and existing viruses. It keeps databases free from current and future viruses by automatically scanning and repairing file attachments in Notes® mail and database documents. For more information, see:

<http://enterprisesecurity.symantec.com/products/products.cfm?ProductID=143&EID=0>

- ▶ **InterScan VirusWall** provides high-performance, comprehensive Internet gateway protection against viruses and malicious users of code. The optional eManager plug-in offers administrators additional tools to block spam, filter content, and schedule e-mail. For more information, see:

<http://www.trendmicro.com/en/products/gateway/isvw/evaluate/overview.htm>

- ▶ **InterScan WebManager** blocks unproductive Universal Resource Locators (URLs), monitors and manages Web usage, and scans Web traffic for viruses and other malicious code at the gateway. For more information, see:

<http://www.trendmicro.com/en/products/gateway/iswm/evaluate/overview.htm>
- ▶ **ServerProtect** software provides real-time antivirus scanning for Linux servers. It detects and removes viruses from files and compressed files in real time, before they reach the end user. Administrators can use a Web-based console to manage virus outbreaks, virus scanning, virus pattern file updates, and notifications. For more information, see:

<http://www.trendmicro.com/en/products/file-server/sp-linux/evaluate/overview.htm>
- ▶ **ScanMail for Lotus Notes®** detects and removes viruses hidden in databases and e-mail in real time, before they spread to the desktop. It provides automatic virus signature file updates, centralized virus activity logs, and integration with the Lotus Notes administration client. For more information, see the following site:

<http://www.trendmicro.com/en/products/email/smln/evaluate/overview.htm>
- ▶ **FireWall-1** enables enterprises to define and enforce a single, comprehensive security policy that protects all network resources. Its architecture delivers a highly scalable solution that integrates all aspects of network security. For more information, see:

<http://www.checkpoint.com/products/protect/firewall-1.html>
- ▶ **VPN-1 Pro** is the foundation of Check Point VPN-1 solutions, the most comprehensive set of products and technologies for remote access, intranet, and extranet VPNs. VPN-1 Pro software protects the privacy of information over the Internet while securing critical network resources against unauthorized access. For more information, see:

http://www.checkpoint.com/products/connect/vpn-1_pro.html

High-availability solutions

High-availability solutions are important to raise the availability of a system built on Linux. Such solutions include:

- ▶ **IBM Tivoli System Automation for Linux** supports the need for a high-availability solution that can also reduce complexity and costs. Failures, outages, and even slowdowns can result in direct business losses (the competition is just a mouse-click away) and may damage a company's credibility. The IBM autonomic computing initiative helps address these issues and more by using technology to make technology self-managing. IBM Tivoli System Automation for Linux provides self-healing capabilities for business applications by detecting failing IT components and then repairing or working around those failures.

For more information, see:

<ftp://ftp.software.ibm.com/software/tivoli/datasheets/ds-sys-auto-linux.pdf>

- ▶ **LifeKeeper for Linux** is a software application that ensures the continuous availability of applications by maintaining system uptime. LifeKeeper maintains the high availability of clustered Linux systems by monitoring system and application health.

For more information, see:

<http://www.steeleye.com/products/linux/>

- ▶ **VERITAS Cluster Server** eliminates both planned and unplanned downtime, facilitates server consolidation, and effectively manages a wide range of applications in heterogeneous environments. VERITAS Cluster Server supports up to 16 node clusters in SAN and traditional client/server environments.

For more information, see:

<http://www.veritas.com/products/category/ProductDetail.jhtml?productId=clusterserver>

Backup solutions

Although Linux controls backup and operating tape devices, additional backup solutions exist that make operations easy. They can also schedule backup and recovery events. Additional backup solutions include:

- ▶ **IBM Tivoli Storage Manager:** The latest version of IBM Tivoli Storage Manager server software is now available on Intel-based Linux systems. By extending IBM Tivoli Storage Manager's support for Linux from file systems and applications to include Tivoli Storage Manager servers, organizations using Linux environments can benefit from the same award-winning data protection as organizations using Tivoli Storage Manager servers on other platforms. IBM Storage Manager underscores IBM's commitment to provide robust, reliable software solutions for the mid-market through the large enterprise.

For more information, see:

<http://www.ibm.com/software/tivoli/solutions/linux/>

- ▶ **BrightStor ARCserve Backup for Linux** brings data protection technology to organizations deploying Linux. Beyond simple data protection, BrightStor ARCserve Backup uses multiple data verification methods to enable maximum data integrity and recovery capabilities.

For more information, see:

<http://www.ca.com/Solutions/Product.asp?ID=3370>

- ▶ **NetVault Enterprise Edition** is a comprehensive backup and restore solution for heterogeneous UNIX, Windows NT/2000, Linux, and Netware enterprise environments. It can rapidly add and configure new servers, devices, and clients, and can control them from a central location. As an organization grows, a modular architecture allows NetVault Enterprise Edition to combine with other BakBone plug-in modules, such as disaster recovery.

For more information, see:

http://www.bakbone.com/products/backup_and_restore/netvault_ee/

5.6.3 Customer scenario: Journyx

Journyx, an IBM Business Partner and a leader in Web-based time- and expense-tracking solutions, provides time-, project-, and expense-tracking software to more than 10,000 organizations. Founded in 1996, the Austin, Texas-based company, employing approximately 50 people, already owns a 70% market share in its industry.

Journyx's flagship product, Timesheet, that is integrated with IBM WebSphere Application Server and IBM DB2 Universal Database, enables companies to efficiently track and manage project time and associated expenses over an intranet or the Internet.

A long-time DB2 user, Journyx based Timesheet on DB2 because of the data management system's consistently solid performance. The latest version of Timesheet leverages DB2 for Linux on the back end to maximize data access and management options.

In addition to Web-based applications, Journyx offers application hosting services to companies that don't want to manage the expense-reporting process in house. Although Timesheet supports a variety of databases, Journyx uses DB2 for Linux for all of its application hosting sites because of the robust data-management system of DB2 for Linux, which offers the flexibility, scalability, and pricing that support our business model.

Journyx hosts its Timesheet application on nearly a dozen xSeries servers running Linux at its collocation facility. The company viewed Linux as a secure, robust, and cost-effective operating environment that would help it deliver reliable service. The xSeries servers also provide the high performance and reliability that Journyx needs. Journyx also uses an XML integration code that allows customers to customize the integration of multiple systems, including most well-known budget management, payroll automation, and accounting systems.

PricewaterhouseCoopers (PwC), the world's largest professional services organization, sold its PwC Consulting service to IBM and is a subscriber to the hosted Timesheet service. A Journyx customer for over two years, PwC relies on

Journyx to accurately manage the timecard submission and approval process for a 900-employee department.

The process begins when an employee logs on to a secure Web page, served up by WebSphere Application Server, on the PwC intranet. From the Web page, the employee completes the timecard and clicks a Submit button to route it to the backend server. From there, the Journyx application collects the data, processes it, and routes it back to PwC, using TCP/IP, with a copy stored in the DB2 database. Using the Journyx application, PwC managers access, review, and approve the timecards online. The data is then sent for payroll purposes, using an internal PwC application and a PeopleSoft HR module.

Timesheet has given PwC a reliable data access and management solution that saves the company time and money. PwC has gained benefits from Journyx Timesheet, which is generating an average saving of \$20,000 weekly or more than \$1 million annually.

By maximizing its operational efficiencies with Linux and DB2 and by managing the system implementation process, Journyx passes substantial cost savings on to its customers. According to Journyx, most customers realize 100% payback within three to six months after installing Timesheet. They also eliminate the three to ten percent error rate associated with manual timecard processing routines.

For more information about this success story, see the following site:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5F7T8L?OpenDocument&Site=db2software>

5.6.4 Customer scenario: Flamenco Networks

Flamenco Networks is a Web Services networking company that enables business partners, even those with different enterprise resource planning systems, to rapidly and securely exchange business documents over the Web. Customers pay a service fee to access Flamenco's network management site, where they can obtain "lightweight" code and upgrades to network management and configuration tools that allow them to leverage existing Internet protocols and conduct e-business worldwide.

To help ensure fast, 24-x-7 access, Flamenco needed to build on a highly redundant, high-performance server platform. Flamenco also required a vendor that offered comprehensive support and services for Red Hat Linux. Flamenco opted for IGS and its commitment to open-standards-based technology.

Also critical to Flamenco's success is SteelEye's LifeKeeper solution, known throughout the industry for ensuring high availability on clustered Linux server

systems. By monitoring system and application health, maintaining client connectivity, and providing uninterrupted data access regardless of where the clients reside, LifeKeeper provides ease and speed of deployment. It also enables swift, out-of-the-box integration without application modification.

In September 2001, after 18 months in development, Flamenco Networks went live with its network management site. Supporting it are two xSeries 342 machines in parallel, two xSeries 340 database servers, and an IBM FAST HA Storage Server. The IBM machines give Flamenco Networks the flexible, highly available, and reliable solution.

IBM system engineers helped Flamenco match which versions of Linux were certified with its relational database software. They also helped install the Red Hat drivers for the storage server-critical points during preparation for production.

With the scalable architecture of xSeries servers, Flamenco can easily add Web servers or database servers as necessary. And the IBM configuration allows Flamenco to push new software onto the system without the risk of outage, helping the company establish a solid reputation. With IBM hardware and SteelEye's clustering software, Flamenco can deliver quality services to the customers at a reduced cost.

For more information about this success story, see:

<http://www.ibm.com/software/solutions/webservices/partners/flamenco.html>

5.7 Infrastructure solutions

Infrastructure solutions continues to offer the largest saving opportunity as customers move to Linux. The typical solutions in this area are Web servers, mail servers, file/print servers, and firewall servers. Most of the Linux distributions for the xSeries server contain such open-source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail, Squid, and PostgreSQL.

Samba

Samba is an open source/free software suite. It provides seamless file and print services to Server Message Block (SMB) and Common Internet File System (CIFS) clients. It allows end users to access and use files, printers, and other commonly shared resources on a company's intranet or on the Internet.

Samba is often referred to as a Network File System (NFS) and can be installed on a variety of operating system platforms, including Linux, most common UNIX platforms, OpenVMS, and OS/2. Samba supports the following UNIX and

UNIX-compatible machines using Windows NT or 2000-compatible file/print server software. It can be used freely under GNU General Public License (GPL).

Samba has the following features:

- ▶ **Cost cutting:** There is a difference in cost between introducing the Windows 2000 Advanced server and the Red Hat Linux 8.0 professional for Intel-Architecture server.

- ▶ **Configuration and management:** Samba's configurations file is written in the only text file (smb.conf). In the Linux-plus-Samba environment, Samba can be set up and managed remotely using Telnet.

There is an easy configuration tool for Samba named SWAT. This tool can set up Samba from a Web browser.

- ▶ **Performance:** Performance data exists that compares Samba for the Intel-Architecture server to Windows. For more information, see:

<http://www.pcmag.com/article/0,2997,s%253D1474%2526a%253D16554,00.asp>

Look for Samba on the Web at:

<http://www.samba.org/>

Apache

Apache is the most famous Web server (HTTP daemon) in the world. According to research by netcraft, Apache has more than 60% of the market share for Web servers. You can learn more about this research on the Web at:

<http://www.netcraft.com/survey>

Apache has two versions: Apache 1.3 and Apache 2.0. Table 5-6 shows what changed from Apache 1.3 to Apache 2.0.

Table 5-6 Changes from Apache 1.3 to Apache 2.0

Change	Details
Adopting Multi Processing Module (MPM)	MPM can communalize other modules on all platforms
Adopting Apache Portable Runtime (ARP)	Development by ARP and C language enables it to perform the same program also on other platforms
Correspondence to a multithred	To request processing by the thread generation from the process generation
Support of a multi-protocol	Apache 2.0 can easily support HTTP, HTTPS, and a new protocol

Change	Details
Support of IPv6	At Apache 2.0, Web server can be made to correspond to IPv6 only by base installation
Adopting new build system	Installation of autoconf and libtool is necessary to install Apache 2.0

To learn more about Apache, see:

<http://www.apache.org/>

Sendmail

Sendmail is the most common Mail Transfer Agent (MTA) for UNIX. Sendmail is configured by modifying `sendmail.cf` (or `sendmail.m4`). Sendmail can be set up by detailed configuration, but it is complex for system administrators to modify `sendmail.cf`. The wrong configuration can cause a security hole.

You can find Sendmail on the Web at:

<http://www.sendmail.org/>

Qmail

Qmail is a secure, reliable, efficient simple message transfer agent. It is designed for typical Internet-connected UNIX hosts. As of October 2001, Qmail is the second most common Simple Mail Transfer Protocol (SMTP) server on the Internet and has by far the fastest growth of any SMTP server. The features of Qmail include:

- ▶ **Secure:** Security isn't a goal, but an absolute requirement. Mail delivery is critical for users. It cannot be turned off, so it must be completely secure. Qmail divides its functions into components. The executive user of each function is defined separately so that, if one component is damaged, other components are not influenced.
- ▶ **Reliable:** Qmail's straight-paper-path philosophy guarantees that, after a message is accepted into the system, it is never lost. Qmail also optionally supports `maildir`, a new, super-reliable user mailbox format. `Maildir`, unlike `mbox` files, is not corrupted if the system crashes during delivery. Even better, a user can safely read their mail over the NFS and any number of NFS clients can deliver mail to them at the same time.
- ▶ **Efficient:** On a Pentium under BSD/OS, Qmail can easily sustain 200,000 local messages per day. These are separate messages injected and delivered to mailboxes in a real test. Although remote deliveries are inherently limited by the slowness of DNS and SMTP, Qmail overlaps 20 simultaneous deliveries by default.

You can find Qmail on the Web at:

<http://www.qmail.org/top.html>

Postfix

Postfix software was developed as Sendmail-compatible MTA. Postfix has many features such as higher compatibility with Sendmail, security, easy maintenance, quick processing, and so on. Recently Postfix was adopted as default MTA by some Linux distributions. Postfix has the following merits:

- ▶ **Sendmail compatibility:** Postfix is similar in function and use as Sendmail. This makes education for new users and teaching them to set up new software for Postfix unnecessary. Of course, a minimum knowledge of MTA is essential. Shifting Postfix from Sendmail can be carried out without stopping its services.
- ▶ **Easy configurations:** It is too difficult for an average user to modify `sendmail.cf` manually. Usually a configuration tool, such as CF, is used. But the tools do not always follow the latest MTA, and there have been some incidents because of the bugs in the tools. The configuration file of Postfix can be modified easily with normal MTA skills.
- ▶ **Stability:** Postfix can run with parallel processing, so it is stable for business use.

You can learn more about Postfix on the Web at:

<http://www.postfix.org/>

Netfilter

Linux uses iptable or ipchains as a Netfilter of open source software. Generally, the Linux distribution of kernel-2.2x includes ipchains and that of kernel-2.4x include iptables. Netfilter analyzes the inside of TCP and IP header and filters the packet by referring to a filtering rule. When it filters, it checks:

- ▶ Source IP and destination IP
- ▶ Source port and destination port
- ▶ Protocol
- ▶ Interfaces
- ▶ ICMP type and code
- ▶ Checking SYN packet
- ▶ Checking fragment packet

In addition, iptable analyzes the following points:

- ▶ Stateful inspection
- ▶ TCP flags
- ▶ MAC address

- ▶ Limit of matching packet
- ▶ Owner of packet producer

To learn more about Netfilter, see:

<http://www.netfilter.org/>

Squid

Squid is the most common proxy server in Linux. Instead of a client, a Squid server accesses the server on the Internet. The Squid server relays the data and passes it to a client. Some advantages of using Squid are:

- ▶ You no longer need a client to access the Internet directly, and you can control access to the Internet.
- ▶ You can decrease network traffic by having Squid cache data.
- ▶ Since all traffic passes through Squid, you can match and investigate data.

You can find out more about Squid on the Web at:

<http://www.squid-cache.org/>

5.7.1 Customer scenario: weather.com

The Weather Channel uses weather.com to service Internet surfers worldwide. As a top-10 Internet Web site during weather crises, this site receives up to 40 million hits per day. Landmark Media Corporation owns weather.com, and it hopes to shift the server focus to an open-source system environment.

The customer found itself paying too much for operating system licensing on outdated equipment whose provider did not offer any services. The customer wanted to save money while, at the same time, updating its products to a higher level.

To replace its existing system, weather.com chose to implement an IBM solution for its robust, scalable products, as well as its superior ATS support, responsiveness, and sales support. The customer chose to replace its Sun E4500 database servers running Solaris with six xSeries 350 machines running SuSE V7.1 Linux software. These xSeries servers hold all of weather.com's data.

To replace its Sun 420R image and content servers, weather.com chose 100 xSeries 330 servers running Tomcat on Linux. This system routes and load balances the site's Internet hits and updates the static information on the Web site. The customer also chose four additional xSeries 340 servers to function as primary domain controller (PDC) and firewall servers.

Since the IBM implementation, weather.com has experienced a 70% cost savings, as well as lower incremental growth costs. The IBM solution also provided the customer with robust and scalable products that helped it meet its business needs. In addition, the customer is happy with its Linux purchase for its low TCO and its leadership in the open-source arena.

For more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5B3NDN?OpenDocument&Site=default>

5.8 xSeries architected solutions

IBM provides a comprehensive solution for modular computing through the blade architecture. The blade architecture is built from blade servers and BladeCenters.

5.8.1 Blade server

The rapid spread of the Internet over several years has pushed the server system that supports it. The scale of e-business expanded quickly. The typical system complex is the high-density mounting of servers, built from the 1U rack-mounted server. In 2001, to meet the demand for super-high-density mounting, a new server appeared, the blade server.

The term “blade” means a flat section, like a knife blade. The blade server is a long, slender circuit board with condensed server functions. Equipped with 6 to 20 blades in a case called the “enclosure” that is 2U-7U size, it achieves a density greater than the 1U server.

BladeCenter

In October 2002, IBM announced the Intel Xeon processor two-way SMP called the blade server *BladeCenter HS20* and delivered the *BladeCenter* enclosure, which was 7U size. BladeCenter and BladeCenter HS20 are designed to save space in a high-density rack environment. BladeCenter can store 14 BladeCenter HS20 blades with a maximum of six BladeCenters in the 42U rack. BladeCenter HS20 supports Red Hat Linux and SuSE Linux.



Figure 5-4 BladeCenter

Characteristics of the BladeCenter marketplace include:

- ▶ **Server consolidation:** BladeCenter is a great solution for a centralization-type server consolidation.
- ▶ **HPC/grid computing:** BladeCenter is extremely good for HPC environments with the highly efficient blade Intel Xeon processors running the following workloads:
 - Linux cluster
 - Genome
 - Engineering Design Automation (EDA)
 - Grid
- ▶ **Application server:** Since BladeCenter has the same performance and function as a general-purpose server (for example, a SAN environment), the BladeCenter becomes the optimum application server and can provide considerable space savings in the following environments:
 - File server
 - Notes server
 - WebSphere Application Server

For more information, see the following Web sites:

http://www.pc.ibm.com/us/eserver/xseries/bladecenter_family.html

<http://www.ibm.com/servers/eserver/blades/>

5.8.2 Customer scenario: Turning Stone Casino

Operated by the Oneida Indian Nation, Turning Stone Casino Resort opened its doors in Verona, New York in 1993. One year after it opened, Turning Stone launched a unique, totally cashless gaming environment that electronically tracks and records results. Patrons use cards, similar to ATM cards, to keep tabs on their account and cash out winnings.

In 1999, Turning Stone realized that using open-source technology could help improve its offerings. The company decided to shift its systems, including the

network, backend systems, and player terminals to an open IT environment. Through this redevelopment effort, Turning Stone hoped to expand its gaming selection, offer patrons a tailored experience, and improve information technology efficiency to help reduce operating costs.

Turning Stone first replaced its aging backend gaming systems with an IBM eServer solution, including pSeries H70 and H50 models and iSeries 720 machines. The H70 servers use high-availability cluster multiprocessing technology for reliability and improved manageability, and use H50 servers for test clustering. The iSeries servers provide a stable platform for Turning Stone's property management and financial applications.

The next step was to find a suitable platform to host the casino's custom protocol translation application, which converts player terminal information coded in proprietary protocol to a universal gaming protocol that the backend systems can process. The casino is a 24-x-7 operation, making any downtime costly. Turning Stone's goal of increasing network uptime, reliability, and high availability were key issues for this solution. After careful consideration, the casino purchased 19 xSeries 135 servers.

Taking advantage of the cable-chaining design of the xSeries 135 machines, the customer now avoids purchasing expensive switches. The advanced Light Path Diagnostics feature simplifies troubleshooting by guiding the technician to faulty components and allowing the technician to replace components without ever powering down the server. The customer also plans to connect the xSeries serially to a collection of service processors, which allows Turning Stone's IT personnel to manage the entire system remotely, saving management costs. By migrating to an open-source environment, Turning Stone has broadened its selection of player terminals from approximately seven games to more than 25 games.

For more information about this success story, see:

http://www.ibm.com/servers/eserver/pseries/solutions/success/turning_stone.html

5.9 xSeries services solutions

Specific types of service vary by geographical area, but it generally includes implementing Linux solutions for new customers and providing support for existing customers. The following sections pertain to service in the United States. For more information, see:

<http://www.ibm.com/services/e-business/linux.html>

5.9.1 Implementing Linux

The implementation of Linux service ranges from installation support to the construction of clusters. Installing and configuring Linux are important business opportunities, as are setting up Apache and Samba. However, since many companies offer those basic services, the price of such services (and the profit from providing them) is decreasing. High-added-value service, such as cluster construction and server consolidation, will be an even more important business opportunity for IBM in the future.

The following sections explain the typical types of service available for Linux on xSeries.

Linux Solution for e-business

The Linux Solution for e-business provides rapid deployment of a customer's Linux-based e-business environment. This solution is tailored to fit the specific needs of the customer by allowing the customer to select from a wide variety of middleware and applications that can be installed and configured on the base of an IBM-supported Linux operating system.

The Linux Solution for e-business service capability consists of two phases. Each phase consists of a set of selectable elements.

- ▶ **Phase 1, Technology enablement:** This phase involves installing, customizing, and demonstrating both base and selected additional technologies in an environment isolated from the customer organization's existing production systems.
- ▶ **Phase 2, Technology enablement:** This phase involves the following elements:
 - Assisting the customer organization in connecting the newly installed e-business server to its existing systems
 - Establishing the Linux Solution for e-business Server connectivity to a public or private TCP/IP network
 - Connecting one of the customer organization's existing database or transaction manager subsystems to the corresponding subsystem on the Linux Solution for e-business server or implementing collaboration and Internet servers.

For more information, see:

http://www.ibm.com/services/e-business/linux_3.html

High-performance clusters services

For customers who design, implement, and support Linux clusters for Internet and enterprise-wide, mission-critical applications, IBM is there to help.

To learn about these services, see:

http://www.ibm.com/services/e-business/linux_6.html

Server consolidation service

The driving force behind server consolidation is the need to reduce TCO. Many customers have a mishmash of hardware and operating systems that are expensive to operate. Also, more modern applications, such as enterprise-resource-planning software, further complicate matters. Using an architecture-independent approach, IBM helps customers reduce the TCO.

For more information, see:

http://www.ibm.com/services/e-business/linux_7.html

5.9.2 Support

Depending on customer need, IBM offers Internet and voice support, ranging from answering usage questions to identifying problems. This support is available 24-hours-a-day, 7-days-a-week. IBM Global Services also provides consulting, planning, and implementation services for Linux. IBM consultants can help you evaluate whether Linux is appropriate for your particular environment.

For more information, see:

http://www.ibm.com/services/e-business/linux_8.html



pSeries and Linux

Linux on pSeries is a key part of the IBM Linux strategy. During the past few years, IBM has been developing the ability to integrate Linux capabilities into the pSeries server. pSeries servers can run Linux alone or can run both IBM AIX and Linux in logical partitions (LPAR) of the same server.

This chapter provides information about Linux enablement for the pSeries. It discusses the ways that you can integrate Linux applications into commercial pSeries environments.

6.1 How Linux runs on the pSeries server

The pSeries server is one of IBM's key solutions to fit a variety of customer requirements. pSeries servers offer leading edge performance and can help you confidently manage your business's growth and risks.

To meet your requirements, pSeries products offer a full range of high-performance servers complete with highly functional, state-of-the-art software. IBM will provide more services for Linux for pSeries in response to growing interest from customers to use Linux as their operating system in this platform.

There are some key concepts worth reviewing before you see how Linux and AIX systems and applications run on the pSeries. For more information, see 1.3, "Linux: The operating system" on page 9.

For more information, see IBM's Linux for pSeries Web site.

<http://www.ibm.com/servers/eserver/pseries/linux/index.html>

6.1.1 Linux-capable pSeries model structure

The Linux-ready pSeries server has models that range from entry deskside models through to high-end enterprise models. Rack mounted models offer physical consolidation possibilities. Figure 6-1 shows all-Linux capable models.

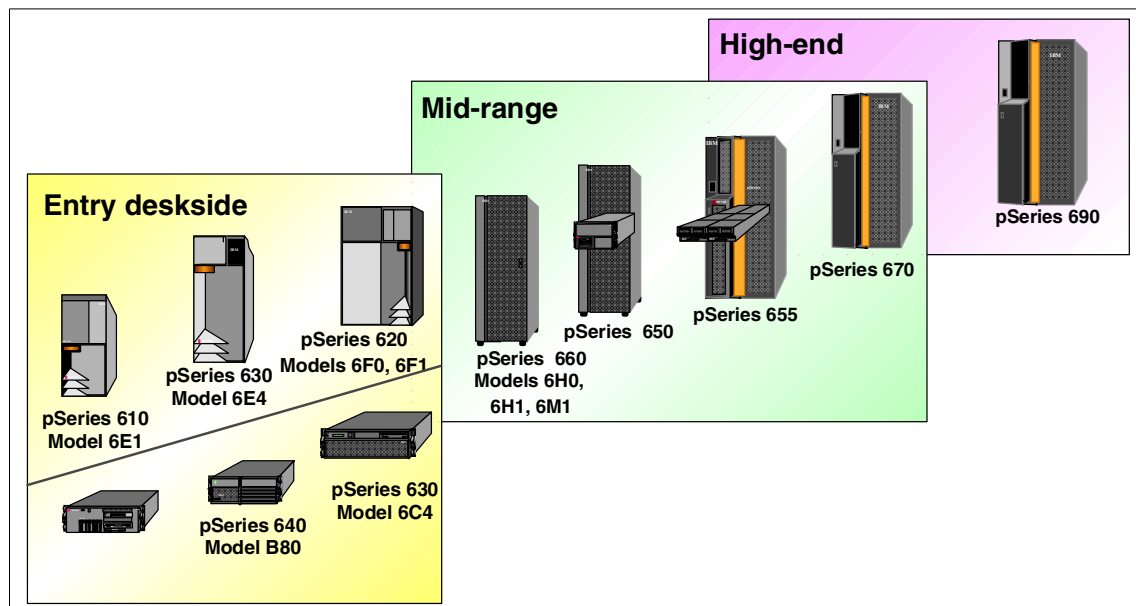


Figure 6-1 pSeries models

For white papers and technical details on each model, visit:

<http://www.ibm.com/eServer/pSeries>

The following pSeries white papers are a must-read for anyone who wants further information about Linux on pSeries:

- ▶ *IBM @server Linux for pSeries: An overview for Customers*
http://www.ibm.com/servers/eserver/pseries/linux/whitepapers/linux_pseries.html
- ▶ *Linux for IBM @server pSeries: Facts and Features*
http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf
- ▶ *LPAR for Decision Makers*
http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/lpar_decision.pdf
- ▶ A collection of pSeries white papers:
http://www.ibm.com/servers/eserver/pseries/library/wp_systems.html

For more information about pSeries logical partitioning, refer to the IBM ITSO Redbooks site on the Web at:

<http://www.redbooks.ibm.com>

6.1.2 pSeries operating system scenarios

The pSeries offers outstanding server performance and a great server consolidation platform. These servers provide the ability to run AIX or Linux as single operating systems or multiple LPARs running any combination of AIX and Linux. Figure 6-2 shows an overview of the internals of the AIX and Linux operating systems, and how applications can run on pSeries in three different scenarios.

Scenario I shows the traditional AIX-only-based pSeries environment. All AIX versions prior to Version 4.3.3 or 5L only run applications compiled for AIX. That means Linux applications cannot be run, unless they are explicitly ported to AIX.

Scenario II shows how AIX 4.3.3 and AIX 5L™ incorporate Linux in the pSeries by means of the *AIX Toolbox for Linux Applications*. The AIX Toolbox for Linux Applications provides the capability to recompile open source applications to run as AIX applications. In most cases, it does this without porting and using the same open source build environment as Linux.

Scenario III shows the fully Linux-enabled scheme. It has a Linux for pSeries kernel. Customers can integrate a broad spectrum of Linux applications into their

pSeries servers. The customer can generate Linux code for any Linux application available for the pSeries.

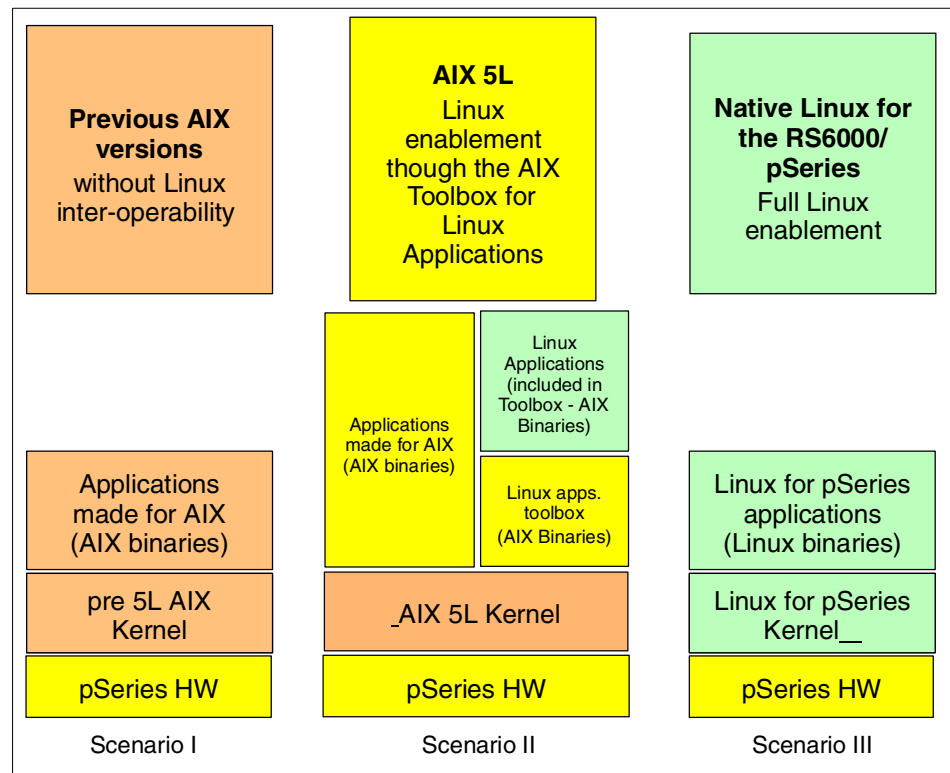


Figure 6-2 How UNIX operating systems and applications run on the pSeries

6.1.3 AIX 5L and Linux

You can run many Linux applications on AIX, a mission-critical, highly scalable operating system through Linux application interoperability to AIX 5L.

IBM will implement an AIX Affinity with Linux in two phases. The first phase is the release of the AIX Toolbox for Linux Applications, which is on the Web at:

<http://www.ibm.com/AIX>

The second phase is IBM's planned inclusion of additional Linux compatible application programming interfaces (APIs) and header files in AIX 5L Version 5.1.

AIX Affinity with Linux uses an API approach to provide Linux application interoperability. This approach is *not* an environment or an additional layer or

wrapper on or in which to run Linux applications. It is the integration of Linux compatible APIs and header files into AIX 5L.

Therefore, recompiled Linux applications are treated as AIX applications and have access to all the reliability, scalability and availability of AIX. The result is tighter integration of the application to the operating system than can be achieved with an Application Binary Interface (ABI) approach.

AIX was developed using UNIX industry standards. As such, there is a high degree of compatibility at the API level between AIX and Linux. This degree of similarity allows many Linux applications to be recompiled and run on AIX 4.3.3 and AIX 5L Version 5.0 today using the AIX Toolbox for Linux Applications. AIX 5L Version 5.1 plans to add even more interoperability between AIX and Linux by including APIs that are presently not similar between AIX and Linux. This will result in an even higher degree of Linux application compatibility. Therefore, much of the functionality of AIX 5L Version 5.1 already is available in AIX 4.3.3 and AIX 5L Version 5.0, with full implementation of IBM's Linux affinity plans available in AIX 5L Version 5.1.

You can learn more about AIX Affinity in the white paper at:

http://www.ibm.com/servers/aix/products/aixos/linux/affinity_linux.pdf

The uses of AIX Toolbox for Linux Applications include:

- ▶ Build and package Linux applications for use on AIX
- ▶ Run Gnome and KDE desktops
- ▶ Run other popular software commonly found in Linux distributions
- ▶ Manage open source software using the popular RPM package management system
- ▶ Develop new applications for AIX using GNU and Linux application development tools

Figure 6-3 shows the AIX toolbox structure running over the AIX kernel.

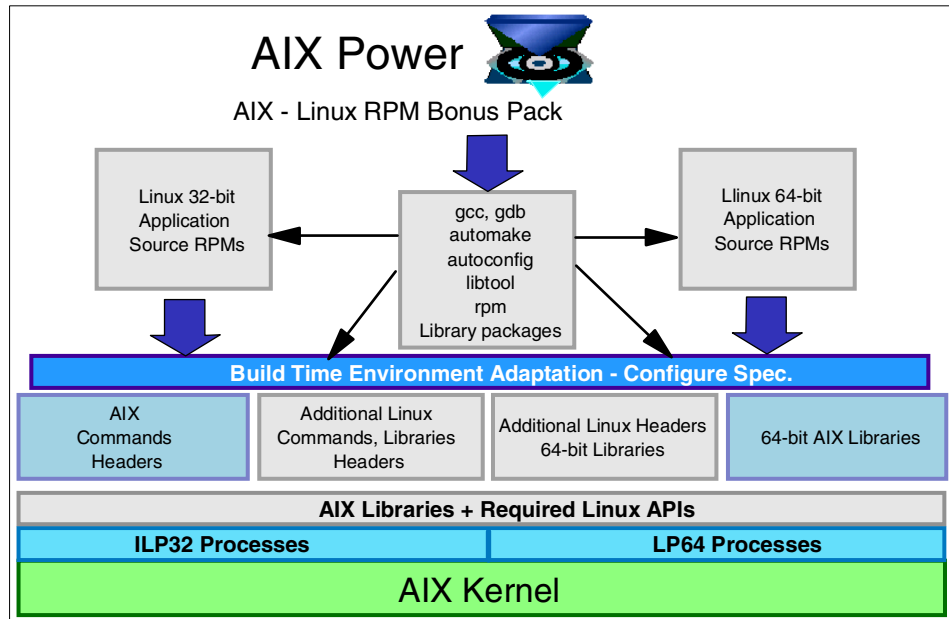


Figure 6-3 The AIX Toolbox for Linux Applications

AIX Toolbox for Linux Applications contains a wide variety of software, including:

- ▶ **Application development:** gcc, g++, gdb, rpm, cvs, automake, autoconf, libtool, bison, flex, gettext
- ▶ **Desktop environments:** Gnome and KDE-2
- ▶ **GNU base utilities:** gawk, m4, indent, sed, tar, diffutils, fileutils, findutils, textutils, grep, sh-utils
- ▶ **Programming languages:** guile, python, tcl/tk, rep-gtk
- ▶ **System utilities:** emacs, vim, bzip2, gzip, git, elm, ncftp, rsync, wget, lsof, less, samba, zip, unzip, zoo
- ▶ **Graphics applications:** ImageMagick, transfig, xfig, xpdf, ghostscript, gv, mpage, Gimp
- ▶ **Libraries:** ncurses, readline, libtiff, libpng, libjpeg, slang, fnlib, db, gtk+, qt
- ▶ **System shells:** bash2, tcsh, zsh
- ▶ **Window managers:** Enlightenment, sawfish

For the latest information about AIX-5L, visit the following Web site:

<http://www.ibm.com/servers/aix/os/52desc.html>

AIX 5L also introduces many new features. You can learn about them in the IBM Redbook *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765.

Exploring the details of running Linux applications on AIX as well as other built in functionality is beyond the scope of this book. We recommend that you consult the redbook *Linux Applications on pSeries*, SG24-6033.

You can find more information about the AIX Toolbox for Linux Applications on the Web at:

<http://www.ibm.com/servers/aix/products/aixos/linux/index.html>

AIX 5L is intended to meet the need for providing interoperability between Linux applications (as well as other Linux machines) and the pSeries that run AIX. Bringing Linux applications to the pSeries with AIX 5L makes interoperability a reality. In addition, you can retain the considerable benefits and functionality derived from AIX.

6.1.4 Options to run Linux on the pSeries server

Customers have many choices with the pSeries servers. They can run Linux applications under AIX or run Linux as an operating system per server (similar to unpartitioned Intel servers). Or Linux can run as a partition among other Linux or AIX partitions. The variations are endless. This gives you a great deal of flexibility to manage and possibly optimize your hardware resources.

This section covers both non-partitioned and partitioned Linux servers.

Linux on the pSeries

We assume that you have selected a Linux application that does not suit the AIX-5L approach. Plus you want to run the application on a separate server. This can be indicative of the real environment where a dedicated server is installed to run specific a application.

Linux applications can run in pSeries LPAR hardware by simply installing the Linux operating system.

While the AIX Toolbox for Linux Applications offers an excellent way to make Linux applications available on pSeries servers, there may be cases when it is preferred to run them under a Linux operating system instead of AIX.

The process is simple:

1. You select the Linux application that is supported by one of the Linux Distribution Partners (LDPs). Linux applications are frequently available as a demo or trial download for testing.

2. You size and select any pSeries models that supports Linux and meets your needs from one of the Linux Express models, the pSeries 630 or 650 models.
3. Then you order the hardware, Linux distribution, Linux application, and necessary support contracts.

These Linux Express models are not shipped with the Linux operating system pre-loaded. Also, AIX is not required for this solution.

This approach also applies to any other Linux applications that a pSeries customer may run that are not available with the AIX Toolbox for Linux Applications.

Linux in LPAR on a supported pSeries server

The LPAR capabilities of the pSeries server make it possible to run one or more instances of Linux along with AIX. This provides a convenient way to begin developing and deploying Linux operating system-ready applications as desired, while retaining the enterprise-ready capabilities of AIX for mission-critical or highly-scalable workloads. LPAR also allows large pSeries servers to be partitioned to run Linux only workloads.

First lets look at some terminology.

- ▶ **LPAR:** This is the shortened version of the familiar term *logical partitions*. It is widely used in zSeries, iSeries, and pSeries servers to indicate that a server can run multiple separate operating system instances.
- ▶ **Static LPAR:** This term applies to a pSeries server or partition. It indicates that the server does not support the rearrangement of resources without powering down. The same restriction applied to early releases of iSeries LPAR. All pSeries Linux partitions are static.
- ▶ **Dynamic LPAR (DLPAR):** This term indicates that a partition accepts certain rearrangements of resources. Partitions running AIX and AIX-5L applications can have dynamic movement of processors and memory between partitions.
- ▶ **Hypervisor:** This is the underlying control mechanism that sits between the physical hardware and the operating systems. The hypervisor owns all system resources and creates partitions by allocating these resources and sharing them.

AIX is not required to partition the machine, and it is not required to install and run Linux. Nevertheless, AIX provides much greater features for hardware analysis, intermittent fault detection, and error logging.

For more information, visit the pSeries Web site at:

<http://www-1.ibm.com/servers/eserver/pseries/linux/>

AIX and Linux can run in separate partitions at the same time on an LPAR-enabled pSeries server in any combination (that is, zero or more Linux partitions along with zero or more AIX partitions). This allows you to consolidate workloads from several separate servers onto a single system.

Since partitioning in an LPAR-supported pSeries is controlled by the firmware-based hypervisor and the Hardware Management Console (HMC), AIX is not required to run Linux.

Figure 6-4 shows a typical logically partitioned pSeries server. It shows the major components the HMC linked to both the partitions and the hypervisor. This includes the internal components that the HMC can control, processors, memory, and I/O adapters.

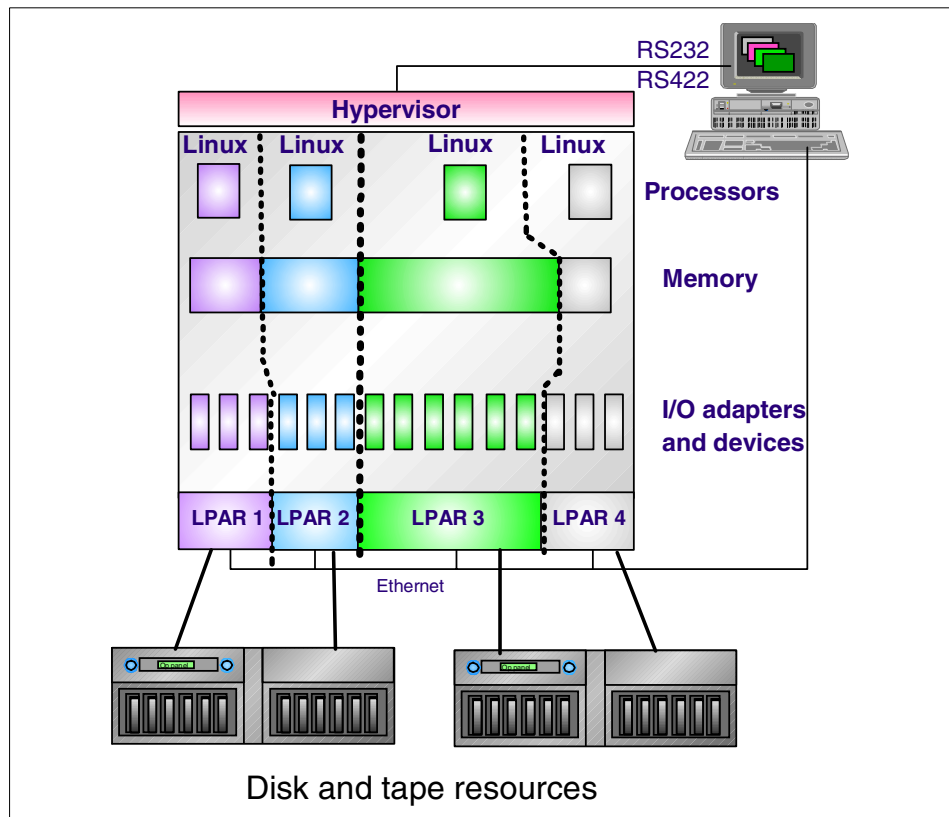


Figure 6-4 Linux for the pSeries server on the LPAR

The logical partitioning functions allow the rearrangement of resources to and from a partition, as well as to move resources between two partitions, without

operating system reboot. The new Dynamic LPAR function in pSeries can dynamically move resources for AIX partitions. It is currently not supported by Linux. However, Linux partitions and dynamic LPAR can run on the same server.

For more information about LPAR-supported pSeries servers, see:

<http://www.ibm.com/servers/eserver/pseries/lpar/resources.html>

In the example in Figure 6-4, any Linux partitions can be an AIX partition.

6.1.5 Linux and its relationship with AIX

Linux can relate with AIX in two ways. There may be system environments in which both types of relationships are present at the same time.

- ▶ The ability to run Linux applications as a value inside AIX 5L
- ▶ The ability to have Linux-based pSeries frontend systems interacting with mission-critical AIX-based pSeries servers

Linux has already been helping AIX, as AIX 5L has gained more momentum due to the AIX Toolbox for Linux Applications. By allowing Linux applications to be easily compiled for AIX with the AIX Toolbox for Linux Applications package, IBM is expanding the available application portfolio. In addition, AIX can now take advantage of the growing Linux skill base.

Furthermore, Linux and AIX are complementary operating systems. AIX is the strategic, proven, mission-critical operating system for the pSeries. Linux is a highly portable operating system that supports all IBM platforms. IBM expects to see many installations running Linux (on xSeries or pSeries hardware) as the front end to mission-critical AIX systems running DB2 and other enterprise applications.

As you can see, AIX is the strategic operating system for the pSeries server. It contains such enterprise features as scalability and Remote Access Security (RAS) that will take years to materialize on Linux. Many AIX capabilities and features from other IBM operating systems are being added to Linux. Over time, Linux will mature into a complete enterprise operating environment.

6.2 Linux for pSeries hardware enablement

Both 32-bit and 64-bit versions of Linux for pSeries are being provided to optimize customer choices and exploit pSeries hardware capabilities. The RS/6000 B50 and 150 are 32-bit PowerPC systems supported by the Linux for pSeries 32-bit kernel and application environment.

The 64-bit POWER3 and POWER4 systems have a 32-bit or 64-bit kernel and support a 32-bit or 64-bit application environment (depending on the model and Linux distribution).

I/O device and adapter support and availability

Many adapters and devices can be attached to pSeries systems. While some of the devices (for example, PCI adapters) have Linux drivers for Intel and AIX drivers, they cannot be used directly in Linux for pSeries.

However, Linux for pSeries has base device support. Be sure to review the latest version of the *Linux for pSeries Facts and Features* white paper on the Web at:

http://www.ibm.com/servers/eserver/pseries/hardware/linux_facts.pdf

Important: Some of the hardware devices listed on any LDP site may not be fully enabled or certified. We encourage you to look at each LDP's Linux hardware support database and compare it with the IBM Support matrices.

It is possible that the kernel will not boot with an unsupported graphics card in the machine. Unplug the card and use a serial console to install SuSE Linux. For more information, see the SuSE Linux hardware support database at:

http://support.suse.de/sdb/en/html/01h_ppc_hardware_ibm.html

We encourage you to periodically look for updated information about the hardware platforms in which Linux for PPC is intended to be enabled and has limited or full support. See:

<http://www.ibm.com/servers/eserver/pseries/linux>

6.2.1 Linux for pSeries scalability

Linux kernel 2.4 has proven to scale well to up to eight processors in a symmetric multiprocessing (SMP) system depending on the workload. This makes it a good match for previous RS/6000 systems such as the 170, 270, pSeries systems such as the 610, 620, 630, 640, 650, 655, 660, and a 1- to 8-way LPARs on the 670 and 690. Linux for pSeries can also scale up to 64 GB in memory and 2 TB file system size.

For details on Linux for pSeries scalability and enablement, see:

<http://www.ibm.com/services/learning/conf/us/pseries/preconf.html>

Also refer to the following Web site for a presentation about Linux for pSeries:

[http://www.ibm.com/partnerworld/pwhome.nsf/vAssetsLookup/LunchLearnLinuxforpSeries2002_08_27.pdf/\\$File/LunchLearnLinuxforpSeries2002_08_27.pdf](http://www.ibm.com/partnerworld/pwhome.nsf/vAssetsLookup/LunchLearnLinuxforpSeries2002_08_27.pdf/$File/LunchLearnLinuxforpSeries2002_08_27.pdf)

6.2.2 Linux for pSeries's Reliability, Availability, and Serviceability

A key attribute of the pSeries is its critical Reliability, Availability, and Serviceability (RAS) features. RAS is a concept used not only by IBM, considering that RAS features are the same. However, implementation may vary depending on the platform and environment on which you are working. For more information about RAS for the pSeries, go to:

http://www.ibm.com/servers/eserver/pseries/hardware/whitepapers/p690_ras.html

Linux for pSeries is rapidly developing RAS capabilities. However, many pSeries RAS features are only fully realized when running AIX.

The following pSeries RAS features are supported when running Linux:

- ▶ Chipkill and ECC memory
- ▶ Journalled file system (several available under Linux). Does not include AIX JSF.
- ▶ PCI Extended Error detection
- ▶ Redundant, hot-plug power, and cooling
- ▶ Error logging
- ▶ Service processor (partial)
- ▶ Boot-time processor and memory deallocation

Some of the pSeries RAS features that are not currently supported in Linux include:

- ▶ Hot-swapping of disk drives (can be done via command line procedures)
- ▶ Hot-plug PCI and memory
- ▶ Disk mirroring
- ▶ First Failure Data Capture
- ▶ Dynamic Processor Deallocation
- ▶ Dynamic logical partitioning (static LPAR is supported)
- ▶ PCI Extended Error recovery
- ▶ Error reporting to Service Focal Point
- ▶ Error log analysis
- ▶ Remote support
- ▶ High Availability Cluster Multiprocessing (HACMP): Alternative third-party and open source solutions may be available

6.2.3 Other sources of information

Through its Linux Technology Center, IBM is working with the open source community on a variety of projects to enhance the value of Linux for customers.

The Linux Technology Center has over 200 people devoted to developing and improving open source.

You can find the Linux Technology Center on the Web at:

<http://ltc.linux.ibm.com>

IBM also participates in several industry-led efforts, such as:

- ▶ Linux Standard Base
<http://www.linuxbase.org>
- ▶ Free Standards Group
<http://www.freestandards.org>
- ▶ Open Source Development Network
<http://www.osdn.com>
- ▶ Open Source Development Lab
<http://www.osdlab.org>

6.3 Consolidation

pSeries Linux-enabled servers provide the capability to run several logical partitions that can be made up of combinations of Linux and AIX on a single server. Logical partitioning on the pSeries provides a cost-effective platform for enterprise workloads while supporting open source standards and solutions. This makes the initial software power base a lot broader in the sense that you can incorporate open source software's rich variety of solutions into the system, while preserving the key capabilities of the AIX at the same time.

For example, consider a typical service provider or Web server environment. It is typically architected as a two- or three-tier model. In most installations, there are front-end systems (typically thin servers) to handle caching, proxy, domain name server (DNS), etc. There may then be a second tier of small systems to perform Web application serving using WebSphere or competing products in conjunction with an Enterprise Resource Planning (ERP) or Customer Relationship Management (CRM) products such as SAP.

The third tier of servers runs AIX or Linux on a large SMP system to provide the backoffice and DBMS functions that require high performance and scalability. In many cases, the first and possibly second tiers are running Linux or Windows NT. This results in more servers and the need for more staff and expensive software to manage multiple platforms. The Linux LPAR configuration with IBM Software, shown in Figure 6-5, is intended to minimize these complex tasks and to offer a more robust, secure solution.

This scenario shows all open source applications being used to consolidate what would normally be separate servers into multiple logical partitions.

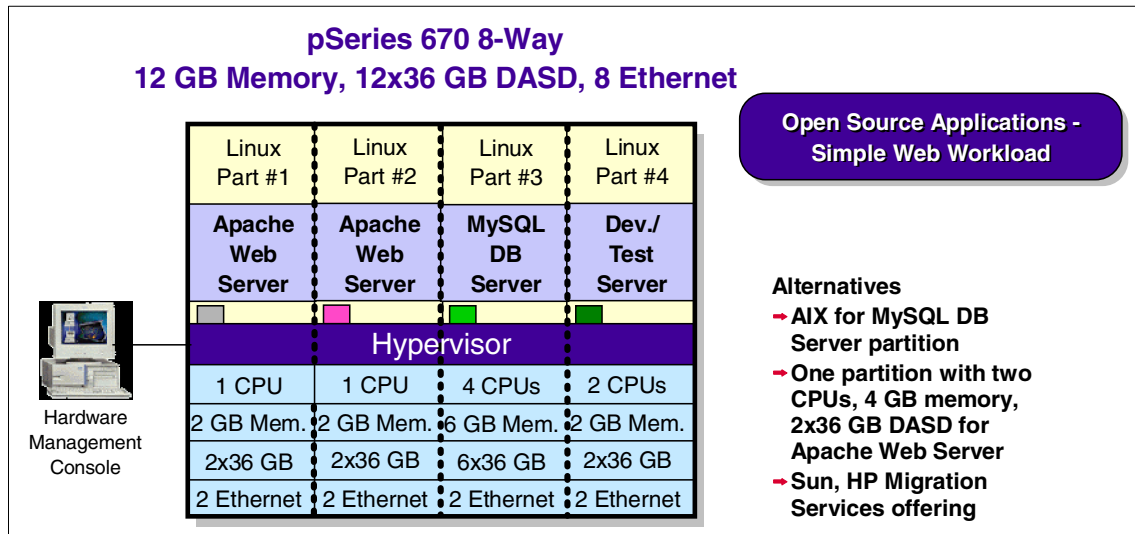


Figure 6-5 Web serving consolidation on Logical Partitions using Linux for pSeries

This Web serving consolidation can be scaled to a complex Web workload using Software Group products. DB2 can be used in partition 3 to accomplish that. Database management system (DBMS) scalability can be achieved using AIX instead of Linux in partition 3. Availability of Software Group products is currently an IBM statement of direction, so Linux for pSeries Software Group scalability is expected to grow.

6.4 Linux for pSeries clusters

None of IBM's software that has been announced for the IBM Linux Intel cluster products, such as Cluster 1300 and Cluster 1350, has been ported to Linux on pSeries. IBM has announced a statement of direction to port its Cluster Systems Management (CSM) technology to Linux for pSeries.

Cluster 1600 is IBM's key UNIX clustering solution for the pSeries. It has increased its scalability and support. Since it also designed to work with CSM, integration of Linux for pSeries-based servers, such as pSeries Model 630, is expected. You can learn more about Cluster 1600 on the Web at:

<http://www.ibm.com/servers/eserver/clusters/index.html>

The POWER4 family of processors should provide excellent Linux performance. POWER4 Linux clusters will surely meet present and future market needs. See Chapter 9, “Linux clusters” on page 243, for more information about Linux clusters.

Development of high availability Linux clusters for the distributed pSeries Linux environment is already available and supported.

Beowulf clustering technology and other open source and some commercial products can be used to cluster pSeries systems running Linux to provide compute or high availability clusters. To learn more about Beowulf, see:

<http://www.beowulf.org>

Myricom has the Myrinet switch available for Linux for pSeries. Myricom supplies and supports Myrinet software for Linux 2.4 kernel for IA32, IA64, Alpha, and PowerPC. The Myrinet switch can be used as a high-speed interconnect to cluster systems of pSeries machines running Linux. Gigabit or 10/100 Ethernet connections can also be used. For more information, see:

<http://www.myri.com/scs/>

6.5 Linux for pSeries application solutions

With Linux for pSeries, you can incorporate the boost of open source applications into your environment. Linux for pSeries makes it possible to enrich the capability of any pre-existing pSeries AIX environment with the power and benefits of open source software. IBM and other third-party companies are currently working on providing applications for Linux on the pSeries.

6.5.1 IBM Software for Linux on pSeries

pSeries servers offer a reliable foundation for Linux-based computing. IBM is working closely with leading Linux distributors to help ensure maximum performance and functionality.

IBM Java Developer Kit (JDK) Linux

This JDK is a development kit and runtime environment that contains the IBM just-in-time compiler, enhanced with a unique mixed mode interpreter and a re-engineered Java 2 virtual machine. Like earlier IBM Java versions, the IBM Developer Kit for Linux passes Sun's Java compatibility test and provides stability and performance when you deploy your enterprise e-business applications. For more information, see:

<http://www06.ibm.com/developerworks/java/jdk/linux140/?dwzone=java>

IBM JDK for Linux has been tested with SLES7 for pSeries in a 64 bit environment. For details about the tested Linux distributions for the IBM JDK Linux, refer to:

<http://www06.ibm.com/developerworks/java/jdk/linux/tested.html>

IBM is developing plans to support WebSphere Application Server, DB2, and many Tivoli products and compilers on Linux for pSeries. For latest updates and new software offerings, see:

<http://www.ibm.com/eserver/pseries/linux>

IBM XL Fortran V8.1 for Linux on pSeries

The XLF compiler is now available for Linux on pSeries. XLF for Linux on pSeries is based on the mature compiler technology of XLF for AIX, and currently supports SLES 8 for pSeries.

FORTRAN 77, Fortran 90, and Fortran 95 standards

XLF V8.1 for Linux on pSeries is a native and highly optimized compiler that exploits pSeries architecture and fully complies with FORTRAN 77, Fortran 90, and Fortran 95 standards. It partially supports draft Fortran 2000 standard features. It also supports IBM and industry extensions from various compiler vendors.

Supported draft Fortran 2000 features include:

- ▶ Allocatable components
- ▶ IEEE floating point exception handling
- ▶ Pointer with INTENT attribute

OpenMP

OpenMP is an industry specification that describes a common set of APIs. They support SMP programming in C/C++ and Fortran on all architectures, including UNIX and Windows NT platforms. OpenMP is a portable, scalable model, which gives shared-memory parallel programmers a simple and flexible interface for developing parallel applications for platforms, ranging from the desktop to the supercomputer.

XLF V8.1 for Linux on pSeries provides full support for the OpenMP Fortran API V2.0. For more information about OpenMP Fortran API V2.0, visit:

<http://www.openmp.org/specs/>

SMP programming

Applications may be parallelized automatically (through the use of a compiler option) or parallelized explicitly (through the use of SMP directives).

Use the `xl_f_r`, `xl_f90_r` or `xl_f95_r` command to compile XL Fortran SMP programs. For the compiler to recognize the SMP directives and activate parallelization, you must also specify the `-qsmp` option.

32-bit and 64-bit support

Use XLF for Linux on pSeries to develop both 32-bit and 64-bit mode applications. XLF V8.1 for Linux on pSeries provides 64-bit pointer and 64-bit addressability. High optimization levels (`-O5`, `-O4` and `-qhot`), interprocedural analysis, and parallelization are supported for 64-bit mode applications.

Exploiting the PowerPC architecture

XLF for Linux on pSeries supports POWER3, POWER4, and PowerPC architectures. Use compiler options `-qarch` and `-qtune` to instruct the compiler to generate code specific to an architecture. This allows the compiler to take advantage of machine-specific instructions to further improve performance. The `-qarch` option specifies the target architectures on which the resulting programs will run. The `-qtune` and `-qcache` options refine the degree of platform-specific optimization performed.

Asynchronous I/O

XLF for Linux on pSeries supports asynchronous I/O. Asynchronous I/O can be used to improve the performance of applications that perform input and output of large amounts of data. Asynchronous I/O makes it possible for other program statements to be executed while the data transfer is taking place.

Symbolic debugger support

Use `gdb` or other symbolic debugger for your programs.

Port or migration support

The XLF compiler helps you port or migrate source code among Fortran compilers. It provides full FORTRAN 77, Fortran 90, and Fortran 95 language support and selected language extensions (such as intrinsic functions, data types) from many different compiler vendors.

The `xl_f`, `xl_f_r`, `f77`, and `fort77` commands provide maximum compatibility with existing FORTRAN 77 programs. The default options provided with the `xl_f90` and `xl_f90_r` commands give access to the full range of Fortran 90 language features. The default options provided with the `xl_f95` and `xl_f95_r` commands give access to the full range of Fortran 95 language features.

IBM XL Fortran Run-Time Environment (RTE) V8.1 for Linux on pSeries

XLF RTE V8.1 for Linux on pSeries provides a run-time environment to support library modules, asynchronous I/O, threadsafe I/O, and I/O services and utilities.

You can find complete information on the Web at:

http://isource.ibm.com/cgi-bin/goto?it=usa_annred&on=203-053

IBM VisualAge C++ support with VisualAge C++, V6.0 for Linux on pSeries

The VisualAge C++, V6.0 for Linux on pSeries software is a highly productive and powerful development environment for building C and C++ applications.

1998 C++ and 1999 C international standards

VisualAge C++, V6.0 for Linux on pSeries builds on official industry standards, so your code can be ported easily among a variety of implementations.

1999 C standard

The ISO/IEC 9899:1999 International Standard (also known as C99) specifies the form and establishes the interpretation of programs written in the C programming language. This is the second edition of the standard. This update enhances the capability of the C language, provides clarifications, and incorporates technical corrections to the previous edition of the standard. The language features supported by the C compiler within VisualAge C++, V6.0 for Linux on pSeries are consistent with the specification described in the ISO/IEC 9899:1999 International Standard. The compiler works with a run-time environment that supports the standard to form a conforming implementation.

The C99 Standard offers additional language functions over the first edition (C89) of the standard, including:

- ▶ Long Long
- ▶ Bool
- ▶ Pragma operator
- ▶ Initializer allow con
- ▶ Empty new argument
- ▶ Macro with variable number of args
- ▶ Inline function specifier
- ▶ Complex

The `-qlanglvl` compiler option allows you to choose from the available language levels that you can use to compile your code within VisualAge C++. Your code can conform to the latest C99 Standard, use extensions to C99, or conform to the previous C89 Standard. VisualAge C++ offers you the flexibility to compile your code to the language level that you need.

1998 C++ standard

The ISO/IEC 14882:1998 International Standard (also known as Standard C++) specifies the form and establishes the interpretation of programs written in the

C++ programming language. This international standard promotes the portability of C++ programs among a variety of implementations.

ISO/IEC 14882:1998 is the first formal definition of the C++ language. The C++ compiler within VisualAge C++, V6.0 for Linux on pSeries supports the 1998 ISO/IEC Standard. In addition, VisualAge C++ supports several of the latest approved clarifications from the standards committee bringing the C++ compiler within VisualAge C++, V6.0 for Linux on pSeries to the latest level of the standard.

OpenMP

The OpenMP is an industry specification that describes a common set of APIs that support multiplatform SMP programming in C, C++, and FORTRAN on all architectures, including UNIX platforms. Jointly defined by a group of major computer hardware and software vendors, including IBM, OpenMP is a portable, scalable programming model. It provides SMP programmers with a simple and flexible interface for developing shared-memory parallel applications for platforms, ranging from the desktop to the supercomputer.

VisualAge C++, V6.0 for Linux on pSeries provides full support for the OpenMP Industry Specification V1.0. For more information about the OpenMP Specification, visit the OpenMP Web site at:

<http://www.openmp.org>

Compiler optimization

VisualAge C++, V6.0 for Linux on pSeries includes the benefits of IBM's compiler optimization technology. Through compiler options, directives, and pragmas, programmers are offered a number of ways to optimize their code. Using the -O2, -O3, -O4, and -O5 options, you can obtain optimization levels that run from comprehensive low-level optimizations to more extensive optimizations that include interprocedural analysis, loop optimization, and automatic machine tuning. Other options for optimization include:

- ▶ -qarch, -qtune, and -qcache that allow you to optimize your application to the POWER3 and POWER4 architecture.
- ▶ -qhot allows for high-order transformations, which are optimizations that specifically improve performance of loops through such techniques as interchange, fusion, and unrolling.
- ▶ -qipa invokes interprocedural analysis (IPA) that enables the compiler to optimize across different files (whole-program analysis). It can result in significant performance improvements.
- ▶ -qpdf (that is profile-directed feedback) uses results from sample program execution to fine tune optimizations in conditional branches and frequently executed code sections.

- ▶ `-qsmp` generates threaded code that is necessary to exploit shared-memory parallel processing leading to high-performance code.

32-bit and 64-bit support

You can use VisualAge C++, V6.0 for Linux on pSeries to develop 32-bit and 64-bit applications. VisualAge C++ provides 64-bit pointer and 64-bit addressability support for SMP and serial code. High optimization levels, interprocedural analysis, and SMP are supported for 64-bit code.

Ease of porting

VisualAge C++, V6.0 for Linux on pSeries has features that help you port existing code from AIX or other UNIX-based systems as well as from code originally written for GNU C/C++. For example, by compiling with the `-qinfo=por` option, you can filter the diagnostic messages emitted by the compiler to show only those that pertain to portability issues.

Also a comprehensive subset of features related to GNU C/C++ is supported by VisualAge C++, V6.0 for Linux on pSeries. Other GNU C/C++ features are recognized (accepted and ignored). This help you minimize the number of changes that you need to make to your GCC source code to ensure that it compiles successfully using VisualAge C++, V6.0 for Linux on pSeries.

Binary compatibility and coexistence

VisualAge C++, V6.0 for Linux on pSeries supports the C++ Abstract Binary Interface (ABI). This means that the compiler produces binary or object files that are compatible with those created with the GCC compiler.

VisualAge C++, V6.0 for Linux on pSeries includes the same headers as those used by a GNU compiler residing on the same system. This allows VisualAge C++, V6.0 for Linux on pSeries to achieve binary compatibility with GCC-compiled code. The VisualAge compiler optimizes the program while maintaining binary compatibility with objects produced by GCC V3.2.

Additional coexistence features include:

- ▶ IBM built-in functions for POWER3 and POWER4 architectures coexist with GNU C built-ins.
- ▶ Compilation uses the GNU C and C++ header files from the Linux distribution.
- ▶ Compilation uses the GNU assembler.
- ▶ Linking uses the GNU linker.
- ▶ Compiled programs use the GNU C and C++ run-time libraries.
- ▶ Compiled object files are compatible with g++ objects.
- ▶ Debugging uses the GNU debugger, gdb.
- ▶ Profiling uses the GNU profiler, gprof.

GNU debugger, gdb

VisualAge C++, V6.0 for Linux on pSeries supports debugging using the GNU debugger, gdb. This symbolic debugger, which comes with the system, supports full debugging capability for code without optimization and partial debugging for code at the -O2 optimization level.

Documentation and online help

VisualAge C++, V6.0 for Linux on pSeries provides product documentation in HTML and Portable Document Format (PDF). The following PDF documents are provided:

- ▶ *Getting Started with VisualAge C++ for Linux on pSeries* describes the different ways the compiler can generate executable code, the accepted file types for input and output, and useful tips for porting an existing application. The book also includes instructions for installing the program and provides an introduction to optimizing the performance of your application.
- ▶ *VisualAge C++ for Linux on pSeries C/C++ Language Reference* contains information about the C and C++ programming languages as supported by IBM, including the new portability and standardization extensions.
- ▶ *VisualAge C++ for Linux on pSeries Compiler Reference* contains information about the various compiler options, pragmas, macros, and built-in functions, including those used for parallel processing.
- ▶ *VisualAge C++ for Linux on pSeries Programming Tasks* contains information about programming using VisualAge C++ for Linux on pSeries not covered in other publications.

You can find these documents on the Web at:

http://isource.ibm.com/cgi-bin/goto?it=usa_annred&on=203-056

6.5.2 Other Linux for pSeries software applications

This section provides examples of pSeries Linux application support in 2003:

- ▶ Absoft: Pro Fortran for POWER Linux
<http://www.absoft.com/newppcproductpage.html>
- ▶ Bison Schweiz AG: Bison Solutions 2.6
<http://www.bison-group.com/sprachen/english/>
- ▶ Droplets, Inc.: User Interface V2.2
<http://www.droplets.com/product/>
- ▶ Micro Software Designs LTD: CUMIS Plus
<http://msd-tt.com/products.php3>

- ▶ Platform Computing: LSF 5.1 for distributed and grid computing solutions
<http://www.platform.com/>
- ▶ Point Solutions LTD: PAC Server
<http://www.point-solutions.co.uk/>
- ▶ Samba: Samba 2.2 File and Print Server
<http://us1.samba.org/samba/samba.html>
- ▶ Wired Minds GmbH: Mindsuite SLM 2.1
http://www.wiredminds.de/index_e.phtml
- ▶ Myricom provides and supports software for the PPC Linux Kernel to do Linux for pSeries clustering. The Myrinet switch is compatible with Linux for pSeries. See Chapter 9, “Linux clusters” on page 243, for more details. You can also refer to the following Web site:
<http://www.myri.com/>

6.6 Linux for pSeries infrastructure and architected solutions

Infrastructure and architected solutions for Linux for pSeries are under development by IBM and third parties. IBM has made statements of direction to port Software Group and infrastructure solutions to Linux for pSeries.

pSeries services solutions are also currently under development by IBM and third parties. Linux for pSeries services are offered along with IBM’s Linux-ready express configurations.

See Chapter 10, “IBM Services for Linux” on page 275, for more information and details on Linux for pSeries services and support. See 10.5, “Platform services” on page 287, and 10.5.3, “pSeries services” on page 290.



iSeries and Linux

This chapter provides a brief overview of the Linux solution available on the iSeries server. It includes information about the following topics:

- ▶ How Linux runs on the iSeries server
- ▶ Why you should use Linux on the iSeries server
- ▶ Linux and its relation to LPAR on the iSeries server
- ▶ iSeries Linux integration
- ▶ Linux support on specific iSeries models
- ▶ Linux distributions supported on iSeries
- ▶ iSeries customers and Linux
- ▶ iSeries and Linux as a server consolidation environment
- ▶ Linux application solutions available on iSeries
- ▶ Linux on iSeries service solutions

For additional and current information about the iSeries server and Linux, go to:

<http://www.ibm.com/servers/eserver/iseries/linux/>

7.1 How Linux runs on the iSeries server

The iSeries server provides the ability to run Linux in a logical partitioned environment. The primary partition must be running OS/400® V5R1 or later, which provides the support required to boot the Linux kernel in a secondary partition. The Linux operating system needs to be installed and run in a secondary partition.

Partitioning on the iSeries server supports multiple images of Linux and OS/400. It is available across the iSeries product line. For example, on a 1-way iSeries server, up to nine partitions can be supported. On a 32-way iSeries server, up to 31 Linux partitions can be supported. You can move processor and memory resources independently between partitions with a granularity of one hundredth (1/100) of a processor and 1 MB of memory. Shared processor resources can be dynamically moved between Linux and OS/400 partitions.

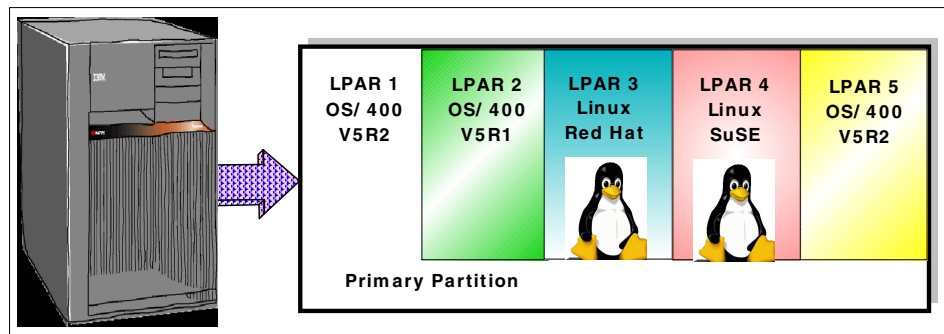


Figure 7-1 iSeries with multiple partitions, two running Linux

Beyond consolidating Linux servers, the iSeries server enables OS/400 applications to be extended with Linux. Linux applications can access DB2 Universal Database (UDB) as well as OS/400 programs and services through one of 16 high-speed virtual Ethernet LANs that connect the partitions. The iSeries server provides integration between Linux applications and OS/400 facilities. Linux applications can access data stored in DB2 UDB in OS/400 and can access files stored in the OS/400 integrated file system (IFS).

The ability to integrate Linux and OS/400 applications on the iSeries server draws together the primary strengths of both operating environments. For example, the database capabilities of OS/400 and the Web serving qualities of Linux provide a powerful, consolidated computing environment.

7.1.1 Why Linux on iSeries

You may choose to implement Linux on iSeries for the following reasons:

- ▶ **Total cost of ownership (TCO) or total cost of acquisition:** With shared processor and virtual I/O support, there is a minimal hardware requirement for a Linux environment on the iSeries server. There are also space and power consumption savings. The iSeries server can consolidate Linux and Windows NT server farms. This can translate into a minimal dollar investment to run Linux applications on the iSeries server.

iSeries solutions are designed to deliver cost advantages in IT environments. CIOview Corporation provides a tool to assess TCO. With the tool, TCOnow!, you can estimate the TCO for selected iSeries solutions compared to competitive alternatives. To use the tool, go to:

<http://www.ibm.com/servers/solutions/serverconsolidation/tco/>

- ▶ **Consolidation:** The iSeries server can consolidate several servers depending on the iSeries model. At the high end of the product line, iSeries supports up to 31 Linux partitions. At the low end, nine Linux partitions are supported on a 1-way server. The iSeries server also supports up to 32 OS/400 logical partitions, up to 48 Integrated xSeries servers, and up to 60 connected xSeries servers for Windows Server integration.

Examples of Linux consolidation scenarios in an LPAR environment include:

- *Workload consolidation:* Partition the iSeries server to support Linux workloads that normally require separate servers. Such applications as Enterprise Resource Planning (ERP), Web servers, and business intelligence are excellent candidates.
- *Foot print consolidation:* Partition the iSeries sever to support Linux workloads from different parts of the world.
- *Production and test:* Partition the iSeries server to support Linux production and test workloads.
- *High availability:* Partition and cluster the iSeries server for a Linux high availability solution.
- *Service Bureau, Application Service Provider, Business Recovery Service:* Partition the iSeries server to support a different Linux-based business in each partition.

iSeries provides customers with the ultimate flexibility to allow:

- ▶ **Capacity Upgrade on Demand (CUoD):** iSeries processors can be activated temporarily or permanently and dynamically allocated to Linux partitions.

- ▶ **Integration:** The iSeries server provides integration between Linux applications and OS/400 facilities. Linux applications can access data stored in DB2 UDB in OS/400 and files stored in the OS/400 IFS.
- ▶ **Performance:** A single large iSeries server with multiple Linux environments can deliver a lower cost per user than multiple systems.
- ▶ **Application flexibility:** The iSeries server supports Java, Domino, UNIX, C, COBOL, RPG, WebSphere, and Windows server applications via xSeries connected servers. Also, with unique virtual disk support, the iSeries server provides an excellent Linux test or development environment.
- ▶ **Resource sharing:** The iSeries server can share resources between OS/400 and Linux partitions. This can reduce the cost of managing and supporting a heterogeneous environment. LPAR supports sharing and dynamically moving processor resources between OS/400 and Linux.

iSeries offers Virtual Ethernet LAN facilities to support partition-to-partition communication. With Virtual Ethernet, you can define up to 16 networks for application communication. These connections can provide 1 Gb performance and do not require LAN adapters, switches, or physical networks.

The iSeries provides storage area network (SAN) facilities for Linux partitions. Linux partitions can access disk, tape, CD-ROM, and DVD resources in OS/400 partitions. With virtual I/O, Linux can leverage the performance availability and manageability of the advanced iSeries storage architecture. iSeries can dynamically add, move, or delete disk space assigned to Linux partitions. The storage spaces for Linux are carved from the OS/400 IFS. This centralization of storage resources enables customers to extend their OS/400 disk subsystem, management skills, and best practices to Linux.

For example, a full OS/400 system backup includes the virtual storage spaces used by the Linux partitions. Daily backups performed by Linux support file-level save/restore and can use the tape devices owned by OS/400.

With shared processor and virtual I/O support, there is a minimal hardware requirement for a Linux environment. This translates into a minimal dollar investment to run Linux applications on the iSeries server.

- ▶ **Reliability:** The iSeries server is a very reliable server and can provide industry leading single server availability.
- ▶ **Fault or failure isolation:** With the iSeries server LPAR technology, Linux workloads or applications can be isolated and result in easier fault isolation, less down time, and faster recovery.
- ▶ **Ease of use and systems management:** Linux running in an iSeries server environment can be easier to manage than multiple systems. The iSeries server provides a fast pipe between hardware resources and applications. All of the standard systems management functions, such as security, backup and

recovery, change and configuration management, network management, and workload management, can be more difficult and time consuming in a multisite, multisystem environment.

- ▶ **Security:** The OS/400 operating environment is a secure operating environment and offers an industry leading secure environment.

iSeries Linux services, listed in Table 7-1, are available for iSeries customers. The Web links take you to platform-based information about the services that are available for iSeries.

Table 7-1 iSeries Linux information matrix

Information	Location on the Web
Linux for iSeries applications	http://www.ibm.com/servers/eserver/iseries/linux/apps.html
Linux for iSeries Education	http://www.ibm.com/servers/eserver/iseries/linux/educ.html
Linux for iSeries planning	http://www.ibm.com/servers/eserver/iseries/linux/plan.html
Linux for iSeries Developer Resources	http://www.ibm.com/servers/eserver/iseries/linux/resc.html
Linux for iSeries Reference Library	http://www.ibm.com/servers/eserver/iseries/linux/library.html
Linux for the iSeries support	http://www.ibm.com/servers/eserver/iseries/linux/support.html

7.1.2 Linux and LPAR

With OS/400 V5R1 and later versions, IBM has enabled Linux to run in multiple secondary partitions. Each secondary partition has its own Linux environment and can support different Linux distributions. An OS/400 V5R1 or later version is required to run in the primary or managing partition. The primary partition is used to create partitions and move resources between partitions.

OS/400 V5R1

OS/400 V5R1 introduced a new concept for logical partitioning, which allows partitions to share processors. OS/400 V5R1 supports up to four partitions per processor. This opens up the uni-processor iSeries Model 270 and 820 to be partitioned into four partitions.

In addition, a 2-way iSeries Model 270 or 820 can support up to eight partitions (2 x 4), and a 4-way iSeries Model 820 can support up to 16 partitions (4x4).

iSeries servers with eight or more processors support a maximum of 32 partitions. On select iSeries Models 270, 820, and 840 with SSTAR processors, these shared processor partitions can run OS/400 V5R1 or Linux. On iSeries Models 820, 830, and 840 with ISTAR processors, Linux requires a dedicated processor.

With LPAR on OS/400 V5R1, the minimum processor allocation requirement to create a new partition is one-tenth (1/10) of a processor. This allows the primary partition processing requirements to be reduced to a minimum with the remaining processing units used to support secondary partitions.

An advantage of LPAR is resource movement among partitions. Processor, memory, and I/O resources can be moved among OS/400 and Linux partitions. On select iSeries models, one-hundredth (1/100) of a processor can be moved between partitions. On iSeries models that support only dedicated processors, processor movement among partitions is supported in whole processor increments. With OS/400 V5R1, moving resources to or from a Linux partition requires the Linux partition to be restarted.

A Linux environment on iSeries can be set up simply and inexpensively. The minimum resource for a Linux partition is one-tenth (1/10) of a processor on iSeries servers that support Linux shared processors. In addition, a minimum of 64 MB of memory is needed for Linux. The largest Linux partition is determined by the support provided in the Linux distribution. The 2.4 64-bit kernel supports up to 24 processors and up to 256 GB of memory.

OS/400 V5R2

In OS/400 V5R2, LPAR was enhanced for Linux. With a V5R2 primary partition, shared processor resources can be dynamically moved between Linux and OS/400 V5R2 partitions. In addition, the number of Linux partitions has increased on iSeries servers with SSTAR and POWER4 processors, for example:

- ▶ On a one-processor server, up to nine Linux partitions are supported.
- ▶ On a two-processor server, up to 19 Linux partitions are supported.
- ▶ On a four-processor server, up to 31 Linux partitions are supported.

For V5R2, Linux is supported on the new Model 800, 810, 825, 870, and 890 servers. The 890 supports up to 32 processors and 31 Linux partitions.

iSeries storage virtualization is also enhanced in V5R2. Storage spaces that are allocated from OS/400 can be dynamically added to Linux partitions. Storage virtualization enables Linux to leverage the advanced iSeries storage architecture.

iSeries Navigator is enhanced in V5R2 to support Linux partition management and resource movement. A wizard is available that makes creating a Linux partition even easier.

In January 2003, IBM announced additional enhancements for Linux. With the iSeries Models 870 and 890 Enterprise Edition, an extra processor is activated for Linux at no additional charge. In addition, Education and Services Vouchers are included in the Model 825, 870, and 890 Enterprise Edition servers that can be redeemed for the Linux on iSeries Implementation class and the Linux on iSeries Integration QuickStart service.

Linux I/O flexibility

The Linux environment on iSeries supports two types of I/O, virtual and direct. With *virtual I/O*, such I/O resources as disk, tape, and CD-ROM are owned by an OS/400 partition. OS/400 shares the resources with Linux. The partition sharing the resources is called the *host*. The host can be the primary partition, but it does not have to be. One partition can be the host for multiple Linux partitions. The shared resources are managed by OS/400. For example, disk space is allocated from OS/400's disk drives for Linux to use. OS/400 provides RAID protection and some backup and restore facilities for the Linux environment.

The Virtual Ethernet LAN support provided in OS/400 V5R1 and V5R2 extends to Linux. Virtual Ethernet provides 16 1 Gb Ethernet communication paths between partitions (OS/400 to Linux and Linux to Linux) without requiring additional hardware resources.

With *direct I/O*, the I/O resources are owned by Linux. For example, disk drives, LAN, or WAN adapters are under the control of one Linux environment. OS/400 cannot use these resources. A Linux server in a partition can take advantage of both virtual and direct I/O.

7.1.3 iSeries Linux integration

The iSeries server offers many points of integration that support Linux applications leveraging OS/400 applications and data.

A PowerPC Linux Open Database Connectivity (ODBC) driver provides access from a Linux application running on iSeries to the OS/400 DB2 UDB for iSeries database. This ODBC driver is included in select Linux distributions.

IBM has enhanced OS/400 NetServer with Samba support for the exchange of files between OS/400 and Linux systems. With this support, Linux can print to OS/400 print queues.

The IBM Developer Kit for Linux, Java 2 Technology Edition, is available for iSeries Linux. This Java virtual machine (JVM) in combination with the iSeries Toolbox for Java or its OpenSource version, JTOpen, allows developers to access DB2 UDB for iSeries data via Java Database Connectivity (JDBC) and leverage OS/400 programs and services. The Developer Kit and JTOpen are included in select Linux distributions.

7.1.4 Linux on iSeries models

The iSeries server supports Linux on models that support logical partitioning (LPAR) and OS/400 V5R1 or later. Table 7-2 shows the iSeries server models by processor feature, whether it supports LPAR, and the number of Linux partitions that are supported.

Table 7-2 The latest iSeries models

Model 800	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2463	V5R2	Yes	Yes	Yes	V5R2	1
	2464	V5R2	Yes	Yes	Yes	V5R2	1
Model 810	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2466	V5R2	Yes	Yes	Yes	V5R2	1
	2467	V5R2	Yes	Yes	Yes	V5R2	1
	2469	V5R2	Yes	Yes	Yes	V5R2	2
Model 825	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2473	V5R2	Yes	Yes	Yes	V5R2	3/6
Model 870	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2486	V5R2	Yes	Yes	Yes	V5R2	8/16
Model 890	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	197	V5R2	Yes	Yes	Yes	V5R2	24
	198	V5R2	Yes	Yes	Yes	V5R2	32
	2487	V5R2	Yes	Yes	Yes	V5R2	16/24
	2488	V5R2	Yes	Yes	Yes	V5R2	24/32
	2497	V5R2	Yes	Yes	Yes	V5R2	16/24
	2498	V5R2	Yes	Yes	Yes	V5R2	24/32

Table 7-3 Model 270

Model 270	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	2248	-	No	No	-	-	1
	2250	-	No	No	-	-	1
	2252	-	No	No	-	-	1
	2253	-	No	No	-	-	2
	2301	-	No	No	-	-	1
	2302	-	No	No	-	-	1
	2422 DSD	-	No	No	-	-	1
	2423 DSD	-	No	No	-	-	1
	2424 DSD	-	No	No	-	-	2
	2431	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2432	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2434	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2452 DSD	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2454 DSD	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
(1) Indicates that dedicated processors are required.							

Table 7-4 Model 820

Model 820	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	150	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	151	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	152	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2395	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2396	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2397	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2398	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2425 DSD	V5R2	Yes	No	-	V5R1	1
		V5R1	Yes	No	-	V5R1	
	2426 DSD	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2427 DSD	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2435	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2436	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2437	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2438	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2456 DSD	V5R2	Yes	Yes	Yes	V5R1	1
		V5R1	Yes	Yes	Yes	V5R1	
	2457 DSD	V5R2	Yes	Yes	Yes	V5R1	2
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2458 DSD	V5R2	Yes	Yes	Yes	V5R1	4
		V5R1	Yes	Yes	Yes	V4R5(1)	

(1) Indicates that dedicated processors are required.

Table 7-5 Model 830

Model 830	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	153	V5R2	Yes	Yes	No	V5R1	8
		V5R1	Yes	Yes	No	V5R1	
	2349	V5R2	Yes	Yes	No	V5R1	4/8
		V5R1	Yes	Yes	No	V5R1	
	2351	V5R2	Yes	Yes	No	V5R1	1/8
	Ã	V5R1	Yes	Yes	No	V5R1	
	2400	V5R2	Yes	Yes	No	V5R1	2
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2402	V5R2	Yes	Yes	No	V5R1	4
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2403	V5R2	Yes	Yes	No	V5R1	8
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
(1) Indicates that dedicated processors are required.							

Table 7-6 Model 840

Model 840	Feature #	Primary Partition	Supports LPAR	Supports Linux	Linux Shared Processor	Minimum Requirement Secondary	n-way
	158	V5R2	Yes	Yes	Yes	V5R1	12
		V5R1	Yes	Yes	Yes	V5R1	
	159	V5R2	Yes	Yes	Yes	V5R1	24
	Å	V5R1	Yes	Yes	Yes	V5R1	
	2352	V5R2	Yes	Yes	Yes	V5R1	8/12
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2353	V5R2	Yes	Yes	Yes	V5R1	12/18
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2354	V5R2	Yes	Yes	Yes	V5R1	18/24
		V5R1	Yes	Yes	Yes	V4R5(1)	
	2416	V5R2	Yes	Yes	No	V5R1	8/12
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2417	V5R2	Yes	Yes	No	V5R1	12/18
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2418	V5R2	Yes	Yes	No	V5R1	12
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2419	V5R2	Yes	Yes	No	V5R1	18/24
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2420	V5R2	Yes	Yes	No	V5R1	24
		V5R1	Yes	Yes	No	V4R5(1)	
		V4R5(1)	Yes	No	-	V4R5(1)	
	2461	V5R2	Yes	Yes	Yes	V5R1	24
		V5R1	Yes	Yes	Yes	V4R5(1)	

(1) Indicates that dedicated processors are required.

7.1.5 iSeries Linux distributions

Linux for the iSeries server is available from leading Linux distributors Red Hat, SuSE, and Turbolinux. These Linux distributions are based on Version 2.4 of the Linux kernel for PowerPC. Starting in January 2002, distributions became available for iSeries based on the 64-bit PowerPC kernel. The latest Linux distributions support 64-bit and 32-bit applications. The Linux distributions include the device drivers and other support required to run on an iSeries server.

Beyond the kernel, these Linux distributions include the operating system, middleware, and applications. For example, they include the Apache, Tux, and Zope Web servers; the Sendmail, Postfix, and exim mail facilities and servers;

the Samba and NFS file servers; the lprng and cups print servers; and the TCP wrappers and iptables firewalls. The selected distributions also contain key integration technology with OS/400 including ODBC and JDBC support to access DB2 UDB in OS/400.

The current distributions are:

- ▶ Red Hat Linux 7.1 for iSeries, 64-bit, supports 32-bit applications
- ▶ SuSE Linux Enterprise Server 8, supports 64-bit and 32-bit applications
- ▶ Turbolinux 7 for iSeries, supports 32-bit applications

7.2 Consolidation

With iSeries LPAR technology and support for multiple application environments, e-business technologies, and operating systems, you have the flexibility to choose the best applications for your needs. In a single server, you can run your core business and e-business applications, such as Lotus Domino on OS/400 or Linux e-mail servers, Web servers, or firewalls. You can centrally manage them while they are running completely independent of one another.

The iSeries server can host and integrate objects and data from another operating system with its own core operating system, OS/400, making it versatile for mixed IT environments. The iSeries supports these operating environments:

- ▶ OS/400 with up to 31 secondary partitions (run OS/400 natively).
- ▶ Linux for PowerPC running natively within an iSeries logical partition. Linux applications can supplement OS/400 applications. You can run up to 31 separate Linux server partitions, all on the same iSeries server.
- ▶ Microsoft Windows servers via the Integrated xSeries Adapter and the Integrated xSeries Server.
- ▶ UNIX via the OS/400 Portable Application Solutions Environment (OS/400 PASE). OS/400 PASE is an environment that is both independent and tightly integrated with OS/400. UNIX applications can be ported to and executed in the OS/400 PASE environment. They can also access objects in OS/400, all with little or no modifications to the original UNIX application.

Figure 7-2 shows how the iSeries works as a consolidation server.

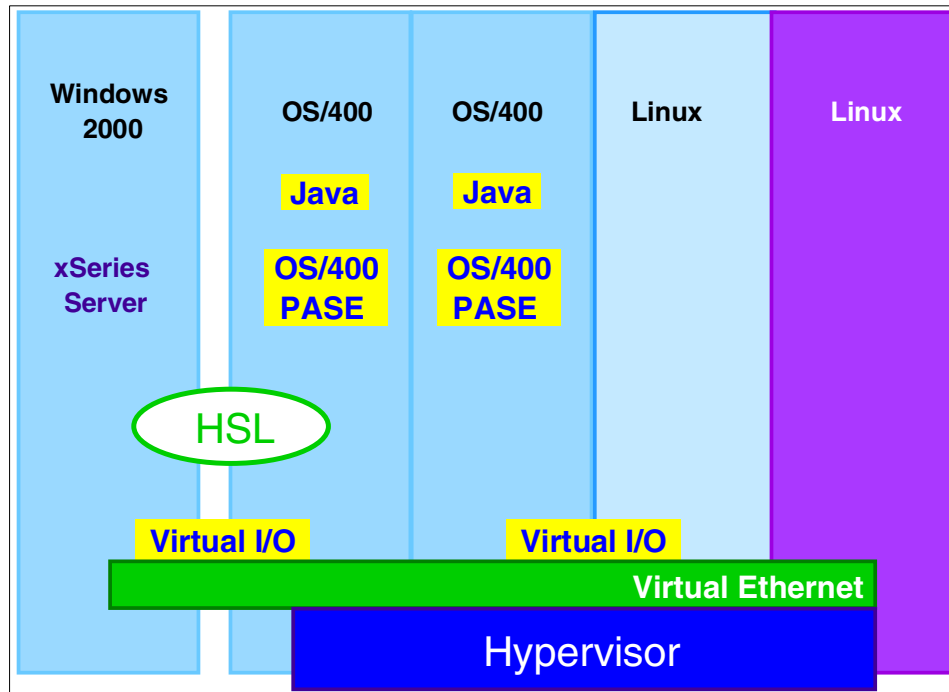


Figure 7-2 iSeries as a consolidation server

7.3 Linux clusters

You can create Linux clusters on the iSeries server. This is possible by using such open source tools as Open Source Cluster Application Resources (OSCAR).

For additional information about Linux clustering technology and solutions, see Chapter 9, “Linux clusters” on page 243.

7.4 Performance

Linux on iSeries workloads are now available in the IBM @server Workload Estimator. Specifically, Samba file serving and Apache Web serving workloads can now be defined. Workload Estimator estimates the size of the Linux partition that is needed to support these workloads. You can access the Workload Estimator on the Web at:

<http://www.ibm.com/servlet/EstimatorServlet>

In addition, a Workload Estimator plug-in is now available for Bynari software. You enter the mail serving workload characteristics (such as number of mail users and the number of messages sent per hour) into the plug-in. Then, the Workload Estimator tool launches and estimates the size of the Linux partition needed to support the mail workload. You can find this plug-in on the Web at:

<http://bynari.net/index.php?id=1314>

7.5 Linux in a distributed enterprise

Running Linux on the iSeries server combines the strengths of Linux and OS/400 for an integrated solution. Linux delivers excellent open source solutions, while OS/400 is a premier integrated platform for business solutions. Linux enables a new stream of e-business applications for the iSeries platform that complements its strength as an integrated core business solution. Linux applications benefit from the iSeries platform's ability to provide resource flexibility, reliability, and connectivity to other applications on a single server.

When deploying Linux on an iSeries server, you can use the same base configuration each time. Linux is run in an LPAR, so you always need an OS/400 primary partition. When you have a hosted Linux partition, the primary partition can be the hosting partition. Therefore, the base scenario is an iSeries server with a heterogeneous combination of Linux and OS/400 partitions.

It is possible to deliver a reliable all-in-one-box integrated solution to end users in a distributed enterprise. Figure 7-3 shows an example of such an implementation, in which each operating system is used for its intended purpose. OS/400 can be the best platform for business applications. Through the years, it has proven to be quite reliable. Linux on the other hand, has proven to be a stable operating system for (inter) networking in general, such as providing firewall functions.

In this example, LPAR is leveraged to support numerous operating environments on one iSeries server. Typically such environments as firewalls and Internet servers are installed on their own servers. Figure 7-3 shows that one iSeries server can support five different servers. It also shows how virtual and direct I/O in a Linux environment can be leveraged. The firewall that interfaces with the Internet has a dedicated LAN adapter. The other resources needed by this partition (disk, tape, CD-ROM, and LAN) are virtual. Virtual LAN is also used to provide a direct connection to the Web-serving partitions.

Two Web-serving partitions running a commerce or business-to-business application can provide a load balancing and high availability environment. The Web-serving partitions are then connected to a second firewall via a different

virtual LAN connection. The multiple virtual LAN connections provide another level of security between the outside and inside environments.

The second firewall uses virtual I/O disk, tape, CD-ROM, and LAN resources. It does not have any direct I/O devices except for the native LAN. It is possible to connect the direct LAN adapter to the first OS/400 partition. The advantage of this approach (with the direct LAN adapter connected to the Linux partition) is that all system users must pass the firewall. This means that you can even stop access to the core business software for some parts of the internal network or log activity to that system.

The OS/400 partition runs the line-of-business applications and contains the data on which the Web applications are based. These five partitions can be set up on a small iSeries server with two processors (for example, a Model 810). The OS/400 and Linux partitions can share these two processors, and resources can be moved between the partitions based on the workload demands. This same solution can be set up on a large iSeries server with additional processors.

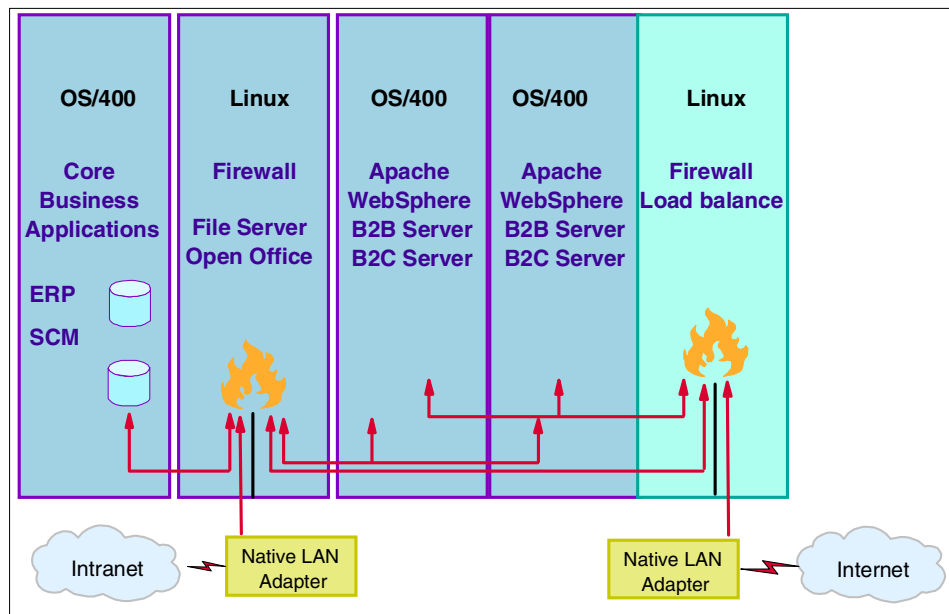


Figure 7-3 iSeries all-in-one integrated solution

7.6 Linux application solutions

PowerPC Linux applications run out of the box on the iSeries server and run pSeries and other PowerPC platforms. The Linux distributions for iSeries contain such notable open source applications as Apache, Samba, IMAP, OpenLDAP, Perl, SendMail MTA, Squid, and PostgreSQL.

Linux applications are sourced from many areas. Intel-based Linux applications need to be recompiled for PowerPC. You can download open source applications that are available on the Internet and compile them for PowerPC. Applications can also be developed for PowerPC. Compilers and languages are available with the Linux for iSeries distributions (for example, C, PHP, and Perl). You can also develop applications on PowerPC workstations or servers and simply move them to run on the iSeries. Cross-compilers are available that support writing applications on Intel systems and compiling them to run on PowerPC platforms.

Solution providers have announced products for Linux on the iSeries server. These offerings include products from Symantec, Sage, Bynari, Mapics, and eOne. Some examples are listed. You can find a complete list on the Web at:

<http://www.ibm.com/servers/eserver/iseries/linux/apps.html>

To assist solution providers with bringing additional solutions to market, iSeries Test Drive for Linux is available that provides remote access to Linux partitions on iSeries. For additional information, see:

<http://www.iseries.ibm.com/developer/factory/testdrive/index.html>

Infrastructure applications

Among the infrastructure applications announced for iSeries Linux solutions are:

- ▶ **extend5:** This is Acucorp's integrated development environment (IDE) for COBOL, which includes ACUCOBOL-GT, the company's flagship product that supports iSeries Linux. Large companies and ISVs have written or extended their COBOL code using ACUCOBOL-GT. This IDE provides software developers with the capability to move their existing ACUCOBOL-GT applications to Linux for iSeries without recompiling.
- ▶ **CM_SAFE Server for iSeries:** This application, from GMx Solutions, provides configuration file management capabilities for open systems hosts, with server support available for iSeries Linux partitions and agent support available for all major UNIX variations and platforms.
- ▶ **Enterprise Firewall for iSeries for Linux:** This solution from Symantec provides iSeries customers fast and secure firewall protection to help secure the network and mission-critical applications against unwanted intrusions. This application is due to general availability during 2003.

- ▶ **ORION:** This multi-platform managed availability solution is provided by Vision Solutions. ORION integrates and manages an entire enterprise's data and application availability needs across multiple applications, databases, and servers from a single enterprise view. ORION allows you to manage data sharing over five industry leading databases residing on multiple server operating systems. It also allows you to manage data and application resiliency on the iSeries platform, Linux servers from SuSE, Red Hat, and Turbo Linux, and Windows 2000.

The ORION engine is architected entirely in Java and XML and provides high portability across operating systems and high scalability in the enterprise. ORION provides an intuitive central interface that gives system operators and administrators maximum autonomy and control over their entire enterprise availability needs.

Line-of-business applications

The line-of-business applications announced for iSeries Linux solutions include:

- ▶ **MAPICS ERP solutions for iSeries Linux:** This solution include two ERP foundations, plus supply chain management, collaborative commerce and maintenance, and calibration management functionality. This is the first manufacturing-centric ERP application to support Linux on this platform.
- ▶ **Sage Line 500 for iSeries Linux:** This business management application, from Sage, enables mid-sized businesses to take a single view of their finance, distribution, manufacturing, and service operations.
- ▶ **Magic eMerchant V2 for iSeries Linux:** This solution from Magic Software is a customizable e-commerce framework designed to deploy an e-business CRM solution.

Web-enablement applications

Examples of Web-enablement applications announced for the iSeries server Linux solutions are:

- ▶ **eOneCommerce:** This solution from eOne Group is a flexible e-business software application and Internet publishing tool that runs in Linux for iSeries.
- ▶ **Insight Server:** This software application from Byarni delivers lightweight or enterprise-level Linux-based messaging services. Insight Server is compatible with Microsoft's suite of Outlook clients, providing similar functionality of an Exchange Server through the use of standards-based protocols.
- ▶ **CoSORT data sorting and ETL acceleration package:** This solution provides high volume select-sort-join-convert-aggregate-report tasks under Linux on iSeries.

- ▶ **Garoon:** This solution from Cybozu includes key collaboration applications that provide a “ready-to-run” intranet solution for small and medium-sized workgroups. Workgroup members can access e-mail, schedules, contacts, tasks, announcements, memos, files, projects, company forms, and much more from a personalized home page.
- ▶ **DI-Atlantis:** This solution from Dimensional Insight, Inc. is a multidimensional software package designed to transform data for intuitive access and analysis. It allows you to give users the information they need in a form they can understand.

7.6.1 Small office example using OpenOffice

The iSeries server offers a consolidated server for the small office. OS/400 runs the line-of-business applications while Linux runs OpenOffice for personal productivity applications. OpenOffice is an open source product that provides a word processor, spreadsheet, and presentation builder. Based on the popular StarOffice product, OpenOffice can read and write the respective Microsoft Office files. Sun provided StarOffice to open source resulting in OpenOffice.

OpenOffice can run on a server, such as the iSeries, and send the graphical user interface (GUI) to X-Windows clients. This is a great solution for the small office environment. OpenOffice is installed. After it is on the server, users can customize the OpenOffice environment, such as the default printer.

OpenOffice is available for all types of client systems. You can download it from the Web at:

<http://www.openoffice.org/>

7.6.2 Java

iSeries supports current levels of Java with the JVM and the IBM Toolkit for Java. For additional information and downloads, see:

<http://www.ibm.com/servers/eserver/iseries/toolbox/>

7.6.3 IBM Software Group products

IBM Software Group has announced a statement of direction to provide WebSphere and DB2 for Linux on the iSeries server. For additional information about Software Group products, see 8.1, “Overview of Linux and IBM Software availability” on page 200.

7.7 Infrastructure solutions

For infrastructure solutions, see 7.6, “Linux application solutions” on page 193.

7.8 iSeries solutions

The iSeries is the industry's leading business server for mid-market companies. It features the intelligent integration of key e-business functions, innovative technologies, application flexibility, superior performance, and the new tools required for managing e-business. iSeries Linux solutions are available for you regardless of the size or complexity of your business.

This section gives you the following examples of Linux solutions that were architected on the iSeries server:

- ▶ Application service provider (ASP): Aitana SBS
- ▶ File server, consolidation: Banco do Brazil
- ▶ Firewall, consolidation, e-commerce: Churchill China

Aitana SBS

Aitana SBS, based in Valencia, Spain, licenses its software solutions and offers them as an ASP. Recently Aitana began to offer Linux as a platform for customers wanting the reliability and cost efficiency of Linux with the price/performance of the IBM @server. Aitana operates four iSeries machines, with Linux LPARs on the two largest 270 models for development, testing, and production.

The challenge for Aitana SBS was to meet the growing need for software and hosting solutions based on Linux. The Aitana SBS architected solution included iSeries servers running Linux and xSeries servers. Aitana SBS specializes in collaboration solutions based on IBM Lotus Domino and other solutions based on the WebSphere Application Server and IBM Universal Database. Aitana SBS chose the Red Hat Linux distribution for their solution.

After transitioning to their iSeries Linux architected solution, Aitana SBS estimates the following business benefits for their ASP customers:

- ▶ Up to a 20% improvement in application speed
- ▶ Up to a 50% improvement in operating system availability
- ▶ Reduced administrative costs

For the complete story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAVO-5FEVNR?OpenDocument&Site=software>

Banco Do Brasil

Banco Do Brasil is a banking concern with a network of branches in 30 countries or regions and corresponding banking services with nearly 1,000 top banks in 145 countries or regions. The challenge for Banco Do Brasil was to consolidate its European operations. Banco Do Brasil decided to implement a plan to consolidate its European operations onto three iSeries servers in its London office. The goal of this plan was to increase IT capacity and performance, generate savings on licensing fees and administration costs by switching to Linux, provide flexibility with scalable systems, and provide an infrastructure capable of adapting to its future business needs.

The architected solution for Banco Do Brasil was to use iSeries servers to:

- ▶ Run Banco Do Brasil's core banking application. The system has increased capacity and LPARs to enable the bank to handle its entire European operation from a single point.
- ▶ Run Dedicated Server for Domino (DSD) for workflow, mail, and shared applications throughout its European operation.
- ▶ Run it as a Linux file server, as well as xSeries connected directly to iSeries servers using Integrated xSeries Adapter technology.

After transitioning to their iSeries Linux architected solution, Banco Do Brasil experienced the following benefits:

- ▶ Cost savings, improved efficiency
- ▶ Improved and simplified administration
- ▶ Improved performance
- ▶ Improved quality and accessibility of information

To learn more about this story, see:

<http://www.ibm.com/servers/eserver/iseries/news/pdf/bancodobrasil.pdf>

Churchill China

Churchill China is a major manufacturer and distributor of high-quality ceramic and china tableware, with a 200-year old craft-led heritage. Their challenge was to develop a robust and scalable commerce for an e-business supply chain management (SCM) system. Formidable overseas competition meant that Churchill China had a crucial business need to adapt to the world of e-business.

The architected solution for Churchill China was developed using Wizz400, a Clover product, running in a LPAR environment and a Linux-based firewall system, also running in an LPAR. The iSeries server chosen was a powerful iSeries Model 820.

Wizz400 is middleware that is an integral part of Churchill China's e-business system. Churchill China's Web site turned into an interactive e-business tool that can extract live data, display accurate stock and account information, and collect order data.

After transitioning to their iSeries Linux architected solution, Churchill China experienced the following benefits:

- ▶ Decreased costs through consolidation of three servers to one
- ▶ Adaptability that includes multilingual capabilities
- ▶ Adaptable and flexible architecture that allows further development to improve customer service
- ▶ Increased accuracy of orders, stock
- ▶ Improved customer service
- ▶ Improved and simplified administration
- ▶ Increased performance and scalability

To learn more about this story, see:

<http://www.ibm.com/servers/library/casestudies/churchillchina.html>

7.9 iSeries service solutions

Support for Linux on iSeries has two key components. Support for the integration of selected Linux distributions on iSeries is provided by the iSeries support organization. Integration support includes help with installing Linux successfully, creating partitions for Linux, and sharing iSeries resources with the Linux partition. This integration support is provided with an iSeries OS/400 SupportLine contract.

Support for the Linux distribution itself is a separate option. This support is available from SuSE, Turbolinux, Red Hat, IBM Global Services, and other Linux service providers.



IBM Software for Linux

As part of IBM's investment in, and strategy around, the Linux operating system, IBM Software Group (SWG) is continually enlarging its software catalog for Linux. Under the DB2, Lotus, Tivoli, and WebSphere brands, IBM Software comprises the most comprehensive stack of middleware vital to successful deployments of data management, group collaboration, systems management, and Web services.

This chapter provides an overview of IBM Software on Linux through the following topics:

- ▶ An overview Linux and IBM Software availability
- ▶ WebSphere brand software for Linux
- ▶ IBM Data Management software for Linux
- ▶ Tivoli brand software for Linux
- ▶ Lotus brand software for Linux
- ▶ Sample Linux solution deployments
- ▶ Customer references
- ▶ IBM Software for Linux resources

8.1 Overview of Linux and IBM Software availability

During the last few years, Linux has become the fastest growing general computing operating system in the market. It is now acknowledged to be an effective business solution platform.

Middleware is often the prominent driving force of business solutions and is indispensable to an infrastructure. Middleware products are required to be available, flexible, and contain multiple functions. Furthermore, the market demands that they conform to open standards, such as J2EE and Web services. As a result, much of the IBM middleware catalog, including DB2, WebSphere, Lotus and Tivoli, is now available for the Linux platform.

Business comes with risks. Therefore, it is important to select a software platform with proven reliability. It is also important to choose middleware that can execute well on many platforms, including Linux. Both the operating system and applications that provide a solution to the customer's needs must be available now. Linux and IBM Software meet these criteria and offer a very powerful solution that includes:

- ▶ **Flexibility:** Linux allows you to adopt a variety of hardware and software that are appropriate to each business task. Linux can run on a variety of platforms along with IBM middleware, making it a good candidate for a multiplatform environment. Also, new applications are easily ported to an existing system.
- ▶ **Reliability:** The extensibility and availability of Linux and IBM Software have already been proven in customer environments. Combining IBM Software with the Linux operating system results in improved reliability, productivity, and customer satisfaction.
- ▶ **Cost performance:** IBM Software running Linux results in a reduced cost of the operating system, improved management of distributed systems, and a wide selection of compatible hardware. Also, running IBM Software on Linux greatly contributes to decreasing your total cost of ownership (TCO).

You can find the IBM Software for Linux home page on the Web at:

<http://www.ibm.com/software/is/mp/linux/software/index.shtml>

Note: The *SWG Products, Solutions and Strategies Reference Guide*, G325-2132, provides an excellent overview of each of our software products. While this document is not specific to Linux-based products, it provides a wealth of information about each of IBM's software offerings. To obtain a copy, go to the following Web site, select your country or region location, and search for the publication number:

<http://www.elink.ibm.com/public/applications/publications/cgibin/pbi.cgi>

IBM Software for Linux availability

IBM Software products for Linux are supported on a per architecture, per distribution basis. You can find the most current list of IBM Software products for Linux availability on the Web at:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

Note: IBM Software that is listed as available for IBM Intel-based xSeries systems is also available for Intel-based systems from vendors other than IBM.

As a general statement of direction, IBM Software brands will support two Linux distributions: the Linux “enterprise” distributions from Red Hat (Red Hat Advanced Server) and the UnitedLinux consortium (SuSE SLES, Turbolinux TLES, SCO Linux Server, and Conectiva CLEE). Other distributions may be supported on an exception basis.

For a detailed description of each Linux distributor, see 1.9, “The Linux distributions” on page 18.

Software availability may also be affected by:

- ▶ Demand and target market for products on specific hardware architectures: For example, a market may demand a middleware product on IBM Intel-based xSeries servers but not on pSeries servers.
- ▶ Requirements on the Linux operating system by the software for performance or functionality: For example, many third-party software products require an “enterprise” version of Linux, such as SuSE Linux Enterprise Server, rather than a “desktop” version of Linux, such as the SuSE Linux Office Desktop.
- ▶ Market availability of IBM hardware platforms
- ▶ Market availability of Linux distributions
- ▶ Product release schedules of Linux distributors
- ▶ Internal IBM brand management decisions
- ▶ IBM middleware product introduction and obsolescence

Note: The absence of an IBM Software product in the software availability matrix is not an indication that the product will never be available for Linux. If you need an IBM Software product that is not listed in the software availability matrix, contact your local IBM sales representative.

8.2 IBM WebSphere brand products on Linux

WebSphere is infrastructure software for dynamic e-business. It delivers a proven, secure, and reliable software portfolio. WebSphere Application Server, WebSphere Commerce Suite, WebSphere MQSeries, and other product groups and development tools employ the latest Web technology to enable smooth extension of e-business to meet market needs. Figure 8-1 shows the three categories into which the WebSphere software platform can be divided.

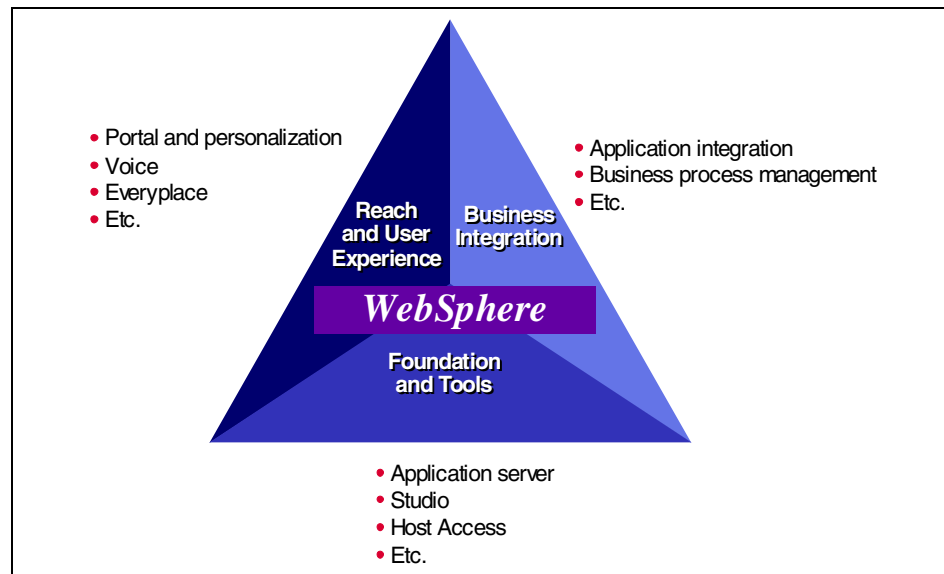


Figure 8-1 WebSphere platform overview

Each category is further explained here:

► **Reach and user experience**

This simplifies the process of creating a single interactive user experience. It allows you to deliver user-centric interactions for customers, partners, and employees alike through portal interfaces across all of your business processes. Commerce offerings enable full transactional support and integrate with existing business systems. Expanding these user interactions for mobile access, providing a common experience throughout, delivers real value to e-business in our global and mobile marketplace.

► **Business integration**

This provides the ability to integrate data, applications, processes, and people. It gives business the power to collaborate with supply chain partners, leverage intelligence, and enhance customer relationship management.

Leveraging existing capital investments is another benefit, since you are integrating your current applications instead of installing all new ones. WebSphere's leadership in business integration is extended through security, performance, and scalability enhancements.

► **Foundation and tools**

Success is creating a foundation with flexibility and scalability for future growth. Through this category, WebSphere delivers business agility and scalability by leveraging Web services and existing software assets. Comprehensive Java 2 Enterprise Edition (J2EE) development and deployment capabilities are all necessities when you build for integration.

The foundation infrastructure, WebSphere Application Server Version 5, provides a rich e-business application deployment environment with a set of application services. These services include enhanced capabilities for transaction management in a heterogeneous environment, comprehensive Web services support, increased security, performance, availability, connectivity, and scalability. In addition, IBM WebSphere Studio provides an open, integrated application development environment that cuts development costs via unprecedented gains in productivity and quality.

You can learn more on the WebSphere for Linux home page at:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/websp.shtml&S_TACT=__S_TACT__&S_CMP=__S_CMP__

The value of WebSphere on Linux

The WebSphere product set is:

- **Flexible:** WebSphere is based on open standards for Java and Web services. This includes J2EE, Extended Markup Language (XML), Universal Description, Discovery, and Integration (UDDI), Simple Object Access Protocol (SOAP), and Web Services Description Language (WSDL). It supports the complexity of interlinking enterprises.
- **Reliable:** WebSphere provides a transactional environment for running Java servlets and Enterprise JavaBeans (EJB) on top of Apache Web servers
- **Cost effective:** WebSphere reduces application development costs by reusing code. It optimizes operating costs because it can run on Linux from PC servers to server clusters to mainframes.

For more information on the IBM WebSphere Family, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

8.2.1 WebSphere product availability on Linux

The following products are among the list of WebSphere branded products on Linux:

- ▶ WebSphere Application Server
- ▶ WebSphere Business Components Composer
- ▶ WebSphere Business Components Studio
- ▶ WebSphere Commerce
- ▶ WebSphere Edge Server
- ▶ WebSphere Everyplace® Connection Manager
- ▶ WebSphere Extended Personalization
- ▶ WebSphere Host on Demand
- ▶ WebSphere MQ for Linux
- ▶ WebSphere MQ Everyplace
- ▶ WebSphere MQ Integrator
- ▶ WebSphere MQSeries
- ▶ WebSphere Performance Pack Cache Manager
- ▶ WebSphere Portal Server
- ▶ WebSphere Site Analyzer
- ▶ WebSphere Studio Application Developer
- ▶ WebSphere Studio Device
- ▶ WebSphere Studio Homepage
- ▶ WebSphere Studio Application Monitor for z/OS
- ▶ WebSphere Studio Site Developer
- ▶ WebSphere Telecom Application Server and Toolkit
- ▶ WebSphere Transcoding Publisher

It is important to note that all versions of all products may not be available on all hardware platforms or for all versions and distributions of Linux.

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

8.2.2 Competitive software

Despite the excellent functionality, performance, and price of WebSphere products, there are many open source and proprietary competitors.

Generally speaking, WebSphere products perform and scale better than the competitive offerings. In addition, IBM provides industry leading support for WebSphere products, which is not often available for open source products.

Alternatives to some WebSphere products include:

▶ **IBM HTTP Server (powered by Apache)**

The evolving product of an open source project, Apache is the most widely used Web server in the world. It is available at no cost from the Apache project.

The IBM WebSphere Application Server is totally compatible with the IBM HTTP Server (powered by Apache). In fact, Apache is the basis for the IBM HTTPD server that is provided with the WebSphere Application Server. Apache is included as part of every Linux distribution. Support for Apache may be available from third-party organizations.

You can find additional information about Apache at:

<http://httpd.apache.org/>

▶ **Tomcat**

The evolving product of an open source project, Tomcat is the servlet container that is used in the official reference implementation for the Java Servlet and JavaServer Pages technologies. It is available at no cost from the Jakarta Tomcat project and is often included as part of most Linux distributions.

For additional information about Tomcat, see:

<http://jakarta.apache.org/tomcat/>

▶ **JBoss**

Written entirely in Java, JBoss is a J2EE application server and EJB container. It is available at no cost from the JBoss Group LLC. Note that JBoss is not currently J2EE certified, although it is claimed to conform to EJB specifications. You can purchase support for JBoss from JBoss Group LLC.

For additional information about JBoss, see:

<http://www.jboss.org/>

▶ **JRun**

JRun is a proprietary J2EE application server that was developed by Macromedia. You can learn more about JRun on the Web at:

<http://www.macromedia.com/software/jrun/>

▶ **Commercial competitors**

The other Linux-enabled competitors to WebSphere Application Server include Oracle Application Server, Sun ONE Application Server (formerly iPlanet), and BEA WebLogic.

► **Other competing technologies**

It is not unusual for Linux-based Web servers to interface via CGI to scripts written in Perl, PHP, etc. These are adequate and freely available alternatives for low volume Web serving. Note that compilers and interpreters for these scripting languages are always included in Linux distributions.

8.2.3 Linux-related WebSphere resources

This section provides information on where to find additional WebSphere resources.

Evaluation copy of WebSphere software

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation, are available via the Web. You can also obtain evaluation copies of WebSphere Advanced Server, WebSphere Studio Site Developer, WebSphere MQ, and Linux-enabled IBM Software.

To learn more about how to obtain your own evaluation copies, see:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

WebSphere Developer Domain

IBM provides a valuable set of online technical developer resources via the WebSphere Developer Domain Web site. You can find this Web site, which often includes new technology previews and WebSphere developer experiences, at:

<http://www.software.ibm.com/wsdd/>

Linux-specific items available from the WebSphere Developer Domain include:

- **ThreadAnalyzer technology preview:** ThreadAnalyzer gathers and analyzes thread dumps from WebSphere Application Server, Version 5. It can obtain a thread dump or open an existing thread dump. ThreadAnalyzer is supported on Solaris, AIX, Windows, and Linux on Intel and 390. You can download the ThreadAnalyzer technology preview from the Web at:

http://www.software.ibm.com/wsdd/downloads/thread_analyzer.html

- **Developing Enterprise Java Applications Using DB2 Universal Database (UDB) Version 8:** This technical article summarizes enhancements in DB2 UDB Version 8 for Linux, UNIX, and Windows for Java programming. This includes integration with WebSphere, development tools, XML technology,

and information about the new JDBC Type 4 driver. You can find this article on the Web at:

<http://www.software.ibm.com/dmdd/library/techarticle/0209hutchison/0209hutchison.html>

- ▶ **WebSphere Commerce Suite Pro Edition for Linux: Linux Programmer's Guide:** This version of the Programmer's Guide includes new information about how to enable a store running in the WebSphere Test Environment to work with a remote Payment Manager. It also includes new information about the transaction scope of controller commands (in PDF format). You can download this guide from the Web at:

http://www.ibm.com/software/webservers/commerce/wcs_pro/ProgrammersGuide_zlinux.pdf

You can find an updated list of Linux-related items on the WebSphere Developer Domain at:

<http://www.software.ibm.com/webapp/dd/ViewServlet.wss?viewType=Library&devDomain=wsdd&sortBy=Posted&start=1&keyword=Linux&prodfam=0&topic=0&format=0&count=10>

Relevant Linux-related Redbooks

The current list of relevant WebSphere on Linux Redbooks and Redpapers includes:

- ▶ *Linux for WebSphere and DB2 Servers*, SG24-5850
- ▶ *Linux Web Hosting with WebSphere, DB2, and Domino*, SG24-6007
- ▶ *CICS Transaction Gateway V5 The WebSphere Connector for CICS*, SG24-6133
- ▶ *IBM WebSphere V4.0 Advanced Edition Handbook*, SG24-6176
- ▶ *Linux Application Development Using WebSphere Studio 5*, SG24-6431
- ▶ *WebSphere Application Server Test Environment Guide*, SG24-6817
- ▶ *IBM WebSphere Portal V4.1 Handbook Volume 1*, SG24-6883
- ▶ *IBM WebSphere V5.0 for Linux, Implementation and Deployment Guide - WebSphere Handbook Series*, REDP3601
- ▶ *WebSphere Commerce V5.4 for Linux, Infrastructure and Deployment Patterns*, REDP0214
- ▶ *Tivoli NetView® Performance Monitor for TCP/IP 1.4 Using Linux as the Web Server*, REDP3602
- ▶ *e-commerce Patterns for z/Linux Using WebSphere Commerce Suite V5.1 Patterns for e-business series*, REDP0411

- ▶ *WebSphere Commerce V5.4 for Linux, Infrastructure and Deployment Patterns*, REDP0214
- ▶ *Deploying Apache on IBM @server BladeCenter*, REDP3588
- ▶ *IBM WebSphere Portal V4.1.2 in a Linux Environment*, REDP0310
- ▶ *WebSphere Application Server V4 for Linux, Implementation and Deployment Guide*, REDP0405
- ▶ *Deploying Apache on IBM @server BladeCenter*, REDP3588

The list of Linux-related IBM Redbooks and Redpapers is continually being updated. You can find these documents on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

8.3 IBM Data Management products on Linux

IBM has the richest set of data management products on the market. Many of the components of the three major data management brands, that is DB2, Informix, and U2, are available on Linux.

8.3.1 DB2 Product Family

IBM's DB2 database software is the worldwide market share leader in the industry. It marks the next stage in the evolution of the relational database. DB2 is the industry's first multimedia, Web-ready relational database management system that delivers leading capabilities in reliability, performance, and scalability. It is the database of choice for customers and partners who are developing and deploying critical solutions.

You can find detailed product information for DB2 on Linux on the Web at:

<http://www.ibm.com/db2/linux>

The IBM DB2 product can add significant value to and benefit a customer business in the following ways:

- ▶ Lowest TCO
- ▶ Broad platform support
- ▶ Seamless Connectivity to Other Data
- ▶ World wide 24-x-7 service and support
- ▶ IBM and independent software vendor (ISV) application and tool enablement
- ▶ Reliability, Scalability, Availability (RAS) and high performance
- ▶ DB2 self management and smart database technology reduces the complexity

Specific value propositions for DB2 for zSeries

zSeries customers can extend and leverage their current investment in hardware. Hundreds or even thousands of Linux operating systems can coexist with other operating systems on a single zSeries under VM. With VM, a customer can offer a complete Linux server environment to each application developer and still run host production systems on the same zSeries server.

You can deploy new applications on Linux and still connect to their existing DB2 data using DB2 Connect. You can also use IBM DB2 Connect on Linux for zSeries to consolidate all gateway connectivity processing at the host. This eliminates the need for middle-tier hardware.

Customers can take advantage of the features specific to UNIX environments on DB2 UDB on Linux. With DB2 UDB and DB2 Connect on Linux for zSeries, it is possible for vendor applications that are currently available for UNIX customers to leverage the zSeries platform.

Specific value propositions for DB2 for xSeries (Intel-Architecture server)

xSeries customers can reduce their TCO. DB2 UDB on Linux for xSeries can run on low-cost hardware. Also, customers are not locked into a proprietary operating system. DB2 Enterprise Extended Edition (EEE) for Linux allows customers to add extra capacity as needed for Web serving and data warehousing at a lower price than other clustering solutions that are currently available.

Moreover, DB2 is the performance leader on Linux. DB2 is among the gtp 2 in a recent industry benchmark, TPC-H, from 22 November 2002. For more information, see:

http://www.tpc.org/tpch/results/tpch_perf_results.asp

DB2 product availability on Linux

The list of DB2-branded products on Linux includes:

- ▶ DB2 UDB (Workgroup, Extended Edition, Enterprise Server Edition, Personal)
- ▶ DB2 Express (special offering for small and medium businesses (SMB))
- ▶ DB2 Connect
- ▶ DB2 Everyplace
- ▶ DB2 Intelligent Miner™ Scoring
- ▶ DB2 Net Search Extender
- ▶ DB2 Relational Connect

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

DB2 UDB support for Linux

DB2 UDB for Linux runs on a variety of hardware platforms. In addition to the Intel 32-bit hardware, DB2 for Linux today runs on the zSeries server and is currently in beta testing for iSeries. IBM has announced support for 64-bit Intel Itanium2-based systems and has indicated its intention to support the PowerPC-based pSeries platform. IBM has also publicly demonstrated DB2 for Linux running on 64-bit AMD Opteron-powered systems.

On the 32-bit x86 platform, DB2 requires systems equipped with Intel Pentium, Xeon, or AMD Athlon or later processors. DB2 Version 7.2 is supported on any Linux distribution with 2.4 kernel, glibc 2.1.2, and libstdc++ 2.9.0 or later. With DB2 Version 8, the Linux Validation Program is introduced to ensure better inter-operability of DB2 on Linux platforms and to enhance customer experience.

IBM is working closely with the Linux community, Linux Distribution Partners, and independent software and hardware vendors to test DB2 using various Linux kernels, distributions, third-party products, hardware and other components that interact with DB2. As a part of the validation program, before a product is supported with DB2, it undergoes rigorous testing. This gives users confidence to run DB2 with other products right out of the box.

Contrary to the general Linux distribution support statement for IBM Software, DB2 UDB is supported for a large number of Linux distributions. The Linux distributions that are successfully validated for use with DB2 UDB include Red Hat Linux Professional 7.2, 7.3, and 8.0; Advanced Server 2.1; SuSE Professional 7.3, 8.0, and 8.1; SuSE Linux Enterprise Server 7 and 8; SCO Linux 4.0; Turbolinux Enterprise Server 7.0 and 8.0; and other distributions powered by UnitedLinux 1.0. See the following validation Web site for a complete and current table of the supported platforms and distributions:

<http://www.ibm.com/db2/linux/validate>

Note: To run Java applications and the graphical tools that come with DB2 (such as Control Center), JDK 1.3.1 is required. For convenience, it is shipped on the DB2 Version 8.1 CDs. In most cases, it is automatically installed with DB2 installation. Refer to the validation Web site (above) for the most current list of supported products.

8.3.2 Informix product family

The Informix data management product set is widely deployed on various distributed servers, such as UNIX and Windows. It provides the backbone of data management for many small to medium customers.

You can find detailed information about Informix products on the Web at:

<http://www.ibm.com/software/data/informix/>

Informix product availability on Linux

The Informix brand of products on Linux includes:

- ▶ Informix 4GL Compiler (plus Interactive Debugger)
- ▶ Informix 4GL Rapid Development System (plus Runtime)
- ▶ Informix C-ISAM
- ▶ Informix Client Software Development Kit (CSDK)
- ▶ Informix Connect Runtime
- ▶ Informix D4GL
- ▶ Informix Dynamic Server - IDS (Workgroup, Enterprise and Unlimited)
- ▶ Informix ESQL (Cobol, C)
- ▶ Informix Excalibur Text Search Datablade
- ▶ Informix Extended Parallel Server (XPS)
- ▶ Informix Geodetic Datablade
- ▶ Informix Image Foundation Datablade Module
- ▶ Informix Internet Foundation 2000 (IIF2K)
- ▶ Informix NET
- ▶ Informix OnLine
- ▶ Informix Red Brick™ Warehouse
- ▶ Informix SE Personal Edition
- ▶ Informix Server Administration
- ▶ Informix Spacial Datablade Module
- ▶ Informix TimeSeries Datablade
- ▶ Informix TimeSeries Realtime Loader
- ▶ Informix Web Datablade

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

Informix support for Linux

Linux is available from multiple distributors including Red Hat, SuSE, SCO, and Turbolinux. Informix moved from supporting each distributor's version, to

supporting the Linux component versions to allow the Linux community to define the distributors that support Informix products.

Regardless of the Linux distributor, a formal recertification of Informix product releases on the Linux operating system is only required when a major Linux upgrade is formally released. This is indicated by the second decimal point number of the kernel or glibc version changing. For example, a certification from kernel version 2.2.5 to 2.2.6 is not required. However, a formal certification from kernel 2.2.5 to 2.4.0 is required. A certification from glibc 2.0.7 to 2.0.8 is not required, but it is required when the version changes from 2.0.7 to 2.1.0.

The Informix product roadmap offers a complete and current table of the supported platforms and distributions. You can find it on the Web at:

http://www.ibm.com/software/data/informix/pubs/roadmaps/linux_rmap.html

8.3.3 U2 product family

The U2 data management family products are similar to those found in a Pick operating system. U2 offers high performance and scalable data management environments for embedding in vertical applications. Over 3.9 million users rely on IBM U2 Data Management solutions.

You can find detailed information about IBM U2 products on the Web at:

<http://www.ibm.com/software/data/u2/>

U2 product availability on Linux

The U2-branded products on Linux include:

- ▶ IBM Cobol Direct Connect
- ▶ MITS
- ▶ IBM RedBack
- ▶ IBM SystemBuilder Plus
- ▶ IBM Unidata
- ▶ IBM Universe

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

8.3.4 Competitive software

Despite the excellent functionality, performance, and price of IBM Data Management products, there are many open source and proprietary competitors.

Generally speaking, IBM Data Management products perform and scale better than competitive offerings. In addition, IBM provides industry leading support for DB2, Informix, and U2 products, which is not often available for open source products.

Alternatives to some IBM Data Management products include:

- ▶ **MySQL:** The evolving product of an open source project, MySQL is claimed to be the world's most popular SQL database. MySQL is included as part of *every* Linux distribution.

MySQL is developed on the following Web site:

<http://sourceforge.net/projects/mysql/>

It is supported and commercially available from MySQL AB (Sweden). You can learn more about MySQL on the Web at:

<http://www.mysql.com/>

- ▶ **PostgreSQL:** The evolving product of an open source project, PostgreSQL is claimed to be the world's most advanced open source database. PostgreSQL is included as part of many Linux distributions. It also powers the .ORG domain name registry.

PostgreSQL is developed at the following site:

<http://www.postgresql.com/>

- ▶ **Commercial competitors:** Oracle is investing heavily in Linux-enablement of its product set. Of particular interest is the potential to migrate Oracle database users from Sun servers to Linux on Intel using Oracle 9iRAC for horizontal scalability.

Sybase has ported their Adaptive Server Enterprise and SQLAnywhere products to Linux as well.

8.3.5 IBM Data Management resources

The following sections provide information about other resources that are available for data management.

DB2 UDB and DB2 Migration Kit: Evaluation copy

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration and support documentation, are available via the Web. You can download evaluation copies of DB2 UDB and the DB2 Migration Kit, as well as other Linux-enabled IBM Software, from the Web at:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

IBM UniVerse 10, Personal Edition

Users of the Personal Edition of IBM's Universe Nested Relational Database, Version 10 (UV10PE) can download the product from the following Web site:

<https://www.ibm.com/software/data/u2/support/u2techconnect/uv10pe.asp>

IBM DB2 Migration Toolkit for Linux

The IBM DB2 Migration Toolkit (MTK) V1.1 helps you migrate from Sybase and Microsoft SQL Server (Versions 6 and 7) to DB2 UDB Version 7.2 databases on any supported DB2 UDB workstation platform. With the Migration Toolkit, you can automatically convert data types, tables, columns, views, indexes, stored procedures, and triggers into equivalent DB2 database objects.

The Migration Toolkit provides database administrators (DBAs) and application programmers with the tools needed to automate previously inefficient and costly migration tasks. You can reduce downtime, eliminate human error, and cut back on person hours and other resources associated with traditional database migration by using the features in the Migration Toolkit.

You can download the IBM DB2 Migration Toolkit from IBM at:

https://www.software.ibm.com/d1/db2_03/migtkt-p

A beta version of the IBM DB2 Migration Toolkit for Oracle 8 is available for download from IBM at:

<http://www14.software.ibm.com/webapp/download/home.jsp?s=p>

Tech Guide: DB2 for Linux

This IBM sponsored white paper from Newmediary, Inc. describes DB2 products on Linux, the market for databases on Linux, and comparisons with open source alternatives. You can download it free (after registering) from the Web at:

http://www.software.ibm.com/dmdd/library/techarticle/0202applied/0202applied.html?isrc_d=99999999

Data Management Developers Domain

IBM provides a valuable set of online technical developer resources via Developer Domain Web sites. These Web sites, which often include new technology previews and developer experiences, are listed according to product family name.

You can find the DB2 Developer Domain on the Web at:

<http://www.boulder.ibm.com/dmdd/>

And you can find the Informix Developer Zone on the Web at:

<http://www.boulder.ibm.com/dmdd/zones/informix/>

Relevant Linux-related Redbooks

For more information on Linux, consult the following Redbooks and article:

- ▶ *e-Business Intelligence: Leveraging DB2 for Linux on S/390*, SG24-5687
- ▶ *e-Business Intelligence: Data Mart Solutions with DB2 for Linux on zSeries*, SG24-6294
- ▶ *Up and Running with DB2 for Linux*, SG24-6899
- ▶ *Web Click Stream Analysis using Linux Clusters*

<http://www.boulder.ibm.com/dmdd/library/techarticle/lurie/0111lurie.html>

For the latest Linux-related IBM Redbooks and Redpapers, see:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

8.4 IBM Tivoli products on Linux

Tivoli offerings consist of four feature solution sets:

- ▶ **Performance and availability management:** IBM Tivoli Performance and Availability Management solutions provide the central nervous system for your complicated e-business landscape. They constantly gather information about your hardware, software and network devices, and, in many cases, cure problems before they actually occur.
- ▶ **Configuration and operation management:** The IBM Tivoli Configuration and Operations solution delivers reliable services to your end users and makes them more productive. This suite of applications distributes software reliably and securely, manages the change and control of IT assets, automates workflow through the enterprise, and remotely controls systems and applications.
- ▶ **Security management:** IBM Tivoli security management solutions address two critical e-business challenges, automated identity management and security event management. The IBM Tivoli identity management solution helps you quickly realize return on investment (ROI) by bringing users, systems, and applications online fast, while effectively managing users, access rights, and privacy preferences throughout the identity life cycle. The IBM Tivoli security event management solution helps you actively monitor, correlate, and quickly respond to IT security incidents across your e-business.
- ▶ **Storage management:** IBM Tivoli Storage Management protects your organization's data from hardware failures and other errors by storing backup

and archive copies of data on offline storage. Scaling to protect thousands of computers running a dozen operating system platforms, its intelligent data movement and store techniques, and complete automation, reduce administration costs while increasing service levels. Tivoli storage products provide a combination of scalability, intelligent data technology, disaster preparation, and broad platform and application support through one centralized, automated solution.

Figure 8-2 illustrates the relationship of these solution sets.

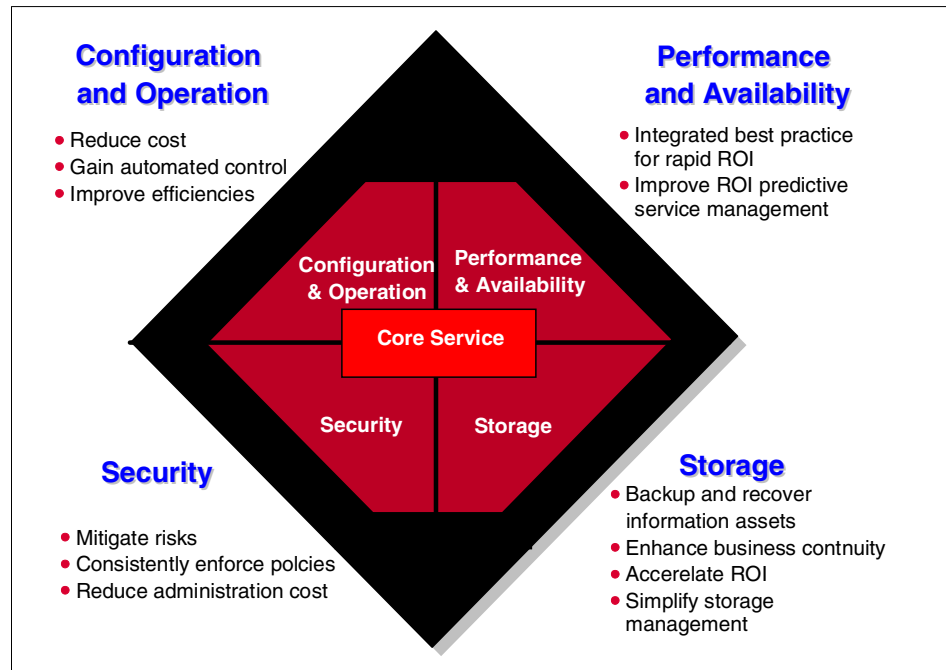


Figure 8-2 Tivoli IT management solutions

Why Tivoli for Linux

IBM Tivoli Software on Linux helps companies to embrace the future by allowing them maximum flexibility in their choice of new technologies and maximizing company ROI. Tivoli software on Linux provides companies' intelligent systems management capabilities and industry best practices. This helps to speed time to deployment, maximizing business efficiently and improving the bottom line.

In particular, Tivoli on Linux provides the following benefits:

- ▶ **Reduces the TCO:** Linux helps lower the cost of ownership. Tivoli lowers the cost of operations.

- ▶ **Enhances security:** Linux is attractive but has occasional security concerns. Tivoli helps address these concerns.
- ▶ **Meets service requirements:** Service-level expectations do not lessen with Linux. Tivoli enhances your ability to meet or exceed the service level requirements.
- ▶ **Offers advanced management:** As Linux moves from a Web serving and print serving solution to a critical business applications platform, Tivoli assists in managing not just the IT resources but the business applications as well.
- ▶ **Support heterogeneous platforms:** Linux may not be the only platform in your business. Tivoli offers consistent management and licensing across heterogeneous platforms.
- ▶ **Adheres to open standards:** Tivoli software from IBM supports the most open standards. This is an important consideration for Linux-based and other open source deployments. The end result is unprecedented interoperability between hardware and software components of your infrastructure.

For further information about Tivoli products on Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/tivoli.shtml&S_TACT=__S_TACT__&S_CMP=__S_CMP__

Tivoli Linux support

Tivoli's current strategy is to run under Red Hat, SuSE, and Turbo Linux running kernel 2.4 and server architectures of the xSeries and the zSeries. Tivoli has issued a statement of direction to provide products on the iSeries and pSeries servers in the future.

8.4.1 Tivoli product availability on Linux

The Tivoli brand of products on Linux includes:

- ▶ Tivoli Access Manager for e-business (formerly known as Policy Director/Tivoli Policy Director for Application Servers)
- ▶ Tivoli Access Manager for Operating Systems (formerly known as Policy Director for Operating Systems (function of Tivoli Security Manager))
- ▶ Tivoli Analyzer for Lotus Domino
- ▶ Tivoli Business Systems Manager
- ▶ Tivoli Configuration Manager
- ▶ Tivoli Decision Support for OS/390
- ▶ Tivoli Distributed Monitoring
- ▶ Tivoli Enterprise™ Console
- ▶ Tivoli Identity Manager
- ▶ Tivoli Inventory
- ▶ Tivoli Management Framework

- ▶ Tivoli Monitoring for Transaction Performance
- ▶ Tivoli Monitoring (formerly known as Tivoli Distributed Monitoring)
- ▶ Tivoli Monitoring for Databases
- ▶ Tivoli Monitoring for Messaging and Collaboration
- ▶ Tivoli Monitoring for Web Infrastructure (for Apache HTTP Server, IBM WebSphere Application Server, WebLogic Server)
- ▶ Tivoli NetView
- ▶ Tivoli NetView for z/OS (MSM Agents, NMC Server)
- ▶ Tivoli Point of Sale Manager
- ▶ Tivoli Remote Control
- ▶ Tivoli Risk Manager
- ▶ Tivoli SANergy™
- ▶ Tivoli SecureWay® Policy Director
- ▶ Tivoli Service Level Advisor
- ▶ Tivoli Software Distribution
- ▶ Tivoli Storage Manager
- ▶ Tivoli Storage Manager for Enterprise Resource Planning (formerly known as Data Protection for ESS for R/3)
- ▶ Tivoli Storage Manager for Mail (formerly known as Data Protection for Domino)
- ▶ Tivoli Storage Resource Manager
- ▶ Tivoli Switch Analyzer
- ▶ Tivoli System Automation for Linux
- ▶ Tivoli User Admin
- ▶ Tivoli Workload Scheduler

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

8.4.2 Competitive software

Generally speaking, Linux distribution systems management tools are primitive. They reflect the UNIX command line single utility model. There are virtually no open source products that compete against the complete Tivoli stack. However, Linux distributors have software distribution software that may compete with or, in some cases, complement the Tivoli Software Distribution product:

- ▶ **Linux Distribution Partner (LDP) software distribution**

With the recent move to Linux enterprise server distributions, Red Hat and SuSE (in particular) now employ a required product update mechanism controlled by their corporate servers (for example, Red Hat Network) or servers behind the customer firewall. Additionally, Ximian markets the Red Carpet distribution manager to provide software distribution to Linux systems.

► **Commercial products**

Many recognizable names in the systems management product space support their products on Linux. In particular, these include BMC, Computer Associates, HP, Veritas, and Mercury Interactive.

8.4.3 Tivoli resources

This section provides resources for Tivoli.

Evaluation Copy of Tivoli Access Manager

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration and support documentation, are available via the Web.

You can download an evaluation copy of the Tivoli Access Manager, as well as other Linux-enabled IBM Software, from the following Web site:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

Tivoli Developer Domain

IBM provides a valuable set of online technical developer resources via the Tivoli Developer Domain Web site. This Web site, which often includes new technology previews as well as developer experiences, is located at:

<http://www06.ibm.com/developerworks/tivoli/>

Relevant Linux-related Redbooks

The current list of relevant Tivoli on Linux Redbooks includes:

- *IBM Tivoli Monitoring Version 5.1: Advanced Resource Monitoring*, SG24-5519
- *A Practical Guide to Tivoli SANergy*, SG24-6146
- *Linux System Administration and Backup Tools for IBM @server xSeries and Netfinity*, SG24-6228
- *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844
- *Tivoli NetView Performance Monitor for TCP/IP 1.4 Using Linux as the Web Server*, REDP3602

The list of Linux-related IBM Redbooks and Redpapers is continually being updated. You can find these documents on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

8.5 Lotus products on Linux

Lotus software sets the standard for truly innovative software and services that reflect the brand's unique representation of the new ways individuals and businesses work together to achieve success. Lotus software is further redefining the concept of conducting business through practical knowledge management, e-business, and other groundbreaking ways of connecting ideas, thinkers, buyers, sellers, and communities around the world via the Internet. Lotus software is marketed in more than 80 countries or regions worldwide through direct sales and extensive business partner channels.

For detailed information about Lotus products, see:

<http://www.lotus.com>

8.5.1 Lotus product availability on Linux

The following Lotus products are available on Linux:

- ▶ Lotus Domino
- ▶ Lotus Domino Server
- ▶ Lotus Workflow™
- ▶ Lotus Learning Management System

For the most current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

Lotus Domino server

The Domino family provides a multiplatform foundation for collaboration and e-business, driving solutions from corporate messaging to Web-based transactions and everything in between. Domino is the leading messaging and collaborative application platform available for Linux today. Domino for Linux provides browser-based access for collaborative applications and messaging. POP3, HTTP, IMAP and LDAP access are all available.

Fully functional Domino solutions are available now on Linux. Key solutions of Domino for Linux involve:

- ▶ Sales Force Automation
- ▶ Project Management
- ▶ Customer Relationship Management
- ▶ Interactive Web based forms
- ▶ Resource Management

Why Domino for Linux

When you invest in Domino for Linux, you gain:

- ▶ Stability of the operating system
- ▶ Lower cost of ownership
 - Reliability
 - Availability
 - Manageability
- ▶ Powerful remote management
- ▶ An alternative to Microsoft Exchange servers on Windows operating systems
- ▶ High-level of interest in migrating from a Domino on Windows NT environment to Domino on Linux environment

Note: SuSE sells a product called the *SuSE Linux Groupware Server* with Lotus Domino for Intel and AMD systems. It comes prepackaged with Domino R5. You can learn more about this offering on the Web at:

<http://www.suse.de/en/index.html>

Other Lotus products

Lotus is planning to use the WebSphere portal-based delivery platforms for several Lotus products that are not currently available for Linux. Contact your local Lotus sales representative for additional information.

Lotus has announced plans to deliver the iNotes™ Web Access product for browsers other than Microsoft's Internet Explorer. These browsers are popular with Linux users.

There are no plans to port the Lotus SmartSuite® desktop productivity products to Linux.

8.5.2 Competitive software

In the area of e-mail products, there are two primary open source competitors and two commercial competitors:

▶ Sendmail

The evolving product of an open source project, Sendmail has been the most widely used mail transport agent in the UNIX and Linux market. It is included as part of many Linux distributions.

A commercially enhanced and supported Sendmail product is available from Sendmail Corporation. You can learn more about Sendmail on the Web at:

<http://www.sendmail.com/>

► **Postfix**

Postfix is rapidly gaining favor in the Linux market. It is an evolving product of an open source project. Postfix is included as part of some Linux distributions. In fact, SuSE resells it as part of a packaged mail server product. You can find more details on Postfix on the Web at:

<http://www.postfix.org/>

► **Bynari InsightServer**

Bynari provides a Linux (and Solaris) functional replacement for Microsoft Exchange. To learn more about Bynari and its products, see:

<http://www.bynari.net>

► **Ximian Evolution**

Unlike the other three products in this list, each of which was an e-mail server, Evolution is a functional replacement (including look and feel) for Microsoft Outlook client software. This product is free and is now shipped as part of many Linux distributions. You can use it as a client for several e-mail servers. However, using it with an existing Microsoft Exchange server requires the purchase of adapter software from Ximian. To learn more, see:

<http://www.ximian.com/>

8.5.3 Lotus resources

This section lists various resources for Lotus.

Evaluation copy of Lotus Domino

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration and support documentation, are available via the Web. You can download an evaluation copy of Lotus Domino, and other Linux-enabled IBM Software, from the following Web site:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

Lotus Developer Domain

IBM provides a valuable set of online technical developer resources via the Lotus Developer Domain Web site. This Web site, which often includes new technology previews as well as Lotus developer experiences, is located at:

<http://www.lotus.com/1dd>

Relevant Linux-related Redbooks

Relevant Lotus on Linux Redbooks and Redpapers include:

- ▶ *Lotus Domino R5 for Linux on IBM Netfinity Servers*, SG24-5968
- ▶ *Lotus Domino 6 for Linux*, SG24-6835
- ▶ *Online Briefing: Mastering Domino for Linux*, REDP0039

You can find a complete list of Linux-related IBM Redbooks and Redpapers on the Web at:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

8.6 Additional IBM Software

In addition to primary brand software, IBM has made several of its other software products available on Linux.

8.6.1 Product availability on Linux

Non-specific brand IBM Software for Linux includes:

- ▶ AFS® (Andrew File System)
- ▶ Application Workload Modeler for Linux
- ▶ IBM Bank Teller Business Components
- ▶ IBM BookManager® Bookserver
- ▶ IBM Communications Server for Linux
- ▶ IBM Cluster Systems Management (CSM) for Linux
- ▶ IBM Developer Kit for Linux
- ▶ IBM Director
- ▶ IBM Directory Integrator
- ▶ IBM Directory Server
- ▶ IBM Everyplace Wireless Gatekeeper
- ▶ IBM General Parallel File System
- ▶ IBM HTTP Server
- ▶ iSeries ODBC Driver for Linux
- ▶ IBM LANDP®
- ▶ IBM Object REXX
- ▶ IBM Screen Customizer
- ▶ IBM Small Business Suite
- ▶ StreamServer Business Communications Platform
- ▶ IBM VisualAge Generator Server
- ▶ IBM VisualAge Pacbase Linux Server
- ▶ IBM VisualAge Smalltalk Enterprise
- ▶ IBM Workstation APL2® for Multiplatforms

For current availability and detailed descriptions of IBM Software products for Linux, see:

http://www.ibm.com/cgi-bin/software/track3.cgi?file=/software/is/mp/linux/software/pdfs/IBMSoftwareOnLinux.pdf&S_TACT=__S_TACT__&S_CMP=__S_CMP__

8.6.2 Additional resources

Consult the following resources for more information.

Evaluation copy of IBM Directory Server

Free evaluation copies of certain Linux-enabled IBM middleware products, complete with installation, configuration, and support documentation, are available via the Web. You can download an evaluation copy of IBM Directory Server and other Linux-enabled IBM Software from the Web at:

<http://www.ibm.com/developerworks/offers/linux-speed-start>

Redbooks

For more information, consult the following IBM Redbooks:

- ▶ *Linux Clustering with CSM and GPFS*, SG24-6601
- ▶ *Building a Linux HPC Cluster with xCAT*, SG24-6623

8.7 IBM Integrated Platform for e-business

Integrated Platform for e-business provides intensive testing, one-stop shopping, and a single point of contact for support for IBM's core e-business products and SuSE's Linux Enterprise Server. This combination allows you to:

- ▶ Deploy advanced e-business applications, capitalizing on the latest industry standards like Java, WebSphere, Linux, and IBM @server technology.
- ▶ Integrate flexible solutions with traditional workloads and mission-critical databases.
- ▶ Save time and resources to design the infrastructure. It's already created, configured, and tested.
- ▶ Tailor the system to customer needs using the breadth of IBM knowledge and experience in e-business.
- ▶ Position customer business for the future IBM @server with the latest advances while leveraging existing technology.

- ▶ Keep TCO low by managing a customer's business centrally and exploiting the economies of server consolidation.
- ▶ Benefit from a total solution for e-business integration.

The Integrated Platform for e-business is currently available for xSeries and zSeries hardware platforms. The following software and support are included in the offering:

- ▶ WebSphere Application Server Advanced Edition
- ▶ IBM Directory Server (LDAP) V4.1
- ▶ One year of software support and IBM Global Services (IGS) Linux line support

You can find detailed information about the Integrated Platform for e-business on the xSeries server on the Web at:

<http://www.ibm.com/servers/solutions/linux/integrated/components.html>

For detailed information about the Integrated Platform for e-business on zSeries, see:

http://www.ibm.com/servers/eserver/zseries/os/linux/integrated/index.html?ca=zseries&me=W&met=drive_from_eserver_IP

8.8 IBM Software on Linux deployment examples

The IBM Solution Scenario Profile Database (SSPD) is a unique reuse mechanism that contains synopses of both solutions tested by IBM development labs and solutions created for customers. Each entry describes the problem being solved and how the pieces fit together. It also includes a diagram, list of products used, and supporting material.

Use the fully searchable database to find which levels of products were tested together. You can also find details on specific solutions from which you can base your solution.

The following sections offer a brief look into examples taken from the SSPD. For detailed information about the Solution Scenario Profile Database, see:

<http://sspd.ibm.com/>

8.8.1 Server consolidation using WebSphere Application Server on the zSeries server

A customer needed a platform-independent Revenue Accounting System (RAS) application to provide general access from the Internet and employee access

from the intranet and Internet. They also needed to modernize their internal-only application, the Flight Schedule Inquiry System, to provide employees both intranet and Internet access to improve their ability to check work schedules from anywhere in the world.

The solution was in using IBM WebSphere Application Server, Advanced Edition V3.5; WebSphere Studio, Advanced Edition; VisualAge for Java; and Linux for zSeries. This allowed the customer to quickly update business applications, consolidate workloads onto IBM mainframes from a variety of servers, and realize savings and increased availability and reliability.

This solution employed the following IBM products:

- ▶ Hardware:
 - zSeries 900
 - zSeries running Linux
- ▶ Software:
 - WebSphere Application Server, Advanced Edition
 - VisualAge for Java

For a full description and diagrams of this scenario, see:

<http://w3-1.ibm.com/support/techdocs/atsmastr.nsf/PubA11Num/TD100650>

8.8.2 Computer services using Domino and WebSphere on xSeries and iSeries

A customer who offered hosting services was experiencing problems with server performance. Since they needed to handle many operating systems to support their customers' needs, they found it difficult with their Windows 2000 Domino systems.

To achieve their goals for upgrading their environment, this customer invested in several xSeries and iSeries servers and Linux to support their internal and hosting applications. Linux allowed this customer to work with their clients who wanted to test and run production middleware, as well as maximize cost efficiency. Overall, this solution allowed the customer to reduce costs, increase performance, and offer a better and more affordable service.

This solution employed the following IBM products:

- ▶ Hardware
 - xSeries 200 and 220
 - iSeries 270

- ▶ Software
 - Lotus Domino
 - Apache
 - Lotus Sametime®
 - QuickPlace
 - DB2 Version 7.1
 - WebSphere Application Server, Advanced Edition
 - Lotus LearningSpace® Core
 - Firewall
 - Logical partition (LPAR)

For a full description and diagrams of this scenario, see:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/TD100883>

8.8.3 DB2 UDB, DB2 Connect, and WebSphere Portal

This company developed and deployed its first Web implementation on non-IBM Windows NT hardware. This included two applications for member stores to reorder items via the Web. The company planned to develop many more applications, but was concerned with the ability of its Windows NT environment to support them.

The customer turned to IBM for help. They converted their existing Windows NT applications to zSeries servers running SuSE Linux V7. They consolidated their two applications from 40 dual-processor Windows NT machines for development and production to run on two zSeries LPARs. The software environment now uses WebSphere Application Server, Advanced Edition, Version 3.5 on Linux to run custom applications for their e-commerce needs. Developers use WebSphere Commerce Studio V5.1, which includes VisualAge for Java V3.5, to develop the new applications and deploy them to Linux for zSeries. And, the company's users migrated from Microsoft and Lotus cc:Mail to Lotus Notes.

The customer used WebSphere Commerce Suite and WebSphere Application Server (both on Linux for zSeries) to implement an initiative for a robust portal for independent grocers, distributors, and retail stores. This plan allows these groups to access important business-critical information through the portal. This includes ordering products, finding current pricing, and conducting forms management, all within the portal environment.

This solution employed the following IBM products:

- ▶ Hardware
 - zSeries

- ▶ Software
 - DB2, DB2 UDB, and DB2 Connect
 - IBM HTTP Server
 - Lotus Domino, Lotus iNotes, Lotus QuickPlace, and Lotus Sametime
 - VisualAge for Java
 - WebSphere Application Server
 - WebSphere Application Server, Advanced Edition
 - WebSphere Application Server, Advanced Edition Version 3.5
 - WebSphere Commerce Studio
 - WebSphere Commerce Suite
 - WebSphere Commerce Studio Version 5.1

For a full description and diagrams of this scenario, see:

<http://w3.ibm.com/support/techdocs/atsmastr.nsf/PubAllNum/TD100764>

8.9 IBM and Linux solution customer success

Several customers have experienced great success in implementing an IBM and Linux solution for their diverse needs. Here are a few of those customers and links to where you can learn about their stories:

- ▶ **WebSphere: Diveo Broadband Networks, Inc.**

Diveo Broadband Networks transitioned to a WebSphere for Linux architected solution. In doing so, they achieved increased affordability and reliability, increased availability, and savings on maintenance costs.

For details and more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-4Z8VXP?OpenDocument&Site=linuxatibm>

- ▶ **DB2: Grede Foundries, Inc.**

Grede Foundries transitioned to a DB2 for Linux architected solution. This brought them improved reliability, scalability, performance, stability, and cost savings.

For details and more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5BFTMG?OpenDocument&Site=linuxatibm>

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-4XHRGT>

► **Tivoli: Tamkang University**

Tamkang University deployed a Linux-based solution. This solution automated backup and recovery and optimized storage allocation. It also gave them improved reliability, capacity, and performance.

For details and more information about this success story, see:

<http://www.ibm.com/software/success/cssdb.nsf/CS/NAV0-5942EY?OpenDocument&Site=linuxatibm>

8.10 Additional IBM Software for Linux resources

IBM provides its employees, partners, and customers with an array of resources designed to assist and improve the deployment of Linux solutions. This section highlights some of those resources.

8.10.1 Evaluation and demonstration software

To get started with IBM Software Group products, the following sections provide information about tools that are available for evaluation or download.

IBM Speed Start Your Linux Application Program

In an effort to provide hands-on experience with IBM middleware on Linux, IBM provides customer evaluation copies of several of its major software products for Intel-based Linux systems.

Among the products currently available are:

- IBM DB2 Universal Database Enterprise Server Edition (ESE) V8.1 for Linux
- IBM DB2 Migration Toolkit V1.1 for Linux
- IBM Directory Server V5.1 for Linux
- IBM Lotus Domino 6 for Linux (English)
- IBM Tivoli Access Manager Base 4.1 for Linux
- IBM WebSphere Application Server V5 for Linux
- IBM WebSphere Studio Site Developer V5 for Linux
- IBM WebSphere MQ V5.3.1 for Linux
- IBM WebSphere SDK for Web Services for Linux

You can find a description of these products and links to freely available evaluation software and installation documentation on the Web at:

<http://www.ibm.com/developerworks/offers/linux-speed-start/>

Free technical support is available for these applications via a Web forum at:

http://www.ibm.com/developerworks/forums/dw_forum.jsp?forum=160&cat=5

The Linux Software Evaluation Kit

If downloading evaluation software from the Speed Start Your Linux Application Web site is too cumbersome, you can obtain the software on CD. It actually comes as a set of four CDs titled the Linux Software Evaluation Kit (SEK).

The Linux SEK is limited to Intel-based Linux implementations. It is free of charge and includes limited Web forum technical support.

For detailed information about the Linux Software Evaluation Kit, see:

<http://www.ibm.com/developerworks/offers/linux-speed-start/?l=sslaitlob>

Additional Linux-enabled software for free download

IBM provides a large number of trial, beta, and free software products for Linux that are downloadable from the Web. You can find a complete list of this IBM Linux-enabled software, which is available for download free of charge, at:

http://www.software.ibm.com/webapp/download/search.jsp?q=Linux&k=any&cat=&sb=n&go=y&sr=1&rs=&S_TACT=&S_CMP=&s=&pf=&dt=&x=7&y=9

8.10.2 IBM Software on Linux training

IBM Learning Services offers an extensive catalog of world class training courses. Its portfolio of training services includes on-site training, consulting, development, delivery, transformation, learning management, and outsourcing.

The courses about IBM Software on Linux include:

- ▶ **DB2 for Linux:** CF201 - *DB2 Admin Workshop for Linux*
http://www.ibm.com/services/learning/MainServlet.wss?pageType=course_description&courseCode=CF201
- ▶ **WebSphere on Linux:** SW247 - *WebSphere System Administration on Linux*
<http://www.software.ibm.com/wsdd/education/enablement/curriculum/sw247.html>

Additional courses are being prepared. For an updated list of Linux courses offered in the United States by IBM Learning Services, see:

<http://www.ibm.com/services/learning/us/catalog/linux/>

8.10.3 IBM Redbooks

IBM Redbooks are written and published by the IBM International Technical Support Organization (ITSO). The ITSO develops and delivers skills, technical know-how, and materials to technical professionals, including IBM, Business Partners, Customers, and the marketplace. The ITSO's value-add information

products address product, platform, and solution perspectives. They explore integration, implementation, and operation of realistic customer scenarios.

The ITSO partners with IBM Divisions, Lotus Development and Tivoli Systems, in the process of IBM Redbook development. The ITSO is part of the IBM Global Technical Support organization within IBM Global Sales and Distribution.

IBM Redbooks, RedDrafts, and Redpapers are available to the public for printed copy purchase or freely downloadable as PDF files. For Linux-related Redbooks, see:

<http://publib-b.boulder.ibm.com/cgi-bin/searchsite.cgi?query=linux>

8.11 Internal IBM Software for Linux resources

The following sections provide information on Linux resources that are available to IBM Marketing and Technical Support. We provide them for the IBM audience. This information is very useful when discussing or implementing Linux solutions with IBM Business Partners and Customers.

Important: This section highlights various resources that are *internal* to *IBM only*. Many of the Web links that are provided are to IBM intranet sites.

8.11.1 Linux sales resources

This section relates to resources that are available for IBM Sales Representatives.

IBM Linux Cyber Sales Kit

The IBM Linux Cyber Kit is produced by the Linux Integration Center and SWG Linux Marketing. The purpose of this kit is to provide Brand Specialists and Software and IGS IT Architects and Specialists with technical Linux pre-sales support information. This sales kit is updated as new resources become available.

Currently this cyber sales kit is available in the Americas. You can find it on the Web at:

<http://lic.austin.ibm.com/LinuxCyberKit.HTM>

IBM Software Group Competitive Project Office

IBM's Software Group maintains a competitive project office that provides white papers, articles, and presentations regarding Software Group competitive products. You can find this on the Web at:

<http://elib.austin.ibm.com/cpweb.nsf/pages/cpo-home>

TCO tools

The primary field sales tool for calculating overall TCO of IBM products (in some cases with comparisons to competitive products) is the CIOView TCONow! tool, which is licensed to IBM.

An example of the TCO analysis that can be performed using TCONow! is DB2 versus Oracle. You can find TCONow! on the Web at:

<http://d25web1.torolab.ibm.com/db2info/dminflib.nsf/dm/tconow>

Additional TCO tools include the TCO analysis of WebSphere versus BEA WebLogic and Oracle AppServer. You can find these tools on the Web at:

<http://w3-3.ibm.com/software/websphere/ticompet.nsf/cb2d4c1e6e2eec0f85256a6c007178f6/5b3f2a7345a9e1538725699d0051db1f?OpenDocument>

IBM Software Group Linux Playbook for 2003

You can find the IBM Software Group Linux Sales Playbook on the Web at:

<http://w3-3.ibm.com/software/sales/salesite.nsf/printthis?OpenForm&ParentUNID=77FE28C087E3D37987256C240049BF76&Seq=1&TableRow=1.1>

Pre-Sale Solution Assurance Review (SAR)

Certain IBM products and solution offerings require a pre-sale SAR prior to proposal to a customer. Currently, WebSphere Advanced Server Version 3.5 or 4.0, when proposed on zSeries Linux, requires such a review. You can find detailed information about this solution review on the Web at:

<http://w3-1.ibm.com/support/assure/assur30i.nsf/WebIndex/SA470>

For all IBM Software products that require a SAR, regardless of operating system, see:

<http://w3-1.ibm.com/support/assure/assur30i.nsf/WebIndex/SA331>

Request for availability of IBM Software on Linux

IBM employees can generate a request for IBM Software that is not currently supported on Linux. To learn about this process, see:

<http://w3-3.ibm.com/software/sales/salesite.nsf/salestools/9A910D3F3A18276C87256BB3005ADCED?opendocument&highlight=0,software+requirements>

The Solution Scenario Profile Database

The SSPD provides samples of “reusable” solutions. For the top ten reusable Linux solutions, including several that are based on IBM Software, see:

http://swtest.ibm.com/s_dir/swtest/swtest.nsf/0eb60ad63cda042e85256c310050aef6/38858db3556c66a885256c7500687247?OpenDocument

8.11.2 Product information

This section offers links to sources of IBM Software product information on the IBM intranet.

Availability of IBM Software on Linux

The internal site for IBM’s Software for Linux availability matrix is maintained by the Linux Integration Center. You can locate this matrix on the Web at:

<http://lic.austin.ibm.com/Whitepapers/linuxmatrix.html>

Availability of IBM Software for Linux and national language support (NLS)

You can find the IBM Software for Linux using double-byte character set (DBCS) implementation for NLS on the Web at:

http://lic.austin.ibm.com/Whitepapers/DBCSNLS_Linux.html

IBM SWG Products, Solutions and Strategies Reference Guide

The *SWG Products, Solutions and Strategies Reference Guide*, G325-2132, is an invaluable resource for anyone faced with a cross-selling opportunity who needs to discuss the IBM Software portfolio. The guide includes sections on each major brand including DB2, Lotus, Tivoli, and WebSphere. It also includes details on primary SWG solutions. For most offerings, the guide provides an overview, top reasons businesses choose to buy, key messages, customer references, and a competitive overview.

This is a great resource for IBMers who need to gain a quick understanding of the products, solutions and strategies being developed and promoted within the IBM Software Group. You can find this guide on the Web at:

<http://w3.ibm.com/software/sales/salesite.nsf/swgsites/SWG+Reference+Guide>

WebSphere Resource Finder

You can find WebSphere-related books, demos, white papers, ISVs, brochures, etc. by using this search facility. You can locate it on the Web at:

http://9.242.192.151:8080/examples/servlet/typcat?typ=*&cat=*&pag=1

DB2 for Linux Sales Kit

This site includes DB2 on Linux value propositions, positioning, benchmark results, and references:

<http://d25web1.torolab.ibm.com/db2info/DMINFLIB.NSF/f3bb8c878318bf7a852567b600741dde/bb08076da9ab4cd885256a4c003edf89?OpenDocument>

8.11.3 Technical resources

There is a wealth of Linux technical resources within IBM. This section provides access points to these resources.

Linux Technical Support Resources

IBM employees can find the most complete list of Linux technical support resources at:

<http://w3.linux.ibm.com/linux/sales/technical.shtml>

IBM employees can find a list of Linux sales and technical contacts at the following locations:

► Americas:

<http://w3.linux.ibm.com/linux/sales/files/linuxwhoswhoamericas.lwp>

► EMEA:

<http://w3.linux.ibm.com/linux/sales/files/EMEA-LinuxWhosWho04-2002.lwp>

► AP:

http://w3.linux.ibm.com/linux/sales/files/w3_linux_sales_sup_ap_who_Mar2002.lwp

For an organization chart of IBM's Software Technical Sales management team for the Americas (not specific to Linux), see:

ftp://w3.ibm.com/support/gtss/ftss_by_brand.pdf

Worldwide Linux Advocates

IBM has developed the Linux Advocates program as a complement to the Linux Integration Center. If a Linux sales opportunity does not meet the requirements to be a candidate for Linux Integration Center (LIC) assistance, a Linux advocate should be considered for engagement.

A Linux Advocate is an IT Architect or IT Specialist who is level 4 skilled in IBM Software brand products (DB2, WebSphere, Lotus, or Tivoli) on Linux.

The Linux Advocate is a valuable resource by:

- ▶ Assisting the technical and sales team with Linux-based IBM Software product technical questions and issues
- ▶ Influencing the technical/sales community so that more Linux-based solutions are recommended to our customers
- ▶ Providing Linux leadership on their assigned accounts
- ▶ Being up to date on the evolution of Linux, especially as it relates to the deployment of IBM Software for customer solutions

Linux advocates have been assigned in every worldwide geographic region. For an updated list of Linux advocates, refer to the Brand Focal Points in the Linux Cyber Sales Kit at:

<http://lic.austin.ibm.com/LinuxCyberKit.HTM#Focal>

SWG presales technical support Web sites

The DB2 and Tivoli brands maintain a technical support Web site. While these sites are not specific to Linux, they have links to various Linux resources:

- ▶ DB2 Presales Technical Support

<http://w3.ibm.com/software/sales/salesite.nsf/printthis?OpenForm&ParentUNID=77FE28C087E3D37987256C240049BF76&Seq=1&TableRow=1.1>

- ▶ Tivoli Presales Technical Support

[http://w3.ibm.com/software/sales/saletool.nsf/salestools/Tivoli+software\\$presales_support](http://w3.ibm.com/software/sales/saletool.nsf/salestools/Tivoli+software$presales_support)

Technical Sales Support Guides

Technical Sales Support Guides have been produced for each sector and geography. While not specific to Linux, these online documents describe the processes and contacts available to field sales personnel who require technical support for software sales.

You can find the Americas Technical Sales Support Guides on the Web at:

<http://w3-1.ibm.com/support/americas/techsuptguide.html>

Linux Integration Center ‘How To’ white papers

As a result of their customer engagement experiences, the Linux Integration Center frequently creates excellent “how to” white papers that summarize the best practices of their engagements. Here is a list of the current collection:

- ▶ *Installing and Configuring of WebSphere Application Server Advanced Edition Version 4.0.2 with Oracle9i Database release 9.0.1 on Linux (Intel architecture)*

This paper describes the basic concepts and steps to install and configure WebSphere Application Server Advanced Edition 4.0.2 with Oracle9i Database, release 9.0.1 (Oracle9i) on Linux Intel architecture. The topics covered in this paper include hardware and software requirements, pre-installation configuration of the Linux system, installing Oracle9i, post-installation configuration of the system and the Oracle9i Database, installation of WebSphere Application Server Advanced Edition Version, testing the installation, and tips and techniques.

You can locate this paper on the Web at:

<http://lic.austin.ibm.com/Whitepapers/WAS402withOracle9i.pdf>

► *Beginners Guide to Installing and Testing WebSphereMQ on Linux*

This paper introduces WebSphereMQ, what it is, how it works, and how to install the basic WebSphereMQ components on a Linux on Intel/xSeries system. Its purpose is a quick start guide for people who are not familiar with WebSphereMQ.

This paper is located on the Web at:

<http://lic.austin.ibm.com/Whitepapers/WebSphereMQ.pdf>

► *Performance Tuning Linux on a Lotus Domino 5.0.6 Server*

This white paper documents the procedure for performance tuning Linux on a Lotus Domino xSeries server. It is important to note that the server hardware configuration described in this paper is designed to achieve optimal performance growth. It does not represent the outcome of an average production server.

<http://lic.austin.ibm.com/Whitepapers/DomPerf.pdf>

► *Tips and Techniques for Running WebSphere Commerce Suite Version 5.1 on Linux/390*

This paper is a supplement to the installation guide provided for the WebSphere Commerce Suite Pro Edition for Linux for zSeries and S/390 Version 5.1. This paper covers such topics as booting the Linux image on the VM environment, installing and configuring WebSphere Commerce Suite for Linux/390, creating a WebSphere Commerce Suite instance, and creating databases on the OS/390 DB2 subsystem.

You can find this paper on the Web at:

<http://lic.austin.ibm.com/Whitepapers/WCStips.pdf>

For an updated list of Linux Integration Center “how to” white papers, see:

http://lic.austin.ibm.com/library_docs/howto.html

iSeries sizing guide for the Apache Web server

This site provides an in-depth description of sizing an iSeries server to host an Apache Web server. Since IBM's WebSphere HTTPD server is essentially the Apache Web server, this document should be useful to sizing for WebSphere Web serving as well. You can find this information on the Web at:

<http://ca-web.rchland.ibm.com/perform/linux/V5R1LinuxSambaSizing.pdf>

8.11.4 Evaluation and demonstration software

In addition to the external evaluation and demo software distribution points, there are internal sources. The following sections list some of these sources.

Linux-based e-business demos using VMware

Several Linux-based e-business in Action demonstrations using IBM middleware have been developed. These live-action demonstrations are easy to install and use. These demonstrations have the following characteristics:

- ▶ They are complete, stand-alone demo systems. They are complete with Linux operating system, preconfigured IBM middleware, demo applications, and data.
- ▶ They are packaged as VMware Workstation system images.
- ▶ They run in virtual machines as applications on your Windows system.
- ▶ They are well documented, including installation video, administration guide, scripted scenarios, and reference information.
- ▶ They are live working demos, which makes them easily extensible and flexible to match customer demonstration requirements.

The current collection of eBiA demos includes:

- ▶ **IBM WebSphere Commerce:** A business to consumer (B2C) e-commerce site for "JK Sports"
- ▶ **IBM WebSphere Portal:** A business to employee portal scenario for "JK Sports", illustrating portal integration
- ▶ **Lotus Workflow, Domino, DECS, and DB2:** A managed workflow demonstration illustrating "JK Enterprises" purchase order submit/approval scenario

You can access e-business in Action demos on the Web at:

<http://edelivery.demopkg.ibm.com>

DEMOpkg

DEMOpkg provides easy access to software and demonstrations. It covers all IBM Software brands. The DEMOpkg offerings are packaged in a compact format for easy and quick installation and access. Best of all, DEMOpkgs are available at no charge for IBM Internal Use and customer demonstration purposes.

IBM employees can find detailed DEMOpkg information on the Web at:

<http://w3.demopkg.ibm.com>

You can find Linux-specific demonstration software, on a per hardware basis, using the links on this page.

Obtaining WebSphere software for evaluation emergencies

IBM representatives can obtain Not-for-Resale (NRF) WebSphere Product Software for their customers by using the NFR order form, but only for *emergency* situations. Use the Internal Orders Database for *regular* situations.

For more information about placing an NRF order, see:

<http://w3-3.ibm.com/software/websphere/websites.nsf/FD6FCAF39CEBD60A85256A3F006F4FF1/BC79CDF0F174D28487256CB100524005?OpenDocument>

8.11.5 IBM Software support organizations

IBM has established a core set of organizations to proof and improve customer solutions based on IBM Software, as well as contribute to the evolution of IBM's middleware products.

IBM Linux Integration Centers

The LICs support and propagate the use of IBM middleware products for Linux solutions. These presales support centers, located in Austin (Texas), Boeblingen (Germany), and Singapore, provide assistance to a variety of industry customers, including financial services, retail and manufacturing, to identify key areas within a customer's infrastructure to implement Linux. The LICs provide IBM hardware, software, and services. They can also simulate computing environments to provide detailed consultations, architectural guidance, and educational briefings surrounding Linux and open source software.

Skilled in a broad spectrum of operating systems, LIC team members can quickly prepare an assessment of a customer's infrastructure and develop a step-by-step blueprint for transitioning to a Linux environment on IBM Software platforms. The team members include system architects, database administrators, project managers, and operating system specialists.

The LIC teams provide the following valuable services:

- ▶ Assist installation, configuration, and product integration during Proof of Concepts
- ▶ Guide architectural designs by incorporating best practices as well as using the experience gained from prior Linux migrations
- ▶ Find solutions to customer problems through replicating the customer's environment and debugging problems
- ▶ Interface with the owner of each product to assess requirements and resolve issues
- ▶ Provide engagement-based support for qualified Linux Middleware opportunities by offering:
 - Technical consulting
 - Product integration
 - Deployment planning
- ▶ Drive Linux as an e-business application server
- ▶ Integrate middleware and applications (ISV or customer)
- ▶ Document migration paths
- ▶ Identify supported hardware and software stack environments
- ▶ Gather performance benchmarks
- ▶ Demonstrate data conversion
- ▶ Create and leverage reusable intellectual assets

To qualify as a candidate for assistance from a Linux Integration Center, a sales opportunity must:

- ▶ Have significant SWG middleware content
- ▶ Allow the LIC to build collateral or assets for reuse
- ▶ Have the potential for a customer to become a referenceable account
- ▶ Support current IBM or SWG strategic initiatives
- ▶ Be an xSeries or zSeries opportunity with significant SWG middleware
- ▶ Include migration from SUN to Linux and IBM middleware
- ▶ Include migration from Oracle to DB2
- ▶ Include migration from BEA WebLogic to IBM WebSphere
- ▶ Have complete OMNotes Entry identified with LINX opportunity code

For detailed information about the Americas IBM Linux Integration Center, see:

<http://lic.austin.ibm.com/americas>

The IBM internal Web site for the Linux Integration Center is located at:

<http://w3.ibm.com/linux/lic>

For information about the EMEA IBM Linux Integration Center, see:

<http://lic.austin.ibm.com/emea>

For information about the ASEAN IBM Linux Integration Center, see:

<http://lic.austin.ibm.com/asean>

Global e-business Solutions Center (GeSC)

The Global e-business Solution Center is a part of IBM's Global Technical Sales Support team. Its mission is to provide architectural and ISV solutions integration technical sales support leadership to help IBM teams win and customers succeed with integrated e-business solutions.

The GeSC provides complex technical sales support for:

- ▶ End-to-end solution architecture and ISV integration solution selling
- ▶ Scalability/performance of cross platform end-to-end solutions

To carry out this mission, the GeSC focuses on the following major support approaches:

- ▶ Complex customer sales engagements
- ▶ Sales call technical assistance
- ▶ Solution and architecture workshop
- ▶ Solution proof of concept
- ▶ Solution scalability/performance assessment

As part of this responsibility, the GeSC provides proactive second-level support directly to the Sector/Solutions Sales Teams in the Americas. For EMEA and Asia Pacific, the GeSC is focused on leveraging its capabilities to help build geography-based, strong technical sales support skills and backing them up on critical, high-impact customer situations.

For information about the Global e-business Solution Center, see:

<http://w3.ibm.com/support/stss/gesc.html>

Solution Partnership Centers: Third-party application porting

Ten worldwide Solution Partnership Centers (SPCs) offer easy on-site and remote access for developers to IBM resources. SPCs are state-of-the-art facilities equipped with the latest IBM @server and software technologies, staffed by IBM technical consultants.

To learn more about IBM's Solution Partnership Centers, see:

<http://www.developer.ibm.com/spc/index.html>

Linux Competence Center (LCC)

The Linux Competence Center in Poughkeepsie, New York, is the result of close cooperation between the eServer Briefing Center, the New Technology Center, and the eTP Design Center. While the LCC is focused primarily on Linux under zVM server consolidation solutions, it can provide demonstrations of WebSphere with servlets connecting to CICS, DB2, WebSphere MQ, and IMS.



Linux clusters

A cluster is a collection of interconnected, complete systems that work together to provide a single, unified computing capability. The resulting cluster can either be a parallel system or a distributed system.

This chapter discusses:

- ▶ Clusters
- ▶ Linux cluster components
- ▶ Linux cluster history
- ▶ Linux cluster solutions
- ▶ The future of Linux clusters

9.1 An introduction to clusters

This section reviews the major cluster concepts and types. It covers the benefits of a cluster and the role that clusters play in emerging computer trends.

9.1.1 Core cluster concepts

As mentioned earlier, a cluster can be roughly defined as a system (either parallel or distributed) that is built from interconnected servers. It is viewed as a single entity to application or user resources.

The following sections explain the benefits and the traditional and emerging roles of clusters.

Cluster benefits

Clusters are justified by several advantages:

- ▶ **Redundancy:** Multiple physical cluster nodes avoid the fatal flaw of traditional massively parallel processing (MPP) and symmetric multiprocessing (SMP) systems. In these systems, multiple processors are stored in a single physical computer creating a single point of failure. Additionally, because cluster nodes are traditionally based on commodity hardware, the replacement cost for a failed node is typically less than repairing a similar component on an MPP system.

Figure 9-1 shows an example of an SMP solution where processors and other computing resources are grouped in a single package.




Figure 9-1 SMP system with shared storage and memory

Figure 9-2 shows an example of a grouping of separate servers in an MPP arrangement.




Figure 9-2 MPP system with an individual processor and memory

- ▶ **Scalability:** Clusters can scale into the low thousands of heterogeneous nodes while remaining competitive on price and performance against MPP systems.
- ▶ **Availability:** Because clusters have multiple interconnected physical nodes and resources (as shown in Figure 9-3), they can ensure better data continuity and recovery operations.




Figure 9-3 An availability cluster

- ▶ **Resource sharing:** Clusters can be the optimum solution for computing environments that need to dynamically assign and manage processing and storage resources in a multiuser environment.

Traditional cluster roles

Originally, clusters were part of a larger parallel computing heritage that emphasized the speed of calculations above all else. Over time, the use of clusters has expanded into other areas. The following three areas are the traditional computing roles played by clusters.

High Performance Computing (HPC)

HPC clusters are designed to use parallel computing to apply more processor power to the solution of a problem. There are many examples from scientific computing where using multiple low-cost processors in parallel to perform large numbers of operations can be highly cost effective while providing massive amounts of computing power. This is referred to as *parallel computing* or *parallelism*. Table 9-1 shows the application areas and some examples of HPC.

Table 9-1 HPC application areas and examples

Application area	Application examples
CAE (Computational Fluid Dynamics)	<ul style="list-style-type: none"> ▶ Fluent ▶ STAR-CD ▶ PowerFlow (from Enterprise X-Architecture)
CAE (Crash)	<ul style="list-style-type: none"> ▶ LS-DYNA ▶ PAM-CRASH ▶ Radioss
CAE (Structures)	<ul style="list-style-type: none"> ▶ Nastran ▶ Ansys ▶ Abaqus
Petroleum (Reservoir Modeling)	<ul style="list-style-type: none"> ▶ Eclipse ▶ VIP
Life Sciences (Ab Initio Chemistry)	<ul style="list-style-type: none"> ▶ Gaussian ▶ Gamess ▶ Jaguar
Life Sciences (Molecular Modeling)	<ul style="list-style-type: none"> ▶ Charmm ▶ Amber

Typically, parallelism in a cluster is achieved through message passing (Figure 9-4). Message parsing involves an application running on one node that takes a large set of data, breaks it up into smaller chunks, and passes those chunks to the other nodes in the cluster. As each node finishes processing its

respective chunk, it returns that chunk to the node of origin, which reassembles the information into a result set.

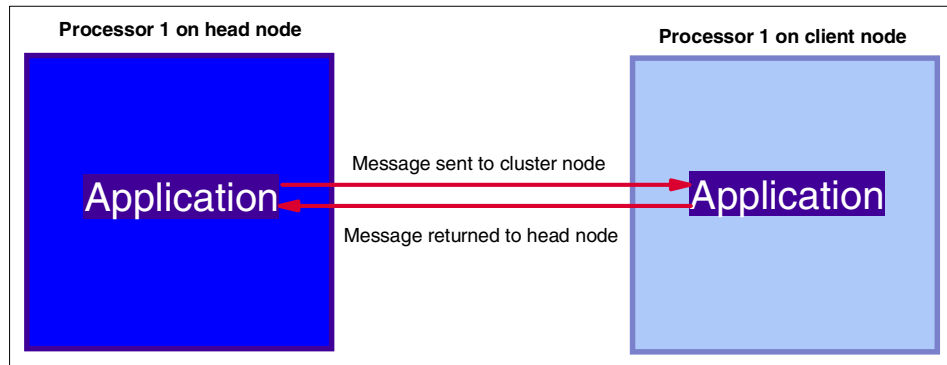


Figure 9-4 A simplified view of cluster message passing

The bottleneck in most clusters is not the processor, but rather the communication infrastructure used to pass messages from processor to processor to include the data bus of each node. The data bus in most commodity computer hardware operates at a noticeably slower speed than the processor (or even the network infrastructure) does. Because of this, clusters typically do not scale much beyond a few thousand nodes, even though the most powerful high performance commodity clusters are usually between 1,024 and 2,048 nodes.

Creating an architecture for this kind of cluster brings its own challenges. These include:

- ▶ How to install and maintain the operating and application environment on all nodes
- ▶ How to proactively manage these nodes through issuing commands and gracefully handling failures
- ▶ The requirement for parallel, concurrent, and high performance access to the same file systems
- ▶ Inter-process communication between the nodes to coordinate the work that must be done in parallel

The goal is to provide the image of a single system by managing, operating, and coordinating a large number of discrete computers. Often in this environment a user interacts with a specific node to initiate or schedule a job to be run. The application, in conjunction with various functions within the cluster, then determines how this job is spread across the various nodes of the cluster to take advantage of the resources available to produce the desired result.

High Throughput Computing (HTC)

While science needs have driven HPC and parallel operations (and consequently, cluster technology), some areas of research place a higher emphasis on throughput (quantity) of data operations than on the speed of those data operations.

Note: A common analogy used to explain the difference between HPC and HTC is the fighter jet versus the commercial airliner jet. While a fighter jet can carry a single individual to a destination faster than an airliner jet can, the airliner jet can carry multiple (more than 200) individuals at a time, on the average, providing a more efficient transportation medium.

Typically these types of operations can take weeks or even months of computational time to solve. These operations are also distinguished from HPC from the lower coupling requirements that are needed. A simple definition of coupling is the required level of data sharing among nodes at any one time. Because many HTC environments don't meet the strict definition of parallel computing (either through design or in execution), they are often referred to as distributed computing environments.

Common examples of HTC applications include protein folding and radio astronomy, among others.

To learn about protein folding, see:

<http://folding.stanford.edu/>

To learn about radio astronomy, see:

<http://setiathome.ssl.berkeley.edu/>

A variety of HTC applications run on common client operating systems. You can learn about these applications on the Web at:

<http://www.techimo.com/teams.html>

Blue Gene research project: This goal of this IBM research project is to build a massively parallel computer to be applied to the study of biomolecular phenomena such as protein folding. The project was expanded by the addition of Blue Gene/L. Blue Gene/L is also part of IBM's research initiative in autonomic computing to design computer systems that are self-healing, self-managing, and self-configuring. To learn more about this research, see:

http://www.research.ibm.com/bluegene/press_release.html

Figure 9-5 shows an example of protein folding for a typical HTC application. Some considerations of HTC clusters include:

- ▶ Defining the minimum level of coupling required by the HTC applications
- ▶ Determining how best to allocate computing resources among multiple users of the cluster
- ▶ Resolving conflicts that may arise if resource ownership is distributed (that is, different entities or individuals own their own nodes that participate in the cluster)




Figure 9-5 Protein folding simulation: Typical HTC application

High-availability cluster

High-availability clusters (Figure 9-6) are typically built with the intention of providing a resilient environment through redundancy. They provide a computing environment where the failure of one or more components (hardware, software, or networking) does not significantly affect the availability of the application or applications being used.

In the simplest case, two computers may be configured identically with access to shared storage. During normal operation, the application environment executes on one system, while the other system simply stands by ready to take over running the application in the case of a failure.




Figure 9-6 High availability example: A Web server farm

When a failure occurs, the second system takes over the appropriate resources (storage, networking address, and so forth). This process is typically called *failover*. The second system then completely replaces the failed system. The end users do not need to know that their applications are running on a different physical machine.

As you may expect, actual high availability clusters can be complex to design, especially if shared, coherent data storage is required. In fact, two subcategories of high availability clusters are file storage and database cluster, each with their own particular requirements and demands. Understanding all possible failure types and building in redundancy and the processes required to provide a smooth, efficient, and transparent failover is a difficult endeavor.

Some of the abilities of high availability clusters include:

- ▶ Providing an administrative switch over between nodes for maintenance, upgrade, or other non-failure situations
- ▶ Maintaining a consistent environment between all nodes
- ▶ Recognizing when a failure has occurred
- ▶ Implementing the failover process so that all required resources are available to the standby node
- ▶ Providing transaction and data integrity in case of a failure

All three areas, despite their differences, attempt to provide some level of abstraction, allowing users, and to an extent applications, to see the cluster as a single resource instead of individual computers.

Emerging roles

Over time, clusters are being used for more targeted use. Variations on the three traditional cluster roles are discussed in the following sections.

Server consolidation

Clusters can be an ideal platform for server consolidation because they provide a single point of control for software installation, hardware control, system monitoring, and application maintenance. Parallel or Web-based applications can take advantage of the scaling (see Load balancing in the following section) and price-performance capabilities of a cluster. Clusters can also make an ideal front end in a multitiered architecture.

Some considerations include:

- ▶ Centralizing management of resources that previously were scattered across multiple sites
- ▶ Greater economy of scale from reducing physical redundancy
- ▶ Data integration provides higher availability, improved recoverability, and easier management of enterprise data
- ▶ Application integration provides an opportunity for business process synergy

Load balancing

Load balancing (and, by extension, horizontal scaling) involves multiple servers providing a single interface to a set of resources that can arbitrarily grow (or shrink) in size over time. The most common example of this is a Web server farm. Here, a single Universal Resource Locator (URL) (the interface) is used to access a Web site (the resource). Requests coming in through that interface can be allocated across a large set of servers. This provides higher capacity and the

ability to manage the end user experience through such functions as load balancing.

Figure 9-7 shows an example of how load balancing was used in the 1998 Winter Olympics.

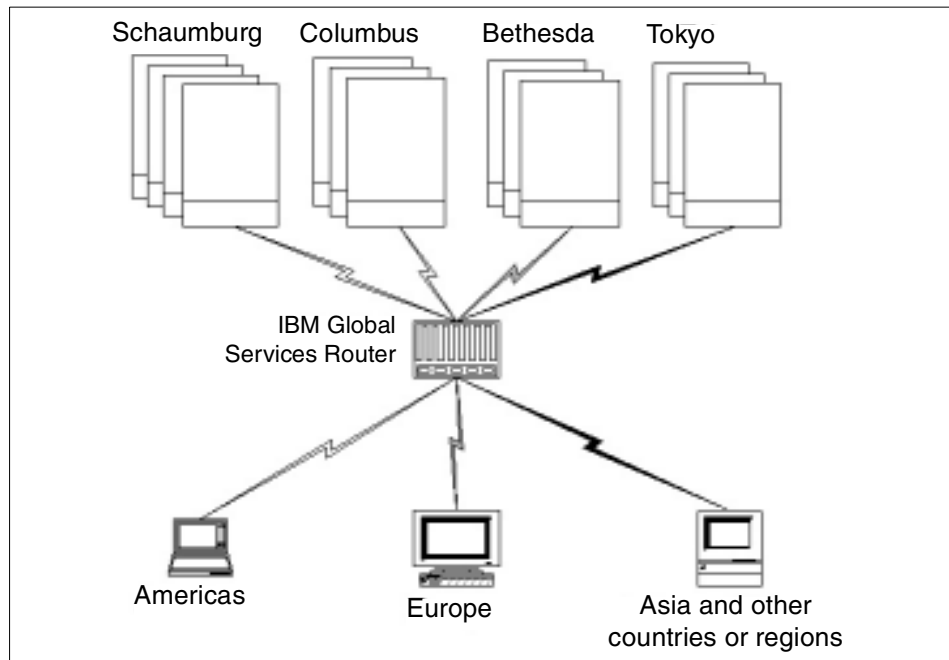


Figure 9-7 Load balancing for the 1998 Winter Olympics

Scaling considerations include:

- ▶ Establishing baseline estimates of high and low resource usage limits
- ▶ Determining which elements of an application must be reworked (if any) to take advantage of a horizontally-scalable cluster

On-demand computing

The technologies that will eventually support on demand computing are still in their infancy. These technologies include grid and autonomic computing. There is little doubt that clusters will have a role to play. See 9.5, “The future of Linux clusters” on page 271, for more information.

9.1.2 IBM cluster offerings

IBM has several product-based cluster offerings. This section discusses some of them.

Parallel Sysplex

Each node in the Parallel Sysplex can have a single CPU or consist of a symmetric multiprocessor. Each node runs a separate copy of the operating system.

The processors in an SMP node run a single operating system, OS/390 (previously known as MVS™) on the IBM S/390. They have shared main memory. The nodes in the Parallel Sysplex have shared disks, which you can access directly from each node. You can configure large complexes of shared disks, using the Enterprise Systems Connection (ESCON®) switch.

SP™

The SP system is a general-purpose scalable parallel system. It is based on *share-nothing architecture*. Share-nothing refers to the fact that each processor has its own cache and disk subsystem.

Generally available SP systems range from 2 to 128 processor nodes. Each processor node is functionally equivalent to a stand-alone RS/6000 workstation or server.

Large SP systems with up to 512 nodes have been delivered and are successfully being used today. Each processor node contains its own copy of the standard AIX operating system and other standard RS/6000 system software.

A set of new software products designed specifically for SP allows the parallel capabilities of SP to be effectively exploited. The SP system is a state-of-the-art horizontally scalable environment.

Cluster 1600

The IBM Cluster 1600 extends and enhances IBM's innovative, proven AIX and RS/6000 SP clustering technologies. The Cluster 1600 includes both the legacy SP system and new clusters made up of pSeries servers.

Cluster 1600 includes the Parallel System Support Programs (PSSP) for AIX 5L software. This software provides a comprehensive suite of applications for the installation, operation, management, and administration of the RS/6000 SP, attached servers, and Clustered Enterprise Servers (CES) from a single point of control.

PSSP 3.4 concentrates on the connection to the SP Switch2 PC Interface (PCI) of RS/6000 servers S80, H80, and M80 and pSeries 6H1, 6H0, 6M1, and pSeries 690 as SP-attached servers using the SP Switch2 PCI Attachment Adapter (F/C 8397). PSSP 3.4 supports Winterhawk1, Winterhawk2, and Silver SP nodes within the SP Switch2 PCI Attachment Adapter connection to the SP Switch2 PCI.

9.2 Linux cluster components

Much focus is given to computers (or nodes) that make up a cluster. Just as important are the communication and storage technologies that work to support the cluster, and the management software that keeps the cluster running properly.

The cluster applications themselves provide the high availability and resiliency necessary in a cluster. Throughout this section, we match general cluster component areas with specific components of Cluster 1350. This is similar in architecture to the earlier Cluster1300, but uses the Intel Xeon processor.

9.2.1 Nodes

Clusters (especially high performance clusters) typically have a large number of computers. In general, most of the nodes are configured identically. The goal of identical configuration is that the individual tasks that make up a parallel application will run equally well on the node to which they are assigned. However, some nodes in a cluster often have physical and logical differences.

The following sections discuss logical node functions and the physical node types.

Node functions

A *cluster* is two or more (often many more) computers working as a single logical system to provide services. Although, from the outside, the cluster may look like a single system, the internal workings to make this happen can be quite complex. The logical functions that are explained in the following sections are found in most clusters.

These are logical functions. In some cases, multiple logical functions may reside on the same physical node. In other cases, a logical function may be spread across multiple physical nodes. Do *not* underestimate the *importance* of cluster management. Cluster management is imperative when coordinating the activities of large numbers of systems.

Compute node

The compute node is where the real computing is performed. The majority of the nodes in a cluster are typically compute nodes. Based on the scheduling of the cluster, the compute node that is given one or more tasks to execute is part of the overall solution.

Note: The compute node in the IBM Cluster 1350 is the xSeries 300 Model 330. The Model 330 is designed for space efficiency and usability in a rack environment. It is a 1U server with one or two Intel Pentium III processors. The system also has a 133 MHz front-side bus, integrated Ultra160 SCSI interface, two integrated 10/100 Ethernet NICs, and two 64-bit/33 MHz PCI slots (one full-length and one half-length).

Control node

Control nodes provide services that help other nodes in the cluster work together to obtain the desired result. Control nodes can provide two sets of functions:

- ▶ Infrastructure services like DHCP, DNS, and other similar functions for the cluster: These functions enable the nodes to easily be added to the cluster and to ensure that they can communicate with the other nodes.
- ▶ Scheduling what tasks are to be done by what compute nodes: For example, if a compute node finishes one task and is available to do additional work, the control node may assign that node the next task that requires the work.

Note: Either an xSeries 300 Model 330 or 342 can be used as a cluster control node.

Storage node

For most applications that are run in a cluster, compute nodes must have fast, reliable, and simultaneous access to the storage system. This can be accomplished in a variety of ways depending on the specific requirements of the application.

In some cases, storage may be direct-attached to compute nodes. In other cases, the storage requirements may be filled through one or more network shared devices. In this case, storage nodes enable and control access to the storage subsystem.

Note: The storage node of Cluster 1350 is the xSeries 300 Model 342. The Model 342 is a 3U server with one or two Intel Pentium III processors. One to three disks per node can be configured with 18, 36, or 72 GB disks.

With an optional ServeRAID™ 4L adapter, mirroring of disks is possible. The node includes an integrated Advanced Systems Management processor for hardware control and management.

Install node

In most clusters, the compute nodes (and other nodes) may need to be reconfigured or reinstalled with a new image relatively often. The installation node provides the images and the mechanism to easily and quickly install or reinstall software on the cluster nodes.

Management node

Clusters are complex environments. Therefore, the management of the individual components is very important. The management node provides many capabilities, including:

- ▶ Monitoring the status of individual nodes
- ▶ Issuing management commands to individual nodes to correct problems or to provide simple (yet useful) commands to perform such functions as power on/off
- ▶ Handling events or alarms originating from the various nodes in the cluster

Do *not* underestimate the *importance* of cluster management. It is imperative when trying to coordinate the activities of large numbers of systems.

Note: Each Model 330 node has an Advanced Systems Management processor (also known as a *service processor*) that allows remote node power on/off/reset capability, monitors node environmental conditions (fan speed, temperature, power, etc.), and allows remote Power-On Self-Test/Basic Input Output System (POST/BIOS) console, power management, and SNMP alerts. The Model 342 nodes must include an optional service processor card to provide the same functionality.

User node

The individual nodes of a cluster are often on a private network that you cannot access directly from the outside or corporate network. Even if they are accessible, most cluster nodes are not necessarily configured to provide an optimal user interface. The user node is the one type of node that is configured to provide that interface for users (possibly on outside networks), who may gain access to the cluster to request that a job be run, or to access the results of a previously run job.

Node types

Now that you understand the types of logical functions that typically exist in a cluster, we discuss the node types (from a hardware perspective) that may exist in a cluster. Generally nodes within a cluster are categorized as one of three types, each of which is explained in the following sections.

Head nodes

Head node is a generic term. The head node aids in controlling the cluster, but can be used in additional ways. Head nodes generally provide one or more of the following logical node functions:

- ▶ User node
- ▶ Control node
- ▶ Management node
- ▶ Installation node

In a small cluster, typically no more than 16 nodes, all of these functions can be combined in one head node. In larger clusters, the functions are best split across multiple machines for security and performance reasons.

Compute nodes

The compute nodes form the heart of the cluster. The user, control, management, and storage nodes are all designed to support the compute nodes.

Most computations are actually performed on the compute nodes. These nodes are logically grouped, depending on the needs of the job and as defined by the job scheduler.

Storage nodes

Often when discussing cluster structures, a storage node is defined as a third type of node. However, in practice, a storage node is often a specialized version of either a head node or a compute node.

The reason that storage nodes are sometimes designated as a unique node type is that the hardware and software requirements to support storage devices might vary from other head or compute nodes. Depending on your storage requirements and the type of storage access you require, this may include special adapters and drivers to support RAID-5, storage devices attached via channel storage, and others.

9.2.2 Networks

Nodes within a cluster have many reasons to communicate. After all, they are cooperating to provide an overall result. The needs defined by the different communication types tend to dictate the various network technologies that work to serve those needs.

Communication types

Some of the key communications that must occur are explained in the following sections.

Interprocess communication

The nodes in a cluster work together to provide an overall solution. Therefore, the task that each node performs is often interrelated with the tasks that the other nodes perform.

Just as in a multi-processing or multi-threaded single system, there is almost always a need for communication between processes to coordinate processing and handle concurrent access to shared resources. The interprocess communication network is often referred to as the *cluster interconnect*. It may share a physical connection with the storage or management networks.

Management operations

With a large number of nodes, it is important to recognize when an individual node may be failing. It is also important to proactively issue commands that may prevent an error from occurring.

It is not feasible to manage multiple nodes of a cluster individually in a stand-alone manner. There must be a way to communicate from one or more management nodes to any other node in the cluster.

Note: IBM Cluster 1350 provides a service processor network. It uses the newer RSA adaptors, which are one for every 11 nodes. The service processor is built in and has its own Ethernet port and cable. Because of this, only one Ethernet cable per eight nodes is required, as opposed to using a PCI slot and Ethernet cable per node.

Software installation

For a variety of reasons, individual nodes often need to have their software reinstalled. This may be due to reconfiguration for a specific task to be performed by the cluster, or simply to replace the operating environment when an individual node has suffered a disk failure.

Either way it is important that system images can be efficiently defined and distributed. Since this operation may be done to certain nodes while the rest of the cluster is operational and performing work, it is important that pushing a new operating environment to some nodes does not adversely affect the performance of the other nodes.

Storage access

In most clusters, individual nodes need to read or write data to a shared storage device. For high performance computing, it is contrary to the intention of the cluster if this storage access proved to be a bottleneck and resulted in other nodes waiting to access the device. Therefore, it is important to have both an

efficient network for fast access to storage and a file system that provides efficient, concurrent access to the data.

Common network technologies

Depending on the actual application and the performance requirements, these various communications should be carried out often on different networks. Therefore, you typically have more than one network and network type linking the nodes in your cluster. Common network technologies used in clusters are explained in the following sections.

Fast Ethernet

Fast Ethernet and TCP/IP are the two standards used for most networking. They are simple and cost-effective, and the technology is always improving. Fast Ethernet works well with lower speed Intel-based machines used for Beowulf nodes, and can operate at either half-duplex or full-duplex. With full-duplex, data can be sent and received at the same time. Full-duplex transmission is deployed either between the ports on two switches, between a computer and a switch port, or between two computers. Full-duplex requires a switch, since a hub will not work.

Gigabit Ethernet

Gigabit Ethernet uses a modified version of the American National Standards Institute (ANSI) X3T11 Fibre Channel standard physical layer (FC-0) to achieve one Gigabit per second raw bandwidth. Gigabit Ethernet supports multi- and single-mode optical-fiber and short-haul copper cabling. Fibre is ideal for connectivity between switches and servers and can reach a greater distance (from 500 meters to 10 kilometers today) than copper.

Myrinet™

Myrinet is a high-performance, high-availability, packet communication and switching technology. Myrinet links may be up to 200m in length on multimode fiber and operate full-duplex at a data rate of 2+2 Gigabits per second. Myrinet switches are simple, low-latency, cut-through devices that can be combined to provide networks up to thousands of hosts and network-bisection data rates in Terabits per second. Myrinet interfaces include a processor and firmware to offload packet processing from the host computer.

Myrinet software support for Linux and for other operating systems provides multiple software interfaces. This includes the native Myrinet API, TCP/IP and UDP/IP, MPI, VI, PVM, and Sockets. All Myrinet software support is open source.

For more information about Myrinet, see the following Web site:

<http://www.myri.com/>

Fibre Channel

Fibre Channel is a data storage technology that can provide cluster nodes with shared access to data. Fibre Channel is a loop architecture as opposed to a bus architecture like SCSI. The loop structure enables the rapid exchange of data from device to device, achieving a maximum data transfer rate of 100 MB/sec. Fibre Channel is a common technology used to provide clusters with rapid and shared data access.

Note: Cluster 1350 supports the IBM FAStT Fibre storage subsystem and the IBM TotalStorage Enterprise Storage Server.

Internet Small Computer Systems Interface (iSCSI)

iSCSI is a standard that allows SCSI information to be encapsulated and transported over a TCP/IP network. The benefit of this approach is that organizations can take advantage of their existing network infrastructure to provide access to data storage resources anywhere on the network using the same I/O commands they use for local data. iSCSI can allow Linux clusters to access remote data stores without requiring additional infrastructure.

As the iSCSI standard matures, several developments are expected. The first is the introduction of Host Bus Adapters (HBAs). HBAs improve performance and throughput by offloading the TCP/IP overhead from client operating systems. The second is the deployment of 10-Gigabit Ethernet, which dramatically improves the data rates at which iSCSI can operate.

Cluster 1350 does not support iSCSI. It is mentioned here for completeness.

Note: IBM TotalStorage IP Storage 200i is an iSCSI solution that provides support for 109.2 GB up to 3.52 TB of RAID storage.

InfiniBand

InfiniBand is a technology that is designed to provide a channel-based switched fabric that is capable of handling any kind of communication between processors and I/O devices. Because of the high bandwidth designed into the standard (2.5 Gbits, 10 Gbits, and 30 Gbits/sec.), InfiniBand is expected to gradually replace the existing Peripheral Component Interconnect (PCI) shared-bus approach used in most of modern Intel-based computers, and is seen as an ideal communication media for the data-intensive demands of clusters.

Cluster 1350 does not support InfiniBand. It is mentioned here for completeness.

Note: IBM offers several InfiniBand products, including the InfiniBlue 4X Host Channel Adapter (HCA). This allows current PCI-based Linux servers to connect to an InfiniBand network.

9.2.3 Support software

Software binds the hardware infrastructure together. It creates and manages the abstraction layer that allows users and applications to view the collective cluster resources as a single computing entity. There are three main cluster software categories.

Operating system

All nodes in a cluster typically run the same operating system. For large clusters, it is important that this operating system be economically viable, as well as easily deployable across the nodes. This implies that it should be easily installable in an automated fashion, and support remote access for management and configuration.

Linux meets these requirements. In fact, the low cost of obtaining Linux and the rich function for developing and deploying applications that run on Linux have been the catalyst for increased interest in using clusters to meet the needs of a variety of computing problems. By using relatively low cost Intel-based hardware and Linux, many businesses that previously could not justify the expense of a cluster are now adopting the technology at a frenetic pace. As we previously mentioned, there are other considerations when building a cluster, such as how to manage a large number of nodes.

Note: To provide a common base on which the complex mixture of hardware and software that is required to form a cluster can be tested and consistently deployed, Red Hat and SuSE will be supported on Cluster 1350 as of first quarter 2003. Additional Linux distributions and versions will be considered for the future.

File system

Depending on the actual application, many clusters require a large number of nodes to have concurrent and very high-speed access to the same files. In these cases, it is important to provide a shared file system that provides data protection and the performance that is typically required for cluster-based applications.

Note: IBM provides the General Parallel File System (GPFS) for Linux. GPFS is a proven file system that has been used for years in AIX clustering environments. It is available for Linux and xSeries 300 hardware.

Management tools

When given the hardware previously discussed, you can fairly easily construct a cluster. This is only the beginning of putting a cluster into production.

One of the most difficult parts of deploying a cluster is having the right tools in place to manage a large number of nodes as a single entity. This includes:

- ▶ Managing node and node group information
- ▶ Monitoring and controlling hardware
- ▶ Running remote commands on multiple machines
- ▶ Monitoring system events
- ▶ Managing and synchronizing configuration files for all machines
- ▶ Providing security
- ▶ Diagnosing problems
- ▶ Installing software

Note: IBM provides cluster management through two utilities: xCAT and Cluster Systems Management (CSM) for Linux. xCAT is a more mature product and uses the xSeries service processor. This maturity is a great strength. xCAT is used by all three NCSA TF systems.

CSM includes technology derived from PSSP. While it is the standard system management software for Cluster 1350, it can also be used for cluster systems management on any open, Intel processor-based server.

9.2.4 Applications

Few off-the-shelf HPC applications exist, for the simple reason that HP clusters tend to be focused on a specific, unique task. HTC solutions also tend to be custom-written and focus on massive calculation projects. HA clusters typically rely on middleware as the software layer that requires the most customization to ensure reliability, recoverability, and even scalability.

The two HA subtypes, storage and database clusters, each have their own special concerns that may require application customization. For storage clusters, more emphasis is placed on the capabilities of the distributed file system itself. However, some customization of specialized retrieval libraries (layered between the application and the distributed file system) is common.

With database clusters, customization is driven by the primary concerns of availability and recovery. However, scalability is increasingly becoming a concern.

You can approach the load balancing clusters in two ways:

► **Server farm**

The traditional way of scaling loads is to provide a server farm that sits behind a server that spreads all incoming requests across the various nodes of the cluster. The advantage to server farms is the low amount of customization typically required to implement this type of load balancing.

► **Process migration/Single System Image (SSI) cluster**

Using kernel-level extensions, process migration clusters can seamlessly move execution processes to and from any node in the cluster. The benefit of this approach is that applications do not have to be rewritten to take advantage of this approach to load balancing.

An SSI cluster expands the simple process migration model to create unified cluster-wide inter-process communication (IPC) objects, such as pipes, fifos, semaphores, message queues, shared memory, sockets, and signals. This allows programs, processes, and users to work together regardless of the node in the cluster on they are executing.

Either type of load balancing (server farm and process migration) may still require a distributed file system or additional customization to deal with data concurrency issues.

Figure 9-8 shows how all the various cluster components come together. Notice that the top four layers are all software components. This abstraction layer serves to allow the creation of a cluster environment from a heterogeneous mix of hardware. Because an effective cluster can be made up of any type or brand of server, interconnect, or storage subsystem, clusters can take advantage of commodity hardware or the existing IT resources of an organization.

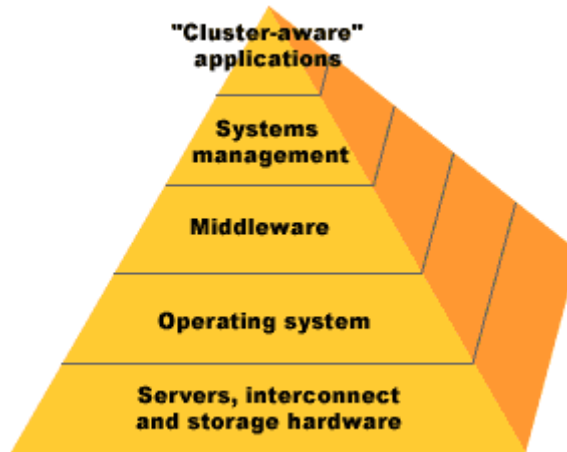


Figure 9-8 Layout of cluster components

Note: IBM provides the following cluster-enabled applications:

- ▶ DB2 Universal Database (UDB) Enterprise-Extended Edition (EEE)
- ▶ WebSphere Application Server
- ▶ WebSphere Commerce Suite
- ▶ WebSphere Edge Server
- ▶ Telecom Application Server
- ▶ WebSphere Everyplace Suite

9.3 Linux cluster history

The current state of Linux cluster solutions is best understood through its history. Modern Linux clusters are conceptually indebted to open source software, the Intel processor, and the demands of the HPC community. There were clusters of non-Intel/non-Linux machines before (and after) the now-famous Beowulf cluster program started in 1994 at NASA. However, the idea of a cluster wasn't truly feasible or economical until the commodity nature of PC-based hardware was combined with the open source nature of the Linux operating system.

The Beowulf program

The Beowulf program started with a simple idea: Use 16 inexpensive Intel-based PCs and open source parallel programming libraries to create a virtual supercomputer environment. Fast Ethernet was too slow to keep up with the bus speed of the individual nodes (bus speed continues to be a major roadblock even in today's clusters). Because of this, Donald Becker (a NASA contractor) rewrote a network card driver to support *channel bonding*. Channel bonding is a process that combines the four network cards in each into a single, virtual connection that is fast enough to permit parallel processing software running on each computer to communicate as if they were part of a single system.



Figure 9-9 Wiglaf, the first cluster in the Beowulf program

All the software that made up the system was open source, including the Linux kernel, the NIC driver, and the parallel processing libraries. Considering this, the NASA team (Becker, Tom Sterling, John Salmon, and other CESDIS members) continued and developed the ability to tune the cluster to the specific characteristics of the system. To learn more about the team's achievements, see:

http://sdcd.gsfc.nasa.gov/DIV-NEWS/CESDIS.11_97.Becker.award.html

Finally, because all of the hardware involved was standardized and inexpensive, additional nodes could be added (or replaced) as necessary. Today, the phrase *Beowulf cluster* refers to a class of cluster where:

- ▶ The majority of the components are commodity off-the-shelf (COTS).
- ▶ The operating system is open source (Linux or BSD variants).
- ▶ There is emphasis of the cluster on HPC and as an Industry Standard Message Passing interface.

Emergence of high availability clusters

Starting in the 1990s, Linux was increasingly used in an infrastructure role providing basic (and essential) services including routing, firewall/IDS, Web, and DHCP/DNS. The need for a higher level of availability also increased along with the penetration of Linux into businesses and government.

The Linux high availability project started with the purpose of bringing failover and takeover capabilities to the Linux operating system. The central component of the Linux high availability project is the *heartbeat program*. The heartbeat program works on a two-node model: two identically-configured computers, with one designated as the “primary”, and one as the “secondary”. The secondary node monitors the “heartbeat” or continual signal sent by the primary through a

network connection or serial cable. When the primary node fails (and fails to send the heartbeat signal), the secondary activates its services and takes ownership of primary's IP address, maintaining the cluster service.

While Linux high availability provides a single backup for essential infrastructure servers, the two-node limitation limited its effectiveness. As a result, administrators increasingly began to combine the heartbeat capabilities of Linux high availability with the availability and load balancing capabilities of another project, the *Linux Virtual Server (LVS)*. LVS works by assigning the public IP address of the cluster to a load-balancing server. The load balancing server then retransmits requests for services to the least-busy node of the cluster. When the load balancing LVS is mirrored with Linux high availability, the result is a highly available and scalable cluster for Web-based services.

Linux load balancing

In the mid 1990s, Professor Amnon Barak of the Hebrew University in Jerusalem ported the code (originally developed for BSD UNIX) of the Multicomputer Operating System for High Performance Cluster Computing (MOSIX) to Linux. MOSIX uses special algorithms to automatically (and optimally) move processes around the cluster. A process can be started from any node and can be moved to or from any other node. Because MOSIX is implemented as a kernel extension, Linux applications don't have to be rewritten to take advantage of its load balancing capabilities.

In 2002, MOSIX began offering an optional parallel I/O (MOPI) component. This component attempts to improve I/O operations by migrating parallel processes to the nodes that hold that data to be used by the processes.

In 1998, some members of the original Beowulf program formed the Scyld Computing Corporation. Scyld Beowulf added extensive cluster installation and management software to the collection of open source cluster tools then in existence. It also incorporated the idea of process migration. Similar to MOSIX, Scyld Beowulf uses an extension to the Linux kernel called Bproc to migrate cluster processes from the master node to other nodes in the cluster. In addition, all process management can be performed from the master node. As a result, individual cluster nodes are not required to have duplicate copies of cluster application code. This simplifies cluster administration and maintenance and frees node storage space for application data and cache.

There is a major difference between MOSIX and Scyld. MOSIX (and openMosix) transparently migrates processes under kernel control based on a least loaded algorithm, with open file handles. Scyld migrates processes under user control, closing file handles. This makes the process faster and leaner to migrate, but necessitates reopening files that you want to access.

In late 2001, Professor Barak decided to discontinue the open source license of MOSIX. Moshe Bar, a former developer of the Linux MOSIX code decided to continue open source developments under the name *openMosix*. openMosix has incorporated several features that are not found in the original MOSIX, including support for user-mode Linux (UML), among others. It continues to be one of the most popular Linux cluster software components.

Note: UML is part of the Linux kernel that allows multiple virtual machines, each running their own version of the Linux kernel, to run on top of a single host kernel. Because the host kernel provides each virtual machine with simulated hardware, they can duplicate nearly all programs that are available on the host itself. UML is used by developers to test and debug Linux kernels and applications in a secure environment.

Management consolidation

While various open source cluster software components were maturing in the 1990s, management tools still lacked maturity and integration. In April of 2000, the Open Cluster Group, a collection of companies (IBM, Intel, Dell, SGI, Veridian, MSC.Software Corporation) and research organizations (Oak Ridge National Laboratory (ORNL), Indiana University, University of Illinois National Center for Super Computing Applications (NCSA)) released their first project: the *Open Source Cluster Application Resources* (OSCAR).

The goal of the OSCAR project is to identify the best open source utilities for building, programming, and administering HPC clusters. Once identified, these programs are bundled together and distributed as a single package. OSCAR has been an extremely popular open source cluster package and is used by a wide range of organizations.

Note: IBM provides tools for CSM to work with OSCAR through its alphaWorks® site. For more information, see:

<http://www.alphaworks.ibm.com/tech/ect4linux>

Another approach to cluster management and infrastructure is the *Single Systems Image project*. SSI attempts to provide a general-purpose Linux clustering environment using the best open source components for availability, scalability and manageability. SSI (and the Cluster Infrastructure (CI) project it relies on) has its origin in Hewlett Packard NonStop Clusters for UnixWare technology. The CI code is responsible for maintaining membership of and communication between the various cluster nodes. SSI creates the appearance of a single system through shared process, IPC, device, and networking space, a clustered file system (through OpenGFS), load leveling (using the MOSIX algorithm), and a single management interface.

Note: SSI was voted the best open source project of the 2002 LinuxWorld Conference and Expo.

9.4 Linux cluster solutions

The IBM Linux cluster solution is the Cluster 1350. IBM has other cluster solutions, which may include Linux alongside other operating systems as a combined cluster package.

Note: IBM offers other non-Linux clusters solutions. See the following Web site for more information:

<http://www.ibm.com/servers/eserver/clusters/hardware/>

9.4.1 Cluster 1350

The Cluster 1350, like the earlier Cluster 1300, is an integrated cluster with nodes of Intel processors. Unlike Cluster 1300, Cluster 1350 uses single or double Intel Xeon processors in each node. It can scale to 512 computer nodes and 32 storage nodes. Cluster 1350 is designed as a solution for industrial, financial services, life sciences, governmental, and educational organizations.

You can read more about Cluster 1350 at:

<http://www.ibm.com/servers/eserver/clusters/hardware/1350.pdf>



Figure 9-10 xSeries Linux Cluster

9.4.2 ClusterProven®

The IBM ClusterProven program sets specific criteria for validating end-to-end solutions that meet industry standards for high availability on every IBM platform. Developers who achieve ClusterProven validation earn the use of the ClusterProven mark. This allows them to use ClusterProven in marketing their solution to customers and receive co-marketing assistance from IBM.

To see a list of ClusterProven applications, go to:

<http://www.software.ibm.com/solutions/isv/igssg.nsf/CPSolutionWeb>

9.4.3 Solution examples

Cluster 1300 and Cluster 1350 are the primary current Linux cluster packages offered by IBM. However, companies still choose IBM hardware and middleware to support their custom cluster solutions. This section offers examples of Linux clusters in the real world.

High-performance computing: NCSA

The National Center for Supercomputing Applications (NCSA) is a research organization based in the United States. Its mission is to create high-performance computing infrastructures for scientists. IBM has installed two Linux clusters for the NCSA. The first cluster includes more than 600 xSeries running Linux and Myricom's Myrinet cluster interconnect network. The second cluster is one of the first to use Intel's next generation 64-bit Itanium processor and will run SuSE.



Figure 9-11 The NCSA Linux cluster built on the xSeries server

It would take more than one-and-a-half million years for a person to tabulate with a calculator the number of calculations that the NSCA supercomputer handles in a single second. The NSCA needs immense speed and power to simulate the violent collision of black holes and the resulting gravitational waves.

You can learn more about the IBM NCSA cluster on the Web at:

<http://www.ncsa.uiuc.edu/About/NCSA/>

Plus, you can find more information about Linux clusters on the NCSA Web site at:

<http://www.ncsa.uiuc.edu/News/Access/Stories/ClusterUpdate/>

Experimental Itanium-based cluster used to study fluid dynamics problems

As of May 2001, NCSA's resident Linux cluster using Intel's 64-bit Itanium architecture was a mere 64-processor collection of machines on loan from IBM. Its purpose is to give NCSA cluster development team the chance to get codes up and running on the Itanium architecture before delivery of the 320-processor terascale Itanium-based cluster.

As it turns out, this experimental cluster is proving to be much more than a tool for benchmarking and troubleshooting. It is showing that breakthrough science can be done on even relatively small Itanium-based clusters.

A team of scientists at the University of Minnesota's Laboratory for Computational Science and Engineering (LCSE) are using the Itanium-based Linux cluster to create detailed computer simulations that show the behavior of turbulent gas flows. This data is now being visualized on a one billion cell grid. The LCSE team expects to perform even more detailed simulations consisting of an eight billion cell grid. These simulations will test theories about the behavior of turbulent fluids first gleaned from an earlier eight billion cell simulation done at Lawrence Livermore National Laboratory in 1998. That simulation was done over one week using a 3,840-processor computer. However, follow-up work was never done because the Livermore machine was relegated to a classified environment.

Cluster run aids gas dynamics research

The 8-billion-cell turbulence simulation on the NCSA Itanium cluster gives the LCSE team the chance to compare new, more detailed data to their 1998 data. Visualization and analysis of the 1998 simulation of the unstable mixing layer in gas flows revealed an unexpected correlation between the behavior of the turbulent flow on small scales and some aspects of the large-scale velocity field.

This suggests a new method for modeling the effects of small-scale motions upon large-scale ones in gas dynamics simulations. The new modeling method requires high resolution, since these newly-found relationships may not be evident in lower resolution simulations.

The LCSE team, led by LCSE Director Paul Woodward, uses the gas dynamics code Piecewise-Parabolic Method (PPM). Because the PPM application is extremely scalable, Woodward expects to see better performance on NCSA Itanium clusters as they grow larger.

Unlike their simulations on the Livermore machine, the researchers can save all the data from this large simulation. This is made possible by the increased availability of disk space gained by using commodity drives and the archival storage system. The data is then transferred over the Abilene network to the LCSE terascale data analysis and visualization system for detailed analysis.

The goal of the research is to develop an accurate model of the behavior of unresolved fluid turbulence in computer simulations. Achieving this goal would make possible more accurate simulations of gaseous flows in astrophysics, meteorology, and aerospace applications. If the model can be extended to liquid flows, it can also help develop more accurate simulations of turbulent flows in oceanography and in engineering applications.

HA Clusters Computing: BBDO Dusseldorf

BBDO Dusseldorf is a German-based company that provides complete Web hosting and infrastructure solutions. To provide Web-hosting for its clients, BBDO INTERACTIVE has chosen an IBM solution including DB2 Universal Database and xSeries servers running the open source Linux operating system. BBDO Dusseldorf maintains high availability of client Web sites by using LifeKeeper™ for DB2 clustering software from IBM Business Partner SteelEye.




Figure 9-12 BBDO Dusseldorf high availability cluster on the xSeries server

You can read more about the BBDO Dusseldorf cluster at:

<http://www.eweek.com/article2/0,3959,131649,00.asp>

Grid computing: National Digital Mammography Archive (NDMA)

The NDMA is a distributed computing infrastructure that connects hospitals at the University of Chicago, the University of North Carolina, and the Sunnybrook and Women's College Hospital in Toronto. The purpose of the NDMA is to deliver computing resources as a utility-like service over a secure Internet connection. This enables up to thousands of hospitals to store mammograms in digital form, giving medical personnel near-instantaneous access to patient records.

The NDMA portal systems run almost exclusively on IBM hardware including 16 distributed Netfinity servers running Linux. You can learn more about the NDMA cluster on the Web at:

<http://nscp.upenn.edu/>

9.5 The future of Linux clusters

Clusters are now used as the building block of other emerging IT models and technologies. In this way, clusters are becoming components for larger projects, as opposed to the traditional view of a cluster as a stand-alone system that performed a specific computing role.

9.5.1 Emerging trends

There are three major emerging trends that rely extensively on the features provide by Linux. These trends are explained in the following sections.

Grid computing

Grid computing is the ability, using a set of open standards and protocols, to gain access to applications and data, processing power, storage capacity and a vast array of other computing resources over the Internet (not just within a physical cluster). The grid user looks at a grid and sees essentially one, large virtual computer.

In this way, a grid can be seen as a *metaccluster*. A metaccluster is a collection of clusters that are interconnected with wide area network (WAN) technologies to provide seamless access to the combined resources of the grid. Grids are distinguished from conventional distributed computing by their focus on large-scale resource sharing, innovative applications, and in some cases, high-performance orientation.

You can learn more about grid computing on the Web at:

<http://www.ibm.com/grid/index.shtml>

Autonomic computing

Autonomic computing is an attempt to manage the complexity of IT systems by using the inspiration found in one of the most complex systems of the human body, the autonomic nervous system. The focus of autonomic computing is to design and build computing systems capable of running themselves. It requires, adjusting to varying circumstances, and preparing their resources to handle most efficiently the workloads placed upon them.

To read more about autonomic computing, see:

<http://researchweb.watson.ibm.com/autonomic/>

On demand computing

On demand computing describes an environment where users own IT assets but can dynamically increase and decrease them based on need. The main advantage for users is variable pricing based on usage, similar to the way energy utility companies provide and bill for their services. In addition, the nature of on-demand computing allows users to outsource the management of such a demand-driven system to third parties.

You can read more about on demand computing on the Web at:

http://www.ibm.com/services/ondemand/index_flash.html

9.5.2 The role of Linux clusters in emerging trends

Linux clusters will serve grid computing by supporting the intensive computational and storage demands imposed by the distributive nature of the

grid. In addition, Linux clusters provide the least expensive way of providing the fault tolerance demands of any autonomic system. Linux clusters are the most cost-effective solution for providing the scalability requirements of on-demand computing.

Another future view of Linux clustering is as an *HPC appliance*. The HPC appliance will be based on super thin nodes, with little or no persistent storage. Nodes will be added to the cluster transparently, with a simple boot, which will make the node amenable to a migrated process. Since the nodes are thin, version skew will be eliminated, as will traditional cluster installation (putting an image on every node).

The HPC cluster of the future should be “open from the ground up”, because it will be based on LinuxBIOS technology. The ramifications of putting Beowulf 2 technology, such as Scyld, on top of a LinuxBIOS platform will be very close to the HPC appliance.



IBM Services for Linux

This chapter introduces you to the IBM services that help Linux customers gain the greatest advantage of a Linux-powered IBM *@server* solution. The chapter is intended to help those who are interested in Linux from technical specialists to sales, brand, marketing, and product area professionals. It provides the first fully consolidated reference to IBM's continuously growing information base on Linux.

The resources portrayed here provide a good, consolidated starting point to more detailed information. It also offers a scope of the tremendous acceptance and commitment Linux has today. We recommend that you read and learn about these resources to help you stay current with Linux's commitments and new services, which continue to grow rapidly.

The information in this chapter was correct at the time this book was written. However, some of the information may change due to normal variations in the life cycle of each product. This is true for products other than Linux as well. This means that IBM may not support a specific product even though it may appear to be supported in this chapter. Therefore, we encourage you to investigate such topics as end-of-service dates, product withdrawals, and other support restrictions while gathering information from this chapter.

Important: The services portrayed here are for IBM customers only and may not be available in some specific countries or regions. Contact your local IBM Global Services (IGS) representative for details.

10.1 IBM Services

IBM's main focus is on the creation, development, and manufacturing of the industry's most advanced information technologies. This includes computer systems, software, networking systems, storage devices, and microelectronics.

IBM Global Services translates these advanced technologies into business value for customers and helps making information technology (IT) easy to acquire and manage. To learn more about IBM Global Services, see:

<http://www.ibm.com/services/>

10.2 IBM Global Services

IBM helps customers optimize the business value of Linux-based solutions through cost-effectiveness, choice, and competitive advantage. IBM has experienced Linux consultants and architects engaged in thousands of open-source projects worldwide. With this, IBM can help customers in all industries enhance resiliency, scalability, stability, and security with their Linux deployments.

IBM Global Services is the largest part of IBM with over 140,000 employees in 164 countries or regions. It is widely recognized as the largest services company in the world.

10.2.1 IBM Support Line

IBM offers “enterprise-grade” support for Linux via the Internet and voice support 24 hours-a-day, 7 days-a-week. This support ranges from answering usage questions to identifying problems, all revolving around the customer needs.

Recent analyst studies have shown that among the most important concerns related to Linux implementations, support is an important consideration for both solution and vendor selection. IBM offers world-class support for Linux as a standard offering, with several options including the ability to customize support according to the skills and experience of the customers current IT staff.

Support Line is IBM's support offering for e-infrastructure software. It is an established worldwide proven service which covers IBM and non-IBM Software products. Support Line offers:

- ▶ Years of experience in world-wide support for IBM and non-IBM operating systems
- ▶ World-wide support structure
- ▶ The right tools

- ▶ The right people
- ▶ The right alliances

IBM Operational Support Services: Support Line for Linux

This fee service provides remote Linux technical assistance to answer usage, installation, and technical questions, to address product compatibility and interoperability concerns, and to isolate and diagnose problems. IBM Linux engineers are available 24 hours-a-day to provide defect support (fixes).

Support Line provides consistent, cross-platform Linux support for IBM @server platforms and Intel/AMD OEM hardware. With this service, customers receive telephone access and electronic access to IBM services specialists.

Support Line is a single, integrated, remote support solution for Linux that provides:

- ▶ An 800 toll-free number, voice and electronic support
- ▶ Support for SuSE, Red Hat, and UnitedLinux
- ▶ 8 a.m. to 5 p.m. support Monday through Friday in your time zone
- ▶ Full shift (24-x-7 optional coverage)
- ▶ Response time: 2 hours primeshift, 2 to 4 hours offshift
- ▶ Available on zSeries, iSeries, pSeries, and xSeries servers, plus Intel/AMD OEM hardware
- ▶ Defect support by the IBM Change team to the source code level
- ▶ Interpretation of product documentation
- ▶ Configuration samples of IBM and multivendor database searches
- ▶ Planning information for software fixes

IBM Support Line for Linux offers:

- ▶ 24-x-7 enterprise-level remote support
- ▶ Fast and accurate problem resolution
- ▶ Access to IBM skilled services specialists
- ▶ Answers the “how-to” questions
- ▶ Resolutions to complex defect issues
- ▶ Support across four major distributions of Linux
- ▶ Remote assistance available through electronic and telephone access
- ▶ Electronic support and problem submission that optimizes your time and gives full online tracking of open support items
- ▶ Support Line options

IBM Operational Support Services: Support Line for Linux Clusters

This fee service provides remote Linux cluster technical assistance to answer usage, installation, and technical questions, to address product compatibility and

interoperability concerns, and to isolate and diagnose Linux cluster problems. It provides specialized Linux cluster support with open source and OEM. It offers remote problem identification and technical support for your Linux clustered environments. This service complements your in-house skills by providing comprehensive knowledge of Linux cluster environments.

This fee service provides:

- ▶ How to, usage, defect, fix information
- ▶ An 800-toll free number and electronic Access
- ▶ 8-x-5 or 24-x-7 coverage
- ▶ Defect support by the IBM Change team to the source code level
- ▶ Response time: 2 hours primeshift, 2 to 4 hours offshift
- ▶ Available on Cluster 1300 and IA/32

Integrated cluster support is available for:

- ▶ IBM hardware
- ▶ Compute nodes
- ▶ System management and storage nodes
- ▶ OEM hardware bundled with cluster:
 - Myrinet and Cisco cluster/network switches
 - KVM switches and terminal servers
 - Power management modules
- ▶ IBM, OEM, and open source software, including:
 - Cluster Systems Management (CSM)
 - General Parallel File System (GPFS)
 - Message Passing Interface (MPI) modules
 - Linux, and other open source software as supplied

IBM Operational Support Services: Advanced Support

IBM's Advanced Support provides the highest level of remote support available through IBM. The customer receives 24-x-7 coverage, 365 days-a-year. This service is tailored to meet unique customer needs for continuous, business-critical Linux system operation.

IBM Operational Support Services: Custom Linux Support

IBM offers remote Linux engineering support for creating modified versions of Linux. This services provides full remote support once you implement this custom distribution throughout your enterprise or customer base. By leveraging IBM's vast industry of Linux knowledge, IBM helps you to avoid serious pitfalls that can occur when you create your own distribution. We help you understand the necessary interdependent packages and required modules.

This service offers:

- ▶ Technical consulting: Development-level Linux support for Linux kernels and operating systems
- ▶ Tailored solutions to meet your development needs and schedules
- ▶ Development-level consulting in determining the appropriate Linux components, open source code, Linux tools, and applications for custom builds
- ▶ Compatibility and interoperability development support
- ▶ Assistance with customize Linux images/builds
- ▶ Assistance with custom Linux integration build
- ▶ Assistance with custom integration build testing
- ▶ 8 a.m. to 5 p.m. Monday through Friday availability
- ▶ Dedicated project team for specific skill area
- ▶ Collaboration with IBM development teams, Linux Technology Center (LTC), and the open source community
- ▶ Help to maximize mission-critical Linux system availability
- ▶ A focus on diagnostics, remote diagnostics, management, monitoring, and security
- ▶ Proactive development-level technical support based on an understanding of your mission-critical environment

IBM has the critical mass to deliver support teams in multiple worldwide locations. It has used some of its best talent in multiple locations to create the Change team. IBM can draw on the skills of over 200 key members of the Linux Technology Center found in more than 20 locations worldwide.

You can learn more about the Support Line at:

<http://www.ibm.com/services/its/us/supportline.html>

10.3 Linux consultancy

Linux acceptance by the marketplace is rapidly moving forward. It has gained credibility and is being deployed widely for Web serving and scientific and technical computing clusters in a relatively short time. It is now being ported as well to mission-critical areas such as workload consolidation, retail points-of-sale, and financial branch automation.

IBM is committed to providing the same level of support normally associated with enterprise computing environments as Linux continues to move in to those key business and industry areas. IBM's Linux service offerings are designed to help Linux customers to achieve better flexibility and cost-to-benefit rates.

IBM Global Services offers one of the industry's most comprehensive portfolios of Linux consultative and support offerings, from planning and design to implementation and technical support. IGS also offers a full portfolio of Linux Services and has been doing so since February 2001. Over 300 IBM consultants skilled in Linux are available worldwide to help customers design, build, and enhance their Linux solutions.

For information about IBM's services and support for Linux, refer to the following IBM Linux Services and support Web sites.

Strategy and design

- ▶ Linux Strategy Workshop, a customized engagement that helps customers initiate or expand Linux

<http://www-1.ibm.com/services/its/us/linuxworkshop.html>

- ▶ Grid Innovation Workshop and other grid computing offerings

http://www.ibm.com/grid/grid_offerings.shtml

- ▶ Server consolidation survey

<http://www.ibm.com/services/its/us/techassess.html>

Systems migration: Infrastructure and applications

- ▶ Linux application services
 - Portfolio assessment
 - Linux application porting: Solaris-to-Linux, OS/2-to-Linux, Windows-to-Linux, HP3000 MPE-to-Linux
 - Application development

Contact Jeff Gibbs, Linux Services principal, by telephone at (602) 357-4360 or by e-mail at <mailto:jefgibbs@us.ibm.com> for more details.

- ▶ Linux migration and implementation services

- Cluster implementation services

<http://www.ibm.com/services/its/us/drmk1m09.html>

- File/print workload consolidation on zSeries

<http://www.ibm.com/services/its/us/drmk1m01.html>

- Site/rollout services
<http://www.ibm.com/services/its/us/siteserv.html>
- IBM DB2, WebSphere, Tivoli implementation services
<http://www.ibm.com/services/its/us/swsupport.html>
- IBM Migration Services
<http://www.ibm.com/services/its/us/migration.html>
- ▶ Linux performance management, testing, and scalability services
<http://www.ibm.com/services/its/us/perftest.html>
- ▶ Linux security services
 - Architectural review: For those who are still considering moving core applications to Linux
 - System Security Assessment: Design-focused, intended to reveal vulnerabilities during or following construction
 - System Security Services: Implementation project actually designed to harden an existing security environment
 - Ethical Hacking: After-the-fact vulnerability testing

Contact Jeff Gibbs, Linux Services principal, by telephone at (602) 357-4360 or by e-mail at <mailto:jefgibbs@us.ibm.com> for more details.
- ▶ Linux Virtual Services
<http://www.ibm.com/services/e-business/hosting/mgdhosting/linux.html>

Support and learning

- ▶ Linux technical training services
<http://www.ibm.com/services/learning/us/catalog/linux/>
- ▶ IBM Support Line for Linux
<http://www.ibm.com/services/its/us/supportline.html>

10.4 IBM Services and Software

IBM is rolling out services to help customers rapidly deploy and optimize Linux-based e-business environments. IBM services for Linux covers Software Group products such as WebSphere, Data Management (DB2), Lotus, and Tivoli. IBM is also delivering migration services to help customers port applications to a wide variety of IBM and non-IBM platforms running Linux.

10.4.1 WebSphere

Many customers need WebSphere as a front end to their legacy applications. Customers need to access legacy applications (for example, CICS) through a browser interface to provide users (internal or external) with Web-enabled access to multiple existing applications. They also want solutions that do not require a long development cycle and that support an open-standards compliant integrated infrastructure. Customers want to leverage existing transactional applications, database assets, and existing investment in hardware platforms that have a superior scalability characteristic on which to run their new application.

Every customer with an IBM mainframe capable of having Integrated Feature for Linux (IFL) presents a good fit for WebSphere for Linux. IFL-capable systems are the 9672 G Series (G5 and G6), Multiprise 3000, zSeries 800, and zSeries 900 vendors.

IBM provides a variety of services for WebSphere that go from migration to specific training in all WebSphere products. For WebSphere Software Services, see:

<http://www.ibm.com/software/ad/vaws-services/websphere.html>

For more information about WebSphere training and technical enablement, see:

<http://www.software.ibm.com/wsdd/education/enablement/>

IBM also provides support for the wide range of WebSphere products made to fit the Linux environment. Table 10-1 lists references where you can find support for WebSphere products that run on Linux.

Table 10-1 WebSphere for Linux product services Matrix

WebSphere product	Product detail	Support
WebSphere Application Server	Advanced Edition 4.0 for Linux on zSeries	http://www.ibm.link.ibm.com/cgi-bin/master?xh=a3KRKds8D5usfr1USenGnN9332&request=announcements&parms=H%5f201%2d366&xfr=N#div19
WebSphere Edge Server	IBM WebSphere Edge Server V2.0 for Linux multiplatform	http://www.ibm.com/software/webservers/edgeserver/support.html
WebSphere Commerce Suite	WCS Business Edition for Linux Intel and Linux on the zSeries	http://www.ibm.com/software/webservers/commerce/wc_be/support.html
	WCS Professional Edition for Linux Intel	http://www.ibm.com/software/webservers/commerce/wc_pe/support.html

WebSphere product	Product detail	Support
WebSphere Personalization	WebSphere Personalization 4.0 for Linux on the zSeries	http://www.ibm.com/software/webervers/personalization/support.html
WebSphere MQ (formerly MQ Series)	MQ V 5.2 for Linux on Intel and zSeries	http://www.ibm.com/software/ts/mqseries/support/index.html
WebSphere Portal	WebSphere Portal 4.1 for the Linux Intel Red Hat or SuSE 7.1	http://www.ibm.com/software/webervers/portal/support.html
WebSphere Studio Application Developer	WSAD V4.0 for Red Hat/SuSE 7.1 or Higher on Intel	http://www.ibm.com/software/ad/studioappdev/support/
WebSphere Transcoding Publisher	Red Hat Linux V7.1 or SuSE 6.4 on Intel	http://www.ibm.com/software/webervers/transcoding/support.html

10.4.2 DB2 for Linux

DB2 for Linux Beta Version 8.1 is now available. You can learn more about this release on the IBM DB2 for Linux Web site at:

<http://www.ibm.com/software/data/db2/linux/validate/>

You should also refer to Chapter 8, “IBM Software for Linux” on page 199, for DB2 for Linux hardware and software requirements.

IBM offers the following services related to DB2 for Linux:

- ▶ **Remote DB2 Installation:** This service is supported for DB2 for Linux running under a logical partition (LPAR), a guest system (VM), or a stand-alone Linux. For more information, see:

<http://www.ibm.com/services/its/us/drmk1m12.html>

IBM can also help customers to migrate to DB2 for Linux. Learn about a variety of options and helpful information on the Web at:

<http://www.ibm.com/software/data/db2/migration/>

- ▶ **DB2 for Linux support:** For technical support, product information, and all other DB2 for Linux support-related issues, go to:

<http://www.ibm.com/software/data/db2/udb/support.html>

For information about DB2 consulting services and other services, refer to:

<http://www.ibm.com/services/>

10.4.3 Tivoli

Customers' growing use of Linux systems within their Tivoli-managed environments has extended IBM's commitment to scale Linux services and support to Tivoli products.

IBM announced Tivoli Linux enablement since the summer of 2002 for its security software and Web management products. For more information, see:

<http://www.ibm.com/software/tivoli/solutions/linux/>

Linux enablement was also extended to Tivoli Storage Management, Configuration, Operations, Performance, and Availability products. For details on the current Linux supported Tivoli products, refer to:

<http://www.ibm.com/software/tivoli/>

You can access services for Linux supported Tivoli products, which include consulting, training, and certification, on the Tivoli services Web site at:

<http://www.ibm.com/software/tivoli/services/>

Tivoli's Customer Support is quite extensive and complete. Some services are available only for registered users. You can find Tivoli Customer Support on the Web at:

<http://www.ibm.com/software/sysmgmt/products/support/>

Product support matrices for Tivoli

Table 10-2 through Table 10-9 on page 287 describe Tivoli product support matrices for Linux for each area of Tivoli.

IBM Tivoli Security Management

Table 10-2 lists product detail for Tivoli Security Management.

Table 10-2 IBM Tivoli Security Management product matrix for Linux

Product	Product detail
Tivoli Access Manager for e-business	Tivoli Access Manager for SuSE SLES7 for zSeries

IBM Tivoli Storage Management

Table 10-3 lists product details for Tivoli Storage Management.

Table 10-3 IBM Tivoli Storage Management product matrix for Linux

Product	Product detail
Tivoli Storage Manager	TSM 5.15 Server for Linux on the x86 Platform
	TSM 5.15 Clients for Red Hat 7.1 and 7.2 x86
	TSM 5.15 Clients for SuSE 7.1, 7.2 and 7.3 x86

IBM Tivoli Configuration and Operations

Table 10-4 lists product details for Tivoli Configuration and Operations.

Table 10-4 IBM Tivoli Configuration Manager product matrix for Linux

Product	Product detail
IBM Tivoli Configuration Manager	TCM Server for SuSE 7.2 on x86
	TCM Server for Red Hat on zSeries
	TCM Gateway for SuSE 7.2 on x86
	TCM Gateway for SuSE 7.0 and Red Hat on zSeries
	TCM Web Gateway for Red Hat 7.1 on x86
	TCM Endpoint on Red Hat 7.1 7.2 and SuSE 7.2 on x86
	TCM Endpoint on Red Hat and SuSE 7.0 on zSeries

IBM Tivoli Workload Scheduler

Table 10-5 lists product details for the Tivoli Workload Scheduler.

Table 10-5 Tivoli Workload Scheduler product matrix for Linux

Product	Product detail
IBM Tivoli Workload Scheduler	TWS V8.1 Job Scheduler Console on Red Hat 7.1 for x86*
	TWS V8.1 Job Scheduler Connector on Red Hat 7.1 for x86

IBM Tivoli Performance and Availability Management products

IBM Tivoli Performance and Availability Management Products for Linux consists of:

- ▶ IBM Tivoli Monitoring
- ▶ IBM Tivoli Enterprise Console®
- ▶ IBM Tivoli NetView
- ▶ IBM Tivoli Service Level Advisor

The details for each product are provided in the following tables. Table 10-6 lists product details for Tivoli Monitoring.

Table 10-6 IBM Tivoli Performance and Availability product matrix for Linux

Product	Product detail
IBM Tivoli Monitoring	IBM Tivoli Monitoring Server, Gateway and Endpoint for Red Hat 7.0 and 7.1 on Intel
	IBM Tivoli Monitoring Server, Gateway and Endpoint for Turbo Linux 6.1 and 6.5 on Intel
	IBM Tivoli Monitoring Server, Gateway and Endpoint for SuSE Linux 7.1 and 7.2 on Intel

IBM's Tivoli Enterprise Console consists of the following components:

- ▶ Event Server
- ▶ Gateway
- ▶ Endpoint Adapter
- ▶ UI Server
- ▶ Event Console

Table 10-7 outlines the support availability for each Tivoli Enterprise Console component on Linux.

Table 10-7 Tivoli Enterprise Console Product Matrix for Linux

Supported version by platform	Tivoli Enterprise Console component				
Red Hat Linux for Intel 7.1 and 7.2	Event Server	Gateway	Endpoint Adapters	UI Server	Event Console
SuSE Linux for Intel 7.0 and 7.1	Event Server	Gateway	Endpoint Adapters	UI Server	Event Console
Turbo Linux for Intel 7.0	Event Server	Gateway	Endpoint Adapters	UI Server	Event Console

Table 10-8 lists product details for IBM Tivoli NetView.

Table 10-8 IBM Tivoli NetView product matrix for Linux

Product	Product detail
Tivoli NetView	Tivoli NetView 7.1.2 for Red Hat or SuSE 7.1 on Intel

Table 10-9 lists product details for IBM Tivoli Service Level Advisor.

Table 10-9 TIBM Tivoli Service Level Advisor for Linux product matrix

Product	Product detail
Tivoli Service Level Advisor	Service Level Advisor SLM component for Linux Red Hat or SuSE 7.1 on Intel
	Service Level Advisor Reports Server for Linux Red Hat or SuSE 7.1 on Intel

10.4.4 Lotus

To support and enhance your Notes and Domino environment, Lotus offers a full range of professional services, including consulting, education, and customer support. See IBM's Lotus Software site for more information:

http://www.lotus.com/lotus/products.nsf/fa_prohomepage

Table 10-10 outlines the Lotus Family services and support matrix for Linux.

Table 10-10 .Lotus support matrix for Linux

Product	Product detail	Service and support
Lotus Notes Domino	Lotus Notes Domino V 6 Red Hat Linux, Version 7.2 or SuSE Linux, Version 8.0 on Intel	http://www.lotus.com/products/rnext.nsf/873769A79D9C5B2285256A0800720B96/478B755CB7C8CB9C85256C40005EF2A?OpenDocument&cwesite=domfamserv

10.5 Platform services

The xSeries, zSeries, iSeries, and pSeries platforms work with Linux as an operating system today, among the other operating systems that exist in the marketplace. Just as software-based services for Linux are needed, so is a well-founded, platform-based service and support structure which is still growing.

The IBM Linux Support Line provides Operational Support Services and a Premier Remote technical support service. For more information, see:

<http://www-1.ibm.com/services/its/us/supportline.html>

IBM provides technical support for the major distributions of the Linux operating system running in the xSeries, zSeries, iSeries, and pSeries platforms, as well as some non-IBM applications that operate in a Linux environment. IBM helps answer how-to questions, performs problem source determination, and provides mechanisms for a solution. In addition, by leveraging partnerships with the key distributors of the Linux operating system, IBM provides defect-level support for the Linux operating system. Remote assistance is available through toll-free telephone or electronic access depending on the country or region.

IBM provides services for all currently supported xSeries, zSeries, iSeries, and pSeries platforms in a different degree of variety and complexity through its Linux Portal. For more information, see:

<http://www.ibm.com/linux/>

As Linux continues to grow, so will IBM's commitment to enhance current services or provide new ones.

IBM also provides Learning Services for customers who are interested in knowing more about Linux and making it their key operating environment for e-business. IBM's Linux education and training services offer a full curriculum and help to establish a personal Linux education plan. For details about IBM's Linux Learning services and Linux Training, you can visit:

<http://www.ibm.com/services/learning/us/catalog/linux/>

10.5.1 zSeries services

The IBM Managed Hosting for Linux on the zSeries service is intended to provide cost-effective managed server capacity on demand. If you are new to e-business on demand, you can find out more information on the Web at:

<http://www.ibm.com/services/ondemand/thinking.html>

With IBM Managed Hosting - Linux virtual services, the customer can tap into managed server capacity without the up-front expense of buying the physical hardware. Instead of the physical Web, database, and application servers that businesses currently rely on, virtual servers on the zSeries running Linux can be leveraged. This means that availability and reliability gain a boost while IT infrastructure is greatly simplified.

If you are interested in finding out more about this service, refer to:

<http://www.ibm.com/services/e-business/hosting/mgdhosting/linux.html>

Table 10-11 lists various resources for platform-based information about services that are available for the zSeries server.

Table 10-11 zSeries Linux information matrix

Information	Location on the Web
Linux for zSeries applications	http://www.ibm.com/servers/eserver/zseries/solutions/s390da/linuxisv.html
Linux for zSeries operating system download	http://www.ibm.com/servers/eserver/zseries/os/linux/dist.html
Linux for zSeries Development tools	http://www.ibm.com/servers/eserver/zseries/os/linux/ldt/
Linux for zSeries Library	http://www.ibm.com/servers/eserver/zseries/os/linux/press.html
Linux for zSeries Education	http://www.ibm.com/servers/eserver/zseries/os/linux/ed.html

10.5.2 iSeries services

IBM provides a variety of services for customers interested in consolidating their servers on the iSeries server. To learn more about these services, see:

<http://www.ibm.com/servers/eserver/series/linux/adv.html>

Table 10-12 lists various references for platform-based information about services that are available for the iSeries server.

Table 10-12 iSeries Linux information matrix

Information	Location on the Web
Linux for iSeries applications	http://www.ibm.com/servers/eserver/series/linux/apps.html
Linux for iSeries Education	http://www.ibm.com/servers/eserver/series/linux/educ.html
Linux for iSeries planning	http://www.ibm.com/servers/eserver/series/linux/plan.html
Linux for iSeries Developer Resources	http://www.ibm.com/servers/eserver/series/linux/resc.html
Linux for iSeries Reference Library	http://www.ibm.com/servers/eserver/series/linux/library.html
Linux for the iSeries support	http://www.ibm.com/servers/eserver/series/linux/support.html

10.5.3 pSeries services

The pSeries server makes a reliable foundation for fast 64-bit native Linux servers. See Chapter 6, “pSeries and Linux” on page 155.

IBM provides hardware and software support service for the new pSeries 630 6C4 and 630 6E4 Linux-ready express configurations that are now available.

For information about pSeries 630 6C4, see:

http://www.ibm.com/content/home/store_IBMPublicUSA/en_US/eServer/pSeries/entry/6306C4.html

For information about pSeries 630 6E4, see:

http://www.ibm.com/content/home/store_IBMPublicUSA/en_US/eServer/pSeries/entry/6306E4.html

IBM will provide more services for Linux for pSeries in response to growing interest from customers to use Linux as their native operating system in this platform.

For more information, see the IBM Linux for pSeries Web site at:

<http://www.ibm.com/servers/eserver/pseries/linux/index.html>

10.5.4 xSeries services

IBM is leveraging enterprise scalability and availability at attractive prices through xSeries Linux enablement. IBM’s Linux combination with the xSeries has proven successful. Many resources and information are available to users who are interested in using Linux on the xSeries as their key environment to meet their needs.

Table 10-13 lists various references that provide platform-based information about services that are available for the xSeries server.

Table 10-13 xSeries Linux information matrix

Information	Location on the Web
Linux for the xSeries information	http://www.pc.ibm.com/us/eserver/xseries/index.html
eServer Blades Center	http://www.ibm.com/servers/eserver/blades/
Linux clustering on xSeries	http://www.ibm.com/servers/eserver/clusters/
Linux for xSeries ISV resources	http://www.ibm.com/developerworks/offers/linux-speed-start/isv.html

10.5.5 Consolidated platform matrix

For more valuable information about IBM Linux for all IBM @server services, as well as latest information, updates, new services available, etc., refer to the IBM Linux dedicated Web site for each platform as listed in Table 10-14.

Table 10-14 IBM Web site matrix for each Linux platform

Platform	Location on the Web
zSeries	http://www.ibm.com/servers/eserver/zseries/os/linux/index.html
iSeries	http://www.ibm.com/servers/eserver/iseries/linux/index.html
xSeries	http://www.ibm.com/servers/eserver/xseries/linux/index.html
pSeries	http://www.ibm.com/servers/eserver/pseries/linux/index.html

10.6 Examples of Linux services

IGS has been successful in delivering services to our new Linux Clients. This section shares a few examples of such situations.

10.6.1 Example #1: A billion dollar financial institution

This customer was running Red Hat 6.2 on a VA Linux hardware cluster solution. The customer's financial application would crash and core dump on 3% of the runs. There was no pattern to when the failure would occur. The problem plagued the customer's IT department for over two months. No other Linux service providers could identify the problem.

The same problem was reported sporadically on the kernel mailing list for over two years with no one able to identify the problem, nor offer a solution. Any attempts to instrument the kernel to isolate the problem changed the timing of the code and the problem would disappear (masked). The only possible way to find the root cause was via detailed code inspection that involved IA32 assembly code at their most fundamental level. Through a series of test cases and code inspection, IBM isolated the problem to a bug in `switch_mm()`, which resulted in a process running with a 'stale' `ldt`.

IBM recommended a fix to the kernel mailing list, and it accepted by Linus Torvalds. He responded in a personal note back to the IBM Linux Change team and released in the v2.4.8 Linux kernel.

10.6.2 Example #2: A multi-million dollar advertising company

This customer was running Red Hat 7.1. They experienced a situation where an embedded Linux NFS client would error on its NFS mounts from the server at random points in time, but within a 24-hour period of booting. IBM reproduced the problem using the exact hardware and a copy of the software that the customer was running. Using iterative debug kernels and creating kernel dumps on error, IBM isolated the problem to NFS kernel code that was erroring with hard links beyond the exported directory.

IBM provided a work around immediately given to the customer, while a patch was developed and consensus in the open source community was gained for a long-term code fix. Agreement was made and the fix will be included in the latest 2.4.x kernel. Previous threads to the kernel mailing list indicated that others had run into this problem previously, but no one could isolate it to the root cause.

10.6.3 Example #3: One of the largest network services providers

When this customer, running Red Hat 6.2, upgraded their hardware, they ran into a problem where the system hung 15 minutes after boot time. The original problem statement indicated that no software changes were made and that only processor upgrades were completed from one release to another. Working through the weekend to investigate the problem, IBM found that a boot parameter was set to hard code memory boundaries. This was in conflict with reserved space for the BIOS.

Further investigation demonstrated that the boot parameter did not need to be set because the system properly probes memory size and sets it accordingly. In Red Hat 7.1, code was added to check BIOS memory allocation and not to override it with boot-time configuration settings.



Desktop Linux

This chapter introduces you to the Linux architecture. In particular, it examines the kernel that is related to the rest of a complete distribution. It describes the components that you see on the desktop and reviews two of the major office suites that run under Linux.

This chapter also presents some of the Linux hybrids that are available, which may help you in transitioning from Windows to Linux. Finally it weighs the advantages and disadvantages of each of the three solutions in deploying Linux in a business.

11.1 Linux desktop environments

A few short years ago, Linux was primarily a command-line environment. The Linux graphical user interfaces (GUIs) have become more and more sophisticated in an effort to become more user friendly and to attract a broader audience. The command-line environment is still available and much preferred by developers and those accustomed to a UNIX system or Linux commands that were derived from the UNIX system. This section reviews a few of the more popular desktop environments.

Behind all the other desktop environments lies the foundation software for all Linux GUIs: the X Window System often abbreviated as “X”. The X Window System is simply the environment that allows graphics to be displayed on your system's hardware. The *window manager* is responsible for how X looks and how it interacts with you and your applications.

Users who do not want all the extra enhancements use X by itself. However, most use the X Window System indirectly through a window manager, such as ICEWM, or a desktop environment, such as the K Desktop Environment (KDE) or GNOME. Each window manager has a unique look and often comes packaged with its own desktop themes. Each desktop environment has its own individual characteristics as well as its own collection of GUI-based Linux applications.

As shown in Figure 11-1, desktop environments run on top of window managers which are built upon the X Window System. X, in turn, requires lower level operating system utilities to function.

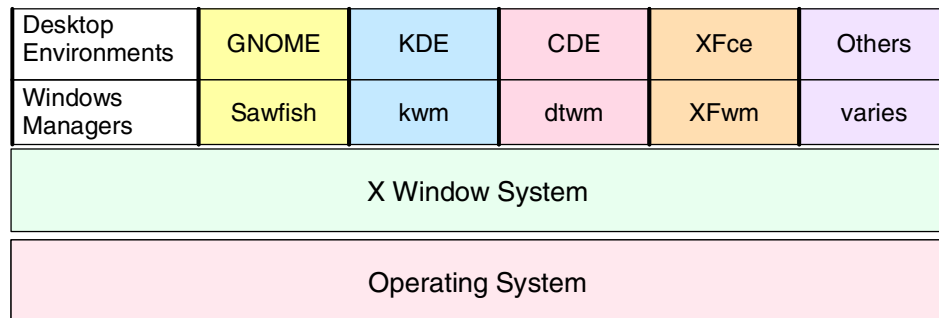


Figure 11-1 Relationship of X, window managers, and desktop environments

Linux distributions generally come packaged with both GNOME and KDE GUI desktop environments. These enhanced user interfaces move from the traditional X Window System interface to ones that are easy to use and familiar to users of other environments. Table 11-1 compares some of the common functions in each

desktop. These desktops have a similar feel and touch. Which one to use is a matter of personal preference.

Table 11-1 Linux open source desktop environments

Function	GNOME	KDE	OpenOffice
Word processor	AbiWord	Kword Kugar - Business	OpenWriter
Spread sheet	Gnumeric	KSpread	OpenCalc
Presentation	Achtung	KPresenter	Impress
Browser	Galeon	Konqueror	
Vector graphics	Sodipodi Sketch	Kontour	OpenDraw
Roster graphics	GIMP	Krita	

11.1.1 KDE and GNOME

A complaint from users who have migrated to KDE or GNOME desktops is the performance when one of these desktop GUIs is launched. The architecture of running a desktop environment involves both a window manager and a desktop environment, which is an implementation of that window manager over the X Window System base.

KDE

The KDE project is a large open group of developers. This group consists of several hundred software engineers from all over the world committed to free software development. The KDE project is a free software project whose code is available under the LGPL/GPL licenses. This means that everyone is free to modify and distribute KDE source code.

KDE brings these advantages:

- ▶ An integrated help system
- ▶ Consistent look and feel of KDE applications
- ▶ Standardized menu and toolbars
- ▶ Internationalization: KDE is available in more than 50 languages
- ▶ A number of KDE applications, including the KOffice Suite

The KDE distribution consists of various run-time libraries, base components (window-manager, desktop (Figure 11-2), Konqueror browser), KMail, graphics applications, and utilities.

You can learn more about KDE on the Web at:

<http://www.kde.org/>

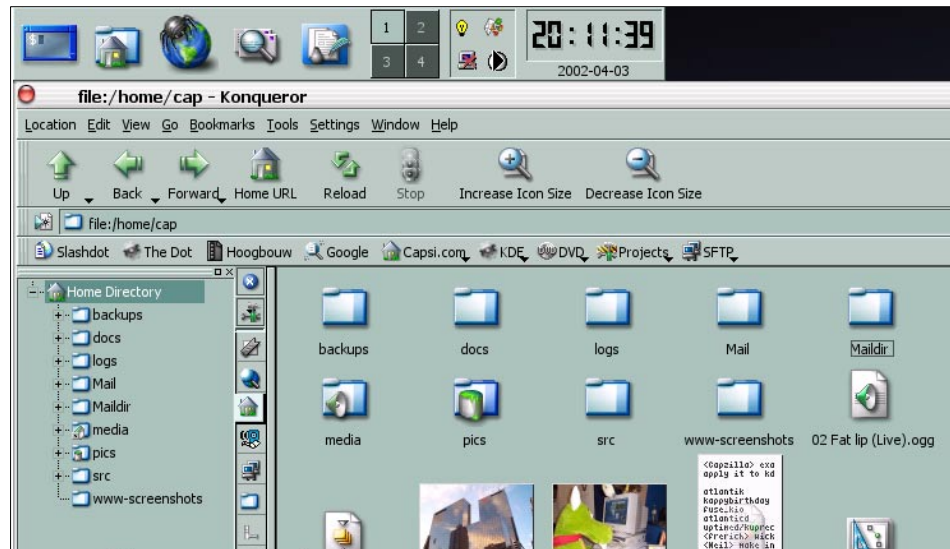


Figure 11-2 Part of the KDE 3.1 desktop

GNOME

GNOME is part of the GNU project and is open source software. The purpose of the GNOME project is to provide a user friendly suite of applications and an easy-to-use desktop that is entirely open source.

The major components of GNOME are:

- ▶ **GNOME desktop:** An easy to use window-based environment for users
- ▶ **GNOME development platform:** A collection of tools, libraries, and components to develop powerful applications on Linux
- ▶ **GNOME Office:** A set of office productivity applications

Miguel de Icaza started the GNOME project in 1997. By 1998, the project had over 200 programmers contributing code. Miguel and Nat Friedman co-founded Ximian, the for-profit commercialization on the GNOME desktop, but even that occurred more than two years earlier. The predecessor company was Helix Code, which was renamed Ximian.

GNOME ranges in scope from utilities to low-level development libraries to end-user applications. The applications include e-mail, an Internet browser, graphics, and an office suite.

For more information, see:

<http://www.gnome.org/>

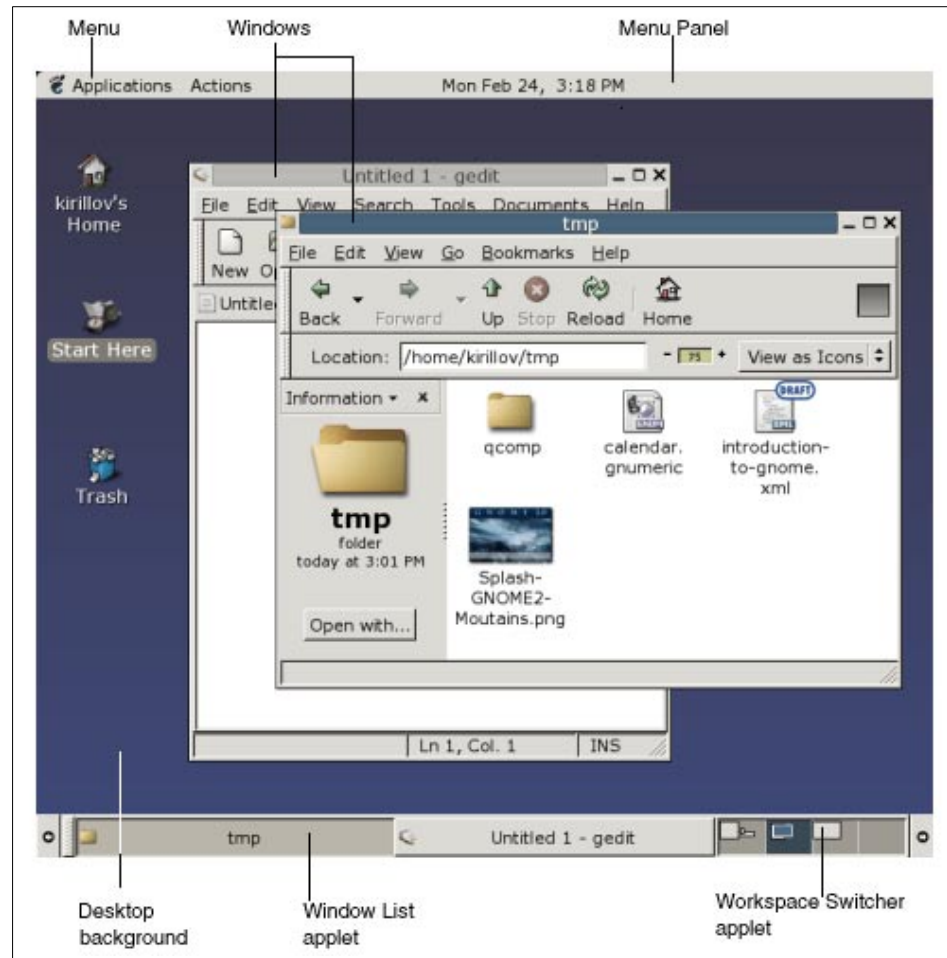


Figure 11-3 GNOME desktop

Other desktops

A variety of additional window desktops are available. ICEWM and Enlightenment are two of them. ICEWM is a desktop environment built primarily for those who are comfortable with the Windows task bar and start menu. ICEWM provides these features with a few unique additions. For example, you can add a text entry field to your task bar that acts like a command prompt. This is a handy way to launch applications that do not appear in your start menus or to launch applications with custom command-line switches.

Enlightenment is a window manager for X. Its design goal is to be as configurable as possible. Enlightenment's design aim is to become a desktop shell. That means it will manage your application windows, launch applications, and manage your files.

Some of the features of the current release include:

- ▶ Fully configurable window borders
- ▶ Theme support
- ▶ Menus
- ▶ Translucent moving of windows
- ▶ Window groups
- ▶ Multiple virtual desktops
- ▶ Auto-scrolling menus
- ▶ KDE hint support
- ▶ GNOME hint support
- ▶ Window auto-cleanup support
- ▶ Graphical online help

11.1.2 StarOffice and Open Office

Although it is possible to run a variety of office suites under Linux, only two major suites are currently in use: StarOffice and its follow-on product OpenOffice. Both of these office suites can open and save most documents originally created in the Microsoft Office format. They can be saved in their original format or in the OpenOffice or StarOffice file format.

StarOffice, from Sun, Inc., is available from selected retailers for about \$80 U.S. StarOffice runs on multiple operating systems, including the Solaris operating environment, Microsoft Windows, and Linux. This office suite contains full-featured applications with word processing, spreadsheet, presentation, graphics, and database capabilities. StarOffice is also fully compatible with other office suites, including Microsoft Office.

In Asian markets, StarOffice is marketed under the brand name StarSuite. Apart from the brand name, there are no major functional differences between the two products. Sun's StarSuite 6.0 and StarOffice 6.0 files are interchangeable and compatible. The primary functionality difference of StarSuite 6.0 Office Suite is the support for Unicode and linguistic support for Chinese (Traditional and Simplified), Japanese (Kanji), and Korean.

OpenOffice shares the same code base as StarOffice 6.0. This office suite contains a word processor, spreadsheet, presentation software, equation editor, and a drawing program. It opens most major formats such as Microsoft Office, saves to PDF, has comprehensive help, and offers spell checking in 15 languages.

You can find a list of specific features on the Web at:

http://www.openoffice.org/dev_docs/source/features.html

What are the differences between StarOffice 6.0 and OpenOffice 1.0? StarOffice 6.0 software is a commercial product aimed at organizations and consumers. OpenOffice.org 1.0 is aimed at users of free software, independent developers, and the open source community.

StarOffice includes licensed-in, third-party technology such as:

- ▶ Spellchecker and thesaurus
- ▶ Database component (Software AG Adabas D)
- ▶ Select fonts including Windows metrically equivalent fonts and Asian language fonts
- ▶ Select filters, including WordPerfect filters and Asian word processor filters
- ▶ Integration of additional templates and extensive clip-art gallery

The other major difference not listed above is the technical support available from StarOffice.

11.2 Windows: Linux hybrids

The following sections look at a few of the available Linux *hybrids*. These are variations of the standard distributions. The most common feature among these is the need to ease transition of a Windows user to Linux. The Linux hybrids in this section include Lindows, Wine, CrossOver, and some demo Linux products.

11.2.1 A look at Lindows

Lindows is a Debian GNU/Linux-based distribution whose main emphasis is to make Linux easy to use. It targets non-technical users versus the traditional Linux audience. Lindows' focus is on delivering applications that provide the same features as a user's favorite Windows application. Its features include the ability to view and print major file formats including Microsoft Word, PowerPoint, Excel, and MP3.

LindowsOS enables you to:

- ▶ Surf the Web (pop-up ad blocking, tabbed Web browsing)
- ▶ Spell check, send, and receive e-mail (POP, IMAP)
- ▶ Instant message with others (AIM, Yahoo, ICQ, MSN)
- ▶ Play MP3 files and digitize your own CDs
- ▶ Play games (more than 100 games included)

- ▶ Use Microsoft file-compatible word processing, spreadsheet, and presentation software with Bitstream Fonts.
- ▶ Perform photo and graphics editing
- ▶ Print to more than 800 different printers
- ▶ Manage administrators and multiple users on one computer
- ▶ Update your operating system and applications automatically
- ▶ Connect to PCMCIA, USB 1.0, USB 2.0, Firewire, and wireless devices
- ▶ Connect and share data over a Microsoft Windows network
- ▶ Built-in pre-configured firewall for maximum security

At the time this redbook was written, Lindows is marketing their product on pre-installed PCs for \$199. It is available in from large retailers on an online order basis (not in stores).

Click-N-Run Warehouse

In addition to the emphasis on ease of use, the key selling point for Lindows is its Click-N-Run Warehouse. Linux applications have long been as easy to use as their Windows counterparts, but installing them has not been easy. Lindows provides a wide variety of applications, over 1,600 according to the Lindows Web site. You can install these applications by simply clicking a button. Most of this software is freely available, although Lindows offers some commercial applications, such as StarOffice.

LindowsOS users can simply visit the Click-N-Run “Warehouse” and browse through dozens of software titles. When they find an application they want, they simply select the Click-N-Run button and the software is automatically downloaded and installed on the user’s LindowsOS computer. In addition, the selected software title is saved in the user’s MyLindows.com account. This allows them to click and run the application on other machines they may own or new machines they may purchase.

11.2.2 CrossOver: Running Windows programs on Linux

CodeWeavers is the developer of CrossOver. This organization provides products and services that help customers transition from the Windows software base to Linux. They accomplish this primarily by creating programs that allow you to use Windows software on Linux. CodeWeavers is the leading corporate backer of the Wine project, which we describe in the following section.

CrossOver Office is capable of running a range of Windows software. CodeWeavers supports the following applications:

- ▶ Microsoft Office 97 and 2000
- ▶ Microsoft Word
- ▶ Microsoft Excel

- ▶ Microsoft Powerpoint
- ▶ Microsoft Outlook
- ▶ Microsoft Internet Explorer
- ▶ Microsoft Visio
- ▶ Lotus Notes
- ▶ Quicken

11.2.3 Wine ‘is not an emulator’

Wine is an open source software initiative that systematically re-implements the Win32 API under Linux. Wine makes it possible for Linux PCs to run Windows applications as if they were native. Although some users call Wine a Windows emulator, the company has gone to great lengths to dispel this notion. Even its name, Wine means “Wine Is Not an Emulator”.

Wine implements Windows Win32 and Win16 APIs on Linux. In other words, it acts as a compatibility layer. Wine does not require Microsoft Windows. Its code is 100% Microsoft-free, but it can optionally use native system dynamic link libraries (DLLs) if they are available.

In mid 2002, Wine consisted of over 1 million lines of C code, written by more than 300 developers from dozens of countries or regions around the world. Wine implements more than 90% of the calls in popular Windows specifications such as ECMA-234 and Open32.

Wine is still under development. It can run some programs flawlessly, while others are not fully supported. For further details on these programs, see the Wine Web site at:

<http://www.winehq.com>

11.2.4 Trying Linux without installing it

There are a few products that allow you to experience Linux without installing it permanently. This section documents a sample of the demo CDs that are available. Some of these are quick installations of Linux, while others enable you to run Linux by booting the CD (without installing it).

WinLinux

WinLinux is a Linux system that installs like a Windows application, automatically detecting and configuring most of your hardware devices.

You do not have to put your existing data and files at risk to install WinLinux 2003 because it shares free hard drive space with Windows. That means you do not have to create additional hard drive partitions. You can start WinLinux 2003 just

as you start any other Windows application. After it is installed, the next time you restart your computer, you can choose between starting Windows or WinLinux.

WinLinux works independently from Windows. After you install it, you can remove Windows from your computer keeping your original files.

DemoLinux

DemoLinux is dedicated to the production of CD-ROMs (the DemoLinux CDs) that allow you to run Linux without installation, disk partitioning, and other hassles that usually prevent people from trying Linux. The site has three downloadable ISO images of Linux. Version 1 is based on Mandrake, version 2 is based on Debian, and version 3 is the latest at the time of this writing and is also based on Debian.

For further details, visit the DemoLinux Web site at:

<http://www.demolinux.org>

Knoppix

Knoppix is a bootable CD with a collection of GNU/Linux software, automatic hardware detection, and support for many graphics cards, sound cards, SCSI and USB devices, and other peripherals. KNOPPIX can be used as a Linux demo, educational CD, rescue system, or adapted and used as a platform for commercial software product demos. It is not necessary to install anything on a hard disk. Due to on-the-fly decompression, the CD can have up to 2 GB of executable software installed on it.

For further details or to download the software, see:

<http://www.knoppix.org>

SuSE evaluation CD

You can find a “live evaluation” CD from SuSE on the Web at:

<http://www.linuxiso.org>

Click the **SuSE 8.1 live** evaluation image.

11.3 Desktop deployment options

We assume that you have already determined you are a good candidate for a Linux conversion. This means your users do not have applications that solely depend on a Windows operating system. You also have users that do not strictly depend on Microsoft Office.

Now you need to consider what solution is the best fit for a migration. There are several options to consider when deciding how to deploy Linux desktops:

- ▶ What kind of users do I have?
- ▶ How many applications are required for my daily business?
- ▶ What resources and time do I want to dedicate to administrating my workstations?
- ▶ How much control am I willing to relinquish to the end users?
- ▶ What security issues are involved and how do I manage them?

11.3.1 Thin client or PDA

A thin client is a network station, smaller than a desktop PC. It is similar to a dumb terminal because it is dependent on a server for applications and for performing the bulk of the data processing load. A thin client is a network computer without a hard disk drive. A fat client includes a disk drive and may operate independent of a server.

Possible operating systems include Windows Terminal Server, Citrix Metaframe, or Linux Terminal Server Project (LTSP).

When looking at thin clients, there are several considerations relating to their deployment:

- ▶ Deploy for users with limited number of applications
- ▶ Cost of workstation: Low
- ▶ TCO: Low (least expensive)
- ▶ Cost of software maintenance – installation and updates: Low
- ▶ Type of users: General system users
- ▶ Client type: Specific function, transaction
- ▶ Number of applications: Low
- ▶ Type of applications: Transaction-based, word processors, browser, e-mail
- ▶ Application choice for end user: None
- ▶ Probable misuse: Low vulnerability
- ▶ Exposure to viruses: Low vulnerability
- ▶ Security of workstation: Good, all data is stored on the server



Figure 11-4
Thin client

11.3.2 Fat client: Quality desktop

A fat client is a client desktop PC that performs the bulk of the data processing load. It contains a hard drive where it stores both applications and data. Although the data may be stored on the server, the client executes the applications and performs calculations.



Figure 11-5 Quality desktop

Possible operating systems include Windows 2000, XP, or the Linux Desktop.

When looking at fat clients, there are several considerations relating to their deployment:

- ▶ Deploy for users with high application requirements
- ▶ Cost of workstation: High
- ▶ TCO: High (most expensive)
- ▶ Cost of software maintenance – installation and updates: High
- ▶ Type of users: Advanced users, MIS/Tech support/developers, dual PC
- ▶ Client type: Technical, business projects, MIS, IT
- ▶ Number of applications: High
- ▶ Type of applications: Technical, animation, engineering, business, office, project
- ▶ Application choice for end user: Flexible
- ▶ Probable misuse: High vulnerability
- ▶ Exposure to viruses: High vulnerability
- ▶ Security of workstation: Average, data may be stored on a local PC

11.3.3 Inexpensive client

The inexpensive client has the same characteristics of the fat client. However, it is a low-cost alternative that may be bought in bulk quantities. It may not have the latest processor, memory, or hard drive capacity as a high quality desktop PC.



Figure 11-6 Inexpensive desktop

Possible operating systems include Windows or Linux Desktop.

When looking at inexpensive clients, there are several considerations relating to their deployment:

- ▶ Deploy for general users
- ▶ Cost of workstation: Low
- ▶ TCO: Moderate
- ▶ Cost of software maintenance – installation and updates: Moderate to high
- ▶ Type of users: General system users

- ▶ Client type: Office, documentation research
- ▶ Number of applications: High
- ▶ Type of applications:
- ▶ Application choice for end user: Flexible
- ▶ Probable misuse: High vulnerability
- ▶ Exposure to viruses: High vulnerability
- ▶ Security of workstation: Average, data may be stored on a local PC
- ▶ Misuse: Unauthorized download of programs, Web browsing, Web listening (streaming audio/video), e-mail, chat, etc.

11.3.4 Packaging and updating Linux

Although there are multiple methods to update libraries or add patches to a running Linux system, the method used by most distributions is Red Hat Package Manager (RPM). Some distributions have update techniques specific to their distribution, such as the *deb* package files used in the Debian distribution.

Most distributions also have an automatic update feature that downloads and installs the package. To update the package, you download it, unpack it, and compile the code so it becomes part of the binary (executable). With the tools now available, this is often done for you by the distribution's update and installation utility.

Red Hat and RPM

RPM and the Debian package system *Delect* are the most frequently used methods for handling packages under various Linux distributions. RPM may only be used if the program you are installing is supplied in the RPM format. It is included with several Linux distribution's including Red Hat, SUSE, and Caldera.

Under Red Hat, you can choose to install program packages via a GUI interface or from a terminal prompt using the RPM command. RPM allows you to install, uninstall, upgrade, query, and verify RPM packages.

SuSE and YOU

SuSE has an automatic update feature via the YaST2 Online Update utility (YOU). This applies to products sold with one-year maintenance.

First you register on the SuSE Web site, point your YaST utility to the Web URL as the installation source, and then start the program. For the YaST Online Update to work, you must have a direct connection to the Internet. The utility automatically downloads the required updates for your configuration and installs them.

You can use the same utility to update the SuSE Linux Office Server (SLOS).

Related publications

The publications listed in this section are considered particularly suitable for a more detailed discussion of the topics covered in this redbook.

IBM Redbooks

For information on ordering these publications, see “How to get IBM Redbooks” on page 309.

- ▶ *IBM Tivoli Monitoring Version 5.1: Advanced Resource Monitoring*, SG24-5519
- ▶ *e-Business Intelligence: Leveraging DB2 for Linux on S/390*, SG24-5687
- ▶ *AIX 5L Differences Guide Version 5.2 Edition*, SG24-5765
- ▶ *Linux for WebSphere and DB2 Servers*, SG24-5850
- ▶ *Lotus Domino R5 for Linux on IBM Netfinity Servers*, SG24-5968
- ▶ *Linux Web Hosting with WebSphere, DB2, and Domino*, SG24-6007
- ▶ *Linux Applications on pSeries*, SG24-6033
- ▶ *CICS Transaction Gateway V5 The WebSphere Connector for CICS*, SG24-6133
- ▶ *A Practical Guide to Tivoli SANergy*, SG24-6146
- ▶ *IBM WebSphere V4.0 Advanced Edition Handbook*, SG24-6176
- ▶ *Linux System Administration and Backup Tools for IBM @server xSeries and Netfinity*, SG24-6228
- ▶ *e-Business Intelligence: Data Mart Solutions with DB2 for Linux on zSeries*, SG24-6294
- ▶ *Linux Application Development Using WebSphere Studio 5*, SG24-6431
- ▶ *Tivoli Storage Manager Version 5.1 Technical Guide*, SG24-6554
- ▶ *Linux Clustering with CSM and GPFS*, SG24-6601
- ▶ *Building a Linux HPC Cluster with xCAT*, SG24-6623
- ▶ *WebSphere Application Server Test Environment Guide*, SG24-6817
- ▶ *Lotus Domino 6 for Linux*, SG24-6835
- ▶ *Disaster Recovery Strategies with Tivoli Storage Management*, SG24-6844

- ▶ *IBM WebSphere Portal V4.1 Handbook Volume 1*, SG24-6883
- ▶ *Up and Running with DB2 for Linux*, SG24-6899
- ▶ *Online Briefing: Mastering Domino for Linux*, REDP0039
- ▶ *WebSphere Commerce V5.4 for Linux, Infrastructure and Deployment Patterns*, REDP0214
- ▶ *IBM WebSphere Portal V4.1.2 in a Linux Environment*, REDP0310
- ▶ *WebSphere Application Server V4 for Linux, Implementation and Deployment Guide*, REDP0405
- ▶ *e-commerce Patterns for z/Linux Using WebSphere Commerce Suite V5.1 Patterns for e-business series*, REDP0411
- ▶ *Deploying Apache on IBM @server BladeCenter*, REDP3588
- ▶ *IBM WebSphere V5.0 for Linux, Implementation and Deployment Guide - WebSphere Handbook Series*, REDP3601
- ▶ *Tivoli NetView Performance Monitor for TCP/IP 1.4 Using Linux as the Web Server*, REDP3602

You can also find a complete listing of Linux-related IBM Redbooks on the Redbooks portal at:

<http://publib-b.boulder.ibm.com/redbooks.nsf/portals/Linux>

Other publications

These publications are also relevant as further information sources:

- ▶ *SWG Products, Solutions, and Strategies Reference Guide*, G325-2132
- ▶ *SWG Software Reference Guide*, Z325-5422

Referenced Web sites

These Web sites are also relevant as further information sources:

- ▶ IBM Linux Web site:
<http://www.ibm.com/linux/>
- ▶ The Linux Technology Center:
<http://www.ibm.com/linux/ltc>

How to get IBM Redbooks

You can order hardcopy Redbooks, as well as view, download, or search for Redbooks at the following Web site:

ibm.com/redbooks

You can also download additional materials (code samples or diskette/CD-ROM images) from that site.

IBM Redbooks collections

Redbooks are also available on CD-ROMs. Click the CD-ROMs button on the Redbooks Web site for information about all the CD-ROMs offered, as well as updates and formats.

Index

Numerics

- 10-Gigabit Ethernet 260
- 31-bit environment option 75
- 64-bit environment option 75

A

- Advanced Power Management (APM) 58
- AIX 5L 158
- AIX and Linux 161
- AIX Toolbox for Linux Applications 157
- antivirus scanners 6
- Apache 145
- application solution 133
 - iSeries 193
- ARPAnet 51
- asynchronous I/O 171
- AT&T 51
- autonomic computing 272
- availability 245

B

- backup solutions 141
- Bastille 27
- BBDO Dusseldorf 271
- Beowulf 264
- Beowulf cluster 265
- Beowulf cluster technology 169
- Beowulf program 265
- binaries 26
- binary code 49
- binary file 49
- BIND (Berkeley Internet Name Daemon) 13
- blade server 149
- BladeCenter 149
- BladeCenter HS20 149
- Blue Gene research project 248
- Blue Gene/L 248
- Bourne Again Shell (Bash) 10
- BSD license 52
- BSD UNIX 266
- Business Intelligence (BI) 16

C

- Caldera 29
- CES (Clustered Enterprise Servers) 253
- channel bonding 265
- Click-N-Run Warehouse 300
- client 17
 - fat 304
 - inexpensive 304
 - thick 17
 - thin 17, 303
- closed software 50
- cluster
 - applications 262
 - architecture 247
 - autonomic computing 272
 - availability 245
 - benefits 244
 - Beowulf 264
 - bottlenecks 247
 - Bproc 266
 - Cluster 1350 254, 268
 - Cluster 1600 253
 - Clustered Enterprise Servers (CES) 253
 - ClusterProven 268
 - commodity hardware 244, 247
 - components 254
 - coupling 248
 - distributed computing environment 248
 - emerging roles 251
 - failover 250
 - Fast Ethernet 259
 - FASTT200 260
 - Fibre Channel 260
 - file system 261
 - General Parallel File System (GPFS) 262
 - Gigabit Ethernet 259
 - grid 271
 - grid computing 272
 - heartbeat 265
 - high availability 126
 - high availability cluster 265
 - High Performance Computing (HPC) 246
 - High Throughput Computing (HTC) 248, 262
 - high-availability computing (HAC) 249

- high-performance computing 128
- history 264
- iSeries 190
- Itanium-based 269
- kernel extension 266
- Linux 126
- Linux high availability 265
- Linux Virtual Server (LVS) 266
- load balancing 127, 251, 263, 266
- load leveling 267
- management 258
- management tools 262
- massively parallel processing (MPP) 245
- message passing 246
- Moshe Bar 267
- MPP (massively parallel processing) 244
- Multicomputer Operating System for High Performance Cluster Computing (MOSIX) 266
- Myrinet 259
 - Myricom 269
- National Center for Supercomputing Applications (NCSA) 269
- networks 257
- nodes 254
- NonStop Clusters for UnixWare 267
- on demand computing 252, 272
- Open Cluster Group 267
- OpenGFS 267
- openMosix 267
- operating system 261
- parallel I/O (MOPI) 266
- Parallel Sysplex 253
- Parallel System Support Programs (PSSP) 253
- parallelism 246
- process migration 263
- protein folding 248
- pSeries 168
- radio astronomy 248
- redundancy 244
- resource sharing 246
- role of Linux clusters in emerging trends 272
- roles 246
- scalability 245
- Scyld Computing Corporation 266
- server consolidation 251
- server farm 251, 263
- Single System Image (SSI) 247, 263
- SP 253
- SteelEye Lifekeeper 271
 - support software 261
 - symmetric multiprocessor 253
 - symmetrical multiprocessing 244
 - user-mode Linux (UML) 267
- Cluster 1350 128, 254, 268
- cluster interconnect 258
- Cluster Systems Management (CSM) 262
- Clustered Enterprise Servers (CES) 253
- ClusterProven 268
- Code Versioning System (CVS) 55
- collaboration 15
- commodity off-the-shelf (COTS) 265
- common network technology 259
- Common Public License (CPL) 57
- communication types 257
- competitive software
 - data management 212
 - WebSphere 204
- compiler 48
- compute node 254, 257
- Conectiva 28
- consolidation 122, 167
 - iSeries 179, 189
- control node 255
- COTS (commodity off-the-shelf) 265
- coupling 248
- CrossOver Office 300
- CSM (Cluster Systems Management) 262
- customer references for IBM Software 225
 - Diveo Broadband Networks, Inc. 228
 - Grede Foundries, Inc. 228
 - Tamkang University 229
- Customer Relationship Management (CRM) 16

D

- data management
 - competitive software 212
 - services for Linux 283
- data storage 13
- database 14
- DB2
 - for Linux Migration Services 283
 - support for Linux 283
- DB2 Product Family 208
- DB2 UDB support for Linux 210
- DB2 Universal Database 133
- DB2 Universal Database Enterprise-Extended Edition (EEE) 264

- Debian Linux 44
- DemoLinux 302
- deployment options for desktop 302
- desktop deployment options 302
- DHCP 13
- direct I/O 183
- distributed computing environment 248
- distributed enterprise 191
- distribution 2, 19
 - iSeries Linux 188
- distro 19
- DNS (domain name server) 13
- domain name server (DNS) 13
- dynamic LPAR (DLPAR) 162

E

- E*TRADE Financial 131
- Eclipse 58
- e-mail antivirus scanners 6
- emerging cluster roles 251
- enterprise application 137
- enterprise content management (ECM) 15
- Enterprise Resource Planning (ERP) 16
- Enterprise Systems Connection (ESCON*) 253
- Enterprise X-Architecture 116–117
- eOne Group 17
- eOneCommerce 17

F

- failover 250
- Fast Ethernet 259
- FASTT200 260
- fat client 304
- Fibre Channel 14, 260
- file and print serving 13
- file structure 10
- file system 261
- firewall 13
- Free Software Foundation (FSF) 51
- Free Standards Group 27, 58
 - Linux Standard Base 58
- FSF (Free Software Foundation) 51

G

- gas dynamics research 270
- GB18030 standards 26
- General Parallel File System (GPFS) 262

- General Public License (GPL) 22, 26, 51
- Gigabit Ethernet 259
- GNOME 21, 57, 294, 296
- GNOME desktop 296
- GNOME development platform 296
- GNOME Office 296
- GNU 55
- GNU C Compiler (gcc) 10, 55
- GNU General Public License 22
- GPFS (General Parallel File System) 262
- GPL (General Public License) 22, 26
- grid computing 271–272

H

- HAC (high-availability computing) 249
- HACMP (High Availability Cluster Multiprocessing) 166
- hardware abstraction layer 10
- hardware enablement for pSeries 164
- Hardware Management Console (HMC) 163
- head node 257
- heartbeat program 265
- high availability 12
- high availability cluster 265
- High Availability Cluster Multiprocessing (HACMP) 166
- High Performance Computing (HPC) 246, 264
 - application examples 246
- High Throughput Computing (HTC) 248, 262
- high-availability cluster 126
- high-availability computing (HAC) 249
- high-availability solutions 140
- high-performance clusters services 153
- high-performance computing 128
- HMC (Hardware Management Console) 163
- host 183
- Host Bus Adapters (HBAs) 260
- HPC (High Performance Computing) 246, 264
- HPC appliance 273
- HTC (High Throughput Computing) 248, 262
- hypervisor 162–163

I

- I/O device, adapter support and availability 165
- IBM
 - Global Services 276
 - Linux Consultancy 279
 - open source 57

- open source licenses 57
 - open source software 57
 - services for Linux 275
 - Services introduction 276
 - IBM Director 122
 - IBM Global Services 276
 - IBM HTTP Server (powered by Apache) 205
 - IBM JDK Linux 169
 - IBM Linux Integration Centers 238
 - IBM Public License (IPL) 57
 - IBM Software
 - DB2 208
 - for Linux 200
 - Lotus Domino 220
 - Tivoli 215
 - IBM Support Line 276
 - IBM TotalStorage Enterprise Storage Server 260
 - IBM XL Fortran V8.1 170
 - IDS (Intrusion Detection Services) 13
 - IFS (integrated file system) 178
 - inexpensive client 304
 - InfiniBand 260
 - InfiniBlue 4X Host Channel Adapter (HCA) 261
 - Informix product family 211
 - infrastructure 12
 - infrastructure applications for iSeries 90, 193
 - infrastructure solution 144
 - Inner Circle 56
 - install node 256
 - Integrated Development Environment (IDE) 58
 - integrated file system (IFS) 178
 - Integrated xSeries Adapter 189
 - Integrated xSeries Server 189
 - integration
 - iSeries 180
 - iSeries Linux 183
 - Intel Itanium processor 269
 - interprocess communication 258
 - Intrusion Detection Services (IDS) 13
 - iSCSI 14, 260
 - iSeries 177
 - application solutions 193
 - consolidation 189
 - I/O 183
 - infrastructure applications 90, 193
 - Integrated xSeries Adapter 189
 - Integrated xSeries Server 189
 - iSeries Linux information matrix 181, 289
 - iSeries services for Linux 289
 - Linux integration 183
 - Linux services for the iSeries 289
 - models 184
 - performance 180
 - Portable Application Solutions Environment (PASE) 189
 - reliability 180
 - security 181
 - service solutions 198
 - Software Group products 195
 - solution providers 193
 - total cost of ownership 179
 - iSeries distributions
 - Red Hat 189
 - SuSE 189
 - Turbolinux 189
 - iSeries solutions 196
 - iSeries Test Drive for Linux 90, 193
 - ISV and application solutions 133
 - Itanium 269
 - Itanium-based cluster 269
- J**
- J.D. Edwards 17
 - jam session 5
 - Java 48
 - iSeries 195
 - JBoss 205
 - JDK Linux 169
 - JRun 205
- K**
- K Desktop Environment (KDE) 21, 58, 294–295
 - KDE (K Desktop Environment) 21
 - kernel 2, 10, 44–45, 55
 - Knoppix 302
 - Knowledge discovery (KD) 15
 - knowledge management (KM) 15
- L**
- LCS (Linux Virtual Server) 266
 - legacy applications 282
 - Li18nux standard 26
 - license
 - closed 50
 - free software 51
 - shareware 50

- software 49
- X 52
- Windows 299
- line-of-business applications 16
 - iSeries 91, 194
- Linus Torvalds 45, 54
- Linux
 - AIX 5L 158
 - application solutions 132
 - commodity off-the-shelf (COTS) 265
 - high availability 58, 265
 - how it runs on the pSeries server 156
 - hybrids 299
 - I/O flexibility 183
 - IBM Software 200
 - IBM Software on the pSeries server 169
 - Integrated Feature for 282
 - iSeries integration 183
 - load balancing 266
 - on the pSeries 670 162
 - on the pSeries 690 162
 - on the pSeries LPAR 162
 - operating system 9
 - pSeries sources of information 166
 - RAS capabilities 166
 - rating or ranking 61
 - relationship with AIX 164
 - role of emerging clusters 272
 - scalability on the pSeries 165
 - updates and packaging 305
 - Web site platform matrix 291
 - WebSphere 202
- Linux cluster 126
- Linux distribution 19
- Linux for S/390 (31-bit) 75
- Linux on xSeries IBM Software 132
- Linux on zSeries (64-bit) 75
- Linux Solution for e-business 152
- Linux Standards Base 27
- Linux Technology Center (LTC) 57
- Linux Virtual Server (LVS) 266
- Linux, UNIX, and Windows 59
 - side-by-side comparison 62
 - strengths and weaknesses 64
 - total cost of ownership 65
- LinuxBIOS Project 11
- Linux-capable pSeries model structure 156
- Linux-related WebSphere resources 206
- LinuxWorld Conference and Expo 268

- load balancing 251, 263, 266
- load leveling 267
- load-balancing cluster 127
- Lotus
 - IBM's Linux services for 287
 - Lotus Products' Linux support matrix 287
- LPAR 162–163
 - iSeries 182
- LTC (Linux Technology Center) 57

M

- management node 256
- management operations 258
- management tools 262
- manager jam 5
- massively parallel processing (MPP) 244–245
- metacluster 272
- middleware 14
- MIT 52
- monitoring 6
- MOPI 266
- MOSIX (Multicomputer Operating System for High Performance Cluster Computing) 266
- MPP (massively parallel processing) 244–245
- Multicomputer Operating System for High Performance Cluster Computing (MOSIX) 266
- Myricom 169
- Myrinet 169, 259

N

- NASA 264
- National Center for Supercomputing Applications (NCSA) 269
- National Digital Mammography Archive (NDMA) 271
- NCSA (National Center for Supercomputing Applications) 269
- NDMA (National Digital Mammography Archive) 271
- Netscape 52
- networks 257
- node 254
 - compute 254, 257
 - control 255
 - function 254
 - head 257
 - install 256
 - management 256

- storage 255, 257
- type 256
- user 256
- non-IBM middleware 136
- NonStop Clusters for UnixWare 267

O

- on demand computing 252, 272
- Open Cluster Group 267
- open source 21, 51
 - IBM 57
 - IBM licenses 57
 - IBM software 57
 - initiative 52
 - software 48
- Open Source Cluster Application Resources (OS-CAR) 57, 190, 267
- Open Source Definition (OSD) 52
- Open Source Developer Lab (OSDL) 57
- Open Source Initiative (OSI) 58
- OpenLDAP 58
- openMosix 267
- OpenOffice 195, 298
- OS/400 178
 - DB2 UDB 178
 - IFS 178
 - V5R1 178, 181
 - V5R2 182
- OSCAR (Open Source Cluster Application Resources) 190, 267

P

- parallel computing 246
- parallel I/O 266
- Parallel Sysplex 253
- Parallel System Support Programs (PSSP) 253
- parallelism 246
- partitioning 181
 - logical 178
- PCI hot plug 58
- PDA 303
- performance monitoring 7
- performance of iSeries 180
- Personal Digital Assistants (PDAs) 18
- personal Linux education plan 288
- plug-in 58
- PowerPC 164
 - iSeries 188–189, 193

- PowerPC architecture 164, 171
- PowerPC processor 164
- PPC-32 58
- PPC-64 58
- process migration 263
- protein folding 248
- pSeries
 - consolidation 167
 - hardware 164
 - introduction 155
 - Linux application solutions 169
 - Linux clusters 168
 - Linux for pSeries offerings 290
 - Linux running on 156–157
 - operating system scenarios 157
 - running Linux 156
 - scalability 165
 - services for Linux 176
 - software for Linux 169
 - solutions on Linux 176
- PSSP (Parallel System Support Programs) 253

R

- RAID 14
- RAS (Reliability Availability Serviceability) 166
- rating or ranking Linux, UNIX, and Windows 62
- Red Hat 23
- Red Hat Package Manager (RPM) 23, 305
- Redbooks Web site 309
 - Contact us xx
- Redflag 43
- redundancy 244
- reliability, iSeries 180
- router 13
- RPM (Red Hat Package Manager) 23, 305
- RS/6000 253

S

- Samba 58, 144
- SAP 17
- scalability 245
- scalability on the pSeries 165
- Scyld Computing Corporation 266
- security 137
 - iSeries 181
- server consolidation 12, 122, 167, 189, 251
- server consolidation service 153
- server farm 263

- service processor 256
- services
 - and IBM Software 281
 - IBM Managed Hosting 288
 - IBM's Linux Learning services 288
 - IBM's Linux Services and support Web sites 280
- shareware software 50
- shell 10
- Single System Image (SSI) 247, 263
- Single Systems Image project 267
- SLES (SuSE Linux Enterprise Server) 37
- SMP (symmetric multiprocessing) 244
- software for Linux 169
- Software Group products 195
- software installation 258
- solution providers for iSeries 193
- source code 48
- SP 253
- SSI (Single System Image) 247, 263
- StarOffice 298
- static LPAR 162
- SteelEye LifeKeeper 271
- storage access 258
- storage node 255, 257
- Stream Control Transmission Protocol 58
- Supply Chain Management (SCM) 16
- support
 - IBM's Linux Services and support Web sites 280
 - IBM's Linux Support Line 287
- Support Line 276
- SuSE 37
- SuSE Linux Enterprise Server (SLES) 37
- symmetric multiprocessing (SMP) 244
- symmetric multiprocessor 253
- systems management 15

T

- TCOnow! 179
- TCSH shell 10
- Telecom Application Server 264
- thick client 17
- thin client 17, 303
- Tivoli 15, 133
 - IBM Tivoli Workload Scheduler 285
 - Linux supported Tivoli products 284
 - on Linux 215

- Tivoli Access Manager for e-business 284
- Tivoli Configuration and Operations Linux support 285
- Tivoli Configuration Manager 285
- Tivoli Enterprise Console 286
- Tivoli Enterprise Console Product Matrix for Linux 286
- Tivoli for Linux 284
- Tivoli Monitoring 286
- Tivoli NetView 287
- Tivoli Performance and Availability Management Products 286
- Tivoli Security Management Linux support 284
- Tivoli Service Level Advisor 287
- Tivoli Storage Management Linux support 285
- Tivoli Storage Manager 285
- Tivoli for Linux services 284
- Tomcat 205
- total cost of ownership 65, 179
 - findings 65, 68
- TotalStorage Enterprise Storage Server 260
- TotalStorage IP Storage 200i 14, 260
- TPC-H 209

U

- U2 product family 212
- UML (user-mode Linux) 267
- UnitedLinux 25
- Universal Database 14
- UNIX 15, 51, 54, 266
 - rating or ranking 61
- Update Xpress 122
- USB 58
- USENIX 58
- user node 256
- user-mode Linux (UML) 267

V

- virtual I/O 183
- VMware ESX Server 124

W

- Web serving 13
- Web-enablement applications for iSeries 91, 194
- WebSphere 15, 202
 - competitive software 204
 - product services matrix for Linux 282

- resources for Linux 206
- services for Linux 282
- software platform 132
- WebSphere Application Server 17, 264
- WebSphere Commerce Suite 264
- WebSphere Developer Domain 206
- WebSphere Edge Server 264
- WebSphere Everyplace Suite 264
- Why Linux 179
 - iSeries 179
- window manager 294
- Windows 17
 - rating or ranking 61
- Wine 301
- WinLinux 301
- workload consolidation 122

X

- X Window System 294
- X-Architecture 116
- xCat 262
- Ximian 296
- XL Fortran V8.1 170
- xSeries
 - architected solutions 149
 - services solutions 151
 - xSeries Linux information matrix 118, 290
 - xSeries services for Linux 290
- xSeries 300 262

Y

- YaST2 Online Update (YOU) 305
- Yellow Dog Linux 42
- YOU (YaST2 Online Update) 305

Z

- zSeries
 - zSeries Linux information matrix 74, 289
 - zSeries services for Linux 288



Linux Handbook: A Guide to IBM Linux Solutions and Resources

(0.5" spine)
0.475" <-> 0.875"
250 <-> 459 pages



Linux Handbook

A Guide to IBM Linux Solutions and Resources



Redbooks

Discover the IBM Software solutions for Linux-based servers

Understand how Linux operates on each IBM @server platform

Learn about IBM resources for implementing Linux solutions

IBM has a strong commitment to Linux as an operating system and Linux-based solutions. This IBM Redbook gives a broad understanding of IBM products and how they relate to the Linux operating system. It tries to encapsulate the IBM Software Solutions that are available for Linux-based servers. It also describes the Linux enablement throughout the IBM @server™ product line.

This IBM Redbook offers enough information to give you a high-level understanding of Linux. Therefore, it is not meant to be a highly technical document. It also provides pointers to further sources of information.

This redbook is intended for IBM representatives, Business Partners, and Customers who are involved in considering and planning Linux-based software solutions. You will find this redbook particularly helpful if you are planning to test or implement Linux in an enterprise.

INTERNATIONAL TECHNICAL SUPPORT ORGANIZATION

BUILDING TECHNICAL INFORMATION BASED ON PRACTICAL EXPERIENCE

IBM Redbooks are developed by the IBM International Technical Support Organization. Experts from IBM, Customers and Partners from around the world create timely technical information based on realistic scenarios. Specific recommendations are provided to help you implement IT solutions more effectively in your environment.

For more information:
ibm.com/redbooks