

A Practical Approach to Cloud IaaS with IBM SoftLayer Presentations Guide

Daniel Aguado
Thomas Andersen
Aram Avetisyan
Jeff Budnik
Mihai Criveti
Adrian Doroiman
Andrew Hoppe
Gerardo Menegaz
Alejandro Morales
Adrian Moti
Marie Joy Salazar
Sebastian Szumczyk



In partnership with
IBM MEA University Program



International Technical Support Organization

**A Practical Approach to Cloud IaaS with IBM SoftLayer:
Presentations Guide**

February 2016

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

First Edition (February 2016)

This edition applies to IBM SoftLayer cloud infrastructure at the time this guide was developed in July 2015.

© Copyright International Business Machines Corporation 2016. 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
IBM Redbooks promotions	xiii
Preface	xv
Authors	xv
Now you can become a published author, too!	xvii
Comments welcome	xviii
Stay connected to IBM Redbooks	xviii
Unit 1. Understanding the IaaS approach	1
1.1 What you should be able to do	2
1.2 References	2
1.3 Teaching topics	2
1.4 What is Cloud Computing?	3
1.5 Definition of Cloud Computing	4
1.6 Essential characteristics	5
1.7 Deployment models	6
1.8 Service models	7
1.9 Cloud service models in practice	8
1.10 Cloud service models – Technology mapping	10
1.11 Cloud adoption and migration is workload driven	11
1.12 Infrastructure as a Service	13
1.13 Example IaaS use case: Website publishing scenario	14
1.14 Platform as a Service	15
1.15 Software as a Service	16
1.16 Implications of Cloud Computing	17
1.17 Business benefits of Cloud Computing	18
1.18 Development benefits of Cloud Computing	19
1.19 Industry impact	20
1.20 IaaS: Common storage models	21
1.21 Overview of storage options	22
1.22 IaaS: SoftLayer public cloud	24
1.23 IaaS products and services overview (SoftLayer)	25
1.24 Key takeaways: Cloud Computing	26
Unit 2. Platform architecture for cloud computing	27
2.1 What you should be able to do	28
2.2 Teaching topics	28
2.3 Typical IaaS offerings	29
2.4 What an IaaS provider could offer	30
2.5 Location location location	31
2.6 SoftLayer presence	32
2.7 The datacenter	33
2.8 Network architecture	34
2.9 Triple network architecture	35
2.10 The Orchestration Management System	36
2.11 Sample IaaS datacenter	37

2.12 Checkpoint	38
2.13 Checkpoint (2)	39
2.14 Management web interface	40
2.15 Management web interface capabilities	41
2.16 Monitor/control/change instances from the web interface	42
2.17 Monitor/control/change instances	43
2.18 Account administration	44
2.19 Adding users to your account	45
2.20 Defining permissions for users	46
2.21 Default permission templates	48
2.22 Defining instance access	49
2.23 Checkpoint	50
2.24 Checkpoint (2)	51
Unit 3. Server offerings in cloud computing	53
3.1 What you should be able to do	54
3.2 References	54
3.3 Teaching topics	54
3.4 Compute Nodes in IaaS Cloud	55
3.5 IaaS cloud server offerings (SoftLayer)	56
3.6 Virtual Servers - Public node	57
3.7 Virtual Servers - Public node (2)	58
3.8 Virtual Servers - Private node	59
3.9 All virtual servers:	60
3.10 Bare metal servers	61
3.11 Bare metal servers with hourly billing	62
3.12 Bare metal servers with monthly billing	63
3.13 Bare metal servers with monthly billing (continued)	64
3.14 All server offerings:	65
3.15 Checkpoint questions	66
3.16 Checkpoint questions (2)	67
3.17 Ordering servers	68
3.18 Provisioning server instances	69
3.19 Inspecting server details	70
3.20 Upgrading or downgrading server instances	71
3.21 Accessing server instances	72
3.22 Canceling server instances	73
3.23 Checkpoint questions	74
3.24 Checkpoint questions (2)	75
Unit 4. Storage options in cloud computing	77
4.1 What you should be able to do	78
4.2 References	78
4.3 Teaching topics	78
4.4 Introduction to storage: Storage types	79
4.5 Introduction to storage: DAS	80
4.6 Introduction to storage: DAS (2)	81
4.7 Introduction to storage: SAN	83
4.8 Introduction to storage: SAN (2)	84
4.9 Introduction to storage: NAS	86
4.10 Introduction to storage: NAS (2)	87
4.11 Introduction to storage: Object Storage	88
4.12 Introduction to storage: RAID arrays	89

4.13 iSCSI: Components	91
4.14 iSCSI addressing	92
4.15 iSCSI access control	93
4.16 NFS components	94
4.17 NFS addressing	95
4.18 NFS access control	96
4.19 SoftLayer storage offerings: DAS	97
4.20 SoftLayer storage offerings: Block storage	99
4.21 SoftLayer storage offerings: File storage	100
4.22 SoftLayer storage offering: Storage options	101
4.23 SoftLayer storage offering: Storage options (2)	102
4.24 SoftLayer storage offering: Storage options (3)	103
4.25 SoftLayer storage offering: Object storage	104
4.26 Overview	105
4.27 Check point: Questions	106
4.28 Check point: Questions and answers	107
4.29 Recap	108
Unit 5. Networking options in cloud computing	109
5.1 What you should be able to do	110
5.2 References	110
5.3 Teaching topics	110
5.4 What is cloud computing?	111
5.5 The network is critical to Cloud Computing	112
5.6 Networking had to change	113
5.7 Attributes of Cloud Networking	114
5.8 Example: Data center switch network architecture	115
5.9 Virtual Ethernet switch in a virtualized server environment	116
5.10 Networking overview	117
5.11 Networking overview (2)	119
5.12 Networking overview (3)	120
5.13 Learning about public networks	122
5.14 Learning about public networks (2)	123
5.15 Learning about public networks - Carriers	124
5.16 Understanding the primary features	125
5.17 Understanding bandwidth	126
5.18 Learning about private networks	127
5.19 Learning about private networks (2)	128
5.20 Learning about private networks (3)	129
5.21 Learning about management network	130
5.22 SoftLayer network architecture	131
5.23 Reviewing SoftLayer network tooling	132
5.24 Using Looking Glass, SoftLayer's IP backbone	133
5.25 Using SoftLayer Looking Glass, SoftLayer's IP backbone	134
5.26 Using SoftLayer Looking Glass, SoftLayer's IP backbone (2)	135
5.27 Using other networking tools	136
5.28 Using other networking tools (2)	137
5.29 Using other networking tools (3)	138
5.30 Using other networking tools from SoftLayer Control Panel	139
5.31 Managing VPN connections to SoftLayer	140
5.32 General concept of VPN	141
5.33 Managing VPN connections to SoftLayer	142
5.34 Managing VPN connections to SoftLayer (2)	143

5.35	Managing VPN connections to SoftLayer (3)	145
5.36	Managing VPN connections to SoftLayer (4)	146
5.37	Direct Link use case	147
5.38	Direct Link use case (continued)	149
5.39	Direct Link use case (continued) (2)	151
5.40	Recap	152
5.41	Checkpoint questions	153
5.42	Checkpoint questions (2)	154
5.43	Checkpoint questions (3)	155
5.44	Checkpoint questions (4)	156
5.45	Checkpoint questions (5)	157
5.46	Checkpoint questions (6)	158
5.47	Checkpoint questions (7)	159
5.48	Checkpoint questions (8)	160
Unit 6. Managing a simple IaaS environment		161
6.1	What you should be able to do	162
6.2	Teaching topics	162
6.3	Basics of cloud infrastructure and components	163
6.4	Cloud service model overview	164
6.5	Platform architecture	165
6.6	Locations of platform architecture	166
6.7	Types of servers	167
6.8	Server options	168
6.9	Storage types and protocols	169
6.10	Other storage types	170
6.11	Network types	171
6.12	Basics of cloud infrastructure and components	172
6.13	Conclusion of recap	173
6.14	Technologies enabling transformation of infrastructure	174
6.15	Overview of technologies	175
6.16	Virtualization	176
6.17	Hypervisor	177
6.18	Provisioning	178
6.19	Managing cloud infrastructure	179
6.20	Cloud infrastructure IaaS key points	180
6.21	Managing cloud infrastructure in IaaS	181
6.22	Checkpoint	182
Unit 7. Server offerings in cloud computing: Advanced topics		183
7.1	What you should be able to do	184
7.2	Teaching topics	184
7.3	Introducing the image template concept	185
7.4	Image template types in SoftLayer	186
7.5	Image template types in SoftLayer: Standard images	187
7.6	Image template types in SoftLayer: Flex Images	188
7.7	Image template types in SoftLayer	189
7.8	Creating image templates in SoftLayer	190
7.9	Creating image templates in SoftLayer (2)	191
7.10	Sharing, finding, and deploying the image templates	192
7.11	Sharing, finding, and deploying the image templates (2)	193
7.12	Sharing, finding, and deploying the image templates (3)	194
7.13	Sharing, finding, and deploying the image templates (4)	195

7.14 Provisioning scripts	196
7.15 Usage scenarios in SoftLayer	197
7.16 Usage scenarios in SoftLayer (2)	198
7.17 Checkpoint questions	199
7.18 Checkpoint questions (2)	200
7.19 Checkpoint questions (3)	201
7.20 Checkpoint questions (4)	202
Unit 8. Storage options in cloud computing: Advanced topics	203
8.1 What you should be able to do	204
8.2 References	204
8.3 Teaching topics	204
8.4 Backup and recovery: Concepts	205
8.5 Backup and recovery: Concepts (2)	206
8.6 Backup and recovery solutions in SoftLayer	207
8.7 Snapshots and replication of Endurance storage	209
8.8 Dedicated storage	210
8.9 Dedicated storage: OS NEXUS QuantaStor	211
8.10 Check point: Questions	212
8.11 Check point: Questions answered	213
Unit 9. Networking options in cloud computing: Advanced topics	215
9.1 What you should be able to do	216
9.2 References	216
9.3 Teaching topics	216
9.4 SoftLayer network topology	218
9.5 IP addresses in SoftLayer	219
9.6 IP addresses in SoftLayer (2)	222
9.7 Five steps to start using IPv6 in SoftLayer	223
9.8 Separating devices and subnets with VLANs	225
9.9 VLAN spanning	226
9.10 VLAN spanning (2)	227
9.11 Load balancing fundamentals	228
9.12 Load balancing fundamentals (2)	229
9.13 Load balancing options in SoftLayer	230
9.14 Load balancing options in SoftLayer (2)	231
9.15 Load balancing solutions in SoftLayer use case	232
9.16 SoftLayer Content Delivery Network	233
9.17 EdgeCast CDN locations	235
9.18 Example SoftLayer CDN users	236
9.19 SoftLayer Content Delivery Network use case	237
9.20 SoftLayer Content Delivery Network implementation guide	238
9.21 SoftLayer Content Delivery Network use case	239
9.22 SoftLayer Content Delivery Network use case (2)	240
9.23 SoftLayer Content Delivery Network use case (3)	241
9.24 SoftLayer Content Delivery Network use case (4)	242
9.25 SoftLayer Content Delivery Network use case (5)	243
9.26 SoftLayer Content Delivery Network use case (6)	244
9.27 SoftLayer Content Delivery Network use case (7)	245
9.28 SoftLayer Content Delivery Network use case (8)	246
9.29 SoftLayer Content Delivery Network use case (9)	247
9.30 SoftLayer Content Delivery Network use case (10)	248
9.31 Vyatta appliance	249

9.32	Recap	251
9.33	Checkpoint questions	252
9.34	Checkpoint questions (2)	253
9.35	Checkpoint questions (3)	254
9.36	Checkpoint questions (4)	255
9.37	Checkpoint questions (5)	256
9.38	Checkpoint questions (6)	257
9.39	Checkpoint questions (7)	258
9.40	Checkpoint questions (8)	259
9.41	Checkpoint questions (9)	260
9.42	Checkpoint questions (10)	261
9.43	Introduction to OSI model	262
9.44	Understanding TCP/IP addressing and subnetting basics	264
Unit 10. Securing cloud environments		267
10.1	What you should be able to do	268
10.2	References	268
10.3	Teaching topics	268
10.4	Why security is important	269
10.5	Data center security	271
10.6	Additional security offerings	273
10.7	Securing communications	274
10.8	Securing instances using firewalls	275
10.9	Securing instances using firewalls (2)	276
10.10	Hardware firewalls	277
10.11	Appliance firewalls	278
10.12	OS firewalls	279
10.13	IaaS provider firewalls	280
10.14	Checkpoint	282
10.15	Checkpoint (2)	283
10.16	Securing instances using OS hardening	284
10.17	Using portal to set up and verify security	286
10.18	Administering firewall from the portal	287
10.19	Verify security in the portal	288
10.20	Checkpoint	289
10.21	Checkpoint (2)	290
Unit 11. Monitoring cloud environments		291
11.1	What you should be able to do	292
11.2	References	292
11.3	Teaching topics	292
11.4	Typical service models responsibilities and typical service access to infrastructure	294
11.5	Typical core infrastructure monitoring and typical infrastructure components monitored	295
11.6	Why monitoring is important	296
11.7	Typical monitoring tool limitations	297
11.8	How typical monitoring tools work	298
11.9	Typical monitoring alert	299
11.10	Typical monitoring response	300
11.11	Practical approach to SoftLayer monitoring	301
11.12	Ordering a monitoring service	303
11.13	Building your own simple monitoring solution	304
11.14	Upgrading a monitoring package	305

11.15	Cancelling a monitoring package	306
11.16	Checkpoint questions	307
11.17	Checkpoint questions (2)	308
Unit 12. Automating cloud management with APIs		309
12.1	What you should be able to do	310
12.2	References	310
12.3	Teaching topics	310
12.4	Introducing the API concept	311
12.5	API advantages	312
12.6	API advantages (2)	313
12.7	Hybrid cloud scenarios that use the API economy	314
12.8	SoftLayer API overview	315
12.9	SoftLayer API overview (2)	316
12.10	The SoftLayer API structure	317
12.11	The SoftLayer API structure (2)	318
12.12	The main SoftLayer library	319
12.13	The Object Store API	320
12.14	The Message Queue API	321
12.15	Using the SoftLayer API	322
12.16	Using the SoftLayer API: Services and methods	323
12.17	Using the SoftLayer API: Data types	324
12.18	Using the SoftLayer API: Service hierarchy sample	325
12.19	Using the SoftLayer API	326
12.20	Using the SoftLayer API (cont.)	327
12.21	Example scenarios of using an API	328
12.22	Checkpoint questions	329
12.23	Checkpoint questions (2)	330
12.24	Checkpoint questions (3)	331
12.25	Checkpoint questions (4)	332
12.26	Checkpoint questions (5)	333
12.27	Recap	334
Glossary		335

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

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 grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, MD-NC119, Armonk, NY 10504-1785, US

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 jurisdictions 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 websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

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

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

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 actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate 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. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

AIX®	IBM®	Redbooks®
Bluemix®	IBM Cloud Managed Services™	Redbooks (logo)  ®
developerWorks®	IBM SmartCloud®	Tivoli®
Global Business Services®	IMS™	Unyte®
Global Technology Services®	Rational®	

The following terms are trademarks of other companies:

Adobe, the Adobe logo, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

CloudLayer, Flex Images, KnowledgeLayer, SoftLayer, and SoftLayer device are trademarks or registered trademarks of SoftLayer, Inc., an IBM Company.

Intel, Intel logo, Intel Inside logo, and Intel Centrino logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Microsoft, Windows, 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 Oracle and/or its affiliates.

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

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

Find and read thousands of IBM Redbooks publications

- ▶ Search, bookmark, save and organize favorites
- ▶ Get personalized notifications of new content
- ▶ Link to the latest Redbooks blogs and videos

Get the latest version of the Redbooks Mobile App



Promote your business in an IBM Redbooks publication

Place a Sponsorship Promotion in an IBM® Redbooks® publication, featuring your business or solution with a link to your web site.

Qualified IBM Business Partners may place a full page promotion in the most popular Redbooks publications. Imagine the power of being seen by users who download millions of Redbooks publications each year!



ibm.com/Redbooks

About Redbooks → Business Partner Programs

THIS PAGE INTENTIONALLY LEFT BLANK

Preface

This IBM® Redbooks® publication is based on the Presentations Guide of the course *A Practical Approach to Cloud IaaS with IBM SoftLayer*, which was developed by the IBM Redbooks team in partnership with IBM Middle East and Africa University Program. This course is designed to teach university students how to build a simple infrastructure as a service (IaaS) cloud environment based on IBM SoftLayer®. It provides students with the fundamental skills to design, implement, and manage an IaaS cloud environment using the IBM SoftLayer platform as an example.

The primary target audience for this course is university students in undergraduate computer science and computer engineer programs with no previous experience working in cloud environments. However, anyone new to cloud computing can benefit from this course.

The workshop materials were created in July 2015. Thus, all IBM SoftLayer features discussed in this Presentations Guide are current as of July 2015.

Authors

This course was produced by a team of specialists from around the world working at the International Technical Support Organization, Raleigh Center.

Daniel Aguado is a Cloud Technical Sales professional in the IBM Cloud business unit in Madrid, Spain. His areas of expertise include IBM SoftLayer and IBM Cloud Managed Services™. Daniel designs solutions for clients based on these cloud offerings. He joined IBM in 2013.

Thomas Andersen is a certified IT specialist, technical team leader, cloud administrator, and SoftLayer solution designer in the IBM Development Support Team (DST) organization in IBM Denmark. Thomas has been with IBM for 20 years filling different roles such as development, services, and infrastructure support. In the last years, his focus has been virtualization and cloud technologies.

Aram Avetisyan is a Cloud Infrastructure Architect in IBM who has over 14 years of experience in the IT industry. He has a rich background and expertise in several technologies such as virtualization, operating systems administration, and disaster recovery. Aram joined IBM in 2011, and is based in the Czech Republic. As a certified instructor, Aram delivers IT courses to IBM employees around the world. Aram is an active blogger and contributor to the cloud community. For his contributions, Aram was awarded VMware vExpert accreditation in 2014 and 2015.

Jeff Budnik is a Certified IT Specialist and Cloud Solution Architect in the IBM DST organization. Jeff has over 18 years of experience in infrastructure services with special focus on cloud computing. As a cloud solution architect, Jeff helps clients to identify requirements and design solutions based on the IBM SoftLayer technology. Jeff joined IBM in 1998 and is based in the United States.

Mihai Criveti is a Technical Sales Leader in IBM Ireland. Mihai supports large sales opportunities for IBM Cloud. He has over 10 years of experience in the IT industry. Mihai holds a degree in Managerial Informatics. His areas of expertise include cloud computing,

business analysis, technical sales, project management, agile software development, DevOps, and UNIX. He has written extensively on cloud computing topics.

Adrian Doroiman is a Project Manager in IBM Romania. He has 12 years of experience in the IT industry. Adrian holds degrees in Computer Science and Economics. His areas of expertise include cloud computing, project management, agile software development, DevOps, and cryptography.

Dr. Andrew Hoppe is a Senior Software Engineer and Technical Architect in IBM Cloud Services, Business Partner Sales team. Andrew's focus is on SoftLayer sales support. He came to IBM with the Rational® acquisition in 2003. Andrew has over 20 years of experience in software design and development, specializing in object-oriented modeling and multitier business system implementation. Andrew is also an active educator, with 15 years of teaching experience at universities in several countries. He has published papers on research conferences and in trade journals, and is the author of several IBM developerWorks® articles, blogs, and technical materials on several IBM cloud topics. He is based in Raleigh, North Carolina, US. Andrew holds a Ph.D. degree from the University of Warsaw, Poland.

Gerardo Menegaz is a Chief Architect in the IBM Global Technology Services® business unit. Gerardo has over 20 years of IT leadership experience formulating strategies and using proprietary technologies in fast-paced, challenging environments. Gerardo has led projects to develop business strategies across multiple industries and new technologies such as cloud computing, mobile, bring your own device (BYOD), big data, analytics, and social media. He also has extensive experience on data center topics such as server consolidation, virtualization and optimization technologies, methodologies, application rationalization, and logical consolidation techniques. Gerardo has published several papers and blogs. He is a graduate of the University of California, Santa Barbara.

Alejandro Morales is an Infrastructure Specialist in IBM Global Business Services® (GBS), IBM Mexico. He has experience in software development and infrastructure support with focus on cloud computing. As an operations technical leader in the IBM DST, Alejandro helps clients with their day-to-day requirements in the IBM SoftLayer IaaS environment.

Adrian Moti is a Project Manager working in Romania for IBM Cloud Managed Services, the private cloud offering of IBM. He has four years of experience in cloud computing and 10+ years experience in IT. He holds degrees in Computer Science and Electronics in Transportation. His areas of expertise include cloud computing, project management, agile software development and customer support in IT.

Marie Joy Salazar is the team leader for the SoftLayer Operations team in the IBM DST, IBM Philippines. Marie Joy was part of the IBM Internship Program in her college days, and she has now been with IBM for five years. She has taken several roles in asset management, systems operations, and infrastructure operations. Her current responsibilities include process improvement and operating system, security, and middleware support for IBM internal and external clients that are hosted in IBM SoftLayer.

Sebastian Szumczyk is a Cloud Advisor in IBM Poland. His primary focus is cloud infrastructure solution design. His current responsibilities include cloud technical sales, technical feasibility studies of complex cloud solution proposals, and transforming client requirements into technical solution designs. His areas of expertise include IBM Tivoli® Storage and System Management, pSeries, IBM AIX®, virtualization, storage, high availability solutions, and cloud solution for enterprise systems. He holds several technical certifications including the prestigious IBM System p Certified Advanced Technical Expert and IBM Architect Accreditation certificates. Sebastian is an author of the IBM Redbooks publications *IBM Information Infrastructure Solutions Handbook*, SG24-7814 and *IBM System Storage Solutions Handbook*, SG24-5250.

The project that produced this course material was managed by **Vasfi Gucer** and **Marcela Adan**, IBM Redbooks Project Leaders, Global Content Services.

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

Lysa Banks
IBM Cloud

Aaron Morris
Network & SoftLayer Automation, IBM Global Business Services

Juan Pablo Napoli
IBM MEA University Programs

David Rhinehart
IBM Cloud Service Technical Sales

Priti Agarwal
Jerry Amal
Yashaswani Aschar
Deana Coble
Shari Deiana
Nagasandesh Narahari Rao
Sreya Sarkar
Shawn Tooley
IBM Redbooks video production team

Mayowa Adeoti
IBM MEA University Programs Internship

Daniel Nussbaummüller
Anja Jessica Paessler
University Program Internship, IBM Germany

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

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 books to be as helpful as possible. Send us your comments about this book or other IBM Redbooks publications in one of the following ways:

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

ibm.com/redbooks

- Send your comments in an email to:

redbooks@us.ibm.com

- Mail your comments to:

IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400

Stay connected to IBM Redbooks

- Find us on Facebook:

<http://www.facebook.com/IBMRedbooks>

- Follow us on Twitter:

<http://twitter.com/ibmredbooks>

- Look for us on LinkedIn:

<http://www.linkedin.com/groups?home=&gid=2130806>

- Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:

<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>

- Stay current on recent Redbooks publications with RSS Feeds:

<http://www.redbooks.ibm.com/rss.html>



Understanding the IaaS approach

This unit covers the following topics:

- ▶ What is cloud computing?
 - Essential Characteristics
 - Deployment Models
 - Service Models
- ▶ Implications of Cloud Computing
- ▶ IaaS: Common Storage Models
- ▶ IaaS: Softlayer Public Cloud

1.1 What you should be able to do

After completion of this unit, you should be able to:

- ▶ Describe the basics of cloud computing, including its essential characteristics, deployment models, and service models.
- ▶ Recognize the implications of changing to a cloud computing approach.
- ▶ Describe common storage models used in IaaS.
- ▶ Identify IaaS concepts in the SoftLayer Public Cloud offerings.

1.2 References

The following publications are useful for further research on the topic presented in this unit:

- ▶ National Institute of Standards and Technology - Special Publication 800-145:
<http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>
- ▶ *IBM CS-101 Introduction to Cloud*
- ▶ Disruptive innovation according to Wikipedia:
https://en.wikipedia.org/wiki/Disruptive_innovation

1.3 Teaching topics

This unit covers these key topics:

- ▶ Essential Characteristics
 - On-demand self-service
 - Broad network access
 - Resource pooling
 - Rapid elasticity
 - Measured service
- ▶ Deployment Models
 - Private
 - Community
 - Public
 - Hybrid
- ▶ Service Models
 - Software as a Service (SaaS)
 - Platform as a Service (PaaS)
 - Infrastructure as a Service (IaaS)
- ▶ Adoption
 - Cloud adoption and migration is workload driven
 - Different workloads are a better fit for different service models and deployment models

1.4 What is Cloud Computing?



Figure 1-1 What is Cloud Computing?

Notes:

(none)

1.5 Definition of Cloud Computing

IBM Digital Sales International Technical Support Organization and Authoring Services IBM

Definition of Cloud Computing*

Model for enabling ubiquitous, convenient, on-demand network access

to a shared pool of configurable computing resources

that can be rapidly provisioned and released with minimal effort or service provider interaction.

* National Institute of Standards and Technology - Special Publication 800-145:
<http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf>

© 2015 IBM Corporation Redbooks Workshop

Figure 1-2 Definition of Cloud Computing

Notes:

The term *cloud* is an abstraction for the complex infrastructure it conceals. The generally accepted definition of cloud computing comes from the National Institute of Standards and Technology (NIST). The NIST definition runs to several hundred words but essentially says that:

"Cloud Computing is a model for enabling **convenient, on-demand network access** to a **shared pool of configurable computing resources** that can be **rapidly provisioned** and released with **minimal management effort** or service provider interaction."

Examples of computing resources include:

- ▶ Networks
- ▶ Servers
- ▶ Storage
- ▶ Applications
- ▶ Services

1.6 Essential characteristics

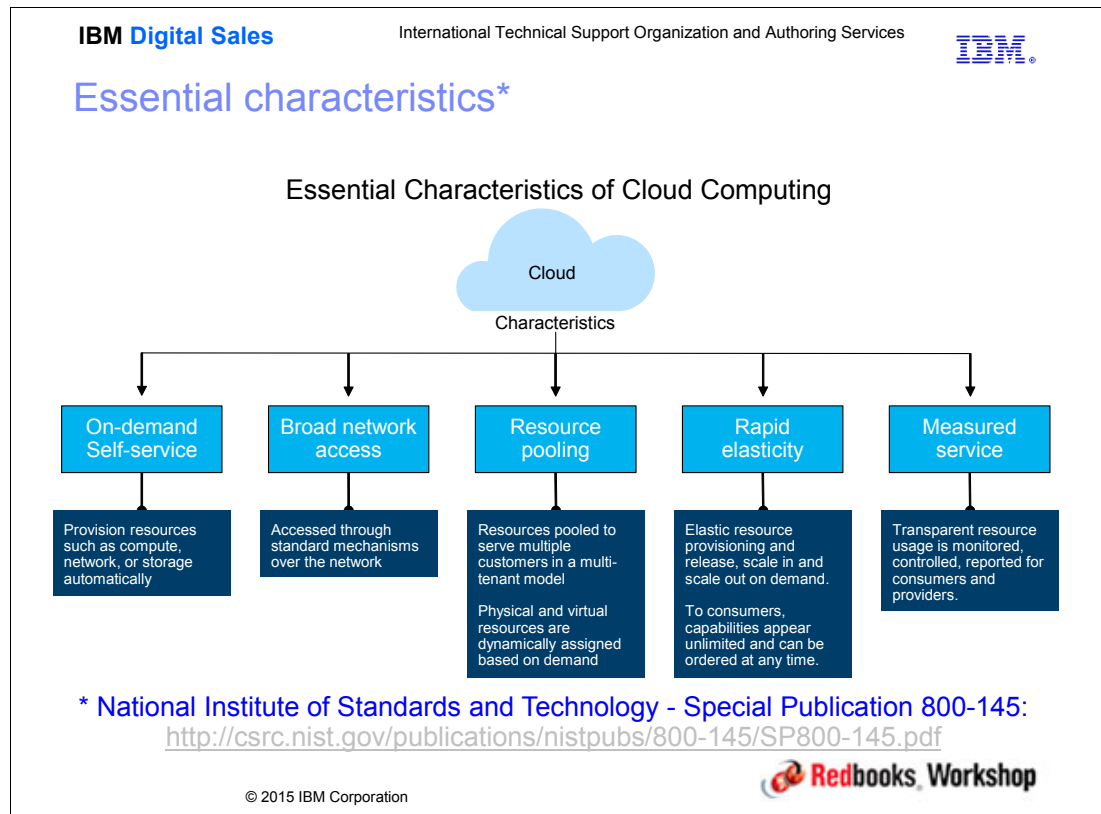


Figure 1-3 Essential characteristics

Notes:
(none)

1.7 Deployment models

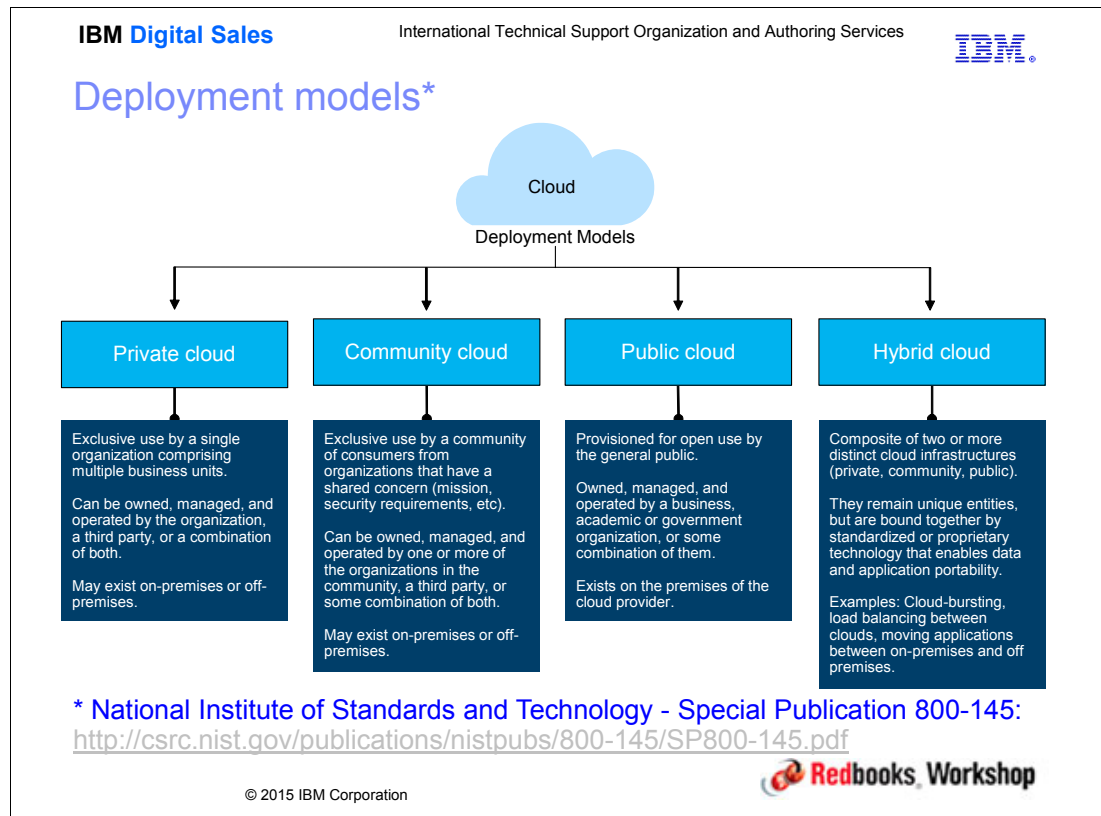


Figure 1-4 Deployment models

Notes:
(none)

1.8 Service models

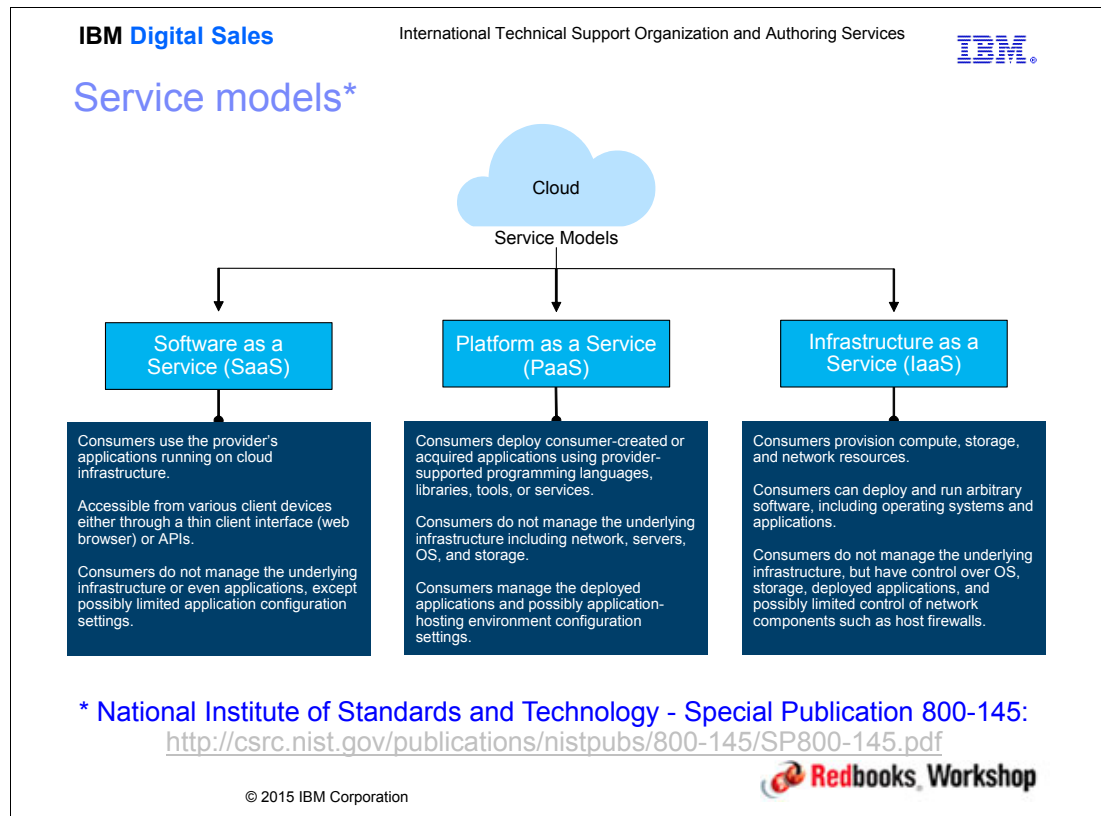


Figure 1-5 Service models

Notes:
(none)

1.9 Cloud service models in practice

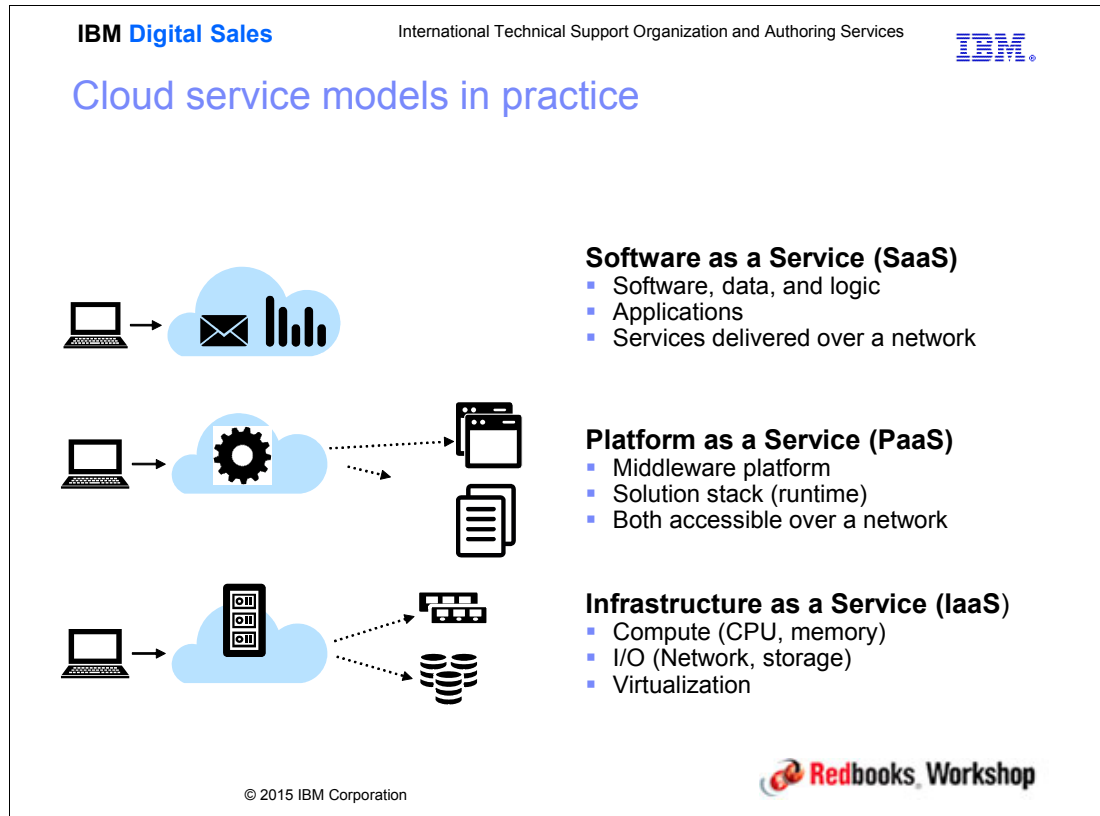


Figure 1-6 Cloud service models in practice

Notes:

- IaaS** Making up the bottom layer of the cloud is the infrastructure services layer (IaaS). In this layer, a set of physical assets such as servers, network devices, and storage disks are offered as provisioned services to consumers. The services at this layer support application infrastructure, regardless of whether a cloud and many more consumers provide that infrastructure. As with platform services, virtualization is an often-used method to provide the on-demand rationing of the resources.
- PaaS** Platform services (PaaS) is the layer in which application infrastructure emerges as a set of services. These services include but are not limited to middleware as a service, messaging as a service, integration as a service, information as a service, and connectivity as a service. The services here are intended to support applications. These applications might be running in the cloud, or in a more traditional enterprise data center. To achieve the scalability that is required within a cloud, the different services that are offered here are often virtualized.
- SaaS** Application services are most familiar to everyday web users. The application services layer hosts applications that fit the SaaS model. These applications run in a cloud and are provided on demand as services to users. Sometimes the services are free and providers generate revenue from things like web ads. At other times, application providers generate revenue directly from the usage of the service. Do these scenarios sound familiar? It probably does because almost everyone uses them. If you use a tax preparation service to file your income taxes online, or use an email service to check your mail, then you are familiar with the top layer of the cloud. These types of applications are just a couple of examples. There are literally

thousands of SaaS applications, and the number grows daily primarily because of Web 2.0 technologies.

There are many applications in the application services layer that are directed to the enterprise community. There are hosted software offerings available that handle payroll processing, human resource management, collaboration, customer relationship management, business partner relationship management, and so on. Examples of these offerings include IBM SmartCloud® For Social Business, Unyte®, Salesforce.com, Sugar CRM, and WebEx.

1.10 Cloud service models – Technology mapping

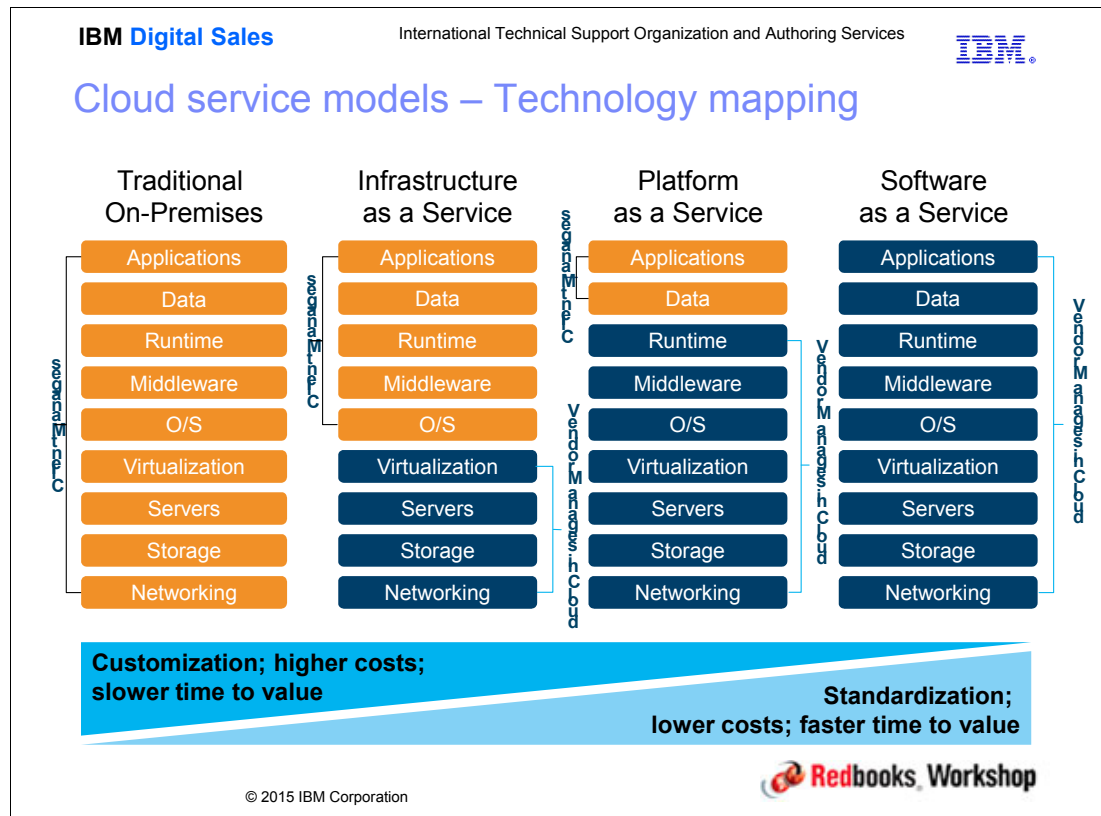


Figure 1-7 Cloud service models – Technology mapping

Notes:

From *IBM CS-101 Introduction to Cloud*

Clients who are not using cloud (Traditional / On-Premises environments) manage the entire technology stack.

As you move higher up the cloud to IaaS, PaaS and SaaS, clients manage less and less infrastructure and focus more and more on their business.

1.11 Cloud adoption and migration is workload driven

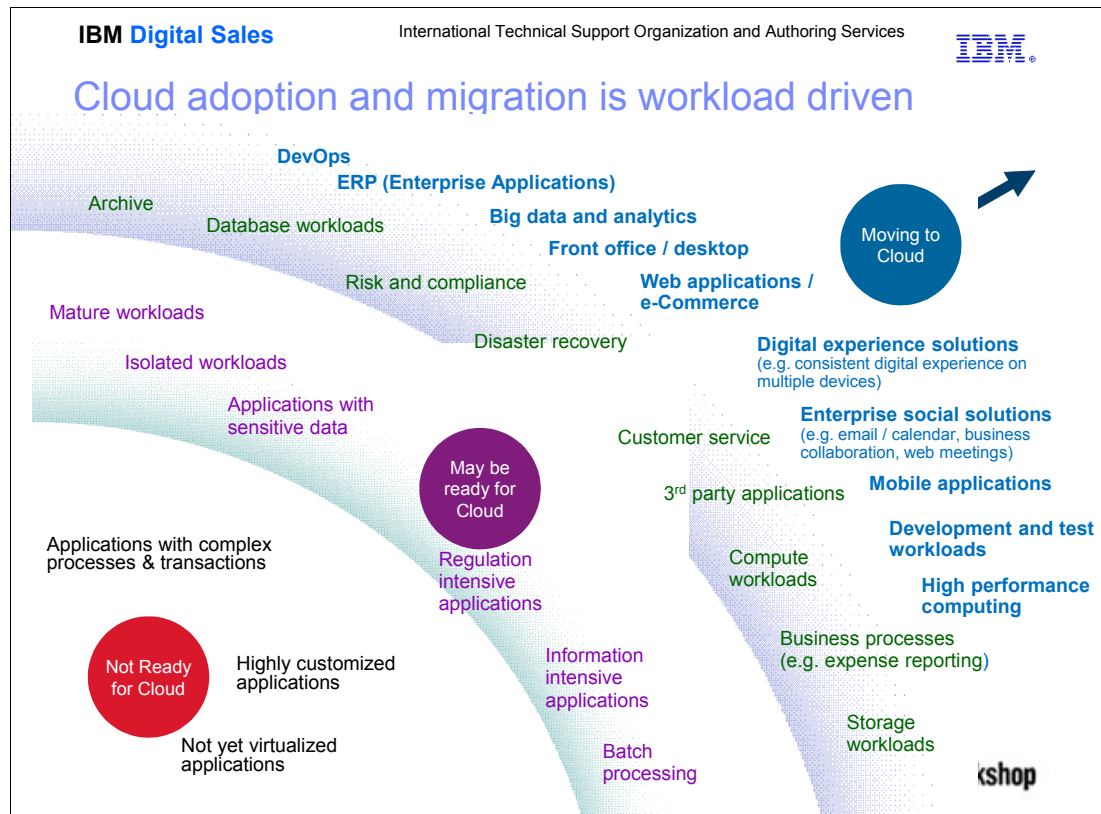


Figure 1-8 Cloud adoption and migration is workload driven

Notes:

From *IBM CS-101 Introduction to Cloud* (including comments below):

Clients are aggressively moving workloads or applications to the cloud. Internet-based applications like email, web applications, and E commerce were the first to be moved. Standardized applications like front office and business process applications also are well suited to cloud. Disaster recovery and archive are also moving and we are seeing that analytics, big data, mobile and social business depends upon cloud deployment models.

What is a workload?

An independent service or collection of code that can be executed.

An abstraction—that is, isolated from the hardware it's running on—focusing on what needs to be done, as opposed to how it's going to be done, in the context of a particular cloud.

Can be a small or complete application, typically combined with other workloads to execute a business process or task.

Why workloads matter:


Companies adopting cloud will likely make decisions about cloud based on workload capacity – cloud architectures and deployment models need to fit.

Companies adopting cloud need to actively manage workloads so they know how their applications are running, what they're doing, and how much an individual department should be charged for its use of services.

Clients are seeing real value by moving their development tests and development operations or DevOps works to the cloud. Cloud can speed the development and deployment for new applications and reduce the time needed for the development process. We also are seeing some clients move workloads to cloud that they were not ready to move just a few years ago. These include information-intensive applications and applications with sensitive data. Network availability, as well as well defined cloud security frameworks, can facilitate moving some of these workloads.


Some applications, however, are not ready for cloud. These include applications with complex processes and transactions that may require excessive reengineering, as well as highly customized applications. In addition, applications that are not yet virtualized may receive few benefits from a cloud model.

1.12 Infrastructure as a Service

IBM Digital Sales International Technical Support Organization and Authoring Services 

Infrastructure as a Service

- An infrastructure service provider makes an entire computing infrastructure available “as a service”.
- Manages a large pool of computing resources and uses virtualization to assign and dynamically resize customer resources.
- Customers subscribe to processing capacity, memory, data storage, and networking resources that are provisioned over a network.




© 2015 IBM Corporation 

Figure 1-9 Infrastructure as a Service

Notes:

Infrastructure as a Service (IaaS) is a way of delivering cloud-computing infrastructure, including servers, storage, network, and operating systems, as an on-demand service. Rather than purchasing servers, software, data center space, or network equipment, clients instead buy those resources as a fully outsourced service on demand.

“Two guys in a Starbucks can have access to the same computing power as a Fortune 500 company” - Jim Deters, Founder - Galvanize

Infrastructure services are built on top of a standardized, secure, and scalable infrastructure. Some level of redundancy needs to be built into the infrastructure to ensure the high availability and elasticity of resources.

Software automation is typically used to provision and de-provision services.

1.13 Example IaaS use case: Website publishing scenario

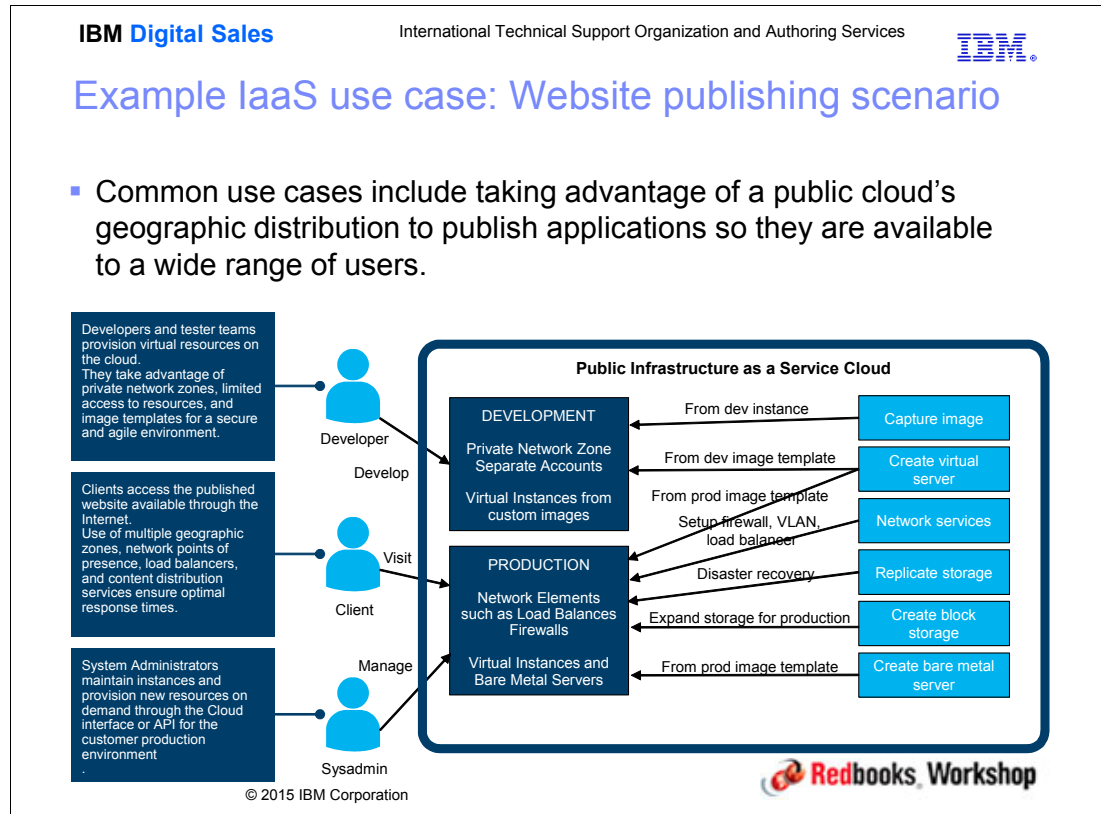


Figure 1-10 Example IaaS use case: Website publishing scenario

Notes:

IaaS can be used for any type of environment (development, test, and production) and can support a wide variety of applications.


This simplified use case shows how a client can leverage IaaS to set up a development environment that provides cheap virtual servers from standard templates that contain development tools, and also maintain a production environment that takes advantage of the cloud's global footprint.

1.14 Platform as a Service

IBM Digital Sales International Technical Support Organization and Authoring Services **IBM**

Platform as a Service

- Service provider supplies the software platform or middleware where the applications run in addition to the underlying infrastructure.
- The customer is responsible for the creation, updating, and maintenance of the application that sits atop the platform.
- Common platform environments include Java Application Servers, NodeJS, Python, PHP, Go, and more, with services commonly provided for SQL and NOSQL data stores.



© 2015 IBM Corporation


 **Redbooks** Workshop

Figure 1-11 Platform as a Service

Notes:

Platform as a service (PaaS) can be used to quickly and easily create and maintain applications without the complexity of maintaining the software stack (such as application servers) or the infrastructure underneath it.

Installing, configuring, licensing, patching and maintaining the software stack is performed by the cloud vendor so that the customer can focus on developing and maintaining their application.

Some of the PaaS offerings on the market include IBM Bluemix®, IBM CMS4Oracle, IBM CMS4SAP, Heroku, Google App Engine and AWS Elastic Beanstalk.

1.15 Software as a Service

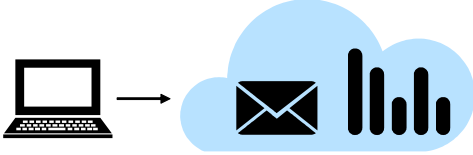
IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Software as a Service

- The service provider is responsible for the creation, updating, and maintenance of the software and application.
- The service user accesses the service through Internet-based interfaces.



© 2015 IBM Corporation

Redbooks Workshop

Figure 1-12 Software as a Service

Notes:

SaaS is software delivery method that provides access to software and its functions remotely (typically as a web-based service). Software as a Service allows organizations to access business functionality at a cost typically less than paying for licensed applications. Because software is hosted remotely, organizations do not need to invest in additional hardware. Software as a Service removes the need for organizations to handle the installation, setup, and maintenance.

Under the SaaS model, the software provider is responsible for the creation, updating, and maintenance of software, including the responsibility for licensing the software. Customers usually rent the software on a per usage basis, or buy a subscription to access it that includes a separate license for each person who uses the software. Upgrades and new features are typically included as part of the on-going application lifecycle.

In this model, the service user only accesses the service itself, and not the platform or the infrastructure the service is running on. The service is usually accessed as a web application or as a wrapped web services application invoked by using web services APIs.

1.16 Implications of Cloud Computing



Figure 1-13 Implications of Cloud Computing

Notes:

(none)

1.17 Business benefits of Cloud Computing

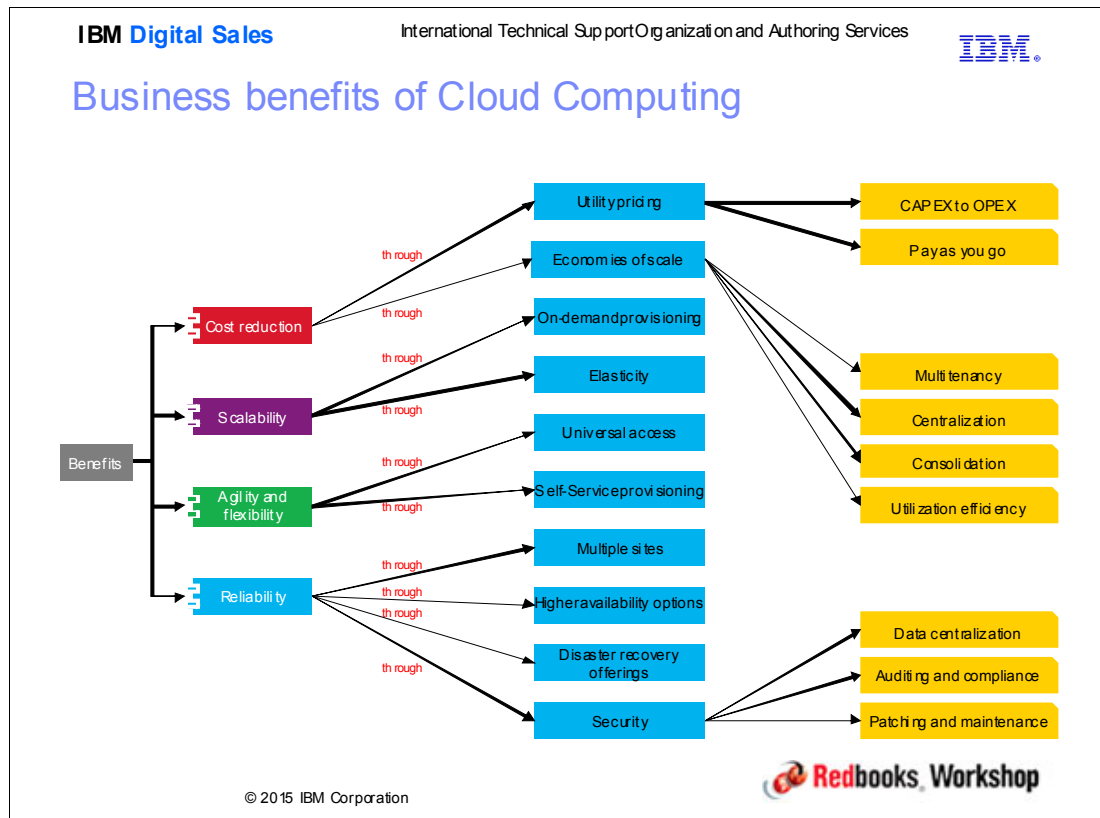


Figure 1-14 Business benefits of Cloud Computing

Notes:

These are potential benefits from a business point of view. Benefits will vary depending on the use case, workload, cloud provider, capabilities, and so on.

1.18 Development benefits of Cloud Computing

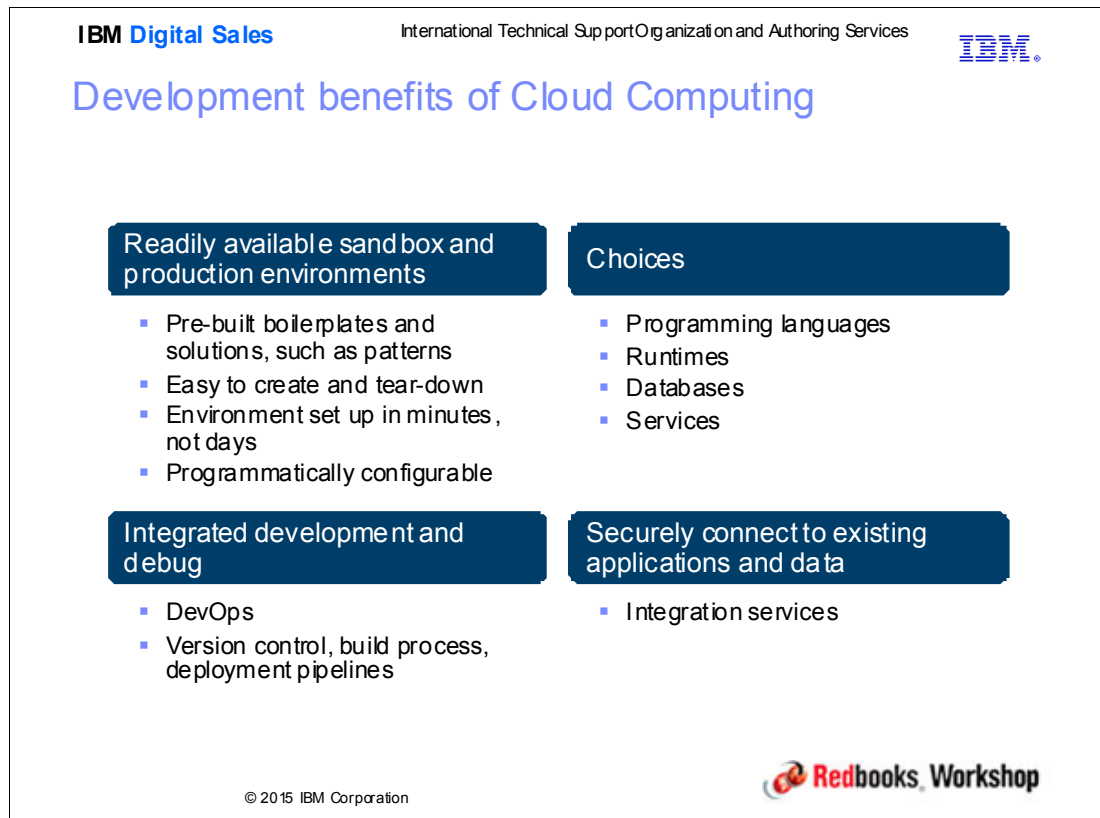


Figure 1-15 Development benefits of Cloud Computing

Notes:

The following are benefits from a developer point of view.

1.19 Industry impact

The slide is titled 'Industry impact' and is part of an IBM Digital Sales presentation. It features two main sections with blue headers and bulleted content. The first section, 'A disruptive change in the IT industry', lists two points: 'New computing model that is different from traditional IT computing models' and 'Based on virtualization, high-speed Internet connectivity'. The second section, 'Demand for dynamic and responsive IT', lists two points: 'Short-duration application lifecycle' and 'Requires new processes, application design, and development environment'. The slide includes the IBM logo in the top right, the text 'International Technical Support Organization and Authoring Services' at the top, and the Redbooks Workshop logo in the bottom right. A copyright notice '© 2015 IBM Corporation' is at the bottom center.

IBM Digital Sales International Technical Support Organization and Authoring Services IBM

Industry impact

A disruptive change in the IT industry

- New computing model that is different from traditional IT computing models
- Based on virtualization, high-speed Internet connectivity

Demand for dynamic and responsive IT

- Short-duration application lifecycle
- Requires new processes, application design, and development environment

© 2015 IBM Corporation Redbooks Workshop

Figure 1-16 Industry impact

Notes:

Reference: https://en.wikipedia.org/wiki/Disruptive_innovation

Companies are embracing new business models and disruptive technologies to help them become more agile, competitive and innovative.

1.20 IaaS: Common storage models

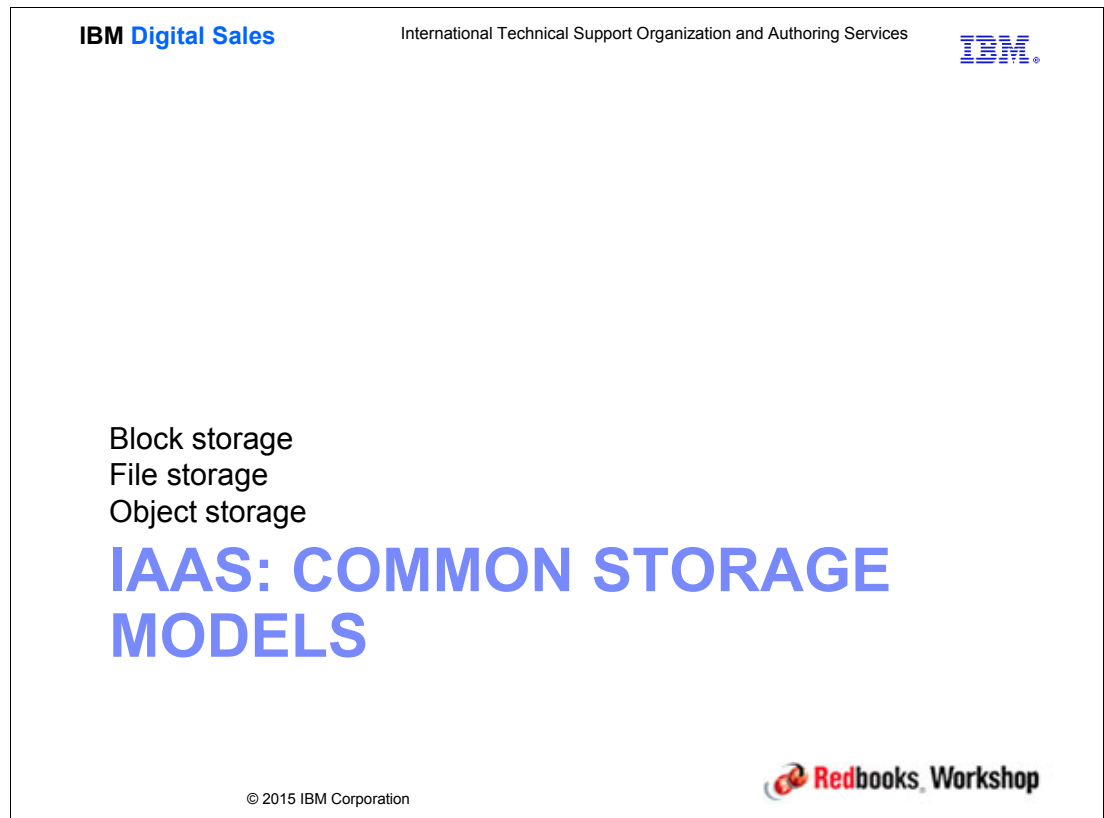


Figure 1-17 IaaS: Common storage models

Notes:

(none)

1.21 Overview of storage options

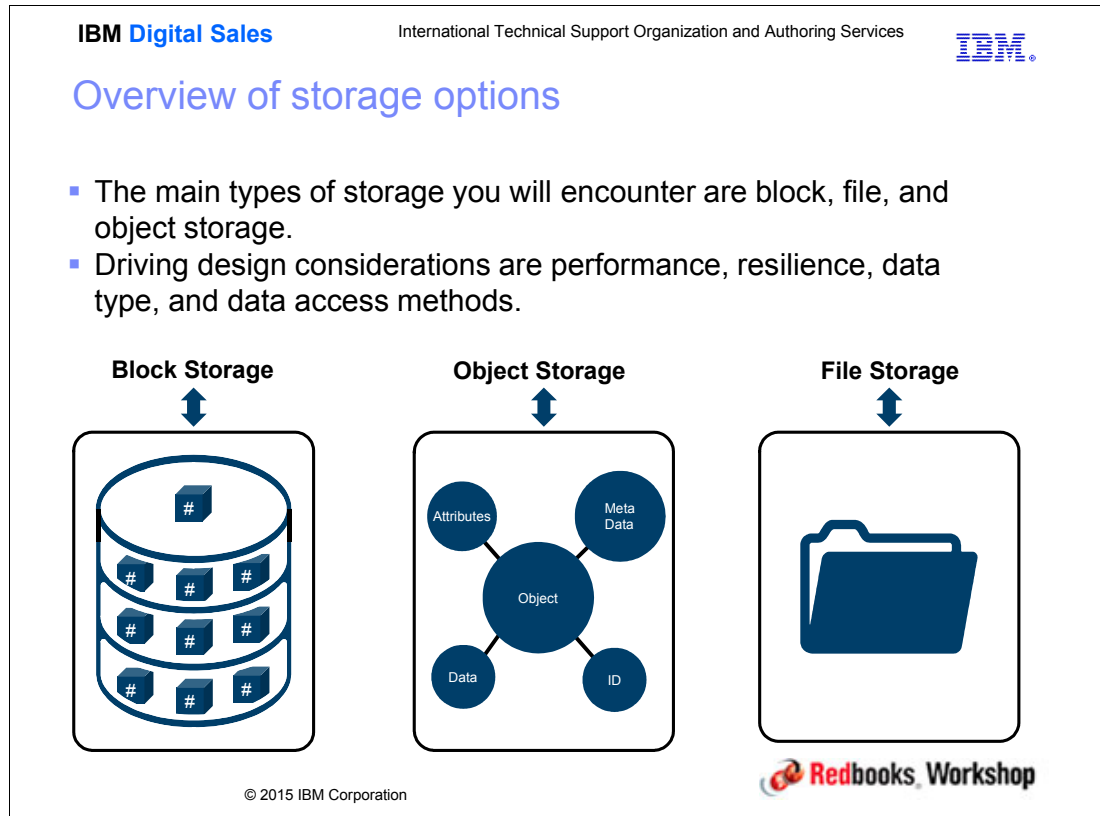


Figure 1-18 Overview of storage options

Notes:

IaaS Storage Considerations There are several options with regards to the storage that CAN be used for the cloud. The important design consideration is the type of data that you are intending to store and the performance that your design requires. The main types of storage are Block, File, and Object.

Block Storage

In IaaS computing, Block Storage refers to what you associate with storage area networks (SANs) where a block (a sequence of bytes and bits) is stored in a data buffer that then reads or writes an entire block at a time. Reading and writing in blocks reduces the processor usage and increases performance. They also support snapshots and replication. Volumes can be provisioned in your desired storage capacity, from 20 GB to 12 TB, and at your desired IOPS tier to support a variety of application needs.

Object Storage

Object Storage is different from Block Storage in that data is stored as objects rather than blocks. This means that each object includes the data, metadata, and a globally unique identifier. Object storage can be implemented at multiple levels, including device, system, and interface level. In each case, object storage seeks to enable capabilities not addressed by other storage architectures.

File Storage

File Storage here refers to network-attached storage (NAS), which is a file-level computer data store connected to a computer network. NAS is specialized for serving files either by its hardware, software, or configuration.

1.22 IaaS: SoftLayer public cloud



Figure 1-19 IaaS: SoftLayer public cloud

Notes:

(none)

1.23 IaaS products and services overview (SoftLayer)

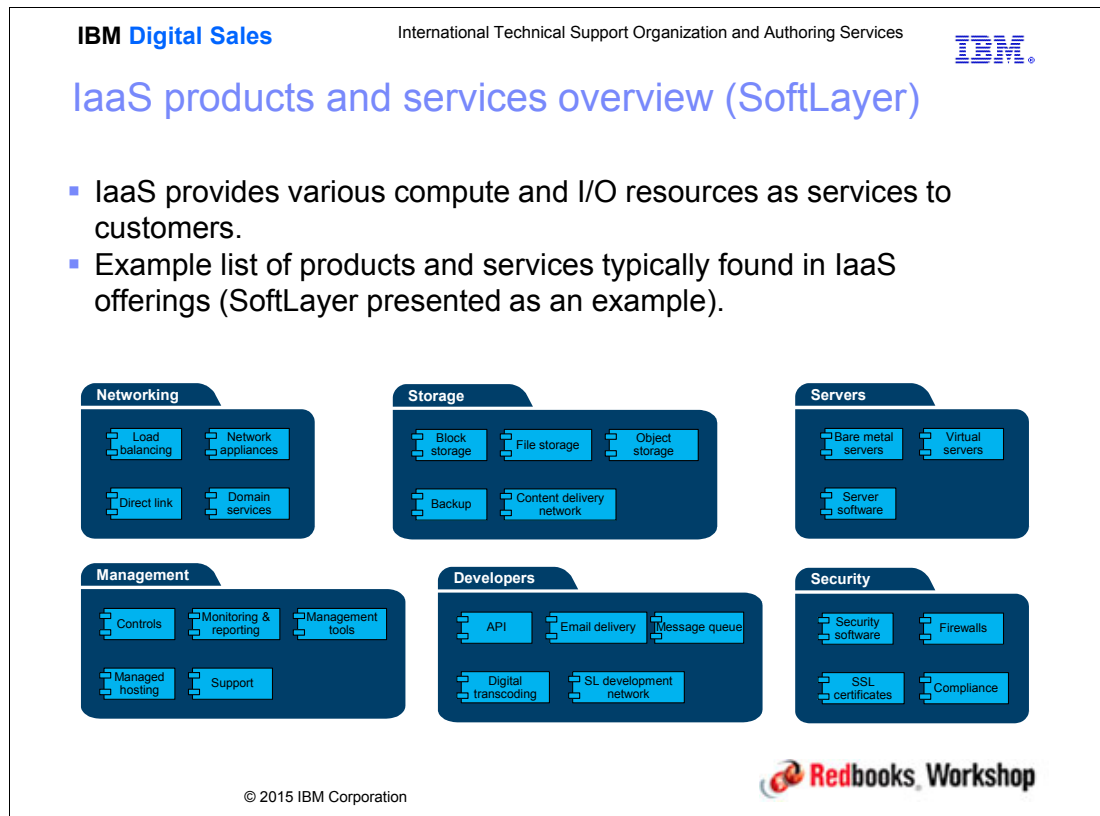


Figure 1-20 IaaS products and services overview (SoftLayer)

Notes:
(none)

1.24 Key takeaways: Cloud Computing

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Key takeaways: Cloud Computing

Essential characteristics	Deployment models	Service models	Adoption
<ul style="list-style-type: none">▪ On-demand self-service▪ Broad network access▪ Resource Pooling▪ Rapid elasticity▪ Measured service	<ul style="list-style-type: none">▪ Private▪ Community▪ Public▪ Hybrid	<ul style="list-style-type: none">▪ Software as a Service (SaaS)▪ Platform as a Service (PaaS)▪ Infrastructure as a Service (IaaS)	<ul style="list-style-type: none">▪ Cloud adoption and migration is workload driven▪ Different workloads are a better fit for IaaS, PaaS, or SaaS in public, private, or hybrid models

© 2015 IBM Corporation


Redbooks Workshop

Figure 1-21 Key takeaways: Cloud Computing

Notes:

(none)



Platform architecture for cloud computing

This unit describes these considerations for platform architecture:

- ▶ Typical IaaS offerings
- ▶ What an IaaS provider should offer
- ▶ Location Location Location
- ▶ SoftLayer presence
- ▶ Network architecture
- ▶ The triple network architecture
- ▶ The data center
- ▶ The Orchestration Management System (OMS)
- ▶ The full overview
- ▶ The management web interface
- ▶ Monitor/Control/Change
- ▶ Account administration
- ▶ Adding users to your account
- ▶ Defining permissions for users
- ▶ Permission templates
- ▶ Defining instance access

2.1 What you should be able to do

At the end of this unit, you should be able to:

- ▶ Identify typical IaaS offerings, and understand what sets SoftLayer apart.
- ▶ Describe the triple network architecture.
- ▶ Identify the benefits of the IBM cloud network.
- ▶ Use the SoftLayer Customer Portal to order, monitor, and control instances.
- ▶ Create users and permissions.

2.2 Teaching topics

This unit covers the following topics:

- ▶ Typical IaaS offerings: What sets SoftLayer apart from other cloud service providers (CSPs) concerning bare metal, virtual, API, and network.
- ▶ The triple network architecture: Describes the unique triple network, and SoftLayer Infrastructure Management System (IBM IMS™).
- ▶ SoftLayer locations: Datacenter location and rapid expansion, data center standardized PoD design, network point of presence (PoP).
- ▶ One Portal to control them all: Describes the Customer Portal.
- ▶ Using the portal to order instances: How to order new instances in the portal.
- ▶ Using the portal to monitor / control instances: How to check health and perform actions.
- ▶ Creating user(s) in the portal: How to create and set up users and permissions.

2.3 Typical IaaS offerings

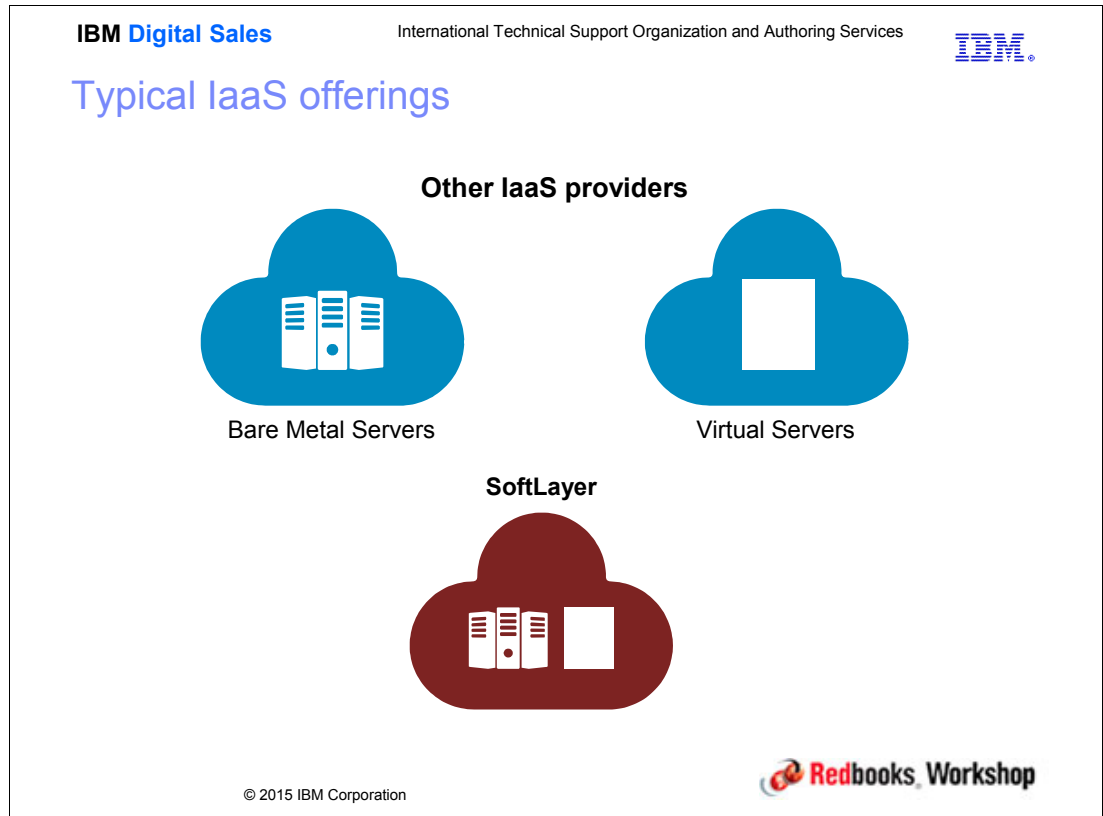


Figure 2-1 Typical IaaS offerings

Notes:

Usually when selecting from the offerings of Cloud Service Providers (CSPs), you can choose either virtual servers or bare metal servers. However, if you wish to mix the two in a single environment, at present SoftLayer is the only provider offering. You can, for example, have a setup consisting of two bare metal servers and five virtual ones.

Note that there are providers of cloud services who are also providers of hosting and outsourcing services, but they do not do so within the same implementation. Currently SoftLayer is the only IaaS provider who uses the same provisioning and services tools in an integrated fashion for both virtual servers and bare metal servers.







2.4 What an IaaS provider could offer

IBM Digital Sales

International Technical Support Organization and Authoring Services


IBM

What an IaaS provider could offer

Servers 	Storage 	Networking 	Security 	Developers 	Management 
Bare Metal Servers ▶	Block Storage ▶	Load Balancing ▶	Software ▶	API ▶	Controls ▶
Virtual Servers ▶	File Storage ▶	Network Appliances ▶	Firewalls ▶	Email Delivery ▶	Monitoring and Reporting ▶
Software ▶	Object Storage ▶	Domain Services ▶	SSL Certificates ▶	Message Queue ▶	Management Tools ▶
	Backup ▶	Private WAN Interface ▶	Compliance ▶	Digital Transcoding ▶	Managed Services ▶
	Content Delivery Network ▶			Developer Community ▶	Support ▶

A full, end-to-end IaaS offering includes advanced management and a full featured set of APIs

Almost everything can be mixed and matched, making the platform very flexible and adaptable to most customer needs.



© 2015 IBM Corporation

Figure 2-2 What an IaaS provider could offer

Notes:

IaaS is more than just servers on a network. It includes additional services to ensure stability, adaptability, and security, which are on most customers' wish lists.

Before choosing a IaaS or cloud provider, ensure that your needs will be met and that the provider offers a means for you to manage as much of the setup yourself. This is important to avoiding wasting time communicating back and forth before you see the results. When using cloud, one of the parameters you look for is how fast you are able to get what you need.

Compared to other IaaS providers, SoftLayer has many possible combinations of the services, so some technical knowledge is recommended when ordering. In addition, you should also have a design plan. You can change your environment later, but while it is easy to add storage, CPU, and memory to instances, changing the network itself is much trickier. Therefore, plan your network carefully before making any changes.

Many of the service shown in the graphic will be described in detail in the later units.

2.5 Location location location



Figure 2-3 Location location location

Notes:

Location is a key parameter when buying real estate, and it should also be a key parameter when choosing your IaaS provider.

If you have customers or potential customers all over the world, look for a provider who has data centers near the key potential markets because internet customers have little patience and will not wait for data to load or web pages to respond.

In the above graphic, the only datacenter is in the United States, and customers in Europe and Africa are unlikely to be happy about the response times.

Additionally, there might be laws dictating where your data should be stored. For example, many European countries have laws that certain data is not allowed to be stored or even transported outside of Europe or their own country.

2.6 SoftLayer presence

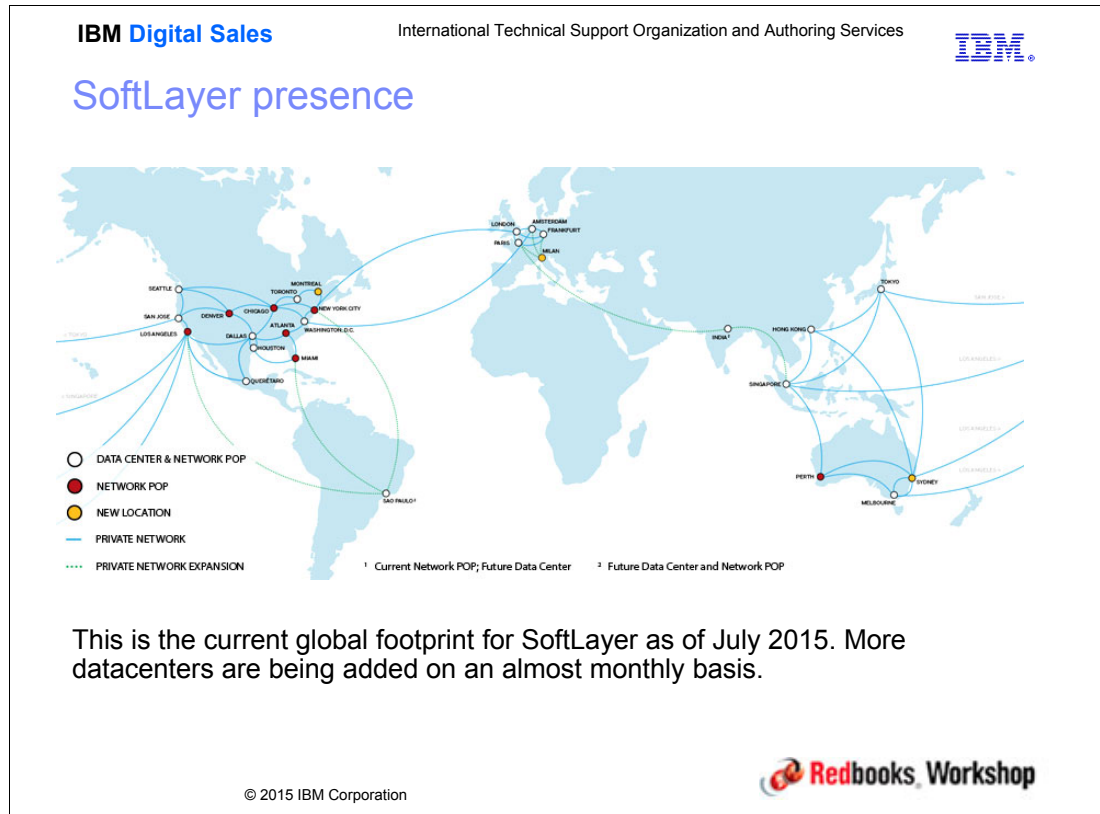


Figure 2-4 SoftLayer presence

Notes:

An example of a IaaS provider who has a global footprint is SoftLayer, which has data centers all over the world and is adding new ones almost monthly. The lines on the picture are not just for show because each all the data centers are connected by a private network. This will be covered in more detail in one of the next slides.

SoftLayer datacenters consist of Point of Presence (PoP) locations and Point of Delivery (PoD) locations.

SoftLayer PoPs are locations which connects SoftLayer PoDs to SoftLayer's global, private, resilient private network, to customer's private Wide Area Networks (via Direct Link) and to the Internet.

Both SoftLayer customers and the end users of SoftLayer customers' applications benefit with SoftLayer's extensive PoP because it reduces the distance that they must traverse on the open Internet before reaching SoftLayer's private, resilient, high speed global network to reach SoftLayer PoDs where a customer's compute and storage reside.

SoftLayer PoDs are where SoftLayer services are delivered from including compute and storage. The next slide shows a SoftLayer PoD.

2.7 The datacenter

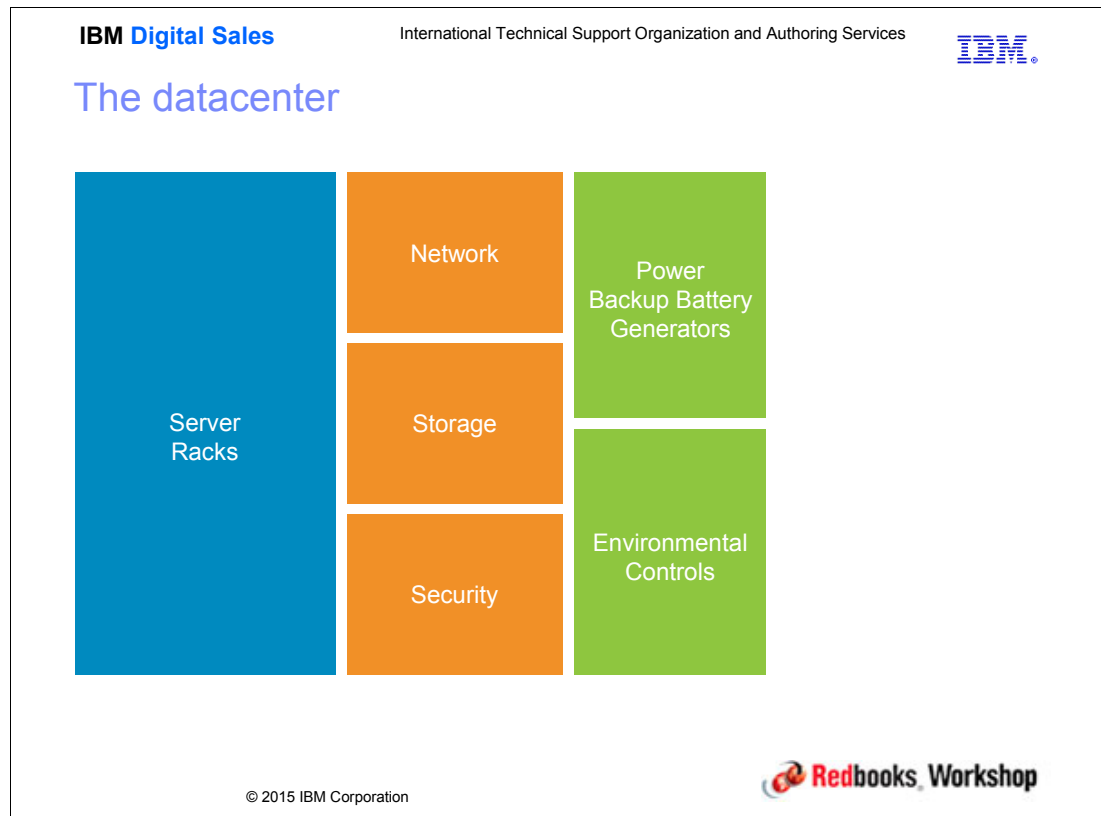


Figure 2-5 The datacenter

Notes:

When choosing an IaaS provider, consider how the datacenter is set up to ensure that you get the uptime and reliability you need and pay for, and that your customers expect.

Make sure that they have backup batteries, generators, and environmental controls to ensure continued operation during a power outage or worse. Make sure that they have contingency plans and that these plans as well as their equipment are reviewed and maintained regularly. You are placing your business in their hands.

A SoftLayer datacenter consists of four or more PoDs. The following are the standard specifications for a PoD as of July 2015:

- ▶ 10000 ft² (930 Sqm)
- ▶ 2 megawatts of power
- ▶ 150 racks
- ▶ 4000 physical nodes
- ▶ N+1 generators (N+1 meaning 1 more than is actually required)
- ▶ N+1 battery backups (N+1 meaning 1 more than is actually required)
- ▶ An additional number of load balancers, firewalls, and storage units

2.8 Network architecture

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Network architecture

The reliability of an IaaS cloud's network is very important because the network is the way in which a cloud customer and their end users interact with their IaaS cloud services.

An IaaS provider could have a network setup that looks like this. Note that this example is from SoftLayer. Other providers might have a different network setup, such as only a public network.

```
graph LR; Customers[Customers] --- PublicNetwork[Public Network]; PublicNetwork --- PrivateNetwork[Private Network]; PrivateNetwork --- ManagementNetwork[Management network]; ManagementNetwork --- ManagementConsole[Management console];
```

- Public network to service customers
- Private network for inter-server communication
- Management network for console access, maintenance, and so on

© 2015 IBM Corporation

Redbooks Workshop

Figure 2-6 Network architecture

Notes:

Note that this example is from SoftLayer. Other providers might have a different network setup, such as only one public network without a private or management network.

2.9 Triple network architecture

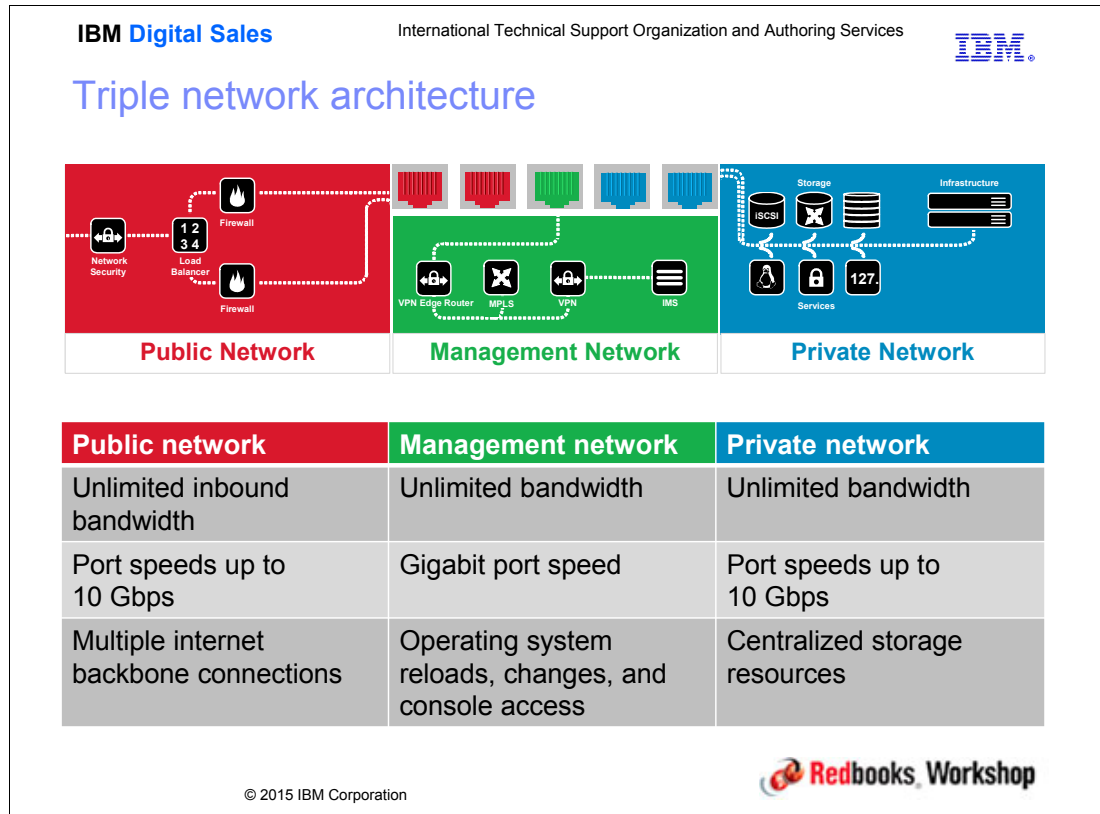


Figure 2-7 Triple network architecture

Notes:

SoftLayer's triple network architecture is very unique not only because it segregates network traffic from the public network (for example Internet, VPN) from the private network (SoftLayer PoD to SoftLayer PoD, Direct Link, customer administrative access) and management network (SoftLayer services) where as other cloud providers lump all this traffic into one network, but also because unlike many IaaS providers SoftLayer does not charge usage fees for SoftLayer customers moving data across the Private Network between a customer's environments in multiple SoftLayer PoDs. The SoftLayer private resilient private network enables lightning fast communication between SoftLayer PoDs.

With the private network that runs between data centers and PoPs, you can transfer data at high speeds to other SoftLayer data centers. This also means that you can have redundant setups in different SoftLayer data centers across the globe where data is synchronized at high speeds at no extra cost. You can also use SoftLayer patch servers and software repositories.

This setup also benefits any clients/customers around the world because no matter in which SoftLayer datacenter your solutions are hosted, the customer only has to reach the nearest datacenter or PoP before Soft Layer's own network takes over, minimizing the number of network hops and handoffs between providers.

The management network allows you, via VPN, to connect to the servers and perform OS reloads, power off/on operations, and monitor your server using keyboard, video, and mouse (KVM) over IP, and get console access to it.

2.10 The Orchestration Management System

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

The Orchestration Management System


Infrastructure Management System provides orchestration and automation


Unique Triple Network Architecture allows seamless communication across distributed environments


x86 Data Center Pod

x86 Data Center Pod

x86 Data Center Pod


Bare Metal Servers


Visual Server Instances


Private Clouds

The Orchestration Management System is a key component in the automation and interaction with IaaS services and services such as:

- Provisioning/deprovisioning
- Billing
- API calls
- Logging
- Alerts

These services are made possible by Infrastructure Management System (IMS) being present and helping with the automation.

© 2015 IBM Corporation

Redbooks Workshop

Figure 2-8 The Orchestration Management System

Notes:

To tie the entire infrastructure together, an IaaS provider will likely use an orchestration management system (OMS).

An OMS handles all interactions with your instances and services from provisioning to restart and logging. Billing and API calls also pass through IMS (Infrastructure Management System), which then handles the automation required to make your orders happen. This is even true for bare metal servers.

SoftLayer is the provider that has the richest set of application programming interfaces (APIs) that allow you to interact directly with the backend system via IMS. The functions available using the API allows you to perform remote server management, monitoring and retrieving information from the various systems such as accounting, inventory, and DNS. Basically, if an action can be performed by using the customer portal, there will be an API for it as well. The customer portal is covered later in this unit, the API in a later unit.

The API uses Representational State Transfer (REST), and many of the most popular programming languages can be used with it (Python, C#, Perl, PHP, and more). This way, SoftLayer customers can control the entire environment from their applications.

2.11 Sample IaaS datacenter

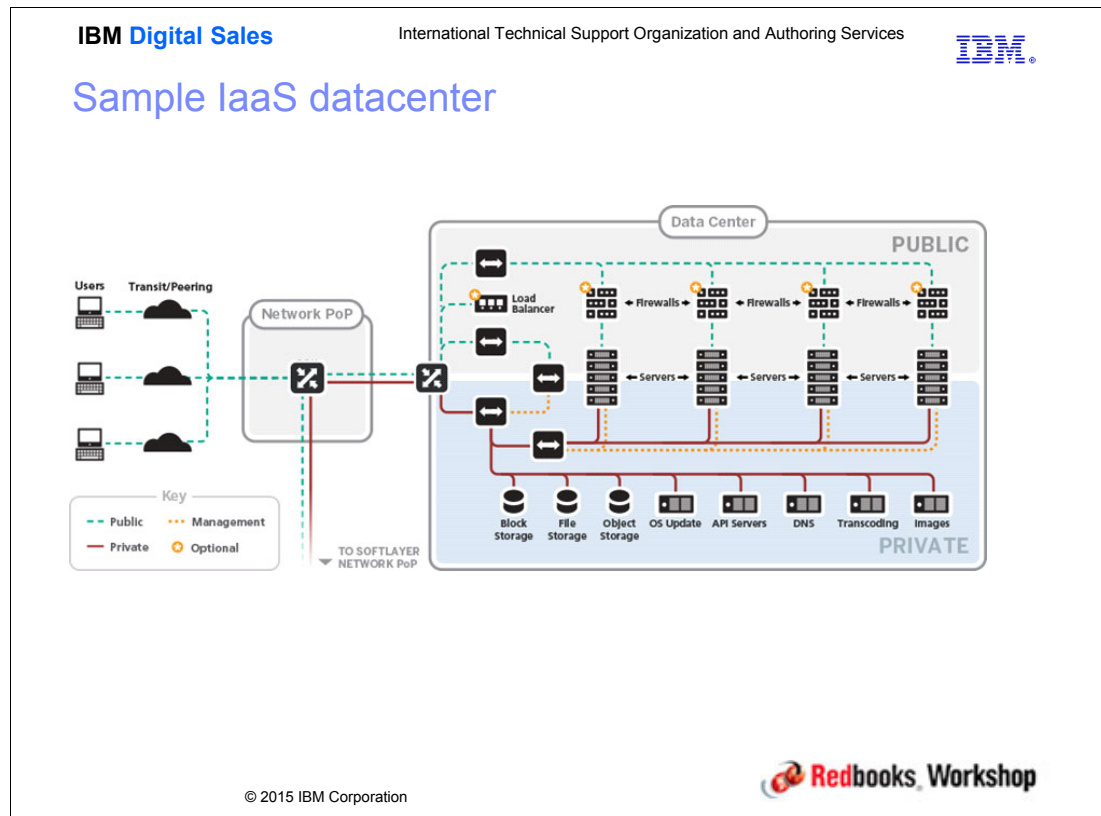


Figure 2-9 Sample IaaS datacenter

Notes:

Now that we covered most of the infrastructure, we can now see the big picture of what an IaaS datacenter looks like when it is deployed. This slide shows how a SoftLayer datacenter is set up. The general structure does not greatly vary from SoftLayer PoD location to location around the world.

You can see that the outside users come in through the PoP through the public network, and that the PoP is connected to the other PoPs and the datacenter.


Inside the datacenter, you can see what is available/accessible on the public network and that you can order firewall and load balancers, both of which will be covered later.

You can also see which parts of the infrastructure are on the private network, such as instance storage, update servers, DNS, API servers, and so on.

2.12 Checkpoint


IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint

- Why do the locations of the IaaS providers' data centers matter?
- How many power generators should an IaaS provider have?
- What are the three network types an IaaS provider can have?
- Which network should console access be delivered on?
- What is the Orchestration Management System?



© 2015 IBM Corporation

Figure 2-10 Checkpoint


Notes:

(none)

2.13 Checkpoint (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint

- Why do the locations of the IaaS providers' data centers matter?
Location = Data Privacy, Latency, Resiliency considerations. Some countries require data to remain in country. Deploying applications closer to the end users can reduce latency and improve end user response times in accessing and utilizing the application. Having diverse locations allows a customer to replicate their data between data centers to mitigate risk from man-made and natural disasters.
- How many power generators should an IaaS provider have?
N+ 1 backup power generators and fail-over battery systems are two ways to provide better reliability of power for their data center..
- What are the three network types an IaaS provider can have?
Public, Private, and Management.
- Which network should console access be delivered on?
Management network.
- What is the Orchestration Management System?
OMS handles all interactions to the servers, be they API calls or server restarts.


© 2015 IBM Corporation

Figure 2-11 Checkpoint (2)

Notes:
(none)

2.14 Management web interface

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Management web interface

HELP CONTACT 2 NOTIFICATIONS 10 OPEN TICKETS IBM - I LOG OUT

SOFTLAYER[®] an IBM Company

Devices Storage Network Security Services Support Account

Search Portal

Account Summary

Current Balance: \$0.00

Estimated Next Balance: \$2,446.43

Current Payment Method: Manual

Scheduled Maintenance & Updates

July 2015

S	M	T	W	T	F	S
28	29	30	1	2	3	4
5	6	7	8	9	10	11
12	13	14	15	16	17	18
19	20	21	22	23	24	25
26	27	28	29	30	31	1
2	3	4	5	6	7	8

August 2015

S	M	T	W	T	F	S
26	27	28	29	30	31	1
2	3	4	5	6	7	8
9	10	11	12	13	14	15
16	17	18	19	20	21	22
23	24	25	26	27	28	29
30	31	1	2	3	4	5

Bandwidth

Order

Devices

Storage

Network

Security

Marketplace

© 2015 IBM Corporation

Redbooks Workshop

Figure 2-12 Management web interface

Notes:

The management web interface is the one place to go to administer your account. You can think of it as the entrance to your data center because you can do the same things here as you would be able to in a physical datacenter. The slide shows the SoftLayer customer portal.

Protect your password to the portal and make sure that when you give users access to the portal that they only have the rights that they need.

Most IaaS providers will have a similar control portal.

2.15 Management web interface capabilities

The screenshot shows the IBM Digital Sales page for SoftLayer. At the top, it says 'IBM Digital Sales' and 'International Technical Support Organization and Authoring Services' with the IBM logo. The main heading is 'Management web interface capabilities'. Below this is a navigation bar with 'SOFTLAYER an IBM Company' and links for 'Devices', 'Storage', 'Network', 'Security', 'Services', 'Support', and 'Account'. The main content area is titled 'Within the customer portal, you can:' and lists four bullet points: ordering/modifying resources, viewing performance details, managing the account, and getting support. At the bottom, there is a copyright notice for 2015 IBM Corporation and a 'Redbooks Workshop' logo.

IBM Digital Sales International Technical Support Organization and Authoring Services IBM

Management web interface capabilities

SOFTLAYER an IBM Company Devices Storage Network Security Services Support Account

Within the customer portal, you can:

- Order, cancel or modify resources on servers. You can also set up firewalls, back up solutions, and basically everything that you would be able to do in a physical data center.
- View performance details and control your bare metal and virtual servers, control the power states of your instances, view bandwidth graphs, test network availability, reboot, and so on.
- Manage your account, set up billing information, see invoices past and present, reset your password, and add users to your SoftLayer account.
- Get support by chatting with Softlayer's support team or submit support tickets, and view valuable technical tutorials and information.

© 2015 IBM Corporation Redbooks Workshop

Figure 2-13 Management web interface capabilities

Notes:

The offerings of the SoftLayer customer portal are shown here, but similar offerings should be available at any IaaS provider on their management web interface

The SoftLayer customer portal offers more than 200 services, which are based on the APIs that are also available to you. This means that you can perform any action available in the portal (and more) from a script, or even make your own portal encapsulating the APIs. APIs are covered in a later unit.

SoftLayer also offers a mobile version of the portal that can run on smartphones where you can perform many of the same tasks as in the main portal. Other providers might have mobile apps as well.


2.16 Monitor/control/change instances from the web interface

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Monitor/control/change instances from the web interface



web01.companya.softlayer.com (Virtual Server)
Public IP: 50.23.4.21 (Houston 2)
Private IP: 10.77.182.8
Current Total: \$12.03 (updated hourly)

[Configuration](#) [Tickets](#) [Usage](#) [Bandwidth](#) [Monitoring](#) [Security](#) [Passwords](#) [Storage](#)

Using the management web interface, you have full control over every aspect of your instance and can perform these tasks for each server:

- Monitor status
- Modify the configuration
- Open support tickets
- Monitor bandwidth usage
- Check audit logs (Security)
- View and edit passwords
- View and edit storage

© 2015 IBM Corporation




Figure 2-14 Monitor/control/change instances from the web interface

Notes:

From the web interface, you have fine-grained control over every aspect of your instances. Because almost everything is automated, you have even more control of the instance than if you had hands on access.

You can add or remove memory and CPU, check the bandwidth used, connect and disconnect the network interfaces, and even modify the speed with which they operate. Some of these actions might require a restart.

If you suspect something is not working properly on the server or just need to be certain that the server is running, you can set up specific monitoring agents. These will be covered in a later unit.

2.17 Monitor/control/change instances

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Monitor/control/change instances

SOFTLAYER

an IBM Company

Devices

Storage

Network

Security

Services

Support

Account

Search Portal

Status:	Active	Location:	Houston 2
Start Date:	2015-07-21 12:46 PM (UTC-04:00)	Last Trans.:	Add Monitoring Agent Config
Reloaded:	N/A		

System

OS: CentOS6.0-64 LAMP for VSI [OS Reload](#)

RAM: 1024 MB [Modify](#)

Processor: 1 CORE [Modify](#)

Network

Public [Order IPs](#)

eth1

Status: ACTIVE

IP Address: 50.23.4.21

Default Gateway: 50.23.4.17

Subnet Mask: 255.255.255.240

Speed: 100 Mbps

Max Speed: 100 Mbps [Modify Max Speed](#)

VLAN: POY2G4 AUTO:hoy02.fcr02a.1700

Private [Order IPs](#)

eth0

Status: ACTIVE

IP Address: 10.77.182.8

Default Gateway: 10.77.182.1

Subnet Mask: 255.255.255.192

Speed: 100 Mbps

Max Speed: 100 Mbps [Modify Max Speed](#)

VLAN: POY2G4 AUTO:hoy02.bcr02a.1720



Figure 2-15 Monitor/control/change instances

Notes:

This slide shows the view that you get when you click the Configuration details for a device in the SoftLayer customer portal. It shows the current configuration and status of your server. You can disconnect or connect the network interfaces, order a reload of the OS, and modify the memory and CPU from this view. Storage is managed in the Storage tab.

You can also order a firewall, not shown in the image, to add security to your server. Firewalls will be covered later.

2.18 Account administration

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Account administration

Account

Place an Order

Sales >

Billing >

Users

VPN Access

Manage >

In the web interface, we can see a full overview of your account and can perform these tasks:

- View quotes and pending orders.
- Cancel orders or your entire account.
- Place orders.
- View invoices.
- Manage users and passwords.
- Manage VPN access and passwords.
- Manage your company profile and contact information.
- Subscribe to alerts from Softlayer to be notified of incidents that could affect your systems.

© 2015 IBM Corporation

Redbooks Workshop

Figure 2-16 Account administration

Notes:

The portal provides more control than just hardware and network topics. The portal also provides good tools. You can perform the following tasks inside the portal:

- ▶ You can place orders or get quotes. You can also see orders waiting for approval and approve or void them.
- ▶ In the billing section, you can see you past and present total invoices for the account and the current balance to be invoiced.
- ▶ You can manage users (adding, deleting, disabling, and changing passwords), and grant and revoke user permissions. This topic is covered in the next slides.
- ▶ You can control who has VPN Access and reset their passwords.
- ▶ You can subscribe to alerts so that you are notified should something happen at the datacenter or network that would affect the availability of your services. You can also change your company profile and contact persons as well as view the audit log for the portal.

2.19 Adding users to your account

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Adding users to your account

Log In Settings

User Editable?: ☐ Yes, user can edit settings

Expire Password In: 90 Days

Restrict Access to IP:

Require Security Questions? ☐ Yes, require security questions on log in

Your Current IP:

Valid passwords must be 8 to 20 characters in length with a combination of UPPER and lower case characters, at least one number, and at least one of the following special characters: _-|@.?!~#\$\$%^&*(){}[]=.

*New Portal Password:

*Confirm Password:

*VPN Password:

*Confirm Password:

The password to the web interface is the key to your server room. Guard it carefully.

© 2015 IBM Corporation

Redbooks Workshop

Figure 2-17 Adding users to your account

Notes:

Adding a user to your account is easy, but remember that giving a user access to the web interface essentially gives that user access to your server room and everything within. You can limit this by using permissions, which are covered in the next slides, but it is still recommended to have strict passwords rules and security policies.

Adding two factor authentication is also an option to add more security, which will be covered in a later unit.

2.20 Defining permissions for users



Figure 2-18 Defining permissions for users

Notes:

In the SoftLayer customer portal, the security permissions are divided into six categories:

- ▶ Support
- ▶ Devices
- ▶ Network
- ▶ Security
- ▶ Services
- ▶ Account

The Support tab allows you to view and edit tickets.

The Devices tabs allows you to manage the hardware/virtual hardware devices you have such as servers, firewalls, and load balancers. Device access is controlled in a separate permission set that is covered in later slides.

The Network tab gives you access to the network settings such as IP addresses, Subnets, VLAN spanning, VPN, and gateways.

The Security tab allows you to request a compliance report.

The Services tab allows you to manage other services offered such as images, licenses, provisioning scripts, and vulnerability scanning.

Account provides access to options related either directly to the account such as information about the company and payment methods, but also the permissions to request services that will incur a charge and to cancel servers.

Other IaaS providers should also have a way of setting permissions in their management web interfaces.

2.21 Default permission templates

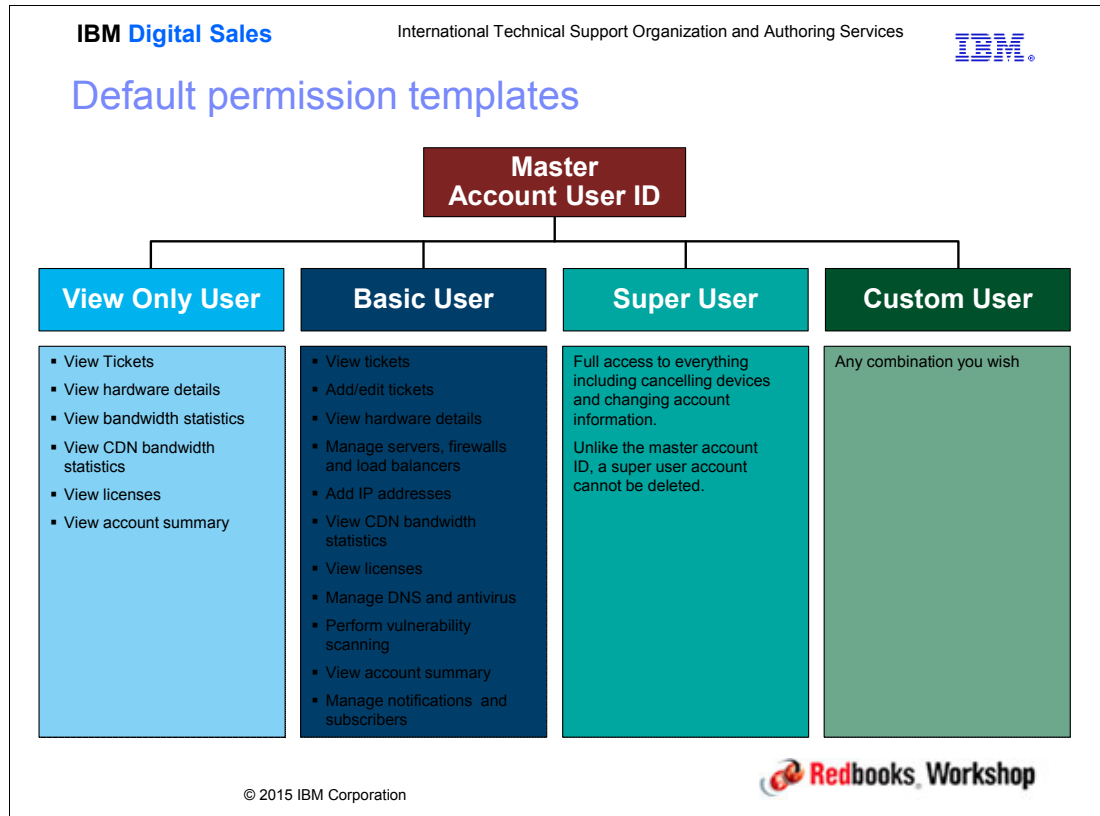


Figure 2-19 Default permission templates

Notes:

The master account user can do everything and cannot be deleted. Generally, do not use this account for daily use. Rather, create users with permissions based on the roles they will perform.

The SoftLayer customer portal has three default permission templates that you can apply. However, but you will likely need to create custom users and use the templates as a base for those users.

- ▶ The View Only User has the privileges to view tickets and basic statistics.
- ▶ The Basic User can manage almost everything concerning servers and devices, but cannot cancel a device. If a basic user orders a service that incurs a charge, someone who can approve charges to the account must approve it before the service is processed.
- ▶ The Super User has the same rights as the master account. Carefully consider whether you really need another user with that much authority. If you do, there should not be many of these for security reasons.

If other IaaS providers do not have templates, you will have to create them or manually set permissions for each user.

2.22 Defining instance access

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Defining instance access

☐ None

☒ All Devices

☐ All Virtual Servers

☐ All Hardware

When new devices are added:

☒ Do NOT add them to this account automatically.

☐ Always add them to this account automatically.

Update Device Access

Cancel

Viewing 1 to 7 of 7

	Device Name	Public IP	Private IP	Notes
<input checked="" type="checkbox"/>	web02.companya.softlayer...	50.23.4.20	10.77.182.15	
<input checked="" type="checkbox"/>	web01.companya.softlayer.c...	50.23.4.21	10.77.182.8	

The access to instances through the web interface or API can be granted per user either per instance or sorted by type:

- All Devices: Access is granted to all instances, both virtual and bare metal.
- All Virtual Servers: Access is granted to virtual servers only.
- All Hardware: Access is granted to all bare metal servers.

The user can access the server if they have a user ID and a password to the server.

© 2015 IBM Corporation

Redbooks Workshop


Figure 2-20 Defining instance access

Notes:

Although the previous permission set could seem to indicate that you have access to instances based on them alone, this is not the case. If using SoftLayer, you will need to set the permissions for the user and decide which instances, if any, that user is able to access. Other IaaS providers might do this differently

The permissions that you set can be even more fine-grained if you use the quick filters. Quick filters allow you to grant permission automatically to the user for any future instances of the same type. Account here means user account.

2.23 Checkpoint

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint

- What can you compare the management web interface to?
- Can you power off your server from the management web interface?
- Can you view your account invoices in the management web interface?
- Can you disconnect your server from the network in the management web interface?
- Can you give a user access to only virtual servers from the management web interface?



© 2015 IBM Corporation 

Figure 2-21 Checkpoint

Notes:

(none)

2.24 Checkpoint (2)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint

- What can you compare the management web interface to?
The server room
- Can you power off your server from the management web interface?
Yes
- Can you view your account invoices in the management web interface?
Yes, both for the account and for individual devices
- Can you disconnect your server from the network in the management web interface?
Yes, you can disconnect both public and private networks, but not the management network
- Can you give a user access to only virtual servers from the management web interface?
Yes it is possible and also grants access to virtual servers provisioned in the future


© 2015 IBM Corporation 

Figure 2-22 Checkpoint (2)

Notes:
(none)



Server offerings in cloud computing

This unit describes:

- ▶ The types of servers available in IaaS cloud offerings, and their primary
- ▶ Characteristics, usage, and advantages
- ▶ What options are available for each server type
- ▶ How you can order, configure, inspect, access, and cancel servers
- ▶ How you can customize the server configuration, and mix and match bare metal servers and virtual servers to build cloud environments that respond to your business needs.

3.1 What you should be able to do

Upon completion of this unit, you should be able to:

- ▶ Recognize and explain the difference between various server types in offerings of IaaS service providers, such as bare metal, public virtual, and private virtual instances with hourly and monthly billing.
- ▶ Order a virtual server from the SoftLayer Customer Portal with hourly billing.
- ▶ Order a SoftLayer bare metal server with monthly billing and local drives configured as RAID.
- ▶ Access provisioned servers and verify their configuration and connectivity.
- ▶ Cancel servers.

3.2 References

The following items are useful for further research:


- ▶ Erl, Thomas; Puttini, Ricardo; Mahmood, Zaigham, *Cloud Computing: Concepts, Technology & Architecture*, Prentice Hall, 2013
- ▶ Virtual Server on the SoftLayer KnowledgeLayer®:
<http://knowledgelayer.softlayer.com/topic/virtual-server-0>
- ▶ Bare Metal Server on the SoftLayer KnowledgeLayer:
<http://knowledgelayer.softlayer.com/topic/bare-metal-server-0>
- ▶ Cabling a SoftLayer Data Center Server Rack:
<https://www.youtube.com/watch?v=jLgvDValxFE>
- ▶ SoftLayer Amsterdam - AMS01 Data Center Tour:
<https://www.youtube.com/watch?v=u0MIg9lggiI>

3.3 Teaching topics

The following topics are covered in this unit:

- ▶ Compute nodes in IaaS cloud
- ▶ Server offerings in SoftLayer IaaS cloud: Virtual servers (public and private node), bare metal servers (with hourly and monthly billing)
- ▶ Ordering and provisioning servers
- ▶ Inspecting and managing servers, upgrading and downgrading
- ▶ Accessing and working with provisioned servers
- ▶ Canceling servers

3.4 Compute Nodes in IaaS Cloud

IBM Digital Sales International Technical Support Organization and Authoring Services 

Compute Nodes in IaaS Cloud

- Premise of cloud computing: Remote “on demand” self-service access to virtualized resources, with “pay per use” billing
- Crucial resources: Computing instances (“compute nodes” – servers)
- Traditionally, cloud compute nodes were virtual servers in multi-tenant environment (“virtual public nodes”)
- SoftLayer extends the traditional concept of cloud computing instances by including these concepts:
 - Dedicated physical servers (“bare metal”)
 - Single-tenant virtual servers (“virtual private nodes”)
 - Available hourly and monthly billing



© 2015 IBM Corporation

Figure 3-1 Compute Nodes in IaaS Cloud

Notes:

Cloud computing is the buzzword of recent years, and it changes the way companies run their IT divisions. It came to life not as a single new idea, but by joining existing ideas that were known for years. This synergy was the source of progress, which is how much of the progress in IT happens.

The following are the ideas that combined to form the cloud:

- ▶ Remote access to computers: People realized that “non-personal”, corporate computers are mostly accessed remotely using Internet protocols that were invented in the 1970s.
- ▶ Virtualization: Machines became so powerful that they could efficiently “pretend” to be other machines. Virtualization was first used by IBM in the 1960s
- ▶ The Metering and “pay per use” billing model, used in the utility industry.

As opposed to customer clouds that are mainly used for storage, in business environments the computing instances (also called compute nodes) are the crucial cloud resources. Traditionally, those were virtual machines provisioned in the provider’s data centers.

IBM SoftLayer pioneered extending the concept of compute node types, adding bare metal servers to the traditional virtual servers, and making virtualization a choice, not a mandate.

3.5 IaaS cloud server offerings (SoftLayer)

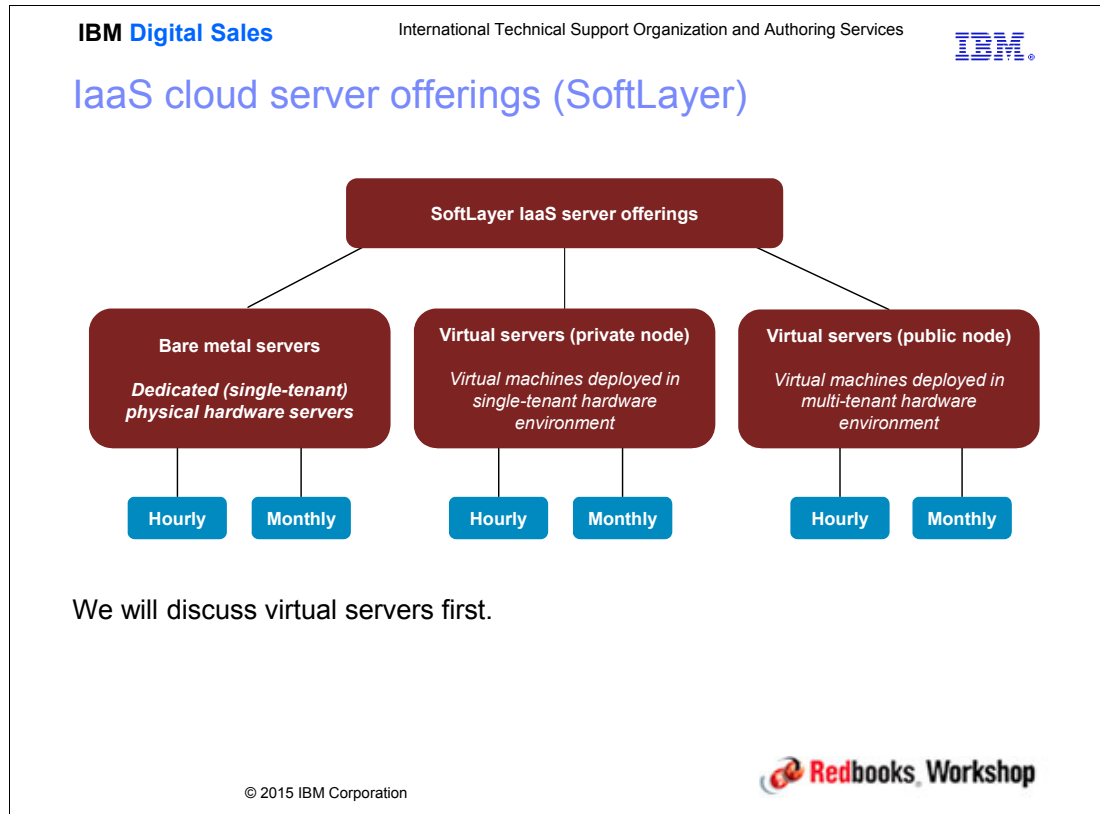


Figure 3-2 IaaS cloud server offerings (SoftLayer)

Notes:

As you can see when ordering SoftLayer servers, there are three types: Bare metal, virtual (private node), and virtual (public node).

Let's start with virtual servers. Those are traditional virtual machines created and run using virtualization mechanisms by a hypervisor running on a host machine.

What is the difference between public and private nodes? It has to do with the concept of tenancy.

A private node is run on a host machine that is dedicated to one customer of the cloud provider (single tenant).

A public node is run on a host machine that is shared between multiple customers (multi-tenant).

3.6 Virtual Servers - Public node

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Virtual Servers – Public node

- Virtual machines are deployed in a multi-tenant hardware environment
- The hypervisor, Citrix Xen®, is fully managed by SoftLayer
- Up to 16 2-GHz cores
- Up to 64 GB RAM

COMPUTING INSTANCE

1 x 2.0 GHz Core

Public Node | Private Node |

☒ 1 x 2.0 GHz Core

☐ 2 x 2.0 GHz Cores

☐ 4 x 2.0 GHz Cores

☐ 8 x 2.0 GHz Cores

☐ 12 x 2.0 GHz Cores

☐ 16 x 2.0 GHz Cores

Can be provisioned in 15 minutes or less

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-3 Virtual Servers - Public node

Notes:

Public virtual servers are deployed in a multi-tenant environment. It is the most traditional model in cloud computing. They can use up to 16 2-GHz cores, and up to 64 GB of RAM.

The Virtual Server cores are “virtual” cores, which are half of a physical hyper-threaded Intel core.

Linux instances are usually provisioned faster than Windows instances, and they are all up and running in the matter of a few minutes.

3.7 Virtual Servers - Public node (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Virtual Servers – Public node

Advantages:

- Fast provisioning
- Affordable solution for deployments without stringent performance or compliance requirements, where resource sharing is OK
- Hourly billing offers flexibility
- Deployments can be automatically scaled up and down
- Typical use is for LAMP-based web servers

Because public virtual servers share physical resources of the host hardware with other public virtual servers, you may observe fluctuations in performance (the “noisy neighbor effect”).

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-4 Virtual Servers - Public node (2)

Notes:

This is the most cost-effective and option, with fastest provisioning.

Two levels of billing granularity are available: Hourly and monthly. Some features available only with monthly billing.

Public virtual servers are used for Autoscaling groups that can be defined in the SoftLayer customer portal.

Because of multi-tenancy, your public virtual server's share of the host machine resources might vary depending on usage of public virtual servers of other SoftLayer customers sharing the host. This is unpredictable. SoftLayer will not allow your public virtual server to be starved to death, but its performance may go up and down within some limits.

Typical usage of public virtual servers is for LAMP (Linux, Apache, MySQL, PHP) deployments that can tolerate lower resource levels.

3.8 Virtual Servers - Private node

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Virtual Servers – Private node

- Virtual machines deployed in single-tenant hardware environment
- Fewer CPU choices and higher prices than public nodes
- When you order your first virtual private server, you get a dedicated physical host machine to house your virtual server.
- Subsequent private virtual servers that you order are provisioned on the same host, so you are the only tenant on that host.
- You get single-tenant private cloud fully managed by cloud provider.
- Other features such as public virtual nodes.

COMPUTING INSTANCE

1 x 2.0 GHz Core

Public Node | Private Node |

☐ Private 1 x 2.0 GHz Core

☐ Private 2 x 2.0 GHz Cores

☐ Private 4 x 2.0 GHz Cores

☐ Private 8 x 2.0 GHz Cores

No “noisy neighbor effect”, so it is better suited for deployments with more stringent performance or compliance requirements because there is no host resource sharing between customers.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-5 Virtual Servers - Private node

Notes:

Private Virtual servers are single-tenant. The host that runs your virtual server is running exclusively on your virtual servers, with no virtual servers of other customers. This option has fewer CPU choices, and is more expensive.

The hypervisor and below is still managed by SoftLayer, so you end up with a private cloud managed by the IaaS provider.

3.9 All virtual servers:

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

All virtual servers:

- Can be provisioned with Windows Server OS, or several types of Linux.
- First and second disk can be local or SAN, plus up to three extra SAN disks (for a total of 5).
- Local disks can be 25/100-300 GB, and SAN disks can be up to 2 TB each.
- Can have uplink port speeds of up to 1 Gbps.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-6 All virtual servers:

Notes:

For all virtual servers, there is a choice of operating systems, including multiple versions of Windows Server and Linux, and even an operating system for Vyatta network appliance (discussed in Advanced Networking unit).

Virtual servers can have up to 5 (virtual) disks. The first two can be local.

SoftLayer updates, features, and prices change periodically to stay competitive in the market and include new models and technologies as they arrive.

3.10 Bare metal servers

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM®

Bare metal servers

- Dedicated (single-tenant) physical hardware servers.
- Premier SoftLayer server offering, and a differentiator in the industry.
- Range of Intel Xeon processors, updated with new offerings.
- You have access to a piece of hardware that your IaaS cloud provider will house, power, cool, and monitor for you, and you pay for what you use. No need for capital expenditure to own hardware.
- Come with hourly or monthly billing.

© 2015 IBM Corporation


Redbooks Workshop

Figure 3-7 Bare metal servers

Notes:

SoftLayer extends the cloud computing paradigm by introducing bare metal servers, which are physical (as opposed to virtual) servers that dedicated to one customer, as an alternative and complement to virtual server offering.

You get all the advantages of physical machine without owning one, and you can pay per use.

Bare Metal servers are provisioned as if they were Virtual Servers, using the same customer interface (web portal and API) and are fully integrated with SoftLayer virtual servers and services.

3.11 Bare metal servers with hourly billing

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Bare metal servers with hourly billing

Dedicated (single-tenant) physical hardware servers:

- Choice of predefined configurations.
- Provisioned in separate racks, with separate public and private VLANs.
- CPU options: with single four-core 3.4 GHz, single four-core 3.5 GHz, dual six-core 2.0 GHz, dual eight-core 2.0 GHz and dual eight-core 2.9 GHz processors. .
- Up to 128 GB RAM.
- Up to 4 1-TB SATA disks or up to 4 x 800 GB SSD disk drives.
- Available with OS - Windows Server, and several types of Linux OS.
- Can be provisioned with VMWare (unmanaged) hypervisor instead of OS, allowing you to create and manage your own private VMWare clouds.
- Uplink port speeds of up to 1 Gbps.
- RAM/disks cannot be upgraded .
- Provisioned typically in less than an hour.

Hourly billing allows affordable short-term deployments while providing a dedicated hardware environment.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-8 Bare metal servers with hourly billing


Notes:

Bare metal hourly servers are pre-configured and are in separate racks. They are not always available in all data centers due to capacity reasons.

The configuration options are limited to just a few processor choices and associated pre-configurations.

What you gain are faster provisioning times and flexibility of hourly billing, while still getting a dedicated piece of hardware.

3.12 Bare metal servers with monthly billing

IBM Digital Sales International Technical Support Organization and Authoring Services 

Bare metal servers with monthly billing

Dedicated (single-tenant) physical hardware servers:

- Wide range of CPU options: Single, dual, or quad 1.60-3.50 GHz processor, with up to 48 cores.
- Up to 3,072 GB of RAM.
- Up to 36 internal disk drives (up to 6 TB SATA, 600 GB SAS, 1.2 TB SSD, with options depending on the data center).
- Can be provisioned with Citrix Xen, Parallels, VMWare (unmanaged), Microsoft HyperV hypervisor instead of OS, allowing you to create and manage your own private clouds of different types.
- Can be provisioned with no operating system allowing you to install almost any x86 operating system or any x86-based virtual appliance.
- Specialty processors with GPU support available for high performance workloads.
- Uplink port speeds of up to 10 Gbps.


© 2015 IBM Corporation 

Figure 3-9 Bare metal servers with monthly billing

Notes:

When you choose monthly bare metal servers, the list of CPU options is greatly expanded. You also see many more options for other configuration elements, like RAM and disks, although some CPU choices place limits on RAM or disks.

The list of operating systems is also expanded. In addition to Linux and Microsoft, you can get operating system for OSNexus storage appliance, and more choices for hypervisor, including popular Citrix Xen, and the No OS option under Other that lets you boot your own OS using the Intelligent Platform Management Interface (IPMI).

3.13 Bare metal servers with monthly billing (continued)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM.

Bare metal servers with monthly billing (continued)

- Intel TXT (Trusted Execution Technology) is available on selected CPUs. It helps workload security, ensuring that your database is running on trusted hardware in a known location.
- Has enhanced system add-ons such as backup agents, control panel (cloud management) software, database software, intrusion detection and protection, advanced monitoring, and dual power supply.
- Provisioned typically in no more than 2-3 hours.
- You can also select CPUs with GPU support and Intel® TXT technology. NVIDIA Grid K2 Graphic Card and NVIDIA Tesla K80 Graphic Card are supported.

Allows for deployments with most stringent performance and compliance requirements. The typical use is for database/custom storage servers.

© 2015 IBM Corporation


Redbooks Workshop

Figure 3-10 Bare metal servers with monthly billing (continued)

Notes:

You can also select CPUs with GPU support and Intel TXT technology. SoftLayer is an ideal platform for gaming and other graphic intensive applications.

System add-ons list is also enhanced.

Bear in mind that SoftLayer periodically adds new configuration elements as they appear on the market, and removes obsolete ones. Therefore, the list in the slide might be different from the one that you see.


All these monthly Bare Metal choices give you freedom in implementing a wide range of business use cases. In addition to typical uses such as database servers, you can use them for your own custom storage solutions, heavy-duty web and application servers, Big Data solutions, private and hybrid clouds, DevOps (development and operations) environments, and vertical industrial solutions like customer relationship management (CRM).

Trusted Execution Technology (TXT) is a technology that is available in select Bare Metal Servers to secure data through a series of encryption keys, launch verified process, and securely boot systems once it verifies that all processes and programs are acting in a predictable manner. It is ideal for customers looking to lock down data that may otherwise be vulnerable and run processes that may deliver such data in a secure manner.

3.14 All server offerings:

IBM Digital Sales

International Technical Support Organization and Authoring Services



All server offerings:

- Can be deployed in any of the data centers (subject to capacity).
- Have unlimited public inbound bandwidth, and metered (hourly) or packaged/unlimited (monthly) outbound bandwidth.
- Free unlimited private network bandwidth.
- Network options include secondary public IP addresses (IPv4 and IPv6).
- System and service add-ons include hardware and software firewalls, antivirus and spyware protection, and monitoring packages with response options.
- Can be mixed and matched in a solution to take advantage of the strengths of each type in different subsystems (tiers).

© 2015 IBM Corporation




Figure 3-11 All server offerings:

Notes:


All types of servers discussed so far can be mixed and matched in an IaaS solution architecture as per your needs. Typically, a web application runs in one or more virtual servers, which can be auto-scaled as explained in later units. Database and other storage solutions can run on bare metal servers.

Development and testing phases can use hourly billing, and production environments can switch to monthly billing.

You can deploy clusters of servers for massively parallel scalable solutions.

Private network traffic is free and can span data centers, so you can design your own replication and disaster recovery schemes.

3.15 Checkpoint questions

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

- What is the difference between virtual server (public node) and virtual server (private node)?
- Who manages the hypervisor running your virtual server (private node)?
- Can two customers share a bare metal server?
- Does bare metal server need a hypervisor? Can it have a hypervisor installed?
- Which server type will be provisioned faster: Virtual or bare metal?



© 2015 IBM Corporation

Figure 3-12 Checkpoint questions


Notes:

(none)

3.16 Checkpoint questions (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint questions

- What is the difference between virtual server (public node) and virtual server (private node)?
A: Multi-tenant vs. single-tenant.
- Who manages the hypervisor running your virtual server (private node)?
A: Your IaaS cloud provider.
- Can two customers share a bare metal server?
A: No, they are dedicated.
- Does bare metal server need a hypervisor? Can it have a hypervisor installed?
A: They don't need a hypervisor if customer wants just to run an OS. However, there are options to install one and run a private cloud.
- Which server type will be provisioned faster: Virtual or bare metal?
A: Virtual

© 2015 IBM Corporation




Figure 3-13 Checkpoint questions (2)

Notes:
(none)

3.17 Ordering servers

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Ordering servers

- Servers can be ordered from the Order window on the customer portal home page, by clicking the **Devices** link.
- You are presented with Order Devices window, which shows types of servers available.

Devices

Find information on ordering devices below. All device order links will redirect to another site.

Bare Metal Servers

Sometimes you need the raw horsepower of bare metal. SoftLayer dedicated servers give you options from entry-level single processor servers to quad proc, hex-core, and even GPU-powered workhorses.

[Hourly](#) [Monthly](#)

Virtual Server (private node)

Single-tenant environment with SoftLayer managed hypervisor. Ideal for applications with stringent resource requirements.

[Hourly](#) [Monthly](#)

Virtual Server (public node)

Multi-tenant environment with SoftLayer managed hypervisor. Ideal for rapid scalability and higher-cost effectiveness.

[Hourly](#) [Monthly](#)

- Click the server type you want, and select hardware, software, and services options. Then click the **Continue Your Order** link.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-14 Ordering servers

Notes:

You can also order devices from the **Devices** → **Device List** window. You end up with the same window showing server types available.

3.18 Provisioning server instances

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Provisioning server instances

- The Order Summary and Billing window shows all the selected options with pricing.
- You can also specify:
 - Public and private VLAN (if your account has more than one provisioned).
 - Provisioning scripts that run automatically after the server is provisioned.
 - Secure Shell (SSH) keys, allowing for more secure login.
 - User metadata, which is server-specific data that can be passed to provisioning scripts.
 - The host and domain name for your server .
- Accept the SoftLayer Master Service Agreement.
- In Devices tab, click **Device List**, and watch for your server become active.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-15 Provisioning server instances

Notes:

VLANs will be explained in more detail in networking units, but you might see VLAN choices when ordering servers, so let's spend a minute on them.

When a first server is provisioned in an account, the account gets a pair of VLANs: Public and private (unless you specify Private VLAN only deployment when selecting Public Bandwidth, which is different from private virtual server, and deploys your server without public VLAN access).

If your account has only one pair of VLANs, all subsequent servers are provisioned in them.

You can purchase additional VLANs by entering a ticket. If you do, a choice of VLANs appears in Order Summary and Billing.

You can specify provisioning scripts, and user metadata typically used to parametrize them.

The host and domain names that you enter are only used for internal naming of your servers, and are not registered with the Domain Name System (DNS).

Provisioning times vary depending on the server type, the software being installed, and other factors.

3.19 Inspecting server details

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Inspecting server details

- The Device List in the Devices tab shows all servers on your account.
- In the **Actions** drop-down menu for a device, you can reboot the device, power it off/on, rename it, upgrade or downgrade it, or cancel the device. If you power down the device, you will be still be paying for it because it consumes data center resources.
- Click the device name to see device details.
- In Device Details window, the details are organized in tabs. The **Actions** menu and the various tabs provide many additional actions that you can start for a device.

Device Details



db01.companya.softlayer.com
Public IP: 173.193.90.102 (Houston 2)
Private IP: 10.77.182.5

[Configuration](#) [Tickets](#) [Bandwidth](#) [Remote Mgmt](#) [Monitoring](#) [Security](#) [Passwords](#) [Storage](#) [Firewall](#)

[Modify Device Configuration](#)

© 2015 IBM Corporation




Figure 3-16 Inspecting server details

Notes:

In Device List, you see all servers on your account that you have permission to see, showing their name, device type, location, public and private IP addresses, start date, and a limited set of Actions.

Click the device name to see device details, organized into set of tabs. These are discussed in detail in Unit 3 Exercise 1.

You can also modify the device configuration in this view, and there is a more comprehensive set of Actions to start, including creating images and reloading servers from them, port control, and so on. These topics are discussed in future units.

3.20 Upgrading or downgrading server instances

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Upgrading or downgrading server instances

- To upgrade or downgrade a server, select the **Upgrade/Downgrade** action on the Device List window. You can also modify elements of your server configuration in the Device Details window.
- You can upgrade or downgrade CPU, RAM, disks, and network options.
- You will have to specify the time window of the upgrade.
- Most upgrades will shut down the server, and you will need to make SoftLayer aware of the current administrator password.

© 2015 IBM Corporation

Redbooks Workshop

Figure 3-17 Upgrading or downgrading server instances

Notes:

Upgrade and downgrade capabilities depend on the server type. For example, bare metal hourly servers are pre-configured and upgrades are limited. Be aware that most upgrades require a server shut-down.

You can request upgrades to be performed right away, or at a later date.

3.21 Accessing server instances

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Accessing server instances

- To access a server running Windows from a Windows computer, you can use Windows Remote Desktop program.
- To access the command prompt on the server from your workstation, you can also use any of the Secure Shell tools available, such as OpenSSH in Linux, or System X or PuTTY on Windows.
- If you generate SSH keys and transfer them to your server when ordering, you can use them to log in securely.
- You can also set up a VNC server on your server, and install, run, and access Linux GUI environments from clients on Windows.
- The initial root/administrator password for your server is available in the Passwords tab in Device Details. This password is not automatically kept in sync with password updates on the device.

© 2015 IBM Corporation


Redbooks Workshop

Figure 3-18 Accessing server instances

Notes:

Once your server is up and running, you can access it using these methods:

- For Windows, use the Remote Desktop program.
- For Linux, the most basic access is command-line. Use an SSH client of your choice to open an terminal window on your server. If you specify a public SSH key during provisioning, you can access your server with your private key. This is a more secure method than using passwords. SSH clients are discussed in the Appendix.

If you want to run one of Linux GUI environments, you need to use a graphical desktop sharing system like VNC, running the server daemon on your server, and a client on your workstation. You will need to install the environment on your server.

If you change the root password on the device, update the password stored in the portal accordingly to enable system updates, reloads, and so on.

3.22 Canceling server instances

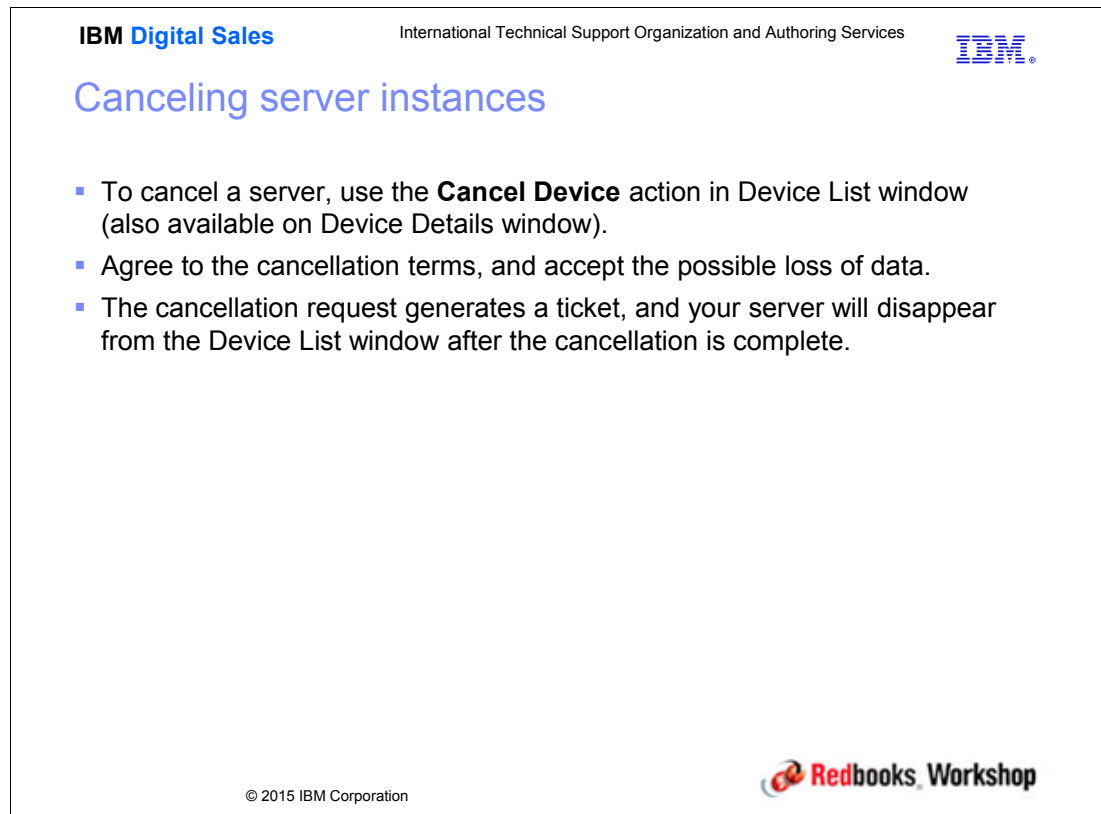


Figure 3-19 Canceling server instances

Notes:


If you do not need your server anymore, in an IaaS cloud environment like SoftLayer, you can just cancel them and they are gone. Server cancellation is the last action in list of **Actions** for each server in the Device List window.

Hourly servers are scheduled for immediate cancellation, and you will stop being charged.

Monthly servers will stay up until your billing anniversary, which is usually the 1st day of next month, and then they are cancelled.

Cancellation involves ticket creation, but it is done automatically.

3.23 Checkpoint questions

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

- What tool do you use to access the command line on a SoftLayer server running Linux?
- Do you need a root password to access a server running Linux?
- If you power down a virtual server, do you stop being charged for it?
- If you cancel a monthly server, when do you stop being charged for it?


© 2015 IBM Corporation 

Figure 3-20 Checkpoint questions


Notes:

(none)

3.24 Checkpoint questions (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint questions

- What tool do you use to access the command line on a SoftLayer server running Linux?
A: Secure Shell (SSH) client
- Do you need a root password to access a server running Linux?
A: No, if you use public key authentication.
- If you power down a virtual server, do you stop being charged for it?
A: No, it still consumes data center resources. You need to cancel it to stop being charged.
- If you cancel a monthly server, when do you stop being charged for it?
A: At your next monthly billing anniversary.

© 2015 IBM Corporation




Figure 3-21 Checkpoint questions (2)

Notes:
(none)



Storage options in cloud computing

The first part of this unit introduces you to the basics of storage technology. It provides an overview of various types of storage such as direct-attached storage (DAS), network-attached storage (NAS), storage area network (SAN), and Object Storage.

Some of the storage types and protocols are discussed in greater detail such as DAS, iSCSI, and NFS. More advanced topics such as RAID arrays are also introduced.

The second part of this unit covers the SoftLayer storage offerings. You will learn details about the block and file storage SoftLayer offers, and about the performance and endurance offerings of both block and file storage. It also covers object storage in SoftLayer and the use cases for it.

The last part of this unit introduces you to these exercises:

- ▶ Ordering and mounting Performance Block storage
- ▶ Ordering and mounting an Endurance File storage share
- ▶ Ordering Object Storage

4.1 What you should be able to do

On completion of this unit, you should be able to do the following:

- ▶ Explain storage basic concepts
- ▶ Describe the SoftLayer storage offerings
- ▶ Be able to make an educated decision when ordering storage from SoftLayer

4.2 References

The following items are useful for further research:

- ▶ SoftLayer Cloud Storage:
<http://www.softlayer.com/cloud-storage>
- ▶ SoftLayer KnowledgeLayer:
<http://knowledgelayer.softlayer.com/>
- ▶ Which storage solution is best for your project?
<http://blog.softlayer.com/tag/san>

4.3 Teaching topics

- ▶ Introduction to storage:
 - DAS – Direct Attached Storage
 - SAN – Storage Area Network
 - NAS – Network attached storage
 - Object Storage
 - RAID arrays
- ▶ SoftLayer Storage offerings:
 - Available local disk types
 - Block Storage
 - File Storage
 - Endurance
 - Performance
 - Object storage

4.4 Introduction to storage: Storage types

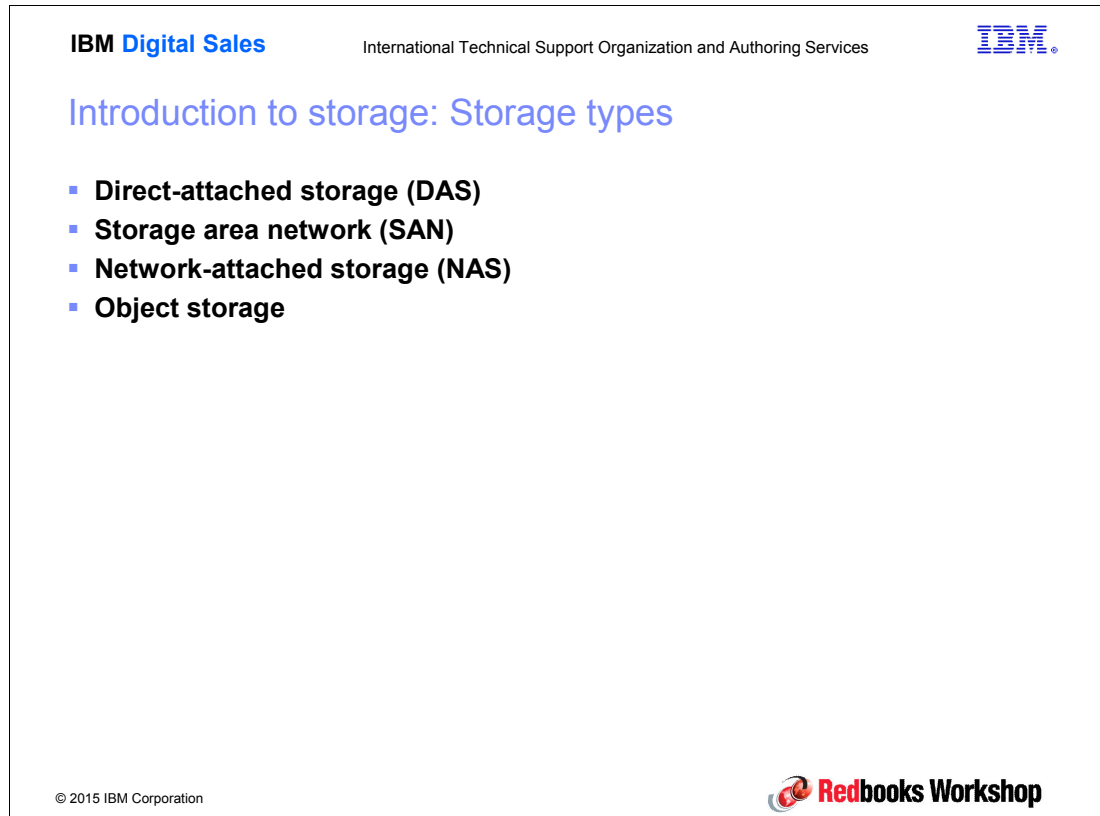


Figure 4-1 Introduction to storage: Storage types

Notes:

The terms *storage type* and *storage protocols* are often confused. You can find a lot of documents where they are used interchangeably. For sake of simplicity in this course, the storage types are DAS, SAN, NAS, and Object Storage. The storage protocols are NFS, FC, iSCSI, FCoE, and CIFS.

This unit discusses direct-attached storage (DAS) which is basically local disk on system, and storage area network (SAN) which is a remotely access block storage. It also describes NAS, which is remotely accessed file storage, and object storage which is a unit approach to storing data and is popular in Cloud Computing.

4.5 Introduction to storage: DAS

IBM Digital Sales

International Technical Support Organization and Authoring Services



Introduction to storage: DAS

Direct-attached storage (DAS) is digital storage directly attached to the computer accessing it using the SATA, SAS, or USB interface. Examples of DAS include hard drives, optical disc drives, and storage on external drives directly attached to the system.



© 2015 IBM Corporation



Figure 4-2 Introduction to storage: DAS

Notes:

As mentioned, DAS is basically the local disk of the system. Direct-attached storage (DAS) is digital storage directly attached to the computer, accessing it using the SATA, SAS, or USB interface. Examples of DAS include hard drives, optical disc drives, and storage on external drives directly attached to the system.

Basically, any storage that is directly attached to your server using your internal Storage BUS is considered DAS.

4.6 Introduction to storage: DAS (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Introduction to storage: DAS

Storage Interfaces

SATA	<ul style="list-style-type: none">• Commonly used in desktops and laptops• Used to connect HDDs, SSDs, and CD/DVD drives• Has low performance
SAS	<ul style="list-style-type: none">• Available in most modern server systems• Used to connect HDDs and SSDs• Provides good performance• Provides backward compatibility with SATA

Disk Types

SATA HDD	Due to ATA technology specifics, SATA HDDs are slower (7200 rpm)
SAS HDD	SAS drives perform much better than SATA drives (10000 or 15000 rpm)
SSD	Does not contain any mechanical component. Uses NAND-based memory. Faster than HDDs.

© 2015 IBM Corporation




Figure 4-3 Introduction to storage: DAS (2)

Notes:

These tables explain the storage interfaces and disk types that can be used.

Although disk types are presented in DAS section, they are applicable for external storage systems as well. Understand that any external storage system is a server machine with a lot of local disks.

The following are the available storage interfaces:

- ▶ **SATA (Serial ATA):** The standard computer bus interface for connecting hard drives (HDDs), solid state drives (SSDs), and CD/DVD drives to the computer. Disks with the SATA interface are mostly used in modern desktop systems, but can be used in server systems as well.
- ▶ **SAS (Serial Attached SCSI):** A point-to-point serial protocol that moves data to and from computer storage devices such as hard drives and tape drives. SAS replaces the older Parallel SCSI bus technology that first appeared in the mid-1980s. SAS, like its predecessor, uses the standard SCSI command set.

SAS is considered the more enterprise-ready solution than SATA, but it is also more expensive.

These SATA and SAS interfaces interact with three main types of disks:

- ▶ **SATA Hard Disk Drives** are usually cheaper and, due to ATA technology specifics, run at lower speeds (usually 7200 RPM).

- ▶ SAS compatible hard disk drives are usually faster than SATA drives. You can have SAS drives working on 10000 rpm and even on 15000 rpm.
- ▶ Solid-State drive (SSD) is a solid-state storage device that uses integrated NAND-based flash memory chip to store data persistently. It contains no actual disk, nor a drive motor. SSDs are fast and currently quite expensive. This is the reason why at these moment SSD drives are typically used as supporting mechanism for backend HDDs, such as for read cache or write cache.

4.7 Introduction to storage: SAN

Introduction to storage: SAN

SAN (storage area network) is a dedicated, high-speed network that provides access to consolidated, block-level data storage. SAN is usually accessible to servers so that they can connect to storage and that storage appear like locally-attached devices to the operating system. A SAN does not provide file abstraction, only block-level operations.



© 2015 IBM Corporation

Redbooks Workshop

Figure 4-4 Introduction to storage: SAN

Notes:

SAN is a dedicated network that is used to provide storage access to servers. The storage which is accessed over SAN is block level, which means it will appear as a raw device to the operating system on the server. You will have to format it with a compatible file system before you can use it. This is the main difference between SAN solutions and other storage solutions.

4.8 Introduction to storage: SAN (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Introduction to storage: SAN

SAN Protocols

FC	<ul style="list-style-type: none">• Often used in enterprise solutions.• High-speed: 4, 6, or 16 Gbps• Uses fiber optic network• Implementation cost: High
FCoE	<ul style="list-style-type: none">• Often used in converged solutions• High-speed• Uses 10G IP network• Requires specific hardware• Implementation cost: Medium
iSCSI	<ul style="list-style-type: none">• Used by most cloud providers• Speed depends on speed of network• Works over usual 1G or 10G IP network• Does not require specific hardware• Implementation cost: Low

© 2015 IBM Corporation


Redbooks Workshop

Figure 4-5 Introduction to storage: SAN (2)

Notes:

SAN supports several protocols. The following are the most commonly used protocols:

- ▶ Fibre Channel (FC) is a high-speed network technology primarily used to connect to computer data storage. Fiber channel is commonly used for enterprise solutions in modern data centers. It can provide high-speed access to storage (4, 8, or 16 Gb/s).
- ▶ Fibre Channel over Ethernet (FCoE) is a transport protocol (similar to TCP used in IP networks) that predominantly transports SCSI commands over Fibre Channel networks.

Fibre Channel SAN uses optical network for communication. That makes it very fast, but also very expensive, so it is mostly used in high-end enterprise solutions.

FCoE encapsulates FCP packets into usual Ethernet packets, but to use FCoE efficiently you will need to run at least a 10 Gb network, and your network equipment must support FCoE.

In modern IT, 10 Gb Ethernet cards are becoming common, so FCoE has become more popular. Basically, customers do not want to pay for additional Fibre Channel HBA if they can use 10 Gb NIC card for storage accessibility.


- ▶ Internet Small Computer System Interface (iSCSI) works on top of TCP, and allows the SCSI command to be sent end-to-end over local area networks (LANs), wide area networks (WANs), or the Internet.

The benefit of iSCSI is that it does not have any specific HW requirements. It can work across any LAN. And although it will definitely benefit from a 10 Gb network, it is not

mandatory. In addition, it is also quite simple to implement. All modern operating systems have a way to run as iSCSI initiator.


iSCSI is one of the main approaches used by cloud providers to deliver block-based remote storage. SoftLayer uses iSCSI for its block storage offering as well, so it will be covered in more detail later in this unit.

4.9 Introduction to storage: NAS

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Introduction to storage: NAS

Network-attached storage (NAS) is a file-level storage system connected to a network providing data access to a heterogeneous group of clients. NAS is specialized for serving files, rather than block level storage.




© 2015 IBM Corporation

Figure 4-6 Introduction to storage: NAS

Notes:

Network-attached storage (NAS) is a file-level storage system connected to a network providing data access to a heterogeneous group of clients. NAS is specialized for serving files, rather than block level storage. The two most common NAS solutions are SMB/CIFS and NFS. NAS, unlike SAN, provides file-level storage.

4.10 Introduction to storage: NAS (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Introduction to storage: NAS

NAS Protocols

SMB/CIFS	<ul style="list-style-type: none">• Standard File sharing mechanism in Microsoft operating systems• Uses IP network
NFS	<ul style="list-style-type: none">• Widely used by all *NIX-based operating systems• Uses IP network

© 2015 IBM Corporation


Redbooks Workshop

Figure 4-7 Introduction to storage: NAS (2)

Notes:

The most common protocols used with SAN are:

► **SMB/CIFS:**

Server Message Block (SMB) is a file sharing protocol which was invented by IBM in 80s. Directories which were made available over network are called Shares

Common Internet File System (CIFS) is a so-called dialect of SMB. Basically CIFS is implementation of SMB created by Microsoft. Currently CIFS is considered the default/native file sharing mechanism for Microsoft Operating systems.

► **Network File System (NFS)** is a distributed file system protocol allowing a user on a client computer to access files over a network much like local storage is accessed. NFS is supported by default in many operating systems, especially *NIX based ones.

Because it is easy to implement, NFS is used by many Cloud providers as a default protocol for file-based storage shares. SoftLayer uses NFS in its file storage offering, so it will be covered in more detail later.

4.11 Introduction to storage: Object Storage

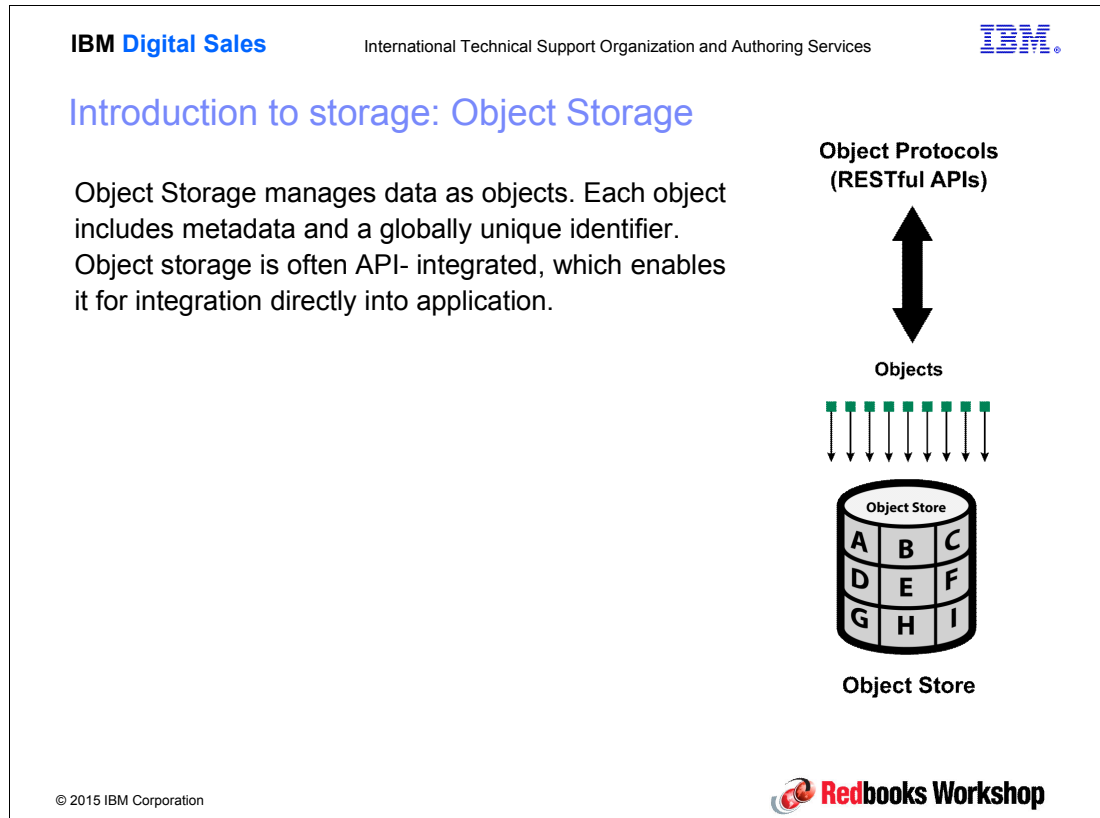


Figure 4-8 Introduction to storage: Object Storage

Notes:

Object storage manages data as objects. Each object includes metadata and a globally unique identifier. Object storage is often API- integrated, which enables it for integration directly into application.

Object storage can be used to store files like Virtual Machine images, backups, and archives as well as photos and videos. Object Storage can be integrated with CDN, as described in Unit 9.

Most cloud-based storage available on the market uses an object storage architecture, including Amazon S3, Google Cloud Storage, and OpenStack Swift.

4.12 Introduction to storage: RAID arrays

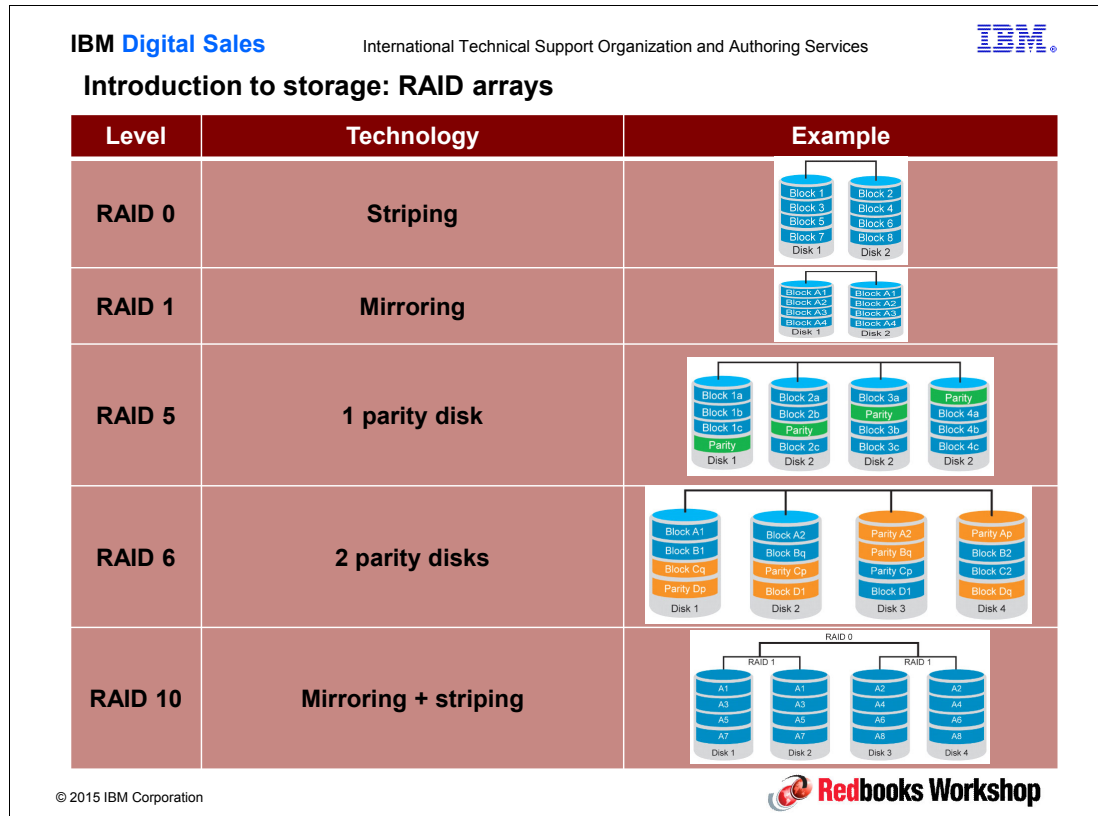


Figure 4-9 Introduction to storage: RAID arrays

Notes:

This slide gives a broad overview of RAID levels. A detailed explanation is beyond the scope of this presentation.

- ▶ RAID 0 (also known as a stripe set or striped volume) splits (“stripes”) data evenly across two or more disks, without parity information, redundancy, or fault tolerance. RAID 0 provides good performance for both read and write, but no redundancy.
- ▶ RAID 1 consists of an exact copy (or mirror) of a set of data on two or more disks. Because the data is mirrored on all disks belonging to the array, the array can only be as big as the smallest member disk. RAID 1 performs well on reads because reads can be served by any member of the mirror, but write performance remains at the single disk level.
- ▶ RAID 5 is a RAID configuration that uses disk striping with parity. Because data and parity are striped across all of the disks, no single disk is a bottleneck. Striping also allows users to reconstruct data in case of a disk failure. Reads and writes are more evenly balanced in this configuration, making RAID 5 the most commonly used RAID method.
- ▶ RAID 6 extends RAID 5 by adding another parity block. It uses block-level striping with two parity blocks distributed across all member disks. RAID 6 does not have a performance penalty for read operations, but it does have a performance penalty on write operations because of the processing associated with parity calculations.

- RAID 10 is a combination of RAID 1 and 0, and is often denoted as RAID 1+0. It combines the mirroring of RAID 1 with the striping of RAID 0. It's the RAID level that gives the best performance, but it is also costly, requiring twice as many disks as other RAID levels.

4.13 iSCSI: Components

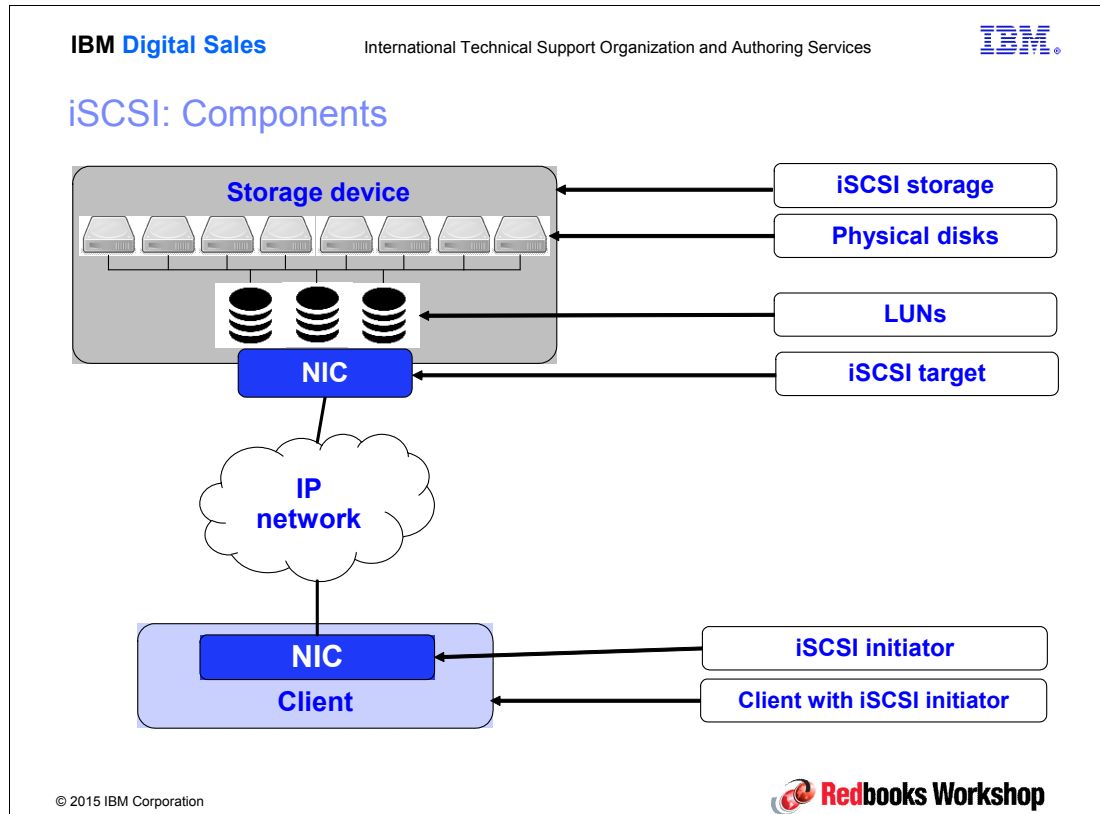


Figure 4-10 iSCSI: Components

Notes:

This diagram shows the components of basic iSCSI SAN:

- ▶ The iSCSI storage device can be either a dedicated storage system or a server with a storage appliance such as OSNexus Quantastor installed. Quantastor is covered in Unit 8. The storage device has its local disk built into some kind of RAID array, LUNs, or Storage volumes that are created on that RAID array.
- ▶ The NICs of the storage system act as iSCSI targets, which means that the NIC is the point of access for iSCSI initiators.
- ▶ A server accesses the storage device over the usual IP network. The NIC of the server acts as the iSCSI initiator. The iSCSI initiator transmits SCSI command to the iSCSI target. This initiator can be either hardware based or software based.

A software-based initiator is software that enables something, usually NIC, to act as the iSCSI initiator. There is an implementation of Software iSCSI initiator for every OS. Because it is fully software based, all processing of traffic is handled by the system CPU.

A hardware iSCSI initiator is a specialized physical device that can offload SCSI commands processing.

You can have several iSCSI initiators and iSCSI targets per system.

4.14 iSCSI addressing

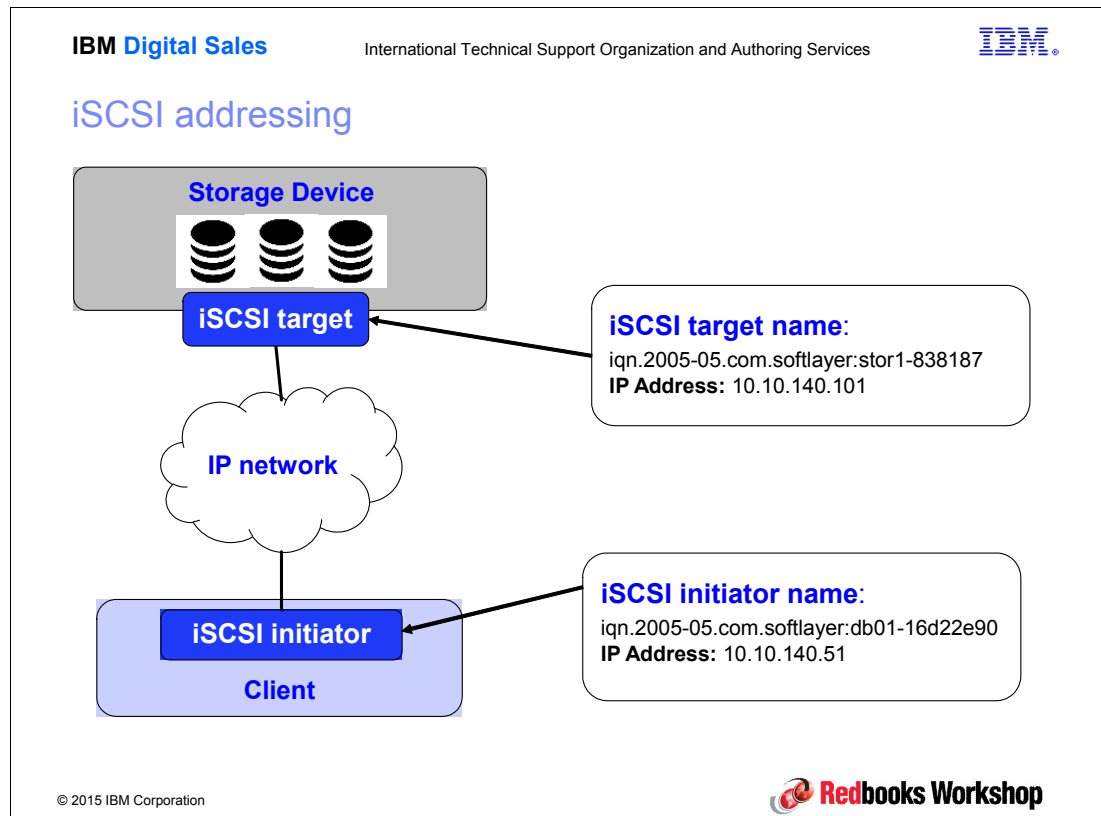


Figure 4-11 iSCSI addressing

Notes:

Although iSCSI uses IP network for protocol itself, higher level naming is used to address the objects within the protocol. The most common naming format is iSCSI qualified name (IQN). IQN has the following format:

- ▶ Literal IQN (iSCSI Qualified Name)
- ▶ Date (yyyy-mm) that the naming authority took ownership of the domain
- ▶ Reversed domain name of the authority (e.g. com.ibm, com.softlayer)
- ▶ Optional “:” prefixing a storage target name specified by the naming authority

You can see examples on the slide.


You might also encounter these formats:

- ▶ Extended Unique Identifier (EUI), which has the following format:
eui.{EUI-64 bit address} (e.g. eui.02004567A425678D)
- ▶ T11 Network Address Authority (NAA), which has the following format:
naa.{NAA 64 or 128 bit identifier} (e.g. naa.52004567BA64678D)

4.15 iSCSI access control

IBM Digital Sales

International Technical Support Organization and Authoring Services



iSCSI access control

- **ACL**

Access control lists can be used on storage devices to control which iSCSI initiator can access certain iSCSI targets. ACL are based on IQNs.
- **CHAP**

Challenge Handshake Authentication Protocol (CHAP) can be used to allow iSCSI initiator to prove its identity to iSCSI targets. It is also possible to configure bidirectional CHAP for better security.

© 2015 IBM Corporation




Figure 4-12 iSCSI access control

Notes:

Access control lists (ACLs) can be used on your storage device to control which iSCSI initiator can access certain iSCSI targets. ACLs are based on IQNs.

Challenge Handshake Authentication Protocol (CHAP) can be used to allow iSCSI initiator to prove its identity to iSCSI Target. It is also possible to configure bidirectional CHAP for better security.

4.16 NFS components

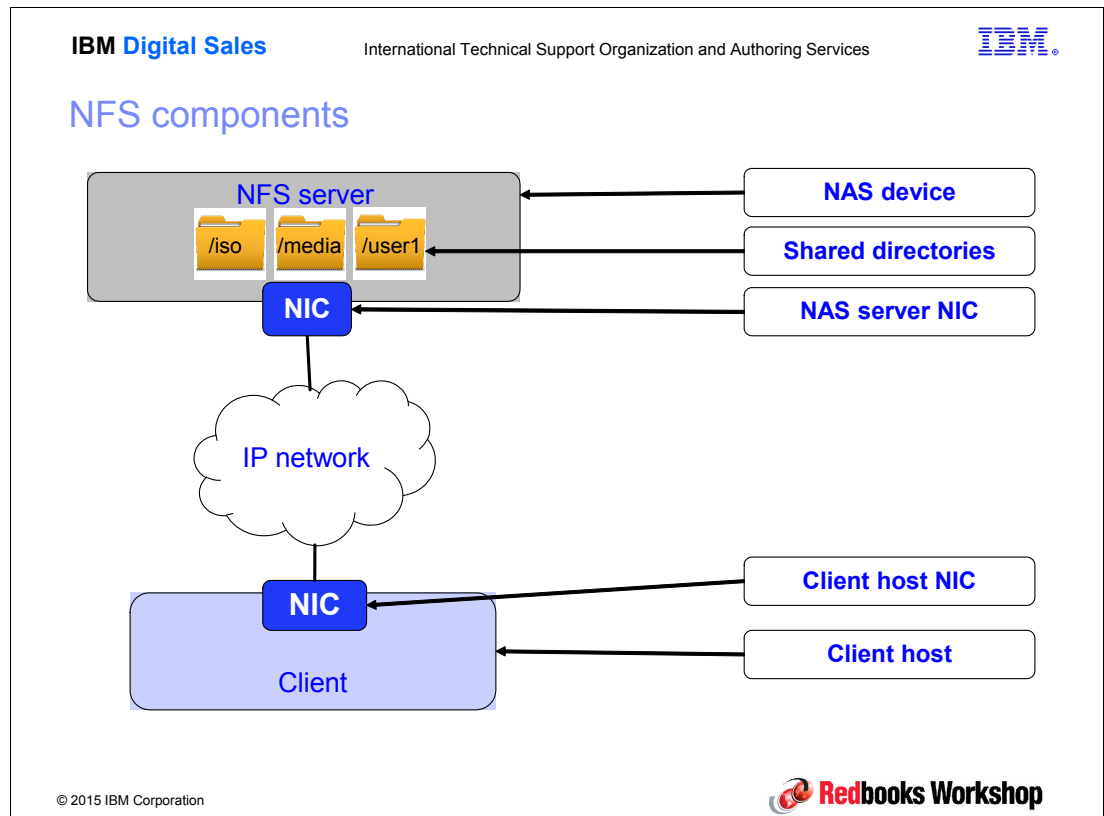


Figure 4-13 NFS components

Notes:

This diagram shows the components for NFS storage. The NFS server can be either a dedicated storage device, a specialized storage appliance installed on usual server, or even a Linux machine configured to act as an NFS server. The server shares the directories shown, and the client accesses the shares over the usual IP network.

4.17 NFS addressing

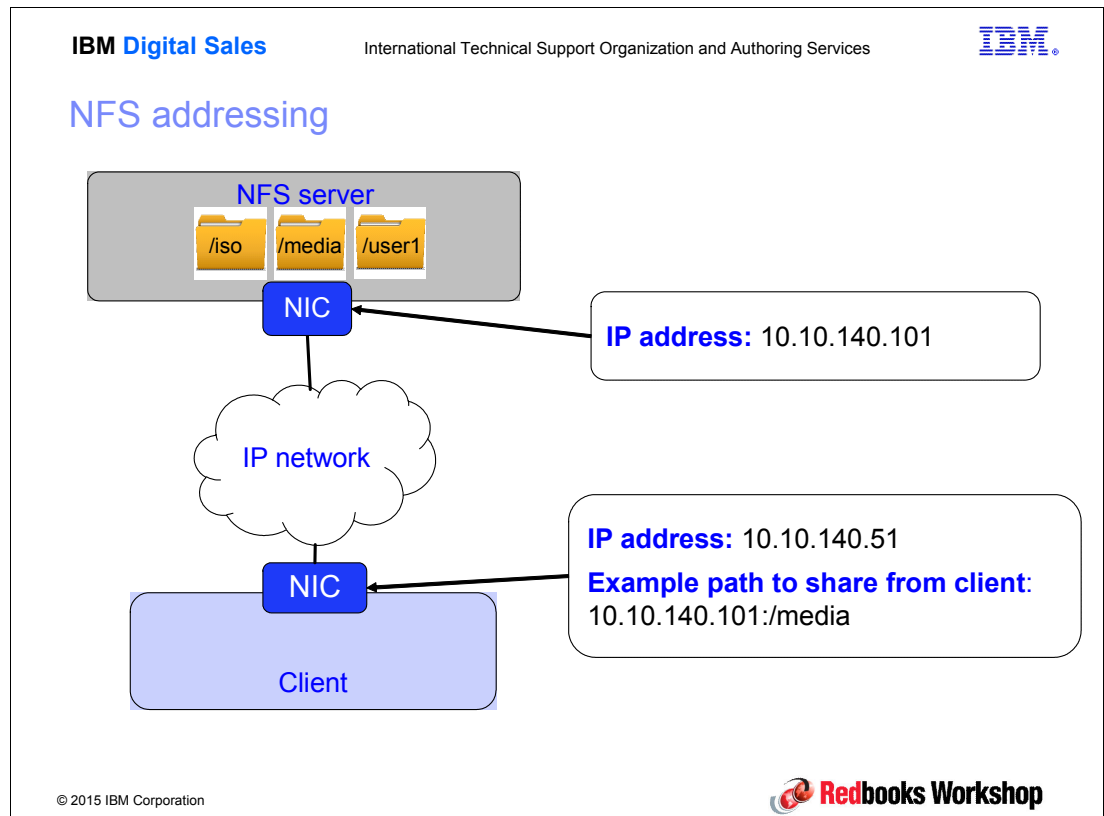


Figure 4-14 NFS addressing

Notes:

There is no high level addressing in case with NFS. All communication is done using IPs and host names.

The share should be mounted on client so that files in it can be accessed. The path to the share has following convention:

IP_OF_NFS_SERVER:/SHARE_NAME

As an example, for /iso it would be: 10.10.140.101:/iso.

4.18 NFS access control

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

NFS access control

The most common way to control access to NFS shares are IP-based or subnet-based ACLs.

© 2015 IBM Corporation


Redbooks Workshop

Figure 4-15 NFS access control

Notes:

The most common way to control access to NFS share is IP based or subnet based ACLs. More advanced access control mechanisms are available such as integration with Kerberos, Microsoft Active Directory, or LDAP.

4.19 SoftLayer storage offerings: DAS

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer storage offerings: DAS

- **For Bare-Metal servers**
 - SAS
 - SATA
 - SSD
- **For Virtual Servers**
 - SAN
 - Local disk
 - Portable storage volume

© 2015 IBM Corporation


 Redbooks Workshop

Figure 4-16 SoftLayer storage offerings: DAS

Notes:

When ordering a bare-metal server from SoftLayer, you can select between three main disk types:

- ▶ SAS disks
- ▶ SATA disks
- ▶ SSDs

The selection should be done based on your needs. For most cases, use more reliable disks like SAS, but SAS is more expensive. For example, highly loaded databases could make a good use of SAS disks, build into a RAID 10 array, with some SSD disks acting as the cache. However, a Mail archive server could use cheaper SATA disks build into a very reliable RAID6 array. Therefore, the decision on which disk should be used depends on business case.

In addition, during order process you can configure a RAID array on local disks. This is useful to have availability configured before the operation system is installed. You can order a bare-metal server with 4, 6, 12, 24, or 36 disk slots. This can help you to build large and complex solutions.

SoftLayer virtual servers can be deployed with primary storage based on local disk or SAN, and with portable storage volumes as secondary storage. Whether your application needs higher disk I/O, resiliency, or long-term flexibility, you can match your virtual server's storage to its application.

Portable storage volumes are secondary storage solutions that are exclusively available on SoftLayer's virtual servers. They can be connected to one virtual server at a time, and are an ideal solution when looking to transfer data between virtual servers in any data center on the SoftLayer network. Portable storage volumes are useful for database applications that require access to raw, unformatted block-level storage, and for moving large data sets between virtual servers.


There is no difference from the operating system perspective between these three configurations.

4.20 SoftLayer storage offerings: Block storage

IBM Digital Sales International Technical Support Organization and Authoring Services **IBM**

SoftLayer storage offerings: Block storage

- iSCSI-based block storage
- Accessible over SoftLayer private network
- Running on reliable storage infrastructure



© 2015 IBM Corporation **Redbooks Workshop**

Figure 4-17 SoftLayer storage offerings: Block storage

Notes:

SoftLayer has its own naming convention for storage offerings. SoftLayer's block storage offering, as you can understand from its name is a SAN offering. According to its description, it works over private networks. SoftLayer uses iSCSI for its block storage offering.

4.21 SoftLayer storage offerings: File storage


IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer storage offerings: File storage

- NFS based file storage
- Accessible over SoftLayer private network
- Running on reliable storage infrastructure



© 2015 IBM Corporation


 Redbooks Workshop

Figure 4-18 SoftLayer storage offerings: File storage

Notes:

The file storage offering provides NAS storage. The File storage offering of SoftLayer provides NFS-based volume, which can be mounted to your systems.

4.22 SoftLayer storage offering: Storage options



Figure 4-19 SoftLayer storage offering: Storage options

Notes:

SoftLayer has two options for both block and file storage:

- ▶ Endurance
- ▶ Performance

Some of the characteristics of these offerings are described in the next two slides.

4.23 SoftLayer storage offering: Storage options (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer storage offering: Storage options

Endurance is a new class of block and file storage from SoftLayer. It brings an effective feature set to help you fulfill availability requirements of your data. At the same time, it provides consistent performance baseline.

Endurance has these notable features:

- Snapshots: Capture point-in-time data snapshots non-disruptively
- Replication: Automatically copies snapshots to a partner SoftLayer data center
- Concurrent access: Allows multiple hosts to simultaneously access both block and file volumes.

© 2015 IBM Corporation



Redbooks Workshop

Figure 4-20 SoftLayer storage offering: Storage options (2)

Notes:

Endurance storage provides advanced availability features such as snapshots and replication. These features are described in detail in Unit 8. Having these features enabled does have a negative effect on performance.

4.24 SoftLayer storage offering: Storage options (3)

IBM Digital SalesInternational Technical Support Organization and Authoring Services

SoftLayer storage offering: Storage options

Performance is an option of SoftLayer's Block and File Storage offerings that is designed to support high I/O applications requiring predictable levels of performance. This predictable performance is achieved through the allocation of protocol-level input/output operations per second (IOPS) to individual volumes.

Performance storage does not support snapshots and replication.


© 2015 IBM Corporation

Figure 4-21 SoftLayer storage offering: Storage options (3)

Notes:

Performance storage is designed to fulfill high IO demand. If you are planning to run an application with predictable IO demand, Performance storage is good option. SoftLayer provides tools to determine which Storage option suits your needs. For more information, see:

<http://www.SoftLayer.com>

4.25 SoftLayer storage offering: Object storage

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer storage offering: Object storage

Object Storage in Softlayer is a redundant and highly scalable cloud storage service that allows users to easily store, search, and retrieve data across the Internet, with optional CDN connectivity, or across SoftLayer's global private network. It is based on the OpenStack Swift platform.

© 2015 IBM Corporation


Redbooks Workshop

Figure 4-22 SoftLayer storage offering: Object storage

Notes:

Object storage in SoftLayer can be hosted in many SoftLayer data centers around the world and integrated with CDN. In SoftLayer, you can use Swift APIs or one of the language clients to control your Object Storage objects.

Object storage in SoftLayer is based on OpenStack Swift, which is an open source object storage implementation developed by the OpenStack project. Swift functions as a distributed, API-accessible storage platform that can be integrated directly into applications or used to store files like VM images, backups, and archives as well as photos and videos.

4.26 Overview

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Overview

	Object	File	Block
Units	Objects (include object ID, data, and meta data).	Files.	Blocks, simply a series of 0s and 1s.
Access method or protocol	API.	NFS, CIFS.	Direct Attachment, FC, FCoE, iSCSI.
Common Use case	Static data.	Shared file data.	Frequently changing and transaction data.
Strength	Scalability and distributed access.	Simplified access and management of shared file, including directories for organization.	High performance. Data is stored without concept of format or type. Relies on applications or file system for data location.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 4-23 Overview

Notes:

This table shows a side-by-side view of the three storage types that have been discussed. For more information about what to consider when selecting a storage solution for your project, see *Which storage solution is best for your project?* at:

<http://blog.softlayer.com/tag/san>

4.27 Check point: Questions

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Check point: Questions

- Is Storage in iSCSI SAN on a block or file level?
- How big will a RAID 0 array consisting of two 500 GB disks be?
- Which protocol is used for the Softlayer File Storage offering?

© 2015 IBM Corporation


 Redbooks Workshop

Figure 4-24 Check point: Questions

Notes:

(none)

4.28 Check point: Questions and answers

IBM Digital Sales International Technical Support Organization and Authoring Services **IBM**

Check point: Questions and answers

- Is storage in iSCSI SAN on a block or file level?
 - Block Level
- How big will a RAID 0 array consisting of two 500 GB disks be?
 - 1000 GB
- Which protocol is used for SoftLayer File Storage offering?
 - NFS

© 2015 IBM Corporation




Figure 4-25 Check point: Questions and answers

Notes:

(none)

4.29 Recap

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Recap

We now know about:

- The basics of storage
- The available storage types
- SoftLayer storage offerings
- Endurance and performance storage

In the lab coming up, we will:

- Order and mount block storage
- Order and mount Endurance storage
- Order object storage



© 2015 IBM Corporation

 **Redbooks Workshop**

Figure 4-26 Recap

Notes:
(none)



Networking options in cloud computing

This unit provides you with a basic understanding of cloud computing networking in general. It describes SoftLayer's network topology, options, and configuration items. This unit includes and overview of the available networking tools.

5.1 What you should be able to do

After completing this unit, you should be able to:

- ▶ Describe the basics of the network in Cloud Computing.
- ▶ Explain the network strategy and topology of SoftLayer.
- ▶ Define and list the characteristics of public, private, and out-of-band management networks.
- ▶ List the available network tools.
- ▶ Explain the function of VPN connections in cloud computing
- ▶ Identify the uses of SoftLayer Direct Link.

5.2 References

The following websites are useful for further research:


- ▶ Network details about SoftLayer:
<http://www.softlayer.com/network>
- ▶ Networking details about SoftLayer:
<http://www.softlayer.com/networking>
- ▶ SoftLayer Knowledge Layer
<http://knowledgelayer.softlayer.com/>

5.3 Teaching topics

The following topics are covered in this unit:

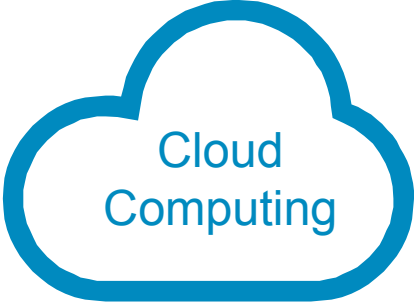
- ▶ The network in Cloud Computing
- ▶ SoftLayer network strategy and topology
- ▶ Definition and characteristics of:
 - Public networks
 - Private networks
 - Out-of-Band management network
- ▶ Network tools
- ▶ VPN connections in Cloud Computing
- ▶ Direct Link use case
- ▶ Recap
- ▶ Checkpoint
- ▶ References

5.4 What is cloud computing?

IBM Digital SalesInternational Technical Support Organization and Authoring Services

What is cloud computing?

To define cloud computing, you can say that is the use of “outsourced” computing resources that can be employed or accessed through **networking** or the internet.



Cloud computing is a model for enabling convenient, on-demand **network access** to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.


© 2015 IBM Corporation

Figure 5-1 What is cloud computing?

Notes:
(none)

5.5 The network is critical to Cloud Computing

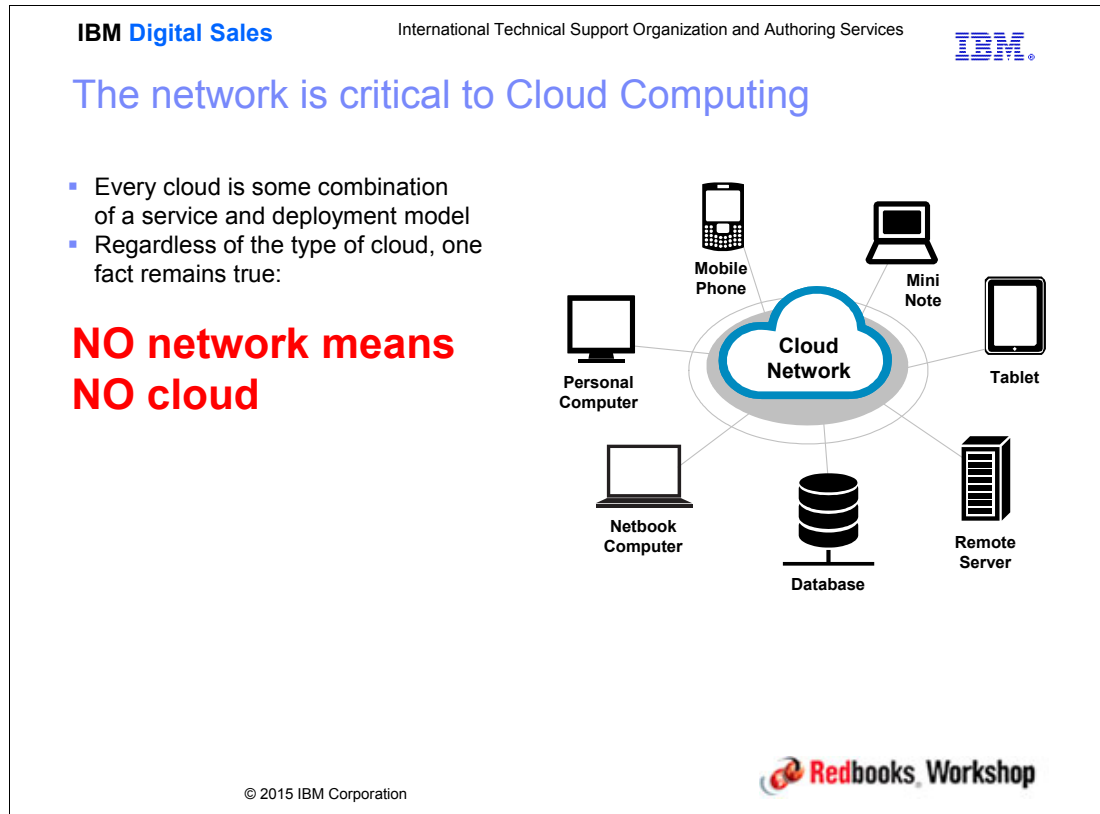


Figure 5-2 The network is critical to Cloud Computing

Notes:

From a networking standpoint, each service model requires the cloud provider to expose part or all of the network, and provide more or less networking capabilities to cloud users.

Each service model requires cloud users to understand and design more or less of the network to which they are exposed.

The network is most exposed in the IaaS model, and least exposed in the SaaS model.

Without networks, users cannot access their cloud services. Without networks, applications, data, and users cannot move between clouds. Without networks, the infrastructure components that must work together to create a cloud cannot.

5.6 Networking had to change

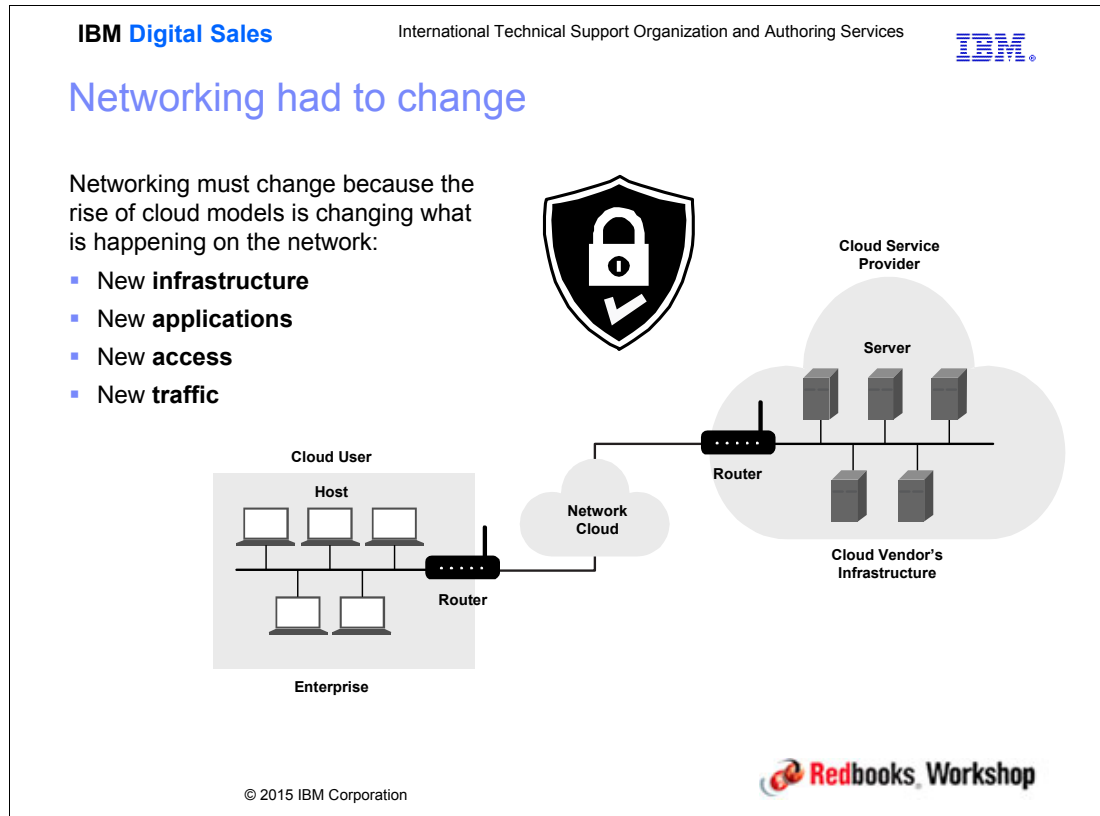


Figure 5-3 Networking had to change

Notes:

- ▶ New infrastructure: Everything is becoming virtualized, infrastructure is becoming programmable, and servers and applications have mobility.
- ▶ New applications: Data-intensive analytics, parallel and clustered processing, telemedicine, remote experts, and community cloud services.
- ▶ New access: Mobile device-based access to everything and virtual desktops.
- ▶ New traffic: Predominantly server-to-server traffic patterns and location-independent endpoints on both sides of a service or transaction.

5.7 Attributes of Cloud Networking

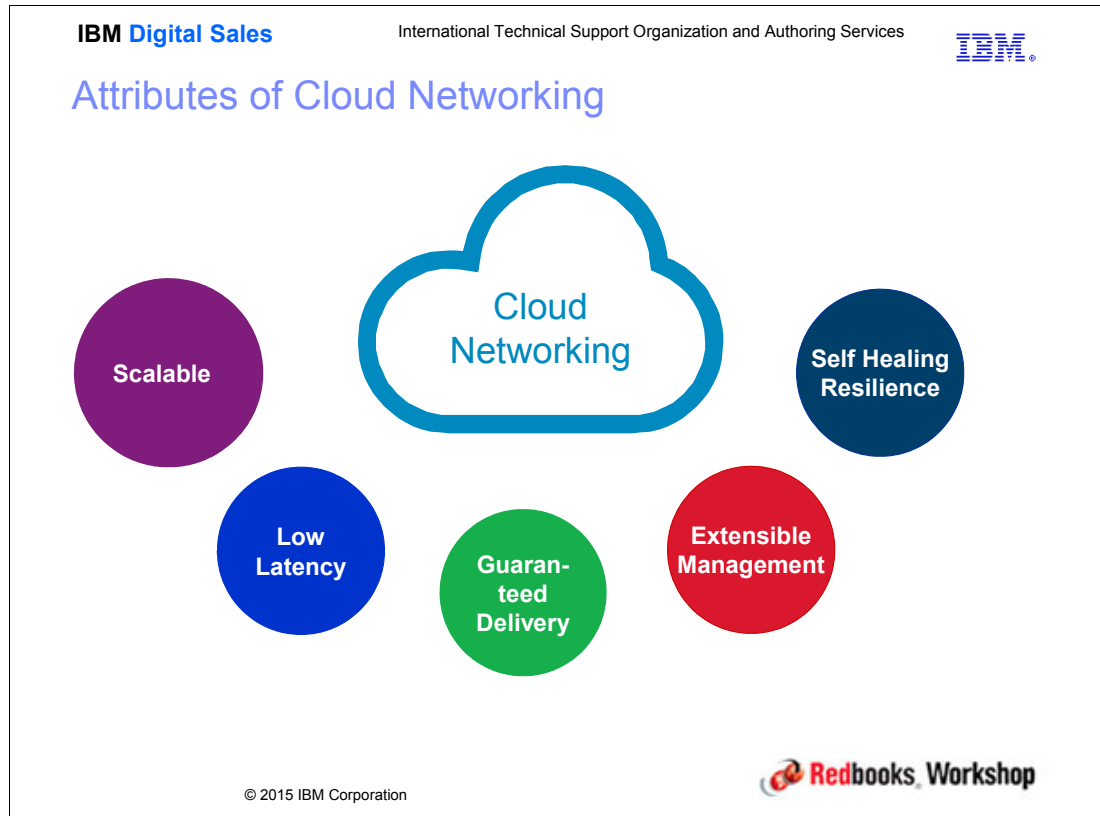


Figure 5-4 Attributes of Cloud Networking

Notes:

1. **Scalability:** The cloud network must scale to the overall level of throughput required to ensure that it does not become a bottleneck. This means that the cloud networking fabric must handle throughputs that will soon reach trillions of packets.
2. **Low Latency:** The cloud network must deliver microsecond latency across the entire network fabric because low latency improves application performance and server utilization. For latency sensitive applications, 10-Gigabit Ethernet is a major improvement.
3. **Guaranteed Performance:** The cloud network must provide predictable performance to service many simultaneous applications in the network, including video, voice, and web traffic.
4. **Extensible Management:** Real-time upgrades and image/patch management in a large cloud-network is a daunting challenge to network administrators. A vastly simpler approach is required to handle networks of this size, which automates provisioning, monitoring, maintenance, upgrading, and troubleshooting.
5. **Self-Healing Resilience:** Cloud networks operate 24x7, so downtime is not an option. This requires a network architecture that offers self-healing and the ability for transparent in-service software updates. On most switches, any software fault results in a reload, resulting in seconds or even minutes of downtime.

5.8 Example: Data center switch network architecture

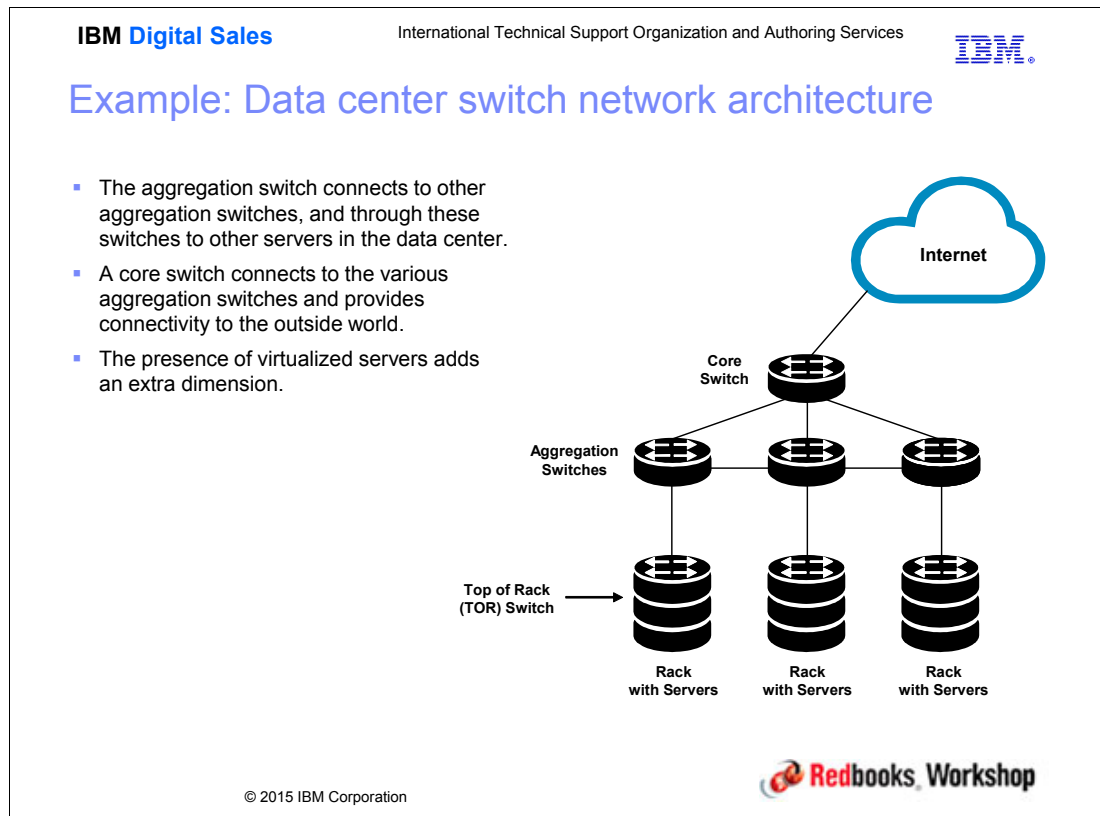


Figure 5-5 Example: Data center switch network architecture

Notes:

The most common network architecture for enterprises is the three-layer architecture with access, aggregation or distribution, and core switches. The data center requires a slightly different variation of this layering, as proposed by some vendors. The data center consists mainly of servers in racks interconnected through a Top-of-Rack (TOR) Ethernet switch that connects to an aggregation switch, sometimes known as an End-of-Rack (FOR) switch.

The aggregation switch connects to other aggregation switches and through these switches to other servers in the data center. A core switch connects to the various aggregation switches, and provides connectivity to the outside world, typically through Layer 3 (IP). It can be argued that most intra-data center traffic traverses only the TOR and aggregation switches. Therefore, the links between these switches and the bandwidth of those links need to account for the traffic patterns.

The presence of virtualized servers adds an extra dimension. Network connections to physical servers need to involve “fatter pipes” because traffic for multiple VMs is multiplexed onto the same physical Ethernet connection. This result is expected because you have effectively collapsed multiple physical servers into a single physical server with VMs. It is common to have servers with 10-Gbps Ethernet cards in this scenario.

5.9 Virtual Ethernet switch in a virtualized server environment

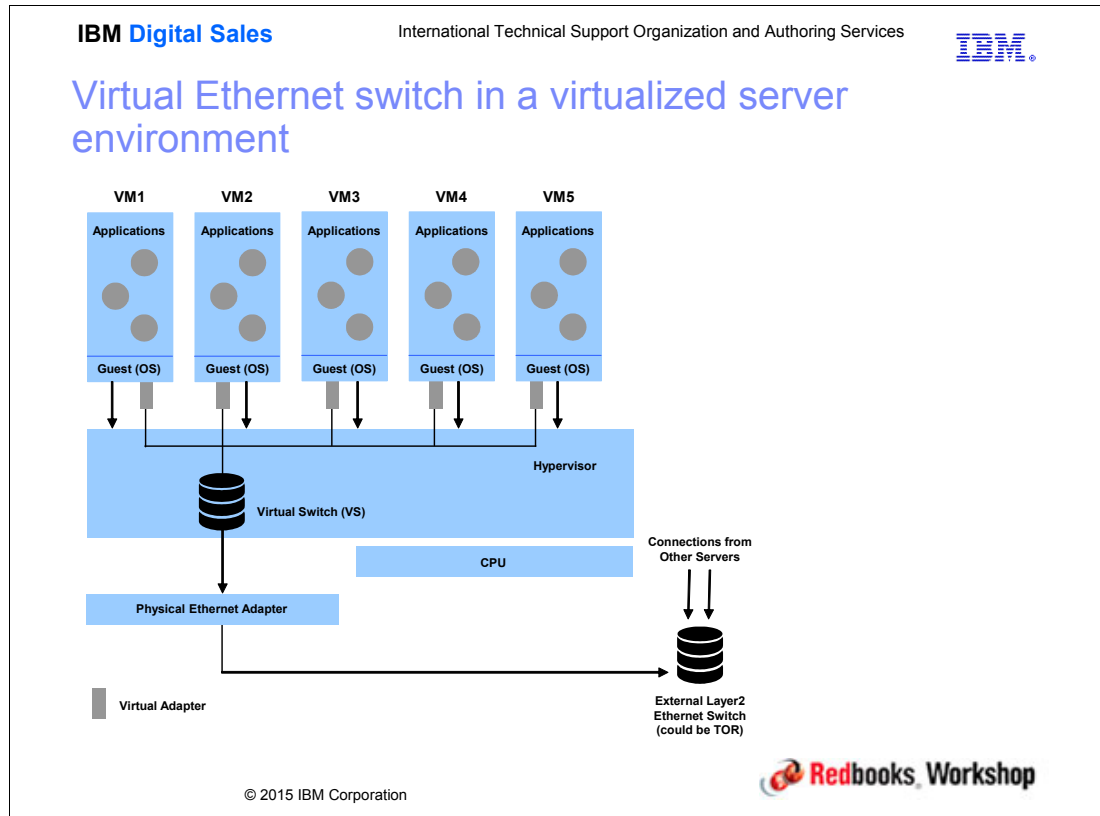


Figure 5-6 Virtual Ethernet switch in a virtualized server environment

Notes:

In an environment with physical servers, switches are used to connect servers to other servers. Firewalls and application-delivery controllers are other types of equipment that you can use in a data center for connection to external clients. With a virtualized environment, you can move some or all of these functions to inside a server.

You can use the Virtual Switch to switch between virtual machines (VMs) inside the same physical server and aggregate the traffic for connection to the external switch. The Virtual Switch is often implemented as a plug-in to the hypervisor. The VMs have virtual Ethernet adapters that connect to the Virtual Switch, which in turn connects to the physical Ethernet adapter on the server and to the external Ethernet switch. To the network manager, the virtual switch can appear as a part of the network. Unlike physical switches, the Virtual Switch does not necessarily have to run network protocols for its operation. It also does not need to treat all its ports the same because some of them are connected to virtual Ethernet ports. For example, it can avoid destination address learning on the ports that are connected to VMs. It can function through appropriate configuration from an external management entity.

5.10 Networking overview

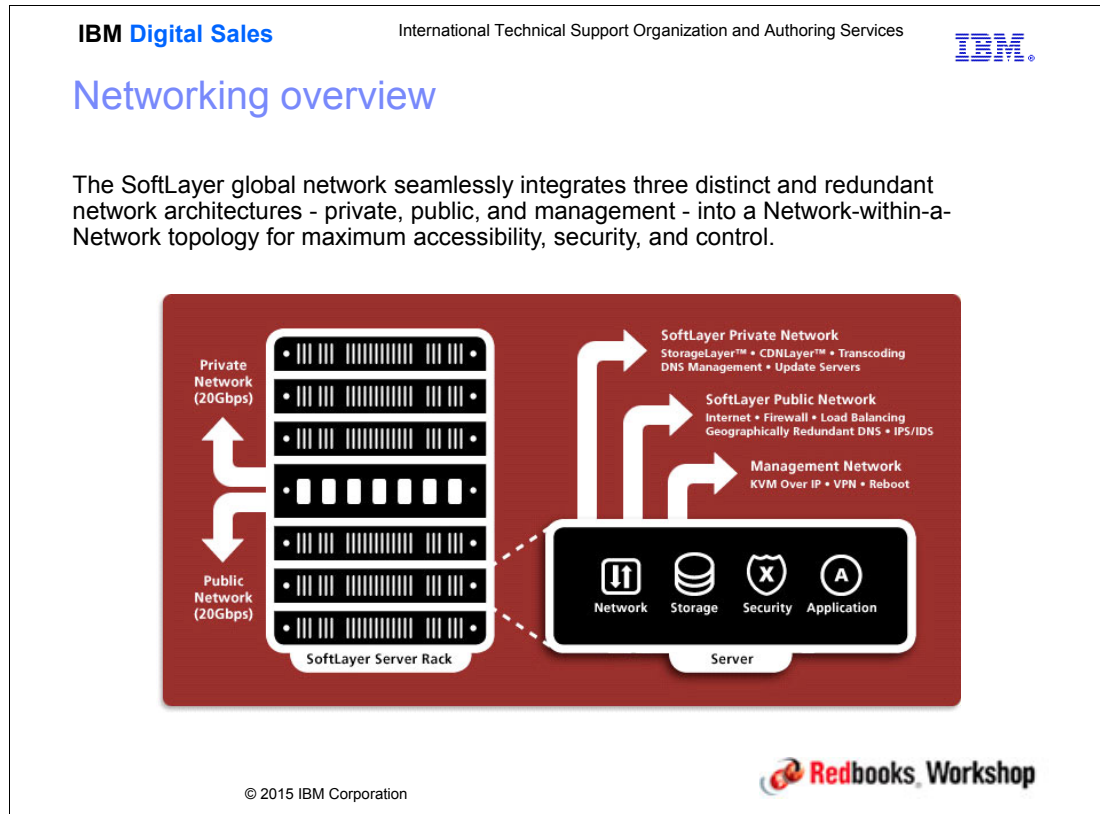


Figure 5-7 Networking overview

Notes:

SoftLayer has a worldwide footprint, with data centers currently in Amsterdam, London, Dallas, Houston, Atlanta, New York, Chicago, Denver, San Jose, Seattle, Los Angeles, Singapore, Washington D.C, Tokyo, Hong Kong S.A.R. of the PRC, Frankfurt, Paris, and Stockholm.

These centers are built with SoftLayer's unique PoD data center design concept. This allows them to provide functions that are independent with distinct and redundant resources and fully integrate all of their compute, storage, and services components in their network architecture. All of this together allows for seamless inter-data center capabilities with all these different services.

The backbone of these PoDs is the network rack and server design. This unit looks at how the rack is designed from a networking point of view, and then at the actual individual servers. This slide shows the overall network architecture from a rack point of view.

SoftLayer as a whole uses a three-network architecture. So when you're coming into the rack, you use the SoftLayer public and private networks, and also use the services that are built into the SoftLayer management network.

This unit focuses mainly on the public and private networks because that's where much of your data traffic is for your day-to-day customer interactions.

Each rack within a SoftLayer data center has at least 10 gigabits per second of redundant (20 gigabits per second total usage) bandwidth for both the private network and public network. This allows for a throughput of 10 gigabits per second throughout the SoftLayer network down to each rack.

The management network, although it is a separate network, uses the private network bandwidth for access to the systems.

5.11 Networking overview (2)

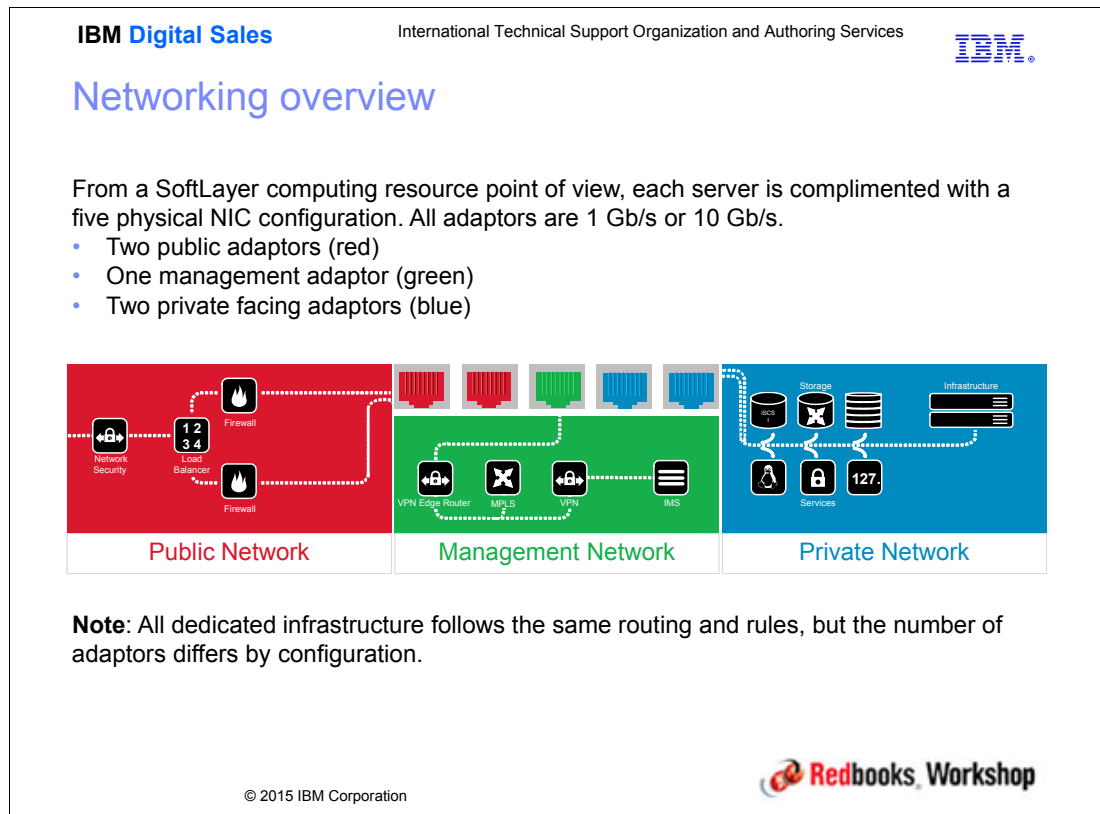


Figure 5-8 Networking overview (2)

Notes:

This image is an architectural representation of the backend of each server that represents the SoftLayer ecosystem. You have a redundant connection to the public network, a redundant connection to the private network, and one connection for the management network. Therefore, each server has a five-NIC setup: Two NICs to the public network, two NICs to the private network, and one to the management network.

SoftLayer is designed so that users do not have to worry about any of the networking components being locked into a switch that fails or anything along those lines. You have a redundant path going into each rack, and another redundant path going into each server.

While all SoftLayer dedicated infrastructure follows the same routing and rules, the number of adapters can differ by chosen configurations in certain compute resources.

5.12 Networking overview (3)

IBM Digital Sales

International Technical Support Organization and Authoring Services


IBM

Networking overview

The following outlines SoftLayer's SLAs for service and power for its networks:

- **Public network:** SoftLayer will use reasonable efforts to provide a service level of 100% for the public network.
- **Private network:** SoftLayer will use reasonable efforts to meet the service level of 100% for the private network.
- **Customer Portal:** SoftLayer will use reasonable efforts to meet the service level of 100% for access to the Customer Portal.
- **Redundant infrastructure:** SoftLayer will use reasonable efforts to meet the service level of 100% for access to the power and HVAC services provided to customers.

SoftLayer's geographically diverse PoPs provide seamless, direct, private, and high-speed access to the backbone network, bringing connectivity closer to the end user. You can choose the SoftLayer PoP location closest to your office or end users. High-speed metro-WAN services and cross connects from providers including Equinix and Telx are also available.



LEGEND: ● Multi-Region Data Centers ○ Single Region Data Centers ● Point of Presence

Connected by 20 Gbps Network

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-9 Networking overview (3)

Notes:

These are the SoftLayer SLAs for service and power for its networks.

As the network capabilities of SoftLayer are explored, you can see what SoftLayer has for SLAs for service and power for its networks. SoftLayer is committed to use reasonable efforts to provide a service level of 100% for the public network, private network, and Customer Portal. The public and the private network SLA are up to the machine levels.

SoftLayer also provides an SLA for redundant infrastructures to use reasonable efforts to meet a 100% access to the power and HVAC services provided to customers.

SoftLayer attempts to put each PoP into the most heavily used co-location site in the city in which the PoPs are deployed. This configuration allows for the easiest peering and transit connections with Telcos to expand the SoftLayer network. A co-location site is a data center where equipment, space, and bandwidth are available for rental to retail customers.

An example of a PoP use is if a customer resides in Germany and they are attempting to access a server in Amsterdam.

- ▶ Instead of going straight to Amsterdam, the customer would go directly to the Germany POP and use SoftLayer's 10 gigabits per second network to make the connection over the physically longer network distance, thus reducing latency.
- ▶ The time that clients use whenever they connect to a system on the SoftLayer network is reduced, which is the main reason for the POPs.

The following are the current POP locations:

- ▶ Amsterdam, Netherlands
- ▶ Atlanta, GA
- ▶ Chicago, IL
- ▶ Dallas, TX
- ▶ Denver, CO
- ▶ Frankfurt, Germany
- ▶ Hong Kong S.A.R. of the PRC
- ▶ Houston, TX
- ▶ London, England
- ▶ Los Angeles, CA
- ▶ Mexico City, Mexico
- ▶ Miami, FL
- ▶ New York, NY
- ▶ San Jose, CA
- ▶ Seattle, WA
- ▶ Singapore, Singapore
- ▶ Tokyo, Japan
- ▶ Washington, D.C.

5.13 Learning about public networks

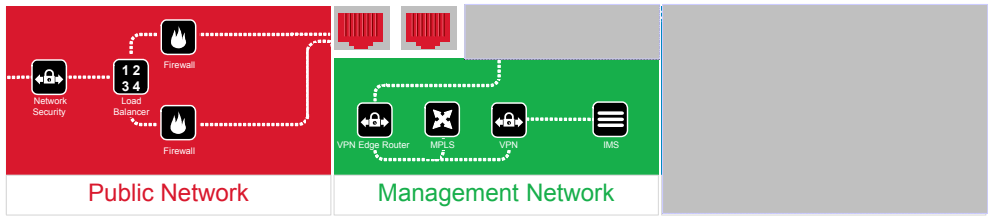

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Learning about public networks

- Public network definition
- Understanding primary network features



© 2015 IBM Corporation

Redbooks Workshop

Figure 5-10 Learning about public networks


Notes:

This section describes what the public network definition is, and how SoftLayer uses network carriers to connect out and expand the public network for customers. It also covers some of the features that are available on the public network.

5.14 Learning about public networks (2)

IBM Digital Sales

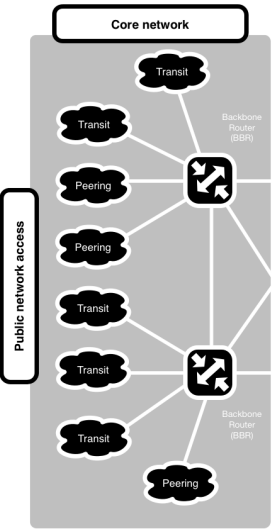
International Technical Support Organization and Authoring Services



Learning about public networks

The SoftLayer network provides all customers with over 2,000 Gbps of connectivity between the data centers and PoPs.

The core public network handles public traffic to hosted websites or online resources. SoftLayer's public network features multi-homed connectivity with bandwidth from independent carriers, combining more than 20 x 10 Gbps connections to create one of the industry's fastest networks.



© 2015 IBM Corporation


 **Redbooks Workshop**

Figure 5-11 Learning about public networks (2)

Notes:

As shown in the diagram, the core network handles public traffic to hosted websites or online resources.

SoftLayer uses multi-homed connectivity with bandwidth from independent peering and transit carriers, combining more the 20 x 10 Gbps connection to create one of the industry's fastest networks.

- **Peering:** When two or more autonomous networks interconnect directly with each other to exchange traffic. This is often done without charging for the interconnection or the traffic.
- **Transit:** When one autonomous network agrees to carry the traffic that flows between another autonomous network and all other networks. Because no network connects directly to all other networks, a network that provides transit will deliver some of the traffic indirectly through one or more other transit networks. A transit provider's routers announce to other networks that they can carry traffic to the network that has bought transit. The transit provider receives a "transit fee" for the service.

5.15 Learning about public networks - Carriers

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Learning about public networks - Carriers

Carriers and capabilities			
Level3 (transit)	Comcast (transit)	JPNAP (peering)	SIX (peering)
NTT America (transit)	AMS-IX (peering)	LINX (peering)	Starhub (peering)
Pacnet (transit)	Any2LAX (peering)	NL-IX (peering)	Telefonica (peering)
PCCW (transit)	DE-CIX (peering)	NOTA (peering)	TIE (peering)
Telstra (transit)	Equinix IX (peering)	NYIX (peering)	Time Warner (peering)
TeliaSonera (transit)	HKIX (peering)	RMI X (peering)	
<ul style="list-style-type: none">Juniper and Cisco 10 G networkCisco Guard DDoS protection		<ul style="list-style-type: none">Arbor Peakflow traffic analysisArbor TMS DDoS protection	

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-12 Learning about public networks - Carriers

Notes:

This is the current list of transit and peering carriers that SoftLayer uses to expand network across data centers and PoPs.

SoftLayer uses both Cisco and Juniper technology to drive the network, and Cisco Guard, Arbor Peakflow traffic analysis, and Arbor TMS for DDOS protection.

5.16 Understanding the primary features

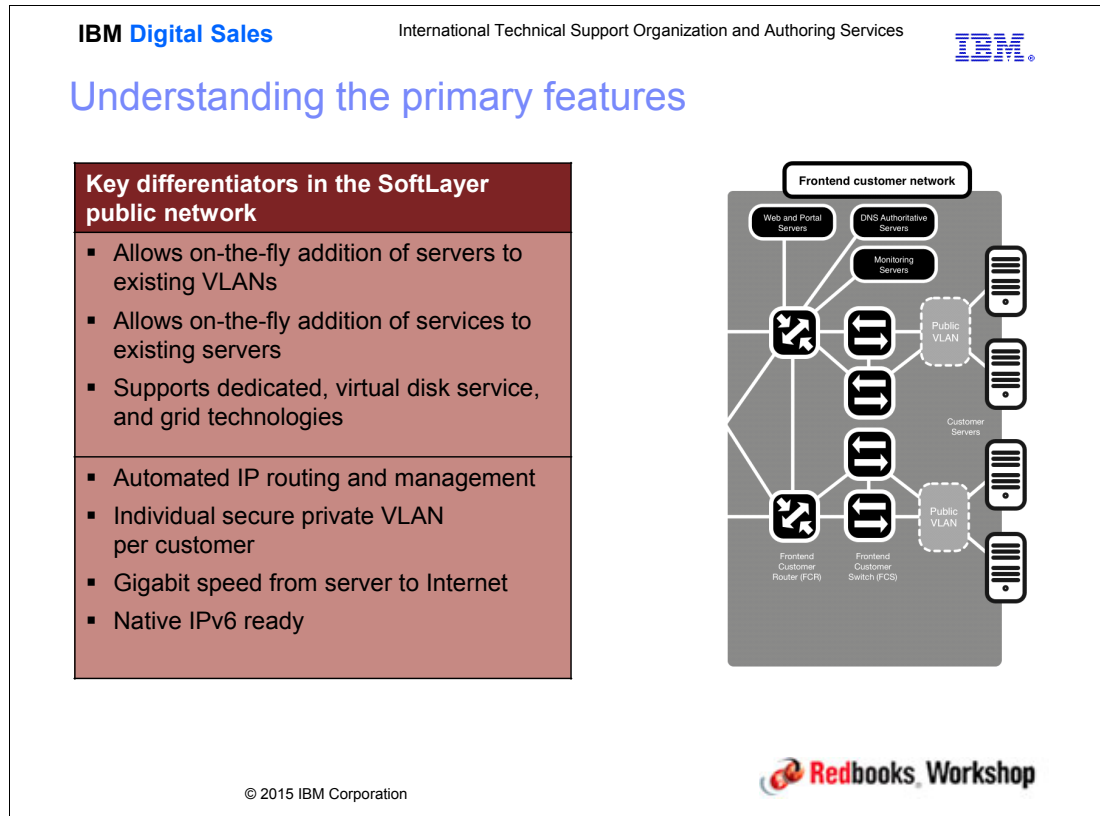


Figure 5-13 Understanding the primary features

Notes:

Continuing from the core network on the previous slide, next is the front end customer network.

- ▶ This section of the public network allows addition of servers to existing public VLANs.
- ▶ With full automated IP routing and management, each VLAN is secure in their customer environment through strict network access control lists (ACLs).
- ▶ Each server has fully gigabit capable speeds from server to the Internet.
- ▶ The entire SoftLayer network is fully IPv6 ready.
- ▶ IPv6 addresses can be requested when ordering a compute resource in the Customer Portal.

5.17 Understanding bandwidth

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Understanding bandwidth

- Unlimited inbound bandwidth
- Metered and unmetered outbound bandwidth available
- Optional bandwidth pooling

Outbound bandwidth options
Each virtual instance comes with 250 GB of outbound bandwidth per month
Each dedicated server comes with 500 GB of outbound bandwidth per month
You can pre-purchase an additional 20 TB of bandwidth
1 Gbps or 10 Gbps unmetered inbound public bandwidth
Overages are metered

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-14 Understanding bandwidth

Notes:

Public bandwidth in the SoftLayer network is handled through unlimited inbound bandwidth and metered and purchased unmetered outbound bandwidth. Therefore, any bandwidth you have coming into the data centers from a public network is completely unlimited. You can use it as much as you want. You can use it at 1 gigabit. If you're using specific dedicated servers, you can even use it as a 10-gigabit connection.

SoftLayer also allows for the option to purchase bandwidth pooling for a small monthly fee. Bandwidth pooling allows for a customer's account to take all of the credited outgoing bandwidth from virtual instance or dedicated server purchases and pool it together for use by any system that is using outgoing bandwidth.

An example scenario of this is if you order three dedicated servers for backend services (databases or application servers) that you never want on the public network, and two virtual instances that will act as your web servers handling all outgoing traffic. If you pool your bandwidth, you can disable the public network ports on the three dedicated servers, and use the 2 terabytes (500 gigabytes x three dedicated + 250 gigabytes x two virtual) of credited outbound traffic for the two web servers.


Pooling can be enabled by opening a ticket in the SoftLayer portal and requesting bandwidth pooling be turned on. You can get more information and pricing details at:

<http://www.softlayer.com/info/pricing>

5.18 Learning about private networks

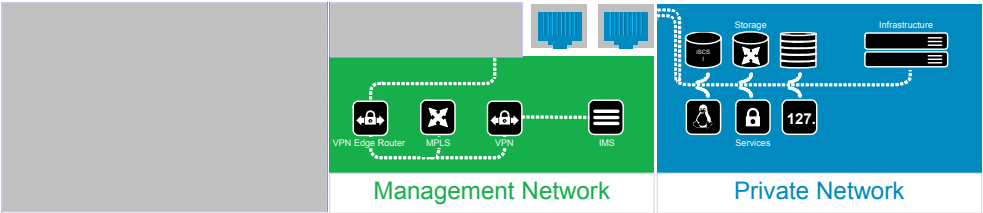
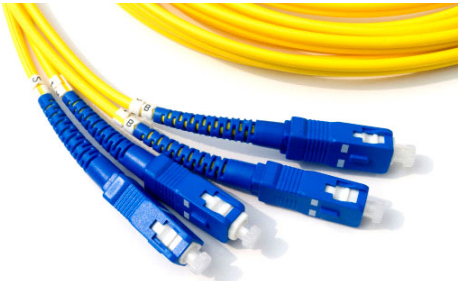
IBM Digital Sales

International Technical Support Organization and Authoring Services



Learning about private networks

- Private network definition
- SoftLayer network carriers and features



© 2015 IBM Corporation




Figure 5-15 Learning about private networks

Notes:

Customers interact with the public network. This section looks at the options available to you as you go through the private network. This is also one of the more powerful options that SoftLayer has.

5.19 Learning about private networks (2)

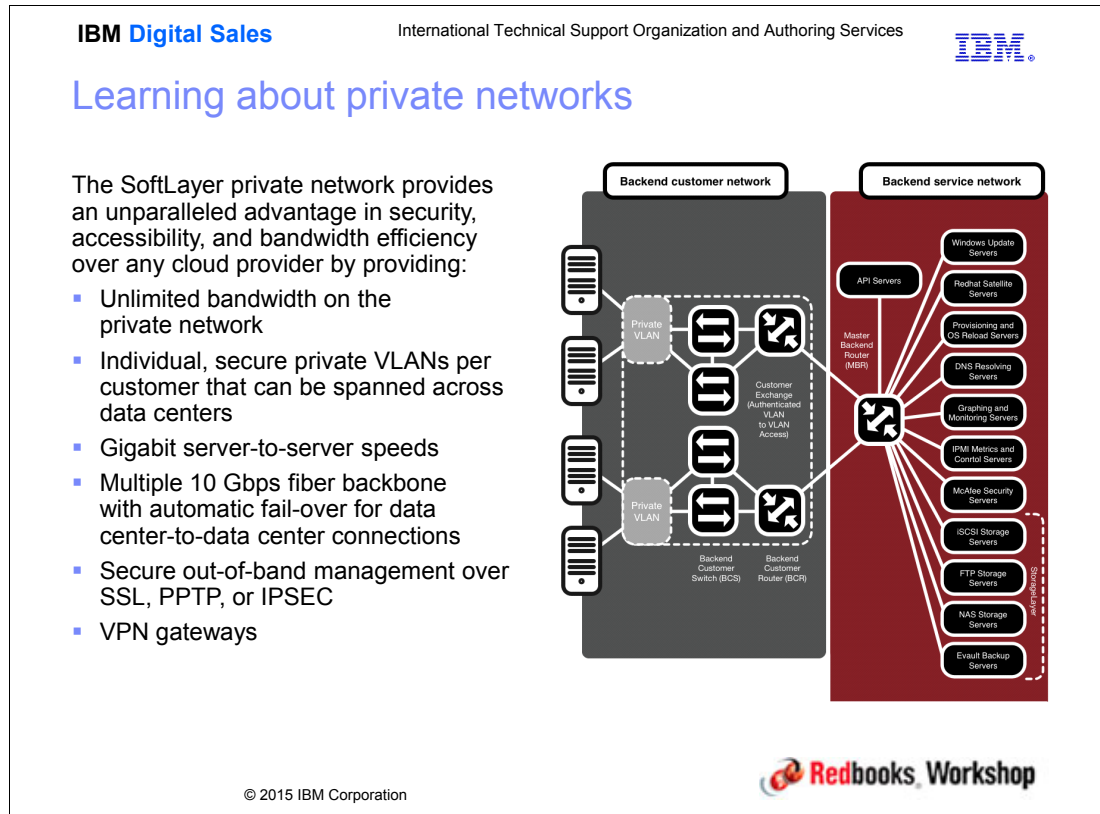


Figure 5-16 Learning about private networks (2)

Notes:

SoftLayer provides you with a private network and then services that are connected to that private network so that you can expand the capabilities that you currently have with SoftLayer offerings. That capability allows you to take any server that you deploy into SoftLayer and, using the span VLANs, connect to any of the other servers you have in data centers inside of SoftLayer.

For example, if you have servers in Washington, DC, servers in Dallas, and servers in San Jose data centers, you can use the private network to be able to move all your data between those servers in each data center at no extra charge with unlimited bandwidth. The fully 10-gigabit network goes across all these data centers and allows you to sync across all of them. This is also completely private, so you're not exposing any of your traffic out to the public network

You can use the private network as secure transit for DR centers, to back up for data on another set of servers, or use it just to sync servers for a large-scale rollout across the world

As you can see from the diagram, each of these private networks and each of these VLANs that you have deployed are redundant because they are backed into the backend routers that connected to each data center. This configuration creates a complete failover system that allows you to keep going even if a switch fails

5.20 Learning about private networks (3)

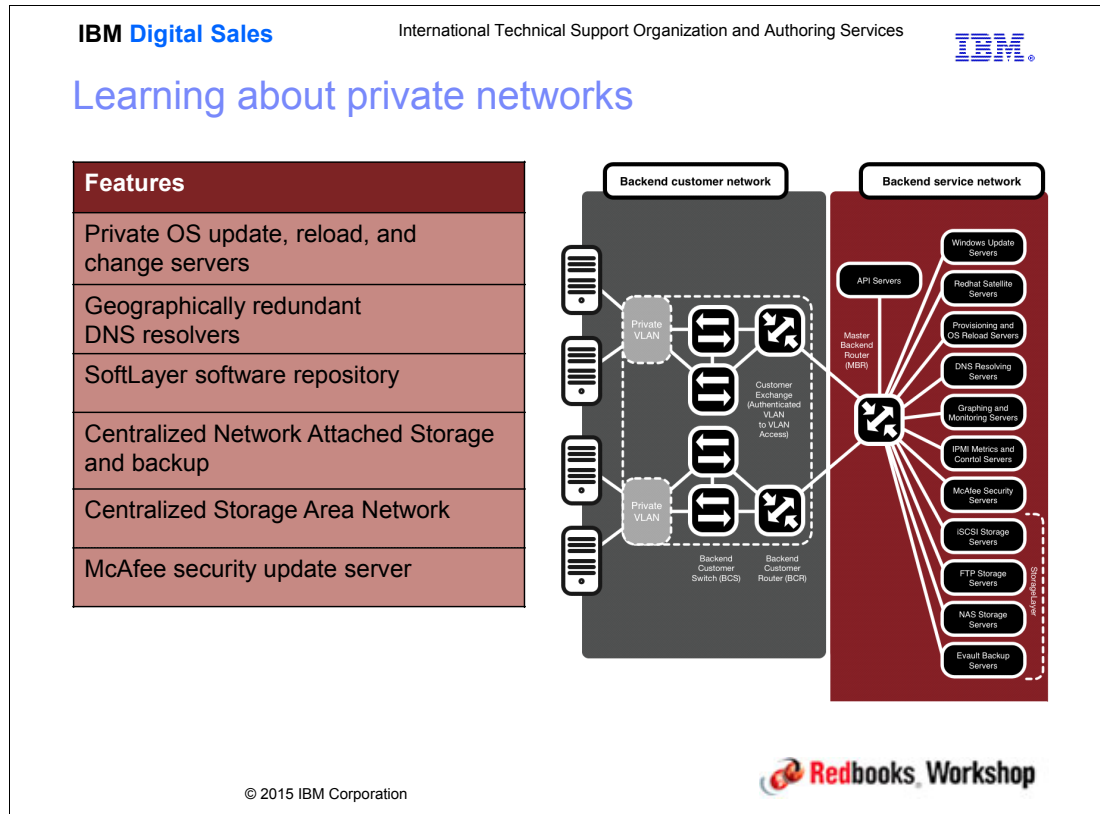


Figure 5-17 Learning about private networks (3)

Notes:

SoftLayer provides you with Windows or Red Hat updates, depending on your operating system. Updates for other Linux instances can be done through repositories that SoftLayer sets up on the private network. You do not have to expose any of your resources to be able to update them, or work out a gateway to update those machines. You can use the SoftLayer backend services to keep them at the highest level of security patches that is available.

Use redundant DNS resolvers if you are using full DNS resolution and you still need machines that have to resolve host names on the private network. The resolvers allow you to resolve any host names you have across your systems.

Centralized network attached storage and backup provide a centralized storage area network across all of your private network. This means that you are not locked to your local region. If, for example, you have an iSCSI in Dallas and you have a system in Washington, DC, you can temporarily connect them to be able to move data.

For Windows operating systems, MacAfee antivirus is available in the security update server, which is also available on the private network. If you have private systems that you want to be able to keep certain secure levels, antivirus, you can keep them on the private network and still be able to manage them securely through private network access.

5.21 Learning about management network

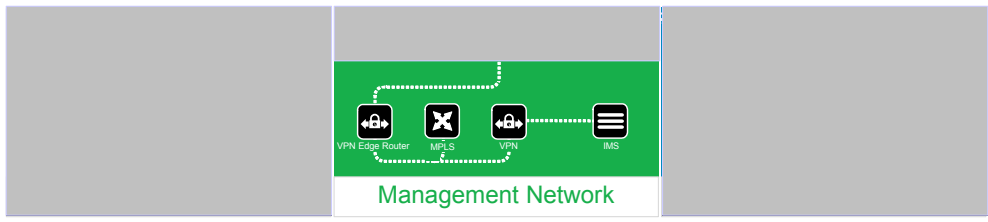

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Learning about the management network

- Out-of-Band Management network definition
- Out-of-Band Management network features



© 2015 IBM Corporation

Redbooks Workshop

Figure 5-18 Learning about management network

Notes:

When it comes to managing your server, you want an unencumbered network connection that will give you direct, secure access when you need it. Splitting out the public and private networks into distinct physical layers provides significant flexibility when it comes to delivering content. However, SoftLayer saw a need for one more unique network layer. If your server is targeted for a denial of service attack or a particular ISP fails to route traffic to your server correctly, you are effectively locked out of your server if you do not have another way to access it. The SoftLayer management-specific network layer uses bandwidth providers that are not included in the public/private bandwidth mix, so you access the server through a dedicated port.

5.22 SoftLayer network architecture

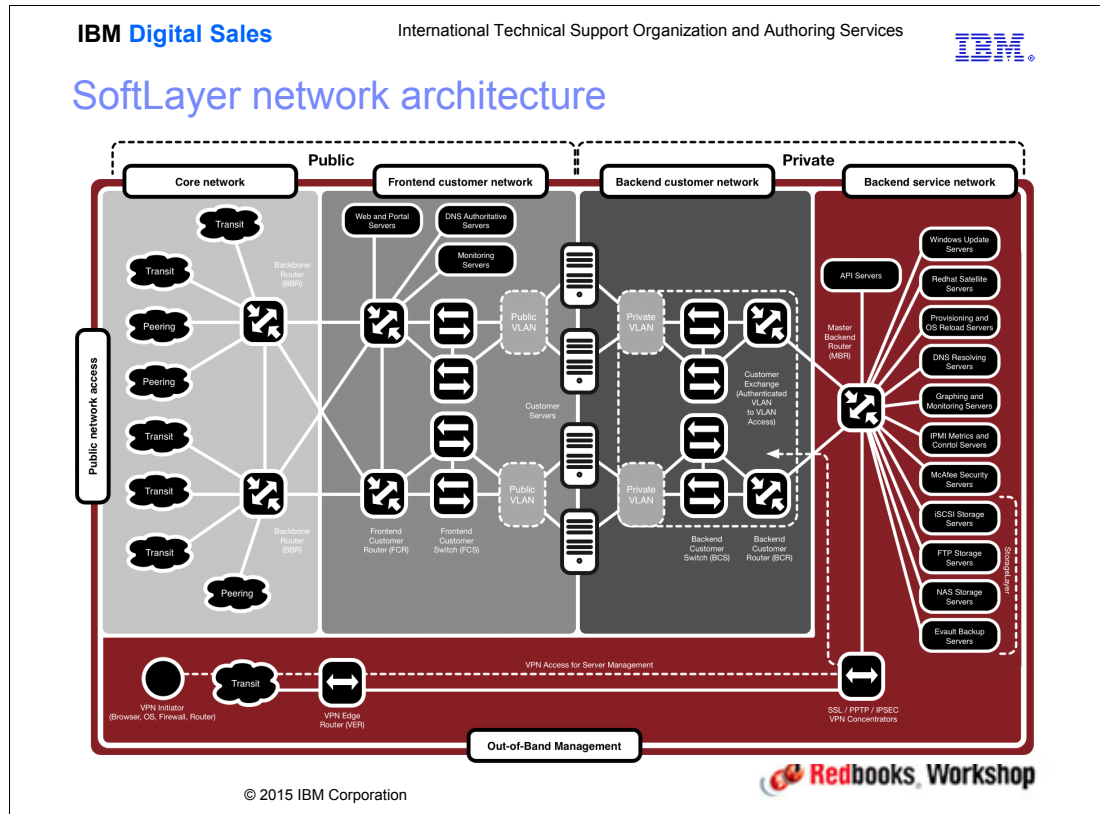


Figure 5-19 SoftLayer network architecture

Notes:

This diagram shows you how the networks and secure connections work together in the the overall SoftLayer network architecture:

- ▶ You can see how you come in from the public network:
 - You come into the core network from different transit and peering connections.
 - The data comes into the public network for SoftLayer and connects into the public VLANs that are available for your servers.
- ▶ On the other side, those servers are also connected to the private network through redundant connections:
 - Those link back to the backend services on SoftLayer.
 - Note also the out-of-bound management.

5.23 Reviewing SoftLayer network tooling



Figure 5-20 Reviewing SoftLayer network tooling

Notes:

You do not have to have an account to use Looking Glass it is available publicly on the Internet.

Go through what's available on this tool and some of the other tools on SoftLayer.

5.24 Using Looking Glass, SoftLayer's IP backbone

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using Looking Glass, SoftLayer's IP backbone

- SoftLayer Looking Glass is an interactive resource that can be used to test latency between SoftLayer data centers, and between a router location and target address, subnet, or hostname.
- Latency tests between a location and its target are conducted using various methods, including the ping and traceroute commands. Users can also test network latency through various test downloads based on size and data center location. Users operate Looking Glass under the premise that only manual queries will be made.
- Use these steps to run a traceroute command on a website:
 1. Enter <http://lg.softlayer.com/> in your browser window.
 2. Read the terms and conditions and click the **I agree** link. Note that clicking the link signifies that you understand that commands and IPs might be logged.

© 2015 IBM Corporation


Redbooks Workshop

Figure 5-21 Using Looking Glass, SoftLayer's IP backbone

Notes:

This slide also begins an activity for students to conduct a latency test using a **traceroute** command.

5.25 Using SoftLayer Looking Glass, SoftLayer's IP backbone

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using SoftLayer Looking Glass, SoftLayer's IP backbone

- Under **Command**, select the **Traceroute** radio button.
- Select your data center using the **Select Location and Router** menu.
- Enter a website in the **Target Address/Subnet or Hostname** field and click **Execute**.

SoftLayer IP Backbone :: Looking Glass

Command:
☐ Ping
☒ Traceroute
☐ BGP Table Lookup
☐ BGP Summary

Select Location and Router
Dallas, TX - Equinix [bbr01.eq01.dal03]
Target Address/Subnet or Hostname
www.ibm.com
Execute

Color Scheme:
White / Black

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-22 Using SoftLayer Looking Glass, SoftLayer's IP backbone

Notes:

Enter the website that you want, and select the locations and the routers in those locations that you want to perform this test against.

You can also select a color scheme.

5.26 Using SoftLayer Looking Glass, SoftLayer's IP backbone (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using the SoftLayer Looking Glass, SoftLayer's IP backbone

6. A Looking Glass Results window displays your traceroute results.

Looking Glass Results

```
bbr01.eq01.dal03> traceroute 129.42.58.216
```

HOST: bbr01.eq01.dal03-re0	Loss%	Snt	Last	Avg	Best	Worst	StDev
1. 213.248.102.173	0.0%	5	0.4	6.5	0.4	28.6	12.4
2. 192.205.37.49	0.0%	5	2.2	4.0	1.7	10.0	3.4
3. 12.122.138.18	0.0%	5	26.3	26.5	25.4	28.1	1.1
4. 12.122.3.222	0.0%	5	27.5	28.9	26.7	32.5	2.2
5. 12.123.140.174	0.0%	5	30.0	32.8	28.2	46.6	7.7
6. 12.122.142.29	0.0%	5	29.5	28.3	25.5	34.2	3.7
7. 12.125.74.74	0.0%	5	27.8	28.1	27.7	29.4	0.7
8. ???	100.0	5	0.0	0.0	0.0	0.0	0.0

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-23 Using SoftLayer Looking Glass, SoftLayer's IP backbone (2)

Notes:

This example shows the results generated by the quick trace route that was entered in the previous slide.

5.27 Using other networking tools

Using other networking tools

You can also test network latency, run test downloads, and perform speed tests.

The **Network Latency** table lets you see the current latency between data centers and PoPs. Hover over a square in the table to see the latency between locations.

SoftLayer IP Backbone - Network Latency

	ATL	CHI	DAL	DEN	HOU	LAX	MIA	NYC	SEA	SJC	WDC	AMS	FRA	LON	HKG	SNG	TOK
ATL	0	31	21	35	30	55	13	20	61	59	14	91	99	88	209	229	167
CHI	31	0	22	26	30	50	45	23	52	52	19	95	106	90	205	244	161
DAL	21	22	0	15	6	31	30	46	42	39	34	114	120	112	194	212	145
DEN	35	26	15	0	24	34	45	46	27	27	44	120	128	116	180	220	135
HOU	30	30	6	24	0	38	23	52	51	43	44	122	126	116	193	214	149
LAX	55	50	31	34	38	0	60	72	26	8	61	139	146	147	160	178	116
MIA	13	45	30	45	23	60	0	34	73	65	26	105	112	100	219	234	173
NYC	20	23	46	46	52	72	34	0	74	74	7	74	78	68	227	267	182
SEA	61	52	42	27	51	26	73	74	0	19	70	147	155	142	171	211	127
SJC	59	52	39	27	43	8	65	74	19	0	70	148	154	142	153	194	109
WDC	14	19	34	44	44	61	26	8	70	70	0	79	86	75	224	263	179
AMS	91	95	114	120	122	139	105	74	146	148	79	0	6	6	301	340	256
FRA	99	106	120	128	126	146	112	78	155	154	86	6	0	13	307	349	263
LON	88	90	112	115	116	147	100	68	142	142	75	6	13	0	295	334	251
HKG	209	205	194	180	193	160	219	227	171	153	224	301	307	295	0	31	44
SNG	229	244	212	220	214	178	234	267	211	194	263	340	349	334	31	0	85
TOK	167	161	145	135	149	116	173	182	127	109	179	256	263	251	44	85	0

Figure 5-24 Using other networking tools

Notes:

Network latency can be tested by running test downloads and performing speed tests. Hover over the square that you are interested in to see the amount of latency.

You can quickly call out two sites and see if the students are able to answer what the latency is between the two sites.

SoftLayer tries to be as transparent as possible when it comes to the speed of network connections. This helps you properly plan what data centers to use and where you need to put your data.

5.28 Using other networking tools (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM®

Using other networking tools

Test Downloads allows you to test the throughput rate for different file sizes from SoftLayer data centers. Click the test that you want to run and download the .zip file.

SoftLayer IP Backbone - Test Downloads

Dallas, TX	Seattle, WA	Washington, DC	Houston, TX	San Jose, CA
<ul style="list-style-type: none">10MB Test100MB Test500MB Test	<ul style="list-style-type: none">10MB Test100MB Test500MB Test	<ul style="list-style-type: none">10MB Test100MB Test500MB Test	<ul style="list-style-type: none">10MB Test100MB Test500MB Test	<ul style="list-style-type: none">10MB Test100MB Test500MB Test
Singapore, Singapore <ul style="list-style-type: none">10MB Test100MB Test500MB Test				
Amsterdam, NL <ul style="list-style-type: none">10MB Test100MB Test500MB Test				

© 2015 IBM Corporation




Figure 5-25 Using other networking tools (2)

Notes:

SoftLayer allows you to download different size test files from all of the data centers so you can test your throughput. You can use 10-, 100-, or 500-megabyte test files.

5.29 Using other networking tools (3)

IBM Digital Sales

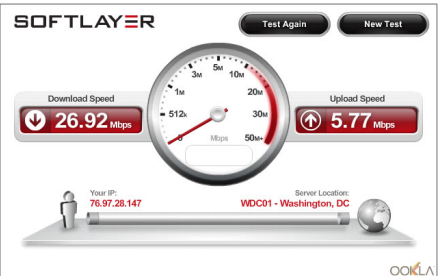
International Technical Support Organization and Authoring Services

IBM

Using other networking tools

You can test the network speed between your workstation and a data center. Use these steps to perform a speed test:

1. Click the **Launch Test Site** link under **SoftLayer IP Backbone – Speed Test Site**.
2. Click the data center icon that you want to test. You can hover over the icon to see the name of the data center.
3. The download and upload speed between your workstation and selected workstation are tested and the results displayed.



SOFTLAYER Test Again New Test

Download Speed 26.92 Mbps

Upload Speed 5.77 Mbps

Your IP: 76.97.28.147 Server Location: WDC01 - Washington, DC

Last Result:
Download Speed: 26922 kbps (3365.3 KB/sec transfer rate)
Upload Speed: 5766 kbps (720.8 KB/sec transfer rate)
Monday, January 13, 2014 3:26:52 PM

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-26 Using other networking tools (3)

Notes:

(none)

5.30 Using other networking tools from SoftLayer Control Panel

IBM Digital Sales International Technical Support Organization and Authoring Services IBM

Using other networking tools from the SoftLayer Control Panel

Network Tools

Ping
Sends out an echo request to a specific server on the network. It can be used to check communication links or to check whether the specific server exists or if it is running.

Traceroute
Used to determine the path a packet is traveling across the Internet. It looks up each machine along the path to a destination host and displays the corresponding name and IP address for that hop. With each name is the number of milliseconds to get to the destination.

NSLookup
Used to resolve a hostname to an IP address or vice versa.

Whois
Used to locate registered information on a domain name such as registrant contact information, nameservers, and domain status.

Check DNS
Used to check if DNS entries have propagated to DNS servers.

```
[www.softlayer.com]
Any DNS changes made to softlayer.com at this time can take up to
0 days to propagate across the Internet.

If the last update to this domain occurred before the following
date/time, the update should be propagated:
Tue, 14 Jul 2015 09:24:37 -0500

Authoritative Nameservers:
ns2.softlayer.net (67.228.255.25)
ns2.softlayer.net (2607:f0d0:0:15:2::1)
ns1.networklayer.com (67.228.254.22)
ns1.networklayer.com (2607:f0d0:0:15:10::1)
ns1.softlayer.net (67.228.254.24)
ns1.softlayer.net (2607:f0d0:0:15:1::1)
ns2.networklayer.com (67.228.255.22)
ns2.networklayer.com (2607:f0d0:0:15:20::1)
```

© 2015 IBM Corporation Redbooks Workshop

Figure 5-27 Using other networking tools from SoftLayer Control Panel

Notes:

A customer with a valid SoftLayer account can access the following networking tools for additional debugging purposes:

- ▶ The Ping tool requests an echo ICMP from the selected server and is used to check communication links and to ensure the specified server is active. Ping requests can be sent to both SoftLayer and external devices
- ▶ The Traceroute tool determines the path that a packet of information is traveling across the Internet by mapping the path to a destination. The results returned include the corresponding name and IP address for each hop, and the number of milliseconds the packet takes to get to the destination.
- ▶ The NSLookup tool resolves hostnames to IP addresses and vice versa, and can be performed on any hostname or IP address. For queries on a hostname, all associated IP addresses will be returned. For queries on an IP, the corresponding hostname will be returned.
- ▶ The Check DNS tool allows users to check the last time that entries have been propagated to DNS servers, as well as the standard propagation time for the selected domain and its authoritative name servers.

5.31 Managing VPN connections to SoftLayer



Figure 5-28 Managing VPN connections to SoftLayer

Notes:

(none)

5.32 General concept of VPN

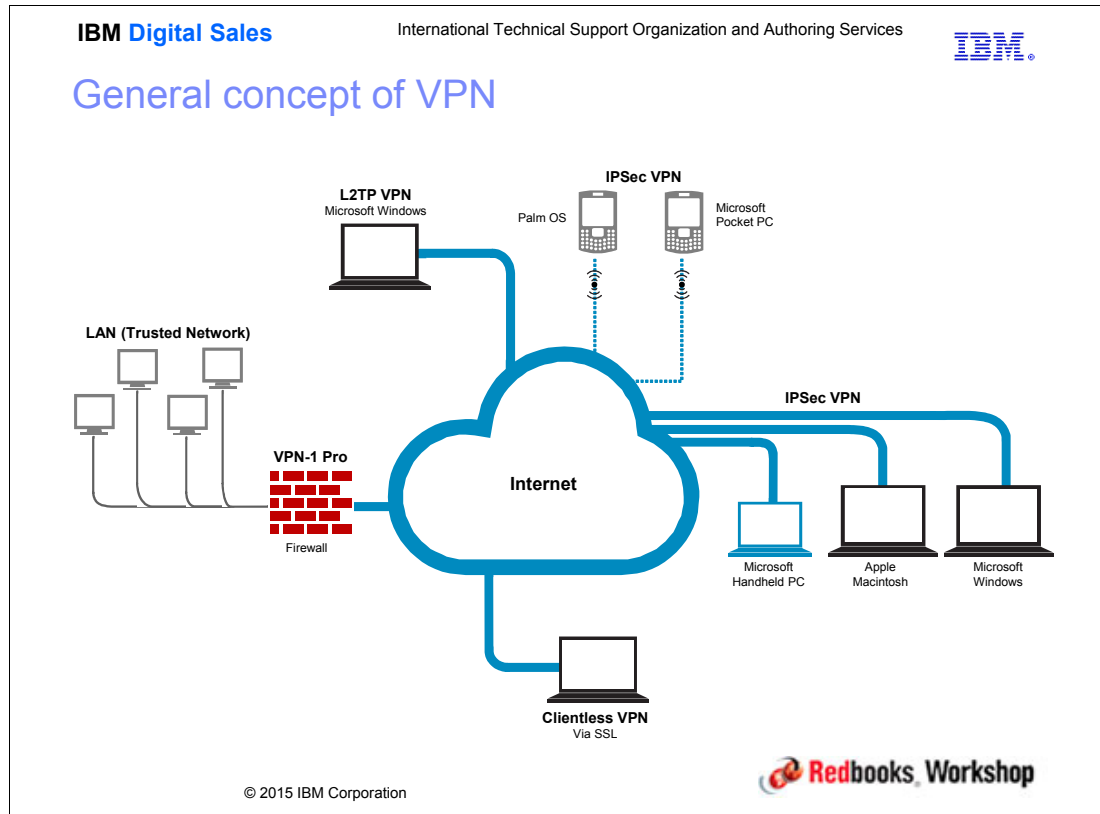


Figure 5-29 General concept of VPN

Notes:

Virtual private networks (VPNs) allow users to securely access a private network and share data remotely through public networks. Much like a firewall protects your data on your computer, VPNs protect it online. And while a VPN is technically a wide area network (WAN), the front end retains the same functionality, security, and appearance as it would on the private network.

For this reason, VPNs are hugely popular with corporations as a means of securing sensitive data when connecting to remote data centers. These networks are also becoming increasingly common among individual users. Because VPNs use a combination of dedicated connections and encryption protocols to generate virtual point to point connections, even if snoopers did manage to siphon off some of the transmitted data, they would be unable to access it because of the encryption. In addition, VPNs allow individuals to spoof their physical location (the user's actual IP address is replaced by the VPN provider), allowing them to bypass content filters.

5.33 Managing VPN connections to SoftLayer

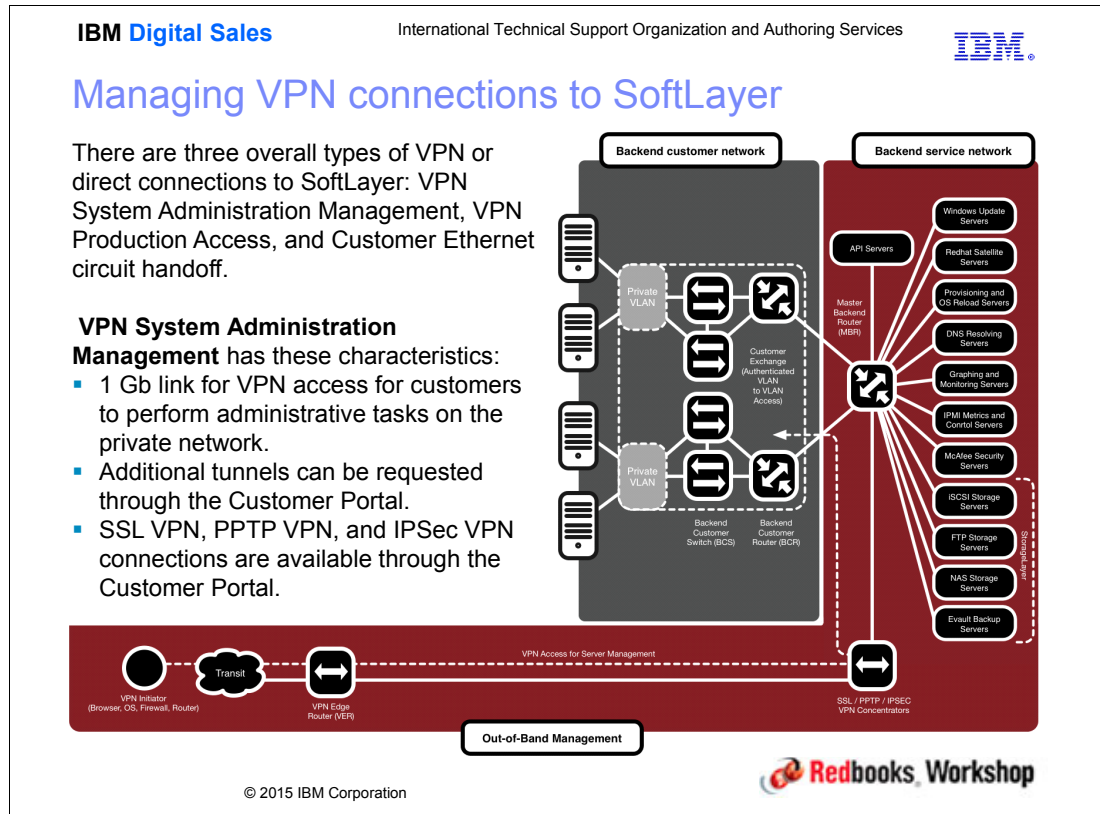


Figure 5-30 Managing VPN connections to SoftLayer

Notes:

This next section talks about connecting securely into SoftLayer. There are three overall types of VPN or direct connections into SoftLayer. The SoftLayer VPN offering for System Administration management uses the Management network to connect.

In the diagram, note the VPN connections coming in across the bottom and linking back into the customer's private network. This is defaulted to a 1 gigabit link that is available to connect through SSL VPN, PPTP VPN, and IPSec VPN.

Additional tunnels can be requested as user access or site-to-site IPsec VPN access. However, these additional tunnels are still only for system Admin use.

5.34 Managing VPN connections to SoftLayer (2)

IBM Digital Sales

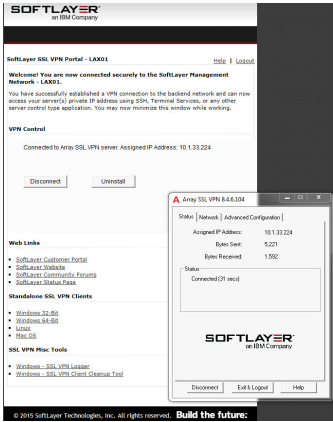
International Technical Support Organization and Authoring Services

IBM

Managing VPN connections to SoftLayer

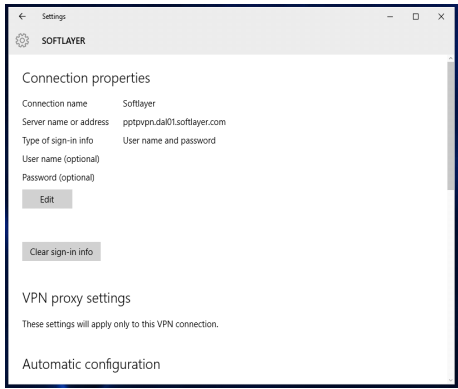
SSL VPN

- No special VPN client software (uses web browser plug-ins)



PPTP VPN

- Widely supported VPN solution (the only VPN client embedded in Windows)



© 2015 IBM Corporation

Redbooks Workshop

Figure 5-31 Managing VPN connections to SoftLayer (2)

Notes:

The **Secure Sockets Layer (SSL) VPN** technology has been growing in popularity. A big advantage of SSL VPNs is that you do not need special VPN client software on the VPN clients because they use the Web browser as the client application. Thus, SSL VPNs are known as "clientless" solutions. This also means that the protocols that can be handled by an SSL VPN are more limited. However, this can also be a security advantage. With SSL VPNs, instead of giving VPN clients access to the whole network or subnet as with IPsec, you can restrict them to specific applications. A disadvantage of this is that to use such plug-ins, the client's browser settings will have to be opened up to allow active content. This configuration exposes the browser to malicious applets unless you set it to block unsigned active content and ensure that the plug-ins are digitally signed.

- In SoftLayer, you need to access `vpn.softlayer.com` in order to perform the VPN install and configuration.

The **Point-to-Point Tunneling Protocol (PPTP)**, developed by Microsoft in conjunction with other technology companies, is the most widely supported VPN method among Windows clients, and is the only VPN protocol built into the Windows operating systems. PPTP is an extension of the Internet standard Point-to-Point Protocol (PPP), the link layer protocol used to transmit IP packets over serial links. PPTP uses the same types of authentication as PPP (PAP, SPAP, CHAP, MS-CHAP, and EAP). PPTP establishes the tunnel, but does not provide encryption. It is used with the Microsoft Point-to-Point Encryption (MPPE) protocol to create a secure VPN. PPTP has relatively low overhead, making it faster than some other VPN methods.

Unit 5. Networking options in cloud computing 143

- ▶ Currently supported operating systems are Microsoft Windows, Ubuntu, MacOSm, and even iOS devices.
- ▶ PPTP has been criticized in the past for various security flaws. Many of these problems have been addressed in current versions of the protocol. Using EAP authentication greatly enhances the security of PPTP VPNs. One advantage of using PPTP is that there is no requirement for a Public Key Infrastructure. However, EAP does use digital certificates for mutual authentication (both client and server) and highest security.
- ▶ For more information about configuring this type of connection in SoftLayer, see:
<http://knowledgelayer.softlayer.com/procedure/set-pptp-vpn-windows-7>

5.35 Managing VPN connections to SoftLayer (3)

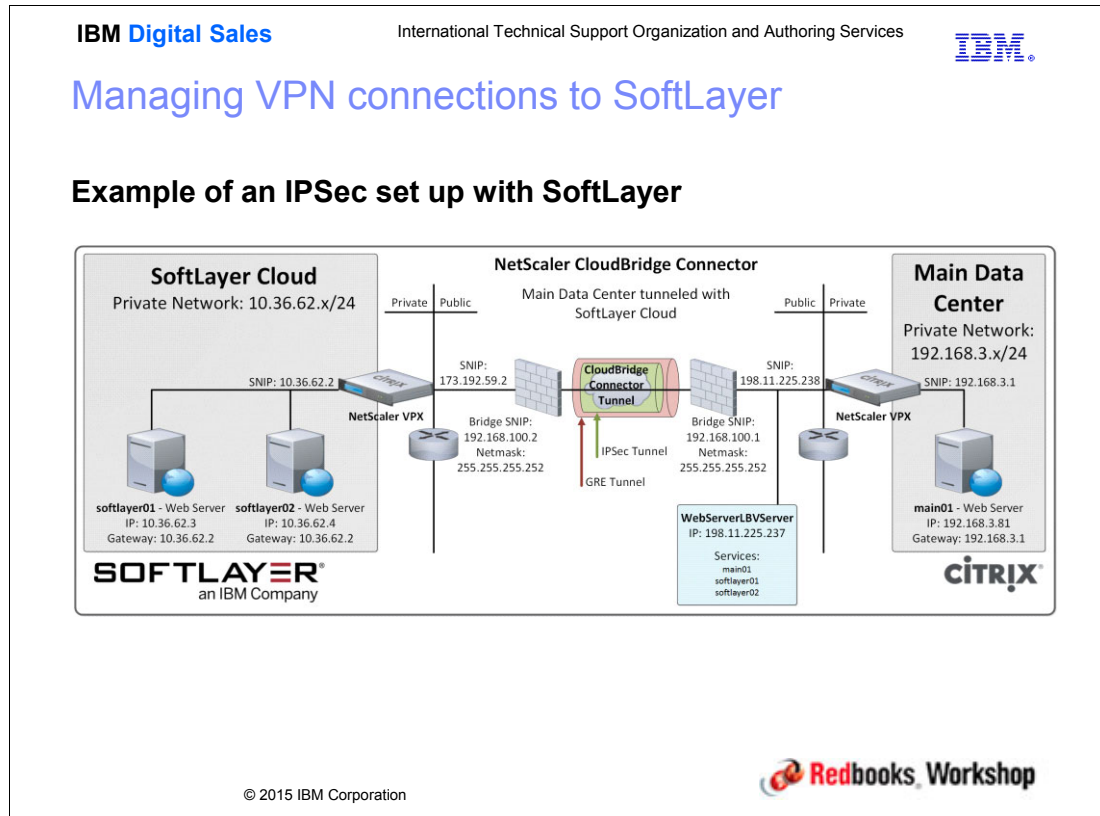


Figure 5-32 Managing VPN connections to SoftLayer (3)

Notes:

Note the following about the IPsec management VPN:

- ▶ There is no inherent redundancy built into this solution.
- ▶ There is just one Internet link to each VPN device in a given city.
- ▶ Unlike the SSL VPN or the PPTP VPN (which are more dynamic), the connection to the SoftLayer IPsec devices are specific to a city and require customer configurations on the user side.
- ▶ The customer can purchase an additional IPsec VPN in a different city if they would like to have a backup connection available.
- ▶ If the customer is using static NATs, they will have to change some configurations on their SoftLayer servers from one set of IPs to another as they move from one IPsec device to another.
- ▶ Reverse static NATs are specific to each city.
- ▶ An I2tp IPsec VPN is used, which requires NAT'd addresses for the server hosted at SoftLayer to initiate a connection with a computer on the other end of the tunnel.

5.36 Managing VPN connections to SoftLayer (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Managing VPN connections to SoftLayer

VPN Production Access

- The recommended solutions for any customer-required production VPN access to the SoftLayer network is either FortiGate security appliance or Vyatta gateway appliance.
- More information on these services and appliances is available in the managed or semi-managed firewall section of the workshop.

© 2015 IBM Corporation


Redbooks Workshop

Figure 5-33 Managing VPN connections to SoftLayer (4)


Notes:

(none)


5.37 Direct Link use case

IBM Digital Sales

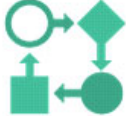
International Technical Support Organization and Authoring Services




Direct Link use case




Fully integrated hybrid environment
Blur the lines between on-prem and off-prem resources. With a high speed, direct network link between your data centers and SoftLayer, you can move data between servers (literally) at the speed of light.



Frequent (and massive) backup and storage
Back up or store huge volumes of data from your corporate data centers on your SoftLayer infrastructure (or the other way around). And because you have a dedicated network connection, your transfer rates will be fast, consistent, and reliable.



Super-secure data
Customers moving sensitive financial, health, or government-regulated data to and from the cloud platform can further ensure its security by completely avoiding exposure to the public internet at all times.



© 2015 IBM Corporation

Figure 5-34 Direct Link use case

Notes:

The third option for a secure connection is a customer Ethernet circuit handoff or a Direct Link.

By using this option, you can plug into SoftLayer's private network with a Direct Link in any of SoftLayer's network points of presence (PoPs), and enjoy fast and secure network connection to and from your servers in any of the SoftLayer data centers around the world. With this connection, you will have unfettered access to your servers on the SoftLayer platform. And because you are connected to SoftLayer's private network, all traffic across your Direct Link and between your servers in all Softlayer data centers is free and unmetered.

Point of presence allows the customer to come into SoftLayer. And going to their telco they can negotiate a rate with their own telco and have them come to a SoftLayer point of presence and actually bring a GBIC (gigabit interface converter), a physical connection for a network, and physically bridge their network to the SoftLayer private network.

This option is available in any point of presence and also any data center location. All data centers have a point of presence location in that city.

You can also do that for private networks. If you wanted to make sure all your data is encrypted as it is moving through the SoftLayer network, you can actually set up VPN points on the private side of things and encrypt your own data so that you can add another level of security to your solution.

Using the Direct Link has the following advantages:

- ▶ Guaranteed latency/bandwidth
- ▶ Move data to and from your on-premises data center across 1Gbps or 10Gbps network connections.
- ▶ Protect your sensitive, business-critical data by controlling every hop of its network path.
- ▶ Enjoy direct, free access to your servers in every SoftLayer data center across the private network.


The following are popular uses cases:

- ▶ Fully integrated hybrid environment
- ▶ Frequent backup and storage
- ▶ Extremely secure data transfers

5.38 Direct Link use case (continued)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Direct Link use case (continued)

Select POP Location for End Point:
Customer is responsible for all costs associated with reaching the POP.
Select POP Location for End Point:

Connection Speed:
Customer is responsible for all cross-connects.
☒ 1G ☐ 10G

Remote Deployment:
Is the remote location an existing build or a greenfield deployment?
☐ New ☒ Existing

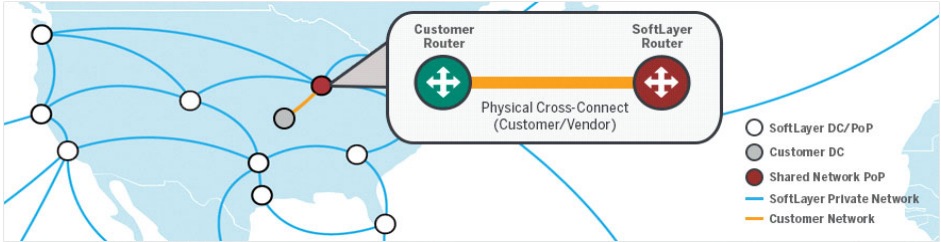
Define IP Scheme for Existing Deployment:

Will you need additional IPs for NAT Services?
SoftLayer does not provide NAT services. IPs must be configured by the customer on the remote side.
☐ No ☒ Yes

How many additional IP addresses do you need?
IPs supplied will originate from 192.168.0.0/22.
/24 - 256

Please describe why you will need these additional IP addresses

- Deploy physical circuit
- Order Direct Link through portal
- Select PoP
- Select connection speed
- Select option for Remote Deployment
- Sign the letter of authorization
- Enjoy your dedicated direct link!



© 2015 IBM Corporation




Figure 5-35 Direct Link use case (continued)

Notes:

To configure Direct Link, you must perform these steps:

1. Deliver and deploy the physical circuit (see previous slides).
2. Order the Direct Link through the SoftLayer portal `control.softlayer.com` under **Network** → **Direct Link**.
3. Select the points of presence (PoP) location for the end point, which is the Data Center where your SoftLayer environment is provisioned.
4. Select the connection speed: 1 Gb or 10 Gb (customer is fully responsible for all cross-connects).
5. Select options for Remote Deployment:
 - New deployment (default): Gives you a layer 3 connection to a single IP address (/30 or /31). The SoftLayer description of these options are NAT (source-nat overload) and Tunneling for BYOIP.
 - Source NAT: Only getting one IP address really forces you to source NAT everything from your remote site. The real life translation is that everything in your remote site can connect to all of your SoftLayer server/services, but there will be limited connectivity from SoftLayer into the remote site. This model only works when your remote site only has users, and no local servers/services.

Unit 5. Networking options in cloud computing 149

- Tunneling for BYOIP: This option is really the only option when you need failover capabilities from a remote site into two SoftLayer PoPs (Data Center Locations). You still rely on the IP addresses that are provided in the SoftLayer environment, but you can bring whatever client-side addressing you want.
 - Existing: When you select that, you get an option to define your existing IP scheme. The key decision is whether you need additional IPs for NAT services. The answer is going to be yes, if you are connecting an existing environment and not using GRE tunneling. What you do need to decide is how many IP addresses that you need for your NAT pool. If you are planning on a huge amount of interaction between the cloud and existing environments, you will need a large pool of IPs.
6. When you receive the letter of authorization, contact your colocation data center provider to request a cross-connect to the SoftLayer PoP.
 7. After the cross-connect for the Direct Link network has been patched to the customer demarcation panel in the colocation facility, connect the customer network switch to that panel.
 8. After the SoftLayer Direct Link cross-connect is patched correctly, use the SoftLayer Customer Portal to submit a support ticket requesting BGP peering IP addresses, a BGP authentication key, and the 802.1q VLAN ID number used for the Direct Link connection.
 9. SoftLayer will assign you an IP address range out of the 172.0.0.0/12 subnet. These will be routable within the SoftLayer environment. So a server sitting on your SoftLayer private VLAN will be able to ping the 172.16.0.10 address.
 10. Configure routing on your local devices (routers, servers).

5.39 Direct Link use case (continued) (2)

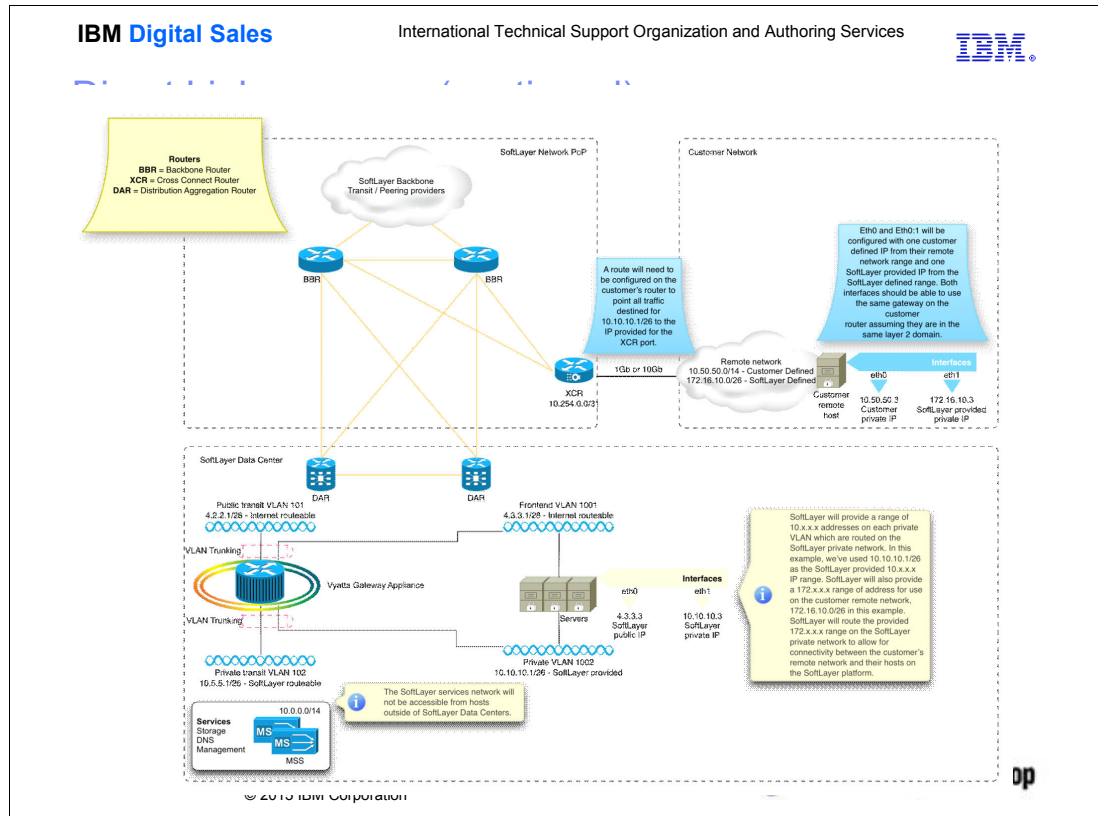


Figure 5-36 Direct Link use case (continued) (2)

Notes:

This diagram is intended to show how a customer can configure communication between their SoftLayer hosts and hosts on their remote network through a direct link with a dual IP scheme on their remote hosts. A customer also has the ability to reIP their existing hosts on their remote network into the SoftLayer provided 172.x.x.x IP range if they prefer not to use a dual IP setup. All IPs used in this diagram are example IPs and will be different on deployment with the exception of the SoftLayer services network (10.0.0.0/14).

Vyatta Gateway Appliance and detailed information about SoftLayer's network topology are covered in Unit 9.

5.40 Recap


IBM Digital Sales

International Technical Support Organization and Authoring Services



Recap

- Network in Cloud Computing
- SoftLayer network strategy and topology
- Network tools in SoftLayer
- VPN connections in cloud computing and SoftLayer
- Direct Link use cases





© 2015 IBM Corporation

Figure 5-37 Recap

Notes:
(none)


5.41 Checkpoint questions

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



1. Briefly explain SoftLayer's tiered network.

© 2015 IBM Corporation




Figure 5-38 Checkpoint questions

Notes:

(none)


5.42 Checkpoint questions (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



CHECKPOINT AHEAD

1. Briefly explain SoftLayer's tiered network.

Answer:

- Public Network: Used to access content from a server over the Internet.
- Private Network: Separated network used for communication between SoftLayer DCs and PoPs without interfering with public network traffic.
- Out-of-Band Management Network: Accessible through VPN allows access to servers or any other devices for administration and maintenance purposes.

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-39 Checkpoint questions (2)

Notes:
(none)


5.43 Checkpoint questions (3)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



2. What are the VPN implementations used in SoftLayer?

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-40 Checkpoint questions (3)

Notes:

(none)


5.44 Checkpoint questions (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



2. What are the VPN implementations used in SoftLayer?

Answer: SoftLayer uses 3 different types of VPN connections:

- SSL VPN – Secure Sockets Layer
- PPTP VPN – Point-to-Point Tunneling Protocol
- IPSec VPN – Internet Protocol Security

© 2015 IBM Corporation

Redbooks Workshop

Figure 5-41 Checkpoint questions (4)

Notes:

(none)


5.45 Checkpoint questions (5)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



3. How many broadband carriers at the minimum are needed to provide links to single SoftLayer's DC or PoP, and why?

© 2015 IBM Corporation




Figure 5-42 Checkpoint questions (5)

Notes:
(none)


5.46 Checkpoint questions (6)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



3. How many broadband carriers at the minimum are needed to provide links to single SoftLayer's DC or PoP, and why?

Answer: SoftLayer needs at least two broadband carriers for a single DC or PoP to avoid single points of failure.

© 2015 IBM Corporation




Figure 5-43 Checkpoint questions (6)

Notes:
(none)


5.47 Checkpoint questions (7)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



4. What is the bundled outbound public bandwidth for SoftLayer's servers?

© 2015 IBM Corporation





Figure 5-44 Checkpoint questions (7)

Notes:
(none)


5.48 Checkpoint questions (8)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint questions



4. What is the bundled outbound public bandwidth for SoftLayer's servers?

Answer: With:

- With monthly billed virtual servers, it is 250 GB/month included.
- With monthly billed bare metals, it is 500 GB/month included.
- With hourly billed servers, it is not included.

© 2015 IBM Corporation




Figure 5-45 Checkpoint questions (8)

Notes:
(none)



Managing a simple IaaS environment

This unit provides information about ordering and deploying a SoftLayer IaaS environment.

6.1 What you should be able to do

After completing this unit, you should be able to:

- ▶ Understand the basic concepts of cloud infrastructure and its components.
- ▶ Understand the technologies that are used to enable IaaS.
- ▶ Know how to manage your cloud infrastructure using the Customer Portal.

6.2 Teaching topics

The following topics are covered in this unit:

- ▶ Review basics of cloud infrastructure and components
 - IaaS overview
 - Platform Architecture
 - Servers
 - Storage
 - Network
- ▶ Technologies enabling IaaS
 - Provisioning
 - Hypervisor
 - Virtualization
 - Security
- ▶ Managing cloud infrastructure
 - Comparison with traditional IT
 - Managing your account in IaaS

6.3 Basics of cloud infrastructure and components

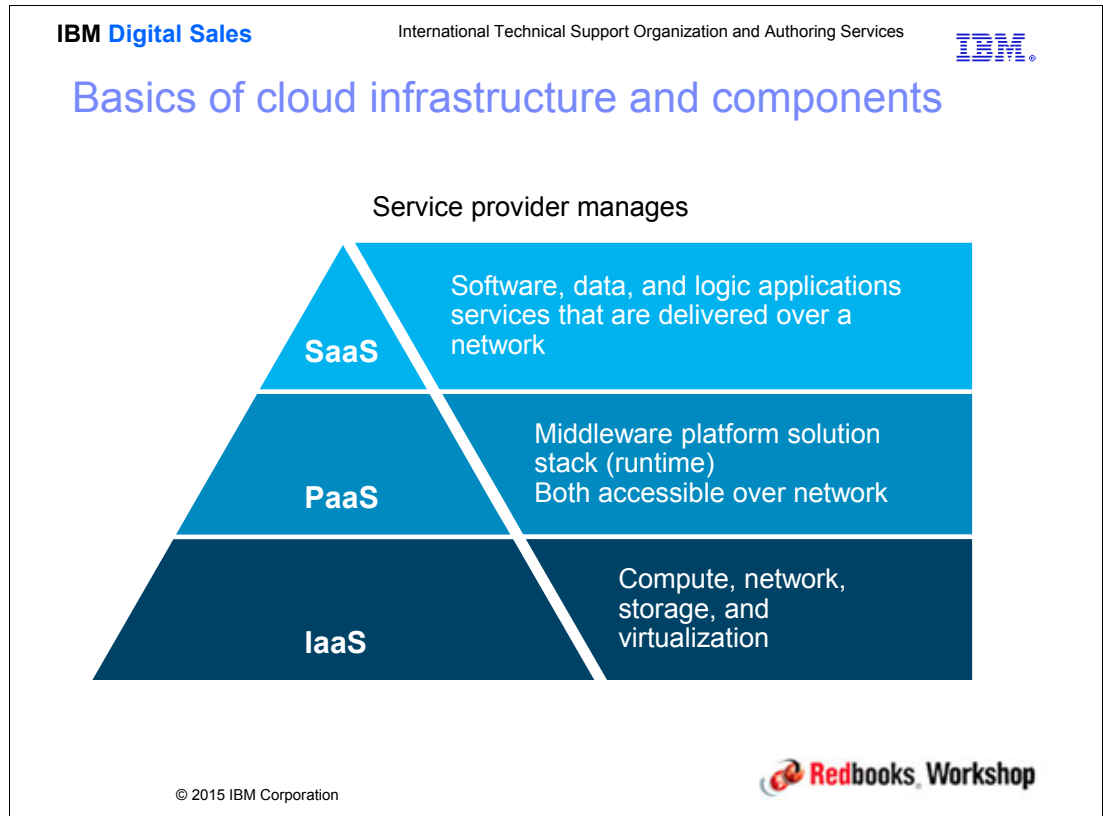


Figure 6-1 Basics of cloud infrastructure and components

Notes:

There are three service models on cloud, which vary in how much the cloud provider is responsible for managing:

- ▶ Software as a Service (SaaS)
- ▶ Platform as a Service (PaaS)
- ▶ Infrastructure as a service (IaaS)

There could be exceptions where the customer can negotiate with the provider to manage more, but most implementations follow one of these three models. This course concentrates on the IaaS service model.

6.4 Cloud service model overview

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Cloud service model overview

IaaS

Compute, network, storage, and virtualization

In IaaS we “rent” our infrastructure (servers, storage, datacenter, network, and the resources needed to maintain them.

This affects our business in several ways:

- No investment in infrastructure.
- Able to scale infrastructure up or down in hours.
- High dependency on the network because the servers are not in house.
- Complete dependency on IaaS provider.
- Cost of infrastructure can vary depending on usage, harder to budget.
- Need a strong design to avoid over-provisioning

© 2015 IBM Corporation

Redbooks Workshop

Figure 6-2 Cloud service model overview

Notes:

With IaaS as used in this course, the customer has little or no need to invest in infrastructure because this is all included in the subscription fee paid to the provider. The customer gains the flexibility to scale up infrastructure within minutes or hours should it be needed, and can scale down immediately after the peak period.

However, not having servers in house make you more dependent on internet access as you cannot reach your servers otherwise. Likewise you are dependent on the IaaS provider because you put your data (and thus your business) in the hands of a third-party company that you need to be able to trust.

Lastly, the financial department will have to cope with the fact that you cannot give them a fixed price for your infrastructure cost because you pay depending on usage. You might need extra capacity during the release of a new version of software. Then, after a few days you scale down back to normal. However, this period where you upgraded the capacity appears on the next bill and will likely cause the financial department to raise questions because the cost differs from the norm.

Again, it is easy to click and provision in IaaS, perhaps too easy so make sure you have a correct design of your environment and do not provision more than you need because it is not changed directly.

6.5 Platform architecture

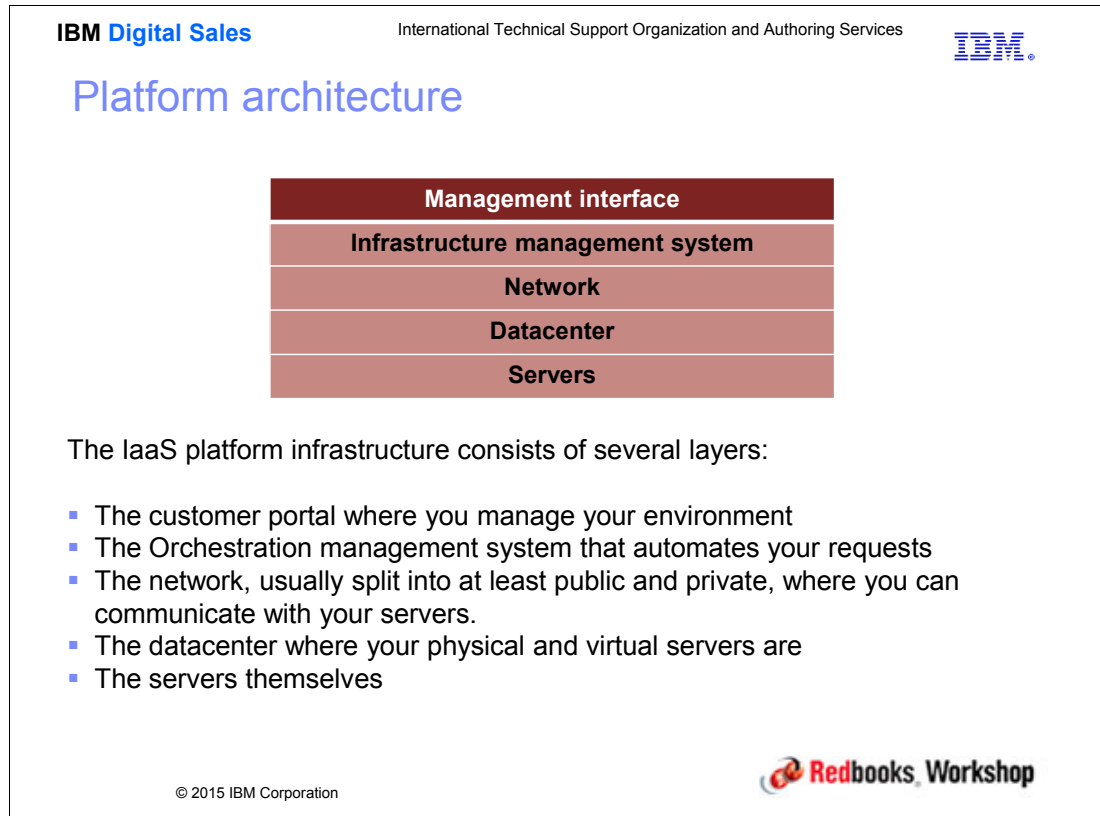


Figure 6-3 Platform architecture

Notes:

The IaaS cloud infrastructure consists largely of a management interface through which you can submit requests to your servers or hosted infrastructure.

Your requests are then intercepted by the automation layer and the infrastructure management system, and your requests are routed through the network to the data center where your servers are hosted.

There might be several networks for the instance (private, public, and management) and depending on the request and type, several networks can be involved.

6.6 Locations of platform architecture

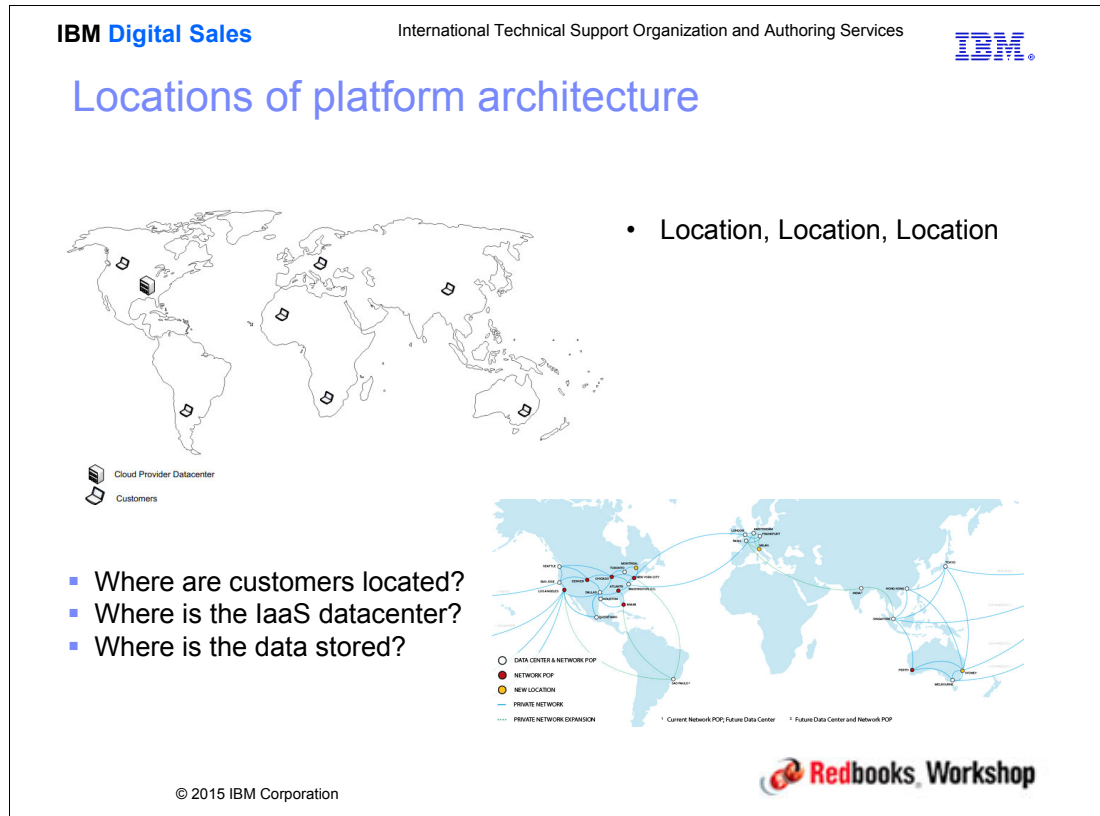


Figure 6-4 Locations of platform architecture

Notes:

The location of the IaaS provider is likely an important issue because stable and fast access to your hosted infrastructure is important. Customers are unlikely to wait for websites to load or to accept unstable connections. In addition, laws in your country or region might dictate where you may host your data and through which countries that the data may travel.

This means that the location of the providers datacenter, and possibly the geographical footprint, is of great importance to your business and choice of provider.

6.7 Types of servers

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Types of servers



IaaS usually involves these types of servers:

- Bare metal servers
- Virtual servers

© 2015 IBM Corporation

Redbooks Workshop

Figure 6-5 Types of servers

Notes:

When you provision servers in your IaaS environment, you can choose virtual server or bare metal server. Not all providers have both and there can be different types of each.

SoftLayer offers the following server types:

- ▶ Bare metal: A physical server dedicated to your business and build according to your provisioning request.
- ▶ Virtual server on a private node: A virtual server dedicated to your business, and running on a physical host that is also dedicated to your business.
- ▶ Virtual servers on a public node: A virtual server dedicate to your business, but running on a host that also serve other customer VMs. Although your VM is isolated from other customers, the resource pool (memory, cpu) on the host is shared with other customers, and their workloads will affect the performance of your server.

It is up to the IaaS provider to dimension the public and private node host so that the VMs running on it can perform, as you as the customer have no influence on that.

6.8 Server options

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Server options



You build your servers as you wish and can pick and choose from many options.

Servers are billed either hourly or monthly.

© 2015 IBM Corporation

Redbooks Workshop

Figure 6-6 Server options

Notes:

When you build your servers, determine how much memory, CPU, and disk you want to add. Specify the network bandwidth and network speed, and order monitoring and firewalls to be in place once the server is up.

Having said that, there are some “rules” set by the IaaS provider so you cannot select 1.5 CPU or 4.5 GB of memory. They have a wide selection, but you might have to order 1 GB memory more than you need. That should be a small compromise to make.

Depending on how you configure your bare metal server, it might take a little longer to get up and running if your specifications do not match the prebuilt ones that the provider supplies.

Naturally, you can remove or add more memory and CPU as needed. This process is faster on a virtual server than a physical one because it requires no physical intervention.

Lastly, some service add-ons are not available with hourly billed servers such as hardware firewalls. Because it is automated, the reason could be that they want a guaranteed longer commitment before setting up some of the more advanced services.

Regardless of whether you select a bare metal or virtual server, it should be ready within a couple of hours at the most and usually minutes (for virtual ones).

6.9 Storage types and protocols



Figure 6-7 Storage types and protocols

Notes:

When choosing storage for IaaS, you can choose among these types:

- ▶ Direct-attached storage (DAS): DAS is storage that is directly attached to the servers just like home computers or laptops.
- ▶ Storage area network (SAN): SAN is a type of network storage that is attached through high speed links. To the operating system, the SAN storage appears as if it was locally attached.
- ▶ Network-attached storage (NAS): NAS is a type of network storage that is specialized for serving files and not as fast as SAN.
- ▶ Object Storage: Object Storage treats data as objects, and can be used to store files like Virtual Machine images, backups, and archives as well as photos and videos.

6.10 Other storage types

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Other storage types

Local Storage options from SoftLayer	
Bare metal	Virtual servers
SAS	SAN
SATA	Local Disk
SSD	Portable Storage Volume

Other storage types offered by SoftLayer:

- iSCSI based block storage
- NFS based file storage
- Endurance, performance, and object storage

Other IaaS providers might have different offerings, but SoftLayer has these storage options available.

© 2015 IBM Corporation

Redbooks Workshop

Figure 6-8 Other storage types

Notes:

The selection of storage should be based on your requirements and usage of the server.

For bare metal the highest speed (and corresponding cost) is achieved with SSD, with SAS and the SATA being less costly, but slower. If speed is not a key requirement, use the type that costs less.

For virtual servers, you can place your virtual disk on local storage of the hHypervisor or on SAN-based storage. You also have the option to place your data disks (except the operating system disk) on portable storage, which enables seamless transfer of data between VMs.

6.11 Network types

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Network types

Network
Public network
Private network
Out-of-band Management network

IaaS providers are likely use three types of networks as part of their infrastructure:

- Public network – Accessible from the internet
- Private network – Accessible from servers/devices in your account
- Out-of-band Management network – Used to manage server and devices and get console access to servers when required

IaaS providers should offer VPN connectivity to the private network.

© 2015 IBM Corporation

Redbooks Workshop

Figure 6-9 Network types

Notes:

An IaaS provider is likely to have three different networks as part of their infrastructure:

- ▶ A public network that can access and is accessible from the internet
- ▶ A private network that is accessible only from server in your account, and if you connect through VPN to the servers.
- ▶ A management network, so called out-of-band, that you cannot modify in any way and is used for maintenance and console access to the servers.

You can usually disconnect private and public network interfaces on your servers as needed. For example, you might want to disconnect a public interface on servers if your account does not need to be accessible from the internet but only from other servers in our account or through VPN. You can access the account from both of these using the private network instead.

6.12 Basics of cloud infrastructure and components

IBM Digital Sales

International Technical Support Organization and Authoring Services



Basics of cloud infrastructure and components



Most IaaS providers will as part of their offering also provide tools to help you monitor the network such as the following:

- Speed test
- Latency of the network between datacenters or between you and your hosted servers.
- Download tests from the providers different datacenters

© 2015 IBM Corporation



Figure 6-10 Basics of cloud infrastructure and components

Notes:

Because the network is essential for both you and your customers to access to your servers and the data stored on them, any serious IaaS provider should offer you the means to test the state of their network.

It could easily be that the nearest data center might not be the one where you or your customer experience the best or most stable connection. In the case of a sudden increase in the time it takes to access your data/servers, it is useful to troubleshoot whether the problem is at the provider's end, your end, or somewhere in-between.

6.13 Conclusion of recap

IBM Digital Sales


International Technical Support Organization and Authoring Services

IBM

Conclusion of recap

This concludes the recap of Units 1 through 5 of this course. If you do not understand some of the terms mentioned in this recap, go back and revisit the relevant unit.

If you start the course from here, be aware that the units after this one assume that the labs in units 1 through 5 have been completed



© 2015 IBM Corporation


 **Redbooks** Workshop

Figure 6-11 Conclusion of recap

Notes:

(none)

6.14 Technologies enabling transformation of infrastructure

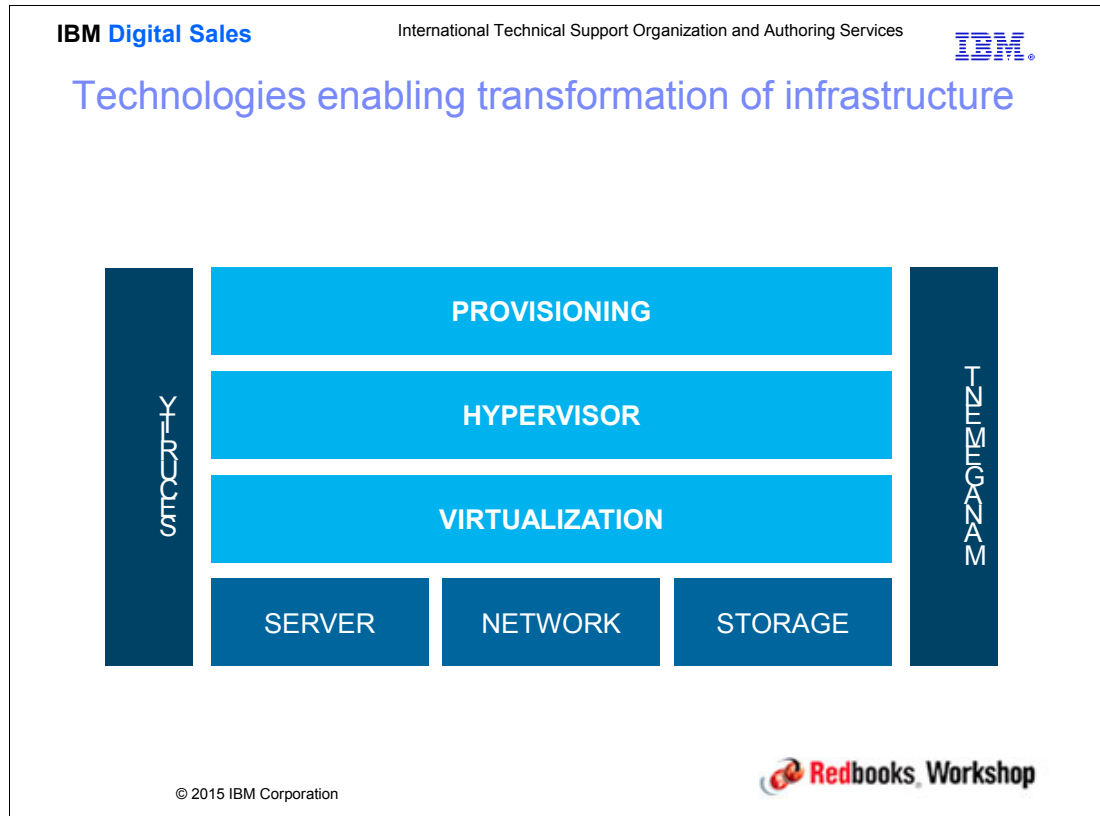


Figure 6-12 Technologies enabling transformation of infrastructure

Notes:

Cloud computing is a group of pre-existing technologies that are enabled to provide services to your clients.

Aside from the need to have a high speed, low cost, and scalable computing environment, some technological power covers the evolution of cloud computing. Nowadays, clouds are supported by a set of primary technology components that are combined together to enable key features and characteristics needed by cloud computing:

- ▶ Virtualization
- ▶ Hypervisor
- ▶ Provisioning

6.15 Overview of technologies

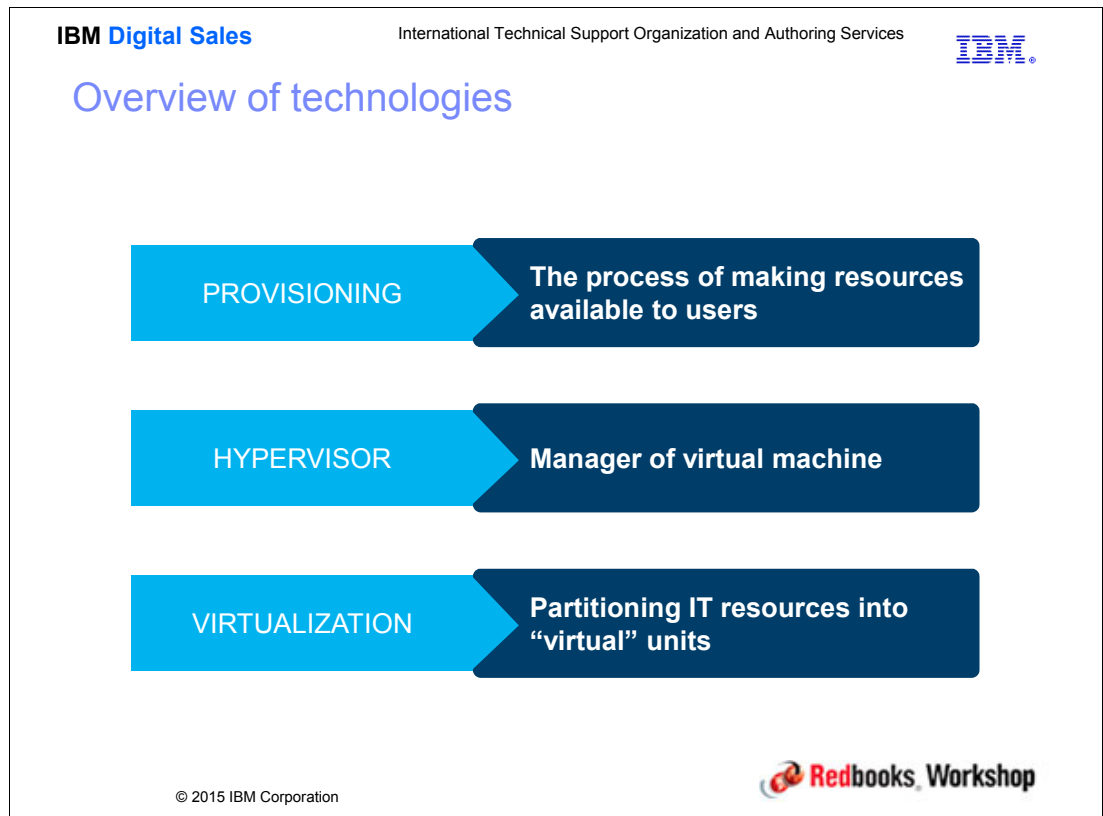


Figure 6-13 Overview of technologies

Notes:

(none)

6.16 Virtualization

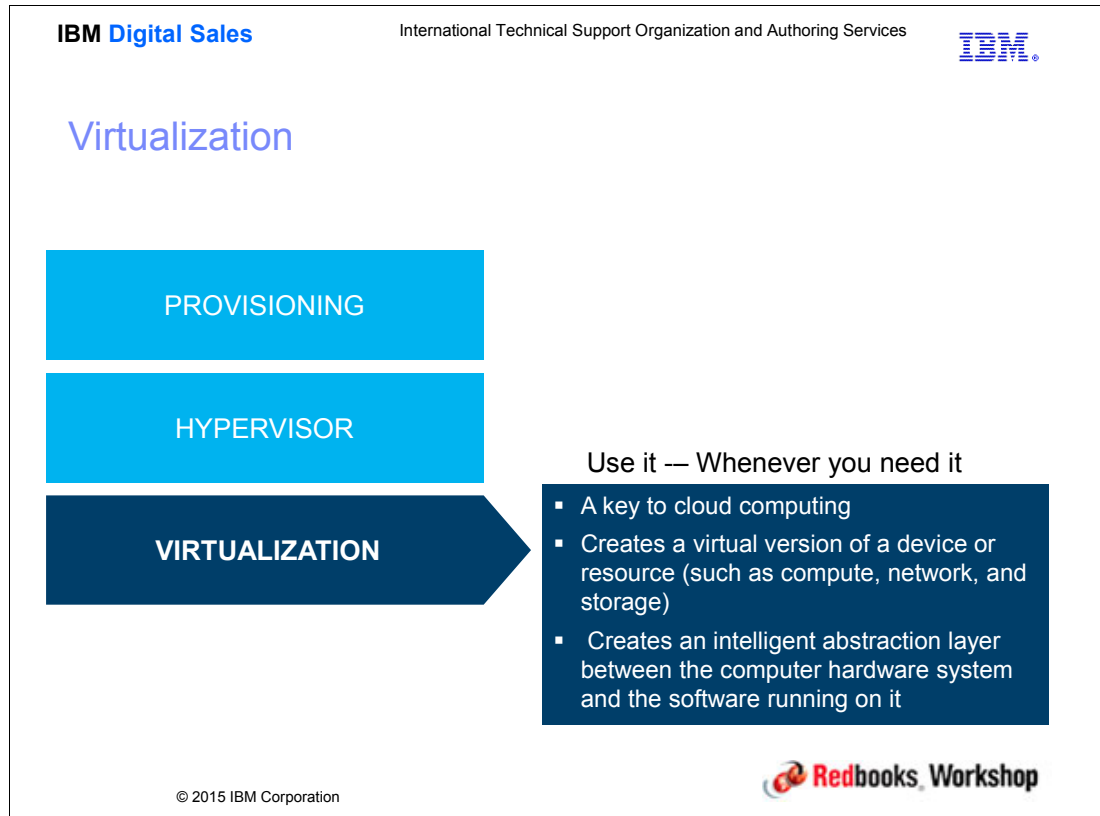


Figure 6-14 Virtualization

Notes:

Virtualization is the key to cloud computing because it enables the technology that allows the creation of an abstraction layer of the computer hardware system and the software running on them. It allows a single machine act as if it were many machines.

The following are the subnets of virtualization:

- ▶ Server virtualization is the masking of server resources. It enables different operating systems to share the same hardware. The operating systems are easy to move between different hardware even if the applications are running.
- ▶ Storage virtualization creates the abstraction layer between the applications running on the servers, and the storage they use to store the data.
- ▶ Network Virtualization is the process of combining hardware and software network resources with network functionality to create a virtual network. It creates logical/virtual networks to separate from underlying network hardware. This is to ensure that it can better integrate and support the virtual environments.

Virtualization is important because you can adjust environment's compute, network, or storage resources anytime if a need arises.

6.17 Hypervisor

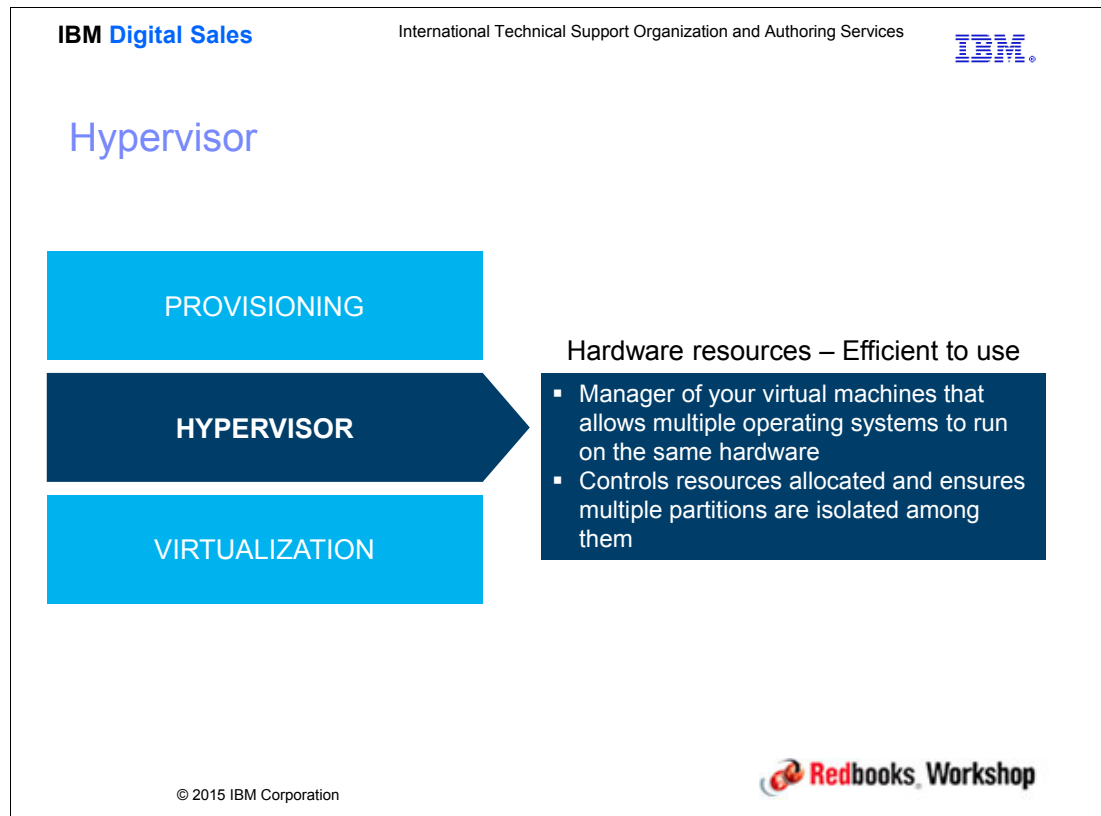


Figure 6-15 Hypervisor

Notes:

Hypervisor is a critical component of a virtual server because it is the virtual machine manager that allows multiple operating systems (virtual machines) to run on a same hardware. It is the foundation for virtualization of server that enables and supervises that partitioned IT resources and ensures isolation among them.

Because the hypervisor allows multiple VMs to run on a same hardware, it helps optimize the use of the resources. This is important to achieve the hardware's maximum productivity for all the VMs hosted to it.

6.18 Provisioning

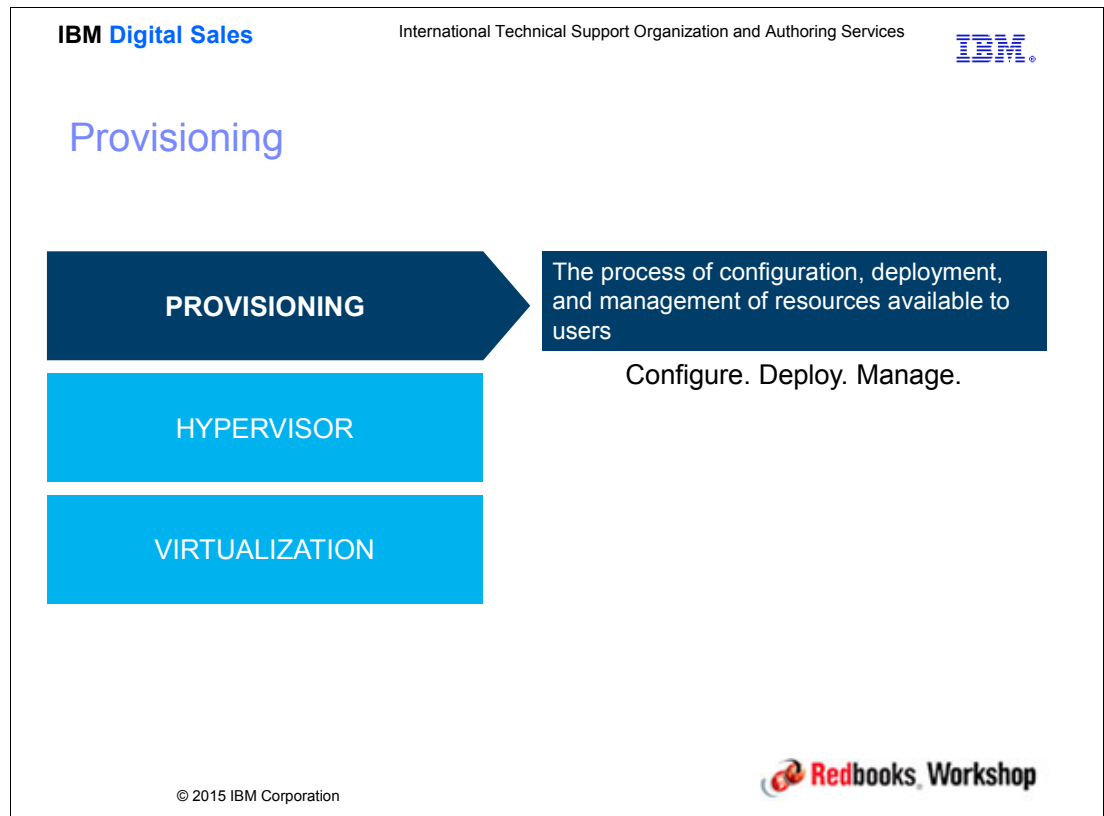


Figure 6-16 Provisioning

Notes:

Provisioning is the process of configuration, deployment, and management of multiple types of IT system resources. The self-service provisioning for cloud computing services allows the users to acquire and remove cloud services anytime.

Configure. Deploy. Manage.

Provisioning helps deploy resources for the application in a small amount of time, configure the resources based on your specifications, and manage the infrastructure whenever you want.

6.19 Managing cloud infrastructure

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Managing cloud infrastructure

Comparison with traditional IT

Traditional IT	IaaS
Traditional data centers consist of a fragmented single purpose tools that are limited in scope	Convenient, on-demand access to an array of products and service offerings (compute, networking, security, storage, and services)
Complex and hard to manage infrastructure that cannot scale easily	Easily provisioned and minimal management
Requires continuous funding to keep environment updated and running	Subscription model, pay for what you use model.
Requires specialized resources to operate and manage environment	Once environment is designed, the provider handles the management.
Big initial investment	No upfront investment needed

© 2015 IBM Corporation




Figure 6-17 Managing cloud infrastructure

Notes:


With the move from traditional IT to IaaS, the fundamentals of how to operate your infrastructure have changed.

With traditional IT, you not only have to invest in hardware and resources to manage and operate that hardware, but you also need to have cooling, power, and likely emergency power to handle unforeseen events. Additionally you need to set aside funding for upgrading, servicing, and updating the infrastructure.

With IaaS, all of these are outsourced to the provider and you simply pay a subscription fee.

Additionally, if you need to upgrade a server for a period of time, you can do so with a few clicks and later remove it again. The only effect is that your subscription fee will be a little higher during that period. With traditional IT, you would have to purchase the upgrade and once the peak period was over it would still be there unused.

6.20 Cloud infrastructure IaaS key points

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Cloud infrastructure IaaS key points

- Low entry fee compared to traditional IT
- Instead of purchasing hardware, floor space, software, and networking equipment, you “rent” it
- Only acquire the resources that you need and when you need them
- Cost savings based on a utility computing basis
- Removal of capital expenditures for hardware and software
- Removal of operational expenses incurred by staffing
- Advanced self-service management tools

© 2015 IBM Corporation




Figure 6-18 Cloud infrastructure IaaS key points

Notes:

These are the key points when considering IaaS versus traditional IT, and part of the reason that startups often use IaaS or one of the other cloud service models.

6.21 Managing cloud infrastructure in IaaS

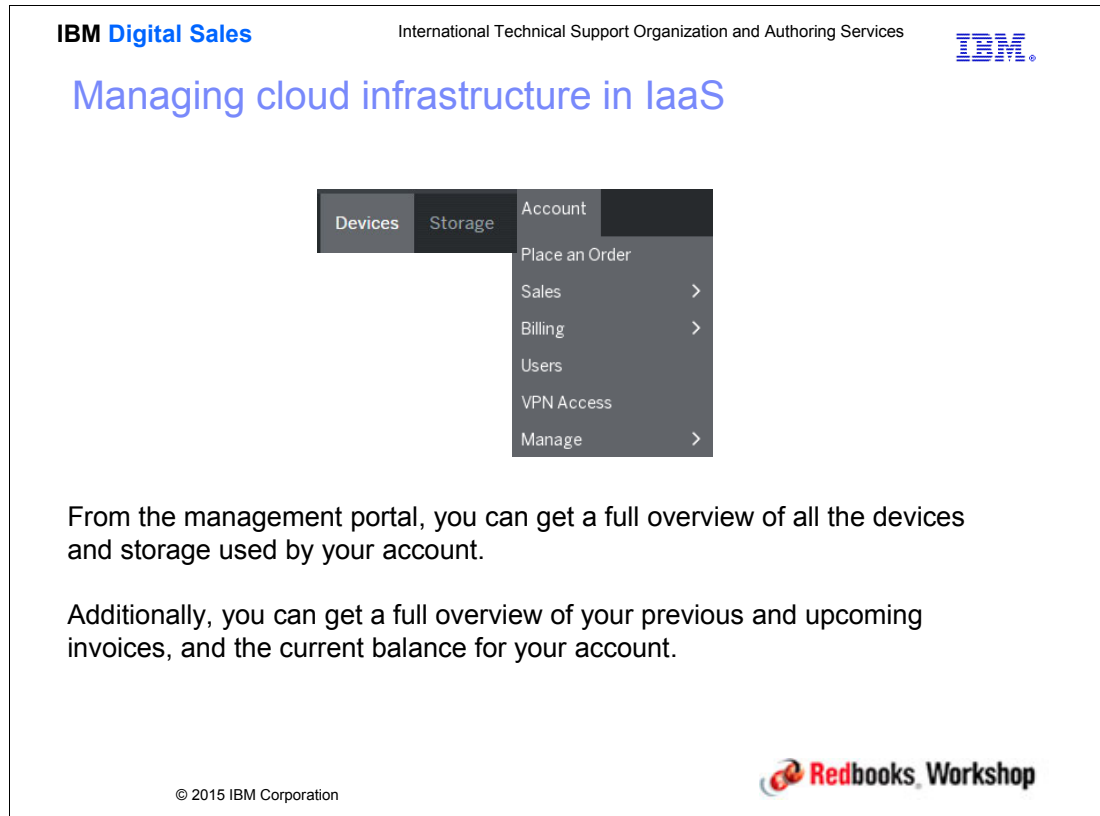


Figure 6-19 Managing cloud infrastructure in IaaS

Notes:


The IaaS management portal, throughout this course referred to as the Customer Portal, is the one place to go to get a full overview of devices that are or have been provisioned. It also displays the amount of storage used and the current balance for your account.

Additionally, you can see past and present invoices and, depending on your account configuration, approve or deny pending requests from users in your account.

6.22 Checkpoint


IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint


- Does IaaS have a low or high entry fee ?
Low, because you pay for what you use and no investments are required.
- Do you need a big staff to manage your IaaS servers ?
No, the provider does most of the managing, and you do yours from the Customer Portal.
- Which place can you see the current balance of your IaaS account ?
The Customer Portal.
- Can you get a full overview of your devices from the Customer Portal ?
Yes, you can get a full overview of devices, storage, and billing from the Customer Portal.



© 2015 IBM Corporation

Figure 6-20 Checkpoint

Notes:
(none)



Server offerings in cloud computing: Advanced topics

This unit provides information about these topics:

- ▶ Image templates
 - Types of image templates
 - Common usage scenarios using image templates
- ▶ Using image templates in SoftLayer
 - Creating, sharing, using and deleting image templates
 - Migrating from virtual servers to bare metal servers
 - Manual and automatic scaling

7.1 What you should be able to do

Upon completion of this unit, students should be able to:

- ▶ Explain the concept of an image template for capture, deployment, and usage scenarios
- ▶ Describe the different SoftLayer image template types
- ▶ Explain how to capture a custom image
- ▶ Create virtual servers and bare metal servers using templates and provisioning scripts
- ▶ Migrate from a virtual to a bare metal server
- ▶ Explain how to scale the environment with images
- ▶ Describe how to automate manual scaling

7.2 Teaching topics

The following topics are covered in this unit:

- ▶ Introducing the image template concept
- ▶ Image template types in SoftLayer
- ▶ Creating image templates in SoftLayer
- ▶ Sharing, finding and deploying the image templates
- ▶ Provisioning scripts
- ▶ Usage scenarios in SoftLayer:
 - Migrating from virtual server to bare metal server
 - Using images to scale an environment

7.3 Introducing the image template concept

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Introducing the image template concept

- Image templates are a method of capturing a device image to quickly replicate its configuration with minimal manipulations
- A point in time image capture can be useful in scaling scenarios
- Image templates are used across various cloud offerings
- There is no common (compatible) implementation across the various cloud offerings

© 2015 IBM Corporation


 Redbooks Workshop

Figure 7-1 Introducing the image template concept

Notes:

Usually there are import/export tools to allow images from a particular cloud offering to be used in a different one.

7.4 Image template types in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Image template types in SoftLayer

SoftLayer currently provides two options for creating image templates, each offering unique features based on operating system and image type:

- **Standard Image** provides an imaging option for all virtual servers, regardless of their operating system.
- **Flex Image** is server-type neutral.

© 2015 IBM Corporation


 **Redbooks Workshop**

Figure 7-2 Image template types in SoftLayer

Notes:

(none)

7.5 Image template types in SoftLayer: Standard images

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Image template types in SoftLayer: Standard Images

- Allow capture of virtual server images and replication of their configuration to other virtual instances.
- Are available on all virtual servers, and do not require a specific operating system for functionality.
- Not available for bare metal servers.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 7-3 Image template types in SoftLayer: Standard images

Notes:

(none)

7.6 Image template types in SoftLayer: Flex Images

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Image template types in SoftLayer: Flex Images

- Flex Images are SoftLayer's platform-neutral imaging system that allows users to capture an image of both bare metal servers and virtual servers and apply that image to either platform.
- Having the freedom to create images across platforms allows users to create bare metal servers from images of virtual servers and vice versa, providing more flexibility in scaling.

© 2015 IBM Corporation

Redbooks Workshop

Figure 7-4 Image template types in SoftLayer: Flex Images™

Notes:

Flex Image is available on machines running out of any SoftLayer data center, worldwide. Flex Image is currently available for use on machines that run one of the following operating systems:

- ▶ CentOS 5 and 6 (7 is not supported yet)
- ▶ Red Hat Enterprise Linux 5 or better
- ▶ Microsoft Windows Server 2003
- ▶ Microsoft Windows Server 2008 R2

Similar to SoftLayer's Standard Image template, Flex Image templates capture an image of a machine and allows you to replicate that machine on another instance. However, Flex Image goes a step further than the Standard Image template because it can be used for replication on both bare metal and virtual servers. In addition, images captured using Flex Image can be used between platforms.

7.7 Image template types in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Image template types in SoftLayer

	Standard Images	Flex Images
Server type - capture	Can be captured only from virtual servers	Can be captured from both virtual and bare metal servers
Server type - provisioning	Can be used to provision only virtual servers	Can be used to provision both virtual and bare metal servers
Operating System	Work with any operating system	Work only with specific operating systems
Migration from bare metal to virtual	Not supported	Supported
Migration from virtual to bare metal	Not supported	Supported
Storage	Object Storage	Object Storage

© 2015 IBM Corporation


Redbooks Workshop

Figure 7-5 Image template types in SoftLayer

Notes:
(none)

7.8 Creating image templates in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Creating image templates in SoftLayer

When creating an image template, the following actions are performed by SoftLayer's Infrastructure Management System (IMS):

- The server is powered off.
- A copy of the data is created.
- The configuration information is recorded.
- The server is started.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 7-6 Creating image templates in SoftLayer

Notes:

(none)

7.9 Creating image templates in SoftLayer (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Creating image templates in SoftLayer

- An image can be captured from both virtual servers and from bare metal servers.
- Standard Images are only applicable to virtual servers. Flex Images support both virtual servers and bare metal servers, not standard images..
- Flex Image templates have specific operating system requirements.
- Image templates can be created from the SoftLayer Customer Portal or using the SoftLayer's application programming interface (API)

© 2015 IBM Corporation


Redbooks Workshop

Figure 7-7 Creating image templates in SoftLayer (2)

Notes:

Image templates can be created in the SoftLayer Customer Portal from the Device List window under the **Devices** menu. The server needs to be turned off when the template is being created. Image templates are charged per gigabyte.

7.10 Sharing, finding, and deploying the image templates

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Sharing, finding, and deploying the image templates

View Storage Usage

Import Image

Filter

Image Templates

Viewing 1 to 25 of 257 | Public Images | Displaying 25 per page | Page 1 of 11

Template Name	Summary	Account / Publisher	Image Type	Create Date	
ArtifactServer Ubuntu 052515 Public	Public image of Artif...		Standard	2015-06-16	Acti...
CohesiveFT VNS3 Cloud Networking Appliance v3...	VNS3 lets you manag...	Cohesive...	Standard	2014-06-04	Acti...
CohesiveFT VNS3 Cloud Networking Appliance v3...	VNS3 lets you manag...	Cohesive...	Standard	2014-06-04	Acti...
100G CentOS 5 32-bit		SoftLayer...	Standard	2011-10-18	Acti...
100G CentOS 5 64-bit		SoftLayer...	Standard	2011-10-18	Acti...
100G CentOS 6 32-bit		SoftLayer...	Standard	2011-10-18	Acti...
100G CentOS 6 64-bit		SoftLayer...	Standard	2011-10-18	Acti...
100G Debian 5 32-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G Debian 5 64-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G Debian 6 32-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G Debian 6 64-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G Fedora 15 32-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G Fedora 15 64-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G RedHat 5 32-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G RedHat 5 64-bit		SoftLayer...	Standard	2011-10-19	Acti...
100G RedHat 6 32-bit		SoftLayer...	Standard	2011-10-19	Acti...

© 2015 IBM Corporation




Figure 7-8 Sharing, finding, and deploying the image templates

Notes:

The diagram presents an example view of the Image Templates window in SoftLayer. The default view on this window shows all the private images associated with the account. Public images can be accessed from the same page.

- ▶ **Private Images:** Private Images are those created by a user on the account or images created on another account that have been shared with the account. By default, all images created are private.
- ▶ **Public Images:** The Image Templates window containing public images displays images of pre-configured machines posted by SoftLayer and are available for use by all SoftLayer customers. The Public Image templates were created with optimal performance in mind and provide a comprehensive list of choices.
- ▶ **Editing image templates:** Details regarding private images can be viewed and edited, while details regarding public images are read-only.

Both Flex and Standard image interaction is identical.

7.11 Sharing, finding, and deploying the image templates (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM.

Sharing, finding, and deploying the image templates

Image Template Details

[View All Images](#)

Image Type: Flex Image Account: 568347 Created: 2015-07-22 12:17 PM (UTC-04:00) Disk Space: 1.08 GB [Actions](#)

*Image Name:

Notes:

Locations:

☐ Amsterdam 1 ☐ Amsterdam 3 ☐ Dallas 1 ☐ Dallas 5 ☐ Dallas 6 ☐ Dallas 9

☐ Frankfurt 2 ☐ Hong Kong 2 ☒ Houston 2 ☐ London 2 ☐ Melbourne 1 ☐ Mexico 1

☐ Milan 1 ☐ Montreal 1 ☐ Paris 1 ☐ San Jose 1 ☐ Seattle ☐ Singapore 1

☐ Sydney 1 ☐ Tokyo 2 ☐ Toronto 1 ☐ Washington 1

Virtual Disks

Device	Capacity	Size on Disk
0 CentOS CentOS 6.5-64	4 GB	1.08 GB

Share Image

This image can be shared with other accounts. Simply type in the account number in the field below and click share.

*Share Image with Account Number:

To unshare, remove the account in the table below.

Account	Shared On
---------	-----------

© 2015 IBM Corporation

Redbooks Workshop

Figure 7-9 Sharing, finding, and deploying the image templates (2)

Notes:

This image shows the following options:

- ▶ Editing the details of a private image
- ▶ Sharing a Private Image across data centers and across SoftLayer accounts

7.12 Sharing, finding, and deploying the image templates (3)

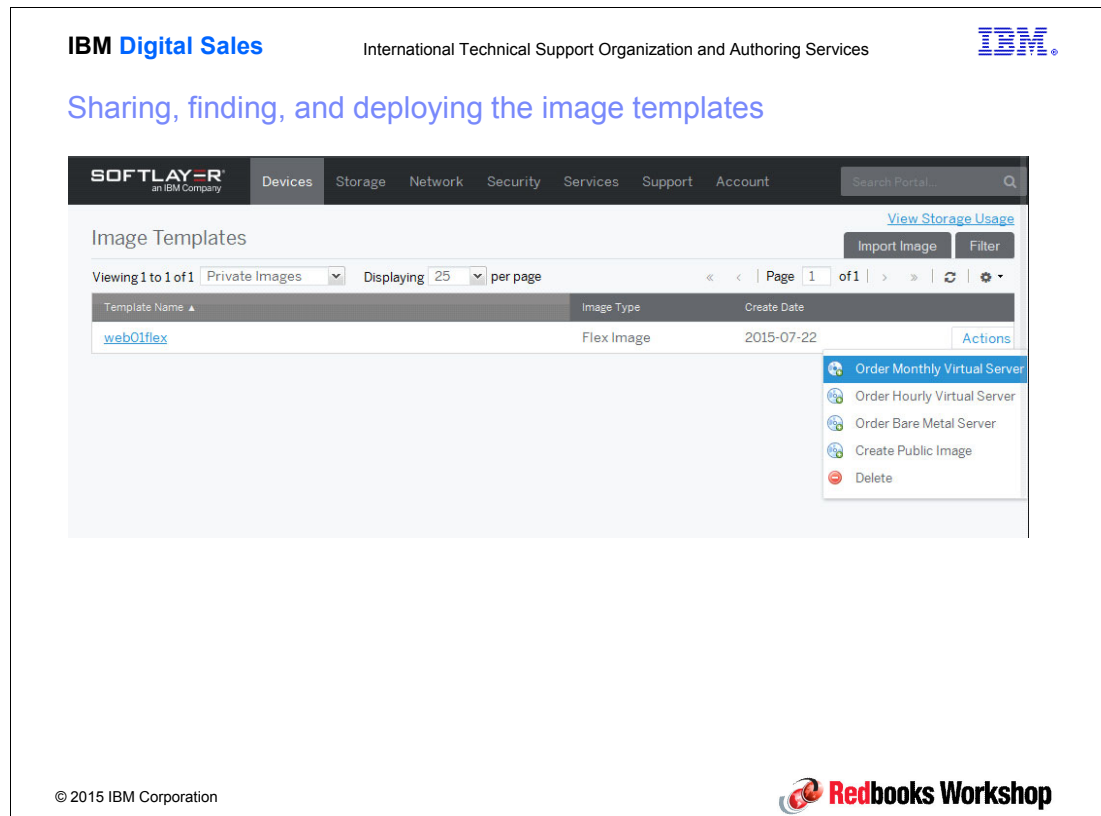


Figure 7-10 Sharing, finding, and deploying the image templates (3)

Notes:

This image shows how to create virtual or bare metals servers from a Flex Image template.

7.13 Sharing, finding, and deploying the image templates (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Sharing, finding, and deploying the image templates

An image template can be deployed on a new or an existing server:

- On a new server, order the template from the Images view or by setting up an Autoscale group.
- On an existing server, select the **Load from Image** option in the Device view in the software Customer Portal.
- Both options can be automated using the SoftLayer API.

© 2015 IBM Corporation


 **Redbooks Workshop**

Figure 7-11 Sharing, finding, and deploying the image templates (4)

Notes:

During the deployment stage of the Image Template, SoftLayer's Infrastructure Management System constructs a new machine based on the data gathered from the selected image, making adjustments for volume. It then restores the copied data and then makes final configuration changes (for example, network configurations) for the new host.

7.14 Provisioning scripts

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Provisioning scripts

A provisioning script is a script that may be downloaded and executed on a device during the provisioning process. Provisioning scripts must have these characteristics:

- Associated with a fully-qualified domain name
- Accessible via the HTTP or HTTPS protocol
- Executable by the operating system

© 2015 IBM Corporation


Redbooks Workshop

Figure 7-12 Provisioning scripts

Notes:

Provisioning can be downloaded to a device during the provisioning process from a URL specified during the time of order creation. For existing accounts, provisioning scripts are managed within the Customer Portal. Additionally, scripts for new accounts or scripts that are not yet tracked on the Customer Portal can be entered manually during the ordering process.

During the provisioning process, scripts associated with an HTTP URL are downloaded to the device and must be manually executed on the device by an administrator after it has been provisioned. Scripts associated with an HTTPS URL are downloaded and executed. If the URL is not associated to an executable script, the script will simply be downloaded and no further action will be taken.

Provisioning scripts should be executable on the provisioned server and the prerequisites (such as the Python interpreter if it is a Python script) should be present.

7.15 Usage scenarios in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Usage scenarios in SoftLayer

Migrating from virtual to bare metal has these considerations:

- Allows for higher performance.
- It is a scenario supported by Flex Image templates.
- A Flex Image template must be captured from a virtual server that has one of the supported operating systems.
- The Flex Image template must be then used to order bare metal servers.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 7-13 Usage scenarios in SoftLayer

Notes:

(none)

7.16 Usage scenarios in SoftLayer (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Usage scenarios in SoftLayer

Environment scale-out can be leveraged with both Flex and Standard Image templates for virtual servers. The scenario is to use image templates and provisioning scripts to clone an existing configuration and provision multiple servers according to the performance requirements.

There already is an automated implementation for scaling virtual servers in SoftLayer, called Autoscale. An Autoscale group works by using an image template and optionally a provisioning script to add a flexible number of virtual servers behind a local load balancer.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 7-14 Usage scenarios in SoftLayer (2)

Notes:

Custom implementations can be developed using the SoftLayer API. A custom implementation can use bare metal servers, global load balancing, and more advanced monitoring triggers when configuring the autoscaling.

7.17 Checkpoint questions

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint questions

- 1. What is the difference between a Standard Image and a Flex Image?**
 - a) A Standard Image template is exclusive to virtual servers, while a Flex Image is device-neutral. Standard Images are available on all virtual servers.
 - b) A Flex Image is exclusive to virtual servers, while a Standard Image template is device-neutral. Flex Images are available on all virtual servers.
 - c) A Flex Image can be viewed and applied to a new machine by any user, while a Standard Image can only be viewed by authorized users.


© 2015 IBM Corporation



Figure 7-15 Checkpoint questions

Notes:
(none)

7.18 Checkpoint questions (2)

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Checkpoint questions

2. What is the difference between using an HTTP or HTTPS URL for your provisioning script?

- a) Using HTTP protocol, the transfer of the provisioning script during the provisioning process can be controlled by the server. Using HTTPS protocol, the information that is being transferred is sheltered to act contrary to possible security flaws.
- b) Using HTTP protocol results in the provisioning script being downloaded to the device and then be manually executed, if necessary, by a user with administrative access. Using HTTPS protocol results in the provisioning script being downloaded and executed, if possible. If the URL is not associated to an executable script, the script will simply be downloaded and no further action must be taken.
- c) Using HTTPS protocol, the files and information downloaded while the provisioning process are encrypted, so you have to decrypt them to use them. Using HTTP, the provisioning script is already decrypted by the server, so you can use it straightaway.


© 2015 IBM Corporation

Figure 7-16 Checkpoint questions (2)

Notes:

(none)

7.19 Checkpoint questions (3)

IBM Digital Sales International Technical Support Organization and Authoring Services **IBM**

Checkpoint questions

3. What is the Auto Scale option in SoftLayer?

- a) It enables you to scale the options you made while ordering the SoftLayer solution.
- b) It enables you to automate the manual scaling process associated with adding or removing virtual servers.
- c) It automates the algorithm's choice for resource-use triggers to aggregate metrics across assets and time.


© 2015 IBM Corporation 

Figure 7-17 Checkpoint questions (3)

Notes:

(none)

7.20 Checkpoint questions (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions

4. What is the function of a Standard Image template?

- a) It gives users the ability to capture an image from a physical or virtual device and create a machine based on the image, so from a physical to a virtual device image and vice versa.
- b) It allows users to capture an image of an existing virtual server regardless of its operating system and create a new virtual server based on the image.
- c) It allows users to recreate lost server configurations with the saved Standard Image template.

© 2015 IBM Corporation


 **Redbooks Workshop**

Figure 7-18 Checkpoint questions (4)

Notes:

Answers:

1. a)
2. b)
3. b)
4. b)



Storage options in cloud computing: Advanced topics

This unit provides information about these topics:

- ▶ Backup and Recovery Basics
- ▶ SoftLayer Backup and Recovery Offerings
- ▶ Snapshots and replication of Endurance Storage offering

8.1 What you should be able to do

After you complete this unit, you should be able to:

- ▶ Explain backup and recovery basic concepts
- ▶ Describe the SoftLayer backup and recovery offerings
- ▶ Make an educated decision when ordering B&R solutions from SoftLayer
- ▶ Use the snapshots and replication functionality of the Endurance storage offering

8.2 References

The following links are useful for further research:

- ▶ Snapshots:
<http://knowledgelayer.softlayer.com/procedure/endurance-snapshots>
- ▶ Replication:
<http://knowledgelayer.softlayer.com/procedure/endurance-replication>
- ▶ CDP:
<http://wiki.r1soft.com/display/CDP/Documentation>
- ▶ QuantaStor:
<http://knowledgelayer.softlayer.com/learning/quantastor-software-defined-storage>

8.3 Teaching topics

This unit covers the following topics:

- ▶ Backup and recovery concepts
- ▶ Backup and recovery solutions available in SoftLayer
- ▶ Snapshots and replication features of Endurance storage
- ▶ Dedicated storage
- ▶ Why you would need dedicated storage
- ▶ Dedicated storage options available in SoftLayer (QuantaStor)

8.4 Backup and recovery: Concepts

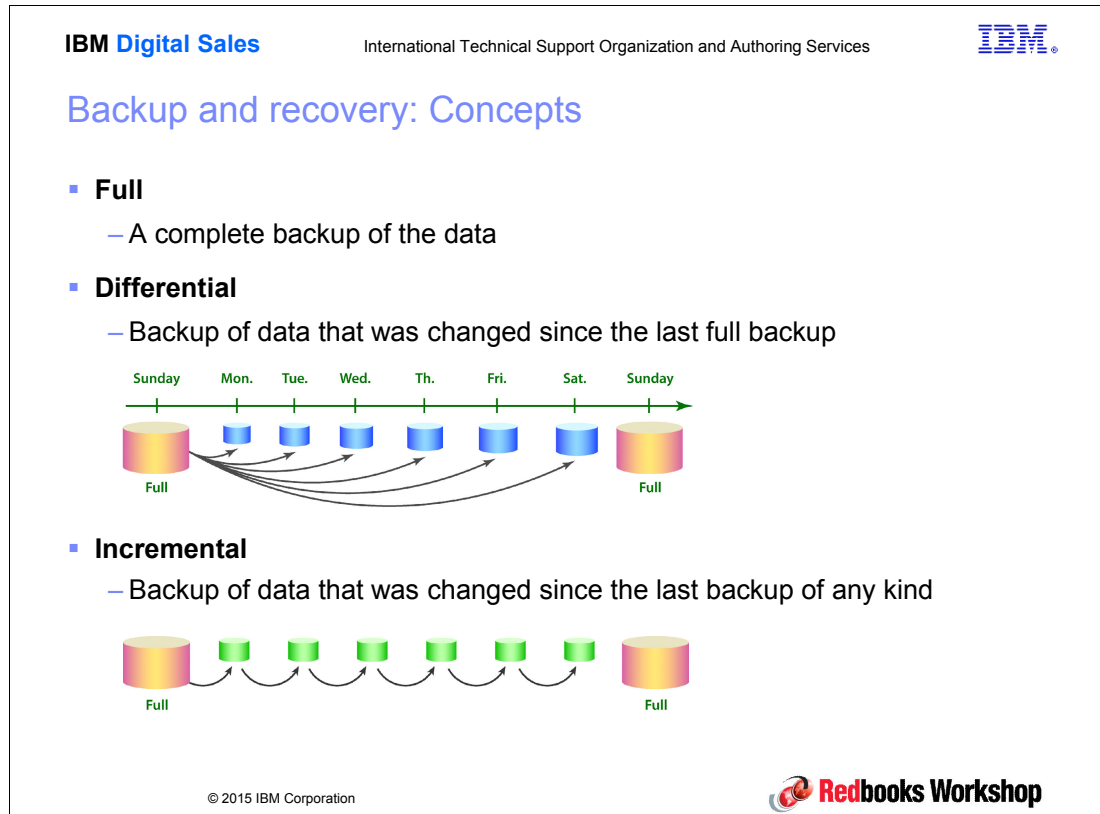


Figure 8-1 Backup and recovery: Concepts

Notes:

In the IT industry, data is the most important asset and must be kept safe. Backup and recovery solutions help you keep it secure. The following are important backup concepts in SoftLayer.

Backup is a process of creating additional copies of the data. If you lose the original version of your data, you can recover your data from the copy.

When working with the Backup and Recovery application, you will often encounter the following terms:

- ▶ **Full backup:** A complete, one to one copy of the existing data.
- ▶ **Differential backup:** Backup of data that was changed since the last full backup. For example, if you did your full backup on Sunday and run a differential backup on Monday, only the data that has changed since Sunday will be backed up. If you will take a backup on Friday, the differential backup backs up the data that was changed between Sunday and Friday.
- ▶ **Incremental backup:** Backup of data that was changed since the last backup of any kind. For example, if you ran a full backup on Sunday and an incremental backup on Monday, an incremental backup on Tuesday will only copy data taht was changed since the Monday backup.

These backup types are often used in combinations. One of the common approaches is to run full backups once per week, and an incremental backup every week day.

8.5 Backup and recovery: Concepts (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Backup and recovery: Concepts

- **Recovery time objective (RTO):** The targeted duration of time and a service level within which a business process must be restored after a disaster (or disruption) to avoid unacceptable consequences associated with a break in business continuity.
- **Recovery point objective (RPO):** The RPO is defined by business continuity planning. It is the maximum targeted period in which data might be lost from an IT service due to a major incident.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 8-2 Backup and recovery: Concepts (2)

Notes:

When thinking about backup strategy you need to consider additional factors, including the following:

- ▶ How fast you will need the data to be recovered in case of failure
- ▶ How old the data restored should be

These factors involve two additional terms:

- ▶ RTO is how fast your data must be recovered after failure.
- ▶ RPO is how old that the data that you restore be so that your business will not be affected. In some cases it can be days, but for some critical databases, like for example billing or banking databases, losing even one hour old data can be critical.

8.6 Backup and recovery solutions in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Backup and recovery solutions in SoftLayer

- **EVault**

EVault Backup is an automated agent-based backup system that is managed by using a centralized web administration console called WebCC. It provides users with a possibility to back up data between servers in one or more data centers on the SoftLayer Network.
- **Idera/R1Soft CDP**

Idera Server Backup provides high-performance disk-to-disk server backup, featuring a central management and data repository. It protects data at block level, and unique disk blocks on the server are stored only once across all recovery points, increasing storage efficiency.
- **Bring your own solution**

In SoftLayer, you are free to install any middleware on your virtual servers and bare metal servers that you wish to. SoftLayer customers can implement their own backup solution on their bare metal servers using software such as IBM TSM, NetBackup, Networker, etc..

© 2015 IBM Corporation


Redbooks Workshop

Figure 8-3 Backup and recovery solutions in SoftLayer

Notes:

EVault Backup is an automated agent-based backup system that is managed through a centralized web administration console called WebCC. It allows you to back up data between servers in one or more data centers on the SoftLayer Network. To use Evault, you complete the following steps:

1. Order an Evault backup volume for the server you want to back up.
2. Install the Evault Client (aka Agent) on your system. Instructions for agent installation are available at <http://knowledgelayer.softlayer.com/topic/evault-backup>.
3. Create a backup job using Evault WebCC.

Another backup solution available from SoftLayer is Idera Server Backup. Idera Server Backup provides high-performance disk-to-disk server backup, featuring a central management and data repository. It protects data at block level, and unique disk blocks on the server are stored only once across all recovery points, increasing storage efficiency.

Idera is ordered as an add-on for a bare metal server. SoftLayer provides the license based on the amount of backup agents that the customer needs. For more information about Idera, see the vendor website at <http://wiki.r1soft.com/display/CDP/Documentation>.

In SoftLayer, you are free to install any middleware on your virtual servers and bare metal servers that you wish to. SoftLayer customers can implement their own backup solution on their bare metal servers using software such as IBM TSM, NetBackup, Networker, etc. This

flexibility allows you to bring your own backup solutions if the solutions provided by SoftLayer do not fit your needs.

Overall, there are plenty of ways you can use to protect your data in SoftLayer. Some other options are covered in the next part of this unit.

8.7 Snapshots and replication of Endurance storage

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Snapshots and replication of Endurance storage

- **Snapshots**

A snapshot represents a volume's contents at a particular point in time. Snapshots enable you to protect your data with no performance impact, minimal consumption of space, and are your first line of defense for data protection. Data can be quickly and easily restored from a snapshot copy if a user accidentally modifies or deletes crucial data from a volume with the snapshot feature.
- **Replication**

Replication uses one of your snapshot schedules to automatically copy snapshots to a destination volume in a remote data center. The copies can be recovered in the remote site in the event of corrupted data or a catastrophic event.

© 2015 IBM Corporation


Redbooks Workshop

Figure 8-4 Snapshots and replication of Endurance storage

Notes:

Snapshot represents a volume's contents at a particular point in time. You can have scheduled snapshots, and in case of need you can rollback to one of previous snapshots. In SoftLayer you can have snapshots scheduled in hourly, daily, and weekly, and can store up to 50 Snapshots. The number of snapshots that you can store also depends on the size of snapshot volume that you order together with Endurance storage.

Snapshot technology can be used as a first line of defense in your Backup and Recover plan. You can easily combine Endurance snapshots with any of external backup solutions. As an extension to the Snapshot functionality, a replication functionality is available for Endurance storage.

Replication uses one of your snapshot schedules to automatically copy snapshots to a destination volume in a remote SoftLayer data center. During a disaster, you can fail over to your replica site and mount the replicated volume to any server in the failover SoftLayer data center.

For more information about snapshots and replication, see the following links:

<http://knowledgelayer.softlayer.com/procedure/endurance-snapshots>

<http://knowledgelayer.softlayer.com/procedure/endurance-replication>

8.8 Dedicated storage

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Dedicated storage

- **Dedicated storage**

Although SoftLayer's storage portfolio covers most business cases, remember that your data is hosted on a shared storage system. Some business cases might require data to be stored on a dedicated storage system.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 8-5 Dedicated storage

Notes:

It is important to understand that having information on a shared storage device does not expose you to any additional risk.

Some examples of cases when you would consider dedicated storage are storing financial, confidential, or personal information. You might also need to consider dedicated storage due to some kind of regulations, such as from a government. Or your company might just want to have better control over their own storage.

For such cases, you can deploy a 2U or 4U server with up to 36 internal drives. Configure it with your choice of hard drives (SATA, SAS, or SSD) and install OS of your choice on it. For example, you can install Linux and share the storage using the NFS or iSCSI protocol. Or you can install Windows and share the disk using CIFS.

SoftLayer provides an option to order an OS NEXUS QuantaStor storage appliance as part of server ordering process.

8.9 Dedicated storage: OS NEXUS QuantaStor

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Dedicated storage: OS NEXUS QuantaStor

- **QuantaStor**

The QuantaStor Storage appliance platform delivers SAN (iSCSI) + NAS (NFS/CIFS) storage on the server hardware of your choice. Designed for IT generalists, QuantaStor appliances configure in minutes and are easy to operate by using an intuitive HTML5 interface.

© 2015 IBM Corporation


Redbooks Workshop

Figure 8-6 Dedicated storage: OS NEXUS QuantaStor

Notes:

QuantaStor is a third-party product developed by OSNexus. It is a customized Linux distribution with an intuitive HTML5 web interface. You can use it to provide iSCSI Volumes, or NFS and CIFS shares. It supports your read and write cache using SSD disks.

For more information about QuantaStor, see:

<http://knowledgelayer.softlayer.com/learning/quantastor-software-defined-storage>

8.10 Check point: Questions

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Check point: Questions

- **True or False:**
Differential backup saves data changed from last full backup?
- **What is Evault WebCC?**
- **True or False:**
You can create daily, weekly, and yearly snapshots for Endurance storage.

© 2015 IBM Corporation


 **Redbooks Workshop**

Figure 8-7 Check point: Questions

Notes:

(none)

8.11 Check point: Questions answered

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Check point: Questions answered

- **True or False:**
Differential backup saves data changed from last full backup?
 - True
- **What is Evault WebCC?**
 - Web Control Center, centralized control panel for Evault.
- **True or False:**
You can create daily, weekly and yearly snapshots for Endurance storage.
 - False: Hourly, daily, and weekly

© 2015 IBM Corporation


Redbooks Workshop

Figure 8-8 Check point: Questions answered

Notes:

(none)



Networking options in cloud computing: Advanced topics

This unit looks in more detail at some advanced capabilities in the SoftLayer network, discusses their purpose, and provides usage examples that can be applied in real life scenarios.

9.1 What you should be able to do

After completing this unit, you should be able to:

- ▶ Describe the network architecture of SoftLayer.
- ▶ Identify the uses of IP addresses in SoftLayer.
- ▶ Explain the use of VLANs and VLAN spanning.
- ▶ Describe the fundamentals of load balancing and load balancing solutions in SoftLayer.
- ▶ Describe the purposes of the SoftLayer Content Delivery Network.
- ▶ Explain the uses of the Vyatta appliance

9.2 References

The following materials are useful for future research:

- ▶ Technical documentation:
<https://developer.ibm.com/marketplace/docs/technical-scenarios/>
- ▶ SoftLayer KnowledgeLayer:
<http://knowledgelayer.softlayer.com/>
- ▶ Content Delivery Network (CDN):
<http://knowledgelayer.softlayer.com/faqs/213#689>
- ▶ Comparison of CDN providers:
<https://www.paessler.com/blog/2010/05/17/monitoring-knowledge/real-world-performance-comparison-of-cdn-content-delivery-network-providers>
- ▶ Hosting a WordPress blog:
<https://developer.ibm.com/marketplace/docs/technical-scenarios/hosting-wordpress-blog-ibm-cloud/>
- ▶ Getting started with the IBM Cloud marketplace:
<https://developer.ibm.com/marketplace/docs/getting-started-2/>
- ▶ SoftLayer CDN:
<http://www.softlayer.com/content-delivery-network>
- ▶ GoDaddy:
<http://www.godaddy.com/>
- ▶ WP Super Cache:
<https://wordpress.org/plugins/wp-super-cache/>

9.3 Teaching topics

This unit covers the following topics:

- ▶ Detailed SoftLayer network architecture
- ▶ IP addresses in SoftLayer
- ▶ VLANs and VLAN spanning
- ▶ Load balancing fundamentals
- ▶ Load balancing solutions in SoftLayer (includes use case)

- ▶ SoftLayer Content Delivery Network (includes use case)
- ▶ Vyatta appliance
- ▶ Recap
- ▶ Checkpoint
- ▶ Networking 101 (optional)
 - Introduction to OSI model
 - Understanding TCP/IP addressing and subnetting basics
 - Netmask quick reference

9.4 SoftLayer network topology

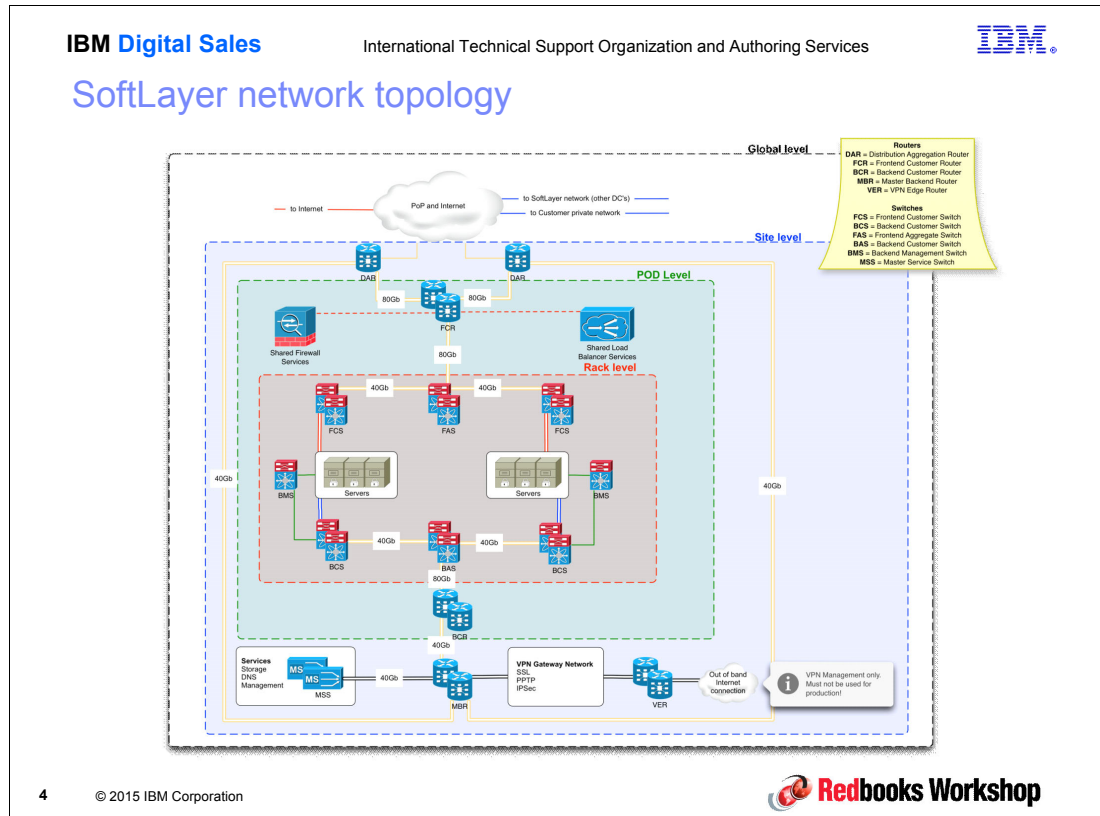


Figure 9-1 SoftLayer network topology

Notes:

This diagram shows the detailed SoftLayer network topology. Softlayer uses a variety of devices in its network topology including, but not limited to Cisco and Juniper network devices, Fortigate security devices, Array Network load balancers, and NetApp storage.

9.5 IP addresses in SoftLayer

IBM Digital Sales International Technical Support Organization and Authoring Services	
IP addresses in SoftLayer	
IP block type	Description
Static IP block	A block of IP addresses that are routed directly to a specific IP on the network.
Portable IP block	<p>Any IP block that can be used on multiple servers within a single VLAN concurrently. Portable IP address are switchable within a VLAN from server to server. There are two types of portable IP blocks:</p> <ul style="list-style-type: none">• Routed to VLAN is a static IP block that is routed to an entire VLAN rather than a specific IP address. Use this method if you have multiple host nodes within the same VLAN. This allows you to migrate a container to a different hardware node within the same VLAN, and not change the IP of the container.• Secondary to VLAN is designed to be used within a virtual environment. It requires that the network, gateway, and broadcast IPs be bound directly to the VLAN rendering these IPs unusable by the customer. This block is used with a virtual machine. To have one usable IP address for a server, you need at least four IP addresses in a block.
Global IP addresses	A Global IP is a static IP address that can be transferred between bare metal servers or virtual servers associated with the account that owns the subnet. Global IPs can be moved to any compatible device on the SoftLayer network.

5 © 2015 IBM Corporation

Redbooks Workshop

Figure 9-2 IP addresses in SoftLayer

Notes:

SoftLayer currently offers two different types of IP blocks: Static and Portable. The different types of IP blocks are designed to be used in different ways. Below is a brief description on each type of block that is offered by SoftLayer as well as a section on using these IP addresses within a virtual machine (VM).

Static IP block

The most popular type of IP block within the SoftLayer network is the Static IP block. A Static IP block is a block of IPs that are routed directly to a specific IP on your network. Every IP address in a Static block is usable on the server. One of the primary benefits of a Static block of IPs is that you do not lose the first two and last IP from the block. Below is an example of a small Static IP block 192.168.0.4/30:

- ▶ 192.168.0.4 - Usable Address
- ▶ 192.168.0.5 - Usable Address
- ▶ 192.168.0.6 - Usable Address
- ▶ 192.168.0.7 - Usable Address

As this example shows, all 4 IPs in this block would be available to the server, while with a portable block, only a single IP from this block would actually be usable on the server due to the network, gateway, and broadcast IPs being bound directly to the VLAN.

Portable IP block

A SoftLayer Portable IP block is considered to be any IP block that can be used on multiple servers within a single VLAN concurrently. SoftLayer currently offers two different types of Portable IP blocks:

- ▶ Routed to VLAN block: A Static IP block that is routed to an entire VLAN rather than a specific IP address.
- ▶ Secondary on VLAN block: Designed to be used within a Virtual Environment.

The primary difference between the two is the number of IPs that are available for use. A Routed to VLAN block, like a static block, provides the user access to all IPs within the block. A Secondary on VLAN block, however, requires that the Network, Gateway and Broadcast IPs be bound directly to the VLAN, rendering them unusable by the user. Use a Routed to VLAN block when you want to use any IP within that block on any server within the VLAN at any time. The Secondary on VLAN block is used with a virtual machine. More information on Secondary on VLAN blocks is provided in the IPs for VMs section.

When ordering a Portable IP block, by default SoftLayer will provide you with a Secondary on VLAN block. If you wish to have this block converted to a Routed to VLAN block for use on your servers within a single VLAN, open a support ticket requesting that it be converted to a Routed to VLAN block.

IPs in PoDs that support HSRP

Remember that PoDs that take advantage of the Hot Standby Router Protocol (HSRP) utilize two more IPv4 addresses (one for the VLAN interface of each participating router) out of every Secondary on VLAN block configured on the VLAN.

Below is an example of a Secondary on VLAN block 192.168.0.4/28 being used for multiple VMs in a HRSP PoD.

- ▶ 192.168.0.0 - Network Address
- ▶ 192.168.0.1 - Gateway Address
- ▶ 192.168.0.2 - Router A VLAN Interface
- ▶ 192.168.0.3 - Router B VLAN Interface
- ▶ 192.168.0.4 - VPS1
- ▶ 192.168.0.5 - VPS2
- ▶ 192.168.0.6 - VPS3
- ▶ 192.168.0.8 - VPS4
- ▶ 192.168.0.9 - VPS5
- ▶ 192.168.0.10 - VPS6
- ▶ 192.168.0.11 - VPS7
- ▶ 192.168.0.12 - VPS8
- ▶ 192.168.0.13 - VPS9
- ▶ 192.168.0.14 - VPS10
- ▶ 192.168.0.15 - Broadcast Address

IPs for virtual machines

Private Clouds are becoming more popular every day. This section covers what type of IP blocks are required to be used in a VM to implement hypervisors which are not managed by SoftLayer, such as Citrix XenServer, VMware, Microsoft Hyper-V and KVM. You may provision a SoftLayer bare metal server with XenServer, VMware or Microsoft Hyper-V using the SoftLayer customer portal. For other hypervisors such as KVM, you should provision a bare metal server with No Operating System and then load your own hypervisor. Remember you are responsible for managing hypervisors which you implement on bare metal servers.

The example provided below is based on Microsoft Hyper-V.

Every VM connected to the SoftLayer network in a virtual environment requires a primary IP address from a portable block of IPs. The reason is that Hyper-V requires each VM to provide a network, gateway, and broadcast address on the same subnet as the primary IP assigned to the VM. One advantage to this network configuration is that a single Secondary on VLAN block can be used for multiple VMs. The following is an example of a Secondary on VLAN block 192.168.0.4/29 being used for multiple VMs.

- ▶ 192.168.0.0 - Network Address
- ▶ 192.168.0.1 - Gateway Address
- ▶ 192.168.0.2 - VPS1
- ▶ 192.168.0.3 - VPS1
- ▶ 192.168.0.4 - VPS1
- ▶ 192.168.0.5 - VPS2
- ▶ 192.168.0.6 - VPS3
- ▶ 192.168.0.7 - Broadcast Address

As the example shows, this Secondary on VLAN block provides five usable IP address out of the eight IP addresses in the block bound across three different VMs. If you want to add more IPs to a VM when all the IPs on the Portable block are used, use a Static block, or a Routed to VLAN Portable block.

To use a Static Block within a VM, first order a new Static IP block from the portal. When you order this block you will be able to select the IP address that you want this block to be routed to. By selecting the IP address that is assigned to the VM, the new block is routed specifically to that VM. You can then bind the new block of IPs directly to that VM and begin using them immediately.

Alternately, if you wish for the new block to be usable by more than one VM, use a Routed to VLAN block. A Routed to VLAN block is available by purchasing a Portable IP block from the portal and selecting the VLAN where the IP address of the VM is. After the IP block is created, it is then available for use on any Server or VM on that VLAN.

Each SoftLayer server (virtual or bare metal) comes with one primary IPv4 address. Additional IP blocks are available with quantities of 1, 2, 4, 8, 16, or 32 IP addresses..

Private IP Addressing is defined by SoftLayer (10.x.x.x).

Global IP addresses

Global IPs provide IP flexibility by allowing users to shift workloads between servers, even ones in different data centers. Global IPs also provide IP persistence by allowing for transitions between servers and VSIs (Virtual Server Instance), such as upgrading from a VSI to a dedicated system without having your IP tied to a particular server or VLAN.

9.6 IP addresses in SoftLayer (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

IP addresses in SoftLayer

Subnet	Description
Primary subnets	A primary subnet is the primary network bound to a VLAN within the SoftLayer network (Each VLAN that SoftLayer creates comes with a primary subnet).
Portable subnets	A portable subnet is a secondary network bound to a VLAN within the SoftLayer network. It will not be touched by SoftLayer systems (such as when you are building your private cloud).
Static subnets	A static subnet is a network bound to a single server. Use static subnets if you have one hardware node per VLAN or in total.
Customer subnets	A subnet that is owned by customer and is put in to use and routed in SoftLayer.

Subnets

[View VLANs](#) | [Order IP Addresses](#)

Filter

Viewing 1 to 13 of 13 | All | **Portable Subnets** | Displaying 25 per page | Page 1 of 1

Subnet	Network	Type	Location	Target	Total IP	Notes (Click cell to edit)
50.23.4.16/...	Pu...	Primary Subnets	Houst...		16	
10.104.0.18...	Pri...	Portable Subnets	Amst...		4	
10.70.210.1...	Pri...	Static Subnets	Amst...		64	
10.77.142.0...	Pri...	Customer Subnets	Houst...	hou02.bcr0...	64	
10.77.182.0...	Pri...	All Subnets	Houst...	hou02.bcr0...	64	

6 | © 2015 IBM Corporation




Figure 9-3 IP addresses in SoftLayer (2)

Notes:

Click **Network** → **IP Management** → **Subnets** to find the four types of subnets configured in your SoftLayer account by ordering IP blocks (which are the equivalent to subnet in SoftLayer terminology). Some new subnet names appear as explained in the table.

9.7 Five steps to start using IPv6 in SoftLayer


IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Five steps to start using IPv6 in SoftLayer

1. Obtain a block of public IPs from SoftLayer's portal.
2. Install IPv6 in your system.
3. Bind your new IPv6 to the application.
4. Add IPv6 addresses to DNS.
5. Make an IPv6 accessibility test.



7

© 2015 IBM Corporation


 Redbooks Workshop

Figure 9-4 Five steps to start using IPv6 in SoftLayer

Notes:

Each type of address (Static, Portable, and Global) can be ordered in SoftLayer either in version 4 or 6. When you connect to the Internet, your device (computer, smartphone, tablet) is assigned an IP address, and any site you visit has an IP address. The IP addressing system that has been used since the beginning of the Internet is called IPv4, and the new addressing system is called IPv6. IPv6 was introduced because the Internet is running out of available IPv4 address space, and IPv6 provides an exponentially larger pool of IP addresses:

- Total IPv4 Space: 4,294,967,296 addresses
- Total IPv6 Space: 340,282,366,920,938,463,463,374,607,431,768,211,456 addresses

An IPv4 address is based on 32 bits, while IPv6 is based on 128 bits. Example of IPv6 address:

2607:f0d0:4545:3:200:f8ff:fe21:67cf

Fortunately, the SoftLayer platform is IPv6 ready, and is already issuing and routing IPv6 traffic. Obtaining a block of public IPs from SoftLayer is as easy as logging into the portal, pulling up the hardware page of a server, and ordering a /64 block of IPv6 IPs.

In addition, most current server operating systems are ready to change to IPv6. This includes Windows 2003 SP1 and most Linux operating systems with 2.6.x Linux kernels. This discussion focuses on Windows and RedHat/CentOS.

To ready your Windows 2003 server for IPv6, complete these steps:

1. In Control Panel, double-click **Network Connections**.
2. Right-click any local area connection, and then click **Properties**.
3. Click **Install**.
4. In the Select Network Component Type window, click **Protocol** and then **Add**.
5. In the Select Network Protocol window, select **Microsoft TCP/IP version 6** and then click **OK**.
6. Click **Close** to save changes to your network connection.

After IPv6 is installed, IIS will automatically support IPv6 on your web server. If a website was running when you installed the IPv6 stack, you must restart the IIS service before the site begins to listen for IPv6 requests. Sites that you create after you enable IPv6 automatically listen for IPv6. Windows 2008 server should have IPv6 enabled by default.

When your Windows server is ready for IPv6, add IPv6 addresses to the server just as you add IPv4 addresses. The only difference is that you edit the properties to the Internet Protocol Version 6 (TCP/IPv6) network protocol.

Now that you have more IPv6 addresses for your servers than what's available to the entire world in IPv4 space, you must bind them to IIS or Apache. This is done the similarly to the way you bind IPv4 addresses.

Add your new IPv6 addresses to your DNS server. If you are using a IPv6-enabled DNS server, simply insert an 'AAAA' resource record (aka quad-A record) for your host.

While your DNS is propagating, test your web server to see if it responds to the IP that you assigned by using square brackets in your browser:
`http://[2101:db8::a00:200f:fda7:00ea]`. This test only works if your computer is on a IPv6 network. If you are limited to IPv4, you will need sign up with a tunnel broker or switch to an ISP that offers IPv6 connectivity.

After about 24 hours, your server and new host should be ready to serve websites on the IPv6 stack.

9.8 Separating devices and subnets with VLANs

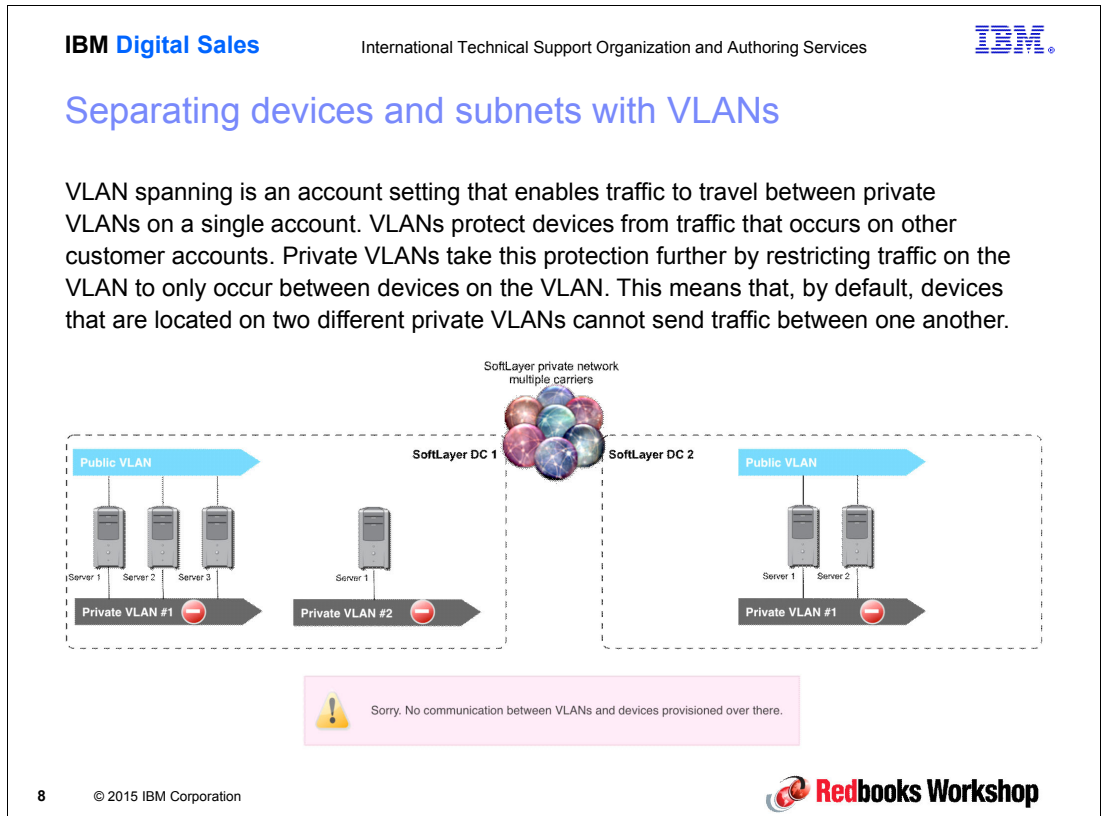


Figure 9-5 Separating devices and subnets with VLANs

Notes:
(none)

9.9 VLAN spanning

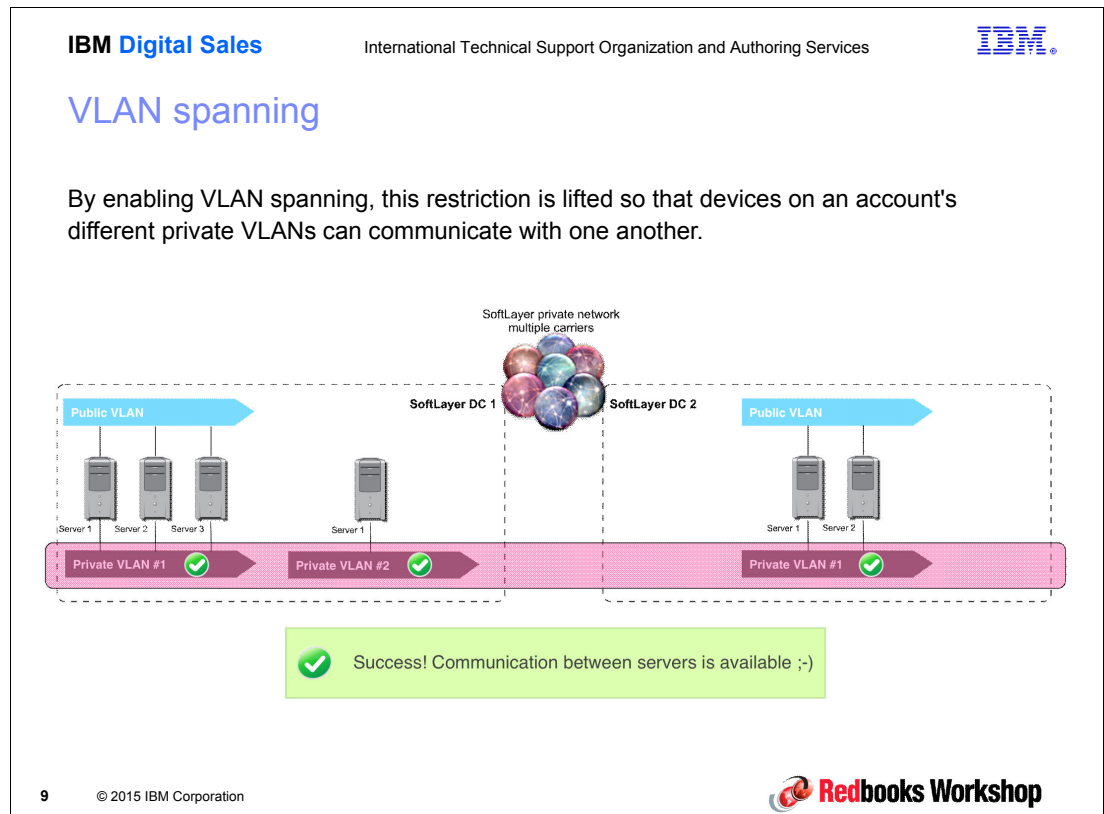


Figure 9-6 VLAN spanning

Notes:
(none)

9.10 VLAN spanning (2)

IBM Digital Sales

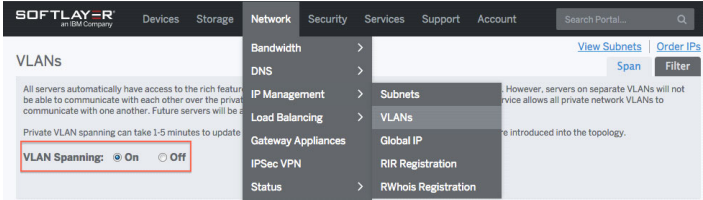
International Technical Support Organization and Authoring Services

IBM

VLAN spanning

To enable or disable VLAN spanning, complete these steps:

1. Access the VLANs screen in the Customer Portal. Refer to Access the VLANs window.
2. Click the Span tab to access the VLAN Spanning window.
3. Click the **On** radio button to enable VLAN spanning. Click the **Off** button to disable VLAN spanning.



After updating VLAN spanning selections, the request can take up to 15 minutes to process. A confirmation of the change will briefly appear below the Span tab. If enabling VLAN spanning, devices will be able to communicate with one another across VLANs using the private network after the update has been processed. If disabling spanning, devices will only be able to connect to one another if they reside in the same VLAN. Cross-VLAN communication will no longer be possible. VLAN spanning settings can be updated at any time by repeating these steps. Toggling between VLAN spanning settings in a short amount of time might result in a delay of settings being applied.

10

© 2015 IBM Corporation




Figure 9-7 VLAN spanning (2)

Notes:
(none)

9.11 Load balancing fundamentals

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Load balancing fundamentals

A load balancer is a device that acts as a reverse proxy and distributes network or application traffic across a number of servers. Load balancers are used to increase capacity (concurrent users) and reliability of applications. They improve the overall performance of applications by decreasing the burden on servers associated with managing and maintaining application and network sessions, as well as by performing application-specific tasks.

Problems

Solution

- Too much traffic to a single web server causes the site go offline.
- Having only one web server is a single point of failure.

Replicate web site content to multiple servers (service pool) and use a network device (LB) in front of these servers to impersonate the address of the website and redirect page requests only to healthy members of the service pool.

11 © 2 hop

Figure 9-8 Load balancing fundamentals

Notes:
(none)

9.12 Load balancing fundamentals (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Load balancing fundamentals

Load balancers are generally grouped into two categories: Layer 4 and Layer 7. Layer 4 load balancers act on data found in network and transport layer protocols (IP, TCP, FTP, UDP). Layer 7 load balancers distribute requests based on data found in application layer protocols such as HTTP.

Requests are received by both types of load balancers and distributed to servers based on a configured algorithm. The following are some industry standard algorithms:

- Round robin
- Weighted round robin
- Least connections
- Least response time

Layer 7 load balancers can further distribute requests based on application specific data such as HTTP headers, cookies, or data within the application message itself, such as the value of a specific parameter.

Load balancers ensure reliability and availability by monitoring the "health" of applications and only sending requests to servers and applications that can respond in a timely manner.

12 © 2015 IBM Corporation


 **Redbooks Workshop**

Figure 9-9 Load balancing fundamentals (2)

Notes:

The following are standard load balancing algorithms:

- ▶ Round robin: One of the simplest methods for distributing client requests across a group of servers. Going down the list of servers in the group, the round-robin load balancer forwards a client request to each server in turn. When it reaches the end of the list, the load balancer loops back and goes down the list again.
- ▶ Weighted round robin: A weight is assigned to each server based on criteria chosen by the site administrator. The most commonly used criterion is the server's traffic-handling capacity. The higher the weight, the larger the proportion of client requests the server receives. If, for example, server A is assigned a weight of 3 and server B a weight of 1, the load balancer forwards three requests to server A for each one it sends to server B.
- ▶ Least connections: Load Balancer passes a new connection to the pool member or node that has the least number of active connections, for example HTTP connections (this is a default method in Citrix NetScaler VPX).
- ▶ Least response time: When Load Balancer is configured to use the least response time method, it selects the service with the least number of active connections and the least average response time.




9.13 Load balancing options in SoftLayer

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Load balancing solutions in SoftLayer

SoftLayer Data Centers		Local load balancing (VIP) SoftLayer Local Load Balancing uses industry-standard techniques including round-robin, lowest latency, least connections, shortest response, and IP persistence to balance traffic among two or more servers. It can be activated and configured in real-time, with servers added to or removed from the balancing pool on-demand with little or no downtime.
		High availability dedicated load balancers Dedicated load balancers are available for environments with higher capacity. These load balancers are also available in a high availability option featuring failover protection and automatic fallback.
Multiple Data Centers		Local and global load balancing with Citrix NetScaler VPX Distribute traffic between your servers in one or multiple SoftLayer data centers with Citrix NetScaler VPX. These multifunction network appliances can perform DNS-based local and global load balancing to give you complete control over how your client traffic will be balanced between your servers.

13

© 2015 IBM Corporation




Figure 9-10 Load balancing options in SoftLayer

Notes:
(none)

9.14 Load balancing options in SoftLayer (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Load balancing solutions in SoftLayer

Local Load Balancing	Dedicated Local Load Balancer with SSL - 2600	Dedicated Local Load Balancer with SSL - 6600	Citrix NetScaler VPX
<ul style="list-style-type: none"> ▪ Multitenant ▪ Proxy: Passes traffic through LB ▪ Single VIP ▪ OSI Layer 4 ▪ SSL offloading optional ▪ Good fit for 2500 connections / sec 	<ul style="list-style-type: none"> ▪ Singletenant ▪ Proxy ▪ One VIP assigned, but can have a few more ▪ HA option ▪ OSI Layer 4 ▪ Up to 15,000 connections / sec ▪ 1 + VIP 	<ul style="list-style-type: none"> ▪ Singletenant ▪ Proxy ▪ One VIP assigned, but can have a few more ▪ HA option ▪ OSI Layer 4 ▪ 150,000 connections / sec ▪ 1 + VIP 	<ul style="list-style-type: none"> ▪ Singletenant ▪ GUI/CLI interface ▪ Private and Public network ▪ HA option ▪ OSI Layer 4 and Layer 7 ▪ Can use static and portable IPs
<ul style="list-style-type: none"> ▪ Shared ▪ Low volumes 		<ul style="list-style-type: none"> ▪ Dedicated ▪ High volumes 	<ul style="list-style-type: none"> ▪ Global and local load balancing ▪ Dedicated ▪ High volumes ▪ Layer 4 and Layer 7 balancing

14

© 2015 IBM Corporation


 Redbooks Workshop

Figure 9-11 Load balancing options in SoftLayer (2)

Notes:
(none)

9.15 Load balancing solutions in SoftLayer use case

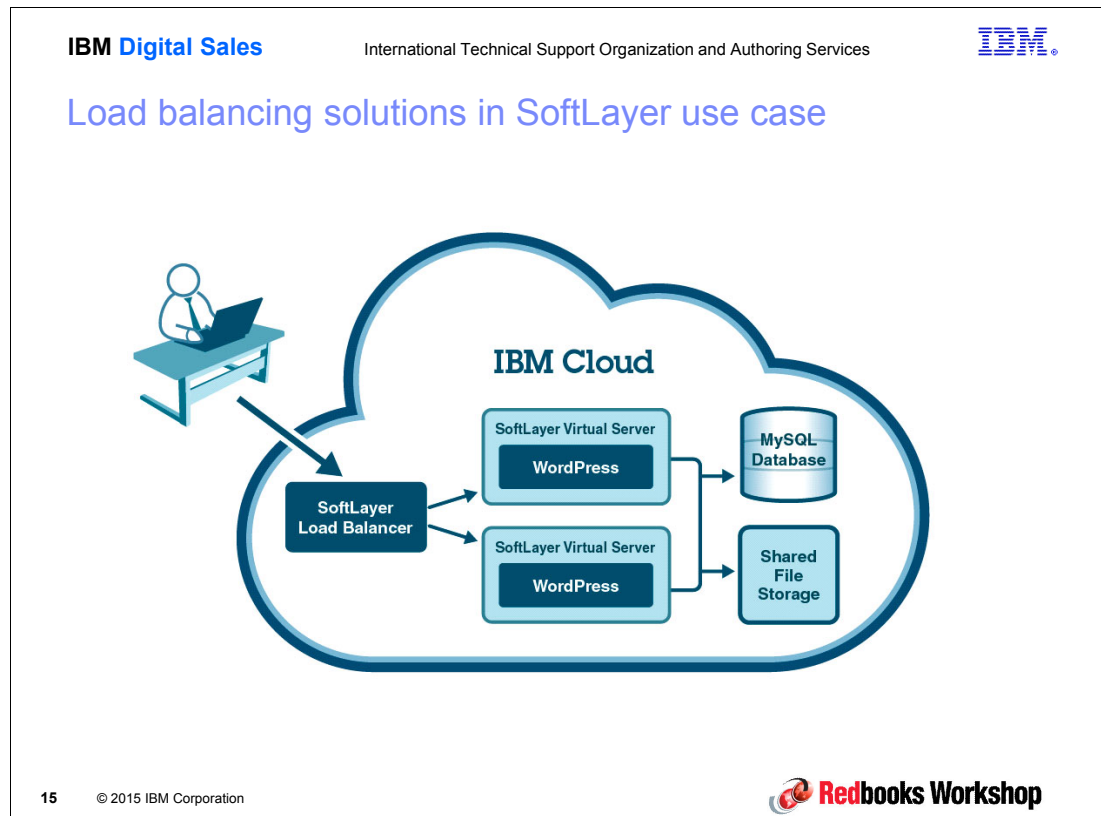


Figure 9-12 Load balancing solutions in SoftLayer use case

Notes:

This diagram illustrates an example of load balancing. Here the load balancing solution is created for Wordpress web application. Two Wordpress servers, which connect to a single, shared MySQL database server and shared file storage, use SoftLayer Load Balancing to distribute requests in the IBM Cloud.

Here Wordpress is just an example, but the general concepts can also be used to load balance other applications that are stateless and use a common, shared data store.

9.16 SoftLayer Content Delivery Network

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network

The SoftLayer Content Delivery Network (CDN) uses EdgeCast, and includes robust tools for digital rights management and content monetization.

Distribution options	
Origin pull	The first time content is request, it is pulled from the host server to the network and stays there for other users to access it.
PoP pull	Customers pre-load content using various methods. The loaded content is pulled from the CDN FTP as opposed to being pulled from the customer's origin location.

Only complete network cache updates are supported. Individual nodes cannot be cleared or deleted.

16

© 2015 IBM Corporation


Redbooks Workshop

Figure 9-13 SoftLayer Content Delivery Network

Notes:

CloudLayer® CDN helps you deliver content to end users faster and more efficiently through a network of 24 cloud-connected nodes running advanced organizing, storing, and streaming software. Rather than serving content directly from your host server, your content is served from a node that is geographically closer to your user, minimizing the distance the data has to travel and thereby avoiding network traffic jams and decreasing latency.

CloudLayer CDN is available in two different content delivery options, letting you optimize your end-user experience and cost:

- **Origin Pull:** Store your content on a SoftLayer server or computing instance, or on another location on the Internet. The content's location is registered with the CDN. When the first user requests the content, it is pulled to the network and delivered from the closest point to that user. The content does not have to be manually uploaded to the CDN. The content then remains on the network for other users to access it. Using Origin Pull requires a bandwidth plan, but no storage charges are incurred.
- **PoP Pull:** Customers can pre-load content by using the EdgeCast API, SoftLayer API, SoftLayer Portal, or the customer's own portal. That content is pulled from the CDN FTP as opposed to being pulled from the customer's origin location. Pay As You Go or monthly pricing is based on bandwidth and storage used, with mix and match rate plans available.

CloudLayer CDN includes robust tools for digital rights management and content monetization, with SoftLayer's renowned ease-of-use and unparalleled level of control. In addition, it can be seamlessly integrated with your SoftLayer dedicated servers, virtual

servers, automated services, and additional CloudLayer services, to create a unified computing solution with unequalled efficiency and interoperability.

Key features and benefits:

- ▶ **Pay As You Go Pricing:** Monthly rates with no long-term commitment let you optimize your cost and performance balance and return on technology investment. Pay only for the CDN performance that you need and use.
- ▶ **Dedicated + Virtual + Cloud Integration:** Seamless integration with your dedicated, virtual, and additional CloudLayer services through SoftLayer's leading private network provides a single computing environment with interoperability and efficiency found nowhere else.
- ▶ **Immediate Scalability:** Add CDN services in real time to accommodate traffic spikes or snags, gaining the level of content accessibility that you need, exactly when you need it.
- ▶ **Ease of Use and Control:** Full access and control by using the SoftLayer Customer Portal and API provides streamlined, simplified management of your solution and a direct connection to all SoftLayer services and tools.
- ▶ **Performance-Neutral Growth:** Spreading content delivery over multiple servers throughout the cloud ensures that increased demand for content does not slow down or compromise delivery.
- ▶ **Secure Content Management:** Secure streamlined tools for managing content and content monetization protects digital rights and maximizes return.
- ▶ **Broader Geographic Reach:** Pushing content to nodes around the world optimizes the speed and reliability of content delivery to users regardless of location.
- ▶ **Higher Quality Content and Rich-Media Websites:** Optimizing your data and content delivery allows you to offer more robust web experiences and richer content without sacrificing performance.
- ▶ **Supported media formats:** Windows Media, Adobe Flash, H.264 over HTML5, Microsoft Silverlight, Smooth Streaming, and HTTP Progressive.
- ▶ **Supported delivery formats:** HTTP, real-time messaging protocol (RTMP), and real-time streaming protocol (Windows Media Services) [RTSP (WMS)].
- ▶ **Supported security methods:** SWF authentication, token authentication, and RTMP encryption.

For more information go to <http://knowledgeclayer.softlayer.com/faqs/213#689>

9.17 EdgeCast CDN locations

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

EdgeCast CDN locations

The EdgeCast CDN provides 24 content delivery nodes around the world, in addition to the SoftLayer global footprint that includes data centers and PoP.

● Content Delivery Node

17 © 2015 IBM Corporation

Redbooks Workshop

Figure 9-14 EdgeCast CDN locations

Notes:
(none)

9.18 Example SoftLayer CDN users

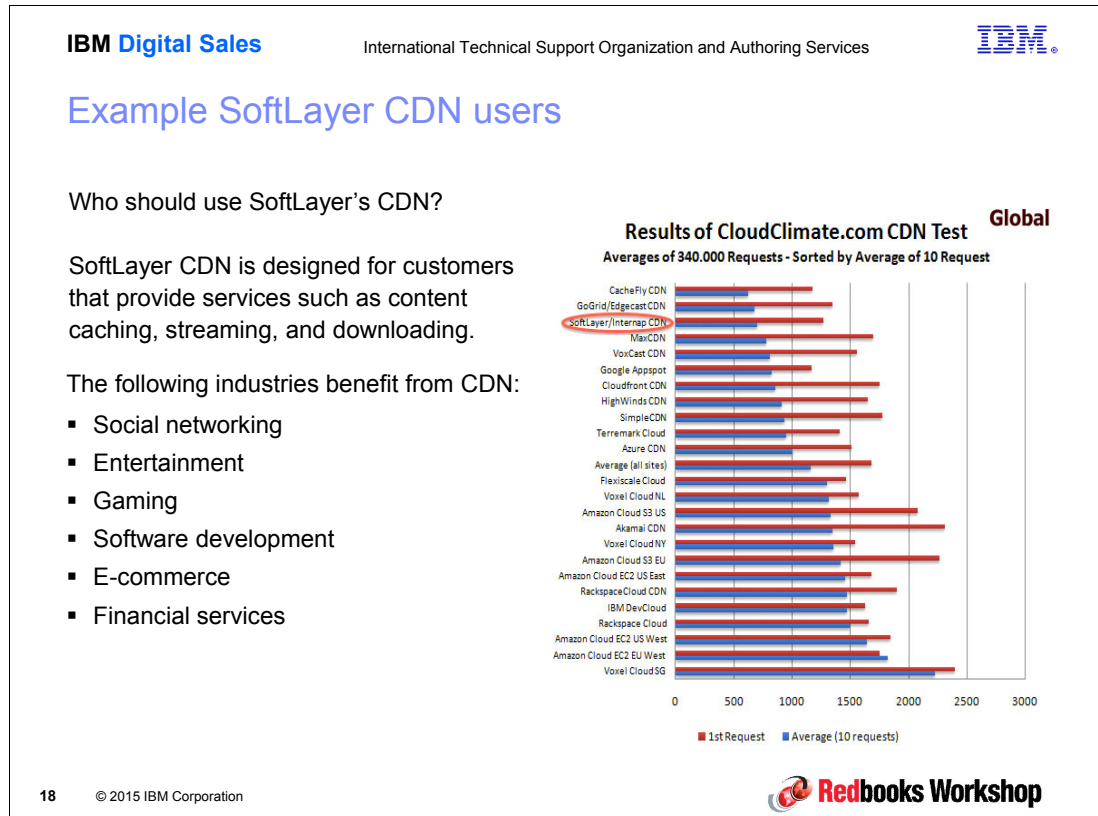


Figure 9-15 Example SoftLayer CDN users

Notes:

For some period of time, visitors of the website <http://www.cloudclimate.com> could measure the “website asset delivery speed” of 24 CDN and cloud providers using their own internet connection. The results for each user were stored in cloudclimate.com database and they gathered the performance results of 340,000 requests. This test provides “real surfer” measurements for the 24 providers because they had actually used the “real world” Internet connections of cloudclimate.com website visitors to run the test requests. SoftLayer CDN was one of the best in test. For more information, visit:

<https://www.paessler.com/blog/2010/05/17/monitoring-knowledge/real-world-performance-comparison-of-cdn-content-delivery-network-providers>

9.19 SoftLayer Content Delivery Network use case

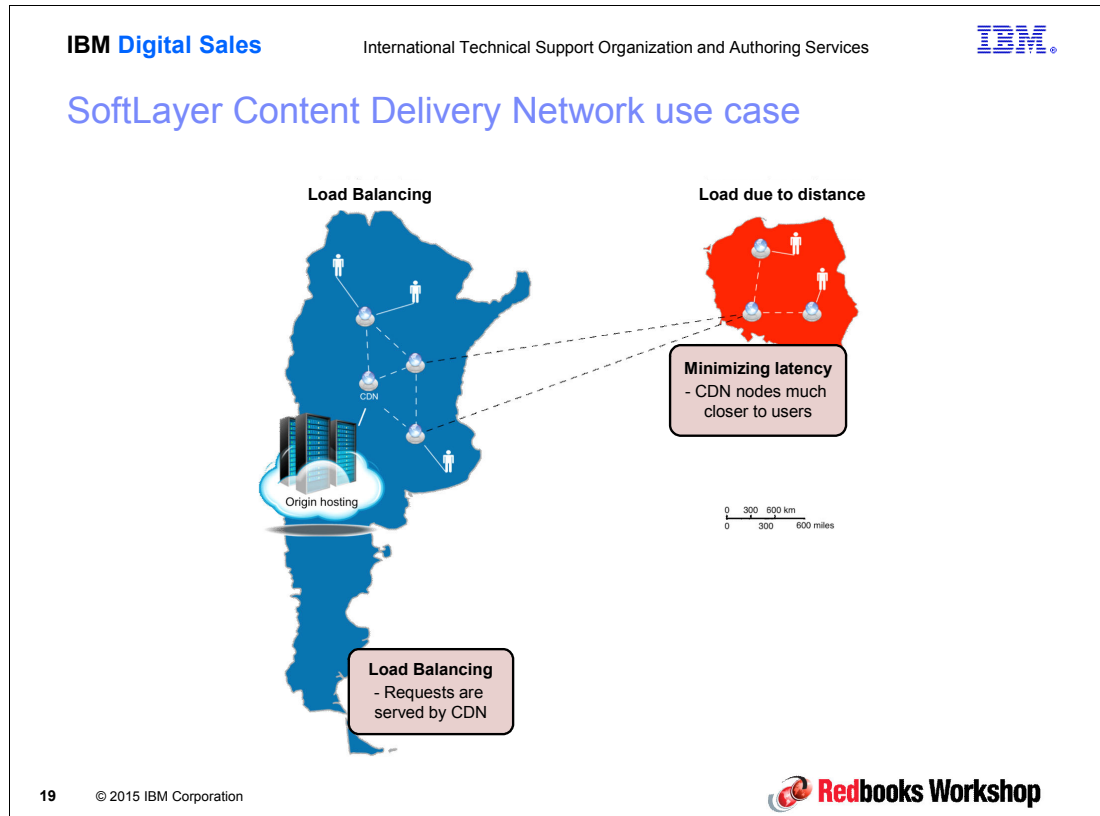


Figure 9-16 SoftLayer Content Delivery Network use case

Notes:

CDNs help get your content closer to your customers reducing the distance AND TIME it takes for your customers to access your content. Used the right way CDNs have a positive impact on your business. Your use of CDNs can start out deployed locally within a country, continent or hemisphere and easily expand around the globe with the expansion of your clientele.

These are the two typical use cases for Content Delivery Network:

- ▶ Reducing Latency (DNS resolution, connection speed, 1st byte time): In a typical, internationally operating B2B manufacturing company or an international financial institution, the main CDN benefit is getting the content much closer to the users spread across the globe, which minimizes the latency and increases performance. Factors like decreased DNS, connection, and first byte times are their main motivators.
- ▶ Externalizing Load Balancing (hit/miss ratio): Another scenario is the heavy traffic sites where the benefit is to outsource the handling of the load and thus minimize the number of requests on their own infrastructure. Therefore, the hit/miss ratio is an important factor when judging the efficiency of the CDN. Hit in this case means the content was served straight from a CDN PoP, and miss means the content had to be fetched by the CDN before it could served it to the user. A miss therefore affects the data center. Such a traffic reduction on the origin can also be used to minimize the effect of denial of service attacks.

9.20 SoftLayer Content Delivery Network implementation guide

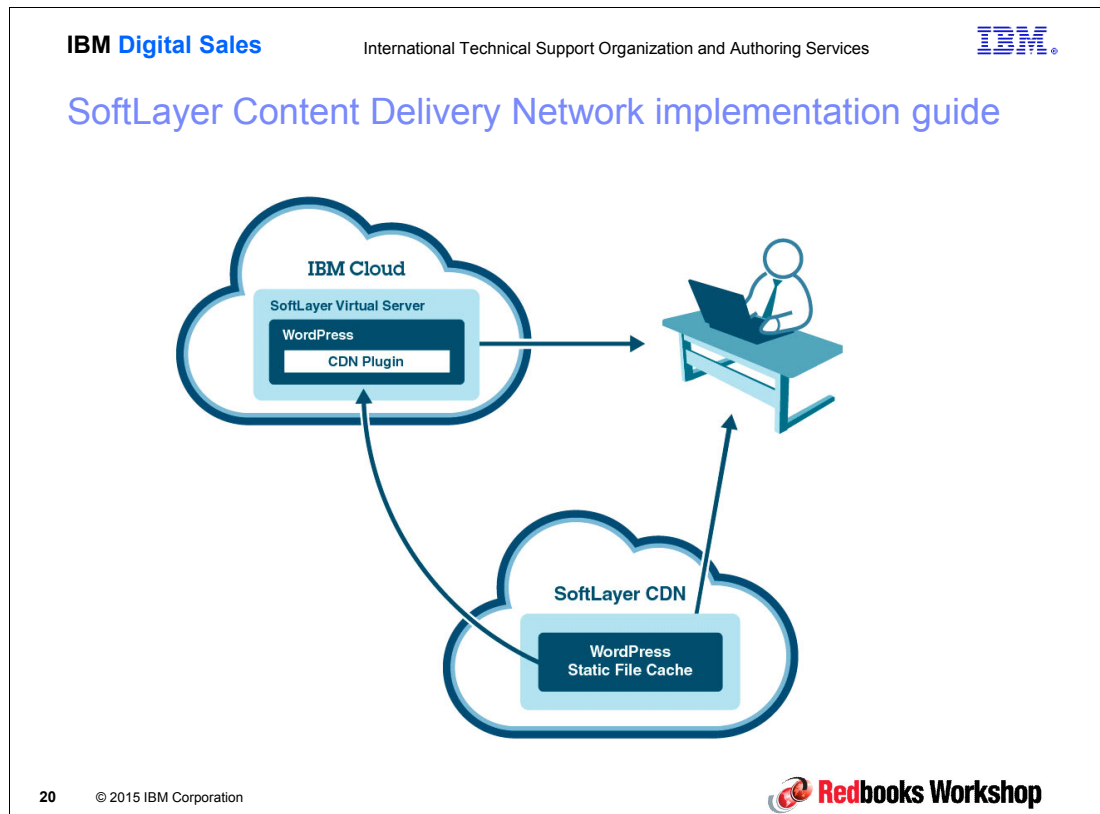


Figure 9-17 SoftLayer Content Delivery Network implementation guide

Notes:

The following diagram illustrates the solution that was created in this example. A WordPress, which is deployed in the IBM Cloud, uses a SoftLayer CDN to deliver its static content. There are two types of stakeholders: Users viewing the content and the SoftLayer customers who run the WordPress site. In this example, “user” is used for the first type and “customer” for the second type. Below you can find out the step by step instructions how to implement CDN for WordPress in SoftLayer.

Step 0: Prepare your environment

1. If you don't have a SoftLayer account already, obtain a SoftLayer account.
2. Ensure that you have permissions to manage the CDN account and CDN File Transfers for your SoftLayer account.
3. If you do not have an IBM ID already, obtain an IBM ID and register for the IBM Cloud marketplace.

Step 1: Set up your WordPress in the IBM Cloud

Follow the steps in the Hosting your WordPress blog in the IBM Cloud to deploy a WordPress blog in the IBM Cloud. For more information, see this website:

<https://developer.ibm.com/marketplace/docs/technical-scenarios/hosting-wordpress-blog-ibm-cloud/>

9.21 SoftLayer Content Delivery Network use case

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

Order Content Delivery Network

Quantity

1 Content Delivery Network

Content Delivery Network Configuration

CDN BANDWIDTH	MONTHLY	SETUP	USAGE
<input checked="" type="checkbox"/> CDN Pay as You Go Bandwidth	\$0.00	\$0.00	\$0.12/GB

Show CDN Bandwidth

Configuration Summary

CDN Bandwidth \$0.00

Standard Options

CDN Storage \$0.00

Subtotals

Quantity: 1

Total (per server): \$0.00

Monthly \$0.00

Setup Total: \$0.00

Continue Your Order

Shopping Cart

You have 0 items in your cart.

[View/Edit shopping cart](#)

21 © 2015 IBM Corporation




Figure 9-18 SoftLayer Content Delivery Network use case

Notes:

Step 2: Order the SoftLayer CDN

There are two ways that a SoftLayer CDN can distribute your content: Origin Pull and PoP Pull.

Typically, sites with heavy traffic loads will benefit from an Origin Pull CDN, because the content is pulled from the host server and users can pull the cache content from the CDN. Alternatively, customers can benefit from a PoP Pull CDN by controlling the content that gets uploaded to a CDN FTP site when it expires. Regardless of the type of CDN chosen, deploying a CDN means getting content to users faster and more efficiently. For WordPress in this article, use an Origin Pull CDN.

Log in to the SoftLayer customer portal and add a CDN service by clicking **Sales** → **Add CDN Service**. There is no configuration at this point, so just place the order.

9.22 SoftLayer Content Delivery Network use case (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

ADMINISTRATIVE

SUPPORT

SECURITY

BARE METAL

SOFTWARE

VIRTUAL

PUBLIC NETWORK

PRIVATE NETWORK

SALES

Virtual Server: **CDN**

CDN Accounts

Account Name	Solution	Status	Note	Created	Action
B06A	Pay as You Go Bandwidth [No storage (origin pull)]	Active		07/29/2013	More
10877	Pay as You Go Bandwidth [No storage (origin pull)]	Origin Pull	zjm	09/02/2014	More
10D23	Pay as You Go Bandwidth [No storage (origin pull)]	Active		09/15/2014	More
10D24	Pay as You Go Bandwidth [No storage (origin pull)]	Active		09/15/2014	More

Order CDN

Note. Double-click on a CDN Note column in order to edit. Hit "enter" key to save and "ESC" to cancel. The maximum length for CDN account note is 30 characters.

CDN Description

How can I get assistance with CDN?

22

© 2015 IBM Corporation

Redbooks Workshop

Figure 9-19 SoftLayer Content Delivery Network use case (2)

Notes:

Step 3: Get your CDN account name

After the order is processed, the new CDN account is displayed in the SoftLayer customer portal (**Public Network** → **Content Delivery Network**). Remember the account name of the newly created CDN account because it will be used in the next step. In this example, the account name is 10D24.

9.23 SoftLayer Content Delivery Network use case (3)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

SettingsDNS Zone FileContacts

Zone File ⓘ
Last updated 9/24/2014 7:32:26 PM MST7 records in this zone

Add RecordDeleteBulk Actions TemplatesMore

Filter List

A (Host) ⓘ
1 Records (0 Selected)

✓	Host	Points To	TTL	Actions
<input type="checkbox"/>	blog	119.81.143.66	1/2 Hour	

AAAA (IPv6 Host) ⓘ
0 Records (0 Selected)

Add Record

CName (Alias) ⓘ
1 Records (0 Selected)

Restore Defaults

✓	Host	Points To	TTL	Actions
<input type="checkbox"/>	cdn	10d24.http.cdn.softlayer.net	1/2 Hour	

23 © 2015 IBM Corporation

Figure 9-20 SoftLayer Content Delivery Network use case (3)

Notes:

Step 4: Update your DNS zone

In addition to your WordPress site domain, you also need a custom domain for your CDN content. This example uses `blog.playvm.com` for the WordPress site domain, and uses `cdn.playvm.com` for the CDN domain. The demo domain name is `playvm.com`, which is registered in GoDaddy. Your CDN domain must be set as a CNAME record that points to your SoftLayer CDN account domain. The SoftLayer CDN account domain is your CDN account name with the suffix of `.http.cdn.softlayer.net`. In this example, the SoftLayer CDN account domain is `10D24.http.cdn.softlayer.net`.

Add the CDN CNAME record to your domain DNS zone in your domain DNS provider. This example uses GoDaddy as the domain provider. Note that the host `blog` is mapped to the IP address `119.81.143.66` and that the `cdn` hostname is mapped to the account domain assigned for your SoftLayer CDN, which is `10d24.http.cdn.softlayer.net`.

9.24 SoftLayer Content Delivery Network use case (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

ADMINISTRATIVE | SUPPORT | SECURITY | BARE METAL | SOFTWARE | VIRTUAL | PUBLIC NETWORK | PRIVATE NETWORK | SALES

Virtual Server: CDN

CDN Accounts

Account Name	Solution	Status	Note	Created	Action
B06A	Pay as You Go Bandwidth [No storage (origin pull)]	Active		07/29/2013	More
10877	Pay as You Go Bandwidth [No storage (origin pull)]	Origin Pull	zjm	09/02/2014	More
10D23	Pay as You Go Bandwidth [No storage (origin pull)]	Active		09/15/2014	More
10D24	Pay as You Go Bandwidth [No storage (origin pull)]	Active		09/15/2014	More

[Order CDN](#)

Note: Double-click on a CDN Note column in order to edit. Hit "enter" key to save and "ESC" to cancel. The maximum length for CDN account note is 31 characters.

[CDN Description](#)
How can I get assistance with CDN?

10D24

[Edit Note](#)
[Bandwidth](#)
[Origin Pull](#)
[Purge Cache](#)
[Upgrade](#)

[Administrative](#) | [Support](#) | [Security](#) | [Hardware](#) | [Software](#) | [Public Network](#) | [Private Network](#) | [Sales](#)

Copyright ©2014 SoftLayer Technologies, Inc. All Rights Reserved. [Legal](#)

SOFTLAYER[®]
an IBM Company

24 © 2015 IBM Corporation




Figure 9-21 SoftLayer Content Delivery Network use case (4)

Notes:

Step 5: Add Origin Pull mapping in your SoftLayer CDN account

After the update of your DNS zone, you can add the Origin Pull mapping in your SoftLayer CDN account. Click **More** to the right of your CDN account and select **Origin Pull** from the menu.

9.25 SoftLayer Content Delivery Network use case (5)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

CDN List | Total bandwidth | Origin Pull | Purge Cache

Virtual Server: CDN

CDN account name: 10D24

Add Origin Pull Mapping

Media Type: HTTP url

Origin Domain: blog.playvm.com * Required.

CNAME Record: cdn.playvm.com

Example origin hosts for "HTTP" and "Flash streaming": "http://www.mycompany.com" or "http://myblog.com/images"

You can also use your CDN FTP space as origin host. For example, you can use **CDN_FTP/http_content** for an origin host. You must provide a valid CNAME record in this case.

See below for more details on the CNAME option. A valid CNAME must point to **10D24.http.cdn.softlayer.net**.

Save Changes

25

© 2015 IBM Corporation




Figure 9-22 SoftLayer Content Delivery Network use case (5)

Notes:

Step 5 (continued)

Select **HTTP url** for the **Media Type**, enter your WordPress site domain (blog.playvm.com in this example) in the **Origin Domain** field, enter your CDN domain (cdn.playvm.com in this example) in the **CNAME Record** field, and save the changes.

9.26 SoftLayer Content Delivery Network use case (6)

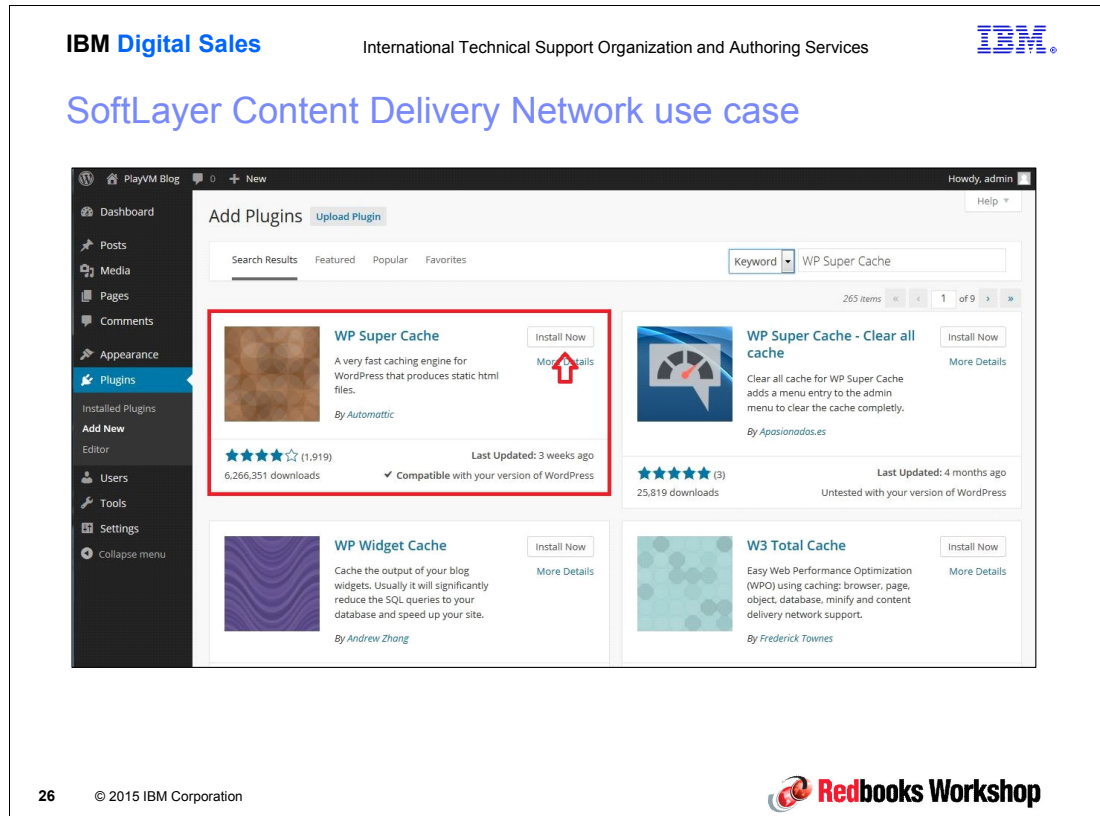


Figure 9-23 SoftLayer Content Delivery Network use case (6)

Notes:

Step 6: Install the WordPress CDN plug-in

To distribute your WordPress static content by using the SoftLayer CDN, you need to rewrite all URIs referencing static content to point to your SoftLayer CDN domain, and you need to configure the WordPress cache policy to use the CDN. Several WordPress plug-ins can help you with this process. This example uses the WP Super Cache plug-in to enable the SoftLayer CDN for your WordPress.

To install the WP Super Cache plug-in, complete these steps in the WordPress admin panel:

1. Click **Plugins** → **Add New**.
2. Search for WP Super Cache.
3. Select the **WP Super Cache** plug-in within the search results, and click **Install Now**.

9.27 SoftLayer Content Delivery Network use case (7)

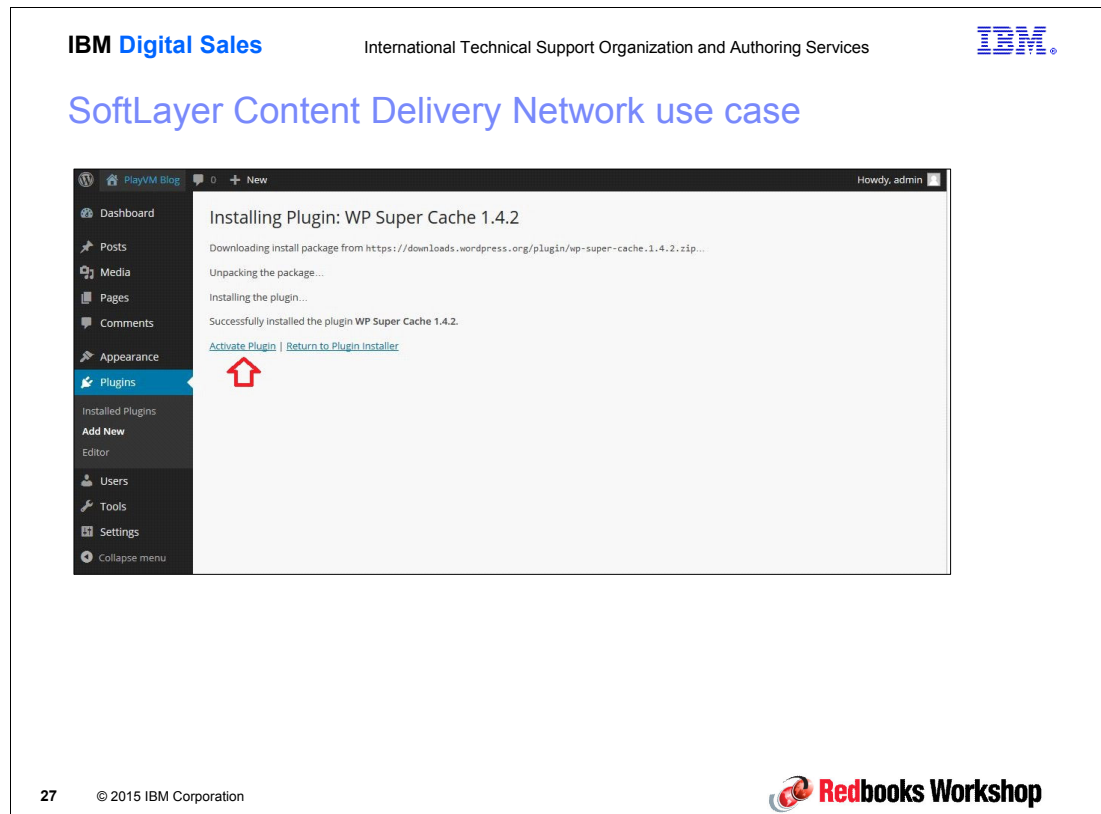


Figure 9-24 SoftLayer Content Delivery Network use case (7)

Notes:

Step 6 (continued)

After the plug-in is installed, click **Activate Plugin** to activate the WP Super Cache plug-in.

9.28 SoftLayer Content Delivery Network use case (8)

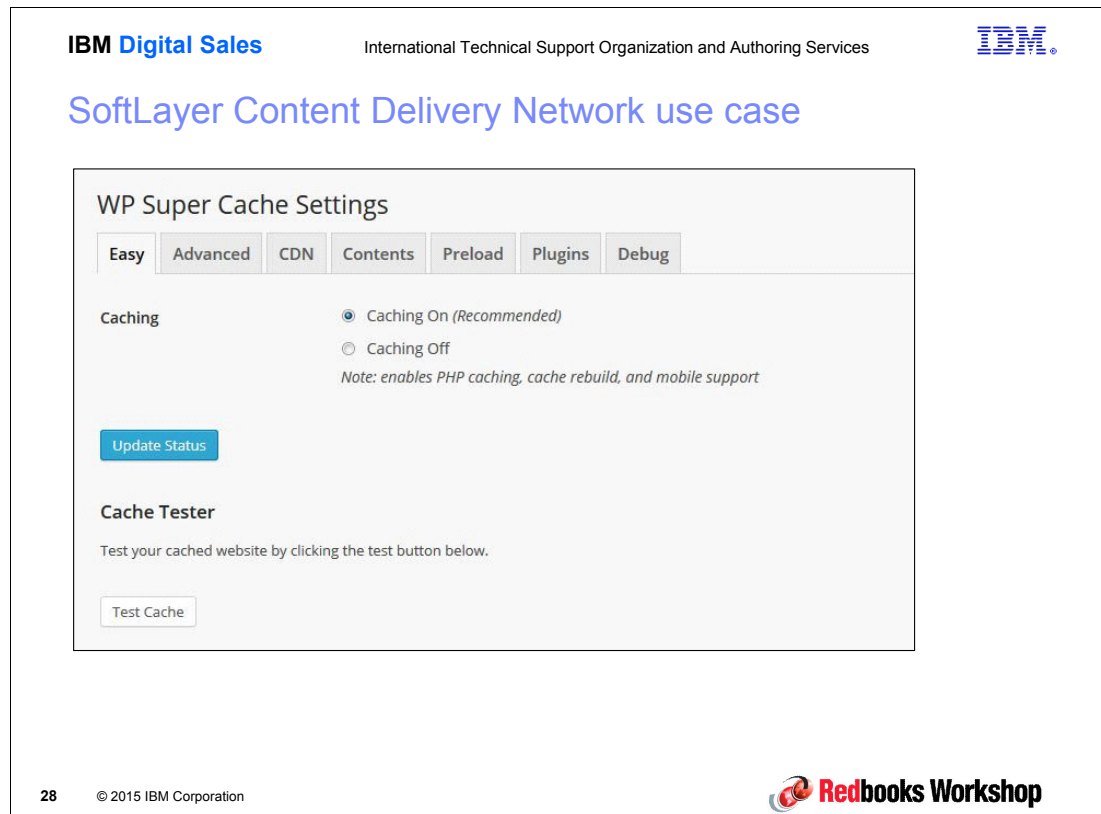


Figure 9-25 SoftLayer Content Delivery Network use case (8)

Notes:

Step 7: Configure the WordPress CDN plug-in

To configure the plug-in, complete these steps:

1. Navigate to **Settings** → **WP Super Cache**.
2. On the Easy tab, select **Caching On** for the **Caching** option and click **Update Status**.

9.29 SoftLayer Content Delivery Network use case (9)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

WP Super Cache Settings

Easy Advanced **CDN** Contents Preload Plugins Debug

Your website probably uses lots of static files. Image, javascript and CSS files are usually static files that could just as easily be served from another site or CDN. Therefore this plugin replaces any links in the `wp-content` and `wp-includes` directories (except for PHP files) on your site with the URL you provide below. That way you can either copy all the static content to a dedicated host or mirror the files to a CDN by [origin pull](#).

WARNING: Test some static urls e.g., `http://cdn.playvm.com/wp-includes/js/prototype.js` to ensure your CDN service is fully working before saving changes.

You can define different CDN URLs for each site on a multisite network.

☒ **Enable CDN Support**

Off-site URL
The new URL to be used in place of `http://blog.playvm.com` for rewriting. No trailing `/` please.
Example: `http://cdn.playvm.com/wp-includes/js/prototype.js`.

Include directories
Directories to include in static file matching. Use a comma as the delimiter. Default is `wp-content, wp-includes`, which will be enforced if this field is left empty.

Exclude if substring
Excludes something from being rewritten if one of the above strings is found in the match. Use a comma as the delimiter like this: `.php, .flv, .do`, and always include `.php` (default).

Additional CNAMEs
These **CNAMEs** will be used in place of `http://blog.playvm.com` for rewriting (in addition to the off-site URL above). Use a comma as the delimiter. For pages with a large number of static files, this can improve browser performance. **CNAMEs** may also need to be configured on your CDN.
Example: `http://cdn1.blog.playvm.com,http://cdn2.blog.playvm.com,http://cdn3.blog.playvm.com`

☒ **Skip https URLs to avoid "mixed content" errors**

[Save Changes](#)

29 © 2015 IBM Corporation




Figure 9-26 SoftLayer Content Delivery Network use case (9)

Notes:

Step 7 (continued)

On the CDN tab, select **Enable CDN Support**, enter `http://cdn.playvm.com` in the **Off-site URL** field, and select **Skip https URLs** to avoid “mixed content” errors. Keep the other configuration defaults and save the changes.

9.30 SoftLayer Content Delivery Network use case (10)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer Content Delivery Network use case

URL	Status	Domain	Size	Remote IP	Timeline
GET blog.playvm.com	200 OK	blog.playvm.com	16.0 KB	119.81.143.66:80	640ms
GET dashicons.min.css?ver=4.0	304 Not Modified	cdn.playvm.com	39.5 KB	72.21.81.253:80	1.03s
GET admin-bar.min.css?ver=4.0	304 Not Modified	cdn.playvm.com	18.3 KB	72.21.81.253:80	1.03s
GET genericicons.css?ver=3.0.3	304 Not Modified	cdn.playvm.com	30.3 KB	72.21.81.253:80	1.05s
GET style.css?ver=4.0	304 Not Modified	cdn.playvm.com	75.6 KB	72.21.81.253:80	830ms
GET jquery.js?ver=1.11.1	304 Not Modified	cdn.playvm.com	93.6 KB	72.21.81.253:80	830ms
GET jquery-migrate.min.js?ver=1	304 Not Modified	cdn.playvm.com	7.0 KB	72.21.81.253:80	830ms
GET admin-bar.min.js?ver=4.0	304 Not Modified	cdn.playvm.com	6.7 KB	72.21.81.253:80	1.06s
GET functions.js?ver=20140616	304 Not Modified	cdn.playvm.com	3.4 KB	72.21.81.253:80	1.05s

30 © 2015 IBM Corporation

Figure 9-27 SoftLayer Content Delivery Network use case (10)

Notes:

Step 8: Check the results

You have completed the SoftLayer CDN support for your WordPress. Open your WordPress site in a Firefox browser. Now the static contents of your WordPress site are fetched from the SoftLayer CDN. In this figure, Firebug was opened in the lower part of the panel. It shows the GET requests from WordPress for a page with static content that is pulled from `cdn.playvm.com`. HTTP response code 304 means that the requested content has not changed since your last request.

References and further reading:

- ▶ IBM Cloud: See the SoftLayer Knowledge Base.
- ▶ Getting started with the IBM Cloud marketplace:
<https://developer.ibm.com/marketplace/docs/getting-started-2/>
- ▶ SoftLayer CDN:
<http://www.softlayer.com/content-delivery-network>
- ▶ GoDaddy:
<http://www.godaddy.com/>
- ▶ WP Super Cache:
<https://wordpress.org/plugins/wp-super-cache/>

9.31 Vyatta appliance

IBM Digital Sales

Vyatta appliance

1. IPSec VPN
2. NAT
3. Firewall
4. Router

International Technical Support Organization and Authoring Services

SOFTLAYER
an IBM Company

Gateway Appliances

Single Processor Multi-Core Servers

Manufacturer	CPU Type	Core Count	CPU Speed	Storage	Ram	Private Network Only	10 Gbps Uplink Support	Starting Price Per Month
Intel	Intel Xeon E3-1270	4 Cores	3.40GHz	Up to 4 drives	4 GB - 32 GB	-	-	\$300.06

Dual Processor Multi-Core Servers

Manufacturer	CPU Type	Core Count	CPU Speed	Storage	Ram	Private Network Only	10 Gbps Uplink Support	Starting Price Per Month
Intel	Intel Xeon E5-2620	12 Cores	2.00GHz	Up to 12 drives	16 GB - 256 GB	-	Yes	\$543.31
Intel	Intel Xeon E5-2650	16 Cores	2.00GHz	Up to 12 drives	16 GB - 256 GB	-	Yes	\$500.31
Intel	Intel Xeon E5-2690	16 Cores	2.90GHz	Up to 12 drives	16 GB - 256 GB	-	Yes	\$672.81

The diagram illustrates the Vyatta appliance architecture. It shows a central 'Vyatta Appliance 1' (represented by a blue cylinder with a network icon) connected to two data centers, 'SoftLayer DC 1' and 'SoftLayer DC 2'. Above the data centers is a cloud icon labeled 'SoftLayer private network multiple carriers'. Below the data centers are two server racks. The left rack contains 'Server 1', 'Server 2', and 'Server 3'. The right rack contains 'Server 1' and 'Server 2'. Arrows indicate traffic flow: 'Public VLAN' (blue arrow) from the left rack to the appliance, 'Private VLAN #2' (grey arrow) from the appliance to the left rack, 'Private VLAN #1' (pink arrow) from the appliance to the right rack, and 'Public VLAN' (blue arrow) from the right rack to the appliance. The appliance is also connected to a 'SoftLayer DC 1' server rack.

31 © 2015 IBM Corporation

Redbooks Workshop

Figure 9-28 Vyatta appliance

Notes:

Vyatta Network Gateway

A network gateway provides tools to manage traffic into and out of one or more virtual local area networks (VLANs). The network gateway serves a customer-configurable routing device in front of designated VLANs. The servers in those VLANs route through the network gateway appliance as their first hop instead of Frontend Customer Routers (FCR) or Backend Customer Routers (BCR).

The general function of a network gateway might seem a little abstract, so here are some real world use cases to see how you can put that functionality to work in your own cloud environment.

Example 1: Complex Traffic Management

This example involves a multi-server cloud environment and a complex set of firewall rules that allow certain types of traffic to certain servers from specific addresses. Without a network gateway, you must manually configure multiple hardware and software firewalls throughout your topology and maintain multiple rules sets. With the network gateway appliance, you can streamline your configuration into a single point of control on both the public and private networks.

After you order a gateway appliance in the SoftLayer portal and configure which VLANs route through the appliance, the process of configuring the device is simple. Define your

production, development, and QA environments with distinct traffic rules, and the network gateway handles the traffic segmentation. If you wanted to create your own VPN to connect your hosted environment to your office or in-house data center, that configuration is quick and easy as well. The challenge of managing several sets of network rules across multiple devices is simplified and streamlined. The figure shows that Private VLANs #1 in two SoftLayer Data Centers are routed to each other. The VLAN #2 is still separated. This is very different from enabling the VLAN spanning in SoftLayer that enables traffics between all VLANs configured on your account.

Example 2: Creating a Static NAT

In this example, you want to create a static network address translation (NAT) so that you can direct traffic through a public IP address to an internal IP address. With the IPv4 address pool dwindling and new allocations being harder to come by, this configuration is becoming extremely popular to accommodate users who cannot yet reach IPv6 addresses. This challenge would normally require a significant level of effort of even the most seasoned systems administrator, but with the gateway appliance, it is a painless process.

In addition to the IPv4 address-saving benefits, your static NAT adds a layer of protection for your internal web servers from the public network, and as described in the first example, your gateway device also serves as a single configuration point for both inbound and outbound firewall rules.

If you have complex network-related needs, and you want granular control of the traffic to and from your servers, a gateway appliance might be the perfect tool for you. You get the control that you want and save yourself a significant amount of time and effort configuring and tweaking your environment on-the-fly. You can terminate IPSec VPN tunnels, run your own network address translation, and run diagnostic commands such as traffic monitoring (**tcpdump**) on your global environment. And in addition to that, your gateway serves as a single point of contact to configure sophisticated firewall rules.

Currently, SoftLayer provides two types of Vyatta Gateway Appliance:

- ▶ Single Processor
- ▶ Dual Processor with 10Gbps uplink support.

To make an order, go to the Network section in SoftLayer's order panel and select **Vyatta**.


9.32 Recap

IBM Digital Sales International Technical Support Organization and Authoring Services 

Recap

- Detailed SoftLayer network architecture
- IP addresses in SoftLayer
- VLANs and VLAN spanning
- Load balancing solutions in SoftLayer
- SoftLayer Content Delivery Network
- Vyatta appliance

SUMMARY




© 2015 IBM Corporation 

Figure 9-29 Recap

Notes:
(none)


9.33 Checkpoint questions

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



1. What types of IP blocks are available within SoftLayer?

33 © 2015 IBM Corporation


 Redbooks Workshop

Figure 9-30 Checkpoint questions

Notes:
(none)


9.34 Checkpoint questions (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



**CHECKPOINT
AHEAD**

1. What types of IP blocks are available within SoftLayer?

Answer: Static and portable

34

© 2015 IBM Corporation


 Redbooks Workshop

Figure 9-31 Checkpoint questions (2)

Notes:
(none)


9.35 Checkpoint questions (3)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



2. What is VLAN spanning?

35 © 2015 IBM Corporation


 Redbooks Workshop

Figure 9-32 Checkpoint questions (3)

Notes:

(none)


9.36 Checkpoint questions (4)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



2. What is VLAN spanning?

Answer: VLAN spanning is an account setting that enables traffic to travel between private VLANs on a single account. By enabling VLAN spanning, devices on an account's different private VLANs can communicate with one another.

36 © 2015 IBM Corporation


 Redbooks Workshop

Figure 9-33 Checkpoint questions (4)

Notes:

(none)


9.37 Checkpoint questions (5)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



3. What are the load balancing solutions in SoftLayer?

37 © 2015 IBM Corporation

 Redbooks Workshop

Figure 9-34 Checkpoint questions (5)

Notes:

(none)


9.38 Checkpoint questions (6)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



3. What are the load balancing solutions in SoftLayer?

Answer:

- Local load balancing
- Dedicated local load balancer
- Citrix NetScaler VPX

38 © 2015 IBM Corporation




Figure 9-35 Checkpoint questions (6)

Notes:
(none)


9.39 Checkpoint questions (7)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



4. What is a CDN?

39

© 2015 IBM Corporation


 Redbooks Workshop

Figure 9-36 Checkpoint questions (7)

Notes:

(none)


9.40 Checkpoint questions (8)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



4. What is a CDN?

Answer: A Content Delivery Network (CDN) is a solution for organizing, storing, and streaming content on the web with optimized flow of content to users.

40

© 2015 IBM Corporation




Figure 9-37 Checkpoint questions (8)

Notes:

(none)


9.41 Checkpoint questions (9)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



5. What are the major functions of the Vyatta appliance?

41

© 2015 IBM Corporation


 Redbooks Workshop

Figure 9-38 Checkpoint questions (9)

Notes:

(none)


9.42 Checkpoint questions (10)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint questions



5. What are the major functions of the Vyatta appliance?

Answer:

- IPSec VPN
- NAT
- Firewall
- Router

42 © 2015 IBM Corporation


 Redbooks Workshop

Figure 9-39 Checkpoint questions (10)

Notes:
(none)

9.43 Introduction to OSI model

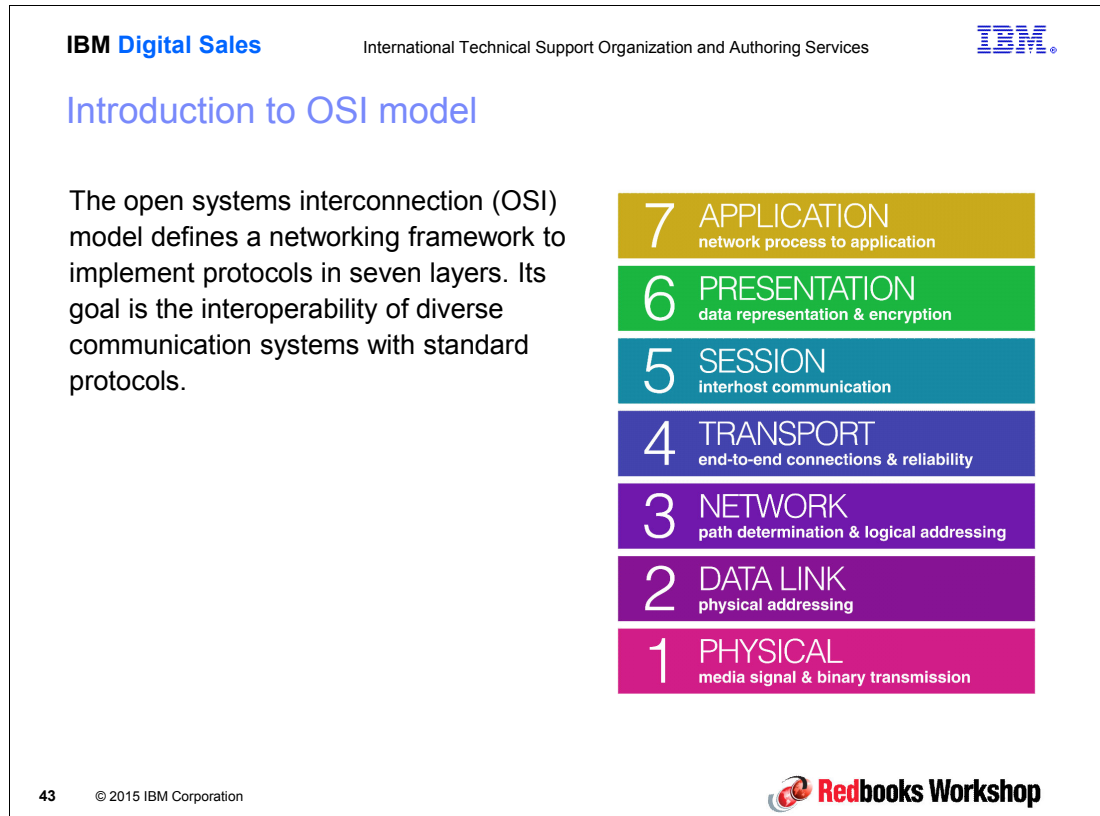


Figure 9-40 Introduction to OSI model

Notes:

The open system interconnection (OSI) model defines a networking framework to implement protocols in seven layers. Control is passed from one layer to the next, starting at the application layer in one station over the channel to the next station until it gets to the bottom layer, and back up the hierarchy. The OSI model doesn't do any functions in the networking process. It is a conceptual framework so you can better understand complex interactions that are happening. In theoretical discussions, the OSI Reference Model helps you understand how networks and network protocols function. In the “real world”, it also helps you figure out which protocols and devices can interact with each other. OSI consists of these layers:

- Physical (Layer 1): This layer conveys the bit stream (electrical impulse, light, or radio signal) through the network at the electrical and mechanical level. It provides the hardware means of sending and receiving data on a carrier, including defining cables, cards, and physical aspects. Fast Ethernet, RS232, and ATM are protocols with physical layer components.

Layer 1 Physical examples include Ethernet, FDDI, B8ZS, V.35, V.24, and RJ45.

- Data Link (Layer 2): At this layer, data packets are encoded and decoded into bits. It furnishes transmission protocol knowledge and management, and handles errors in the physical layer, flow control, and frame synchronization. The data link layer is divided into two sub layers: The Media Access Control (MAC) layer and the Logical Link Control (LLC) layer. The MAC sub layer controls how a computer on the network gains access to the data and permission to transmit it. The LLC layer controls frame synchronization, flow control, and error checking.

Layer 2 Data Link examples include PPP, FDDI, ATM, IEEE 802.5/ 802.2, IEEE 802.3/802.2, HDLC, and Frame Relay.

- Network (Layer 3): This layer provides switching and routing technologies, creating logical paths, known as virtual circuits, for transmitting data from node to node. Routing and forwarding are functions of this layer, as well as addressing, internetworking, error handling, congestion control, and packet sequencing.

Layer 3 Network examples include AppleTalk DDP, IP, and IPX.

- Transport (Layer 4): This layer provides transparent transfer of data between end systems or hosts, and is responsible for end-to-end error recovery and flow control. It ensures complete data transfer.

Layer 4 Transport examples include SPX, TCP, and UDP.

- Session (Layer 5): This layer establishes, manages, and terminates connections between applications. The session layer sets up, coordinates, and terminates conversations, exchanges, and dialogues between the applications at each end. It deals with session and connection coordination.

Layer 5 Session examples include NFS, NetBios names, RPC, and SQL.

- Presentation (Layer 6): This layer provides independence from differences in data representation (e.g., encryption) by translating from application to network format, and vice versa. The presentation layer transforms data into the form that the application layer can accept. This layer formats and encrypts data to be sent across a network, providing freedom from compatibility problems. It is sometimes called the syntax layer.

Layer 6 Presentation examples include encryption, ASCII, EBCDIC, TIFF, GIF, PICT, JPEG, MPEG, and MIDI.

- Application (Layer 7): This layer supports application and end-user processes. Communication partners are identified, quality of service is identified, user authentication and privacy are considered, and any constraints on data syntax are identified. Everything at this layer is application-specific. This layer provides application services for file transfers, e-mail, and other network software services. Telnet and FTP are applications that exist entirely in the application level. Tiered application architectures are part of this layer.

Layer 7 Application examples include WWW browsers, NFS, SNMP, Telnet, HTTP, and FTP.

9.44 Understanding TCP/IP addressing and subnetting basics

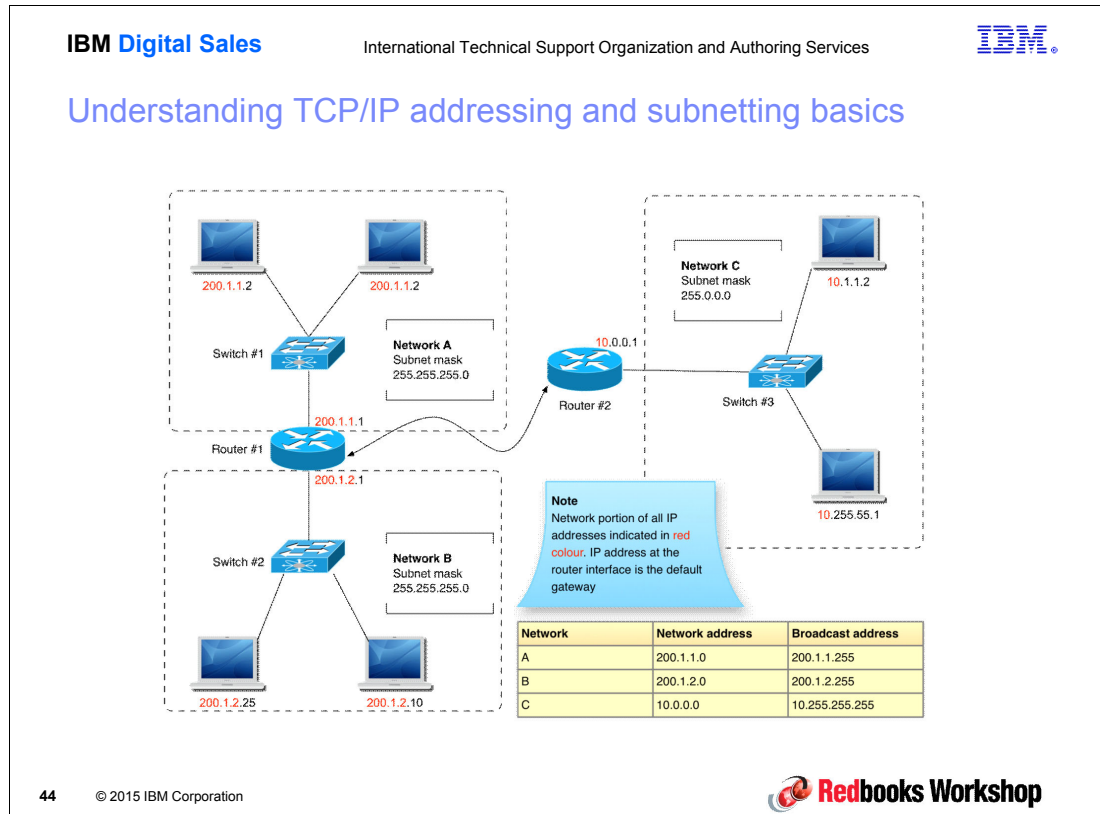


Figure 9-41 Understanding TCP/IP addressing and subnetting basics

Notes:

What is an IP address?

An IP address is used to identify a host (host or any network device) in a network. IP is a 32-bit binary number divided into four octet groups, each octet giving a maximum of 255 in decimal. For easier addressing of these IP address octet, they are written as dotted decimals, such as 10.0.0.1.

Classes of IPv4 address in networking

IP is divided into five classes of network addresses based on the range of the first octet. Out of the total valid addresses in each class, two dedicated IP address are reserved for these items:

- Network address
- Broadcast address

The total number of available IP addresses is therefore $2^n - 2$.

Public and Private IP addresses

To communicate over an internet connection, a device must have a public IP address that is provided by the Internet Assigned Numbers Authority (IANA). The private range of IP addresses are used in an intranet (an internal network that uses internet technology). IANA also provides address for private networks in each class as follows:

- ▶ Class A: 10.0.0.0 - 10.255.255.255
- ▶ Class B: 172.16.0.0 - 172.31.255.255
- ▶ Class C: 192.168.0.0 - 192.168.255.255

What is Subnet Mask?

A subnet mask is a 32 bit address used with an IP in order to identify its network and host portions. For example, if you have an IP address 200.1.1.2 with a subnet mask 255.255.255.0, it means that 200.1.1 is the network portion and last octet is the host portion. Any IP that starts with 200.1.1 goes to the same network (Network A), like 200.1.1.1, 200.1.1.10, 200.1.1.100 up to 200.1.1.254. These IPs therefore do not require a router to communicate with each other.

In Network A, the first IP (200.1.1.0) is used to indicate network address and the last IP (200.1.1.255) is used to send broadcast messages to all host computers in network A. Another IP 200.1.2.2 that has the same subnet mask cannot communicate with Network A without using a router because there is a change in the network part. It belongs to another network with network address 200.1.2.0 (Network B).

Another IP 10.1.1.2 with subnet mask 255.0.0.0 makes you understand that it belongs to the network 10.0.0.0 (Network C), where only the first octet indicates network.

Therefore, subnet masks help you understand which IPs belongs to which network. By default, the following subnet masks are used:

- ▶ Class A: 255.0.0.0
- ▶ Class B: 255.255.0.0
- ▶ Class C: 255.255.255.0

All host bits '0' are a network address.

All host bits '1' are a broadcast address.

Class A Network

Class A network range goes from 1.0.0.0 to 126.255.255.255. The Class A network subnet mask is 255.0.0.0, which means it has eight network bits of which the first bit is fixed as '0'. And hence a total of seven network bits and 24 host bits. The total number of networks is $2^7 - 2 = 126$. Two are subtracted because 0.0.0.0 is the default network and 127.0.0.0 is the loopback IP address used for checking proper functionality (self testing). The total number of hosts per network is $2^{24} - 2 = 16777214$. Two are subtracted for the network and broadcast addresses.

Class A network example:

- ▶ Network address - 1.0.0.0
- ▶ Subnet Mask - 255.0.0.0
- ▶ First host IP address - 1.0.0.1
- ▶ Last host address - 1.255.255.254
- ▶ Broadcast address - 1.255.255.255

Class B Network

The class B network range goes from 128.0.0.0 to 191.255.255.255. The default subnet mask is 255.255.0.0, which means it has 16 network bits of which first two bits are fixed as '10'. It has a total of 14 network bits and 16 host bits, so the total number of networks is $2^{14} = 16384$. The total number of hosts per network is $2^{16} - 2 = 65534$. Two IPs are subtracted, one each for the network and broadcast addresses.

Class B network example:

- ▶ Network address - 128.0.0.0
- ▶ Subnet Mask - 255.255.0.0
- ▶ First host address - 128.0.0.1
- ▶ Last host address - 128.0.255.254
- ▶ Broadcast address - 128.0.255.255

Class C Network

The IP range goes from 192.0.0.0 to 223.255.255.255. The Class C network subnet mask is 255.255.255.0, which means it has 24 network bits of which the first three bits are fixed as '110'. It therefore has a total of 21 network bits and 8 host bits, so the total number of networks is $2^{21} = 2097152$, and the total number of hosts per networks is $2^8 - 2 = 254$. Two IPs are subtracted: One for the network address and the other for the broadcast address.

Class C network example:

- ▶ Network address - 192.0.0.0
- ▶ Subnet Mask - 255.255.255.0
- ▶ First host address - 192.0.0.1
- ▶ Last host address - 192.0.0.254
- ▶ Broadcast address - 192.0.0.255

IPv6 - brief description

IPv6 is the next generation Internet Protocol (IP) address standard that will supplement and eventually replace IPv4, the protocol most Internet services use today. The world ran out of the 4.3 billion available IPv4 addresses, so to allow the Internet to continue to grow and spread across the world, implementing IPv6 is necessary. IPv6 uses a 128-bit address, allowing 2^{128} , or approximately 3.4×10^{38} addresses, or more than 7.9×10^{28} times as many as IPv4. IPv6 addresses are represented as eight groups of four hexadecimal digits with the groups being separated by colons, for example 2001:0db8:85a3:0042:1000:8a2e:0370:7334.



Securing cloud environments

This unit provides information about securing your SoftLayer cloud environments by using physical security, encryption, and firewalls.

10.1 What you should be able to do

After you have completed this unit, you should be able to:

- ▶ Explain how the infrastructure changes when moving the workloads from your server room to the cloud.
- ▶ Articulate the importance of knowing and accepting the security policies that your IaaS provider has chosen.
- ▶ Ensuring that your data is secured not only in the data center, but also when you transfer it.
- ▶ Describe the different types of firewalls and how they can be used.
- ▶ Explain how you can secure the server instance itself.
- ▶ Use the Customer Portal to manage the hardware firewall and perform vulnerability scans on your systems to ensure their integrity.

10.2 References

The following resources are useful for further reference:

- ▶ Compliances and certifications:
<http://www.softlayer.com/compliance>

10.3 Teaching topics

This unit covers the following topics:

- ▶ Why security is important
- ▶ Data center security
- ▶ Additional security offerings
- ▶ Securing communications
- ▶ Securing instances using firewalls
- ▶ Hardware firewalls
- ▶ Appliance firewalls
- ▶ OS firewalls
- ▶ IaaS provider firewall offerings
- ▶ Securing instances by using OS hardening
- ▶ Using the Customer Portal to set up and verify security

10.4 Why security is important

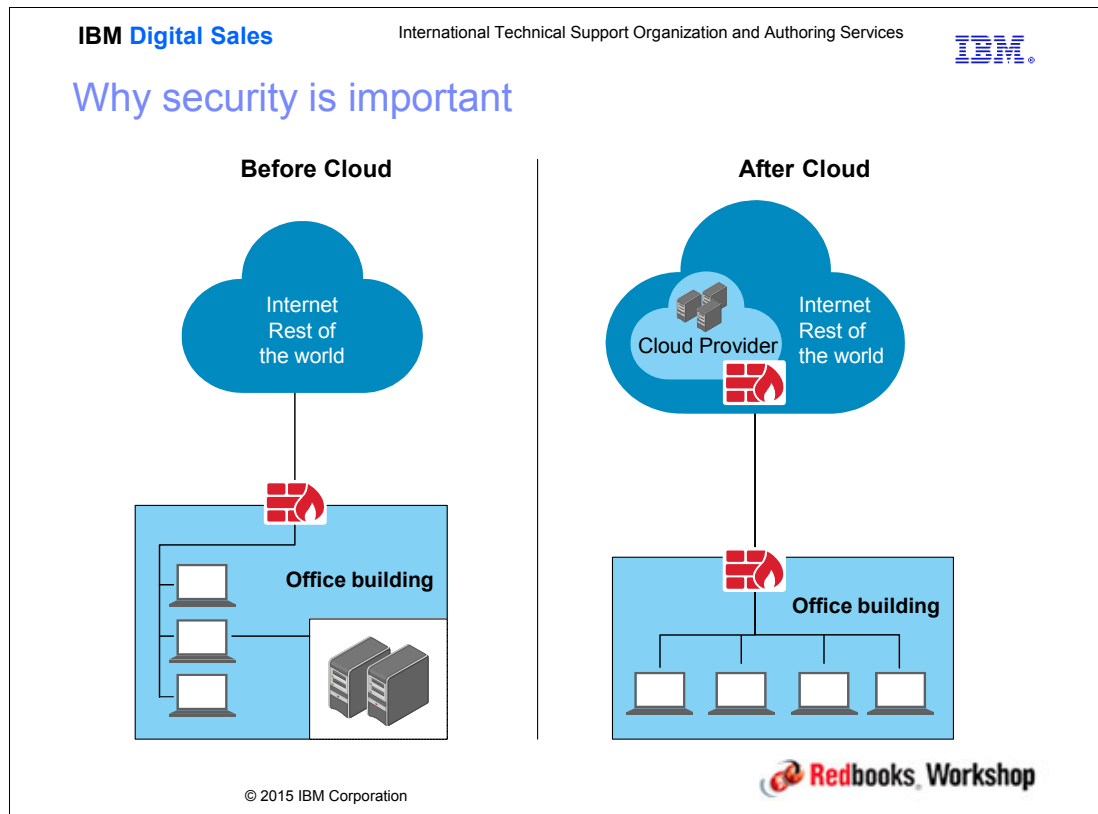


Figure 10-1 Why security is important

Notes:

Security has always been important, and the move to IaaS has not changed that. In fact, before moving to IaaS, your company servers were safely secured in your company server room and accessible only from the company workstation or laptops. They could be completely separated from the Internet. Physical access to the servers required access to the building and again to the server room which typically is a secured area only accessible by few people in the company. Many people would not know where their company's server room is located.

After moving to IaaS a number of things change:

- ▶ The servers are located at an IaaS provider's data center.
- ▶ To access our servers, you now go over the Internet.
- ▶ Access is now through a portal, which allows you to perform the same functions as if you had physical access to the servers.

Your server room now has two doors into it:

- ▶ The physical door at the cloud provider's data center, so you need to ensure that your provider's security procedures are in place and are acceptable to you.
- ▶ The second door is through the Internet, and you must ensure that you secure this door. In this case, this involves securing your user ID and password to the portal which you use to control your servers and IaaS services as well as securing all user IDs and passwords defined on each of your IaaS servers so that you can ensure the integrity of your servers and data.

One thing that has not changed is that you should secure the communication from your servers to the clients accessing them from the Internet. This can be done by installing Secure Socket Layer (SSL) certificates on those servers sending data out over the Internet to your clients because this will encrypt the communication between the server and endpoint.

10.5 Data center security

IBM Digital Sales International Technical Support Organization and Authoring Services **IBM**

Data center security



When moving to IaaS, the physical security is in the hands of the provider. **Choose wisely**

© 2015 IBM Corporation

Redbooks Workshop

Figure 10-2 Data center security

Notes:

Once you move your workload to IaaS, the physical part of security is the responsibility of your IaaS provider. A serious IaaS provider will share their certifications and security measures (to a certain degree). It is up to you to decide if they meet your requirements.

SoftLayer, for example, readily shares that they have physical controls that limit only certified employees to their data center, but does not share how the physical controls work as that would be a breach of protocol.

Additionally, SoftLayer takes the following security precautions:

- ▶ Provides Service Organization Controls (SOC) 1, SOC 2, and SOC 3 reports. These reports evaluate SoftLayer's operational controls with respect to criteria set by the American Institute of Certified Public Accountants (AICPA) Trust Services Principles.
- ▶ Are ISO 27001 compliant. This is a widely-adopted global security standard that outlines the requirements for information security management systems and provides a systematic approach to managing company and customer information based on periodic risk assessments.
- ▶ Are Safe Harbor compliant. Safe Harbor is an important way for U.S. companies to avoid experiencing interruptions in their business dealings with the EU or facing prosecution by European authorities under European privacy laws. Safe Harbor certification ensures that EU organizations know that your company provides "adequate" privacy protection, as defined by the Directive.

Different countries have different standards, so choose an IaaS provider that fulfills the standards required by you and your country.

10.6 Additional security offerings

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Additional security offerings



Now that the physical security is in place, you must make sure that you can access servers and data as securely as if they were in your own data center.

© 2015 IBM Corporation

Redbooks Workshop

Figure 10-3 Additional security offerings

Notes:

Before IaaS, traffic between company workstations and servers were on the internal network. This is not the case anymore, so you also need to consider how to ensure that the information sent between your servers at the IaaS provider's data center and our computers is not intercepted by malicious third parties.

You need to see what your selected provider offers so that you can communicate with your system in their data centers and be certain that only the intended recipients have access to it. Does the IaaS provider offer VPN access? Do they offer encrypted communication between servers or servers and endpoints? SoftLayer does offer these and more such as site to site VPN and client to site VPN.

Although it may seem like all IaaS providers sell the same services, study their offerings and services closer because there might be important differences.

10.7 Securing communications

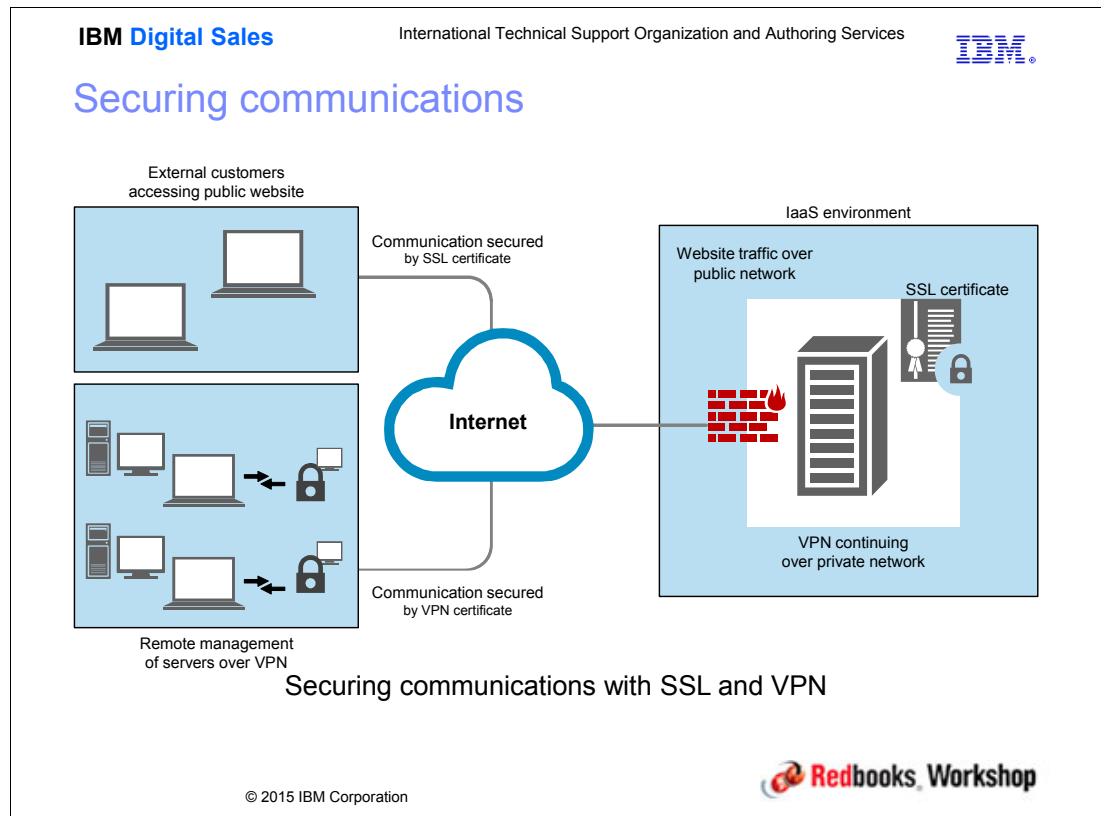


Figure 10-4 Securing communications

Notes:

Because you now must communicate with your servers through the Internet, take measures to ensure that the data you send and receive is suitably encrypted to protect it from malicious parties.

For your customers accessing your public website, there is little change in the normal procedure because you already use SSL certificates to secure communication sent from and to your web servers. Communication now takes place on the public network at the IaaS provider's end. SSL certificates are purchased from certified vendors and are valid for one or more years before they have to be renewed.

Your administrators connect through a VPN connection. After they connect to the IaaS provider's network, they are routed over the private network at the IaaS provider's end.

VPN services should be supplied by the IaaS provider.

10.8 Securing instances using firewalls

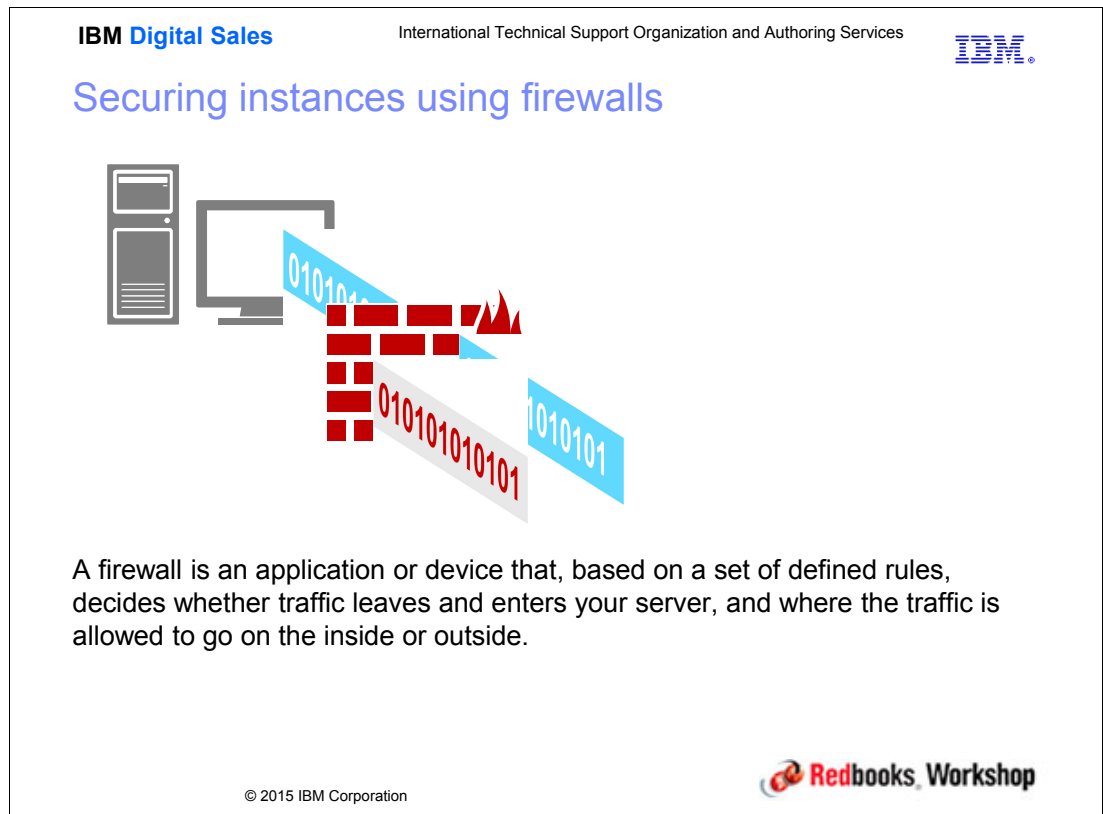


Figure 10-5 Securing instances using firewalls

Notes:

Now that you have ensured that physical security is in place and that communication to and from your devices is secure, ensure that data on your servers can only exit to the network in a way that you specify. Also, ensure that no one can enter your server from the network except along paths that you allow (a website for example).

Most IaaS providers offer firewalls, and all of them use them inside their infrastructure. In addition, firewalls and VLANs are used by all IaaS providers to separate customers inside their offerings.

10.9 Securing instances using firewalls (2)

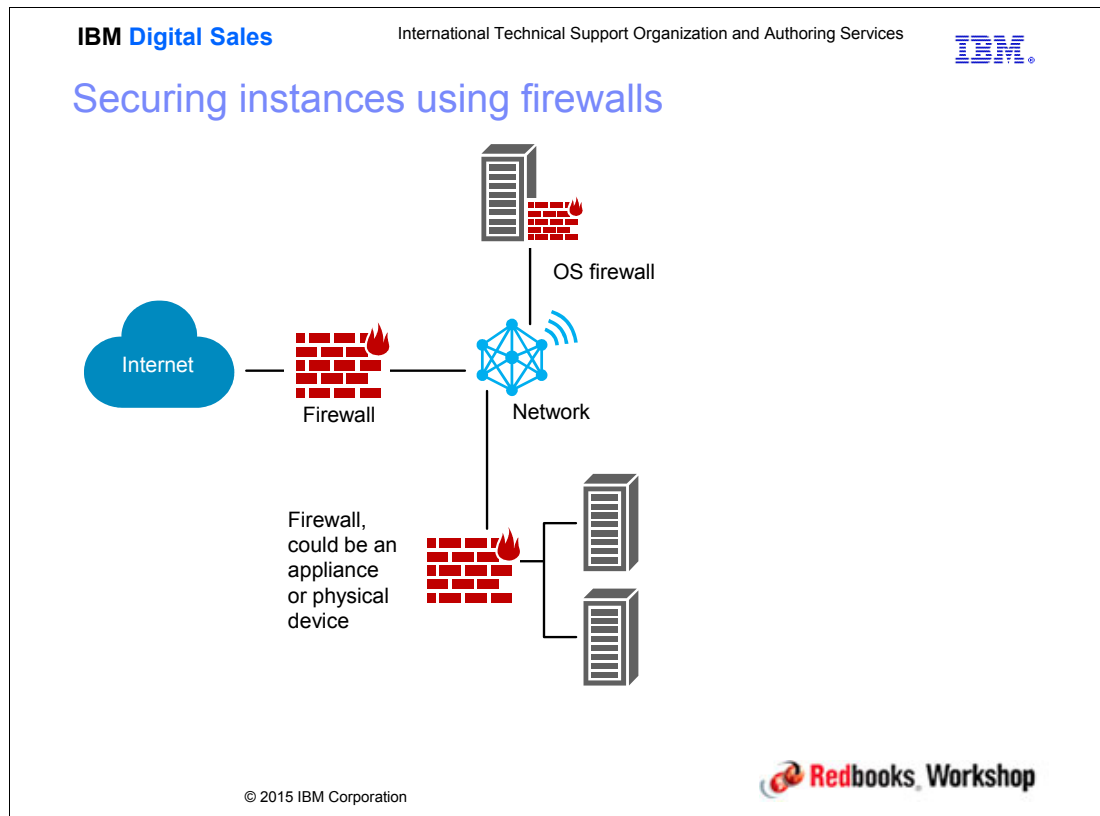


Figure 10-6 Securing instances using firewalls (2)

Notes:

A firewall can be placed on different parts of the network, or on the machine it is to protect/isolate itself. This example shows these different configurations:


- ▶ A system that has an OS firewall, meaning that the firewall is either running as an application on the machine or is part of the operating system itself.
- ▶ Two systems that are “sharing” a firewall, in this case an appliance firewall or a physical firewall device.
- ▶ Another firewall that is protecting/isolating all of the systems from the Internet.

You can have as many firewalls as you like, and each can have different rules. However, be sure to clearly document which rules each firewall has because you could end up having to do a lot of troubleshooting if you have conflicting rules defined.

10.10 Hardware firewalls


IBM Digital Sales

International Technical Support Organization and Authoring Services




Hardware firewalls

A physical or hardware firewall is a network device, not unlike a router, that is placed on the network to isolate or protect devices on a network from another part of the same network or a different network.



Pro	Con
One hardware firewall can protect your entire network	Expensive compared to software firewalls
Run own dedicated CPU and memory	Not as easy to configure
Cannot be disabled by malware	Will expose more data if it fails



© 2015 IBM Corporation

Figure 10-7 Hardware firewalls

Notes:

A hardware firewall is the most advanced and generally also the most expensive firewall, and is typically what you would use in larger installations. It has very advanced configuration options and is very secure.

10.11 Appliance firewalls

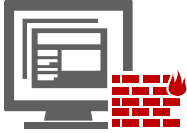
IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Appliance firewalls

An appliance firewall is a server (physical or virtual) that is running an operating system (OS) customized to start up with a firewall application.



Pro	Con
One appliance firewall can protect your entire network	Requires a separate server instance
Has its own CPU and memory	Requires regular maintenance with security patches because it is essentially a server
Easier to set up than hardware firewall	Not as effective resource-wise as a hardware firewall

© 2015 IBM Corporation




Figure 10-8 Appliance firewalls

Notes:

An appliance firewall is a stripped down operating system (usually Linux based) running a firewall application at startup. It can be placed anywhere on the network, and can protect one or more machines or network segments.

Compared to a hardware firewall, which has no OS, appliance firewalls run a real OS that could have its own vulnerabilities and a software-based firewall. It is also more sensitive to load than a hardware firewall. The setup can be anywhere from simple to advanced and is usually set up by using a web interface on the appliance.

An appliance firewall can be much more than just a firewall. An example is the Vyatta gateway appliance available with SoftLayer that can work as a router, firewall, gateway, and VPN server.

10.12 OS firewalls


IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

OS firewalls

An OS firewall is a firewall that runs on top of or embedded in the operating system of the server that is to be protected. It protects only the server itself and should it fail the server will be exposed.



Pro	Con
Cheap or free compared to the alternatives	Protects only one server.
Easy to setup	Consumes resources from the host that it is running on
Can have stricter rules than a firewall protecting many servers.	Often less stable than a hardware firewall


© 2015 IBM Corporation

Figure 10-9 OS firewalls

Notes:

The OS firewall is either built into or running on top of the operating system of the server it is meant to protect. They are easy to setup and work fine. However, they only protect that one server and consume resource from the server they are protecting, which means less resources are available for the workloads running on the servers.

The OS level firewall is the last line of defense because the firewall is directly on top of what it is to protect.

10.13 IaaS provider firewalls

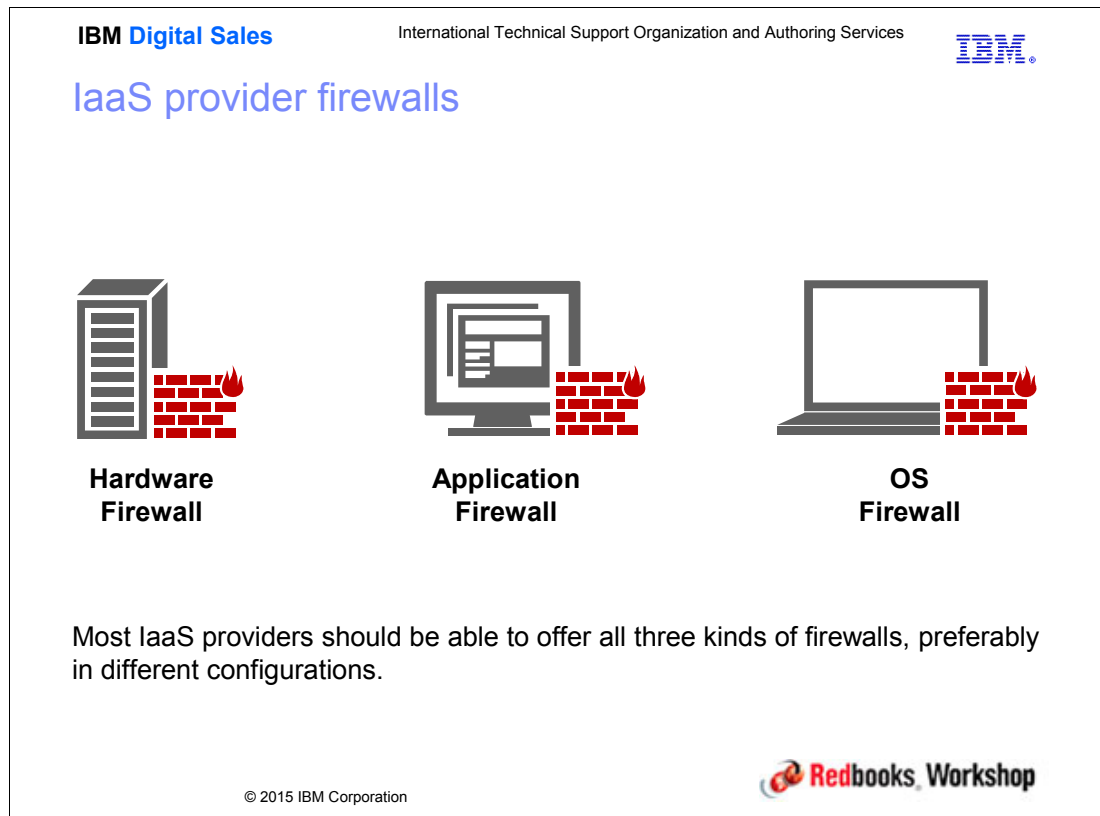


Figure 10-10 IaaS provider firewalls

Notes:

Most IaaS providers can provide all three types of firewalls. For example, SoftLayer offers the following options:

- ▶ **Hardware firewall:** Protect individual servers with hardware firewalls provisioned on demand without service interruptions.
- ▶ **Dedicated hardware firewall:** Protect one, multiple, or all servers that share the same VLAN with a dedicated hardware firewall, provisioned on demand without service interruptions.
- ▶ **High Availability redundant firewall:** Protect one, multiple, or all servers that share the same VLAN, with a secondary physical firewall for failover protection (and automatic fall back when primary firewall is restored).
- ▶ **Fortigate Security Appliance:** Provides complete, granular control over advanced firewall and security features. High availability options are available.
- ▶ **Gateway Appliances:** Software-defined firewall, router, VPN, and more that lets you create and manage virtual routers, firewalls, and VPN devices through user-defined parameters.
- ▶ **Vyatta Network OS Gateway Appliance:** A SoftLayer bare metal server with Vyatta Network OS can be customized, monitored, and tweaked to protect your infrastructure and optimize your network performance.

SoftLayer has even made the hardware firewalls configurable from the Customer Portal.

Note that when you configure the firewalls, make sure not to block the ports used by the IaaS provider to monitor your systems. Contact your IaaS provider to find out which ports are used.

10.14 Checkpoint

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Checkpoint

- Should your cloud provider have security policies?
- What are the changes in your network when moving to IaaS?
- What is a firewall?
- What are the three general types of firewalls?
- Can you have more than one firewall in your setup?

© 2015 IBM Corporation


Redbooks Workshop

Figure 10-11 Checkpoint

Notes:

(none)

10.15 Checkpoint (2)

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Checkpoint

- Should your cloud provider have security policies?
Yes, you need to ensure your servers and data are safe.
- What are the changes in your network when moving to IaaS?
Communication with your servers happens over the internet.
- What is a firewall?
A device that sorts traffic to and from your servers based on a set of defined rules.
- What are the three general types of firewalls?
Hardware, appliance, and OS.
- Can you have more than one firewall in your setup?
You can have as many as you like, but be sure to document the rules.


© 2015 IBM Corporation

Figure 10-12 Checkpoint (2)

Notes:
(none)


10.16 Securing instances using OS hardening

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Securing instances using OS hardening



The operating system of your computer and its configuration is part of the security. Some general rules to secure your server are:

- Keep the operating system updated with relevant security patches
- Do not give user more access than needed
- If adding users to the server, have strict permission rules and, if possible, authenticate using SSH keys, not passwords
- Only use your server for the purpose it was provisioned for
- Only allow inbound and outbound network traffic as required for the operations that it is to perform and are required by the IaaS provider
- Use antivirus/antimalware software to protect your server
- Disconnect any unneeded interfaces

© 2015 IBM Corporation

Redbooks Workshop

Figure 10-13 Securing instances using OS hardening

Notes:

Firewalls will not help you if you do not have a security policy for your servers and the users accessing them. Follow these best practices:

- ▶ Keep the OS updated with the latest security patches. It is also a good idea to have a fixed maintenance window scheduled for such maintenance.
- ▶ Give administrative access to the server only to users who actually need it.
- ▶ Whenever possibly try to avoid passwords as we all have heard about people using the same password for everything and often not even a strong one. Instead it is recommended to use SSH keys where the public key is located on the server and the private key at the person logging in.
- ▶ Do not suddenly start installing software onto the server to do something that only has to be done once and can be done from elsewhere. Remember the more software you put on your server the more potential vulnerabilities you introduce.

As mentioned, you can have multiple firewalls; you can easily have a OS firewall implementing the same rules as the appliance or hardware firewalls. Most operating systems come with such an offering and because it is free, consider adding this extra layer of security.

You will need to make sure the OS on your servers is hardened before you open your server to traffic. Be sure to apply all patches BEFORE you open your server and KEEP it patched on a regular, scheduled basis including applying emergency security patches as they arise. Most IaaS providers should be able to provide OS patch management services..

If you have a server in your IaaS environment that is only accessed by other servers from that environment, disable the public interface of that server and let the other servers use the IaaS provider's private network to communicate with it. This adds an extra layer of security for that server and you will still be able to access it through VPN on the private network.

10.17 Using portal to set up and verify security

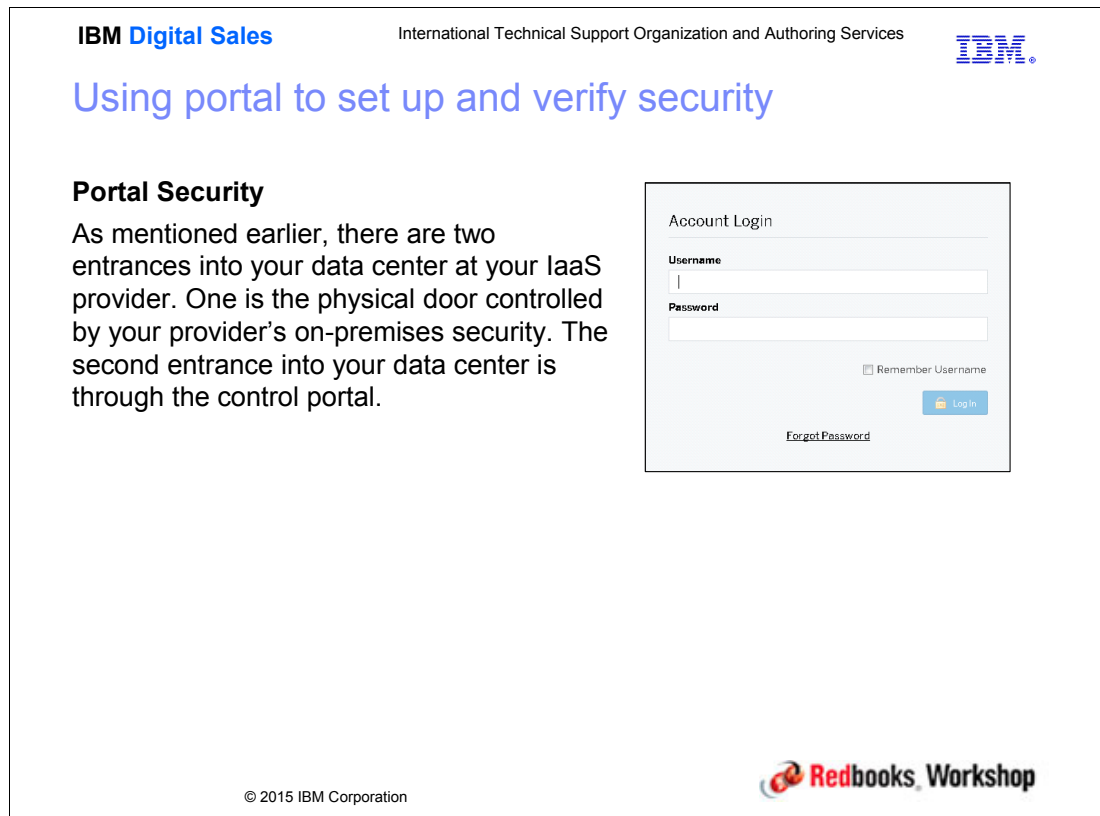


Figure 10-14 Using portal to set up and verify security

Notes:

The portal is your access to the datacenter. Make sure that only those who really need it have access to the portal and that their authorization inside the portal does not give them more access than they need. Consider adding security questions or two factor authentication as extra security measures when logging in to the portal.

A best practice is to have procedures to regularly validate the users' continued need for access to the portal.

SoftLayer provides these security enhancements:

- ▶ Password rules
- ▶ Audit logs for who logged into the portal
- ▶ Enhance security with security questions and phone authentication

You can enable all of these in the portal to increase security.


10.18 Administering firewall from the portal

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Administering firewall from the portal

db01.companya.softlayer.com
Public IP: 173.193.90.102 (Houston 2)
Private IP: 10.77.182.5

Configuration Tickets Bandwidth Remote Mgmt Monitoring Security Passwords Storage Firewall

Status: ? There are no rules associated with this firewall Actions

IPv4 Rules

Action	Source	CIDR	Destination	CIDR	Port Range	Protocol
1: Deny	0.0.0.0	0	Any		22 - 22	TCP
Rule 2: Permit	Any		Any			TCP
Notes:						
Deny	Any		Any		1 - 65535	Any

IPv6 Rules

Action	Source	CIDR	Destination	CIDR	Port Range	Protocol
Deny	Any		Any		1 - 65535	Any

Adding Rules

© 2015 IBM Corporation




Figure 10-15 Administering firewall from the portal

Notes:

The IaaS provider's customer portal might also have features allowing you to add/edit rules to the firewalls protecting your servers. This graphic shows how the hardware firewall in a SoftLayer setup can be modified directly from the portal by adding rules, editing rules, or disabling the firewall entirely. You can only edit the firewall rules if your profile permissions allow it.

Being able to configure the hardware firewall in the portal makes it the easiest of the firewalls to configure because the appliance firewall requires you to go to the appliance and do the setup, and the OS firewall requires editing files on the operating system of the server that you wish to protect.

10.19 Verify security in the portal

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Verify security in the portal

Hosts which were alive and responding during test	1
Number of security holes found	0
Number of security warnings found	7

Host List

Host(s)	Possible Issue
50.23.4.21	Security warning(s) found

[\[return to top\]](#)

Analysis of Host

Address of Host	Port/Service	Issue regarding Port
50.23.4.21	general/udp	Security notes found
50.23.4.21	general/tcp	Security notes found
50.23.4.21	nimcontroller (48000/tcp)	No Information
50.23.4.21	nimspooler (48001/tcp)	No Information
50.23.4.21	https (443/tcp)	Security warning(s) found
50.23.4.21	ssh (22/tcp)	Security notes found
50.23.4.21	http (80/tcp)	Security warning(s) found

© 2015 IBM Corporation


Redbooks Workshop

Figure 10-16 Verify security in the portal


Notes:

Your IaaS provider's portal could, like SoftLayer's does, also have the option to perform vulnerability scanning of your servers. This image shows part of the result of such a vulnerability scan performed on a server, and shows that there are seven warnings.

Further down you will see details of those warnings. It is up to you to decide if they should be acted upon. Remember that the vulnerability scan could suggest a fix for something that is needed for your server to work, so read the scan carefully and do not act blindly upon it.

This is a great way to monitor that your servers are safe, whether or not they are secure, and that the patch you applied has had the wanted effect.

10.20 Checkpoint

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Checkpoint

- Name three ways to harden the OS of a server.
- Why is it so important to secure the access to the portal?
- Which firewall is the easiest to configure from the portal?
- Why are vulnerability scans useful?


© 2015 IBM Corporation

Figure 10-17 Checkpoint


Notes:

(none)

10.21 Checkpoint (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services



Checkpoint

- Name three ways to harden the OS of a server.
Limit user access to the server, disable unneeded interfaces, and keep the OS updated.
- Why is it so important to secure the access to the portal?
The portal is your entrance to the data center and all the servers within.
- Which firewall is the easiest to configure from the portal?
The hardware firewall.
- Why are vulnerability scans useful?
They can help you detect vulnerabilities and verify that those vulnerabilities have been fixed after patching.

© 2015 IBM Corporation




Figure 10-18 Checkpoint (2)

Notes:

(none)



Monitoring cloud environments

This unit covers the methods you can use to monitor your SoftLayer cloud environment.

11.1 What you should be able to do

After completing this unit, you should be able to explain and describe the following topics:

- ▶ The cloud service access to the common cloud infrastructure
- ▶ The common infrastructure components that are monitored in a cloud infrastructure
- ▶ The common monitoring tools' benefits and limitations
- ▶ How the monitoring tools work in the cloud infrastructure
- ▶ The common monitoring alerts and responses
- ▶ The practical approach of SoftLayer monitoring:
 - Monitoring options
 - Monitoring response options
 - Notifications
 - Ordering a monitoring service
 - Building your own simple monitoring solution
 - Upgrading a monitoring package
 - Cancelling a monitoring package

11.2 References

The following links are useful for further research:

- ▶ Adding a monitor to a device in SoftLayer:
<http://knowledgelayer.softlayer.com/procedure/add-standard-monitor-device>
- ▶ SoftLayer server monitoring:
<http://www.softlayer.com/server-monitoring>
- ▶ Monitoring in the SoftLayer KnowledgeLayer:
<http://knowledgelayer.softlayer.com/topic/monitoring>
- ▶ Basic Monitoring - SoftLayer 101:
<https://www.youtube.com/watch?v=EtdPn1LAJHw#t=71>
- ▶ A guide to monitoring your IBM SoftLayer environment:
<http://wedowebosphere.de/blogpost/guide-monitoring-your-ibm-softlayer-environment>

11.3 Teaching topics

This unit describes the following topics:

- ▶ Cloud service models responsibilities and cloud service access to infrastructure
- ▶ Typical core infrastructure monitoring and typical infrastructure components monitored
- ▶ Why monitoring is important
- ▶ Typical monitoring tools limitations
- ▶ How monitoring tools work
- ▶ Typical monitoring alert
- ▶ Typical monitoring response
- ▶ Practical approach of SoftLayer monitoring:
 - Monitoring options
 - Monitoring response options
 - Notifications

- ▶ Ordering monitoring service
- ▶ Building your own simple monitoring solution
- ▶ Upgrade monitoring package
- ▶ Cancelling monitoring package

11.4 Typical service models responsibilities and typical service access to infrastructure

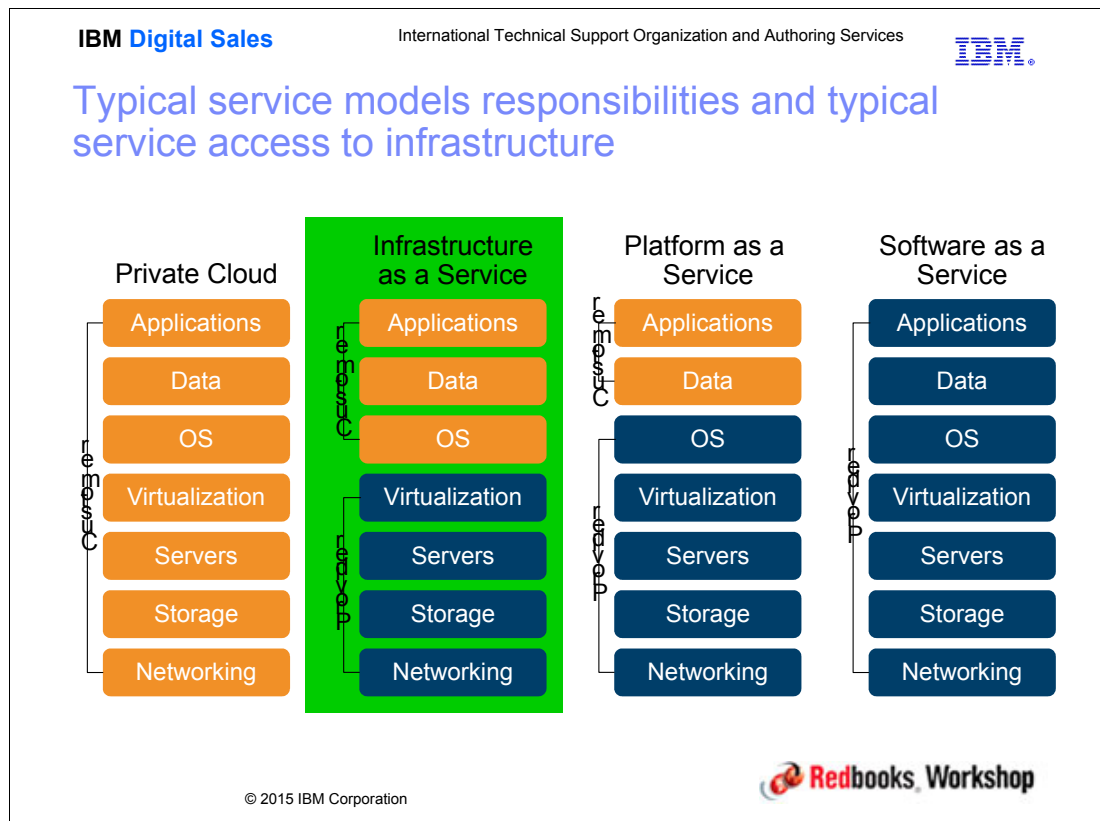


Figure 11-1 Typical service models responsibilities and typical service access to infrastructure

Notes:

A customer is usually responsible for all of the components in a traditional service model or its own private cloud. These responsibilities are shared at different levels in the IaaS, PaaS, and SaaS cloud service models.

An infrastructure as a Service (IaaS) model can be provided by a self-service IaaS or a fully managed IaaS cloud provider. In an IaaS self-service model, the cloud provider is responsible for data centers, which include virtualization, servers, storage, networking, and other infrastructure components. The provider is also responsible for monitoring these components. Customers have responsibility for monitoring the components delivered to them as part of the service. However, the scope of these responsibilities can be negotiated in the contract between customer and provider.

Monitoring options should be provided to the customers by an IaaS Cloud provider, but customers are not required to use these. They can install, configure, and manage their own monitoring software or service solutions.

11.5 Typical core infrastructure monitoring and typical infrastructure components monitored

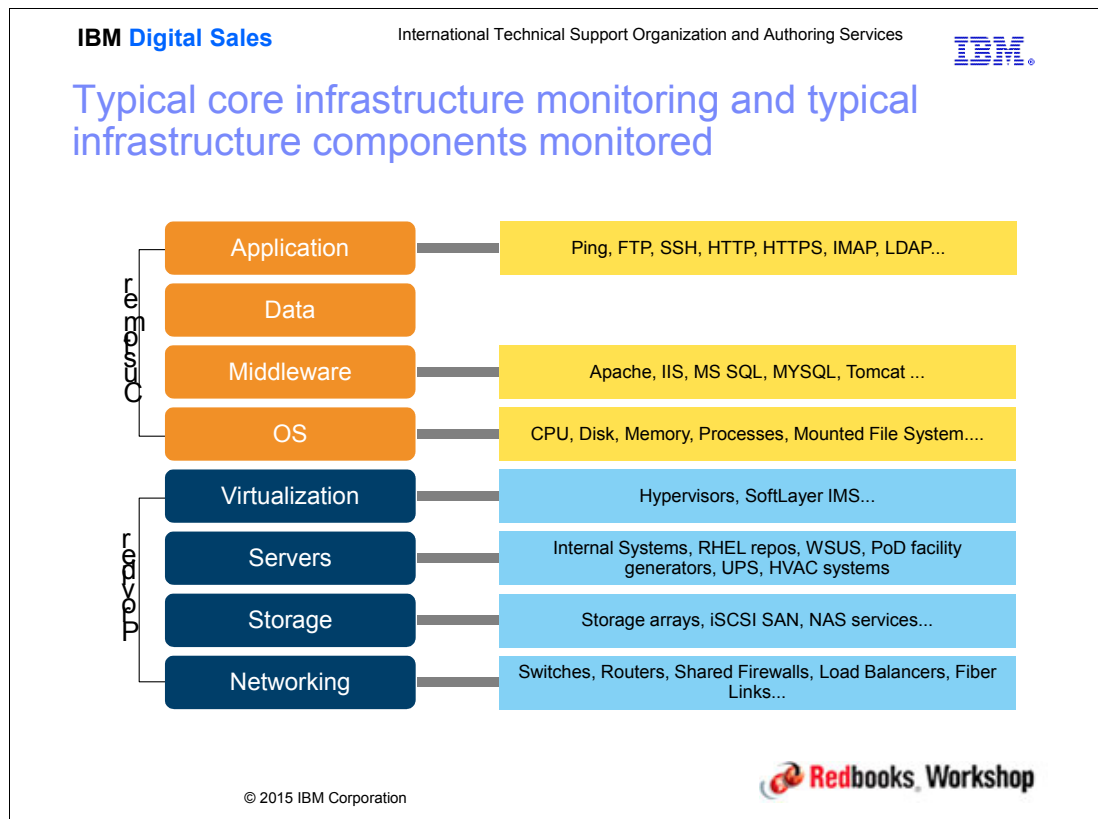


Figure 11-2 Typical core infrastructure monitoring and typical infrastructure components monitored

Notes:

In an IaaS service model, physical and virtualized components, applications, OS processes, and services are part of the cloud provider and customer infrastructure, and need to be monitored.

The monitoring services have various benefits that can help in building your business monitoring solution. The cost depends on your provider.

11.6 Why monitoring is important

IBM Digital Sales International Technical Support Organization and Authoring Services 

Why monitoring is important

Monitoring has these key benefits:

- Visibility into the cloud infrastructure
- Virtual environment management
- Capacity planning
- Optimization



© 2015 IBM Corporation 

Figure 11-3 Why monitoring is important

Notes:

In today's competitive market, a stable IT infrastructure is an important concern. Downtime can cause a business loss in productivity, affect quality of service, and damage a business' reputation.

Server monitoring is vital to help avoid outages and performance problems. It is a preventive measure that helps detect issues that can affect your productivity and foresee future problems.

Monitoring provides these benefits:

- ▶ **Visibility into the cloud infrastructure:** Visibility into the cloud infrastructure provides real-time information about your environment's health and performance. This information enables you to detect and isolate issues on your servers, storage, and network resources before they negatively affect your business.
- ▶ **Virtual environment management:** You can view critical information in easy to understand charts in a single dashboard to detect performance problems and identify their source. Resources can conveniently work with each other using automation.
- ▶ **Capacity Planning:** Allows you to understand the usage patterns and plan resource requirements in advance with the extensive reporting feature.
- ▶ **Optimization:** Find the most suitable value for a function within a given domain.

11.7 Typical monitoring tool limitations

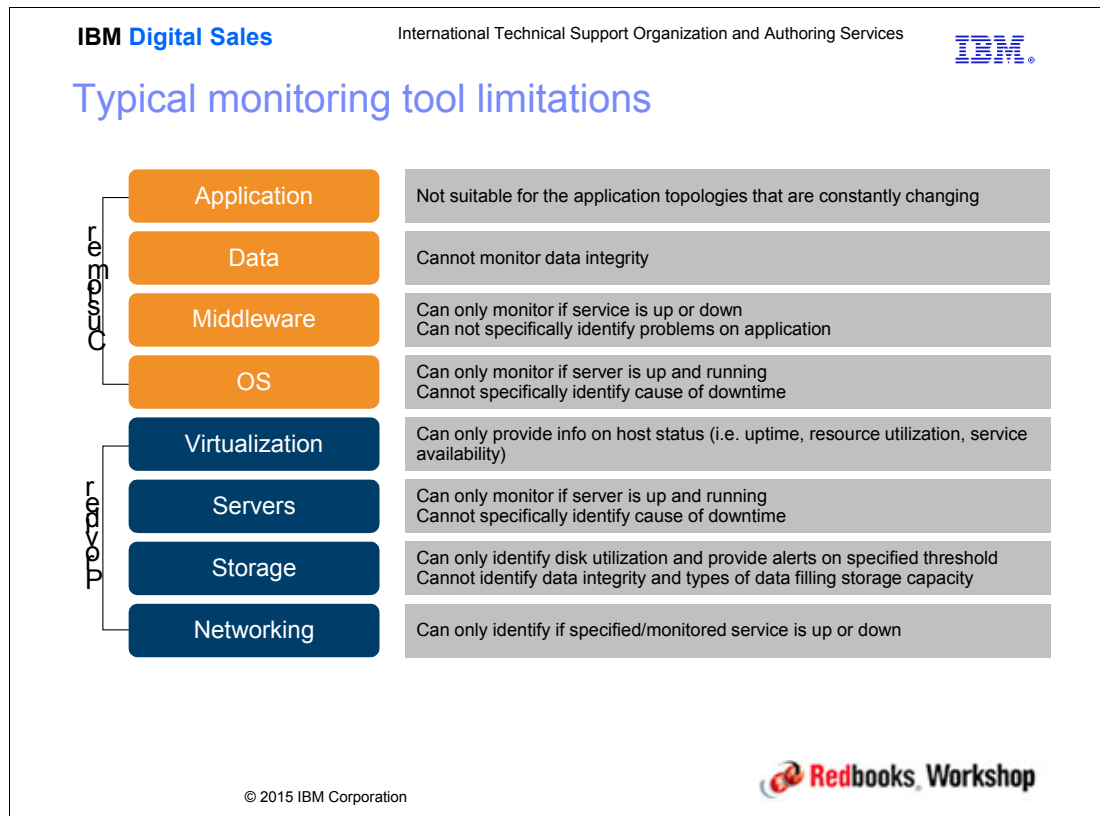


Figure 11-4 Typical monitoring tool limitations

Notes:

Many monitoring tools are available from cloud providers that provide stand-alone monitoring, monitoring as service, and full monitoring solutions with an equivalent cost.

11.8 How typical monitoring tools work

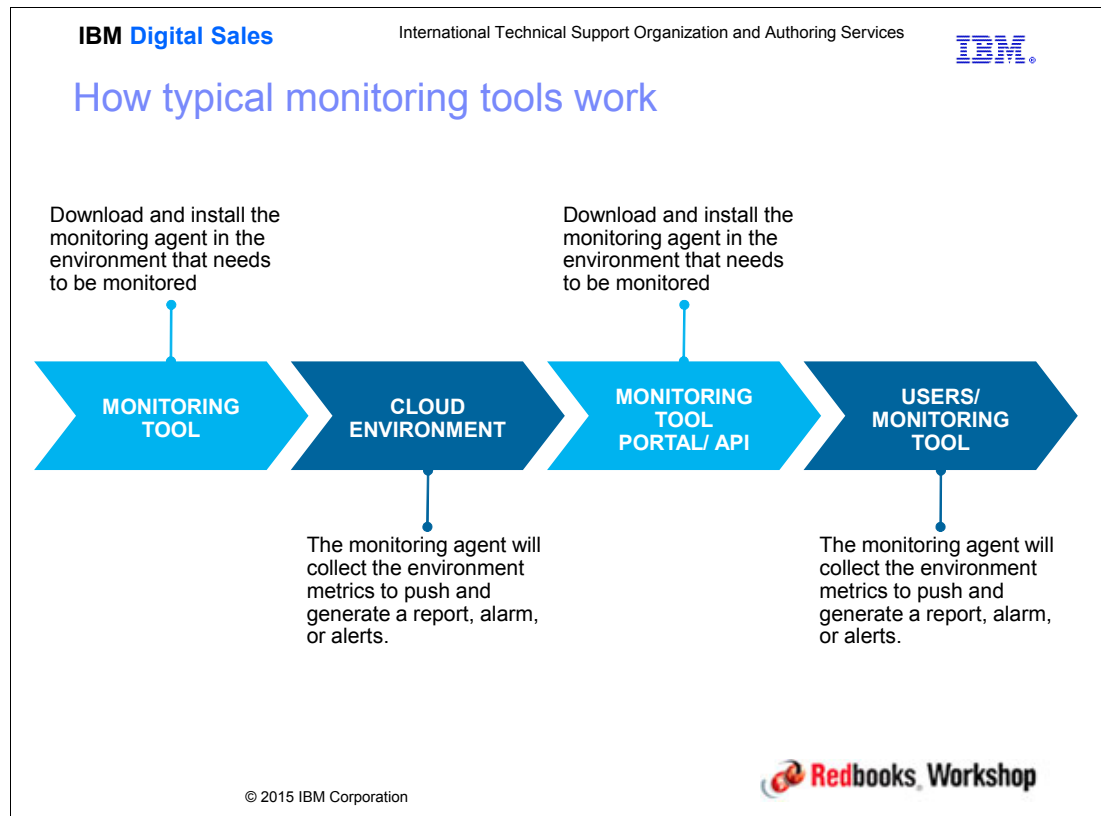


Figure 11-5 How typical monitoring tools work

Notes:

While trends for IaaS are emerging, IaaS providers generally not only offer portals into the operational health of the underlying infrastructure, but also offer APIs to the data that feeds these portals.

11.9 Typical monitoring alert



Figure 11-6 Typical monitoring alert

Notes:

Monitoring tools usually provide detailed graphing, and customizable alarms and alert notifications:

- ▶ **Graphical Report:** Provides a comprehensive visual depiction of your usage to understand the patterns and plan resource requirements in advance.
- ▶ **Customizable Alarms:** Will let you know when a service is outside an expected range. Alarms can be tracked from the monitoring portal, and also be configured to send email alerts.
- ▶ **Notifications:** Alerts can be tracked from the monitoring portals. IT might have an equivalent portal ticket and can be configured to send email notifications.

11.10 Typical monitoring response

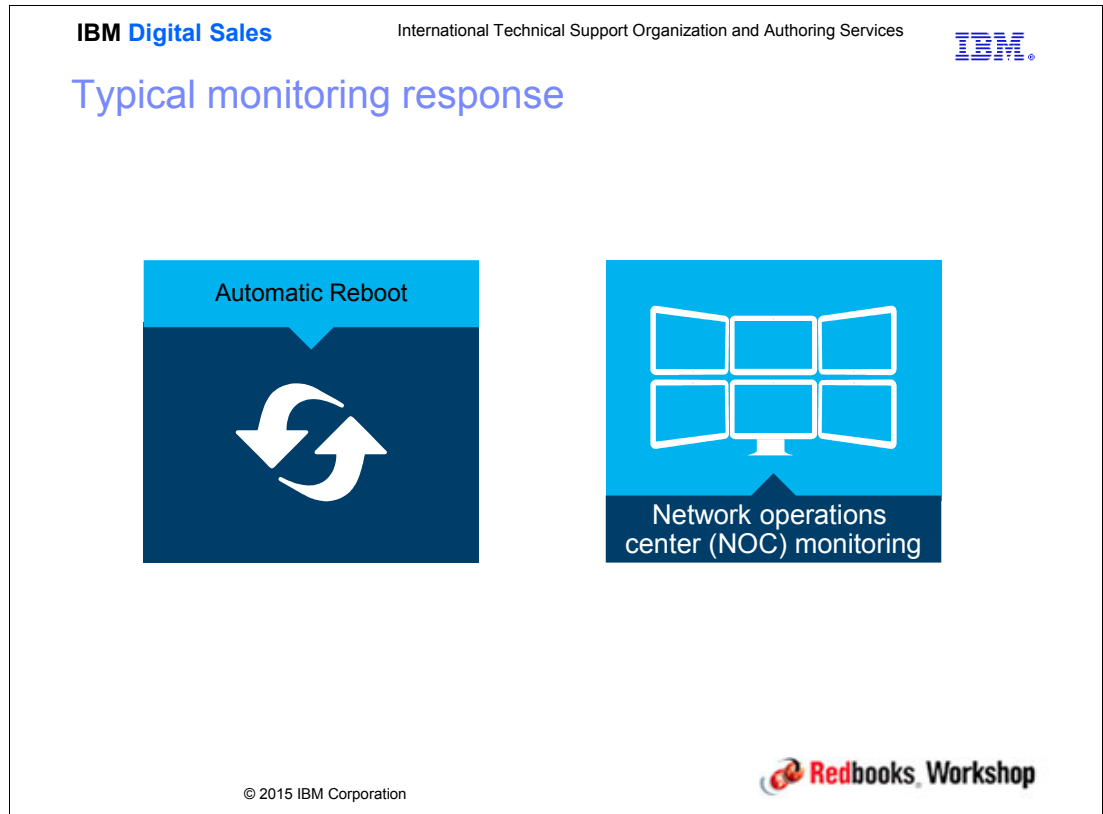


Figure 11-7 Typical monitoring response

Notes:

Automatic server reboot restarts the system if a failure is detected.

Network operations center (NOC) monitoring includes engineers who actively monitor your servers, and provide immediate response and personalized notification of alerts and failures.

Infrastructure monitoring consists of monitoring the servers, network, and the data center environment.

User experience monitoring simulates user behavior and activities to replicate problems and find the most effective solutions.

11.11 Practical approach to SoftLayer monitoring

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Practical approach to SoftLayer monitoring

SoftLayer Monitoring	
Monitoring Options	Monitoring Response Options
Standard Monitoring	Automated Server Reboot
Standard	24/7/365 NOC Monitoring
Host Ping + IPMI + Services (optional)	

W

Windows

L

Linux

W

Windows

L

Linux

Nimsoft Monitoring	
Basic	OS, CPU, MEM, DISK, Process & Services monitoring
Advanced	Basic Package + File System, Network Traffic, Network Time, DHCP, LDAP & SNMP data
Premium	Advanced Package + DNS, Email, IIS, MS SQL, MySQL, Tomcat & URL Response

! NOTIFICATIONS

Email/Ticket Notification
Immediate email/ticket alerts for any urgent issues requiring your attention.

Automated Customer Notification
Automated notification of order confirmations, payment reminders, ticket updates and scheduled maintenance.

© 2015 IBM Corporation


 Redbooks Workshop

Figure 11-8 Practical approach to SoftLayer monitoring

Notes:

An example of an IaaS provider offering monitoring is SoftLayer, which provides two monitoring services: Standard Monitoring and Nimsoft Monitoring to cover various monitoring needs with no extra cost. These monitoring services provide additional features with an extra charge that can help in building your business monitoring solution.

- ▶ Standard Monitoring is available for both physical and virtual servers at no extra charge, and provides basic host ping monitoring so you can set up notifications upon failure and based on Intelligent Platform Management Interface (IPMI) statistics. In addition to this, Standard Monitoring provides an option to monitor TCP service connections with an additional fee.
- ▶ Nimsoft Monitoring allows the monitoring of a wide variety of statistics on Windows and Linux servers. It consists of three levels (packages):
 - Basic (Hardware and OS): This package monitors your OS metrics, such as CPU, memory, disk, processes, and service for no extra charge.
 - Advanced (System Health): This package includes the basic package and has more components like file system, network traffic, and time, DHCP, LDAP, and SNMP data collection. You can choose between hourly and monthly billing.
 - Premium (Application): This includes the advanced package with extra components like DNS, email, IIS, MS SQL, MYSQL, Tomcat, and URL responses. You can choose between hourly and monthly billing.

Monitoring provides the following response options:

- ▶ Automatic server reboot restarts the system if failure is detected.
- ▶ NOC Monitoring includes engineers who actively monitor servers and provide immediate response and personalized notification of alert or failure:
 - Infrastructure monitoring covers servers, network, and the data center environment.
 - User experience monitoring simulates user behavior and activities to replicate problems and find the most effective solutions.

11.12 Ordering a monitoring service

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Ordering a monitoring service

IaaS providers might not provide an automatic monitoring tool in the environment. Customers should order and configure the monitoring package that they prefer to build.

1. During provisioning

ADVANCED MONITORING	HOURLY	SETUP
Monitoring Package - Basic	\$0.000	\$0.00
<input type="radio"/> None	\$0.000	\$0.00
<input checked="" type="radio"/> Monitoring Package - Basic	\$0.000	\$0.00
<input type="radio"/> Monitoring Package - Advanced	\$0.022	\$0.00
<input type="radio"/> Monitoring Package - Premium Application	\$0.044	\$0.00

MONITORING	HOURLY	SETUP
Host Ping and TCP Service Monitoring	\$0.007	\$0.00
<input checked="" type="radio"/> Host Ping and TCP Service Monitoring	\$0.007	\$0.00
<input type="radio"/> Host Ping	\$0.000	\$0.00

RESPONSE	HOURLY	SETUP
Automated Reboot from Monitoring	\$0.007	\$0.00
<input checked="" type="radio"/> Automated Reboot from Monitoring	\$0.007	\$0.00
<input type="radio"/> Automated Notification	\$0.000	\$0.00
<input type="radio"/> NOC Monitoring	\$0.070	\$0.00

2. Post-provisioning

SOFTLAYER-R

HELP CONTACT 2 NOTIFICATIONS 10 OPS

Devices Storage

Device List

Monitoring

Manage

Auto Scale

Account Summary

Current Balance:

Estimated Next Balance:

Device Name Public IP Private IP Location Basic Advanced

ah-dbsvrer.sa5003.sledst.ibm.c... 169.54.172.88 10.253.209.147 Dallas 6 1 Monitor Add Monitoring

© 2015 IBM Corporation

Redbooks Workshop

Figure 11-9 Ordering a monitoring service

Notes:

IaaS providers might not have an automatic monitoring tool installed on the environment. Order and configure the monitoring package that you prefer to build.

There are two ways to order a monitoring service on SoftLayer:

- Pre Provisioning: Monitoring can be selected to be added at the time of instance provisioning.
- Post Provisioning: Monitoring can be added on an already provisioned environment.

Unit 11. Monitoring cloud environments

303

11.13 Building your own simple monitoring solution

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Building your own simple monitoring solution

1. Select a device that you want to monitor.

SOFTLAYER-R[®]
an IBM Company

Account Summary

Current Balance:

Estimated Next Balance

Devices

Storage

Network

Monitoring

Manage

Auto Scale

Device Name	Public IP	Private IP	Location	Basic	Advanced
web01.company.softlayer.com	50.23.4.21	10.77.182.8	Houston 2	Monitor	Actions

Server Info

Add Monitor

IP Address

Monitor Type

Parameter

Notify?

Notify Wait

Add Monitor

Manage Notified Users

Users to Notify: Dela Cruz, Juan (juan.dela)

Add User

Doe, John

johndoecompanya

johndoecompanya@

Redbooks Workshop

Figure 11-10 Building your own simple monitoring solution

Notes:

Monitoring tools that are used by cloud providers usually provided a web portal and API services to help you easily design and build your monitoring solution. This example is from the SoftLayer IaaS provider. For more information, see the SoftLayer KnowledgeLayer at:

<http://knowledgelayer.softlayer.com/procedure/add-standard-monitor-device>

11.14 Upgrading a monitoring package

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Upgrading a monitoring package

1. Select **Modify Monitoring Package**.

Actions

View Agent Status

View Agent Reports

Configure Agent Reports

View Active Alarms

View Closed Alarms

Re-deploy All Agents

Modify Monitoring Package

Cancel Monitoring

2. Select the new Monitoring Package.

Current Monitoring Package: Monitoring Package - Basic

Select Monitoring Package

Monitoring Package - Advanced (System Health) Basic (Hardware and OS) Plus: [53.63 per month]

Network Traffic Monitoring Agent

DHCP Response Monitoring Agent

File and Directory Checking Agent

LDAP Response Monitoring Agent

Mounted File System Monitoring Agent

Network Time Protocol Response Monitoring Agent

SNMP Data Monitoring Agent

Performance Monitoring Agent

Monitoring Package - Premium Application Basic (Hardware and OS) & Advanced (System Health) Plus: [57.25 per month]

Apache Monitoring Agent

DNS Response Monitoring Agent

Email Response Monitoring Agent

JIS Monitoring Agent

MySQL Monitoring Agent

MySQL Monitoring Agent

Tomcat Server Monitoring Agent

URL Response Monitoring Agent

Configuration

Premium Monitoring Package - Linux

Continue Cancel

3. Confirm your order.

Confirm Order

Promo Code: Recalculate

Charges for this order

Charged To

One Time: \$0.00 Initial Payment is \$0.00, no payment required at this time.

Hourly: \$0.044

Tax: \$0.00

Total: \$0.04

Adjustments

Prorated Total: \$0.00

Initial Charge: \$0.00

Order Review

Item: Monitoring Package - Premium Application

Quantity: 1

☒ I have read the Master Service Agreement and agree to the terms therein.

Previous Place Order Cancel

© 2015 IBM Corporation

Redbooks Workshop

Figure 11-11 Upgrading a monitoring package

Notes:

Monitoring tools provided by your cloud provider might allow you to modify your current monitoring package. In SoftLayer, you can easily modify the monitoring package with just a few clicks.

11.15 Cancelling a monitoring package

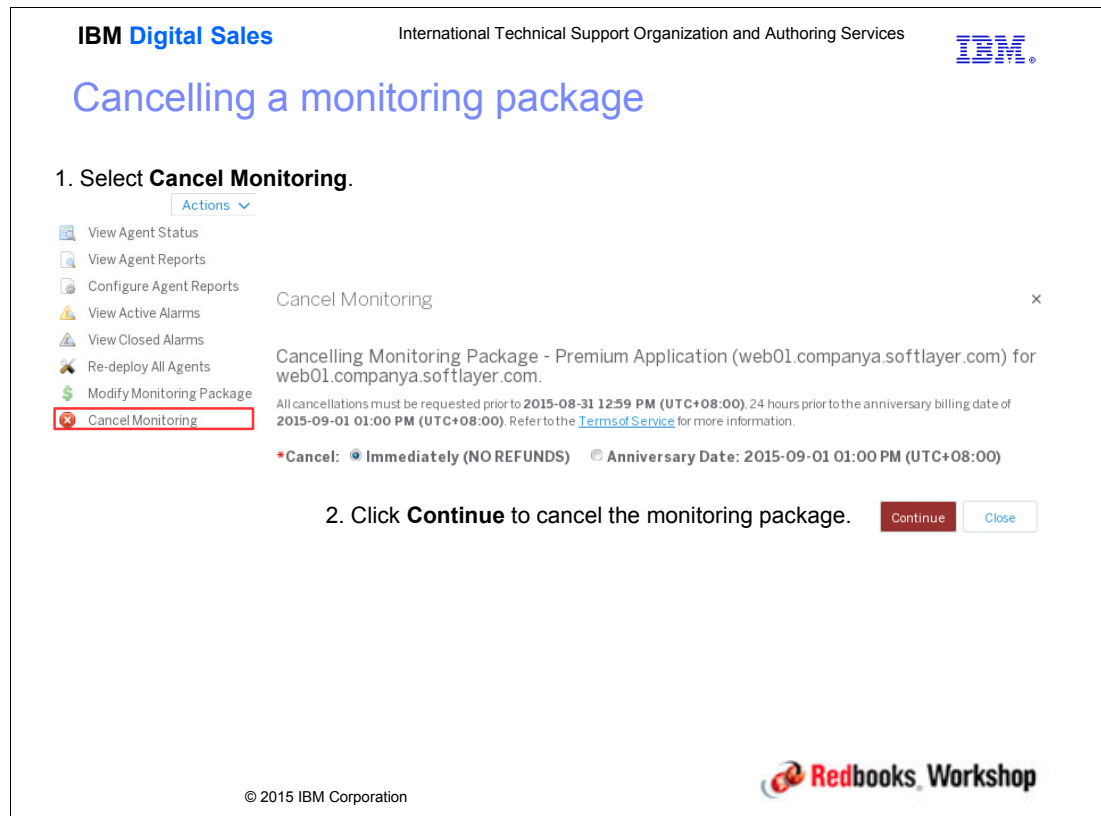


Figure 11-12 Cancelling a monitoring package


Notes:

Monitoring tools provided by the cloud providers might have the flexibility in canceling the monitoring package that you configured.

In SoftLayer, you can easily cancel a monitoring package with just a few clicks.


11.16 Checkpoint questions

IBM Digital Sales
International Technical Support Organization and Authoring Services



Checkpoint questions

- What is a key benefit of monitoring?
 - a) Capacity Planning
 - b) Increasing productivity
 - c) Optimization
 - d) a & c
- It is customer's responsibility to monitor the...
 - a) Switches, routers
 - b) Storage, SAN
 - c) CPU, disk, applications in their servers
 - d) All of the above
- True or false: Automatic server reboot is response to a monitoring alert?
 - a) True
 - b) False






© 2015 IBM Corporation

Figure 11-13 Checkpoint questions

Notes:


- ▶ What is a key benefit of monitoring?
Answer: D
- ▶ It is customer's responsibility to monitor the...
Answer: C
- ▶ Automatic server reboot is a type of monitoring alert?
Answer: True

11.17 Checkpoint questions (2)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

- What monitoring options are available on SoftLayer?
- The SoftLayer provider is based on which cloud service model?
 - a) PaaS
 - b) SaaS
 - c) IaaS



© 2015 IBM Corporation




Figure 11-14 Checkpoint questions (2)

Notes:

- What monitoring options are available on SoftLayer?
Answer: Standard Monitoring and Nimsoft Monitoring
- The SoftLayer provider is based on which cloud service model?
Answer: C



Automating cloud management with APIs

This unit covers the following topics:

- ▶ API values and usage
 - What is the purpose of having an API?
 - What is the value in using an API?
- ▶ The API implementation in SoftLayer
 - Types of APIs
 - Structure of APIs
 - Common usage scenarios

12.1 What you should be able to do

Upon completion of this unit, students should be able to do the following:

- ▶ Explain the concept of an API
- ▶ Explain the value of APIs
- ▶ Explain the SoftLayer API structure
- ▶ Describe the different types of SoftLayer APIs
- ▶ Explain how to use APIs for various tasks in SoftLayer

12.2 References

- ▶ SoftLayer API overview:
<http://sldn.softlayer.com/reference/overview>
- ▶ SoftLayer API structure:
<http://sldn.softlayer.com/reference/overview>
- ▶ Services reference:
<http://sldn.softlayer.com/reference/services/>

12.3 Teaching topics

This unit covers the following topics:

- ▶ Introducing the API concept
- ▶ Advantages of using an API
- ▶ API usage scenarios for cloud
- ▶ SoftLayer API overview
- ▶ The SoftLayer API structure
- ▶ Using the SoftLayer API

12.4 Introducing the API concept

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Introducing the API concept

- **What is an API?**
 - API stands for Application Programming Interface.
 - It is the interface by which an application program accesses a software component.
 - An API exposes an interface in terms of services, inputs, outputs, and underlying types.
 - An API exposes functionalities, or services, that are independent of their implementation.

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-1 Introducing the API concept

Notes:

One important characteristic of APIs is stability. The API functions as a black box of software services. Customers can use the API without understanding what is inside. Function implementation can vary if the interface exposed remains unchanged. There might be multiple implementations of the same function, and the implementation will evolve in terms of non-functional requirements, like performance, maintainability, and serviceability.

An implementation of API can add functions incrementally, refactoring on the way if necessary. This configuration allows the API structure to evolve rather than be implemented based on a static design.

APIs often come in the form of a library that includes specifications data structures, services, and parameters.

In some cases, such as SOAP and REST services, the specification includes only a definition of the remote calls exposed to the customers.

The trend in APIs is moving away from web services based on SOAP towards web resources based on Representational State Transfer (REST) and a Resource-Oriented Architecture (ROA).

12.5 API advantages

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

API advantages

- APIs enable automated management of a range of services.
- An API can facilitate integration of new features in existing applications.
- An API can facilitate common usage scenarios of otherwise distinct applications.

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-2 API advantages

Notes:

A global API is an important selling point for companies because you can call APIs for the offered services from anywhere.

12.6 API advantages (2)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

API advantages

- **Cloud solutions gain these benefits from APIs:**
 - Interoperability
 - The ability to implement a software-defined infrastructure, including scenarios like auto scale or load balancing
 - The ability to monitor and manage the infrastructure to meet budget constraints

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-3 API advantages (2)

Notes:

Different cloud solutions have distinct benefits and drawbacks. All solutions have interoperability in common. For example, CloudStack supports other cloud API models like AWS API, OpenStack API, and VMware vCloud API. Development of APIs for cloud enables applications to be designed and developed on the cloud. Such applications are not written for a static infrastructure, but for an infrastructure that is provided as a service and accessible through APIs.

12.7 Hybrid cloud scenarios that use the API economy

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Hybrid cloud scenarios that use the API economy

- Integration capabilities can exist on multiple service levels:
 - Software as a Service (SaaS) APIs – Application level.
 - Platform as a Service (PaaS) APIs – Service level.
 - Infrastructure as a Service (IaaS) APIs – Infrastructure level.
- A cloud broker service and API economy can be used to integrate the service levels into a coherent offering.

© 2015 IBM Corporation


Redbooks Workshop

Figure 12-4 Hybrid cloud scenarios that use the API economy

Notes:

The following service levels affect the APIs as follows:

- ▶ SaaS: Application level APIs can be used to extend applications and create additional modules that interact with the application.
- ▶ PaaS: APIs are used to access and interact with platforms (such as Java, NodeJS, Python, Ruby, and PHP), the application environment, the tooling (such as version control and build tools), and pipeline, as well as connect and configure services (such as SQL and NOSQL data stores).
- ▶ IaaS: APIs are used to provision, deprovision, and manage infrastructure level services and resources. Compute resources include CPU and memory allocation, I/O resources such as network and storage, and security services such as firewalls and intrusion detection systems.

Cloud broker APIs provide a unified view that uses an existing API economy across multiple cloud providers, platforms, and services to present a single solution. An example is Cloud Marketplace providers that present entire business applications that can use analytics, integration, backup, or resiliency services across multiple providers.

12.8 SoftLayer API overview

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

SoftLayer API overview

- SoftLayer provides Infrastructure as a Service (IaaS).
- The API implementation in SoftLayer provides direct interaction with SoftLayer's backend system.
- The functionality exposed by the API allows users to perform these tasks:
 - Perform remote server management.
 - Perform monitoring.
 - Retrieve information from SoftLayer's subsystems such as accounting, inventory, and DNS.

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-5 SoftLayer API overview

Notes:

The SoftLayer API (SLAPI) is available to all SoftLayer customers at no additional charge.

Object-oriented programming allows you to take full advantage of the capabilities offered in the SoftLayer API.

Use the API to automate tasks that would otherwise take more time and be prone to human errors.

12.9 SoftLayer API overview (2)

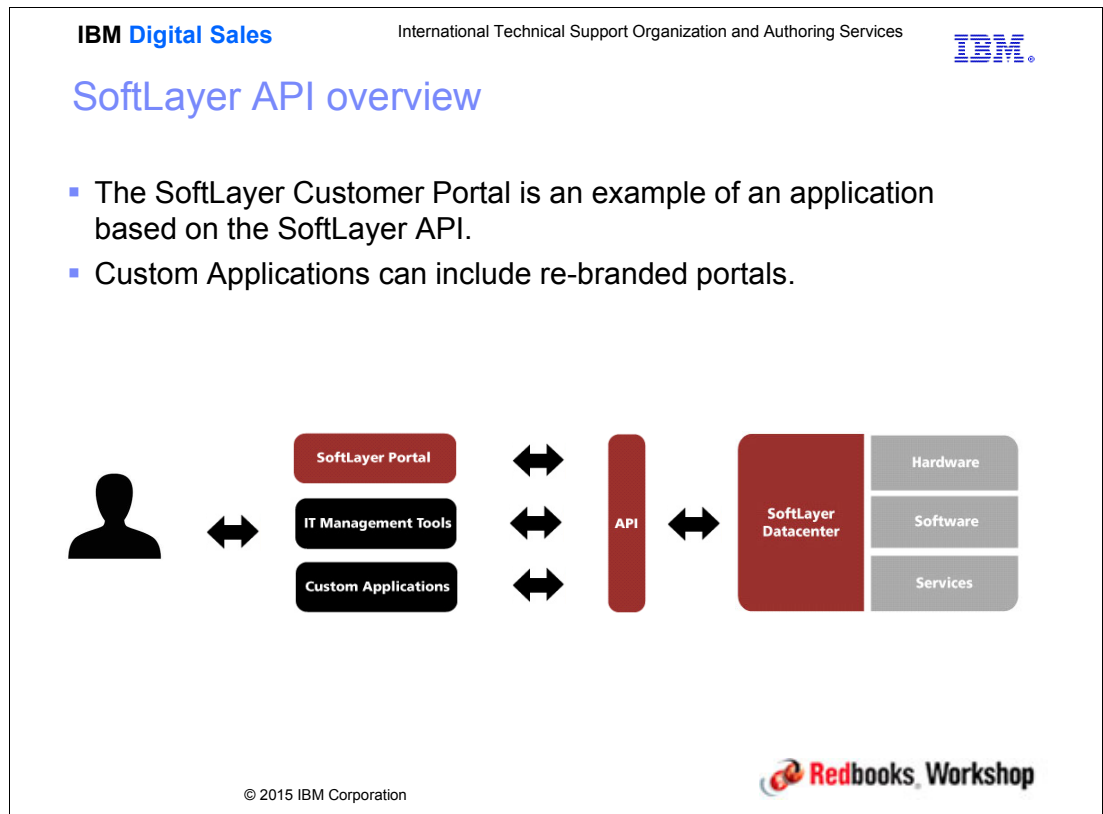



Figure 12-6 SoftLayer API overview (2)

Notes:

(none)

12.10 The SoftLayer API structure

IBM Digital Sales International Technical Support Organization and Authoring Services 

The SoftLayer API structure

SoftLayer has developed an Infrastructure Management System (IMS). The IMS orchestrates and automates the management of SoftLayer's offerings.

The offerings are exposed through the API library.



© 2015 IBM Corporation

Figure 12-7 The SoftLayer API structure

Notes:

(none)

12.11 The SoftLayer API structure (2)

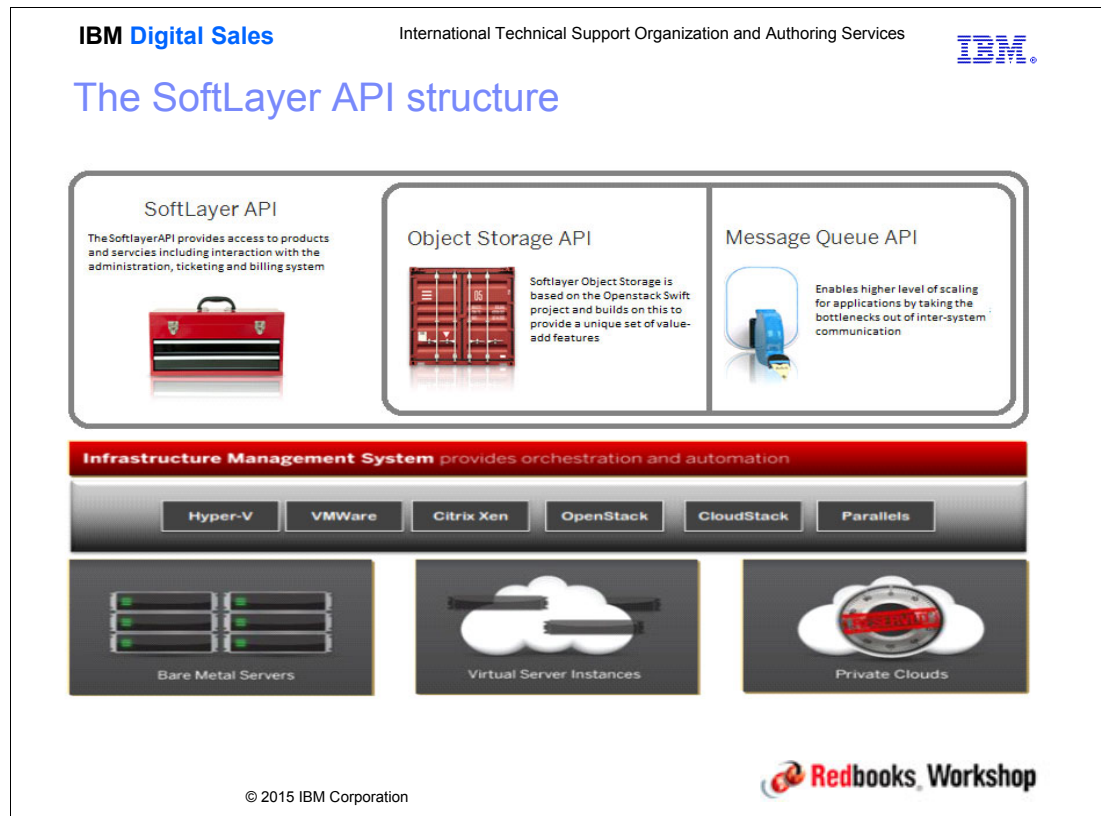


Figure 12-8 The SoftLayer API structure (2)

Notes:

SoftLayer implements a main API library, an Object Storage API, and a Message Queue API.

12.12 The main SoftLayer library

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

The main SoftLayer library

- Is implemented using SOAP protocol and XML-RPC.
- Is available as Representational State Transfer (REST).
- Is available through a command-line interface.
- Supports the following language clients: C#, Perl, PHP, Python, Ruby, and VB.Net.
- Additional languages are supported through the use of standard REST or SOAP clients.

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-9 The main SoftLayer library

Notes:


The main SoftLayer API library was written using two standards for developing web services: SOAP and XML-RPC:

- ▶ SOAP is a widely accepted standard for starting software services over computer networks, and passing structured data in and out of them.
- ▶ XML-RPC can be viewed as a simplified version of SOAP.

REST is a standard for starting software services over the Internet. REST advocates using HTTP or HTTPS protocols for starting web pages with operations such as GET and PUT to maintain and update remote resources over the Internet in a stateless manner. Stateless manner means one call does not know anything about other calls.

SoftLayer also provides a REST interface to their APIs. Simple REST calls can be shown in action through a simple REST client, or even through a web browser.

12.13 The Object Store API

IBM Digital Sales International Technical Support Organization and Authoring Services 

The Object Store API

- Is based on OpenStack Swift API.
- Is implemented using REST.
- Supports the following language clients: Java, Ruby, PHP, and Python.
- Supports hierarchical structures of containers and objects.
- Additional languages are supported through the use of standard REST clients
- Can be used with Content Delivery Network (CDN).


© 2015 IBM Corporation

Figure 12-10 The Object Store API

Notes:
(none)

12.14 The Message Queue API

IBM Digital SalesInternational Technical Support Organization and Authoring Services

The Message Queue API

- Is implemented using REST.
- Is able to contain any data in a message, plus message fields containing key and value pairs.
- Supports the following clients: C#, Java, Ruby, PHP, and Python.
- Supports topics, topic tags, and subscriptions.



© 2015 IBM Corporation

Figure 12-11 The Message Queue API

Notes:

(none)

12.15 Using the SoftLayer API

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using the SoftLayer API

- The SoftLayer API is structured into two hierarchies.
- There is a separate hierarchy each for services and methods:
 - Services are functional groupings of API calls.
 - Methods are specific actions that can be performed on a service.
- Data types are elements describing the offering and allowing its manipulation.

© 2015 IBM Corporation

Redbooks Workshop

Figure 12-12 Using the SoftLayer API

Notes:

The normal scenario in an object-oriented environment is to first create an instance of the SoftLayer_Client type by authenticating with a user name and an API key, and use that to access the various services. Each service provides methods either for changing the infrastructure, in which case the correct data types need to be provided as parameters, or for retrieving information that is then provided as instances of the specific data types. The object returned contains all the local properties if there is no mask provided and only the relational properties provided through a mask field.

12.16 Using the SoftLayer API: Services and methods

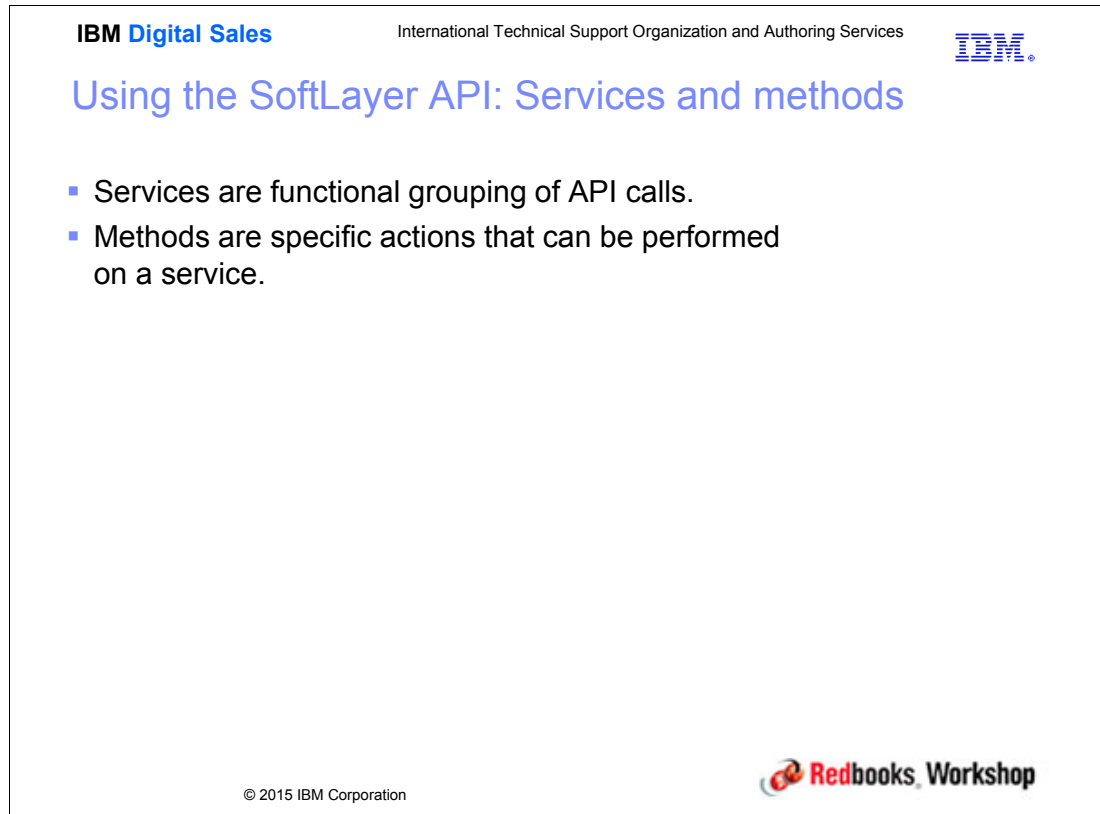


Figure 12-13 Using the SoftLayer API: Services and methods

Notes:

A service is an endpoint associated with internal SoftLayer systems. Each service is a collection of methods, or actions, which can be performed. All SoftLayer services begin with “SoftLayer_” and contain more terms that define the general function that the service provides such as “Hardware”, “Account”, “Billing”, and “Network”. Each service is extended from there with a name that defines the service's specific function within that particular subset. Each service that is associated with the SoftLayer API has a unique name. While some services, such as SoftLayer_Account and SoftLayer_Account_Address, can share a common prefix, their interaction is not necessarily similar. There is no direct inheritance for services of a similar name. Because of this, each service should be approached individually.

While each service offers a unique set of methods, many services offer the getObject method. These methods can be used to retrieve an object of the same type from the API. For example, calling the getObject method on the SoftLayer_Network_Subnet service returns a SoftLayer_Network_Subnet data type object.

A method is a specific action that can be performed for a SLAPI service. Each method returns a scalar or structured data type, and might require specific parameters, permissions, or headers to run. Method parameters should be passed by using the techniques described in each language's or endpoint's documentation. In situations where multiple parameters are required, pass the parameters in the order that they are expected to be received.

12.17 Using the SoftLayer API: Data types

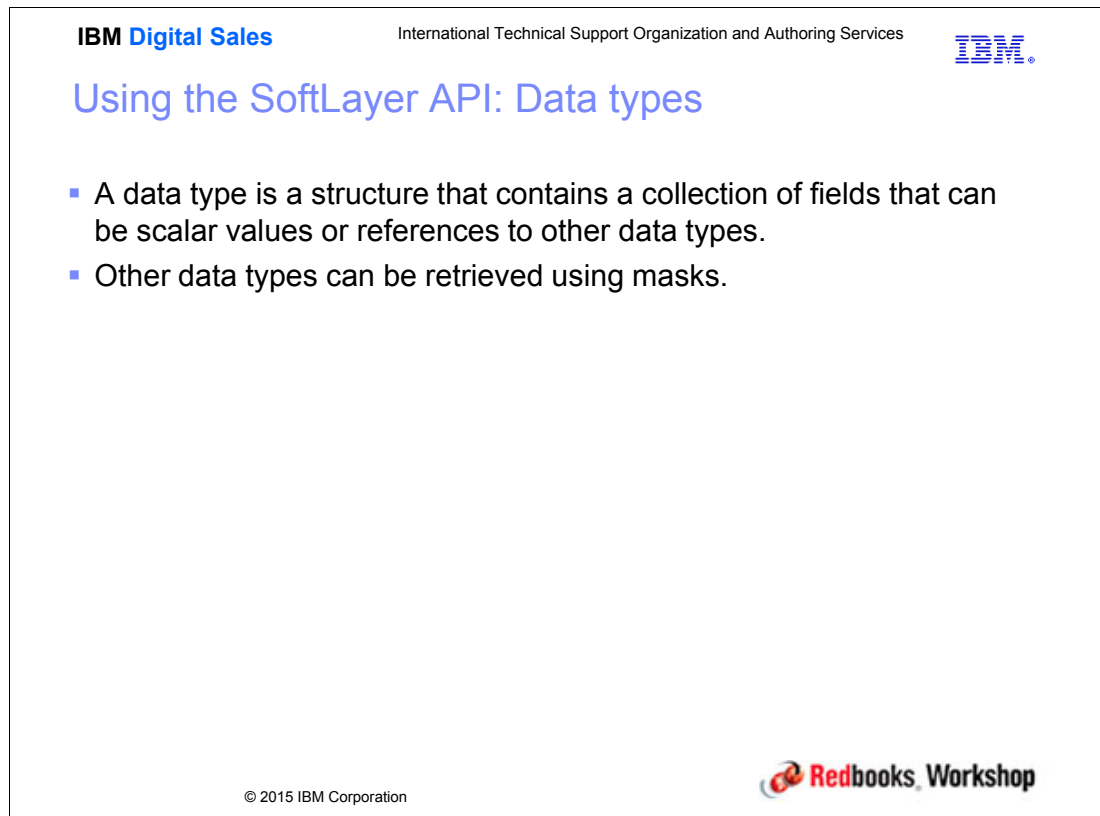


Figure 12-14 Using the SoftLayer API: Data types

Notes:

A *data type* is a structure that contains a collection of scalar values and other data types. In addition to traditional scalar values such as string, bool, and int, the SLAPI also uses complex data types that contain properties that define the objects passed to and returned by methods in the API. Each data type potentially contains a number of local, relational, and count properties.

A *local property* is a direct child of a data type. Local properties are typically returned when getObject() is called. Some local properties are required when creating an instance of this data type when calling createObject().

A *relational property* is an indirect child of a data type. Relational properties are defined in other data types or their properties. For example, the SoftLayer_Account data type has a relational property for hardware. This relational property is an array of SoftLayer_Hardware data types. When tapped with an object mask, this property returns an array that contains a SoftLayer_Hardware object for each hardware device on the account.

A *count property* is a convenience property that can be used to determine the total number of objects that are associated with a property. For example, you can retrieve the total number of VLANs associated with a specific server by using an object mask with included Softlayer_Hardware_Server networkVlanCount.

12.18 Using the SoftLayer API: Service hierarchy sample

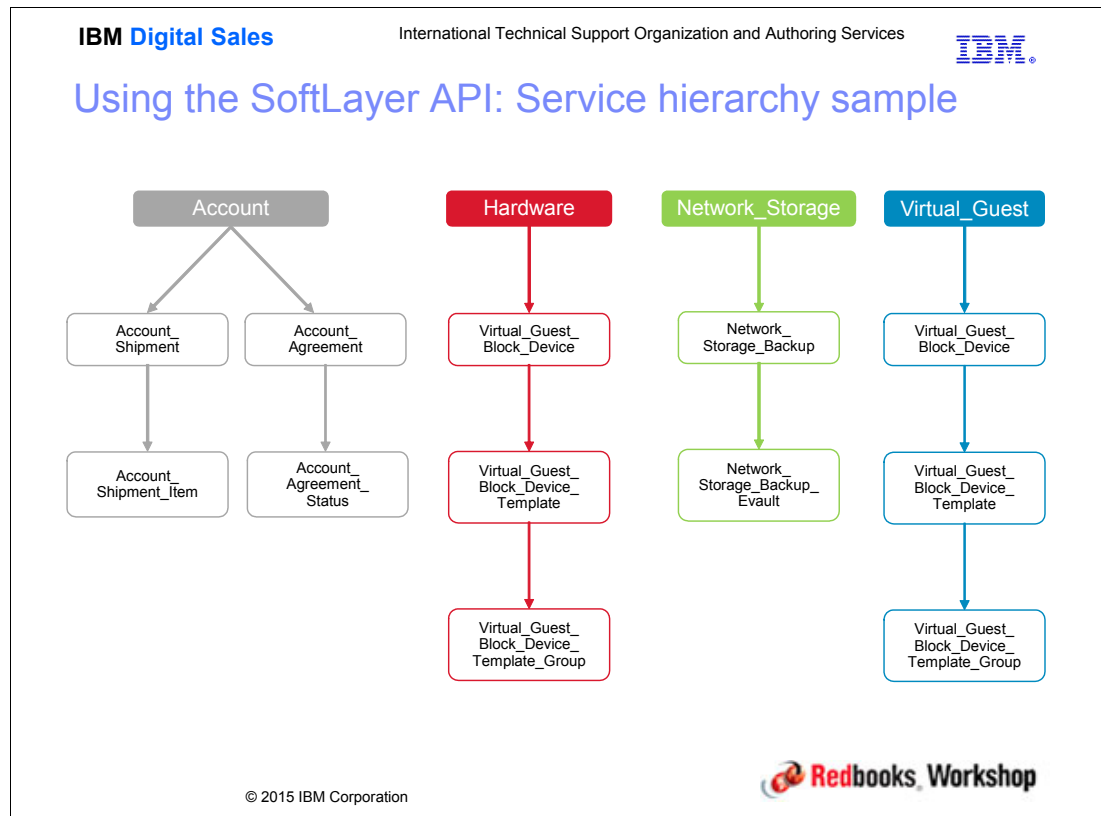


Figure 12-15 Using the SoftLayer API: Service hierarchy sample

Notes:

The SoftLayer API has these service groups, among others:

- ▶ User accounts
- ▶ User billing
- ▶ Virtual server management
- ▶ Hardware management (dedicated/bare metal and other hardware)
- ▶ Product ordering
- ▶ Configuration templates
- ▶ Software components
- ▶ Locations
- ▶ Network (firewalls, gateways, load balancers, subnets, and VLANs)
- ▶ Storage (iSCSI, NAS, and backup)
- ▶ Reboots and reloads
- ▶ Ticketing
- ▶ DNS
- ▶ Security (certificates, keys, and scans)
- ▶ Monitoring
- ▶ Portal customization
- ▶ Auxiliary functions

12.19 Using the SoftLayer API

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using the SoftLayer API

- The SoftLayer API allows developers to define what data is returned from each call through the use of headers.
- The headers allow control over the amount of data returned by the API:
 - Result limits.
 - Define an offset and amount of objects to return. This allows for pagination in the case of large data sets.
 - Object masks.
 - Specify which local properties to return from a method.

© 2015 IBM Corporation


Redbooks Workshop

Figure 12-16 Using the SoftLayer API

Notes:

In addition to the typical create, read, update, and delete actions, the SLAPI allows developers to define how data is returned from each call by using special API call headers. These headers allow an extra level of control over the amount of data returned by the API.

A result limit is a support method that allows you to define an offset and number of objects to return. This methods allow pagination of large data sets.

An object mask allows the user to specify which local properties to return from a method and retrieve information found in both relational and count properties. A map, or “mask” is created to define the specific data to include in the return value. For example, it is possible to gather the IDs for each VLAN on a SoftLayer_Hardware_Server by specifying an object mask for “networkVlans.id” when calling SoftLayer_Hardware_Server::getObject.

Object filters can be used to limit the results that are returned by the API. They differ from object masks in that they determine what data type objects are returned while Object Masks define what properties to retrieve from the returned objects.

12.20 Using the SoftLayer API (cont.)

IBM Digital Sales

International Technical Support Organization and Authoring Services

IBM

Using the SoftLayer API (cont.)

- Retrieve information found in both relational and count properties
- Filters
- Define criteria that must be met by all the objects retrieved

© 2015 IBM Corporation


Redbooks Workshop

Figure 12-17 Using the SoftLayer API (cont.)

Notes:


(none)

12.21 Example scenarios of using an API

IBM Digital SalesInternational Technical Support Organization and Authoring Services

Example scenarios of using an API

- Invoke selected services for white label service providers to implement their own rebranded portals.
- Programmatic upscaling and downscaling in a public or private cloud (horizontal scaling).
- Handle cloud monitoring events, such as re-instantiating servers, rebooting, and OS reloads.
- Programmatic cloud management, including upgrading and downgrading (vertical scaling), adding storage, backing up, and restoring.
- Write cloud-native software applications.




© 2015 IBM Corporation

Figure 12-18 Example scenarios of using an API

Notes:
(none)

12.22 Checkpoint questions

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

1. What does API stand for?

- a) Advanced Programming Interface
- b) Accessible Programming Interface
- c) Advanced Procedural Interface
- d) Application Programming Interface



© 2015 IBM Corporation 

Figure 12-19 Checkpoint questions

Notes:

(none)

12.23 Checkpoint questions (2)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

2. Which of the options below is an advantage of using an API from a customer perspective?

- a) The organization is more agile.
- b) There is less need for project management.
- c) There is more automation.
- d) It is a key selling point.




© 2015 IBM Corporation

Figure 12-20 Checkpoint questions (2)

Notes:

(none)

12.24 Checkpoint questions (3)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

3. What is a desirable characteristic of a cloud API?

- a) Complex structure
- b) Interoperability
- c) Restrict HTTP calls
- d) Horizontal scaling



 © 2015 IBM Corporation

Figure 12-21 Checkpoint questions (3)

Notes:
(none)

12.25 Checkpoint questions (4)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

4. What is the relation between a service and a method in SoftLayer's API implementation?

- a) A method can service any number of objects.
- b) A method is a specific action that can be performed for a service.
- c) A method is one of multiple ways of performing a service.
- d) A method has multiple services.



 © 2015 IBM Corporation

Figure 12-22 Checkpoint questions (4)

Notes:
(none)

12.26 Checkpoint questions (5)

IBM Digital Sales International Technical Support Organization and Authoring Services 

Checkpoint questions

5. What is the difference between an object mask and a filter in the SoftLayer API?

- a) A filter can be reused.
- b) A filter is shorter than an object mask.
- c) A filter is executed on the client side, whereas the object mask is executed on the server side.
- d) A filter restricts the number of objects returned, whereas an object mask restricts the number of fields returned for each object.


© 2015 IBM Corporation 


Figure 12-23 Checkpoint questions (5)

Notes:

Answers:

1. d
2. c
3. b
4. b
5. d

12.27 Recap

IBM Digital Sales International Technical Support Organization and Authoring Services 

Recap

Students should now be able to:

- Explain the concept of an API.
- Explain the value of APIs.
- Explain the SoftLayer API structure.
- Describe the different types of Softlayer APIs.
- Explain how to use APIs for various tasks in SoftLayer.


© 2015 IBM Corporation 

Figure 12-24 Recap

Notes:

(none)

Glossary

ACL Access Control List

API (application programming interface) An interface that specifies how software components should interact with each other

Appliance firewall Custom OS with software to filter traffic

Audit log An automatically created report of who has done what and when

Autoscale The automated scaling option that is provided by SoftLayer for virtual servers

B&R Backup and recovery

Bare metal server In SoftLayer, a physical machine that is made available to the customer as part of a cloud offering

CHAP Challenge-Handshake Authentication Protocol

CIFS Common Internet File System

Citrix NetScaler VPX A multifunctional network and security device offered by Citrix and available in SoftLayer as a service. It provides load balancing, firewall functions, data compression, and more

Client to site VPN VPN connection between one user device and a site

Cloud Computing Model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction

DAS Direct-attached storage

DHCP (Dynamic Host Configuration Protocol) A standardized network to allocate IP addresses to computers

DNS (Domain Name System) Used to resolve human-readable host names into IP addresses

DR Disaster recovery

Firewall Device to filter traffic

Flex Image The SoftLayer platform-neutral imaging system

Global Load Balancing Refers to specially designed DNS server that load balances DNS requests

HA (high availability) A term that refers to a system or component that is continuously operational by avoiding single point of failures

Hardware firewall Physical device to filter traffic

HDD Hard disk drive

Horizontal scaling Upscaling or downscaling by adding or removing servers from an environment

HVAC Heating, ventilating, and air conditioning

Hybrid cloud A configuration of cloud IaaS resources that mixes public and private cloud assets

Hypervisor A piece of computer software, firmware, or hardware that creates and runs virtual machines. Also known as virtual machine monitor (VMM)

IaaS (infrastructure as a service) Cloud service model in which the consumer can provision fundamental computer resources such as processors, storage, and networking resources

IMS (Infrastructure Management System) A software platform for managing the infrastructure developed by SoftLayer

In-Band Management A way to manage computer devices locally through the network itself, using a telnet/SSH connection to them

Intelligent Platform Management Interface (IPMI) A standardized interface that is used by system administrators as a way to manage a computer that is powered off or otherwise unresponsive. It connects to the hardware rather than to an operating system or login shell

Internet Information Services (IIS) A web server created by Microsoft

IP (Internet Protocol) Address A numeric label that is assigned to each device that communicates over the IP protocol. The IP address space is split into public and private spaces. Comes in two versions: v4 (traditional, currently running out of public addresses), and v6 (new standard to accommodate much larger address space)

IPSec (Internet Protocol Security) VPN A suite of protocols that are designed to authenticate and encrypt all IP traffic between two locations

IQN iSCSI qualified name

iSCSI Internet Small Computer System Interface

Kerberos A computer network authentication protocol that works based on tickets to allow nodes communicating over a non-secure network to prove their identity to one another in a secure manner

LDAP (Lightweight Directory Access Protocol) A protocol for accessing and maintaining distributed directory information service

Load balancer A device that distributes network or application requests over a number of servers. A client connects to VIP and the load balancer then forwards the traffic to one of the servers

Local load balancing Refers to load balancer distributing network or application requests locally (in the same network segment)

LUN Logical unit number

Multi-tenant An architecture in which a host machine supports cloud infrastructure for multiple customers

NAS Network-attached storage

NAT (network address translation) Allows a router to modify packets to allow multiple devices to share a single IP address

Network handoff Switching between different network providers

Network hop Each device that the traffic passes on its way to its destination is called a hop

NFS Network File System

NIC (Network Interface Card) A computer circuit board that is installed in computers so that it can be connected to a network

NOC (network operations center) A team that monitors network performance and security 24x7

On-demand self service A service that can be unilaterally requested and acquired by the customer when needed without requiring assistance from the service provider

OpenStack A free and open source cloud computing software platform

OS firewall Application installed on or embedded in the operating system to filter traffic

OSI (open systems interconnection) model A conceptual model that characterizes and standardizes communication of computing systems. It is divided into seven layers of abstraction

Out-of-Band Management A way to manage computer devices using a dedicated channel that is available even when your device is not functioning properly

PaaS (platform as a service) Cloud service model in which the computing platform and solution stack are made available as a service. Customers can develop, test, and deploy their applications on the cloud

Pay-per-use A billing model that monitors usage of a service and charges only for the amounts used. Service usage is monitored, controlled, and reported, providing transparency for both the provider and user

Ping A networking utility to verify that one server can be reached from another

PPTP (Point-to-Point Tunneling Protocol) VPN An implementation of VPN using a specialized client software to establish secure tunnel from the user workstation to a private network

Private cloud A single tenant environment that is dedicated to a single customer that uses virtualization tools to deliver a virtual computing platform

Private Image Customer-created image templates in SoftLayer

Private IP address An IP address within the private address space that allows organizations to create their own private networks

Private VLAN A VLAN that is isolated from public Internet, and uses a private IP address space. It can be accessed through a VPN connection

Provisioning Cloud provisioning is the allocation of a cloud provider's resources to a customer

Provisioning Script Executable file that is automatically deployed by SoftLayer after provisioning and that can be run automatically or manually

Public cloud A multi-tenant environment that uses virtualization tools to deliver a virtual computing platform to multiple customers

Public Images Image templates that are optimized and provided by SoftLayer

Public IP address An IP address within the public address space that can be accessed over the Internet

Public VLAN A VLAN that can be connected to from the public Internet that uses a public IP address space

RAID Redundant Array of Independent Disks

REST (Representational State Transfer) A protocol that is used for communicating with web resources

RPO Recovery point objective

RTO Recovery time objective

SaaS (software as a service) Cloud service model in which software or applications are provided to different customers, or consumers through a network, usually the Internet

SAN Storage area network

SAS Serial-attached SCSI

SATA Serial Advanced Technology Attachment

SCSI Small Computer System Interface

Single tenant An architecture in which a host machine supports cloud infrastructure for a single customer

Site to site VPN Direct VPN line between two sites that allows devices on both sites to communicate securely

SMB Server Message Block

SNMP (Simple Network Management Protocol) A standard protocol for managing devices on IP networks

SOAP A protocol that is used for web services

SSD Solid-state drive

SSH (Secure Shell) A cryptographic network protocol for secure data communication, remote command line login, remote command execution, and other secure network services between a pair of client and server systems

SSH keys A pair of files that contain encryption keys (public and private) for two-key authentication supported by the SSH protocol

SSL (Secure Sockets Layer) VPN An implementation of VPN that allows a user to create a secure tunnel from that user's workstation to a private network such as the SoftLayer Private Network

Subnet A logical grouping of connected network devices for faster data transfer and easier management

Subnet mask A mechanism to split a network into subnets

Vertical scaling Upgrading or downgrading the existing servers

VIP (Virtual Internet Protocol Address) An IP address that is not a primary of a physical network card. In the case of load balancing, the VIP usually refers to the address that the client receives through DNS and connects to

Virtual server An instance of virtual hardware platform that is created and run by using virtualization techniques

Virtualization In computing, recreating a virtual (rather than actual) version of something, including but not limited to a virtual computer hardware platform, operating system (OS), storage device, or computer network resources

VLAN (virtual local area network) A logical grouping of network nodes that is configured as though they were in the same LAN even if they are in separate ones

VLAN trunking Enables the movement of traffic to different parts of the network configured in a VLAN

VPN (virtual private network) A network that is constructed by using public networks (usually the Internet) to connect to a private network such as the SoftLayer Private Network

Vulnerability scan A scan of one or more devices to search for security issues in the configuration or OS/software of a device

Vyatta appliance A multifunctional network device offered by Brocade and available in SoftLayer as a service. It provides software-based virtual network gateway, virtual firewall, and VPN capabilities

Workload Independent service or collection of code that can be run

XML-RPC A Remote Procedure Call (RPC) protocol that uses XML to encode its calls and HTTP as a transport mechanism

Redbooks

A Practical Approach to Cloud IaaS with IBM SoftLayer: Presentations Guide

SG24-8350-00

ISBN 0738441562



(0.5" spine)

0.475" <-> 0.873"

250 <-> 459 pages



SG24-8350-00

ISBN 0738441562

Printed in U.S.A.

Get connected

