CPET 581 Cloud Computing: Technologies and Enterprise IT Strategies

Lecture 6

Cloud Platform Architecture over Virtualized Data Centers

Part -2:

Data-Center Design and Interconnection Networks & Architecture Design of Compute and Storage Clouds

Text Book: Distributed and Cloud Computing, by K. Hwang, G C. Fox, and J.J. Dongarra, published Elsevier/Morgan Kaufmann, 2012.

Spring 2015

A Specialty Course for Purdue University's M.S. in Technology Graduate Program: IT/Advanced Computer App Track

Paul I-Hai Lin, Professor

Dept. of Computer, Electrical and Information Technology

Purdue University Fort Wayne Campus

Ch. 4 - Topics of Discussion

- Cloud Computing and Service Models
- **Data-Center Design and Interconnection Networks**
- **Architectural Design of Compute and Storage Clouds**
- Public Cloud Platforms: Google App Engine, Amazon Web Services and Microsoft Window Azure
- Inter-Cloud Resource Management
- Cloud Security and Trust Management

4.2 Data-Center Design and Interconnection Networks

- Warehouse-Scale Datacenter Design
- Datacenter Interconnection Networks
- Modular Data Center in Shipping Containers
- Interconnect of Modular Data Centers
- Data Center Management Issues

The Architecture of a Small Server Cluster (~ 1000 servers)

interconnected by an Ethernet switch and housed in a warehouse or in a container environment

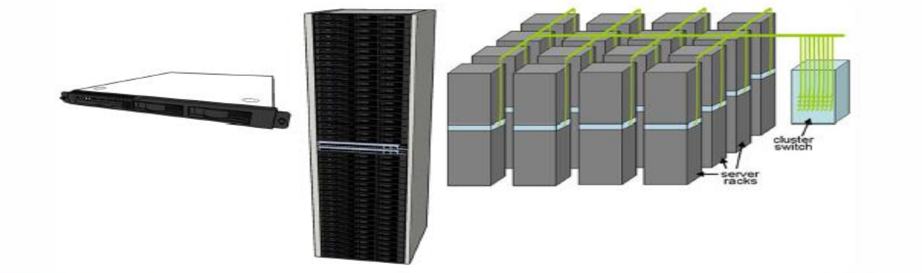

FIGURE 1.1: Typical elements in warehouse-scale systems: 1U server (left), 7´ rack with Ethernet switch (middle), and diagram of a small cluster with a cluster-level Ethernet switch/router (right).

(Courtesy of Luiz Barroso and Urs Holzle, Google Inc., 2009)

Warehouse-Scale Computer and Datacenter (WSC)

- Provides Internet services
  - Search, social networking, online maps, video sharing, online shopping, email, cloud computing, etc.
- Differences with HPC "clusters":
  - Clusters have higher performance processors and network
  - Clusters emphasize thread-level parallelism, WSCs emphasize request-level parallelism
- Differences with datacenters:
  - Datacenters consolidate different machines and software into one location
  - Datacenters emphasize virtual machines and hardware heterogeneity in order to serve varied customers

(Courtesy of Hennessy and Patterson, 2012)

Typical Datacenter Layout

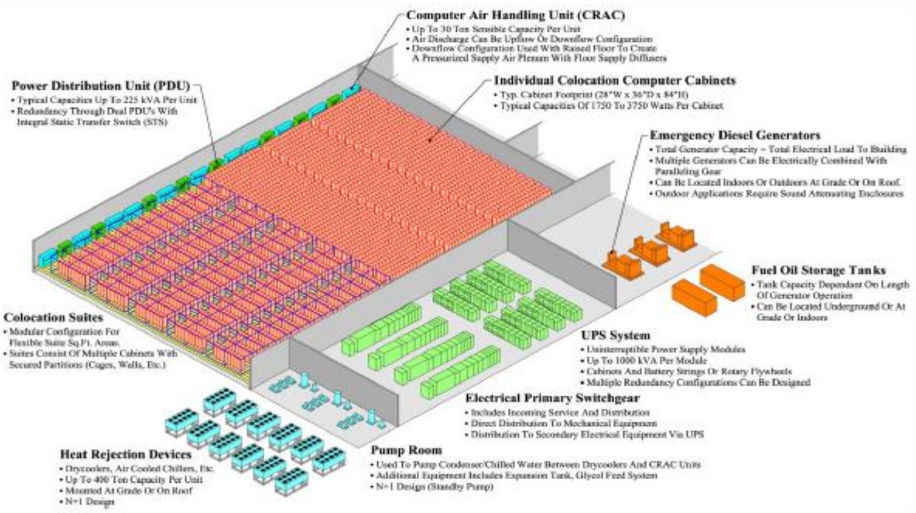

FIGURE 4.1: The main components of a typical datacenter (image courtesy of DLB Associates [23]).

Example 4.8 A huge datacenter that is 11 times the size of a football field, housing 400,000 to 1 million servers.

Power Consumption in Servers

(computer, network switches & routers, cooling supplies, Uninterrupted Power Supply)

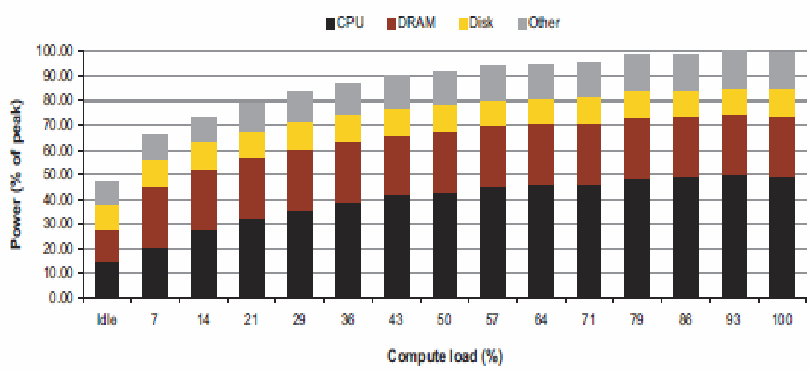

FIGURE 5.8: Subsystem power usage in an ×86 server as the compute load varies from idle to full usage.

Figure 4.9 The cooling system in a raised-floor datacenter with hot-cold air circulation supporting water heat exchange

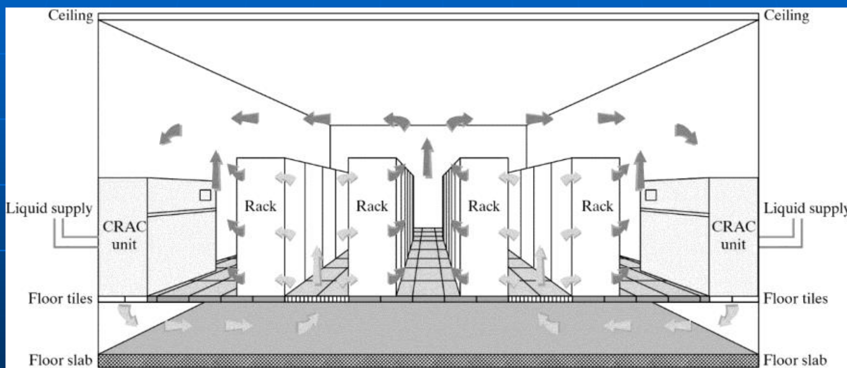

- CRAC – Computer Room Air Conditioning (12 – 14 °C)

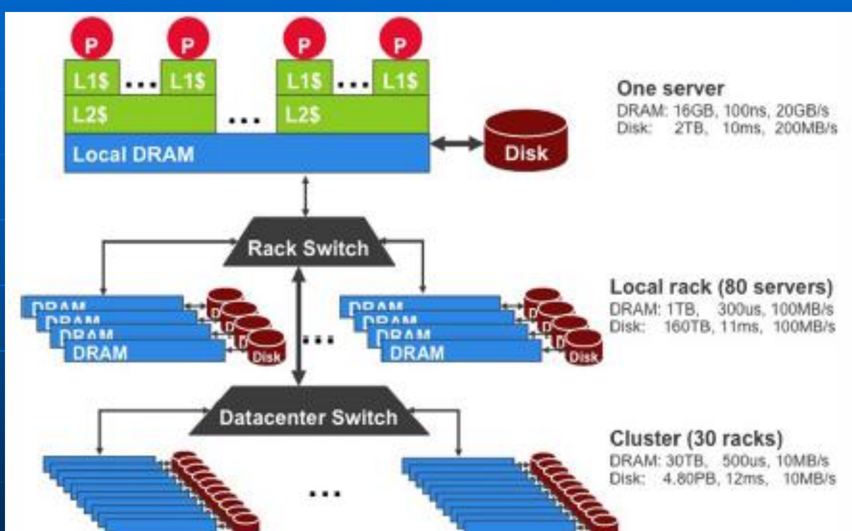

(Courtesy of Luiz Andre Barroso and Urs Holzle, Google Inc., 2009)

Example 4.4 A Fat-free Interconnection Network for Data Center (Figure 4-10)

- Two layers topology with multipath and fault tolerant capability
- Bottom Layer – Server nodes

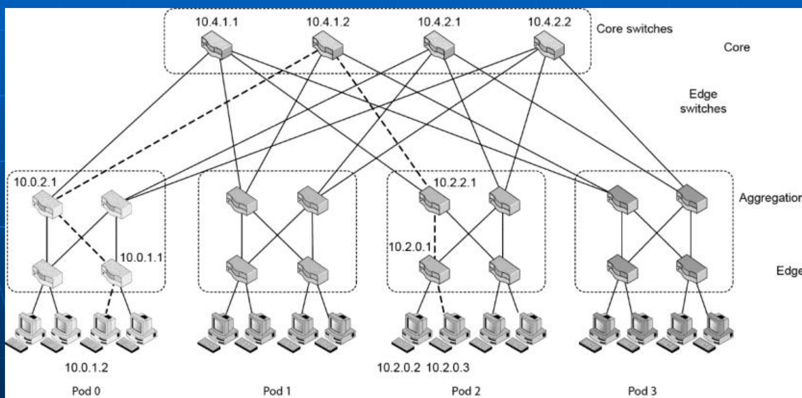

Example 4.5 A Server-Centric Network for Modular Data Center

Server --- O circle; Switch -- Rectangle

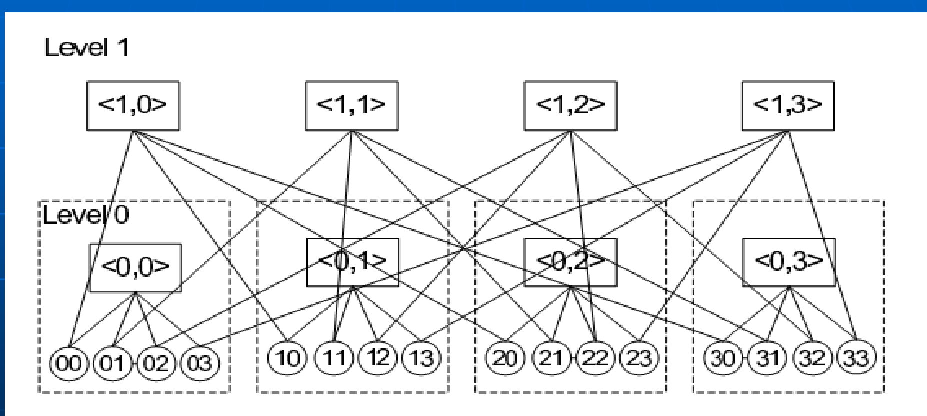

Figure 4.12 BCube: A High Performance, server-centric network for modular datacenters. (Courtesy of C. Guo, et al, *ACM SIGCOMM Computer Communication Review, Oct. 2009.* **[25]**).

Figure 4-11 A modular datacenter built in a truck-towed ICE Cube container, that can be cooled by chilled air circulation with cold water heat exchange

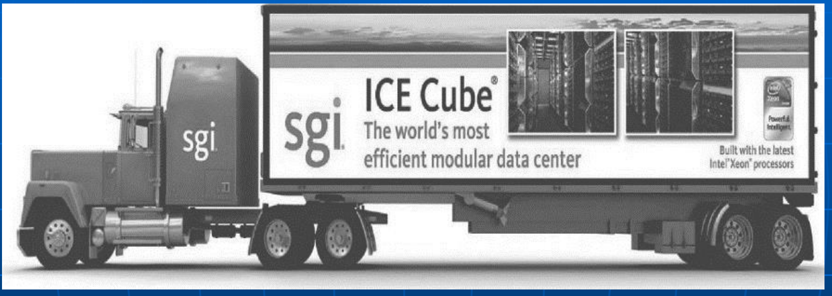

Inter-module Connection Networks

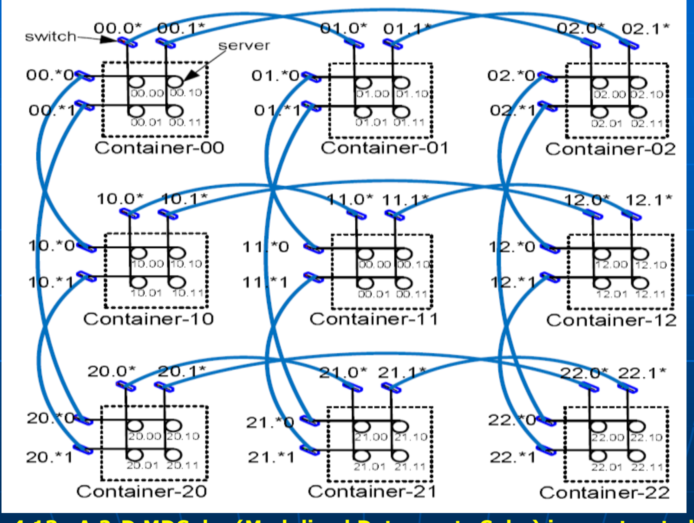

Figure 4.13 A 2-D MDCube (Modulized Datacenetr Cube) is constructed from 9 BCube containers. (Courtesy of . Wu, et al, *ACM CoNEXT'09*, Dec. 2009, [77]).

Larger Datacenter Growth

- One at a time:
  - 1 system
  - Racking & networking: 14 hrs ($1,330)
- Rack at a time:
  - ~ 40 systems
  - Install & networking: .75 hrs ($60)
- Container at a time:
  - ~1,000 systems
  - No packaging to remove
  - No floor space required
  - Power, network, & cooling only
- Weatherproof & easy to transport
- Datacenter construction takes 24+ months
  - Both new build & DC expansion require regulatory approval

Cloud Computing

- Warehouse Scalable Computers (WSCs) offer economies of scale that cannot be achieved with a datacenter:
  - 5.7 times reduction in storage costs
  - 7.1 times reduction in administrative costs
  - 7.3 times reduction in networking costs
  - This has given rise to cloud services such as Amazon Web Services
    - “Utility Computing”
    - Based on using open source virtual machine and operating system software

(Courtesy of Hennessy and Patterson, 2012)

4.3 Architectural Design of Compute and Storage Cloud

- Four Cloud Platform Design Goals
  - Scalability
  - Virtualization
  - Efficiency
  - Reliability
- Cloud-Enabling Technologies (hardware, software, networking)
  - Fast platform deployment
  - Virtual clusters on demand
  - Multitenant techniques
  - Massive data processing
  - Web-scale communication
  - Distributed storage
  - Licensing and billing services

A Generic Cloud Architecture

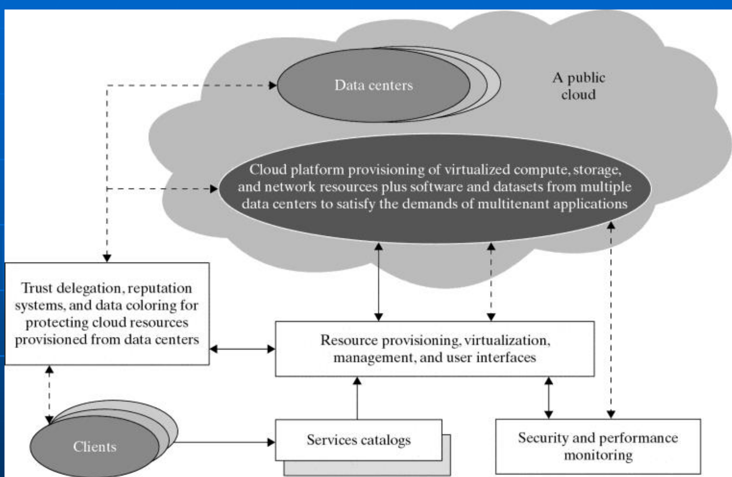

Figure 4.14 A Security-aware platform built with a virtual cluster of VMs, storage, and networking resources

Layered Cloud Architecture Development:
Infrastructure, Platform, and Application

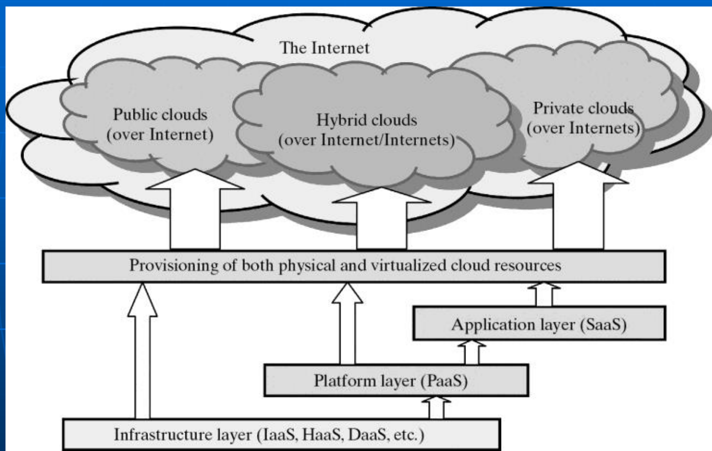

Figure 4.15 Layered architectural development of the cloud platform for IaaS, PaaS, and SaaS applications over the Internet

4 - 19

Market-Oriented Cloud Architecture

- QoS (Quality of Service) based resource allocation mechanisms
  - Users/Brokers
  - SLA resource allocator
  - VMs
  - Physical machinec
- Critical QoS Parameters:
  - Time
  - Cost
  - Reliability, and
  - Trust/security

Prof. Paul Lin

20

Market-Oriented Cloud Architecture

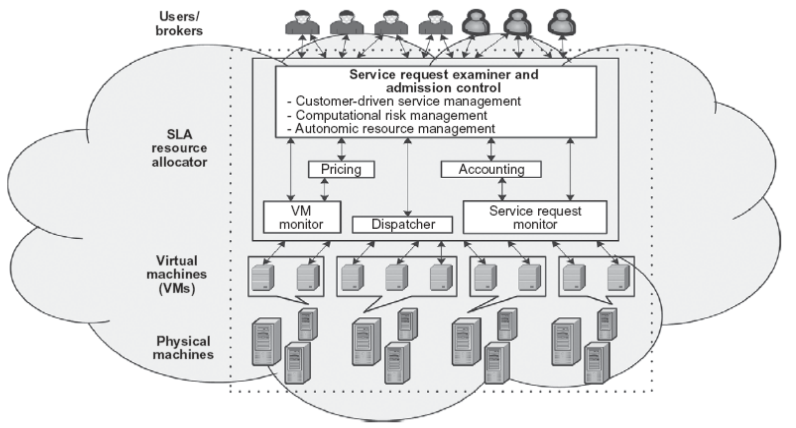

FIGURE 4.16

Market-oriented cloud architecture to expand/shrink leasing of resources with variation in QoS/demand from users.

(Courtesy of Raj Buyya, et al. [11])

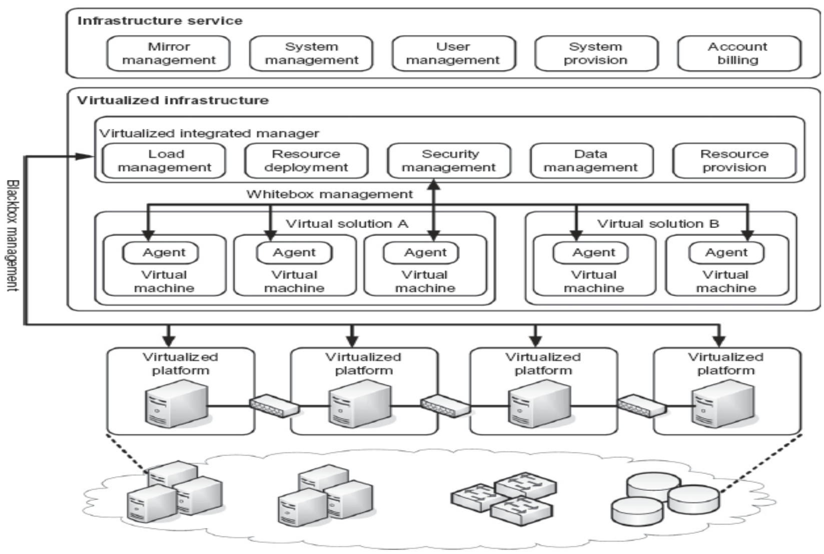

FIGURE 4.17

Virtualized servers, storage, and network for cloud platform construction.

(Courtesy of Zhong-Yuan Qin, SouthEast University, China)

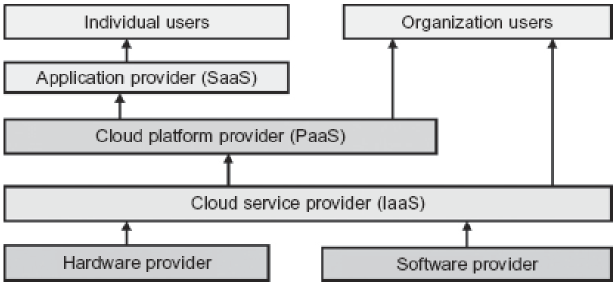

FIGURE 4.19

Roles of individual and organizational users and their Interaction with cloud providers under various cloud service models.

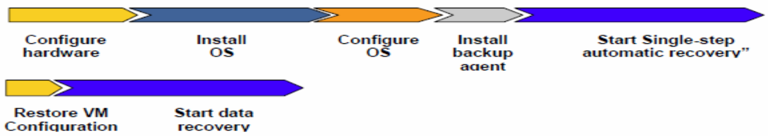

Figure 7.21 Recovery overhead of a conventional disaster recovery between physical machines, compared with that required to recover from live migration of virtual machines

Table 1.7 Three Cloud Computing Platforms and Underlying Technologies [14]

Platform Features	Google Cloud Platform [17]	IBM BlueCloud System [7]	Amazon Cloud Cluster [
Platform Architecture and Technology	Google server clusters dynamically drafted from 450K Google servers, GFS, and datacenters	A clustered server platform to provide a total system to distributed problem solving and decision making	Amazon built an utility cluster (iDataPlex) of 2000 nodes with distributed storage
Target Applications claimed by cloud providers	Upgraded web-scale services, distributed data storage and services based on Software or Platform as a Service (SaaS or PaaS) models	Business applications, academic services, and raw supercomputing, and collaborative computing based on the Platform as a Service (PaaS) model.	To lease CPU time and storage to serve massive number of small users in business applications using an Infrastructure as a Service (IaaS) Model
Programming Model and Software Tools Applied	Extending MapReduce and BigTable for web-scale search and distributed computing	Use open software from Google and Hadoop plus IBM WebSphere 2.0, DB2, PowerVM, and Tivoli software	Use Hadoop EC2 to provide CPU cycles and S3 for storage services in business and e-commerce

Table 4.4 Virtualized Resources in Compute, Storage, and Network Clouds [4]

Provider	AWS	Microsoft Azure	GAE
Compute cloud with virtual cluster of servers	x86 instruction set, Xen VMs, resource elasticity allows scalability through virtual cluster, or a third party such as RightScale must provide the cluster	Common language runtime VMs provisioned by declarative descriptions	Predefined application framework handlers written in Python, automatic scaling up and down, server failover inconsistent with the Web applications
Storage cloud with virtual storage	Models for block store (EBS) and augmented key/blob store (SimpleDB), automatic scaling varies from EBS to fully automatic (SimpleDB, S3)	SQL Data Services (restricted view of SQL Server), Azure storage service	MegaStore/BigTable
Network cloud services	Declarative IP-level topology; placement details hidden, security groups restricting communication, availability zones isolate network failure, elastic IP applied	Automatic with user's declarative descriptions or roles of app. components	Fixed topology to accommodate three-tier Web app. structure, scaling up and down is automatic and programmer-invisible

Cloud Services and Major Providers

Cloud application (SaaS)	Concur, RightNOW, Teleo, Kenexa, Webex, Blackbaud, salesforce.com, Netsuite, Kenexa, etc.
Cloud software environment (PaaS)	Force.com, App Engine, Facebook, MS Azure, NetSuite, IBM BlueCloud, SGI Cyclone, eBay
Cloud software infrastructure: Computational resources (IaaS), Storage (DaaS), Communications (Caas)	Amazon AWS, OpSource Cloud, IBM Ensembles, Rackspace cloud, Windows Azure, HP, Banknorth
Co-location cloud services (LaaS)	Savvis, Internap, NTTCommunications, Digital Realty Trust, 365 Main
Network cloud services (NaaS)	Owest, AT&T, AboveNet
Hardware/Virtualization cloud services (HaaS)	VMware, Intel, IBM, XenEnterprise

FIGURE 4.23

A stack of six layers of cloud services and their providers.

(Courtesy of T. Chou, Active Book Express, 2010 [16])

Conclusion and Summary