Unit III

Architectural design of compute and storage clouds

Cloud platform design goals

- Major design goals of cloud
- Scalability- if one service takes a lot of processing power, storage capacity, or network traffic, it is simple to add more servers and bandwidth
- Virtualization
- Efficiency
- Reliability- data can be put into multiple locations so even if one server in one location fails the user has copy in other locations also.
- security

Cloud management receives the user request, finds the correct resources and then calls the provisioning services which invoke the resources in the cloud.

Enabling technologies for clouds

Key driving forces behind cloud computing are

- Ubiquity of broadband and wireless networking, falling storage costs and progressive improvements in internet computing software

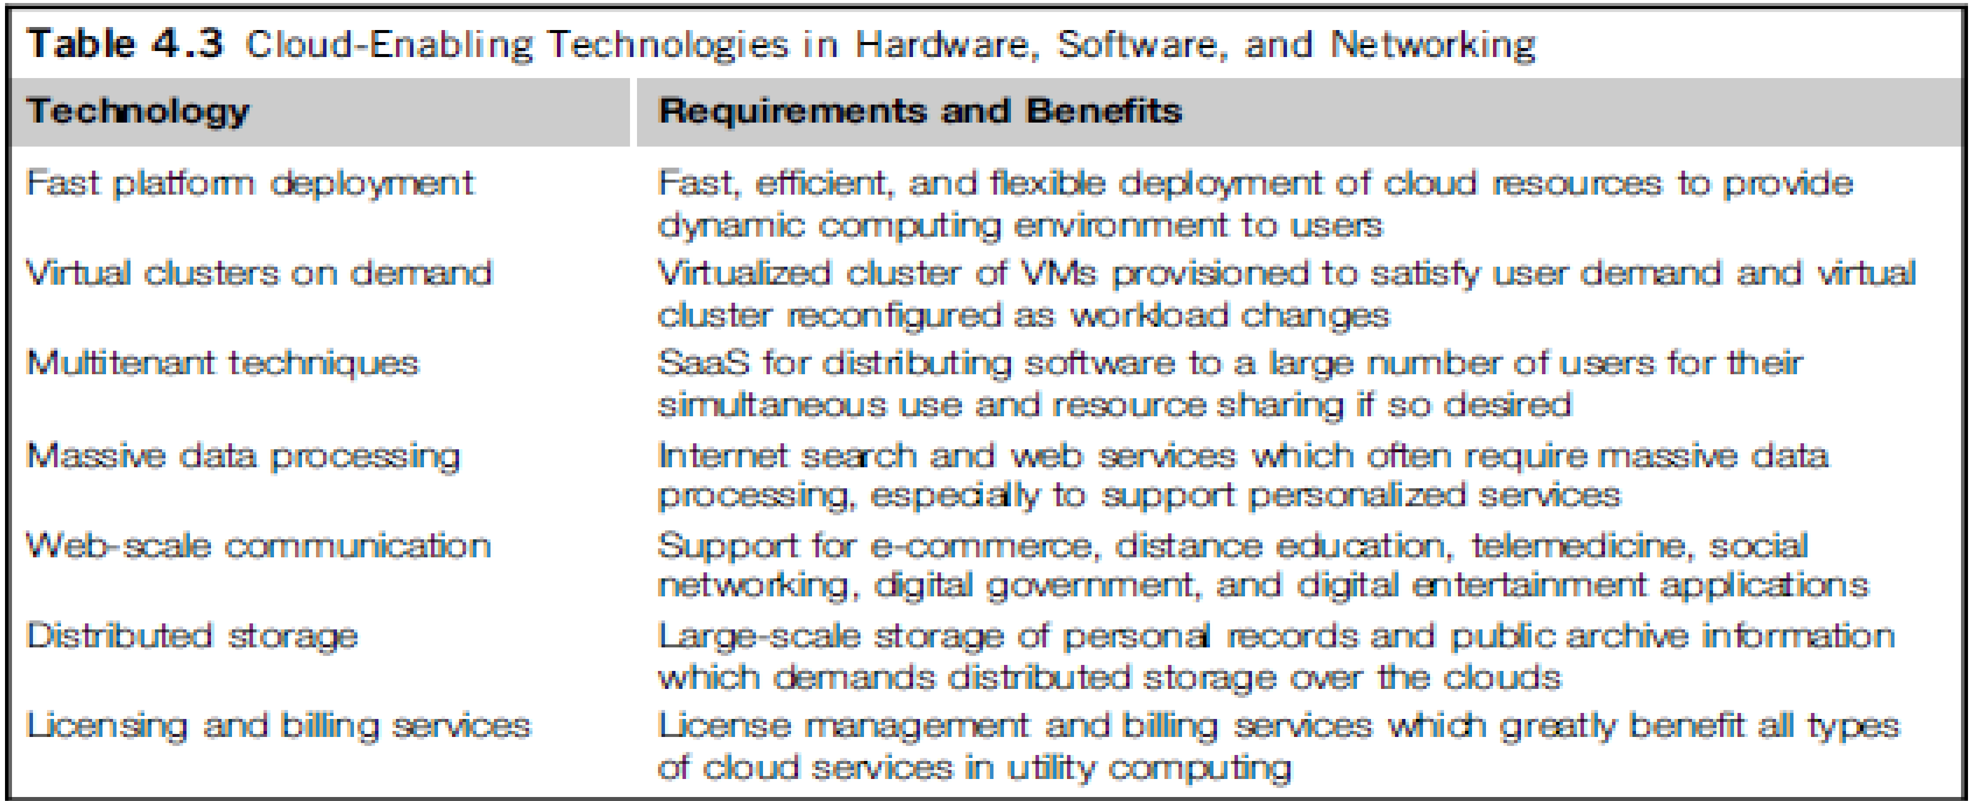

Table 4.3 Cloud-Enabling Technologies in Hardware, Software, and Networking

Technology	**Requirements and Benefits**
Fast platform deployment	Fast, efficient, and flexible deployment of cloud resources to provide dynamic computing environment to users
Virtual clusters on demand	Virtualized cluster of VMs provisioned to satisfy user demand and virtual cluster reconfigured as workload changes
Multitenant techniques	SaaS for distributing software to a large number of users for their simultaneous use and resource sharing if so desired
Massive data processing	Internet search and web services which often require massive data processing, especially to support personalized services
Web-scale communication	Support for e-commerce, distance education, telemedicine, social networking, digital government, and digital entertainment applications
Distributed storage	Large-scale storage of personal records and public archive information which demands distributed storage over the clouds
Licensing and billing services	License management and billing services which greatly benefit all types of cloud services in utility computing

A generic cloud architecture

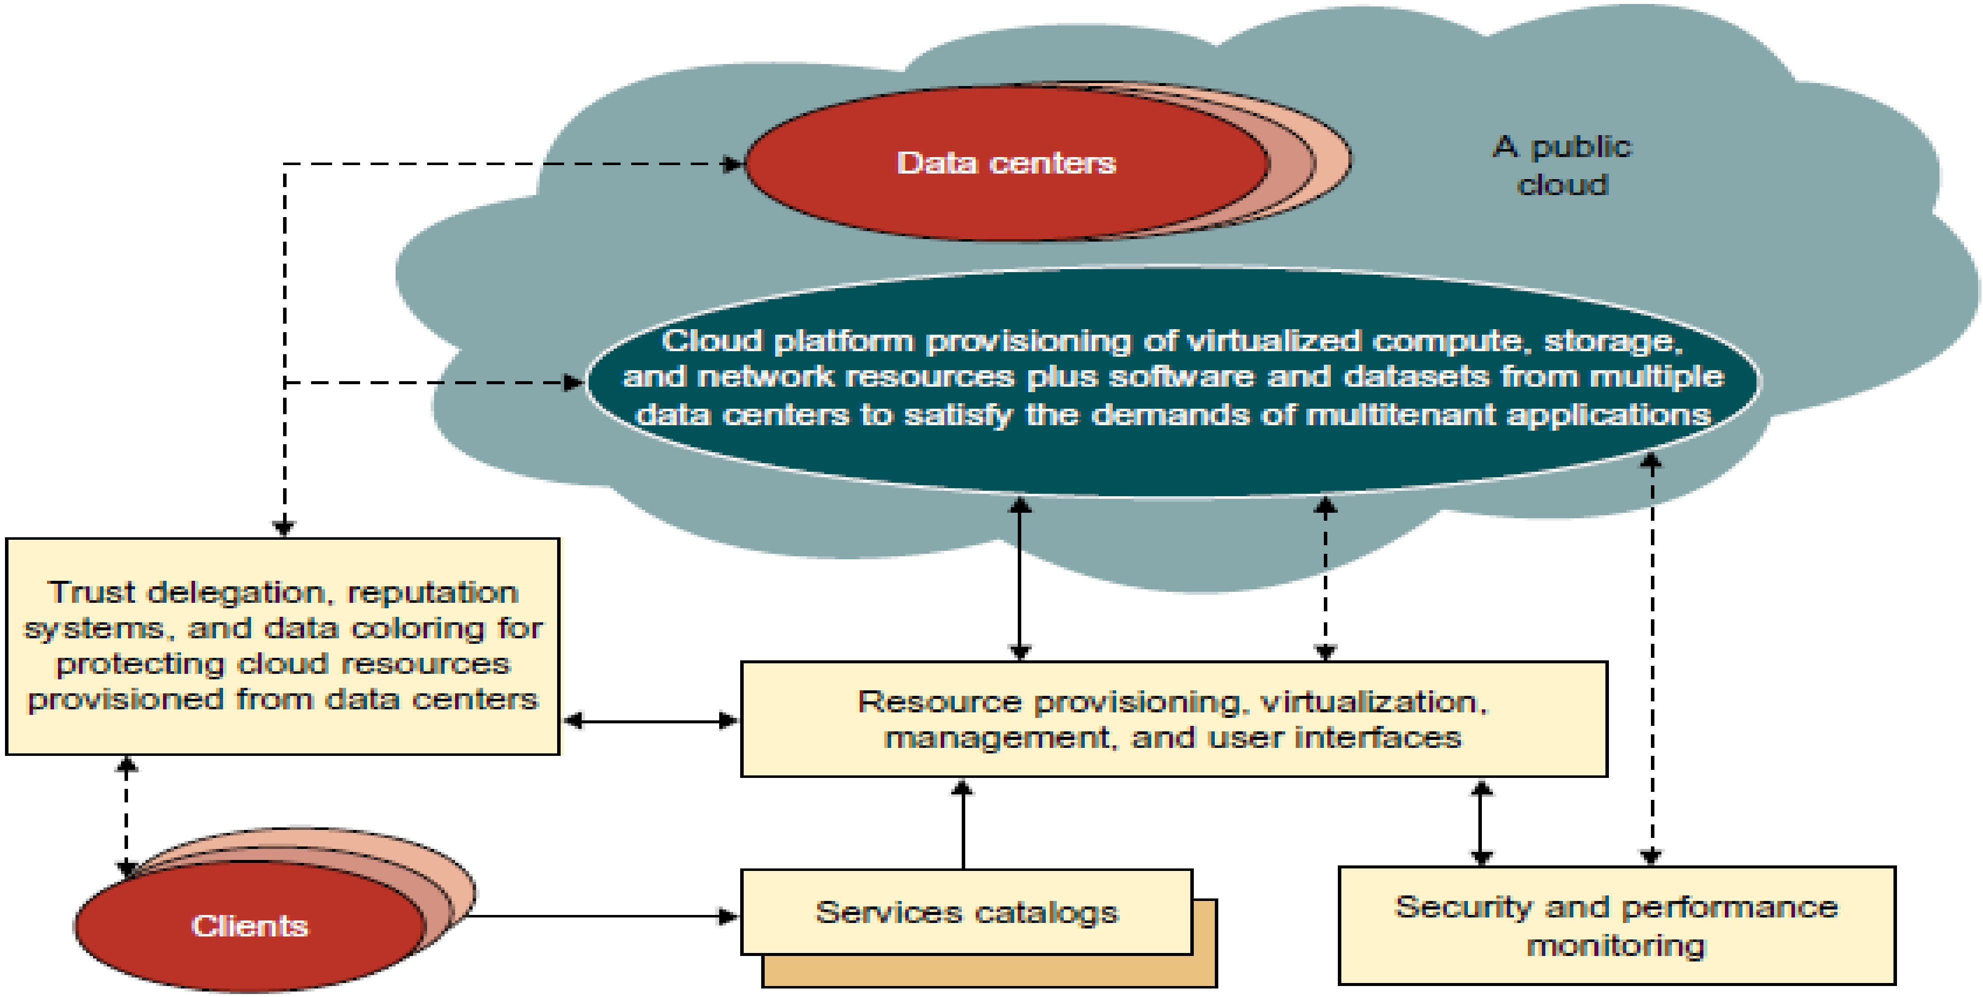
Data centers
A public cloud
Cloud platform provisioning of virtualized compute, storage, and network resources plus software and datasets from multiple data centers to satisfy the demands of multitenant applications
Trust delegation, reputation systems, and data coloring for protecting cloud resources provisioned from data centers
Resource provisioning, virtualization, management, and user interfaces
Clients
Services catalogs
Security and performance monitoring

FIGURE 4.14

A security-aware cloud platform built with a virtual cluster of VMs, storage, and networking resources over the data-center servers operated by providers.

(Courtesy of K. Hwang and D. Li, 2010 [361])

- The internet cloud is envisioned as massive cluster of servers.
- The servers provide collective web services or distributed applications using data-center resources on demand.
- Servers in the cloud may be physical or VMs.
- User interfaces are applied to request services.
- Cloud computing resources in data centers are owned and operated by third-party provider.
- Monitoring and metering units are used to track the usage and performance of provisioned resources.
- The software infrastructure of a cloud platform must handle all resource mgt and do most of the maintenance automatically.
- Private clouds are easy to manage and public clouds are easy to access.
- Now the trend is more towards hybrid cloud.
- Security becomes a critical issue.

Layered cloud architectural development

The architecture of a cloud is developed at three layers:

- infrastructure
- Platform
- Application
- Infrastructure layer is deployed first to support IaaS services.- built with virtualized compute, storage, network resources.
- Based on infrastructure layer, platform layer is build-this provides users with an environment to develop their applications , to test operation flows and to monitor execution results and performance.
- Based on platform layer application layer is built.- formed with a collection of all needed s/w modules for SaaS applications. Eg office mgt work tools (for information retrieval, document processing, calendar and authentication services, CRM, financial transactions, supply chain management.
- Services at the application layer demand more work from providers.

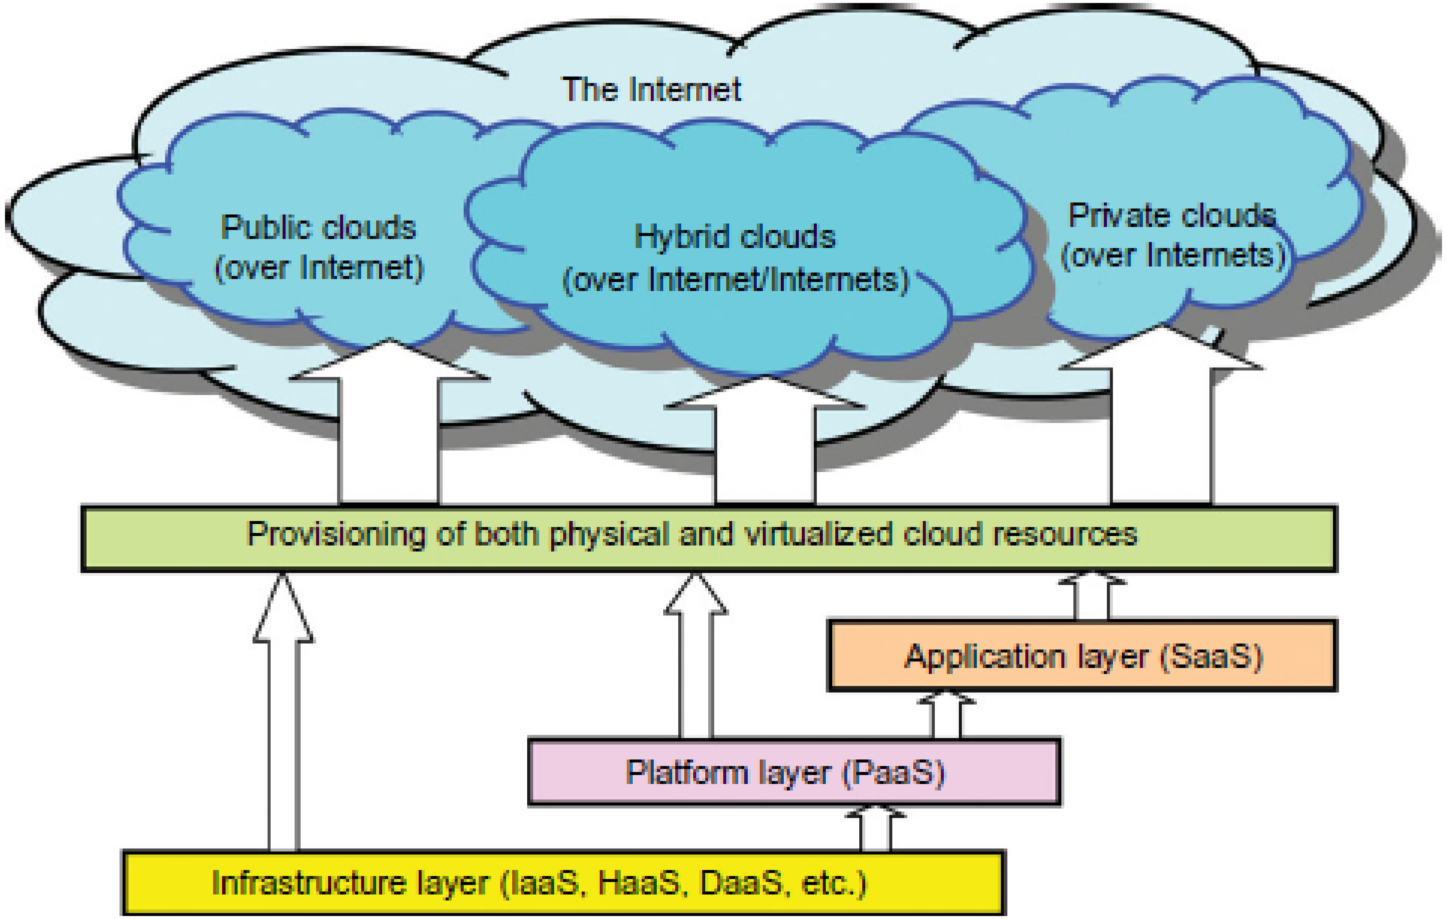
The Internet
Public clouds
(over Internet)
Hybrid clouds
(over Internet/Internets)
Private clouds
(over Internets)
Provisioning of both physical and virtualized cloud resources
Application layer (SaaS)
Platform layer (PaaS)
Infrastructure layer (IaaS, HaaS, DaaS, etc.)

FIGURE 4.15

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

Market-Oriented cloud architecture

- Each customer is to be satisfied with a special or customised QoS level according to the specific SLA.(Service level Agreement)
- The providers cannot deploy traditional system-centric resource management architecture, they must deploy market-oriented resource mgt to regulate the supply and demand of cloud resources.
- Feedback on service can be obtained from customers for improvement.
- Cost can be reduced which increases the profit of cloud providers.
- Users or brokers submit service requests .
- The SLA resource allocator acts as the interface between the provider and user/brokers.
- The decision to accept or reject is taken based on submitted request by service request examiner by ensuring there is no overloading of resources.

The following mechanism helps in doing this efficiently

- It gets the latest status information regarding resource availability –from **VM Monitor mechanism** and workload processing as well as progressing of work– from **Service Request Monitor mechanism**
- **Dispatcher mechanism-** starts the execution of accepted service requests on allocated VM
- **Pricing mechanism** decides on charging of service requests. Eg based on submission time (peak/off-peak), pricing rates (fixed/changing), or availability of resources(supply / demand)
- This helps in prioritizing resource allocations effectively.
- **Accounting mechanism-** maintains the actual usage of resources for final charging.

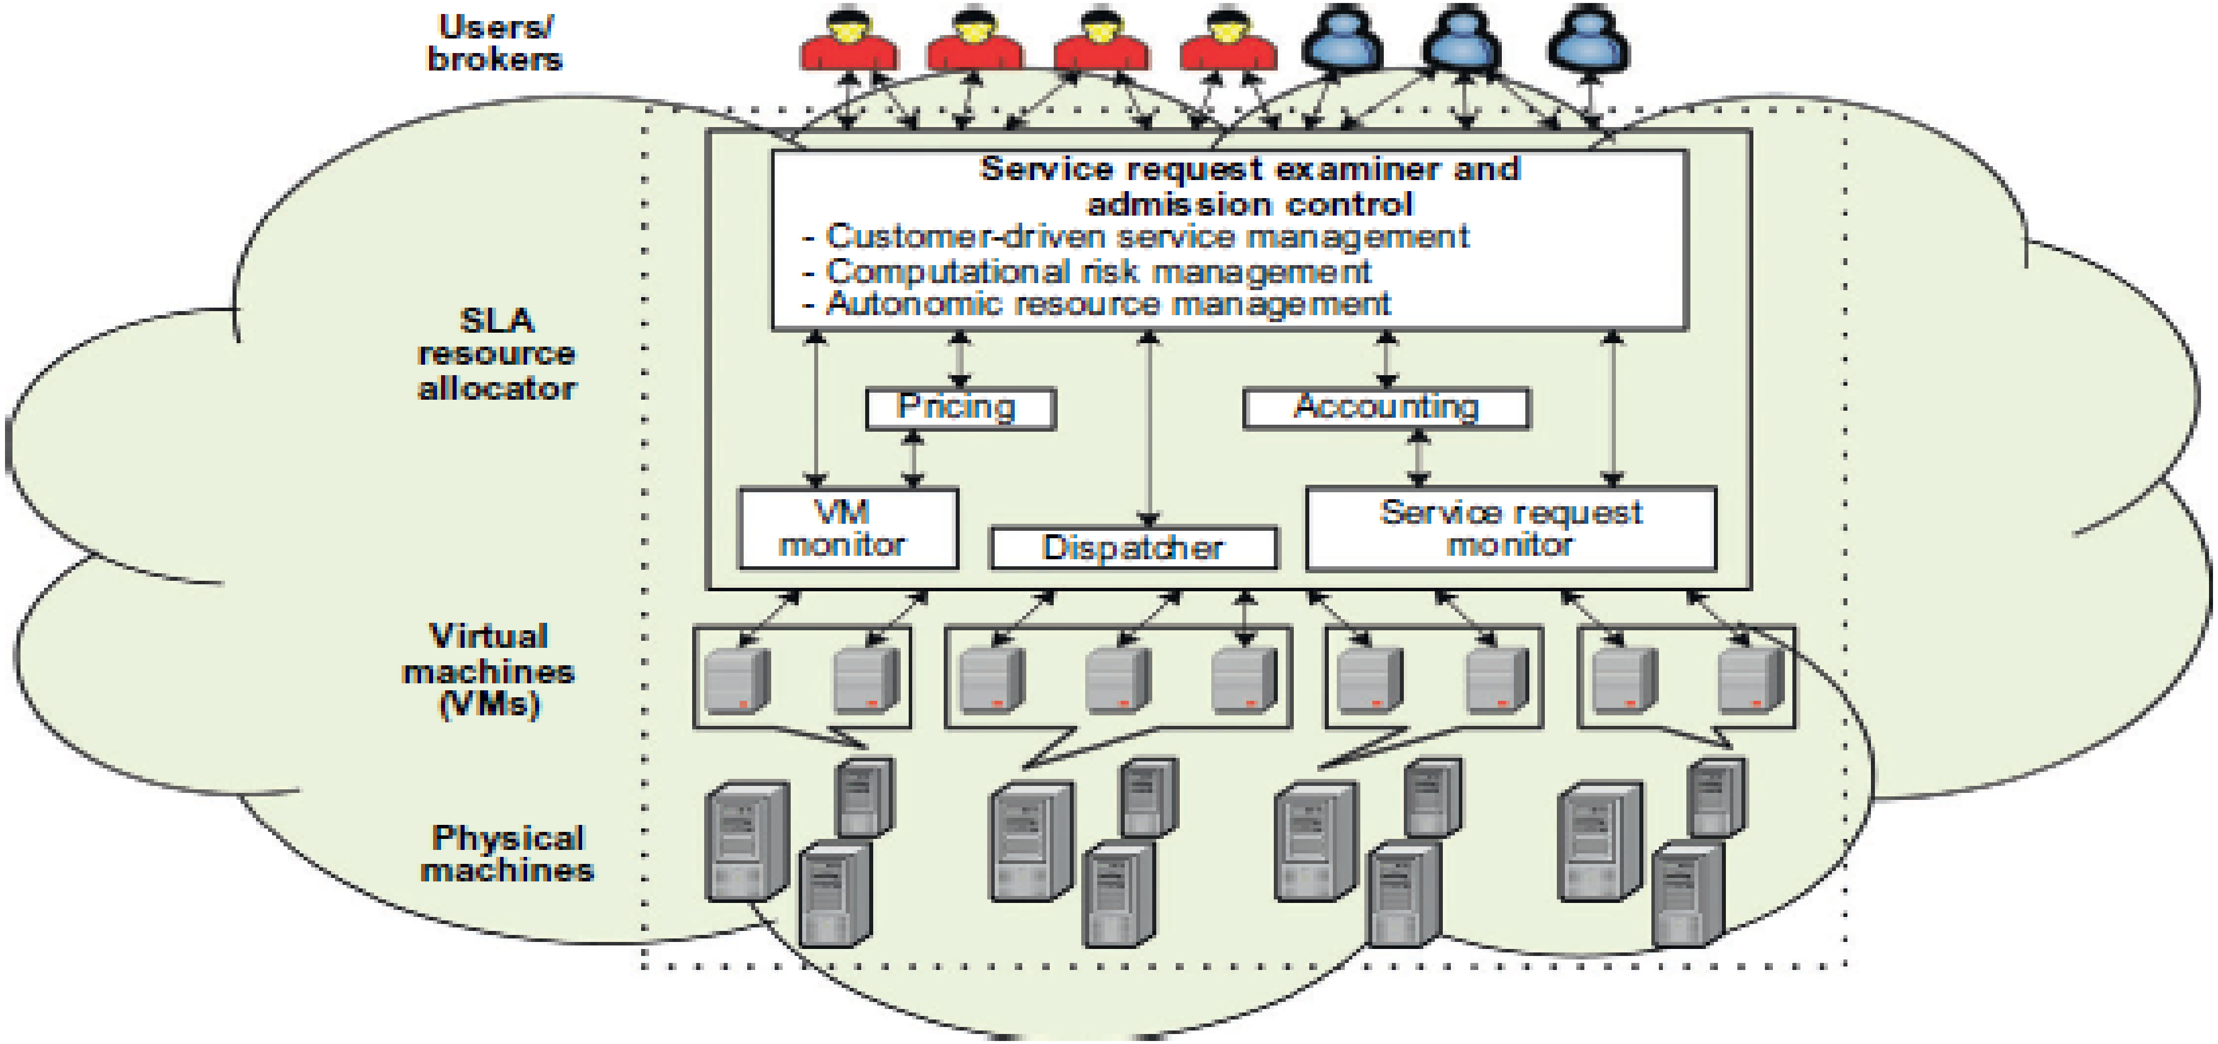
Users/ brokers
SLA resource allocator
Virtual machines (VMs)
Physical machines
Service request examiner and admission control
- Customer-driven service management
- Computational risk management
- Autonomic resource management
Pricing
Accounting
VM monitor
Dispatcher
Service request monitor

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])

Quality of service factors

- The data center comprises multiple computing servers that provide resources to meet service demands.
- Some of QoS parameters are time, cost, reliability, trust/security etc.,
- QoS requirements cannot be static and may change based on business operation and operating environment changes.
- More importance should be given to customers as they pay for what they use.
- Commercial clouds define computational risk mgt tactics to identify, assess, and manage risks involved in the execution of applications with regard to service requirements and customer needs.
- The system must incorporate autonomic resource mgt models that effectively self-manage changes in service requirements to satisfy both new service demands and existing service obligations and leverage VM technology to dynamically assign resource shares according to service requirements.

Virtualization support and disaster recovery

- Cloud users do not need to know and have no way to discover physical resources that are involved while processing a service request.
- Hardware virtualization- virtualization s/w is used as a platform for developing new cloud applications that enable developers to use any operating systems and programming environments they like. The development environment and deployment environment can now be the same, which eliminates some runtime problems.
- In traditional environment sharing of cluster resources depends on user and group mechanism on a system. Sharing is not flexible. An environment that meets one user's requirement often cannot satisfy another user. Virtualization allows users to have full privileges while keeping them separate.
- Users have full access to their own VMs which are completely separate from other users VMs.
- Multiple VMs can be mounted on the same physical server.
- Different VMs may run with different OSes.
- The virtualized resources form a resource pool (CPU, networks etc)
- The virtualized infrastructure (black box in the middle) is built with many virtualizing integration managers. These managers handle loads, resources, security, data and provisioning functions.

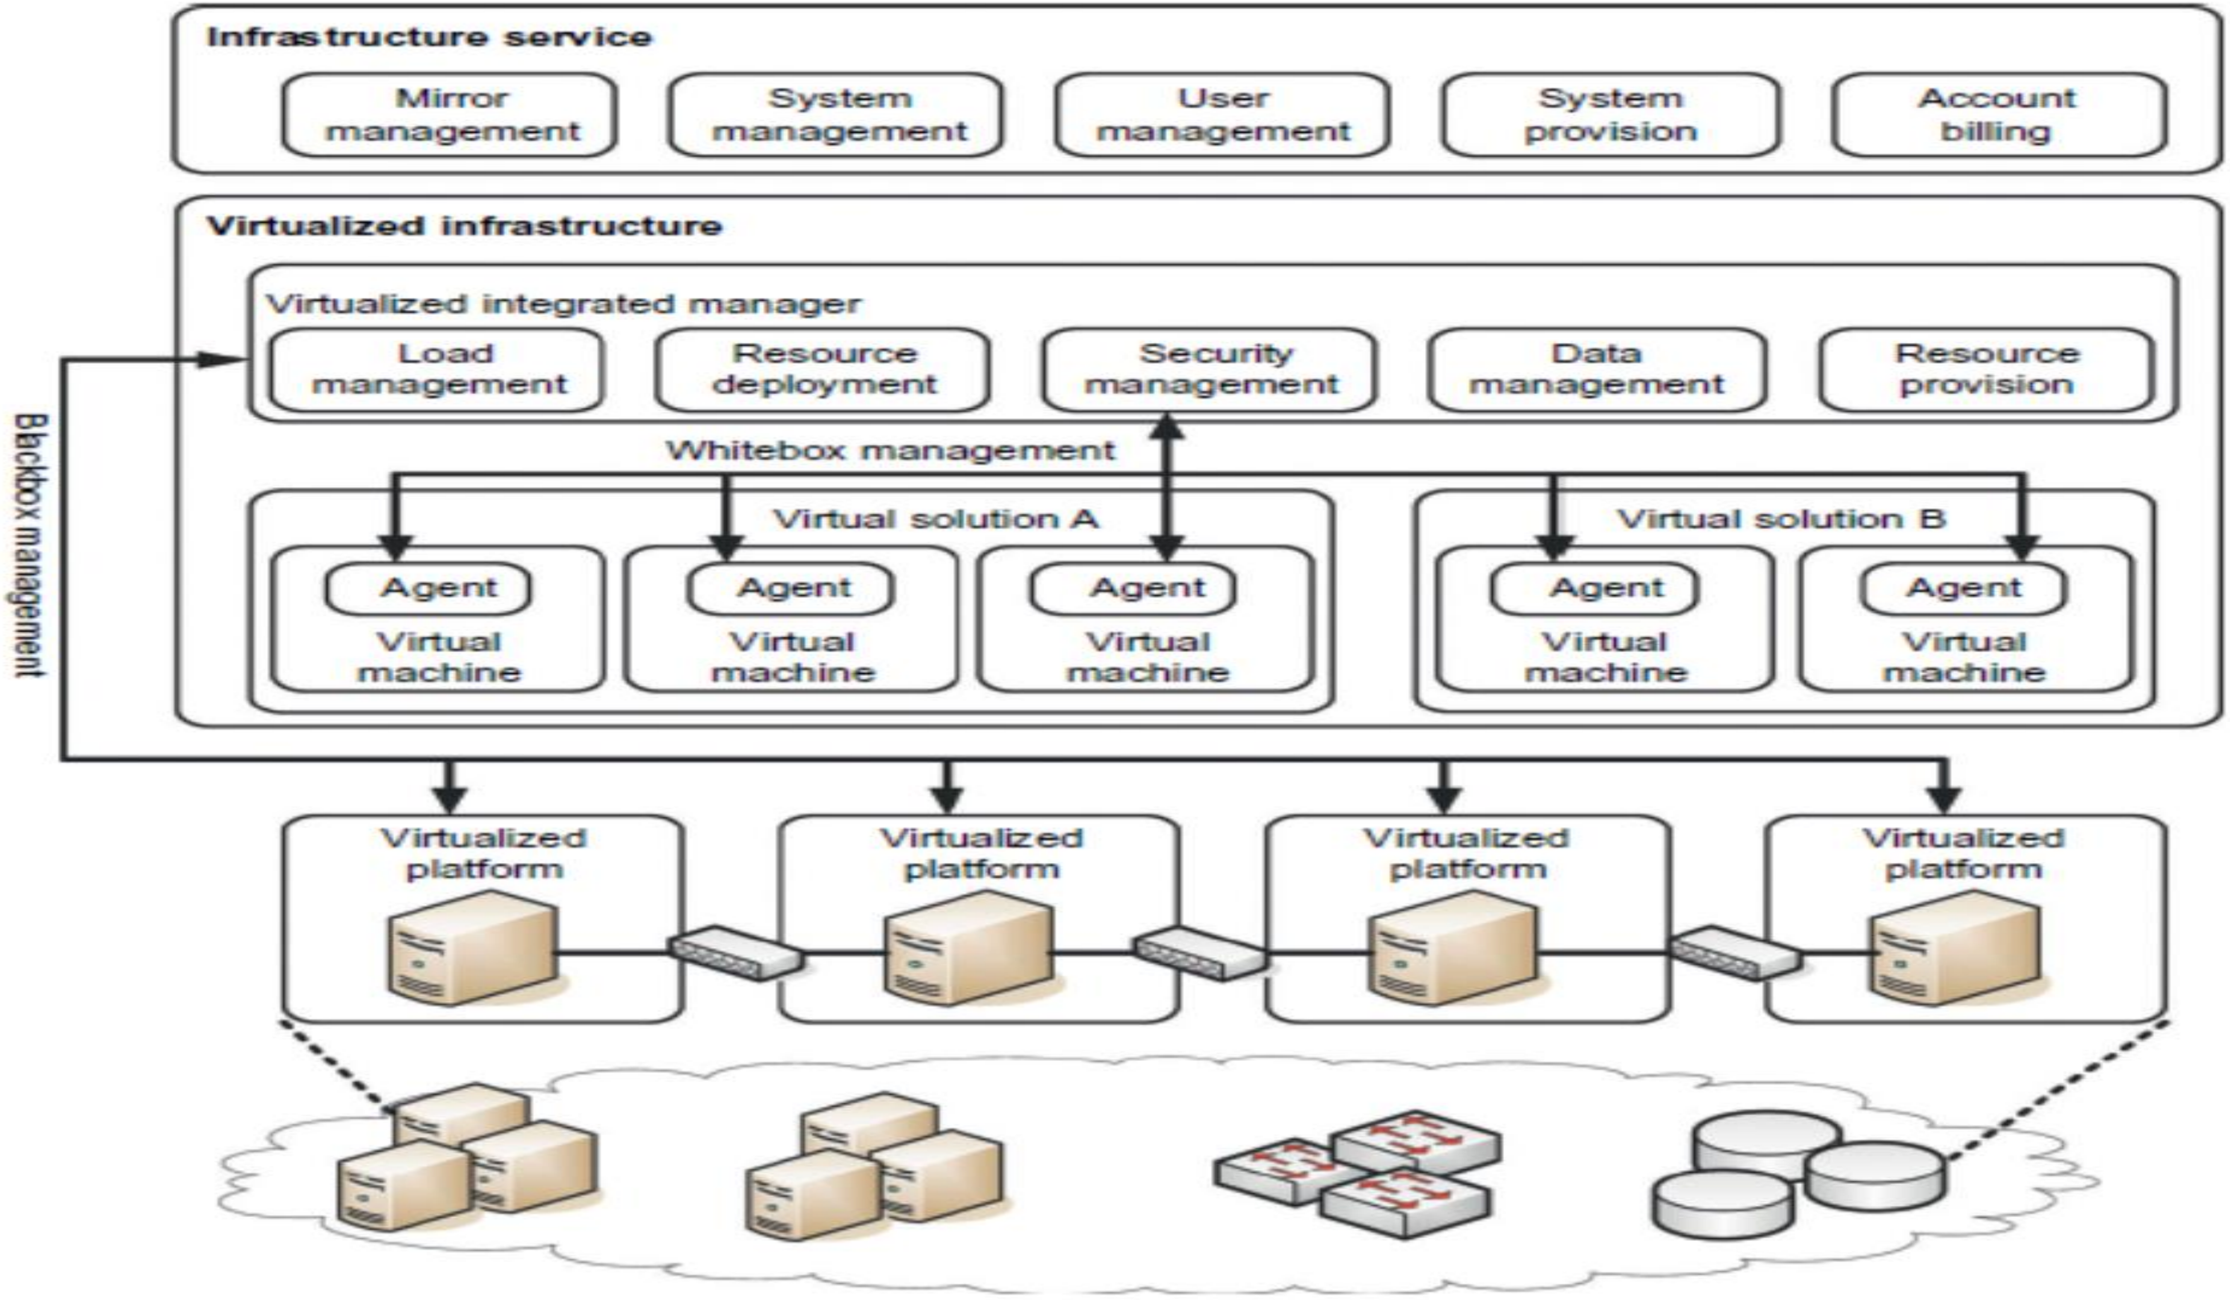

FIGURE 4.17

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

Virtualization support in public clouds

- AWS (Amazon Web Service)- provides extreme flexibility(VMs) for users to execute their own applications
- GAE(Google App Engine)- limited application-level virtualization for users to build applications only based on the services that are created by Google.
- Microsoft –provides programming-level virtualization(.NET) for users to build their applications.
- Refer picture 4.18 and table 4.4(next slide)

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

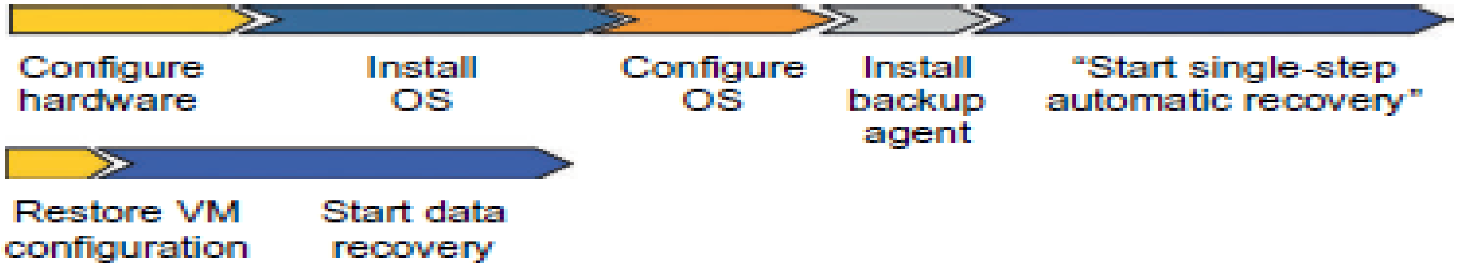

FIGURE 4.18

Recovery overhead of a conventional disaster recovery scheme, compared with that required to recover from live migration of VMs.

- IT power consumption in the US has more than doubled to 3 percent of the total energy consumed in the country.
- Storage virtualization for green data centers.- virtualization leads to less power consumption.
- Virtualization for IaaS-use of VMs in cloud has the following benefits
    - System administrators consolidate workloads of underutilized servers in fewer servers
    - VMs have the ability to run legacy code without interfering with other APIs (*Legacy code* is source *code* that relates to a no-longer supported or manufactured operating system or other computer technology. ... When the manufacturer upgrades a platform (or the platform is superseded), the *code* may no longer work without changes, and becomes *legacy code*.)
    - VMs can be used to improve security through creation of **sandboxes** (is a security mechanism for separating running programs, usually in an effort to mitigate system failures or software vulnerabilities from spreading. It is often used to execute untested or untrusted programs or code, possibly from unverified or untrusted third parties, suppliers, users or websites, without risking harm to the host machine or operating system) for running applications with questionable reliability
    - Virtualized cloud platforms can apply performance isolation, letting providers offer some guarantees and better QoS to customer applications

VM cloning for Disaster Recovery

- Scheme 1-Recover one physical machine by another physical machine
- Scheme 2- Recover one VM by another VM
- Scheme1- traditional disaster recovery is rather slow, complex and expensive. Total recovery time is attributed to the hardware conf, installing and config OS, installing backup agent, and long time to restart the physical machine.
- Scheme2- in this the steps installation and configuration of OS and backup agents are eliminated which greatly reduces the recovery time.
- Cloning of VM is another solution.
- make a **clone** on a remote server for every running VM on a local server. At a time only one needs to be active. A cloud center should be able to activate this clone VM in case of failure of the original VM, taking a snapshot of the VM to enable live migration in a minimal time. Only updated and modified states are sent to the suspended VM to update its state.

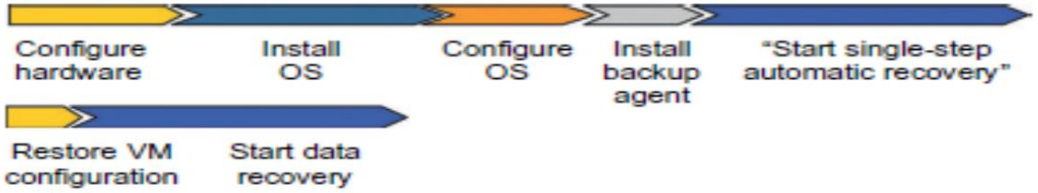

FIGURE 4.18

Recovery overhead of a conventional disaster recovery scheme, compared with that required to recover from live migration of VMs.

Architectural design challenges

Challenge 1- service availability and data lock-in problem

- the management of a cloud service by a single company is often the source of single points of failure. Using of multiple cloud providers, the company has multiple data centers located in different geographic regions which may have common s/w infrastructure and accounting systems. It gives more protection from failures.
- Another threat is distributed denial of service(DDoS) attacks. Criminals threaten to cut off the incomes of SaaS providers by making their services unavailable.
- Provide common standardised API so that the SaaS developer can deploy services and data across multiple cloud providers. This will rescue the loss of all data due to the failure of a single company. This will eliminate data-lock-in problem. This enable “surge computing” in which the public cloud is used to capture the extra tasks that cannot be easily run in the data center of a private cloud.

Architectural design challenges cont..

Challenge 2- Data Privacy and security concerns

- current cloud offerings are public networks, exposing the system to more attacks. These can be overcome by encrypted storage, virtual LANs, and network middleboxes(firewalls, packet filters)
- Laws are levied by countries to keep the data within their boundaries
- In traditional network attacks include buffer overflows, DoS attacks, spyware, malware etc
- In a cloud environment, the attacks are hypervisor malware, guest hopping (**Hyperjacking** is an attack in which a hacker takes malicious control over the hypervisor that creates the virtual environment within a virtual machine (VM) host.[1] The point of the attack is to target the operating system that is below that of the virtual machines so that the attacker's program can run and the applications on the VMs above it will be completely oblivious(unaware) to its presence . In a **guest-hopping attack,** an attacker will try to identify two virtual machines likely to be hosted on the same physical hardware. Assuming the attacker is interested in data from virtual machine A, but is unable to directly penetrate virtual machine A, the attacker will try to penetrate virtual machine B, and then try to gain access to virtual machine A.)and hijacking or VM rootkits, man-in-the-middle attack for VM migration
- Passive attacks- steal sensitive data or passwords
- Active attacks- manipulation of kernel data structures which cause major damage

Architectural design challenges cont..

- Challenge 3- Unpredictable Performance and Bottlenecks
- Sharing of CPU, Memory is ok but I/O sharing is problematic
- Applications are “pulled apart” across the boundaries of clouds which complicate data placement and transport.
- Cloud users and providers have to think about the implications of placement and traffic at every level of the system, if they want to minimize cost.
- Data transfer bottlenecks must be removed, bottleneck links must be widened, and weak servers should be removed.

Architectural design challenges cont..

Challenge 4- Distributed storage and widespread software bugs

- The database is always growing.
- Opportunity is to create a storage system which scales up and down on demand.
- This demands the design of efficient distributed SANs(storage area network)- in this data consistency checking is a big challenge
-

Architectural design challenges cont..

- Challenge 5-cloud scalability, interoperability, and standardization
- The pay-as-you-go model applies to storage and network bandwidth.
- GAE-automatically scales up and down to load increase and decrease and users are charged by the cycles used
- AWS-charges by the hour for the number of VM instances used.
- OVF(open Virtualization Format)-describes an open, secure, portable, efficient, and extensible format for the packaging and distribution of VMs. This format does not rely on the use of a specific host platform, virtualization platform or guest os.
- OVF also defines a transport mechanism for VM templates, and can apply to different virtualization platforms with different levels of virtualization. Virtual appliances must run on any virtual platform.

Architectural design challenges cont..

- Challenge 6-Software Licensing and Reputation Sharing
- Many cloud computing providers rely on open source software as it is ideal for utility computing.
- The company's can consider using both pay-for-use and bulk-use licensing schemes to widen the business coverage.
- One customer's bad behaviour can affect the reputation of entire cloud. For instance, black-listing of EC2 IP address by spam-prevention services may limit smooth VM installation.
- Another legal issue is transfer of legal liability. Cloud providers want legal liability to remain with the customer and vice-versa. This has to be solved at SLA level

Inter-Cloud Resource Management
Extended cloud computing services

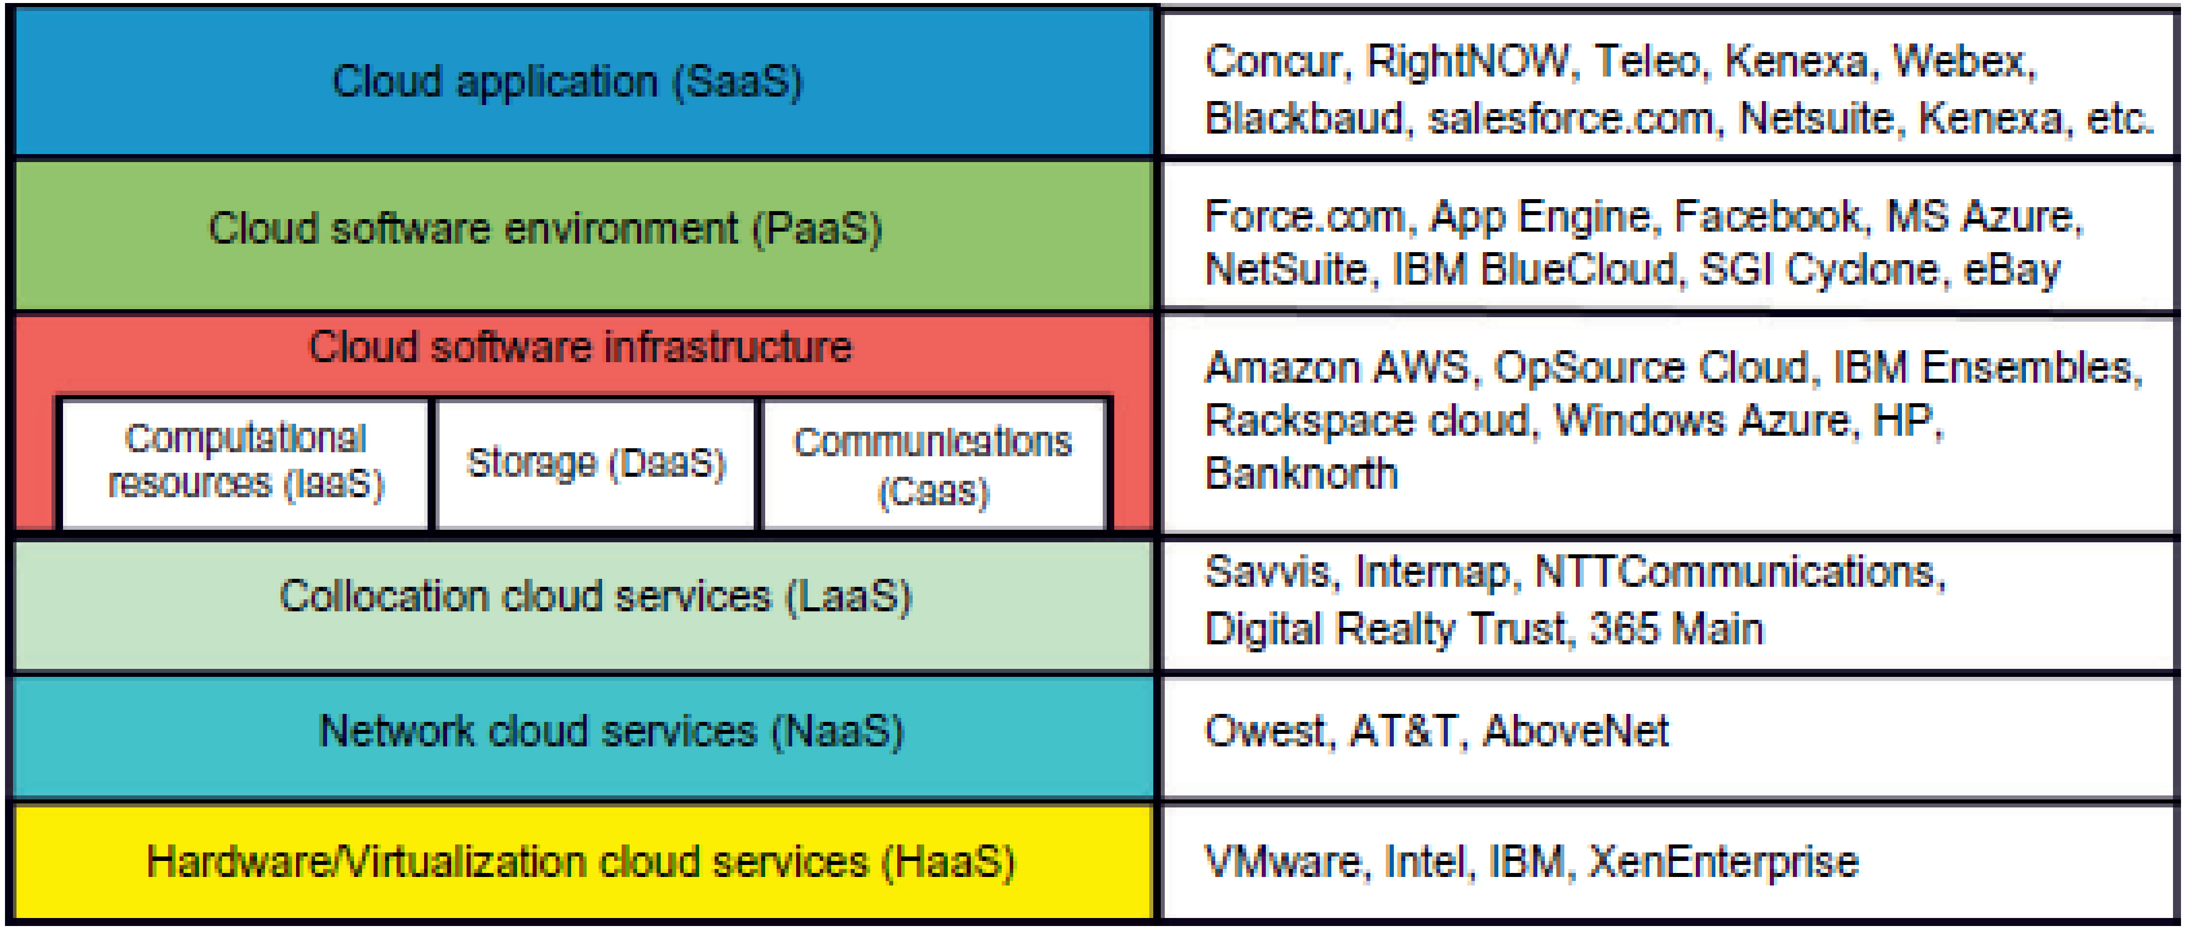

Cloud application (SaaS)	Concur, RightNOW, Teleo, Kenexa, Webex, Blackbaud, salesforce.com, Netsuite, Kenexa, etc.
Cloud software environment (PaaS)	Force.com, App Engine, Facebook, MS Azure, NetSuite, IBM BlueCloud, SGI Cyclone, eBay
Cloud software infrastructure Computational resources (IaaS) Storage (DaaS) Communications (Caas)	Amazon AWS, OpSource Cloud, IBM Ensembles, Rackspace cloud, Windows Azure, HP, Banknorth
Collocation 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.

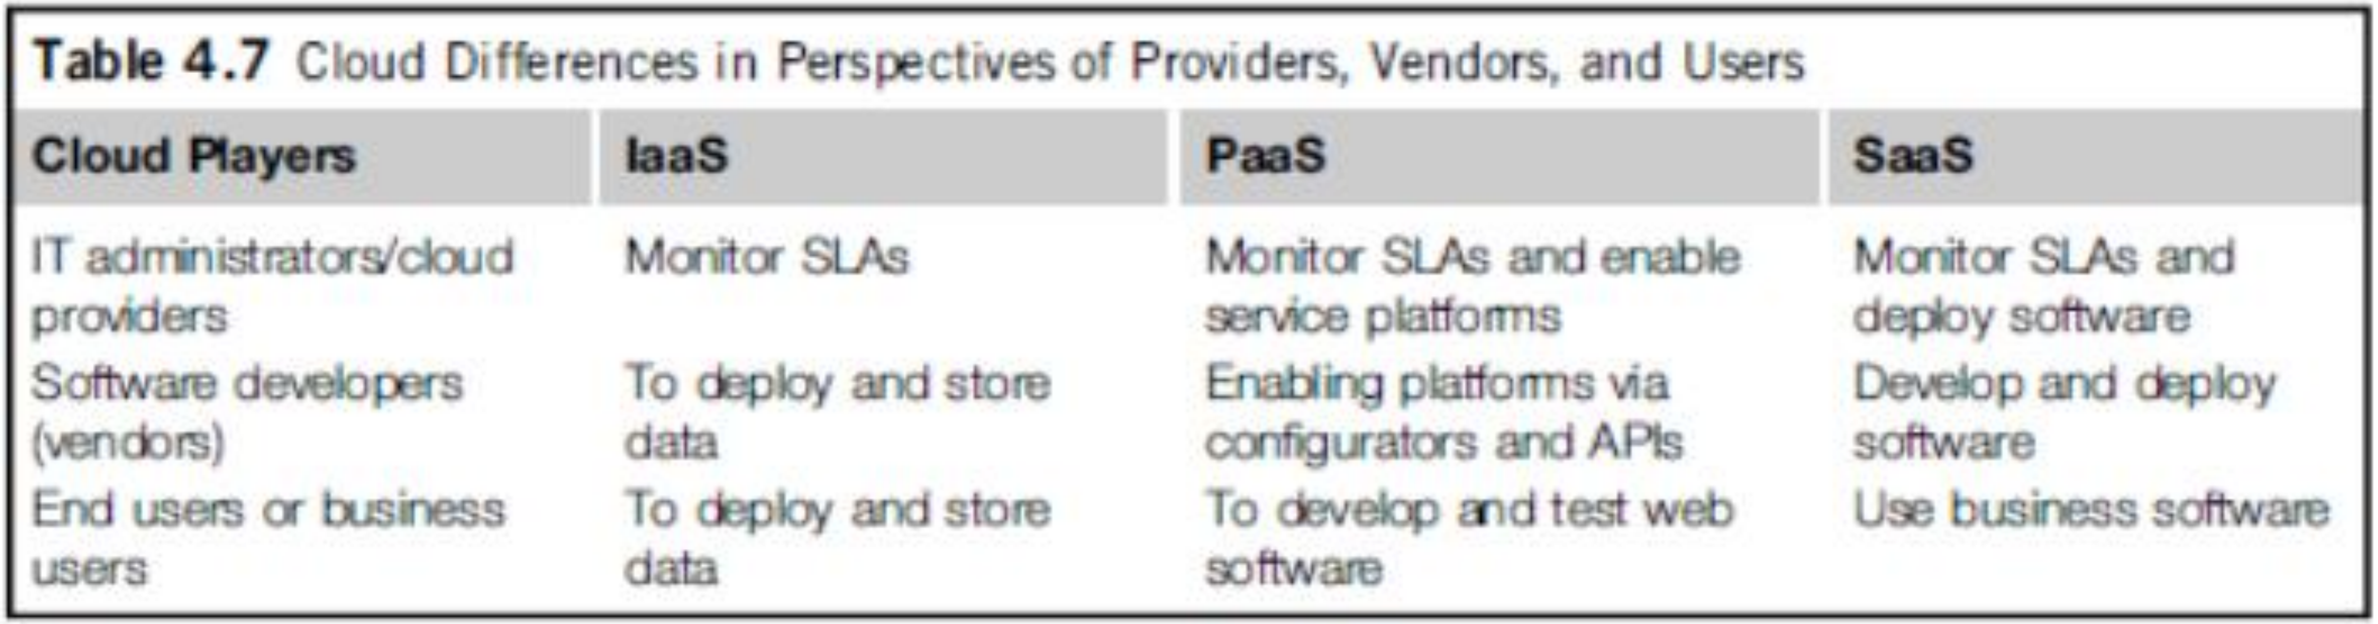

Table 4.7 Cloud Differences in Perspectives of Providers, Vendors, and Users

Cloud Players	IaaS	PaaS	SaaS
IT administrators/cloud providers	Monitor SLAs	Monitor SLAs and enable service platforms	Monitor SLAs and deploy software
Software developers (vendors)	To deploy and store data	Enabling platforms via configurators and APIs	Develop and deploy software
End users or business users	To deploy and store data	To develop and test web software	Use business software

Software vendors’ perspective- application performance on a given cloud platform is most important

Providers’ perspective – cloud infrastructure performance is primary concern

End users perspective – quality of service, including security is important

Cloud service tasks and trends

- Cloud services are introduced in five layers
- Top layer is for SaaS applications- mostly for business apps like CRM, distributed collaboration, financial and human resource mgt.
- The approach in CRM is to widen market coverage by investigating customer behaviors and revealing opportunities by statistical analysis.
- Below that is PaaS- by Google, salesforce.com, facebook
- Below PaaS layer is IaaS-by Amazon, windows Azure, Rack Rack
- Below PaaS- Collocation services-requires multiple cloud providers work together to support supply chain in manufacturing.
- Network cloud services-provide communications such as those by AR&T, Qwest, AboveNet etc.

Software stack for Cloud Computing

- Despite heterogeneous nodes in cloud overall s/w stacks are built from scratch to meet rigorous goals.
- Developer have to consider how to design the system to meet critical requirements such as high throughput, High Availability, fault tolerance
- OS also can be modified
- The S/W is built into various layer, and the upper layer provides interface to lower layers.
- The top layer is mainly for storing massive amounts of data and acts like the file system in traditional single machine.
- Other layers running on top of the file system ae the layers for executing cloud computing applications like database storage system, data query support etc.

Runtime Support Services

- Cluster monitoring- used to collect the runtime status of the entire cluster.
- Cluster job management- the scheduler queues the tasks submitted to the whole cluster and assigns the tasks to the processing nodes according to node availability.
- Runtime support- is s/w needed in browser-initiated applications applied by thousands of cloud customers.
- SaaS model provides the s/w apps as a service, rather than letting user purchase the s/w.
- So from customer side there is no investment.
- From provider side, costs are low

Resource provisioning and Platform Deployment

The emergence of computing clouds suggests fundamental changes in software and hardware architecture.

Provisioning of compute Resources (VMs)

- Providers supply cloud services by signing SLA with end users. They must provide sufficient resources to users. Under provisioning will lead to broken SLAs and penalties. Overprovisioning will lead to resource underutilization which in turn decrease the revenue of providers.
- Efficient VM provision demands efficient installation of VMs, live migration, and fast recovery from failures.
- To deploy VMs, users treat them as physical hosts with customized os for specific applications.
- In Amazon EC2 platform, predefined VM templates are provides from which users can choose.
- IBM’s BlueCloud does not provide any VM templates.
- Power-efficient schemes should be adopted.(for saving thermal power and controlling of heat dissipation from data centers)
- Public and private cloud must promise all these things

- Resource Provisioning Methods
- Fig a)Overprovisioning with the peak load causes heavy resource waste(shaded area)
- Fig b) Underprovisioning of resources results in losses by both user and provider in that paid demand by the user(shaded area above the capacity) is not served and wasted resources still exist for those demanded areas below the provisioned capacity.
- Fig c)The constant provisioning of resources with fixed capacity to a declining user demand could result in even worse resource waste.
- The user may give up service resulting in reduced revenue for the provider.
- Three methods are given:
- Demand –driven method
- Even-driven method
- Popularity – driven method

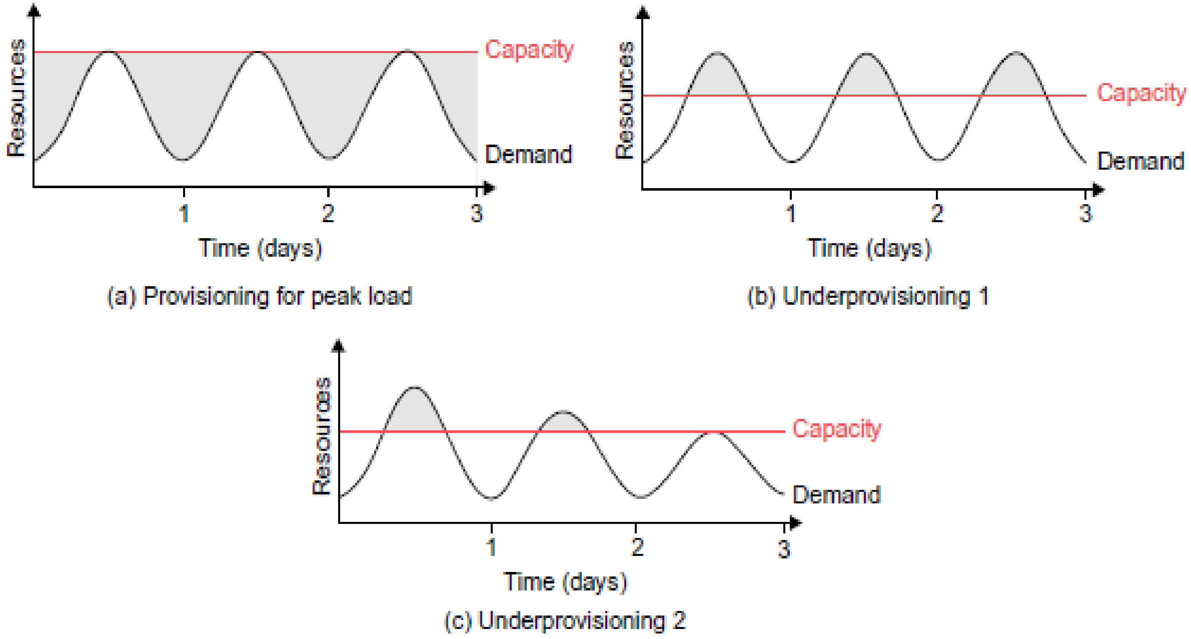

(a) Provisioning for peak load

(b) Underprovisioning 1

(c) Underprovisioning 2

FIGURE 4.24

Three cases of cloud resource provisioning without elasticity: (a) heavy waste due to overprovisioning, (b) underprovisioning and (c) under- and then overprovisioning.

- Demand-driven resource provisioning- provides static resources and been in use by grid computing
- Adds or removes computing instances based on current utilization level of allocated resources.
- It keeps a threshold and when the user uses above this it automatically allocates more resource. But when the user uses it below the threshold the particular resource allocation is reduced.
- Amazon implements this
- In the figure, in the beginning
- Heavy fluctuations of CPU load
- Are encountered. Gradually the
- Utilization rate is stabilized

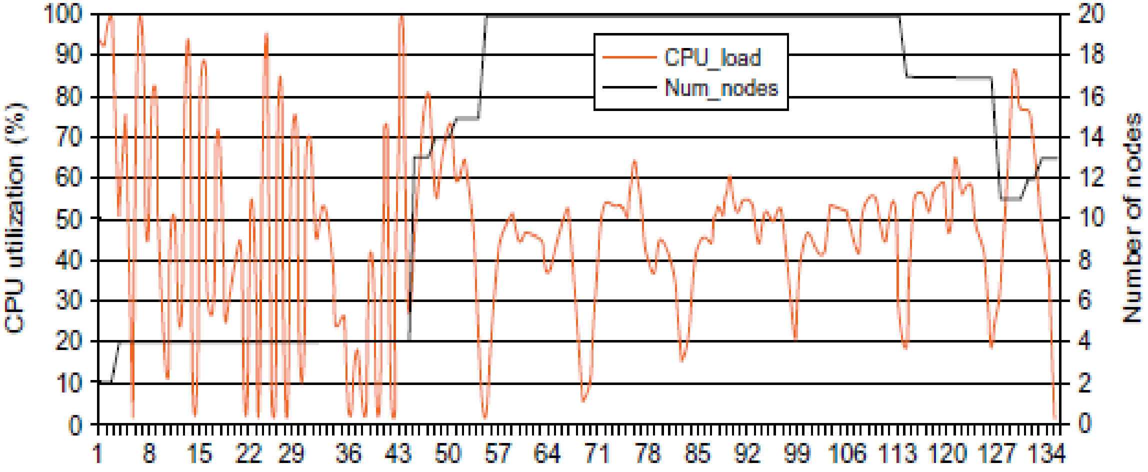

(a) Demand-driven

- Event-driven resource provisioning
- This scheme adds or removes machine instances based on a specific time event.this works well for season time like Christmas time etc.
- During these events, the number of users grows before the event period and then decrease.
- This scheme anticipates peak time . This results in minimal loss of QoS if the event is predicted correctly. Other wise it will be big waste of resources.

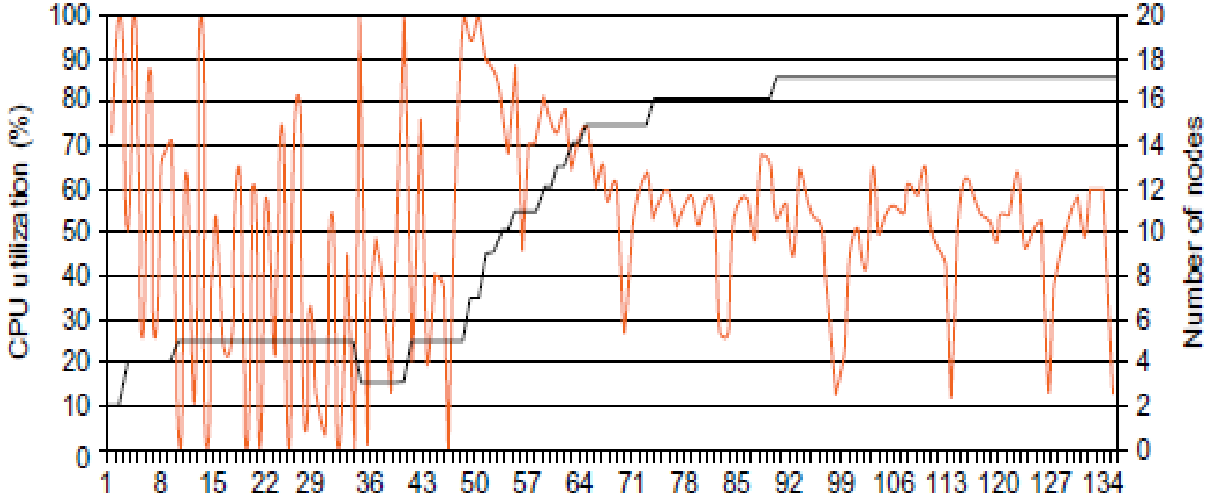

(b) Event-driven

- Popularity-driven resource provisioning
- The internet searches for popularity of certain applications and creates the instances by popularity demand.
- This scheme anticipates increased traffic with popularity.
- This results in minimal loss of QoS if the event is predicted correctly. Other wise it will be big waste of resources

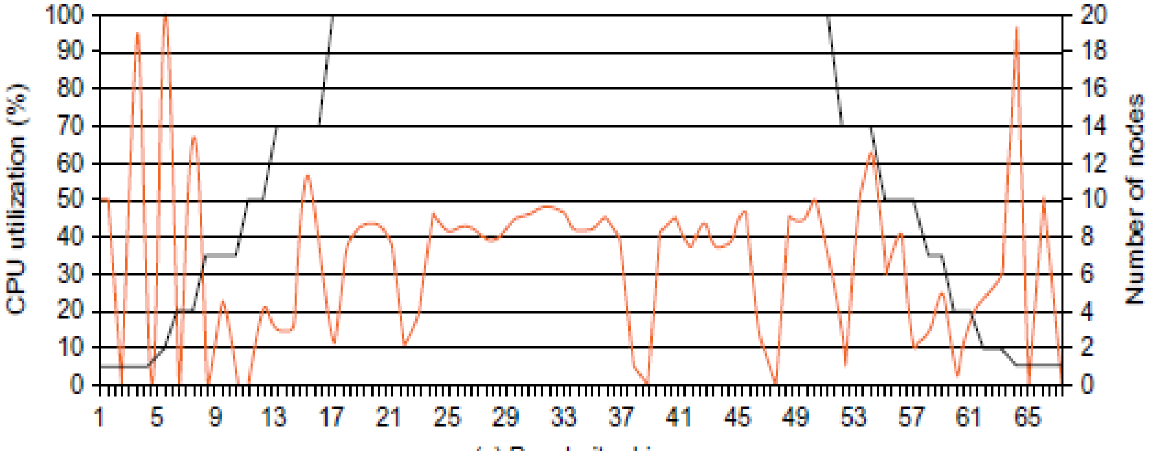

(c) Popularity-driven

Dynamic resource deployment

- The cloud uses VMs as building blocks to create an execution environment across multiple resource sites.
- It uses inter-grid managed infrastructure.
- This lets the users to create cloud environment on top of all participating grid resources.
- An intergrid gateway(IGG) allocates resources from a local cluster to deploy applications in three steps: 1) requesting VMs, 2) enacting the leases 3) deploying the VMs requested.
- Under peak demand, this IGG interacts with another IGG and allocate resources.

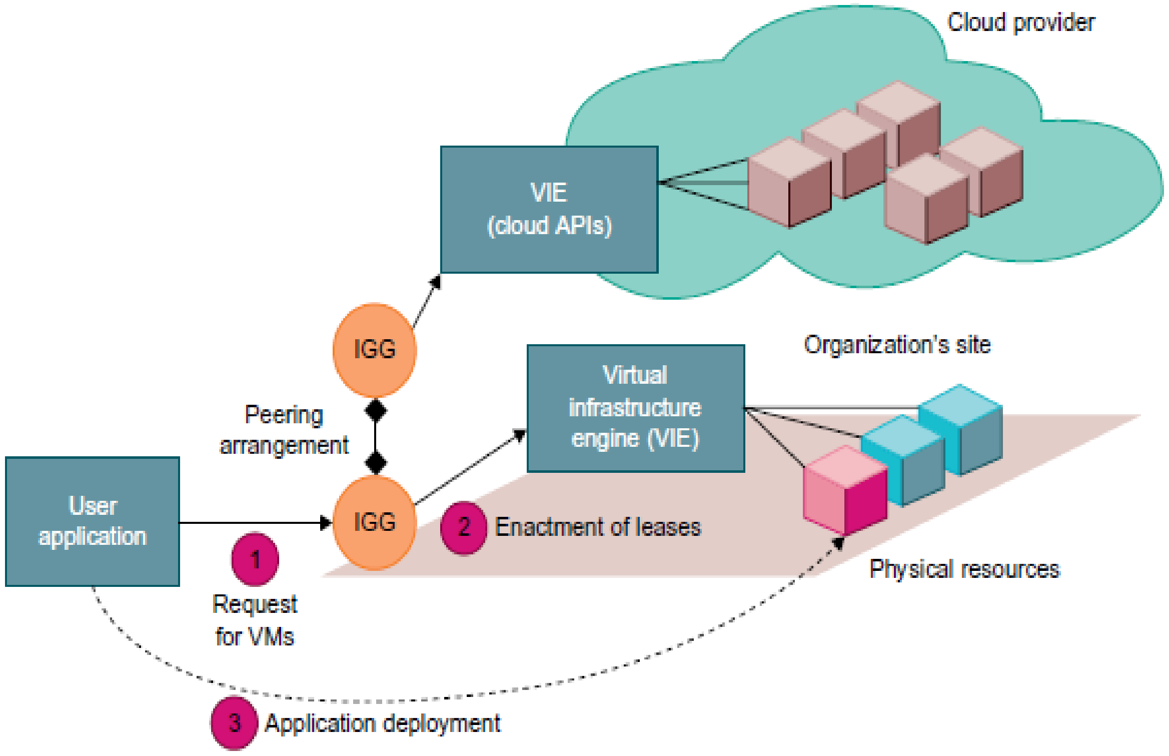

FIGURE 4.26

Cloud resource deployment using an IGG (intergrid gateway) to allocate the VMs from a Local cluster to interact with the IGG of a public cloud provider.

Provisioning of storage resources

- A user may use many emails which may consists of lot of data to be stored.
- A distributed file system is very important for storing large-scale data.
- Google GFS stores web data and some other data such as geographic data for google earth.
- Like this many providers adopt their own way of storing data.
- Like traditional storage, cloud computing also provides some form of structure database.
- For eg webdata which is in semi-structured (HTML)
- If some forms of database capability can be used, in which the developers are trained, then it will be each in cloud application development. Eg for cloud databases are BigTable from Google, SimpleDB from Amazon, SQL service from Microsoft Azure.

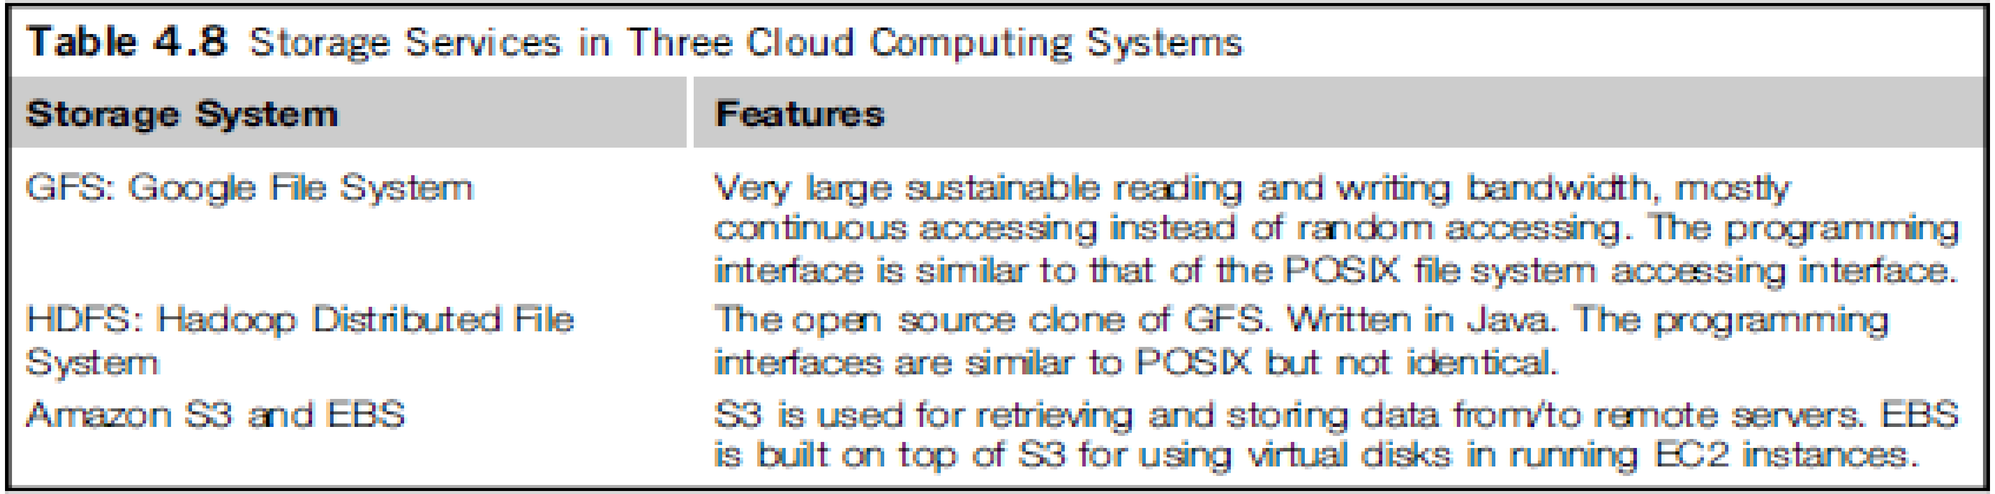

Table 4.8 Storage Services in Three Cloud Computing Systems

Storage System	Features
GFS: Google File System	Very large sustainable reading and writing bandwidth, mostly continuous accessing instead of random accessing. The programming interface is similar to that of the POSIX file system accessing interface.
HDFS: Hadoop Distributed File System	The open source clone of GFS. Written in Java. The programming interfaces are similar to POSIX but not identical.
Amazon S3 and EBS	S3 is used for retrieving and storing data from/to remote servers. EBS is built on top of S3 for using virtual disks in running EC2 instances.

Global Exchange of Cloud Resources

- To support large number of consumers around the world cloud infrastructures providers have established data centers in multiple geographical locations.
- Amazon has data centeres in US(east coast, west coast), Europe etc. but is asking consumer to select the preferred place which if difficult for them to determine in advance.
- Another short coming is Qos expectation of consumer differs with geographical region which is difficult to decide by SaaS providers.
- This leads to opening of data centers in different federation to meet QoS of different customers.
- But this is quite not possible in reality as opening of this much center.
- Hence they would like to make use of services of multiple cloud infrastructure service providers who can provide better support for their specific consumer needs.
- The cloud Exchange (Cex) acts as a market maker for bringing together service producers and consumers.
- It aggregates the infrastructure demands from application brokers and evaluates them against the available supply cuuently published by the cloud coordinators.
- This allows participants to locate providers and consumers with fitting offers.
- Such markets enable services to be commoditized, and thus will pave the way for creation of dynamic market infrastructure for trading based on SLAs.
- An SLA specifies the details of service to be provides in terms of metrics agreed upon by all parties, and incentives and penalties for meeting and violating the expectations.

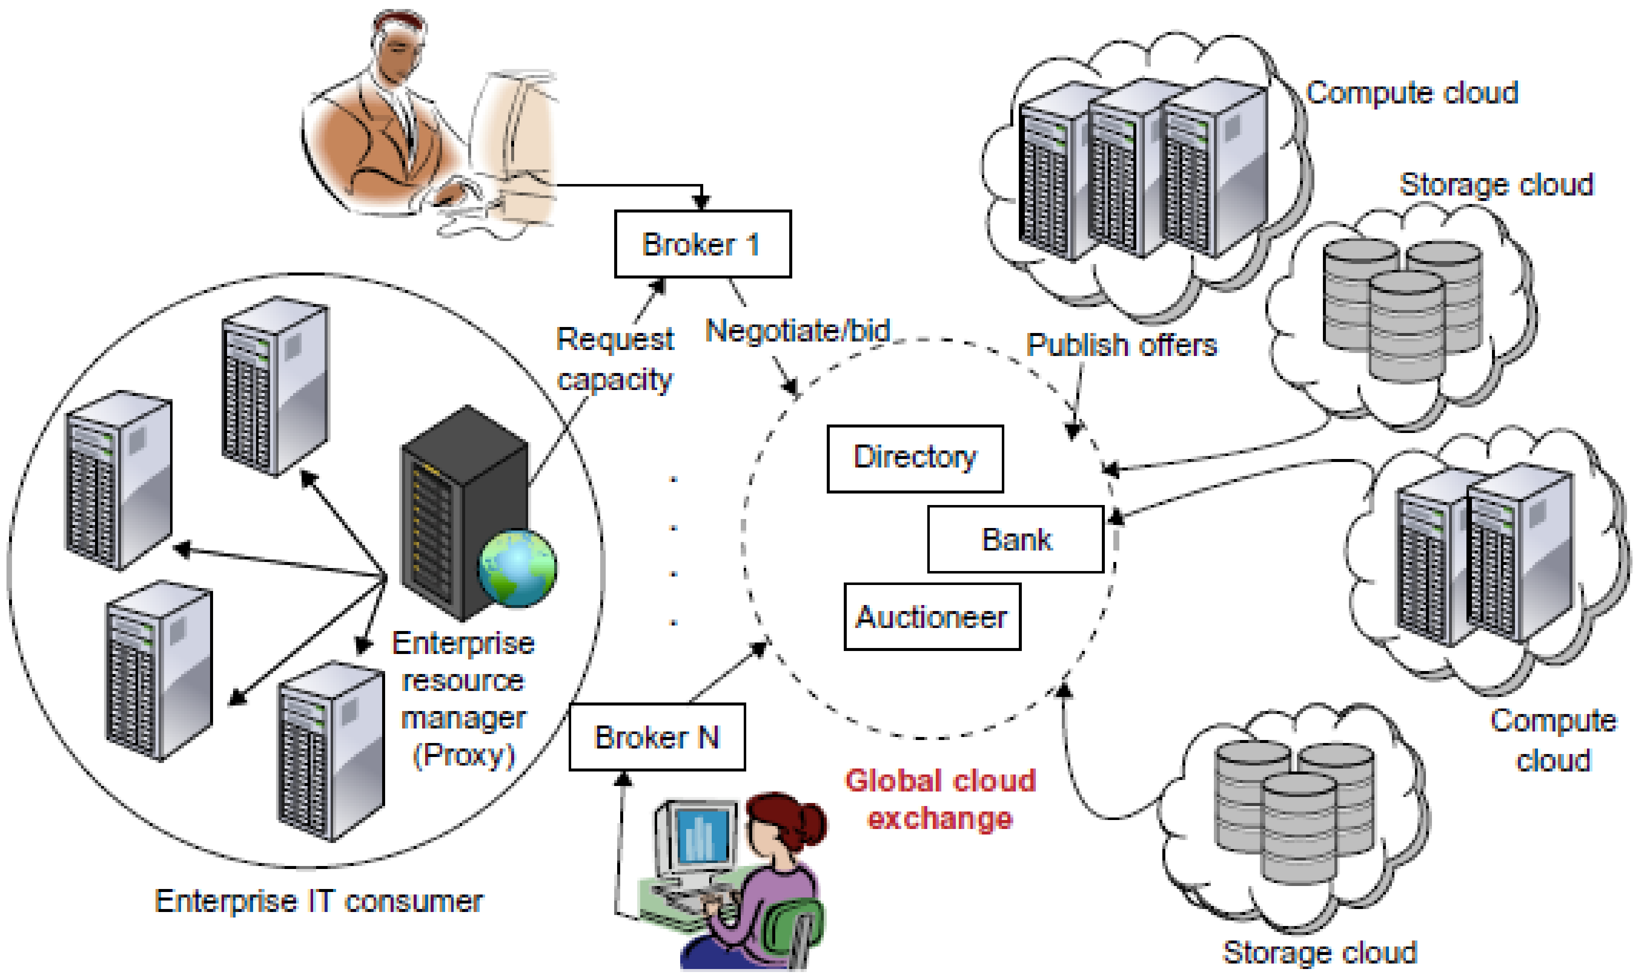

FIGURE 4.30

Inter-cloud exchange of cloud resources through brokering.