Advanced Software Engineering

Required Courses: Software Engineering, Human Computer Interaction
Teacher: Vidakis Nikolaos PhD.
Language: Greek or English

Cloud Computing

Acknowledgement

Some material presented at this lecture is originated at the following sources

- "Cloud Computing – An overview" by Suganth.N
- Amazon Web Services, Jeff Barr, jbarr@amazon.com, Senior Web Services Evangelist, Amazon Web Services
- CLOUD COMPUTING, Hà Lê Hoài Trung 09070473, Võ Lâm Khang 09070446
- An Introduction to SaaS and Cloud Computing, Ross Cooney
- An Introduction and Overview of Cloud Computing, by: Prof Mark Baker
- What is Cloud Computing?, Jimmy Lin, The iSchool, University of Maryland

Cloud Computing - Some terms

- Term cloud is used as a metaphor for internet

Source	Definition
Gartner	"a style of computing in which massively scalable IT-related capabilities are provided "as a service" using Internet technologies to multiple external customers" (Gartner 2008b)
IDC	"an emerging IT development, deployment and delivery model, enabling real-time delivery of products, services and solutions over the Internet (i.e., enabling cloud services)" (Gens 2008)
The 451 Group	"a service model that combines a general organizing principle for IT delivery, infrastructure components, an architectural approach and an economic model – basically, a confluence of grid computing, virtualization, utility computing, hosting and software as a service (SaaS)" (Fellows 2008)
Merrill Lynch	"the idea of delivering personal (e.g., email, word processing, presentations.) and business productivity applications (e.g., sales force automation, customer service, accounting) from centralized servers" (Merrill Lynch 2008)

Definitions

- Cloud computing is using the internet to access someone else's software running on someone else's hardware in someone else's data center. Lewis Cunningham[2]

- A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet. Ian Foster[9]

- A Cloud is a type of parallel and distributed system consisting of a collection of interconnected and virtualized computers that are dynamically provisioned and presented as one or more unified computing resources based on service-level agreements established through negotiation between the service provider and consumers. Rajkumar Buyya[10]

Cloud Computing - Some terms

- Concept generally incorporates combinations of the following
  - Infrastructure as a service (IaaS)
  - Platform as a service (PaaS)
  - Software as a service(SaaS)
- Not to be confused with
  - Grid Computing – a form of distributed computing
    - Cluster of loosely coupled, networked computers acting in concert to perform very large tasks
  - Utility Computing – packaging of computing resources such as computing power, storage, also a metered services
  - Autonomic computing – self managed

Grid Computing

- Share Computers and data
- Evolved to harness inexpensive computers in Data center to solve variety of problems
- Harness power of loosely coupled computers to solve a technical or mathematical problem
- Used in commercial applications for drug discovery, economic forecasting, seismic analysis and back-office
- Small to big
  - Can be confined to a corporation
  - Large public collaboration across many companies and networks

Grid Computing

- Compute grids
  - Batch up jobs
  - Submit the job to the scheduler, specifying requirements and SLA(specs) required for running the job
  - Scheduler matches specs with available resources and schedules the job to be run
  - Farms could be as large as 10K cpus
- Most financial firms has grids like this
- Grids lack automation, agility, simplicity and SLA guarantees
- Most grid solutions are built on
  - Computer Agents
  - Resource Manager
  - Scheduler

Utility Computing

- More related to cloud computing
  - Applications, storage, computing power and network
- Requires cloud like infrastructure
- Pay by the drink model
  - Similar to electric service at home
- Pay for extra resources when needed
  - To handle expected surge in demand
  - Unanticipated surges in demand
- Better economics

Cloud computing – History

- Evolved over a period of time
- Roots traced back to Application Service Providers in the 1990's
- Parallels to SaaS
- Evolved from Utility computing and is a broader concept

Cloud computing

- Much more broader concept
- Encompasses
  - IIAS, PAAS, SAAS
- Dynamic provision of services/resource pools in a co-ordinated fashion
  - On demand computing – No waiting period
  - Location of resource is irrelevant
    - May be relevant from performance(network latency) perspective, data locality
- Applications run somewhere on the cloud
  - Web applications fulfill these for end user
  - However, for application developers and IT
    - Allows develop, deploy and run applications that can easily grow capacity(scalability), work fast(performance), and offer good reliability
    - Without concern for the nature and location of underlying infrastructure
  - Activate, retire resources
  - Dynamically update infrastructure elements without affecting the business

Clouds Versus Grids

- Clouds and Grids are distinct
- Cloud
  - Full private cluster is provisioned
  - Individual user can only get a tiny fraction of the total resource pool
  - No support for cloud federation except through the client interface
  - Opaque with respect to resources
- Grid
  - Built so that individual users can get most, if not all of the resources in a single request
  - Middleware approach takes federation as a first principle
  - Resources are exposed, often as bare metal
- These differences mandate different architectures for each

Cloud versus Grid

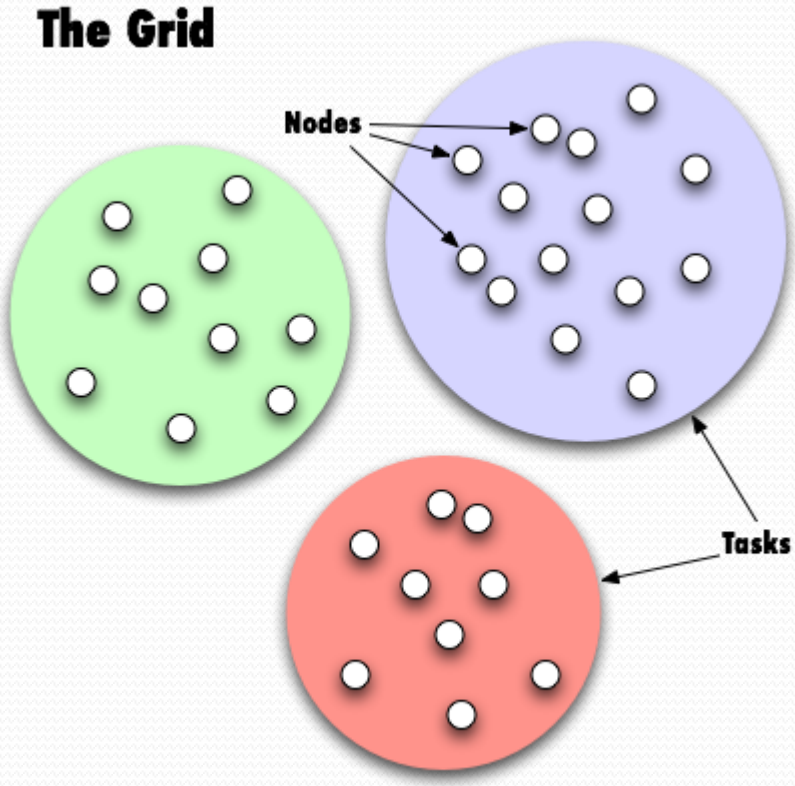

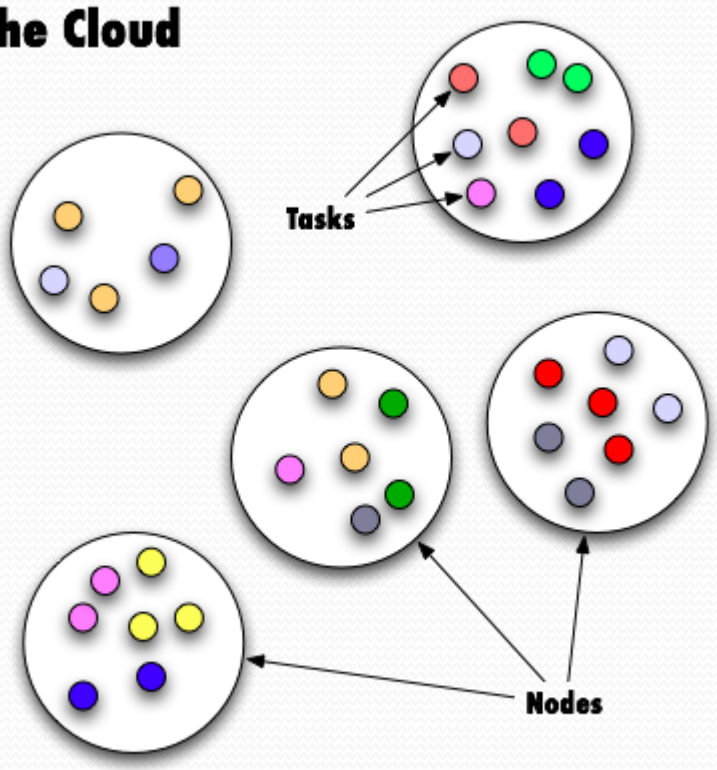

The grid. Different colored jobs belong to different clients. (One of those jobs belongs to the Department of Defense.)

The cloud. Different colored jobs belong to different clients. (One of those jobs belongs to your 18-year-old nephew.)

Images originated at : http://arstechnica.com/business/2009/11/the-cloud-a-short-introduction/2/

Cloud Mythologies

- Cloud computing infrastructure is just a web service interface to operating system virtualization.
  - "I'm running Xen in my data center – I'm running a private cloud."
- Cloud computing imposes a significant performance penalty over "bare metal" provisioning.
  - "I won't be able to run a private cloud because my users will not tolerate the performance hit."
- Clouds and Grids are equivalent
  - "In the mid 1990s, the term grid was coined to describe technologies that would allow consumers to obtain computing power on demand."

Virtualization

- Virtual workspaces:
  - An abstraction of an execution environment that can be made dynamically available to authorised clients by using well-defined protocols,
  - Resource quota (e.g. CPU, memory share),
  - Software configuration (e.g. O/S, provided services).
- Implement on Virtual Machines (VMs):
  - Abstraction of a physical host machine,
  - Hypervisor intercepts and emulates instructions from VMs, and allows management of VMs,
  - VMWare, Xen, etc.
- Provide infrastructure API:
  - Plug-ins to hardware/support structures

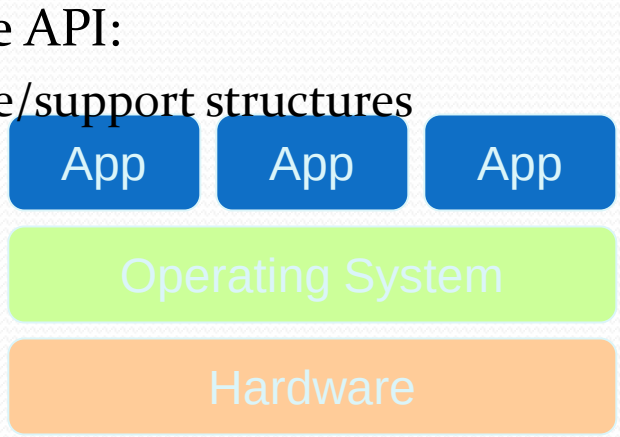

Traditional Stack

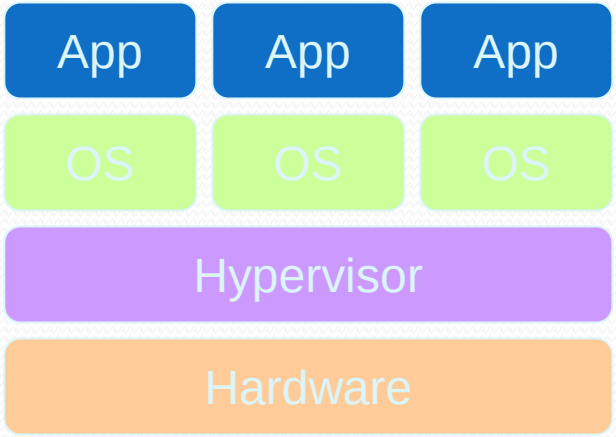

Virtualized Stack

Clouds and Virtualization

- Operating System virtualization (Xen, KVM, VMWare, HyperV) is only apparent for IaaS
  - AppEngine = BigTable
- Hypervisors virtualize CPU, Memory, and local device access as a single virtual machine (VM)
- IaaS Cloud allocation is
  - Set of VMs
  - Set of storage resources
  - Private network
- Allocation is atomic
- SLA
- Monitoring

Virtual Machines

- VM technology allows multiple virtual machines to run on a single physical machine.

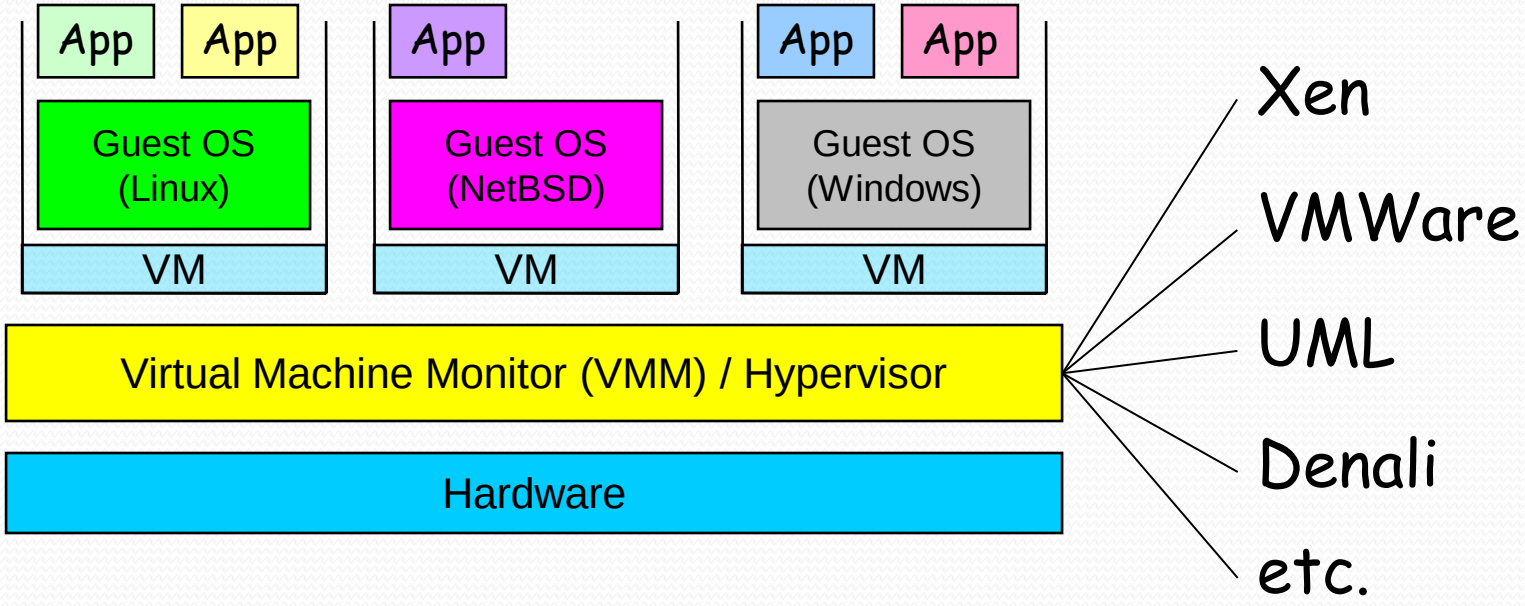

Performance: Para-virtualization (e.g. Xen) is very close to raw physical performance!

Divide and Conquer

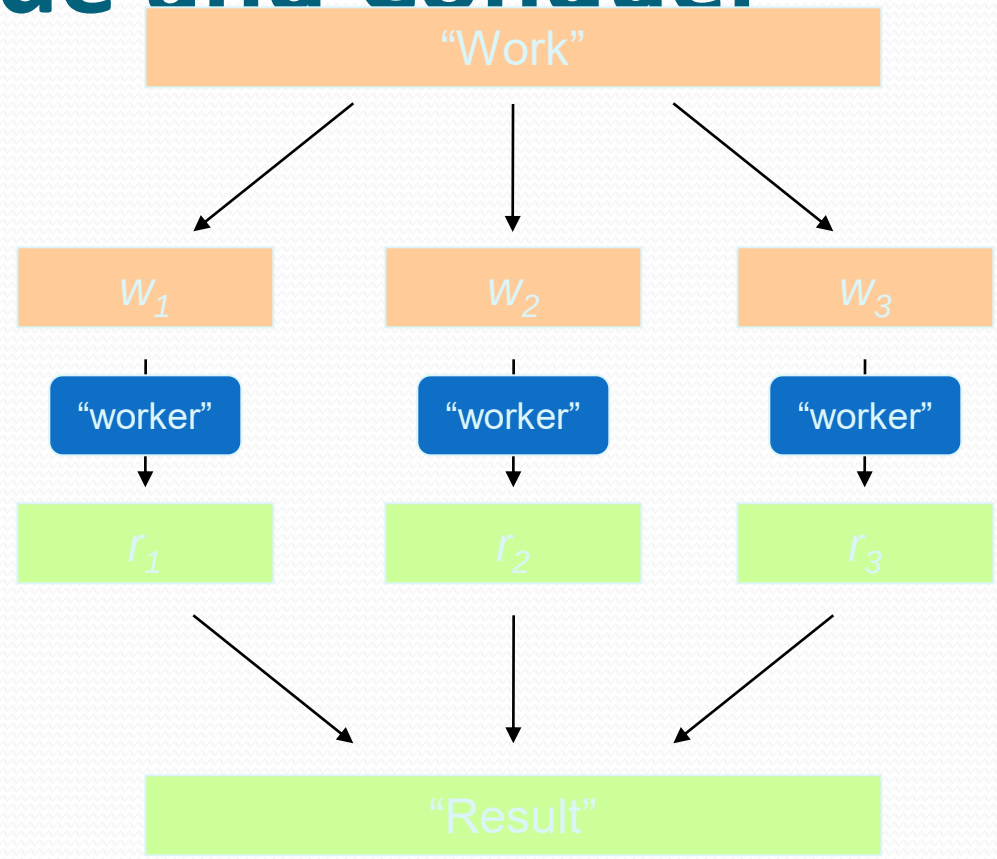

Partition

Combine

Architecture

NIST Visual Model of Cloud Computing Definition

Cloud Architecture

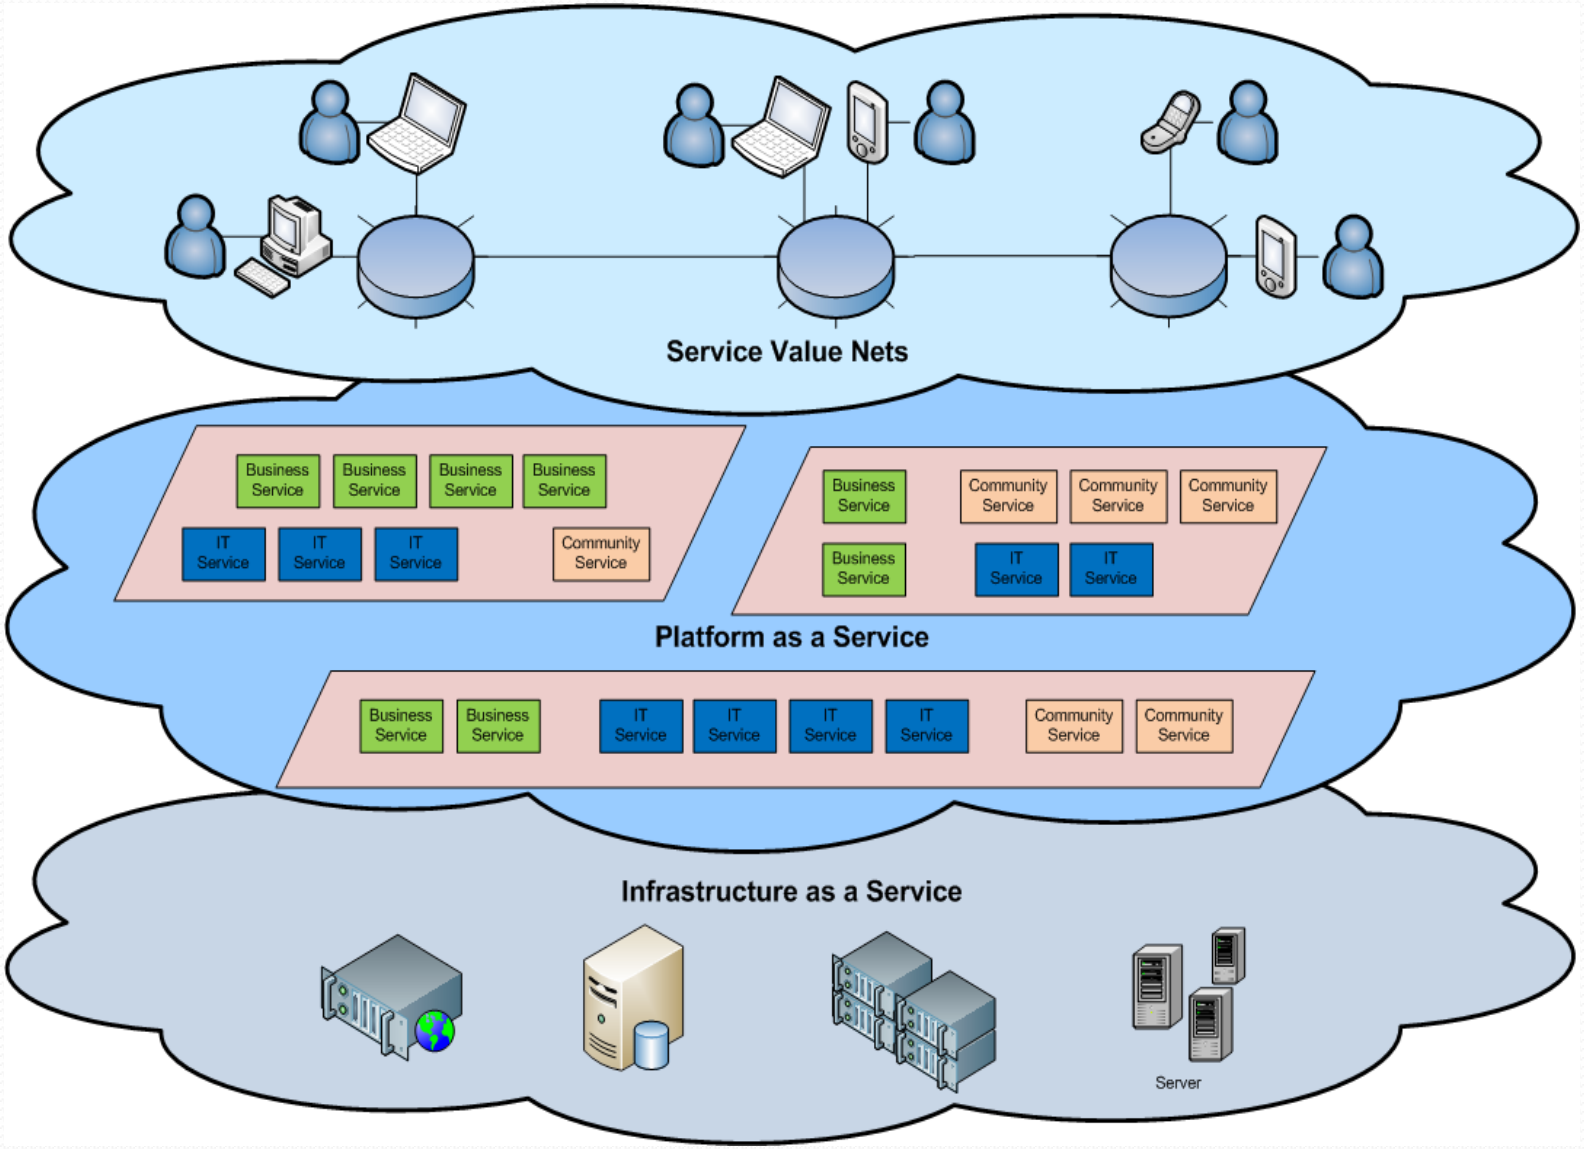

Vidakis Nikolaos
mark.baker@computer.org
Artificial Intelligence and Systems Engineering Lab
Department of Informatics Engineering and Electical Engineering, TEI of Crete

Essential Characteristics[7]

- **On-demand self-service.**
  - A consumer can unilaterally provision computing capabilities such as server time and network storage as needed automatically, without requiring human interaction with a service provider.
- **Broad network access.**
  - Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, laptops, and PDAs) as well as other traditional or cloud-based software services.
- **Resource pooling.**
  - The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand.
- **Rapid elasticity.**
  - Capabilities can be rapidly and elastically provisioned - in some cases automatically - to quickly scale out; and rapidly released to quickly scale in.
  - To the consumer, the capabilities available for provisioning often appear to be unlimited and can be purchased in any quantity at any time.
- **Measured service.**
  - Cloud systems automatically control and optimize resource usage by leveraging a metering capability at some level of abstraction appropriate to the type of service.
  - Resource usage can be monitored, controlled, and reported - providing transparency for both the provider and consumer of the service.

Cloud Service Models

SPI Model (**S**aaS, **P**aaS, **I**aaS)

- Cloud **S**oftware as a Service (**S**aaS)
  - The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure.
  - The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email).
  - The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited userspecific application configuration settings.
- Cloud **P**latform as a Service (**P**aaS)
  - The capability provided to the consumer is to deploy onto the cloud infrastructure consumer created or acquired applications created using programming languages and tools supported by the provider.
  - The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.
- Cloud **I**nfrastructure as a Service (**I**aaS)
  - The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources.
  - Consumer is able to deploy and run arbitrary software, which can include operating systems and applications.
  - The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).

Cloud Software as a Service (SaaS)

- Increasingly popular with SMEs
- No hardware or software to manage
- Service delivered through a browser

Advantages

- Pay per use
- Instant Scalability
- Security
- Reliability
- APIs

Examples

- CRM
- Financial Planning
- Human Resources
- Word processing

Commercial Services:

- Salesforce.com
- emailcloud

Cloud Platform as a Service (PaaS)

- Platforms are built upon Infrastructure, which is expensive
- Estimating demand is not a science!
- Platform management is not fun!

Advantages

- Pay per use
- Instant Scalability
- Security
- Reliability
- APIs

Examples

- Google App Engine
- Mosso
- AWS: S3

Popular services

- Storage
- Database
- Scalability

Cloud Infrastructure as a Service (IaaS)

- Access to infrastructure stack:
  - Full OS access
  - Firewalls
  - Routers
  - Load balancing

Examples

- Flexiscale
- AWS: EC2

Advantages

- Pay per use
- Instant Scalability
- Security
- Reliability
- APIs

Commercial clouds

Amazon Elastic Compute Cloud (Amazon EC2) - *Beta*

Cloud Taxonomy

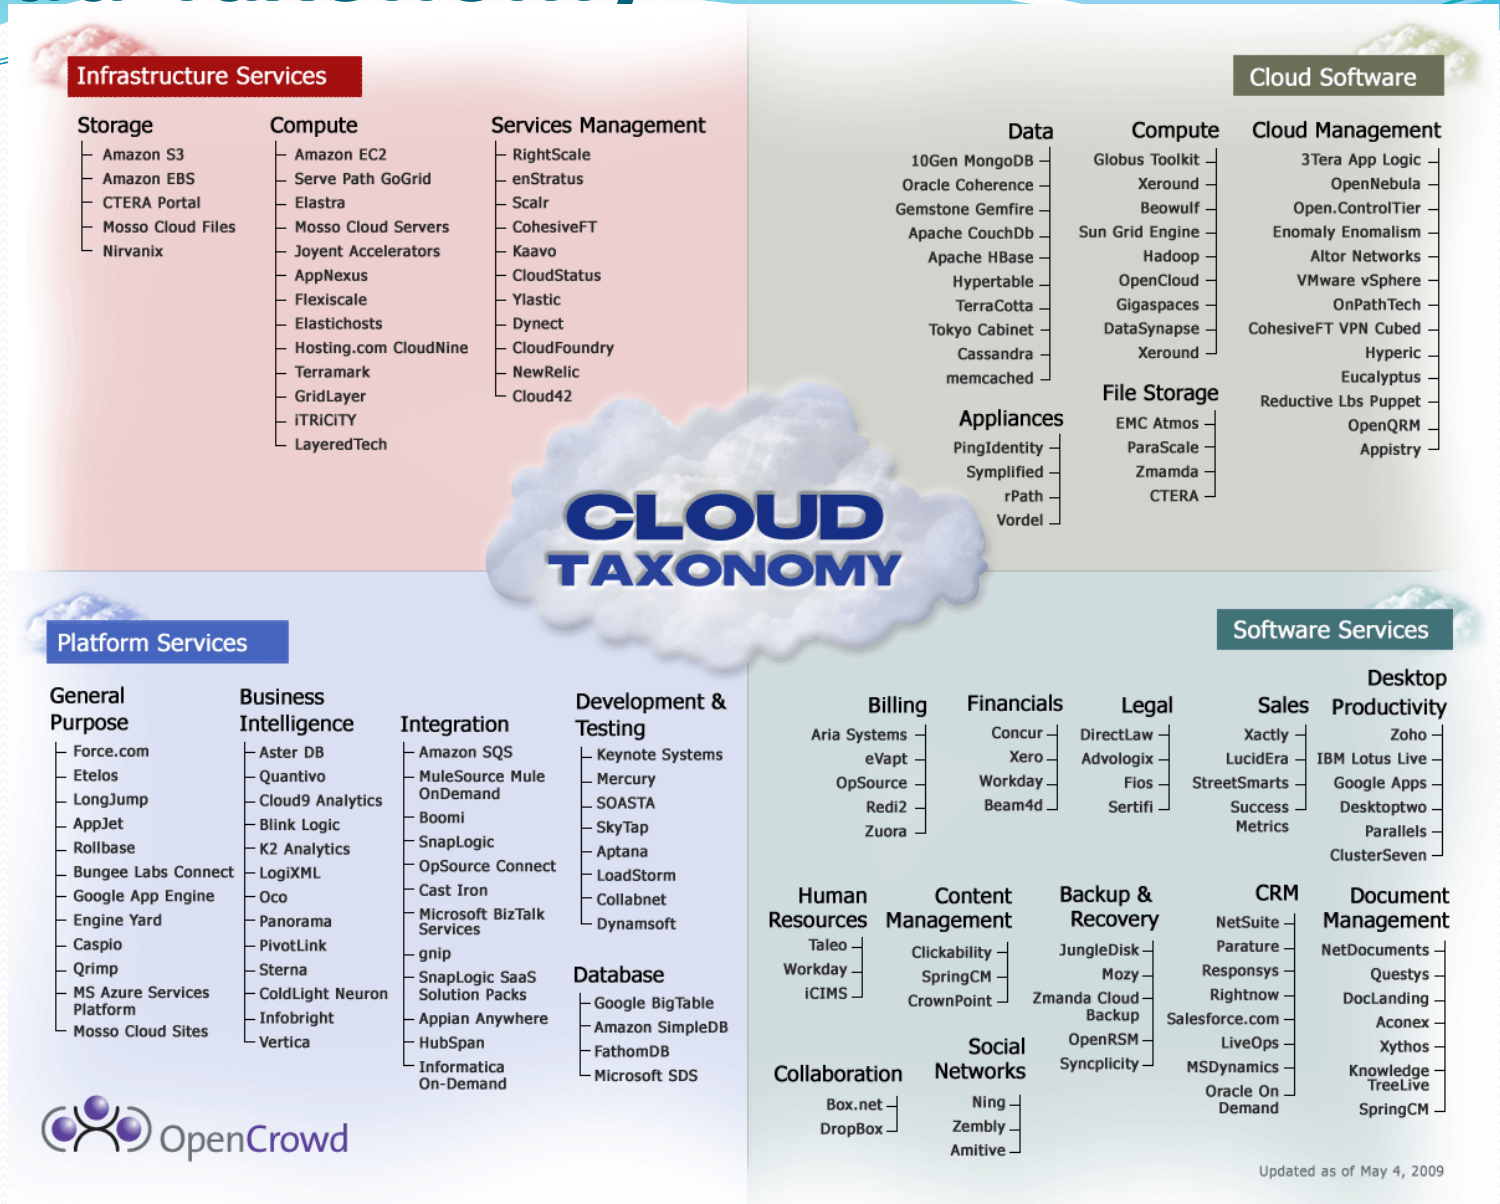

19th May, 09

mark.baker@computer.org

Why cloud computing

- Data centers are notoriously underutilized, often idle 85% of the time
  - Over provisioning
  - Insufficient capacity planning and sizing
  - Improper understanding of scalability requirements etc
- including thought leaders from Gartner, Forrester, and IDC—agree that this new model offers significant advantages for fast-paced startups, SMBs and enterprises alike.
- Cost effective solutions to key business demands
- Move workloads to improve efficiency

How do they work?

- Public clouds are opaque
  - What applications will work well in a cloud?
- Many of the advantages offered by Public Clouds appear useful for “on premise” IT
  - Self-service provisioning
  - Legacy support
  - Flexible resource allocation
- What extensions or modifications are required to support a wider variety of services and applications?
  - Data assimilation
  - Multiplayer gaming
  - Mobile devices

Cloud computing - Characteristics

- Agility – On demand computing infrastructure
  - Linearly scalable – challenge
- Reliability and fault tolerance
  - Self healing – Hot backups, etc
  - SLA driven – Policies on how quickly requests are processed
- Multi-tenancy – Several customers share infrastructure, without compromising privacy and security of each of the customer's data
- Service-oriented – compose applications out of loosely coupled services. One service failure will not disrupt other services. Expose these services as API's
- Virtualized – decoupled from underlying hardware. Multiple applications can run in one computer
- Data, Data, Data
  - Distributing, partitioning, security, and synchronization

Public, Private and Hybrid clouds

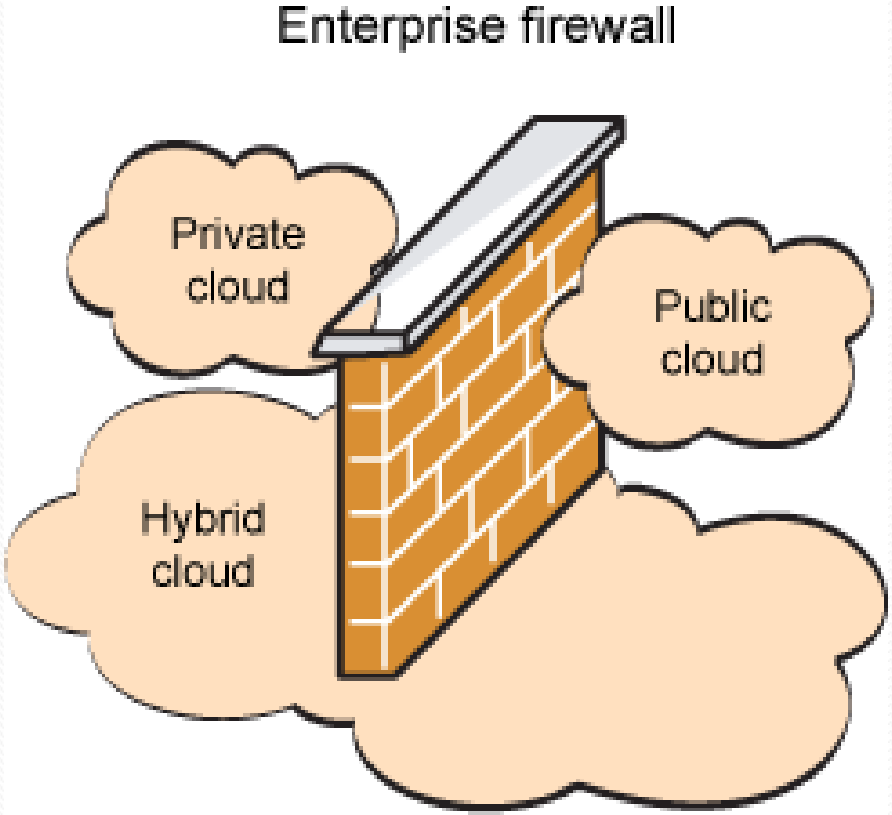

Public clouds

- Open for use by general public
  - Exist beyond firewall, fully hosted and managed by the vendor
  - Individuals, corporations and others
  - Amazon's Web Services and Google appEngine are examples
- Offers startups and SMB's quick setup, scalability, flexibility and automated management. Pay as you go model helps startups to start small and go big
- Security and compliance?
- Reliability concerns hinder the adoption of cloud
  - Amazon S3 services were down for 6 hours

Public Clouds (Now)

- Large scale infrastructure available on a rental basis
  - Operating System virtualization (e.g. Xen, kvm) provides CPU isolation
  - “Roll-your-own” network provisioning provides network isolation
  - Locally specific storage abstractions
- Fully customer self-service
  - Service Level Agreements (SLAs) are advertized
  - Requests are accepted and resources granted via web services
  - Customers access resources remotely via the Internet
- Accountability is e-commerce based
  - Web-based transaction
  - “Pay-as-you-go” and flat-rate subscription
  - Customer service, refunds, etc.

Private Clouds

- Within the boundaries(firewall) of the organization
- All advantages of public cloud with one major difference
  - Reduce operation costs
  - Has to be managed by the enterprise
- Fine grained control over resources
- More secure as they are internal to org
- Schedule and reshuffle resources based on business demands
- Ideal for apps related to tight security and regulatory concerns
- Development requires hardware investments and in-house expertise
- Cost could be prohibitive and cost might exceed public clouds

Community Cloud

- The cloud infrastructure is shared by several organizations and supports a specific community that has shared concerns (e.g., mission, security requirements, policy, or compliance considerations). It may be managed by the organizations or a third party and may exist on-premises or off-premises.

Hybrid Cloud

- The cloud infrastructure is a composition of two or more clouds (private, community, or public) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

Private VS Public Cloud

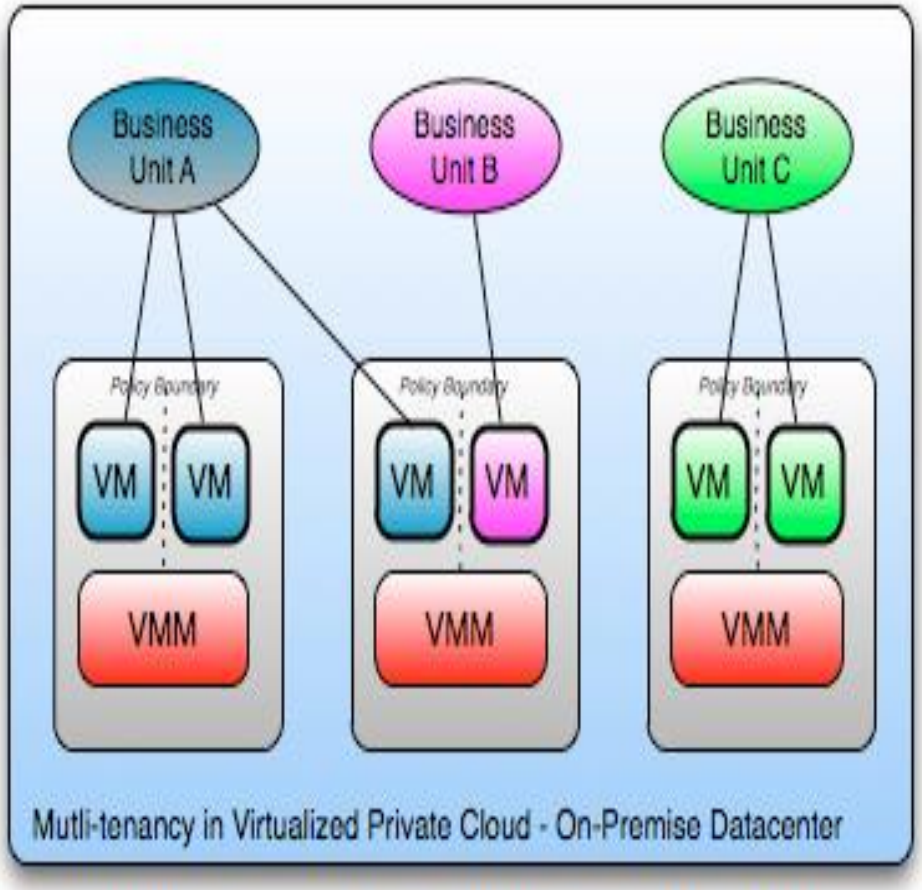

Private Cloud of Company XYZ with 3 business units, each with different security, SLA, governance and chargeback policies on shared infrastructure

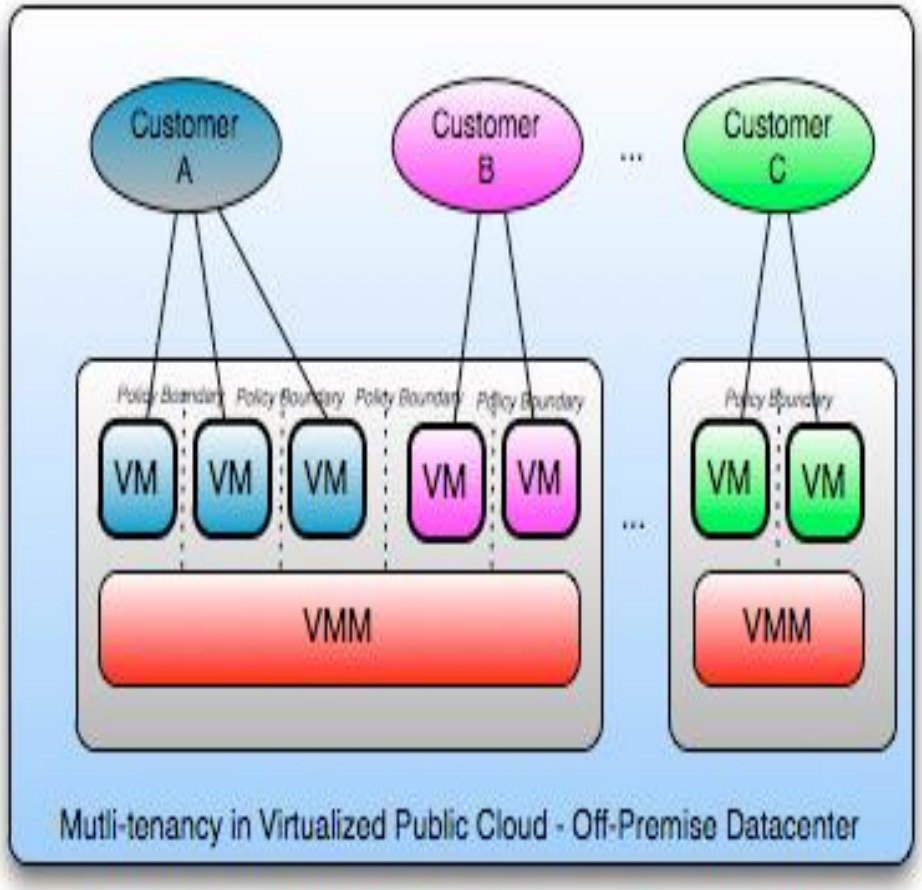

Public Cloud Provider with 3 business customers, each with different security, SLA, governance and billing policies on shared infrastructure

Clouds and SOA

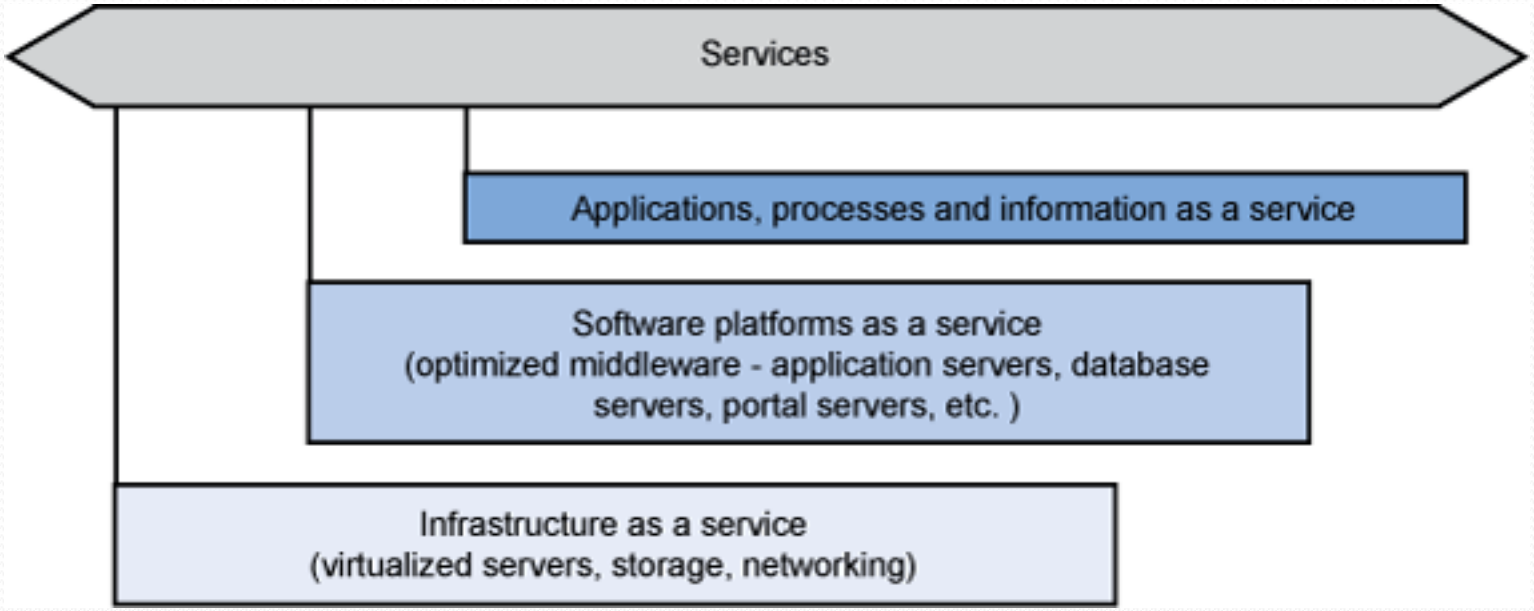

- SOA Enabled cloud computing to what is today
- Physical infrastructure like SOA must be discoverable, manageable and governable
- REST Protocol widely used(Representational State Transfer)

Clouds for Developers

- Ability to acquire, deploy, configure and host environments
- Perform development unit testing, prototyping and full product testing

Open Source Cloud Infrastructure

- Simple
  - Transparent => need to “see” into the cloud
  - Scalable => complexity often limits scalability
  - Secure => limits adoptability
- Extensible
  - New application classes and service classes may require new features
  - Clouds are new => need to extend while retaining useful features
- Commodity-based
  - Must leverage extensive catalog of open source software offerings
  - New, unstable, and unsupported infrastructure design is a barrier to uptake, experimentation, and adoption
- Easy
  - To install => system administration time is expensive
  - To maintain => system administration time is really expensive

Microsoft and Amazon face challenges

- Globus/Nimbus
  - Client-side cloud-computing interface to Globus-enabled TeraPort cluster at U of C
  - Based on GT4 and the Globus Virtual Workspace Service
  - Shares upsides and downsides of Globus-based grid technologies
- Enomalism (now called ECP)
  - Start-up company distributing open source
  - REST APIs
- Reservoir
  - European open cloud project
  - Many layers of cloud services and tools
  - Ambitious and wide-reaching but not yet accessible as an implementation
- Eucalyptus
  - Cloud Computing on Clusters
  - Amazon Web Services compatible
  - Supports kvm and Xen
- Open Nebulous
- Joyent
  - Based on Java Script and Git

Open Source Cloud Ecosystem - Tools

RightScale

- Startup focused on providing client tools as SaaS hosted in AWS
- Uses the REST interface

- Canonical
  - Ubuntu 9.10 (Karmic Koala)
  - Includes KVM and Xen Hypervisors

Open Source Cloud Anatomy

- Extensibility
  - Simple architecture and open internal APIs
- Client-side interface
  - Amazon's AWS interface and functionality (familiar and testable)
- Networking
  - Virtual private network per cloud
  - Must function as an overlay => cannot supplant local networking
- Security
  - Must be compatible with local security policies
- Packaging, installation, maintenance
  - system administration staff is an important constituency for uptake

Open Source Cloud Anatomy .. cntd

- Private clouds are really hybrid clouds
  - Users want private clouds to export the same APIs as the public clouds
- In the Enterprise, the storage model is key
  - Scalable "blob" storage doesn't quite fit the notion of "data file."
- Cloud Federation is a policy mediation problem
  - No good way to translate SLAs in a cloud allocation chain
  - "Cloud Bursting" will only work if SLAs are congruent
- Customer SLAs allow applications to consider cost as first-class principle
  - Buy the computational, network, and storage capabilities that are required

Open Source Clouds contd.

Cloud Computing at UCSB

Open-source implementations of popular cloud systems

Eucalyptus

- Infrastructure-as-a-service (IAAS) framework
- Web services based implementation of elastic/utility/cloud computing infrastructure
- Linux image hosting via virtualization
- Emulates the **Amazon AWS interface** – applications and tools can't tell the difference
- Real use, real users, real impact
  - Distributed with Ubuntu
  - Large international user community

http://eucalyptus.cs.ucsb.edu

Lead: Rich Wolski

AppScale

- Platform-as-a-service (PAAS) framework
- Web services based implementation of **Google AppEngine APIs**
- Runs over Eucalyptus, Amazon EC2, and virtualization layers (Xen/KVM)
- Implements multiple database backends (Hbase, Hypertable, Cassandra, MySQL,…)
- Real use, real users, real impact
  - International user community

http://appscale.cs.ucsb.edu

Lead: Chandra Krintz

- Open, extensible, easy to install/use/maintain, transparent, scalable
- Enables investigation of and experimentation with
  - Real applications in real settings
  - IAAS + PAAS interoperation and integration
- Frameworks for investigation the next generation of ***energy efficient*** distributed systems *technologies, languages, applications, and services*

UCSB

Eucalyptus (Elastic Utility Computing Architecture Linking Your Programs To Useful Systems)

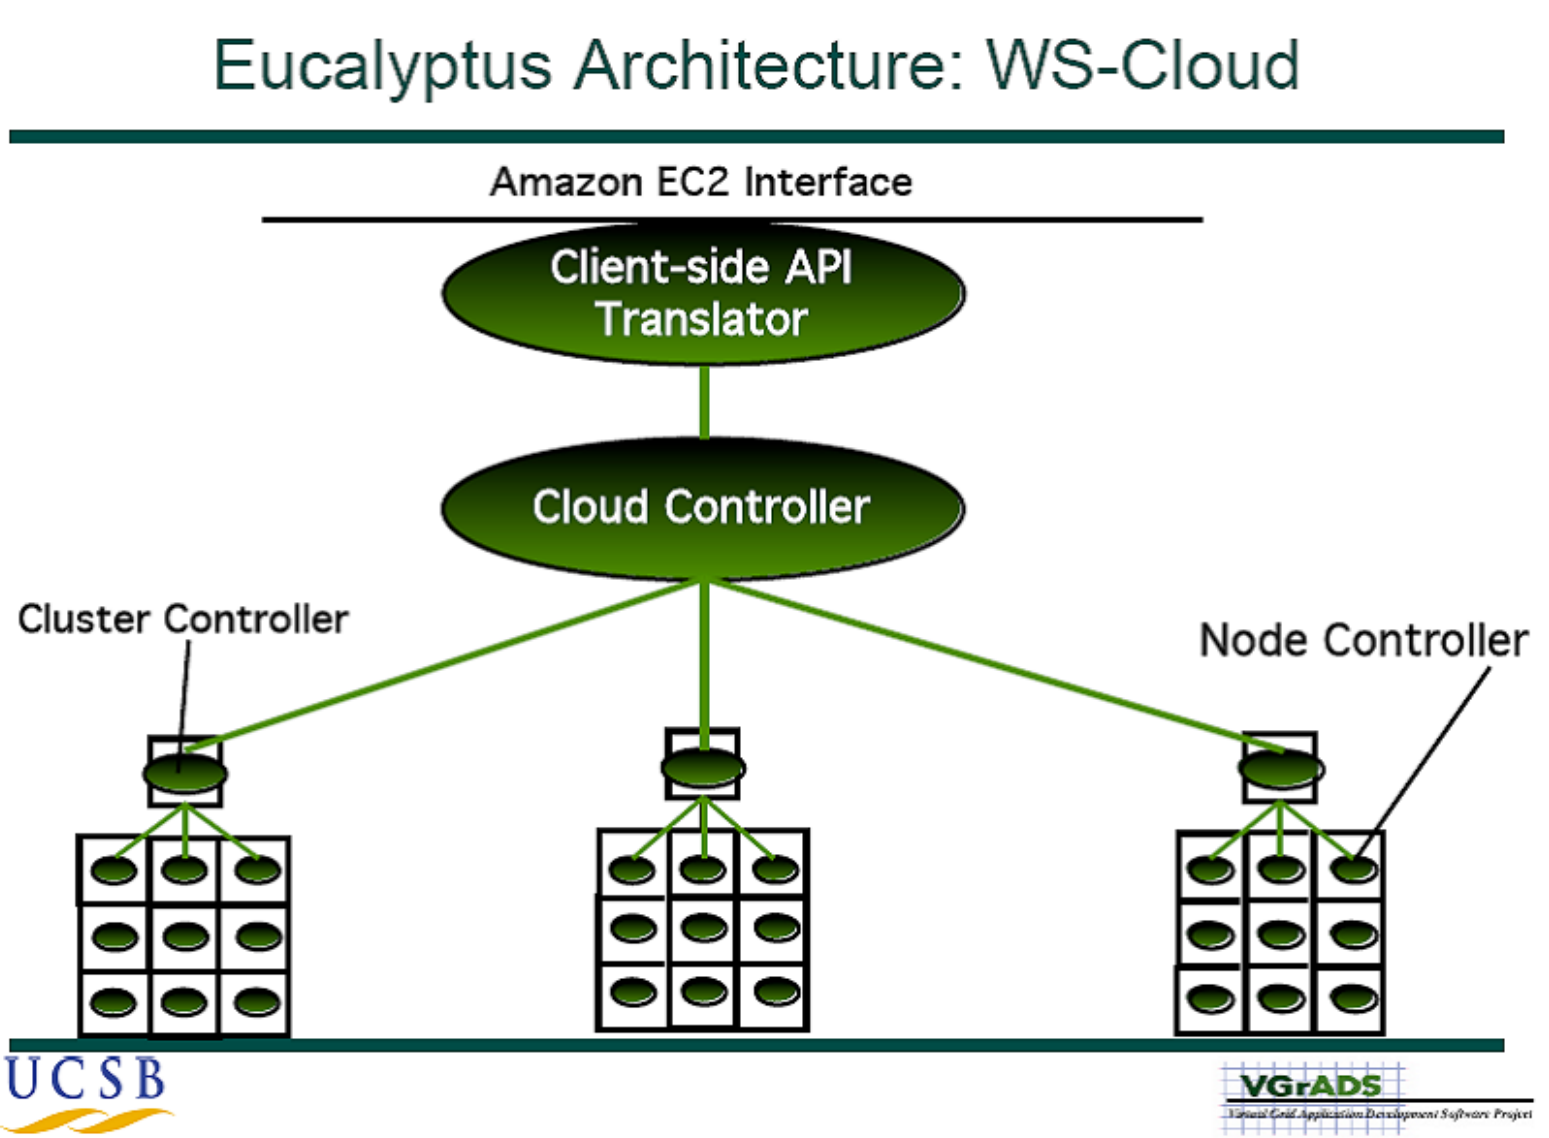

Cloud Performance

- Extensive performance study using HPC applications and benchmarks
- Two questions:
  - *Performance impact of virtualization*
  - *Performance impact of cloud infrastructure*
- Observations:
  - Random access disk is slower with Xen
  - CPU bound can be *faster* with Xen -> depends on configuration
  - Kernel version is far more important
  - No statistically detectable overhead
  - AWS small appears to throttle network bandwidth and (maybe) disk bandwidth -> *$0.10 / CPU hour*

Cloud Infrastructure

- Network operations center

- Physical Infrastructure

- Physical Security

- Cooling

Cloud Infrastructure ..contd

- Power infrastructure, Network Cabling, Fire safety

Advantages of Cloud Computing

- Lower computer costs:
  - You do not need a high-powered and high-priced computer to run cloud computing's web-based applications.
  - Since applications run in the cloud, not on the desktop PC, your desktop PC does not need the processing power or hard disk space demanded by traditional desktop software.
  - When you are using web-based applications, your PC can be less expensive, with a smaller hard disk, less memory, more efficient processor...
  - In fact, your PC in this scenario does not even need a CD or DVD drive, as no software programs have to be loaded and no document files need to be saved.
- Improved performance:
  - With few large programs hogging your computer's memory, you will see better performance from your PC.
  - Computers in a cloud computing system boot and run faster because they have fewer programs and processes loaded into memory...
- Reduced software costs:
  - Instead of purchasing expensive software applications, you can get most of what you need for free-ish!
  - That is right - most cloud computing applications today, such as the Google Docs suite, are totally free.
  - That is a lot better than paying $200+ for similar Microsoft Office software - which alone may be justification for switching to cloud applications.
- Instant software updates:
  - Another advantage to cloud computing is that you are no longer faced with choosing between obsolete software and high upgrade costs.
  - When the application is web-based, updates happen automatically - available the next time you log into the cloud.
  - When you access a web-based application, you get the latest version - without needing to pay for or download an upgrade.
- Improved document format compatibility.
  - You do not have to worry about the documents you create on your machine being compatible with other users' applications or operating systems.
  - Where Word 2007 documents cannot be opened on a computer running Word 2003, all documents can be read!
  - There are potentially no format incompatibilities when everyone is sharing documents and applications in the cloud.

Advantages of Cloud Computing

- Unlimited storage capacity:
  - Cloud computing offers virtually limitless storage.
  - Your computer's current 200 Gbyte hard drive is small compared to the hundreds of Pbytes available in the cloud.
  - Whatever you need to store, you can.
- Increased data reliability:
  - Unlike desktop computing, in which if a hard disk crashes and destroy all your valuable data, a computer crashing in the cloud should not affect the storage of your data.
  - That also means that if your personal computer crashes, all your data is still out there in the cloud, still accessible.
  - In a world where few individual desktop PC users back up their data on a regular basis, cloud computing is a data-safe computing platform!
- Universal document access:
  - That is not a problem with cloud computing, because you do not take your documents with you.
  - Instead, they stay in the cloud, and you can access them whenever you have a computer and an Internet connection.
  - All your documents are instantly available from wherever you are.
- Latest version availability:
  - Another document-related advantage of cloud computing is that when you edit a document at home, that edited version is what you see when you access the document at work.
  - The cloud always hosts the latest version of your documents; as long as you are connected, you are not in danger of having an outdated version.
- Easier group collaboration:
  - Sharing documents leads directly to better collaboration.
  - Many users do this as it is an important advantages of cloud computing - multiple users can collaborate easily on documents and projects.
  - Because the documents are hosted in the cloud, not on individual computers, all you need is an Internet connection, and you are collaborating.
- Device independence.
  - You are no longer tethered to a single computer or network.
  - Changes to computers, applications and documents follow you through the cloud.
  - Move to a portable device, and your applications and documents are still available.

Disadvantages of Cloud Computing

- Requires a constant Internet connection:
  - Cloud computing is impossible if you cannot connect to the Internet.
  - Since you use the Internet to connect to both your applications and documents, if you do not have an Internet connection you cannot access anything, even your own documents.
  - A dead Internet connection means no work and in areas where Internet connections are few or inherently unreliable, this could be a deal-breaker.
  - When you are offline, cloud computing simply does not work.
- Does not work well with low-speed connections:
  - Similarly, a low-speed Internet connection, such as that found with dial-up services, makes cloud computing painful at best and often impossible.
  - Web-based applications require a lot of bandwidth to download, as do large documents.
  - If you are labouring with a low-speed dial-up connection, it might take seemingly forever just to change from page to page in a document, let alone to launch a feature-rich cloud service.
  - In other words, cloud computing is not for the broadband-impaired!
- Can be slow:
  - Even with a fast connection, web-based applications can sometimes be slower than accessing a similar software program on your desktop PC.
  - Everything about the program, from the interface to the current document, has to be sent back and forth from your computer to the computers in the cloud.
  - If the cloud servers happen to be backed up at that moment, or if the Internet is having a slow day, you would not get the instantaneous access you might expect from desktop applications.

Disadvantages of Cloud Computing

- Features might be limited:
  - This situation is bound to change, but today many web-based applications simply are not as full-featured as their desktop-based applications.
  - For example, you can do a lot more with Microsoft PowerPoint than with Google Presentation's web-based offering.
  - The basics are similar, but the cloud application lacks many of PowerPoint's advanced features.
  - If you are a power user, you might not want to leap into cloud computing just yet.
- Stored data might not be secure:
  - With cloud computing, all your data is stored on the cloud.
  - The questions is How secure is the cloud?
  - Can unauthorised users gain access to your confidential data?
  - Cloud computing companies say that data is secure, but it is too early to be completely sure of that.
  - Only time will tell if your data is secure in the cloud.
- Stored data can be lost:
  - Theoretically, data stored in the cloud is safe, replicated across multiple machines.
  - But on the off chance that your data goes missing, you have no physical or local backup.
  - Put simply, relying on the cloud puts you at risk if the cloud lets you down.
- HPC Systems:
  - Not clear that you can run compute-intensive HPC applications that use MPI/OpenMP!
  - Scheduling is important with this type of application – as you want all the VM to be co-located to minimise communication latency!
- General Concerns:
  - Each cloud systems uses different protocols and different APIs… so it may not be possible to run applications between cloud based systems.
  - Amazon has created its own DB system (not SQL 92), and workflow system (many popular workflow systems out there) – so your normal applications will have to be adapted to execute on these platforms.

Clouds – open for innovation

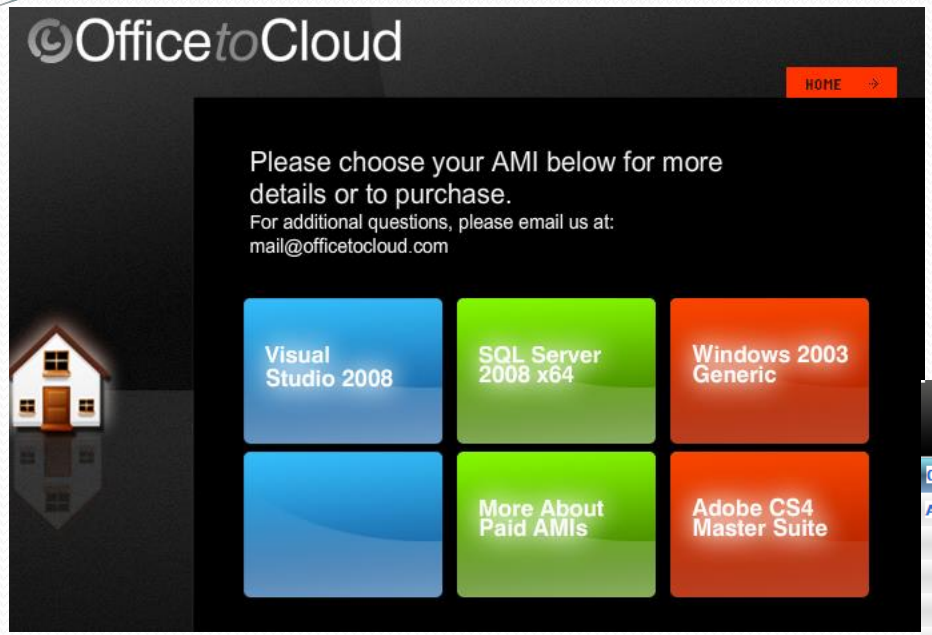

Cloud computing open issues

- Governance
  - Security, Privacy and control
  - SLA guarantees
  - Ownership and control
  - Compliance and auditing
    - Sarbanes and Oxley Act
- Reliability
  - Good servive provider with 99.999% availability
- Cloud independence – Vendor lockin?
  - Cloud provider goes out of business
- Data Security
- Cloud lockin and Loss of control
  - Plan for moving data along with Cloud provider
- Cost?
- Simplicity?
- Tools
- Controls on sensitive data?
  - Out of business
- Big and small
  - Scalability and cost outweigh reliability for small businesses
  - Big businesses may have a problem

Battle in the cloud

- Amazon Web Services
- Google App Engine
  - Free upto 500 MB,
    - Free for small scale applications?
    - Universities?
  - Pay when you scale
- GoGrid
- .. Some more Hosting companies
- Where is HP, IBM, Oracle(+sun) and Dell?

References - Further Reading

[1] *Introduction to Cloud computing White paper,* Sun Microsystems, June 2009.

[2] Lewis Cunningham, *Cloud Computing with Amazon and Oracle,* 2008.

[3] Michael Armbrust et al, *Above the Clouds: A Berkeley View of Cloud Computing,* Electrical Engineering and Computer Sciences University of California at Berkeley, February 2009.

[4] Jinesh Varia, *Architecting for the Cloud: Best Practices,* Amazon, May 2010.

[5] John W. Rittinghouse, James F. Ransome, *Cloud Computing Implementation, Management and Security,* CRC Press, 2010.

[6] Paul T. Jaeger, Jimmy Lin, Justin M. Grimes, *Cloud Computing and Information Policy: Computing in a Policy Cloud?,* Journal of Information Technology & Politics, May 2010.

[7] *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1,* Cloud Security Alliance, December 2009.

[8] *Security Best Practices,* Amazon, January 2010.

[9] Ian Foster, Yong Zhao, Ioan Raicu, Shiyong Lu, *Cloud Computing and Grid Computing 360-Degree Compared,* Grid Computing Environments Workshop, 2008.

[10] Rajkumar Buyya, Chee Shin Yeo, and Srikumar Venugopal, *Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities,* International Conference on High Performance Computing, 2008.

[11] http://www.salesforce.com/cloudcomputing/

[12] http://aws.amazon.com/ec2/

- http://blogs.zdnet.com/Hinchcliffe/?p=488&tag=btxcsim
- http://blogs.zdnet.com/Howlett/?p=558&tag=btxcsim
- http://blogs.zdnet.com/BTL/?p=9560&tag=btxcsim
- http://www.businessweek.com/technology/content/aug2008/tc20080082_445669_page_3.htm
- http://www.ibm.com/developerworks/websphere/techjournal/0904_amrhein/0904_amrhein.html
- http://cloudcomputing.sys-con.com/

Cloud Computing

Google Apps

Cloud computing – Google Apps

- Email, chat.
- **Google App Engine**

Google App Engine

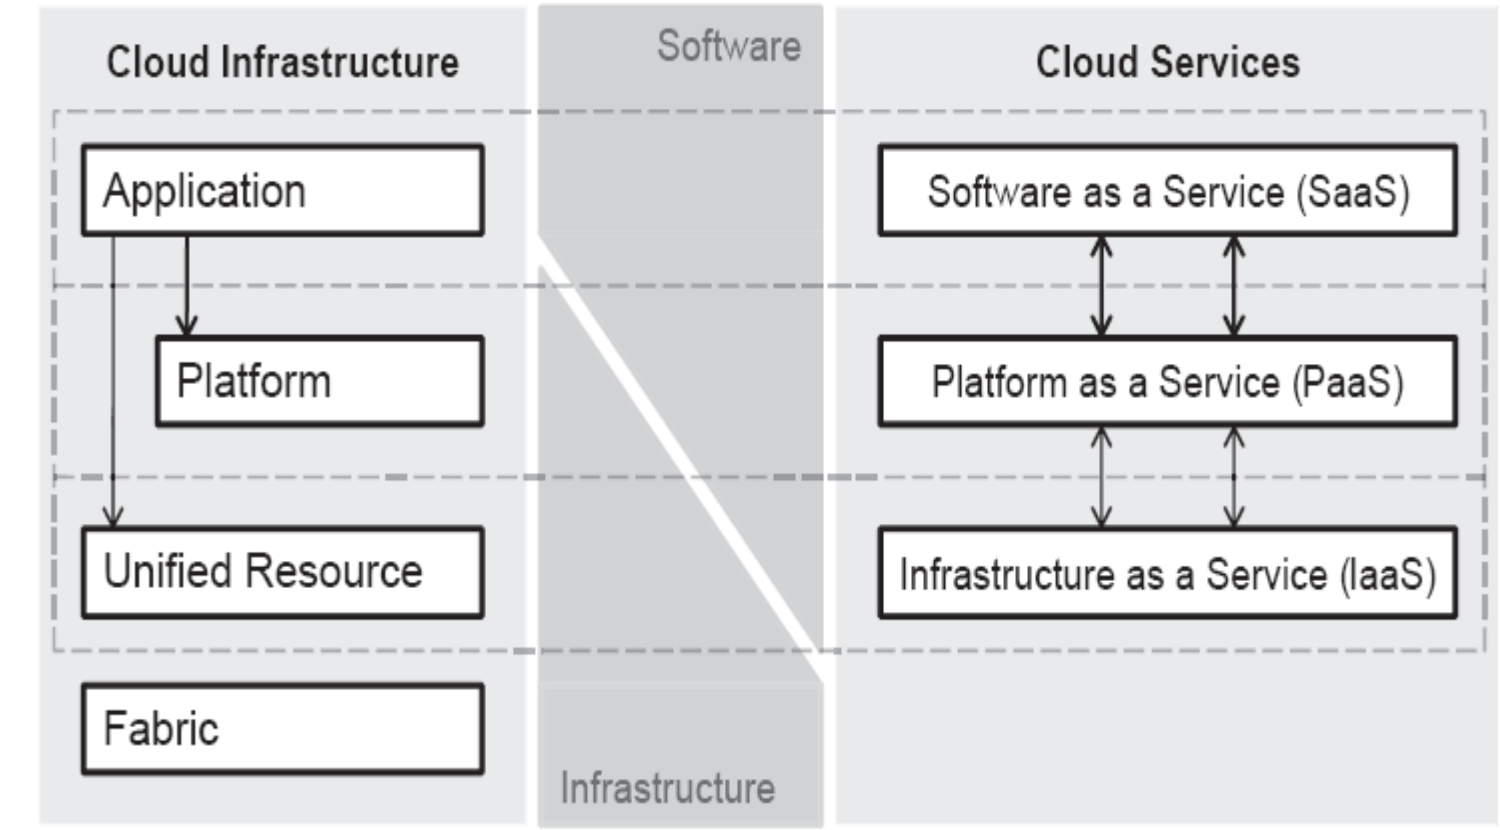

Google App Engine

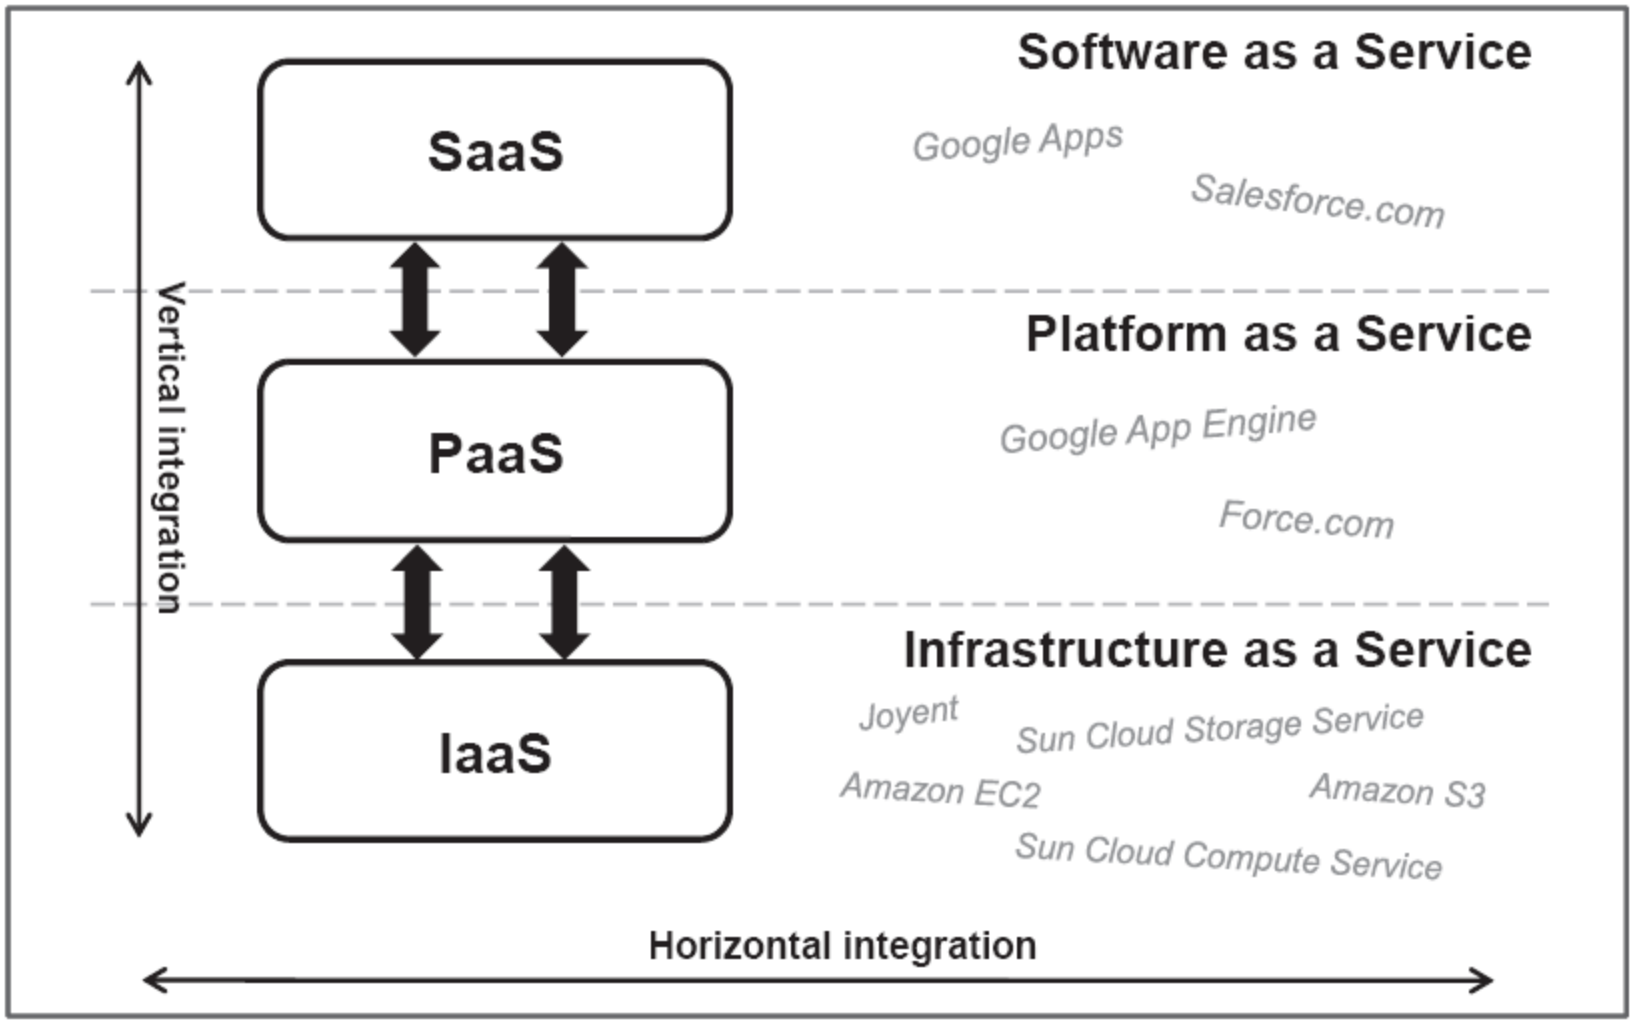

Google App Engine

- Google App Engine?
- Create application.

Google App Engine?

- Google App Engine enables you to build web applications on the same scalable systems that power Google applications. App Engine applications are easy to build, easy to maintain, and easy to scale as your traffic and data storage needs grow.

Easy to build → Write local, upload server

Easy to scale → how many user, how much data

easy to maintain → 10 year (data & application)

Google App Engine?

- Cost ➔ ?
  - Pay only for what you actually use.
  - Exceed the free quota of 500 MB of storage and around 5M pageviews per month.
  - Trial? ➔ 1GB store & 5M pageviews

Create application

- build an App Engine application using standard Java web technologies, such as servlets and JSP.
- create an App Engine Java project with Eclipse →use the Google Plugin for Eclipse for App Engine development. (Use SDK)
- use the App Engine datastore with the Java Data Objects (JDO) standard interface.
- upload your app to App Engine.