

Cloud Computing: *Making the right choices*

Kalpak Shah

Clogeny Technologies Pvt Ltd

IndicThreads.com Conference On
Upcoming Technology



2010: Cloud Computing

Pune, India

About Me

● Kalpak Shah

- Founder & CEO, Clogeny Technologies
- Passionate about economics and technology evolving through cloud computing
- Leading innovation on exciting projects across the cloud computing stack

@ kalpak@clogeny.com



Agenda

- Cloud Computing Overview
- Cloud Spectrum
 - Infrastructure as a Service
 - Platform as a Service
 - Software as a Service
- Ways to “Cloudify”
- Cloud Mobility - federated clouds, cloud bursting
- Private Cloud platforms
- Microsoft Azure - Covering the cloud stack
- Business Considerations & Examples
- Questions

Description
Comparison
Examples



Cloud Computing Overview



- Elasticity
- Programmability
- Pay as you go
- Fully automated



Why do we want to go to the cloud?

- No CAPEX, only OPEX
- Reduced IT administration & managed services costs
- Elasticity: Fast & Inherent scale-up and scale-down
- Variable & unpredictable loads
 - Ideal loads for the cloud
- Commoditization of IT
 - Only storage, only content delivery, only messaging
- Programmability - Automation
- Tools and ecosystem: Pay-as-you-go
- Reduced time-to-market: Focus on core competency
- Buzz... - Variety of options



Infrastructure-as-a-Service (IaaS)

- Pay-as-you-go Virtualized Resources - CPU, Storage, Network
- Infrastructure management services & tools
- Application cannot dynamically scale on-demand
- Local Server moved into the cloud - managing, patching, securing, monitoring is still a responsibility
- Extremely flexible
- Very little vendor lock-in
- Examples: Amazon EC2, Terremark vCloud, GoGrid Cloud, Rackspace Cloud



Considerations for IaaS - 1

- Support & HelpDesk
 - email, phone, 24x7
- Hardware & Performance
- Operating System Support
 - Distribution versions and range of support
 - Imaging of server?
 - Distribution and re-selling of images?
- Software availability and partner network
 - Middleware?
 - Database layer?
 - For example, Amazon EC2 has partnered with IBM, Symantec, Microsoft, Redhat, Oracle, etc. to provide a wide variety of pre-packaged software



Considerations for IaaS - 2

● APIs / Ecosystem

- Bindings in various languages - maintenance, backward compatibility
- Tools and ecosystem

● Licensing

- Big pain from our experience!
- Brings limitations to cloud migration
- Ex: IBM Websphere clustering not available on EC2
- Ex: Windows licenses cannot be migrated from local datacenter to cloud

● Other services

- Databases storage, backups, monitoring, auto-scaling, load-balancing, messaging, etc.



Let's compare IaaS Vendors

	OS Support	Disk Persistence	Choice of resources	Tools / Ecosystem	Hosting	Performance
Amazon AWS	Good but not flexible	Yes & No	Medium	Extensive - by a long way	Ideal	?
Rackspace	Limited	Yes	Limited	Good	Good	?
GoGrid	Limited but growing	Yes	Limited	Medium	Good	?
Terremark	Extensive & Flexible	Yes	Extensive	Very limited	Not good	?

- Performance needs to be tested for various applications across vendors.
- Lot of variation depending on CPU, network or disk bound application.



Platform-as-a-Service (PaaS)

- Scalable platform, hosted framework to develop and deploy applications
- Development language and platform decided by vendor
- No infrastructure management headaches
- No licensing concerns
- Scalability & availability are managed by the vendor
- Little flexibility - limited to choices provided by vendor
- Vendor lock-in concerns
- Examples: Google AppEngine, Microsoft Azure, Heroku



Software-as-a-Service (SaaS)

- Consumer facing - Business oriented
- Business layer is provided - extensions can be written
- Very similar to traditional hosting but with cloud models
- Consumed over the web
- No flexibility
- Almost complete vendor lock-in unless vendor follows standards for data extraction
- Examples:
 - Salesforce
 - Google Apps
 - Box.net
 - Zoho



Storage-as-a-Service

- Storage - attractive to store in the cloud
- Populated & consumed in the cloud
- Enterprise usage increasing
 - No management overhead, backups, disaster recovery
 - Amazon S3 provides all in \$15 cents/GB/month
 - Pricing to the tune of \$5 per month for unlimited storage
- High consumer penetration
- Examples
 - Mozy from EMC
 - Amazon S3
 - Rackspace CloudFiles
 - Carbonite
 - DropBox

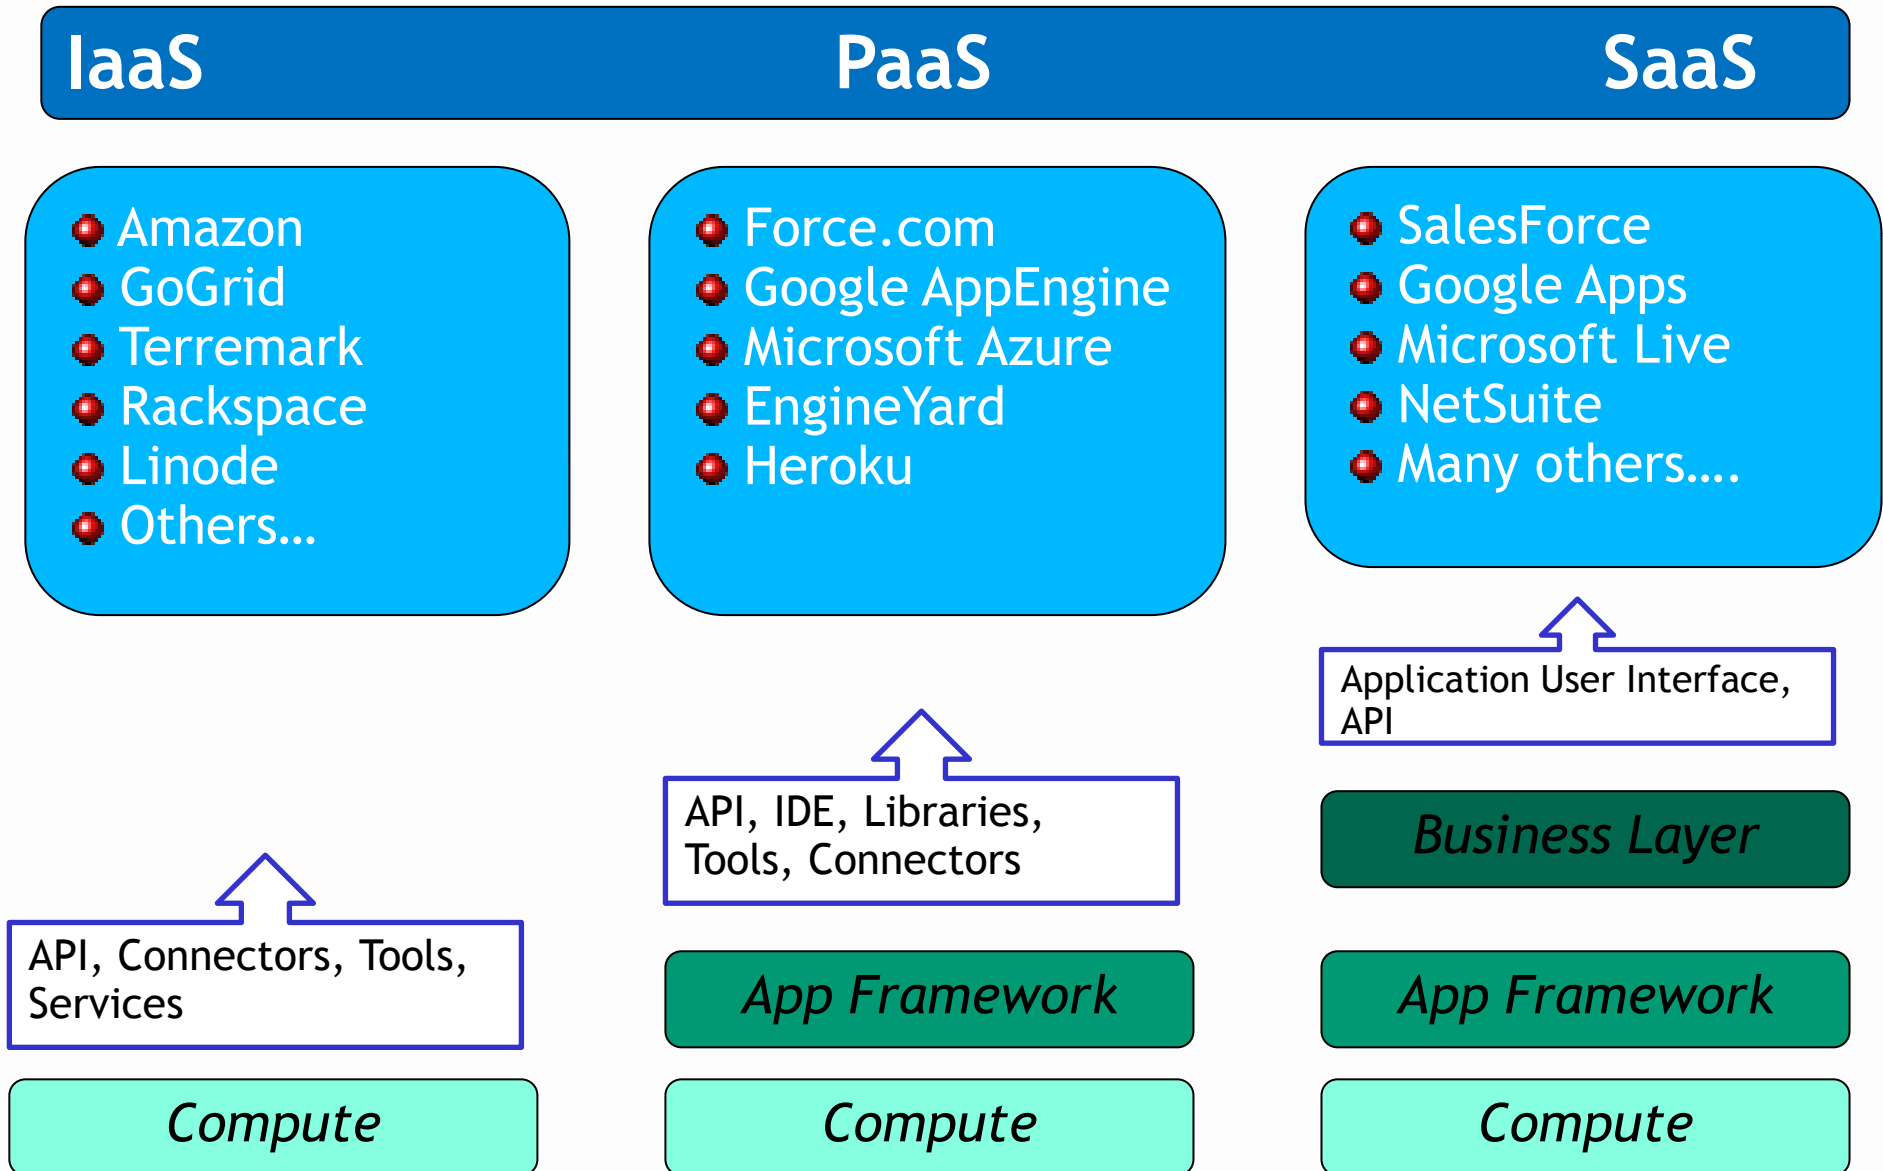


Comparing PaaS & SaaS Vendors


- Development language & available skill-sets
 - Google AppEngine supports Java & Python frameworks
 - Heroku is a Ruby on Rails Cloud platform
 - Microsoft Azure has a PaaS offering
- Cost of development
- Ease of deployment & maintenance
- Size of vendor & ecosystem maturity
- Tools, monitoring, connectors, adaptors - will my development time be reduced?
- Maturity of APIs & bindings available
- Does vendor allow private cloud integration?



The Cloud Spectrum



Cloud Layer - Visibility



Ways to Cloudify - 1

● Public Clouds

- Large datacenters with cloud capabilities
- Provide many services - CPU, storage, CDN, Databases, Messaging, Monitoring, Load balancing
- Large enterprises face security, latency and bandwidth hurdles

● Private Clouds

- Normal on-premise datacenters with cloud capabilities
 - Elasticity, Self-service, Pay-as-you-go, programmability
- Higher utilization levels and management flexibility
- Offerings from VMWare, IBM, Microsoft, Eucalyptus for Infrastructure
- Private setups of PaaS also available from Intalio, Wolf frameworks, etc.
- Cloud bursting is an option



Ways to Cloudify - 2

● Hybrid Clouds

- Best of both public and private clouds
- Disaster recovery and business continuity
- Regulation & compliance issues
- Seamless scalability
- VPN approach - Amazon Virtual Private Cloud
- IBM Private Cloud
 - Ex: Websphere cluster with some nodes in public cloud
- Windows Azure AppFabric
 - On-premise & cloud database remains in sync through use of this layer
 - Shared Windows framework gives unique advantage to Microsoft



Ways to Cloudify - 3

● SaaS

- CRM, accounting, taxation
- Any business or domain

● Virtual Desktops

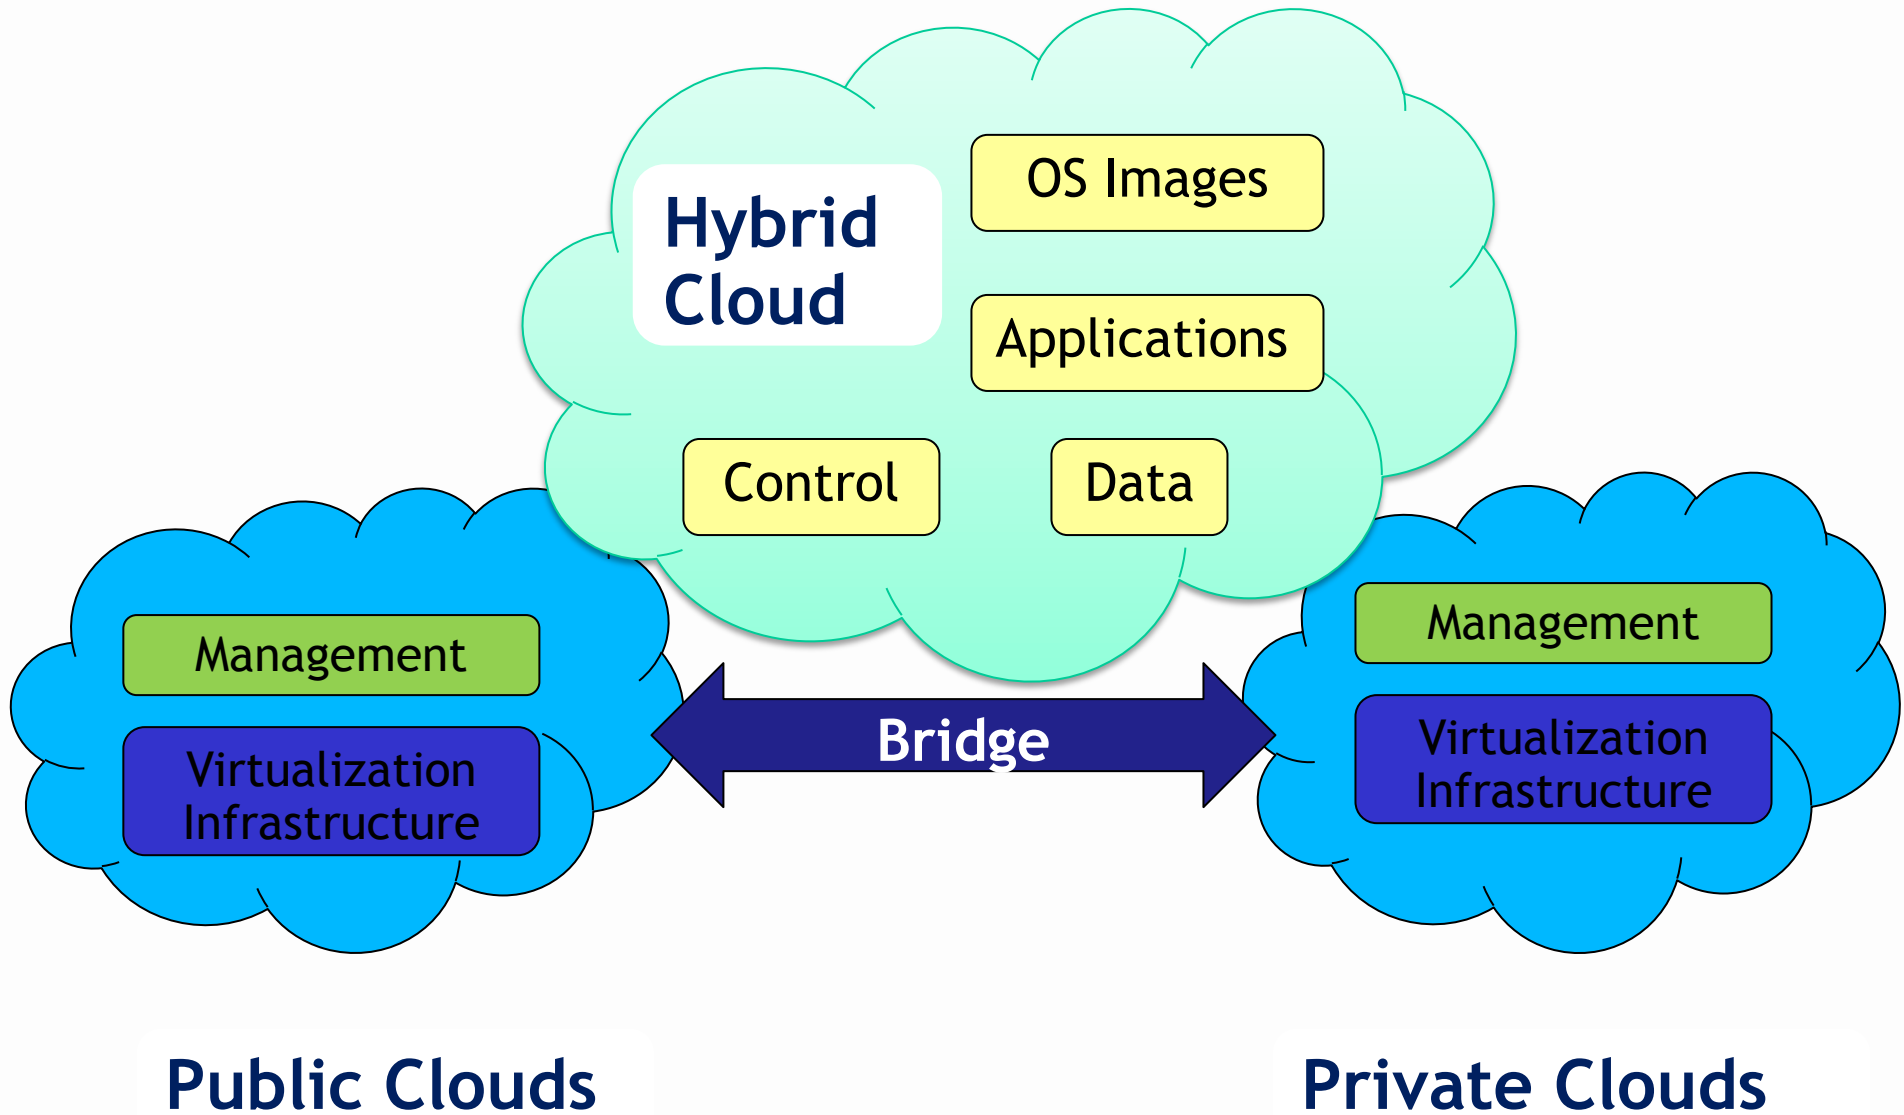
- Client virtualization & application streaming
- Trainings, virtualized desktops, centralized management
- Mixture of IaaS & SaaS

● Cloud Storage

- Backup
- Collaboration



Cloud Mobility - Federated, Bursting



Interesting Private Cloud Platforms - 1

● Eucalyptus

- Open-source IaaS cloud computing platform
- Compatible with Amazon's AWS enabling (almost) seamless movement with AWS
- Supports VMWare, Xen and KVM
- Extremely flexible and easy to administer

● VMWare vCloud

- *“Run, secure and manage applications in the private cloud or have them federated on-demand to partner-hosted public clouds with VMware vCloud”*
- Forging partnerships with public cloud providers like Terremark
- Expensive but worth it!



Interesting Private Cloud Platforms - 2

● Appistry

- *Application* platform for private, public and hybrid clouds
 - Can be hosted on private data center, EC2, Terremark, etc
- **CloudIQ Manager** - Application Deployment and Management Software
- **CloudIQ Engine** - Cloud Application Server for App Scalability and deployment
- **CloudIQ Storage** - Distributed Storage System for Data management across the cloud

● Other platforms:

- Rackspace - OpenStack Compute and OpenStack Storage
- CloudStack - formerly VMOps: Open source and enterprise editions



Microsoft Azure

❖ *Cloud Services Operating System*

- Development, service hosting and service management environment
- Integrates with existing on-premises environment - Windows Framework
- Provides services across the cloud stack: IaaS, PaaS and SaaS
- End-to-end services
- Private Cloud Integration
- SQL Azure, SharePoint, Dynamics CRM
- Visual Studio as development and testing environment
- Supports multiple development languages

Unique positioning in the market!



Cloud Databases

● Relational databases

- Traditional databases like SQL, MySQL, Oracle
- Should be used for quick porting of legacy applications
- Known loads & ability to serve them
- Scaling & performance needs to be managed
- Quicker development time & skills available
- Amazon's MySQL based *Relational Database Service (RDS)* is interesting

● NoSQL databases

- SimpleDB, Hadoop, BigTable
- Non-relational databases
- Supported & managed by the cloud vendor (in some cases)
- Inherent flexibility and scale



Business Considerations - 1

● Licensing

- Still not pay-as-you-go
- Per-CPU pricing, flat fees, limitations
- Can spoil cloud pricing and make it unattractive
- Expected to get better with time

● Service Level Agreements (SLAs)

- Around 99.9% for most vendors but read the fine print...
- Still evolving - sometimes does not meet enterprise demands
- Notification onus on customer, all commodities are not covered, limited liability



Business Considerations - 2

● Pricing

- Complex pricing - difficult to estimate without testing
- No APIs for automation
- Examples:
 - \$0.10 per 1 million I/O requests on Amazon EBS
 - \$0.01 / 10K storage transactions on Azure

● Geography & Government laws

● Location of datacenters

● Compliance

● Protocol Support





Thank You!!

Questions?



Example Cloud Migrations

- Porting existing stand-alone application to cloud platform
 - Reach out to global customers without high sales cost
 - Save normal hosting costs since load would be arbitrary
 - Extremely CPU intensive
 - IaaS or PaaS?
- Analytics application to be written on cloud platform
 - Application being written from scratch
 - Basic load can be estimated - scaling amount and times cannot be estimated
 - Cannot use virtualization locally as it affects performance
 - Client ready for CAPEX to save costs over a period of time
 - Application architecture designed as per plan to utilize public clouds

