

Building an IaaS cloud

Zaeem Arshad

16th July 2012

SANOG 20, Karachi

SANOG

CYBERNET

Who Am I?

- Technical Architect of RapidCompute
- Developer, IP enthusiast, Sysadmin
- You can reach me at
 - zaeem.arshad@rapidcompute.com

What is it?

- Maybe I'm an idiot, but I have no idea what anyone is talking about. What is it? It's complete gibberish. It's insane. When is this idiocy going to stop? We'll make cloud computing announcements. I'm not going to fight this thing. But I don't understand what we would do differently in the light of cloud.

– Larry Ellison, CEO Oracle Corporation

Ok Really what is it?

- Cloud computing relies on **sharing of resources** to achieve coherence and **economies of scale** similar to a utility (like the electricity grid) over a **network** (typically the Internet). At the foundation of cloud computing is the broader concept of **converged infrastructure** and **shared services**

– Wikipedia

Essential Characteristics

- Resource pooling
- Broad network access
- Rapid elasticity
- Measured Service
- On Demand Self Service

– NIST SP800-145

Service Models

- Software as a Service (SaaS)
 - Gmail, Salesforce.com
- Platform as a Service (PaaS)
 - Microsoft Azure, Google AppEngine, Engineyard
- Infrastructure as a Service (IaaS)
 - Amazon EC2, Rackspace, GCE, RapidCompute

Build a cloud – short version

1. Select or develop a provisioning model
2. Buy servers
3. Install hypervisors
4. Buy & connect storage
5. Figure out the networking
6. Some Magic
7. Profit!!

Infrastructure as a Service

- Virtualization is the key!
- But virtualization alone is not enough
- Consider
 - Service provisioning
 - Compute resources
 - Storage
 - Network
 - Security
 - Scalability

Service provisioning

- Plethora of open source options
 - Apache CloudStack
 - Openstack
 - Eucalyptus
 - OpenNebula
- Commercial options
 - VMWare
 - Microsoft/Dell/HP/IBM
 - Too many to list
- Write your own service provisioning layer

Compute Resources

- Hypervisor
 - Citrix Xenserver
 - Opensource Xen and XCP
 - RedHat KVM
 - VMWare
 - Oracle VMM
 - Microsoft Hyper-V
- Maturity
- Feature richness
- Stability

Compute Resources

- Servers and CPUs
- Biggest clouds built on commodity hardware
- Commodity does not mean low-quality
- Define your compute needs
- CPU Frequency is NOT the measure of performance
- Investigate low power cpus if your workload allows

Storage

- SAN, NAS, DAS, distributed storage systems
- SANs deliver the highest reliability
- But they are complex to manage
- IO contention still applies to virtualized infrastructure
- Spindles vs SSD – choose your RAID wisely
- Openstack Swift, GFS, Sheepdog, ZFS, GlusterFS
- EMC, Hitachi, IBM

Network

- Cloud is designed for failure
- That means when virtual machines are migrated, networks are migrated
- Plan L2, L3 and L4 network for high availability
- OpenFlow – Nexus vSwitch – Open vSwitch
- SDN is the future
- Architect your access and distribution for multiple gigs of throughput

Security

- Encryption
- L2 and L3 security
- Use VLANs to segment traffic
- Use firewalls to filter traffic
- Will the security work when we recover from failure?
- ISO 27001:2005

Scalability

- Loose-coupling is the key
- Remember CAP theorem
- ACID vs BASE
- Reduce MTTR
- Automation is the key – 80/20 rule
- Backup and recovery impose limits on physical scalability

Questions 😊

SANDOG

CYBERNET