



SC12
Salt Lake City, Utah

Conference Dates:
November 10-16, 2012

Exhibition Dates:
November 12-15, 2012



Association for
Computing Machinery

Infrastructure Clouds for Science and Education: Infrastructure Cloud Offerings

Pierre Riteau
University of Chicago



IaaS Clouds

- Multitude of IaaS providers available
- Most are following the same model
 - On-demand access
 - Pay as you go
- What makes them different?
- Why choose one rather than the other?

Comparison of IaaS Clouds

- Compute Resources
- Virtualization Technology
- Storage
- Network
- Geographic Distribution
- Pricing

COMMERCIAL CLOUDS

12/11/2012



www.nimbusproject.org

Commercial Providers

- Amazon EC2
- Google Compute Engine
- Rackspace Cloud
- Windows Azure
- ... and many others

Amazon EC2

- Leader of the IaaS clouds
- Launched in August 2006
- Part of Amazon Web Services
- Started with:
 - A single instance type
 - No static IPs
 - No availability zones
- A lot of evolution in six years



Amazon EC2: Compute Resources

- Eight families of instance types:
 - First generation standard instances
 - Second generation standard instances
 - Micro instances (cheapest at \$0.02 per hour)
 - High-memory instances (up to 68.4 GB)
 - High-CPU instances (high CPU to memory ratio)
 - Cluster compute instances (for HPC, in TOP500)
 - Cluster GPU instances (with NVIDIA Tesla GPUs)
 - High I/O instances (backed by SSDs)
- CPU power measured in EC2 Compute Unit

Amazon EC2: Virtualization

- Xen hypervisor (mainly **paravirtualization**)
- Hardware-assisted virtualization (HVM) for:
 - Cluster Compute Instances
 - High I/O instances
 - Second generation (m3) instances
 - Windows instances
- Operating systems (non-exhaustive):
 - Linux
 - OpenSolaris / Solaris Express Community Edition
 - Windows Server 2003 / 2008
 - FreeBSD
 - NetBSD

Amazon EC2: Storage

- Ephemeral: instance storage
 - From 160GB to 3370GB on local disks
 - 2 SSD-based volumes on High I/O instances
- Persistent: Elastic Block Store (EBS)
 - From 1GB to 1TB
 - Boot volumes
 - Replication in same availability zone
 - Snapshots to S3
- AWS Import/Export
 - Move large data in/out of AWS by shipping HDDs

Amazon EC2: Network

- Routed network (L3) between instances
- Attach Elastic IPs to instances
- For Cluster Compute Instances:
 - 10Gb Ethernet
 - Placement groups for low latency
- Amazon Virtual Private Cloud
 - Private network on Amazon's infrastructure
 - Connect to your own network with VPN
- AWS Direct Connect (for large volume users)
 - Direct connection with Amazon's infrastructure
 - Reduced network data charges

Amazon EC2: Geographic Distribution

- North America
 - North Virginia
 - Oregon
 - Northern California
- Europe (Ireland)
- Asia Pacific
 - Singapore
 - Tokyo
 - Sydney (**NEW TODAY!**)
- South America (Sao Paulo)

Amazon EC2: Pricing

- Compute per-instance hour
 - Partial hours are fully billed
- EBS volume & requests
- Data-in free of charge
- Charged for:
 - Data between availability zones
 - Data between EC2 regions
 - Data out

Amazon EC2: More Pricing

- Reserved instances
 - 1 or 3 years commitment
 - Up to 71% saving over on-demand prices
 - Marketplace to sell remainder of reserved instances
- Spot instances
 - Variable price over time
 - Can shutdown instances at any time

Google Compute Engine (GCE)

- Announced at Google I/O in June 2012
- Currently in limited preview
- Runs on Google's infrastructure next to Search, Gmail, Ads...
- Integrates with other Google services
- Targets large compute jobs rather than web applications

GCE: Compute Resources

- Designed for high-performance clusters
- From 1 to 8 CPUs per instance
- Compute power
 - Measured in GCEUs
 - One unit at least as powerful as one ECU
 - Advertises 50% more compute power/\$ than other leading clouds

GCE: Virtualization

- KVM hypervisor
- Platform built for consistent performance
 - Based on cgroups?
- Locked down kernel
- Ubuntu Precise and CentOS 6 images

GCE: Storage

- Local disks
 - Encrypted data
 - Boot from local disk
 - Up to 3.5TB with 8-CPU instances
- Persistent storage
 - Encrypted data
 - Read-only mounts for several instances
 - Less than 3% performance variance
 - Three times more bandwidth than local disk for large read/writes

GCE: Network

- Routed network (L3) between instances
- One VPN per project spanning all VMs
- Can attach IP addresses to VMs

GCE: Geographic Distribution

- Only in USA datacenters for now
 - 3 datacenters on the east coast
- More global presence in the future

GCE: Pricing

- Charging for:
 - Compute time (hourly)
 - Storage
 - Network usage
- Same per-GCEU pricing for all instances
- Will have SLA for enterprise customers

Rackspace

- Rackspace **Cloud Servers**
- Used to rely on Slicehost technology
- *Next generation* infrastructure based on **OpenStack** since 2012
 - Rackspace is one of the founder of the OpenStack project

Rackspace: Compute Resources

- Many instance types:
 - From 1 to 8 CPUs
 - From 512MB to 30 GB memory
 - From 20GB to 1.2 TB local disk
 - From 20 Mbps to 300 Mbps public network
 - From 40 Mbps to 600 Mbps internal network

Rackspace: Virtualization

- Xen and XenServer hypervisors
- Operating systems:
 - Linux
 - FreeBSD
 - Windows
- File injection feature

Rackspace: Storage

- Local storage with RAID 10
 - Persists host failures
- Persistent storage: **Cloud Block Storage**
 - Built for consistent and reliable performance
 - Choice of SSDs or standard disks
 - Free I/O
 - Snapshots to Cloud Files
 - Can only be used from one instance
 - Based on iSCSI

Rackspace: Network

- Bandwidth caps for outbound public and internal networks
- No cap for input bandwidth
- Soon: Rackspace Cloud Networks
 - Layer 2 virtual networks

Rackspace: Geographic Distribution

- USA
 - Texas
 - Chicago
- UK
 - London

Rackspace: Pricing

- Compute-hours
 - From 2.2c to \$1.2 per hour
 - Plus pricing for Windows, SQL Server, etc.
- Bandwidth: 18c per GB out

Windows Azure

- Platform-as-a-Service since 2010
- Released new features in June 2012
 - With virtual machines

Azure: Compute Resources

- Five instance types:
 - From 1 shared CPU to 8 CPU cores
 - From 768 MB to 14 GB of memory

Azure: Virtualization

- Custom hypervisor (not Hyper-V)
- Disk images compatible with VHD format
 - Upload your own VHD images to Azure
- Operating systems:
 - Windows
 - Linux (CentOS, Ubuntu, openSUSE)

Azure: Storage

- Persistent storage:
 - Mount Azure blobs as disks
 - Can be formatted as NTFS for Windows or another FS for Linux
- Continuous storage geo-replication (optional)
 - Replicated within 15 minutes
 - Remote data center at least 500 miles away
 - Stays in same geopolitical region (USA/EU)

Azure: Network

- Virtual Private Networking
 - Connect on-premise datacenter network with the Windows Azure network

Azure: Geographic Distribution

- North America
 - North-central US - Chicago, IL
 - South-central US - San Antonio, TX
 - West US - California
 - East US - Virginia
- Asia
 - East Asia - Hong Kong, China
 - South East Asia - Singapore
- Europe
 - West Europe - Amsterdam, Netherlands
 - North Europe - Dublin, Ireland

Azure: Pricing

- Currently in preview mode
- Price will go up with general availability
- Hourly GA prices:
 - Non-windows from \$0.02 to \$0.68
 - Windows from \$0.02 to \$0.92
- Discounts for 6 or 12-month commitment
- Free data-in
- Charges for data-out

SCIENTIFIC CLOUDS



Scientific Clouds

- FutureGrid clouds
- Purdue Wispy
- Red Cloud
- SDSC Cloud Storage
- Magellan

FutureGrid

- Several clouds technologies available:
 - Nimbus
 - Eucalyptus
 - OpenStack

FutureGrid: Nimbus

- APIs: Nimbus WSRF / EC2 / S3 (Cumulus)
- Flexible memory and CPU allocation (WSRF)
- Hotel (University of Chicago)
 - 41 nodes, 328 cores, Xen
- Foxtrot (University of Florida)
 - 26 nodes, 208 cores, Xen
- Sierra (SDSC)
 - 18 nodes, 144 cores, Xen
- Alamo (TACC)
 - 15 nodes, 120 cores, KVM

FutureGrid: Eucalyptus

- Running Eucalyptus 3.1
- Accessible with euca2ools / EC2 API
- In Indiana and UCSD
- Uses the Xen hypervisor
- Five instance types:
 - India (Indiana)
 - From 1 to 4 CPUs
 - From 512 MB to 9GB of memory
 - Sierra (UCSD)
 - From 1 to 4 CPUs
 - From 256MB to 2GB of memory

FutureGrid: Openstack

- 20-node cluster in Indiana
- Running the Essex release
- EC2 API
- Accessible with euca2ools
- Five instance types
 - From 1 to 8 CPUs
 - From 512 MB to 16384 MB of memory
 - Up to 160 GB of ephemeral storage
- Uses the Xen hypervisor

Purdue Wispy

- Nimbus cloud
- Supports the EC2 and S3 APIs
- 8 VMMs each with:
 - 16-core Intel Xeon E5-2670 processors
 - 32GB of memory
 - Gigabit Ethernet connection.
- Uses the KVM hypervisor

Red Cloud (Cornell University)

- Eucalyptus 2.0.3 cloud (EC2 / AWS API)
- Subscription-based (\$500 for Cornell / \$750 other academic)
 - One core for one year
 - 50 GB of storage included (extra charge for additional storage)
 - **No network charges**
- SaaS with MATLAB and NVIDIA GPUs
- EBS volumes of up to 1TB in 1GB increments
- Infrastructure:
 - 8 servers (96 cores total), 384GB RAM, 10 Gb Ethernet
 - 7.5TB storage accessible via 10Gb iSCSI.
- Five instance types:
 - From 1 to 12 cores, 4GB to 48GB of memory, 20GB to 1TB disk
- Visit them at **SC12 booth 2030**

SDSC Cloud Storage

- Object-based storage system
- Based on OpenStack Swift
- Rackspace and S3 APIs
- Continuous automatic data verification
 - At least two on-site replicas
 - Off-site replication
- Objects accessible via the web
- Pricing:
 - \$32.50/TB/month for UCSD/UC
 - \$56.90/TB/month for academic and non-profit
 - \$65.00/TB/ month for industry partners
 - **No network charges**
- Advertises peak transfer rates of up to 8GB/sec
- Also High Performance NFS and CIFS Storage

Magellan

- Project by DoE
- Goal: allow scientists to experiment with cloud computing
- Two sites:
 - Argonne National Laboratory
 - Lawrence Berkeley National Laboratory
- Experimented with Eucalyptus and OpenStack
- Infrastructure: 750 nodes
 - 500 compute nodes
 - 200 storage nodes
 - Big memory (1-terabyte) nodes
 - 12 management nodes
- Transitioning from testbed to production system

Conclusion

- Many commercial providers competing with Amazon EC2
- OpenStack project backed by many organizations
- Community clouds for academic projects
 - Some free of charge!