HPC USER FORUM
Cloud PANEL

September 2010
Seattle, WA

Panel Participants

- **Moderator**
  - Sharan Kalwani, KAUST
- **Participants**
  - Tom **C**oull, Penguin Computing
  - Bob **G**raybill, Nimbis Services
  - Richard **K**aufmann, HP
  - Josh **S**imons, VMware
  - Deepak **S**ingh, Amazon
  - Christian **T**anasescu, SGI
  - Ryan **W**aite, Microsoft
  - Jeff Broughton, Lawrence Berkeley National Labs

Panel Format

- **3 Questions**
  - Provided ahead of time
- **3 minutes per question for each participant response**
- **Follow-up and Audience after each participant has a chance to share**

Q1. Security in the HPC cloud…..

- Please respond with the audience, briefly – the issues surrounding Security concerns and how is this being addressed?

Q1. Security in the HPC cloud……..

- Security is not “One-Size Fits All”
- Standard security measures are sufficient for many users of Penguin Computing on Demand (POD):
  - Login nodes (virtual/dedicated) isolating workspace
  - ssh, IP matching through the firewall, account permissions
  - Web services with authentication
  - Server-to-server disk transfer
  - Commercial data center security
  - One user per server at a time
  - “Acceptable Use Policy”
- Additional security measures available:
  - Encrypted data transmission or shipment
  - Prolog scripts to reboot compute nodes/clean cache
  - Login nodes with dedicated/user controlled storage

- Tom Coull

Q1. Security in the HPC cloud……..

- *There is significant investment in Cloud Security by the private and public sectors.*

- *Cloud may take many forms*
  - *Public*
  - *Community*
  - *Private*

- *Security concerns may depend level of use*
  - *IaaS, PaaS, SaaS and WaaS*

- *Clouds may actually be more secure due to the ability to centralized monitoring ability*

- *Secure cloud examples exist today*

- Bob Graybill

Q1. Security in the HPC cloud……..

- *Here today as a "hardware head," but…*
  - *A spectrum of offerings is needed where customers can pay for security a la carte*
    - *Example: co-tenancy increases occupancy %age → less idle time → lower costs*
      - *Isolation, therefore, must offset this efficiency by higher prices or committed usage*
    - *Network isolation also costs money.*
    - *PCI Compliance, etc.*

Table 1: Quick Preview of Top 10 Obstacles to and Opportunities for Growth of Cloud Computing.

	Obstacle	Opportunity
1	Availability of Service	Use Multiple Cloud Providers; Use Elasticity to Prevent DDOS
2	Data Lock-In	Standardize APIs; Compatible SW to enable Surge Computing
3	Data Confidentiality and Auditability	Deploy Encryption, VLANs, Firewalls; Geographical Data Storage
4	Data Transfer Bottlenecks	FedExing Disks; Data Backup/Archival; Higher BW Switches
5	Performance Unpredictability	Improved VM Support; Flash Memory; Gang Schedule VMs
6	Scalable Storage	Invent Scalable Store
7	Bugs in Large Distributed Systems	Invent Debugger that relies on Distributed VMs
8	Scaling Quickly	Invent Auto-Scaler that relies on ML; Snapshots for Conservation
9	Reputation Fate Sharing	Offer reputation-guarding services like those for email
10	Software Licensing	Pay-for-use licenses; Bulk use sales

http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf, **Above The Clouds: A Berkeley View of Cloud Computing**

- Rich Kaufmann

Q1. Security in the HPC cloud……..

- With virtualization, a shift from hardware-based perimeter protection to software-based defense in depth at multiple levels:
  - VM
  - Application
  - Virtual Data Center

- Josh Simons

Q1. Security in the HPC cloud……..

- ***For the technical and engineering customers looking to run simulations in the cloud they want assurance that their competition who might also be using the same cloud service can't see their data. SGI is addressing this issue with each customer though the following levels of security:***

- **Level 1- Secure data transfer over the internet.**
  - **Customer use VPN and/or SSH to transfer their files and access Cyclone**
  - **Cyclone is located behind SGI secure firewall on an isolated network.**
- **Level 2- Gated Access within network.**

  SGI Cyclone provides customers with gated access to compute resources via their own login/management node. This reduces what a customer can see within Cyclone to their own data and machine.
- **Level 3- Password Security protocols are strictly enforced**
- **Level 4- File System Security**

  Each customer has their own data storage that is inaccessible from anyone else in Cyclone
- **Level 5- Encryption of data within their file system (additional charge).**
- **Level 6- Dedicated access to full compute resources.**

  System taken off the Cyclone internal network. Customer allowed to be onsite and run jobs.

- Christian Tanasescu

Q1. Security in the HPC cloud……..

- **Cloud security should integrate with existing infrastructure**
  - Authorization/Authentication
  - Compliance auditing
  - Certifications
  - Protection against traditional attacks like elevation of privilege, etc.
- **Services model allows for efficient updates**
  - OS, middleware, and applications patched immediately
- **Example of secure cloud app: Exchange Online**
  - 9 layers of security: IDS, system level firewalls, application level authentication, and more
  - N+1 redundancy
  - 99.9% uptime guarantee
  - 24x7x365 support

- Ryan Waite

Q1. Security in the HPC cloud……..

- *User as sysadmin (root)*
  - *Limited experience => security deficiencies*
  - *Third-party exploits and compromises*
  - *Unintended control over infrastructure*
  - *User authentication outside normal controls*
- *You can leave the cloud outside the firewall*
  - *But, this ignores the issue*
- *Approaches*
  - *Pre-built kernels and images*
  - *Load-time inspection and modification of image*
  - *Containment, Detection and Auditing*

- Jeff Broughton

Q2. Bandwidth to the Service- Who pays?.....

- Typical HPC involves sending massive amounts of data to the service and receiving perhaps X^n times the same back! Network service providers charge an exorbitant premium. So who pays and tell us why it will NOT destroy business case?

Q2. Bandwidth to the Service- Who pays?...

- Costs per GB/month ~ $0.10 to $0.20
- Rule of thumb

  < 250GB, use the Internet (8 hours typical); otherwise “FedEx” (UPS, etc.)

  Check your actual, local bandwidth first!
- For POD
  - FedEx ~ 2TB drive caddy = $20 per shipment (including helping hands, caddy is free) ~ $0.01 per GB
  - FedEx ~ n x 2TB drives (server-to-server) faster/cheaper but requires purchase of drives
- Multiple TB transfers per month per user common

- Tom Coull

Q2. Bandwidth to the Service- Who pays?.....

Nimbis Services

- ***Not all "technical computing" use cases require massive movement of data --- user application dependant***
  - *Academic, OEMs and SME*
- ***Massive data movement challenges may change your work flow***
  - *Remote post processing*
  - *Remote data storage*
  - *Pixel vs data movement*
- ***Last mile bandwidth challenges do exist for SMEs***

- Bob Graybill

Q2. Bandwidth to the Service- Who pays?.....

- Raw internet bandwidth costs money. $10 → ‹ $2 per Mbps, 95th %ile.
  - Distance costs; co-location doesn't.
  - Providers of IaaS services have to recover their bandwidth costs.
- Bandwidth doesn't want to be free!
  - Users want it to be, though.

- Rich Kaufmann

Q2. Bandwidth to the Service- Who pays?.....

- What's with the 95th %ile stuff?
  - Bandwidth providers charge using a byzantine algorithm. Sort 5-10 second quanta of a month by the max(in, out b/w). Pick 95th %ile. Multiply max by the posted rate, e.g. $5.
    - Blended charge of average and peak use
- Heuristic: to get per-bit charge, multiply 95th %ile cost by 2.5.

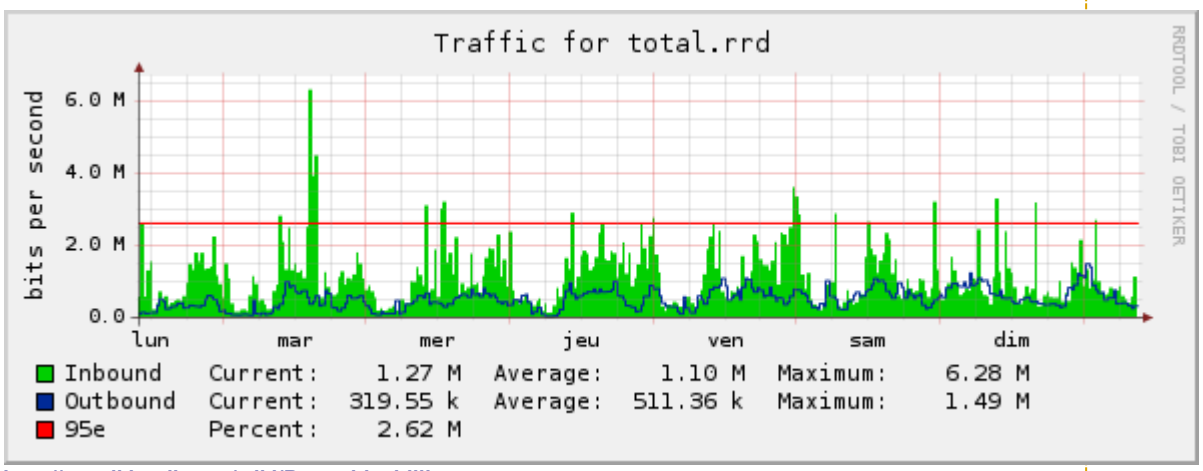

http://en.wikipedia.org/wiki/Burstable_billing

- Rich Kaufmann

Q2. Bandwidth to the Service- Who pays?.....

- **Amazon charges for bandwidth – not clear this must be part of all cloud business models – market will decide**
  - AWS Import/Export -- $80/device + $2.49/DLH
  - …or about $120 to load 2TB
- **Application data requirements are very diverse – some will map well to remote compute environments, some will not**
- **Our national supercomputing centers have been dealing with this issue for decades**

- Josh Simons

Q2. Bandwidth to the Service- Who pays?....

- SGI doesn't currently charge for uploading or downloading customer files to and from Cyclone
- SGI does offer "sneaker net" for extremely large or sensitive files via a secure FedEx service
- SGI business model is simple and transparent, charging for applications and infrastructure
- Extra charges for out of scope requests, special service requests beyond standard SLA

- Christian Tanasescu

Q2. Bandwidth to the Service- Who pays?.....

- **Broad and provocative question but not all HPC follows the pattern described the question**
  - Eclipse: ~2GB input, ~10GB output
  - LS-DYNA (model dependent): ~5GB input, ~10GB output
  - Risk Modeling (finance): <100 GB input (cached), <10GB output
  - Digital Content Creation: 10-100 TB
- **Also presumes cloud *equals* compute**
  - Big data processing requires big I/O with low CPU
  - Some big data is too big or too sensitive to move to the cloud
- **Three models for thinking about using data**
  - Large data ingress and egress (the subject of this question)
  - Data ingress with remote visualization of results stored in the cloud
  - Small data ingress for integration with public cloud data

- Ryan Waite

A2. Bandwidth to the Service- Who pays?.....

- *Compute and storage are (almost) free.*
  - *I/O (bandwidth) is expensive.*

- ***Clouds need capability to store, process and analyze large datasets entirely in the cloud***
  - *Minimize need to import or export raw data*
  - *Promote data exchange amongst users including community datasets.*

- ***Cloud providers must peer directly with the large research networks (e.g. DOE's ESnet)***

- ***Does this sound like a supercomputing center?***

- Jeff Broughton

Q3. Technology Issues or how to satisfy everyone?….

- ***HPC solutions often involve several special hardware platforms and techniques (Cray X, Blue Gene, novel ways of connecting gear, GPCPUs, etc.)? Yet the HPC cloud seems to be forcing everyone to the plain vanilla x86 model.***
- ***HPC ⇒ "Performance, performance, performance"…is key. So how do you convince folks on this dilemma? (and potential performance loss?)***

Q3. Technology Issues or x86 everyone?….

- POD has:
  - Tiered x86 offerings (Intel Harpertown, Nehalem, Westmere) + Intel compilers
  - NVIDIA Tesla C1060/C2050 (soon) GPU Coprocessors
  - AMD quad-socket high-memory nodes
- Majority of users:
  - Just need reasonable pricing
  - Require highly-available compute and storage
  - Underlying technology is secondary
  - Need 10Gig or DDR+ IB, but GigE is often OK
- POD performs at “local” HPC cluster levels (or faster)
  - Compute is “on the metal”
  - Direct-attached, local storage
- HPC as a Service is generally a “step-up” for many users

- Tom Coull

Q3. Technology Issues or x86 everyone?….

Nimbis Services

- *Not all technical computing users require ultimate performance at a large scale (e.g. SMEs)*
- *SMEs are more concerned about ease of access, use and functional results*
- *Many variations of cloud performance points exist today*
- *Cloud providers are starting to provide non- x86 solutions (e.g. GPUs)*
- *Clouds may actually increase the use of specialized computing solutions due the lower barriers of access and cost*
- *Nimbis Services business model is to provide a technical computing "market place"*

- Bob Graybill

A3. Technology Issues or x86 everyone?….

- What's the problem?
  - If there's a price/perf benefit to a non-x86 CPU, and it meets other requirements (reliability, accuracy, etc.)…
    - And there's enough volume…
      - Someone will make it available as a service
- BTW, that's a VERY high jump bar
  - Non-x86 CPUs can make sense in many non-throughput situations
  - Careful apples::apples analysis shows x86 is spot-on for general throughput calculations running "normal" codes
    - Where it's not true… "out of whack" memory subsystem use (e.g. using 1/8th of a cache line, higher than normal b/w profile)
    - GPUs are one example of a successful departure from the x86 hegemony @ scale, esp. Fermi (ECC). What's the second?
- Being different is painful
  - Less so if you're in an appliance (who knows, who cares?!)
  - Less so if you're running captive codes (ditto)

- Rich Kaufmann

Q3. Technology Issues or x86 everyone?….

- x86 market choice for some time
- Cloud = remote compute/storage + v12n
- Q1: Can virtualization work for HPC?
  - Base performance
  - Accelerators: IB, GPGPU, etc.
- Q2: Why bother?
  - Access to cloud
  - Application resiliency (reactive and proactive)
  - Dynamic resource management
    - Increased efficiency
    - Power management
  - Support for heterogeneity, multi-tenancy

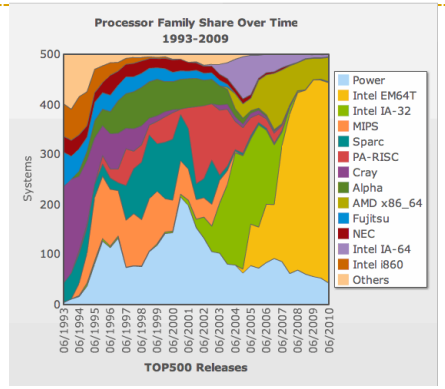

- *Josh Simons*

Q3. Technology Issues or x86 everyone?….

Cyclone Offers Flexible Choice of

- Platform
  - Scale-up (Altix UV), Scale-out (Altix ICE), Hybrid systems with GPU and alternative acceleration options (Tilera)
- Operating System (SUSE, Red Hat, Windows HPC, CentOS)
- Interconnect (NUMAlink, InfiniBand, GigE)
- Topology (hypercube, all-to-all, fat-tree, light weitgh)
- Physicalization and virtualization

Application-Specific Cloud

- 6 HPC segments
- Open source and close source applications

Differentiators

- Significantly broader platform flexibility
- SGI ProPack driven acceleration
- Application expertise as a Service

HPC USER FORUM

- Christian Tanasescu

Q3. Technology Issues or x86 everyone?....

- **“Cloud is Beowulf 2.0”**
  - By any metric, Cluster HPC solutions have been a success and the success has been based on the “plain vanilla x86 model”
    - X86 systems dominate purchasing
    - X86 systems are the majority of Top500 systems
  - Cloud inherits from the Beowulf legacy providing greater scale and broader access than traditional clusters.
  - Cloud solutions will evolve to provide even better HPC support, just as x86 clusters evolved to compete with traditional supercomputers.

- **Cloud further democratizes HPC**
  - Users previously unable to purchase, deploy, and maintain HPC clusters will have access to HPC software and services

- **The top remains the top**
  - There will always be room for very expensive, very custom supercomputing solutions.

HPC USER FORUM

- Ryan Waite

Q3. Technology Issues or x86 everyone?….

- *Clouds aren't a technology, they are a business model*
- *Commercial clouds are luring people to try to run HPC on a commodity web infrastructure.*
  - *Works for apps with low comm. & data needs*
- *A cloud can be tailored for "real" HPC*
  - *Single-tenancy nodes versus VMs*
  - *Simultaneous marshalling of a cluster*
  - *10G Ethernet or IB*
  - *Parallel file systems*
  - *GPUs, FPGAs & other specialized hardware*
  - *X86 or Cray or BlueGene or ….*

- Jeff Broughton

General Audience Q&A and wrap up

- ***Open floor….Time permitting!***