

HPC USER FORUM Cloud PANEL

**April 2011
Houston, TX**

**oohpc
USER FORUM**

Panel Participants

•Moderator

- Sharan Kalwani, KAUST

•Participants

- Bill Bryce UNIVA
- Raymond O'Brien NASA
- Shane Canon Lawrence Berkeley National Labs
- Pavan Pant CloudSwitch.com
- Chris Porter Platform Computing
- Jason Stowe CycleComputing.com

Panel Format

- **3 Questions**

- Provided ahead of time

- **3 minutes (max) per question for each participant response**

- **Follow-up and Audience after each participant has a chance to share**

Q1. Workflow Pre-qualifications HPC cloud...

➤ ***Please list what qualifies for HPC in the Cloud here.....***

Q1. Workflow Pre-qualifications HPC cloud...

There are several general areas that must be addressed to qualify HPC in the cloud they are:

Middleware Software (The HPC Platform)

- DRM systems, OS Provisioning, User Management, Filesystem Management, Configuration Management, Package/Image Management.
- User Tools, MPI, Parallel Shell
- API to access/Control Cloud
- Hypervisor (VMWare, Xen)

Storage

- Local/Remote, NFS, Hypervisor Storage, Snapshots, NAS Appliances, Storage alignment with Compute nodes.

Network

- Ethernet Gig or 10Gige, Infiniband, Low latency & High Bandwidth, VLAN (Isolation), VPN (bursting).

Compute Nodes

- Physical Nodes or Virtual, Multi-Core (8,12,16) nodes.
Sharing vs Dedicated access to nodes.

Any HPC in the Cloud solution has to address all of these areas to be successful.

Proper selection of above components provides a 'HPC Platform in a Cloud' and should deliver:

Scalability – to 1000+ nodes

Reliability – Handle failures of nodes, jobs, etc..

Performance – Physical Performance on Virtual Nodes



Workflow Characteristics

- *Loosely coupled or independent task execution*
- *Able to operate in virtualized environment*
- *Able to operate with heterogeneous set of processors*
- *Ideally, workflow:*
 - *Is tolerant of individual process failures*
 - *Consists of tasks that can be stopped and re-started to allow execution of higher priority work*
 - *Can take advantage of cloud scale up and scale out features*
 - *Can take advantage of a highly scalable object storage capability*

Q1. Workflow Pre-qualifications HPC cloud...LBNL



- ***High-performance interconnect***
- ***High-performance parallel I/O including concurrently writing to a shared object***
- ***Dedicated access to key resources (CPU, memory bandwidth, interconnect)***
- ***Highly scalable***

Q1. Workflow Pre-qualifications HPC cloud...

- ***The definition of Cloud includes virtualization to deliver on the on-demand requirement of cloud computing. Therefore, HPC in the cloud requires that the workloads and deployments be able to run in a virtualized environment:***
 - ***This means CPU virtualization (VMWare, XEN, Hyper-V)***
 - ***This often also includes network and storage virtualization***
- ***In addition HPC in the cloud is currently defined by horizontal scale. Most public clouds today have reasonable, bit limited vertical scale.***

Q1. Workflow Pre-qualifications HPC cloud...

- IaaS providers are not ready for full HPC yet
 - Amazon CCI + GPU instances are moving in the right direction
- Private clouds are the right choice for HPC today
- Private cloud characteristics:
 - Policy-based workload management
 - Flexibility
 - PM &/or VMs
 - Multi OS
 - Accountability
 - Job, user, group, project, OS, instance type

Q1. Workflow Pre-qualifications HPC cloud...



HPC in the Cloud Workloads (Type, Data, Hardware, Security, Access)

- Parallel - Only 1-node or non-low-latency sensitive
- Embarrassingly Parallel – Parameter sampling, Monte Carlo
- Data – Doesn't necessarily depend on total size (web 2.0 use PetaBytes), depends upon deltas and bandwidth to transfer
 - Data aware synchronization/placement is a requirement
 - Massive scales are done but still challenging (PB)
- GPUs workloads supported (e.g. Mol. Dyn. workload cost $\frac{1}{4}$ as much)
- Scheduling w/multi-user support
- Virtualized / Isolated data storage
- Multi-tenant, with virtualized, isolated access/auth
- Programmatic access to Resources (no required phone calls to providers)
- Pay for what you use

Q2. Best HPC in the Cloud example?.....

- **In your experience, which is the most successful HPC in the cloud implementation and what was different about it?**
- **Your citation here.....**

A2. Best HPC in the cloud usage/implementation

- **Amazon EC2 is the best HPC cloud.**
 - EC2 has many limitations but...
 - VPC (Virtual Private Cloud),
 - CCI (Cluster Compute Instances)
 - Dedicated access to hypervisors
 - Multiple instance 'sizes'
 - 'De-facto' standard API
 - Clearly the most popular and used HPC cloud

A2. Best HPC in the cloud usage/implementation

- Examples:
 - NASA JPL
 - Pathwork Diagnostics
 - Cycle Computing
 - Abaqus in EC2 (Univa)



- **NASA Wide Field IR Survey Explorer Project**

- 4 IR cameras have imaged entire sky *each*
- Needed: resolution enhancement, deconvolution
- Project 1: 2000 distant galaxies
 - Increase resolution with processing = 100 CPU hrs *each*.
- Project 2: Some sky areas require huge RAM-based processing.
 - - Set up an 80 GB RAM instance
 - - Finished first phase on Nebula in two days.
 - - WISE scientist: "I am amazed"

A2. Best HPC in the cloud usage/implementer



- **Quotes from Dominic Benford, WISE project scientist:**

“After having access to Nebula for only a few days, I find that I am already able to accomplish more data-intensive calculations than I can do on any of the local servers we have here, and with no difficulties at all!

“The system is easy to access and use, and offers a capability that I absolutely need occasionally but for which I could never justify the expense if it were for my needs alone.

A2. Best HPC in the cloud usage/implementer



- **Additional quote from Dominic Benford, WISE project scientist:**

“With the recent addition of a large-RAM instance, I am now able to conduct calculations that could not be done on our project's large server farm. Nebula has provided me with a tool for science data analysis that far surpasses anything that I could envision in a single-user context. NASA Cloud computing may be the way forward for our data-intensive projects in the future, since only a NASA system could provide the necessary reliability and proprietary controls on our data “

A2. Best HPC in the cloud usage/implementer



- **Cloud is well suited for NASA projects:**

- With HPC workloads possessing the characteristics described earlier
- Constrained by availability of compute resources
- That cannot justify the cost of dedicated infrastructure for key processing which is only required occasionally or for a limited duration

- **Another example:**

- See NASA presentation titled: Proposed Use of the NASA Ames Cloud Computing Platform for Numerical Weather Prediction and the Distribution of High Resolution Satellite Imagery
- By Ashutosh S. Limaye, Andrew L. Molthan, and Jayanthi Srikishen
- 2010 AGU Fall Meeting, San Francisco, CA
- Session IN21D-08: "Research Clouds: Virtualization of Infrastructure, Tools, and Services I"

Available from the NASA Aeronautics and Space Database

A2. Best HPC in the cloud usage/implementation



- **In your experience, which is the most successful HPC in the cloud implementation and what was different about it?**
- **What do we mean by successful?**
 - Best known?
 - Most profitable?
 - **Most useful to the user/customer!**

A2. Is an HPC Center Cloud-like enough?



HPC Centers ?

- **Resource pooling.**
- **Broad network access.**
- **Measured Service.**
- **Rapid elasticity.**
 - Usage can grow/shrink; pay-as-you-go.
- **On-demand self-service.**
 - Users cannot demand (or pay for) more service than their allocation allows
 - Jobs often wait for hours or days in queues



A2. Best HPC in the cloud usage/implementation

- **Automated provisioning (PXE boot of nodes, ROCKS configuration)**
- **Data set sizes reasonable for transport across Internet / WAN (both initial set, and results)**
- **Latency between the DC and the Cloud was not a significant factor – interactive performance, reach-back was not in-line with processing**
- **Provided an outlet for overtaxed internal resources**

A2. Best HPC in the cloud usage/implementation

- **Our most successful HPC cloud deployment was a 1000 core compute node with ROCKS and Sun Grid Engine for use in the Pharma industry. It was a good deployment because:**
 - It provided on-demand scaled compute deployment
 - The application was designed for horizontal scale
 - Had reasonable network and storage bandwidth requirements

A2. Best HPC in the cloud usage/implementation

- **Pfizer**

- IAAS provider: Amazon EC2
- Technical ingredients:
 - Amazon VPC
 - Serial, memory-bound workload
 - Application: RosettaDock
 - Embarrassingly parallel
 - Small data
- <http://biopharmadirectory.com/news/?p=159>

Q2. Best HPC in the Cloud example?.....

CycleCloud has started 2000+ clusters since 2007, for **Pfizer, Schrodinger, Varian, Inc., Fortune 500 Insurance/ Finance, Genomic Health**, many, many others & read about numerous others' work.

My favorite is our CycleCloud cluster announced on 4/5/11:

Single click creates 10000-core cluster in Cloud, for Genentech, in minutes, for \$1060/hour (Cycle + IaaS)

NETWORKWORLD

This story appeared on Network World at <http://www.networkworld.com/news/2011/040611-linux-supercomputer.html>

10,000-core Linux supercomputer built in Amazon cloud

Cycle Computing builds cloud-based supercomputing cluster to boost scientific research.

By [Jon Brodtkin](#), Network World
April 06, 2011 03:15 PM ET

The Register®

Being the hand that feeds IT



Original URL: http://www.theregister.co.uk/2011/04/06/cycle_computing_hpc_cloud/

Cycle Computing fires up 10,000-core HPC cloud on EC2

Only \$1,060 per hour, management included

By [Timothy Prickett Morgan](#)

Posted in Cloud, 6th April 2011, 16:05 GMT

A2. Best HPC in the cloud usage/implementation

- Full cluster environment with Shared File system (TBs), and Scheduler
- Secured data transfer and at-rest using AES 128-/256-bit
- Researcher demand: 80000 compute hours for protein binding simulations
- Run-time: 8 hours on 1250 servers, 10000 cores
- Up-front costs: \$0, no software to buy, pay for what you use
- Operating Cost: \$1060/hr ~\$.106/hr, Total cost: \$8500
- User Effort: Single click, no programming / IT admin (unless you want to)

A2. Best HPC in the cloud usage/implementation

Scientists were able to focus on the science, get immediate results

Great for Peak/Parameter Sweep/Monte Carlo workloads

In General:

- <http://blog.cyclecomputing.com> 2000-core, 80GPU, 4000-core, 10000-core
- Support Any Scheduler (GridEngine, Torque, Condor, as well as licensed)

- Using CycleCloud:

Sign up for an account,
Configure cluster (Filesystem, etc.)
Launch cluster
Log in 5-20 minutes later



Q3. What HPC in the cloud is not?....

➤ *Examples of HPC in the cloud gone awry.....*

A3. What HPC in the cloud is not?

➤ **Not:**

- **A Solution for moving your data in and out of cloud**
- **Applicable to all applications**
 - **Sometimes things don't scale the same way**
- **A solution to license management**
 - **Some commercial applications have 'expensive' license schemes.**
- **Cheaper than your own cluster!**
 - **It is more expensive per hour than internal**
 - **Assuming you need to use the cluster all the time**
- **Guaranteed Secure**
 - **Unless you get the dedicated, storage, network and compute!**
 - **Or you are running on a Cloud with Intel 'Trusted Compute'**

A3. What HPC in the cloud is not?




- A substitute for the highly optimized HEC systems NASA operates to address most large-scale CFD and similar workloads
- Yet readily accepted within NASA as a substitute for project-specific infrastructure due to:
 - Historical reliance on, and comfort with, “owned infrastructure,” despite the inherent constraints
 - A business model for NASA cloud services that is still under development
 - User concerns about the effectiveness of cloud security
 - Current lack of a solution engineering group to assist projects with cloud adoption

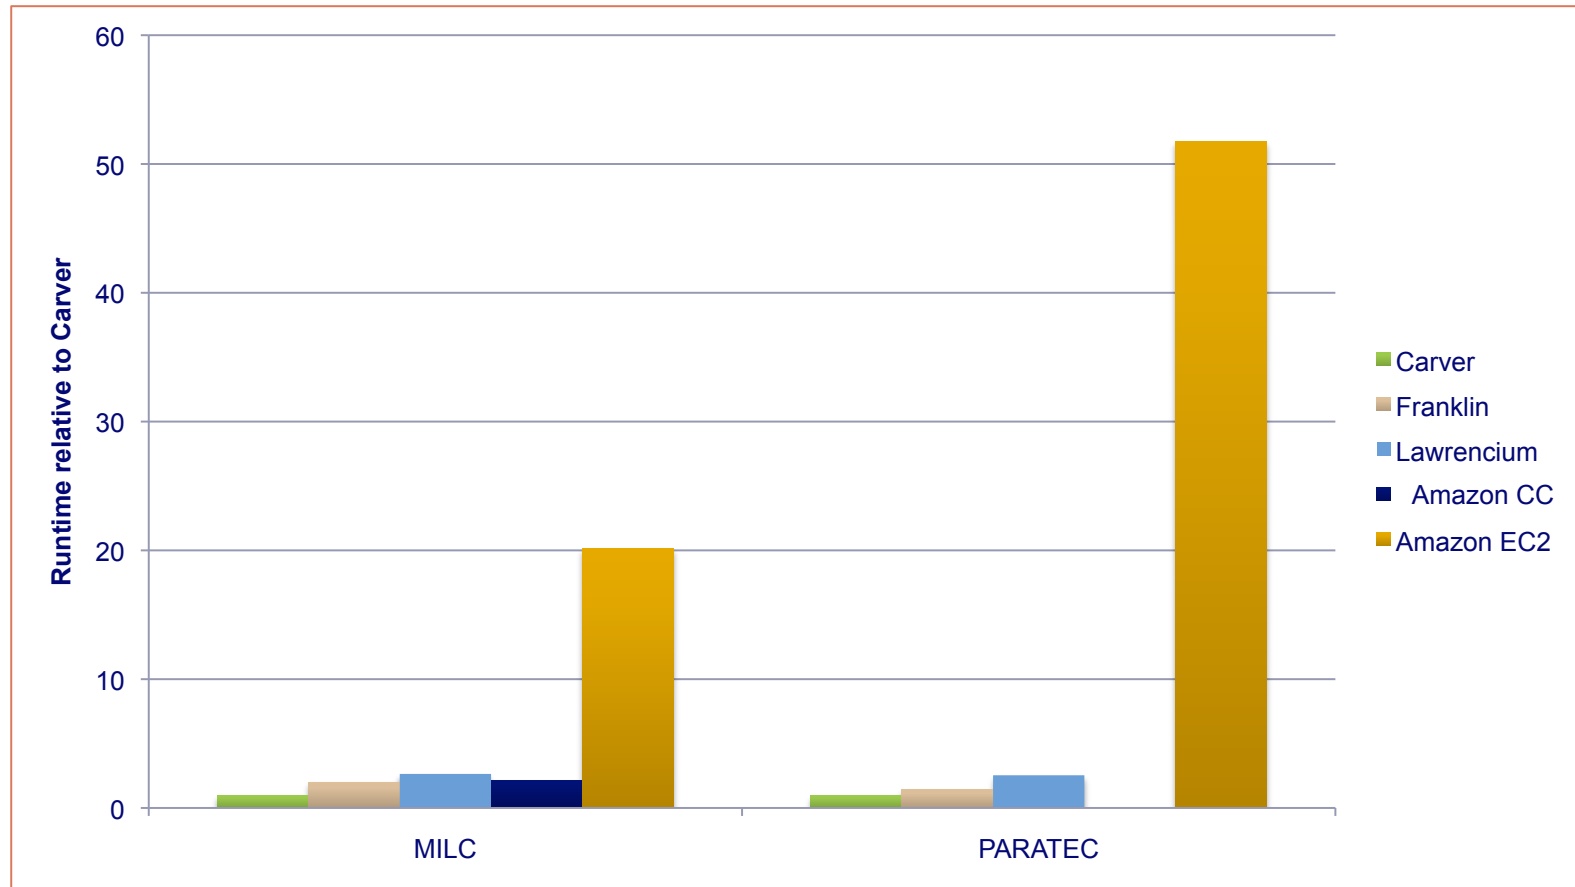
A3. What HPC in the cloud is not?

- ***Business/Web solutions repurposed as HPC***
 - ***Virtualized and shared with other users (unpredictable/variable performance)***
 - ***Commodity network***
 - ***Spread across a data center (or multiple data centers)***
 - ***No simple IO solution***

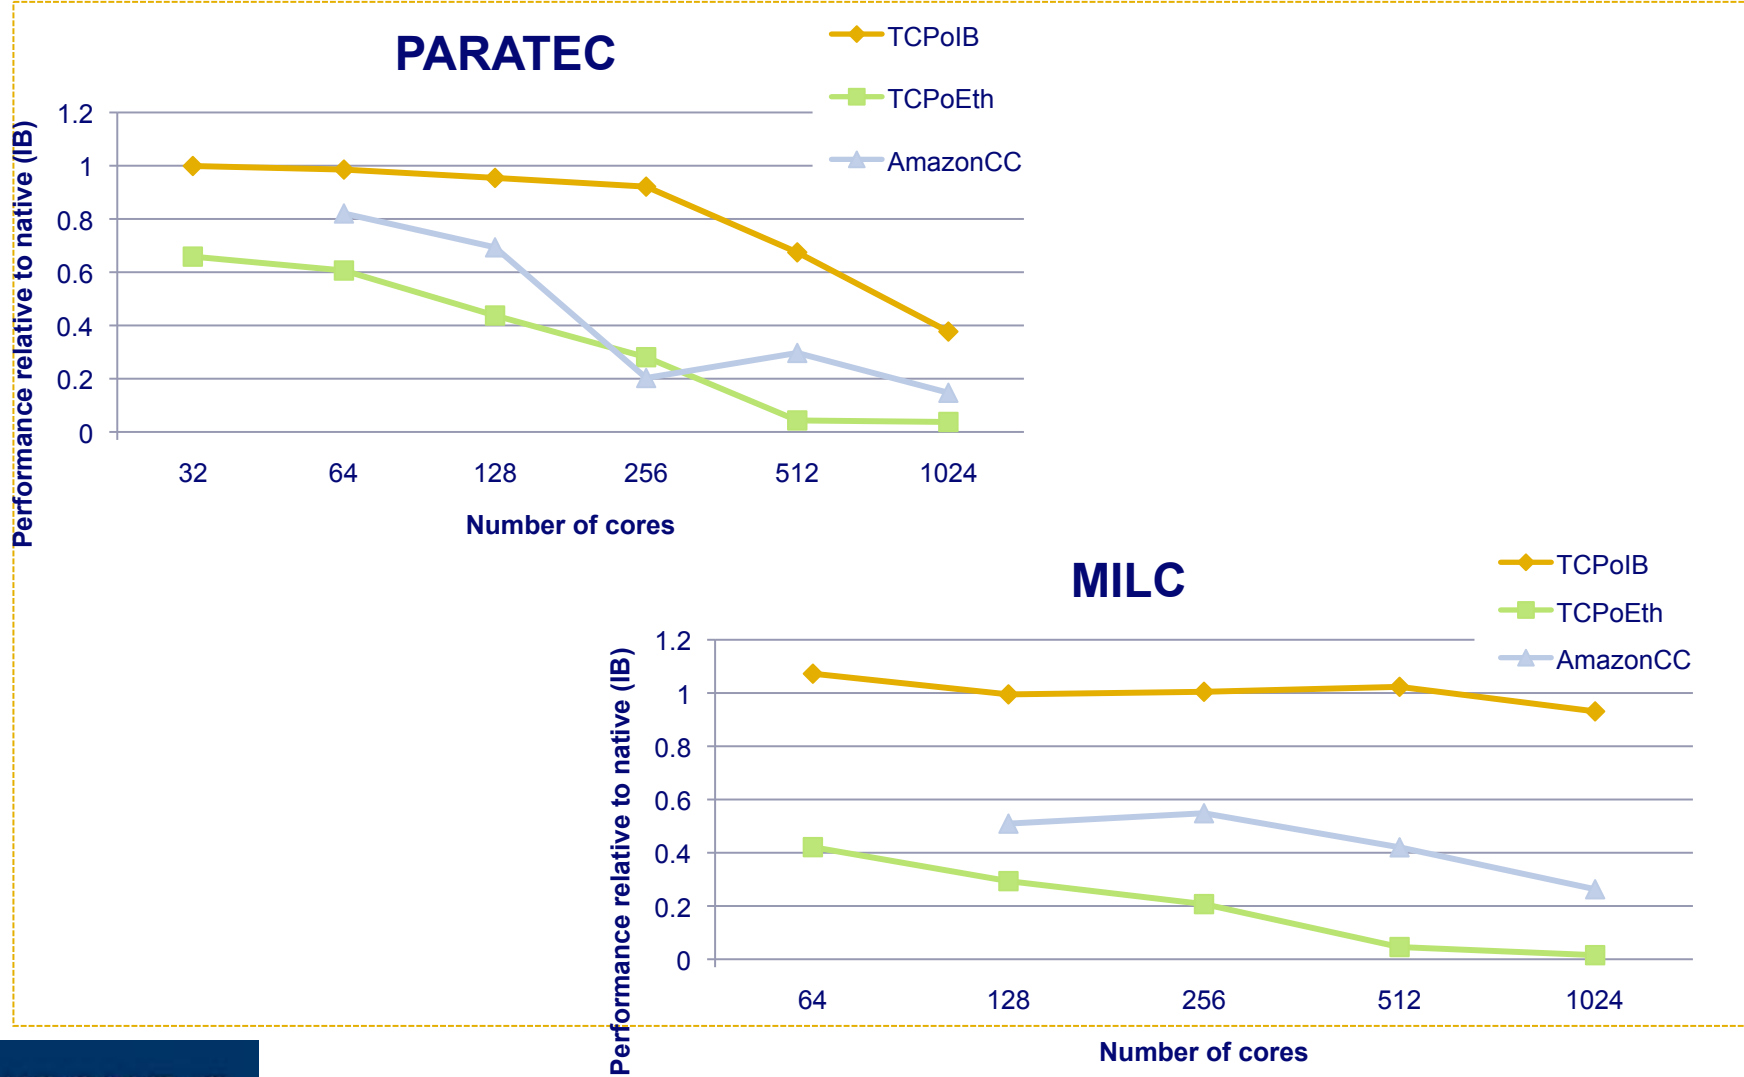
A3. What HPC in the cloud is not?



A3. What HPC in the cloud is not?



A3. What HPC in the cloud is not?



A3. What HPC in the cloud is not?

- *Deploying applications into the Cloud from an environment where the compute, storage, and networking is highly tuned.*
- *Vertical scale applications – i.e. E15K*
- *Specialized/dedicated hardware*

A3. What HPC in the cloud is not?

- ***Applications that require highly tuned Storage back-end (i.e. Tuned SAN)***
 - ***The public clouds today have poor SLA/QOS attributes. This leads to varying performance of the storage and networking infrastructure based on multi-tenant loading.***
 - ***Any applications that are sensitive to this will bottleneck, or potentially fail – and at least they will have unpredictable run-times***
- ***HPC that requires specialized H/W – this is impacted by both virtualization, and cloud vendor hardware selection.***
- ***For virtualization, the issues is if (and how well) the virt. Platform can deliver the functionality. For the vendor selection – the clouds are designing for majority use cases (and may not be focused on HPC)***

A3. What HPC in the cloud is not?

- **A good choice for fine grained parallel jobs**
 - Network latency
- **A good choice for jobs with large data (in/out)**
 - Data charges + time for transfer
- **A good choice for long running parallel jobs**
 - Reliability still in question

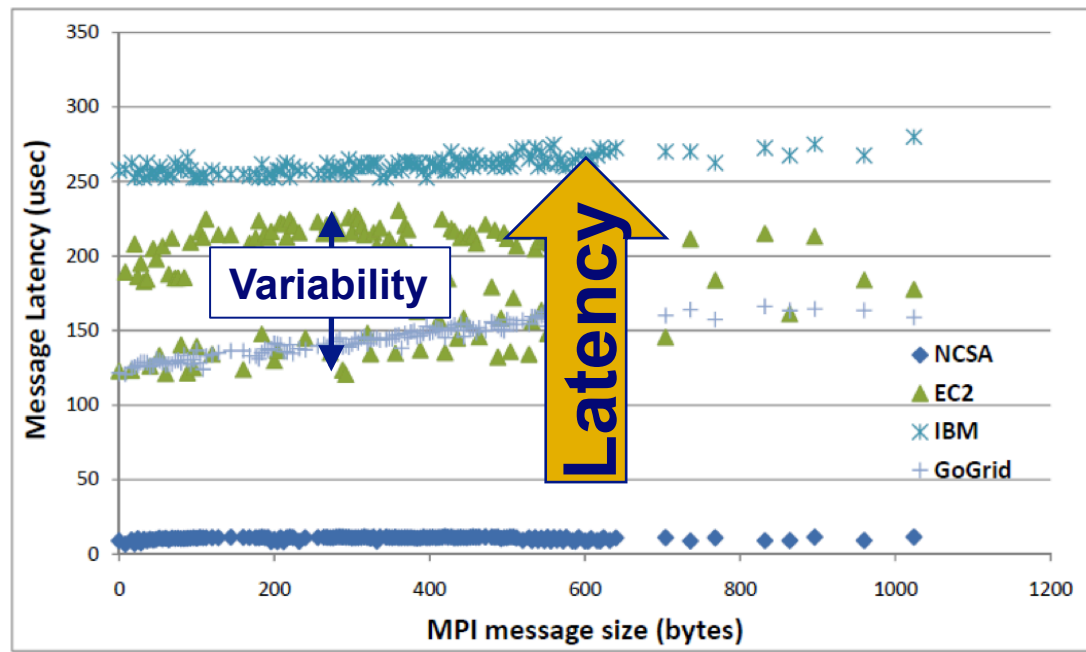


Chart: Q. He, S. Zhou, B. Kobler, dc Duffy, and T. mcglynn,
"Case study for running HPC applications in public clouds", in Proc. HPDC, 2010

Q3. What HPC in the cloud is not?....

➤ HPC in the Cloud is

- Massive Scale, On-Demand, Cost transparency, with no upfront expense, Performance conscious, Data conscious, & “No phone calls required”

Is not:

- Current “Cloud-bursting” methodologies in various schedulers, and various tools
 - Automatically shifting tasks to cloud resources
 - Ignores the data, provisioning a VPN causes far more downloading of data, latency, etc.
 - Bad for your bill, bad for performance (upside?)

A3. What HPC in the cloud is not?

- Non-instantaneous capacity
 - Counter to Cloud agility, no phone should be needed
 - Low-latency, extremely big I/O
- Various “automated”-VPN software layers for servers
 - These hurt performance, are proprietary, encourage security laziness (dangerous), use open source tools
- Shared environments that don't use low-level virtualization or dedicated server/strict auth systems to quarantine data.
- Systems that require Web API programming and don't do managed data placement

General Audience Q&A and wrap up

➤ *Open floor....*