



SC12 Cloud Computing for Science Tutorial: Cloud Challenges

Kate Keahey, John Breshnahan, Patrick Armstrong, Pierre Riteau
Argonne National Laboratory
Computation Institute, University of Chicago

Cloud versus Cloud



Custom user environments!
On-demand access!
Elastic computing!
Growth and cost management!
Capital expense -> operational expense!



Do I need to become a sys admin?
How long does it take to deploy a virtual cluster of a few thousand nodes ?
Elasticity is fine, but how do I use it in practice over many incompatible clouds?
It costs too much! And what if Amazon raises prices?
You can't do HPC on the cloud!

Appliance Management

I want to manage my own VMs!



I don't want to manage my own VMs!



- Users require control
- The emergence of community appliance management and maintenance
- More and better tools

- Users are not skilled in appliance management
- Roles take time to emerge
- The configuration/deployment trade-offs are not obvious

Deployment Performance

- Scientific images have lots of stuff on them
- We typically deploy lots of them at a time


- LANTorrent
- Large deployments often operate on same image
- Caching

On Magellan: a thousand VMs in 10 minutes

11/13/12


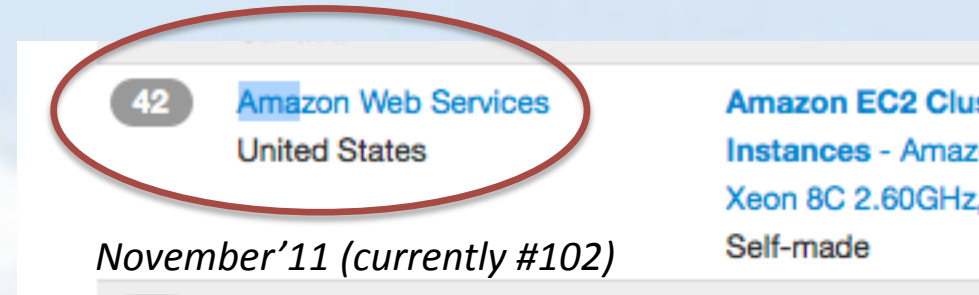


www.nimb



Execution Performance

- In 2011, ~87% of the top500 systems were commodity clusters
- Amazon Cluster Compute
- I/O is an issue, but there is a configuration trade-off
- Aspects of performance:
 - Hardware
 - Configuration
- Availability@price




CPU cost for TCP connection at 1 Gbps
(xen-unstable (03/16/2007) ; PV Linux guest; X86 - 32bit)



How good are clouds for HPC?
“Comparisons: not as Odious as Once Thought”,
see www.scienceclouds.org/blog
Soon to be updated...

What's Not To Like?

- Latency!
- ... and bandwidth
- ... and lack of IB
- ... and anything having to do with I/O
- ... and noise (esp. with multi-tenancy)
- ... and MTBF




Courtesy of Ramakrishnan et al., PMBS'11


How good are clouds for HPC?
"Mahommad and the Mountain",
see www.scienceclouds.org/blog

What's Not To Like?

- ... and overall impact on tightly-coupled applications



PARATEC




MILC

Courtesy of Ramakrishnan et al., PMBS'11

...but Wait, There's More!

- New hypervisors
 - KHVMM (Blue Gene/P)
 - Palacios
- Xen-IB (since ~2006)
- “Symbiotic virtualization”
 - Passthrough I/O
 - Preemption control
 - Optimized paging
- How do we put that into practice?




*CTH (shockwave simulation)
within ~5% of native performance
On Cray XT4 (Sandia)*

From Lange et al., VEE' 11

HPC Cloud BOF Tomorrow @ 5:30!

- Are clouds a viable platform for HPC?



- Join us for a fun discussion in 355-BC!

Security

- Isolation
 - Allows the user high degree of freedom while mitigating risk to the provider
 - Network traffic controlled via virtual routers
- Virtualization in security solutions
 - Short-lived VM deployments
 - Honeypots
 - Support for security training

What's Not To Like?


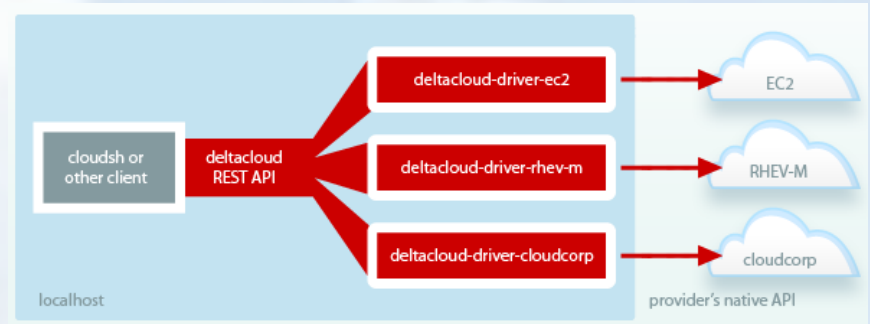
- Layers:
 - Hypervisor, virtual network, control domain
 - Cloud infrastructure
 - VM image management
- More complexity == more things to attack
- New model == new attacks
- Incentives: commercial versus community
- Data Privacy
 - Trusting the provider
 - Trusted computing
 - Homomorphic encryption
- New legal challenges

Cloud Fungibility

- Can't easily move from one provider to another:
 - Cloud infrastructure APIs are incompatible
 - Contextualization across providers is incompatible
 - VM image formats are incompatible
- No easy way to compare performance across cloud providers
- High risk of provider lock-in

What's to like?


- Emergent API standards:
 - Cloud standards: OCCI (OGF), OVF (DMTF), and many more...
 - www.cloud-standards.org
- IaaS API adapters
 - Deltacloud, jcloud, libcloud, and many more...
- VM image standards....
- ...and adapters
 - rBuilder, BCFG2, CohesiveFT, Puppet, and many more...
- Emergent efforts to compare performance



The Bitsource: CloudServers vs EC2 LKC Cost by Instance

Cost, Utilization, and Price

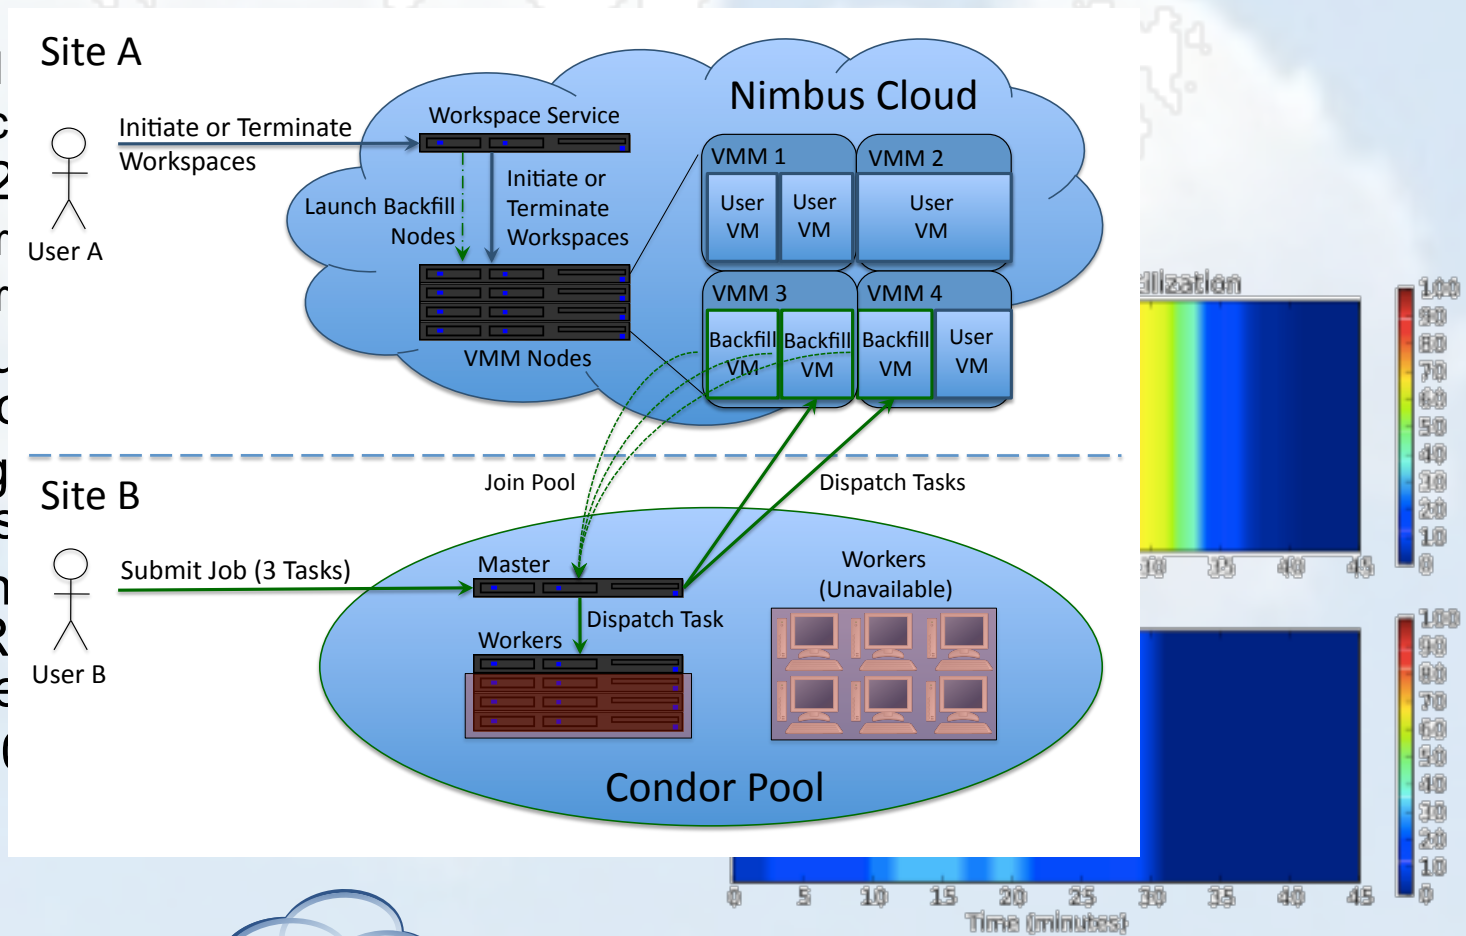
- Most science today is done in batch
 - Not very responsive...
 - ... but very efficient!
- **Challenge:** On-demand catch-22 for providers
 - You can overprovision (Expensive!)
 - Or you can reject requests (Not really on-demand)
- Utilization -> Cost -> Price



*Utilization on the Fusion cluster @ ANL
courtesy of Ray Bair, LCRC*

Cost, Utilization, and Price

- Solution 1:
 - Backfill with volunteer VMs
- Benefits:
 - Up to 1
 - Significant
- Solution 2
 - Spot pr
 - Support
- Open Sou
- contributi
- Preparing
- workloads
- Extension
- Service R
- Nimbus re
- CCGrid 20



Parting Thoughts

- Infrastructure clouds are a disruptive innovation
 - Opportunity: on-demand availability
 - But also many challenges!
- “Crossing the chasm”
 - We are living in the age of early adopters
 - Broadening the set:
 - Complexity and scale
 - New models
 - Performance and features
 - The challenges are in crossing the chasm
- Impact on Applications
 - The ability to absorb resources is critical
 - Fault-tolerance is increasingly a requirement
 - Emphasis on programming models that support them

Questions?

All tutorial slides will be available at:
www.nimbusproject.org/docs/sc12

Acknowledgements



NSF OCI
“FutureGrid”



DOE ASCR
“Infrastructure Strategy to Support Collaboration”