“Clouds” in the U.S. ATLAS Facilities

Michael Ernst
Brookhaven National Laboratory

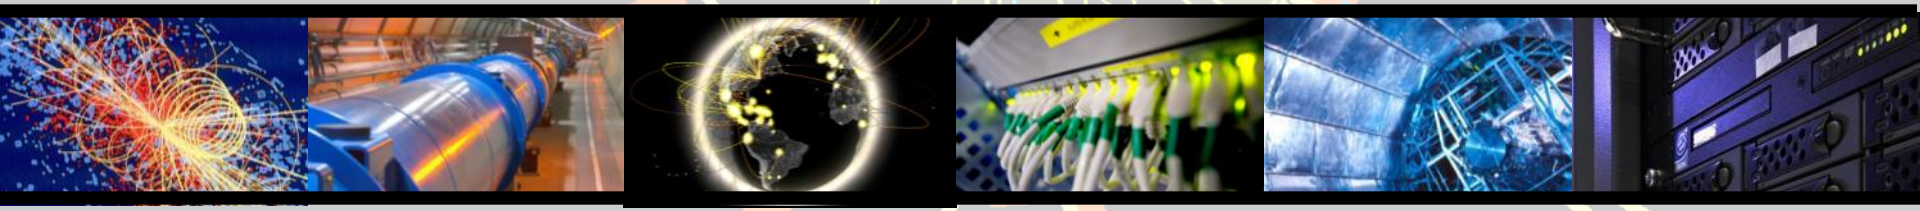

Aspects forming a Cloud System

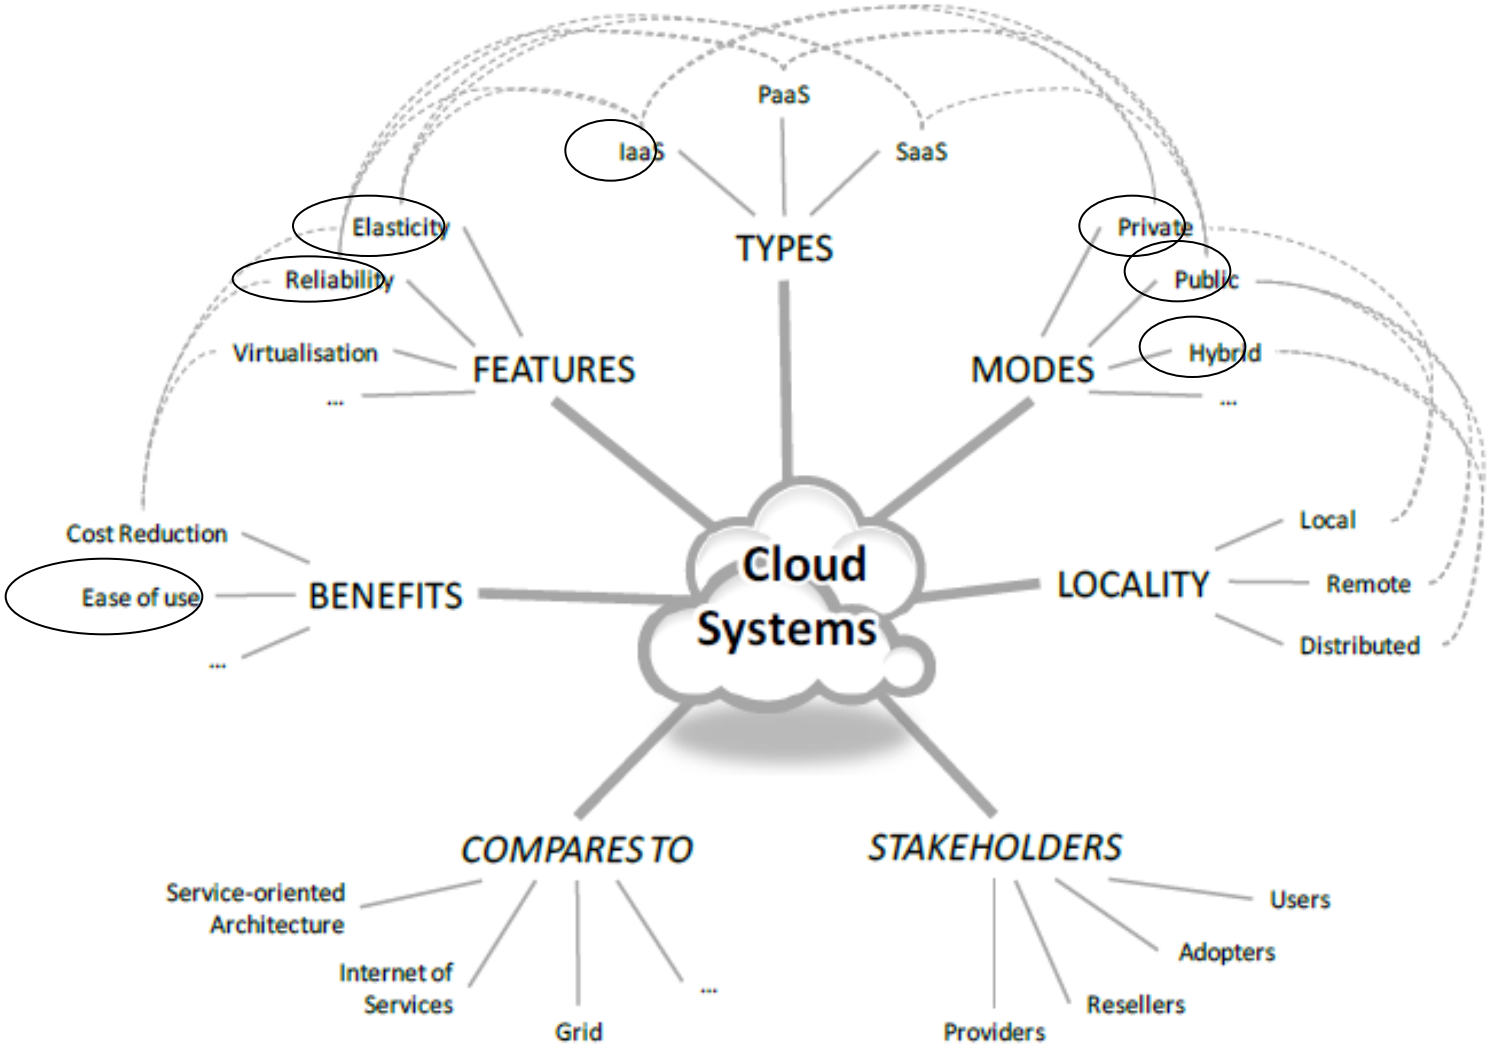

Goal

- Run large Condor pool on multiple cloud platforms and/or providers.
  - Dynamically based on waiting workload.
  - Spread across large area, possibly multi-continental.
  - Include facility OpenStack instance(s).
  - Utilize spot pricing on EC2.
  - Run typical ATLAS production (simulation) workloads on it.
    - » Workload may expand to Analysis once we understand data transfer and cloud storage implications
- Why?
  - Free up high-performance Tier 1 resources for user analysis by moving low I/O work to EC2/GCE and/or academic clouds.
  - Pre-position ATLAS to utilize additional cloud-based resources that might become available.

mernst@bnl.gov

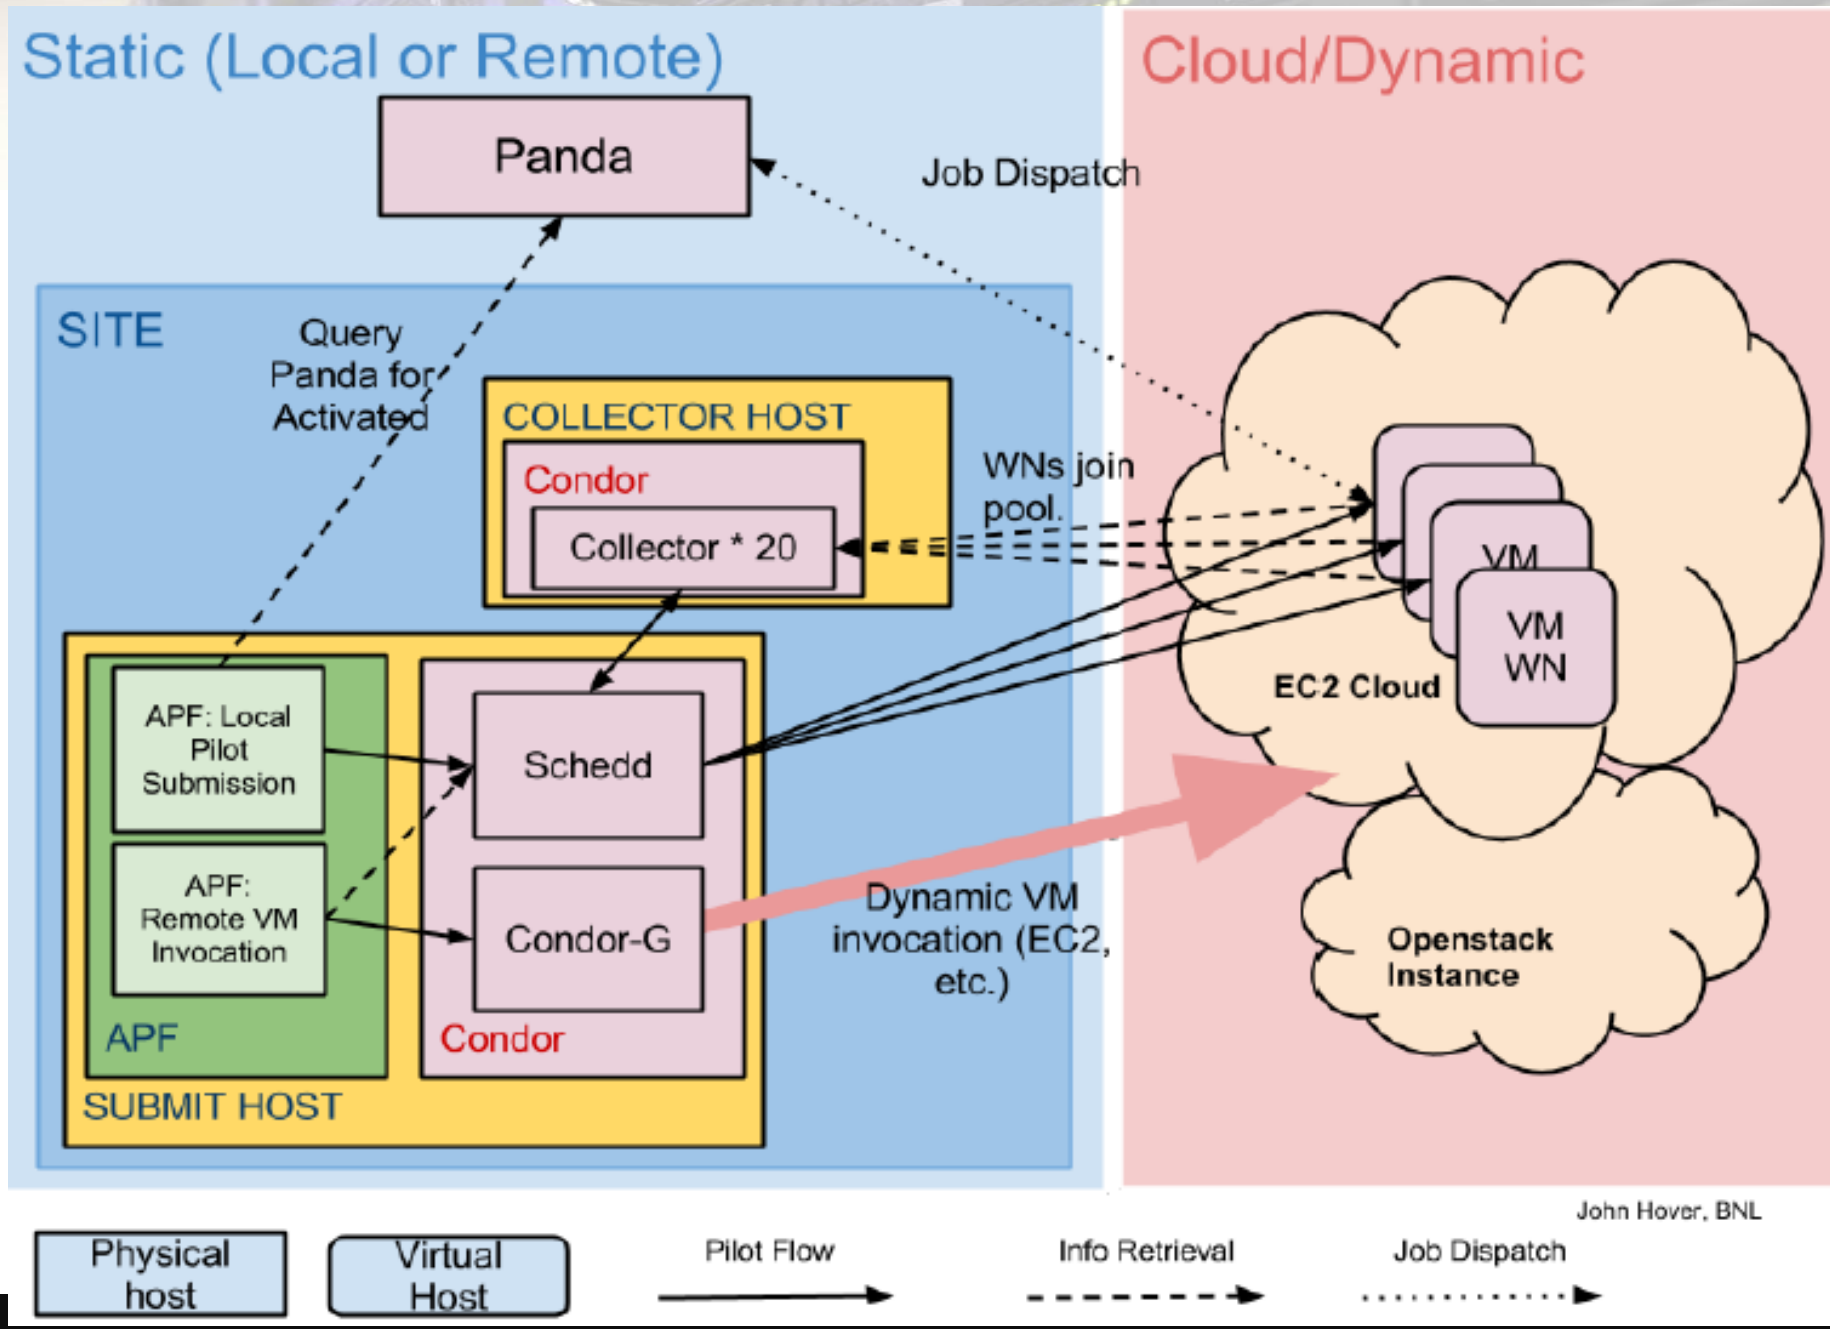
Static (Local or Remote)
Cloud/Dynamic
Panda
Job Dispatch
SITE
Query Panda for Activated
COLLECTOR HOST
Condor
Collector * 20
WNs join pool.
APF: Local Pilot Submission
APF: Remote VM Invocation
APF
Schedd
Condor-G
Condor
SUBMIT HOST
Dynamic VM invocation (EC2, etc.)
VM
VM WN
EC2 Cloud
Openstack Instance
John Hover, BNL
Physical host
Virtual Host
Pilot Flow
Info Retrieval
Job Dispatch

Capabilities

Code based on APF's intrinsic queue/plugin architecture, allows:

- Multiple targets
  - E.g., submission to EC2 us-east-1, us-west-1, us-west-2 in a load-balanced fashion.
- Policy and Cost-aware Resource Scheduling
  - We can preferentially utilize free site clouds (e.g. local Openstack or other academic clouds)
  - Once that is full we submit to EC2 spot-priced nodes.
  - During particularly high demand, submit to EC2/GCE on-demand nodes.
  - Retire and terminate in reverse order.

Amazon EC2 Pricing Options

	Up-Front Costs	Hourly Rates	Availability
On Demand Instances	None	Highest	Not Guaranteed to start. Amazon will not terminate.
Instance Reservations	Varies with Reservation Type and term length. Guaranteed availability	Varies with Reservation Type and term length.	Guaranteed in specific Region / Availability Zone
Spot Requests	None	Variable, typically lowest.	May be terminated by Amazon when spot price exceeds bid

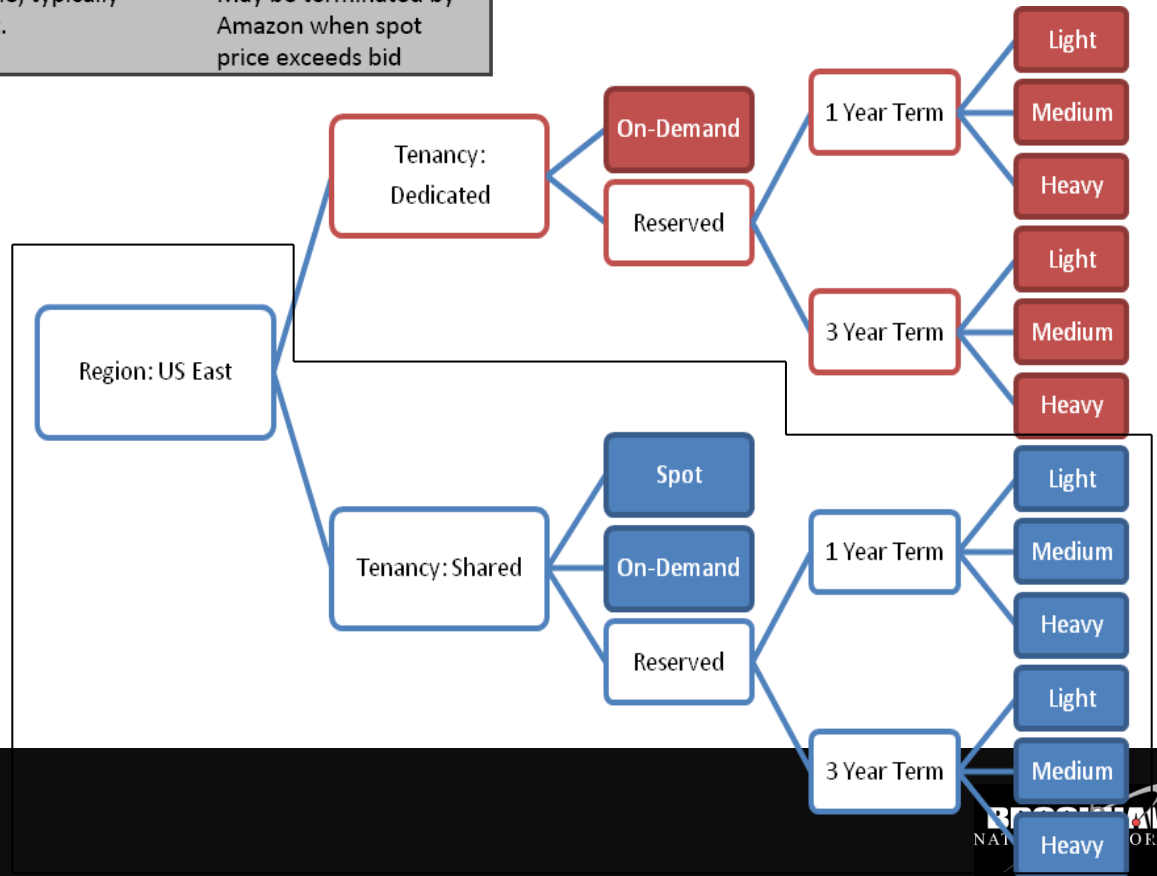

Reserved Pricing Examples

Term	On-Demand		Light		Medium		Heavy	
	Up-Front	Hourly	Up-Front	Hourly	Up-Front	Hourly	Up-Front	Hourly
None	-	$0.24	-	-	-	-	-	-
1 Year	-	-	$276	$0.156	$640	$0.096	$780	$0.064 8,760 Hours
3 Year	-	-	$425	$0.124	$1000	$0.076	$1200	$0.052 26,280 Hours

Rates for Linux Large Instance

Term	Light	Medium	Heavy
1 Year	3286 Hours	4444 Hours	5586 Hours
	137 Days	185 Days	233 Days
	38% of Term	51% of Term	64% of Term
3 Years	3663 Hours	6098 Hours	10694 Hours
	152 Days	254 Days	445 Days
	14% of Term	23% of Term	41% of Term

On-Demand/Reserved Break-Even Point for Linux Large Instance

Term	On-Demand	Light	Medium	Heavy
None	$6,307	-	-	-
1 Year	-	$4,928	$4,443	$4,022
3 Year	-	$3,683	$2,997	$2,566

Three Year Cost of Linux Large Instance

Cost/Core at Tier-1 in 2012

	USATLAS (7328 cores)	RHIC (9504 cores)
Server	$210/yr	$229/yr
Electrical	$6/yr	$8/yr
Space	$9/yr	$9/yr
Staff	$42/yr ($500/server/yr)	$29/yr ($345/server/yr)
Total	$267/yr ($0.030/hr)	$275/yr ($0.031/hr)

Σ $225/yr

- Amortize server over 4 years and use only physical cores (12)
- Assume 2 FTE each for US ATLAS and RHIC ($200k/FTE)
- Except for electrical, all other costs are burdened
- Calculations for calendar year 2012.

- Cost Comparison using the lowest cost for a Reserved EC2 Instance (3-year term) comparable w/ performance of In-house core
  - AWS EC2: $2,566
  - At BNL Tier-1: $ 675

EC2 Spot Pricing and Condor

On-demand vs. Spot

- On-Demand: You pay standard price. Never terminates, but not guaranteed to start.
- Spot: You declare *maximum* price. You pay current, variable spot price. If/when spot price exceeds your maximum, instance is terminated without warning. Note: NOT like priceline.com, where you pay what you bid.

EC2 Spot Market

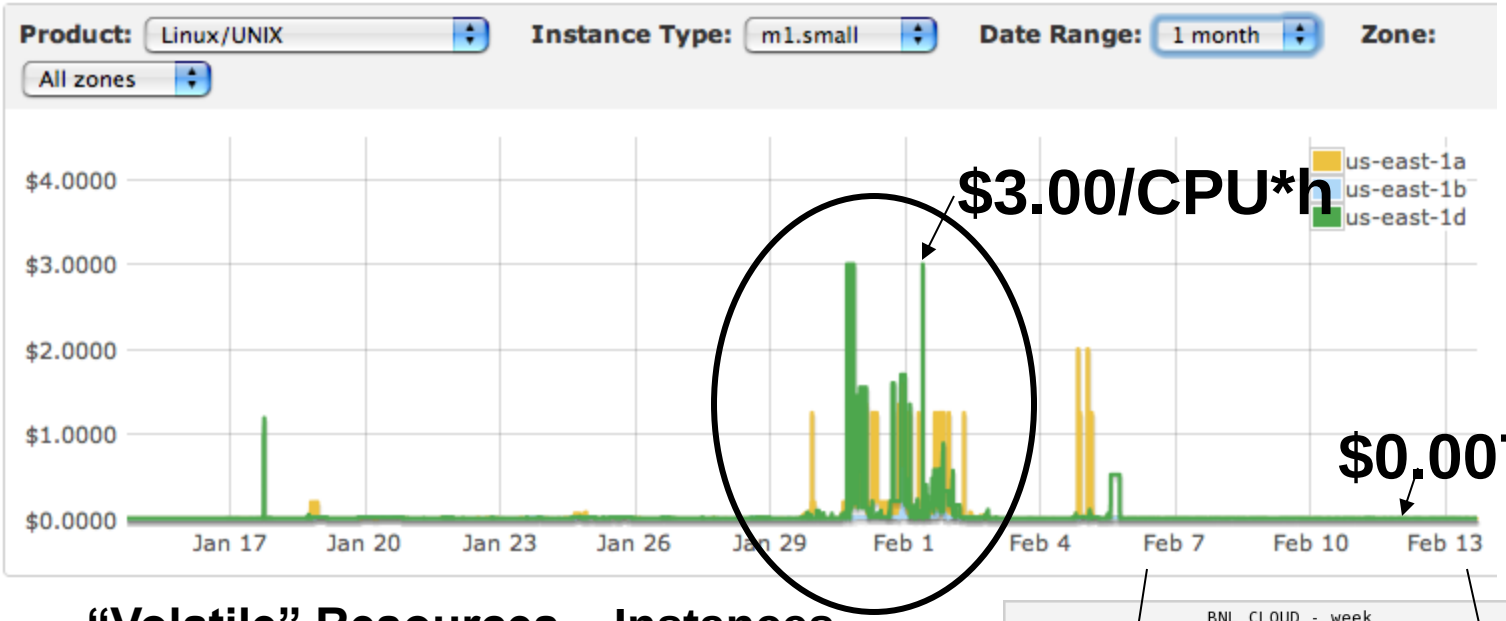

"Volatile" Resources – Instances obtained w/ lower bids are terminated

Issues w/ EC2 spot market

- **User job termination w/o warning**
- **Even at moderately sized requests (e.g. 5k) we found it's sometimes not possible to get the number of instances one wants/needs**

5k concurrent jobs ~$1,000/day

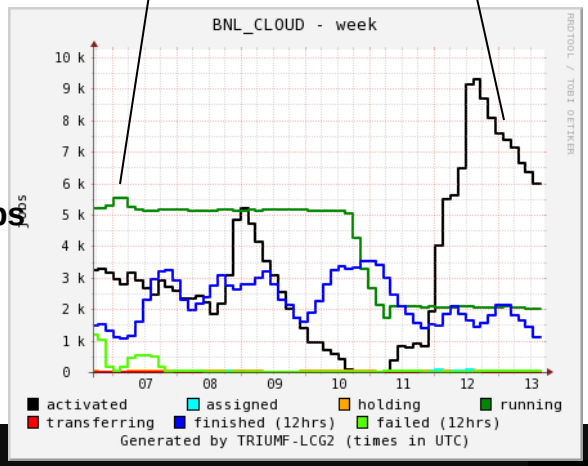

Cost of Network Transfer

Source	Destination	Pricing	Price Per GB
Amazon EC2	Amazon EC2, in the same availability zone, using private IP addresses.	Availability Zone Data Transfer Out	$0.00 – Flat Rate
Amazon EC2	Other AWS Services (like S3) in the same Region.	No Charges	$0.00 – Flat Rate
Anywhere	Amazon EC2	Internet Data Transfer In	$0.00 – Flat Rate
Amazon EC2	Amazon EC2, in a different availability zone in the same Region	Regional Data Transfer Out	$0.01 – Flat Rate
Amazon EC2	Internet	Internet Data Transfer Out	$0.05 – Tiered Pricing

- While transfers into EC2 are “free”, traffic out of EC2 to e.g. the Tier-1 could be rather costly
- 5k Simulation jobs (assume 1 day wall time, produce ~2GB/job, ~300TB/month) => ~$15k/month
  - Even with low-volume simulation cost for I/O is not negligible, would me much higher for analysis

BNL OpenStack Network

Potential I/O Bottleneck
- **multiple gateways**
- **public IP on second interface**

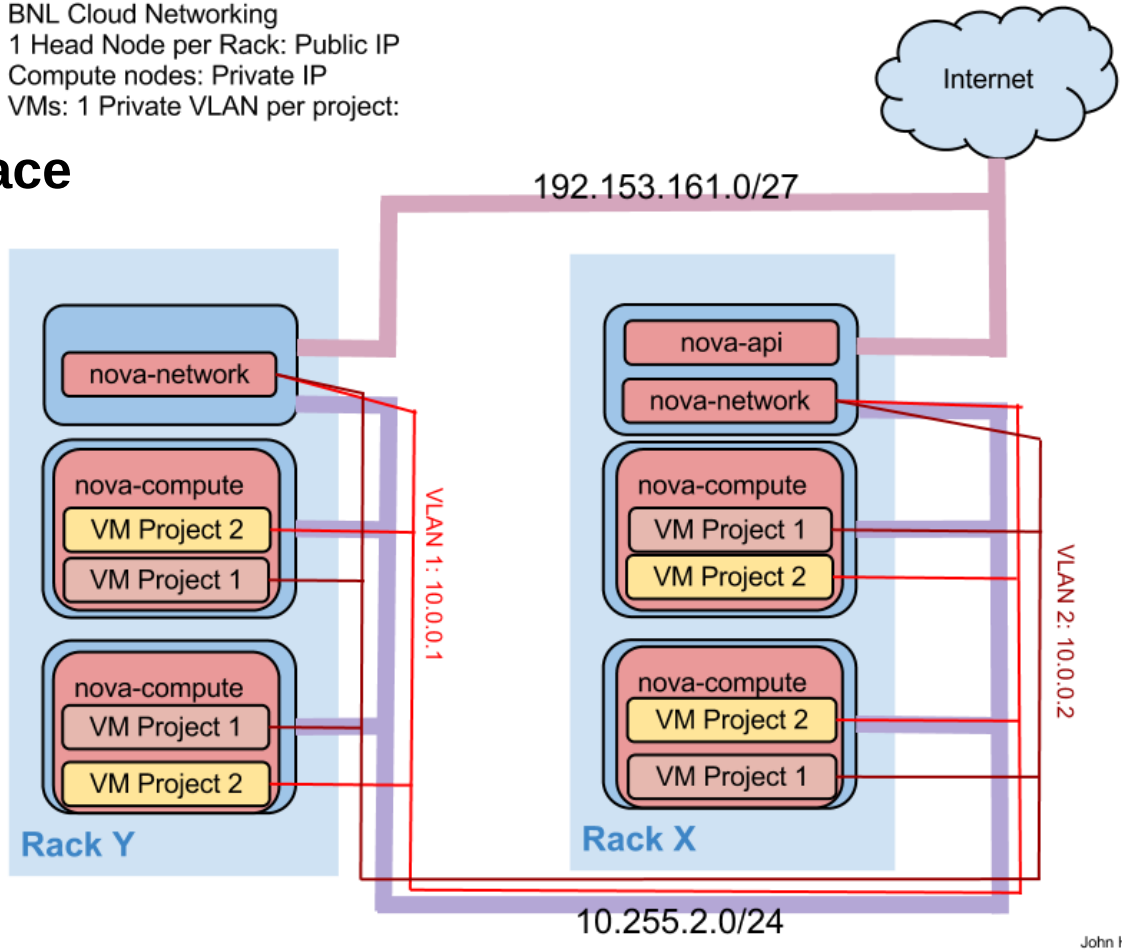

mernst@bnl.gov

Networking in a Virtualized Computing Environment

- "Cloud Networking" is driven by Virtualization and Mobility
  - Different than enterprise networking
  - Needs to be built of a strong S/W foundation
    - Enterprise is about connectivity, cloud & SDN is about optimizing applications
- Disruption requires new kind of S/W answer
  - No more one physical server w/ one connection
  - Thousands of VMs instead
  - Storage no longer SAN, Hadoop clusters
  - Core count increasing – millions of cores in future

➢ With such change, network & I/O is bottleneck

Cloud Principles

- User joins Community the Cloud creates
  - Issue of multi-tenancy at the Network level as much as at CPU/server, DB levels
  - Shared resources must be shared in a way that applications from different users don't interfere
  - Resources have to be partitioned so they are private and secure
    - While IP & Ethernet have virtual network capability, they are limited in terms of number of tenants and how isolated a tenant is

Effects of Networking on Cloud Environments

- Network Latency could be a cloud killer
  - Cloud offers benefits of scalability & flexibility but w/o enough network capacity benefits are small
- Shifting workloads off-premises to clouds burdens network between cloud and our own data centers
  - Applications may suffer from poor performance
  - Network monitoring to understand network needs
- ➢ When considering adding cloud resources network demands must be part of the plan
  - ➢ I.e. WAN b/w remains one of least elastic components of the infrastructure

The Network as a Partner in Cloud Computing

- Amazon's "Elastic IP" is an application-driven approach to integrating the Network and the Cloud
- OpenStack includes Network services as one resource it virtualizes as with CPU/server and Storage
  - OpenStack's Quantum interface defines how a virtual network can be created to "host" CPU and other elements
    - Quantum does not define the technology to create virtual networks
    - Cloud provider is responsible for mapping their technology to virtual network models

OpenStack Quantum

- With Folsom/Grizzly release OpenStack users have OpenStack Quantum, an industry-standard, open API for Cloud Networking orchestration
  - Providing abstraction layer
  - Instrument to solve Network Operations problems in Cloud Computing and virtualized Data Centers
  - Decouples operations from network mechanics
  - OpenStack had API for Compute called “Nova”
    - Networking was hardcoded
    - Quantum makes Networking Subsystem pluggable, modular, standard interface to run any number of virtual networking solutions
    - Quantum allows to create virtual Networks in an industry-independent way, attach those networks to VMs and orchestrate configuration of services

Importance of SDN in Cloud Networking

- SDN is increasingly accepted as path to Cloud Networking
  - Meaning the transformation of networks and services to support the use of cloud computing
  - Navigating the various missions and technology models of SDNs is critical to properly position cloud services and realize advantages of clouds
- ➢ Knowing cloud provider's SDN plans, as well as plans of private cloud S/W stack vendors, is the most critical element in assessing these providers' long-term value

Backup

mernst@bnl.gov

Non-Functional Aspects

- Elasticity
  - Capability of infrastructure to adapt to changing, potentially non-functional requirements
    - Amount of data supported by app, # of concurrent users
  - There is horizontal and vertical scalability
    - Horizontal: # of instances to satisfy changing demands
    - Vertical: Size of instances
  - Cloud scalability involves rapid up- and downscaling

Non-Functional Aspects

- Availability and Reliability
  - Ensure constant operation without disruption
    - No data loss, no code reset during execution, etc
  - Typically achieved through redundant resource utilization
    - Many aspects move from hardware to software-based solutions (e.g. redundancy in the file system vs. RAID controllers, stateless front end servers vs. UPS, etc)
  - Provide redundancy for services and data so failures can be masked transparently
    - Fault tolerance also requires ability to introduce new redundancy, online and non-intrusively

Non-Functional Aspects

- Agility and Adaptability
  - Capabilities that strongly relate to Elasticity
    - On-time reaction to changes (# of requests and size)
    - Changes to environmental conditions (types of resources)
    - Requires resources to be autonomic and enable them to provide self-* capabilities

Elastic Cluster Components

- AutoPyFactory (APF): Coordinates submissions
  - One APF queue observes a Panda queue, submits pilots to local Condor pool.
  - Second APF queue:
    - Observes a local Condor pool, when jobs are Idle,
    - submits WN VMs to IaaS (up to a limit).
    - Notices spot terminations, submits additional VMs.
- Worker Node VMs
  - Condor startds join back to local Condor cluster.
  - Publish EC2 attributes from UserData via ClassAd mechanism.
- HTCondor pool
  - Static Central Manager (Collector, Negotiator, CCB)
  - Local pilot submission

mernst@bnl.gov

EC2 Spot Considerations

Service and Pricing

- Nodes terminated without warning. (No signal.)
  - The HTCondor Team has a solution
- Partial hours are *not charged*.

Therefore, we need to consider:

- Shorter jobs. Simplest approach. Originally ATLAS worked to ensure jobs were *at least* a couple hours, to avoid pilot flow congestion. Now we have the opposite need.
- Checkpointing. Some work in Condor world providing the ability to checkpoint without linking to special libraries. (But not promising.)
- Per-event stage-out (event server).

If we had sub-hour units of work, we could get resources for free!