Running Scientific Workflow Applications on the Amazon EC2 Cloud

G. Bruce Berriman

NASA Exoplanet Science Institute, California Institute of Technology

Ewa Deelman and Gideon Juve

Information Sciences Institute, University of Southern California

Should I Use The Cloud?

- **The answer is….it depends on your application and use case.**
- **Recommended best practice: Always perform a cost-benefit analysis to identify the most cost-effective processing and data storage strategy because Amazon itemizes charges for processing, data transfer and data storage.**
- Amazon offers the best value for compute- and memory-bound applications. However, data storage, and especially data transfer costs, exceed processing costs for data-intensive applications.
- If you need the best performance, HPCs with parallel file systems and high-speed networks are a better choice than Amazon for data-intensive applications.

Where Can I Learn More?

- This poster is based on the paper **Scientific Workflow Applications on Amazon EC2** by G. Juve, E. Deelman, K. Vahi, B. Berriman, P. Berman, and P. Maechling, published by the Cloud Computing Workshop in Conjunction with e-Science 2009 (Oxford, UK). **http://arxiv.org/abs/1005.2718**
- **Cloud Computing in the Age of Data-Intensive Science**, by G. B. Berriman, E. Deelman and G. Juve. An SPIE Newsroom Release (May 25, 2010) **http://spie.org/x40451.xml?highlight=x2418&ArticleID=x40451**
- Visit Bruce Berriman's blog, "Astronomy Computing Today," at **http://astrocompute.wordpress.com**

What Are Our Goals?

The study was designed to answer the question: **How useful is cloud computing for scientific workflow applications?**

Workflow applications are loosely coupled applications in which the output files from one component become the input to the next.

There were three goals:

1. Conduct an experimental study of the performance **of three workflows** with different I/O, memory and CPU requirements on a commercial cloud.
2. Compare the performance of cloud resources with the performance of a typical High Performance Cluster (HPC). The cloud uses **commodity hardware and virtualization** and HPCs use **parallel file systems and fast networks**.
3. Provide an **analysis of the various costs** associated with running workflows on a commercial cloud.

We chose Amazon EC2 as the cloud provider and the NCSA Abe cluster as a high-performance cluster.

http://www.ncsa.illinois.edu/UserInfo/Resources/Hardware/Intel64Cluster

http://aws.amazon.com/ec2/

Computing Resources

Type	Arch	CPU	Cores	Memory	Network	Storage	Price
Amazon EC2							
m1.small	32-bit	2.0-2.6 GHz Opteron	1-2	1.7 GB	1-Gbps Ethernet	Local	$0.10/hr
m1.large	64-bit	2.0-2.6 GHz Opteron	2	7.5 GB	1-Gbps Ethernet	Local	$0.40/hr
m1.xlarge	64-bit	2.0-2.6 GHz Opteron	4	15 GB	1-Gbps Ethernet	Local	$0.80/hr
c1.medium	32-bit	2.33-2.66 GHz Xeon	2	1.7 GB	1-Gbps Ethernet	Local	$0.20/hr
c1.xlarge	64-bit	2.0-2.66 GHz Xeon	8	7.5 GB	1-Gbps Ethernet	Local	$0.80/hr
Abe							
abe.local	64-bit	2.33 GHz Xeon	8	8 GB	10-Gbps InfiniBand	Local	…
abe.lustre	64-bit	2.33 GHz Xeon	8	8 GB	10-Gbps InfiniBand	Lustre	…

Processor and Network Specifications

Processors and OS

- Amazon offers a wide selection of processors. Our choices reflect the range of options to look at cost vs. performance.
- Ran Linux Red Hat Enterprise with VMware
- *c1.xlarge* and *abe.local* are equivalent – estimate overhead due to virtualization
- *abe.lustre* and *abe.local* differ only in file system

Networks and File Systems

- HPC systems use high-performance network and parallel file systems.
- Amazon EC2 uses off-the-shelf hardware.
- Ran all processes on single, multi-core nodes. Used local and parallel file system on Abe.

Amazon Processing Costs

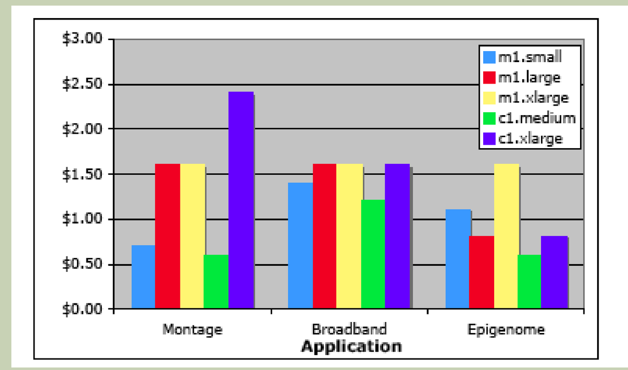

Montage:

- Clear trade-off between performance and cost.
- **Most powerful processor *c1.xlarge* offers 3x the performance of *m1.small* – but at 5x the cost.**
- Most cost-effective processor is *c1.medium* – 20% performance loss over *m1.small*, but 5x lower cost.

Broadband and Epigenome:

- No reason to choose anything other than the most powerful machines.

The Applications

Montage (**http://montage.ipac.caltech.edu**) creates science-grade image mosaics from multiple input images.

Broadband simulates and compares seismograms from earthquake simulation codes.

Epigenome maps short DNA segments collected using high-throughput gene sequencing machines to a reference genome.

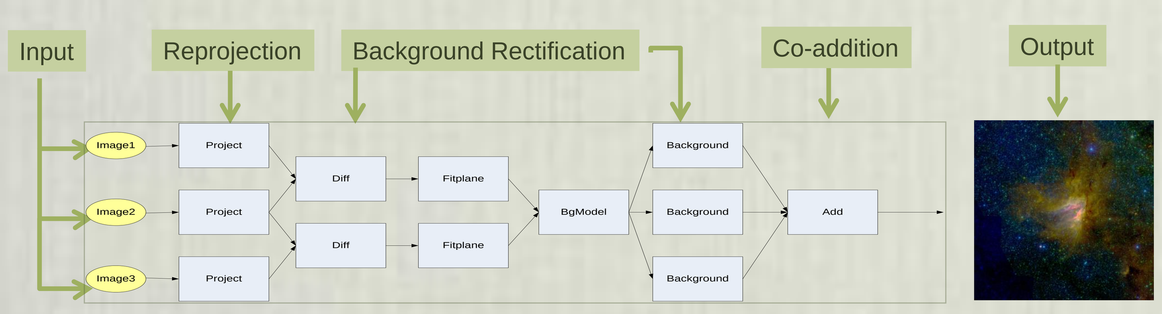

Montage Workflow

Application	I/O	Memory	CPU
Montage	High	Low	Low
Broadband	Medium	High	Medium
Epigenome	Low	Medium	High

Resource Usage of the Three Workflow Applications

Application	Workflow	# Tasks	Input	Output
Montage	8 deg. sq. mosaic of M16, 2MASS K-band	10,429	4.2 GB	7.9 GB
Broadband	4 earthquake sources, 5 sites	320	6 GB	160 MB
Epigenome	Maps DNA sequences to ref. chromosome 21	81	1.8 GB	300 MB

Workflow Specifications for this Study

Runtime Performance

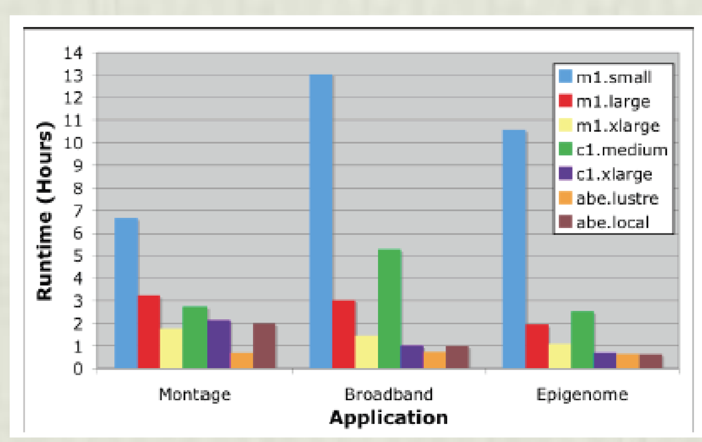

Montage (I/O-Bound)

- Fastest on those machines with the most core and largest memory: *m1.xlarge, c1.xlarge, abe.lustre, and abe.local.* Linux kernel uses the available memory for the file system buffer cache, and this reduces time waiting for I/O.
- The parallel file system on *abe.lustre* offers a big performance advantage of x3 for I/O-bound systems. Cloud providers would need to offer parallel file systems and high-speed networks to compete with the HPC.
- Virtualization overhead on Amazon <8%

Broadband (Memory-Bound)

- Lower I/O requirements – little difference between *abe.lustre* and *abe.local.*
- Amazon can achieve same performance as Abe if there is more than 1 GB memory per core. available
- Poor performance on *c1.medium* – only 1.7 GB of memory. Cores may sit idle to prevent system running out of memory.
- Virtualization overhead on Amazon is very small.

Epigenome (CPU-Bound)

- *c1.xlarge*, *abe.lustre* and *abe.local* give best performance – they are the three most powerful machines (64-bit, 2.3-2.6 GHz)
- The parallel file system on *abe.lustre* offers little performance advantage.
- Virtualization overhead on Amazon is roughly 10% -- the application competes for CPU with the OS.

Data Storage Costs

Data Storage Costs

- Amazon charges for storing Virtual Machines (VM) and user's applications in local disk
- It also charges for storing data in network-attached Elastic Block Storage (EBS).

Item	Charges ($)
Storage of VMs on local disk	$0.15 GB/Month
Storage of data in EBS disk	$0.10 GB/Month

Storage Volumes

Application	Input (GB)	Output (GB)	Logs (MB)
Montage	4.2	7.9	40
Broadband	4.1	0.16	5.5
Epigenome	1.8	0.3	3.3

Storage Costs

Application	Data ($)	VM ($)	Monthly Cost ($)
Montage	$0.95	$0.12	$1.07
Broadband	$0.02	$0.10	$0.12
Epigenome	$0.22	$0.10	$0.32

Montage Storage Costs Exceed Most Cost-Effective Processor Costs

Data Transfer Costs

Transfer Rates

- Amazon charges different rates for transferring data into the cloud and back out again.
- Transfer-out costs are the higher of the two.

Operation	Cost ($)
Transfer In	$0.10/GB
Transfer Out	$0.17/GB

Transfer Costs

- For Montage, the **cost to transfer data out of the cloud is higher** than monthly storage and processing costs.
- For Broadband and Epigenome, **processing incurs the biggest costs.**

Application	Input	Output	Logs	Total
Montage	$0.42	$1.32	<$0.01	$1.75
Broadband	$0.40	$0.03	<$0.01	$0.43
Epigenome	$0.18	$0.05	<0.01	$0.23