Cloud Computing

Alex Crawford

Ben Johnstone

Overview

- What is cloud computing?
- Amazon EC2
- Performance
- Conclusions

What is the Cloud?

- A large cluster of machines
  - Economies of scale [1]
- Customers use a set of machines, platforms, or services for their products
- Virtualization allows for rapid scaling and deployment

History [2]

- Concept originated from telecommunication companies changing to VPN
- 1999: Salesforce.com - delivery of applications via web
- 2002: Amazon launches Amazon Web Services (AWS)
- 2006 : Google Docs, Amazon Elastic Compute Cloud (EC2)
- 2008: Eucalyptus
- 2009: Microsoft Azure

Types of Cloud Services

- Infrastructure as a Service
  - Provides the physical resources
- Platform as a Service
  - Provides a software platform to build on
- Software as a Service
  - Provides a computation service

Examples of Cloud

- Amazon Web Services (IaaS)
- Microsoft Azure (PaaS)
- WolframAlpha (SaaS)

Amazon Web Services

- Variety of services
  - Simple Storage Service (S3)
  - Elastic Compute Cloud (EC2)
  - DynamoDB
  - Virtual Private Cloud
- Clients buy VM time
- Programmatically control VMs (boot, shutdown, etc.)
- Build virtual networks

Amazon Web Services (cont.)

- Amazon backs the VMs with physical machines
- Datacenters around the world
- Uses a redundant network topology

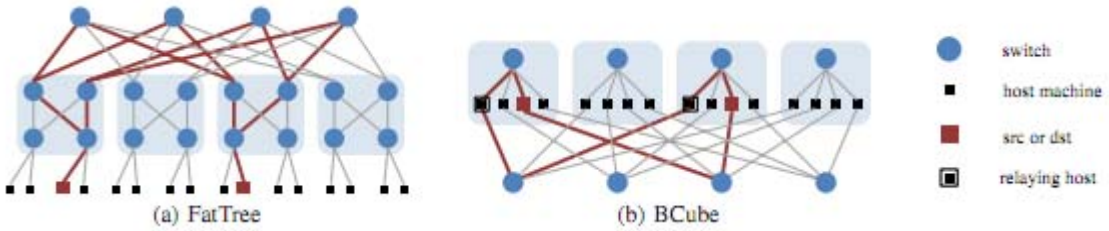

Elastic Compute Cloud

- Started in 2006 [3]
  - 2010: Cluster instances introduced
- Provides resizable compute capacity in the cloud
- Pay for what you use
- Access VM instances through a web interface
- Compute Unit: measure of capacity
  - 1-1.2 GHz 2007 Opteron/Xeon

Instance Types

Name	Memory (GB)	# Compute Units	# Virtual Cores	Storage	I/O Performance
Small	1.7	1	1	160	Moderate
Extra Large	15	8	4	1690	High
luster Compute Quadruple Extra Large	23	33.5	2 x Intel Xeon X5570, quad-core “Nehalem”	1690	Very High (10 Gb Ethernet)
Cluster Compute Eight Extra Large	60.5	88	2 x Intel Xeon E5-2670, eight-core "Sandy Bridge"	3370	Very High (10 Gb Ethernet)
Cluster GPU Quadruple Extra Large	22	33.5	2 x NVIDIA Tesla “Fermi” M2050 GPUs	1690	Very High (10 Gb Ethernet)

Instance costs

Name	Linux/UNIX Cost ($ per hour)	Windows Cost ($ per hour)
Small	0.08	0.115
Extra Large	0.640	0.920
Cluster Compute Quadruple Extra Large	1.300	1.610
Cluster Compute Eight Extra Large	2.400	2.970
Cluster GPU Quadruple Extra Large	2.100	2.600

Virtualization

- Abstracts guest OS from physical HW
- VMs managed by hypervisor
  - Handles resource sharing
- HW assisted and paravirtualization
- For cloud:
  - Configurability
  - Maintenance
  - Availability

Virtualization Issues

- Networking resources shared among VMs
- Each VM can't have its own IP address
  - Solution: NAT
- Oversubscription
  - More resources given to VMs than are physically available
  - High contention for physical resources

So, how is this useful?

- Scalable
  - Small personal machine or large supercomputer
- Configurable
- Available
- You don't buy the hardware
  - No need to buy and maintain your own supercomputer

42nd Fastest Supercomputer

- Built on AWS
- 30,000 processing cores
- 240 TFLOPS
- $1,300/h to operate

Performance vs. NCSA's Abe

- Experiment pitted EC2 vs. Abe [4]
  - Abe specs representative of existing HPC systems
- 3 Programs performed:
  - Montage
    - Astronomical images using data from telescopes
    - Lots of I/O
  - Broadband
    - Generates and compares seismograms
    - High memory cost
  - Epigenomics
    - Maps DNA segments to genomes
    - Heavy computation

Results

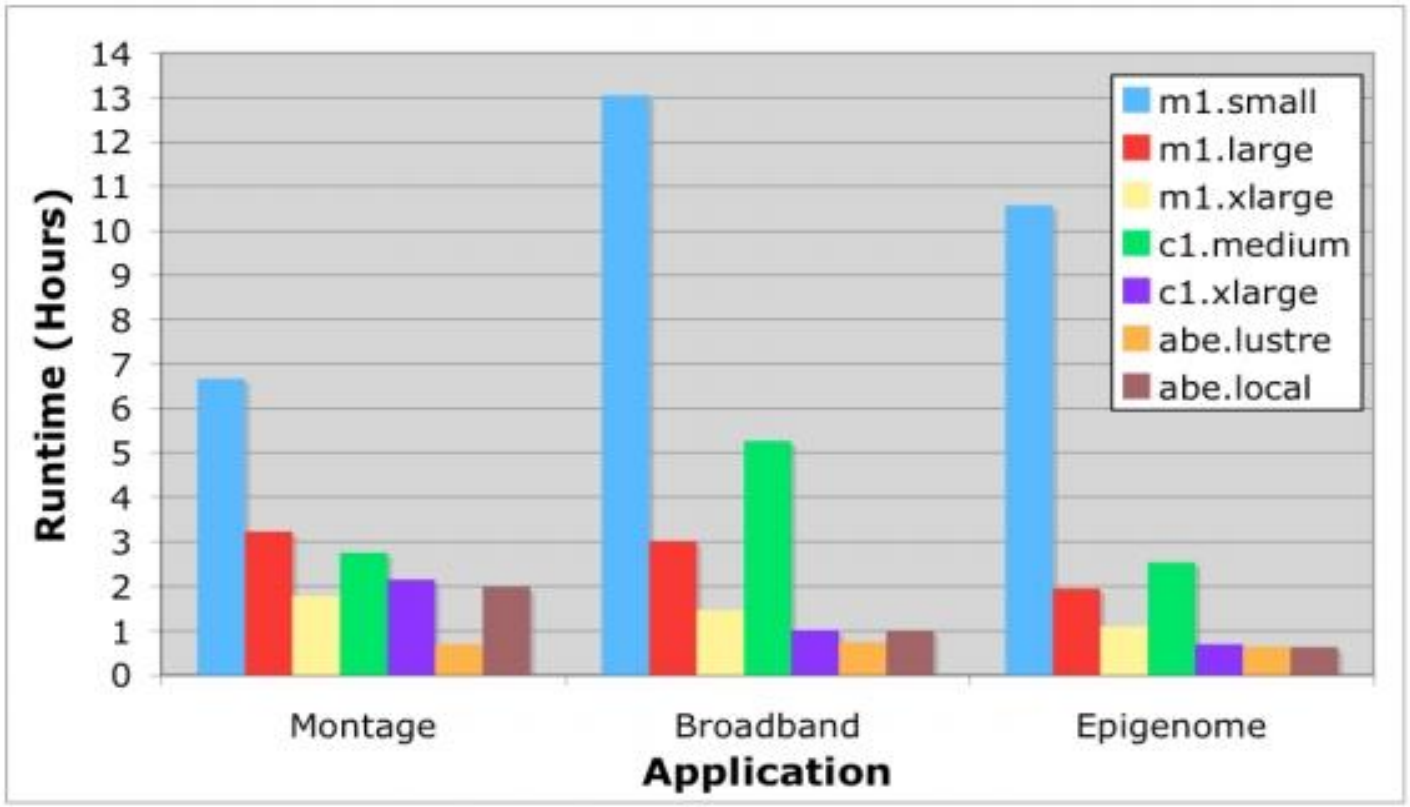
Runtime (Hours)
0
1
2
3
4
5
6
7
8
9
10
11
12
13
14
m1.small
m1.large
m1.xlarge
c1.medium
c1.xlarge
abe.lustre
abe.local
Montage
Broadband
Epigenome
Application

Performance vs Abe (cont.)

- Abe performed better, but EC2 used less powerful resources
  - Similar performance with EC2 c1.xlarge instance

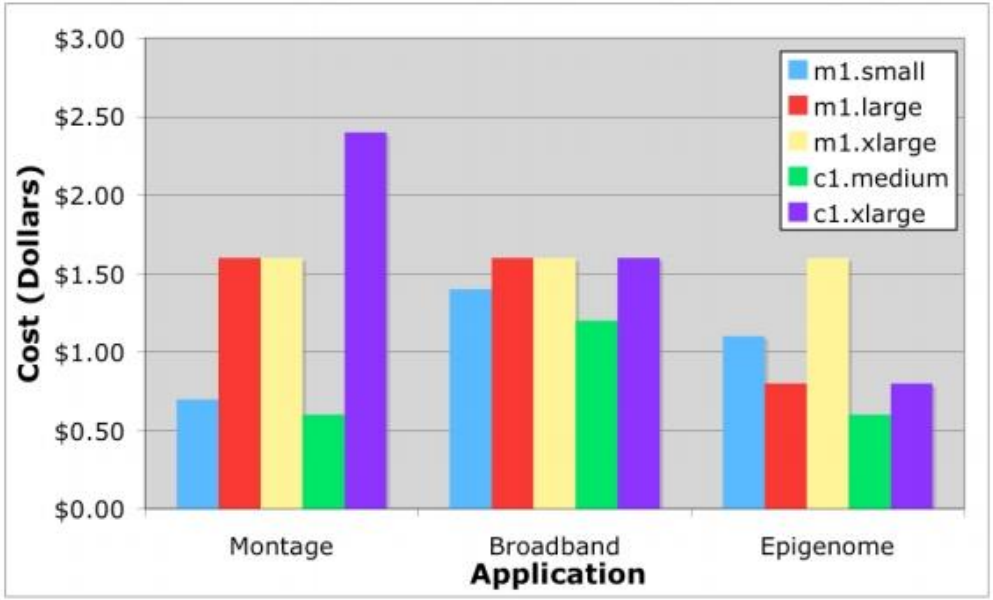

Cluster Compute Performance

- Weather forecasting model [5]
- EC2 Cluster Quad XL vs. Cyberstar
  - Cyberstar: 2 Quad core Xeon 2.67 GHz, 24 GB RAM
  - EC2: 2 Quad core Xeon 2.8 GHz, 23 GB RAM
- Cyberstar won
  - High EC2 virtualization overheads
  - 52% advantage with EC2 hyper-threading
  - EC2 Cost = ~$18

Benchmark	Cyberstar time	cc1.4xlarge time	Advantage
1 node, 8 processes per node	448 min	584 min	**30.3%**

Virtualization Overheads

- CAP3- DNA Sequencing [6]
- More VMs per node means more overheads

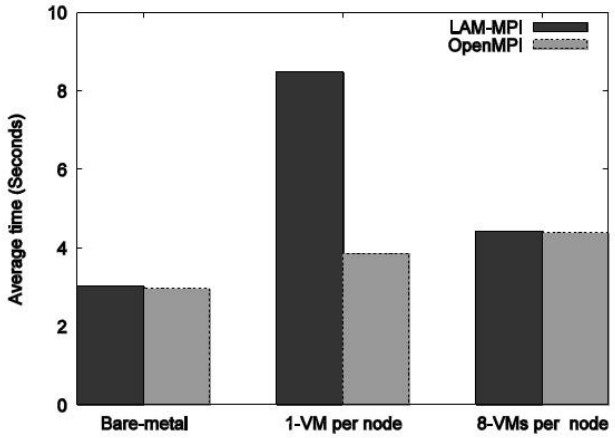

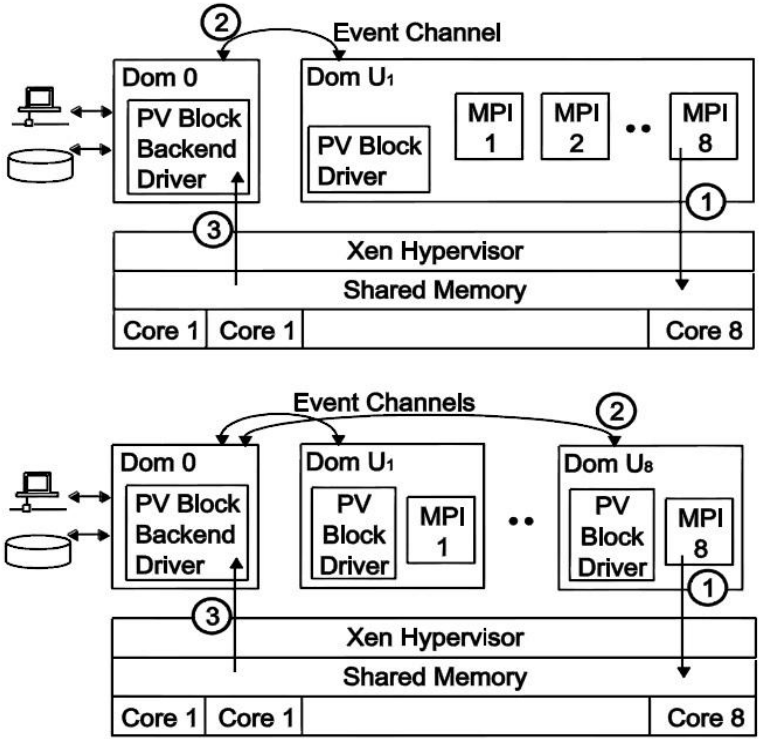

Conclusions

- Cloud computing is good for embarassingly parallel programs
- Performance degredation:
  - Virtualization overheads
  - Sub-optimal interconnects
- New EC2 Cluster instances are promising
- Biggest advantages are scalability and cost

The Future of Cloud Computing

- Large growth expected
  - $3 billion in 2011 to $13 billion in 2014
  - $4 billion IaaS
- Expand to small and medium sized companies [7]
- A power grid for computing [1]

Questions?

References

[1] I. Foster, Y. Zhao, I. Raicu, and S. Lu, *Cloud Computing and Grid Computing 360-Degree Compared,* http://arxiv.org/abs/0901.0131 [Accessed May 19, 2012].

[2] S. Biswas, "A History of Cloud Computing," February 9, 2011, http://www.cloudtweaks.com/2011/02/a-history-of-cloud-computing [Accessed May 19, 2012].

[3]Amazon Web Services, "Amazon Elastic Compute Cloud (Amazon EC2)," http://aws.amazon.com/ec2/ [Accessed May 18, 2012].

[4] G. Juve *et al*, *Scientific Workflow Applications on Amazon EC2,* www.isi.edu/~gideon/publications/JuveG-EC2.pdf [Accessed May 18, 2012].

[5] M. Fenn, J. Holmes, J. Nucciarone, *A Performance and Cost Analysis of the Amazon Elastic Compute Cloud (EC2) Cluster Compute Instance,* rcc.its.psu.edu/education/white_papers/cloud_report.pdf [Accessed May 20, 2012].

[6] J. Ekanyake and G. Fox, *High Performance Parallel Computing with Clouds and Cloud Technologies,* http://grids.ucs.indiana.edu/ptliupages/publications/cloud_handbook_final-with-diagrams.pdf [Accessed 18 May 2012].

[7] S. Biswas, "Where Is Cloud Computing Going? Up, Up And Away!" January 25, 2011, http://www.cloudtweaks.com/2011/01/where-is-cloud-computing-going-up-up-and-away [Accessed May 19, 2012].