Clouds and Panda

Sergey Panitkin

BNL

ATLAS

May 19, 2011

BROOKHAVEN NATIONAL LABORATORY

Outline

- Introduction
- EC2 based activities
- Magellan activities
- Summary

Introduction

- There are several reasons why cloud computing is interesting for ATLAS and other HEP experiments.
- Virtualization technologies as a platform for distributed computing infrastructure in ATLAS.
- Use of commercial and academic clouds as an additional resource for existing Grid based computing infrastructure.
- In a past two years a number of people from Panda group were involved in the studies of virtualization technologies and cloud computing.
- I will not discuss in this talk studies of performance of ATLAS analysis workloads in virtual environments.

Panda in Amazon EC2

- Cloud related activity was started in 2009 during Panda server migration from BNL to CERN
- Learned how to build multiple VM images for EC2 – CernVM based, Scientific Linux, RH, CentOS, etc
  - Learned VM image building technology
  - Tried different type of builds: rBuilder, Modify running VM, build by hand from scratch
  - Tried to build VM based on CernVM
- Prototype AMIs based on SL44 were build and uploaded to EC2
- Panda server on EC2 was built in May 2009 and ran on EC2 until recently
  - We did not go further than simple tests, since migration to CERN was very successful
  - No problems with EC2 stability , server ran for 2 years with no problems

Panda in the clouds

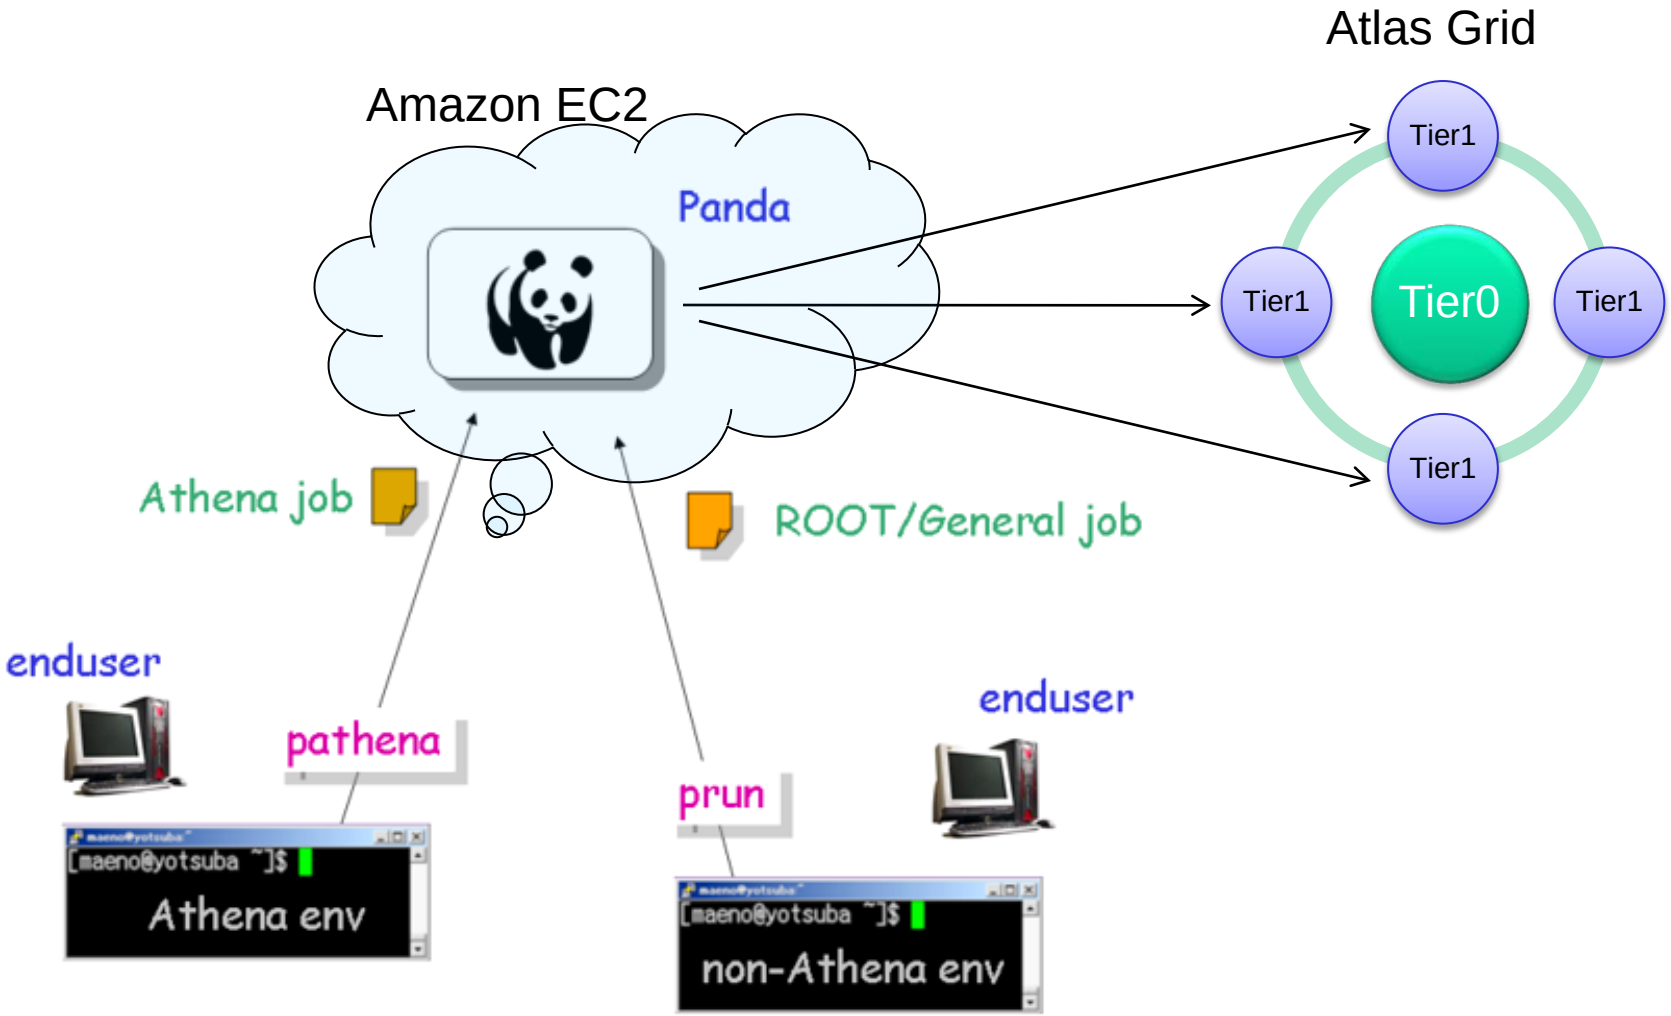

Amazon S3 Storage

- We (Torre) also used Amazon S3 storage for Panda Monitor data ~560GB as of May 2011.
  - That was used as data source by several projects (CassandraDB, user analysis monitoring,etc)
  - Download speed was satisfactory for the needs (maxed out destination's 1Gb network at BNL)
- It would be good to test IO performance using these data

EC2 Costs

Part of actual EC2 Bill for Panda related activities Feb. 2011

		Totals
Amazon Elastic Compute Cloud		
US East (Northern Virginia) Region		
Amazon EC2 running Linux/UNIX		
$0.085 per Small Instance (m1.small) instance-hour (or partial hour)	579 Hrs	49.22
Elastic IP Addresses		
$0.01 per non-attached Elastic IP address per complete hour	91 Hrs	0.91
	»	50.13
Amazon Simple Storage Service		
US Standard Region		
$0.140 per GB - first 1 TB / month of storage used	280.142 GB-Mo	39.22
$0.01 per 1,000 PUT, COPY, POST, or LIST requests	3,856 Requests	0.04
$0.01 per 10,000 GET and all other requests	1,167 Requests	0.01
	»	39.27
Amazon Simple Notification Service		
	»	0.00
Amazon Virtual Private Cloud		
	»	0.00
AWS Data Transfer (excluding Amazon CloudFront)		
$0.100 per GB - data transfer in per month	151.372 GB	15.14
$0.000 per GB - first 1 GB of data transferred out per month	1.000 GB	0.00
$0.150 per GB - up to 10 TB / month data transfer out	317.427 GB	47.61
$0.010 per GB - regional data transfer - in/out/between EC2 Avail Zones or when using public/elastic IP addresses or ELB	0.000006 GB	0.00
		62.75
Bill Summary		
Usage charges and monthly recurring fees during this billing cycle (More Info)		$152.15

Panda server

Panda monitor data

Panda monitor data transfers

EC2 Costs

Part of actual EC2 Bill for Panda related activities Apr. 2011

		Totals
Amazon Elastic Compute Cloud		
US East (Northern Virginia) Region		
Elastic IP Addresses		
$0.01 per non-attached Elastic IP address per complete hour	744 Hrs	7.44
	»	7.44
Amazon Simple Storage Service		
US Standard Region		
$0.140 per GB - first 1 TB / month of storage used	387.825 GB-Mo	54.30
$0.01 per 1,000 PUT, COPY, POST, or LIST requests	4,759 Requests	0.05
$0.01 per 10,000 GET and all other requests	62,627 Requests	0.06
	»	54.41
Amazon Simple Notification Service		
	»	0.00
Amazon Virtual Private Cloud		
	»	0.00
AWS Data Transfer (excluding Amazon CloudFront)		
$0.100 per GB - data transfer in per month	129.790 GB	12.98
$0.000 per GB - first 1 GB of data transferred out per month	1.000 GB	0.00
$0.150 per GB - up to 10 TB / month data transfer out	2,533.798 GB	380.07
		393.05
Bill Summary		
Usage charges and monthly recurring fees during this billing cycle (More Info)		$454.90

Panda server IP address

Panda monitor data

Panda monitor data transfers

Magellan Cloud Platform

- Magellan is a US DOE sponsored project devoted to cloud computing research for scientific applications.
- R&D effort to establish a nationwide scientific mid-range distributed computing and data analysis testbed.
  - Part of the ARRA funding ~$32M
  - Two sites: LBNL and ANL
  - Hardware IBM iDataPlex dx360 M2 cluster
    - Space saving power efficient design
    - Half depth racks, water cooled doors, etc
    - One of the most power efficient systems in the world
  - Allegedly ~10 times cheaper to use than EC2 per CPU hour
    - http://magellan.nersc.gov/?p=878

Magellan cluster@ LBNL

IBM iDataplex chassis using 5,760 processor cores for a theoretical peak performance of 61.5 teraflop/s. QDR InfiniBand interconnect.

Magellan@ LBNL hardware

- Base compute nodes - 560 nodes
  - 2 quad-core Intel Nehalem 2.67 GHz processors per node
  - 8 cores per node (4,480 total cores)
  - 24 GB DDR3 1333 MHz memory per node
- Expanded compute nodes - 160 nodes
  - 2 quad-core Intel Nehalem 2.67 GHz processors per node
  - 8 cores per node (1,280 total cores)
  - 48 GB DDR3 1066 MHz memory per node
  - 1 TB (local) SATA disk per node
- 18 Login and service nodes

Science-oriented Features of Magellan@ NERSC

- Node aggregation into virtual clusters (as opposed to node virtualization into independent systems)
- Provisioning of full, virtual private clusters for individual research projects
- Dynamic provisioning of multiple software environments
- High bandwidth, low-latency interconnect (InfiniBand QDR)
- Global file system, shared with other NERSC systems
- Access NERSC's large tape archive for bulk storage of scientific data

Borrowed from Jeff Broughton talk on Magellan@NERSC

Test on Magellan

- We started working on Magellan@LBNL about 3 month ago
- Goals
  - Learn how to work with Eucalyptus interface on Magellan
  - Perform tests with Panda pilot in VM
  - Tests with Atlas software
- Built a VM with a pilot wrapper (by Jose Caballero) and pilot
- Tested communication with Panda server and ran simple tests ,where pilot in VM talked to Panda server and got payload and ran it.
  - [root@localhost pilot]# ./trivialWrapper.sh --site=TEST3 --queue=TEST3

Panda monitor
Times are in UTC

Panda info and help

Panda jobs

Jobs: 1199139091

Click for help

Jobs - search
States: running, defined, waiting, assigned, activated, finished, failed
Types: analysis, prod, install, test

PandaID, Owner, Working group	Job	Status	Created	Time to start	Duration	Ended/ Modified	VO/Site	Priority
1199139091 Jose Caballero OSG	jobsetID=10254 trivial_test 1298496773_1	finished	2011-02-23 21:32	2 days, 21:18:31	0:00:03	02-26 18:51	osg/TEST3	1000

Magellan@ANL

- Two weeks ago we got an allocation of 100,000 hours until October 2011 on Magellan cluster at ANL.
  - Project name: “ATLAS_DataAnalysis”
- Mostly similar in hardware to the West Coast Magellan cluster with a few notable differences
  - High speed active storage nodes (SSD)
  - GPU servers
  - Open Cloud interface

Magellan@ANL

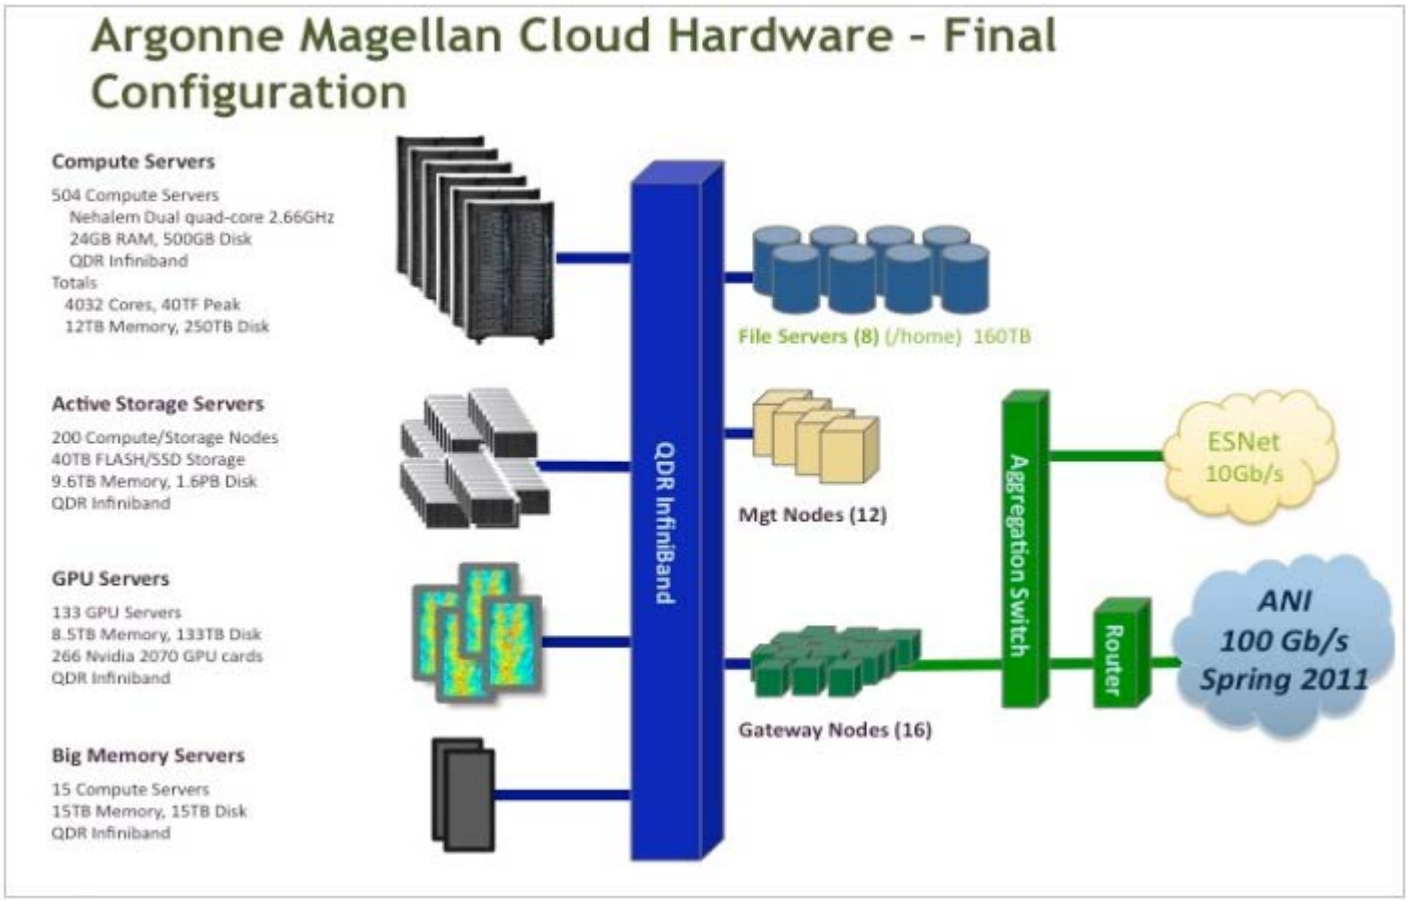

Summary

- Integration of cloud computing technologies with ATLAS distributed computing system and in particular with Panda work load management system is an interesting capability that should be explored
- We have ported Panda server to Amazon EC2.
- Amazon S3 is used for storage of Panda monitor data
- Performed studies of performance of ATLAS payloads in virtual environment
- Started tests on DOE's Magellan cloud platform.
- Looking forward to participating in the ATLAS Cloud R&D project!

Acknowledgements

- Many thanks to people who contributed to these activities at different times Predrag Buncic, Jose Caballero, Kaushik De, Tadashi Maeno, Maxim Potekhin, Torre Wenaus, Yushu Yao

The End