

Eucalyptus Overview

Lavanya Ramakrishnan
Iwona Sakrejda
Shane Canon

Lawrence Berkeley National Lab

June 2011

What are we going to do today?

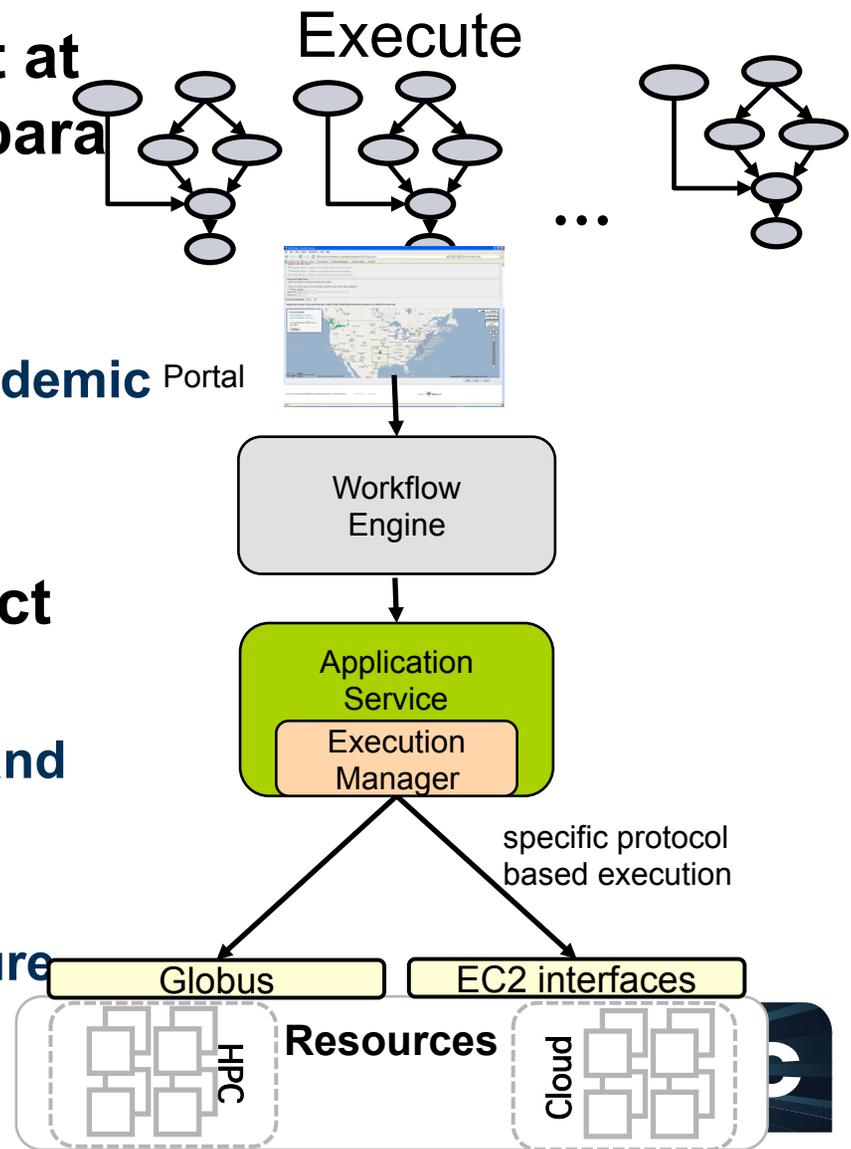
- **Eucalyptus Overview**
- **Demo**
- **Q&A**
- **Preparation for Hands-On**
 - **Sign NERSC User Agreement**
 - **Apply for Eucalyptus accounts on Magellan**

What is Eucalyptus?

- **Elastic Utility Computing Architecture Linking Your Programs To Useful Systems**
- **Open source IaaS implementation**
 - provides a way to provision virtual machines
- **API compatible with Amazon AWS**
- **Multiple Linux distributions**

History

- **Started as a research project at Univ of California Santa Barbara**
- **Scheduling LEAD (weather forecasting workflows)**
 - **controlled environments in academic clusters**
- **Virtual Grid Application Development Software Project (VGrADS)**
 - **novel programming language and runtime system techniques to manage grid workflows**
 - **experiment if same infrastructure works for public clouds**



Timeline

Coding starts	Feb 2008
First release of EC2 interface	May 2008
** Demonstration at Supercomputing	Nov 2008
S3 support	Dec 2008
Commercialization efforts	Jan 2009
Ubuntu Enterprise Cloud (powered by Eucalyptus)	April 2009
Enterprise Edition 2.0	June 2010
Eucalyptus 2.0	Aug 2010
Eucalyptus 3.0 (Expected)	Summer 2011

Distributions and Hypervisors

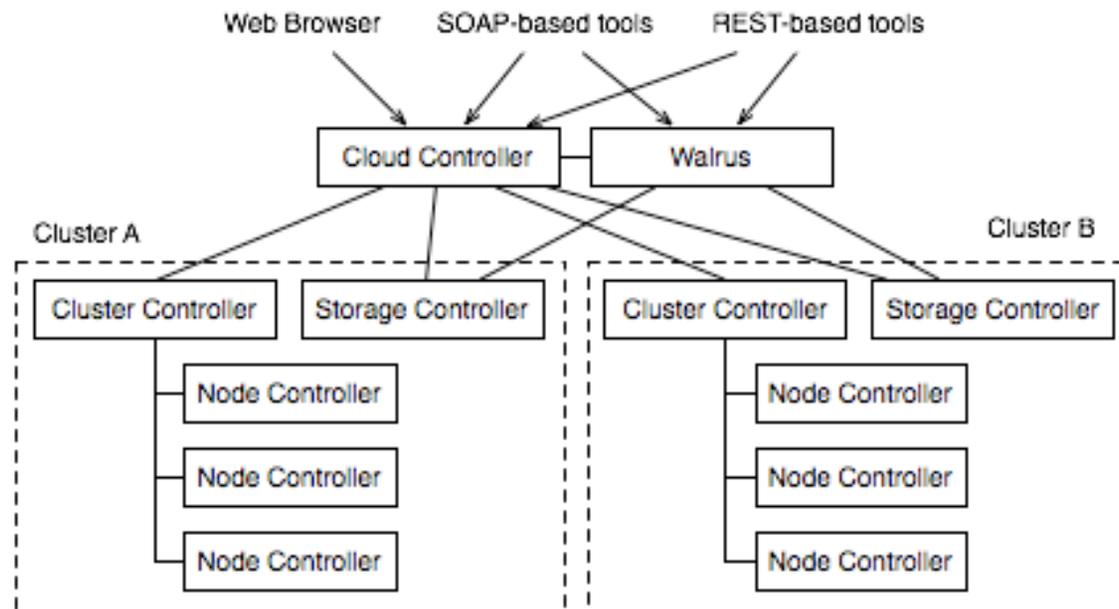
- **Supported on multiple distributions of Linux**
 - including commercial Linux distributions: **Red Hat Enterprise Linux (RHEL) and SUSE**
- **Ubuntu Enterprise Cloud**
- **Supports multiple hypervisors**
 - **KVM, Xen, VMWare, etc**

Comparison Chart

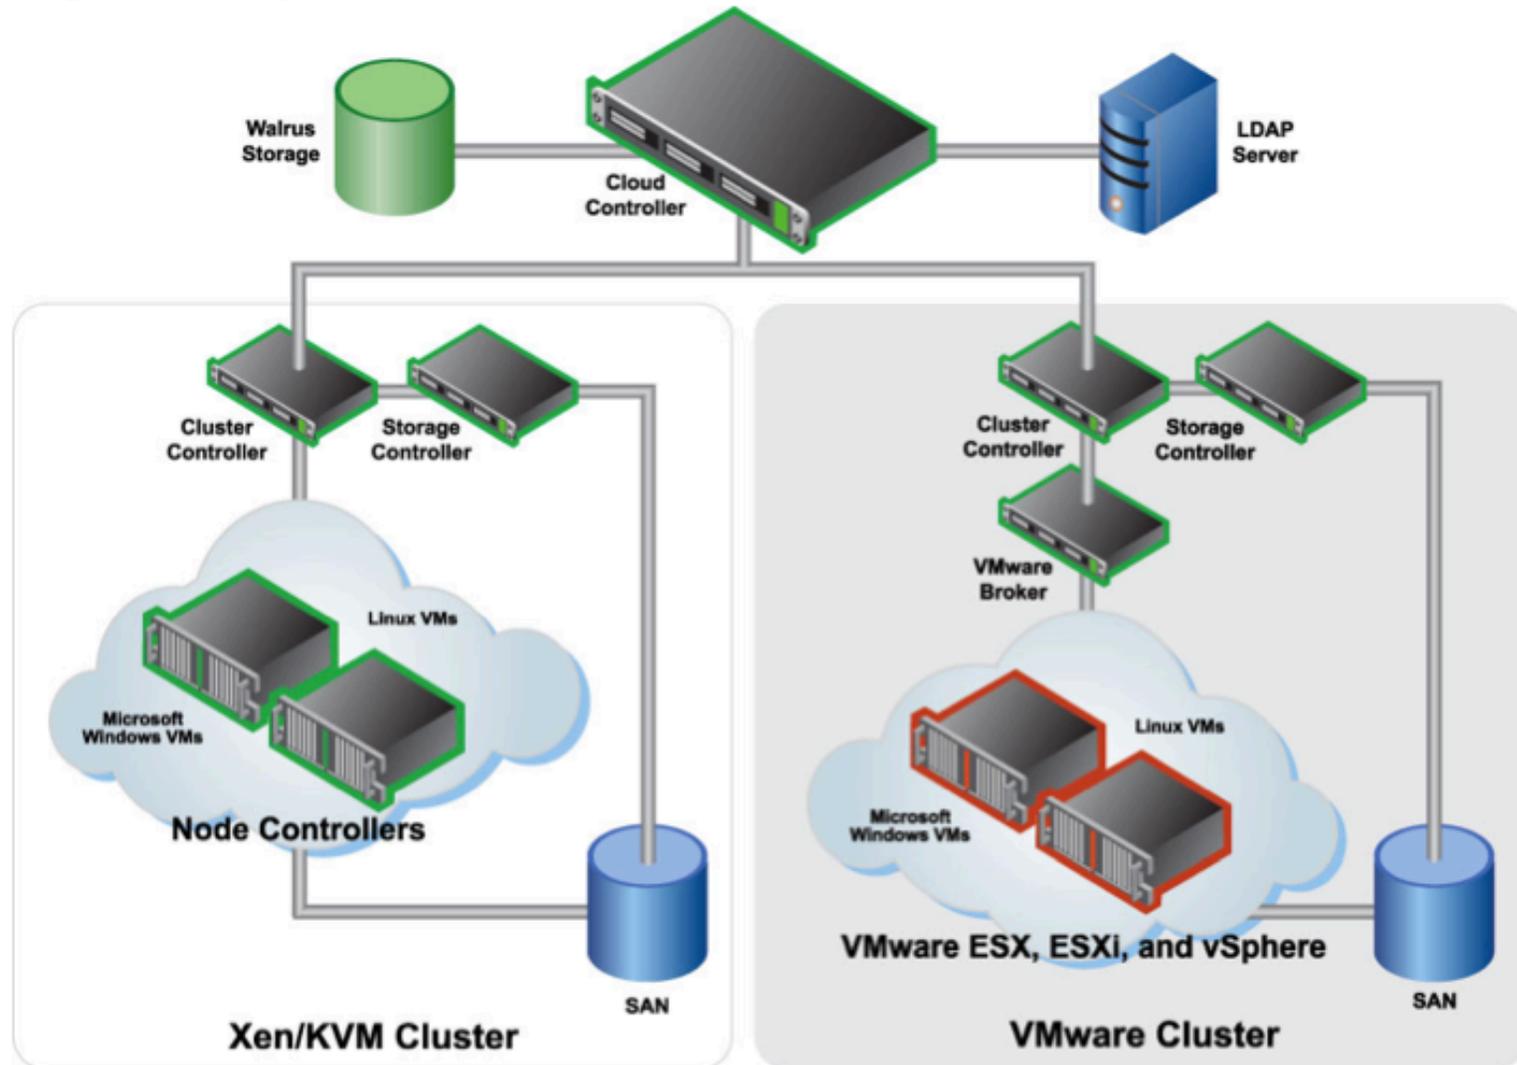
Eucalyptus Feature Comparison	Open Source	Enterprise Edition
Amazon AWS Interface Compatibility	✓	✓
Flexible Clustering and Availability Zones	✓	✓
Network Management, Security Groups, Traffic Isolation	✓	✓
Cloud Semantics and Self-Service Capability	✓	✓
Bucket-Based Storage Abstraction (S3-Compatible)	✓	✓
Block-Based Storage Abstraction (EBS-Compatible)	✓	✓
Xen and KVM Hypervisor Support	✓	✓
VMware Hypervisor Support		✓
Virtual-to-Virtual Image Conversion for VMware		✓
Microsoft Windows Guest Support		✓
Direct SAN Integration		✓
Quota Management and Accounting		✓
User and Group Access Management		✓
High-Performance MySQL Database Backend		✓



Open Source Architecture



Enterprise Architecture



Storage

- **Local File System on the VM**
 - ephemeral
- **Block Store**
 - similar to Amazon EBS
 - only mounted on single VM
 - persistent storage across lifetime of VMs
- **Walrus**
 - similar to S3
 - object store for large objects

Network Configuration

- **Simple Modes**
 - security groups, dynamic assignment of IPs, isolation of network traffic, meta-data service are not available
 - SYSTEM
 - STATIC
- **Advanced**
 - MANAGED
 - MANAGED-NOVLAN

Network Configuration: SYSTEM

- **Assigns MAC address to VM instance**
- **Attaches VM's ethernet to physical ethernet**
- **Uses DHCP to get an IP address**
 - **must be setup DHCP server**

Network Configuration: STATIC

- **Control over VM IP address assignment**
 - map of **MAC addresses/IP address**
- **Eucalyptus controlled DHCP server**
- **Network attachment is managed similar to SYSTEM**

Network Configuration: MANAGED

- **Manages a large pool of (usually) private unroutable IP addresses**
- **Runs its own DHCP server**
- **Also can specify pool of public IPs**
- **Implements security groups**
 - **users can specify group a VM must be in**
 - **users can specify rules of access**

Network Configuration: MANAGED-NOVLAN

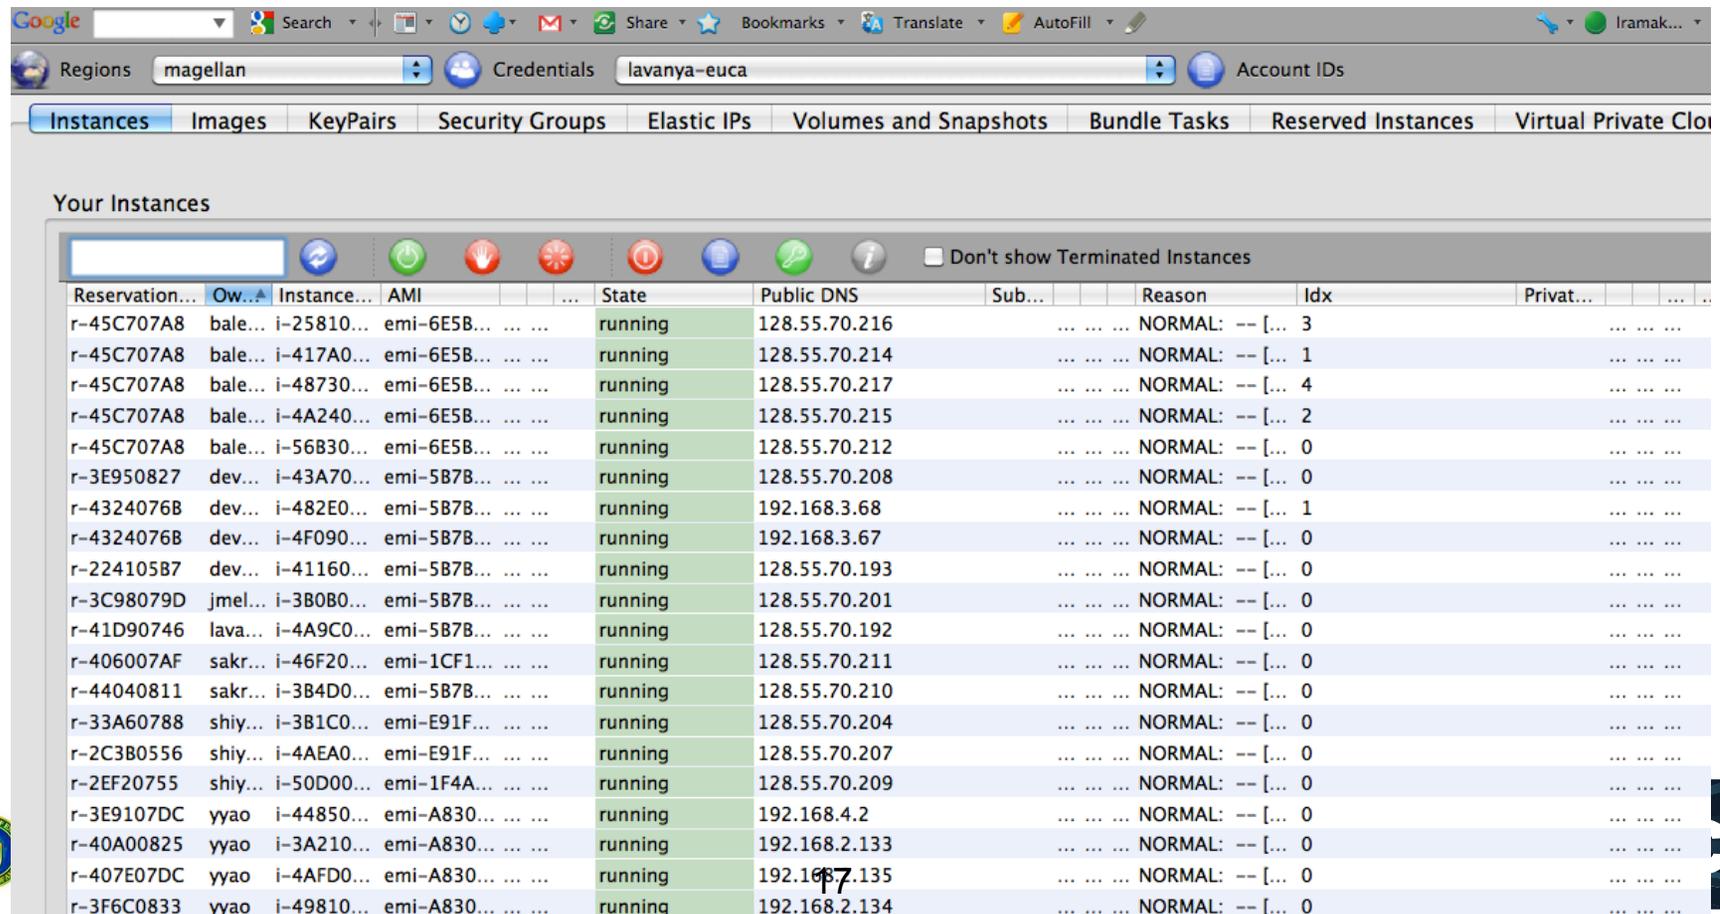
- **Similar to MANAGED**
- **No network isolation**

Security

- **WS-security for authentication**
 - **Encryption of inter-component communication is not enabled by default**
 - Configuration option
- **ssh key generation and installation**
 - **cloud controller generates the public/private key pairs and installs them**
- **User sign-up is web based**

Client Tools

- **Command-line tools**
 - ec2-api-tools, euca-tools
 - S3 tools
- **Browser-based tools**
 - HybridFox



The screenshot shows the AWS Management Console interface. At the top, there are navigation tabs for 'Regions' (set to 'magellan'), 'Credentials' (set to 'lavanya-euca'), and 'Account IDs'. Below these are tabs for various AWS services: 'Instances', 'Images', 'KeyPairs', 'Security Groups', 'Elastic IPs', 'Volumes and Snapshots', 'Bundle Tasks', 'Reserved Instances', and 'Virtual Private Cloud'. The main content area is titled 'Your Instances' and displays a table of EC2 instances. The table has columns for Reservation ID, Owner, Instance ID, AMI, State, Public DNS, Subnet, Reason, and Index. All instances shown are in a 'running' state. A search bar and a 'Don't show Terminated Instances' checkbox are located above the table.

Reservation...	Own...	Instance...	AMI	State	Public DNS	Sub...	Reason	Idx	Privat...
r-45C707A8	bale...	i-25810...	emi-6E5B...	running	128.55.70.216		NORMAL: -- [... 3		
r-45C707A8	bale...	i-417A0...	emi-6E5B...	running	128.55.70.214		NORMAL: -- [... 1		
r-45C707A8	bale...	i-48730...	emi-6E5B...	running	128.55.70.217		NORMAL: -- [... 4		
r-45C707A8	bale...	i-4A240...	emi-6E5B...	running	128.55.70.215		NORMAL: -- [... 2		
r-45C707A8	bale...	i-56B30...	emi-6E5B...	running	128.55.70.212		NORMAL: -- [... 0		
r-3E950827	dev...	i-43A70...	emi-5B7B...	running	128.55.70.208		NORMAL: -- [... 0		
r-4324076B	dev...	i-482E0...	emi-5B7B...	running	192.168.3.68		NORMAL: -- [... 1		
r-4324076B	dev...	i-4F090...	emi-5B7B...	running	192.168.3.67		NORMAL: -- [... 0		
r-224105B7	dev...	i-41160...	emi-5B7B...	running	128.55.70.193		NORMAL: -- [... 0		
r-3C98079D	jmel...	i-380B0...	emi-5B7B...	running	128.55.70.201		NORMAL: -- [... 0		
r-41D90746	lava...	i-4A9C0...	emi-5B7B...	running	128.55.70.192		NORMAL: -- [... 0		
r-406007AF	sakr...	i-46F20...	emi-1CF1...	running	128.55.70.211		NORMAL: -- [... 0		
r-44040811	sakr...	i-3B4D0...	emi-5B7B...	running	128.55.70.210		NORMAL: -- [... 0		
r-33A60788	shiy...	i-3B1C0...	emi-E91F...	running	128.55.70.204		NORMAL: -- [... 0		
r-2C3B0556	shiy...	i-4AEA0...	emi-E91F...	running	128.55.70.207		NORMAL: -- [... 0		
r-2EF20755	shiy...	i-50D00...	emi-1F4A...	running	128.55.70.209		NORMAL: -- [... 0		
r-3E9107DC	yyao	i-44850...	emi-A830...	running	192.168.4.2		NORMAL: -- [... 0		
r-40A00825	yyao	i-3A210...	emi-A830...	running	192.168.2.133		NORMAL: -- [... 0		
r-407E07DC	yyao	i-4AFD0...	emi-A830...	running	192.168.2.135		NORMAL: -- [... 0		
r-3F6C0833	yyao	i-49810...	emi-A830...	running	192.168.2.134		NORMAL: -- [... 0		



What can I do with Eucalyptus?

- **Setup groups and rules**
- **Start and stop VMs**
- **Upload files to S3**
- **Create custom images**

Elastic IPs

- **Static IP addresses allocated to an account**
- **Dynamically associate with an instance**
- **Mask instance or availability zone failures by remapping to another instance**

Security Groups

- **Sets of networking rules applied to a group.**
- **Users specify ingress rules**
 - e.g., ping (ICMP) or SSH (TCP, port 22).
- **“default” security group denies incoming network traffic from all sources**

Custom Images

- **Bundle images**
 - could start from a physical host
 - easier to start from an existing VM instance
 - identify corresponding kernel and ramdisk
 - need admin privileges for registering these
- **Upload images**
- **Register images**

Instance Metadata

- Instance specific metadata and user-specified metadata
- Query a web server using a query API
 - <http://169.254.169.254/2008-08-08/>
 - 2008-08-08 is API version
- User-specified launch data is not encrypted
- E.g., instance id, public IP, private IP, etc

Magellan at NERSC

Scientific Discovery through Cloud Computing

Magellan

In the News

Science

Research

Using Magellan

About NERSC

› Magellan

› In the News

› Science

› Research

› Using Magellan

› Submitting Batch Jobs

› Using Hadoop

› Using Eucalyptus

› Eucalyptus Q&A

› Creating a SOCKS Proxy

Using Eucalyptus Getting Started

Eucalyptus is an open-source implementation of Amazon's popular EC2 cloud platform. Eucalyptus' interfaces are designed to replicate the APIs used on EC2. This includes implementing many of the capabilities of EC2 including Elastic Block, S3, Elastic IPs, etc.

This is a step-by-step guide for getting started using Eucalyptus on Magellan at NERSC. (Consult the [Eucalyptus](#) and [EC2](#) web sites for more in-depth information).

<http://magellan.nersc.gov>

Image Types on Magellan

Name	CPUs	Memory (MB)	Disk (GB)
m1.small	1	256	10
c1.medium	2	256	10
m1.large	2	512	10
m1.xlarge	4	10240	20
c1.xlarge	8	19456	140

Eucalyptus Configuration on Magellan

- **Space reserved for snapshots (GB): 500**
- **Maximum buckets per user:15**
- **Maximum bucket size (MB):200GB**
- **Space reserved for unbundling images (MB): 307GB**
- **5 public IPs per user**
- **5TB reserved for volumes**
- **Max volume size 150GB**

Virtual Cluster Scripts

- **Creates a virtual cluster on top of Eucalyptus**
 - Base setup uses EBS on head node
 - Serves EBS volume using NFS
 - Assigns public IP to head node and uses it as a proxy
 - Event-driven custom scripts can be added
- **Other options without EBS, without master, with EBS on every host is also available**

User Configuration

CLUSTER_TYPE=<nfs,torque>

EBS_VOLUME_ID

Cluster Configuration

CLUSTER=mytorquecluster
IMAGE_ID=emi-1D1A15BA
KERNEL_ID=eki-A86F17CD
RAMDISK_ID=eri-1062190B
IMAGE_USER=root
AVAILABILITY_ZONE=euca
EBS_ATTACH_POINT=/dev/vdb
APPLICATION_SPACE=/apps/
INSTANCE_TYPE=m1.small

Cluster Events

**LOCAL_SCRIPT=/global/common/carver/tig/
virtualcluster/0.1/configure-centos-nfs**

**MASTER_STARTUP=configure-master-
centos-nfs,configure-swap**

**WORKER_STARTUP_FROM_MASTER=confi
gure-slave-centos-nfs-from-master**

**MASTER_REGISTERWORKER_PRE=registe
r-worker-pre**

MASTER_REGISTERWORKER_POST=

**#MASTER_DEREGISTERWORKER=<not
implemented yet>**

Commands

```
$ module load tig virtualcluster
```

```
$ source .cloud/nersc/eucarc
```

```
$ export CLUSTER_CONF=<cluster-conf>
```

```
$ vc-launcher newCluster <noNodes>
```

```
$ vc-launcher addNodes <noNodes>
```

```
$ vc-launcher terminateCluster
```

Terminology

- **Image/Disk image**
- **VM/Virtual Machine**
- **Instance**
- **Block Store/EBS/volume**
- **S3/Walrus**
- **Instance Type**
- **Virtual Cluster**



Demo

Acknowledgements

- **US Department of Energy DE-AC02-05CH11232**
- **Eucalyptus Website**
 - for diagrams!
- **Magellan**
 - Tina Declerck, Scott Campbell, Brent Draney

Questions?

Lavanya Ramakrishnan
LRamakrishnan@lbl.gov



U.S. DEPARTMENT OF
ENERGY

Office of
Science

34



NERSC User Agreement

I have read the NERSC Policies and Procedures and understand my responsibilities in the use of NERSC resources.

Answers to all of the entries below are required, including NERSC Principal Investigator.

Signature:

Print Name:

Citizenship:

Organization:

Email Address:

Work Phone Number:

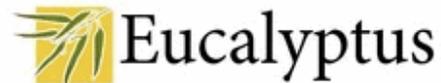
Name of NERSC Principal Investigator for one
of your NERSC project accounts (repositories):

Date:

Please sign the forms and return to us to
receive your login and password

Eucalyptus Account and Credentials

In your browsers go to <https://mageuca.ner-sc.gov:8443/#login>



Version 2.0.1

Please, sign in:

A yellow rectangular login form with the following fields and elements:

- Username: [input field]
- Password: [input field]
- Remember me on this computer
- Sign in button



[Apply](#) for account | [Recover](#) the Password

You will be asked for:

Username: **train<n>**

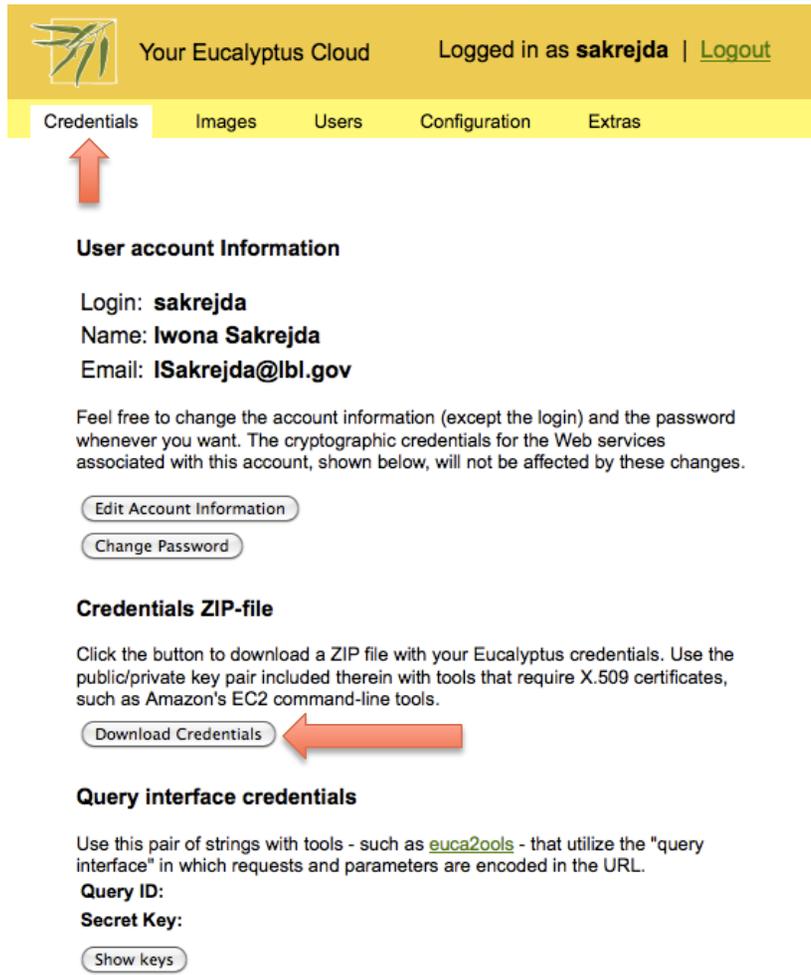
Password: **makeUpYourOwn**

Password, again:

Full Name: **Your Name**

Email address: **for notification**

Credentials



Your Eucalyptus Cloud Logged in as sakrejda | Logout

Credentials Images Users Configuration Extras

User account Information

Login: sakrejda
Name: Iwona Sakrejda
Email: ISakrejda@lbl.gov

Feel free to change the account information (except the login) and the password whenever you want. The cryptographic credentials for the Web services associated with this account, shown below, will not be affected by these changes.

Edit Account Information
Change Password

Credentials ZIP-file

Click the button to download a ZIP file with your Eucalyptus credentials. Use the public/private key pair included therein with tools that require X.509 certificates, such as Amazon's EC2 command-line tools.

Download Credentials

Query interface credentials

Use this pair of strings with tools - such as [euca2ools](#) - that utilize the "query interface" in which requests and parameters are encoded in the URL.

Query ID:
Secret Key:
Show keys

- Go to <https://mageuca.nersc.gov:8443>
- Select “Credentials” from the top bar
- Click on “Download Credentials”
- scp zip file with credentials to carver.nersc.gov
`scp euca2-$USER-x509(3).zip carver.nersc.gov:~/.`
- ssh to carver.nersc.gov
`mkdir ~/.euca`
`mv euca2-$USER-x509(3).zip .euca/.`
`cd ~/.euca`
`Unzip euca2-$USER-x509(3).zip .euca/.`
`chmod 0700 ~/.euca`
`chmod 0600 ~/.euca/*`
- Set up your environment and look at what's new
`source ~/.euca/eucarc`
`printenv`
- Create ssh keys to access your VMs
`cd ~/.euca`
`euca-add-keypair $USER-euca > $USER-euca.private`
`cat $USER-euca.private`