

OSDC PIRE Summer Internship, Sao Paulo, Brazil: Integration of Remote Clusters Under the Same Cloud

Felipe Navarro

Florida International University

Computer Engineering Master's Student

Cloud Computing Engineer for C.I.A.R.A.



Overview

- ❑ Goals and Purpose
- ❑ Component description
- ❑ Implementation Concept
- ❑ Understanding Underlying Hardware and Network Setup
- ❑ Operating System Choice and Installation
- ❑ Eucalyptus Installation and Configuration
- ❑ Testing the Cloud's Operation
- ❑ Issues Encountered
- ❑ Proposed Solutions and Trouble Shooting

Goals and Purpose

- ❑ Implement a Eucalyptus cluster in Brazil as a starting point for the implementation of multiple geographically distributed clusters for the OSDC network
 - ❑ To encourage scientists and researchers from all over the world to be a part of the OSDC community by providing resources at key locations around the globe
- ❑ Integrate the Eucalyptus cluster in Brazil as part of the Eucalyptus cloud in Miami so that they function as a single cloud with distributed clusters
 - ❑ Users would work with clusters closest to their respective geographic locations, but would have access to other clusters within the cloud should additional resources be required
- ❑ Document the process of implementing multiple remote clusters under a single cloud
 - ❑ Increase the ease and efficiency with which future cluster implementations can be completed

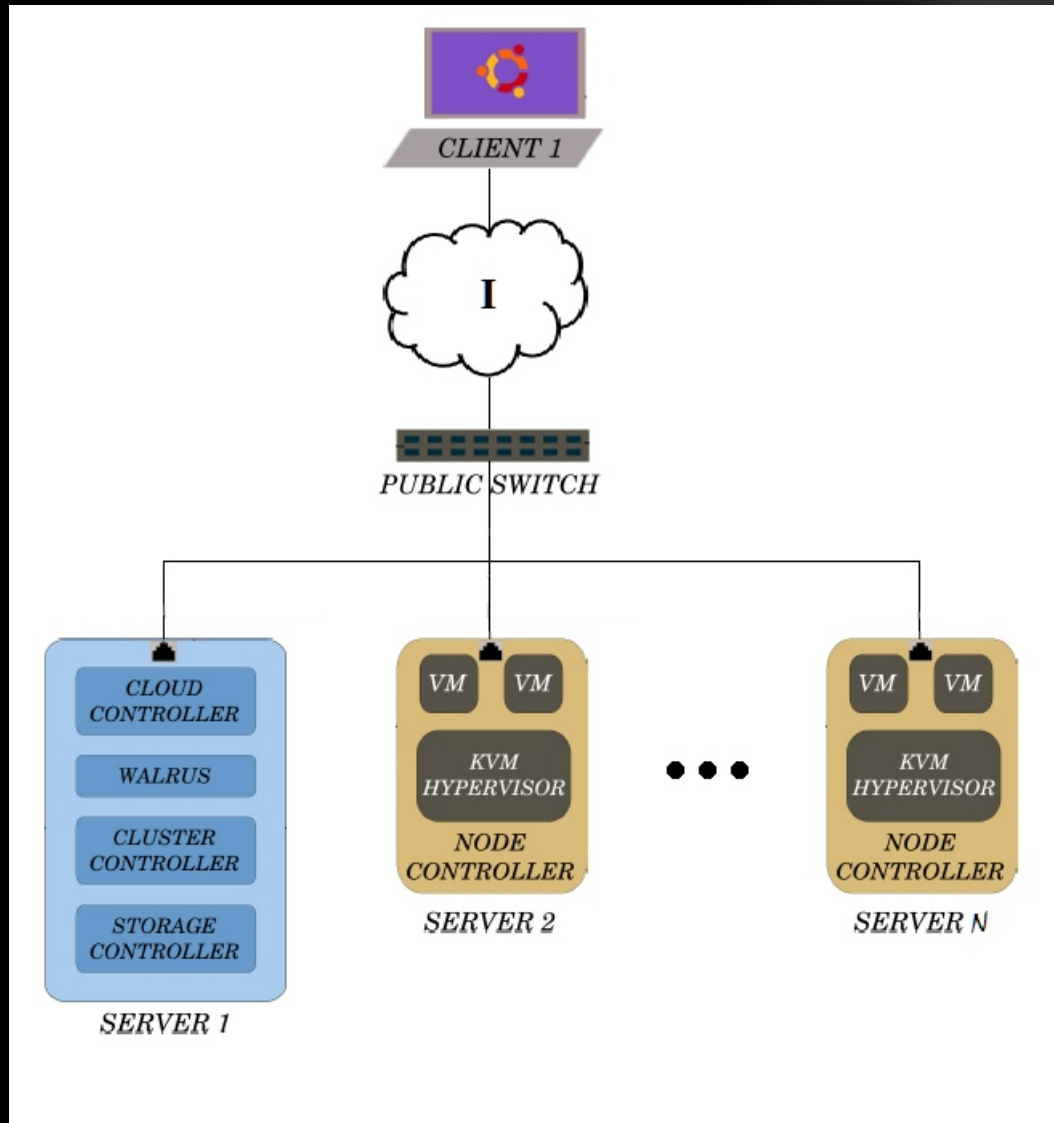
Front-end Component Functions

- ❑ Cloud Controller (CLC):
 - ❑ Monitor availability of resources
 - ❑ Resource arbitration
 - ❑ Monitor running instances
- ❑ Cluster controller (CC):
 - ❑ Received requests to run instances from CLC
 - ❑ Determine which NCs to run instances on
 - ❑ Control the virtual network used by instances
 - ❑ Report NC information to the CLC
- ❑ Node Controller (NC)
 - ❑ Collect resource data and report to CC
 - ❑ Instance life cycle management

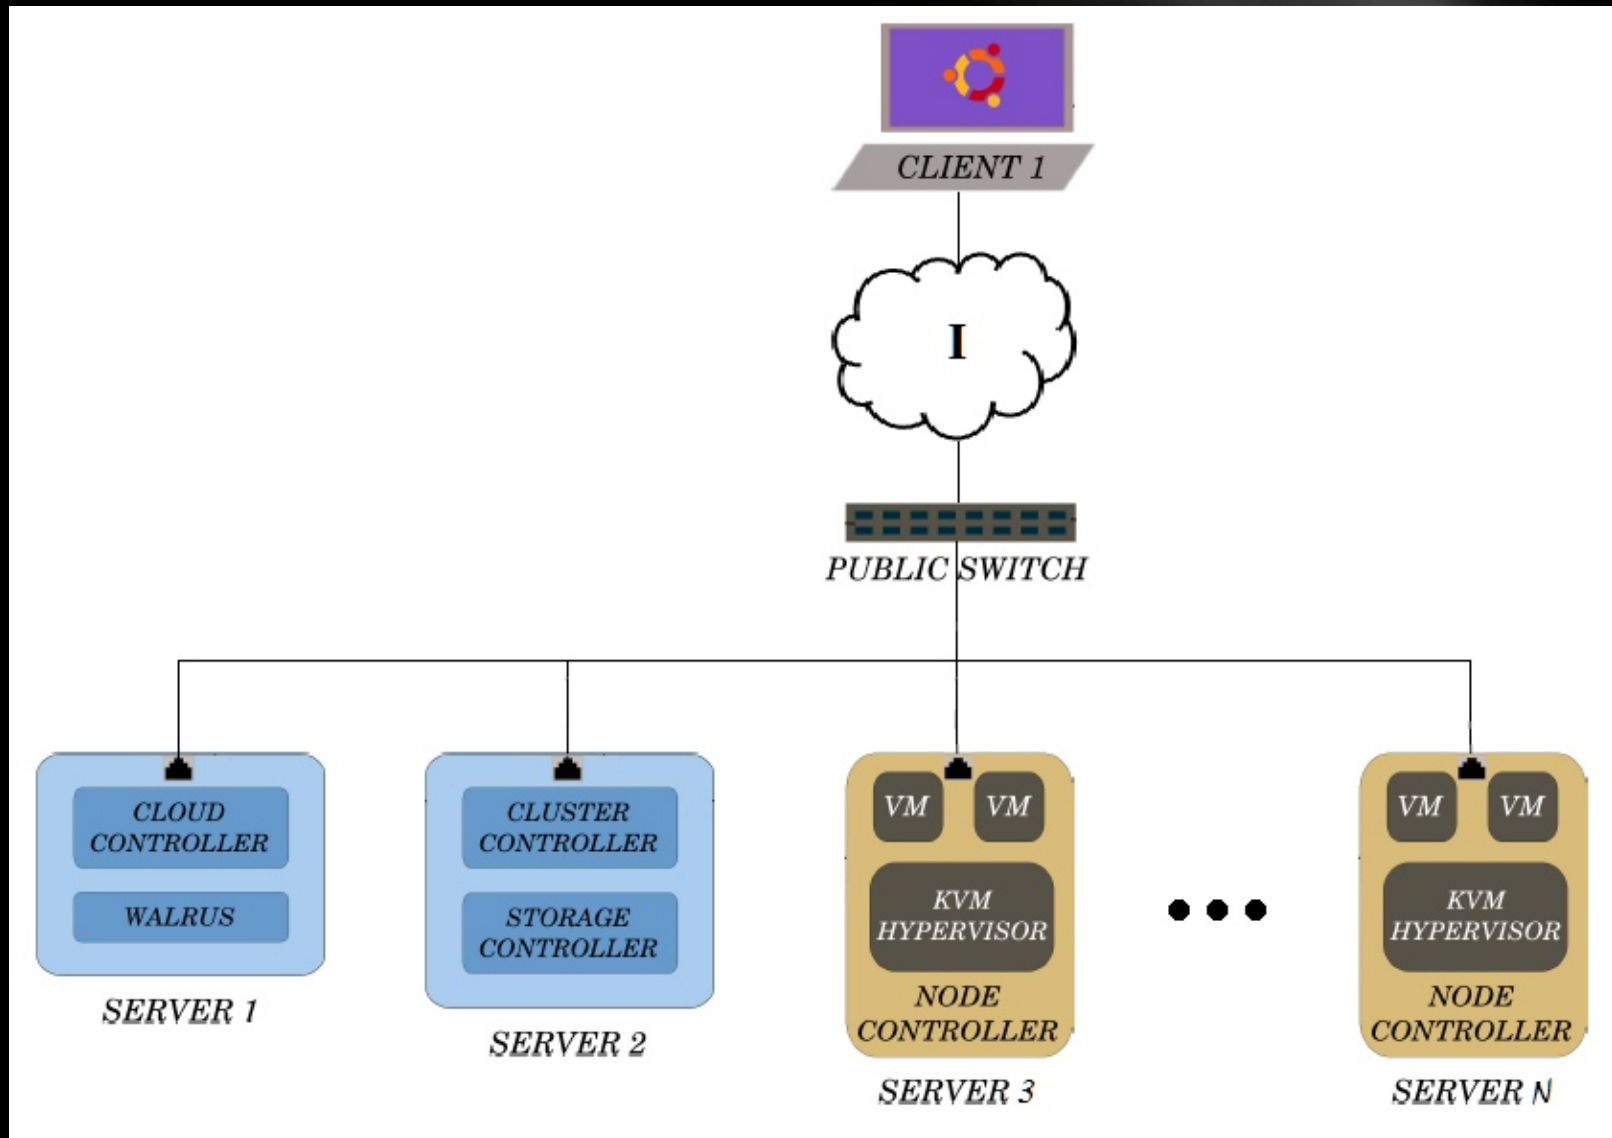
Front-end Component Functions

- ❑ Storage Controller (SC)
- ❑ Creation of persistent Elastic Block Storage devices (Amazon EBS)
- ❑ Creation of snapshot volumes
- ❑ Provide block storage to instances
- ❑ Walrus Storage Controller (WS₃)
 - ❑ Store machine images/snapshots
 - ❑ Storing and serving files using S₃ API

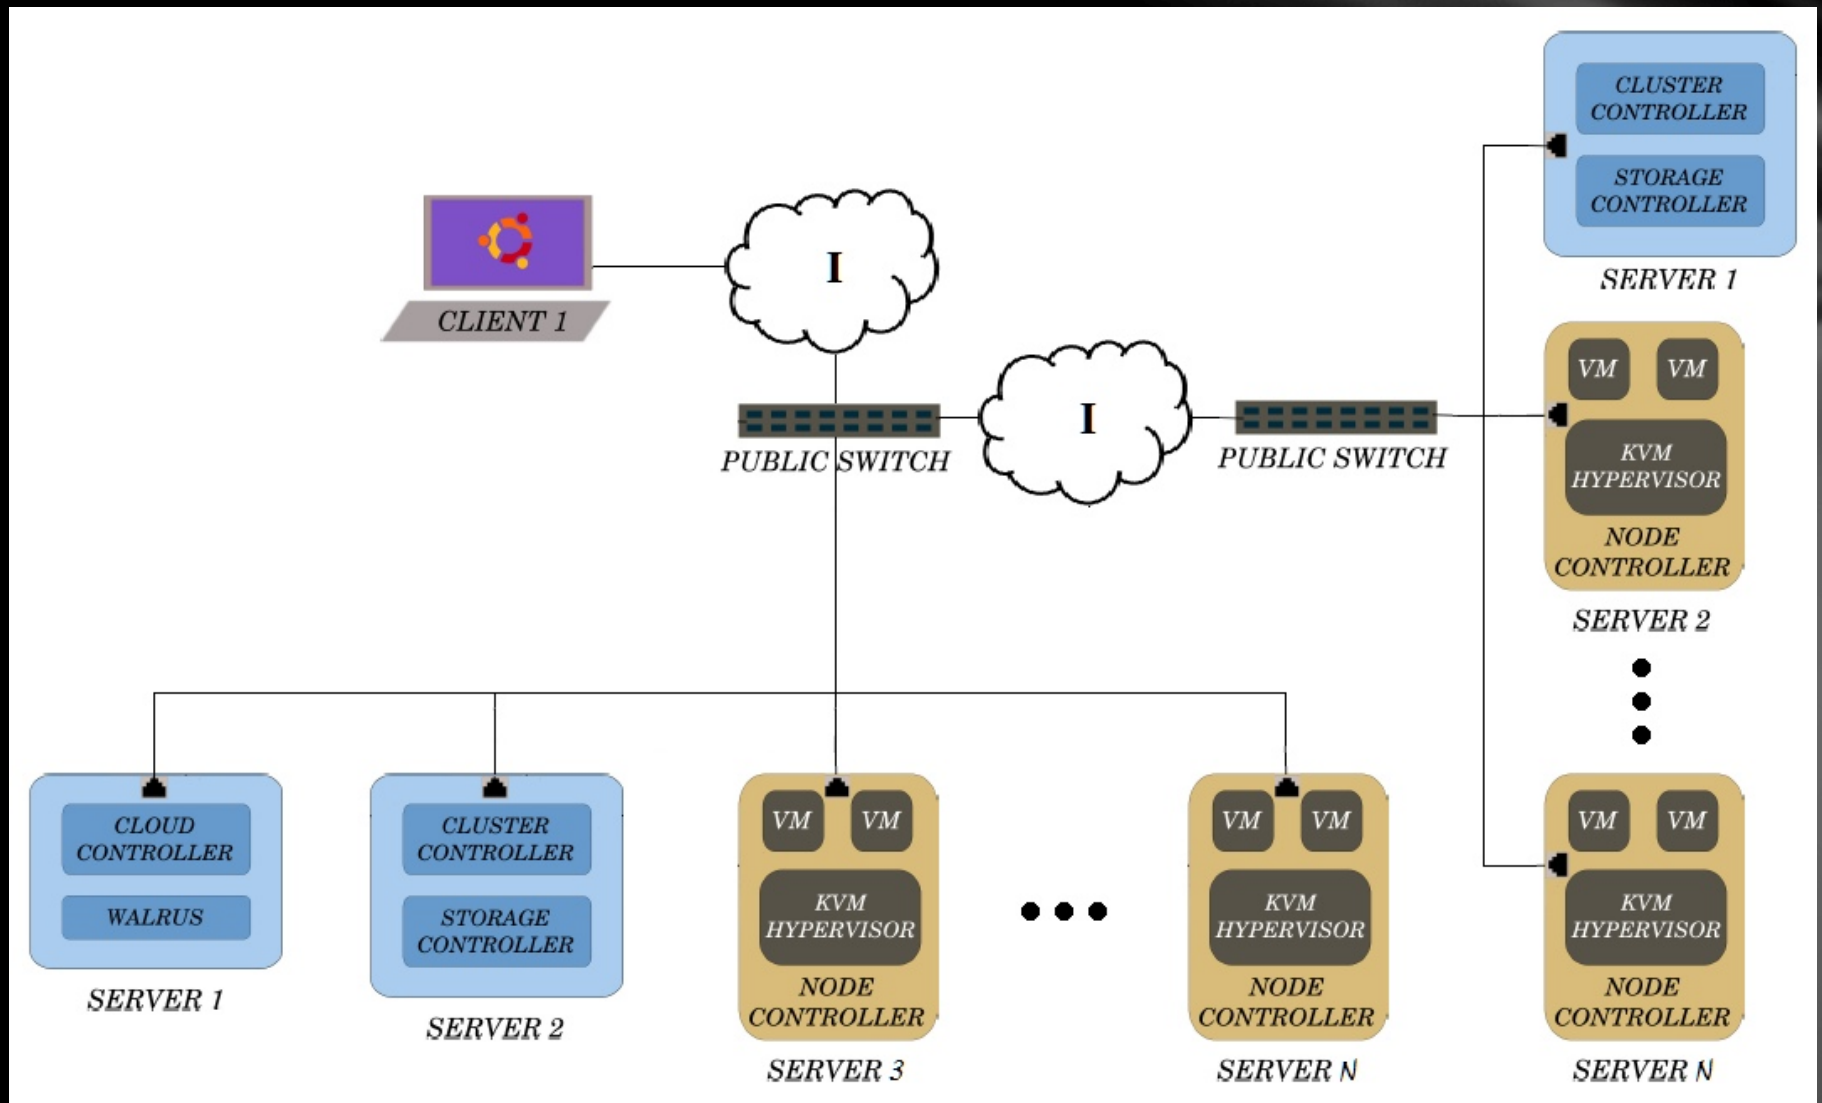
Implementation Concept (1)



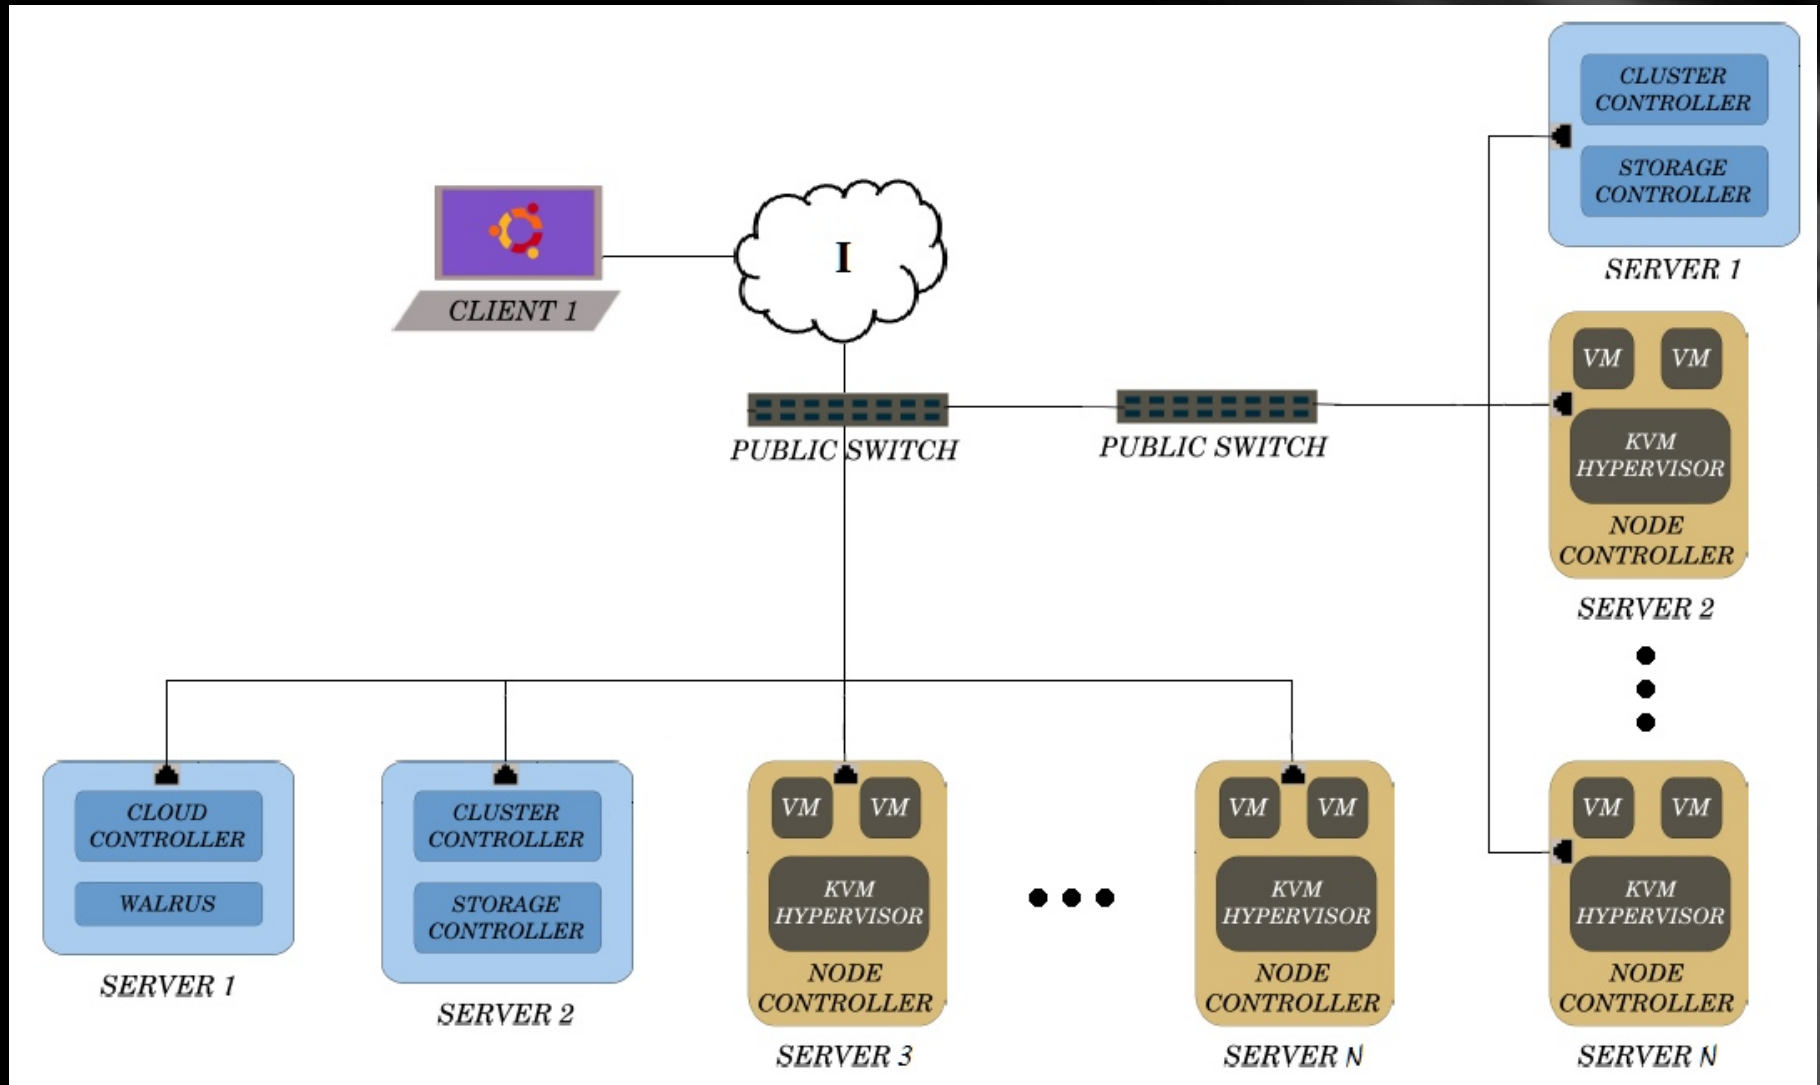
Implementation Concept (2)



Implementation Concept (3)



Implementation Concept (4)



Understanding Underlying Hardware

❑ Hardware:

- ❑ Cluster will initially consist of only 2 servers, a front-end and a compute node.
 - ❑ Front-end: 4-core Intel Xeon @ 2.67 GHz, 4 GB RAM, 250 GB disk, VT-enabled
 - ❑ Compute node: 8-core Intel Xeon @ 2.67, 6 GB RAM, 250 GB disk, VT enabled
- ❑ Once cluster is set up and tested to function properly, more compute nodes to be added to resource pool

❑ Network:

- ❑ DHCP server used for assigning IP addresses rather than manually configuring the network
- ❑ Front end and node connected by cross-cable rather than by a switch

Operating System choice and

❑ Ubuntu 10.04 LTS

- ❑ After early struggles with previous versions of Eucalyptus running on CentOS and Debian, Ubuntu provided faster installation, configuration, and Compatibility with respect to the OS itself, as well as Eucalyptus
- ❑ Ubuntu 12.04 LTS preferred, but installation issues arose with both the operating system as well as compatibility with Eucalyptus

❑ Partitioning of disk space

- ❑ Eucalyptus uses the /tmp directory to store image data when uploading and registering images, therefore majority of system disk space is allocated to the /tmp folder in front end server
- ❑ Eucalyptus uses the /var directory as dynamic disk space for instance creation, therefore majority of system disk space is partitioned to the /var folder in compute nodes

Eucalyptus Installation and

❑ Installation:

- ❑ Initially Eucalyptus Enterprise 3.0.2
- ❑ Eucalyptus combines open source and enterprise into version 3.1

❑ Configuration:

- ❑ DHCP limits available networking modes, attempts to configure DHCP to function under more advanced networking modes does not function
- ❑ This version would be installed and explored. Cluster in Miami to be updated to this version

Eucalyptus Testing

- Ensure all front end components are enabled and properly linked
- Upload an image for instance creation
- Generate a key pair to assign to instances
- Launch an instance using the uploaded image and ensure proper public and private IP address are assigned, as well as running status
- Log into instance using both the public IP address and the private IP address

Issues with integrating remote clusters

- ❑ Eucalyptus employs strict synchronization and timing between components, which causes time out issues between components if there is too much latency between them. i.e. if they are located too far from each other.
- ❑ Remote clusters under the same cloud-controller would not distinguish which IP ranges belong to which clusters. i.e. IP address from Miami could be assigned to virtual machines running in Sao Paulo, and vice versa

Trouble Shooting Issues With Remote

- ❑ Set up clusters as individual clouds and:
 - ❑ Set up a user front end with credentials to all clouds user has permission to access
 - ❑ Develop CLI commands to integrate multiple cluster access
 - ❑ Possibly go as far as a graphical interface for ease of use
- ❑ Upgrade OS/Cloud software to most recent versions
- ❑ Connect Miami and Sao Paulo clusters to Chicago cluster

Acknowledgements

- ❑ Bob Grossman and the University of Chicago
- ❑ Heidi Alvarez and FIU
- ❑ Carlos Moreira and NCC/UNESP
- ❑ Malcolm Atkinson and the University of Edinburgh
- ❑ Everyone in attendance

Questions?