

Thursday, December 2nd  
CloudCom 2010

# A Comparison and Critique of Eucalyptus, OpenNebula and Nimbus

Peter Sempolinski & Douglas Thain

What do you do when your  
boss says:  
“build me a cloud”  
?

# Why Private Clouds?

- Finer control of where VMs are deployed
- VM templates tuned to your environment
  - ex. AFS cell
- Secure proprietary information
- Reuse old hardware
- Save money (maybe)
- Interface with other private clouds

# Open Source Clouds

- Alternative to commercial clouds
- Good For:
  - A Private Company
  - Researchers
  - Anyone who wants to customize




OpenNebula



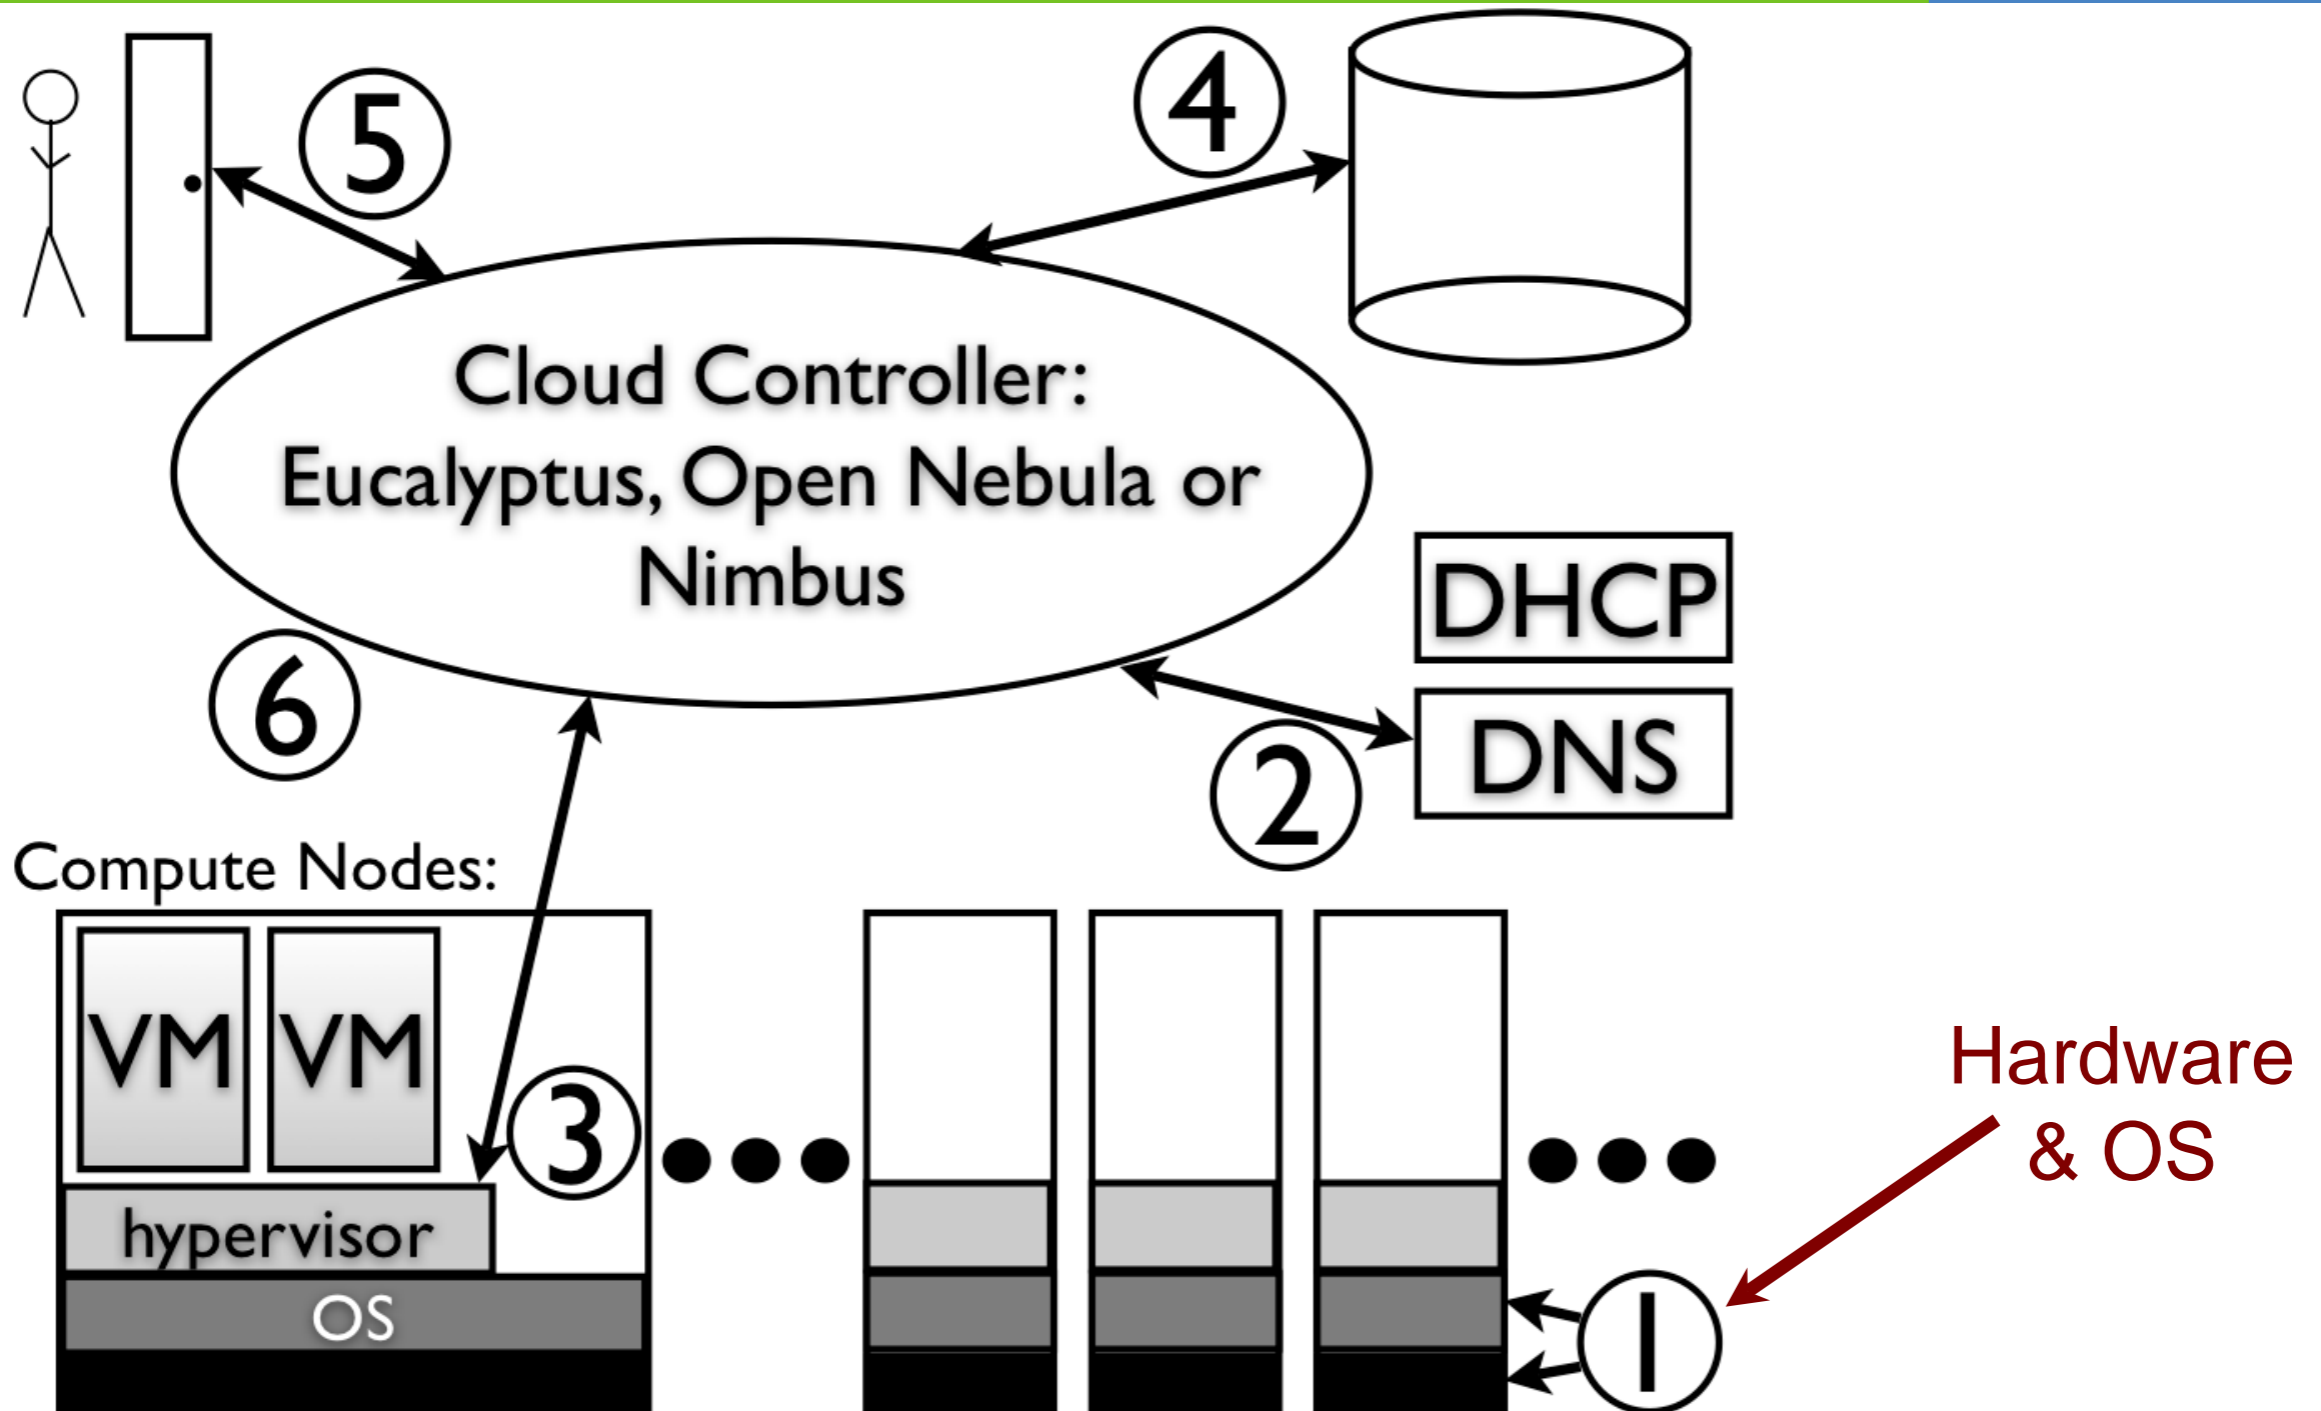
# Starting Observations

- Eucalyptus, OpenNebula and Nimbus are ***Evolving Projects***
- There are many parts to a ***Complete Cloud Computing Software Stack***
- Like all (good) open-source software, a private cloud allows us to ***Customize***

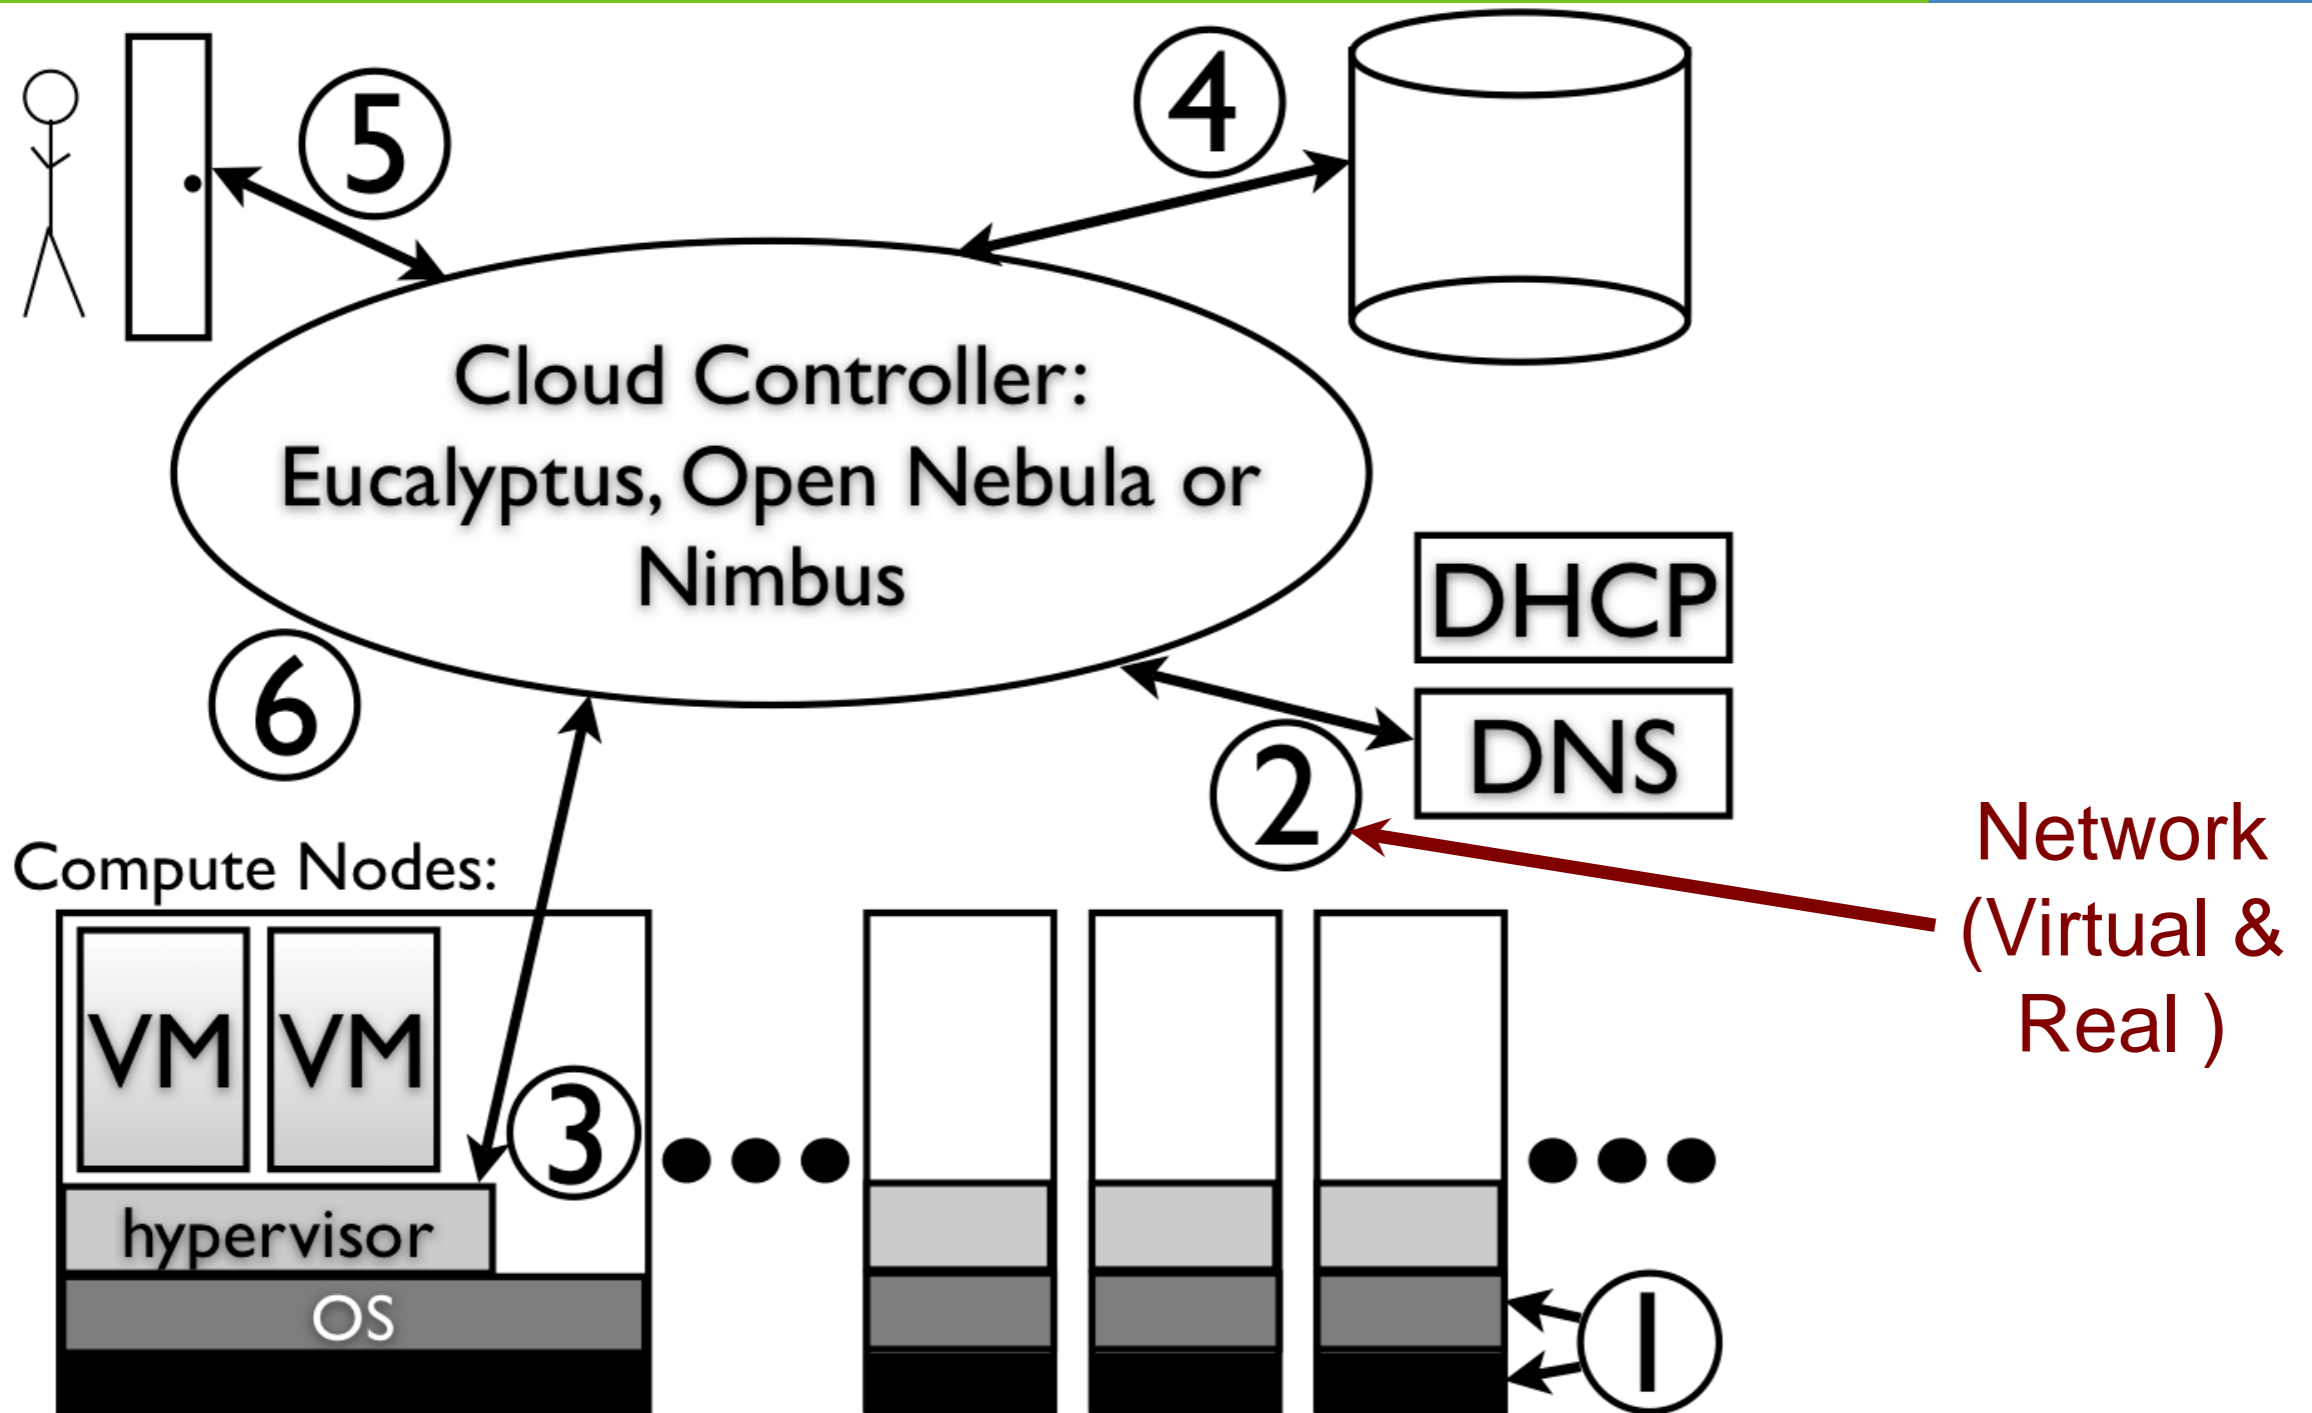
# An Abstract Cloud



# With Parts:

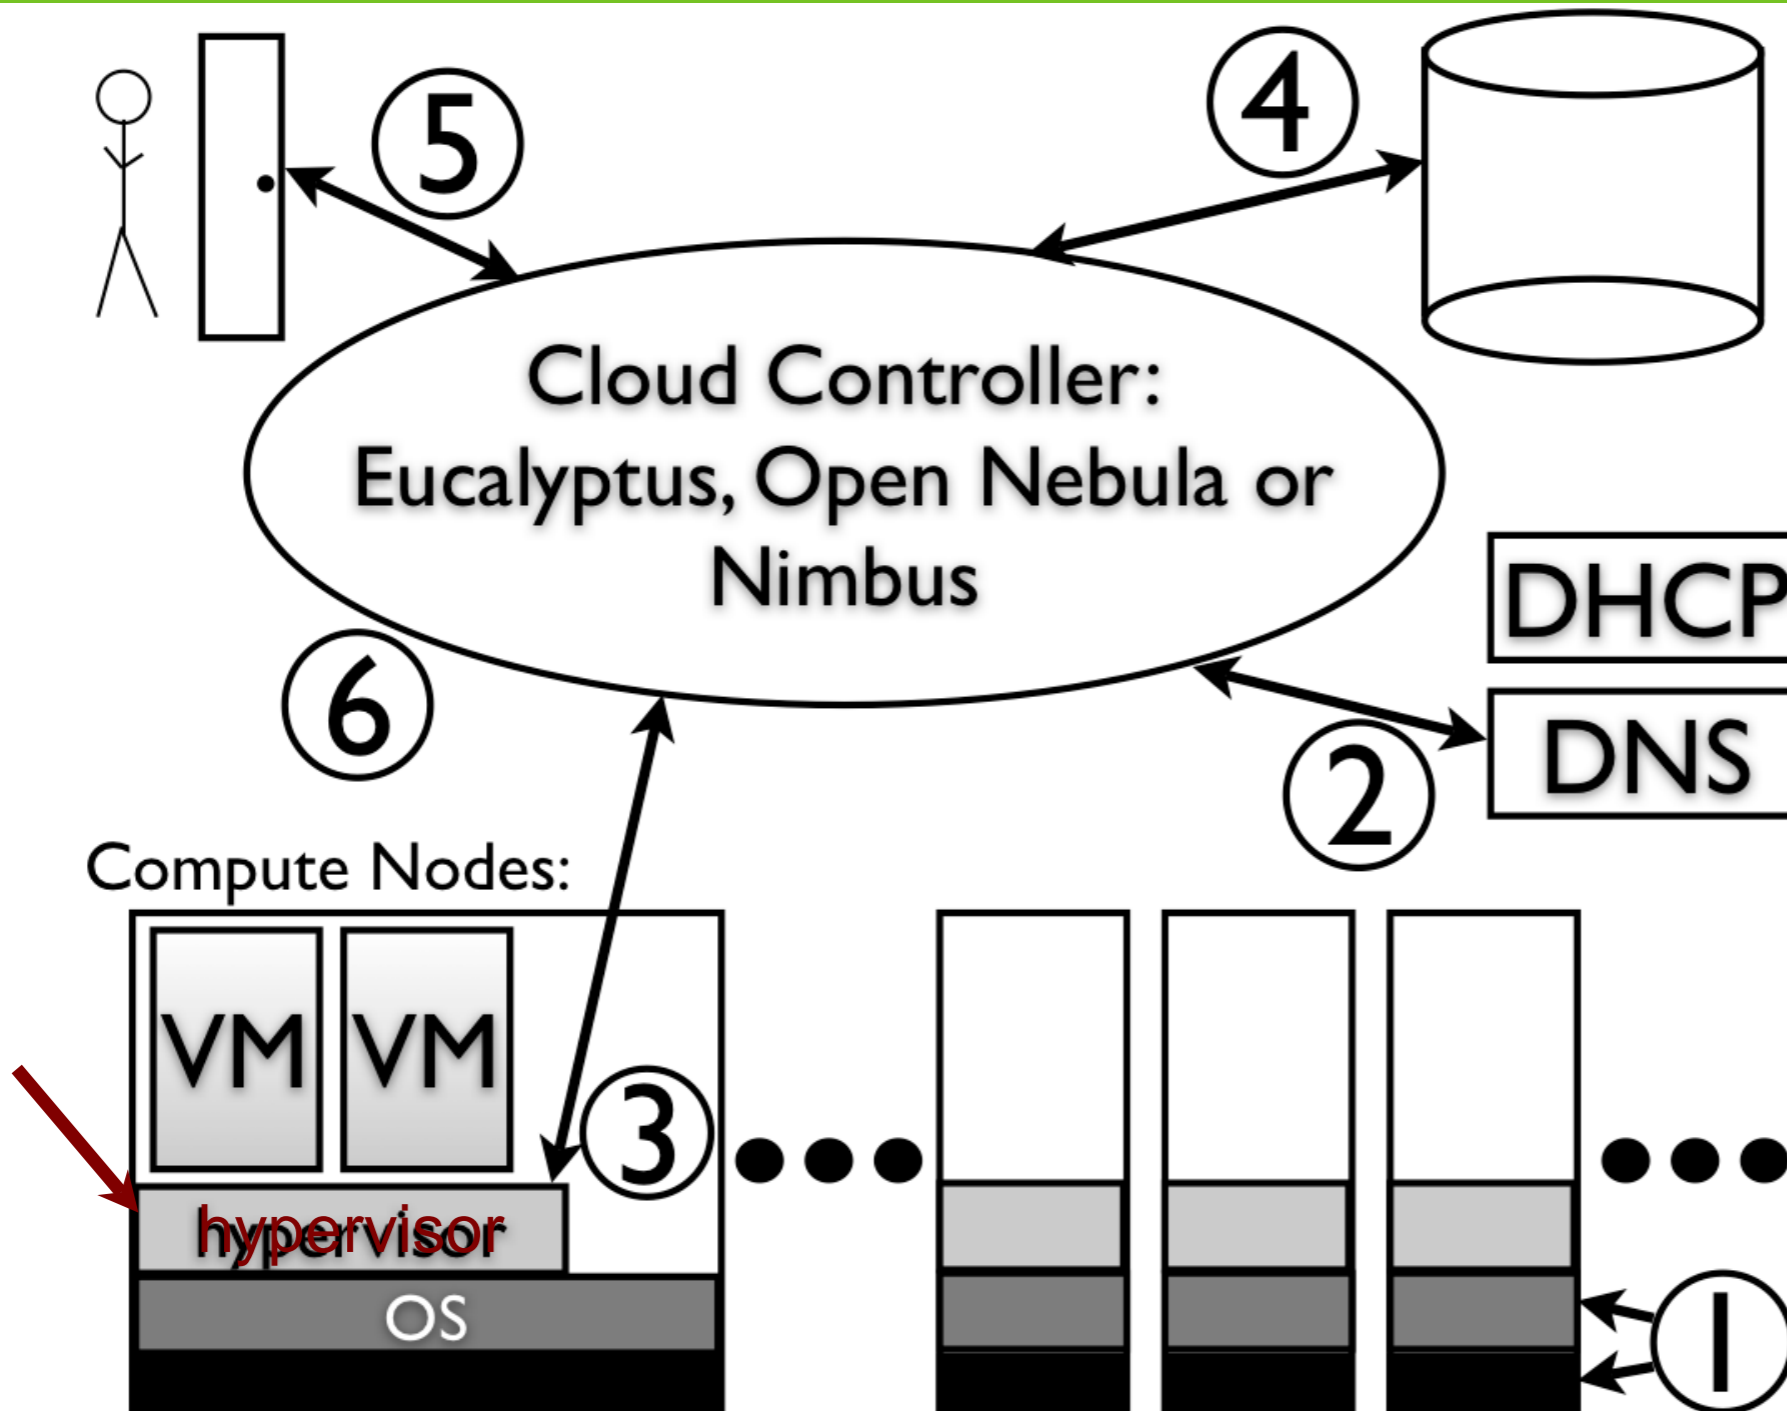


# With Parts:

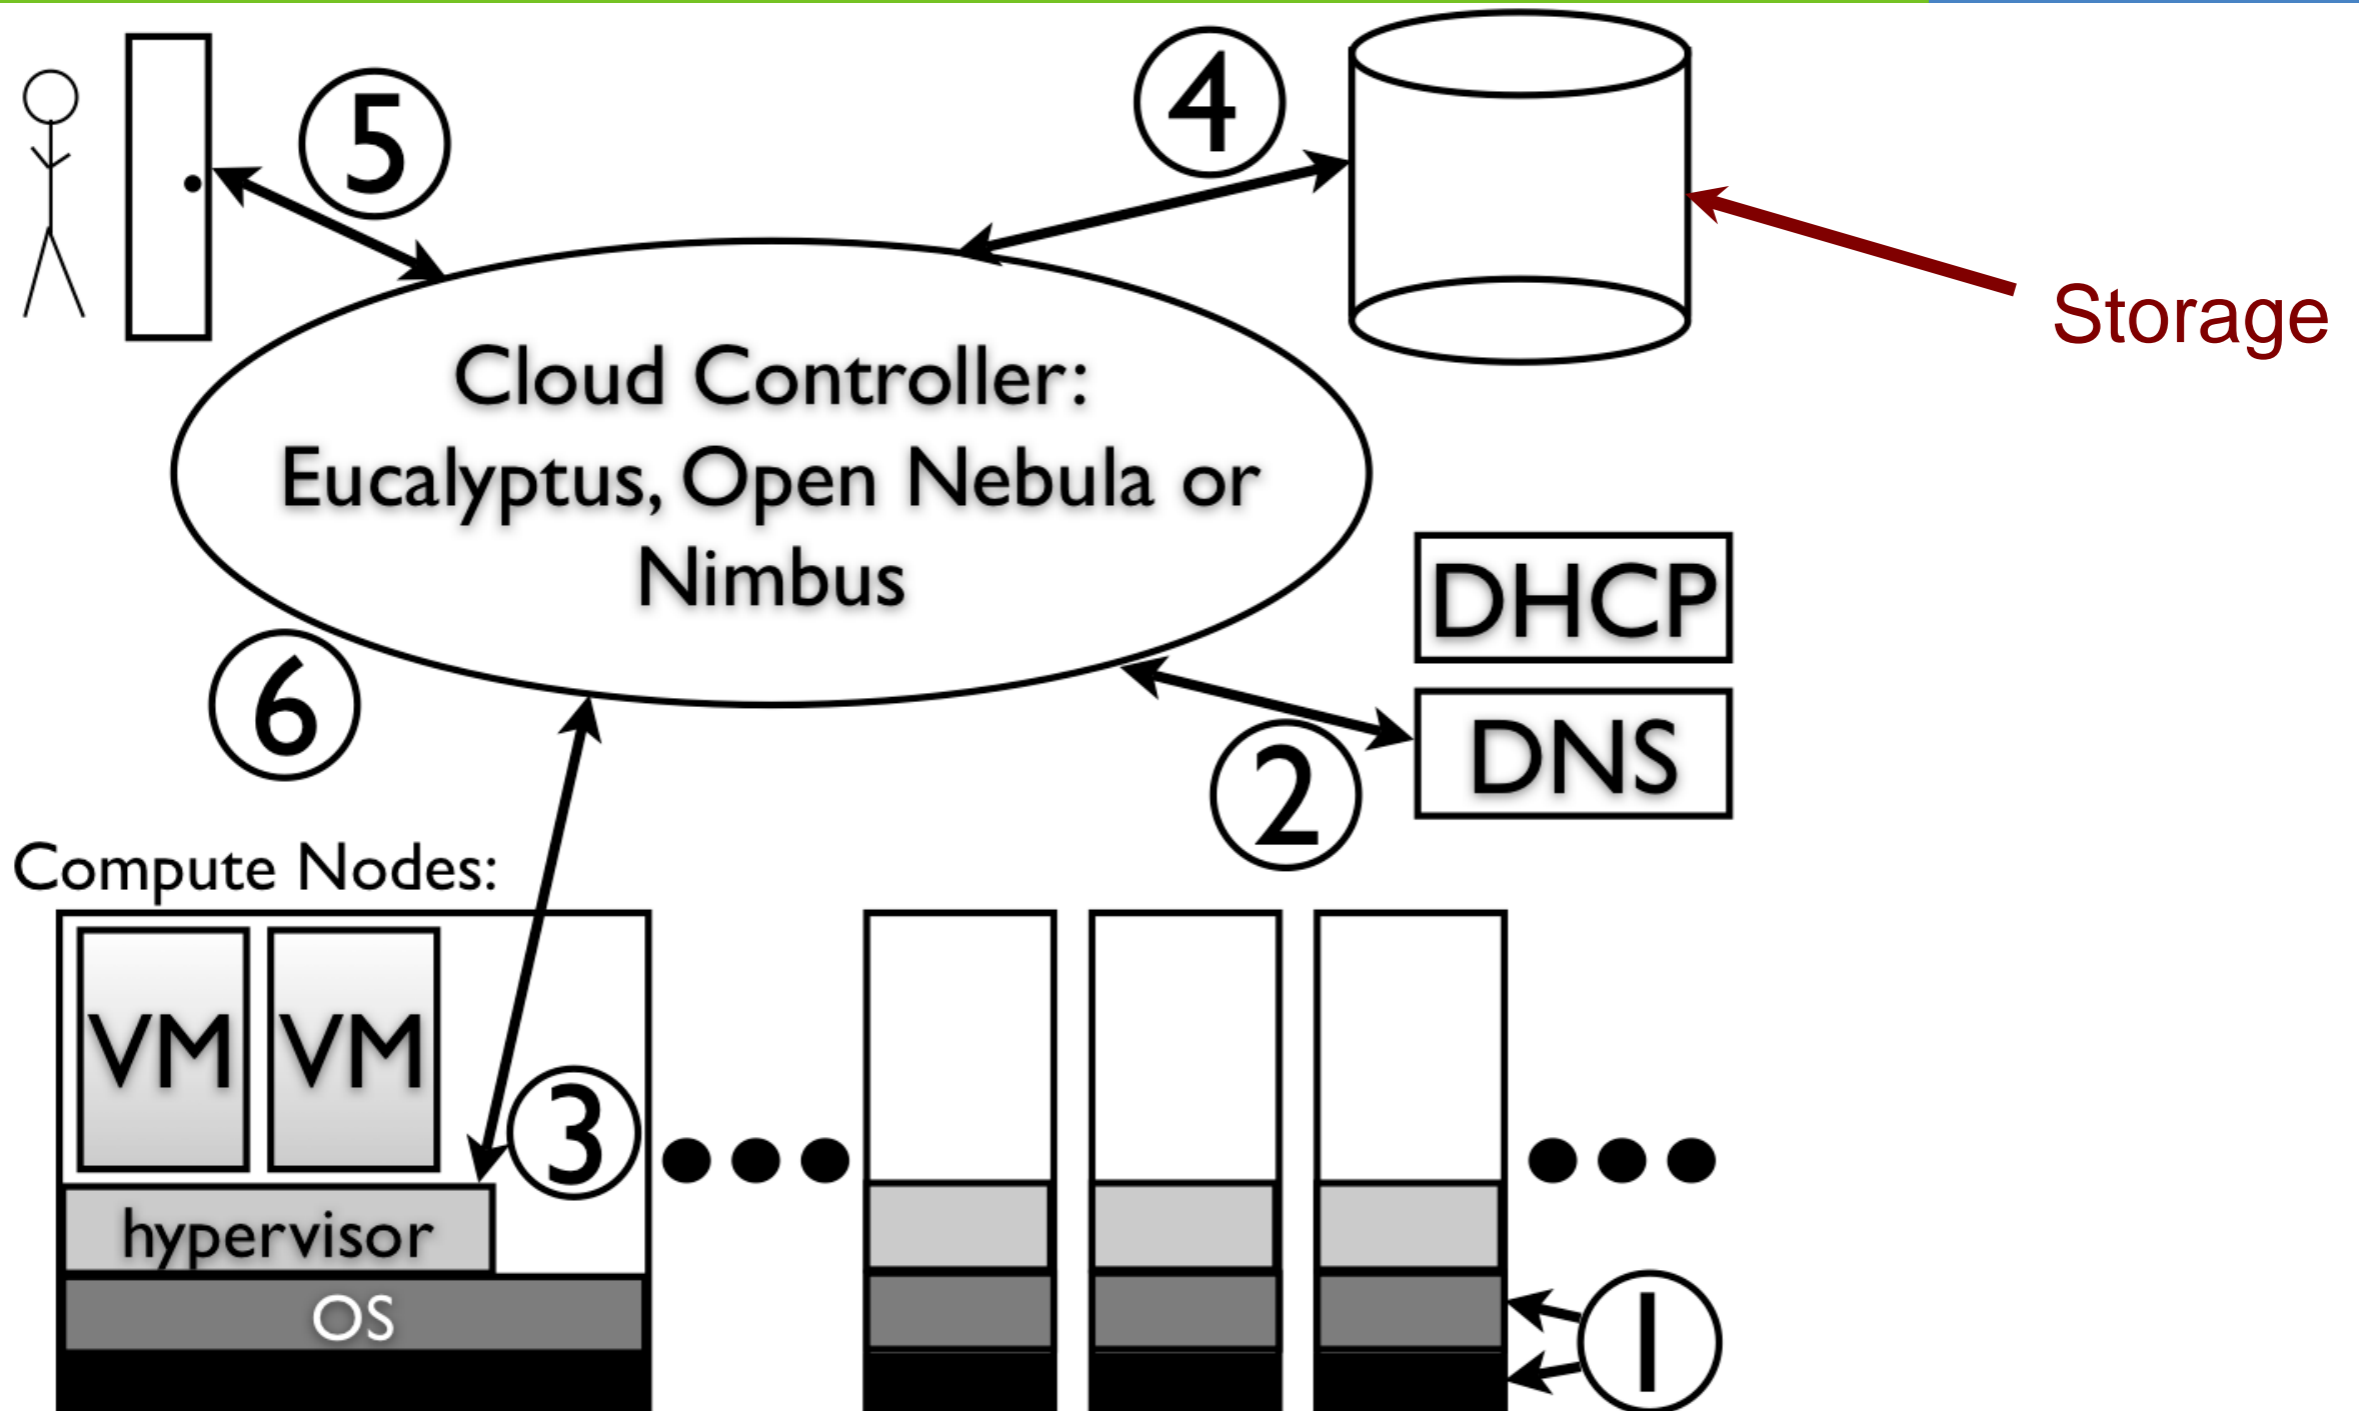




# With Parts:

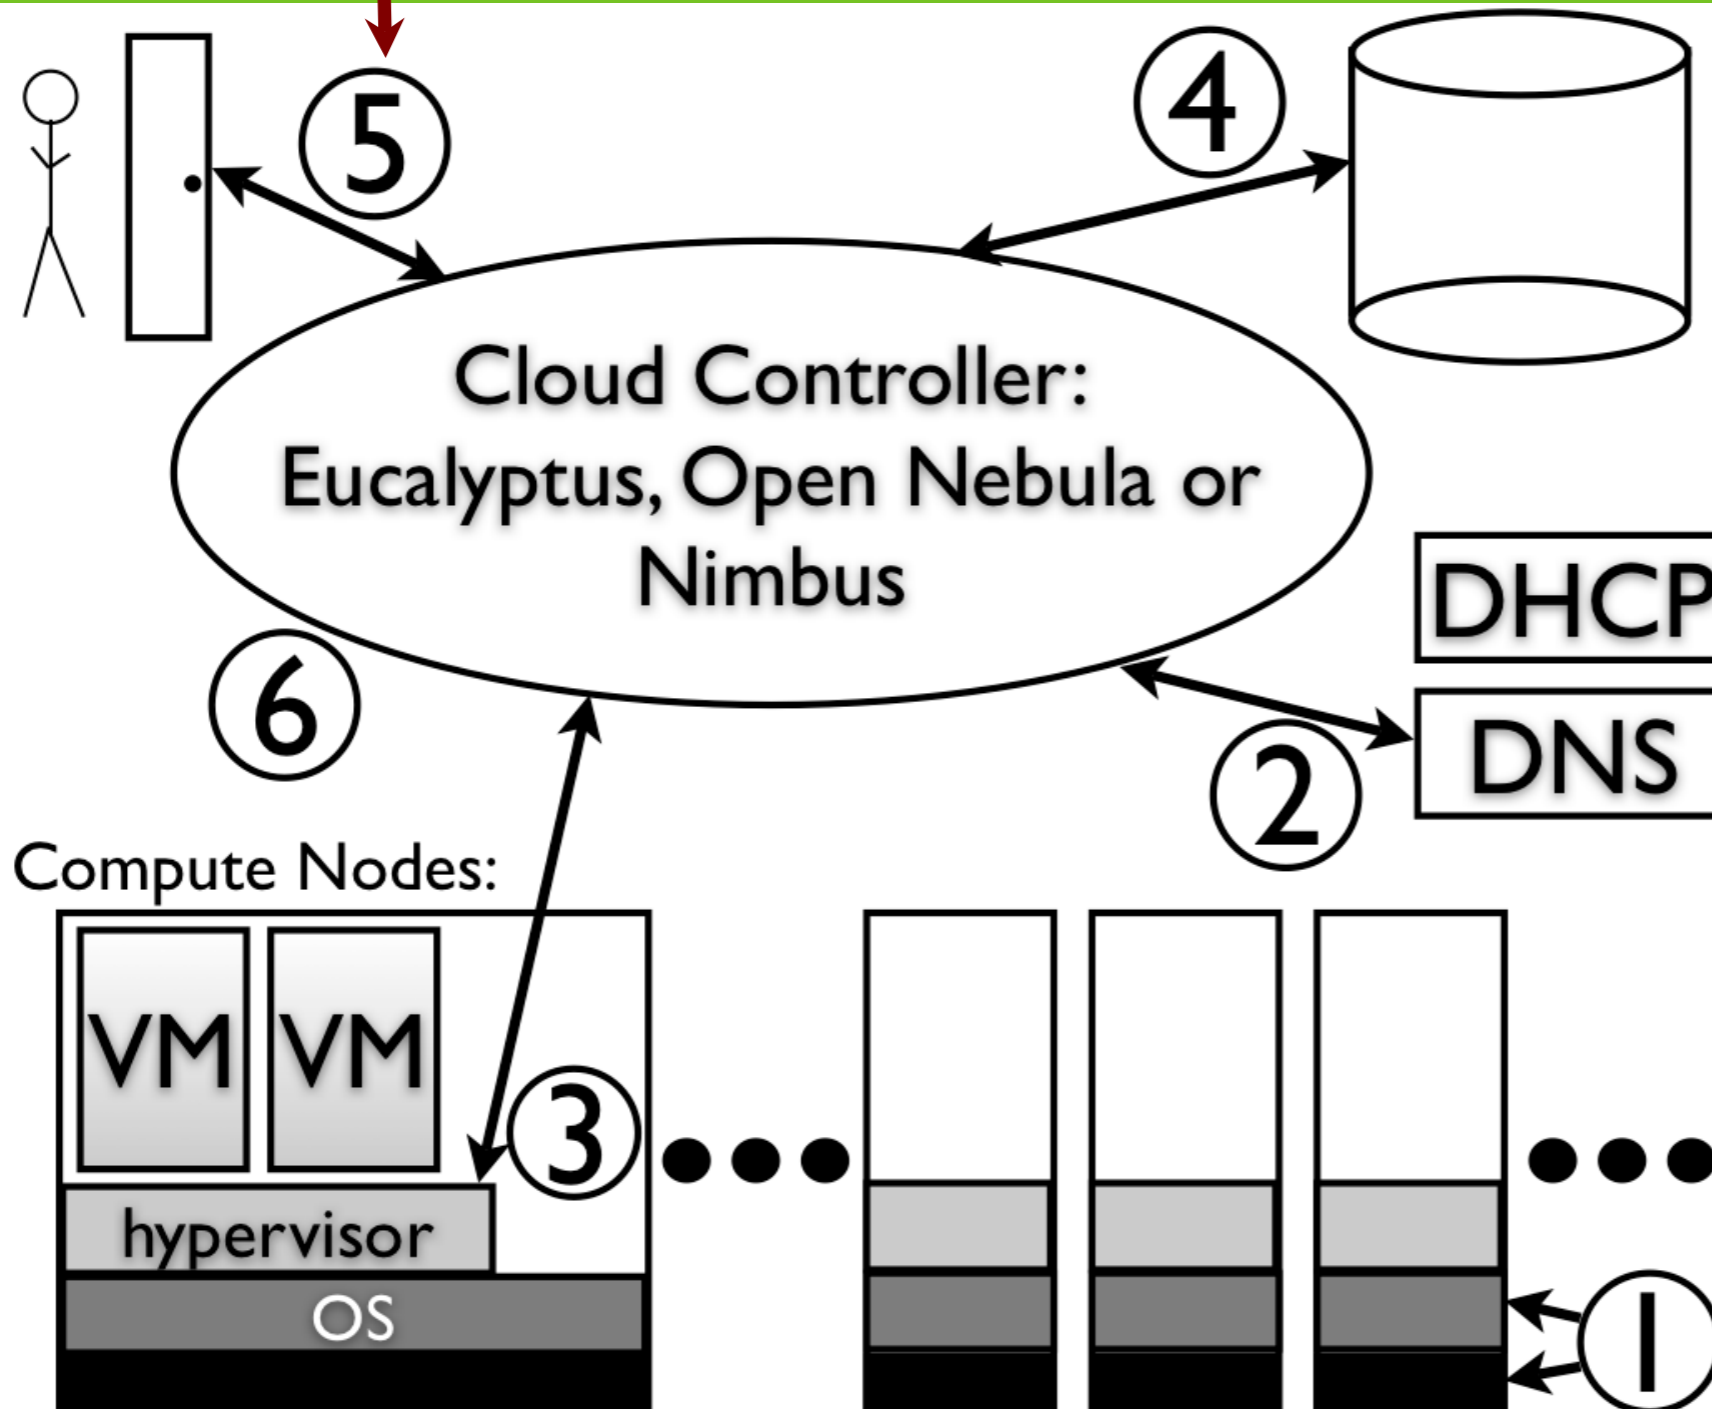


# With Parts:




# With Parts:

User Front-end



# With Parts:

The Software to tie  
all this together



# Quick Summary

- Comparison of:
  - underlying structure
  - guiding philosophy
- Because feature sets change
- We also talk about some difficulties in deployment

# Eucalyptus

- Open-source answer to EC2
- Extra tools for user management
- Multiple clusters
- Our Assessment: Very enterprise oriented

# OpenNebula

- Almost every part is customizable
- Defaults to private cloud where people log into head node
- Either distributed filesystem or scp
- Our Assessment: Very customizable. Defaults to smaller, private scale.

# Nimbus

- Uses globus tools (globus credentials)
- VERY active email forum
- Numerous research projects in cloud sharing & collaboration
- Our Assessment: Lives up to “Science” cloud self-appellation



# Quick Summary

- Eucalyptus: Like EC2, large deployments, user management
- OpenNebula: Customizable, centralized, private cloud
- Nimbus: “Science” cloud, globus, collaboration

# Underlying hardware & OS

- The major issue is compatibility
- Often dictate versions of other software
- example: kvm for RedHat 5 does not simulate scsi disks
- Affects configuration of all 3 clouds

# Network Components

- Assumptions made about network control
- Best case: cloud control has their own subnet range
- Practically, negotiation required with the DHCP, DNS of wider network.

# Hypervisor

- Dictates allowed disk image type
- Greatly affects performance
- Libvirt's abstraction is tricky
  - libvirt is used to spawn VM
  - But, the cloud still needs to account for underlying hypervisor

# Storage

- Eucalyptus & Nimbus use S3-like distributed storage
- OpenNebula has option of shared file system or scp
- Storage consumes both space & time

# Front-End

- Most customizable part
- Interesting algorithmic problem
- Commercial clouds use money
- Grids and condor use preemption
- But, what is fair scheduling for private VMs?

# Summary

- If you want more details on:
  - the parts of a private cloud
  - the underlying ideas and structure of Eucalyptus, OpenNebula, Nimbus
  - Some challenges and problems in open-source clouds
- Please read the paper

Peter Sempolinski  
Computer Science and  
Engineering Department  
University of Notre Dame

Questions?