

# NGS Surgery – Eucalyptus Private Cloud

---

Matteo Turilli

**Oxford e-Research Centre (OeRC)**

Steve Thorn

**National e-Science Centre (NeSC)**

David Fergusson

**National e-Science Centre (NeSC)**



# Outline

---

- Community Driven R&D effort
- SaaS, PaaS and IaaS
- Eucalyptus Vs Nimbus
- User perspective
- EC2 tools
- Eucalyptus Architecture
- Eucalyptus Networking Architecture
- Eucalyptus Data Architecture
- Eucalyptus AAA
- Oxford Infrastructure
- Edinburgh Infrastructure
- Roadmap


# NGS 3 EWP2

---

- “NGS Agile Deployment Environments”
- EPSRC funded, 2 years
- People
  - Matteo Turilli (OeRC, Oxford) [0.75 FTE]
  - Steve Thorn (NeSC, Edinburgh) [0.5 FTE]
  - David Fergusson (NeSC, Edinburgh) [WP Leader]

# Community Driven R&D Effort

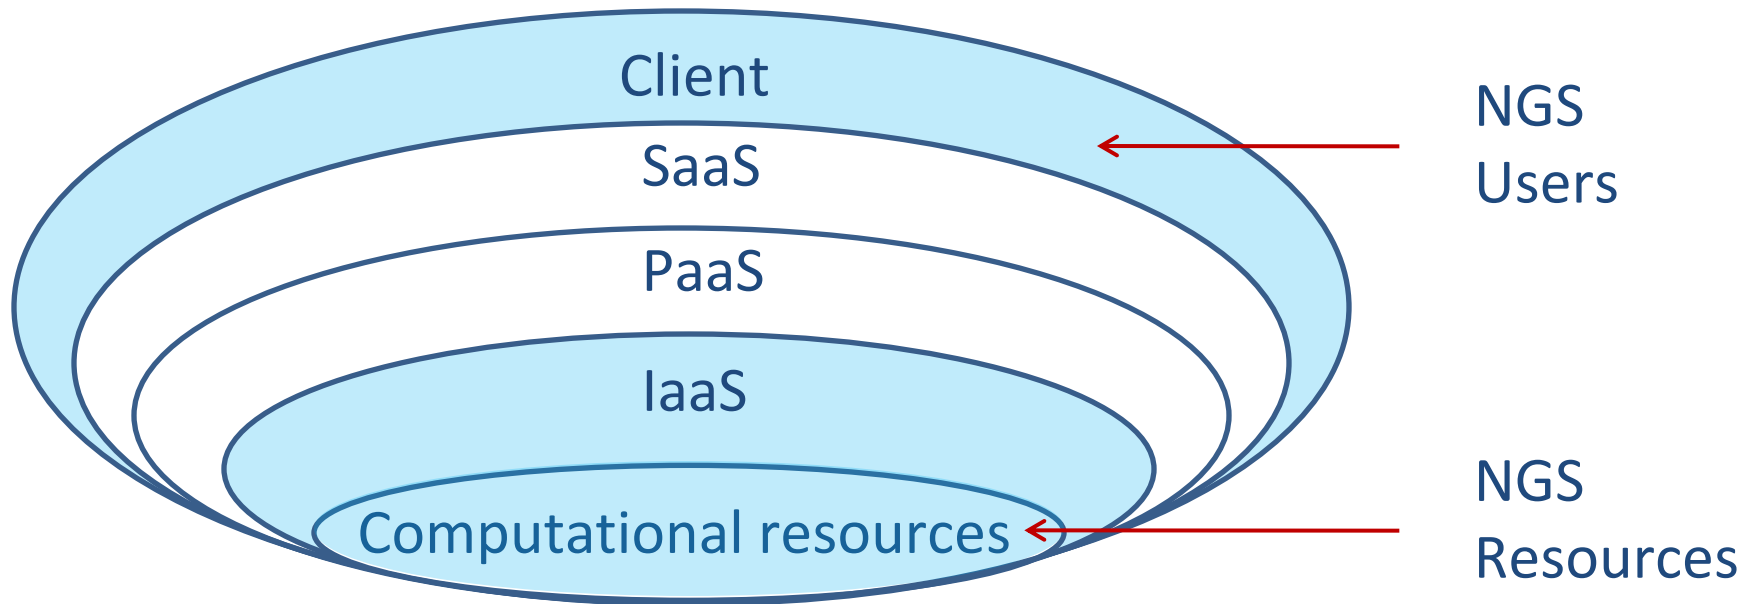
---

- Several NGS users use or plan to use Amazon EC2
  - NGS user forum: users from the bioinformatics community explicitly asked for EC2 compatible cloud solutions
  - NGS training: deploying virtual NGS resources ‘in the cloud’
  - NGS testing and validation: validation of documentation about NGS software stack, R&D
- 
- EC2 compatible, open source solution: Eucalyptus and/or Nimbus
  - Development of a dedicated pool of images
  - Collecting data about feasibility, costs, stability
  - Identify use cases and gather further requirements

# SaaS, PaaS, IaaS and Private Cloud

---

- SaaS: Software as a Service : Google Apps, Office Web Apps
- PaaS: Platform as a Service : Google App Engine, Amazon Web Services
- IaaS: Infrastructure as a Service : Amazon EC2, Eucalyptus, Nimbus



# Eucalyptus

Vs

# Nimbus

## Pros

- Good implementation of EC2 and EBS APIs
- 'Fast' deployment via RPMs
- Modular design
- Xen and Kvm compatible
- Open source and commercial

## Cons

- Early development phase
- Design limitations
- AAA
- Unstable

## Pros

- AAA
- More mature
- Adopted in research context

## Cons

- Limited EC2 implementation
- No native support of EBS
- Globus WS4

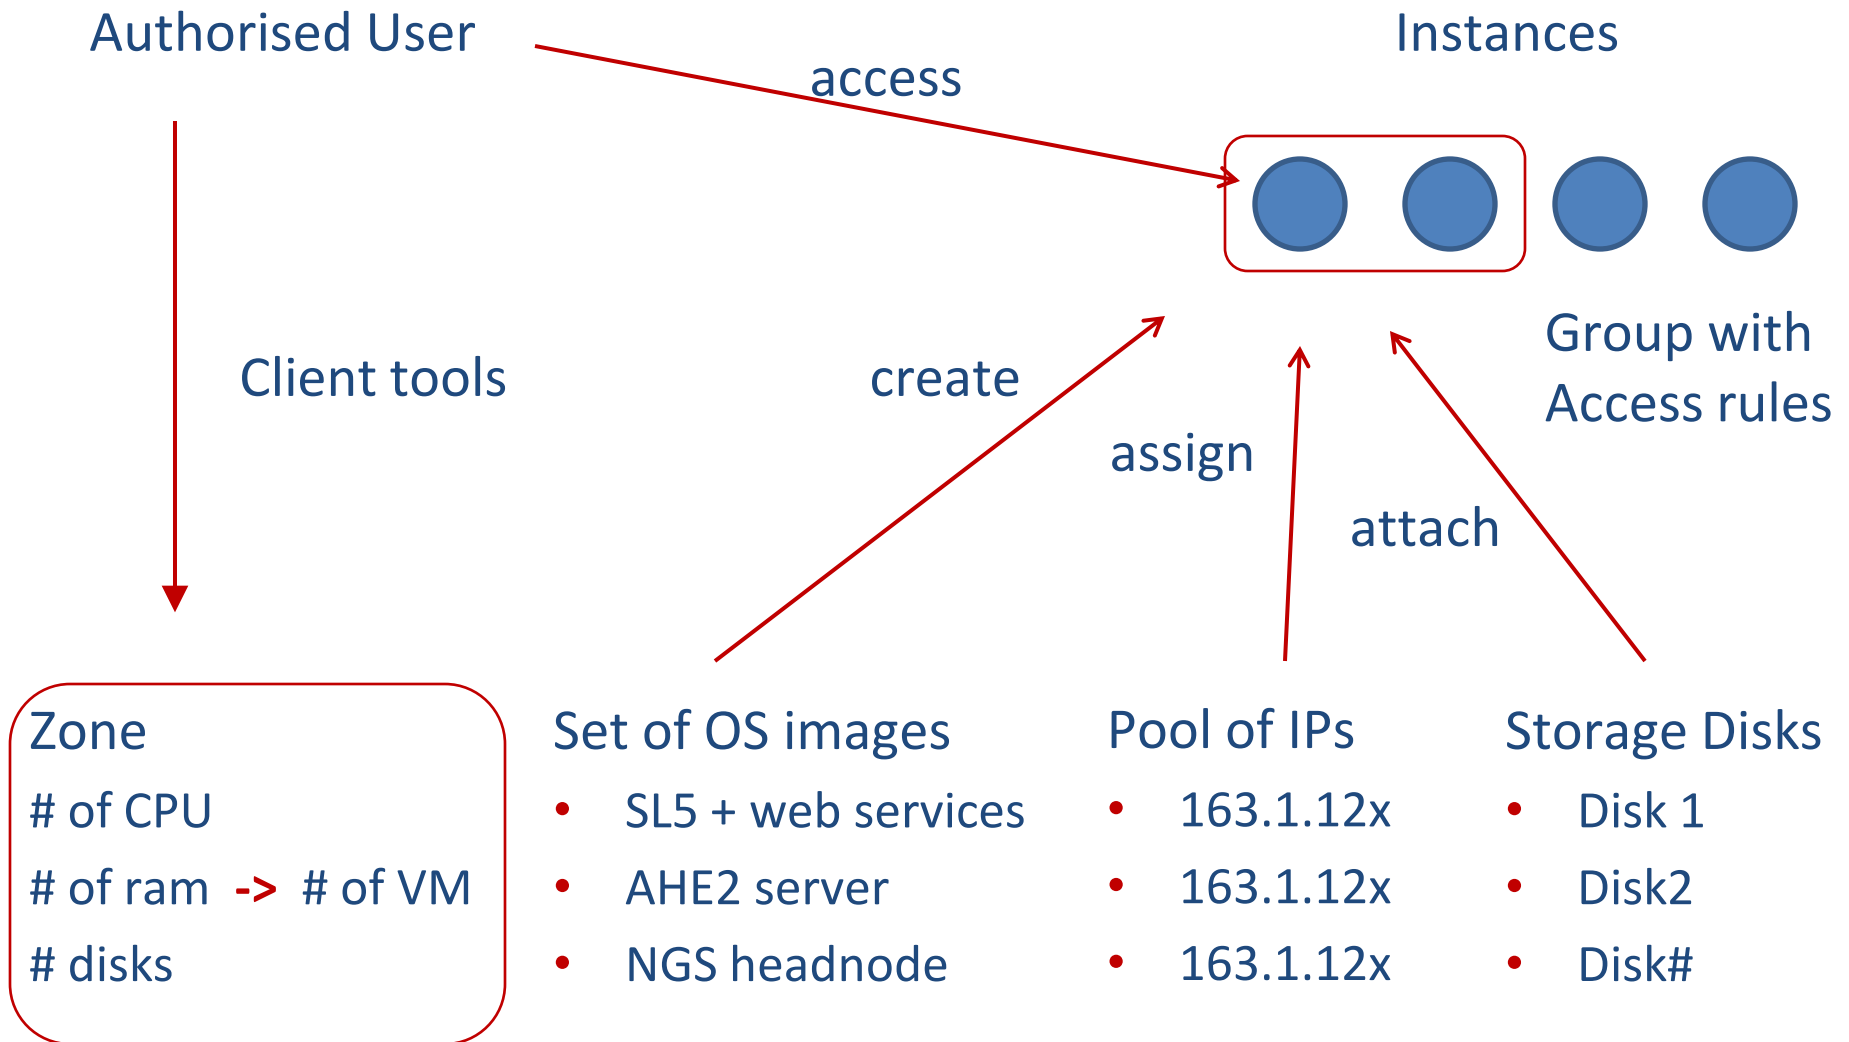
# User Perspective of IaaS

---

## Virtualization platform

- User has (full) control:
  - VM instantiation: start, stop, #VMs
  - OS: create, modify or use existing images
  - Network: dynamic public IP, virtual firewalls, isolation (security groups)
  - Virtual block devices: create, attach, snapshot, copy
  - ...
- Great flexibility vs. extra effort
- Hardware abstracted - SEP

# User Perspective



# EC2 Client Tools

---

- Listing available resources subdivided into zones:
  - `euca-describe-availability-zones`
- Managing credentials:
  - `euca-add-keypair`
  - `euca-describe-keypairs`
- Managing (OS) images:
  - `euca-bundle-image`
  - `euca-upload-bundle`
  - `euca-register`
  - `describe-images`

# EC2 Client Tools

---

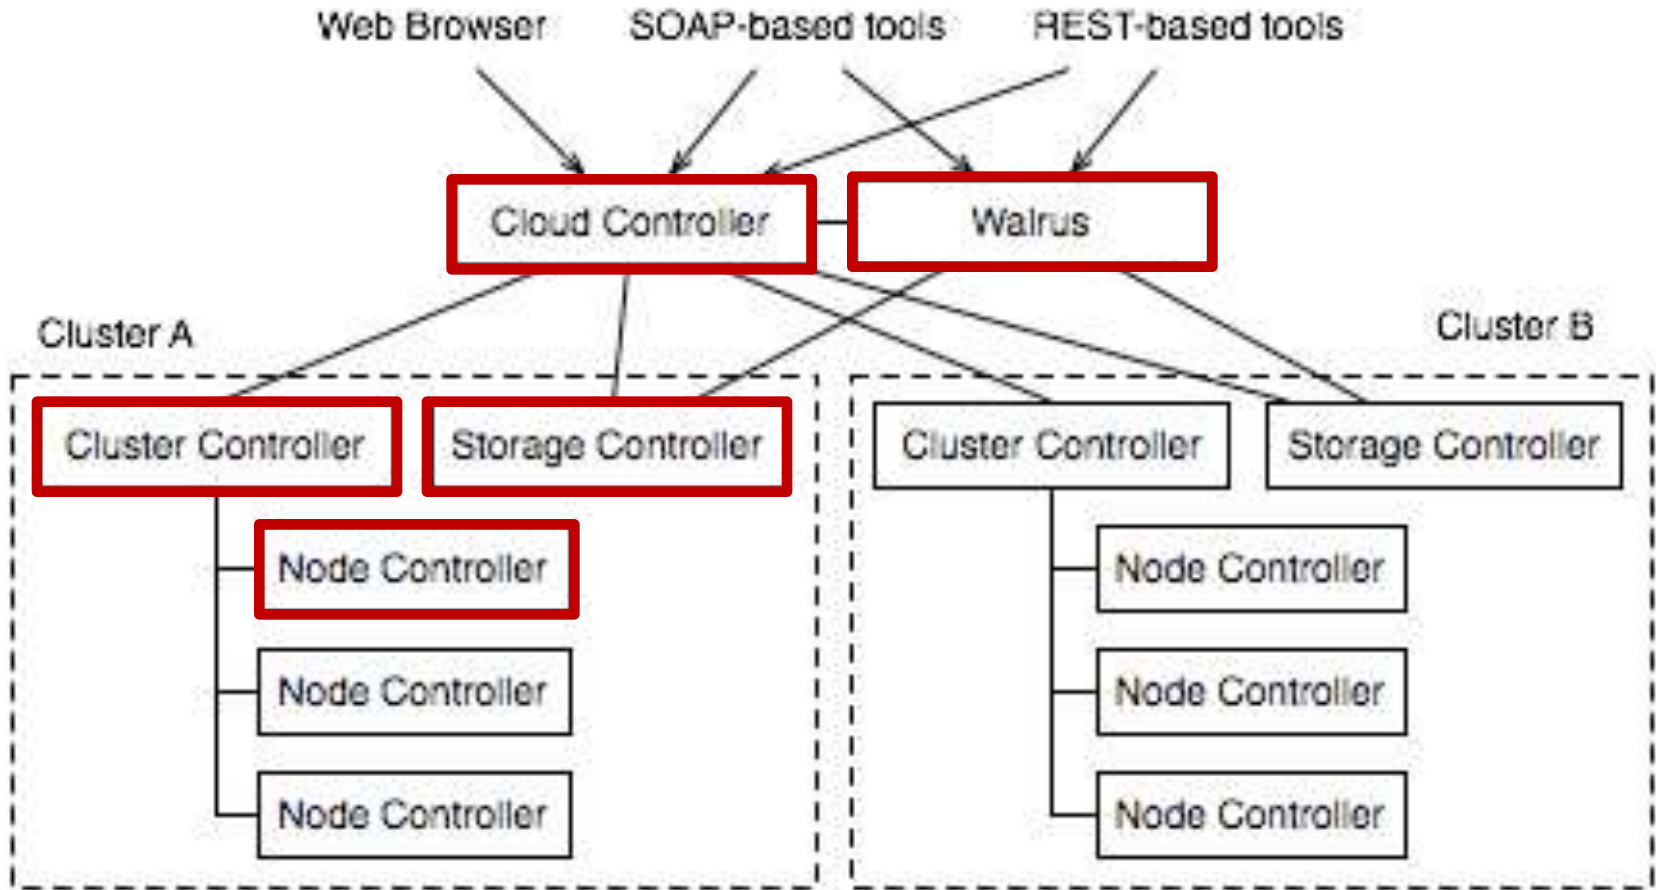
- **Managing instances:**
  - `euca-describe-instances`
  - `euca-run-instances`
  - `euca-terminate-instances`
  - `euca-reboot-instances`
- **Managing elastic IP:**
  - `euca-allocate-address`
  - `euca-associate-address`
  - `euca-disassociate-address`
  - `euca-release-address`

# EC2 Client Tools

---

- **Managing security groups:**
  - `euca-add-group`
  - `euca-authorize`
  - `euca-revoke`
  - `euca-delete-group`
- **Managing EBS**
  - `euca-create-volume`
  - `euca-attach-volume`
  - `euca-detach-volume`
  - `euca-delete-volume`
  - `euca-create-snapshot`

# Eucalyptus Architecture



[http://open.eucalyptus.com/wiki/EucalyptusInstall\\_v1.6](http://open.eucalyptus.com/wiki/EucalyptusInstall_v1.6)


# Eucalyptus Network Architecture

---

## 4 networking modes:

- **SYSTEM:** random MAC address. No elastic IPs, no security groups, no traffic isolation among VMs.
- **STATIC:** managed DHCP and static MAC/IP pairs.
- **MANAGED:**
  - Private, un-routable, large (/8, /16) network for VMs
  - Managed DHCP.
  - Segments of private network managed as a vlan.
  - Pool of public IPs .
- **MANAGED-NOVLAN:** iptables chains instead of vlans, cannot offer VM isolations.

# Managed-novlan



## Head node iptables


FORWARD	ACCEPT	all	--	*	*	0.0.0.0/0	!192.168.0.0/16
FORWARD	ACCEPT	all	--	*	*	192.168.2.0/26	192.168.2.0/26
FORWARD	group_name	all	--	*	*	0.0.0.0/0	0.0.0.0/0
group_name	ACCEPT	all	--	*	*	x.x.x.x/x	192.168.2.0/26
PREROUTING	DNAT	all	--	*	*	0.0.0.0/0	163.1.12y.yy to:192.168.2.2
OUTPUT	DNAT	all	--	*	*	0.0.0.0/0	163.1.12y.yy to:192.168.2.2

# Eucalyptus Data Architecture

---

## 2 data subsystems:

- **Walrus:** enables the creation of 'buckets', repositories of OS, kernel, initrd images. Users can create them via euca-tools. (TBD: explore isolation of buckets).
- **Storage controller:** enables the creation of EBS. EBS can be used to snapshotting VMs. If properly unmounted, EBS are persistent across attachments (to different VMs) or VM reboots.



# Eucalyptus AAA


---

- **Authentication/Authorisation:**
  - **Users:**
    - **login/password:** users connect via https to the cloud controller web interface. Register, get approved, log into the web interface and download a zipped x509 certificate.
    - **x509:** credentials used to interrogate the cloud controller about zone, images, instances, storage blocks, etc.
    - **Key pair:** generated by the users to access running instances. The key is injected into the instance at creation time.
  - **Cluster/Node controllers:** synchronised key pairs.
- **Accounting:** controllers activity is logged. Network usage is not.

# Oxford Infrastructure

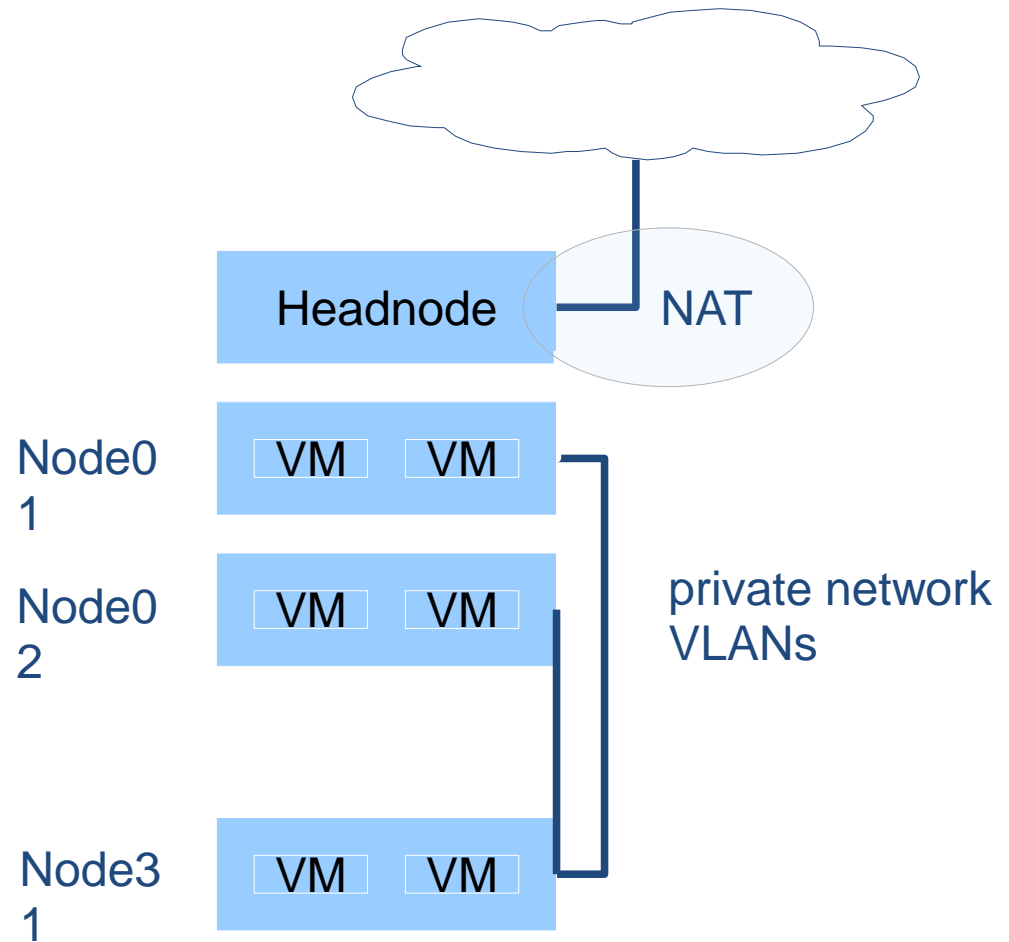
---

- NGS2 Cluster: 1 head, 3 storage, 8 thick, 49 thin nodes
- x86\_64, 32/8GB ram, no HW virtualisation
- FB SAN, 25TB storage
- Theoretical peak load: 175/525
- Linux/BSD/\*nix images
- Managed network mode



# Edinburgh Infrastructure

- 32 x Sun Fire X4100
- Dual-core, 2.8 GHz  
Opteron 8 GB RAM, 70 GB RAID1
- 64 cores
- 1 Headnode (Cloud and Cluster controllers)
- 31 Nodes (Node controller)
- Max 2 VMs per core:  
124 slots (2GB RAM)
- VLANs for VM isolation



# Roadmap

---

- Finalising infrastructures deployment
- Investigating 3<sup>rd</sup> party tools: RightScale, Elasticfox, s3fs, s3curl, s3cmd
- Opening both infrastructures to beta testers/early adopters
- Finalising and deploying training images
- Developing more images
- Evaluating Nimbus for AAA
- Coordinating with eucalyptus developers for AAA integration

# Thank you

---

Matteo Turilli

**Oxford e-Research Centre (OeRC)**

Steve Thorn

**National e-Science Centre (NeSC)**

David Fergusson

**National e-Science Centre (NeSC)**

