

# The StratusLab Cloud Distribution

Workshop on Software Services

Timisoara, Romania

6 June 2011



StratusLab is co-funded by the  
European Community's Seventh  
Framework Programme (Capacities)  
Grant Agreement INFSo-RI-261552



## High-Energy Physicist (<2001)

- E735 (Duke), Fermilab Tevatron
- CDF (Rutgers), Fermilab Tevatron
- ALEPH (CERN), CERN LEP
- ATLAS (UCSC), CERN LHC

## Grid & Cloud Computing (2001—)

- European DataGrid: Integration and deployment
- EGEE, EGEE-II, EGEE-III: Coordinator for user-facing activities
- Quattor: Principal developer for pan language compiler
- StratusLab: Project Coordinator

## SixSq Sàrl (2007—)

- Founded in 2007 by 3 partners, based in Geneva
- Provides consulting and tools for automated systems testing

# StratusLab Project



## Goal

- Create comprehensive, open-source, IaaS cloud distribution
- Focus on supporting grid services

## Information

- 1 June 2010—31 May 2012 (2 years)
- 6 partners from 5 countries
- Budget : 3.3 M€ (2.3 M€ EC)

## Contacts

- Site web: <http://stratuslab.eu/>
- Twitter: @StratusLab
- Support: [support@stratuslab.eu](mailto:support@stratuslab.eu)



CNRS (FR)



UCM (ES)



GRNET (GR)



SIXSQ (CH)



TID (ES)



TCD (IE)

## **Grid and cloud technologies are complementary**

- Uniform security model (grid)
- Sharing of resources, algorithms, and expertise (grid)
- Dynamic allocation of resources (cloud)
- Customized environments (cloud)

## **Only develop new software when necessary**

- Integrate existing solutions if possible
- Practical development → real needs of users

## **Maintain production quality with rapid evolution**

- Use agile and scrum methodologies
- Iterative integration: always maintain working distribution
- Public releases approximately every 6 weeks

# Infrastructure as a Service (IaaS)



## Goal

- Offer remote access to computing resources

## Advantages

- Customized environments
- Rapid access via simple API
- Complete control (root access) with “pay as you go” model

## Disadvantages

- Non-standard interfaces (vendor lock-in)
- Creating new virtual machines is difficult

Software as a Service (SaaS)

Platform as a Service (PaaS)

Infrastructure as a Service (IaaS)




ElasticHosts



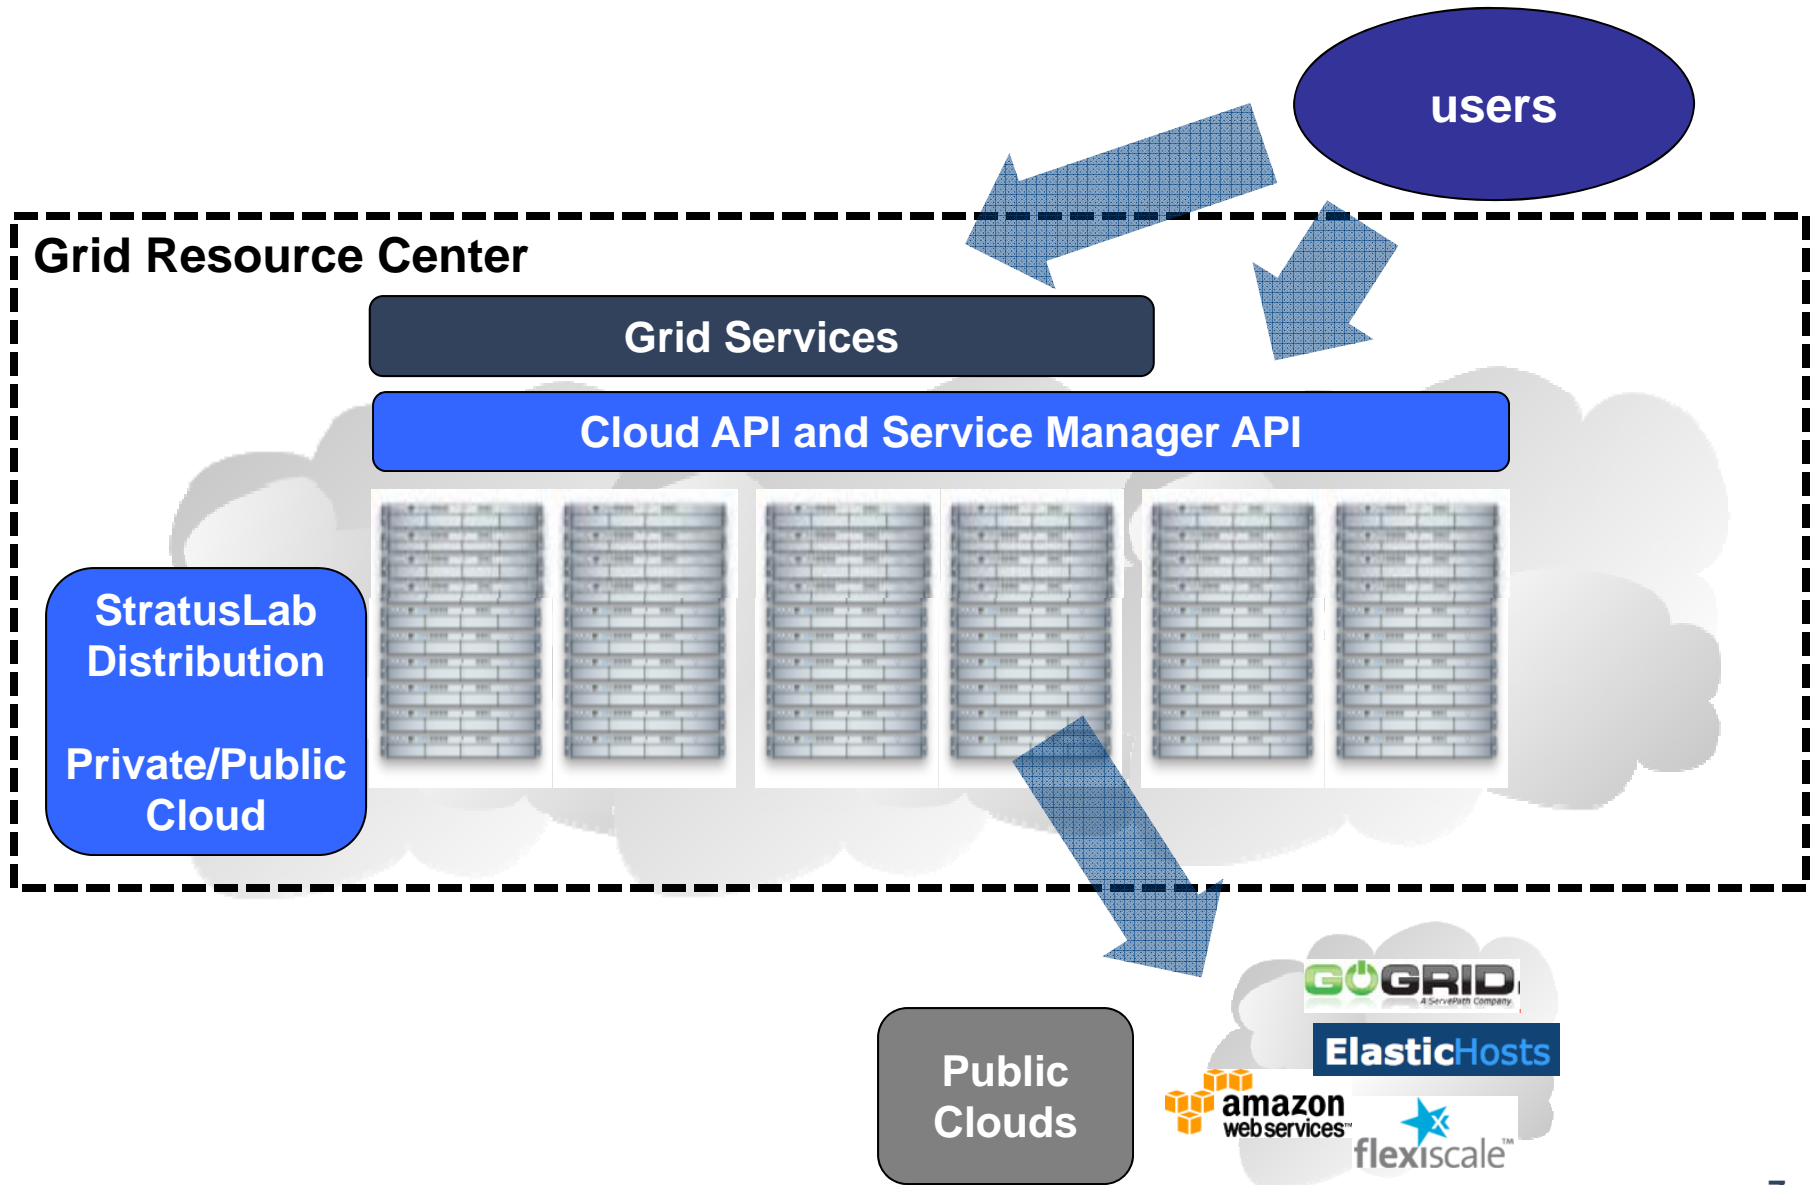
# EGI is Center of DCI Ecosystem



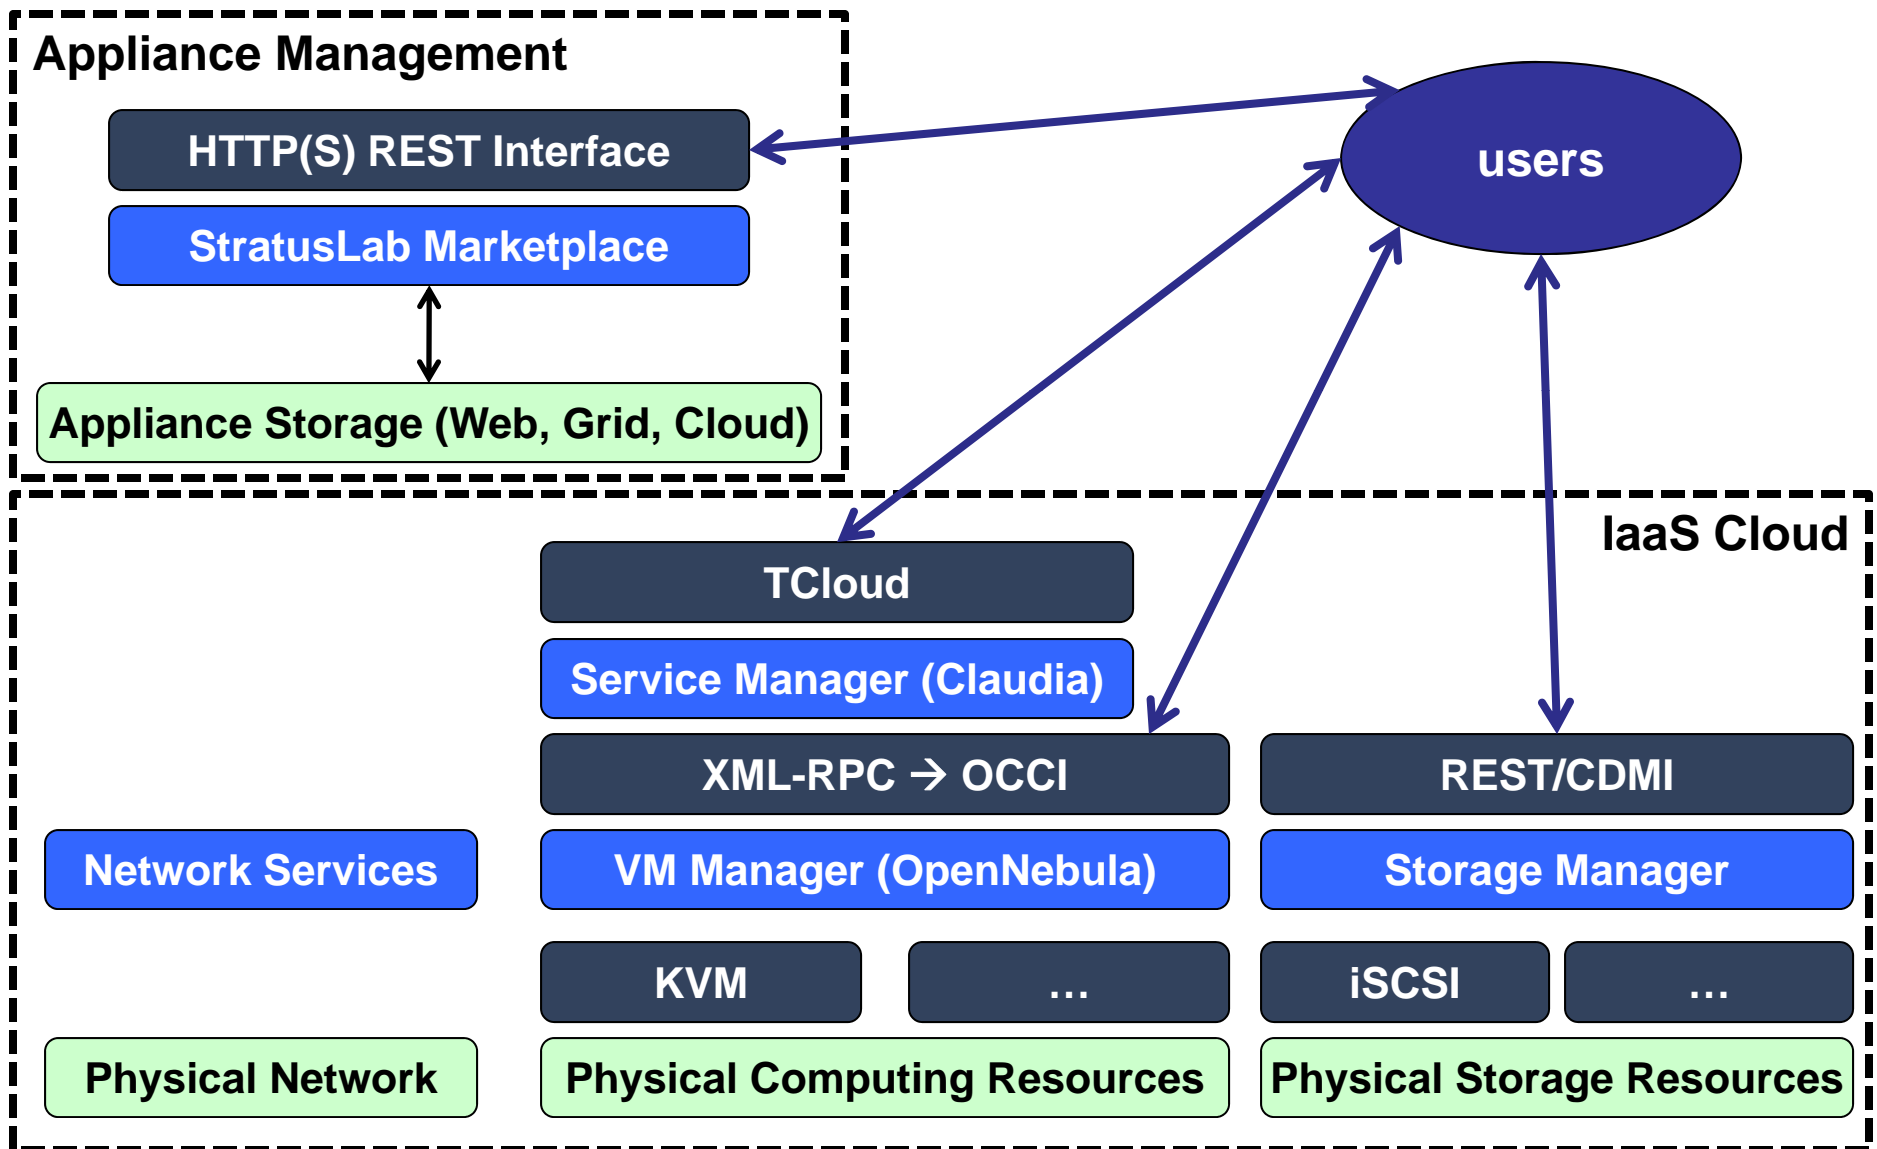
Metric (Jan. 2011)	Yearly Increase
13800 users	38%
288000 CPU cores	19%
117 PB disk	193%
92 PB tape	50%
28 x 10 <sup>6</sup> jobs/month	87%
340 sites	7%
56 countries	8%
217 VOs	24%
30 active VOs	~0%



# Grid and Cloud Integration



# Architecture of StratusLab v1.0



## OpenNebula

- Provides core of virtual machine management (start, stop, kill)
- Plug-in architecture allows use of multiple hypervisors (kvm, ...)

## Enhancements

- Quarantine of stopped images for forensic analysis
- Improved logging of user and resource information
- Ability to pass error messages from plug-ins to user
- Improved fault tolerance
- Improved management of network addresses
- Support for users, groups, and roles (post-1.0)

## Persistent (Read-Write) Disks

- Allows the storage of service state or user data
- Mounted as a disk on VMs
- Disks are persistent and have a lifecycle independent of a single VM
- Can be mounted by single VM at any time
- Only available within a single cloud instance

## Static (Read-Only) Disks

- Useful for distribution of quasi-static databases
- Handled like VM images via Marketplace

## File-based Storage

- Expect grid users to continue using SRM-related tools
- Will implement cloud-based file storage in Y2 of project

## IP Address Classes & Selection

- Public: Internet-accessible services
- Local: Batch systems or parallel calculations
- Private: Slaves in pilot job systems

## Future Services

- IP address reservation
- User specified firewalls
- Dynamic VLANs
- IPv6 use/validation

## **Machine image creation is a barrier to cloud adoption**

- Creating virtual machine images is time-consuming
- Ensuring that machines are secure and correct is difficult
- Sharing existing machines lowers this barrier

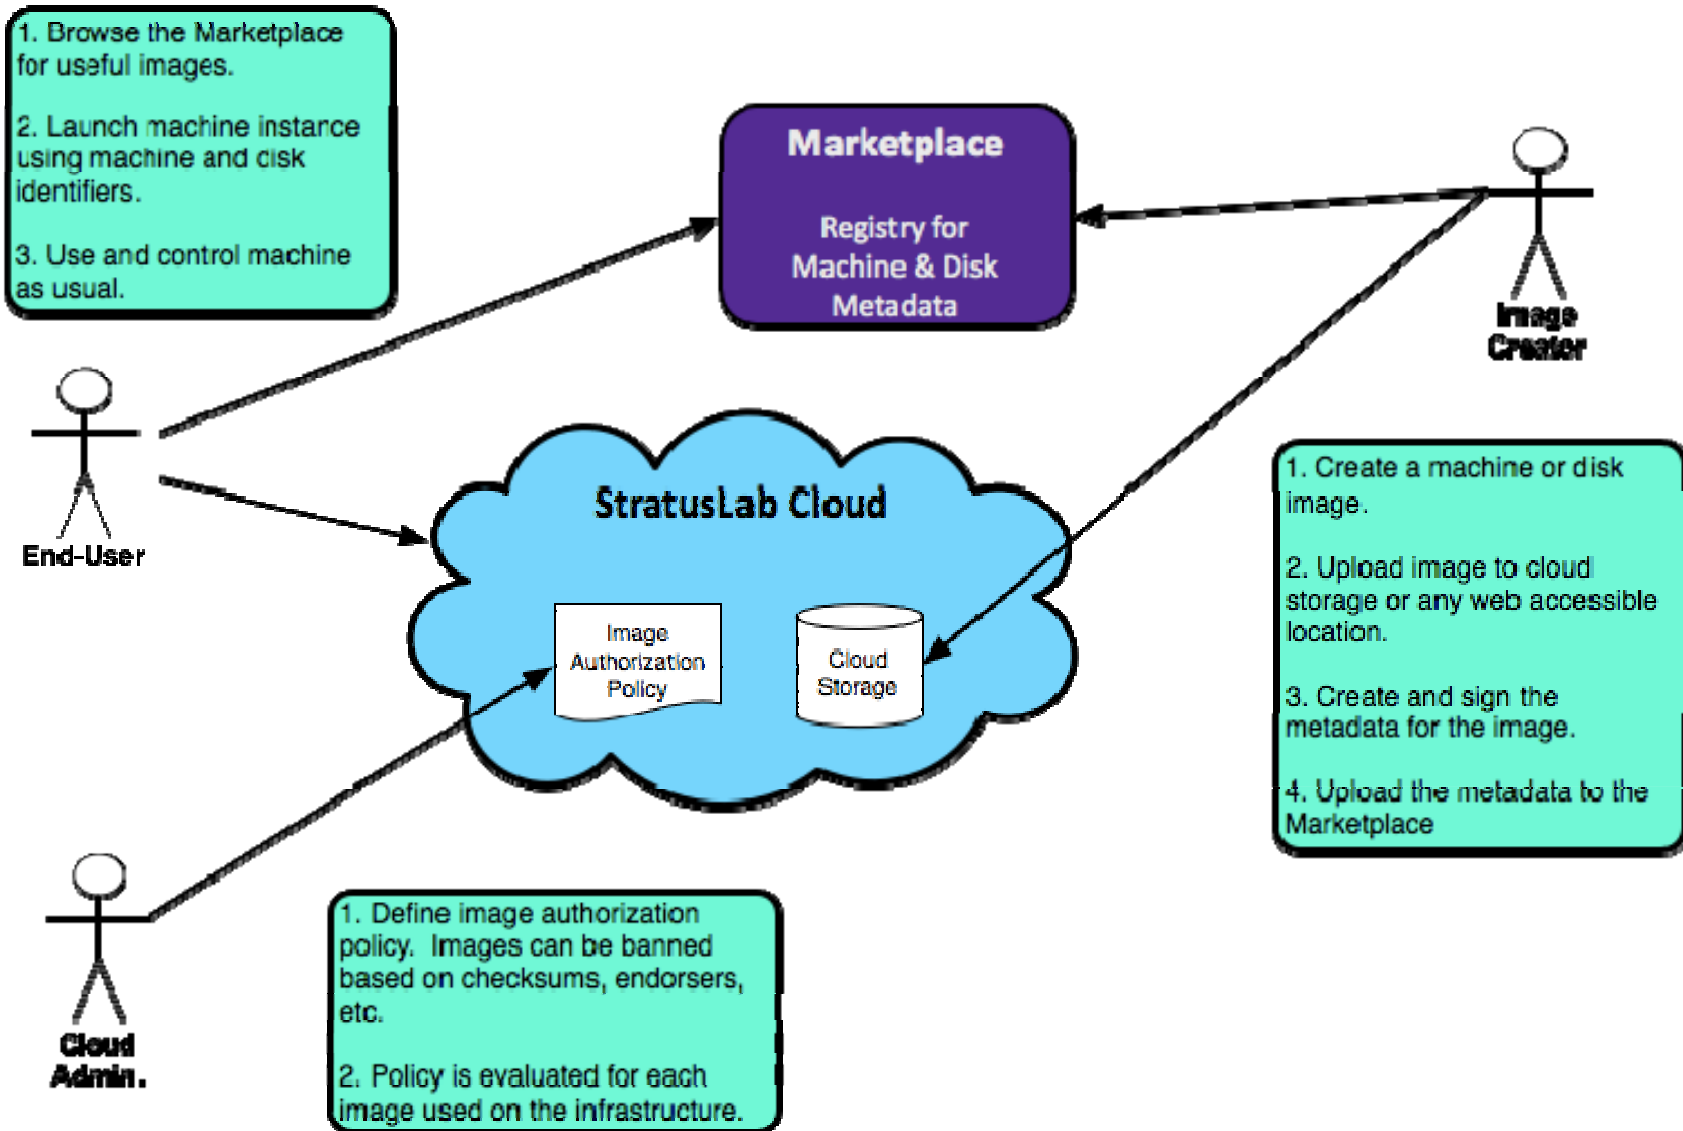
## **Marketplace facilitates sharing of images**

- Registry of metadata for machine & disk images
- Image contents are kept in cloud, grid, or web storage
- Supports trust between creators, users, and administrators

## **Benefits**

- End-users: browse and use existing images for their analyses
- Creators: publicize their work and attract larger user base
- Cloud Admins.: Use metadata to evaluate trustworthiness of images

# Marketplace Workflows



## **Claudia: Service Manager**

- Provides service (ensemble of machines) management
- Dynamic control/scaling based on monitoring feedback
- Currently being integrated with other StratusLab services

## **Authn/Authz**

- Authentication done through common proxy service
- Allows username/password from LDAP or from file
- Allows use of grid certificates and VOMS proxies
- Authorization done in individual services
- Delegation currently not needed/used (will change if machine or disk images are protected)

## Registration Service

- Web service for user registration
- LDAP DB for easy integration with cloud and other services

## Accounting/Monitoring

- Ganglia for monitoring of physical and virtual infrastructure
- Simple scripts to extract accounting information for reports
- No publication of the information for the moment

## Reference Infrastructure

- Open to public (with registration) to provide feedback to project
- Send an email to [support@stratuslab.eu](mailto:support@stratuslab.eu)
- GRNET runs a production grid site over a StratusLab cloud

## StratusLab Client

- Command line scripts in python with few dependencies
- Works on Mac OSX, Windows, and Linux

## Provided Appliances

- Base images: ttylinux, CentOS 5.5, Ubuntu 10.04, OpenSuSE (?)
- Grid: CE, SE, WN, APEL/BDII, UI
- Bioinformatics: Data server and analysis images

## Placement Service

- Needs to find (quickly) appropriate resource for starting a VM
- Only available resource information needed for decision

## Optimization Service

- Load balancing
- Energy optimization
- SLA management
- Very large variety of data to support such services

## Infrastructure

- Messaging (eventually a service) to obtain information
- API to modify locations of VMs

# "Infinite Capacity"



**Cloud usage models really require the illusion of infinite capacity.**

**How can this be achieved with (very) limited resources?**

## **Queuing**

- Spot instances like Amazon?
- "Hidden" use of classical queuing systems, fair share, ...?

## **Federation**

- "Grid-like" federation: expose multiple resources to users
- Transparent federation: cloud administrators manage federation

## **Bursting**

- Use of commercial providers to augment resources

## Storage Abstraction

- Provide disk and file-based abstractions
- Better higher-level abstraction?

## Federation Issues

- Optimized access: always complex interplay between local compute, network bandwidth, and data access pattern
- Security, privacy, and legal issues

## Network Abstraction

- Connectivity is the real need, but needs to be dynamic and simple
- What to expose to users?
  - Firewalls, IPv4 vs. IPv6, bandwidth requirements
- No common API and no common description of network configuration

**Need to have hardware and software that better supports dynamic configuration of networking.**

## PaaS based on Application Patterns

- Hadoop
- Clusters
- Grid
- 3-tier systems

## Scientific PaaS infrastructures

- Data access and data are critical
- How can communities offload effort to standard patterns/services?
- How well can details be hidden and parameterized?

## StratusLab Release(s)

- Nearly complete, functional beta available
- Used to run production grid site over reference infrastructure
- Complete beta (v0.4) ~16 May
- Production release (v1.0) ~10 June

## Continued Evolution

- Further enhance distribution through developments and collaborations
- Looking especially for collaborations on networking aspects
- Software can be tested on reference infrastructure or installed
- Want feedback (positive and negative) from user, administrators, ...

**Copyright © 2011, Members of the StratusLab collaboration: Centre National de la Recherche Scientifique, Universidad Complutense de Madrid, Greek Research and Technology Network S.A., SixSq Sàrl, Telefónica Investigación y Desarrollo SA, and The Provost Fellows and Scholars of the College of the Holy and Undivided Trinity of Queen Elizabeth Near Dublin.**

**This work is licensed under the Creative Commons  
Attribution 3.0 Unported License  
<http://creativecommons.org/licenses/by/3.0/>**

