

## StratusLab Cloud Distribution

ACGRID-III (Hanoi)

1 November 2011



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



## 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)**

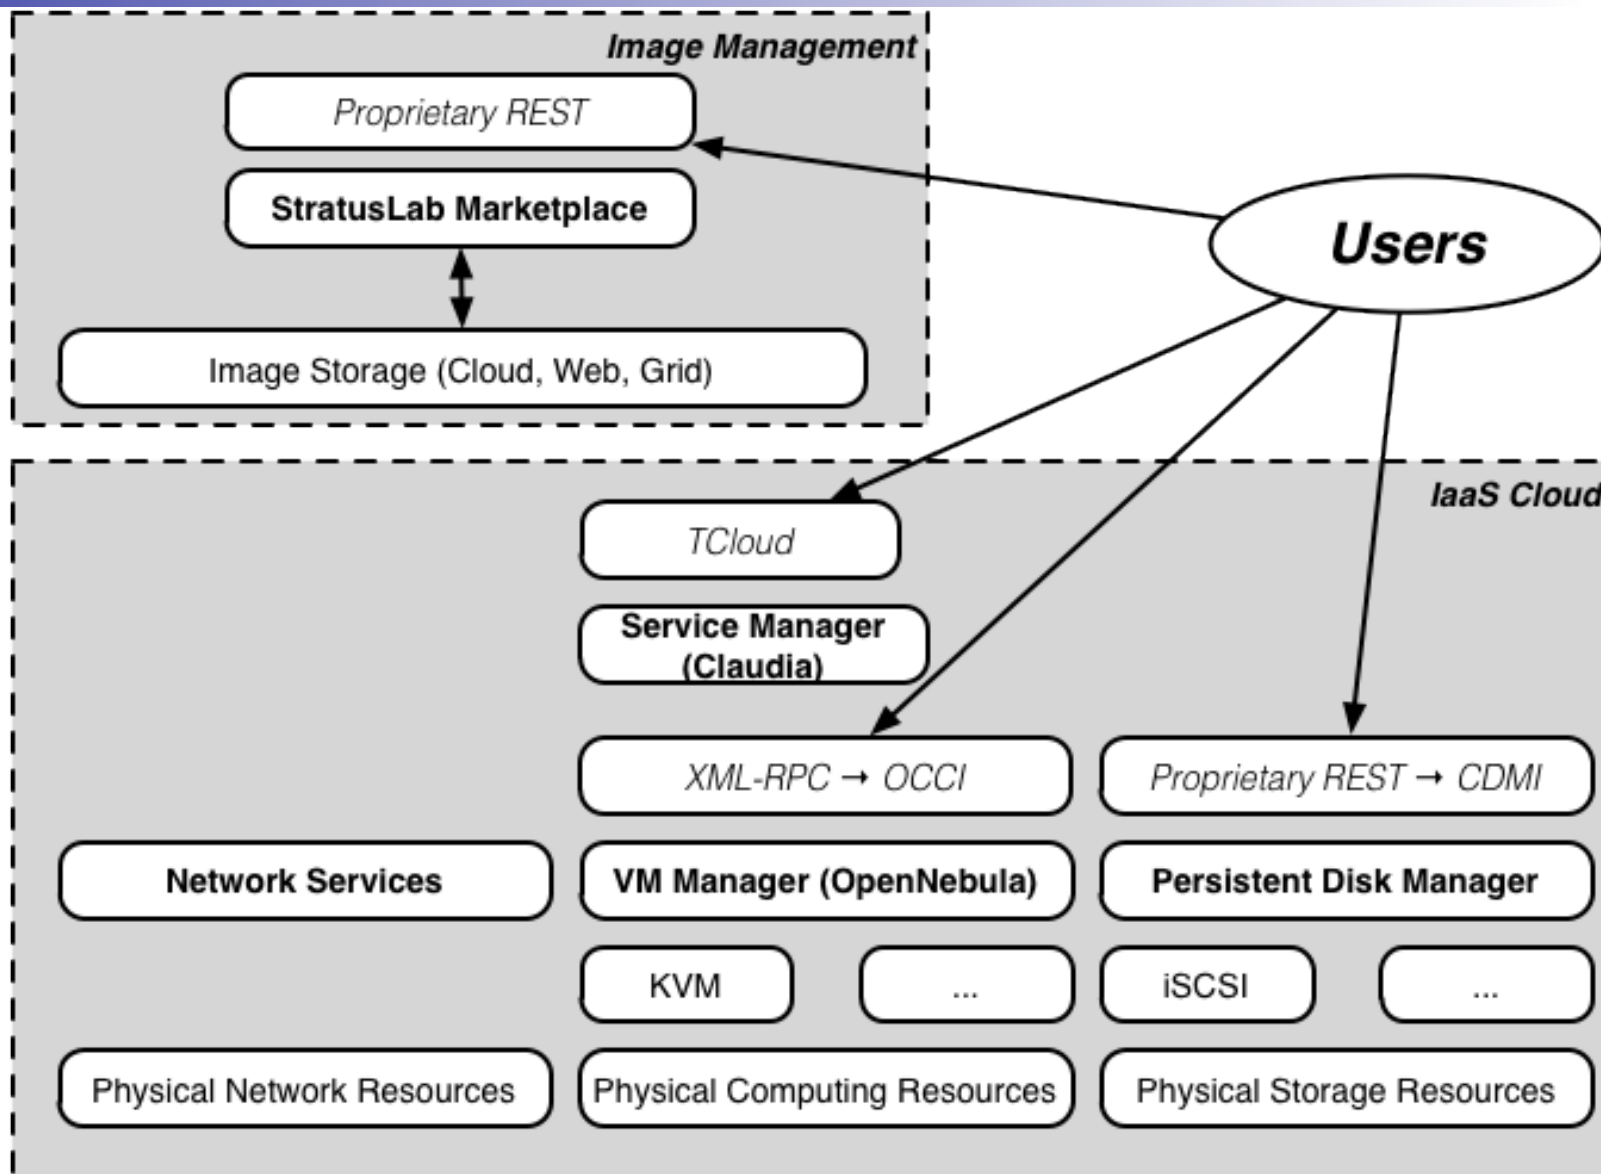
**Infrastructure as a Service (IaaS)**



**ElasticHosts**



# StratusLab Architecture (v1.0)



## OpenNebula ([opennebula.org](http://opennebula.org))

- 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 and shared like VM images via Marketplace

## **Volatile (Read-Write) Disks**

- Useful for temporary (!) data storage
- Data will disappear when VM instance is destroyed

## File-based Storage

- Normal client tools can be installed in VMs
- Access services normally from VM (e.g. tools for SRM)
- Unlikely to be implemented by StratusLab, although will provide CDMI interface to persistent storage service

## Object Storage

- Simple object storage, usually minimal hierarchy and chunked data
- Won't implement this in StratusLab, could take implementations from elsewhere, e.g. OpenStack

## Key-value Pair Database

- Exposes simple API for “database” of key-value pairs (e.g. Cassandra)
- Can deploy VM with persistent disk to provide this service



## 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

## **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

## StratusLab Client

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

## Programming Interfaces

- Most services provide REST interfaces
- Easy to program from any language
- Straightforward resource  $\leftrightarrow$  URL mapping
- Standard interfaces will be implemented (OCCI, CMDI, ...)

## Libraries

- jclouds, ...?

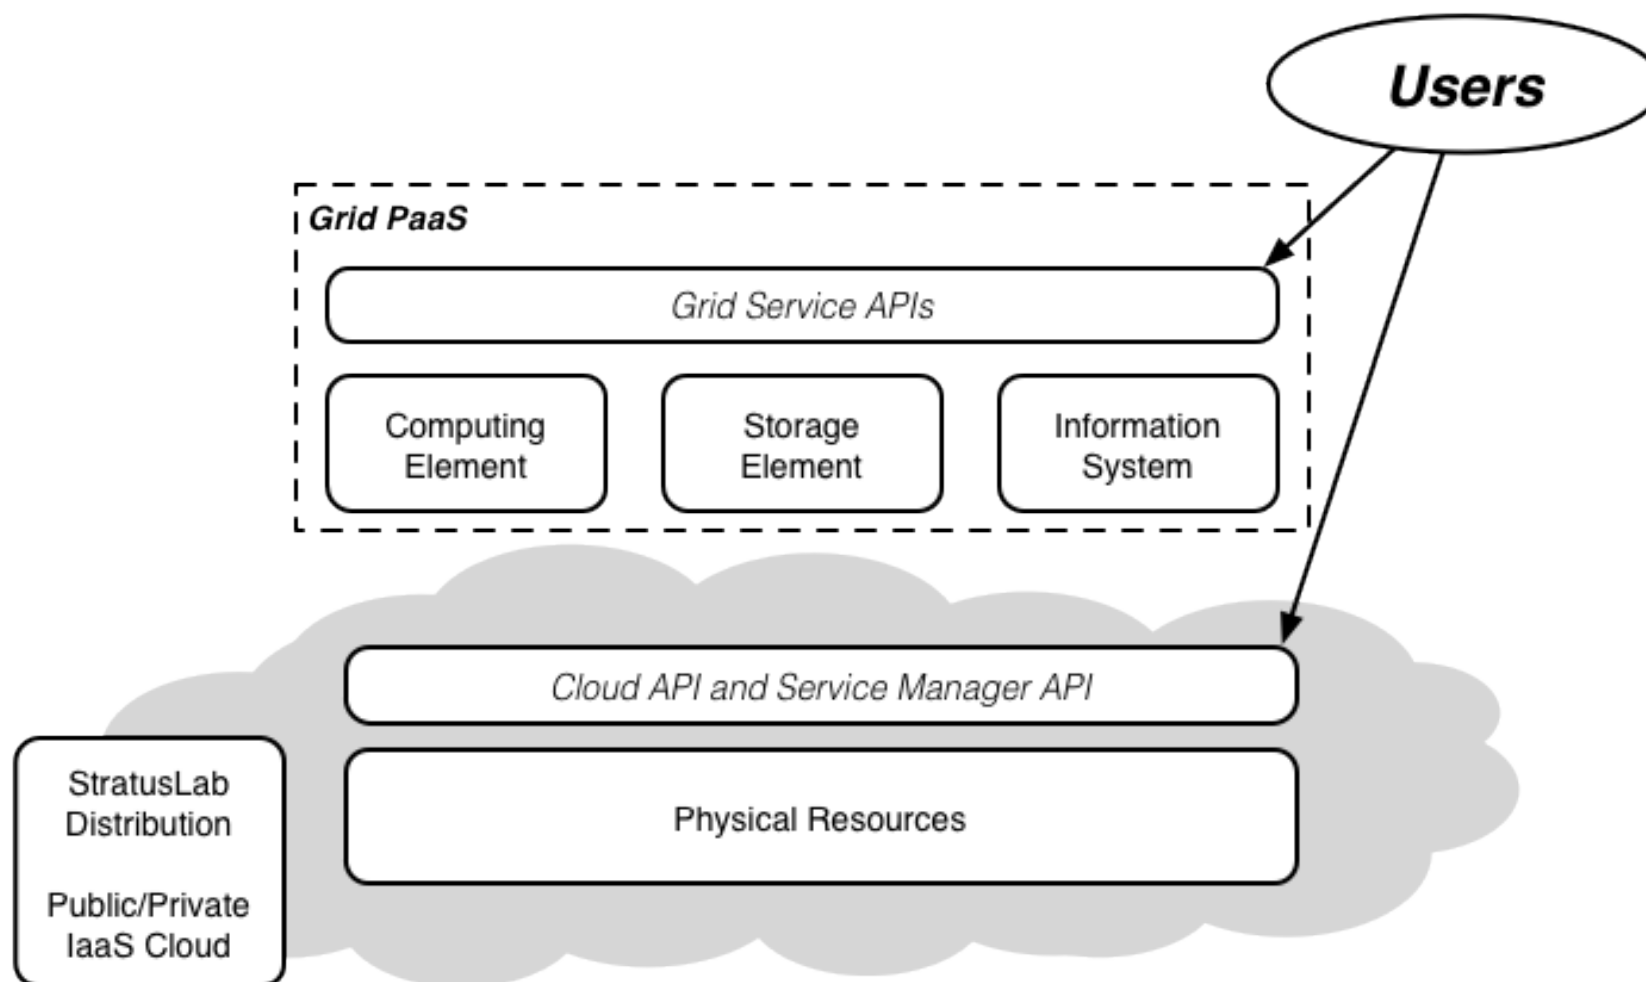
## Appliances

- Provide pre-configured, pre-installed software and services
- Makes it easier to get started quickly using cloud resources
- Can make and publish your own appliances

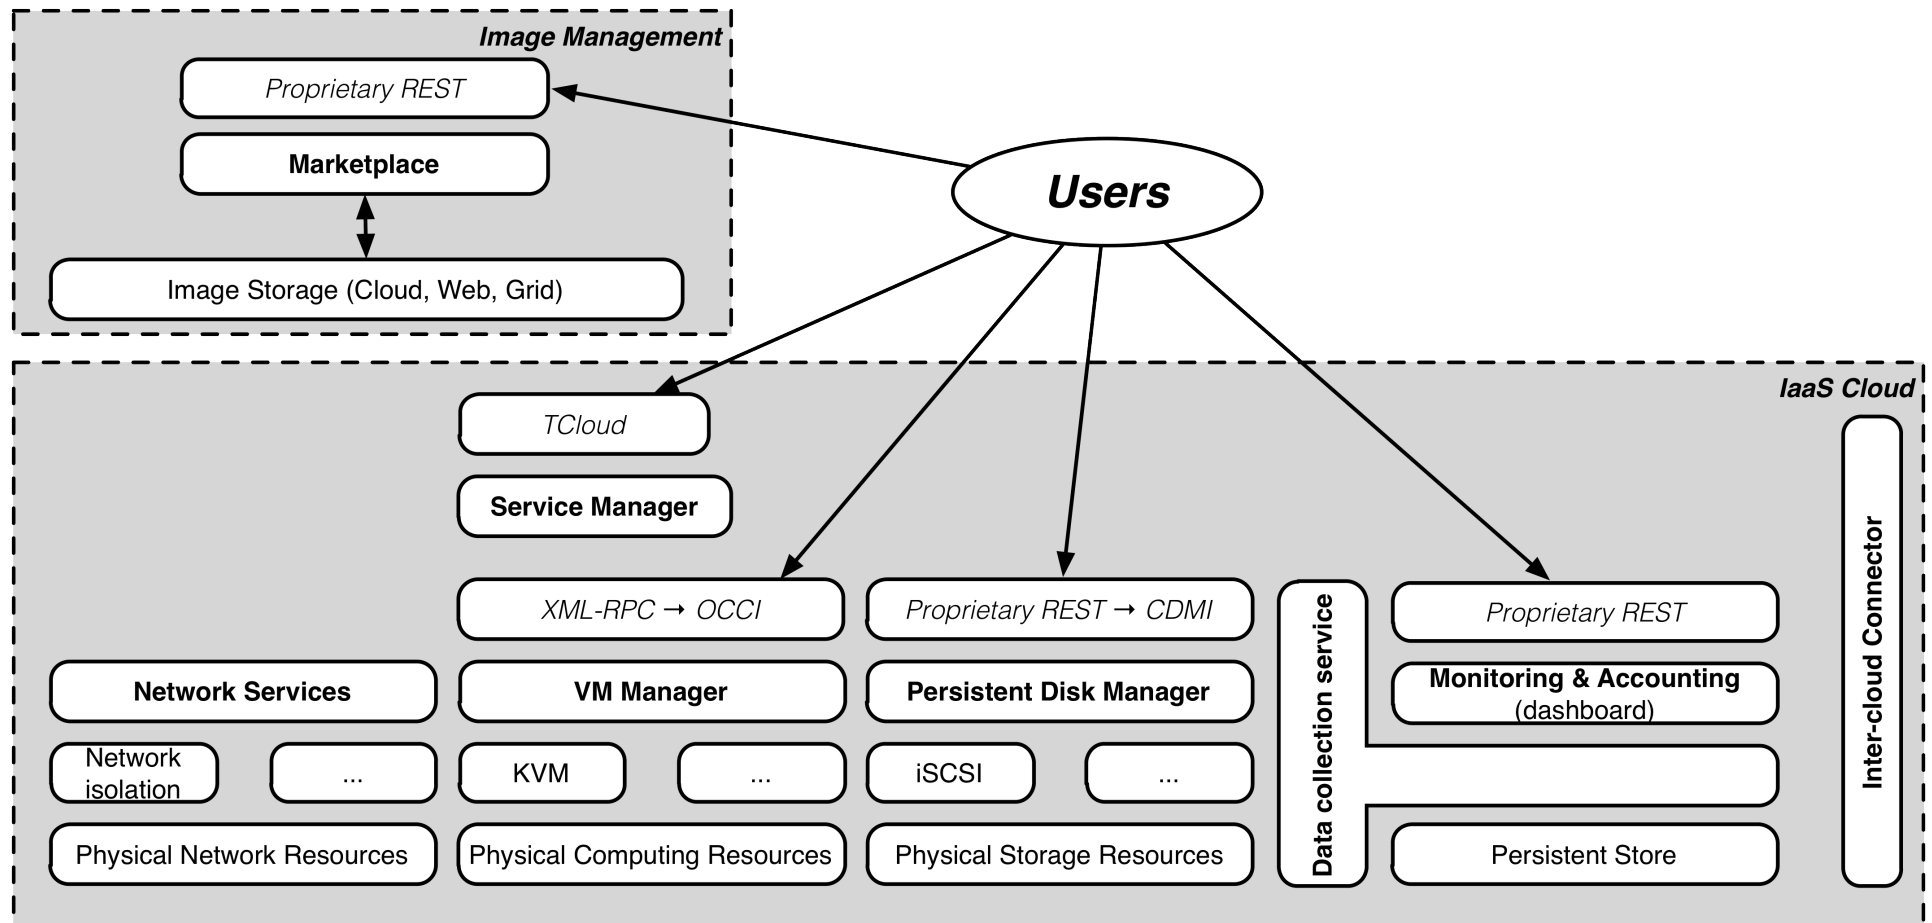
## StratusLab-Provided Appliances

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

# Grid and Cloud Together



# StratusLab Architecture (v2.0)



**Existing releases provide complete IaaS cloud solution.**

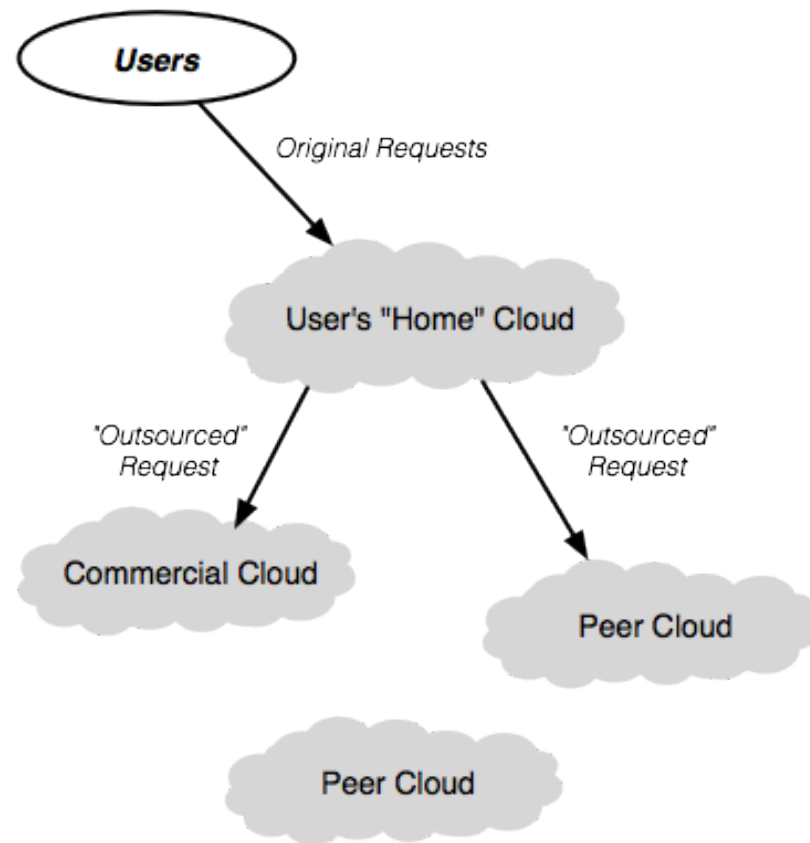
**Future work concentrates on monitoring, accounting, and federation.**



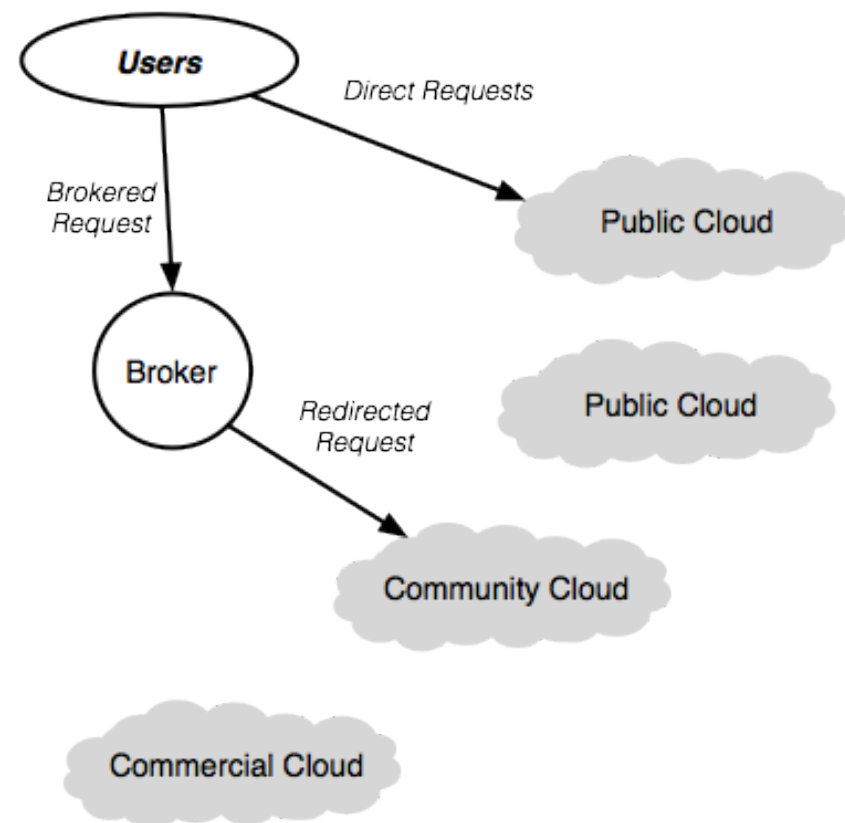
# Hybrid Clouds and “Sky” Computing



## *Peer Federation or Bursting*



## *Brokered Federation*



# Questions and Discussion

**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/>**

