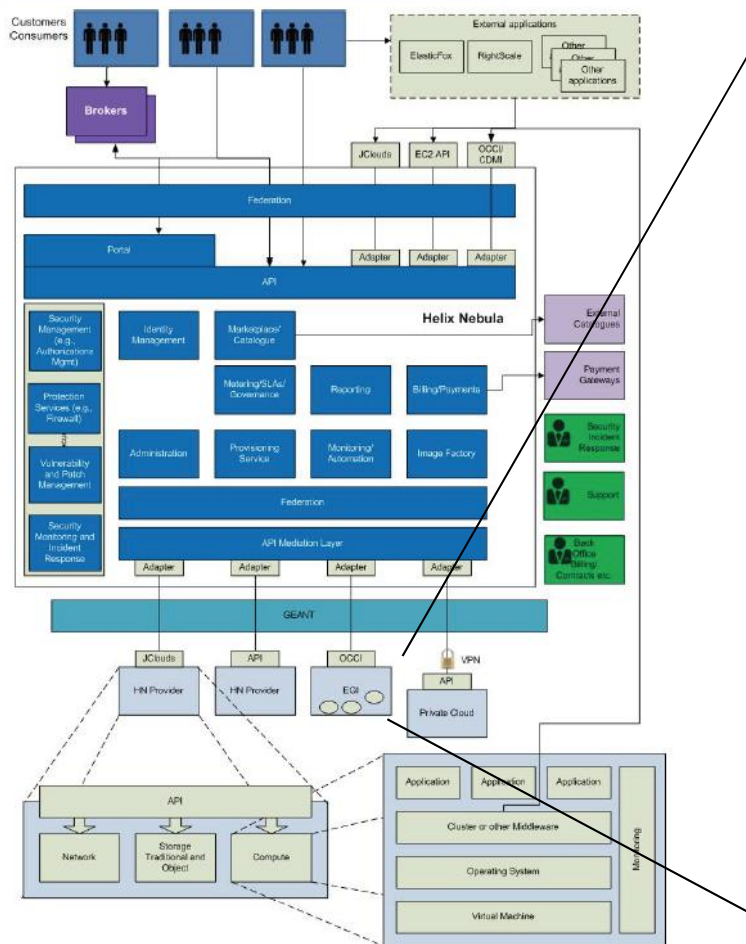


# Deploying Helix Nebula flagship applications in a hybrid scenario

Salvatore Pinto  
Cloud Technologist  
EGI.eu

- Helix Nebula and the EGI Federated Cloud
- The ESA flagship on Helix Nebula
- ESA Proof of Concept on the EGI Federated Cloud
- Next Steps

HELIX NEBULA - THE SCIENCE CLOUD **Tech Arch Series**



The EGI Federated Cloud is federation of institutional private Clouds, offering Cloud Infrastructure as a Service to scientists in Europe and worldwide.

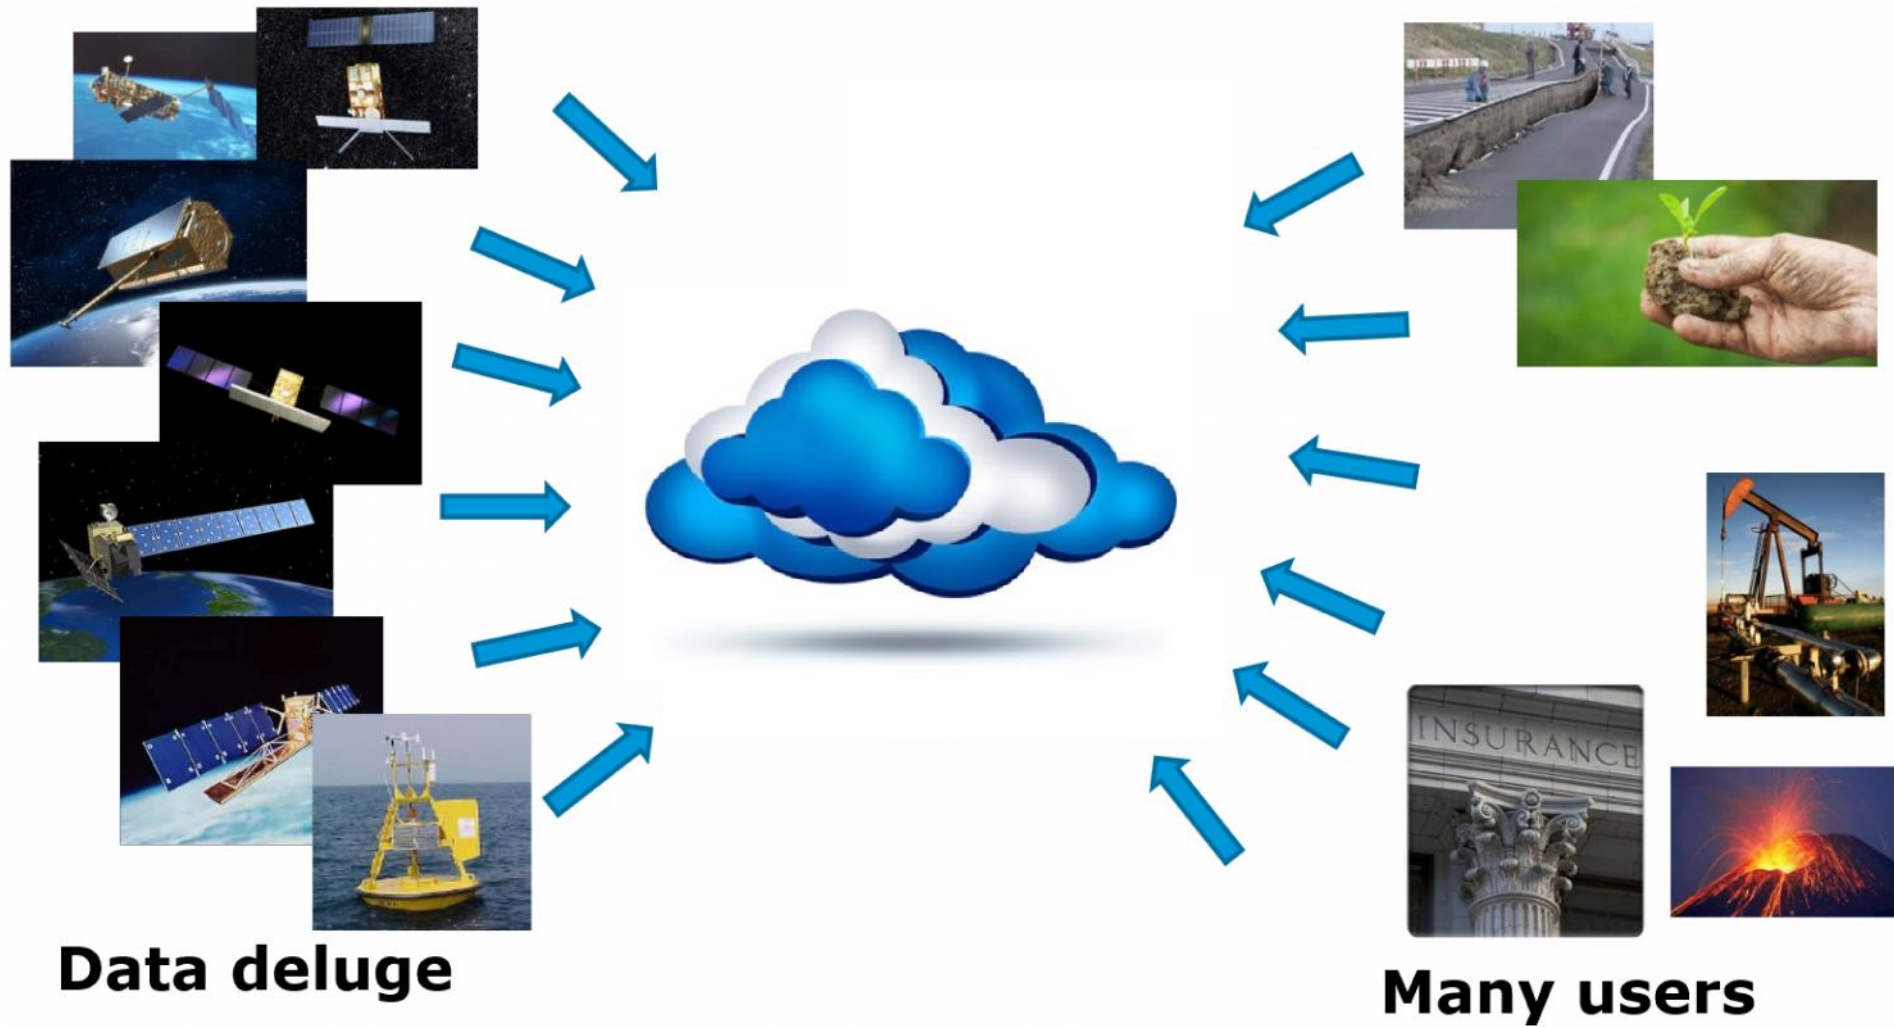
- **Standards and validation:** federation is based on common open-standards – OCCI, CDMI, OVF.
- **Heterogeneous implementation:** no mandate on the cloud technology, the only condition is to expose the chosen interfaces and services.

Considering the Helix Nebula, **Blue Box** approach, the obvious interoperability solution for VM management is to add a OCCI connector to the Blue Box



We concentrated our effort on one of the Blue Box solutions under test, **SlipStream**, which is an Open Source solution and has an EGI-InSPIRE mini-project to support OCCI cloud connectivity.

NOTE: enStratus proposed Blue Box does not support OCCI at the time





# The ESA flagship (SSEP)

ESA flagship is integrated into the ESA's **S**uper**S**ite **E**xploitation **P**latform project, which exploit cloud resources to

1. Provision IaaS resources from different cloud providers as well as private cloud resources (**hybrid multi-sourced cloud**)
2. Allow users (scientists) to search and **access/download data at high performance** and data providers (ESA and other agencies) to upload data.
3. Allow users (scientists) to process data on **dedicated sandboxes**; for algorithm development and small data processing
4. Provide users (scientists) access to **large data processing** where to deploy and exploit their applications via a portal.
5. All nicely integrated with a **common security framework**

More info at: <http://ssep.eo.esa.int>

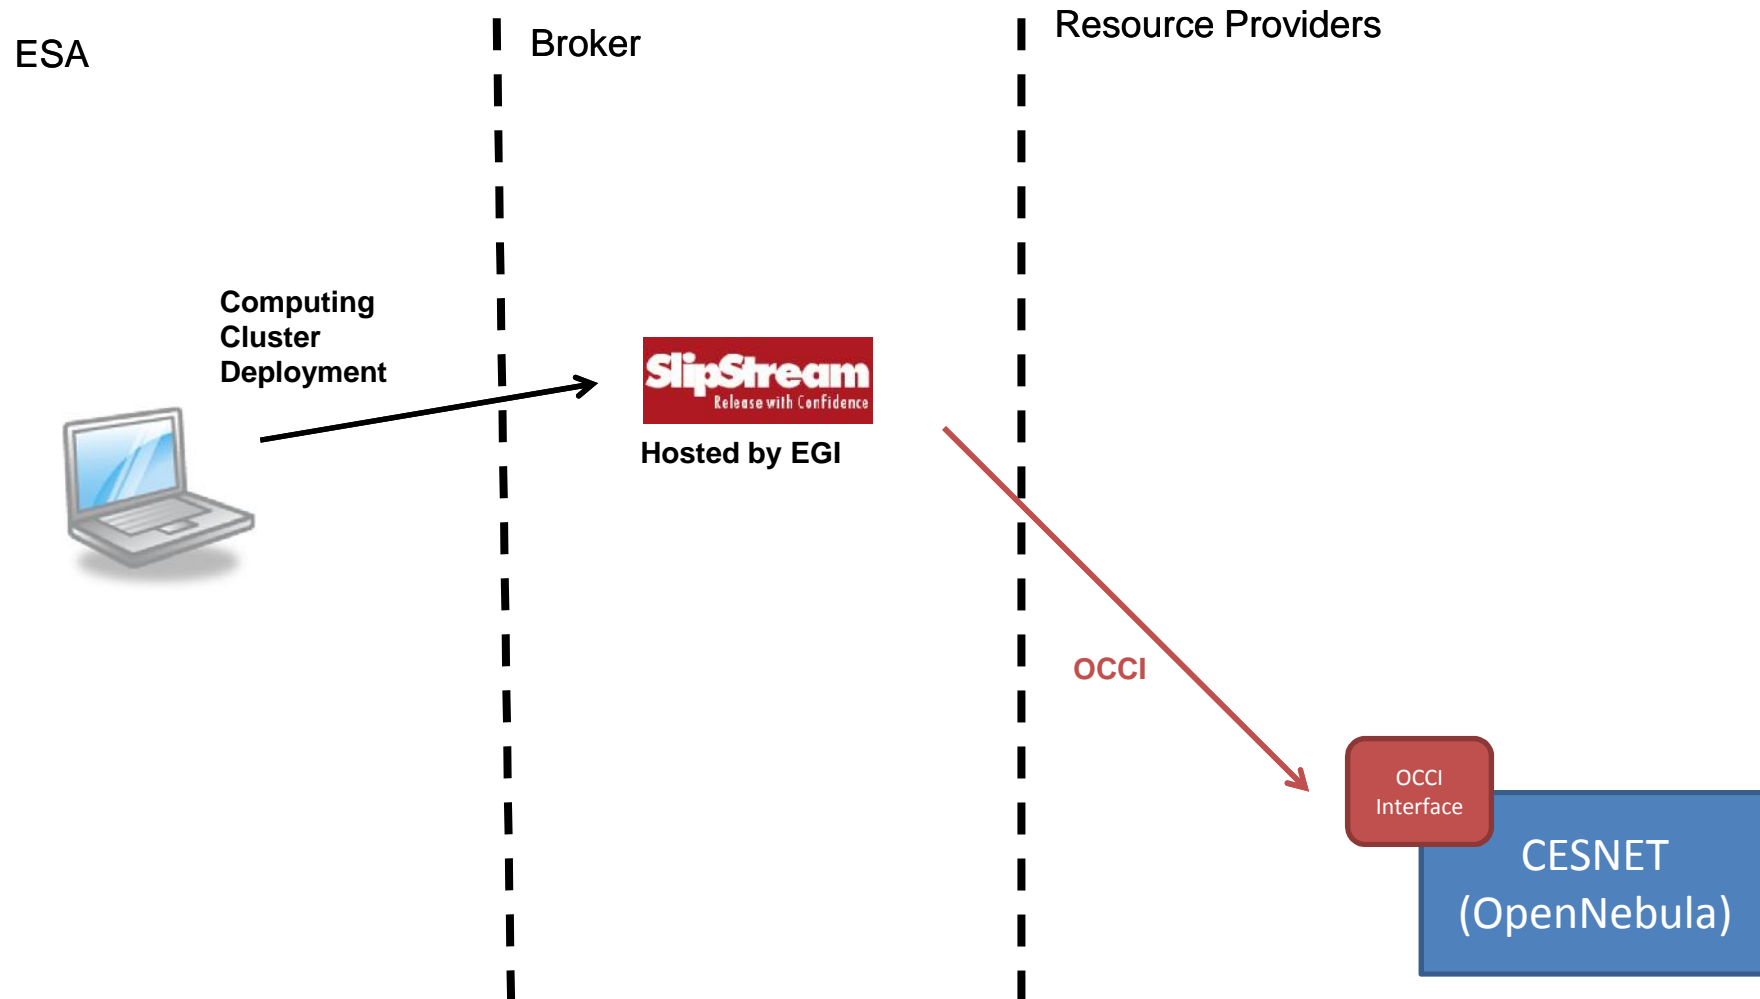


# ESA PoC on EGI Federated Cloud

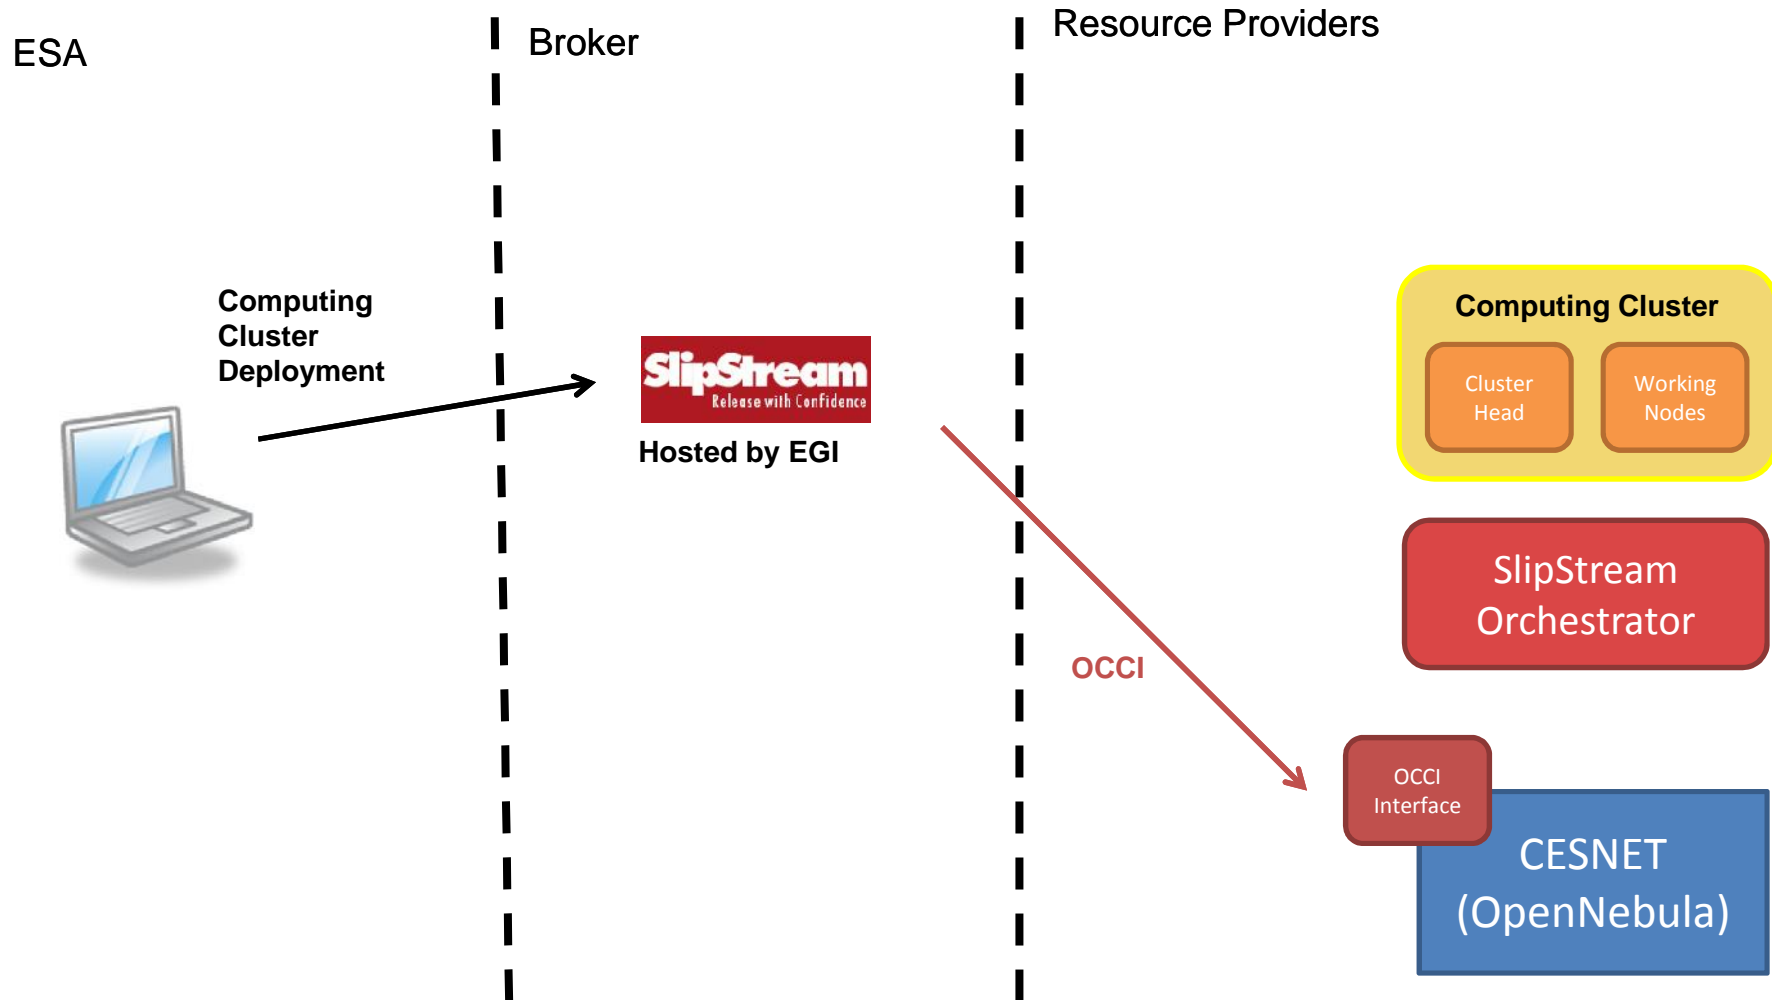
The ESA Proof of Concepts on the EGI Federated Cloud focuses on demonstrating the possibility to provide Processing Services to ESA scientists using the EGI Federated Cloud resources

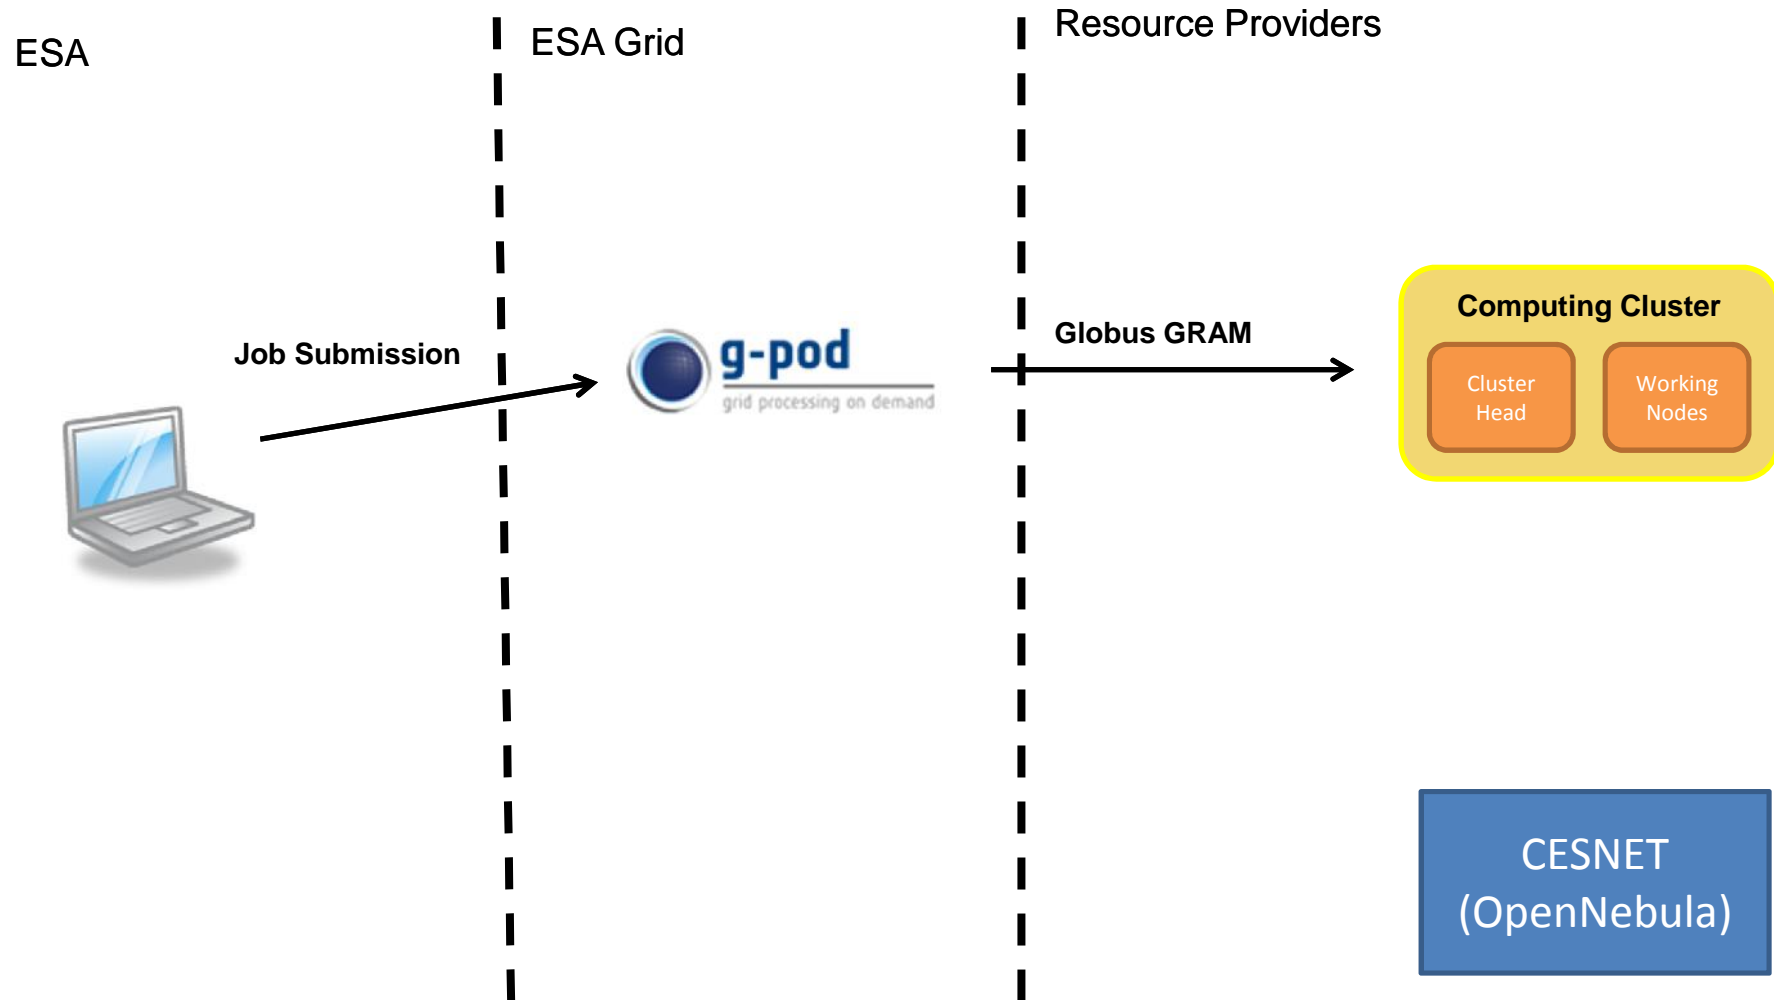
- Participants:
  - User community
    - **ESA Research and Service Support** section: Configuration and execution of the tests
  - Technology providers
    - **Sixsq**: Provided Open Source SlipStream software and OCCI connector
  - Resource providers
    - **CESNET** (OpenNebula): Performance tests and multi-site tests
    - **GRNET** (synnefo): Multi-site tests
    - **CESGA** (OpenNebula): Hosting of the SlipStream server

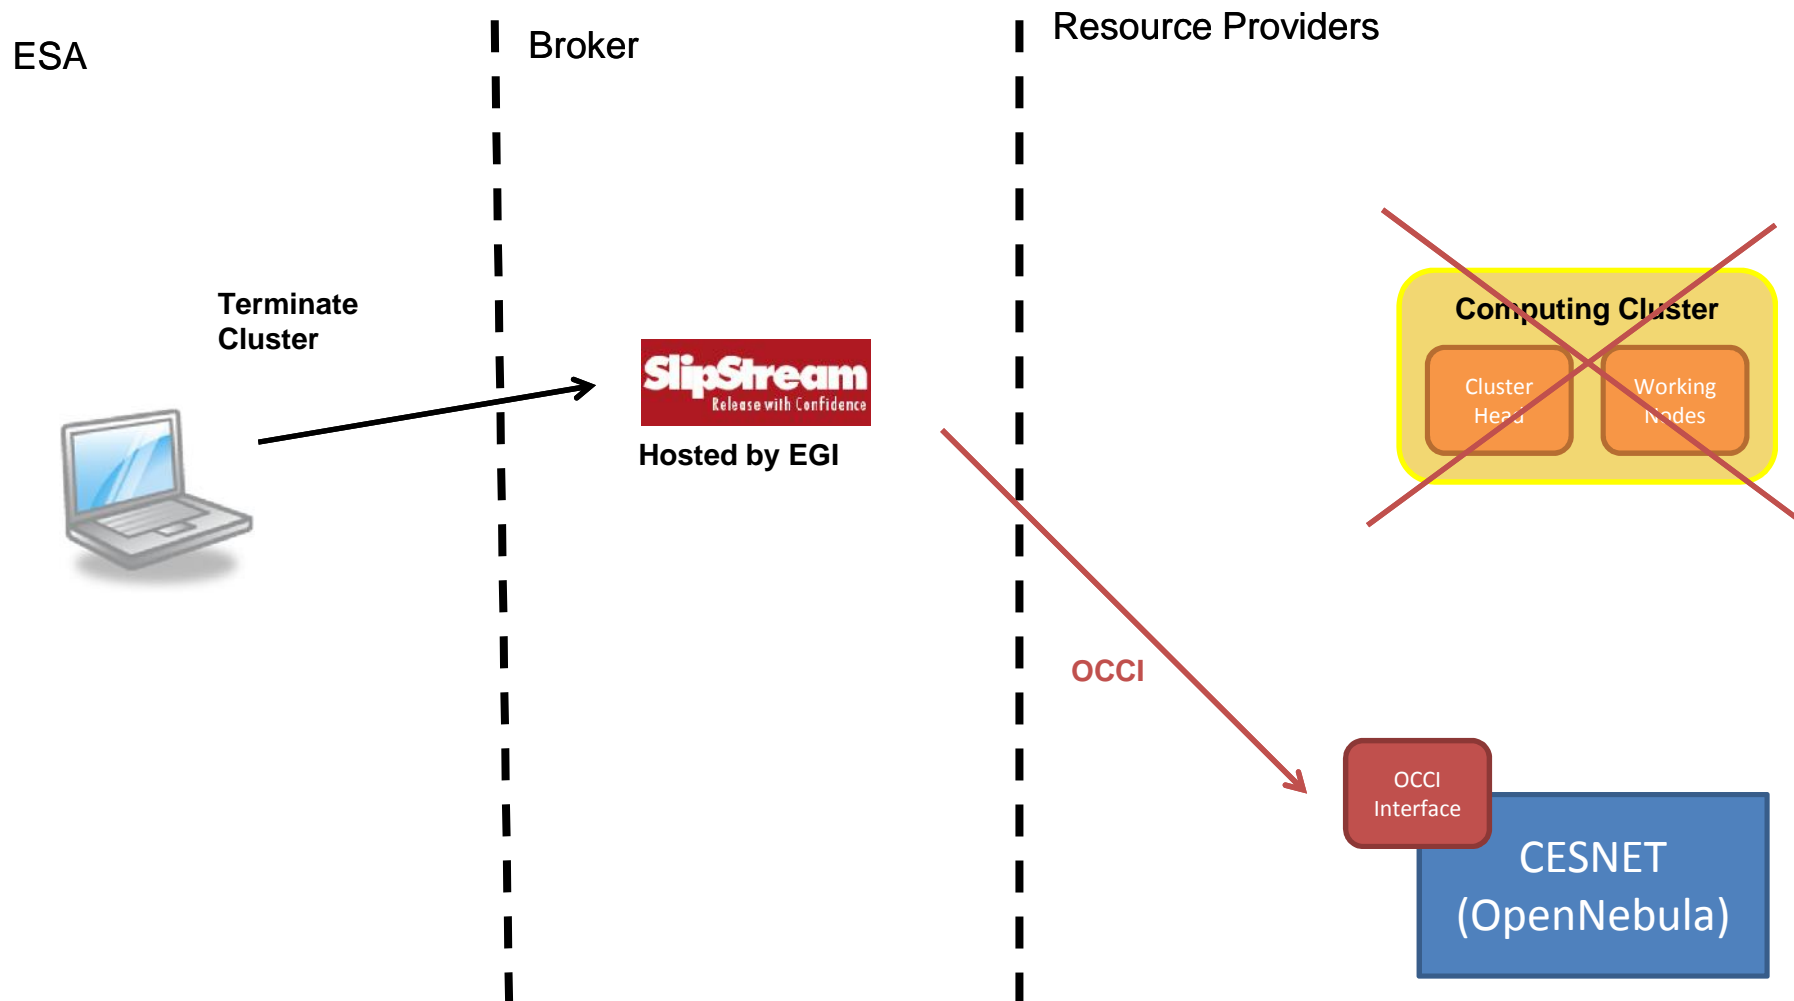




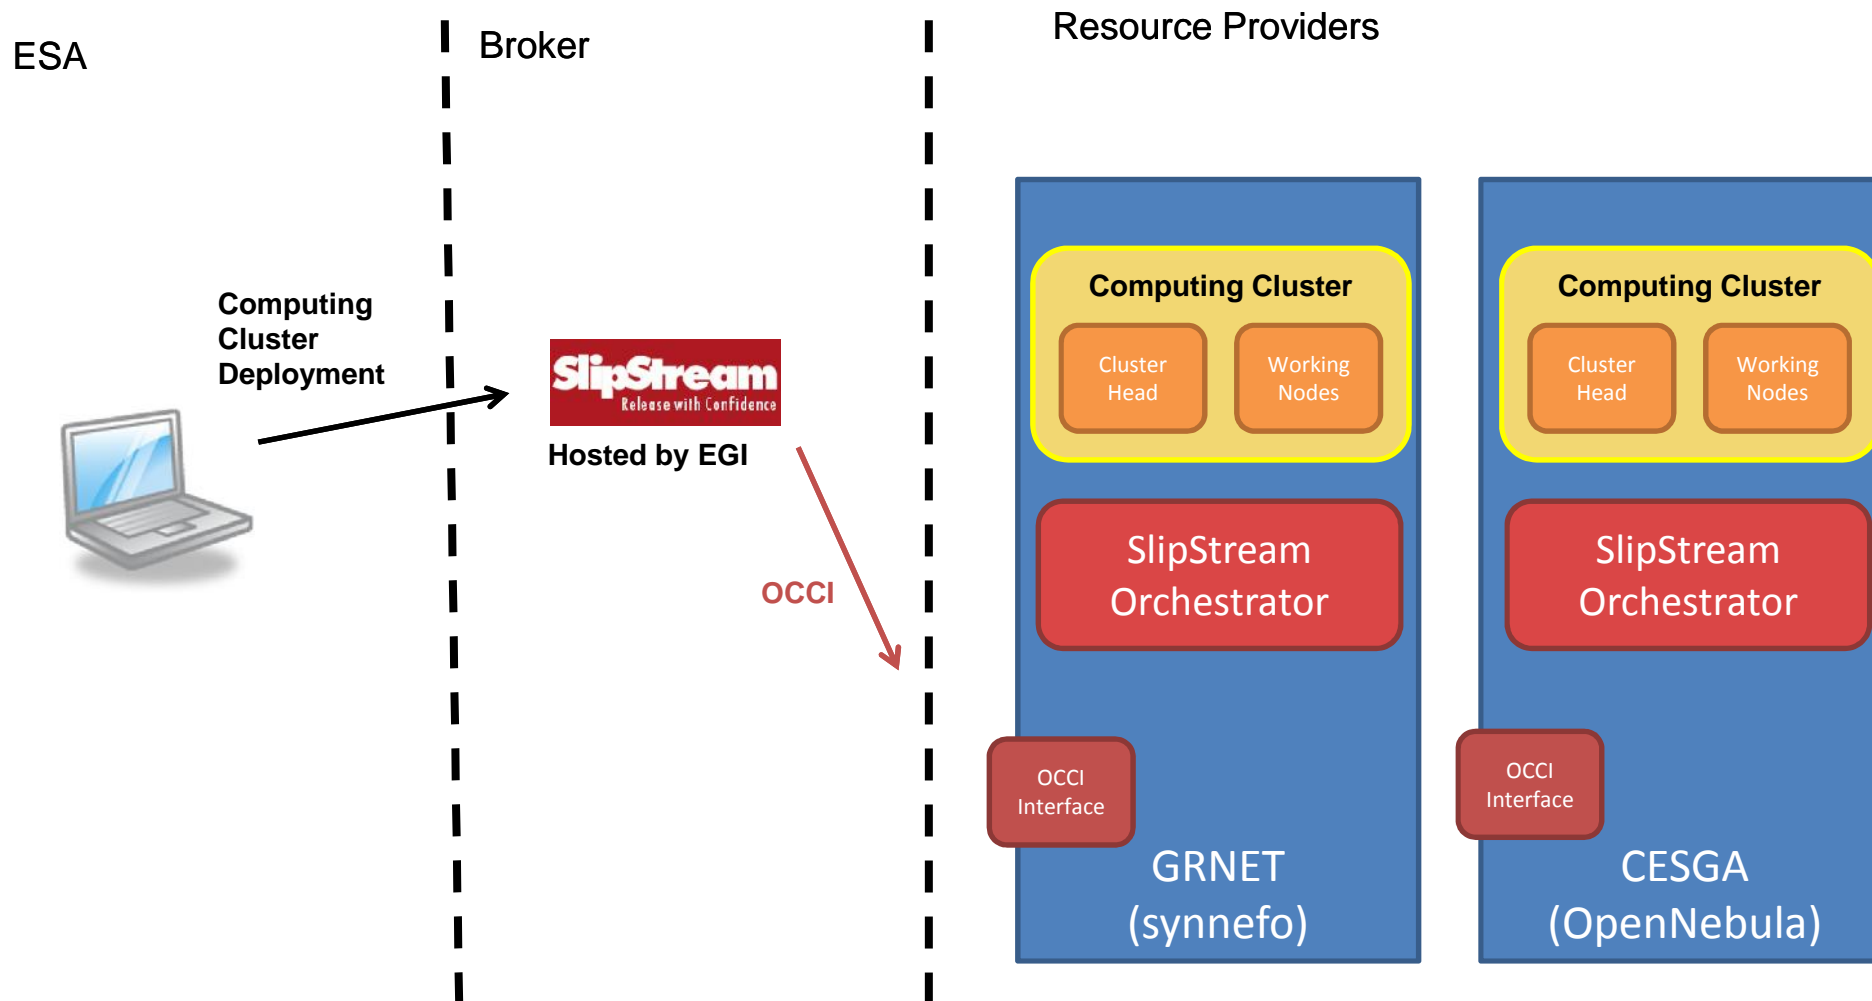


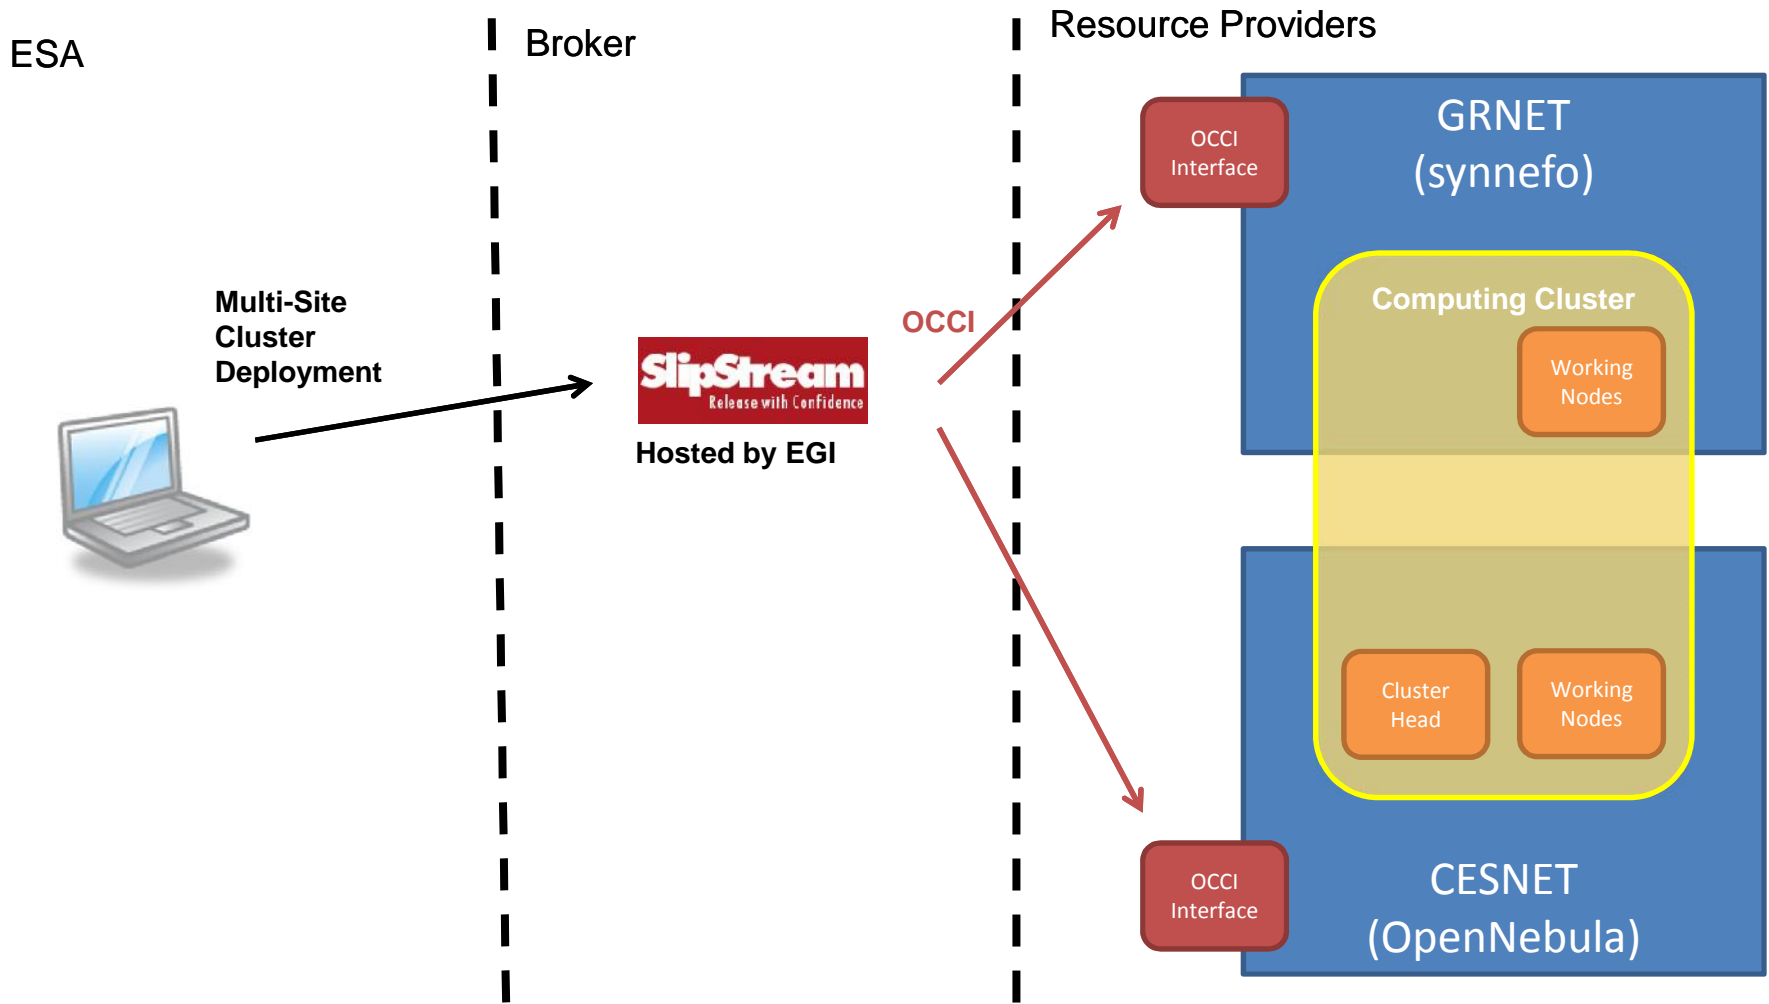






# ESA PoC Scenario 1: Single site Processing Cluster (with different back-ends technologies)








- Automatic Deployment of Processing Cluster
  - Time required for complete startup: **22 minutes**



Time is inline with the expectation, considering that:

- All the software is installed and configured on startup, the image is a basic OS installation
- Cluster is started and then operated for at least one week
- The same procedure, performed manually, usually takes about 1 day of work

- VM creation/deletion using direct OCCI API ✓
- Job submission to the Computing Cluster ✓
  - Average job overhead: 5 seconds
- Computing Cluster end-to-end availability ✓
  - 99.56% during 5 days (31/7021 failures)
- Stress test (10WN Cluster under CPU/RAM/Disk stress) ✓
  - Average job overhead under stress: 5 seconds

- Internal network bandwidth 
  - Average 9.2 Gbit/sec
- Local storage performances 
  - Average: Write 102M/sec, Read 103M/sec
- Network latency (time to download an empty file using SCP) 
  - Averaged for internal network: 0.12 s
  - Averaged ESA to FedCloud: 0.6 s

EGI sample site performances are inline with Helix Nebula commercial providers



- Multi-Site Deployment of Processing Cluster



Cluster Deployment was correctly completed, but network issues in the sites connections caused the job submission tests to fail

Still to be investigated



- Proved **interoperability** between academic and commercial clouds.
- Globus-based computing cluster **deployment in one-click**.
- **Cloud technology agnostic**. SlipStream worked with the same connector and configuration on OpenNebula and synnefo sites.
- Test of site availability, performances, etc... (results shows provided resources are **inline with commercial providers**).

# Next Steps: CERN PoC (October 2013)

CERN



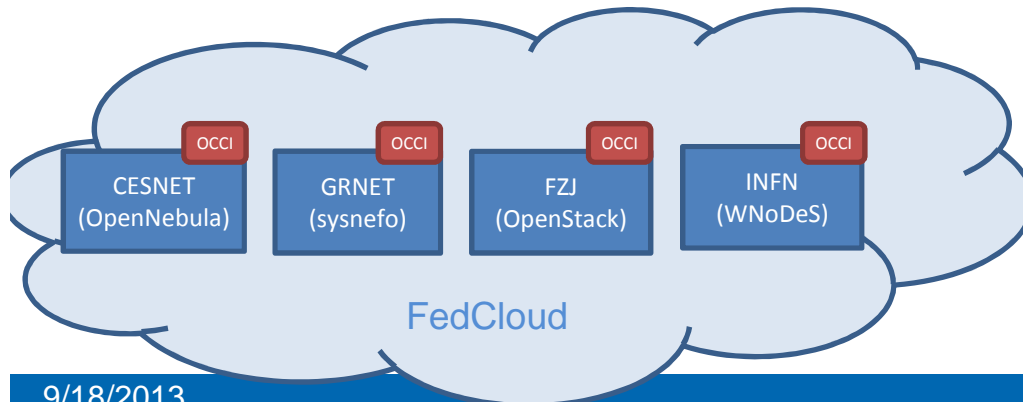
Computing Cluster Deployment

Broker



Resource Providers

OCCI



Commercial Cloud Providers



## Next Steps: SlipStream FULL connector (January 2014)

SlipStream FULL connector, providing

- Support for “**Image Creation**” feature
- Support for EGI **FedCloud Contextualization system** (OCCI User Data extension)
- Support for **attached storage** (block or object storage)
- Better **error catching** (to debug issues in the sites)
- Available by default in **SlipStream 2.0**

Thank you

Questions?

