Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

Models for Integrating Grid and Cloud Technologies

Both cloud and grid technologies have features that make them attractive for technical computing tasks. An improved e-Infrastructure can be provided to European researchers by combining both technologies on a single pan-European platform.

Cloud Features

- Dynamic, instantaneous provisioning of resources.
- Ability to customize execution environment and installed applications.

Grid Features

- Robust security and policy framework permitting global authentication and authorization of users.
- Mechanisms for federating distributed resources to produce a larger, more capable infrastructure for users.

A combined e-Infrastructure:

- Appeals to a larger, more diverse scientific community via customized environments and alternate operating systems.
- Allows VO and user-level services permitting the construction of complete scientific platforms and services.
- Allows sharing of same physical resources with different access modes without partitioning of resources.
- Provides more flexible management of services and resources.

Model 1: Hidden Virtualization

Site administrators use virtualization and cloud technologies but do not make this visible to grid users.

Site administrators benefit from easier system management; users do not directly benefit.

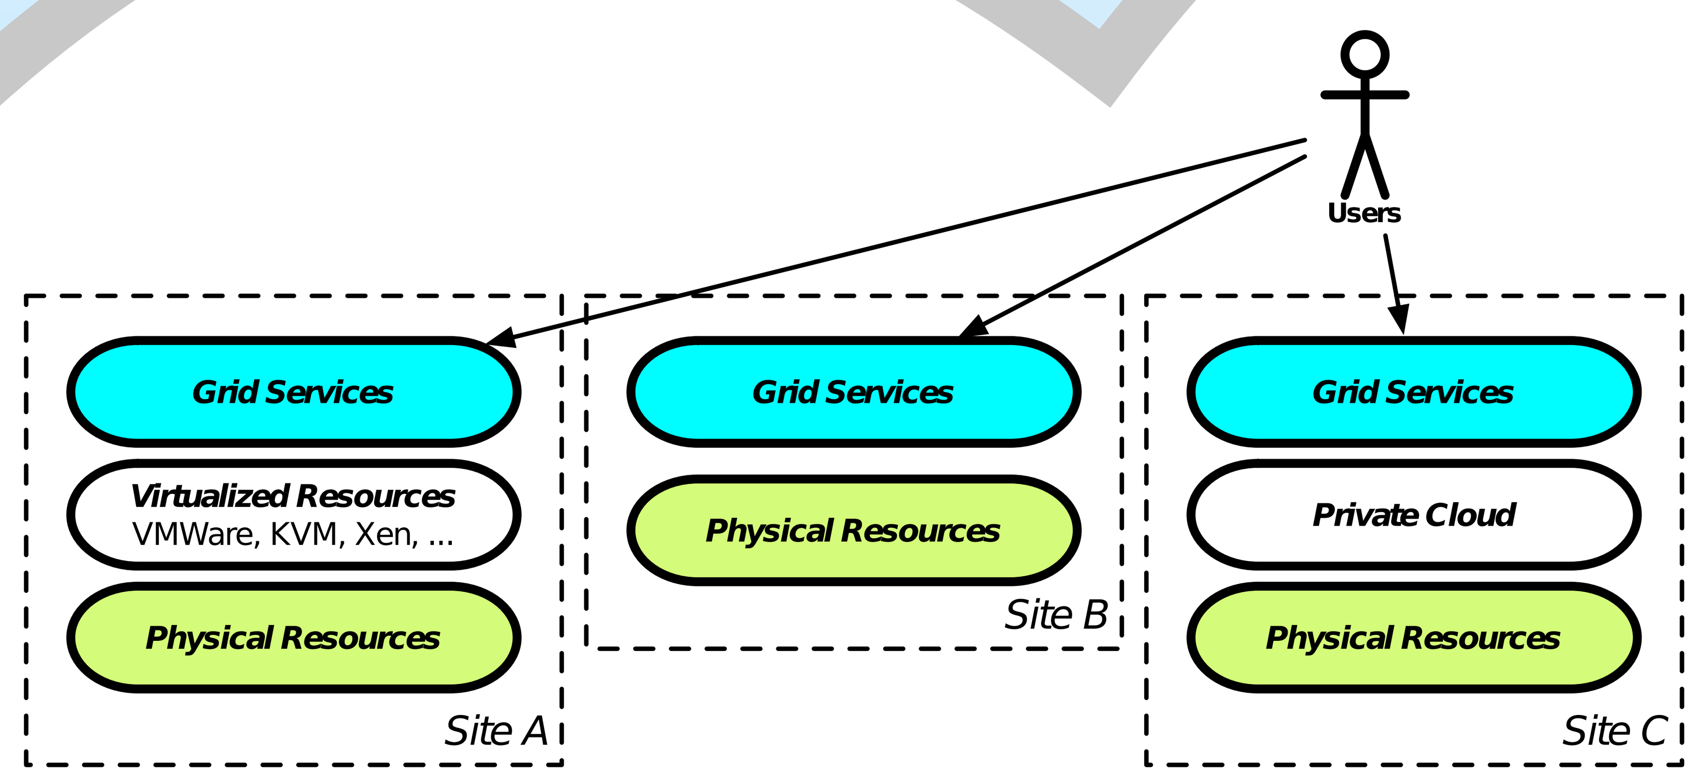

All Supported by StratusLab!

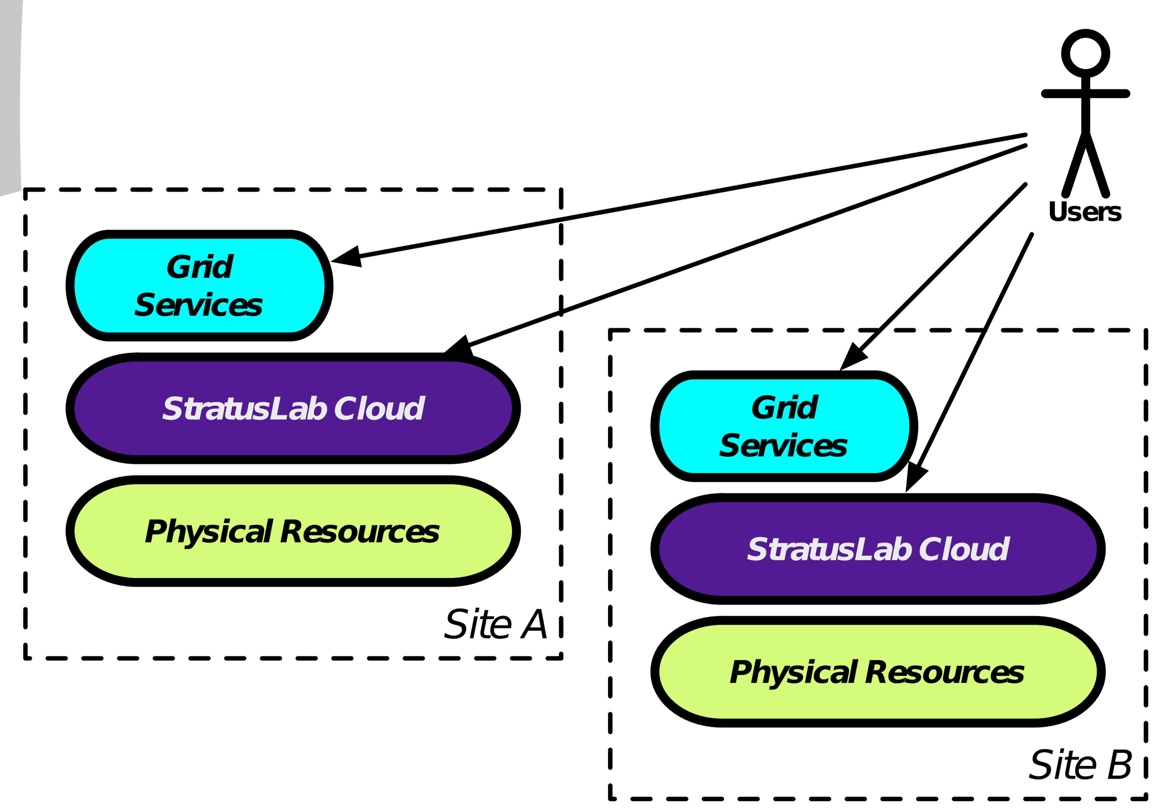

Model 2: Public Cloud and Grid Services

Each site deploys a cloud over its physical resources running the grid services within that cloud with the cloud visible to end-users.

Administrators see the benefits of Model 1 and end-users can take advantage of the customization and dynamic provisioning of clouds.

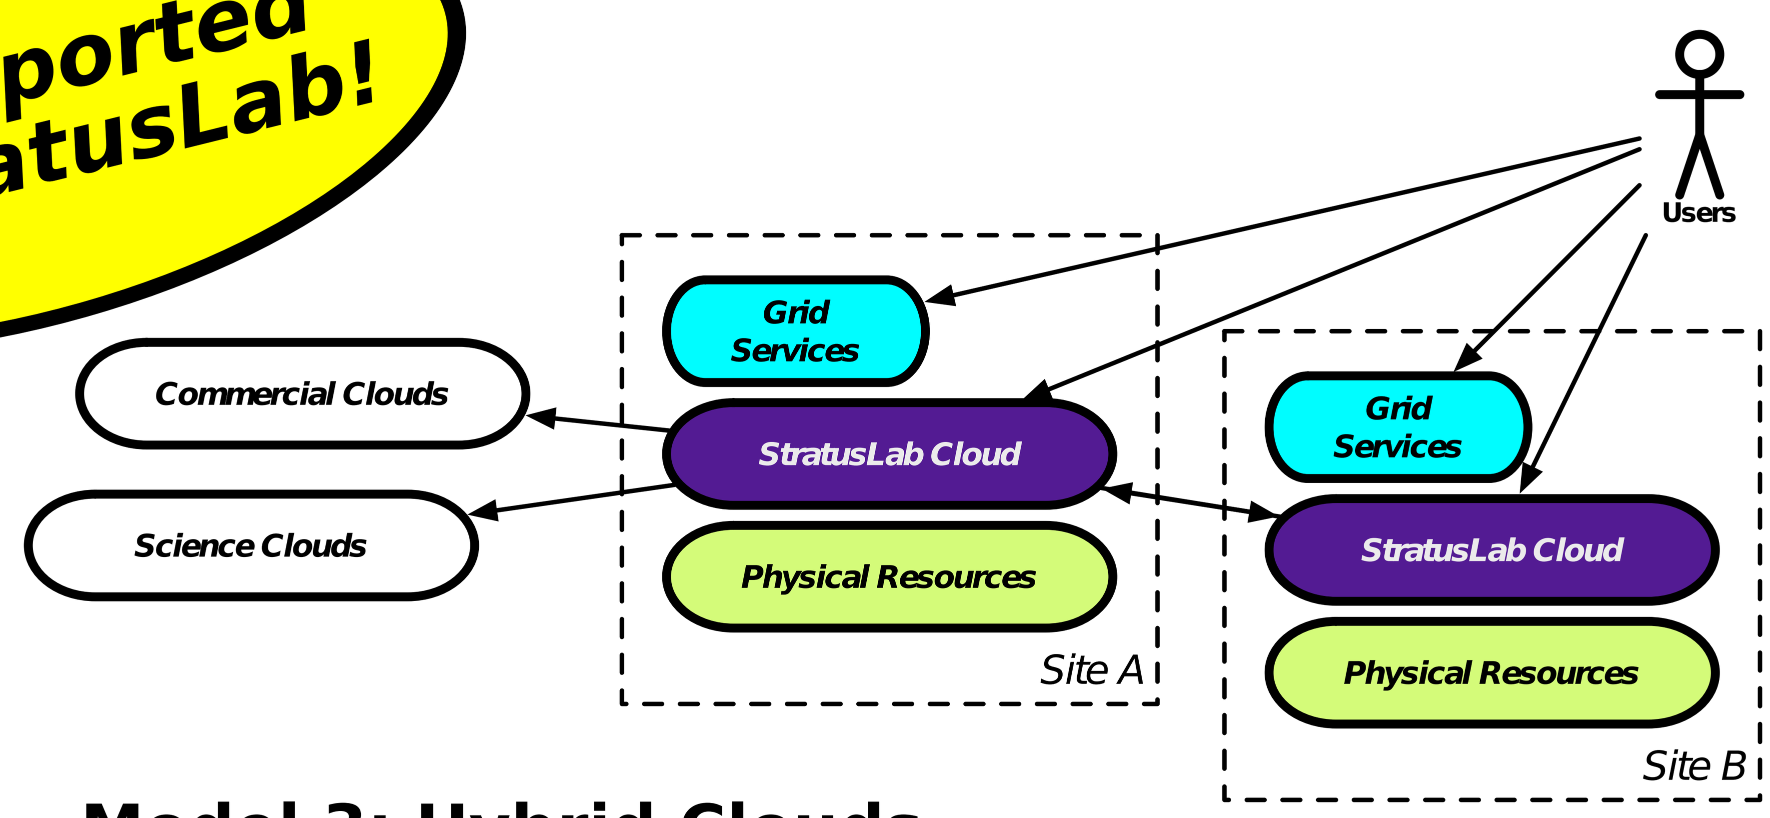

Model 3: Hybrid Clouds

Like Model 2, but clouds can now transparently use resources from other StratusLab clouds, commercial clouds or science clouds.

All the benefits of Model 2, but more resources can be made available and sharing of resources is permitted between sites, not just between users and VOs.

Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique *France*
Universidad Complutense de Madrid *Spain*
Greek Research and Technology Network S. A. *Greece*
SixSq Sàrl *Switzerland*
Telefonica Investigacion y Desarrollo S.A. *Spain*
Trinity College Dublin *Ireland*

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

StratusLab Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

Agile Software Development

StratusLab uses the agile methodology Scrum to manage its development. This allows the project to evolve the StratusLab distribution quickly according to user and system administrator requirements while always maintaining a functioning release.

Sprints last 3 weeks on average, with a public beta release produced every other sprint. Four public releases were made in the first year of the project, culminating in the first production release in June 2011. The latest release is v1.1, finished in mid-September.

Advantages

- New features and improvements are released often
- Continuous feedback from users on working software informs design evolution
- Users influence requirements and their priority based on real use of earlier versions
- Users see progress as new iterative releases are made
- No 'big bang' integration required, significantly reducing risks
- End of sprint demo forces developers to integrate and show their work using a functional system
- Project advancement measured based on facts instead of subjective evaluation

Challenges

- Regular release requires higher level of automation: build procedures, test procedures, and production upgrades.
- Functional increments must be broken down into tasks that can be implemented in a single sprint
- Developer's mindset has to adapt to incremental development, rather than relying on heavyweight up-front design

Product Backlog

Scrum requires that the functional requirements be expressed in the form of user stories, with each user story implementable in a single sprint. All of the unimplemented stories form the '**Product Backlog**', the most important

Releases

Kickoff: 1 June 2010
v0.1: 9 Nov. 2010
v0.2: 17 Dec. 2010
v0.3: 15 March 2011
v0.4: 3 May 2011
v1.0: 12 June 2011
v1.1: 16 Sept. 2011

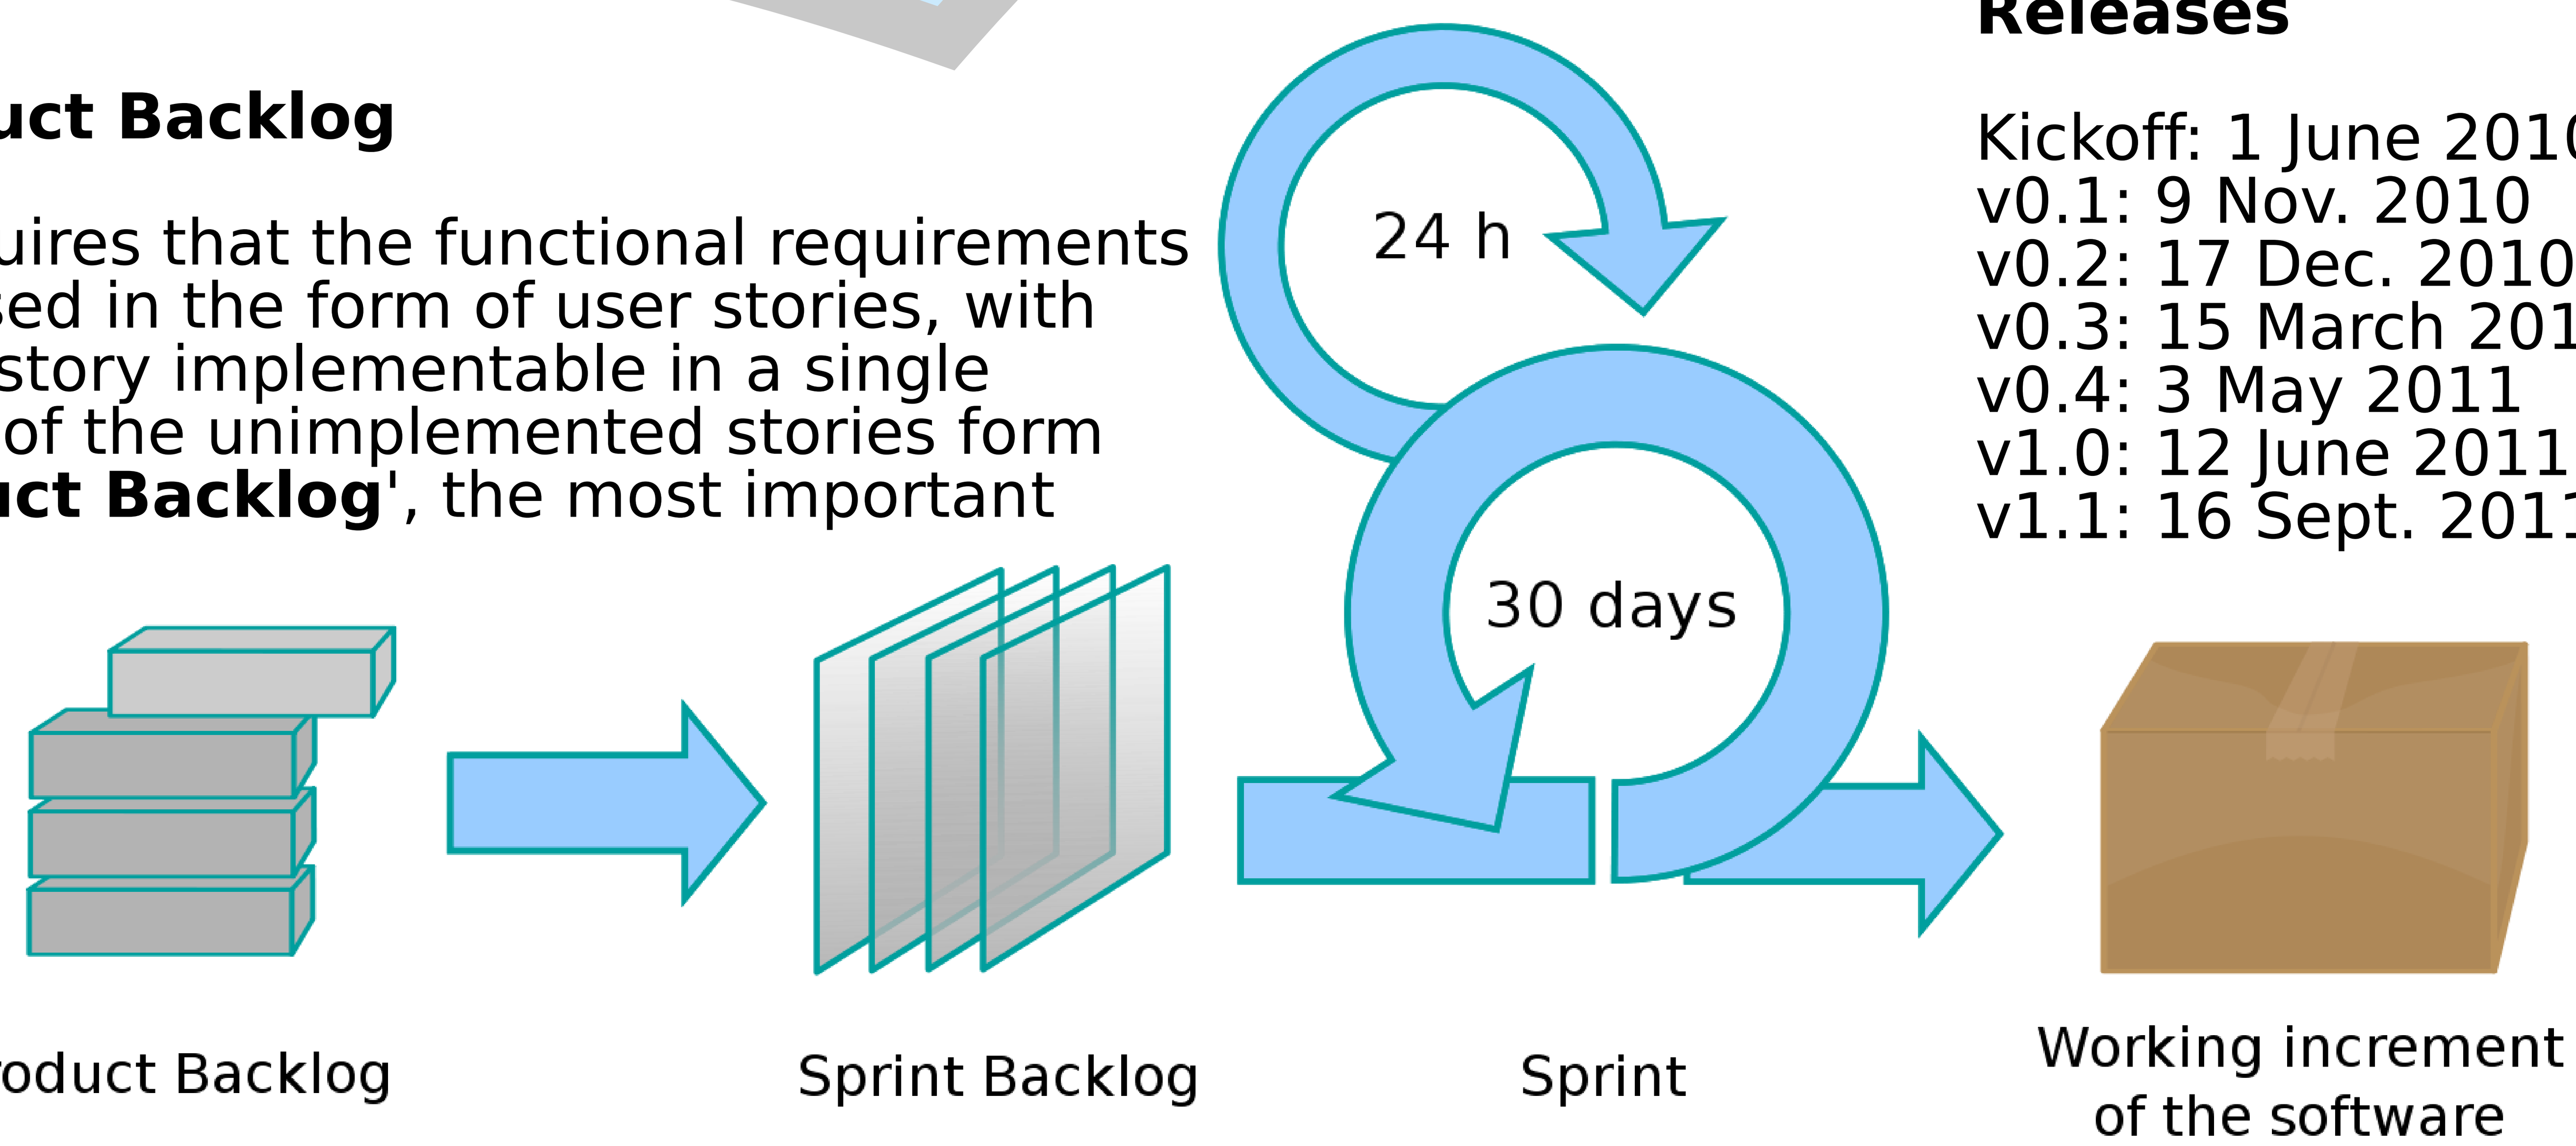

Planning Meetings

Each sprint starts with a planning meeting. During this meeting, user stories from the product backlog are reviewed and selected for the sprint.

The items selected for the sprint become the '**Sprint Backlog**'.

Daily Meetings

In order to ensure a fluid communication during the sprint, we use a 'daily stand-up' meeting. This meeting, which never exceeds 15 minutes, is the place where each team reports on work completed the previous day, what it plans to work on next, and any impediments it has or foresees.

Longer topics that require further discussion are scheduled offline or just after the stand-up.

Demo Meetings

Each sprint concludes with a live demo of each user story implemented during the sprint. This is a powerful way of measuring real progress and ensuring that developments can be released in production.

Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique *France*
Universidad Complutense de Madrid *Spain*
Greek Research and Technology Network S. A. *Greece*
SixSq Sàrl *Switzerland*
Telefonica Investigacion y Desarrollo S.A. *Spain*
Trinity College Dublin *Ireland*

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

StratusLab

Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

StratusLab Production Cloud Service

The Reference Cloud Service, built with the latest version of StratusLab distribution and deployed in GRNET's datacenter, provides public access to a cloud testbed both for evaluation and for production use.

Service Features

- 17 physical nodes (1 Frontend + 16 hosts)
- 256 CPU cores
- 768 GB total memory (48 GB per node)
- 3.6 TB shared storage over NFS
- Support for public IP and local IPs
- Currently running StratusLab v1.1 on Fedora 14.
- Uses OpenNebula 2.2 over kvm for VM management
- Integration with the Claudia Service Manager available

9 months of continuous operations

- More than 3600 VMs instantiated
- Longest running VM is 120 days without interruption
- Hosts the HG-07-StratusLab production grid site, exhibiting high availability and reliability (> 90%)
- 4 major upgrades with minor, limited service interruptions
- VM appliances for gLite grid servicess and various other applications available from the StratusLab Marketplace

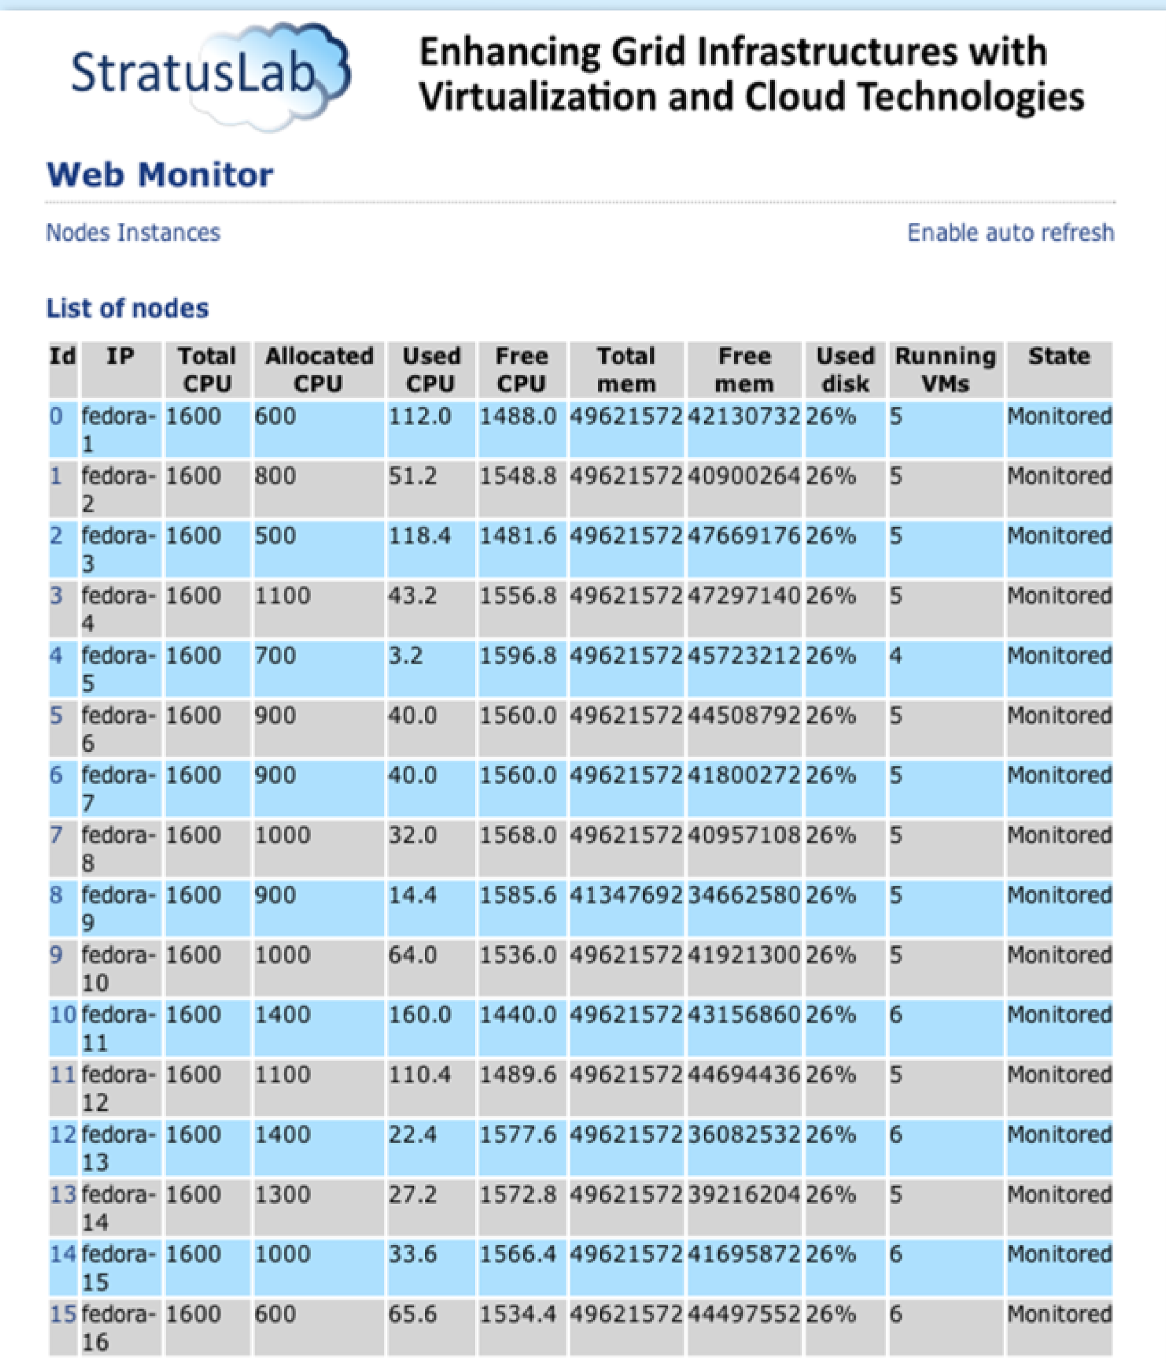

StratusLab — Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

Web Monitor

Nodes Instances — Enable auto refresh

List of nodes

Id	IP	Total CPU	Allocated CPU	Used CPU	Free CPU	Total mem	Free mem	Used disk	Running VMs	State
0	fedora-1	1600	600	112.0	1488.0	49621572	42130732	26%	5	Monitored
1	fedora-2	1600	800	51.2	1548.8	49621572	40900264	26%	5	Monitored
2	fedora-3	1600	500	118.4	1481.6	49621572	47669176	26%	5	Monitored
3	fedora-4	1600	1100	43.2	1556.8	49621572	47297140	26%	5	Monitored
4	fedora-5	1600	700	3.2	1596.8	49621572	45723212	26%	4	Monitored
5	fedora-6	1600	900	40.0	1560.0	49621572	44508792	26%	5	Monitored
6	fedora-7	1600	900	40.0	1560.0	49621572	41800272	26%	5	Monitored
7	fedora-8	1600	1000	32.0	1568.0	49621572	40957108	26%	5	Monitored
8	fedora-9	1600	900	14.4	1585.6	41347692	34662580	26%	5	Monitored
9	fedora-10	1600	1000	64.0	1536.0	49621572	41921300	26%	5	Monitored
10	fedora-11	1600	1400	160.0	1440.0	49621572	43156860	26%	6	Monitored
11	fedora-12	1600	1100	110.4	1489.6	49621572	44694436	26%	5	Monitored
12	fedora-13	1600	1400	22.4	1577.6	49621572	36082532	26%	6	Monitored
13	fedora-14	1600	1300	27.2	1572.8	49621572	39216204	26%	5	Monitored
14	fedora-15	1600	1000	33.6	1566.4	49621572	41695872	26%	6	Monitored
15	fedora-16	1600	600	65.6	1534.4	49621572	44497552	26%	6	Monitored

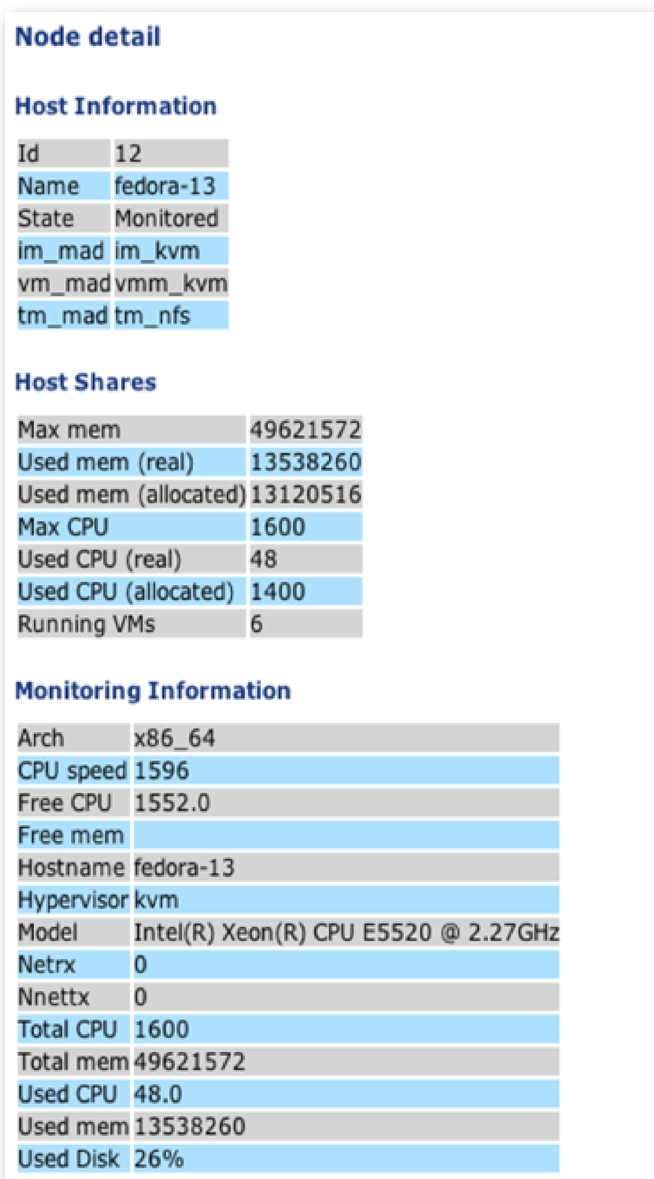

Node detail

Host Information

Id	12
Name	fedora-13
State	Monitored
im_mad	im_kvm
vm_mad	vmm_kvm
tm_mad	tm_nfs

Host Shares

Max mem	49621572
Used mem (real)	13538260
Used mem (allocated)	13120516
Max CPU	1600
Used CPU (real)	48
Used CPU (allocated)	1400
Running VMs	6

Monitoring Information

Arch	x86_64
CPU speed	1596
Free CPU	1552.0
Free mem	
Hostname	fedora-13
Hypervisor	kvm
Model	Intel(R) Xeon(R) CPU E5520 @ 2.27GHz
Netrx	0
Nnettx	0
Total CPU	1600
Total mem	49621572
Used CPU	48.0
Used mem	13538260
Used Disk	26%

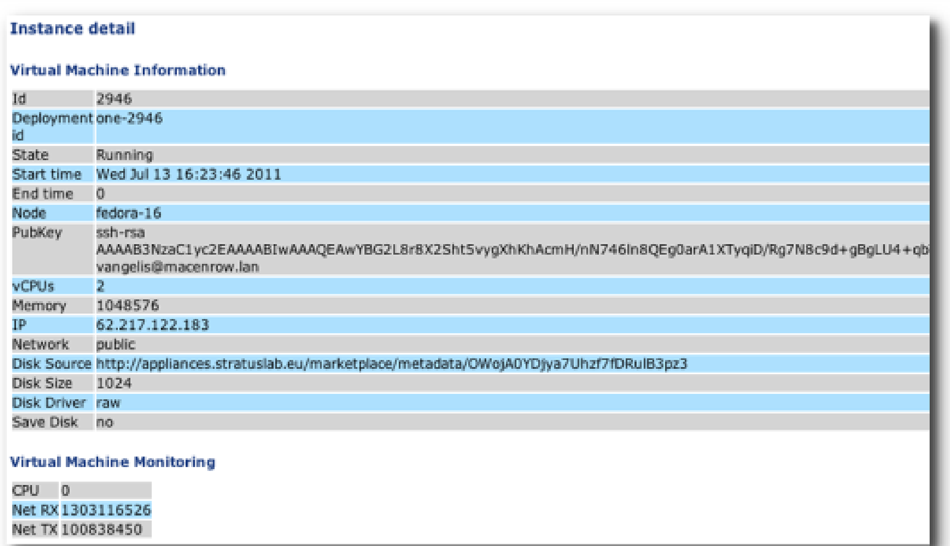

Instance detail

Virtual Machine Information

Id	2946
Deployment id	one-2946
State	Running
Start time	Wed Jul 13 16:23:46 2011
End time	0
Node	fedora-16
PubKey	ssh-rsa AAAAB3NzaC1yc2EAAAABIwAAAQEAwYBG2L8r8X2Sht5vygXhKhAcmH/nN746In8QEg0arA1XTyqiD/Rg7N8c9d+gBgLU4+qb vangelis@macenrow.lan
vCPUs	2
Memory	1048576
IP	62.217.122.183
Network	public
Disk Source	http://appliances.stratuslab.eu/marketplace/metadata/OWojA0YDjya7Uhzf7fDRuIB3pz3
Disk Size	1024
Disk Driver	raw
Save Disk	no

Virtual Machine Monitoring

CPU	0
Net RX	1303116526
Net TX	100838450

Monitoring

Web monitor provides overview of hosting nodes and VMs currently running in the cloud.

Ganglia used for monitoring the physical infrastructure.

How to access and use the service

- Send email to support@stratuslab.eu, providing your details and purpose of usage
- Connect using username/password or your grid certificate (preferable)
- Manage your VMs using the StratusLab command line interface (stratuslab-* tools)
- Prepared appliances and base images (ttylinux, CentOS, Ubuntu, ...) available from public repository operated by TCD
- You can always create your own images and host them in your own repository!)

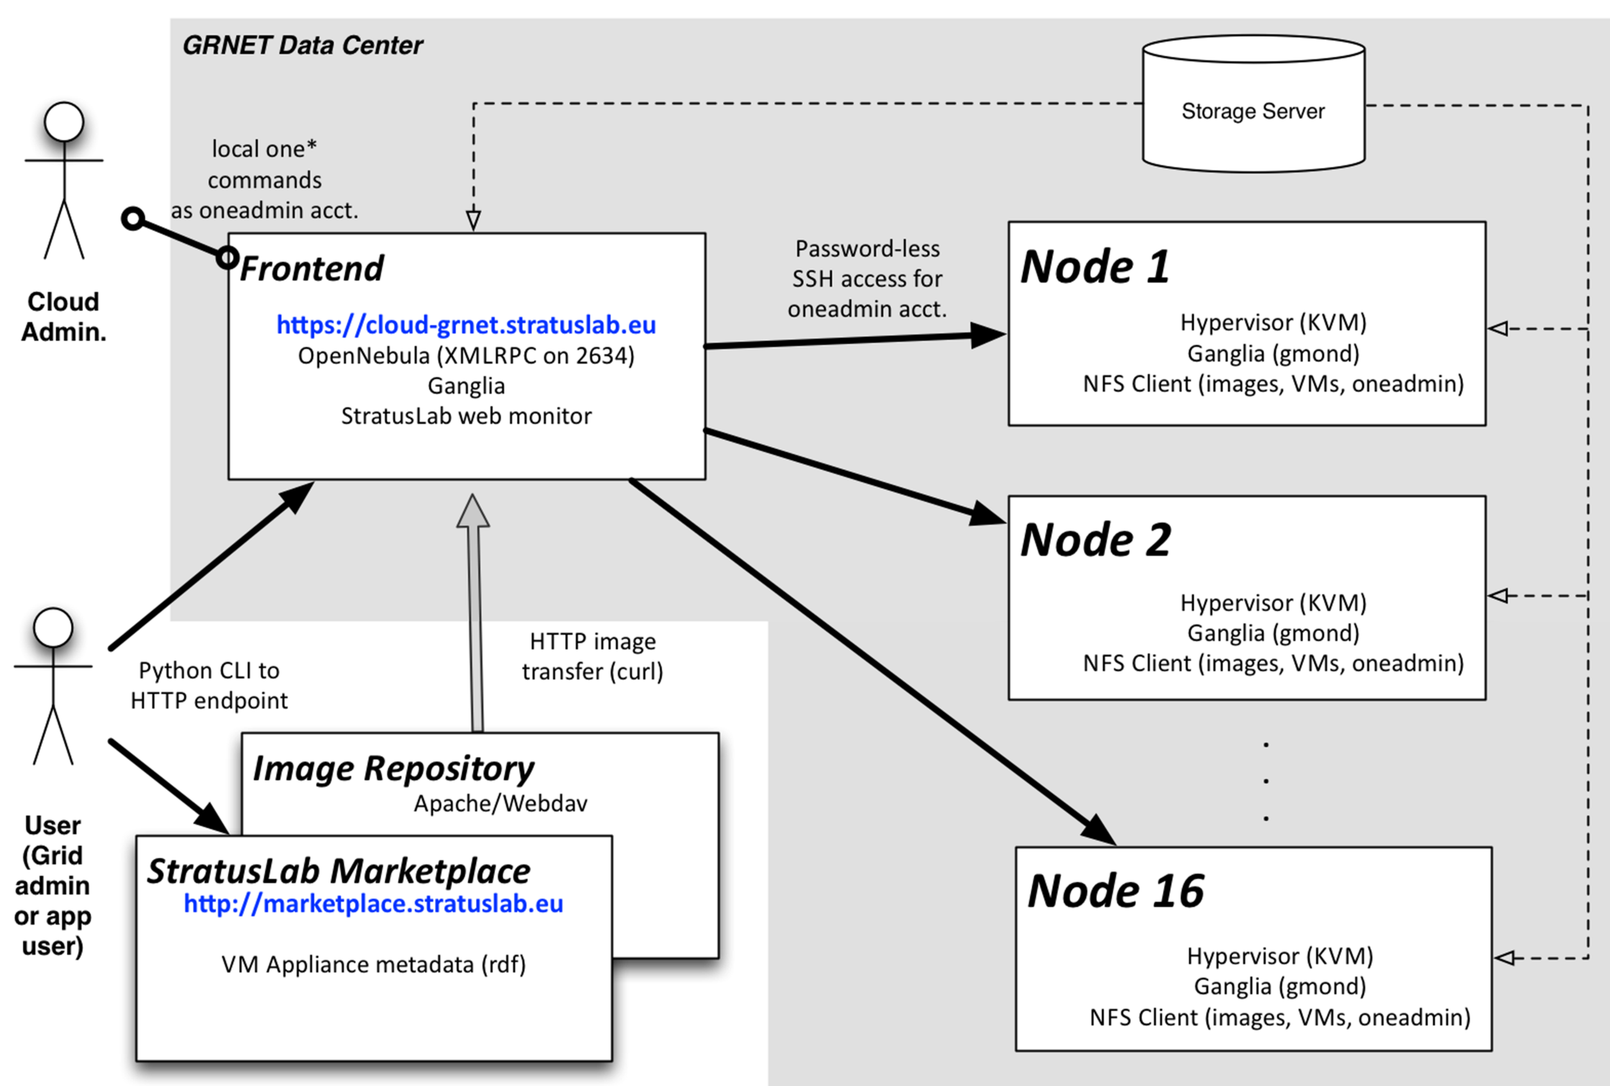

Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique *France*
Universidad Complutense de Madrid *Spain*
Greek Research and Technology Network S. A. *Greece*
SixSq Sàrl *Switzerland*
Telefonica Investigacion y Desarrollo S.A. *Spain*
Trinity College Dublin *Ireland*

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

StratusLab

Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

Marketplace: Registry for Machine & Disk Images

Making secure machine images for virtualized and cloud infrastructures remains a significant hurdle for wider adoption and use of these technologies. The StratusLab Marketplace lowers this barrier by facilitating the sharing and reuse of existing images.

Benefits

End-Users: Scientists, engineers, and other users of StratusLab cloud infrastructures can browse the Marketplace to find existing images that have the operating systems and applications to run their analyses, saving them the time needed to create their own images.

Image Creators: People creating machine and disk images can publish metadata about them in the Marketplace, making them accessible to a wider audience and soliciting feedback to further improve those images.

Cloud Administrators: Those running cloud infrastructures can use image metadata from the Marketplace to evaluate whether they can trust a particular machine or disk image, banning if necessary, images which do not meet the administrator's requirements.

Image Contents
Contents are uploaded into cloud storage or any other web accessible location; the Marketplace does not store the image contents. The link between the metadata and the image is based on the SHA-1 checksum.

Image Metadata
Existing standards have been used to provide a lightweight, extensible metadata format that allows searching and selection based on a variety of attributes. (See example below.)

Standards
XML (W3C)
XML Signature (W3C)
RDF-XML (W3C)
Dublin Core Metadata (DCMI)

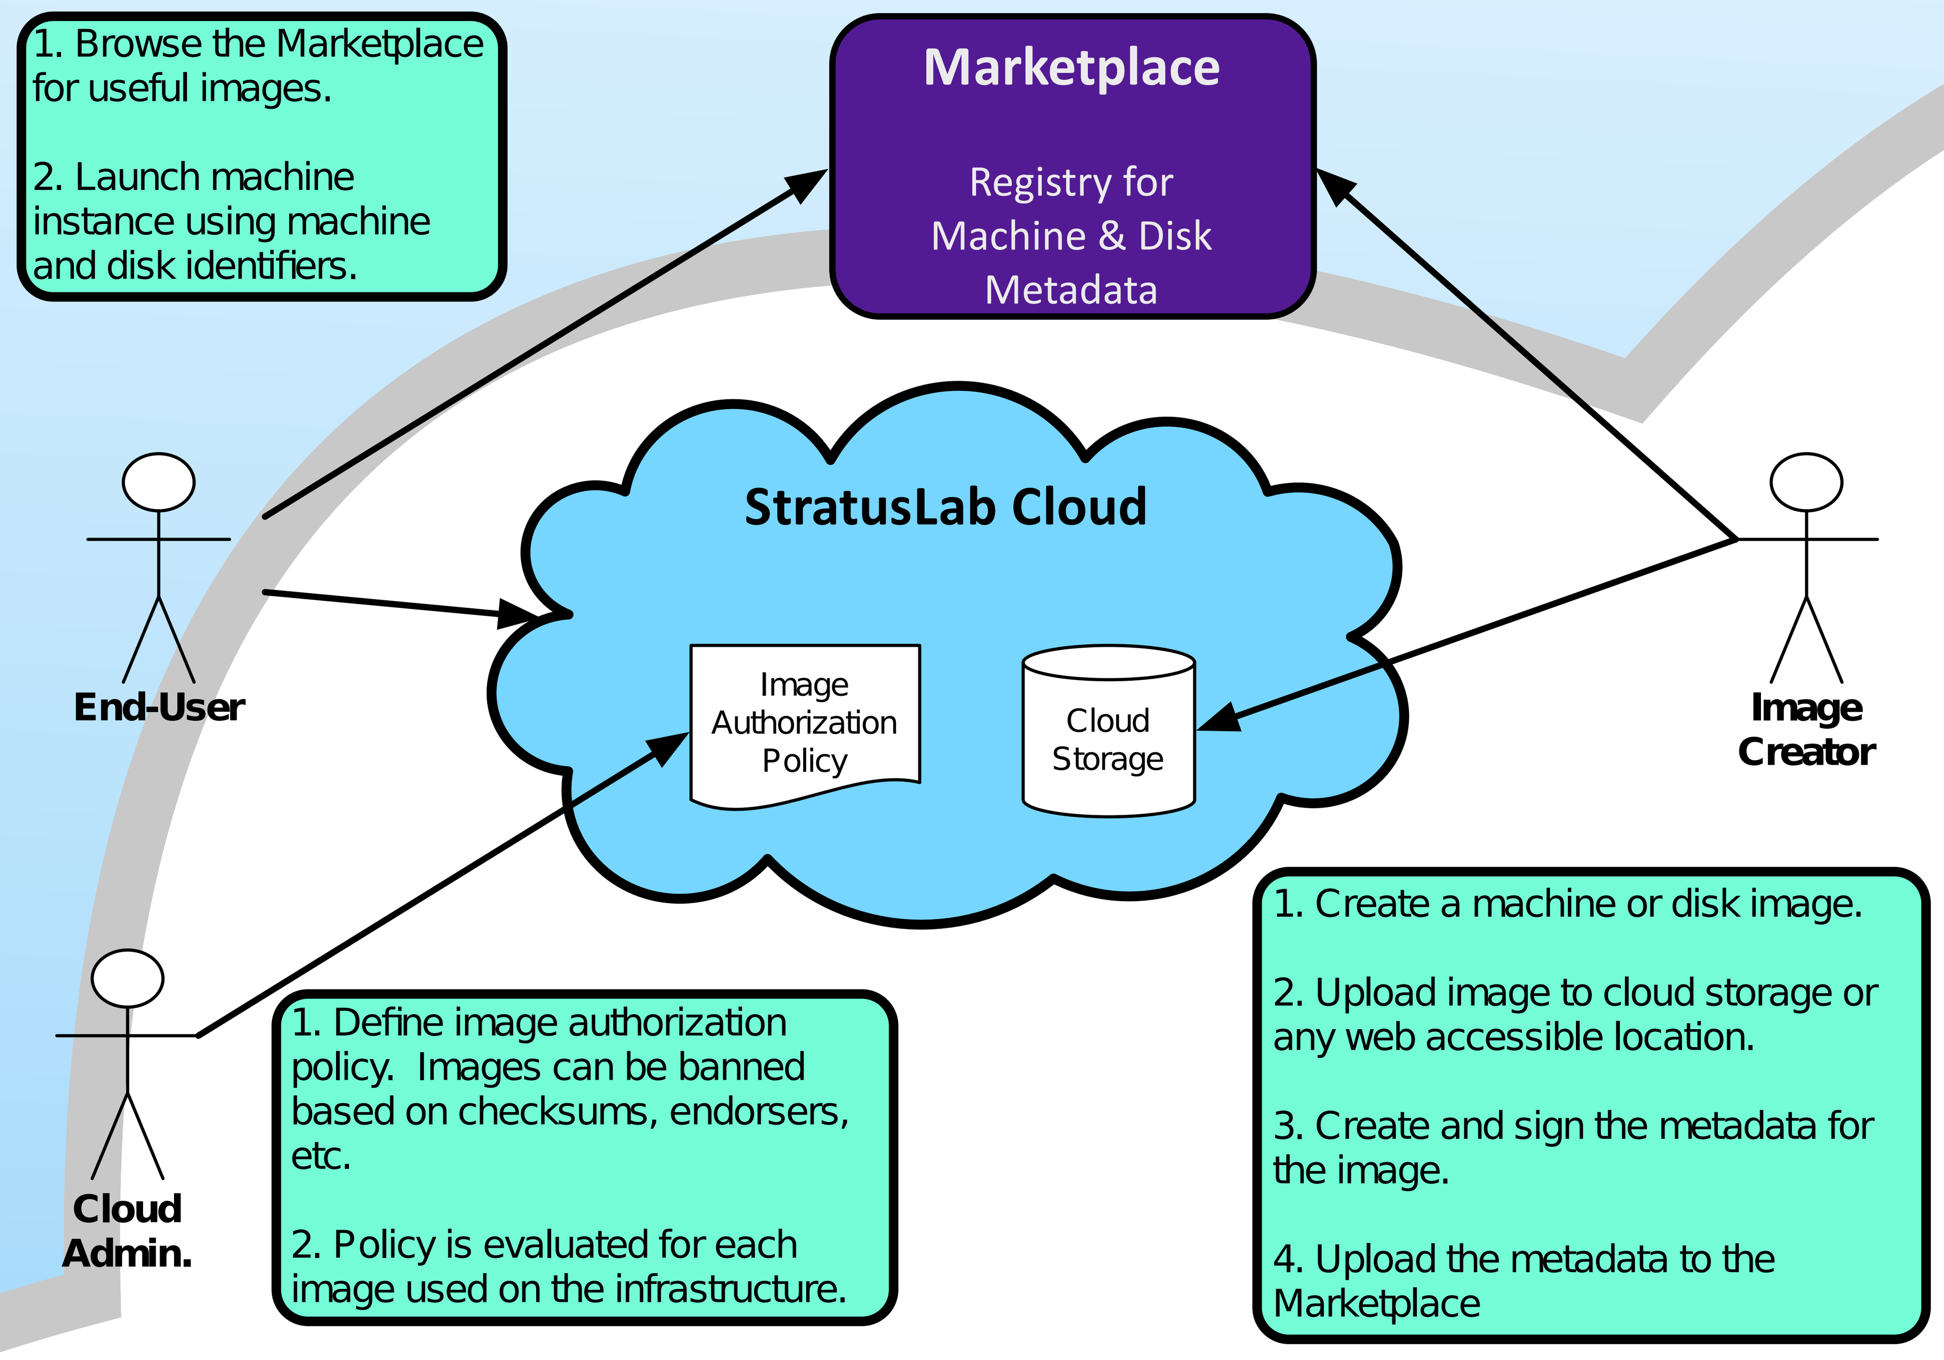

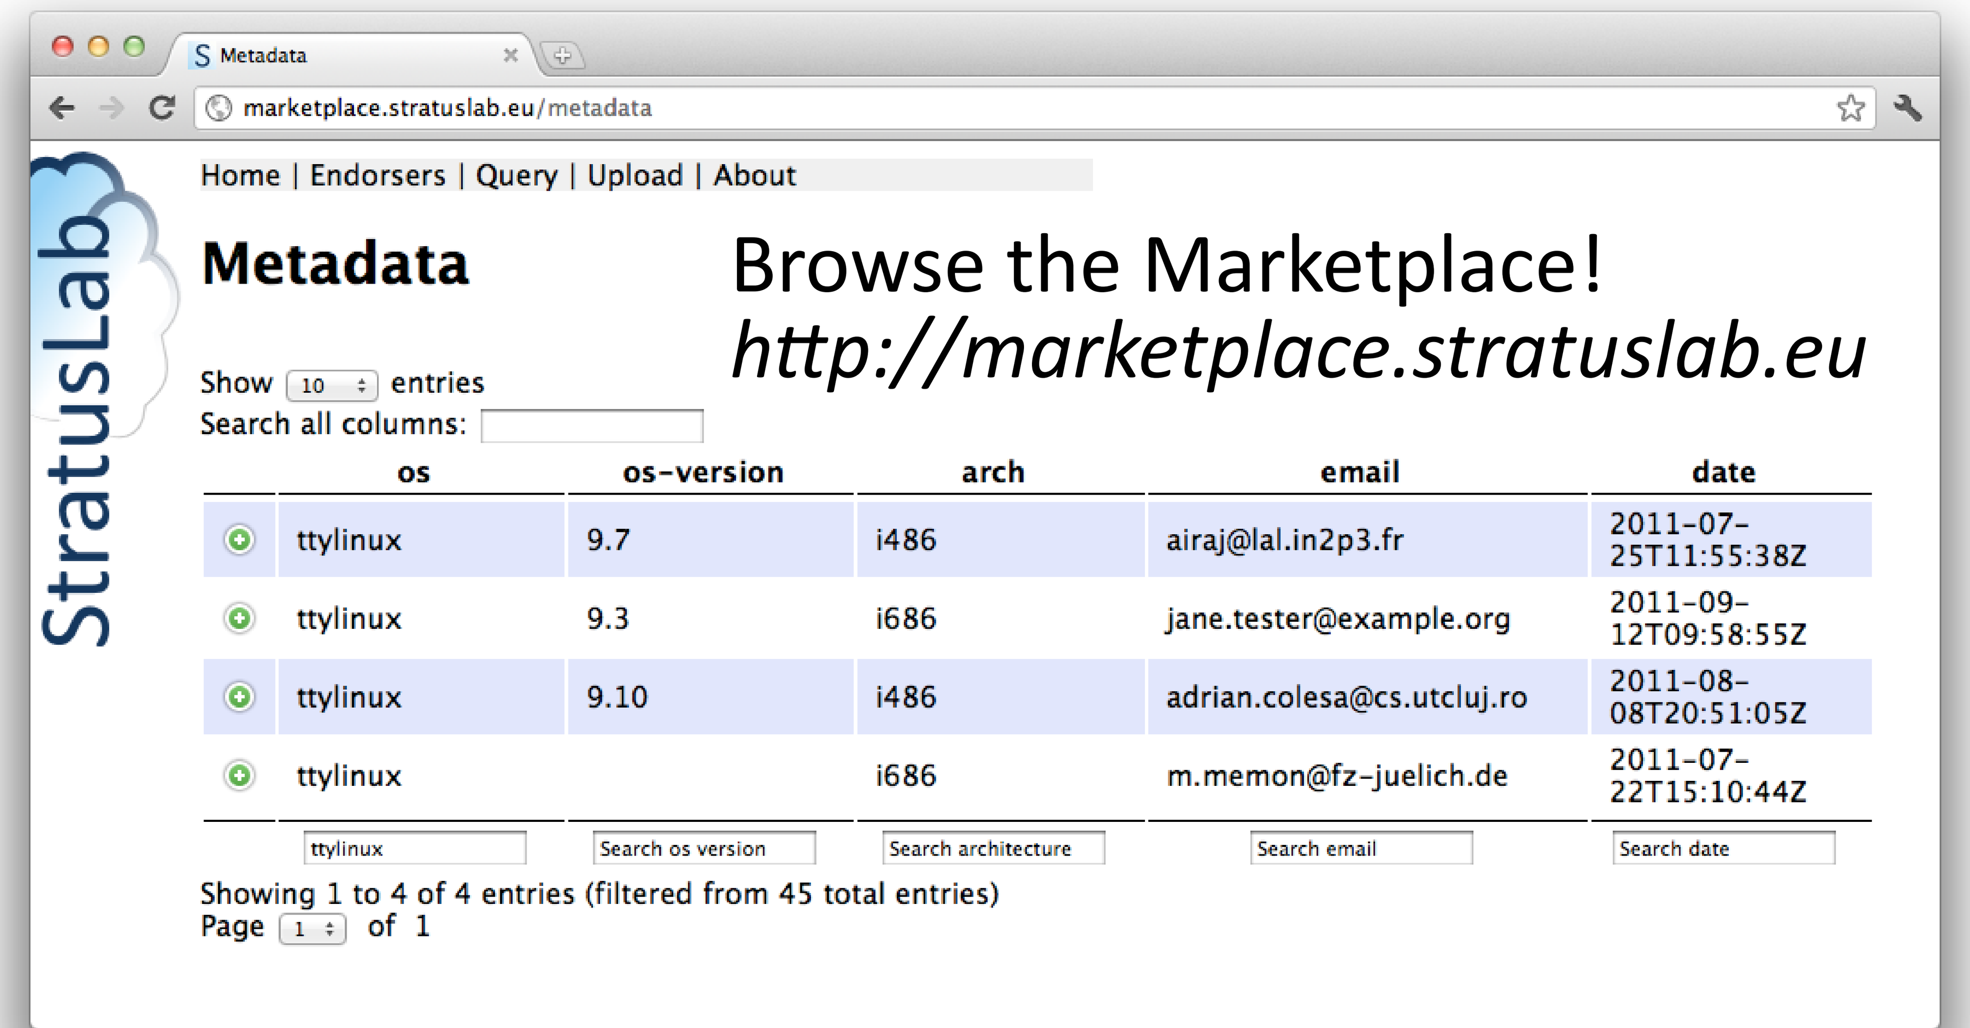

XML Namespaces · Image Identifiers · Image Endorser · General Image Information · Additional Information · Cryptographic Signature


```
<rdf:RDF
  xmlns:dcterms="http://purl.org/dc/terms/"
  xmlns:ex="http://example.org/"
  xmlns:rdf="http://www.w3.org/1999/02/22-rdf-syntax-ns#"
  xmlns:slreq="http://mp.stratuslab.eu/slreq#"
  xmlns:slterms="http://mp.stratuslab.eu/slterms#"
  xml:base="http://mp.stratuslab.eu/">

  <rdf:Description rdf:about="#MMZu9WvwKIro-rtBQfDk4PsKO7_">

    <dcterms:identifier>MMZu9WvwKIro-rtBQfDk4PsKO7_</dcterms:identifier>
    <slreq:bytes>100</slreq:bytes>
    <slreq:checksum rdf:parseType="Resource">
      <slreq:algorithm>SHA-1</slreq:algorithm>
      <slreq:value>c319bbd5afc0a22ba3eaed0507c39383ec28eeff</slreq:value>
    </slreq:checksum>

    <slreq:endorsement rdf:parseType="Resource">
      <created xmlns="http://purl.org/dc/terms/">2011-03-25T21:25:07Z</create
      <endorser xmlns="http://mp.stratuslab.eu/slreq#" rdf:parseType="Resourc
        <email>loomis@lal.in2p3.fr</email>
        <subject>CN=Charles Loomis,OU=LAL,O=CNRS,C=FR,O=GRID-FR</subject>
        <issuer>CN=GRID2-FR,O=CNRS,C=FR</issuer>
      </endorser>
    </slreq:endorsement>

    <dcterms:type>machine</dcterms:type>
    <dcterms:valid>2011-07-23T10:59:42Z</dcterms:valid>
    <dcterms:publisher>StratusLab</dcterms:publisher>

    <dcterms:title>linux-with-my-apps</dcterms:title>
    <dcterms:description>A 32-bit ttylinux...</dcterms:description>
    <slterms:serial-number>0</slterms:serial-number>
    <slterms:version>1.0</slterms:version>

    <slterms:hypervisor>kvm</slterms:hypervisor>
    <slterms:inbound-port>443</slterms:inbound-port>
    <slterms:outbound-port>25</slterms:outbound-port>
    <slterms:icmp>8</slterms:icmp>

    <slterms:os>ttylinux</slterms:os>
    <slterms:os-version>9.7</slterms:os-version>
    <slterms:os-arch>i486</slterms:os-arch>

    <ex:other-info>additional metadata</ex:other-info>
    <ex:yet-more>still more info</ex:yet-more>
    <ex:relatedImages rdf:parseType="Resource">
      <dcterms:identifier>MMZu9WvwKIro-rtBQfDk4PsKO7_</dcterms:identifier>
      <dcterms:identifier>NMZu9WvwKIro-rtBQfDk4PsKO7_</dcterms:identifier>
    </ex:relatedImages>
  </rdf:Description>

  <Signature xmlns="http://www.w3.org/2000/09/xmldsig#">
    <SignedInfo>
      <CanonicalizationMethod Algorithm="http://www.w3.org/..."/>
      <SignatureMethod Algorithm="http://www.w3.org/2000/09/xmldsig#rsa-sha1"
      <Reference URI="">
        <Transforms>
          <Transform Algorithm="http://www.w3.org/2000/09/xmldsig#..."/>
        </Transforms>
        <DigestMethod Algorithm="http://www.w3.org/2000/09/xmldsig#sha1"/>
        <DigestValue>go6VR...</DigestValue>
      </Reference>
    </SignedInfo>
    <SignatureValue>NrTJFF...</SignatureValue>
    <KeyInfo>
      <X509Data>...
        <X509Certificate>...</X509Certificate>
      </X509Data>
    </KeyInfo>
  </Signature>

</rdf:RDF>
```


Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique *France* · Universidad Complutense de Madrid *Spain* · Greek Research and Technology Network S. A. *Greece* · SixSq Sàrl *Switzerland* · Telefonica Investigacion y Desarrollo S.A. *Spain* · Trinity College Dublin *Ireland*

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

StratusLab Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

Persistent Disk Service

The StratusLab cloud distribution contains a service for managing "disk-like" storage areas. Disk-based abstractions allow both read-only distribution of disk images through the StratusLab Marketplace and persistent read-write disks within a particular cloud instance.

Users can continue to use their standard file-based storage services, provided that the clients for these services are available within the running virtual machine. In particular, grid users can continue to use the standard SRM interfaces for files stored on the grid.

Persistent Disks

These may be created and used by a single machine instance within a single cloud. They are useful for saving the persistent state of services.

Disks can be mounted when the machine starts, or added and removed dynamically if the VM supports ACPI.

Using a Persistent Disk

1. Create a disk with a given size via the web or command-line client.

2. Launch a machine instance referencing that disk image.

3. Partition (fdisk) and format (mkfs) the disk via the running virtual machine.

4. Store data to the disk as usual.

5. Dismount the disk or halt the machine instance.

6. Disk with the persistent data is available for use by another machine instance.

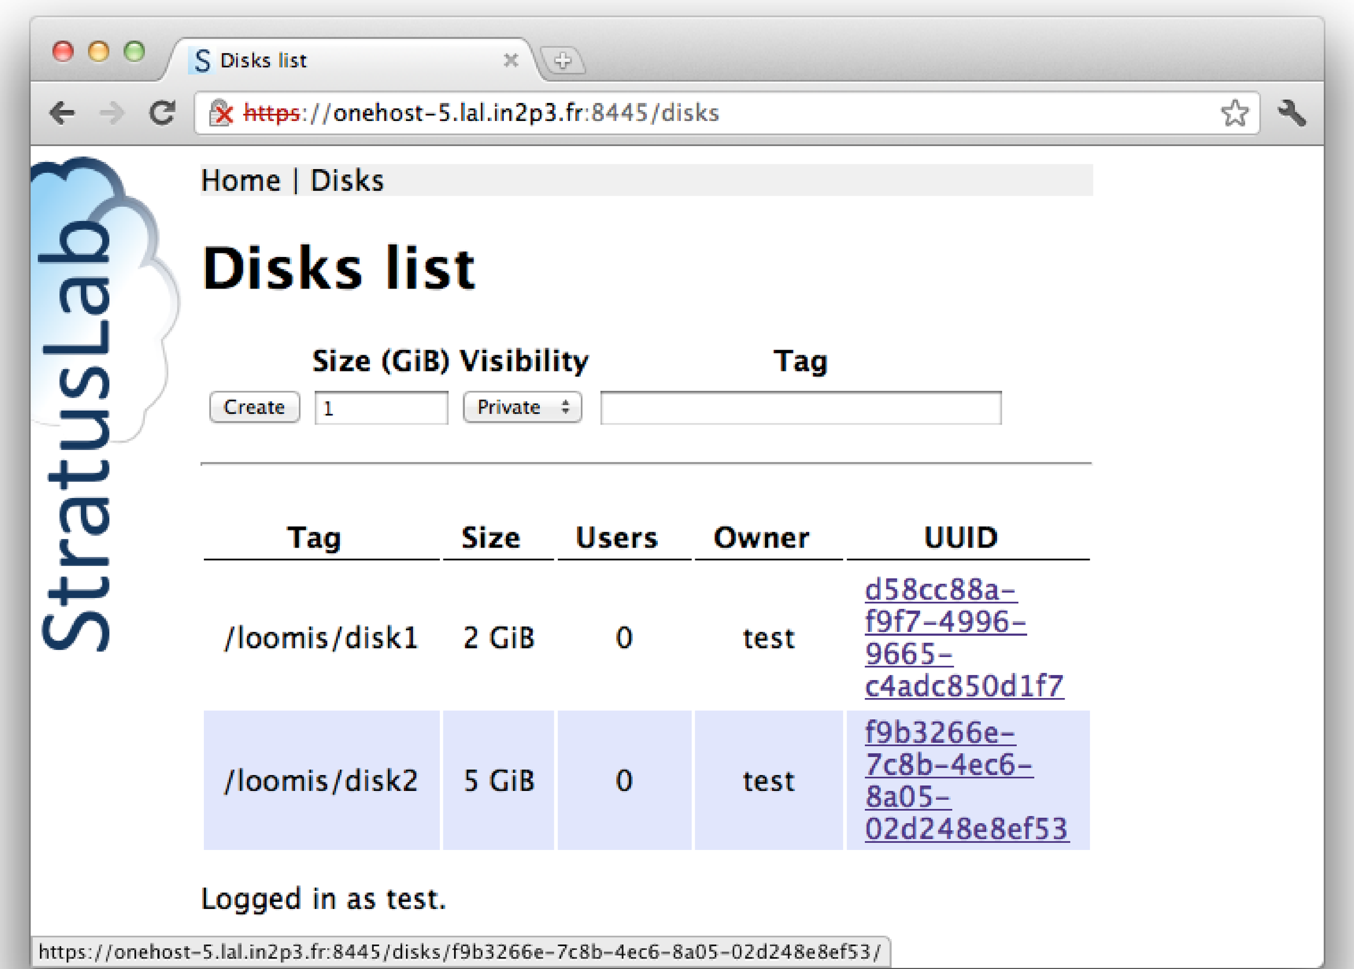

Persistent disks can be created and managed via the web or command line interfaces, including dynamic mounting on running virtual machines.

Shared Read-Only Disks

It is often useful to share datasets via read-only disks. This allows caching of these disks and sharing between different cloud instances. It also permits users to use specific versions of the database for their analyses.

These are supported though the Marketplace and via the StratusLab access and caching mechanisms.

Using a Shared Read-Only Disk

1. Launch a machine giving the Marketplace reference for the image.

2. Machine is launched with the given disk image on the given device (typically /dev/hdd).

3. Mount the disk on the machine and use the data for an analysis.

4. Dismount or halt the machine instance.

The disk is read-only and no changes can be made to the disk image.

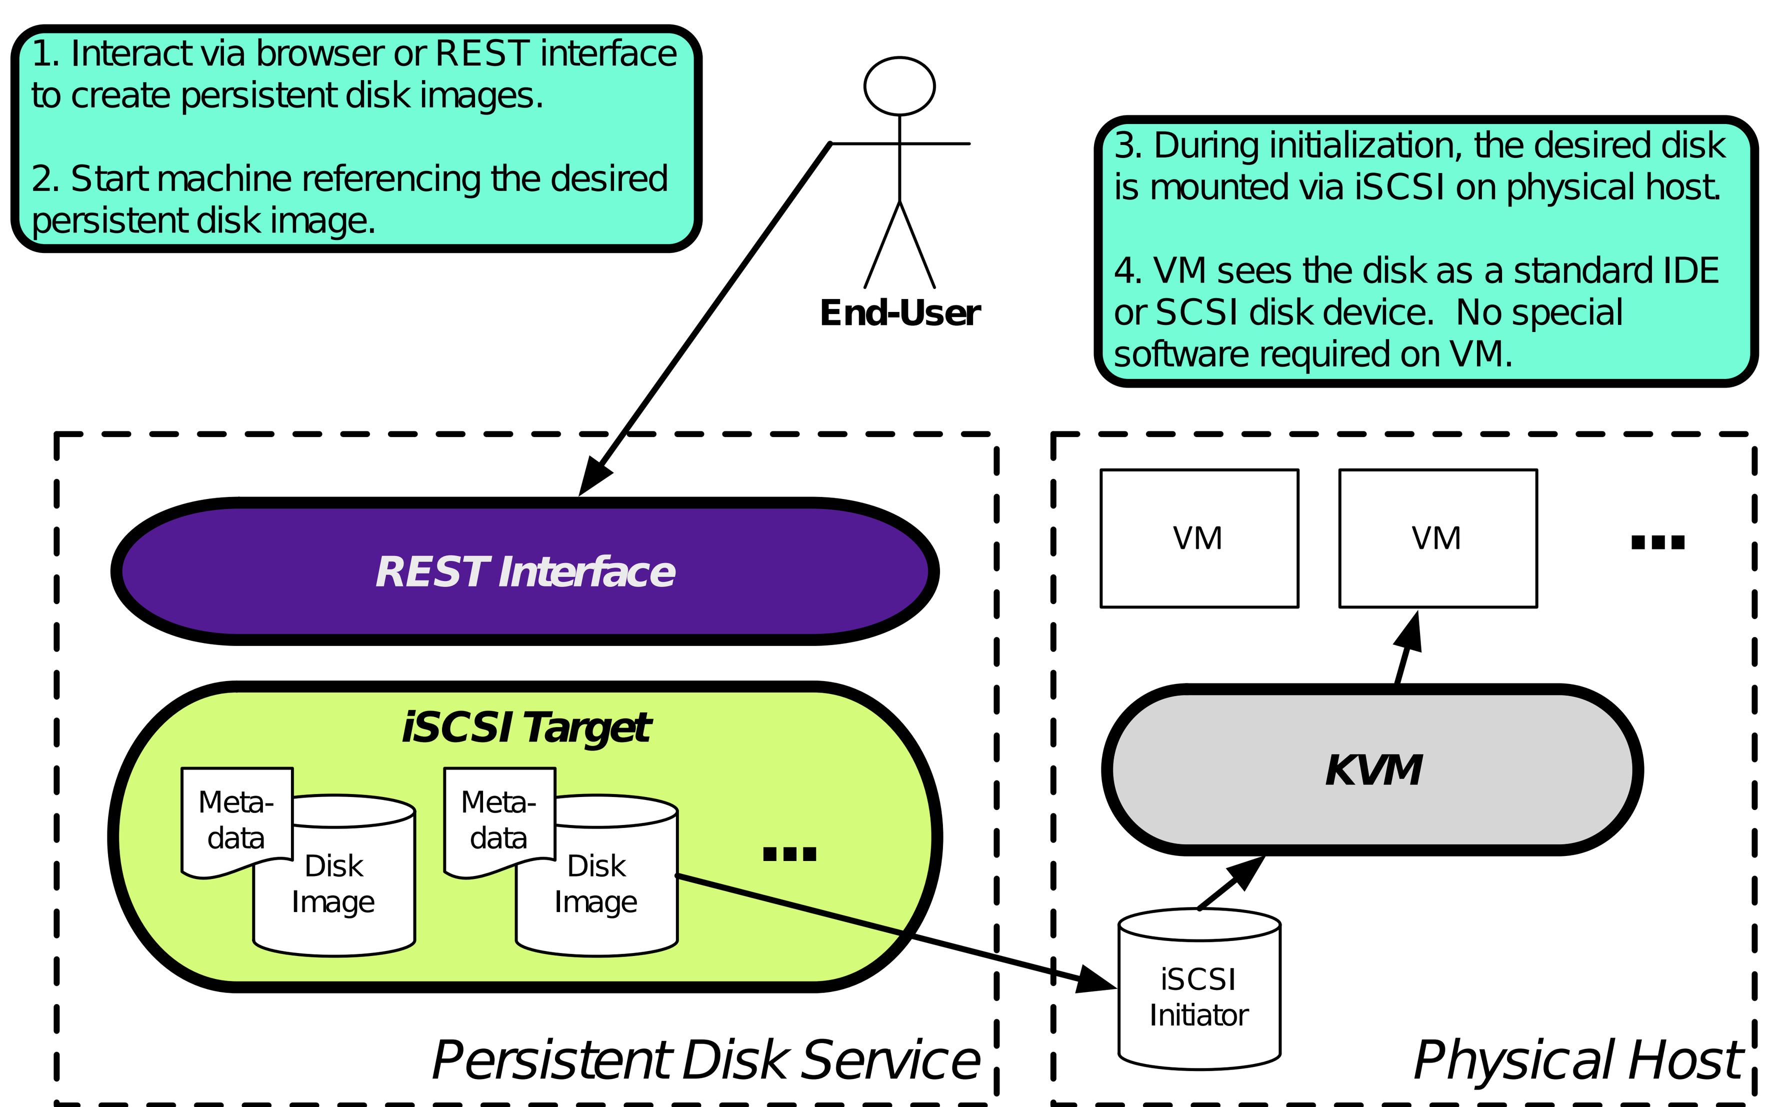

Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique *France*
Universidad Complutense de Madrid *Spain*
Greek Research and Technology Network S. A. *Greece*
SixSq Sàrl *Switzerland*
Telefonica Investigacion y Desarrollo S.A. *Spain*
Trinity College Dublin *Ireland*

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

StratusLab

Bioinformatics Appliances to foster Cloud for Life Science Applications

Several experimental technologies have been improved to such a degree that obtaining data is easy, causing a deluge of data for the Bioinformatics community. The challenge is to be able to analyze these data with the relevant applications, for example, sequencing a whole genome obtained from Next Generation Sequencing (NGS) instrument. Many projects are working on the genome sequence of different organisms, continuously providing new sequences for analysis. Some bioinformatics algorithms like BLAST, FastA or ClustalW are used for that analysis and are usually classified as data-intensive, processing gigabytes of data stored in flat-file databases like UNIPROT, EMBL or PDBseq via a shared filesystem. Others like Abyss, BWA or Ray take the output sequences of sequencing machines and assemble them to get the complete sequence of the studied genome.
In the context of the StratusLab project (EU-FP7, www.stratuslab.org), we have built two bioinformatics virtual appliances: a "Biological databases repository" and a "Bioinformatics compute node".

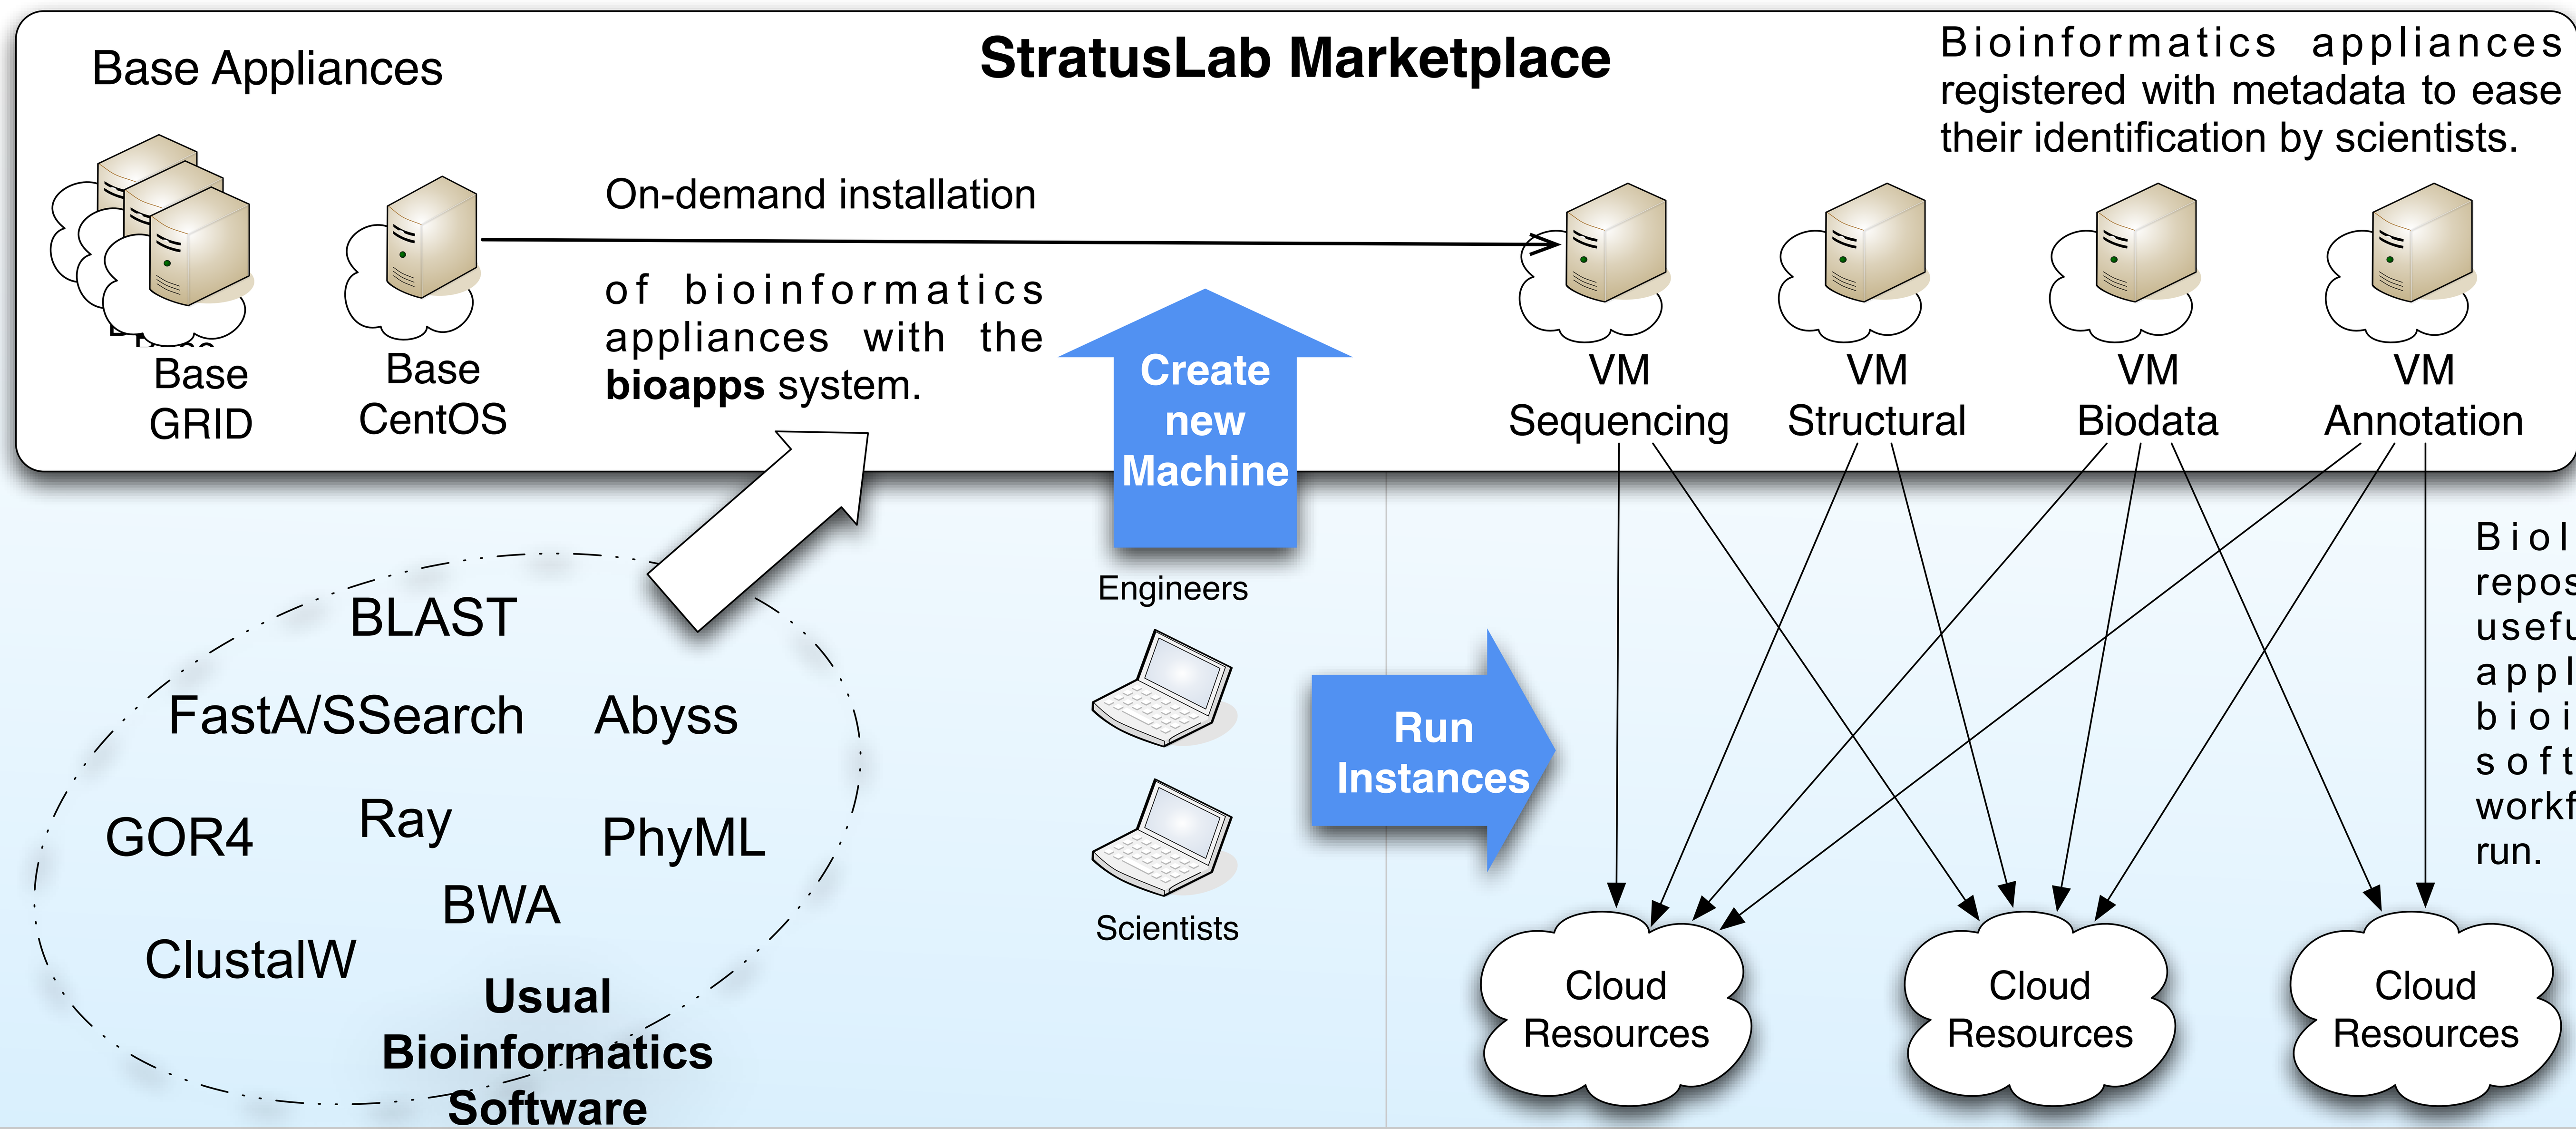

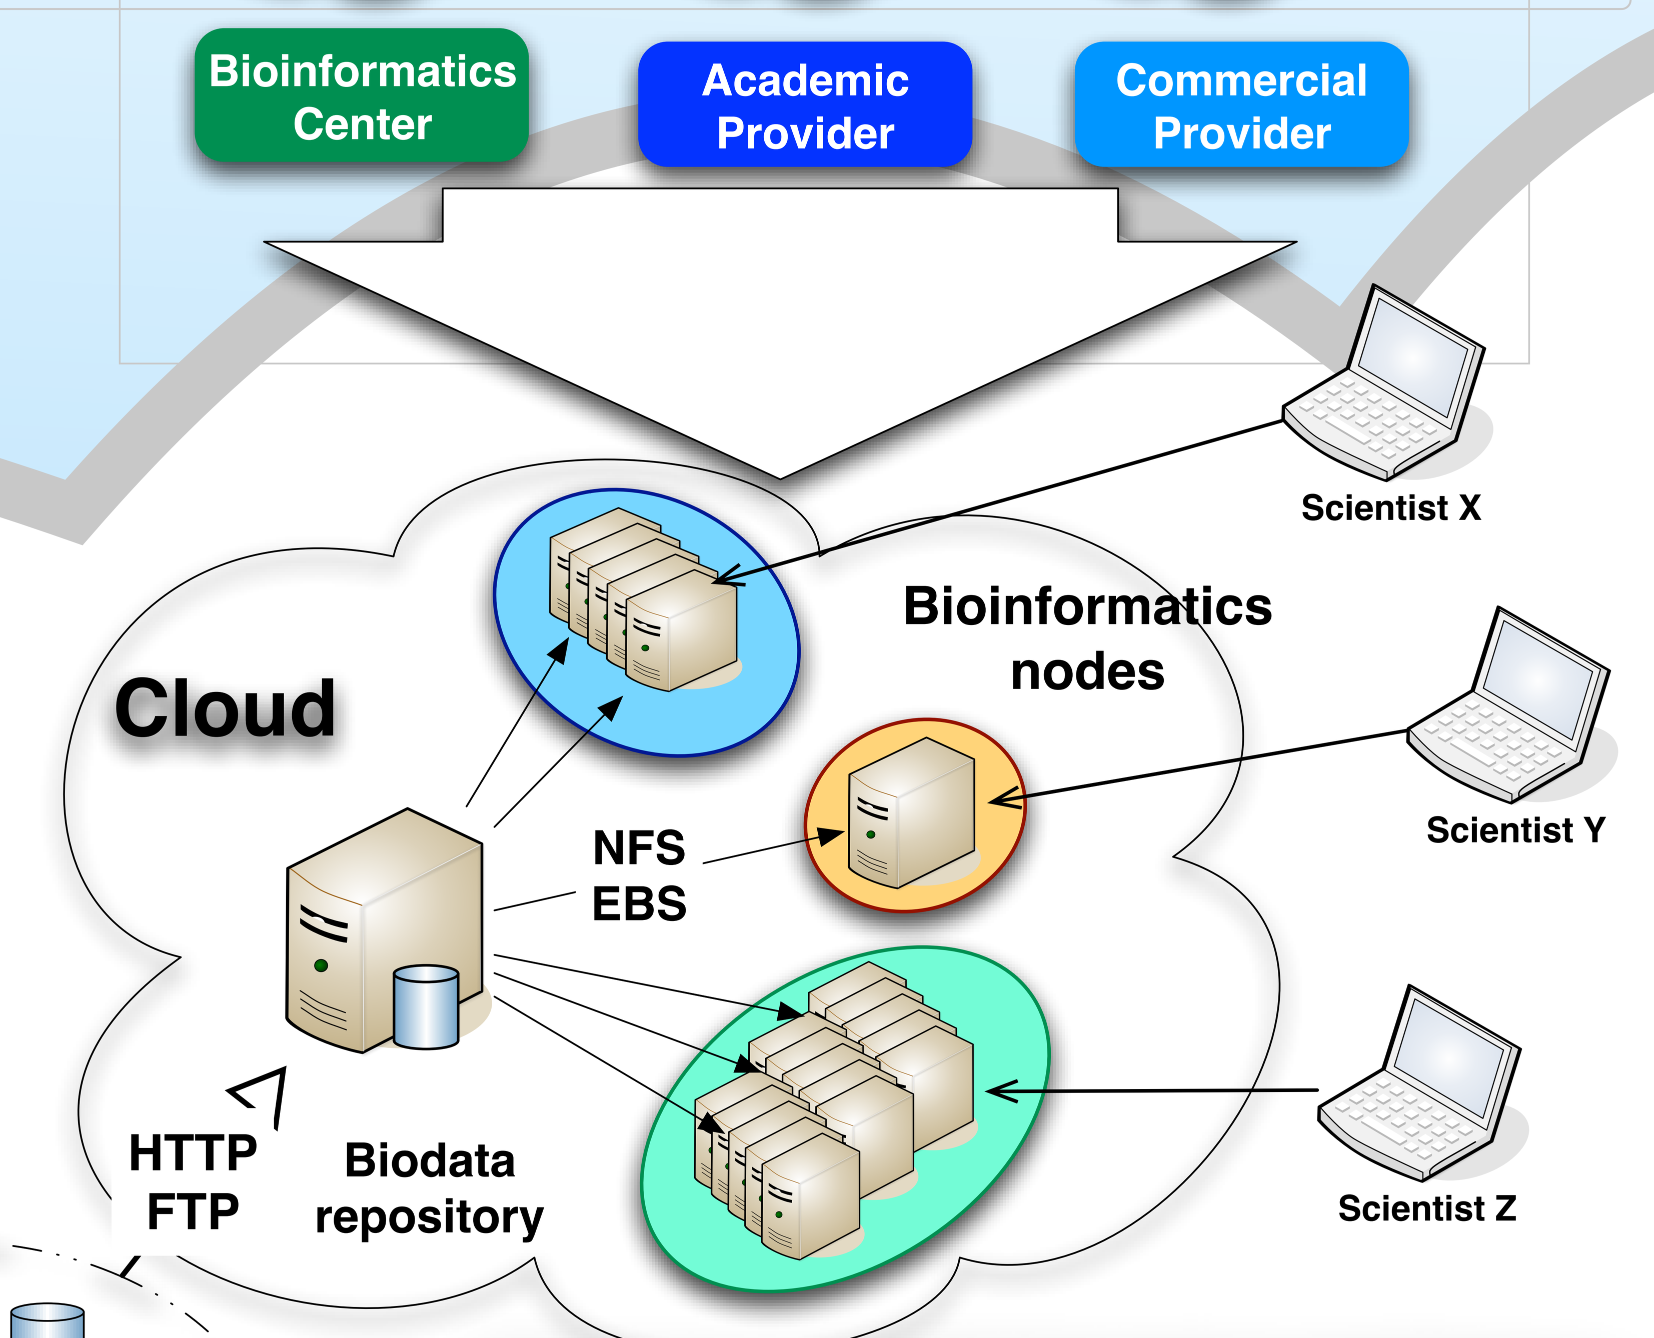

Cloud Advantages for Bioinformatics

Cloud technologies provide scientists with the flexibility to deploy bioinformatics applications on different virtual machines. But clouds have to be connected with public bioinformatics infrastructures, especially the network of international biological databases, to be useful.

Platform as a Service - PaaS

- Provide scientists with Bioinformatics appliances on academic or commercial data centers, e.g. RENABI, or on their own computer or private cloud
- The mission of bioinformatics centers switches from providing services to providing virtual appliances

Infrastructure as a Service - IaaS

- Bioinformaticians can pre-define and deploy clusters or web service infrastructures
- Biologists can deploy different virtual servers to run a complete analysis pipeline

Access to Reference Databases

Access from any compute node is a strong requirement. Such databases contain, for example, protein or gene sequences and associated data, protein structures or complete genomes. We have built a virtual appliance that acts as a proxy between the internet where all the reference databases are published and the cloud where the bioinformatics analyses will be computed. To import and maintain these databases, we use the BioMaJ system (biomaj.genouest.org). This virtual machine stores the data on a local disk (in the future on a cloud-persistent disk), and then exports it with NFS to all the bioinformatics computing machines of the cloud.

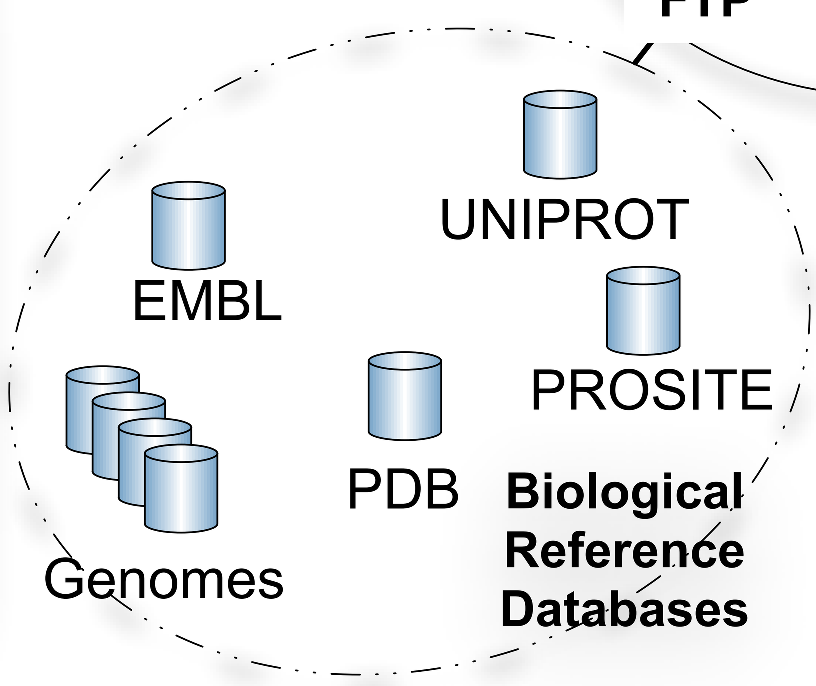

Ready-to-run Bioinformatics Appliances

Appliances containing common bioinformatics software is also an important requirement. We are building such appliances, keeping them up-to-date, and making them available from the StratusLab Marketplace. Scientists with a cloud login can then launch as many instances of these appliances as required by their analysis pipelines. Because bioinformatics applications require access to reference data to process their analyses, the compute node mounts the exported volumes containing the biological databases.

Bioinformatics Cloud Usage

Usage must to be connected with public bioinformatics infrastructures like the French Bioinformatics Network RENABI (www.renabi.fr) and especially its grid infrastructure GRISBI (www.grisbio.fr). The adoption of clouds for bioinformatics applications will be strongly correlated to the capability of cloud infrastructures to provide ease-of-use and access to reference biological databases and common bioinformatics software. In that sense, StratusLab is collaborating with RENABI to help fulfill the requirements of the Bioinformatics community.

Centre National de la Recherche Scientifique, France
Universidad Complutense de Madrid, Spain
Greek Research and Technology Network S. A., Greece
SixSq Sàrl, Switzerland
Telefonica Investigacion y Desarrollo S.A., Spain
Trinity College Dublin, Ireland

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

StratusLab

Enhancing Grid Infrastructures with Virtualization and Cloud Technologies

MapReduce on StratusLab with SlipStream™

MapReduce is a patented software framework introduced by Google in 2004 to support distributed computing with large data sets on clusters of computers. Apache Hadoop is an open source implementation of MapReduce. StratusLab users can easily run MapReduce applications with Hadoop, using the available public services (reference IaaS cloud, Marketplace) and the power of the SlipStream™ platform.

Development Steps

1. Define machine images

- Reference Hadoop appliance from Cluster module.
- Develop necessary logic that prepares environment for passwordless SSH logins and invokes application

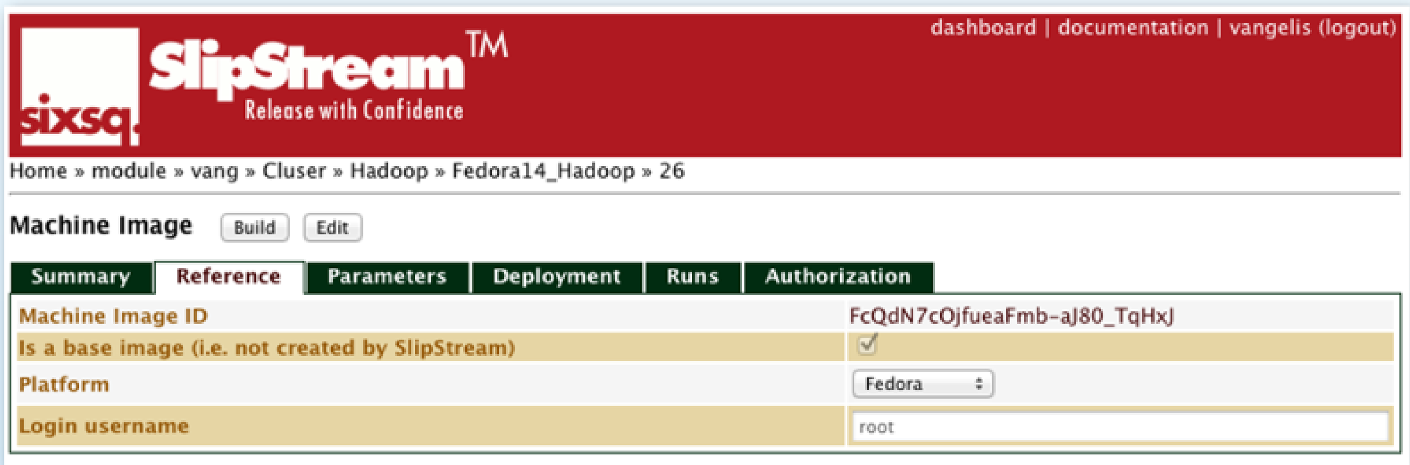

2. Define deployment

- Select instance type per node (master, worker)
- Set worker cardinality
- Define master/worker dependencies

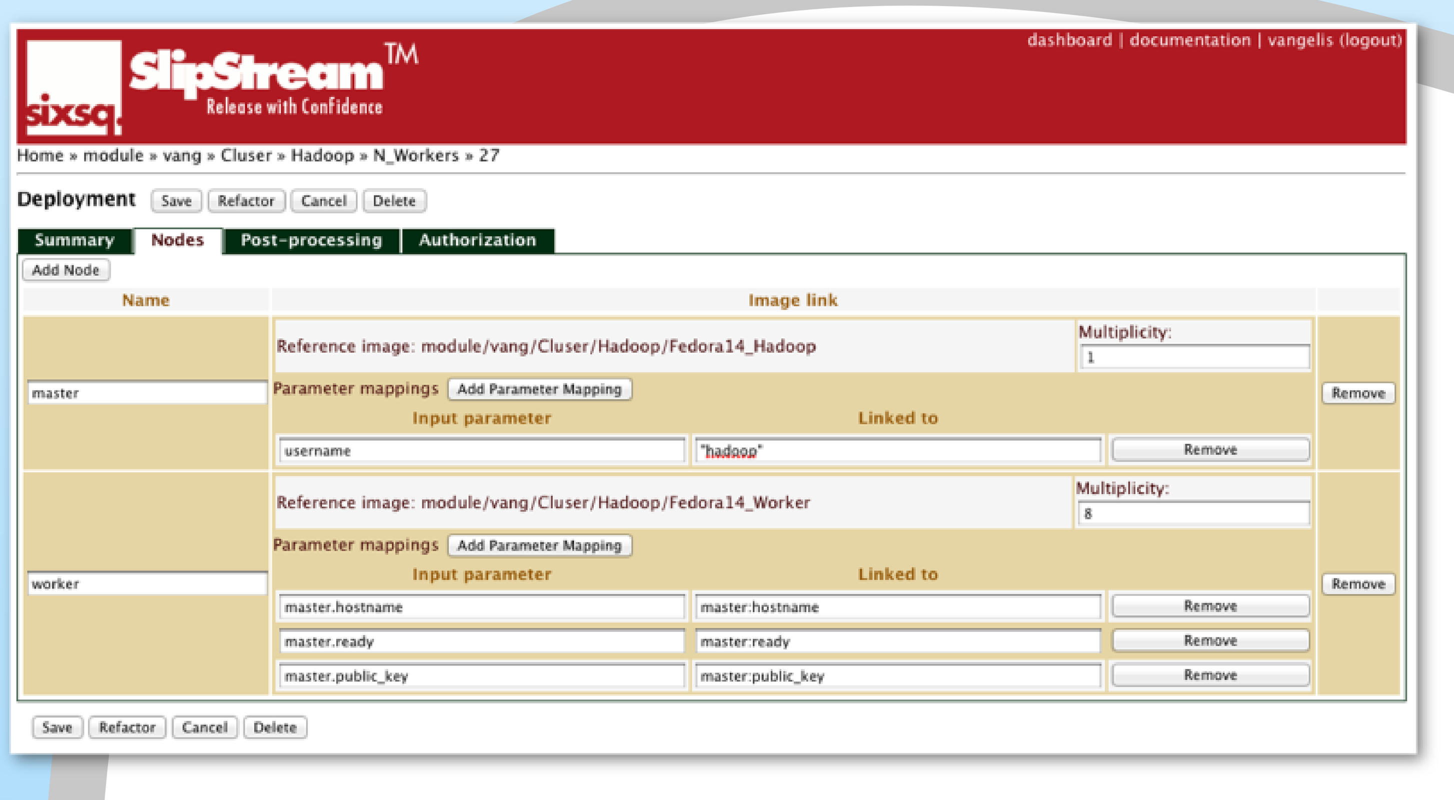

3. Run deployment

- Monitor deployment progress
- Hadoop web console available from master node
- Application output available for off-line retrieval upon termination

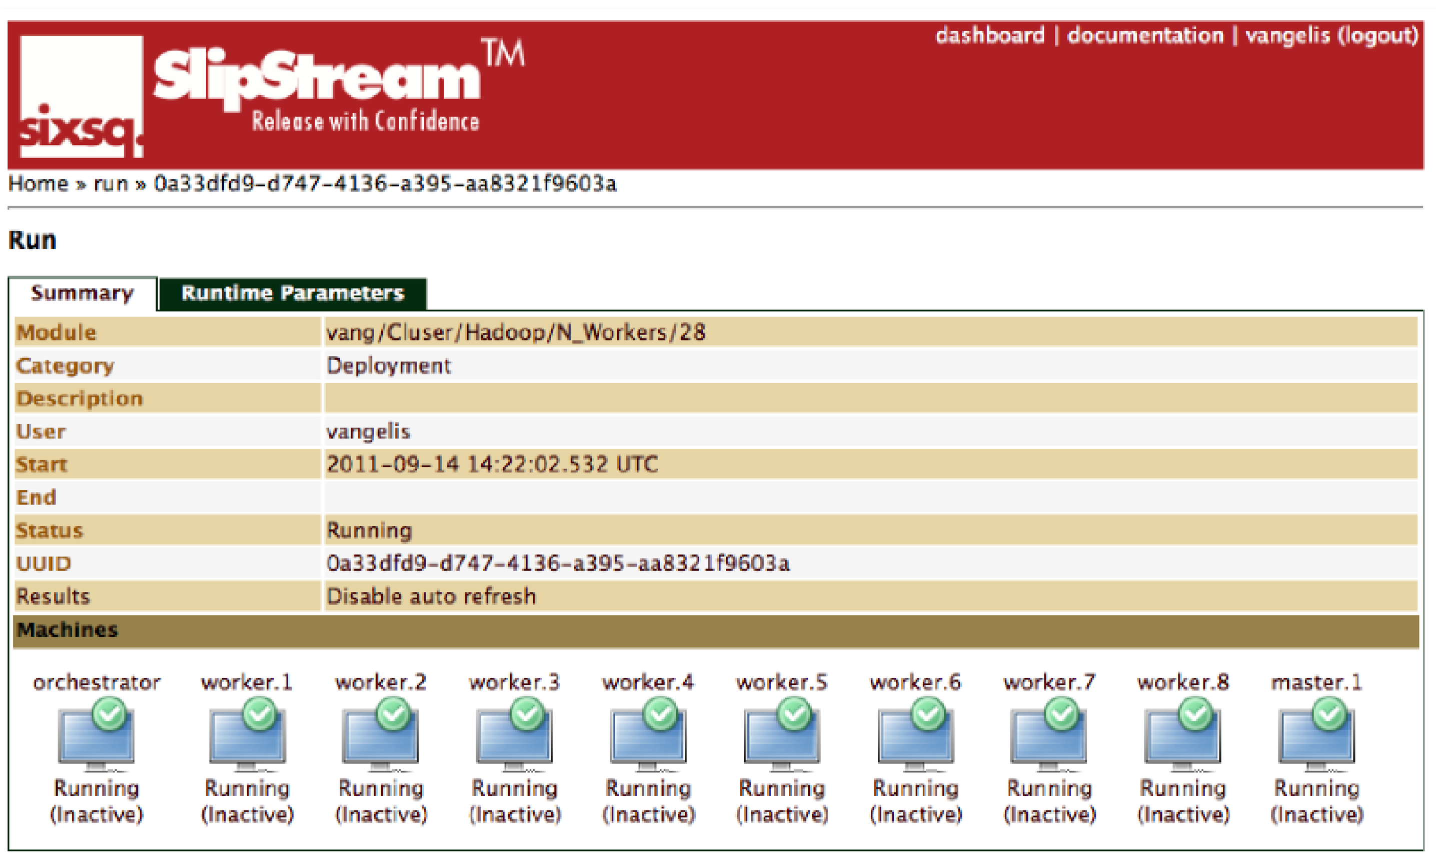

Hadoop appliance available from StratusLab Marketplace

- Fedora 14 base OS
- Hadoop 0.20.203.0
- JDK 1.6 runtime

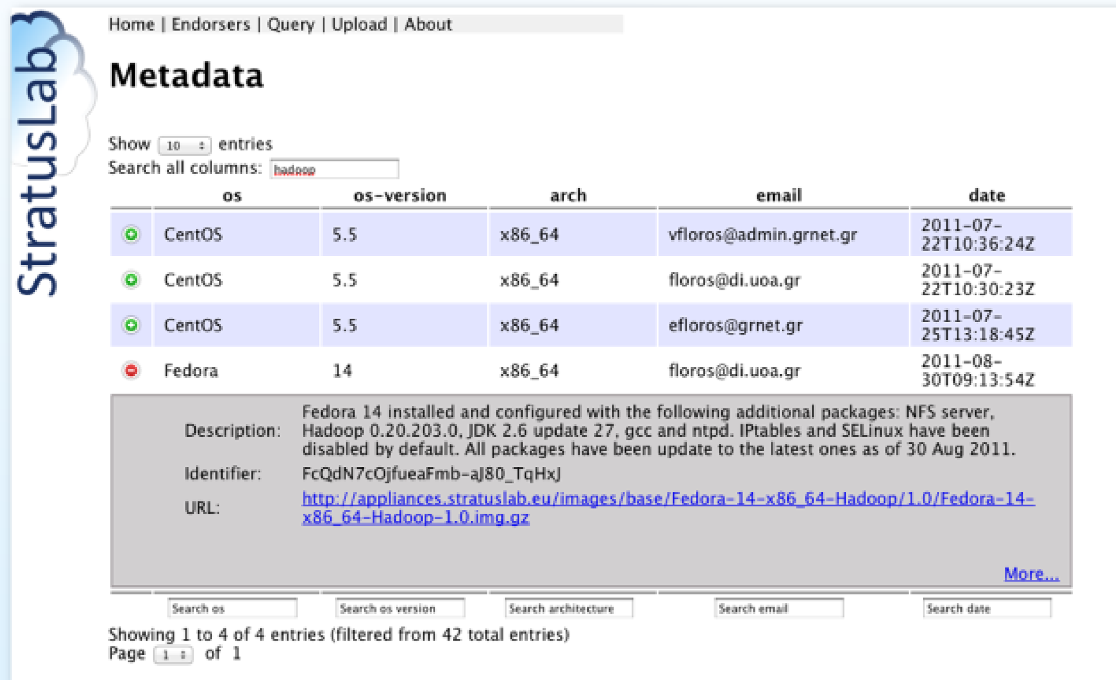

Pre-configured Hadoop Cluster Architecture

- 1 Master node - Running JobTracker and NameNode service
- N workers acting as DataNodes and TaskTrackers
- Other setups can easily be configured

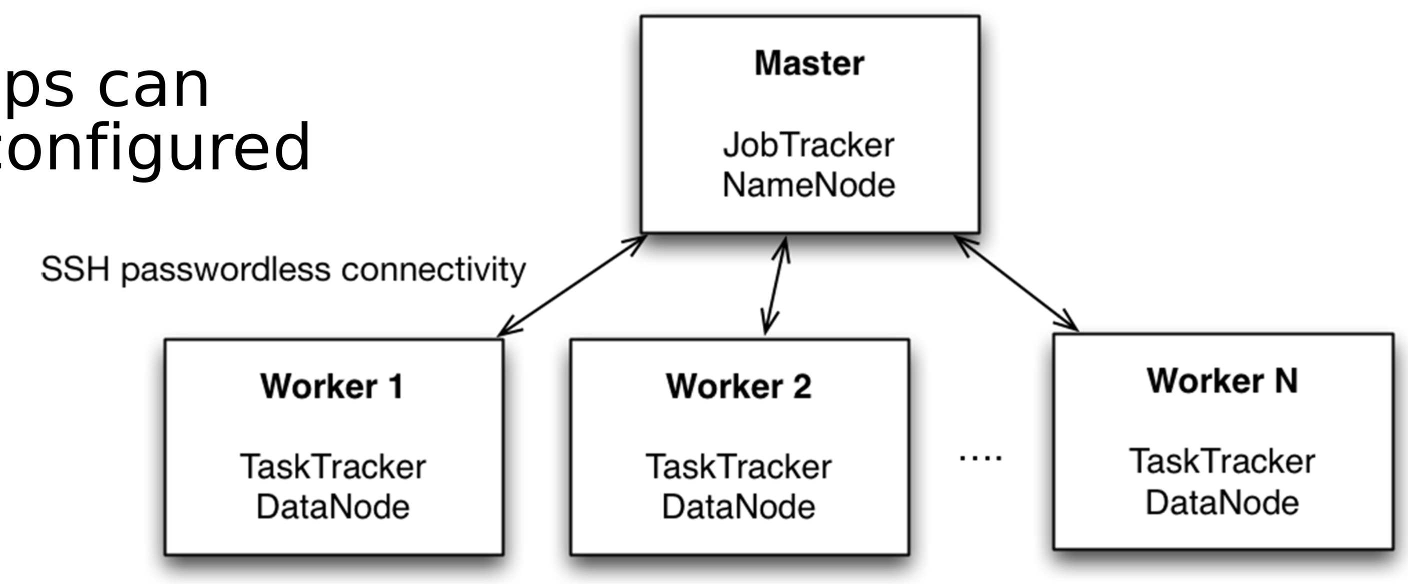

Application monitoring

- Hadoop web console available from master node

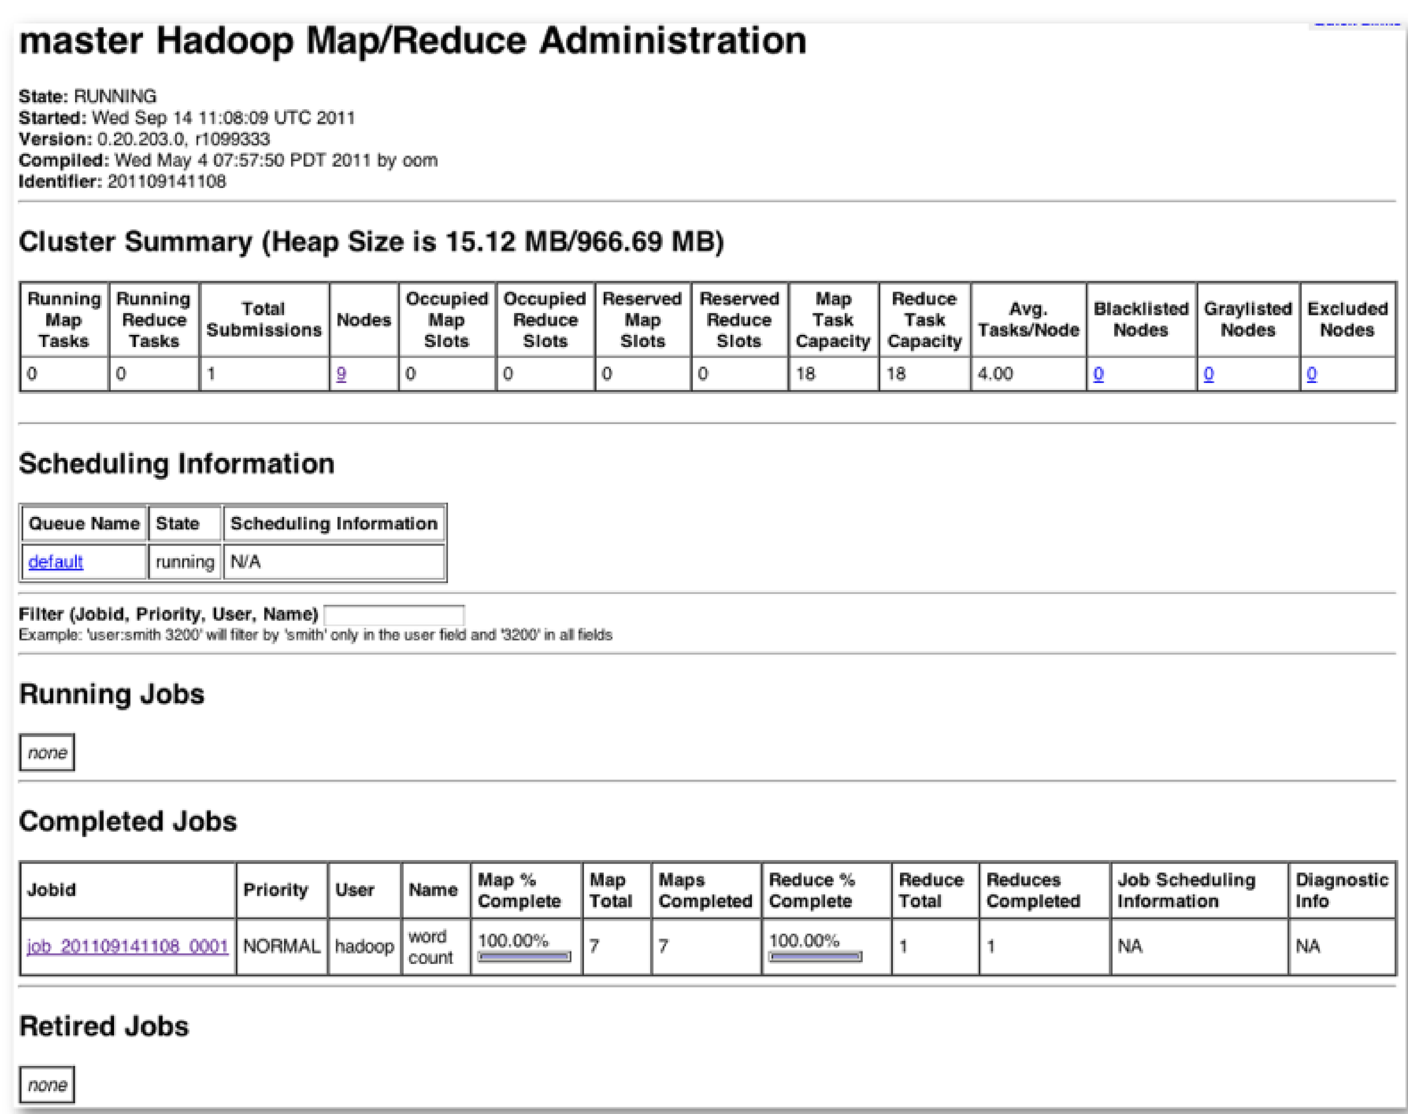

Website: http://stratuslab.eu
Support: support@stratuslab.eu
Contact: contact@stratuslab.eu

Centre National de la Recherche Scientifique, *France*
Universidad Complutense de Madrid, *Spain*
Greek Research and Technology Network S. A., *Greece*
SixSq Sàrl, *Switzerland*
Telefonica Investigacion y Desarrollo S.A., *Spain*
Trinity College Dublin, *Ireland*

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

SlipStream™ Release with Confidence

On-Demand Creation of Runtime Environments on Clouds

SlipStream™, by SixSq, deploys runtime environments in a cloud automatically, on-demand, and releases the resources after use. SlipStream™ users do not have to worry about configuring complex hardware and software to create personal runtime environments, representative of production conditions.

SlipStream™ is now available both on StratusLab* and Amazon EC2 for cloud resources. This allows SlipStream™ to be deployed on private infrastructures. The SlipStream™ deployment at *slipstream.stratuslab.eu* is free of charge for all StratusLab users.

How does it work?

SlipStream™ is accessed via a simple Web 2.0 application. Users can organize their work in project hierarchies. Each SlipStream™ artifact is version controlled, providing complete history of changes.
Two main workflows are available in SlipStream:

1. *Virtual machine (and disk) factory*: Allows a user to specify the content of each virtual machine.

2. *Deployment and service orchestration*: Control of and synchronization between inter-dependent software components.

SlipStream™
Release with confidence

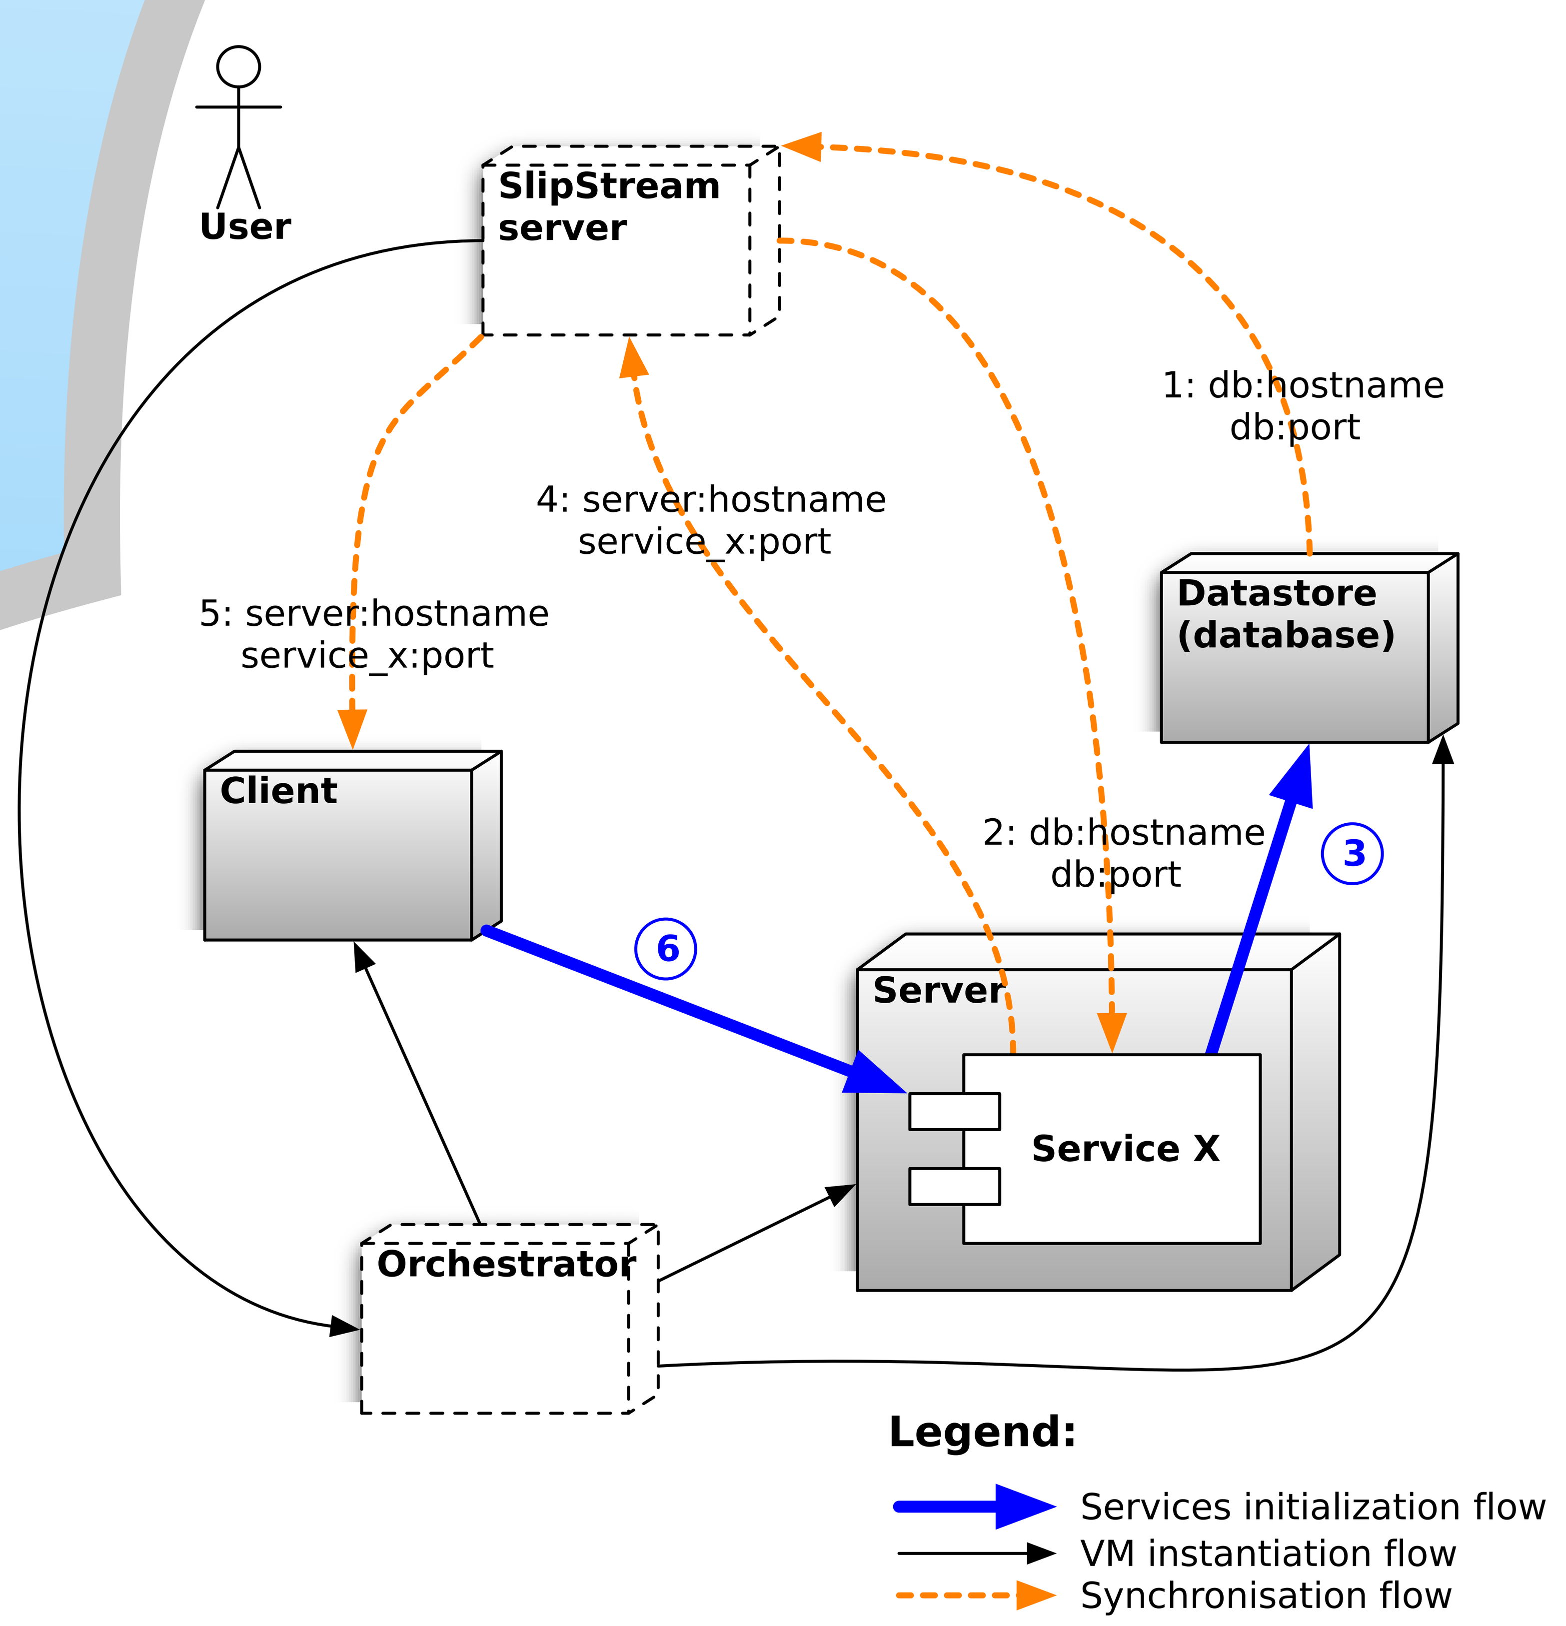

Simple deployment overview

Hybrid Cloud Solutions

SlipStream™ is ready for multi-cloud usage, an important use case for the second year of StratusLab.

Users will soon be able to leverage their virtual machine creation recipes, as well as their deployment models, across different clouds.

Site Administrators the possibility to re-generate virtual machines for different cloud services, on-the-fly.

"SlipStream™ is to software what the robotic assembly line is to manufacturing"

Continuous integration and deployment on clouds case studies

SlipStream™ is currently being evaluated in the context of ground control software systems at the European Space Agency. This work includes streamlining their development and release processes, such that their time-to-market is dramatically reduced.

Using SlipStream™, people can automate the creation of virtual applicances and deployments, such that full sites, can be deployed within minutes. For example, deployment of European Grid Infrastructure (EGI) sites via StratusLab open-source private IaaS cloud distribution.

"Declare war on pre-release stress. Release often and under your own terms. Improve software quality. Reduce time-to-market. These are a few reasons why we hope you'll enjoy using SlipStream™!"

* StratusLab (http://stratuslab.eu/) is developing a complete, open-source IaaS cloud distribution, and is particularly focused on enhancing distributed computing infrastructures such as the European Grid Infrastructure (EGI).

Website: http://sixsq.com
Support: support@sixsq.com
Contact: contact@sixsq.com

SixSq Sàrl
Switzerland