



## **StratusLab Cloud Distribution**

StratusLab Tutorial (Orsay, France)

28 November 2012

# Infrastructure as a Service (IaaS)

## Abstraction

- Access to remote virtual machines
- Aimed at service providers

## Advantages

- Customized environment
- Simple and rapid access
- Access as “root”
- Pay-as-you-go model

## Disadvantages

- Non-standardized and multiple interfaces (vendor lock-in)
- Virtual machine creation is difficult and time-consuming

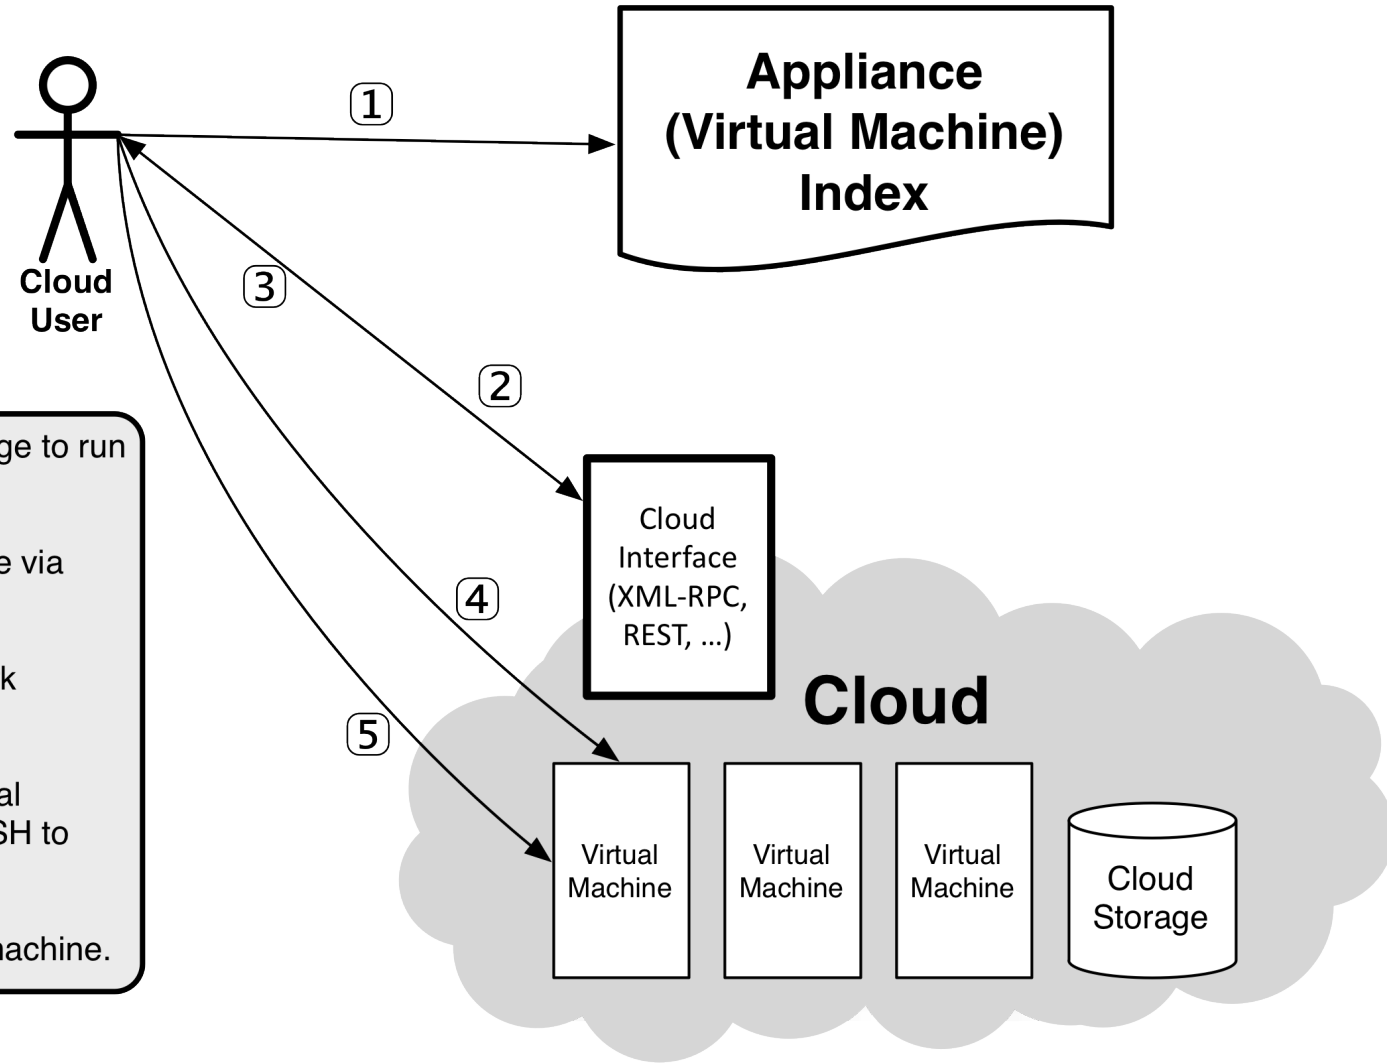
Software as a Service (SaaS)

Platform as a Service (PaaS)

Infrastructure as a Service (IaaS)



# Using an IaaS Cloud



1. Find virtual machine image to run on the cloud.
2. Launch machine instance via the cloud front-end.
3. Obtain machine's network address.
4. Use and control the virtual machine as usual. (E.g. SSH to machine as root.)
5. Shutdown or kill virtual machine.

---

# Why use a cloud?

---

# Customized Environment

Operating system(s) suited to your application

Your applications preinstalled and configured

CPU, memory, and swap sized for your needs



## Deluge of Data

- Huge databases that grow quickly in size and number

## Numerous Applications

- Written for different environments and having different requirements

## Variety of Databases

- Different topics (genetics, proteins, ...), subject to continuous updates

***Challenge: Creating the proper environment to use them simultaneously for a complete analysis.***

# Customized Appliances

## Biocompute

- Common bioinformatics applications and libraries preinstalled
- BLAST, ClustalW2, FastA, ...

## Customized portal

- Easy access to cloud infrastructure
- Streamlined for bioinformatics use

## Biodata

- Gives access to the latest updates of common databases
- SwissProt, Prosite, ...

You are signed in as cblanchet | [Settings](#) | [Home](#) | [Help](#) | [Sign out](#)

**idéeB** Bioinformatics cloud Powered by StratusLab

Shutdown [Go] New Instance Refresh Room for VMs

ID	Name	State	Appliance	CPU%	CPU	Mem. (GB)	Storage	Port translation
<input type="checkbox"/> 1149	test	Running	BioData	2%	1	4		http
<input type="checkbox"/> 1199	upg	Running	ARIA2.3	5%	4	16		ssh
<input type="checkbox"/> 1239	qr7	Running	BioCompute	9%	2	8		ssh http
<input type="checkbox"/> 1246	Test2	Running	BioCompute	5%	2	8		ssh http
<input type="checkbox"/> 1247	Pdisk	Running	BioCompute	3%	4	16	chb	ssh http
<input type="checkbox"/> 1249	portal	Running	Galaxy	5%	2	8		http
<input type="checkbox"/> 1250	portal	Running	Galaxy	10%	2	8		http
<input type="checkbox"/> 1254	Structural bio	Running	IDB NMR	5%	2	8		ssh
<input type="checkbox"/> 1255	tosceni	Running	ARIA2.3	6%	24	16		ssh
<input type="checkbox"/> 1256	Biompute with data	Running	BioCompute	3%	4	16	test-context	ssh http
<input type="checkbox"/> 1257	Genome analysis	Running	Galaxy	5%	8	32		http
<b>11</b>			<b>5</b>		<b>55</b>	<b>140</b>		

**Room for VMs**  
 xsmall 63 / 144  
 small 59 / 104  
 medium 27 / 48  
 large 11 / 20  
 xlarge 3 / 6  
 bigmem 1 / 2  
 xxl 1 / 6  
 htc 0 / 4

**Instance** me (78.57%)

**Cpu** them (18.06%)

**Memory** free (63.74%)

IDB acknowledges co-funding by the European Community's Seventh Framework Programme (INFOS-RI-261552) and the French National Research Agency's Arpege Programme (ANR-10-SBGI-001)

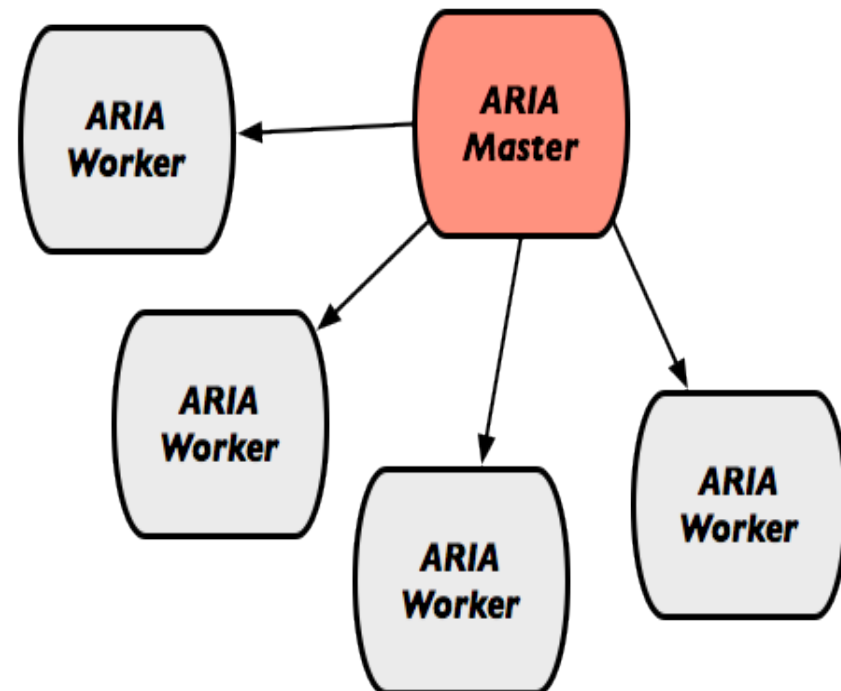
- IDB | Mentions légales -

## “TOwards StruCTural AssignmeNt Improvement”

- Improve the determination of protein structures based on Nuclear Magnetic Resonance (NMR)

### ARIA

- Calculates structures based on NMR
- Variable resource utilization during calculation





---

# Dynamic Provisioning

New storage and compute resources in minutes

Used resources freed just as quickly

Ideally suited to variable workloads

# Dynamic Training Infrastructure



## Benefits

- No need for a dedicated training infrastructure; create it as needed.
- Provides students each with their own personal playground

## NARVAL

- Data acquisition software used by several nuclear physics institutes, like IPNO, INFN, GSI, Ganil, ...
- Understanding the system requires installing, using and customizing it



***Successfully trained 20 NARVAL students using the StratusLab cloud infrastructure.***

# Dynamic Testing Infrastructure

## Developers Dream

- Easily test software in multiple environments
- Deploy full, independent test systems without interference from others



## SlipStream

- Software engineering PaaS from SixSq allowing automated deployment and testing of complete software systems
- Uses the StratusLab cloud distribution for dynamic resource provisioning

---

# Flexible Service Deployment


Deploy user-level, network-accessible services

Create domain-specific analysis platforms



## HEP Experiments

- Have 1000s of users, analyze PB of data, and use 100s of sites around the world
- Develop and maintain their own software services and platforms to manage the data and control access to it



## Grid Infrastructure

- Heavily used to support HEP data management and analysis
- Consists of a large number of complex services


## Grid site in a StratusLab cloud

- Production site maintained for nearly 2 years
- Demonstrates cloud's ability to support complex services


# StratusLab History



Informal collaboration to investigate running grid services on Amazon EC2 (2007)



Project (June 2010 to May 2012) co-funded by EC with 6 partners from 5 countries



*Open collaboration* to continue the development and support of the StratusLab software

Website: <http://stratuslab.eu>  
Twitter: @StratusLab  
Support: [support@stratuslab.eu](mailto:support@stratuslab.eu)  
Source: <http://github.com/StratusLab>

# StratusLab Principles



## **Simple to deploy and simple to use!**

- Command line client for all major platforms
- Web interface for most services
- REST interfaces for programmers

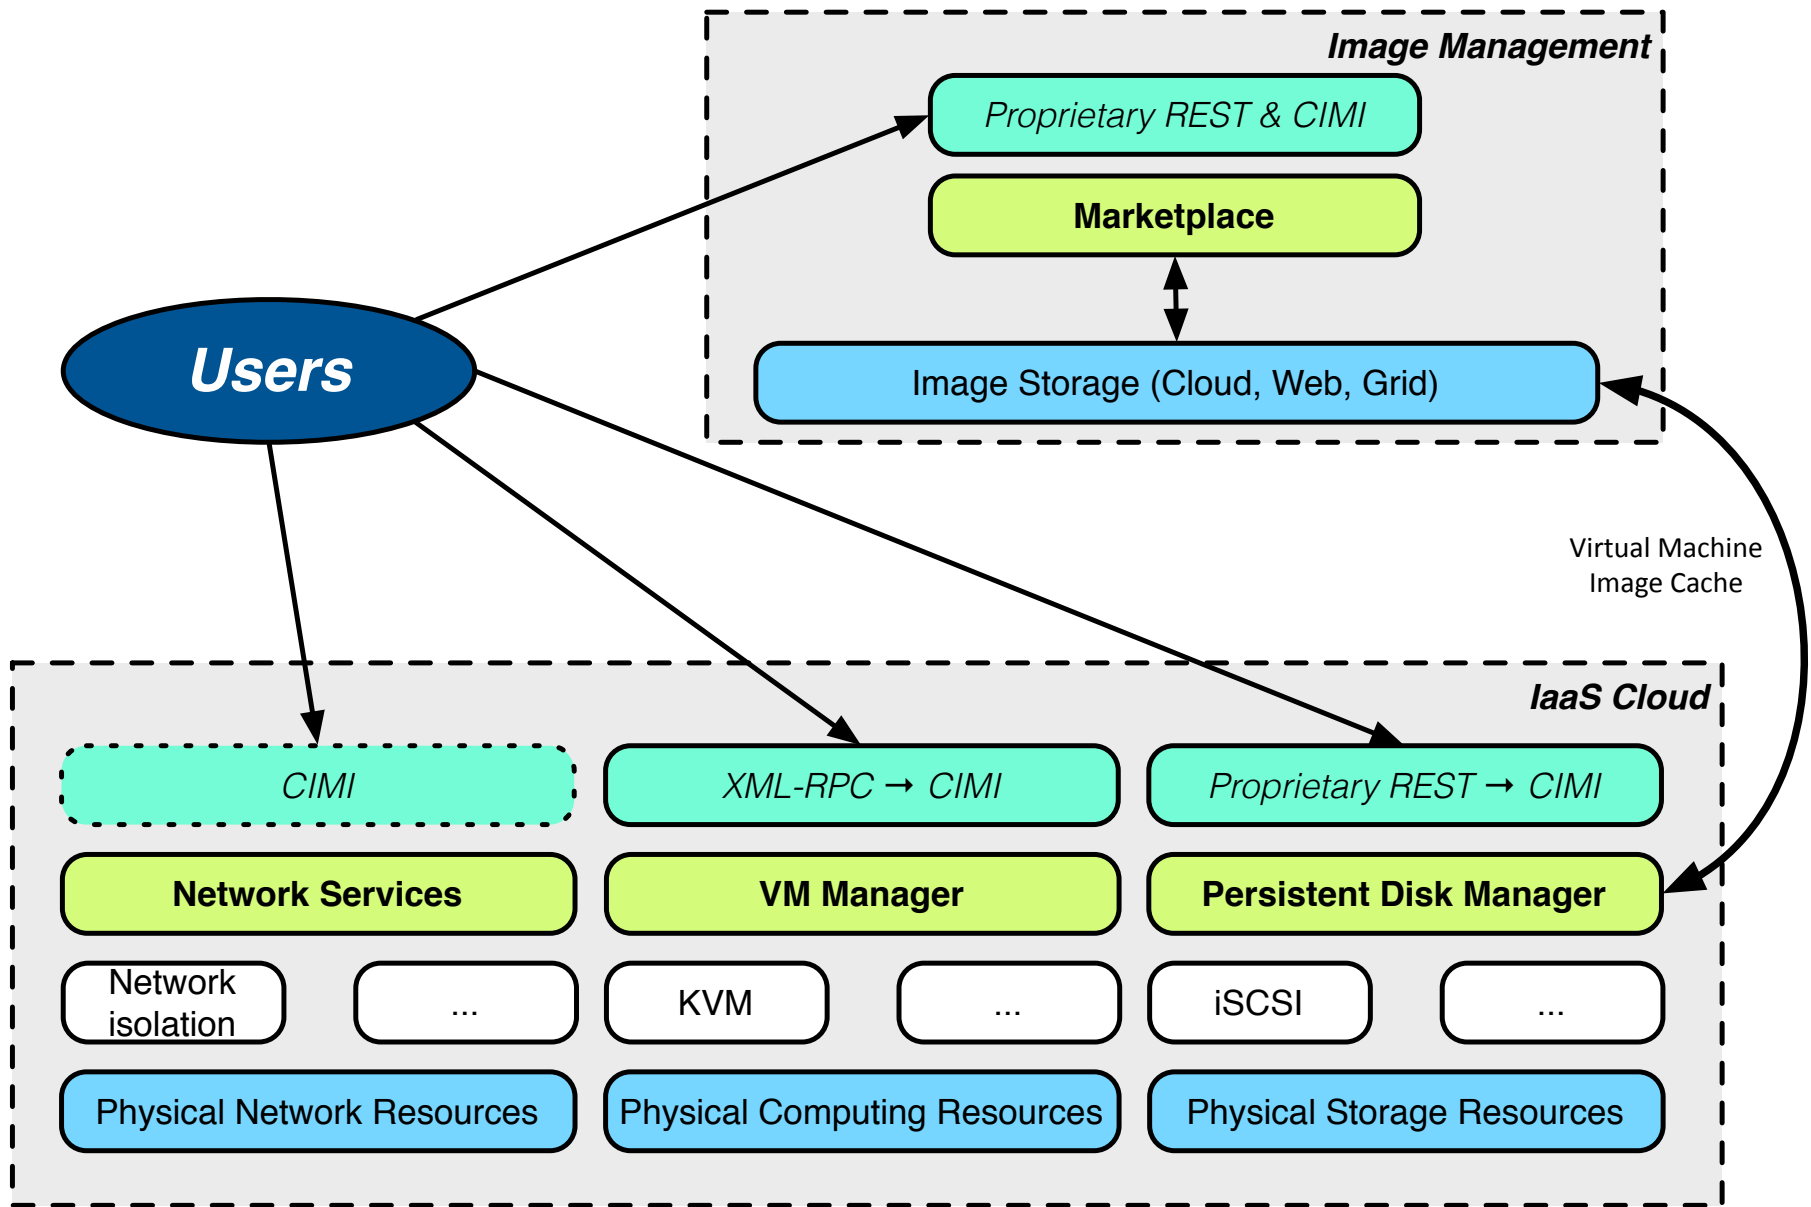
## **Focused, practical development**

- Develop for real needs of users
- Focus on providing IaaS services well
- License (Apache2) that allows academic and commercial use

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

- Iterative integration: always maintain working distribution
- Public releases approximately every 3 months
- Welcome contributions by institutes and individual developers

# StratusLab Architecture







## Virtual Machine Management

- OpenNebula ([opennebula.org](http://opennebula.org)) provides core of VMM (start, stop, kill)
- Plug-in architecture allows use of multiple hypervisors (kvm, ...)

## Features

- StratusLab-specific caching allows low-latency startup of VMs
- Quarantine of stopped images for forensic analysis
- Detailed logging of user and resource information
- Improved error feedback from plug-ins to user
- Integration with StratusLab user management



## 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

# Other Storage Types



## File-based Storage

- Normal client tools can be installed in VMs
- Access services normally from VM (e.g. tools for SRM)
- May provide CDMI interface to StratusLab storage

## Object Storage

- Simple object storage, usually minimal hierarchy and chunked data
- StratusLab works well with Swift from 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

***Unlikely to see StratusLab implementations, but distribution may facilitate co-deployment of such services from others.***



## IP Address Classes & Selection

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

## IPv6

- Software supports IPv6 networking for VMs
- Not available on reference infrastructure yet

## Future Services

- User specified firewalls
- IP address reservation
- Dynamic VLANs



## 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



## Virtual machines with pre-installed/configured services

- Makes it easier to get started quickly using cloud resources
- Good way to package software to avoid installation hurdles

## Utilities for making and publishing customized images

- Security guidelines
- Incremental changes to base image
- Marketplace for publication

## StratusLab supported appliances

- Base images: ttylinux, CentOS, OpenSuSE, Ubuntu, Debian
- Bioinformatics: Data server and analysis images



## 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

# Accessing Services



## StratusLab Client

- Command line scripts in python/java with minimal dependencies
- Fully tested on Mac OSX and Linux
- Core functionality works on Windows
- **Be sure to use client version corresponding to cloud infrastructure**

## Programming Interfaces

- All services except VMM provide proprietary REST interface
- VMM exposes OpenNebula XML-RPC interface
- Transitioning to using CIMI for all services (incl. VMM)
- StratusLab plugin for rOCCI impl. of OCCI is available

## Web Interfaces

- Provided for all services except virtual machine management
- Expect to have complete, unified interface after transition to CIMI





## Reference Cloud Infrastructure

- Allow users to test a StratusLab cloud without having to install one
- Two sites: LAL (Orsay, France) and GRNET (Athens, Greece)
- Registration Service accounts work on both sites
- Problems, ask questions via [support@stratuslab.eu](mailto:support@stratuslab.eu)

## Other StratusLab Infrastructures

- South Africa
- Vietnam
- UK
- Bordeaux
- ...

---

# Questions and Discussion



## StratusLab Website

- Find information for system administrators
- Find information for users

## Code

- <http://github.com/StratusLab>
- Can you find the code related to the various types of resources?
- Any code that you can't understand in a IaaS context?



<http://www.stratuslab.eu>

Copyright © 2012, Members of the StratusLab collaboration.

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

