



## **StratusLab: Darn Simple Cloud**

Charles (Cal) Loomis (CNRS/LAL & SixSq Sàrl)  
FOSDEM'13: Cloud Devroom (3 February 2013)



## What is it?

- Complete Infrastructure as a Service (IaaS) cloud distribution
- Open source with Apache 2 license
- Solid IaaS for direct use and for PaaS/SaaS developers

## Focus: Darn Simple Cloud

- Simple to install, using commodity hardware
- Simple to use, from (most) any client machine
- Scales down as well as up!

# Bridging the Chasm...



## Users

- How to use virtual machines to get my work done?
- How to structure, store, access, and protect data?
- Realize shared infrastructures with customized env. are possible

## Application Developers

- How to use cloud techniques to improve my applications?
- ... and my development workflows?
- Applications can be services with pluses and minuses...

## Data Centers


- Reuse existing (commodity) hardware investments
- Take advantage of (and train) existing system administrators
- How to manage/use a (private, community, public) cloud?

**Significant benefits from cloud even without large scale elasticity!**

# Where did it start?

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


*Identified need for open source cloud distribution.*



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



*Production dist. with academic & commercial deployments.*

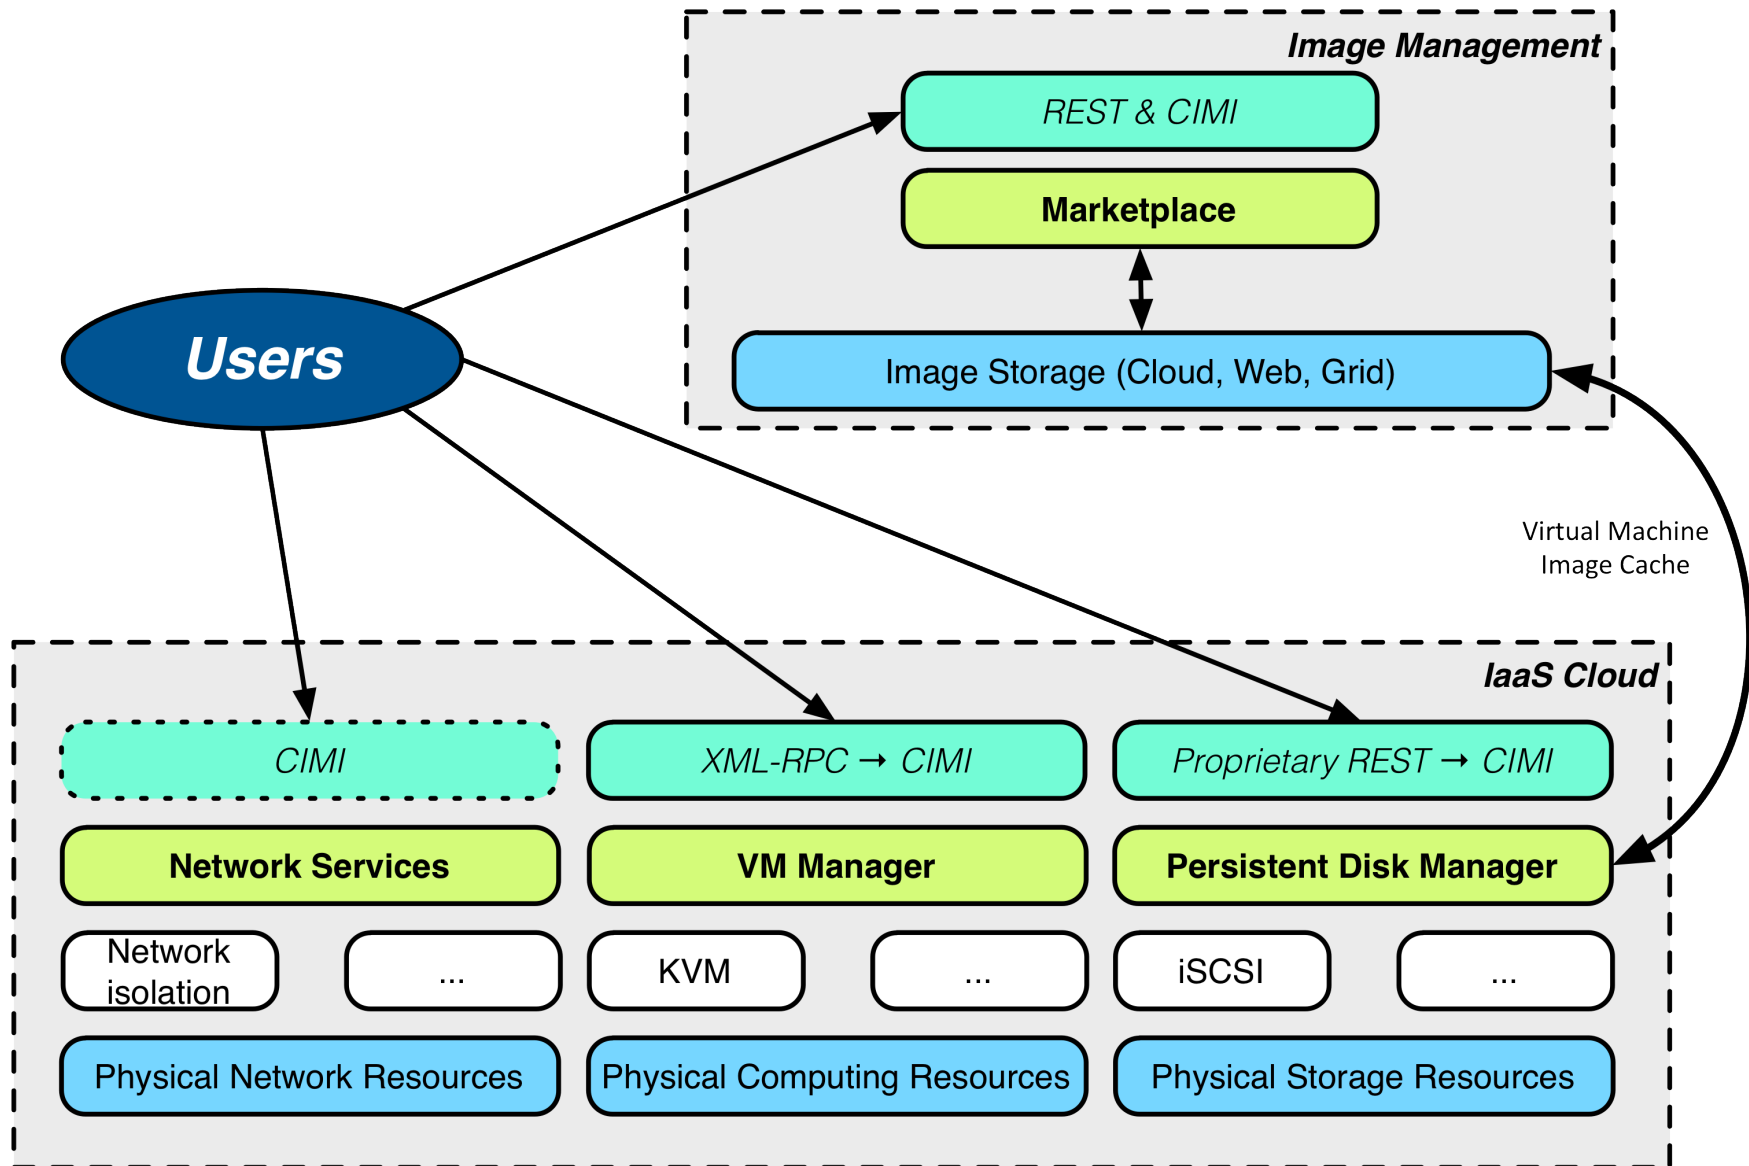


Website: <http://stratuslab.eu>  
Twitter: @StratusLab  
Support: [support@stratuslab.eu](mailto:support@stratuslab.eu)  
Source: <http://github.com/StratusLab>

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



# StratusLab Services



## Post-Project Releases

- V2.1 (16/10): Streamlined release with obsolete components removed  
Better support for virtio drivers to improve IO performance
- V2.1.1 (29/11): Bug fix release (storage upload feature) and better Windows support
- V13.02 (31/1): Support for CloudInit contextualization and bug fixes

## Release Policy

- Quarterly timed releases (13.02, 13.05, ...)
- Intermediate bug fix releases as needed
- Provide (6-month) roadmap for the StratusLab evolution



## Features

- Fast provisioning of virtual machines, with low latency start-up
- Availability of VM images of common operating systems

## Contextualization

- HEPiX & OpenNebula CDRROM contextualization by default
- CloudInit (disk based) also supported

## Implementation

- API: XML-RPC interface of OpenNebula
- OpenNebula (C++, Ruby) with customized hooks
- Hooks primarily for caching, snapshots, and storage access
- CentOS, Ubuntu, OpenSuSE, Debian, Fedora, ScientificLinux images created and supported by StratusLab



## Features

- Volume abstraction for storage service
- Provide users with persistent storage for data
- Serves also as cache of images for VM instances
- (No file-based or object-based storage service)

## Implementation

- API: Proprietary REST interface with CRUD actions
- Java-based service using MySQL database for state information
- Can use iSCSI or shared file system for physical storage
- Can use simple files or LVM volumes for disk content





## Features

- Support 3 specific use cases: public service (public), batch system (local), and BOINC-like worker (private)
- Dynamic configuration of network switches not needed
- Use standard services for (virtual) machine network configuration

## Implementation

- No API: manual, static configuration of network
- Rec. configuration: VLAN for cloud services separate VLAN for VMs
- All classes of IP addresses are optional, can create other classes
- Uses DHCP for VM network configuration
- Users responsible for protecting their machines

# Marketplace & Image Handling



## Priorities

- Mechanism for sharing and trusting images
- Possible to distribute fixed, read-only data sets as well
- Split the storage of image metadata and image contents
- Define roles for creator, user, administrator, and validator




## Implementation

- Marketplace API: Proprietary REST API for create, read, search
- Marketplace acts as image registry and handles only metadata
- Image contents can be located on any public (web) server
- 'Private' images can also be held in cloud storage

# Image Handling Workflow

1. Browse the Marketplace for useful images.
2. Launch machine instance using machine and disk identifiers.

**Marketplace**  
Registry for Machine & Disk Metadata



1. Define image authorization policy. Images can be banned based on checksums, endorsers, etc.
2. Policy is evaluated for each image used on the infrastructure.

1. Create a machine or disk image.
2. Upload image to cloud storage or any web accessible location.
3. Create and sign the metadata for the image.
4. Upload the metadata to the Marketplace

## Command Line Client

- Administrator: simplifies StratusLab installation
- Users: access StratusLab cloud from anywhere

## Administration

- Quarantine for stopped virtual machines
- Monitoring of cloud activity and resources

## Authentication and Authorization

- Supports username/password, certificates, cert. proxies
- Specification in local file and/or LDAP



## Information from Developers

- Web site documentation and support mailing list
- Live tutorials (usually 2-3 per year)
- Workshops (next week of 11 March 2013 in Paris)



## Reference Cloud Services

- (~)Open infrastructures for using StratusLab and providing feedback
- Operated on a first-come, first-served, best-effort basis
- In production 2+ years, with 250+ registered users
- Two sites: LAL (Orsay, France) and GRNET (Athens, Greece)

## Other deployments...

- Academic: France, Ireland, UK, Vietnam, South Africa, ...
- Commercial: Atos, Helix Nebula, ...

## Building on top...

- SlipStream: cloud based systems deployment and testing
- Go-based CLI ...

# Priorities for Evolution



## Interfaces

- Provide Libcloud (python) driver for StratusLab
- *Adopt CIMI as the standard interface to services*
- Provide complete browser interface for all services
- Provide EC2 adaptor to REST interface

## Simplicity, Scalability, & Robustness

- *Direct use of libvirt as VM manager*
- *Distributed database (Cassandra, Couchbase) as information 'bus'*

## Better services for system administrators

- Improved overview and monitoring of infrastructure
- Fine-grained accounts for all resources
- Migration control

# Direct Use of Libvirt



## libvirt

- libvirt model fits very well with StratusLab
- Already using libvirt indirectly behind OpenNebula
- Will significantly simplify the distribution, esp. wrt storage volumes

## Plan

- Develop 'placement service' to replace OpenNebula scheduler
- Bind current VM management code to libvirt directly
- Modify services to take advantage of libvirt storage pools





## Use distributed database as coordination backbone

- Already have experience with Cassandra, looking at Couchbase also
- Will make it easy to scale both front-ends and hypervisor nodes
- Provides easy mechanism for reporting resource utilization
- Similarly for coordination between services

## Plan

- Evaluate the pros/cons of the DB solutions
- Implement prototype service (probably monitoring) as POC
- Move other services to DB, starting with VM management



## CIMI (DMTF DSP0263 <http://dmtf.org/standards/cloud>)

- Resource-Oriented API with reasonable mappings to JSON and XML
- IaaS-focused specification that will likely become international std.
- Defined mapping to OVF schema
- Full coverage of StratusLab services

## Plan

- Migrate services individually
- Keep CLI and Libcloud working throughout migration
- Changes should be transparent to users
- Developers using API will see backward-incompatible changes

**Aside: A working StratusLab plugin for rOCCI is available.**

# Conclusions



## StratusLab Cloud Distribution

- Supported, stable, and production-quality IaaS cloud distribution
- Used for reference cloud service for 2+ years
- Other academic and commercial deployments
- Defined, ambitious roadmap for its continued evolution

## StratusLab Collaboration

- New collaborators welcome! Developers and documenters!
- Weekly phone conference between developers
- Biannual StratusLab workshops (next week of 11 March 2013)

**Try it and provide feedback on software, services, support, ...**

---

# Questions and Discussion

Website: <http://stratuslab.eu/>

Twitter: @StratusLab

Support: [support@stratuslab.eu](mailto:support@stratuslab.eu)

Source: <http://github.com/StratusLab>



<http://stratuslab.eu/>

Copyright © 2013, 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/>).

