CernVM Online and Cloud Gateway

a uniform interface for CernVM contextualization and deployment

George Lestaris - Ioannis Charalampidis

D. Berzano, J. Blomer, P. Buncic, G. Ganis and R. Meusel

PH-SFT / CERN

Background

CernVM: a virtual appliance that serves as a portable environment for developing and running LHC data analysis

- easily deployed in *cloud* or *local* environment by supporting various image types and hypervisors
- see Jakob's talk about **µCernVM**, the next-generation CernVM: http://chep2013.org/contrib/213
- visit: http://cernvm.cern.ch

Background

Virtual machine contextualization: is the process of configuring a VM instance for the needs of various deployment use cases

- we apply a **context**
  - INI file - based on the AMIConfig library
  - defines settings for the various **contextualization plugins** - INI file sections
- using
  - *user-data* field in clouds that support it
  - *HEPIX* contextualization

Motivation

This work has to main goals:

1. make contextualization of local VMs easier
   - no need to write contexts manually
   - reduce the use of the *CernVM web appliance* to configure local virtual machines by providing easier alternatives
2. deployment of CernVM
   - in local environments
   - virtual clusters in the cloud

CernVM Online

contextualization made easy

cernvm-online.cern.ch

Can connect with CERN authentication or create a local account

CernVM Online

Web application: used to *define*, *store* and *share* **contexts**

- Contexts are immutable. They cannot change, they can only be **cloned**
- **Secure contexts** are supported by encrypting context content with user defined passphrase
    - context may contain sensitive information like passwords, keys, etc..

CernVM Online

Marketplace: share contexts for common use cases

Deployment: user can deploy VMs locally with the **CernVM WebAPI**

- browser plugin able to spawn virtual machines automatically using *VirtualBox hypervisor* (is installed by the plugin if not available)
- implemented by Ioannis Charalampidis for Theory division at CERN (http://crowdcrafting.org/app/cernvm/)

Pairing: contextualization of already running CernVM instances *with console access*

Dashboard

CernVM Online

About | Dashboard | Marketplace | Documentation | Downloads | Publications

Commands

- Dashboard
- Pair an instance
- Create Context
- Create Cluster
- Marketplace

Recent Definitions

- T4T-AtomSmasher
- T4T-Copilot-Server-v12
- Condor worker
- T4T-Client-5
- Condor master

Paired VMs

Dashboard

Your context definitions

	Name	ID	Operations	Deployment
	Condor master	2a6143b44dcb461bbc1688f385e4c61c	Clone Publish	WebAPI
	Condor worker	6302bff3261c46cfbac12b282d20f071	Clone Publish	WebAPI
	Local VM context	030cc9b469dd466bade2499b26c8091f	Clone Publish	WebAPI

Create new context...

Deploy VM locally

Publish to marketplace

Your cluster definitions

	Name	ID		
	Condor cluster	1e9227b4d185445e9edda31cdae4fe4d	Clone Publish	

Create new cluster

Your virtual machines

	Machine	CernVM	Context	Operations
	No instances paired yet			

Pair an instance of CernVM

Context creation

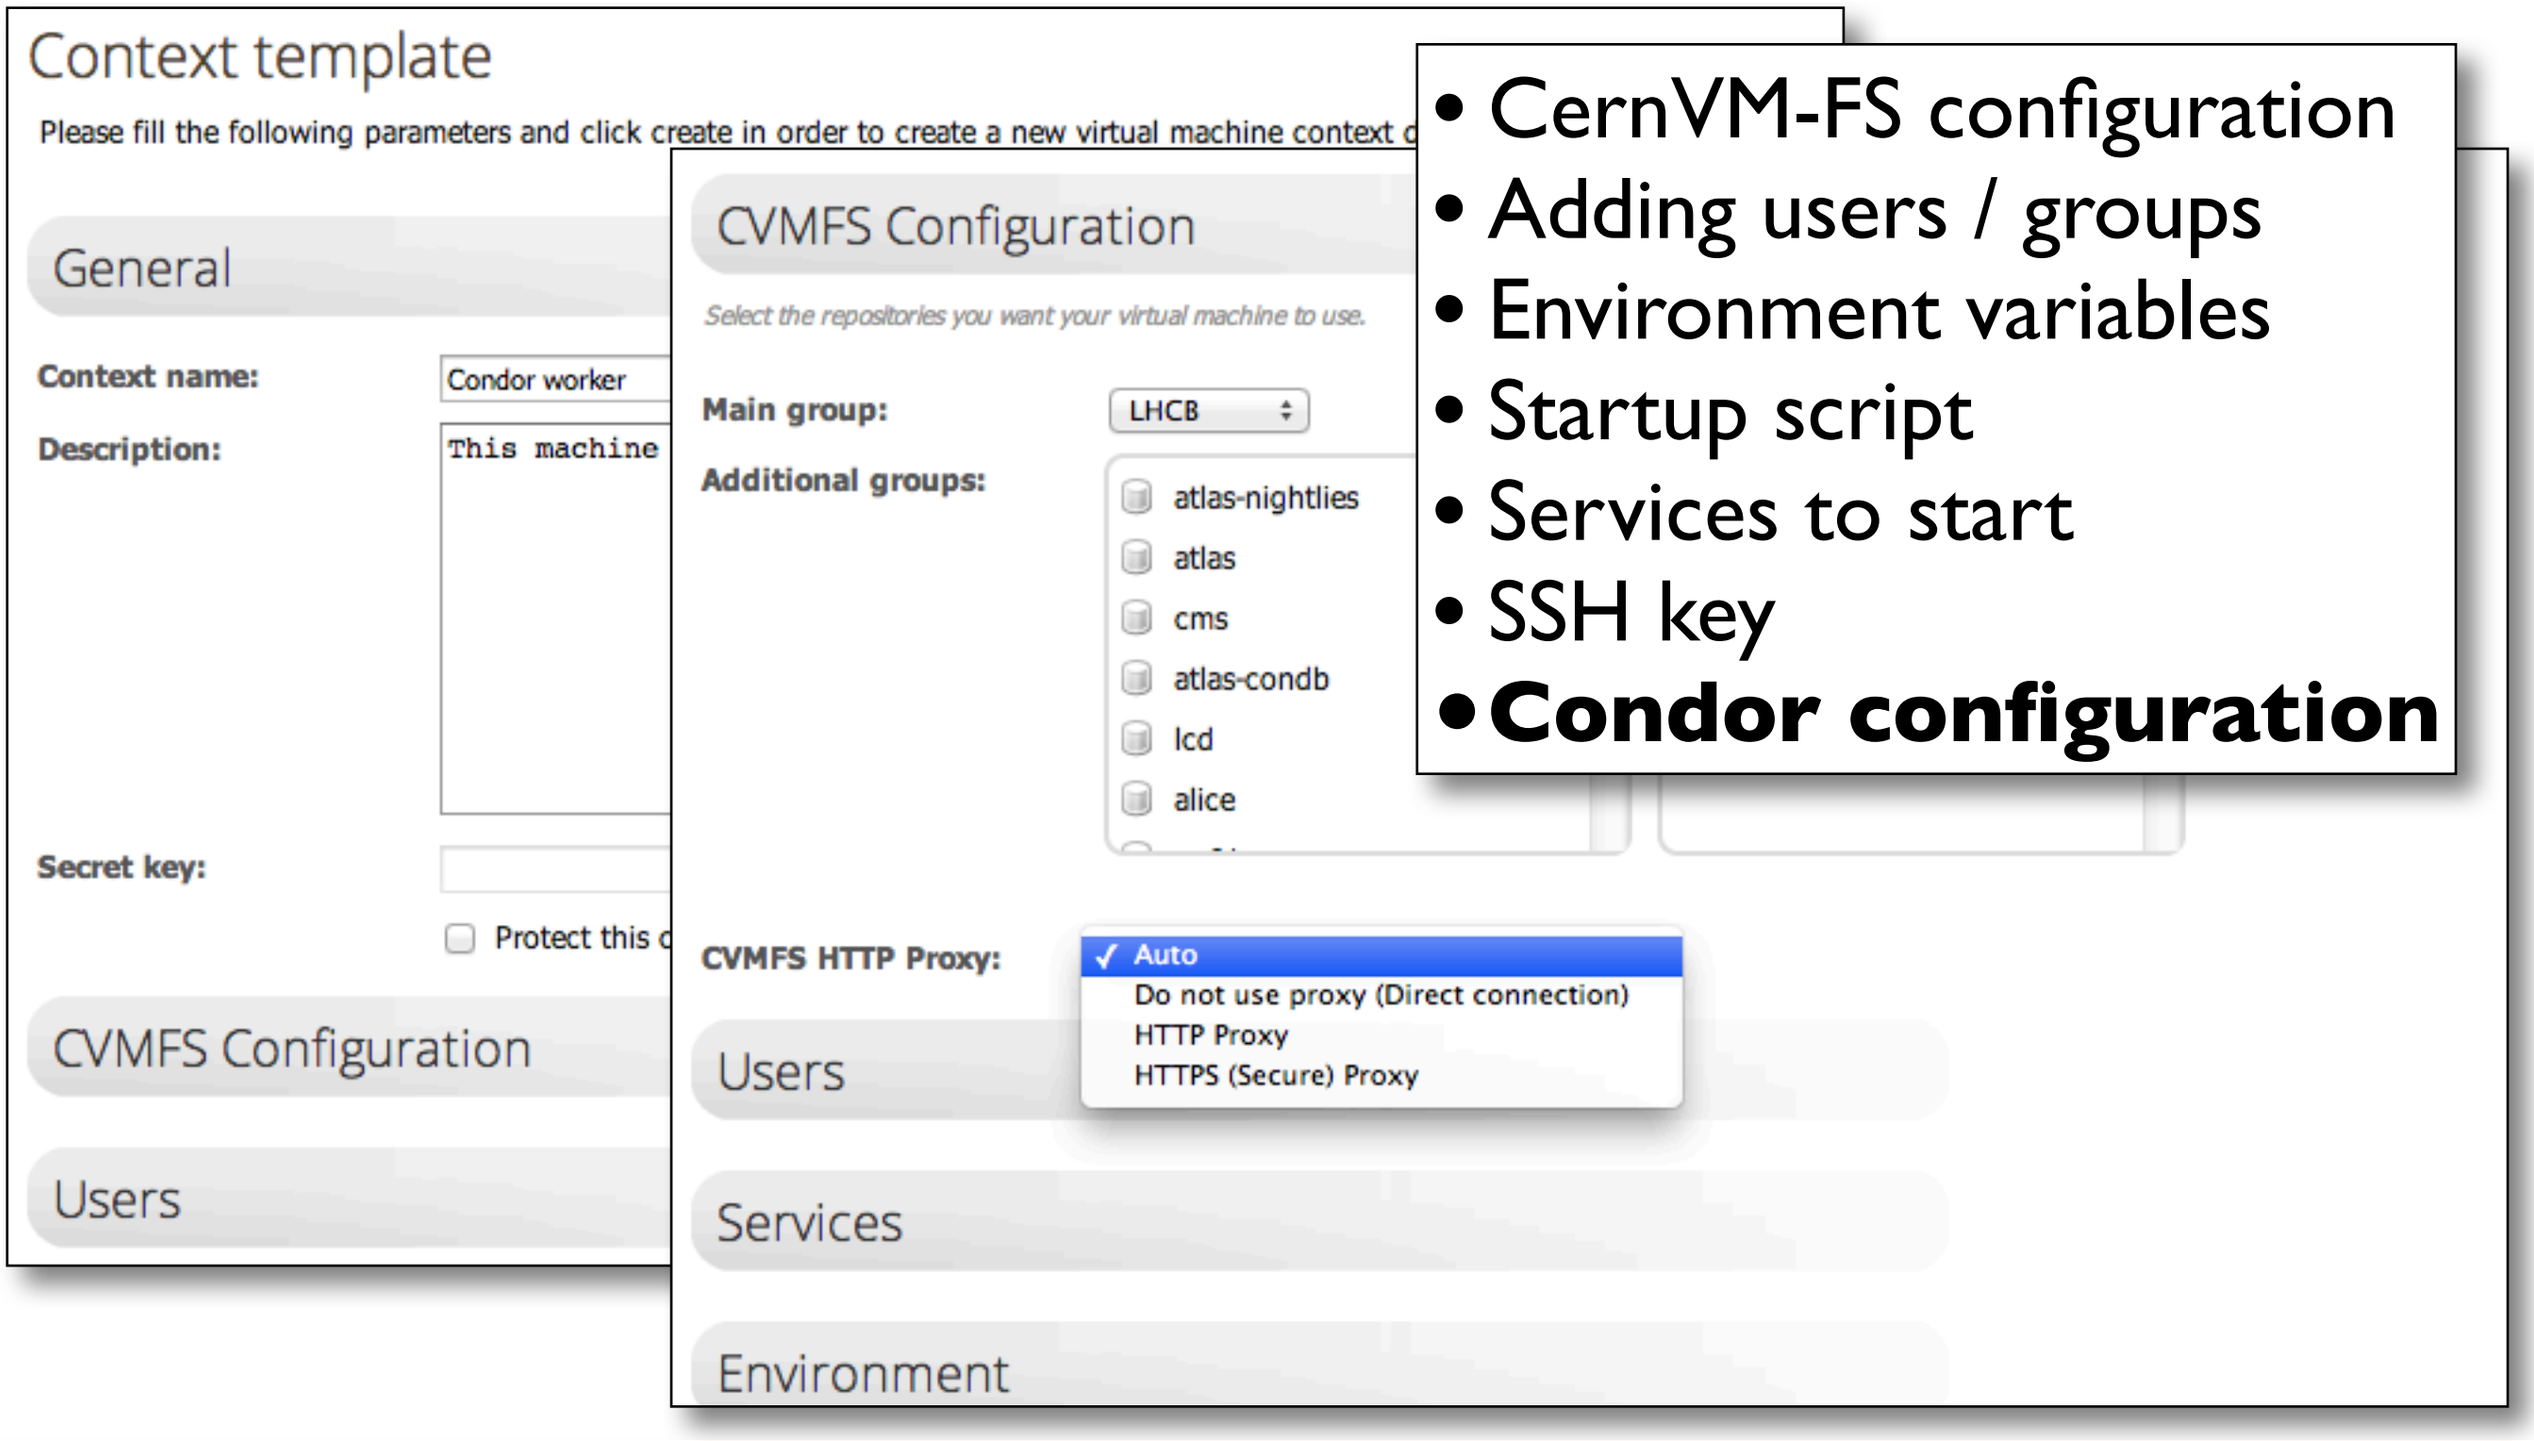

Marketplace

CernVM Marketplace experimental

Pick one of the public contextualization information and pair you CernVM instance.

T4T-Client-5

T4T-AtomSmasher

This is a pilot project from Test4Theory group that allows users to play while helping CERN physicists.

T4T-Copilot-Server-v12

User can pair/clone contexts made by others

Search for contexts →

T4T-AtomSmasher 0

The Virtual Atom Smasher is an interactive game from the Test4Theory group that allows people from outside CERN to help physicists, by playing an interactive game.

Author: icharala

Tags: t4t + - volunteer + -

Access: Open

Pair Clone

ATLAS
CMS
ALICE
LHCb
NA61
Experimental

Virtual machine pairing

Pair instance - Step 2

Wait your virtual machine to boot and put the following key in your virtual machine contextualization screen. You will have the option to contextualize it right after.

1) Obtain PIN from CernVM online

51 85 64

Local VM context VM [Running]

```
Welcome to CERN Virtual Machine, version 3.0.0.0
  based on Scientific Linux release 6.4 (Carbon)
  Kernel 3.7.10-17.cernvm.x86_
IP Address of this VM: 10.0.2.
In order to apply cernvm-online context, use #<PIN> as user
localhost login: #518564_
```

2) Apply PIN

Your virtual machines

	Machine	CernVM	Context	Operations
	128.141.235.236 (6c8e7a1b-a0cc-42cf-b9d2-d04423366c0f)	3.0.0.0	Local VM context	Unmanage

Pair an instance of CernVM

list of paired instances

Deployment with WebAPI beta

1) Select configuration

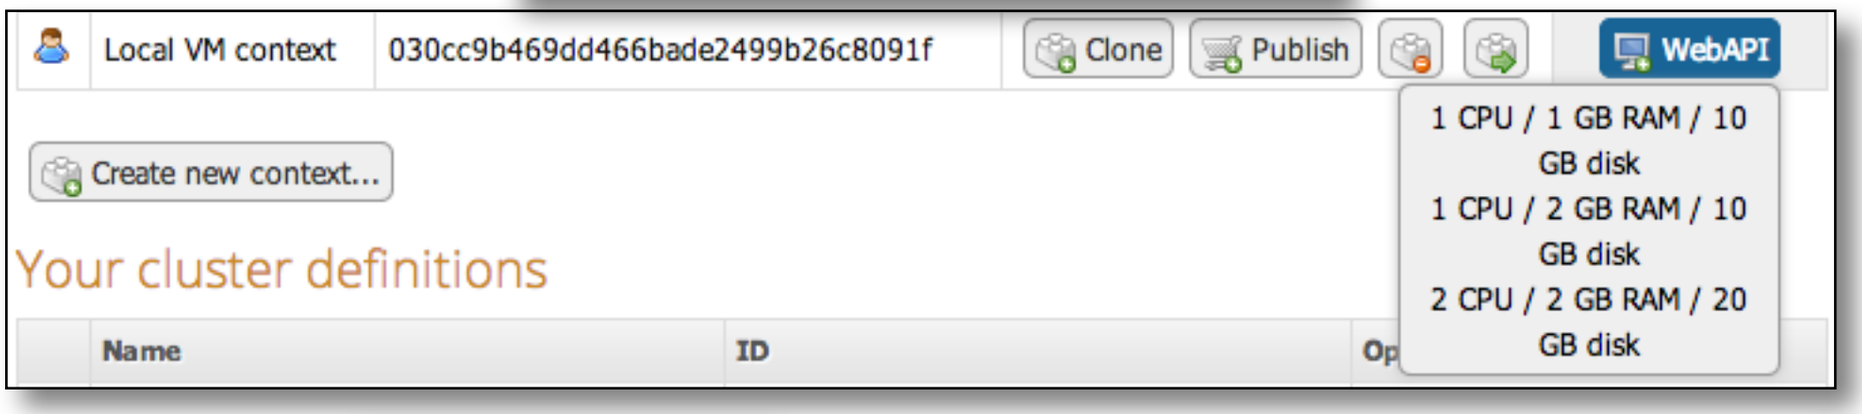

2) Confirm

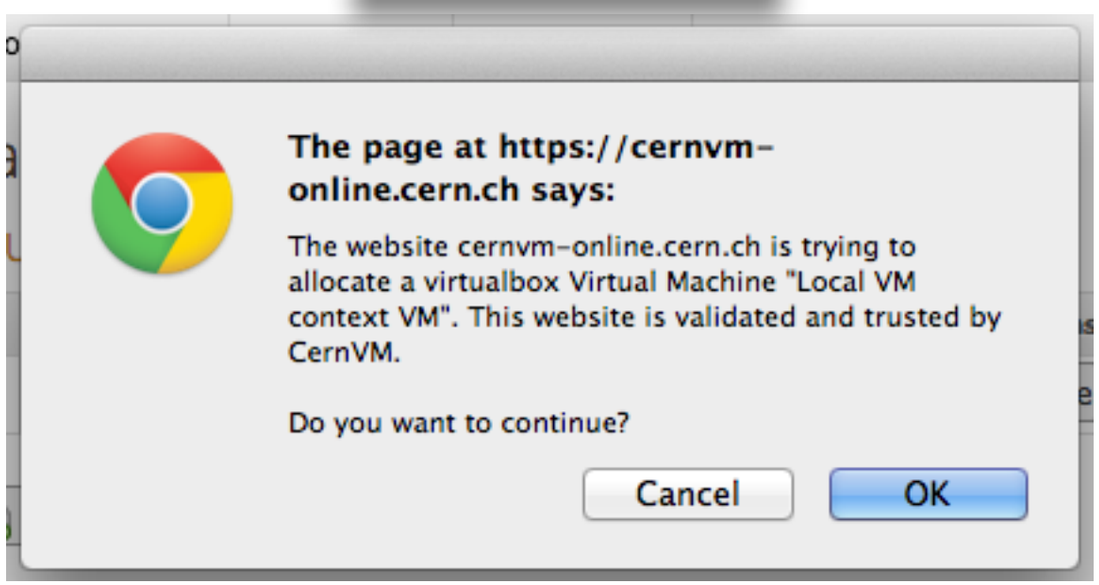

3) Enjoy

WebAPI plugin is available for Chrome, Firefox and IE.

Pairing vs WebAPI

WebAPI **spawns** a VM using user's local PC resources

Pairing is **contextualizing** an existing VM

- requires *console access*
- VM can be local, or remote

Both can be used to avoid having to contextualize manually user's VM with the web appliance.

CernVM Online / Cloud Gateway: interface for contextualization and deployment

CernVM Cloud beta

deploy CernVM virtual clusters

cern.ch/cernvm-cloud

Contact us for a beta tester account

CernVM Cloud Gateway

a distributed system that provides a single interface to use **multiple** and **different** clouds:

- **by cloud type:** OpenStack, CloudStack, OpenNebula, ...
- **private or public:** CERN OpenStack, AWS, ...
- **geographically distributed**

CernVM Virtual clusters

Uses cloud available cloud resources to deploy virtual clusters.

Virtual cluster is a set of VMs able to communicate with each other:

- consists of services
- each service should be deployed in the cloud
- a service defines:
  1. the *context* of the VMs that will implement it,
  2. the VM configuration (*flavor, CernVM version*)
- **fixed services** are deployed once and before any other service
- **scalable services** rely on the fixed services and they can be *scaled up and down*

Virtual cluster example

Master will start first

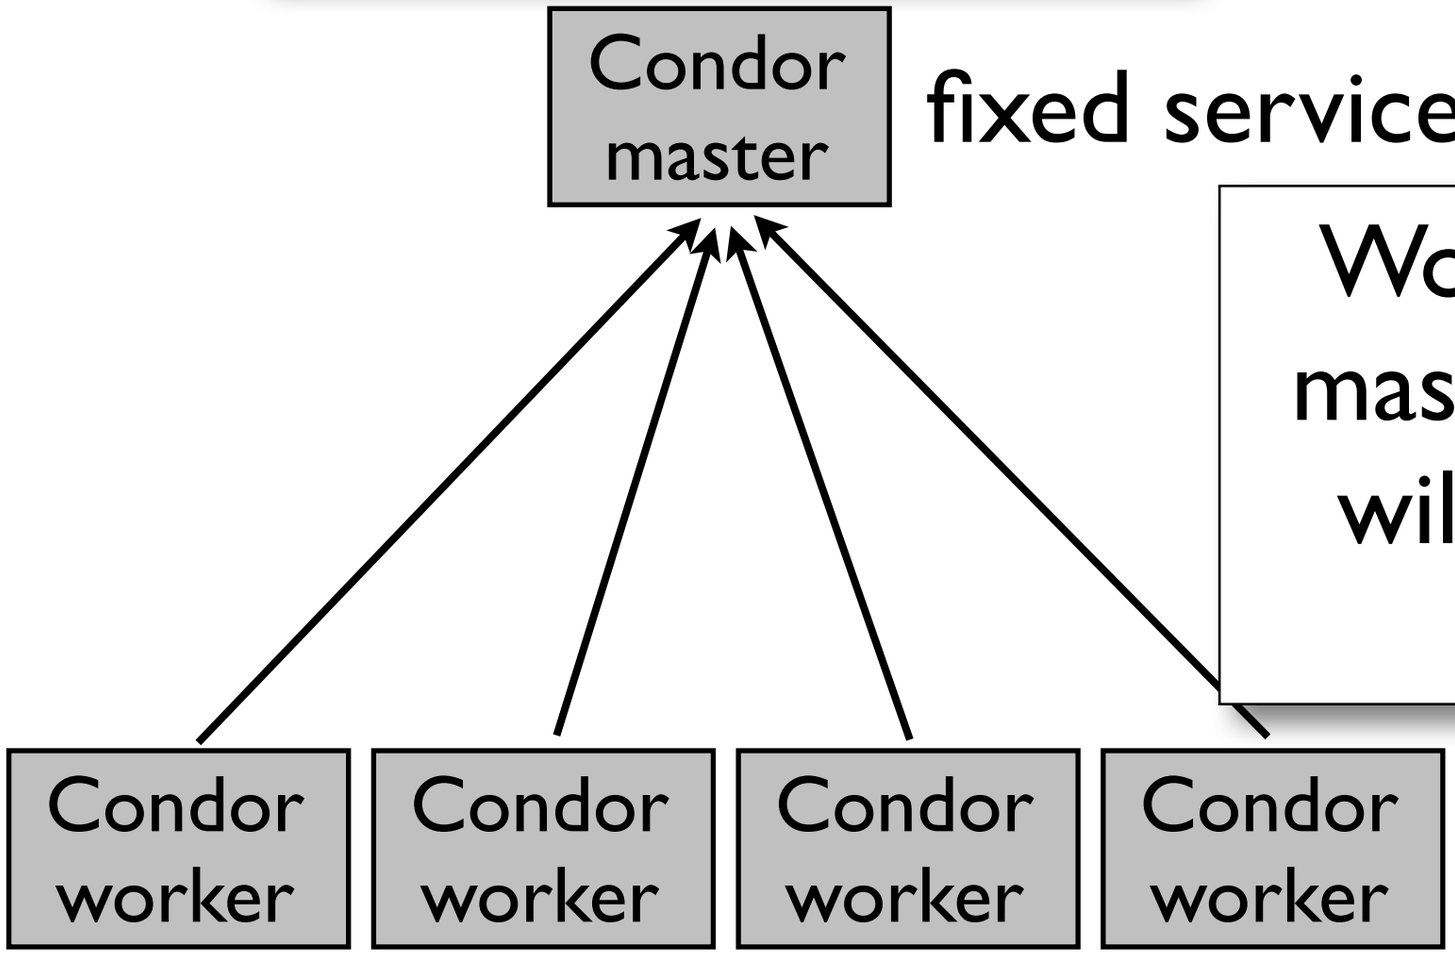

fixed service

Workers will start after master and their context will contain master's IP address

scalable service

There are cases with more than one fixed services: proxy server, VO box, etc

Implementation tools

iAgent Framework: implemented in Perl, it provides ways to develop agents that can expose their functionality through XMPP.

XMPP: communication protocol. Has useful features as:

- **Presence:** to know which agents are online
- **Queuing:** messages will be delivered once agent is up
- **PubSub channels:** many cloud agents listen for messages from the gateway agents

Architecture

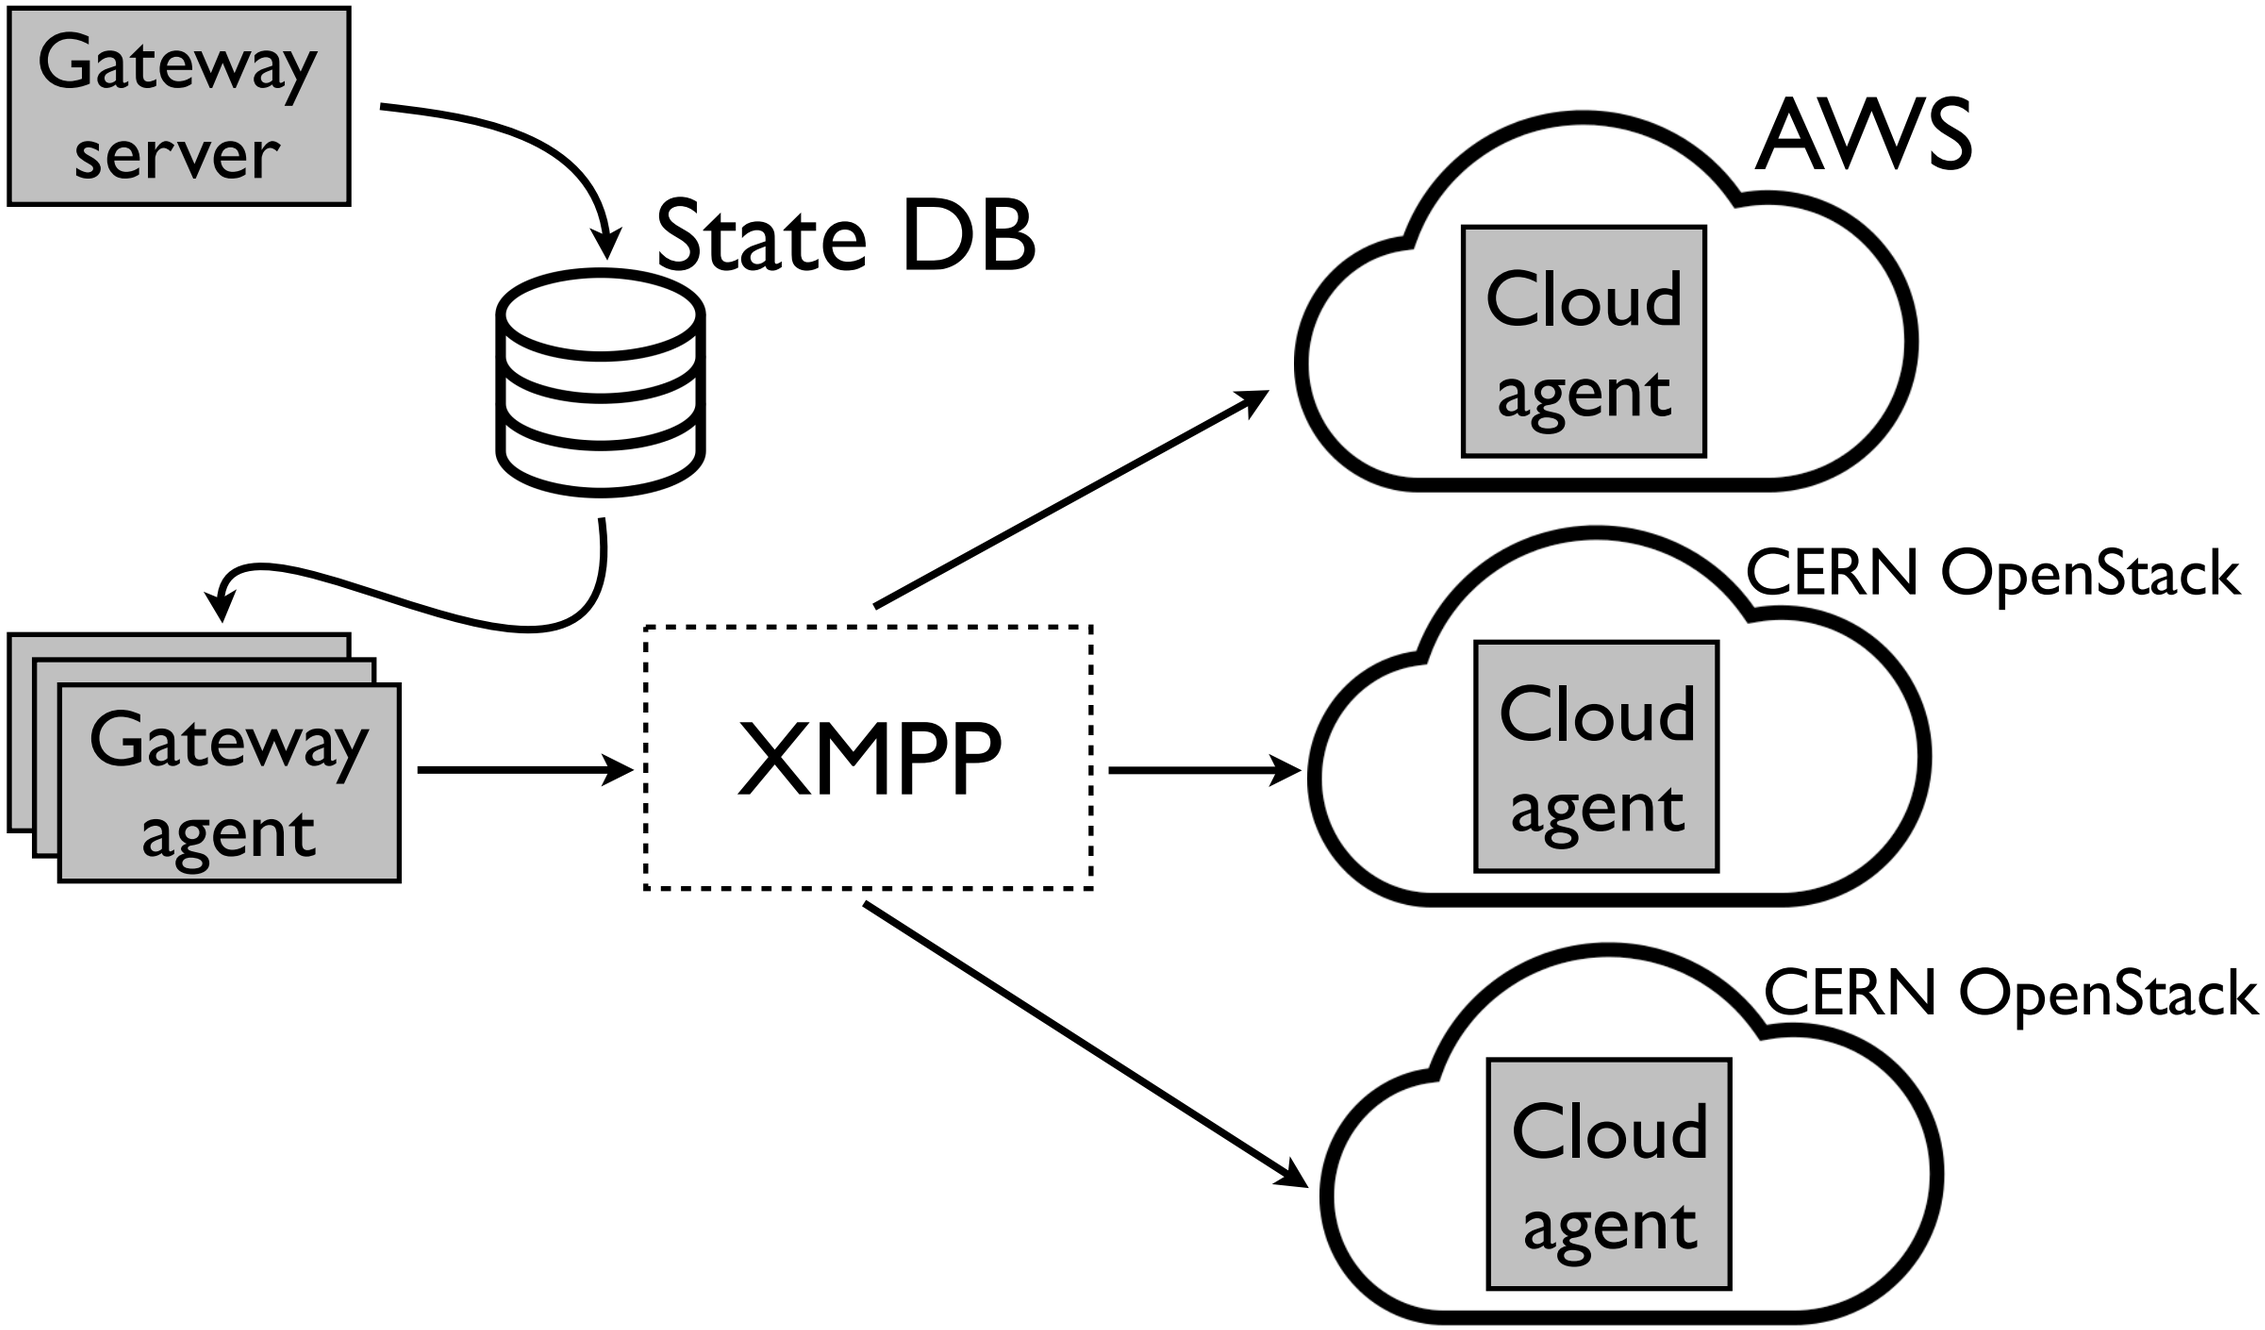

Gateway server / agent

Gateway server: API endpoint and web interface

- Handles client authentication
    - users and groups
- Feeds the system with *user requests*
- **Web interface** works with mobile devices (*responsive*)
- **REST API** provides complete functionality for deploying clusters (http://cern.ch/cernvm-cloud/Wiki/Documentation.html)

Gateway agent: processes the user requests and forwards them to *cloud agents* though XMPP.

Cloud Agent

Associated with a single cloud access credentials

- **Listens** to XMPP PubSub channel for user requests
  - **simple ACL:** administrator can define which users/ groups for cloud gateway can use his cloud instance
  - **soft quota** for managed cloud resources
  - **mapping** of general flavors / templates to cloud specific
- Communicates with *respective cloud* via a **cloud driver** that "speaks" its API

Cloud credentials storage

Current model: cloud credentials are stored in the cloud agents

✓cloud agent can run in a machine managed by the owner of the credentials

✓sensitive credentials are not communicated to the central server

- a cloud agent has to be deployed for each cloud key-pair

CernVM Cloud in action

Defining a cluster

Service definition

Context selection

Cluster services

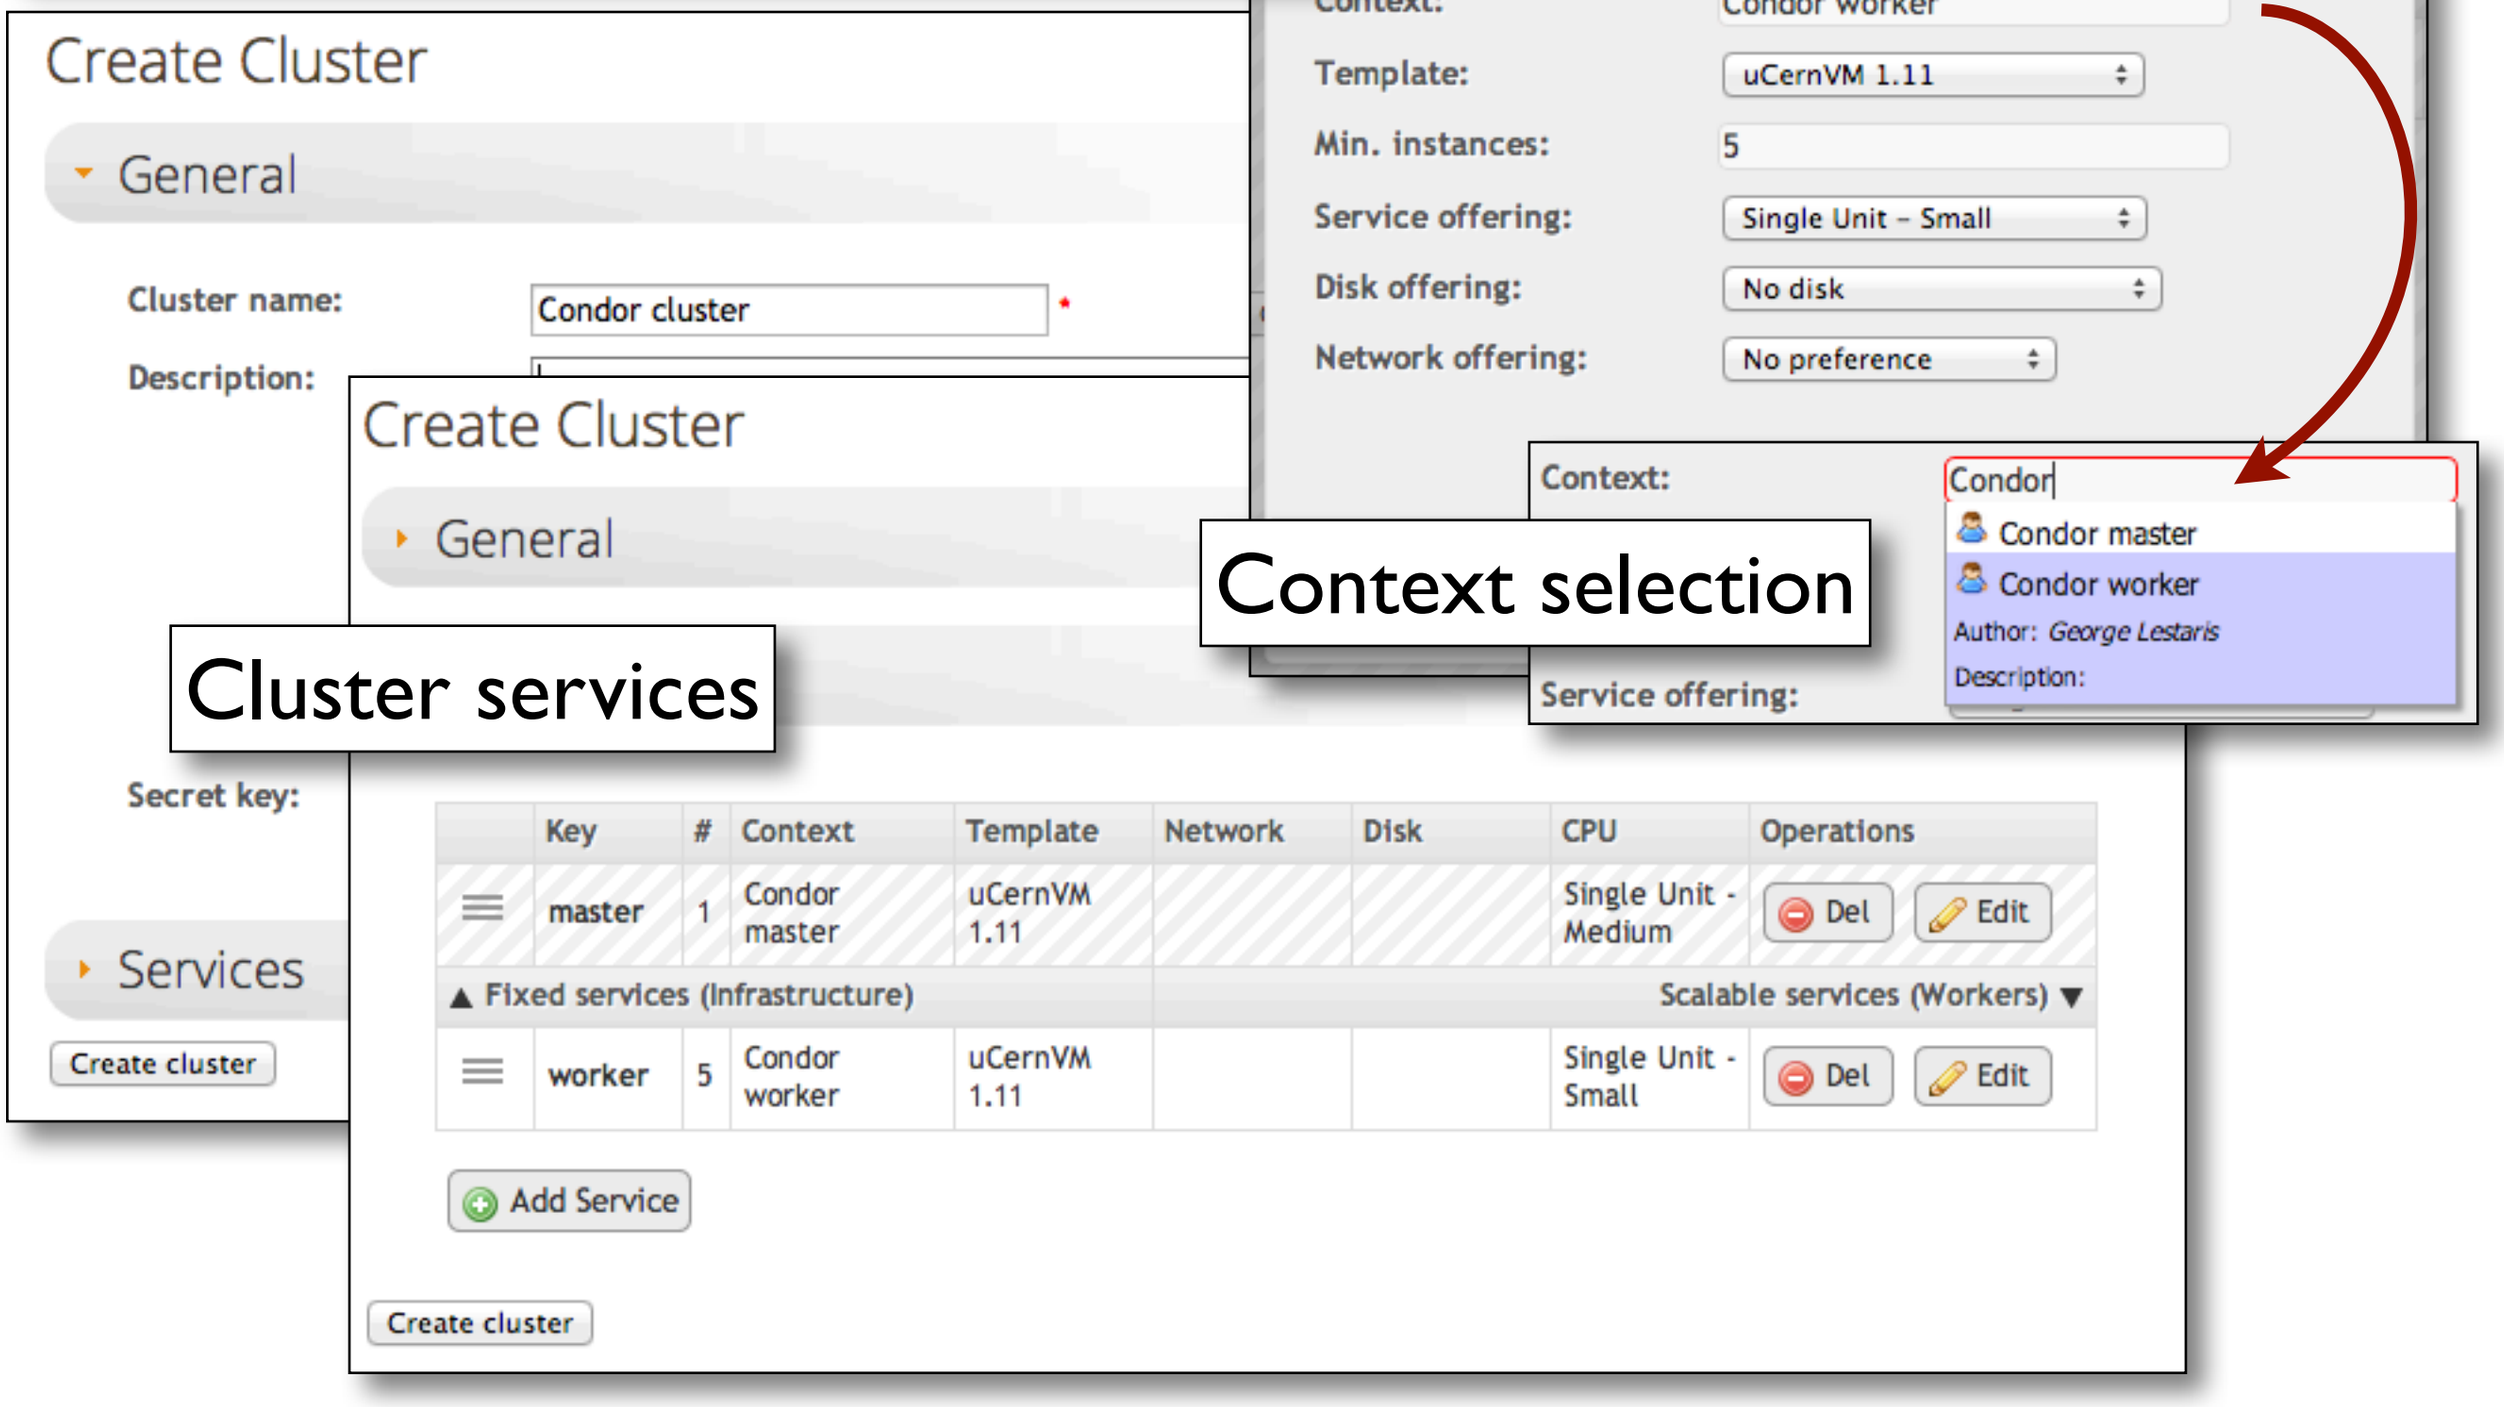

Cluster deployment

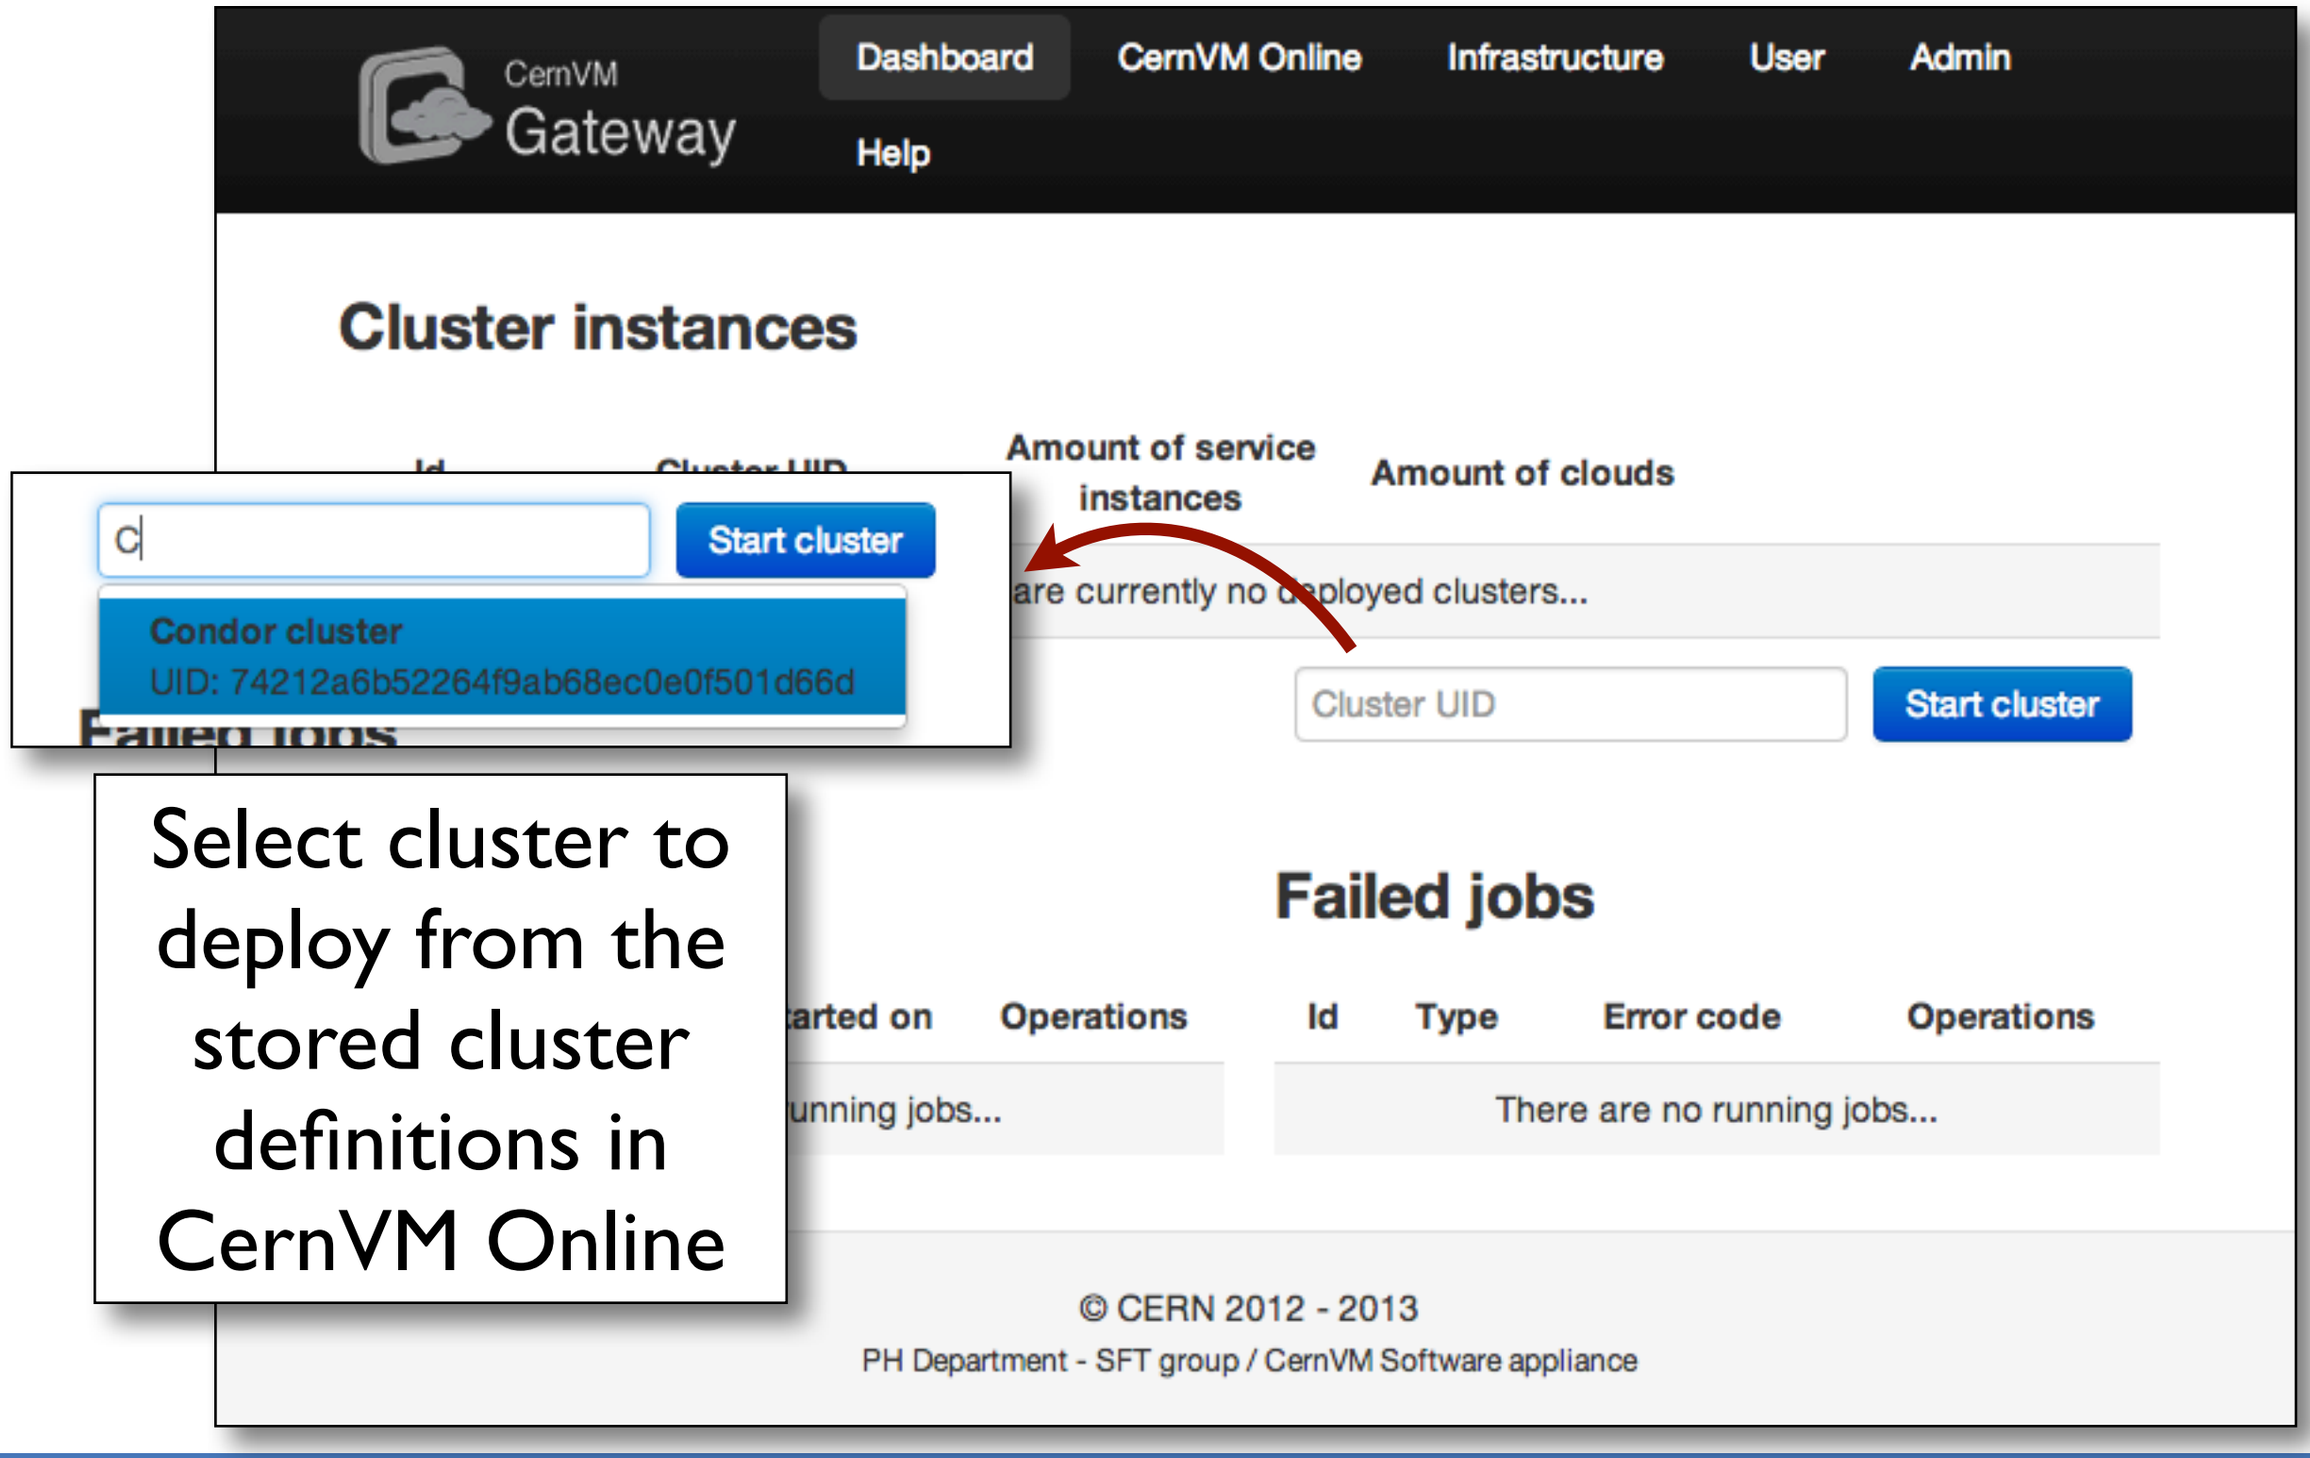

Select cluster to deploy from the stored cluster definitions in CernVM Online

Deployment management

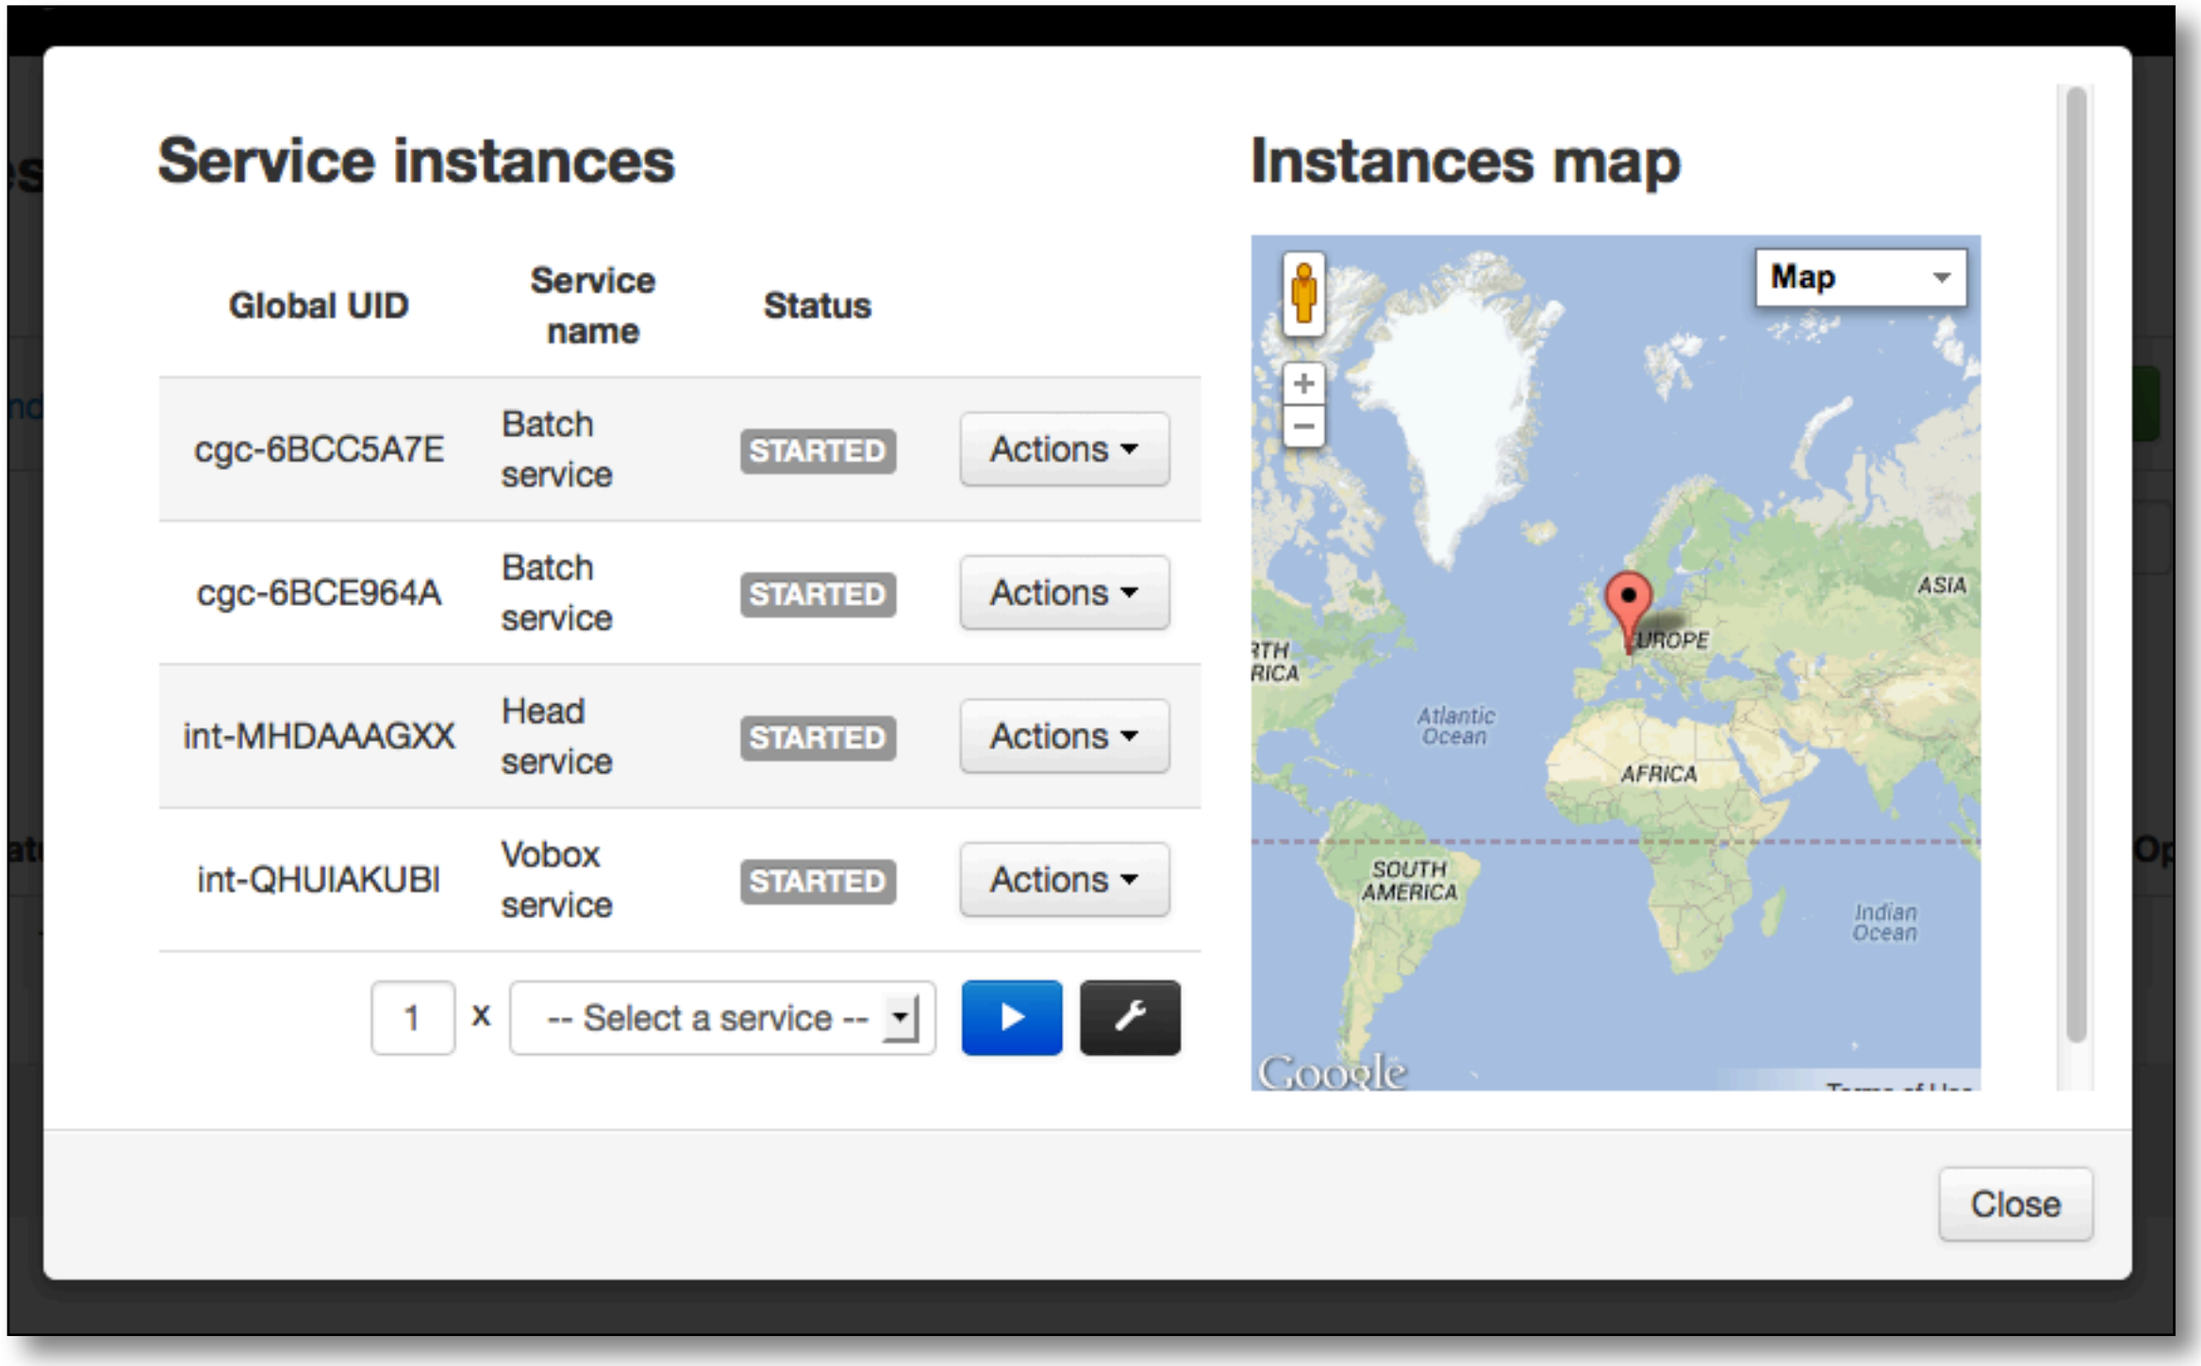

Deployment management

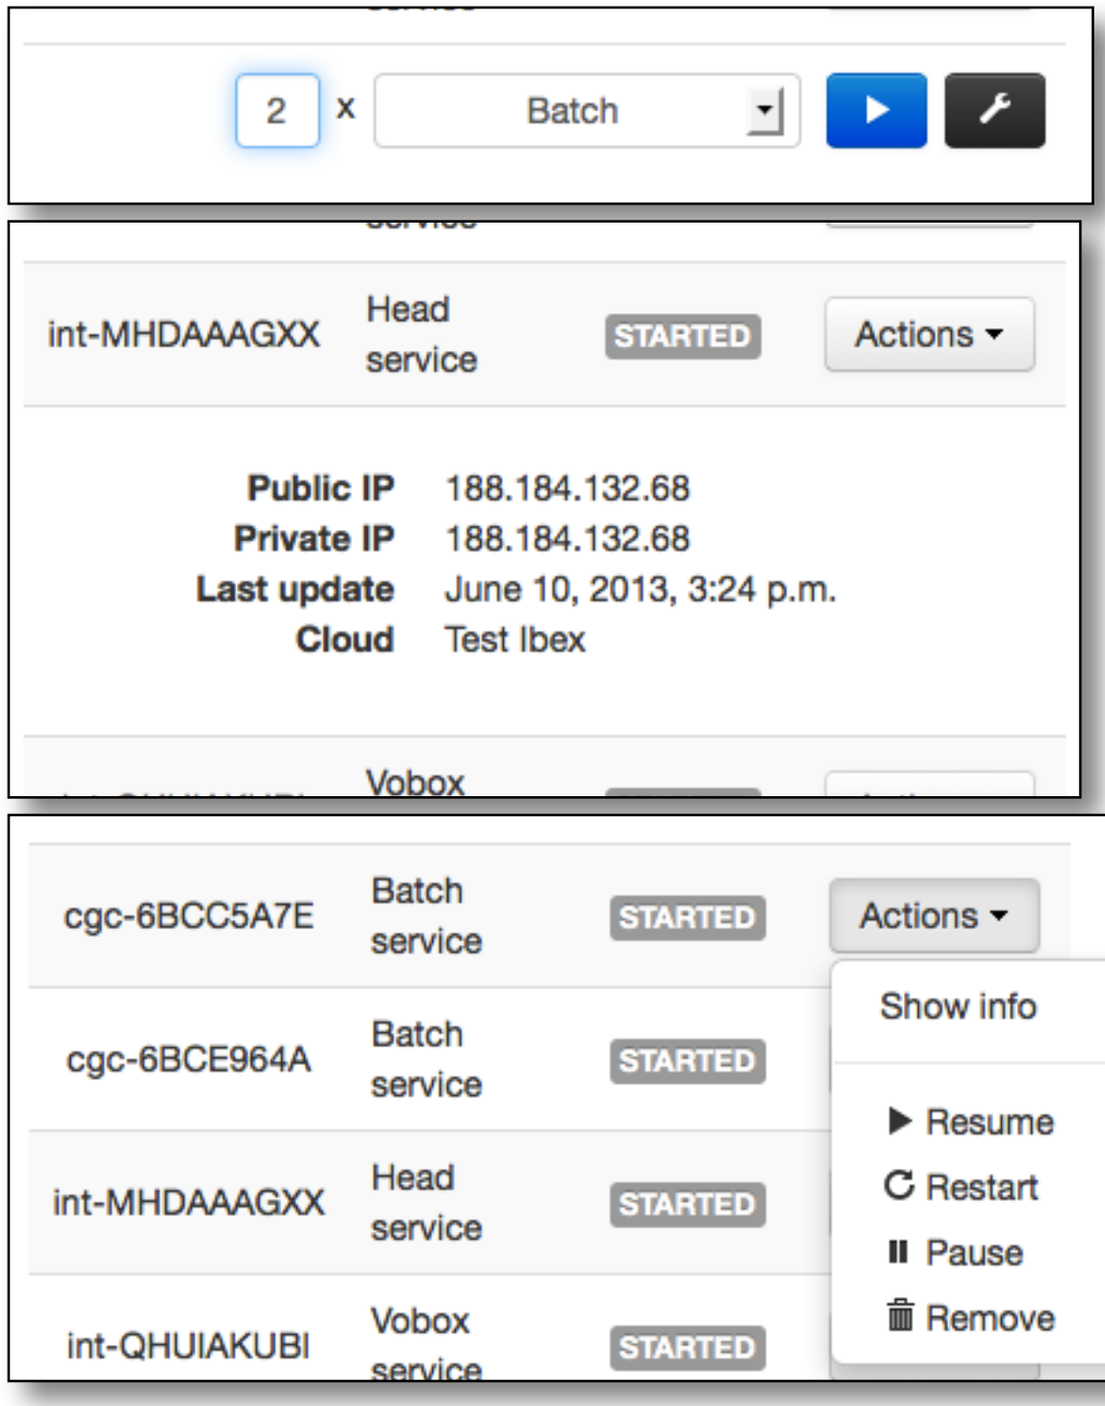

Scaling up services

Per-instance information

VM management

Clusters overflowing clouds

Experimental

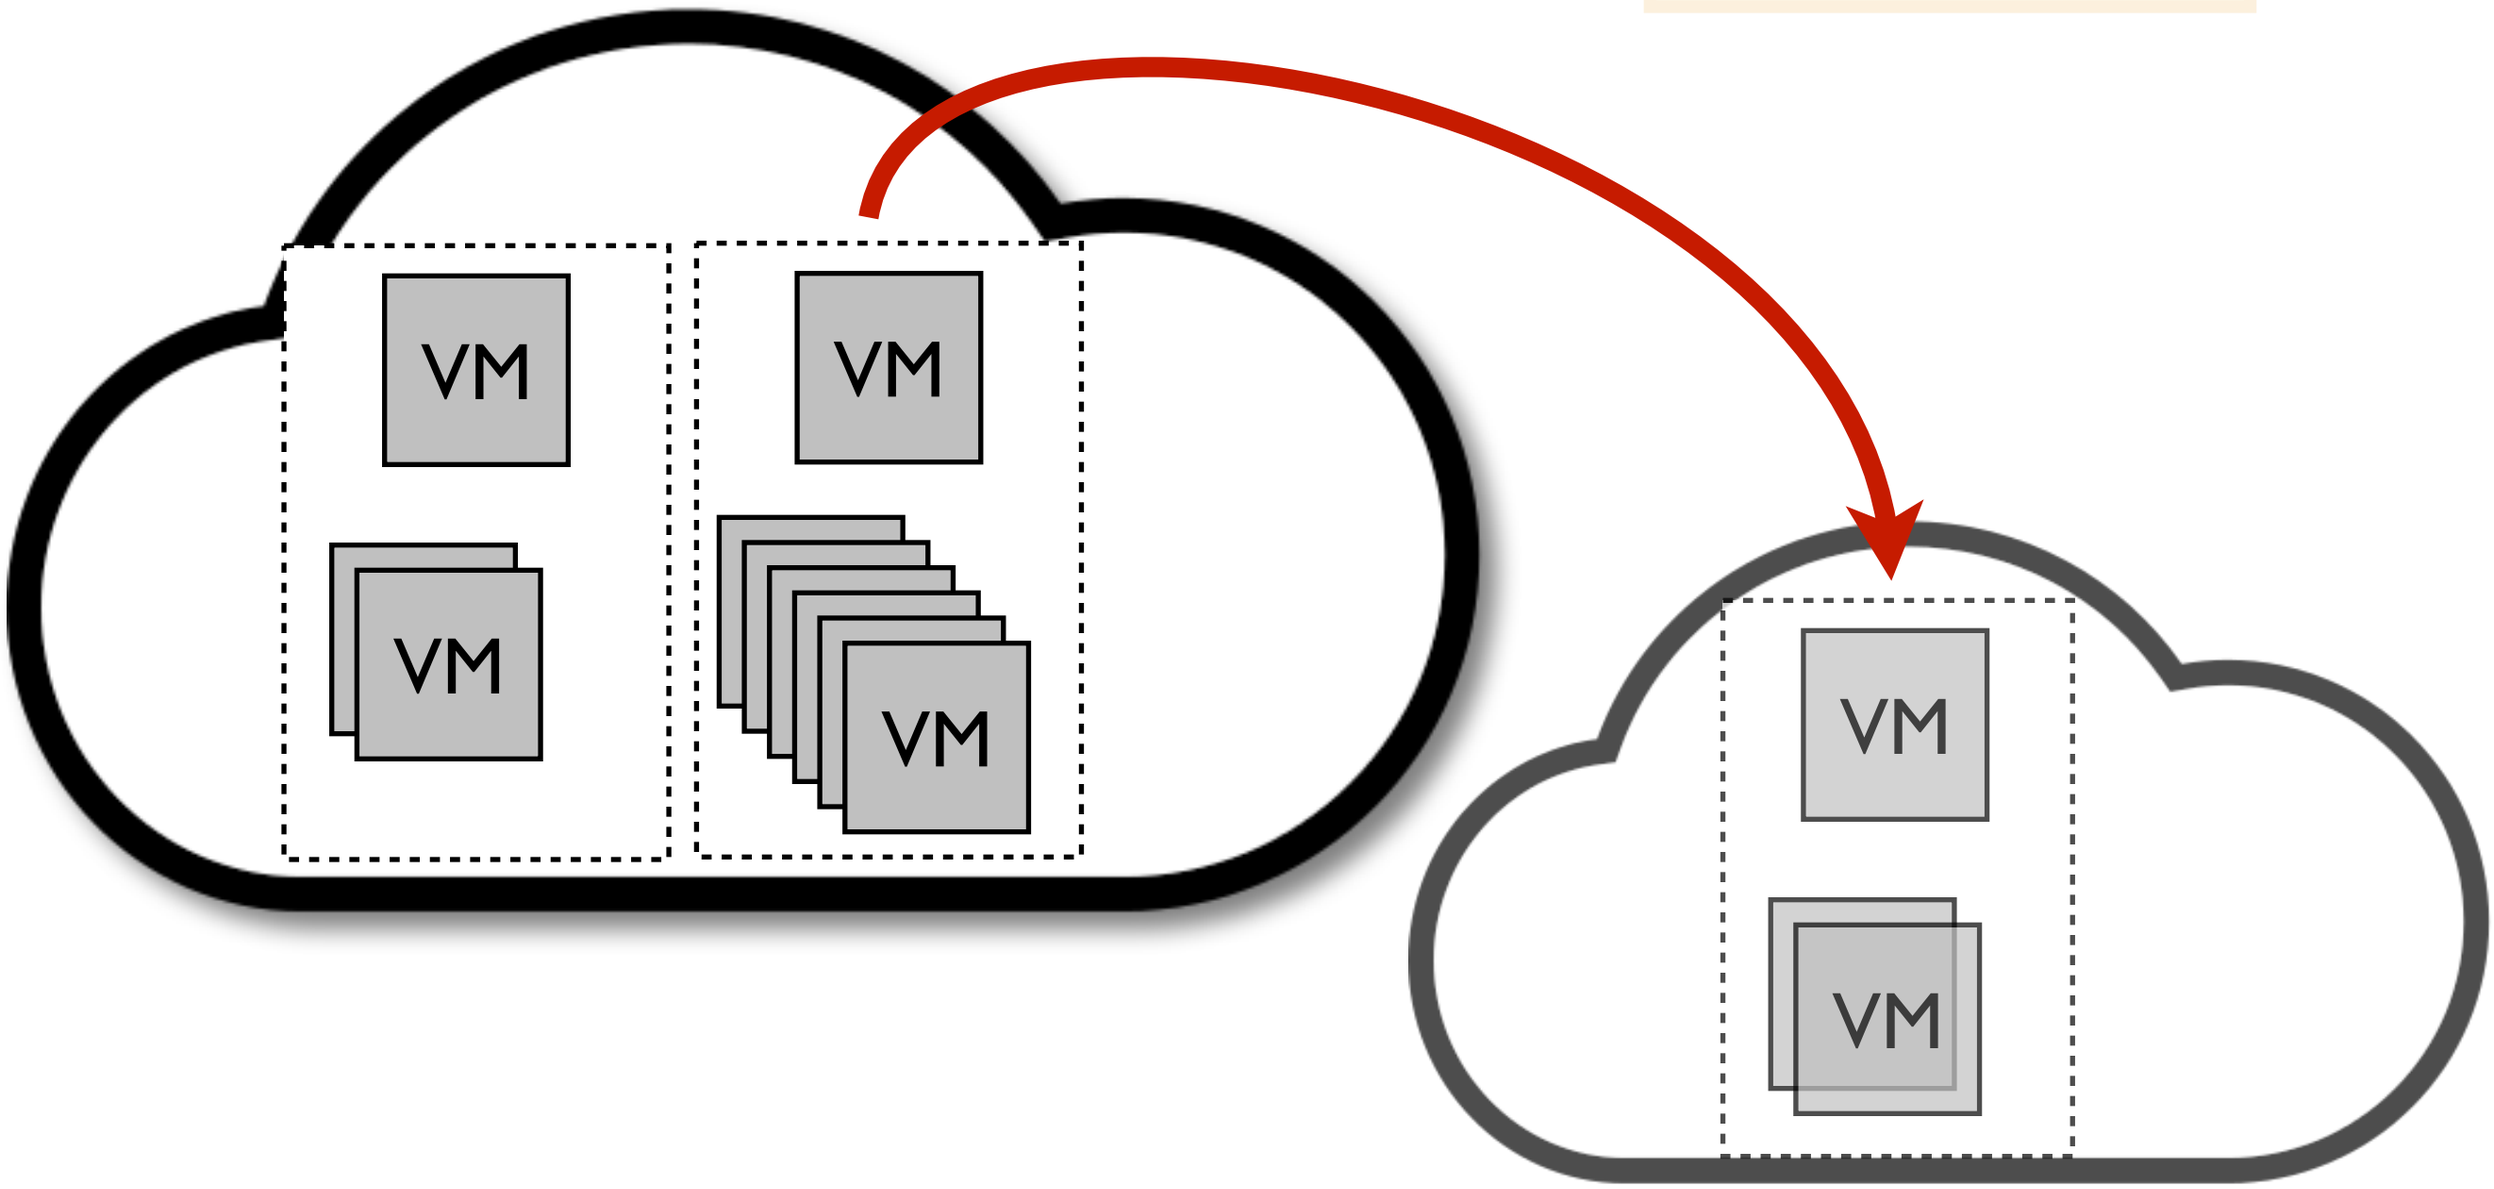

Clusters overflowing clouds

Experimental

- CernVM Cloud can support multiple clouds
- Once a cluster reaches cloud's capacity and user requests to scale it up, it **expands to another cloud**
- In this expansion the fixed services are being *replicated* to the *new* cloud

✓ Cluster VMs will be able to communicate with *fixed services* as they will always be in the local network of the same cloud

Conclusion

CernVM Online: production

- define, store and share *contexts* with a nice web UI
- *Pairing* and *WebAPI* to deploy single VM

CernVM Cloud: beta

- deploy *clusters* defined in CernVM online
- *multiple* & *different* clouds
- *overflowing clouds* and expanding clusters to different clouds
- RESTful API