ADOPTING CLOUD TECHNOLOGIES

Stefano Bagnasco | INFN Torino

OUTLINE

- Leftover: a more realistic IaaS deployment

- Some general musings
- Toolkits & projects
- A simple virtual farm
- Bringing elasticity to the farm
- Flexibility and user-friendliness

A MORE COMPLEX DEPLOYMENT 1

VMs providing **critical services**:

- Run on a cluster of server-class redundant hypervisors
- Public & private IP
- Shared system disks on resilient storage allowing live migration (Services need to run continuously)

VMs providing **computing workforce**:

- Run on a cluster of compute-optimized, less expensive hypervisors
- Locally cached image repository for fast startup (Workers are often reallocated)
- Access to fast storage for data
- Private network only

A MORE COMPLEX DEPLOYMENT 2

- **Image Repository Datastore** holds the OS images

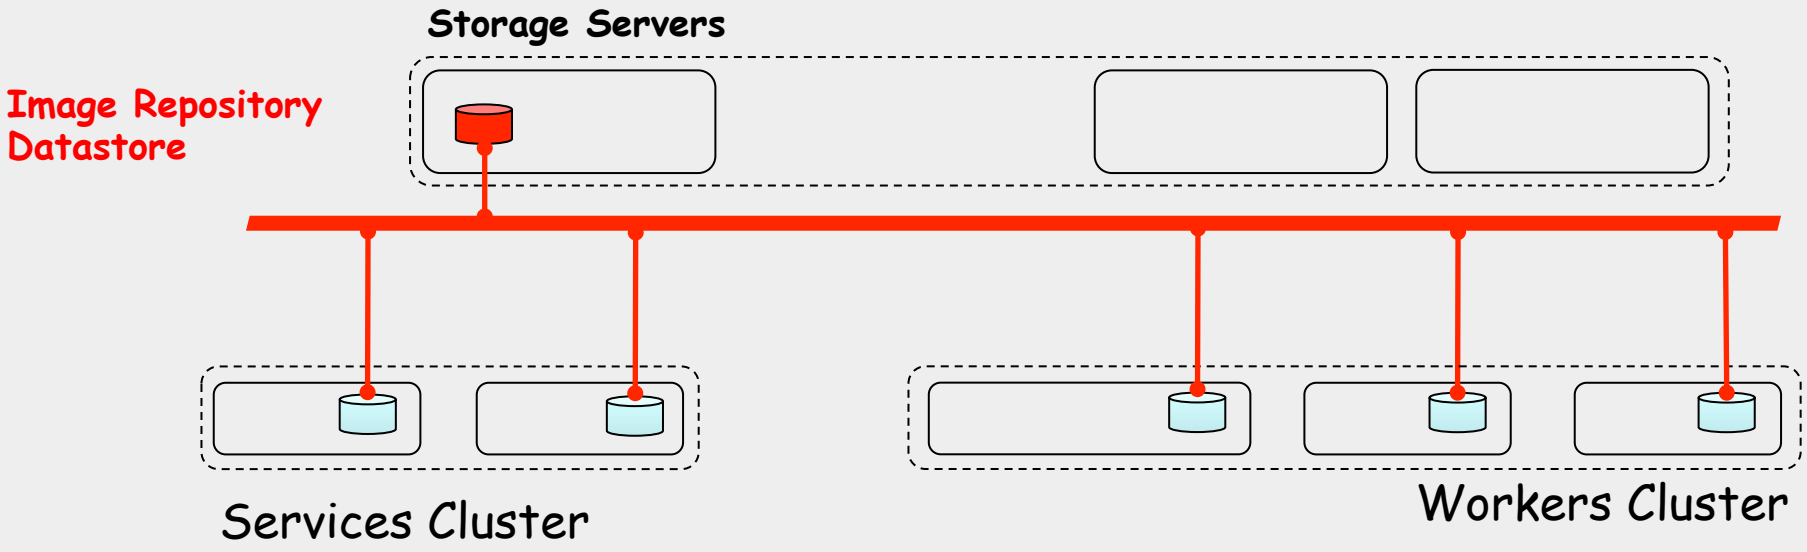

A MORE COMPLEX DEPLOYMENT 2

- **Image Repository Datastore** holds the OS images
- **Services System Datastore** is **shared** to allow live migration

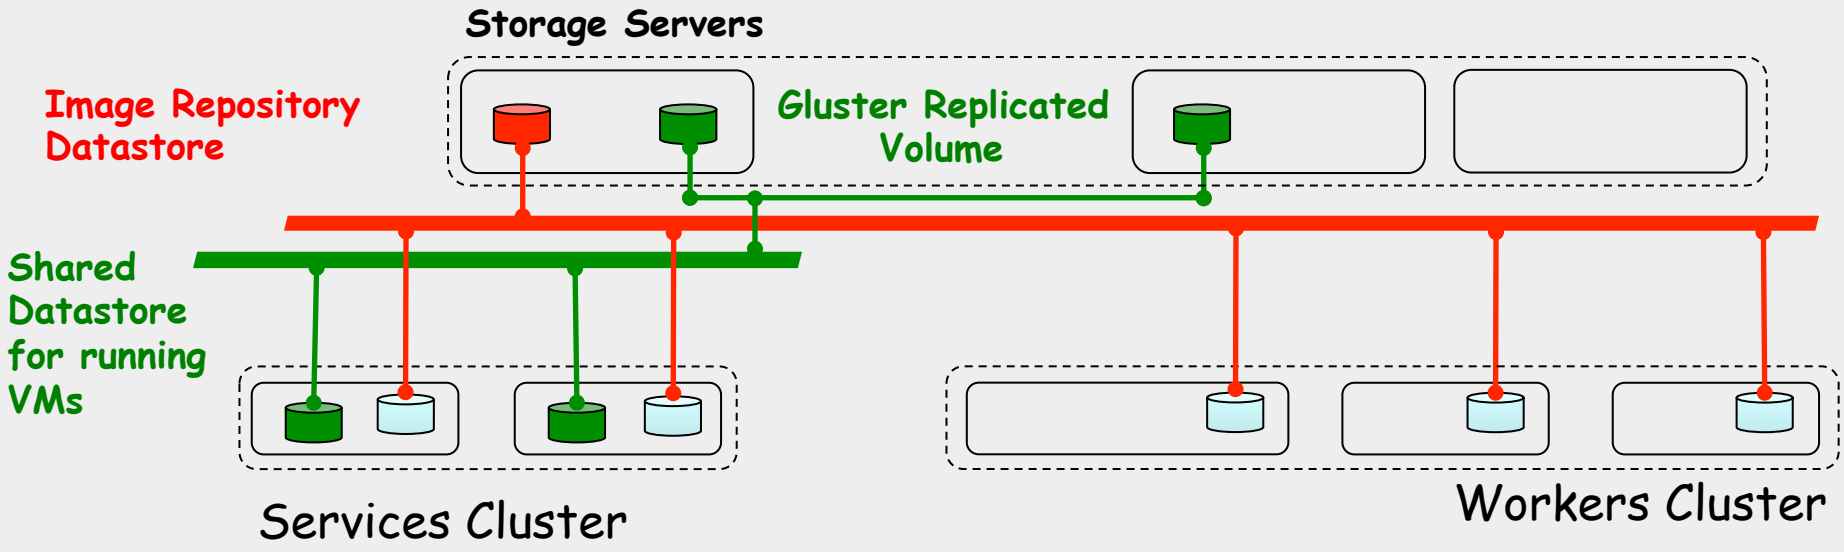

A MORE COMPLEX DEPLOYMENT 2

- **Image Repository Datastore** holds the OS images
- **Services System Datastore** is **shared** to allow live migration
- **Workers System Datastore** is **local** to the hypervisors to increase I/O capacity. Images repository is locally cached on each hypervisor to reduce startup time

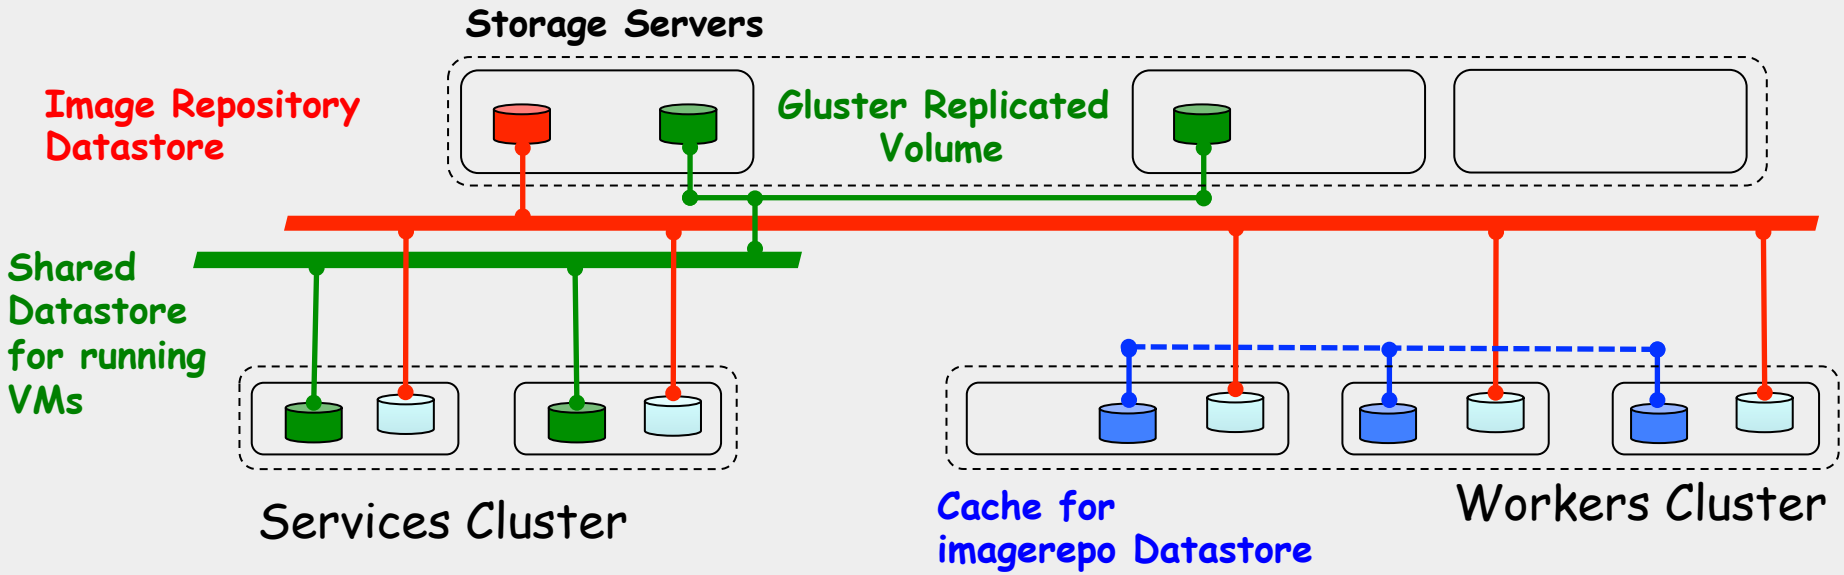

A MORE COMPLEX DEPLOYMENT 2

- **Image Repository Datastore** holds the OS images
- **Services System Datastore** is **shared** to allow live migration
- **Workers System Datastore** is **local** to the hypervisors to increase I/O capacity. Images repository is locally cached on each hypervisor to reduce startup time
- **Persistent Space Datastore** is mounted on the relevant hypervisors using the **iSCSI** Transfer Manager to provide persistent storage to Virtual Farms

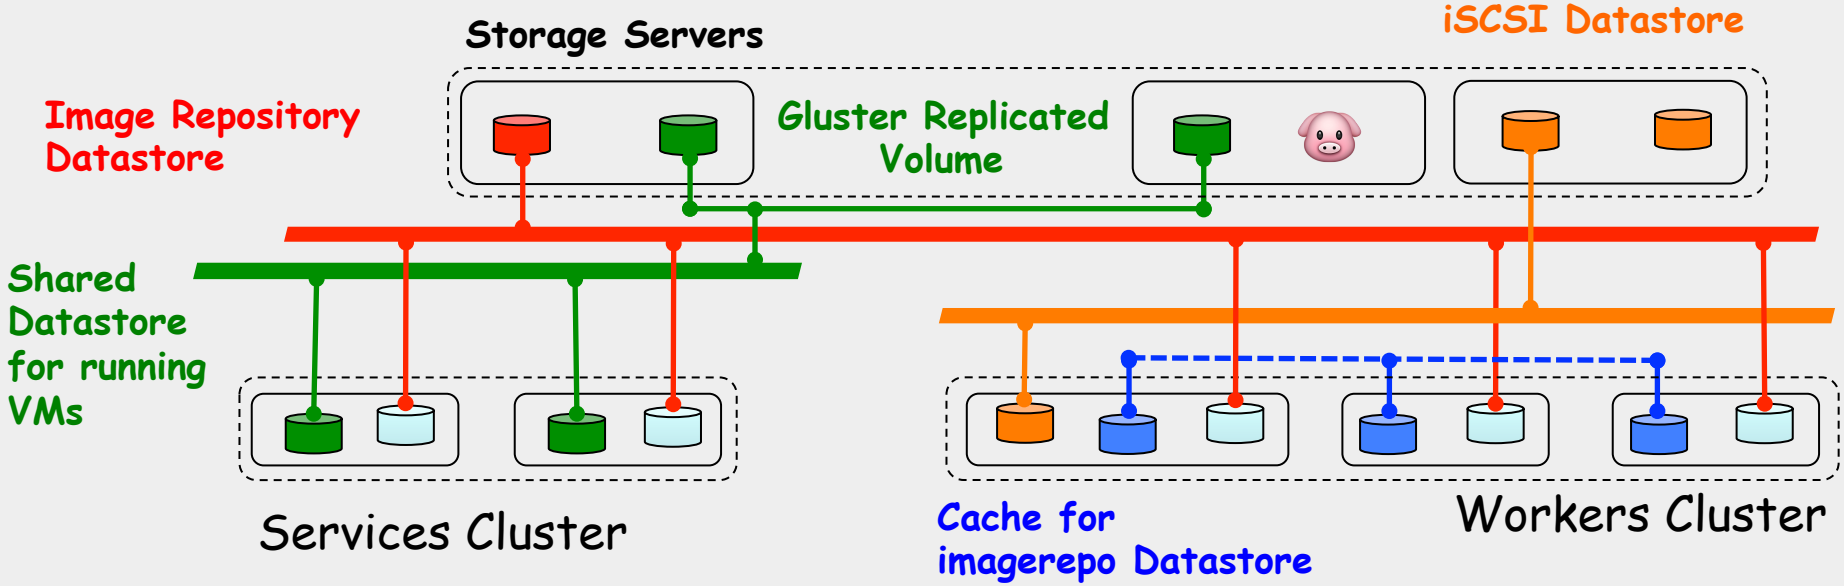

Cloud Computing technologies promise to help the scientific computing community!

What the cloud is:

- A technology to ease resource management, provisioning and sharing
- An industrial standard technology

What the cloud is not:

- A magical “resource multiplier”
- A complete end-to-end scientific computing system

WHY SCIENTIFIC CLOUD COMPUTING?

- Cloud computing is becoming (one of) the industry standard ways to provision computational power
  - We're not any more the only ones working with Big Data, so industry is developing tools
- The technology decouples the infrastructure from the application
  - Scientists will not need any more to worry (much) about infrastructure management
- This is a general trend!
  - 14% of contribution titles at CHEP2015 included the word "cloud"

- Adapting scientific workloads to public clouds
  - Several activities ongoing, specially in large experiments
  - See e.g. CHEP2013 (Amsterdam) and CHEP2015 (Okinawa) conferences for a number of reports
  - https://indico.cern.ch/event/214784/
  - https://indico.cern.ch/event/304944/
- Main issues
  - Data access
  - Prices

- Creating dedicated local and federated cloud infrastructures to cater to the scientific community's computing needs
  - Min bias: using the IaaS model to manage a Computer Centre
    - Small scale example: INFN-Torino
    - Large scale example: CERN Agile Infrastructure
  - Then: adapting experiment's Computing Models to exploit Cloud technologies
    - See e.g. Xiaomei's talk tomorrow
    - Several more in CHEP2013 and CHEP2015 proceedings
  - Next: Building a large scale federated cloud infrastructure
    - Now to complement, eventually to replace the existing Grid infrastructure

A FEW PROJECTS

- A “Grid of Clouds”: the EGI Federated Cloud
- Readily usable toolkit: the CernVM Ecosystem
- An EU-funded development project: INDIGO-DataCloud

THE EGI FEDERATED CLOUD

- A seamless grid of academic private clouds and virtualised resources, built around open standards and focusing on the requirements of the scientific community.
  - https://www.egi.eu/infrastructure/cloud/
- Features both a federated infastructure and an Application Marketplace
  - https://appdb.egi.eu
  - Holds, manages and populates “Virtual Appliances”

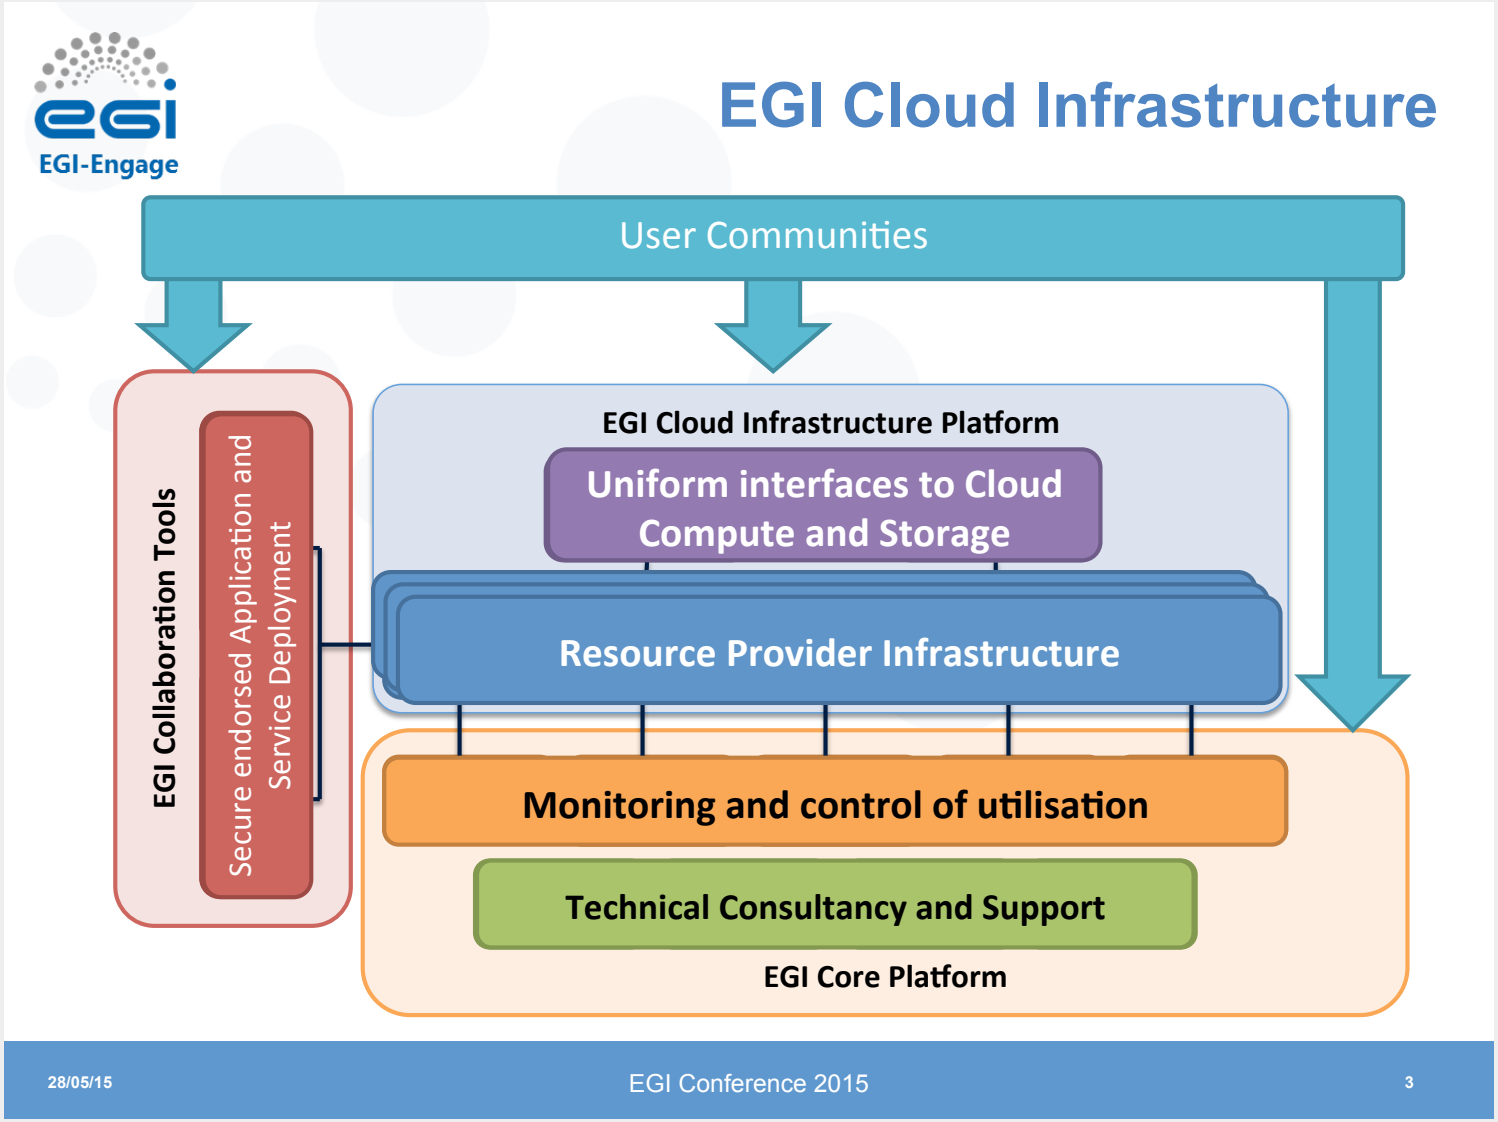
EGI-Engage
EGI Cloud Infrastructure
User Communities
EGI Collaboration Tools
Secure endorsed Application and Service Deployment
EGI Cloud Infrastructure Platform
Uniform interfaces to Cloud Compute and Storage
Resource Provider Infrastructure
Monitoring and control of utilisation
Technical Consultancy and Support
EGI Core Platform
28/05/15
EGI Conference 2015
3

EGI-Engage

VM Management

On demand compute to run any kind of workloads on virtual machines

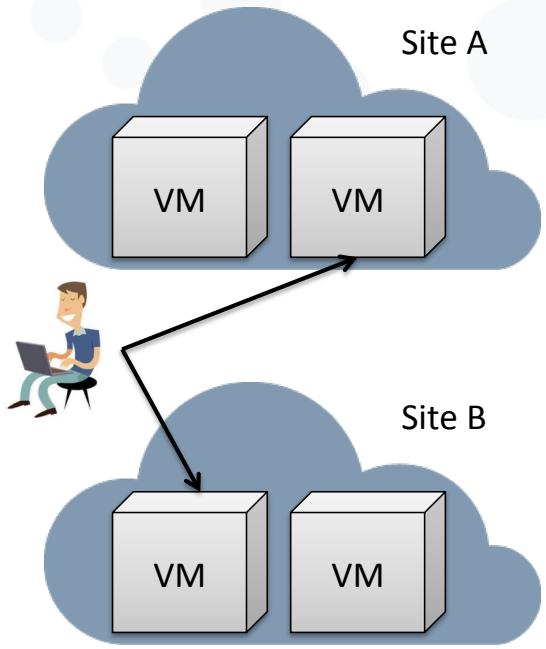

- **OCCI** as standard interface
  - FedCloud **contextualization** extension
  - FedCloud **profile** extension
- Clients
  - ruby and java clients
  - OCCI connectors in brokers
- Servers
  - rOCCI-server
  - OCCI-OS
  - snf-occi

EGI-Engage

VM Image Management

Automatic and secure distribution of endorsed VM images for Virtual Organisations

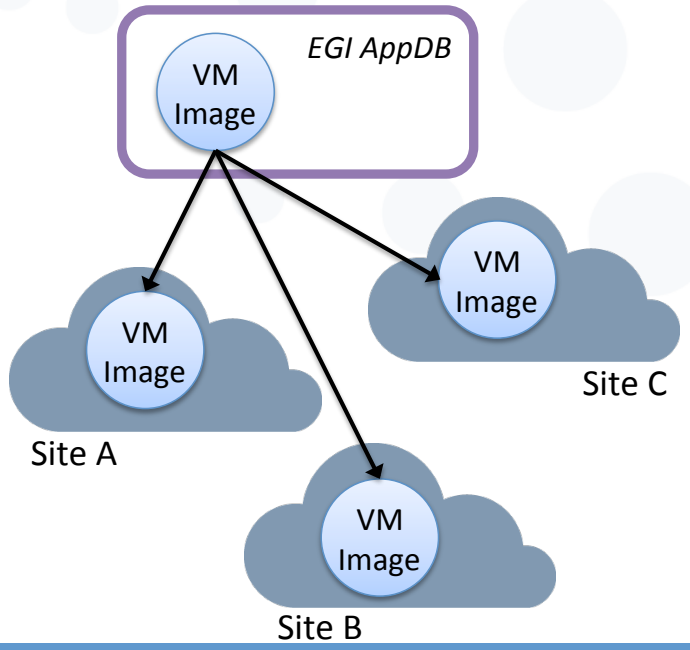

- **Web based** front-end in AppDB
- **OVF** for packaging images
- **HEPiX lists** for distribution
  - Endorsed by VO managers
  - Signed metadata
  - RP subscribe and download
- EGI endorses basic OS images

EGI-Engage

High Level Tools (PaaS & SaaS)

- Extend the IaaS capabilities of the EGI cloud
  - 'Alternatives' of the OCCI client and API
  - More than OCCI
- External contributions (→ support many other clouds too)
- New developments expected (e.g. INDIGO-Datacloud)

Catania Science Gateway	• SaaS • Identity Federation	VMDIRAC	• Abstraction on top of various HPC/HTC/cloud
Slipstream	• PaaS for automating deployments • Helix Nebula	WS-PGRADE	• Workflow development and enactment
COMPSs	• Programming framework for auto-parallelisation	Vcycle	• VM lifecycle manager

VIRTUAL FARM PROVISIONING MODEL

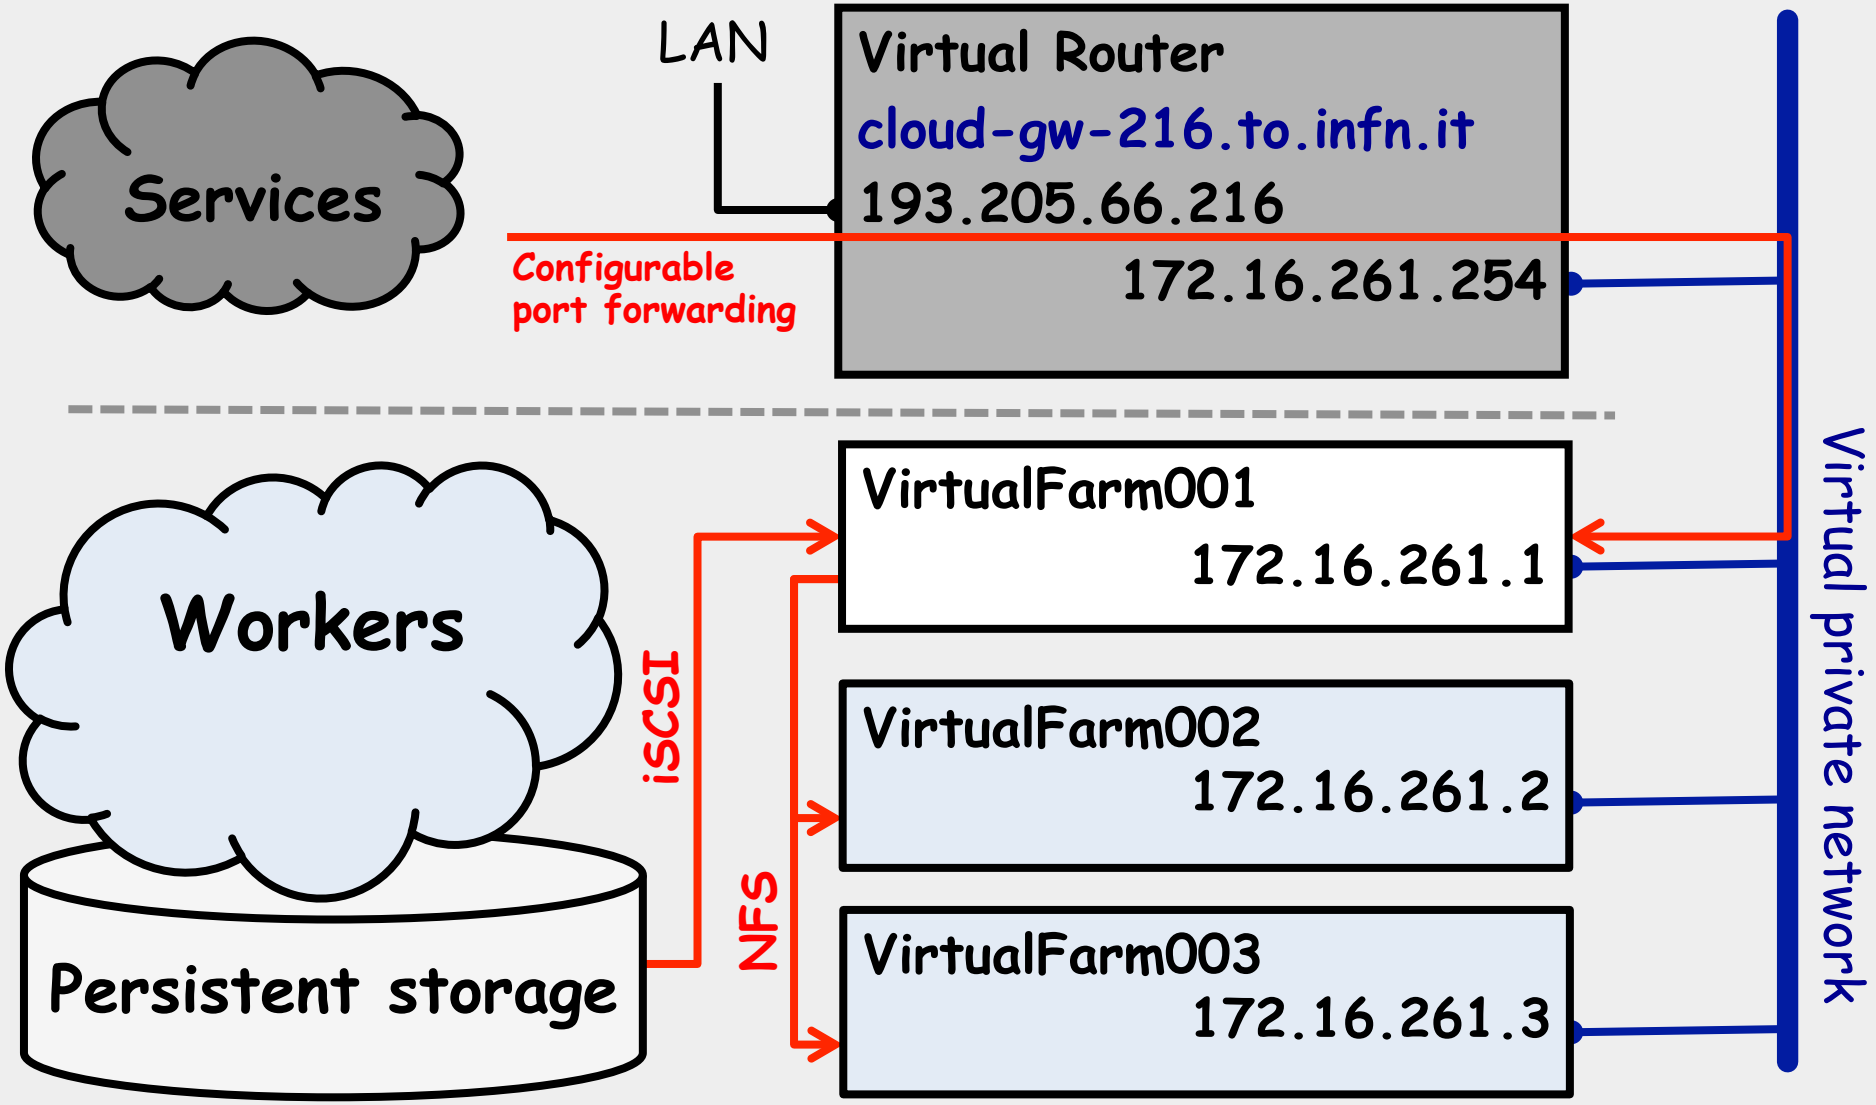

VIRTUAL FARM PROVISIONING MODEL

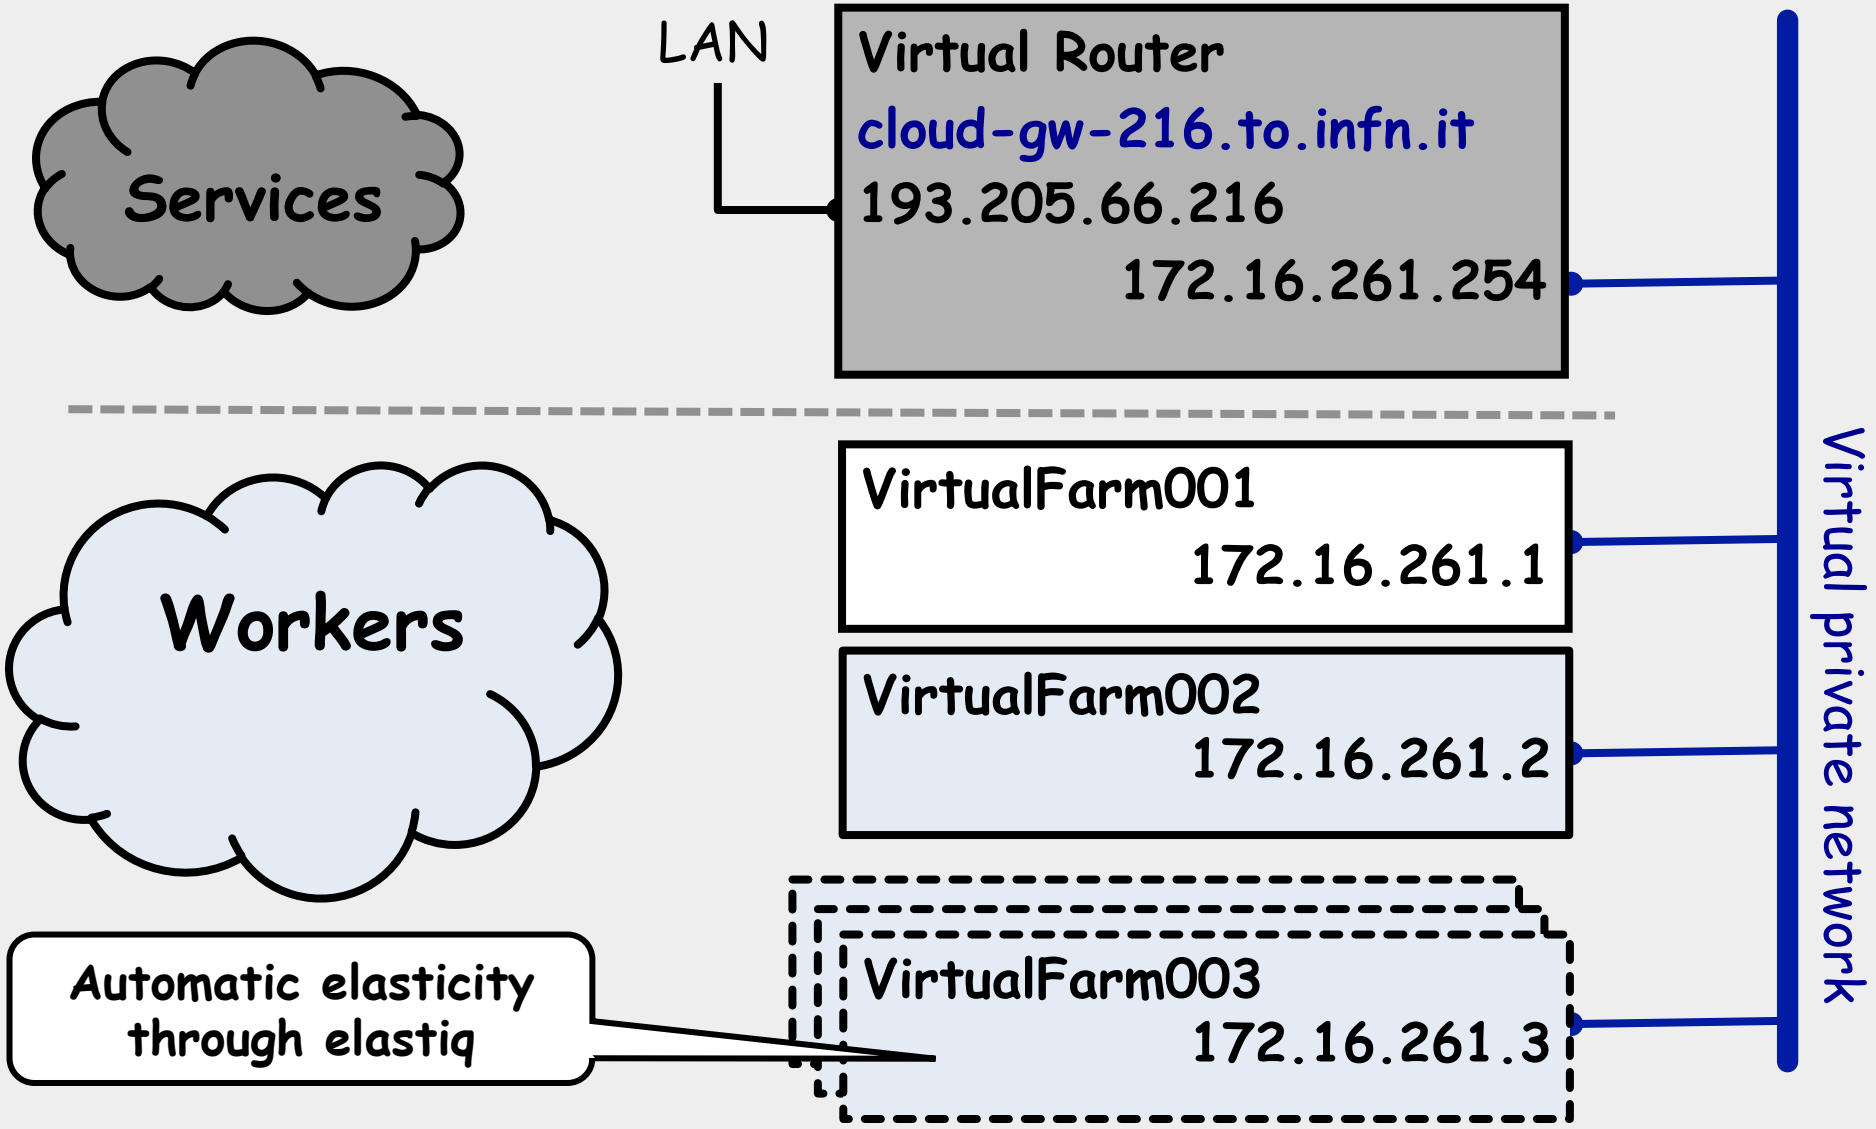

The Virtual Analysis Facility

PROOF+PoD | CernVM | HTCondor | elastiq

What is the VAF?

- A cluster of CernVM virtual machines: one head node, many workers
- Running the HTCondor job scheduler
- Capable of growing and shrinking based on the usage with elastiq

- Configured via a web interface: cernvm-online.cern.ch
- Entire cluster launched with a single command
- User interacts only by submitting jobs
- Elastic Cluster as a Service: elasticity is embedded, no external tools
- PoD and dynamic workers: run PROOF on top of it as a special case

How elasticity works: elastiq

elastiq is a Python app monitoring the queue to make it elastic

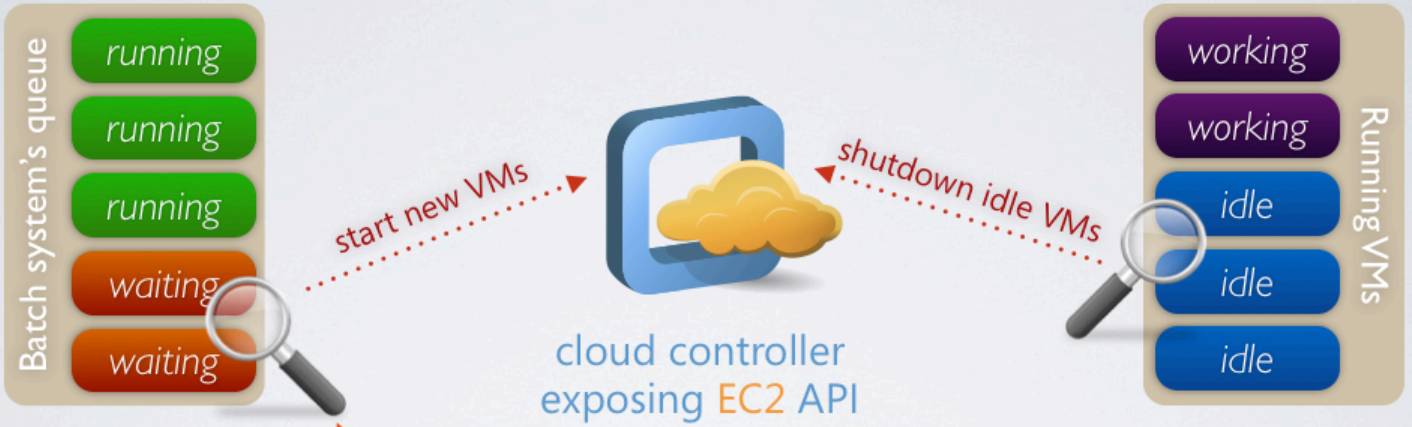

Jobs waiting too long will trigger a **scale up**

Supports **minimum and maximum quota** of VMs

You deploy only the master node: minimum quota immediately launches VMs automatically

Integrated in CernVM 3
source: **github.com/dberzano/elastiq**

Dario Berzano's talk @ CHEP2013

CERNVM ECOSYSTEM

- Virtual Software Appliance for CERN LHC experiments.
  - A complete, portable and easy to configure environment for developing and running LHC data analysis locally and on private and public clouds.
  - http://cernvm.cern.ch/portal/

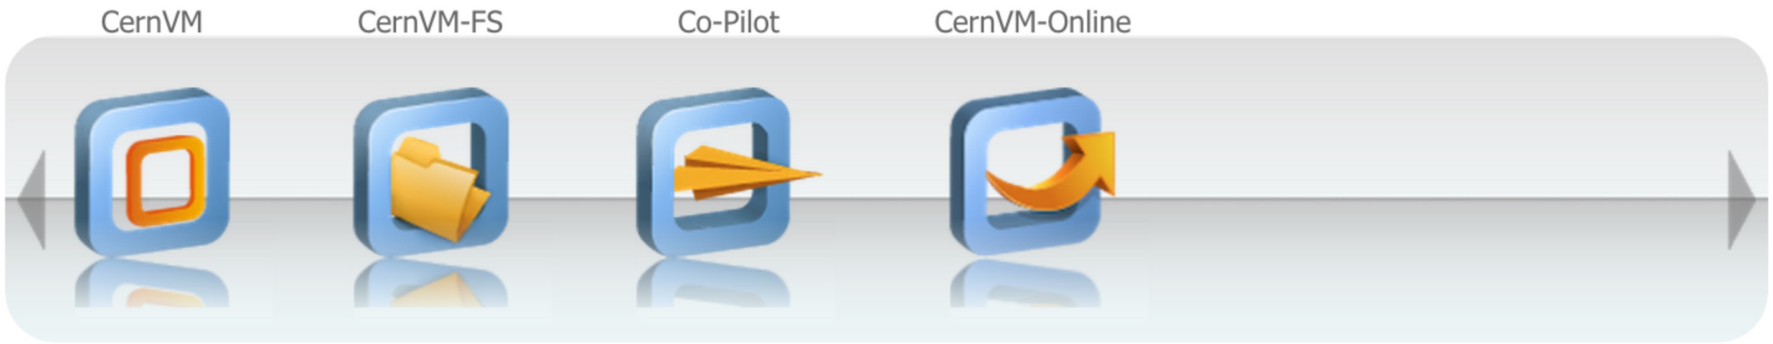

- Virtual Software Appliance for CERN LHC experiments.

CernVM is a baseline Virtual Software Appliance for the participants of CERN LHC experiments.

µCernVM is an even lighter appliance.

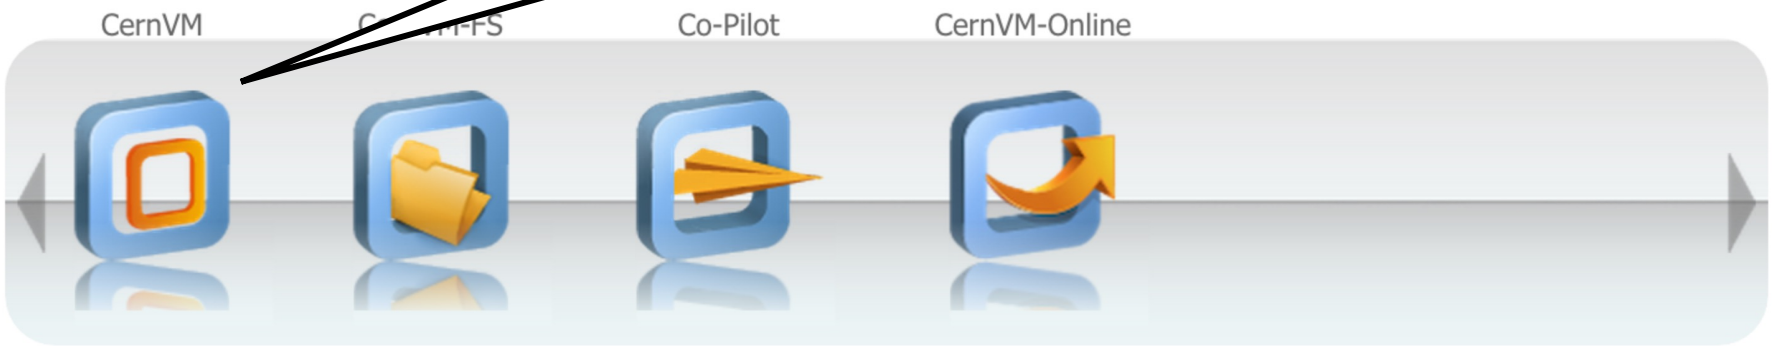

- Virtual Software Appliance for CERN LHC experiments.

CVMFS is a network file system based on HTTP and optimized to deliver experiment software.

(See also yesterday's talk by Tian)

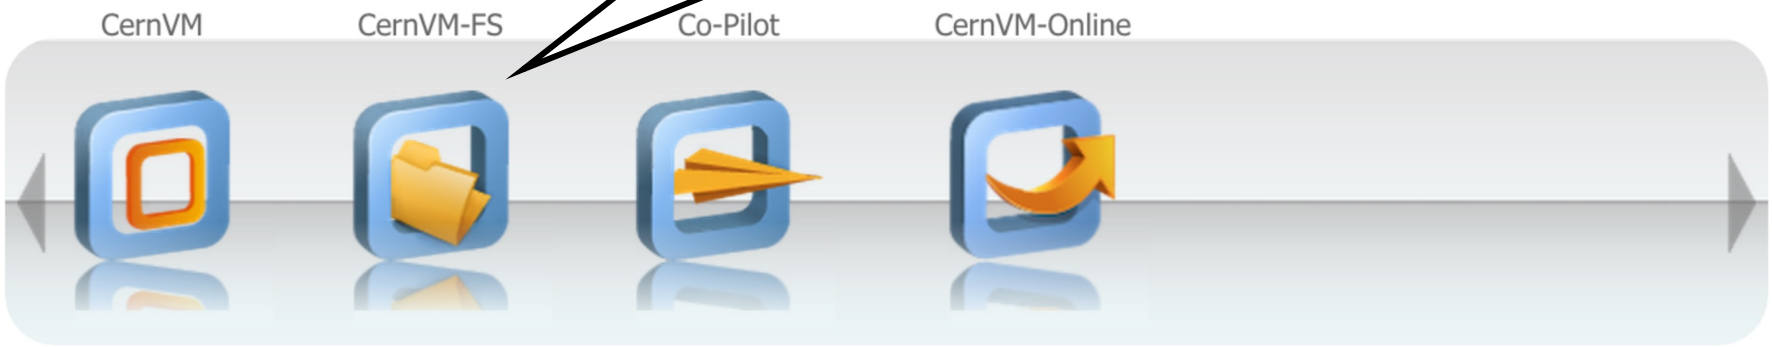

- Virtual Software Appliance for CERN LHC experiments.

Co-pilot is a framework which allows to instantiate a distributed computing infrastructure on top of virtualized computing resources (sometimes called a "Cloud Orchestrator")

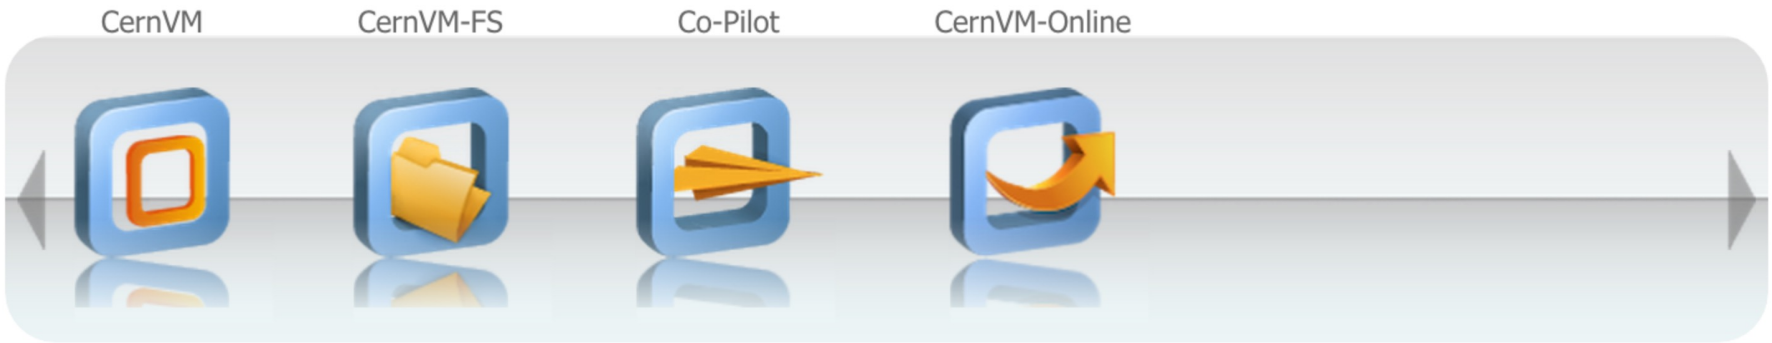

- Virtual Software Appliance for CERN LHC experiments.

CernVM-Online is an easy contextualization mechanism that allows CernVM instances to be contextualized via a simple web interface.

https://cernvm-online.cern.ch/user/login

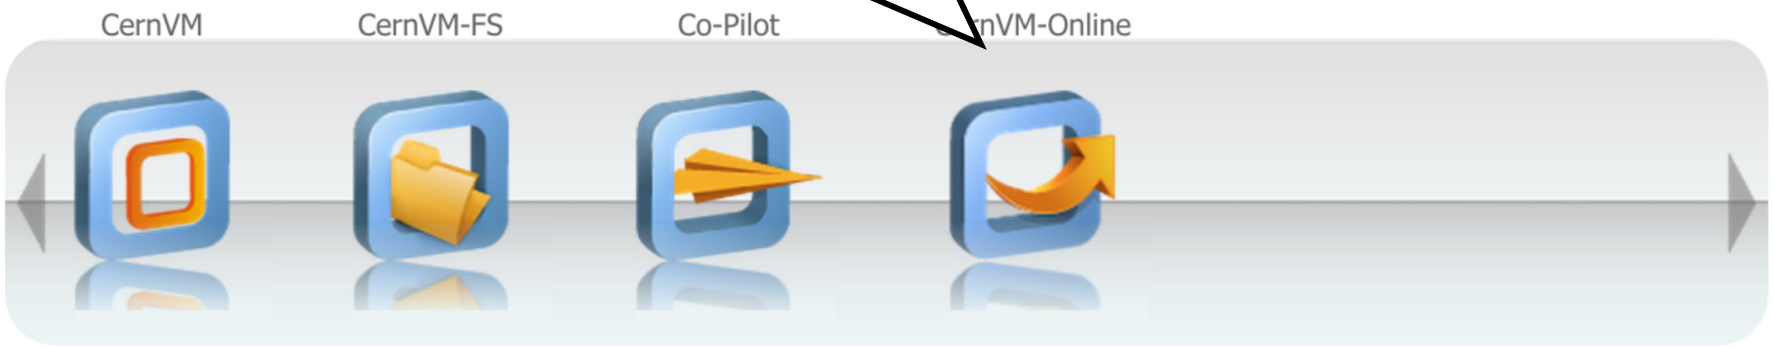

CERNVM ONLINE

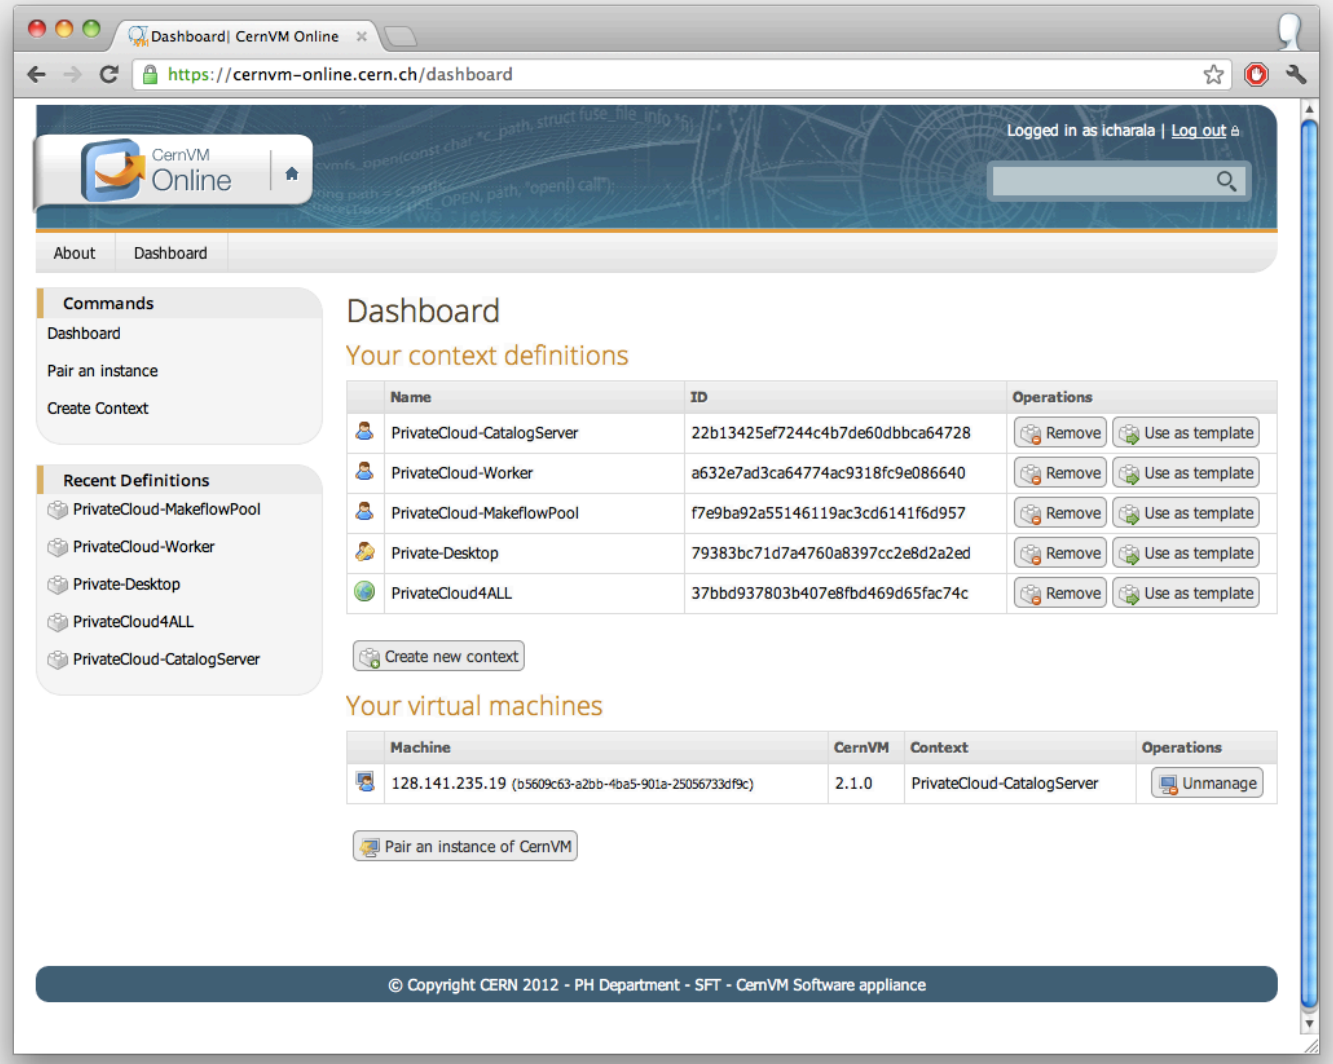

- Develop a data/computing platform targeting scientific communities, deployable on multiple hardware and provisioned over hybrid (private or public) e-infrastructures.
  - Which means cloud infrastructures + HPC clusters
  - https://www.indigo-datacloud.eu/

Key points:

- Based on Open Source solutions, will develop Open Source software.
- Rooted in use cases and support by multi-disciplinar scientific communities, big and small.
- Exploitation of available, general solutions rather than on custom, home-made specific tools or services.
- Possibility to run the software in a hybrid, distributed Cloud environment.

LITERATURE

- S. Bagnasco, D. Berzano, R. Brunetti, S. Lusso, S. Vallero, "Managing a Tier-2 Computer Centre with a Private Cloud Infrastructure", Proceedings of ACAT2013, Beijing, *J. Phys.: Conf. Ser. 1742-6596* **523** 012012 (2014)

- S. Bagnasco, D. Berzano, R. Brunetti, S. Lusso, S. Vallero, "Integrating multiple computing needs via a Private Cloud infrastructure", proceedings of CHEP2013, Amsterdam, *J. Phys.: Conf. Ser. 1742-6596* **513** 032100 (2014)

- D. Berzano, J. Blomer, P. Buncic, I. Charalampidis, G. Ganis, G. Lestaris and R. Meusel, "PROOF as a Service on the Cloud: a Virtual Analysis Facility based on the CernVM ecosystem", proceedings of CHEP2013, Amsterdam, *J. Phys.: Conf. Ser. 1742-6596* **513** 032007 (2014)

Questions?

Stefano.Bagnasco@to.infn.it