Infrastructure as a Service

Jose Castro Leon
CERN – IT/OIS

Cloud Computing

CERN IT Department

- On-Demand Self-Service
- Scalability and Efficiency
- Resource Pooling
- Rapid elasticity

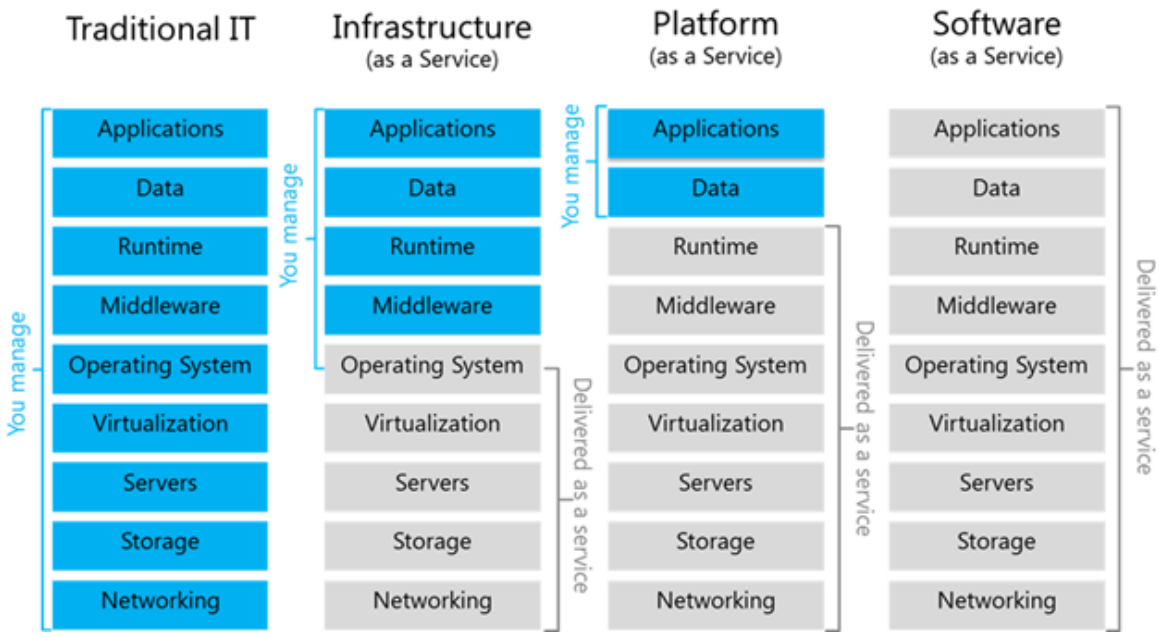

CERN Agile Infrastructure

Infrastructure as a Service

- Objectives
  - 90% of hardware virtualized
  - 300,000 VMs needed
  - Merge the 2 IT evaluation projects in this area
    - Lxcloud & CVI
- Infrastructure as a Service (IaaS)
  - More efficient use of our hardware
  - Better tracking of usage
  - Enable remote management for new data centre
  - Consolidate support to a single service
  - Support potential new use cases (PaaS)

Openstack

- Openstack provides an operating platform, or toolkit, for orchestrating clouds.

- Features
  - Multi Site
  - Scalable
  - Cloud interfaces
    - Position ourselves as another Cloud Provider

Openstack Community

- >150 companies supporting Openstack

- Find the complete list at:

  http://www.openstack.org/community/companies/

Openstack @ CERN

CERN IT Department

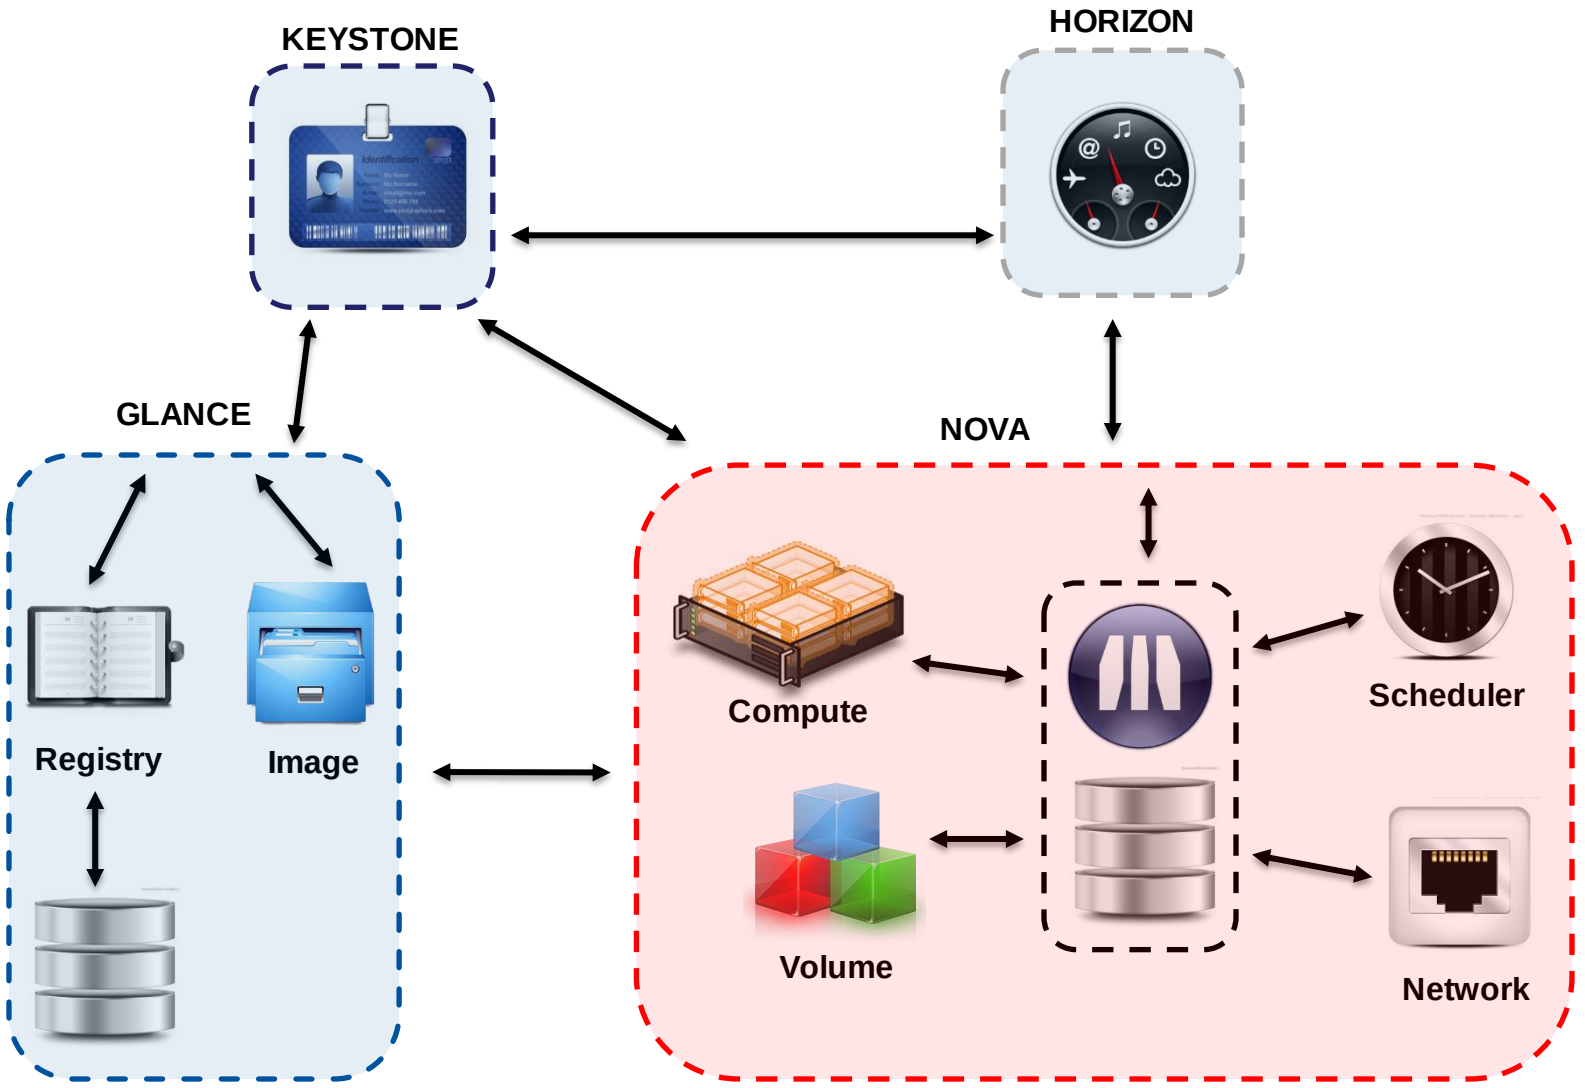

Nova @ CERN

- Nova
  - Instance & Volume Management
    - 2 APIs: OpenStack and Amazon EC2
    - Scheduler decides the Hypervisor
- Nova @ CERN
  - KVM as Hypervisor
  - GlusterFS for High Availability
  - Integration with existing network DB (LanDB)
    - MAC to IP mapping
    - Instance to Network IP subset mapping

Glance @ CERN

- Images Management
- Configurable storage backend
  - Local Storage
  - Openstack Swift
  - GlusterFS
- Multi-format image registry
  - RAW
  - VHD (Hyper-V)
  - VDI (VirtualBox)
  - QCOW2 (Qemu/KVM)
  - VMDK (VMWare)
  - OVF (VMWare, others)

Keystone @ CERN

CERN IT Department

- Cloud Identity Service
  - Authentication
  - Authorization
  - Service Catalog
- Role Based Access Control
- Identity Backends
  - LDAP: Integrate with CERN Active Directory
  - SQL Database possible

Horizon @ CERN

- Graphical UI for Openstack
  - Django application
  - Plug-in architecture
  - Highly integrated with Keystone
  - Self-service for end-users
- Integrate with CERN SSO
  - Shibboleth authentication

Applications

- Different services consume IaaS resources
  - Server consolidation, "user VMs"
  - Batch services
  - Pilot frameworks
  - BOINC
- With different requirements (QoS)
  - Management, Storage, Availability, Flexibility
  - Disaster Recovery
- Scheduler is needed for automation
- Monitor the resource usage

Scheduling Architecture

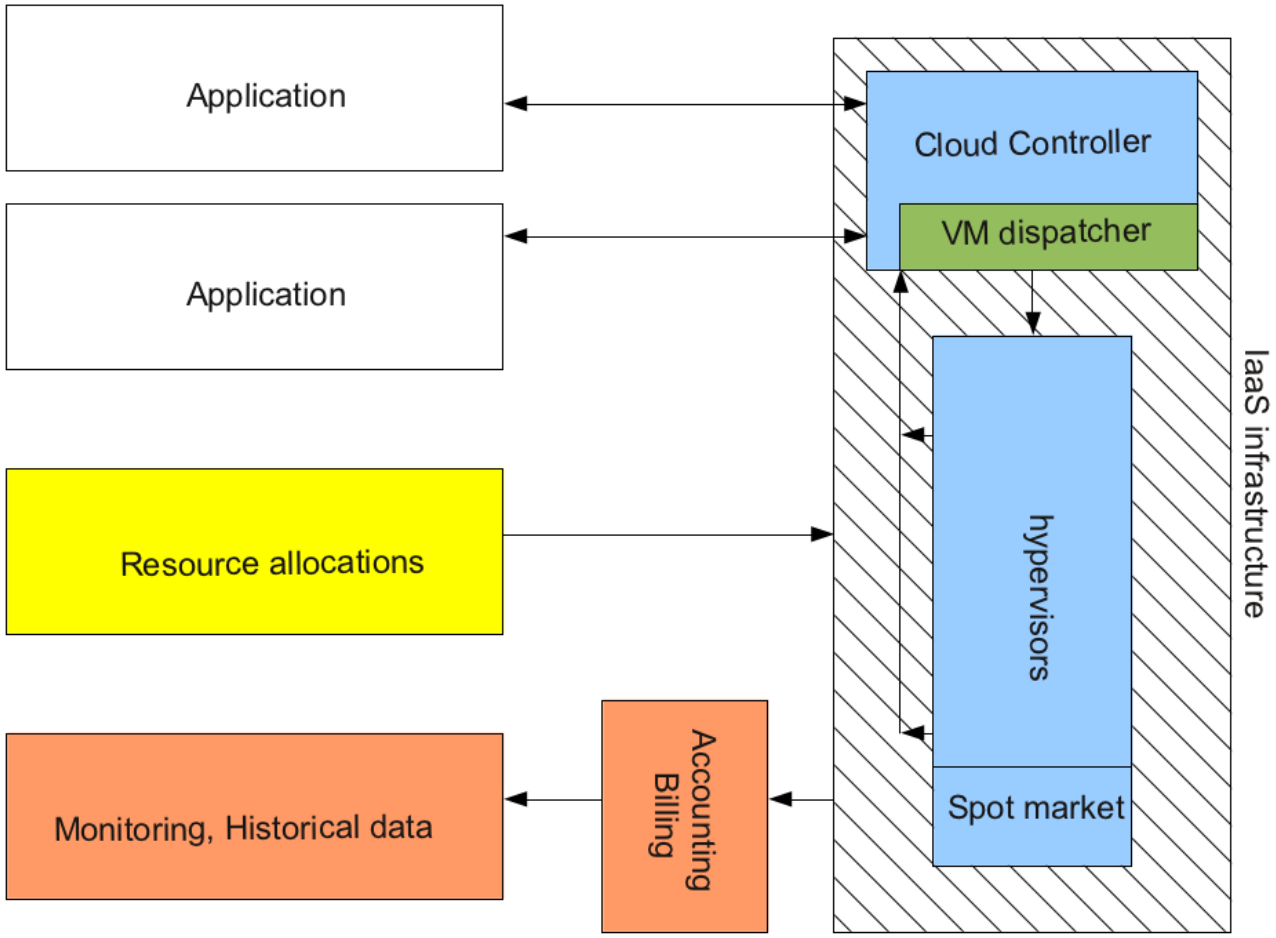

Scheduling & Accounting

- One single entry point for all services
- Decouple applications from the IaaS
- Distributed Computer Centre
- Resource usage tracking
  - Hardware usage (CPU,Memory,Disk)
  - Software usage (Process monitoring)
- Auditing & Billing

IaaS Configuration

- Production nodes based on SLC6 + EPEL
  - Test deployment based on Fedora 16
- All hypervisors are registered in Foreman
  - Automatic hypervisor provisioning
- Hypervisors use Puppet to configure themselves
  - Puppetlabs modules for Openstack
- CERN Customization
  - Network DB registration

Experiences

- Maintain the project linked with upstream
- Openstack
  - Participate in Openstack Design Summit
- Fedora
  - RPMs for F16/F17 are quickly released
    - Essex EPEL6 expected end of April
  - Bugzilla, Wiki, IRC: we participate & benefit
  - Essex Test day (March 08) good experience
- Issues
  - Sometimes ahead of upstream

Summary

- Agile Infrastructure
  - Convergence of our infrastructures and virtual technologies
  - Implement specific modules for internal use
  - Participate in the Community
  - Production ready in time for new Data Centre

Questions

CERN IT Department

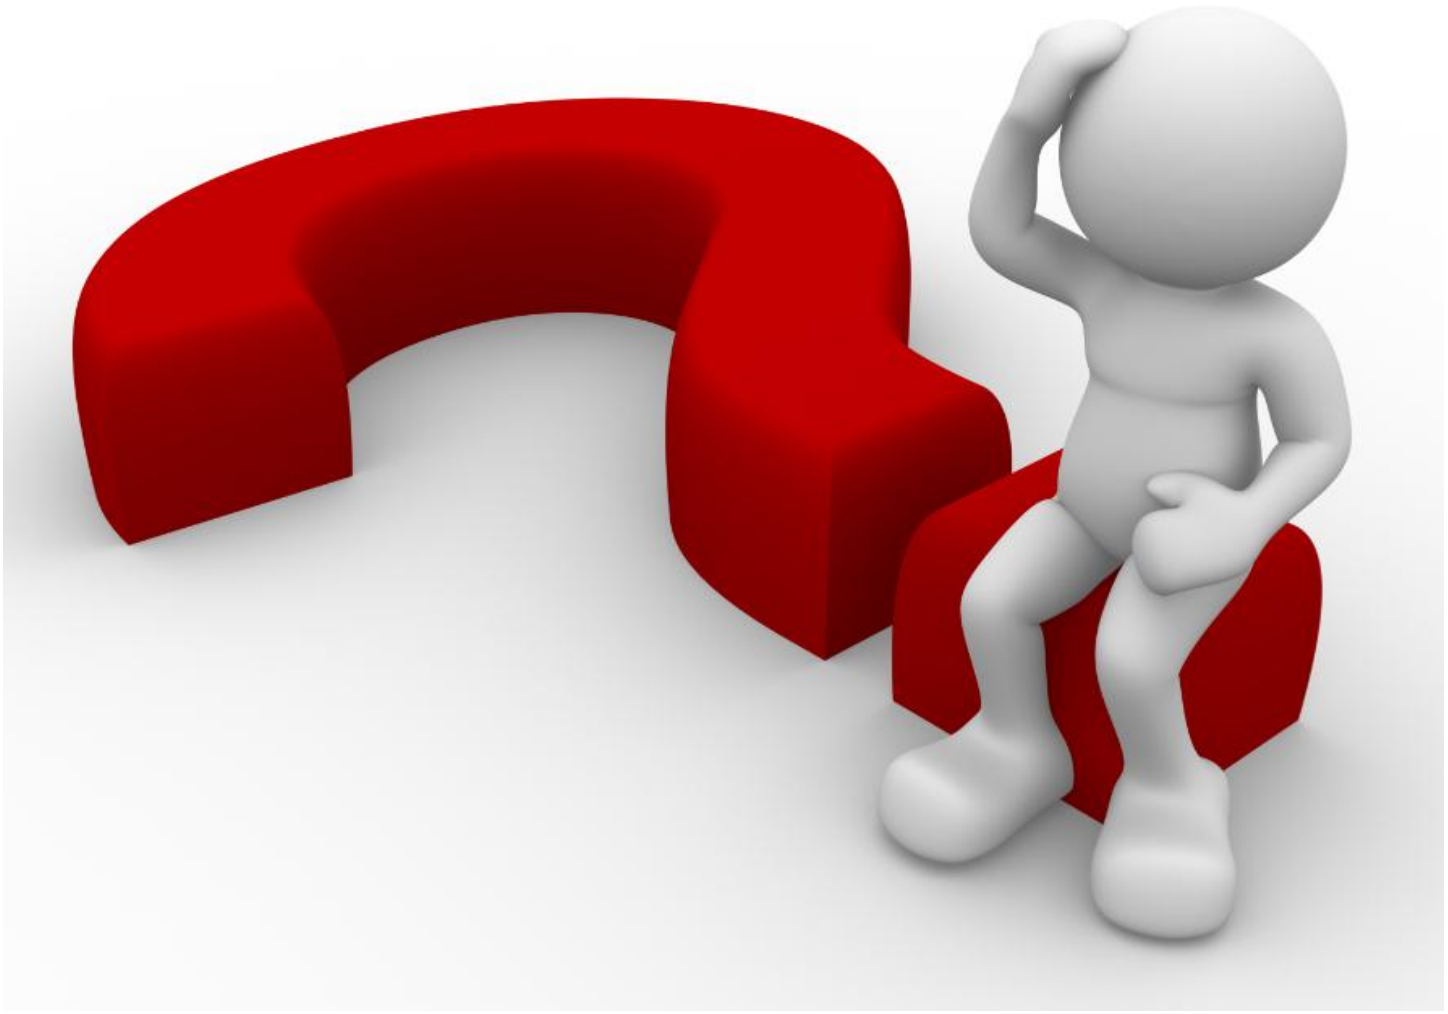

Backup Slides

Openstack Architecture

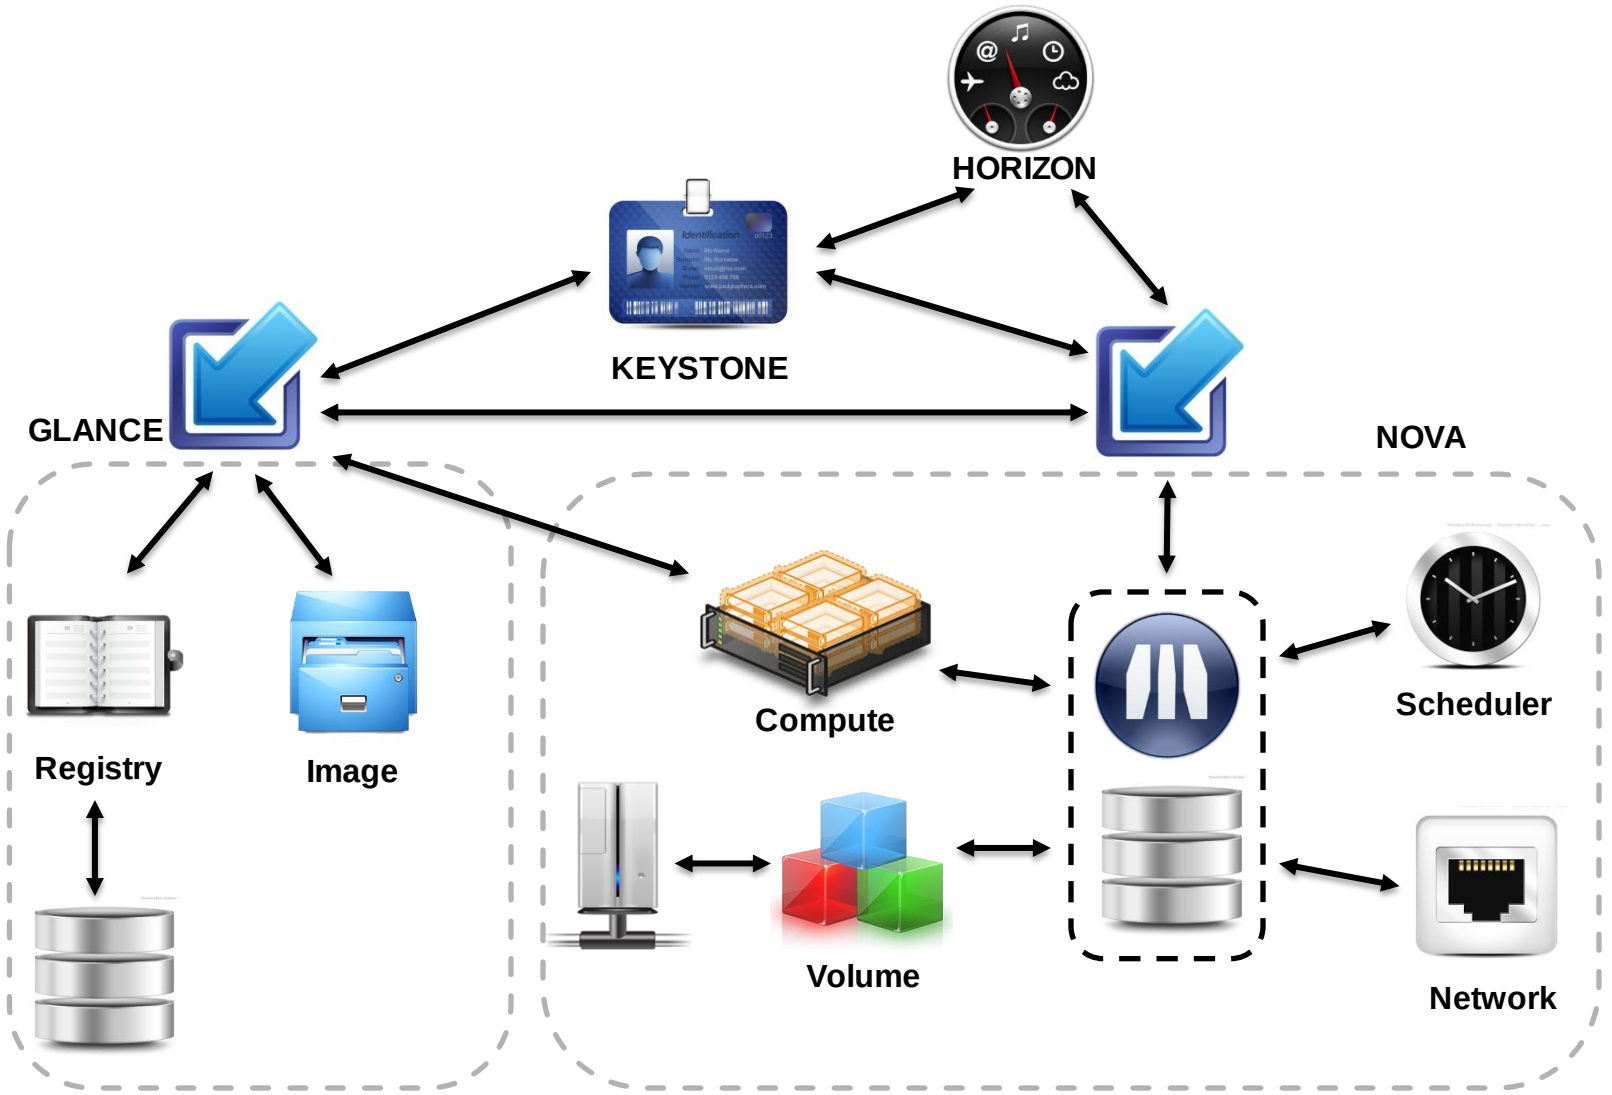

IaaS Configuration Use Case

CERN IT Department

- Hypervisor Configuration Workflow
  1. Machine is physically installed and registered
  2. It's started by the operator
  3. It boots from the network and OS is installed
  4. It is configured by Puppet
  5. It registers itself to the Cloud Controller
  6. It is ready to accept VMs

VM Configuration

- Use the same tools as IaaS
  - Puppet
  - Foreman
- Instance could be managed or not
  - User decides
- Several OS supported
  - Image selection in Horizon
  - Customization of the image

Openstack Architecture

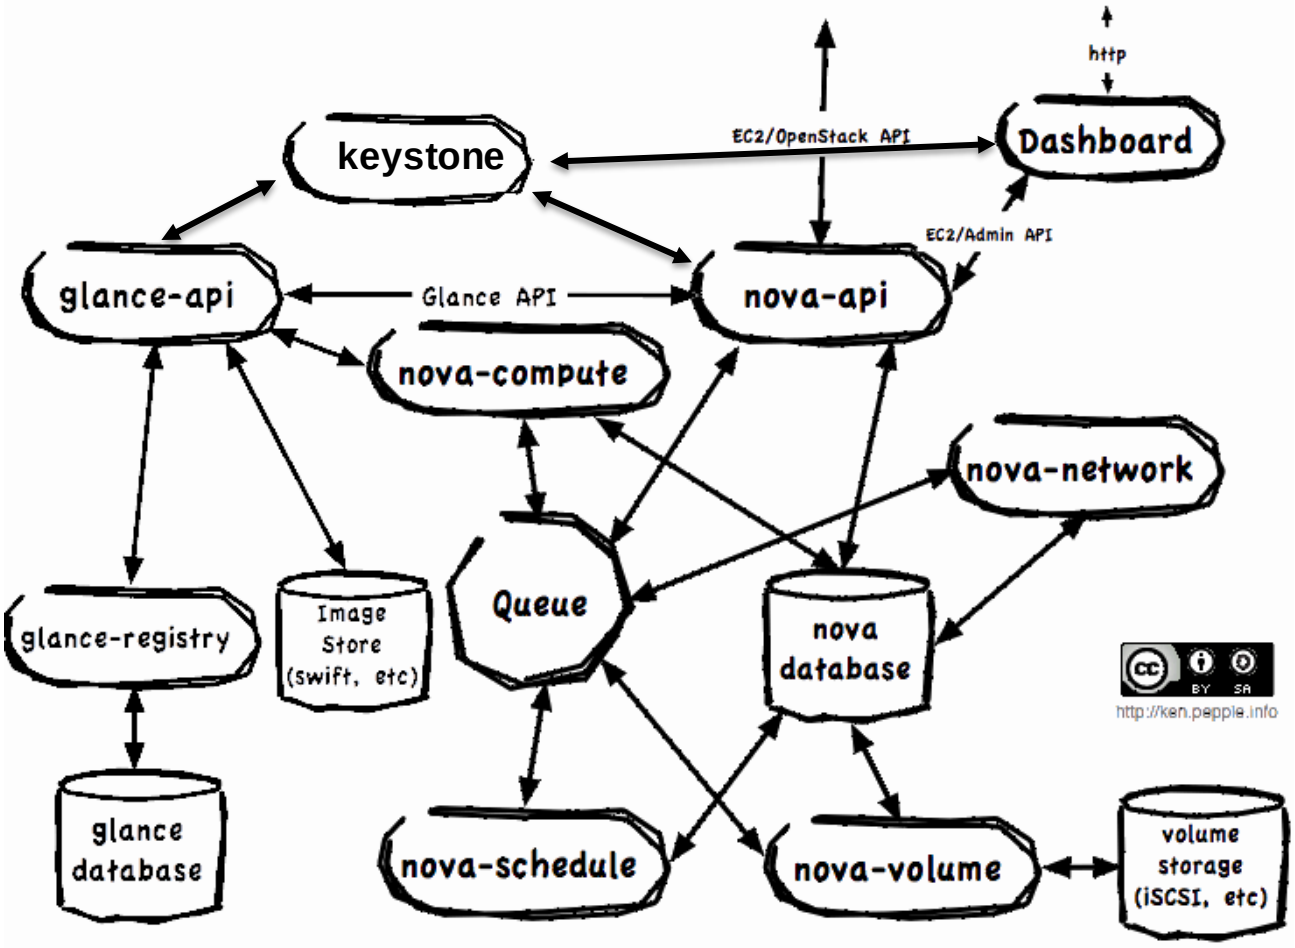

CERN Deployment

CERN IT Department

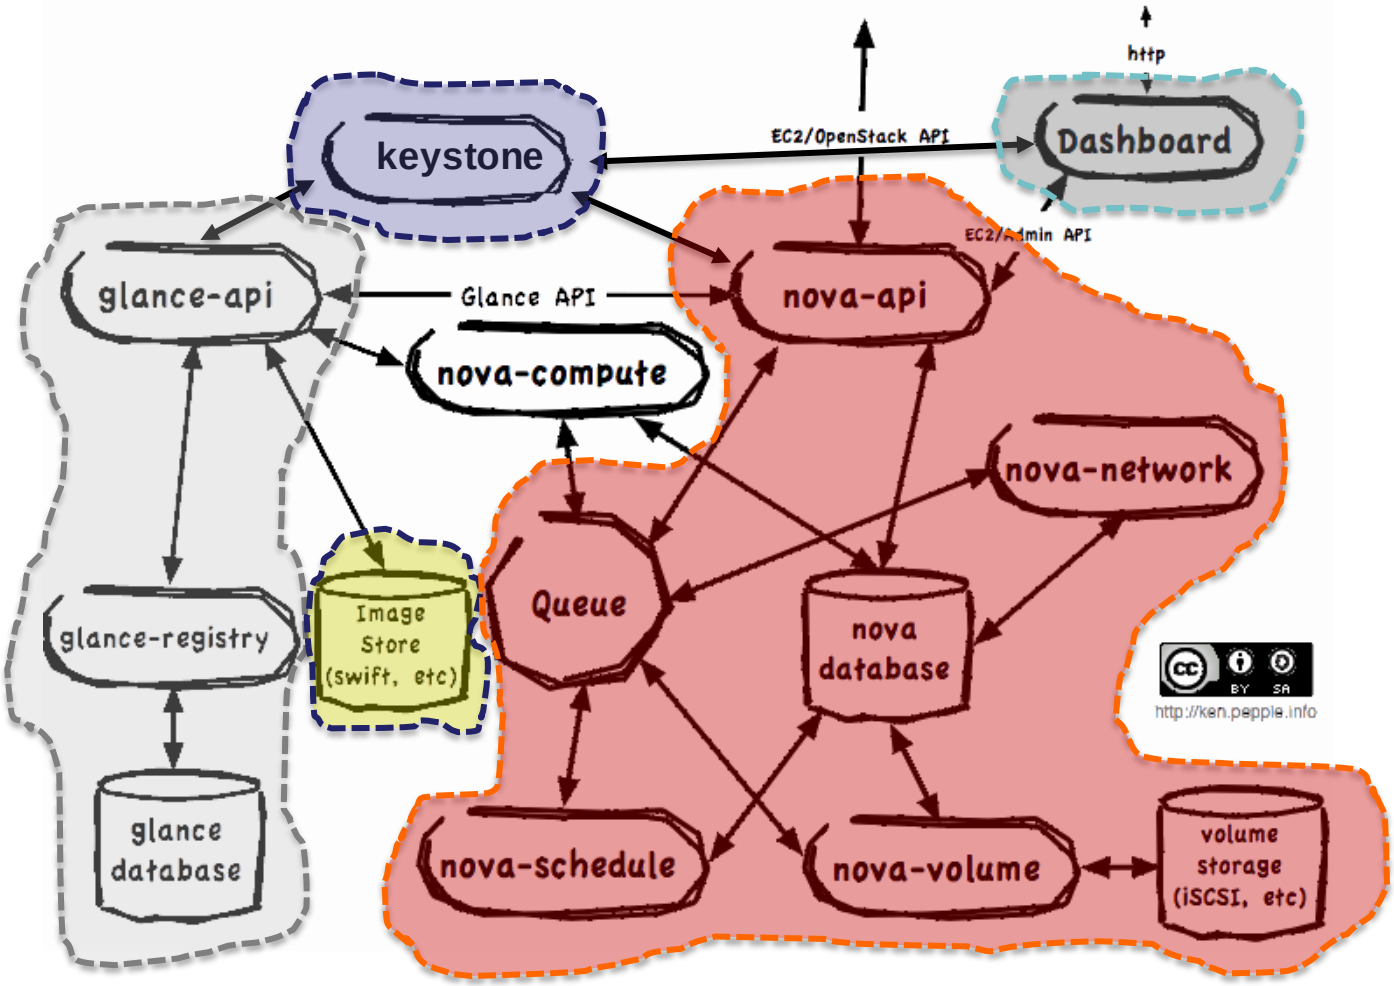

Industry Context - Cloud

- Cloud computing models are now standardising
  - Facilities as a Service – such as Equinix, Safehost
  - Infrastructure as a Service - Amazon EC2, CVI or lxcloud
  - Platform as a Service - Microsoft Azure or CERN Web Services
  - Software as a Service – Salesforce, Google Mail, Service-Now, Indico
- Different customers want access to different layers
  - Both our users and the IT Service Managers

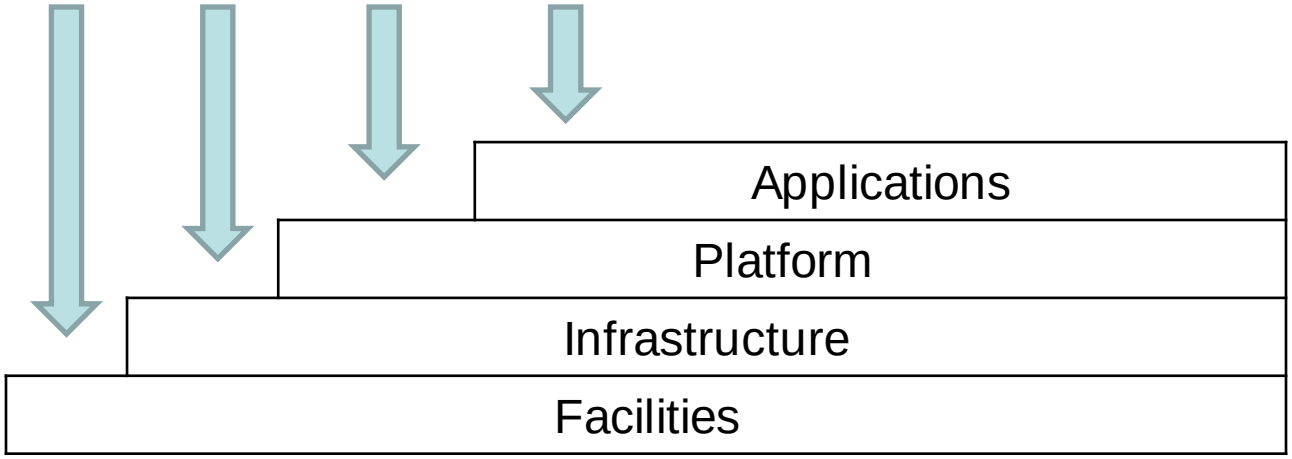

Industry Context - DevOps

CERN IT Department

- DevOps community adopting tool chains with small tools interlinked for particular functions

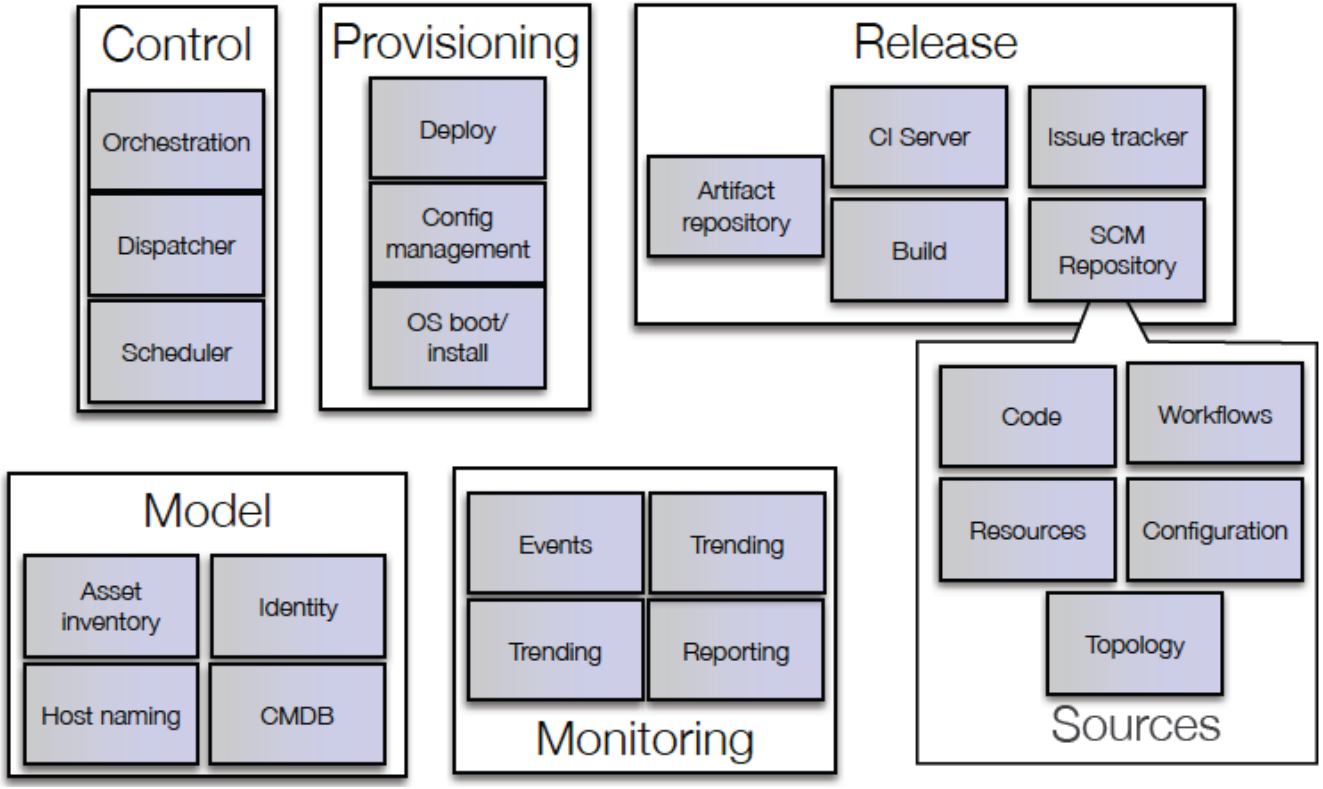

IaaS Approach Benefits

CERN IT Department

- More efficient use of our hardware
  - Use machines throughout their lifetime
  - Reduce impact of hardware failure
  - Pack server workload for day/night, batch and I/O intensive workload
- Better tracking of usage
  - Consistent accounting and provisioning
- Enable remote management for new data centre
  - Only 'smart hands' locally
  - System administrator procedures simplified
- Consolidate support to a single service
  - Reach critical mass for coverage
- Easily exploit work done in industry such as load balancers
  - We are not special
- Support potential new use cases
  - PaaS, cloud APIs, bursting

Current situation

- IT has two projects in this area
  - Lxcloud based on OpenNebula providing EC2 services for physics users to validate concepts and test
  - CERN Virtual Infrastructure on Hyper-V providing long lived server consolidation workloads such as desktop replacements, VO Boxes and other production/test/dev workloads
- Neither solution is felt to be ideal from long term self-sustainability and scalability
  - O(300K) VMs needed
  - Currently two teams of 1.5 FTE each