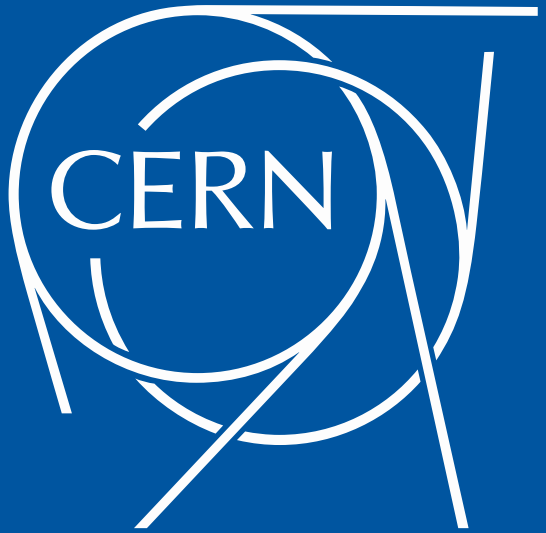
CERN

Federated Identity and OpenStack Cloud Services

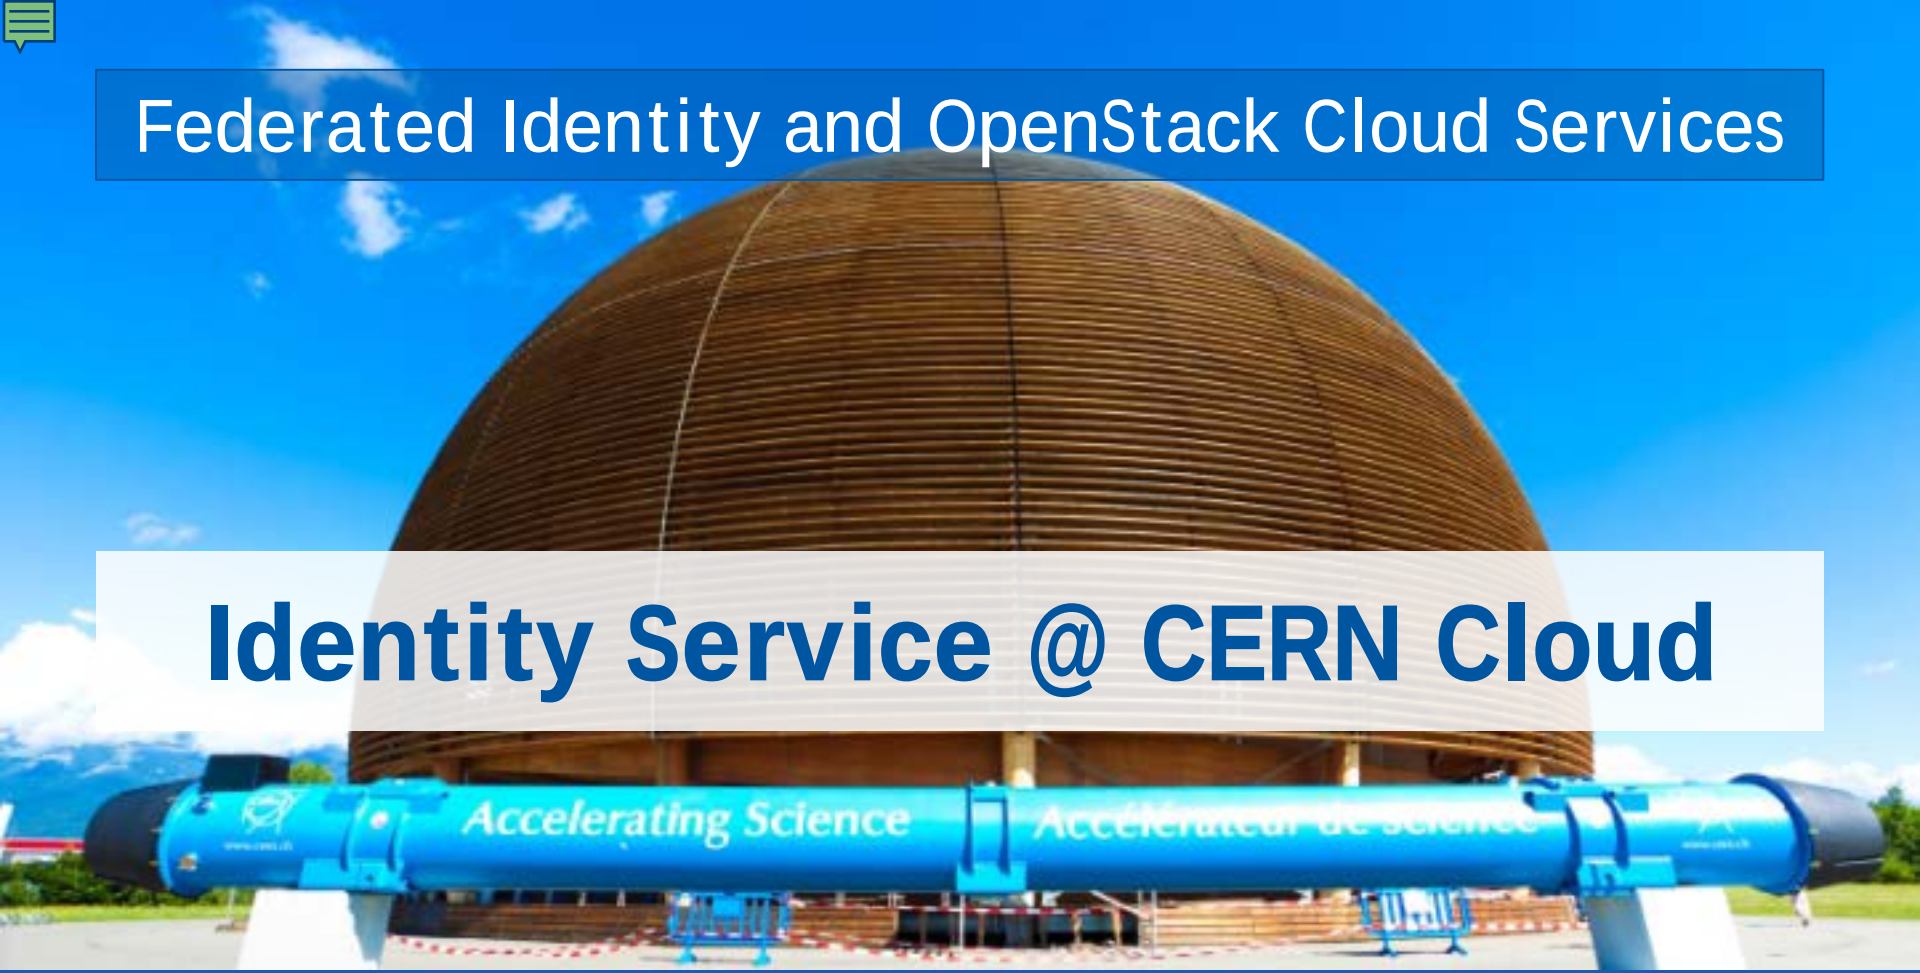

Identity Service @ CERN Cloud

Jose Castro Leon
CERN Cloud Infrastructure

Outlines

- CERN
- Agile Infrastructure
- OpenStack
- Identity Service @ CERN Cloud
  - Architecture
  - Federation
  - Indigo
- Questions

CERN

European Organization for Nuclear Research

- Founded in 1954,
- Today 21 member states,
- World's largest particle physics laboratory
- Located at Franco-Swiss border near Geneva
- ~2'300 staff members, >12'000 users

CERN IT to enable the laboratory to fulfill its mission:

"CERN's mission is fundamental research in physics: pushing back the frontiers of human knowledge."

Agile Infrastructure

Identify new tools needed to build CERN Cloud Infrastructure

Configuration Tools

Cloud Software

Monitoring Tools

Storage Solution

OpenStack Architecture Overview

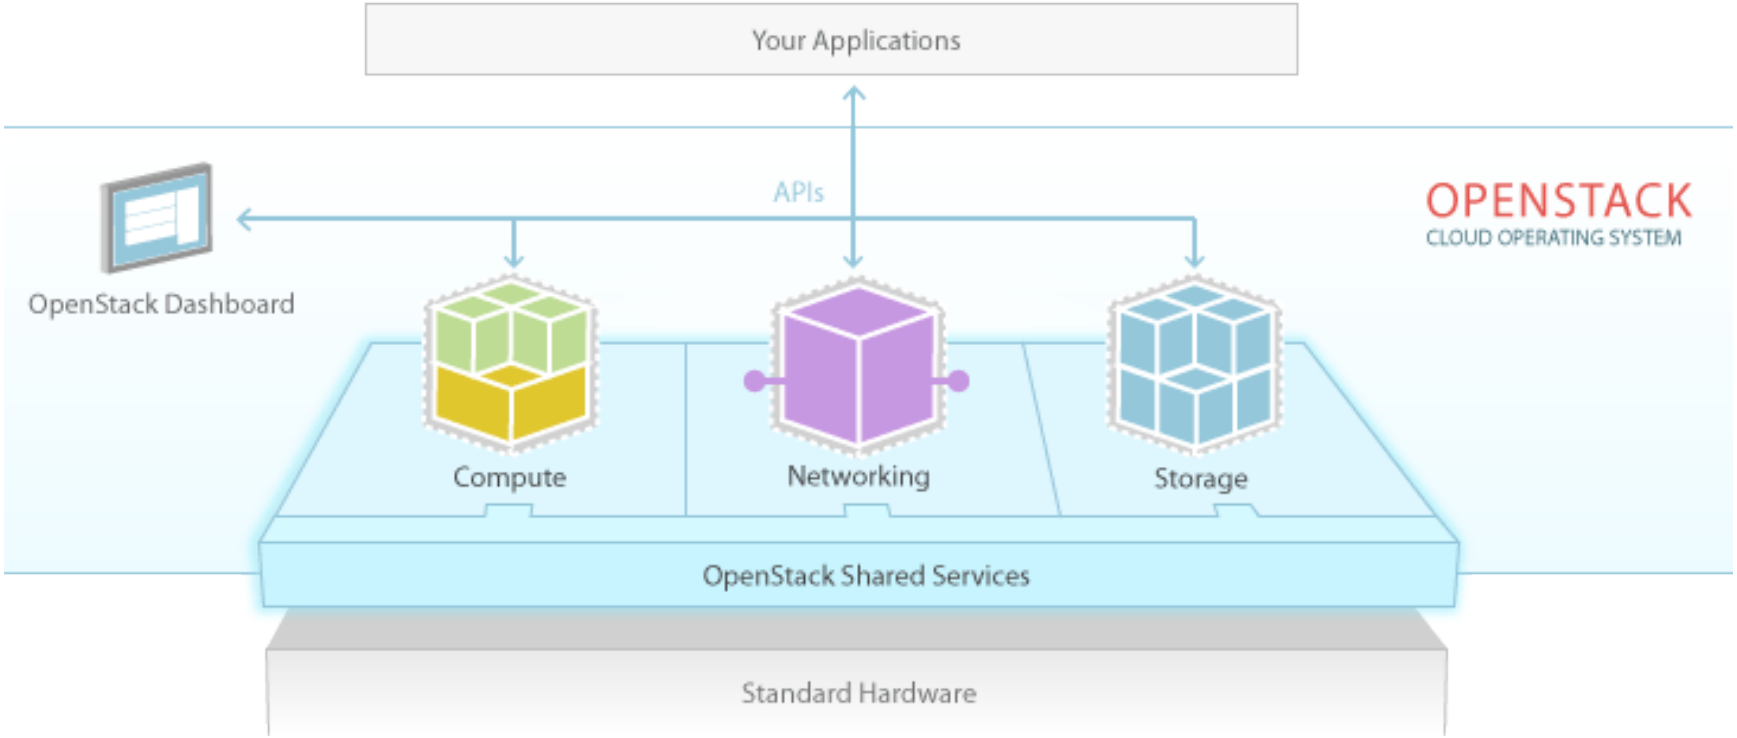

OpenStack projects

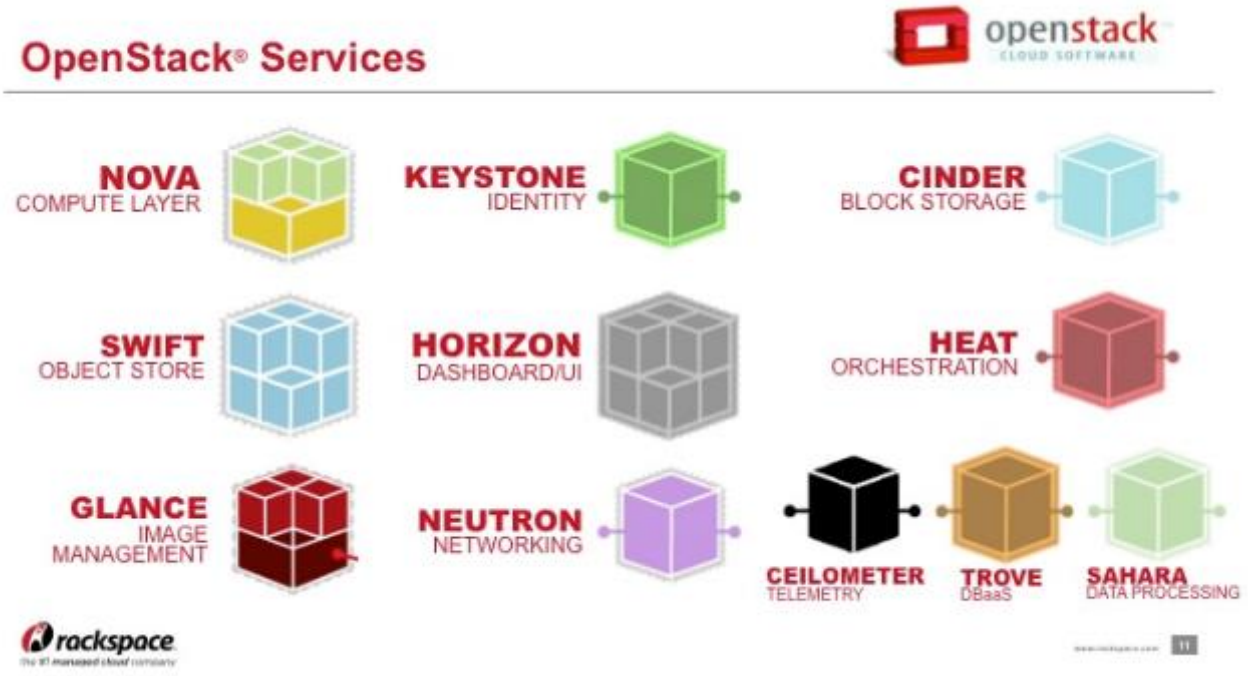

OpenStack @ CERN

- Configuration infrastructure based on Puppet
- Community Puppet modules for OpenStack
- Scientific CERN 6, CERN CentOS 7 Operating Systems
- RDO Community Packages

- Series of (pre)production services increasing features and scale
- Private Cloud extended across computer centres

CERN OpenStack projects

- Modular architecture
- Designed to easily scale out

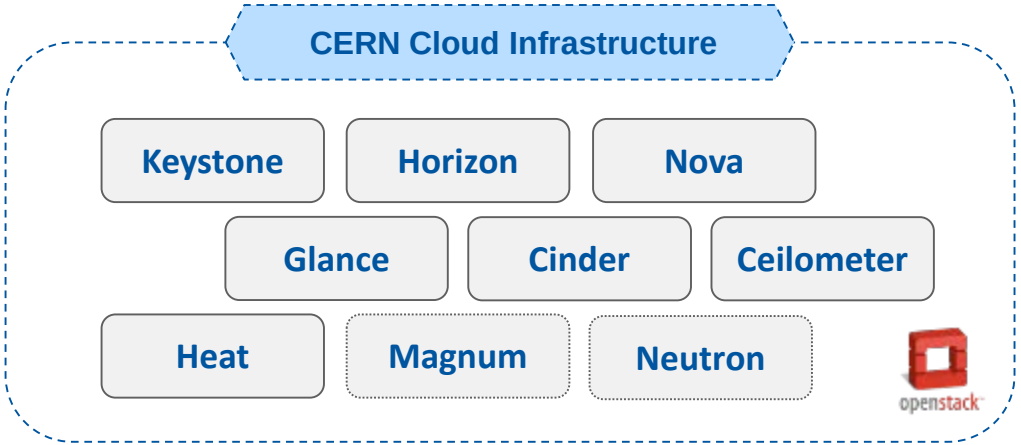

Identity Service @ CERN

- Integrated with CERN AD via LDAP backend
  - CERN's Active Directory infrastructure:
    - Unified Identity across the site
    - 44k users, 30k security groups
    - ~200 arrivals/departures per month

Identity Service @ CERN

- Integrated with Resource Management
  - Self-service subscription for CERN users
    - Created "Personal Project" with limited quota
  - Shared projects created upon request
  - ~2000 users, ~2400 projects

- Lifecycle compliant with standard CERN policies

Identity Service @ CERN

- Several AuthN methods

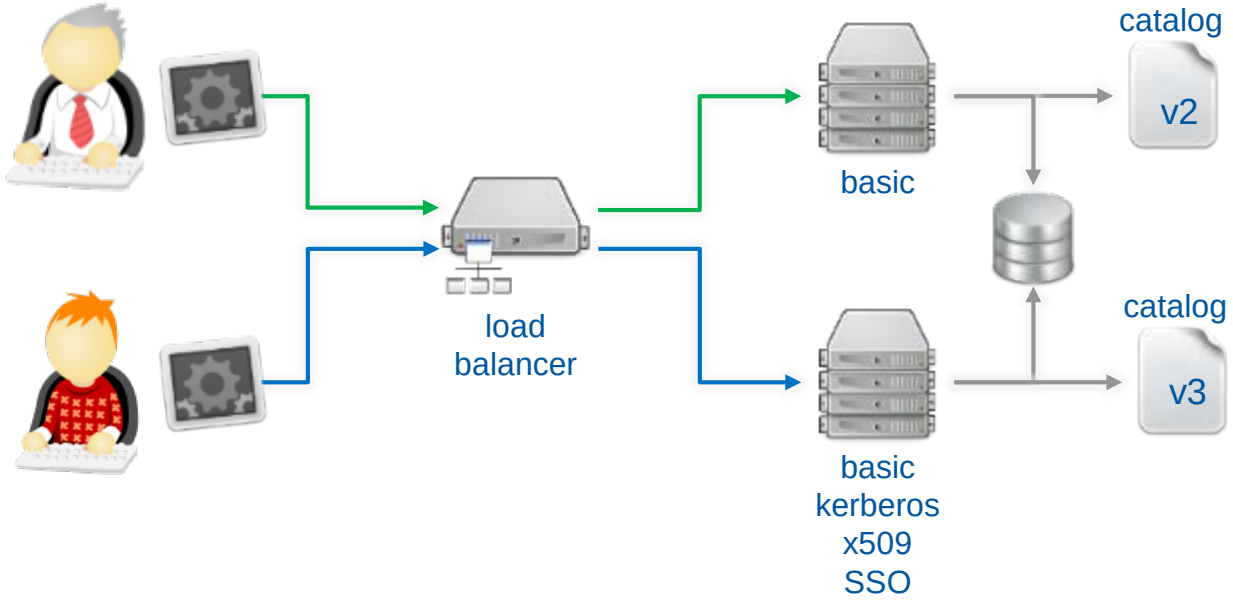

Federation

- Joint project with rackspace
- Hybrid Clouds
  - Seamless cloud bursting (Private ⇔ Public)
- 30 Public Cloud providers plan to offer it
- Online Experiment trigger farms

Federation

- Major upstream contribution
- 1st site to have WebSSO enabled
- Keystone configured as SP on CERN IdP
  - Microsoft ADFS + mod_shibd
  - Leverage current eduGAIN federation
- Explore k2k

Federation

- Resource request
  - Shared project lifecycle
  - e-group ACLs
  - Additional configuration (HW mapping…)
- Mapping rules if needed
- Auditing through CADF

Indigo Datacloud

- Develop an open source platform for computing and data targeted at multi-disciplinary scientific communities provisioned over hybrid e-infrastructures

- Participating in WP3, WP4 and WP5
- Providing resources using eduGAIN federation

Work in progress

- Upstream
  - Use of keystoneauth instead of keystoneclient
- CERN
  - ADFS + kerberos authentication plugin

Summary

- OpenStack Federation is there and working
- All the code is open source and upstream
- Extendible to new Auth methods as needed

Thank you

github.com/cernops

openstack-in-production.blogspot.ch

jose.castro.leon@cern.ch

marek.denis@cern.ch

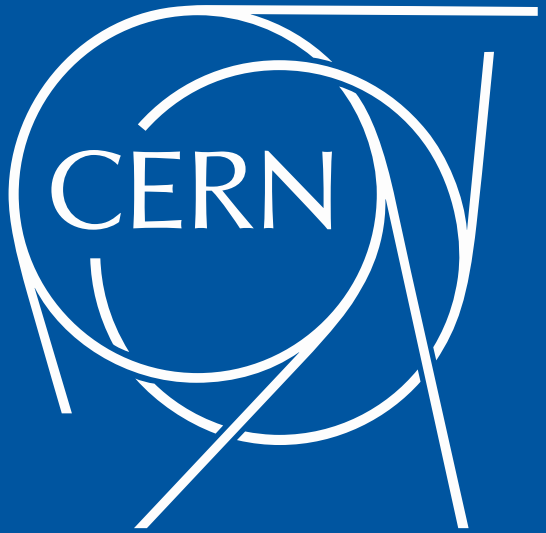
CERN

Backup Slides

CERN IT Infrastructure in 2011

- ~10k servers
  - Dedicated compute, dedicated disk server, dedicated service nodes
  - Mostly running on real hardware
  - Server consolidation of some service nodes using Microsoft Hyper-V/SCVMM
    - ~3400 VMs (~2000 Linux, ~1400 Windows)
  - Various other virtualization projects around
- Many diverse applications (”clusters”)
  - Managed by different teams (CERN IT + experiment groups)

CERN IT challenges in 2011

- Expected new Computer Centre in 2013
- Need to manage twice the servers
- No increase in staff numbers
- Increasing number of users / computing requirements
- Legacy tools - high maintenance and brittle

Why Build CERN Cloud

Improve operational efficiency

- Machine reception and testing
- Hardware interventions with long running programs
- Multiple operating system demand

Improve resource efficiency

- Exploit idle resources
- Highly variable load such as interactive or build machines
- Accountability

Improve responsiveness

- Self-service

Experiments started to use public cloud resources

CERN Integration (Resources)

- Self-subscription service
- Unit of ownership => Project

Type	Request	Affiliation Expired	Account Disabled	Account Deleted
Personal	On Subscription	-	Stop Resources	Delete Resources
Shared	On Request	Promote owner	-	-