

## Agile Infrastructure: an updated overview of IaaS at CERN

Stefano Zilli

on behalf of Cloud Infrastructure Team

HEPiX Spring 2014

# Agenda

- OpenStack
- CERN Private Cloud
- Architecture
- New Features
- Working On ...
- Next Steps ...
- Future Developments
- Questions



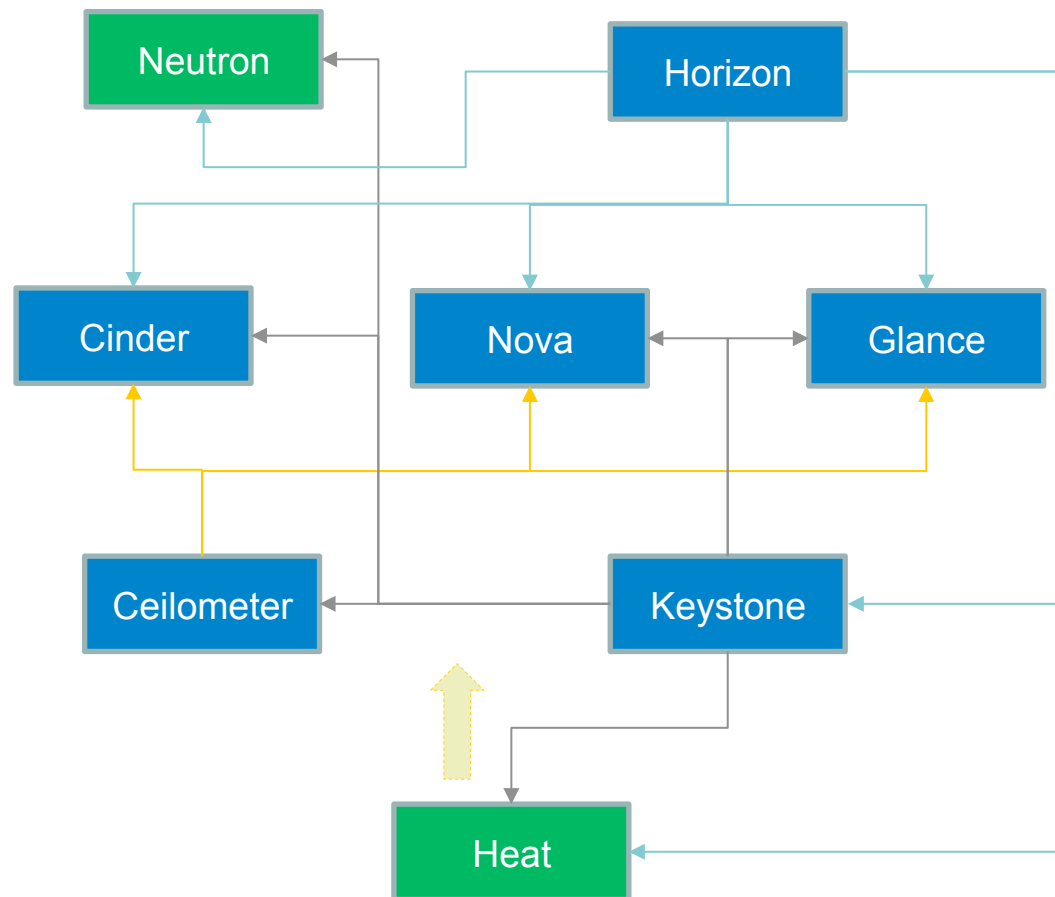
- OpenStack is a collection of open source projects that provides an operating platform for orchestrating clouds in a massively scale.
- Founded by Rackspace Hosting and NASA, OpenStack has grown to be a global software community of developers collaborating on a standard and massively scalable open source cloud operating system.
- All of the code for OpenStack is freely available under the Apache 2.0 license. Anyone can run it, build on it, or submit changes back to the project.



- Six months release cycles:
  - Past releases:
    - Grizzly (April 2013)
    - Havana (October 2013)
  - Current release: Icehouse (April 2014)
  - Next release: Juno (October 2014)
- Lot of services:
  - Compute (Nova)
  - Image Storage (Glance)
  - Identity (Keystone)
  - Dashboard (Horizon)
  - Block Storage (Cinder)
  - Metering (Ceilometer)
  - Orchestration (Heat)
  - Networking (Neutron)
  - Bare Metal (Ironic)
  - Database Service (Trove)
  - Queue Service (Marconi)
  - Object Store (Swift)
  - ...

- 3 pre-production deployments:
  - “Guppy” (June 2012)
  - “Hamster” (October 2012)
  - “Ibex” (March 2013)
- Production deployment:
  - July 2013
  - Initially based on Grizzly
  - Currently on Havana
  - 2850 compute nodes today
  - Target 15000 by the end of 2015
  - EPEL/RDO packages
  - SLC6 Operating System
  - Deployed via Puppet

- Available services:
  - Compute (Nova)
  - Image Storage (Glance)
  - Identity (Keystone)
  - Dashboard (Horizon)
  - Block Storage (Cinder)
  - Metering (Ceilometer)
- Future services:
  - Networking (Neutron)
  - Orchestration (Heat)




- 2 types of projects
  - Personal
  - Shared
- 2 types of roles
  - Owner: user
  - Member: user or e-group
- Quota:
  - Personal: 10 VM, 20 cores, 50GB memory, 10 volumes, 100GB volume space
  - Shared: 50 VM, 50 cores, 100GB memory, 10 volumes, 1TB volume space
- Life cycle for projects


	Affiliation Expired	Account Disabled	Account Deleted
Personal	-	Disabled	Deleted
Shared	Promoted	-	-

# Some Numbers (1/2)

- 2 datacenters:
  - ~2400 hypervisors in Geneva
  - ~450 hypervisors in Wigner




- Number of projects:
  - personal: ~690
  - shared: ~150
  - active (with VMs): ~480







## Some Numbers (2/2)

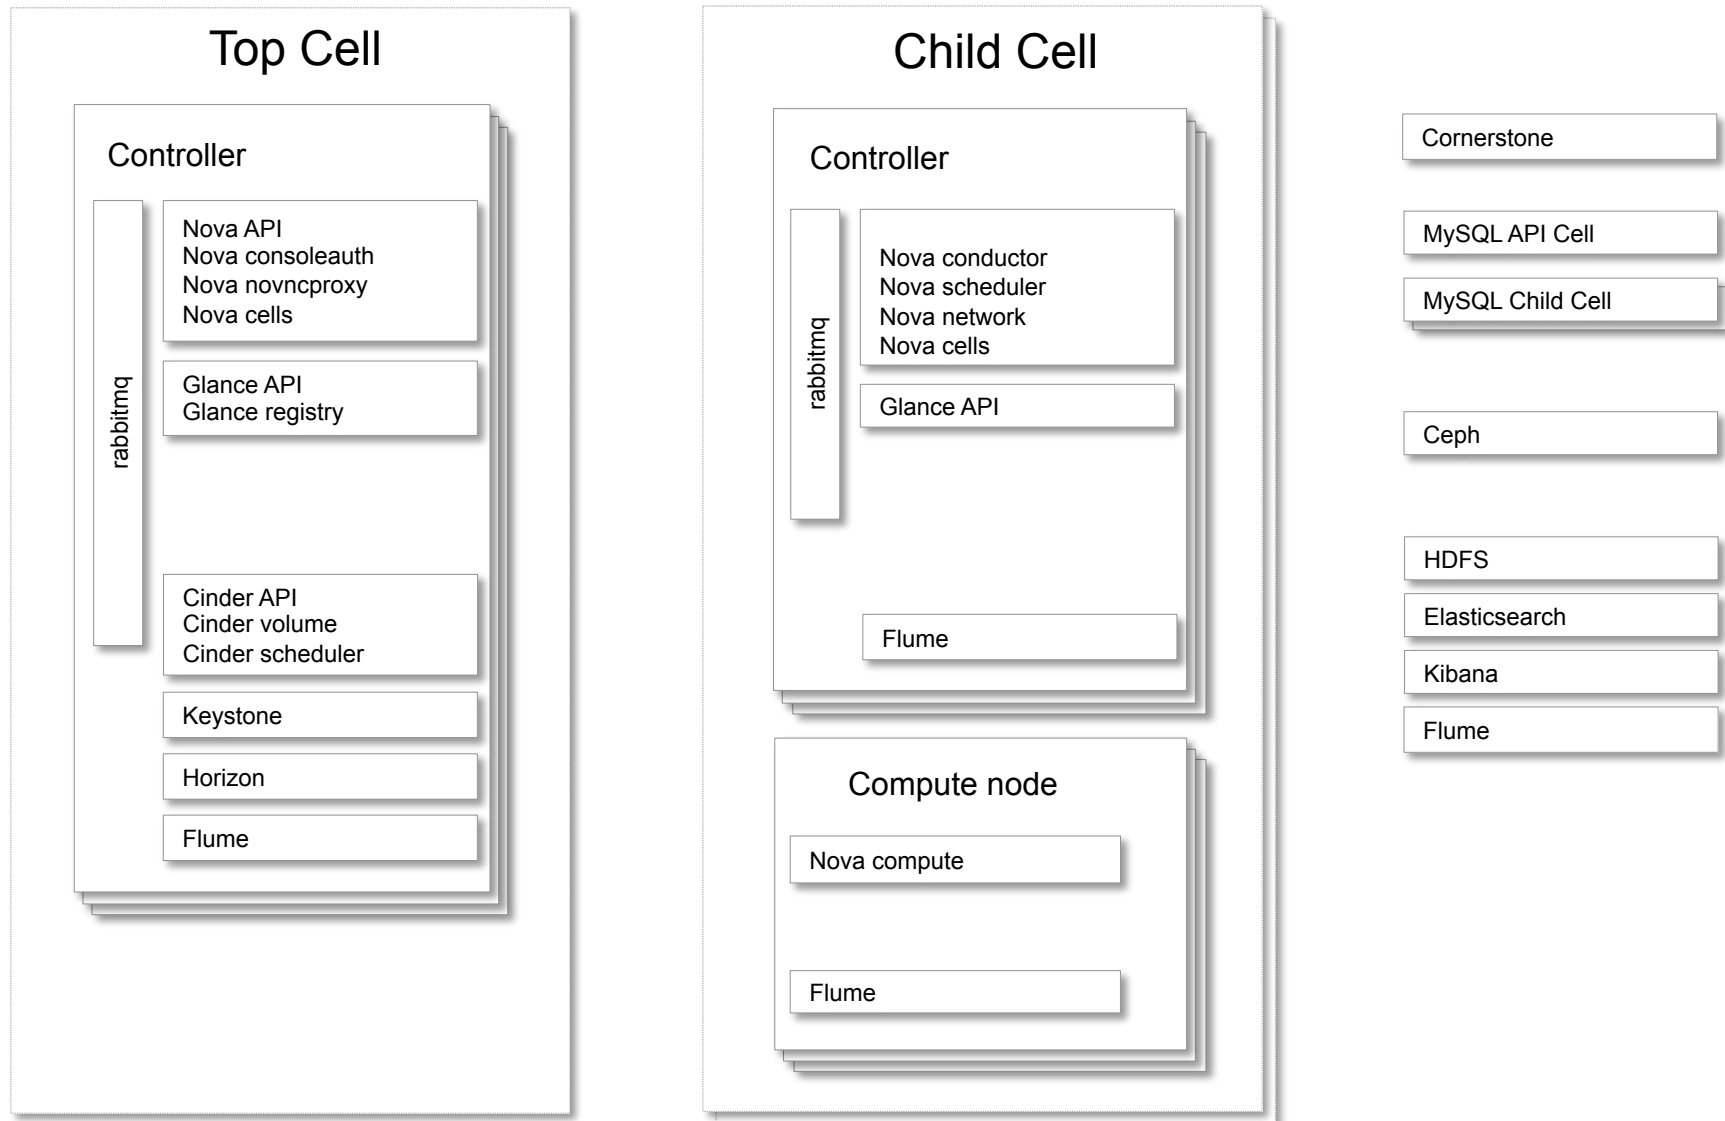
- Number of VMs: ~5800
  - ~2000 Batch System
  - ~700 IT Services
  - ~500 Experiments
  - ...

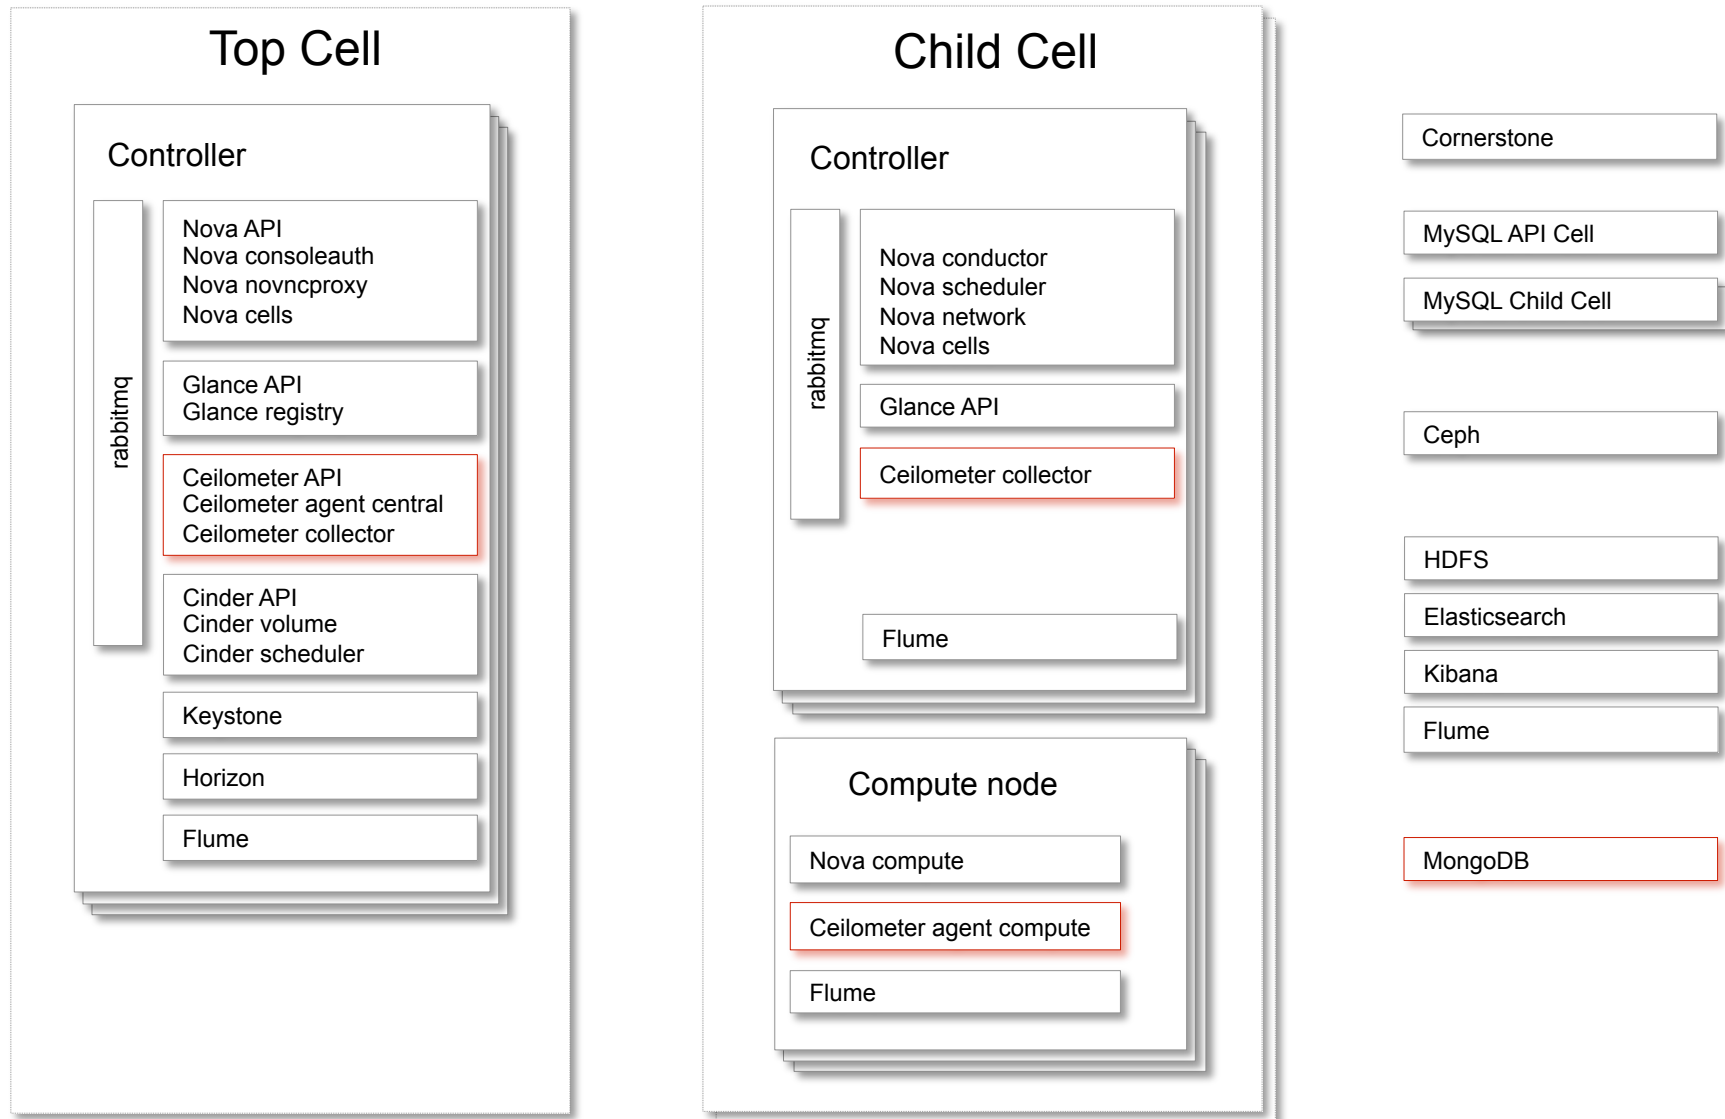


- Number of volumes: ~370
- Total volume size: ~100TB

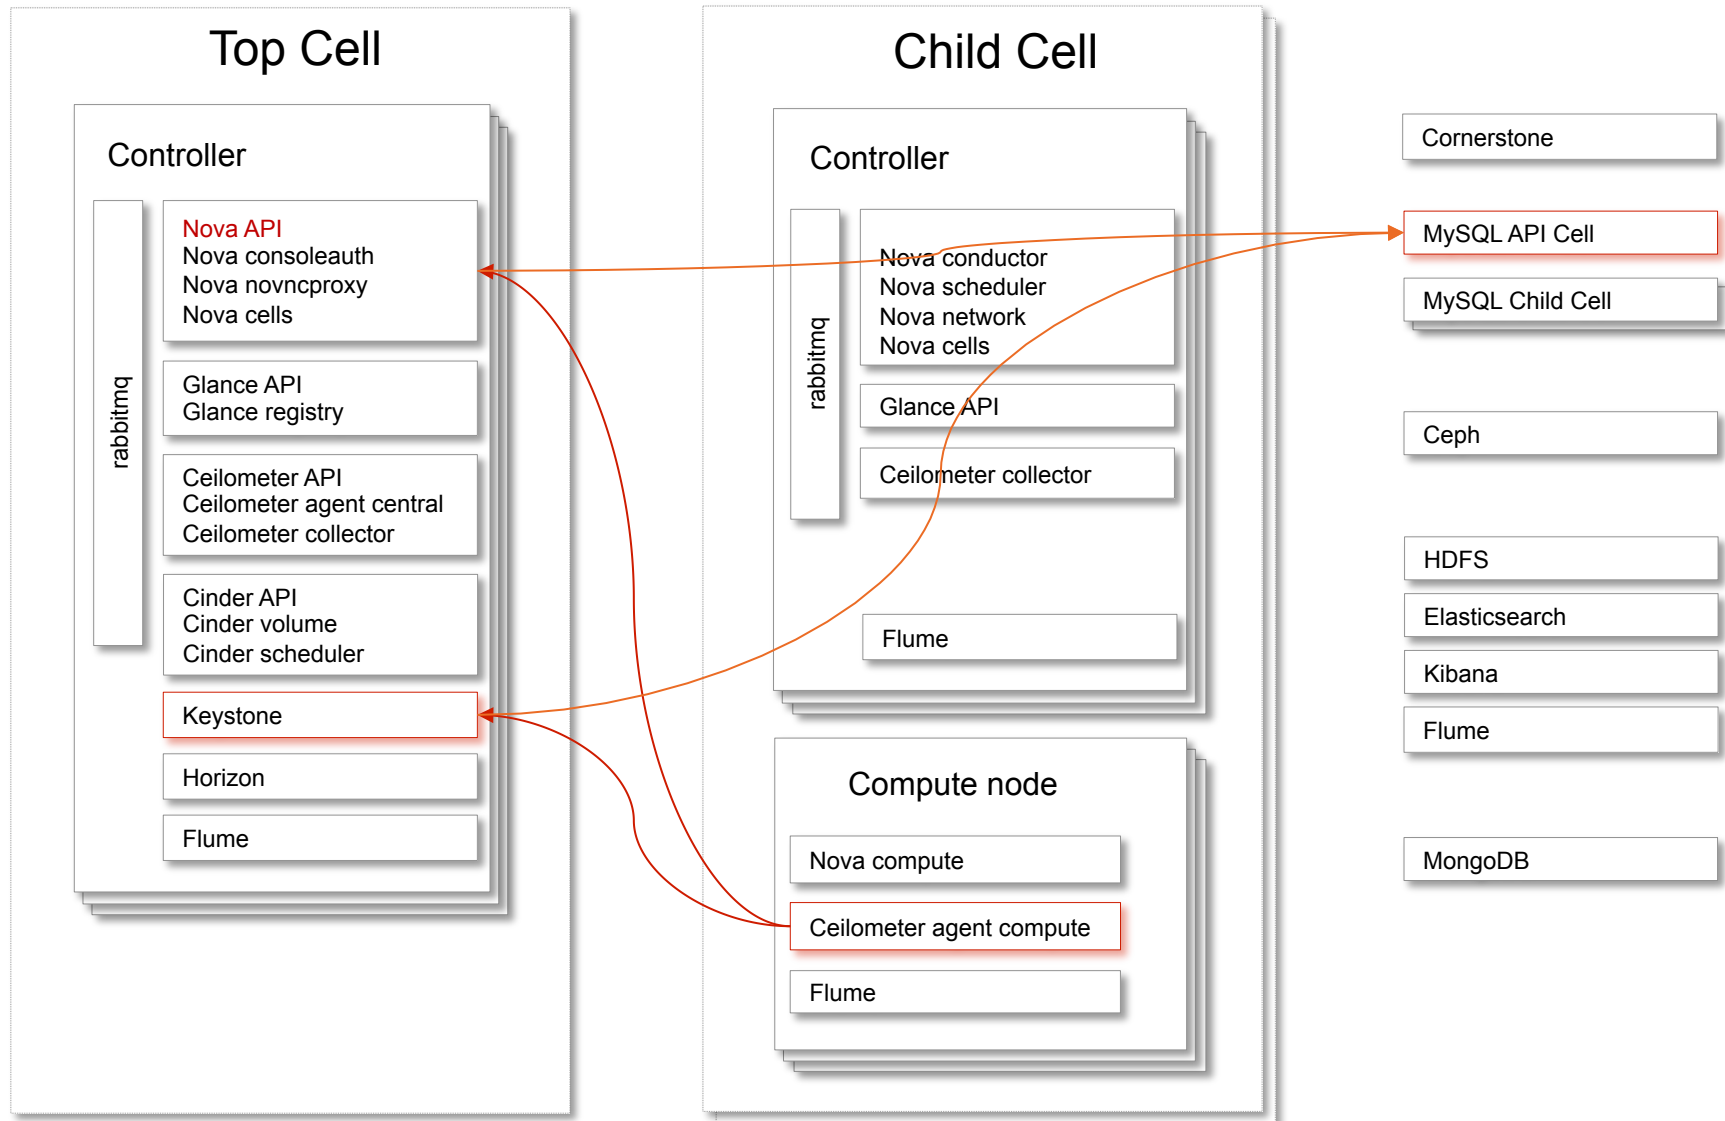


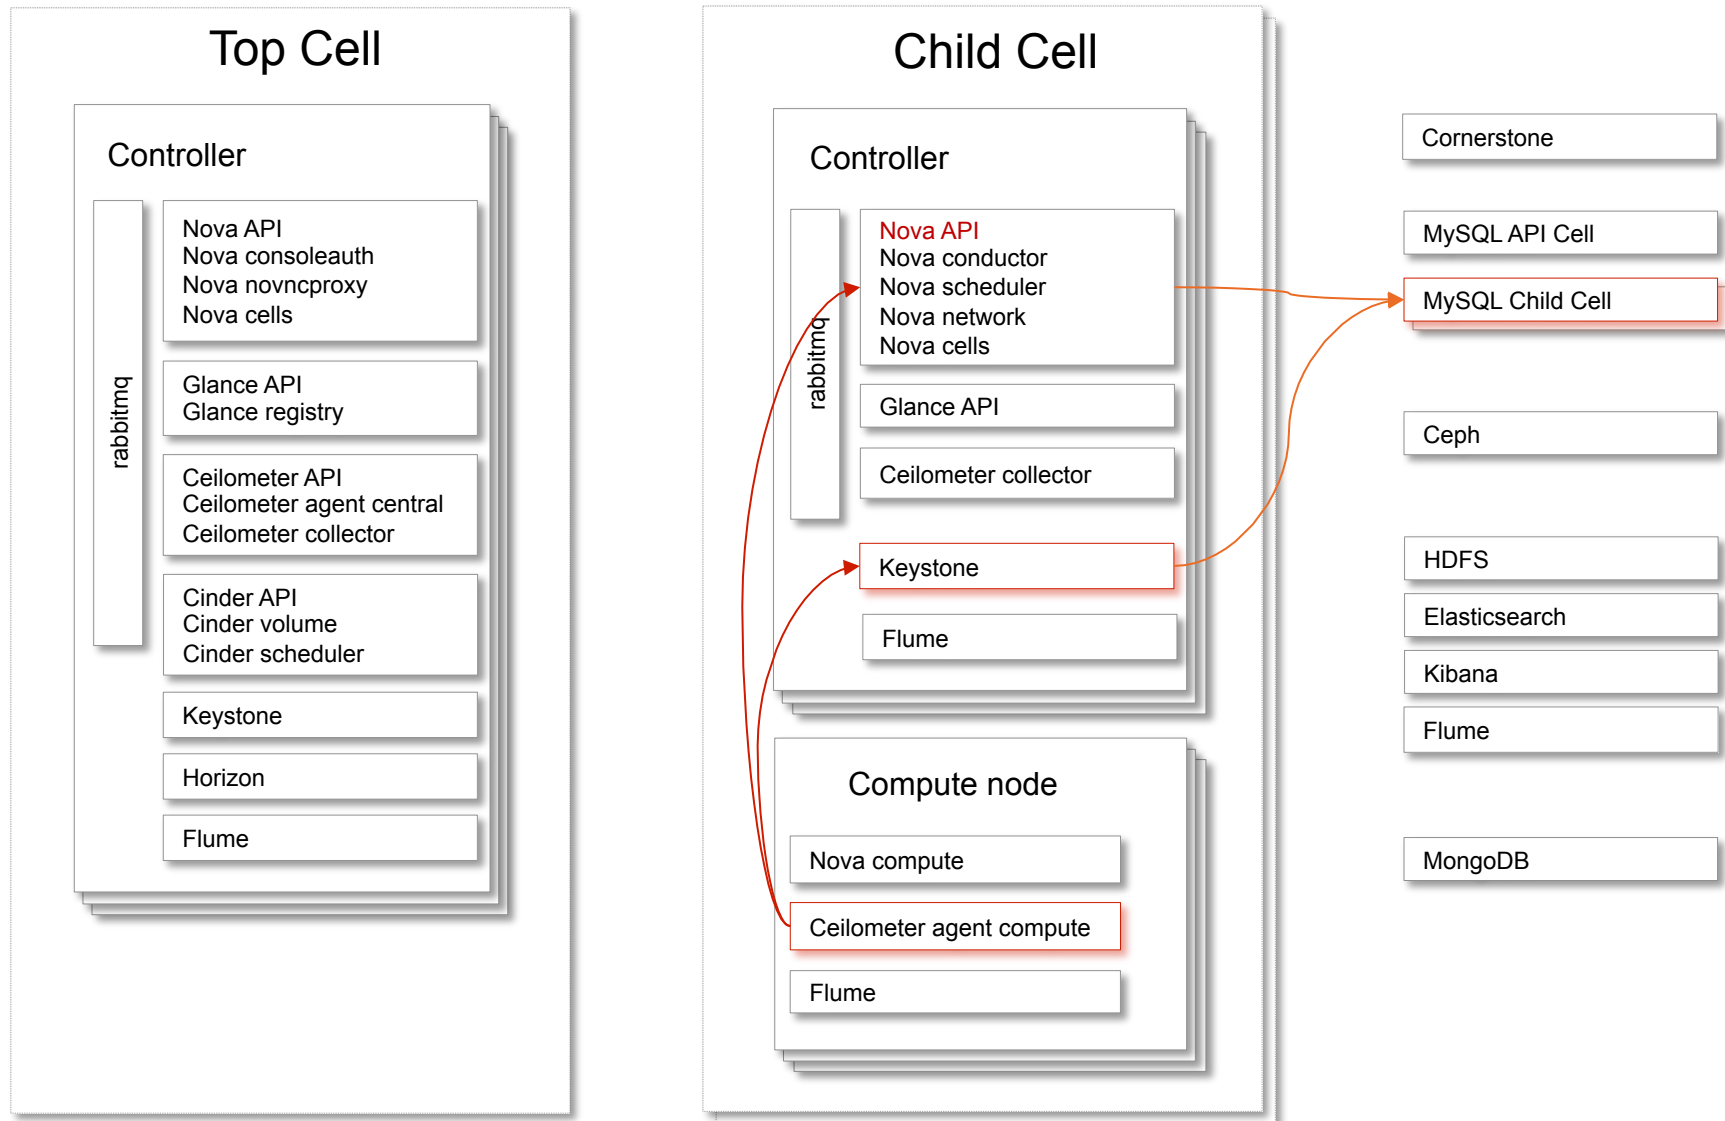






# Architecture (4/5)





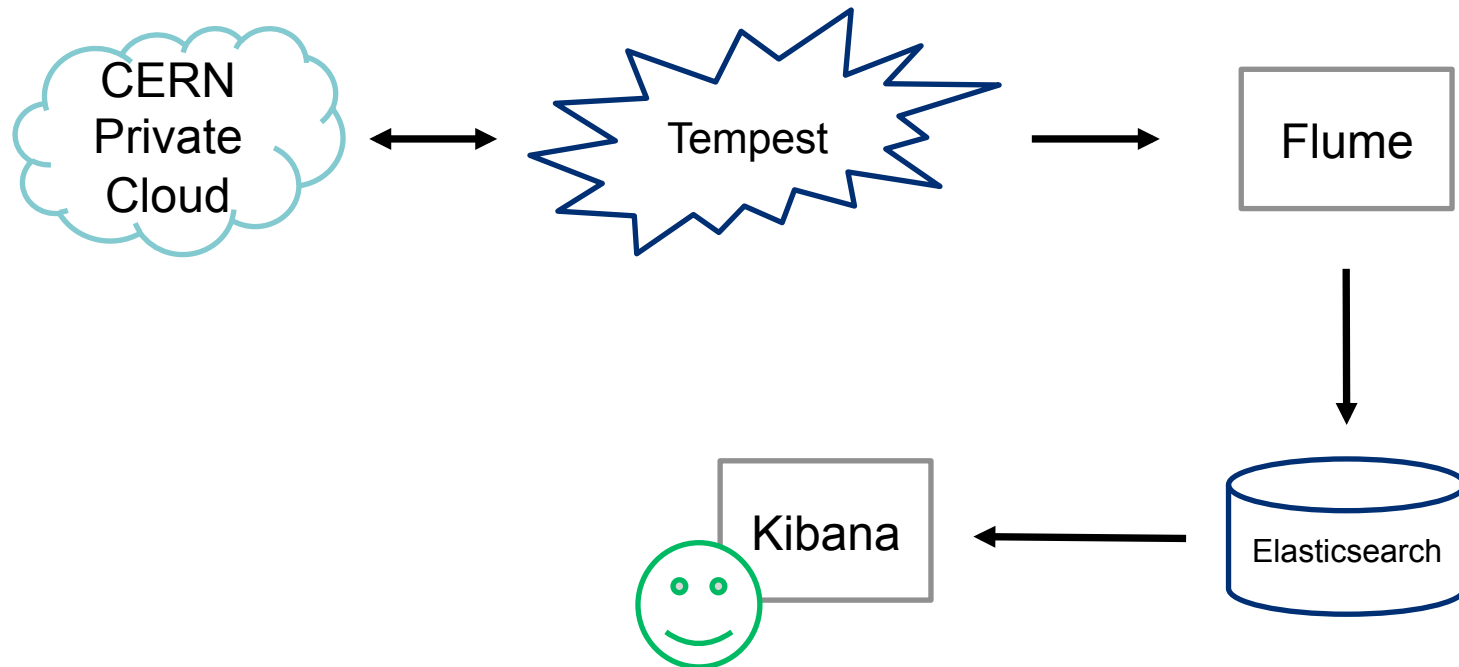


- Nova
  - Live migration
  - IPv6 enabled
  - Update network settings via nova metadata
  - Improvement in cell scheduling
  - More scalability and stability
- Glance
  - Glance on Ceph
  - Shared images
  - New set of images:
    - SLC 5 / 6 CERN specific
    - Windows 8.1 / 2012 R2

- Cinder
  - Cinder volumes GA
    - Personal projects: 10 volumes and 100GB
    - Shared projects: 10 volumes and 1000GB
- Keystone
  - Delegate tasks to project managers:
    - Reassign ownership
    - Add/Remove users
    - Allow access to operators
- Ceilometer
  - Enabled polling on compute nodes
  - Enabled Glance and Cinder notifications
  - Specific role for WLCG accounting user
  - Data TTL set to 3 months



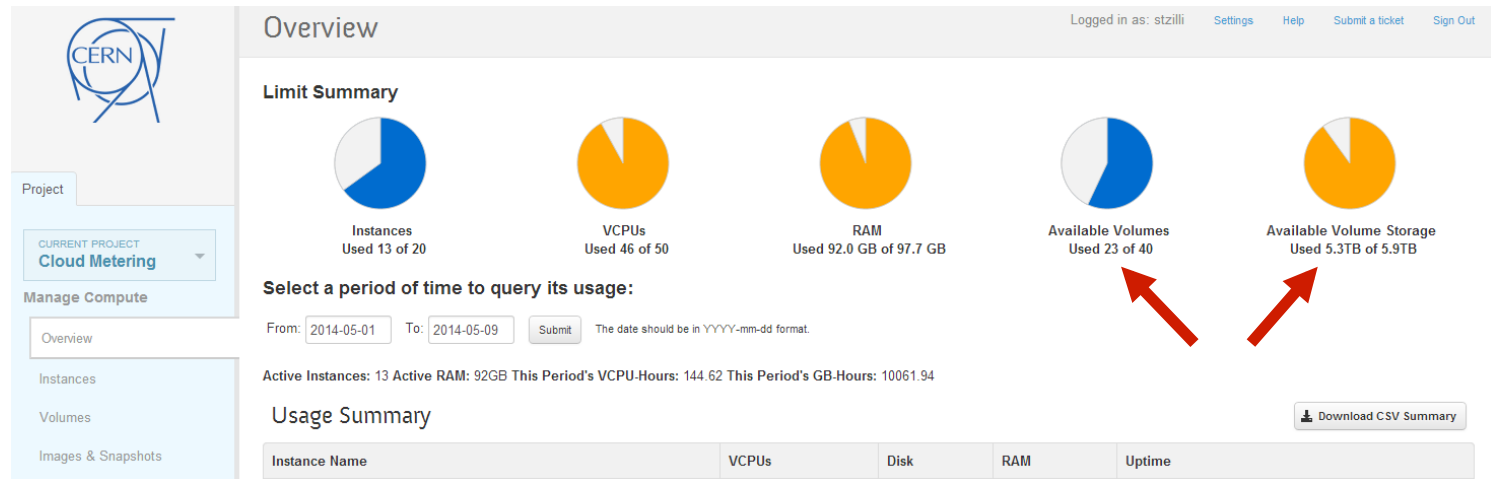
- Monitoring
  - QA validation using Tempest



# New Features (4/6)



- Horizon
  - Pie charts for volume usage



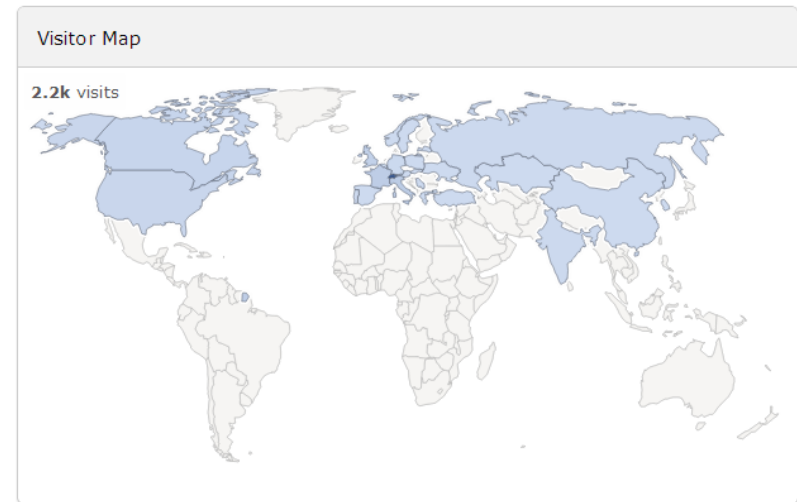


- PowerShell RC for Windows



## — Piwik

- Analytic software
- Usage of the Dashboard
- Statistic plots
  - Visits per region
  - Visits during last day
  - Visits during last month
  - ...






- General
  - Refactor configuration
  - Automate some common operations
  - Better management of hardware issues
  - Migration from CVI to OpenStack
- Nova
  - Feature parity between cell and “normal” deployments
  - Spice console
- Glance
  - Introduce API v2
  - Integrate Image Builder with Jenkins



- Cinder
  - Volumes for Windows
- Keystone
  - Add support to Kerberos
- Ceilometer
  - Alerting (for Heat)
- Monitoring
  - More scenarios for QA
  - Improve notifications in case of problems

- Upgrade to Icehouse
  - Maintain the project linked to upstream releases
  - Deployment delay:
    - Testing
    - CERN patches
  - Upgrade one component a time



- X.509 user certificate authentication
- SSO for Horizon



- Deploy Heat
  - Orchestrates deployment of clusters of VMs
  - Based on load metrics a cluster can grow or shrink
  - Better use of resources
- Deploy Neutron
  - Network as a Service
  - Not yet ready
  - Scalability problems
- Introduce sub-projects
  - Quota is assigned to a top-level project
  - Project managers can create sub-projects
  - Project managers can distribute quota to sub-projects
  - Not yet clear how it will be implemented



- Constantly growing
- Deployed over two sites
- Stable service
- No major issues
- Several projects going on
- Very active community

# QUESTIONS ?

[HTTPS://OPENSTACK.CERN.CH](https://openstack.cern.ch)

[HTTP://CERN.CH/GO/LWN8](http://cern.ch/go/LWN8) [DOC]

[HTTP://CERN.CH/GO/9KTL](http://cern.ch/go/9KTL) [BLOG]

Thank you !