



CERN private Cloud in a nutshell

CERN Cloud Infrastructure





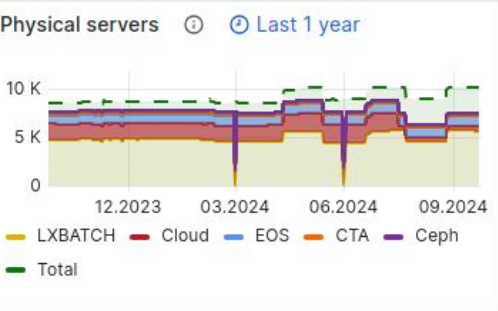
- Infrastructure as a Service
- Production since **July 2013**
- Running on **Redhat Enterprise Linux / AlmaLinux 9**
 - Based on Redhat Distribution of OpenStack (RDO)
- Meyrin and Preveessin Data Centres
- Currently running **Yoga+** release
 - Some services already in Zed release



Openstack services statistics


Users 3512	Projects 4867	Loadbalancers 433	Images 3503	Volumes 3772	Volumes siz 2.40 PB	File Shares 4847	File Shares : 2.62 PB	Object Store 649	Object Store 199 TB	
Servers Physical 10474 Physical in use 10179 Hypervisors 1851 Virtual 15950		Cores Physical 672 K Hypervisors 69.4 K Virtual 105 K			RAM Physical 3.01 PB Hypervisors 123 TB Virtual 250 TB			Batch Servers 5936 Cores 415743 RAM 1.63 PB		

Time series




CERN Cloud Infrastructure - now

IaaS+




IaaS

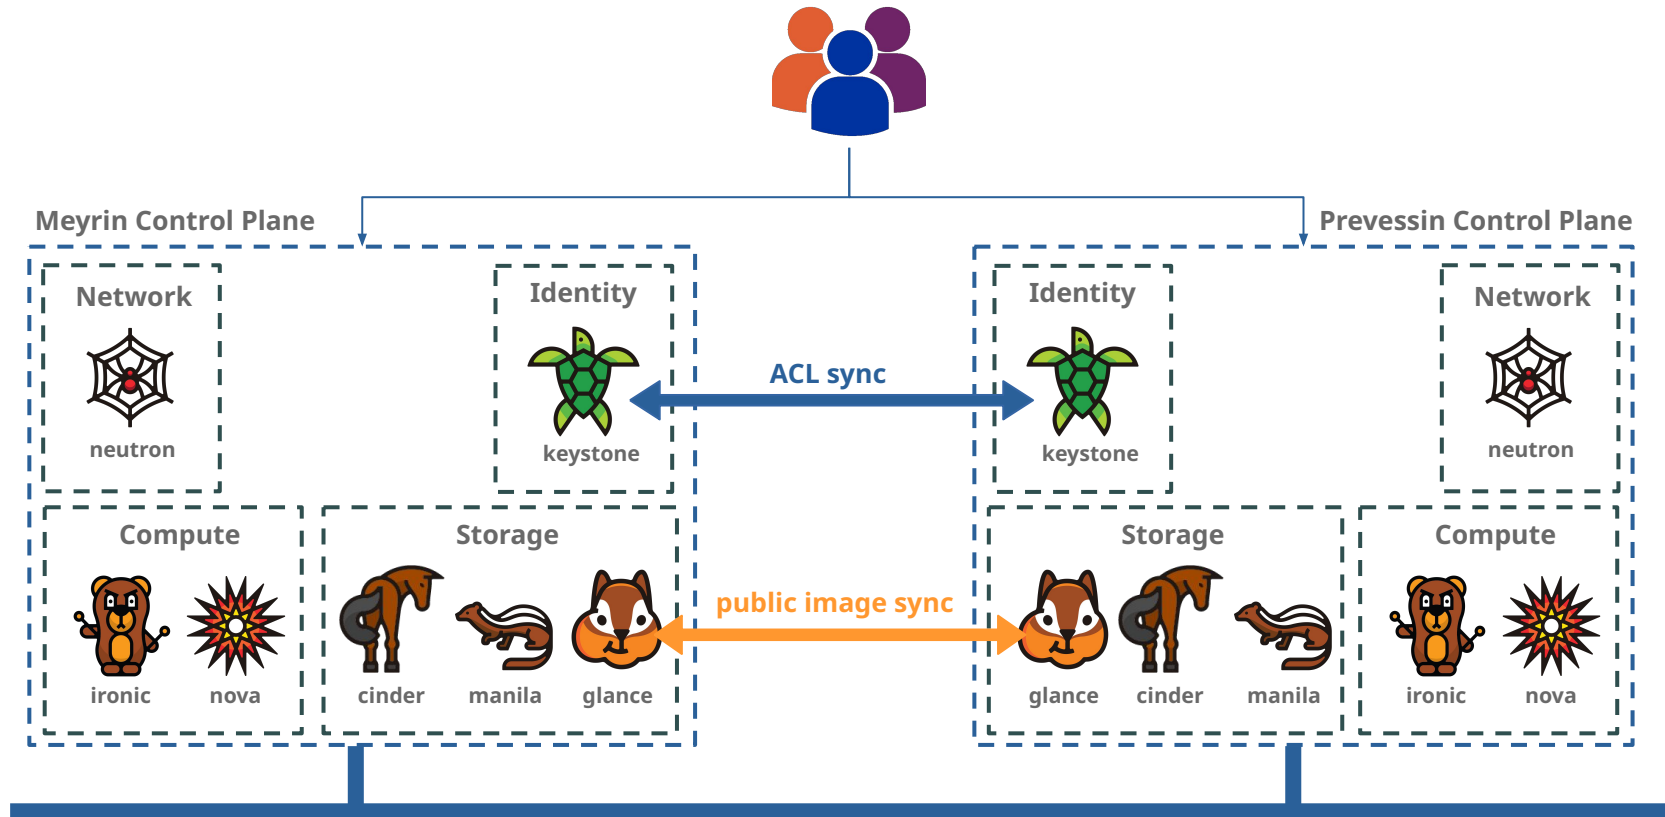


In MDC & PDC

Infra



Offer 2 datacentres under the same Cloud

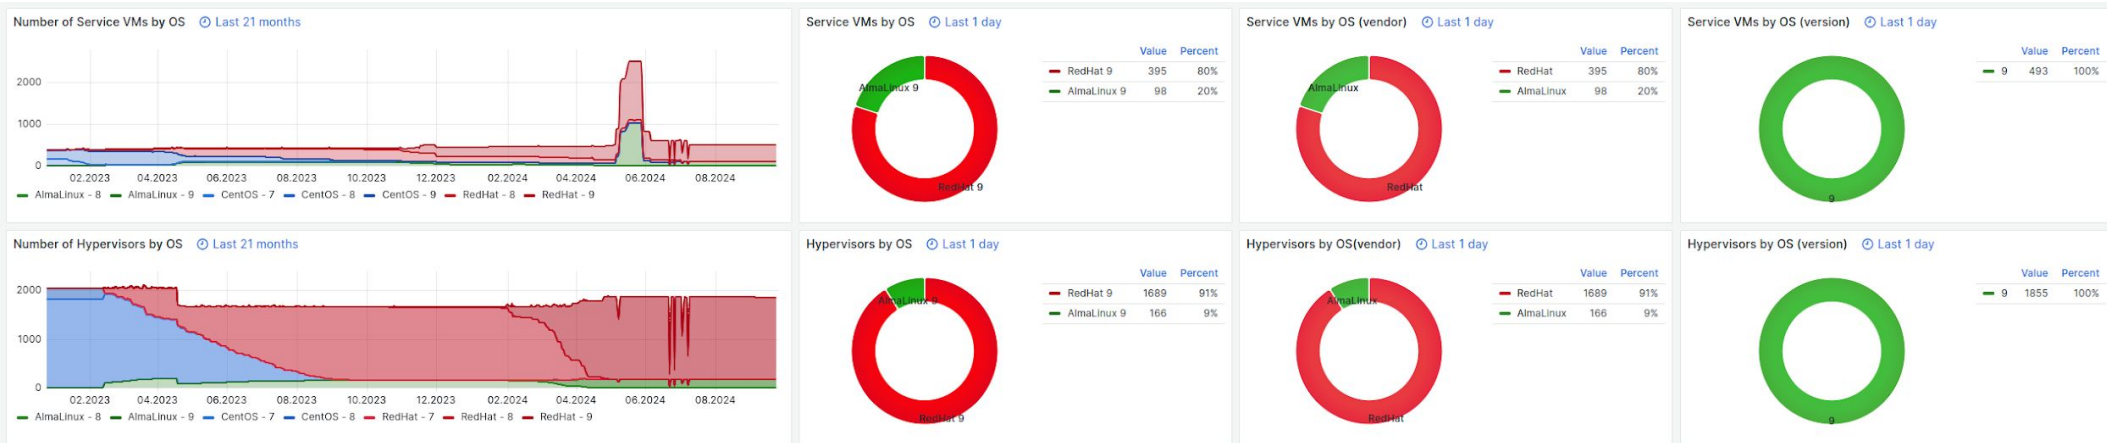


Differences between sites

Feature	Meyrin DC	Prevessin DC
OpenStack version	Yoga+	Yoga+
OS version	RHEL9 and ALMA9	RHEL9 and ALMA9
Availability Zones	3 Compute 3 Storage	1 Compute & Storage
Number of Cells	34	1
Cross Zone attachments	YES	NO
Anti-/Affinity Filters	Host	Host, Rack, Room
Networks	Provider	Provider & Private
SDN Features	Load Balancers	Security Groups Load Balancers Floating IPs
Nodes	6590	3860
Hypervisors	1728	101
Capacity	412TB	76TB
Capacity (on Diesel)	12TB	-
UPS expected lifetime	10min	5min

Service deployment

- From shared to “per microservice” architecture
- All deployed in VMs on our own infrastructure: “eat our own dogfood”
 - Bootstrap procedure and recovery methods
- Puppet managed running on RHEL/ALMA 9

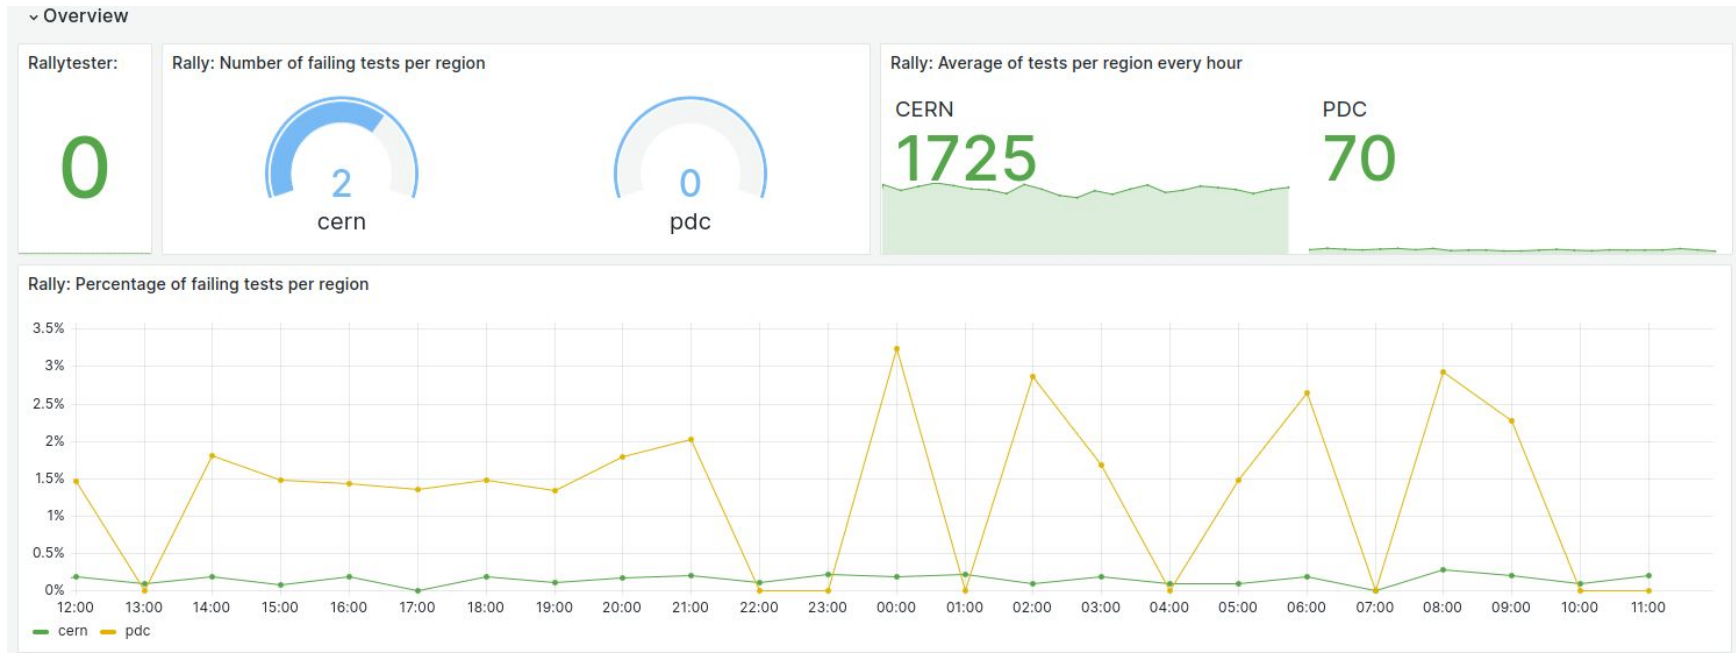


Service operations

- Meyrin Deployment upgraded since **July 2013**
- Per-service upgrade model (A/B or in place)
- Compute + Storage availability zones
- Huge investment on **automation**:
 - Delegate as much as possible administrative tasks (repair team, quota mgmt, end-user)
 - Detect and fix known issues
 - User communication
- Quite some big campaigns:
 - KVM consolidation, Spectre/Meltdown and L1TF, Cold Migration, Migration to 8/9, ...

Continuous probe the Cloud APIs

- Extensive use of automated probe system
- Focus on infrastructure wide issues



Discussion



- More info:
<https://techblog.web.cern.ch/techblog/>
- All our code is **open source** and available on:
<https://gitlab.cern.ch/cloud-infrastructure>