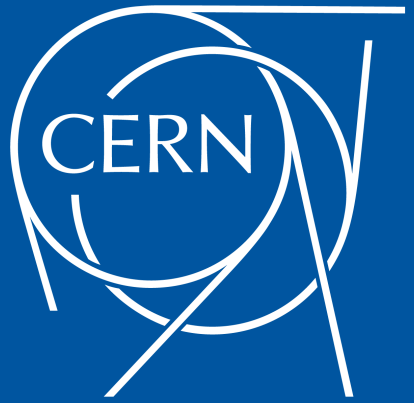
CERN

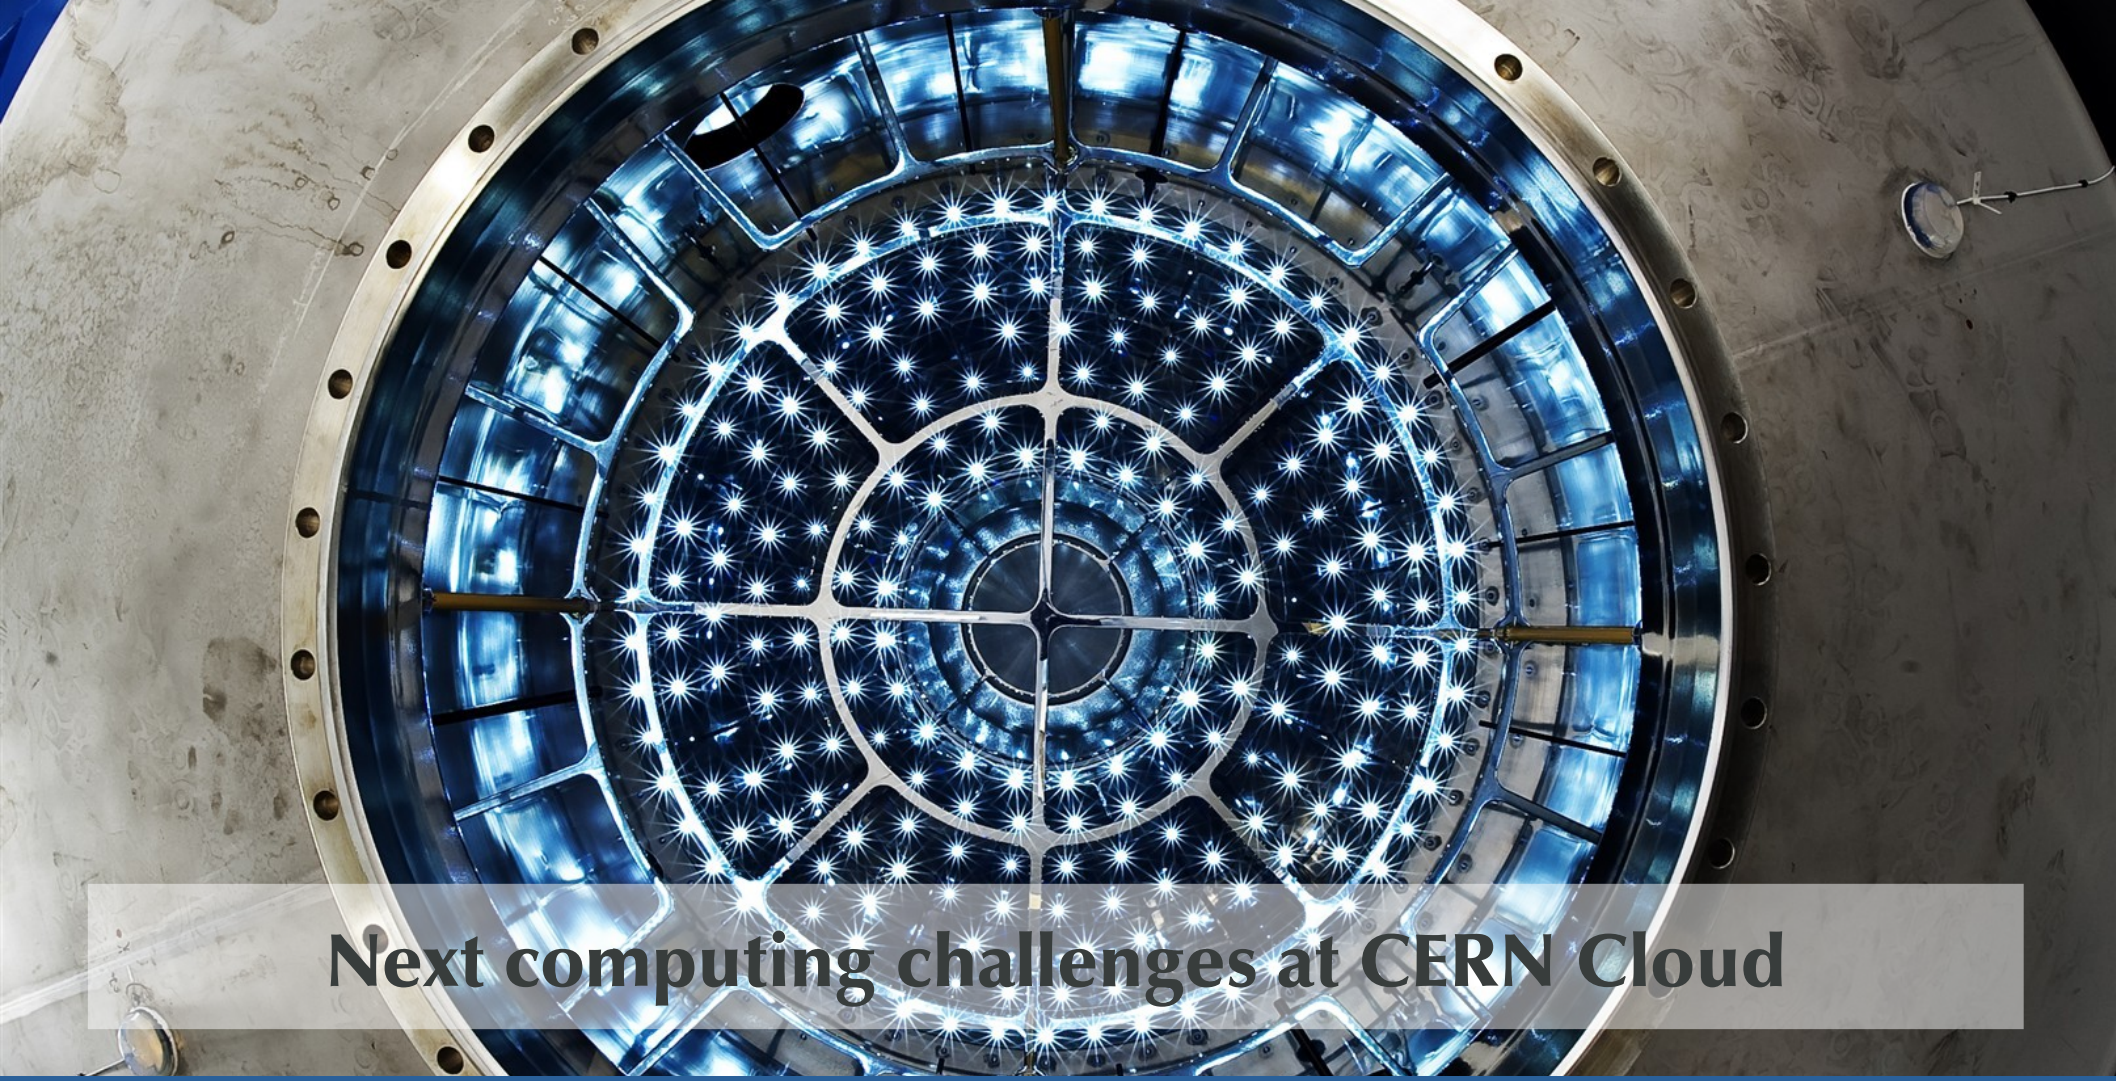

Next computing challenges at CERN Cloud

José Castro León
CERN Cloud Infrastructure

About me

- Technical leader of the CERN Cloud Service
- Joined CERN in 2010 to work into virtualization
- Core Team that built the Cloud Service in 2012
- In the OpenInfra community since 2012

Outline

- Introduction
- CERN Cloud
- Next Computing challenges
  - Approach
  - New Features:
    - GPU
    - SDN

European Organization for Nuclear Research

- World largest particle physics laboratory
- Founded in 1954
- 23 member states
- Fundamental research in physics

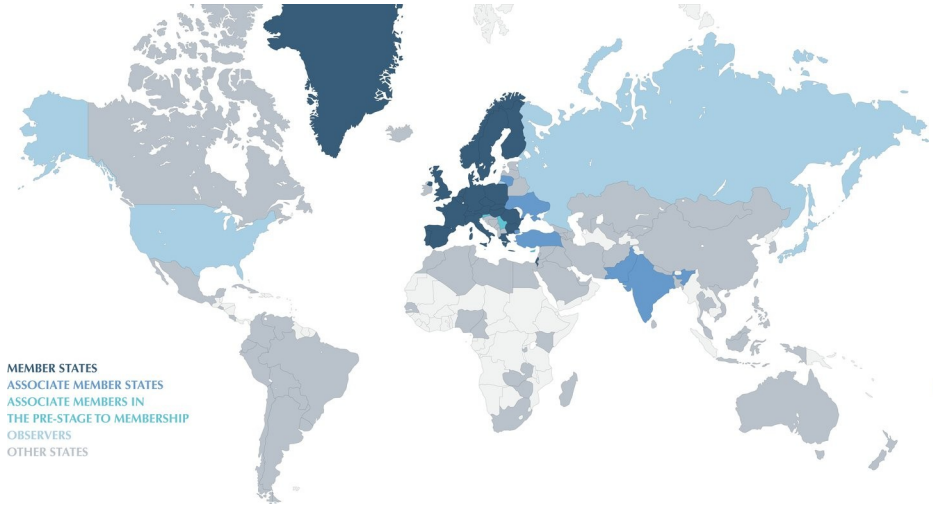

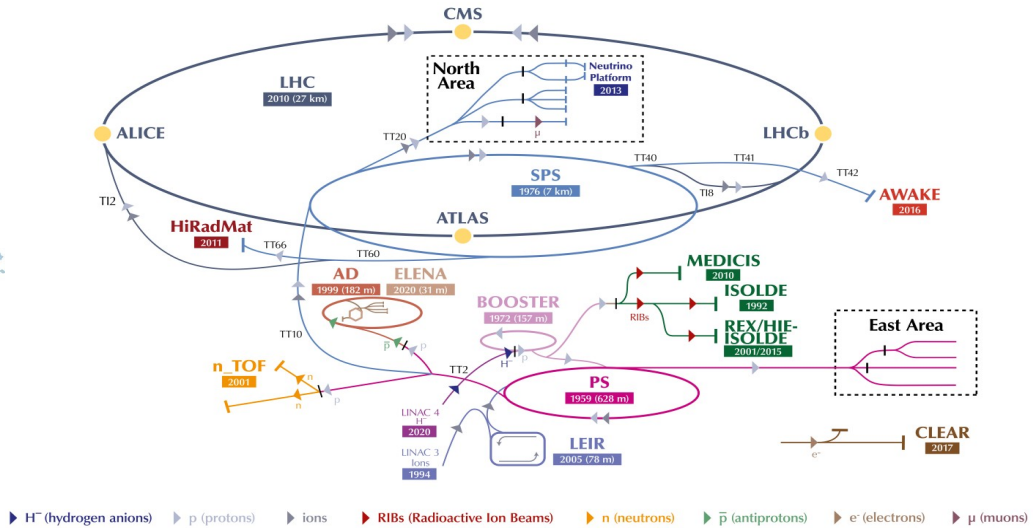

and RUN3 has just started …

CERN Cloud Infrastructure

- Infrastructure as a Service
- Production since **July 2013**
- **CentOS 7** based (adding CentOS Stream 8 soon)
  - Based on RDO, x86_64 architecture
- Geneva Computer centre (adding a new DC)
- Highly **scalable** architecture
  - 48 cells on 5 regions
- Currently running **Stein*** release
  - Some services already in Xena release

Openstack services statistics

Users	Projects	Loadbalancers	Images	Volumes	Volumes si...	File Shares	File Shares...	Object Stor...	Object Stor...
3382	4589	332	4764	7324	3.76 PB	5129	1.10 PB	477	67.7 TB

Servers				Cores			RAM			Batch		
Physical	Physical in use	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Physical	Hypervisors	Virtual	Servers	Cores	RAM
9052	8833	2013	13452	485 K	58.3 K	86.6 K	2.02 PB	379 TB	202 TB	5213	281453	1.07 PB

Time series

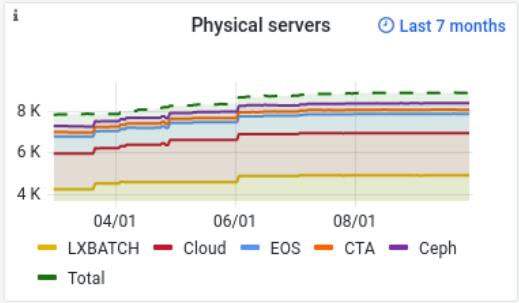

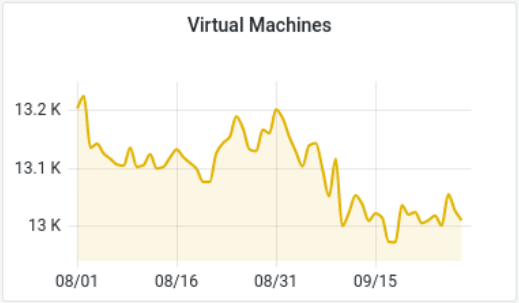

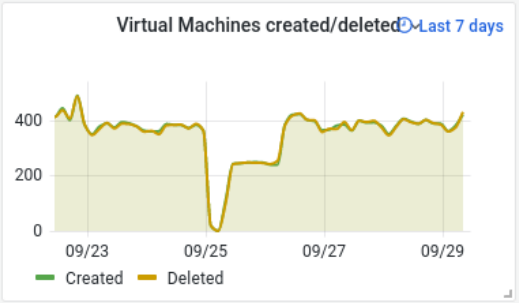

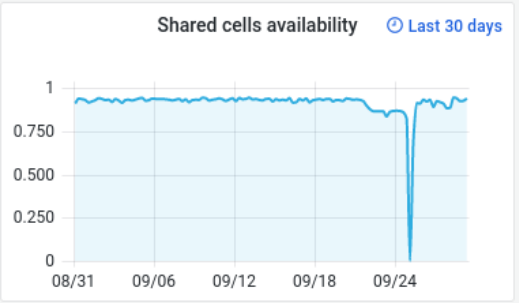

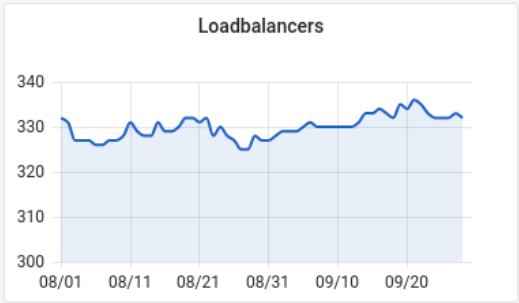

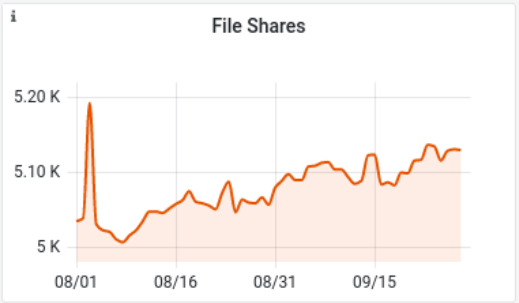

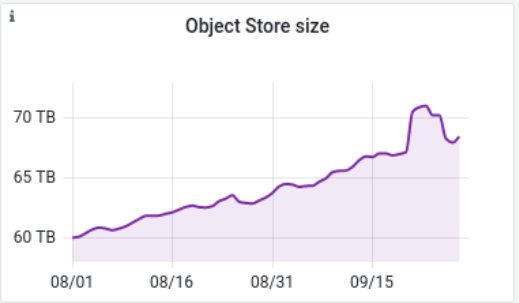

CERN Cloud Infrastructure

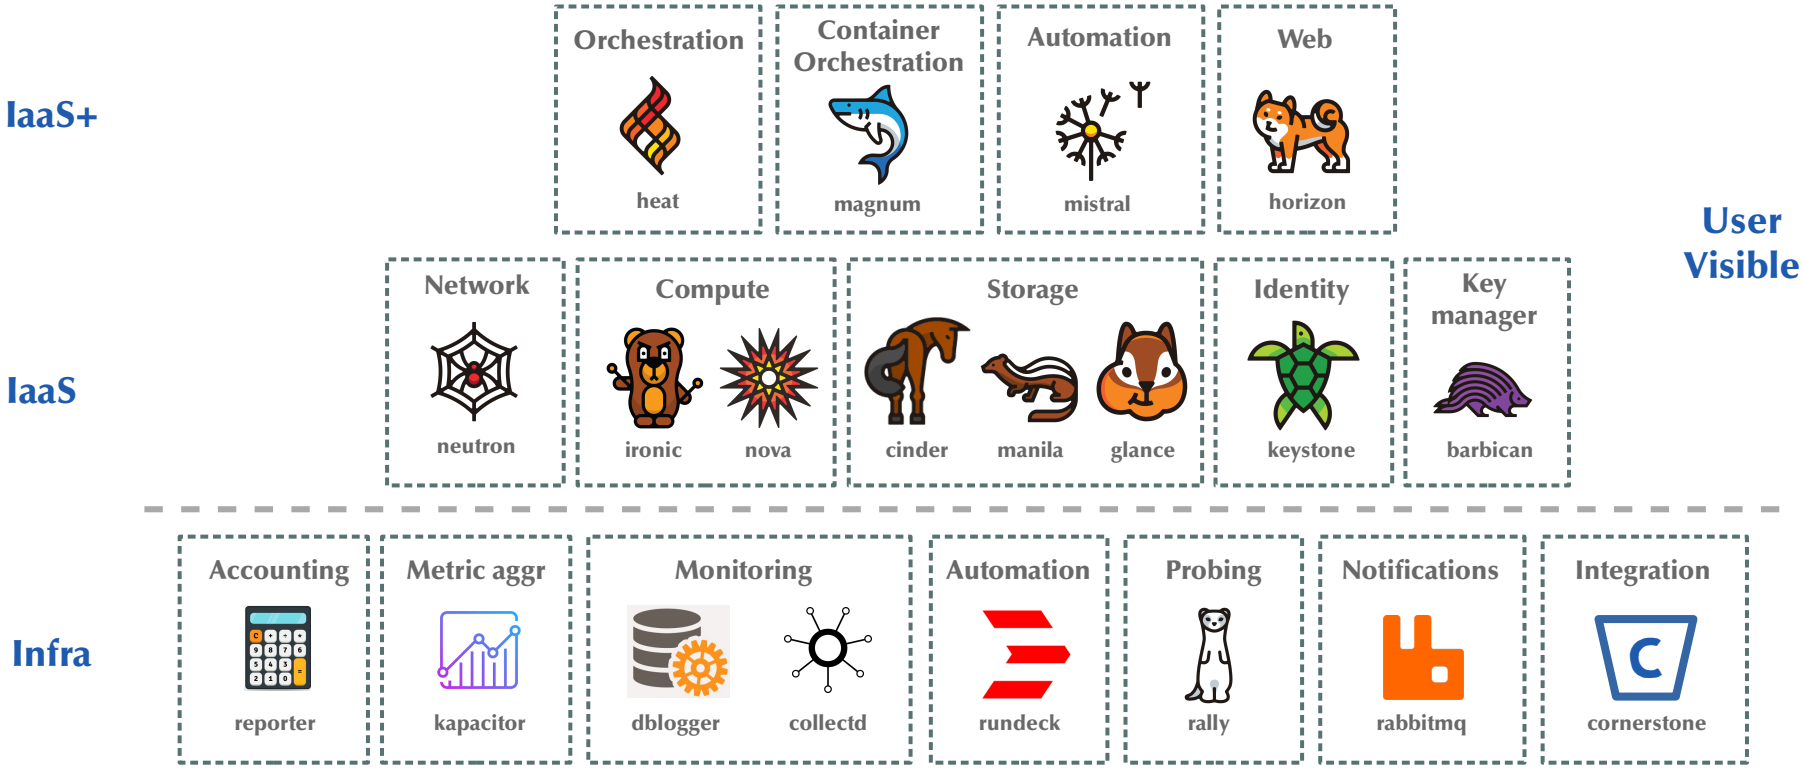

Huge computational challenge

- Continuously increasing computing needs
  - Current model is not enough
- Experiments
  - Increase data throughput to the DC
  - Exponentially increase of CPU resources
  - Increasing usage of GPUs
    - HW → SW → FPGAs → GPUs
- In a difficult context

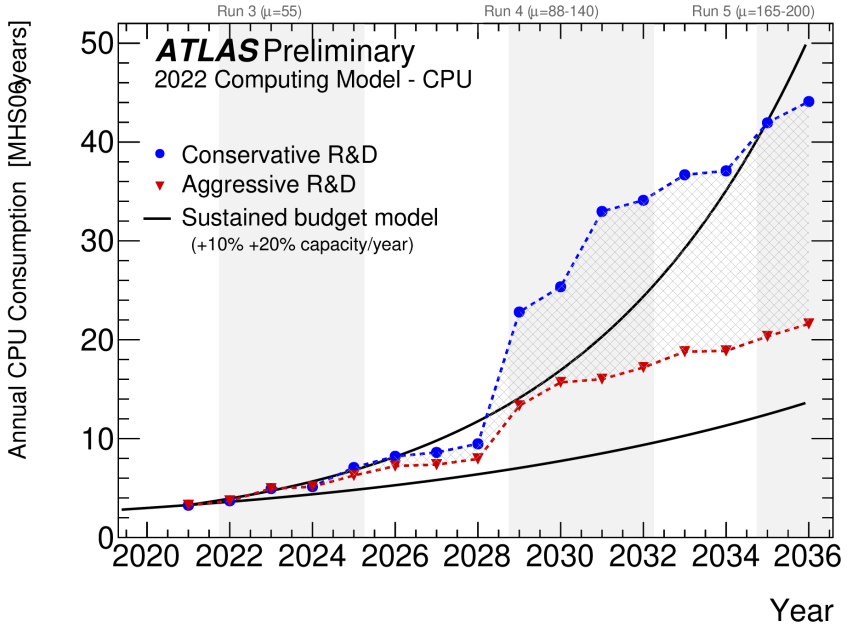

Threefold approach

- Increase RAW **capacity**
  - Current DC in Meyrin (3.5 MW)
  - New datacentre in Prevessin (4 - 12 MW)
- Boost **flexibility**
  - End user and operator perspective
- Enhance computing **performance**
  - Inclusion of heterogeneous resources
    - GPUs, ARM, … and any future architecture

New datacentre in Prevessin

- Currently under construction, delivery by end of 2023
- Provide **extra capacity** for the upcoming LHC and HL-LHC runs
  - 3 floors with up to 4 MW per floor (**12MW**)

- Green field deployment
  - **AvZs considered from the start**
  - Dedicated OS control plane and Ceph Clusters
  - May change **hypervisor disk layout**
  - Introduce **Software-Defined Networking**

(Credit: + IMGS – Rocco Valantines)

Disaster Recovery

- Additional datacentre will be used for **computing** and also IT **services**
  - Focus on **critical** IT services for the organization

- Expose it as an **additional fully independent** region
  - Extra overhead to manage it

- Looking at **replication**, **multi-site** setups
  - Review building blocks available for users

Flexibility boost

- Close the **gap** with upstream
  - Currently under a big cold migration campaign
  - OpenStack (Stein* => Zed) + OS (el7 => el9)
  - Benefit of the **latest** features of the code
- Double down on **monitoring**, automation and probing
  - Handle 2 completely independent environments
- Remove the **boundaries** with network physical topology
  - Add Software Defined Networking

… be closer to upstream

Find the needle in the haystack

- Threshold based alarming on extreme cases
- **Anomaly detection** to find misbehaving nodes

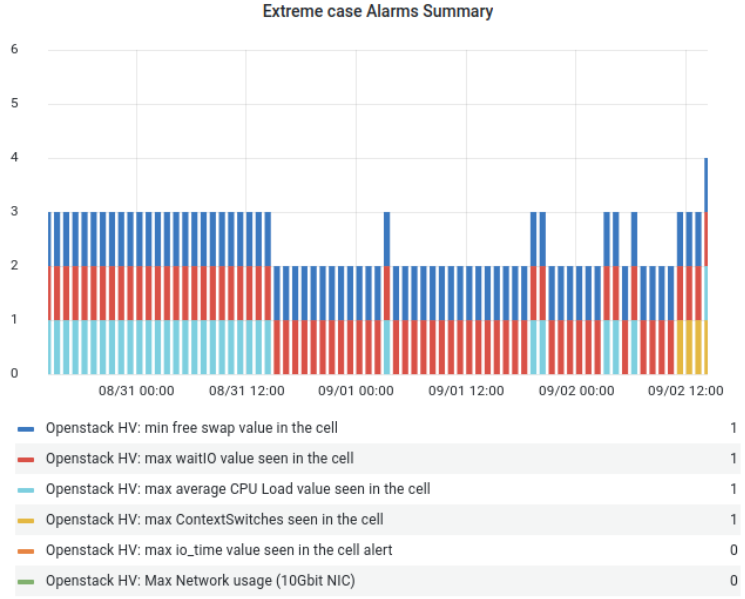

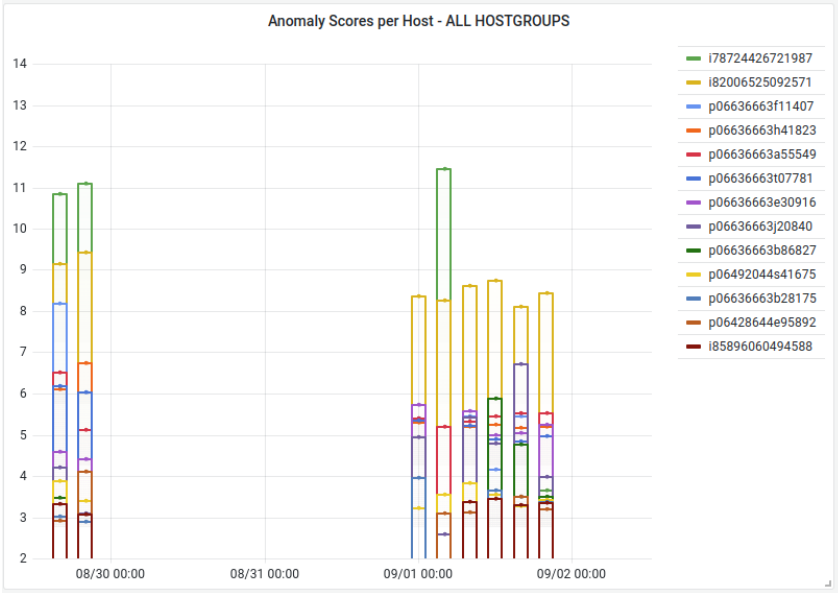

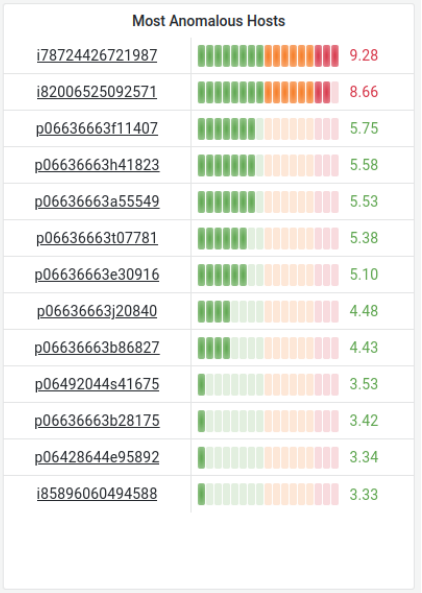

Continuous probing the Cloud

- Use Rally as automated probe system
- Focus on infrastructure wide issues

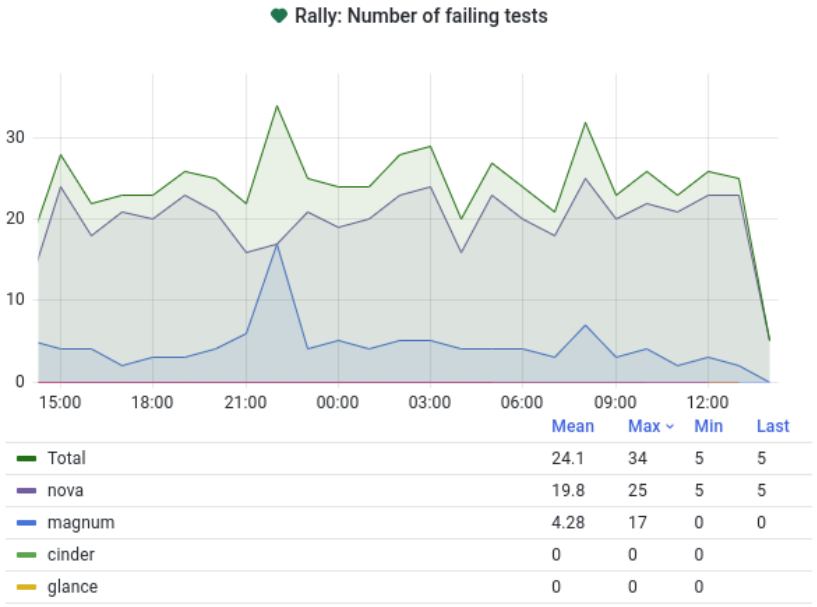

Passing % in time frame per availability zone

availability zone	attach-volume	boot-linux	boot-linux-with-dns	ping-linux	reboot-linux	live-migrate-linux	cold-migrate-linux
cern-geneva-a	100%	100%	100%	100%	100%	95%	93%
cern-geneva-b	100%	100%	100%	100%	100%	95%	91%
cern-geneva-c	99%	100%	100%	99%	99%	100%	91%
gva-critical	98%	100%	100%	100%	100%	98%	87%

Global actions: Passing % in time frame

deployment	authenticate	boot-from-snapshot-linux	boot-from-volume-linux	create-and-delete-image	list-images
global	100%	100%	100%	100%	100%

Cinder actions: Passing % in time frame

deployment	create-and-delete-snapshot	create-and-delete-volume	create-and-extend-volume	list-volumes
cinder	100%	100%	100%	100%

Manila actions: Passing % in time frame

deployment	create-share-and-allow-and-deny-acces	create-share-and-delete	create-share-and-extend	create-share-and-shrink
manila	100%	100%	100%	100%

Software Defined Networking

- Current networking model really tightened to the infrastructure
- Several technologies evaluated or under evaluation
  - OpenDaylight, OpenContrail/Tungsten Fabric, OVN
  - Currently offering LBaaS
- Fully fledged SDN deployment on new DC
  - Virtual Networks, Floating IPs, LBaaS ...
  - Provide maximum flexibility to end users

... be closer to upstream

Enhance computing performance

- Adding more computing resources (performance per watt per dollar)
- Need to consistently provision, monitor and configure them
  - May trigger changes in the whole stack
  - ARM recently added into the portfolio
    - Users can start rebuilding / validating their frameworks
  - Received batch of Nvidia A100 GPUs
    - Currently added into the GPU offering

Trying to efficiently provision and use GPUs

- Many different use cases require access to GPUs with different utilization
  - deep learning, inference, analysis, simulations, GIS, mechanical, …
- 4 different Nvidia models available (T4, V100, V100s and A100)
- Available as vGPU and pci-passthrough (currently looking at MIG)
- Really scarce resources, preparing a lease model
  - Missing quota handling, we **really** need your help here

Conclusion

- Quite some challenges ahead of us
  - Catch up, scale up, provision more and also heterogeneous
- Building the foundation for the years to come
  - The path is already laid out, we just need to walk through it
- Joint effort with the community
  - Quotas for dedicated resources (flavors, custom resources)

… and this won't be possible without…

CERN Cloud and Linux teams

Thank you

More info:

https://computing-blog.web.cern.ch/

All our **open source** code is available on:

https://gitlab.cern.ch/cloud-infrastructure

Thanks again to my colleagues

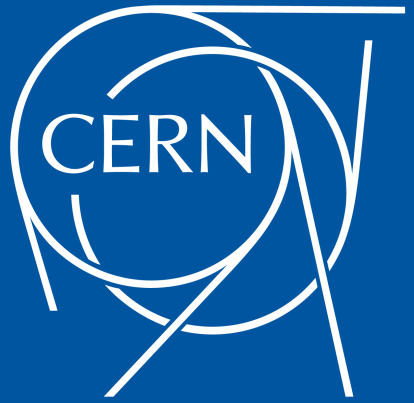
CERN

BACKUP SLIDES

Largest machine on Earth: 27 Km

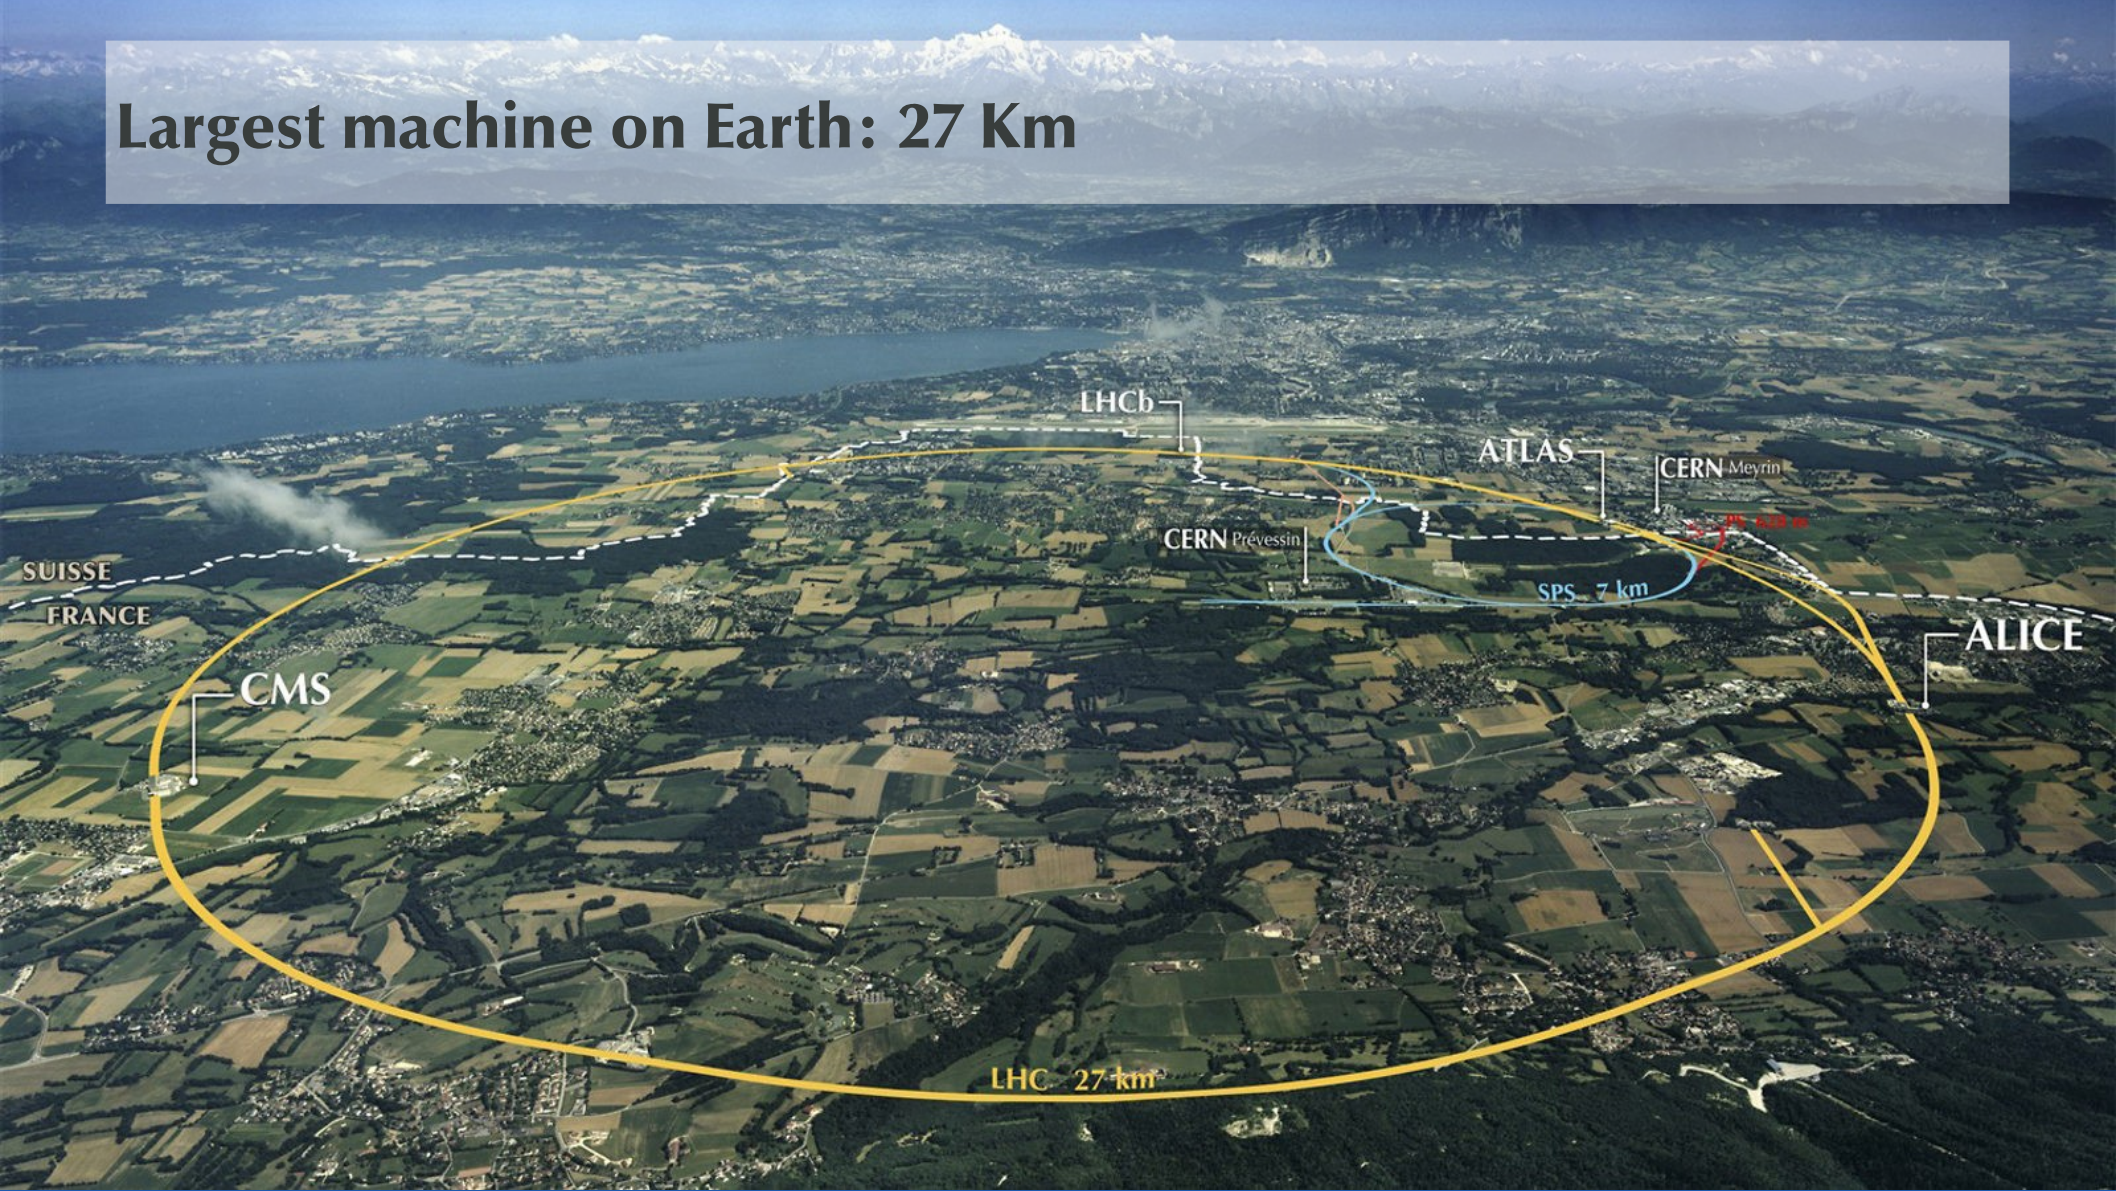

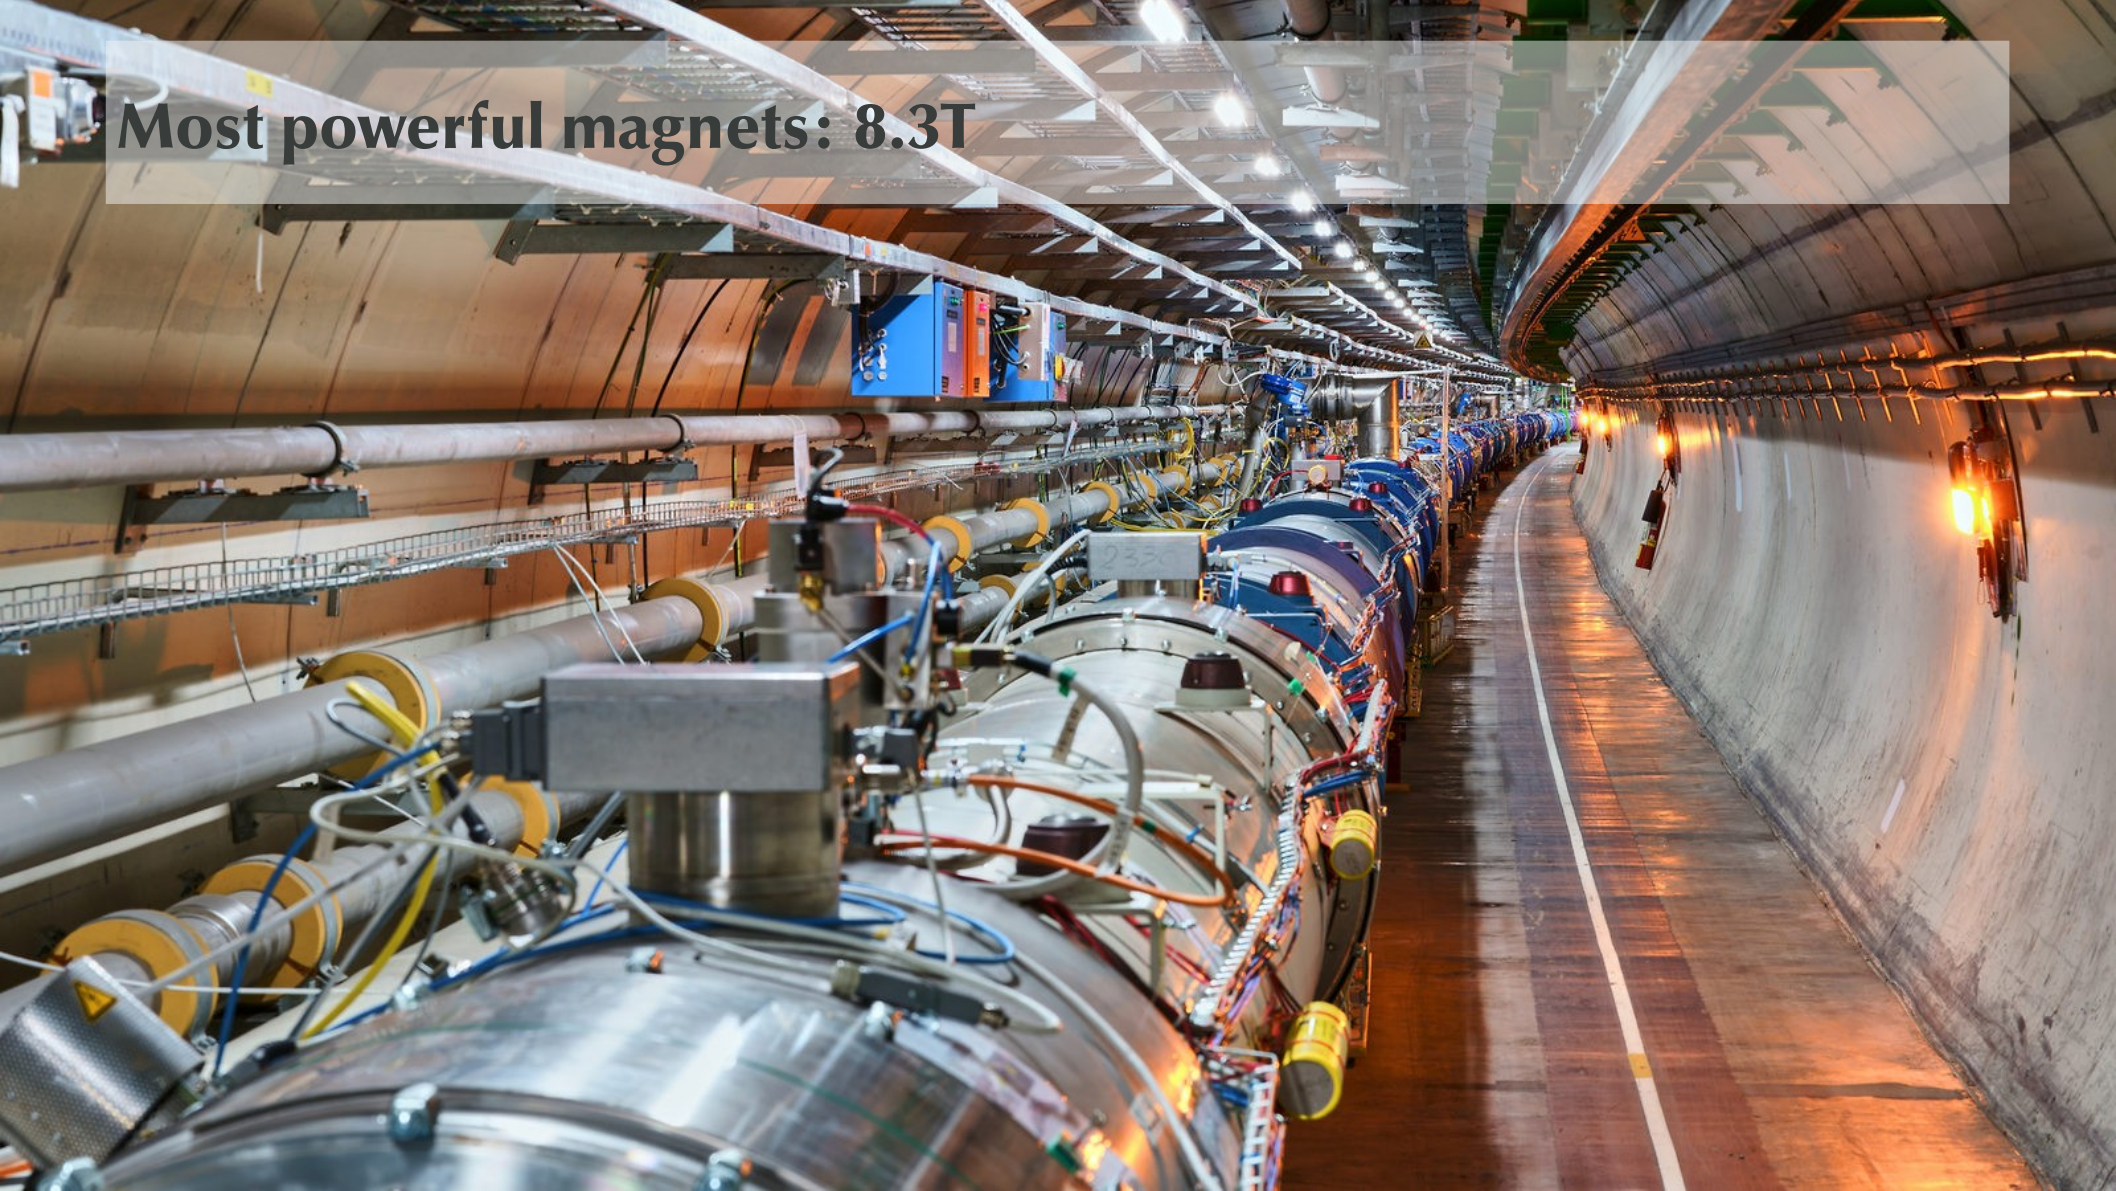
Most powerful magnets: 8.3T

Highest vacuum: 10 times less than on the moon


Coldest temperature: -271°C

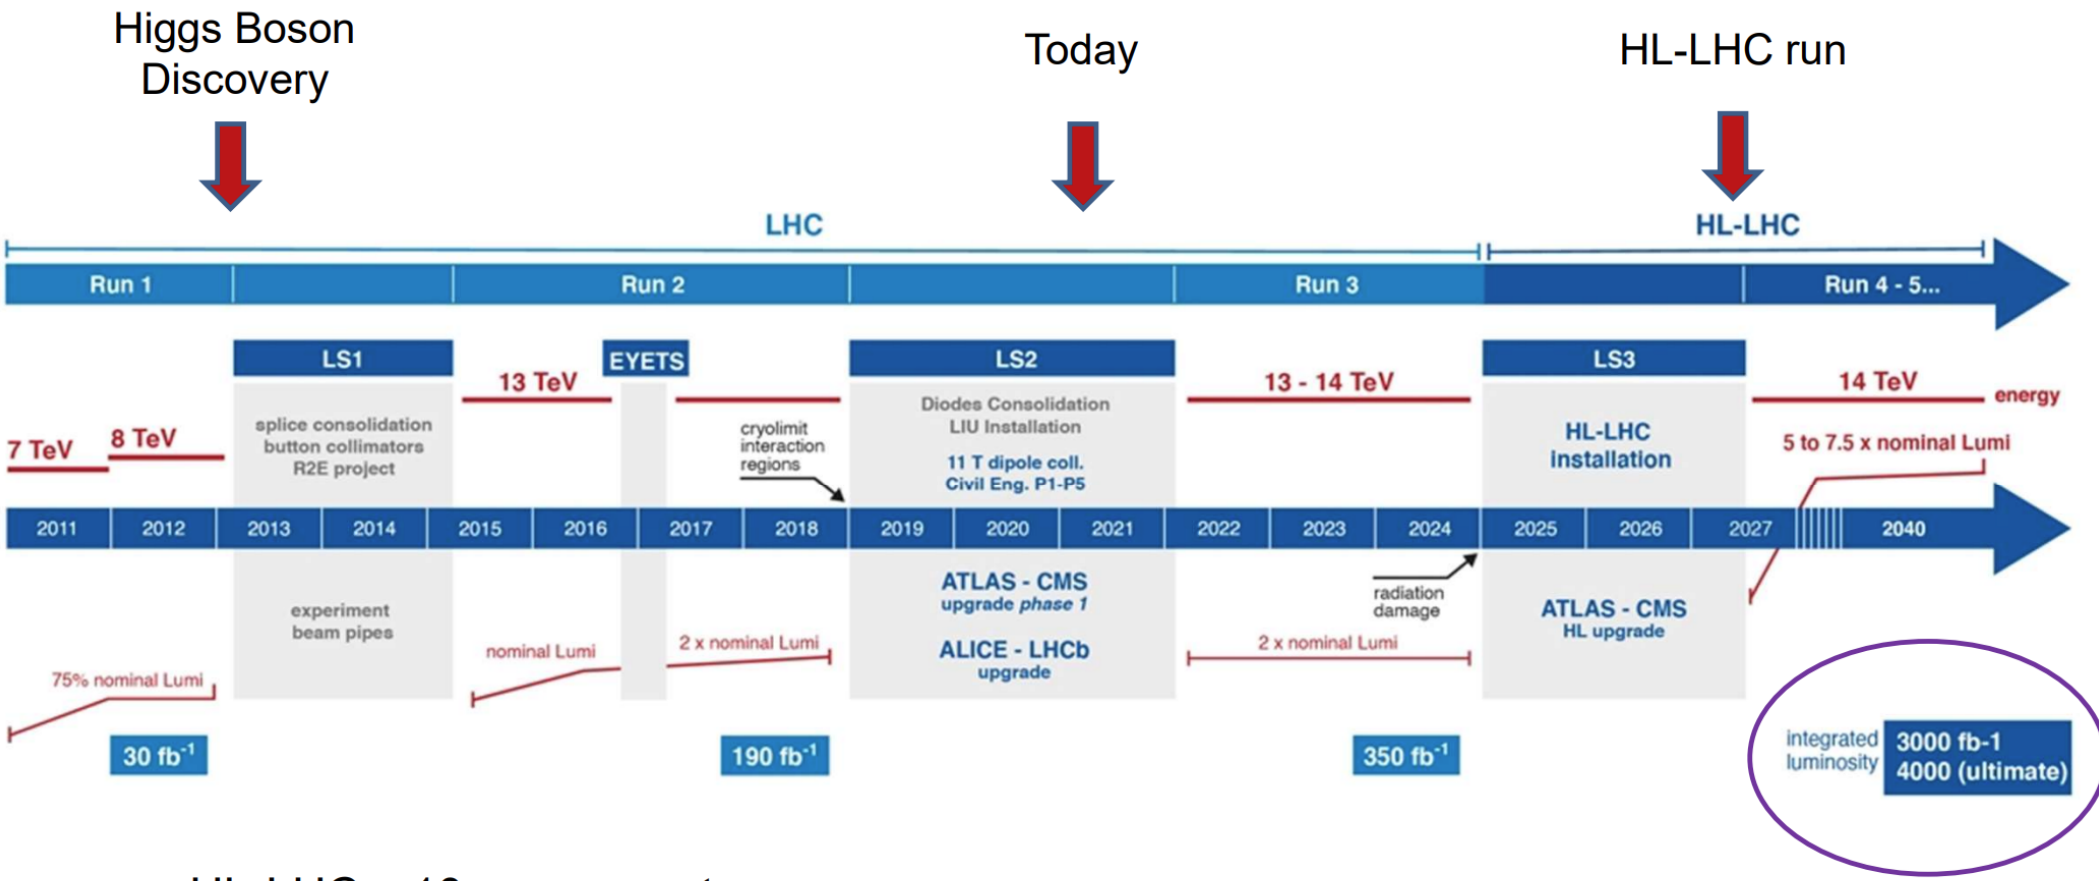

HL-LHC: x10 more events.
Bigger and more complex events (x5)

An unprecedented challenge!

Baremetal provisioning

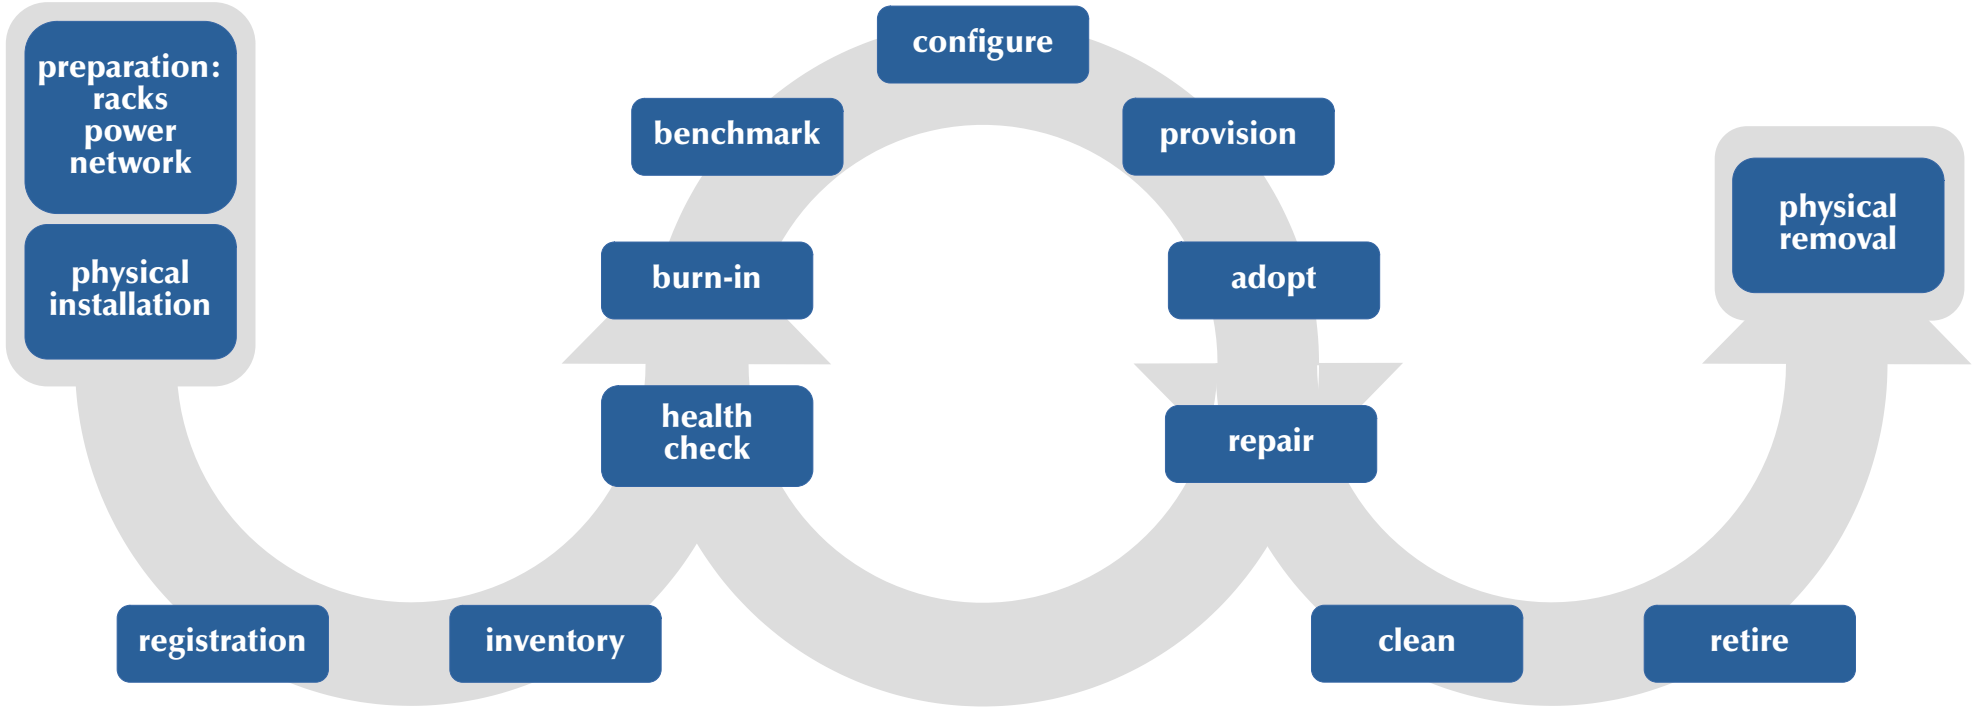

… now we also manage ARM servers

Ironic Service setup and status

Total Nodes	Active Nodes	Total kCores	Total RAM [...]	Error Nodes	Orphans	Service Erro...	Service Warn	Service Info	API Errors	API Warn	API Info
9040	8826	485	1924	1	0	300	2133	1335	130	0	0

Maintenance Modes

	current	percentage
false	9 K	100%
true	34	0%

Power States

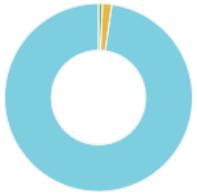

	current	percentage
Value	61	1%
power off	145	2%
power on	9 K	98%

Unusual Provisioning States

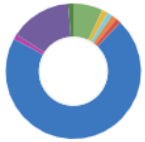

	current	percentag
clean failed	5.8	7
clean wait	1.3	2
cleaning	1	1
deploy failed	1	1
deploying	1	1

API Load distribution

	current	percentage
cci-ironic-t0135542.cern.ch	213	4%
cci-ironic-t0136653.cern.ch	213	4%
cci-ironic-t0595719.cern.ch	212	4%
cci-ironic-t0818444.cern.ch	421	7%
cci-ironic-t2267845.cern.ch	213	4%

Response codes

	current	percentage
Client Errors	4	0%
Information	0	0%
Redirection	0	0%
Server Errors	0	0%
Success	8 K	100%

Ironic Rally Tests

Last 24 hours

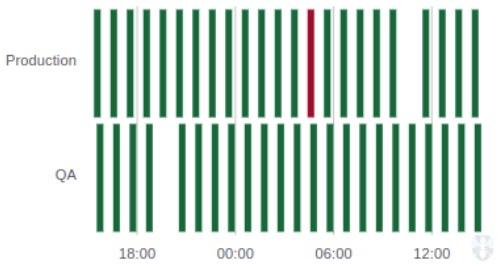

Bare Metal Events

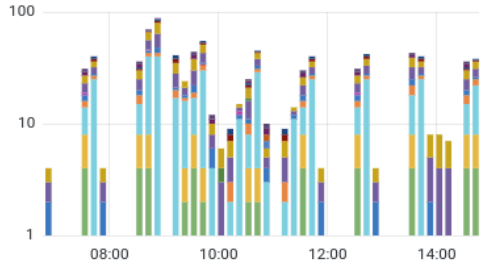

API calls per minute

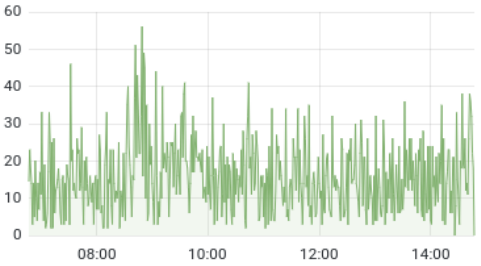

Log Messages per minute