



# Contributing for the Long Haul


(An 'OpenStack at CERN' FOSS case study)

**Arne Wiebalck**

Head of Compute and Devices Group, CERN IT

CERN OSPO Seminar: Making the most of your open source contributions, 21 Feb 2025





## IT: 2.x Data Centers ...

- Meyrin: 3.7MW (1970s)
- Preessin: 4..12MW (since Feb 2024)
- Remote hosting: 2.5MW (2013-2019)
- P8 Containers: 0.5MW (2019-2024)

## ... operating ...

- ~12k servers
- ~1EB of disk storage
- ~1EB of tape
- >100s of services

## ... for three main tasks!


- Initial reconstruction
- Long-term storage
- World-wide distribution (WLCG)

# Why did CERN set up a private cloud in 2012?




## → Situation in 2012

- EU projects finished ← **person power & dev effort dropped!**
- LHC Computing CPU & Storage **needs increasing massively**
- Other deployments have surpassed CERN's scale
- LS1 about to start ← **opportunity for change!**
- “Agile Infrastructure” Project: Config' Mgmt, Monitoring, **IaaS**



## → CERN's working cycle is determined by the LHC

- Multi-year **RUNs**
- Multi-year **Long Shutdowns (LSX)**



# How did CERN set up an Agile Infrastructure?



## TOOLS

### → Resource Provisioning (IaaS)

- Based on OpenStack



### → Configuration Management

- Based on Puppet



### → Centralised Monitoring

- Based on Lemon (sensor)
- ELK stack



## POLICIES

### → Virtualized services!

- Within reason
- Exceptions are expensive!

Faster provisioning  
Increase resource efficiency  
Simplify infra management

### → Puppet managed!

Profit from others  
Auto-document  
Ensure reproducibility

### → Centrally monitored!

Efficient debugging  
Integrated alarming  
Facilitate accounting

# What is OpenStack?

## → Cloud Computing Platform


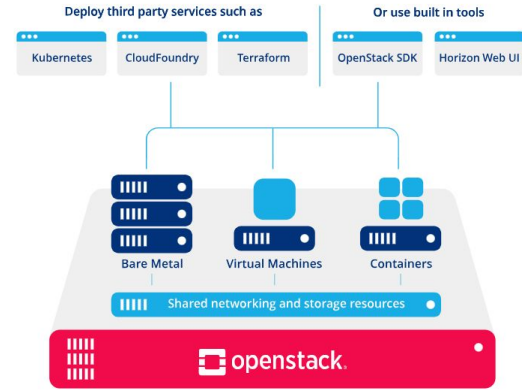
- Manages compute resources and offers them to users
- Primarily developed for virtual machines
- Over time: storage, load balancers, physical servers, ...

## → Started in 2010 by Rackspace and NASA

- Managed by non-profit foundations since 2012

## → One of the most active FOSS projects

- 400k commits from ~9'000 individual contributors!





# Why did CERN build its private cloud on FOSS?



## → Adaptability feasible ... and required!

- Brownfield deployment, many existing services
- Buying compute with fixed procedures


## → CERN's role as Tier-0 site in WLCG

- Sometimes “pave the way” for other sites

## → FOSS aligns with CERN's goals

- “Giving back” ← societal impact!
- Talent management ← member states!

## → (Helps with some data governance questions)




# How do we interact with the OpenStack community?



## → We contribute code / docs!

- We proposed and added features we needed
- From **March 14, 2012 ... to Jan 23, 2025**
- 140'000 LOC, >1000 commits, 38 contributors

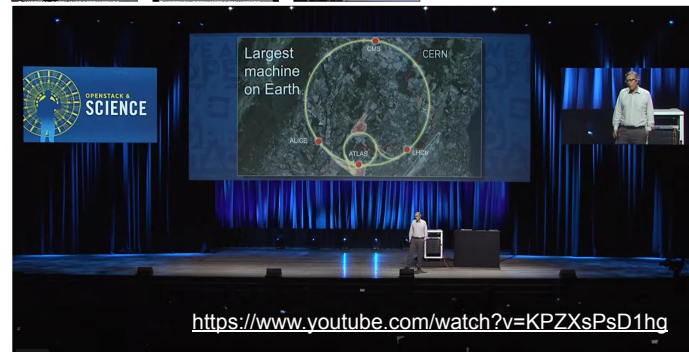


Commits by Module		
#	Module	Commits
1	magnum	340
2	nova	107
3	keystone	103
4	ironic	65
5	ironic-python-agent	59
6	releases	40
7	python-magnumclient	30
8	heat-translator	29
9	rally-openstack	23
10	rally	22

<https://www.stackalytics.io/?release=all&metric=filed-bugs&company=cern>

## → We share *what we do* with it!

- Presentations at summits, user meetings, ...  
"By contributing to open source you help us find the lost 96% of the universe!"
- <https://www.openstack.org/videos/search?search=cern>



<https://www.youtube.com/watch?v=KPZXsPsD1hg>

# How do we interact with the OpenStack community?



→ We share *how* we do it!

→ We attend & organise user meetings!

**CERN TechBlog**  
European Organization for Nuclear Research

## Adoption of in-production physical nodes into Ironic and Nova

Posted on December 16, 2020 | 8 minutes (1639 words) | Arne Wiebalck


When Ironic is chosen to manage bare metal nodes, it is not unusual that there is already an existing in-production fleet of physical servers. This may leave the operator in the sub-optimal situation where tools and workflows need to handle Ironic and pre-Ironic systems. To address this, Ironic supports an “adoption” feature: adoption allows to add nodes which Ironic should regard as in-use, and they can therefore take a slightly different path through the Ironic state machine. While this helps with direct or stand-

<https://techblog.web.cern.ch/techblog/>



**OpenStack DAYS**  
EMBORE

- OpenInfra Day Sweden**  
May 7, 2024  
Volvohallen, Gothenburg, Sweden  
→ Register Now | CPE
- OpenInfra Day Germany x Sovereign Cloud Stack Summit**  
May 14 & 15, 2024  
Villa Elisabeth, Berlin, Germany  
→ Register Now | CPE | Sponsor
- OpenInfra Day Turkey**  
May 20, 2024  
Albert Long Hall Cultural Center, Bogazici University, Istanbul, Turkey  
→ Register Now | CPE | Sponsor
- OpenInfra Day France**  
May 2024  
→ More info coming soon
- OpenInfra Meetup: Switzerland**  
June 6, 2024  
CERN, Geneva, Switzerland  
→ Register Now





# How do we interact with the OpenStack community?



## → We take on community roles!

- Contributor
- Core Reviewer
- Project Team Lead
- User Committee Member
- Technical Committee Member
- Lead of Special Interest Group
- OpenInfra Foundation Board Member

## → We are a Foundation Member

- This shows commitment
- Helps the community & the project



## → We came with the right attitude

- It's a give & take!
- It takes time and effort!

# How did this interaction pay off?



→ We have **established trust & standing** in the community

- CERN's issues are handled ← pretty prompt community support
- CERN's use cases are considered ← helps shaping direction / control
- CERN's deployment is a reference ← reputational benefit for CERN
- CERN's contributors have a professional network ← careers

**Our collaboration with the open source community allowed us to have a platform *for our needs long term.***

Bonus: An **enormous amount of fun and inspiration** from working with so many talented people when furthering CERN's use cases!



# Additional resources



## 10 years of OpenStack at CERN:

<https://www.openstack.org/videos/summits/virtual/10-years-of-OpenStack-at-CERN-From-0-to-300k-cores>

## OpenStack in Production: Mysteries, Challenges, and Achievements:

<https://superuser.openinfra.dev/articles/openstack-production-cern-lightning-talk/>

## CERN Tech Blog:

<https://techblog.web.cern.ch/techblog/>

## CERN Summit videos:

<https://www.openstack.org/videos/search?search=cern>







[www.cern.ch](http://www.cern.ch)