





How to manage **10'000+** servers: Automation in CERN IT

Arne Wiebalck / Benoit Clement
Computing Engineer

CERN Open Days 2019

What is the role of the CERN IT data center?


The heart of the

World Wide LHC Computing Grid (WLCG)

- Initial analysis for data quality assurance
- Permanent storage of all LHC data
- Distribution of the LHC data in the WLCG

Supports also ...

- ... the non-LHC experiments
- ... services for the accelerator complex
- ... infrastructure & administrative services



Where are all the 10'000+ servers?

Distributed over **three data centres!**



Meyrin (CH)

~12800 servers



Budapest (HU)

~2200 servers



LHCb Point-8 (FR)

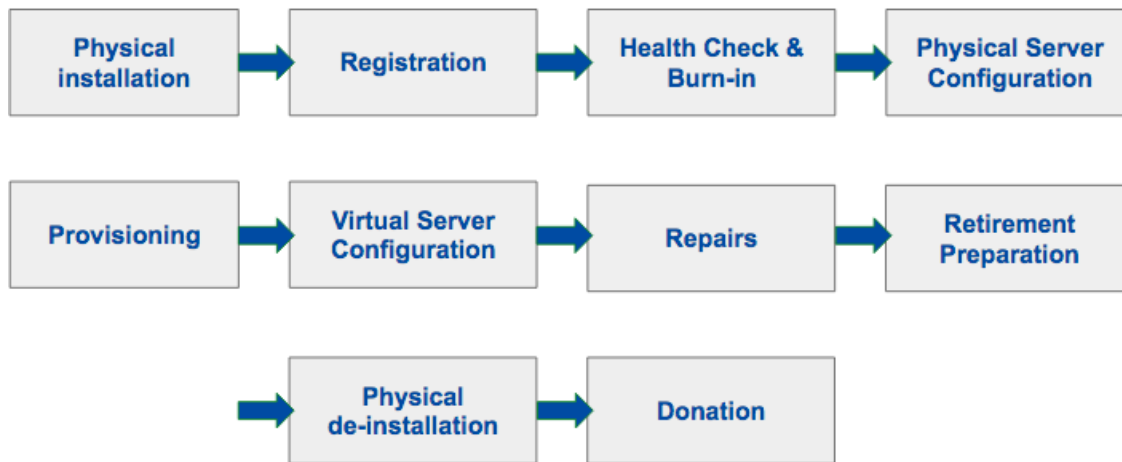
~800 servers



Why do we need automation?

Managing 10'000+ servers is challenging! Servers need to be ...

- ... purchased
- ... rack-prepared
- ... installed & cabled
- ... registered
- ... configured
- ... provisioned
- ... monitored & repaired
- ... eventually retired




Performing these tasks **continuously** and **at scale** is hard!

Purchasing already supports automation!

We buy 100s or even **1000s identical ones** in one go!

This helps with:

- **Physical installation**
 - same dimensions, component placement, connectors, power consumption, ...
- **Testing & initial burn-in**
 - easier to find outliers



Example of a
“slow” CPU



The **same servers** may also
have the **same issues!**

Mitigated by buying large
chunks from different
suppliers.

How are servers physically installed?

Rack preparation and physical installation is manual, but:

Nodes **register themselves** with the data centre infrastructure!

Upon first boot ...

- ... the nodes sent out a discovery request
- ... get access to the network
- ... boot into an initial system which automatically
 - checks the inventory
 - benchmarks CPUs, disks, RAM, network



Hello
my name is

*new server in the
CERN data centre*

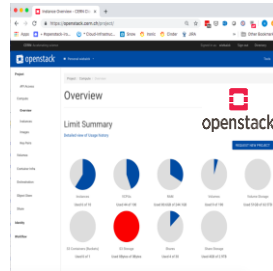

How are the resources provisioned to users?

“Users” == physicists or engineers which need servers for their work.

Via an additional abstraction layer: **virtual servers!**

Virtual servers are basically programs that pretend to be a “normal” server.

Since they are software programs, they can be started quickly on-demand.

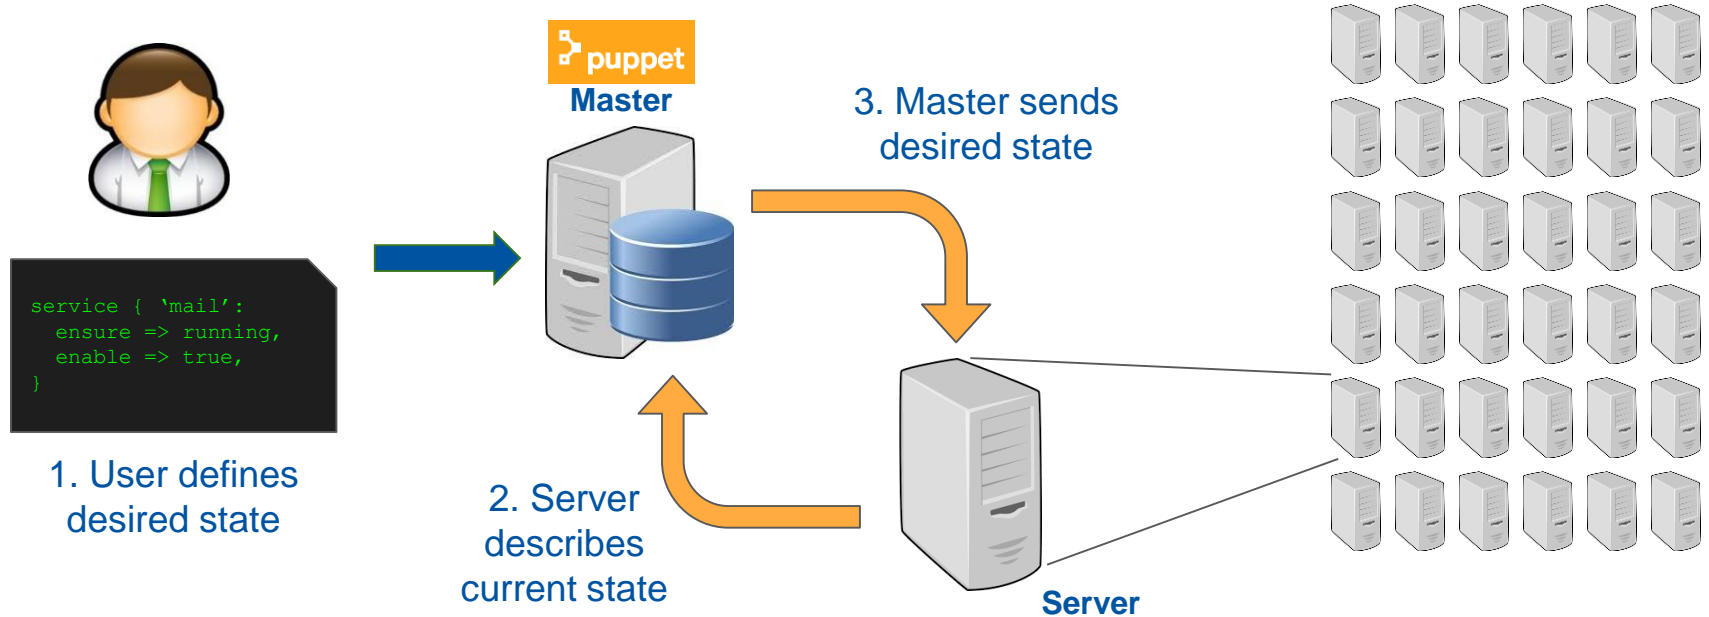


Resource provisioning can be automated:

- Assignment of resources
- Creation of virtual servers

How are the servers installed and configured?

Via a **configuration management system!**




What happens when something breaks? (1/2)

With 10'000+ servers, **things break constantly ...**

Hardware (e.g. disks, memory, ...)

Software (e.g. programs crash ...)

➔ **Centralised Monitoring!**

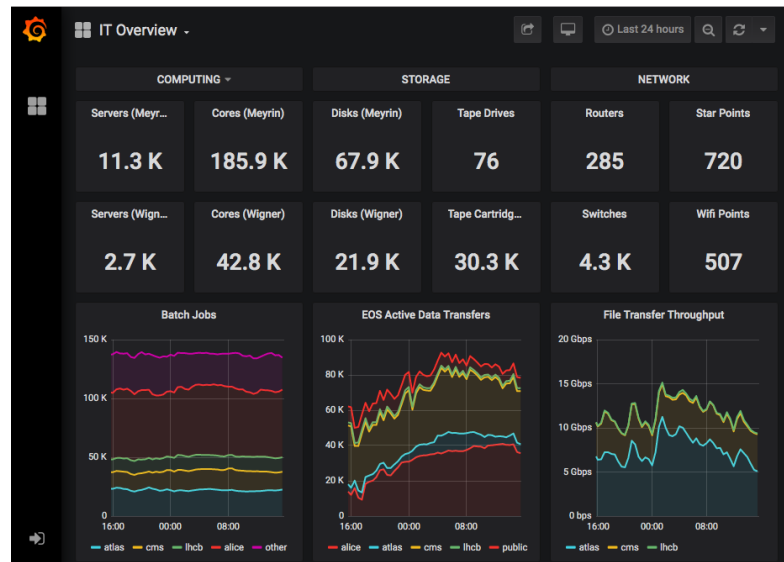


metrics

logs

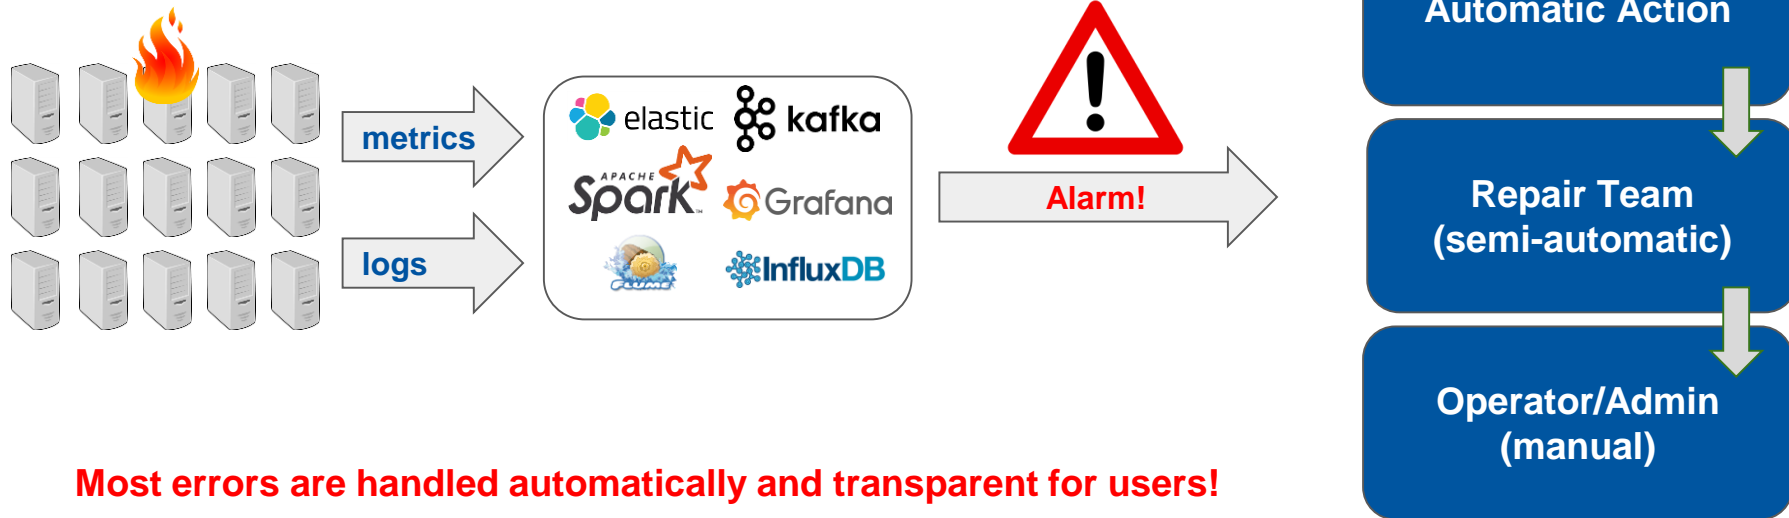


Transport / Storage / Visualization



What happens when something breaks? (2/2)

Automatic failure handling is key!



Most errors are handled automatically and transparent for users!

How do we automate server retirements?

After several years of service, **servers need to be removed.**

Process is similar to the initial burn-in.

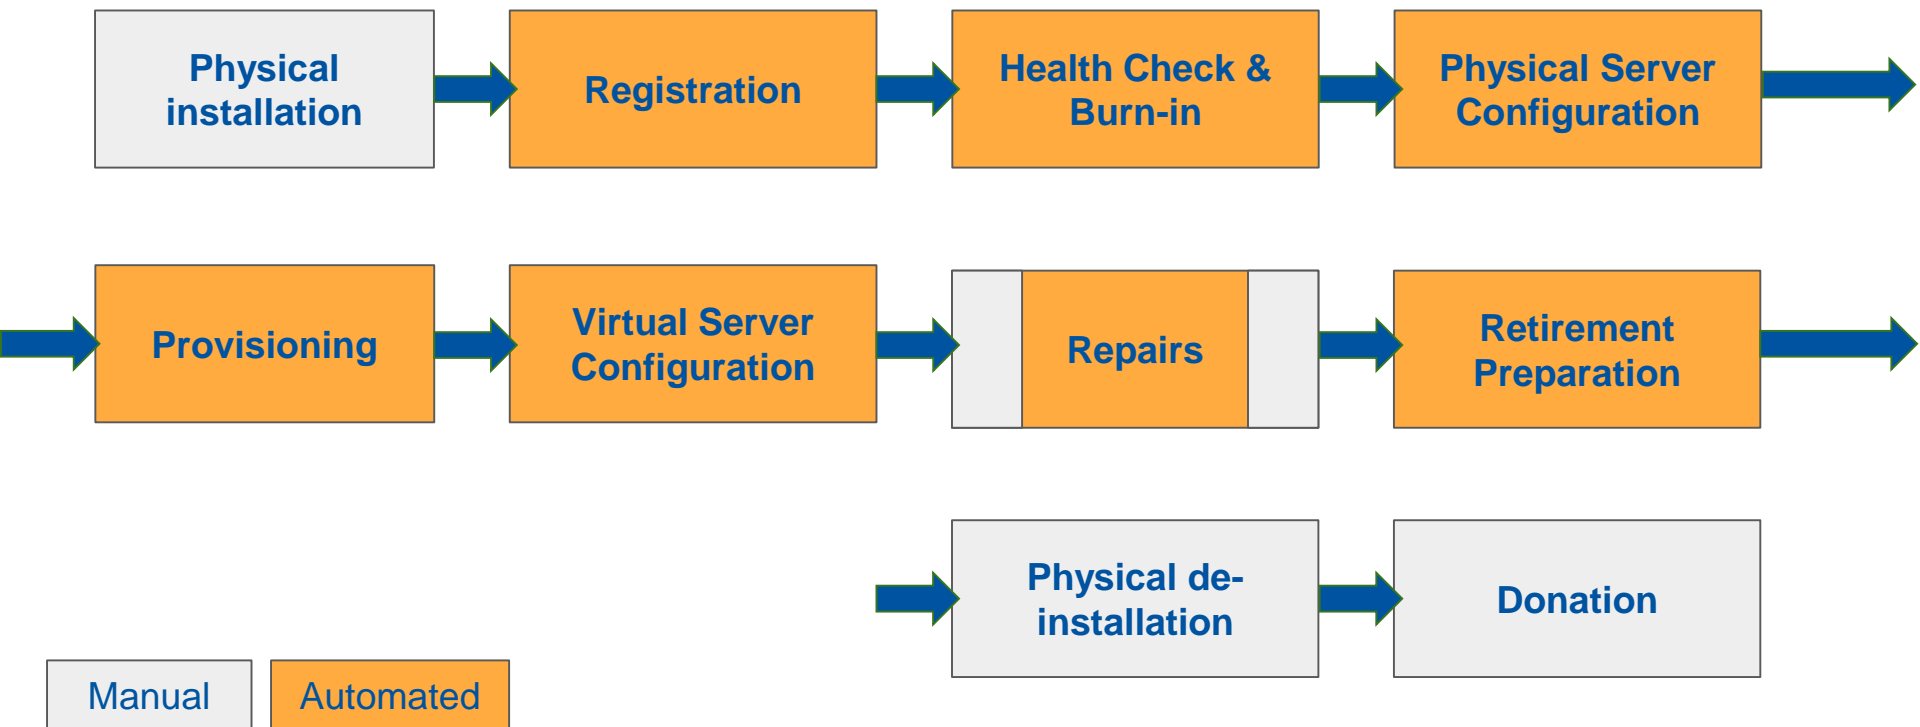
The nodes are marked to boot into a special system ...


- ... which makes sure all data is securely erased
- ... and re-checks the system in case of donation

Physical de-installation is again manual.



Automation for the CERN IT Server Life-cycle





Questions?



www.cern.ch