How do we manage 40K machines in the CERN Computer centre

ZHECHKA TOTEVA (CERN/IT)

CERN – 06/10/2022

Outline

- CERN Computer Centre (CC) in numbers
- Overview of the CERN network and data storage
- Overview of electricity and cooling
- WLCG is couple of numbers
- Configuration management at CERN IT
- CERNMegabus@CERN
  - Architecture
  - Overview of major implemented use cases
  - CERN Computer Centre (CC) power cut management

CERN Computer Centre (CC) in numbers

Computing network

- 250 routers, 4100 switches, 1200 Wi-Fi points
- 35 000 km optical fibre (only ~5 000 less than the equator length)

- Wigner Data centre in Hungary
  - 1200 km distance
  - with three 100 Gb/s

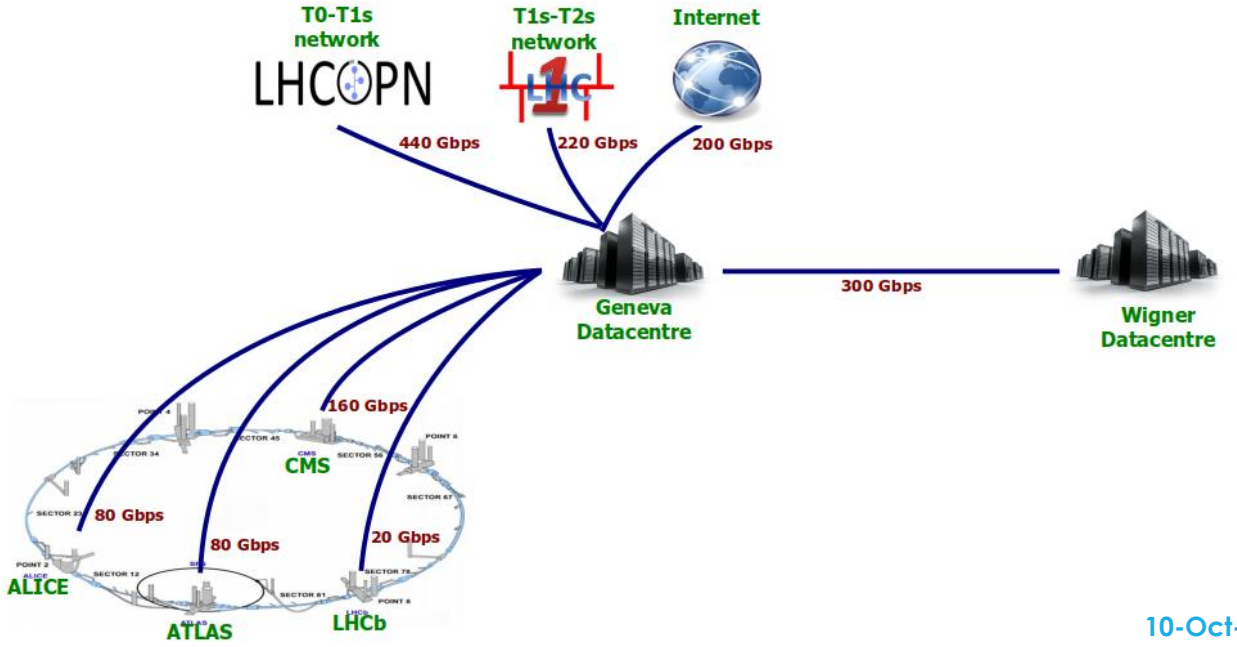

Data storage

- 50 Pbytes / year from LHC
  - + 25 Pbytes / year from non-LHC experiments
- **RECORD**: August 2018: 13.8 Pbytes of data written on tape (of which 11.56 is LHC data)
  - More than 2 PB read/write daily
- Tape drives faster than disks; but slower in mounting (latency)
  - 90 K disk drives (of which 10-15% are SSD, providing less than 10% capacity)
  - SSDs are 5-10 times more expensive than spinning disks

www.cern.ch/eos
www.cern.ch/castor

CASTOR dashboard

Last 1 year

Physics Data in CASTOR

https://filer-carbon.cern.ch/grafana/dashboard/db/castor-dashboard?orgId=1

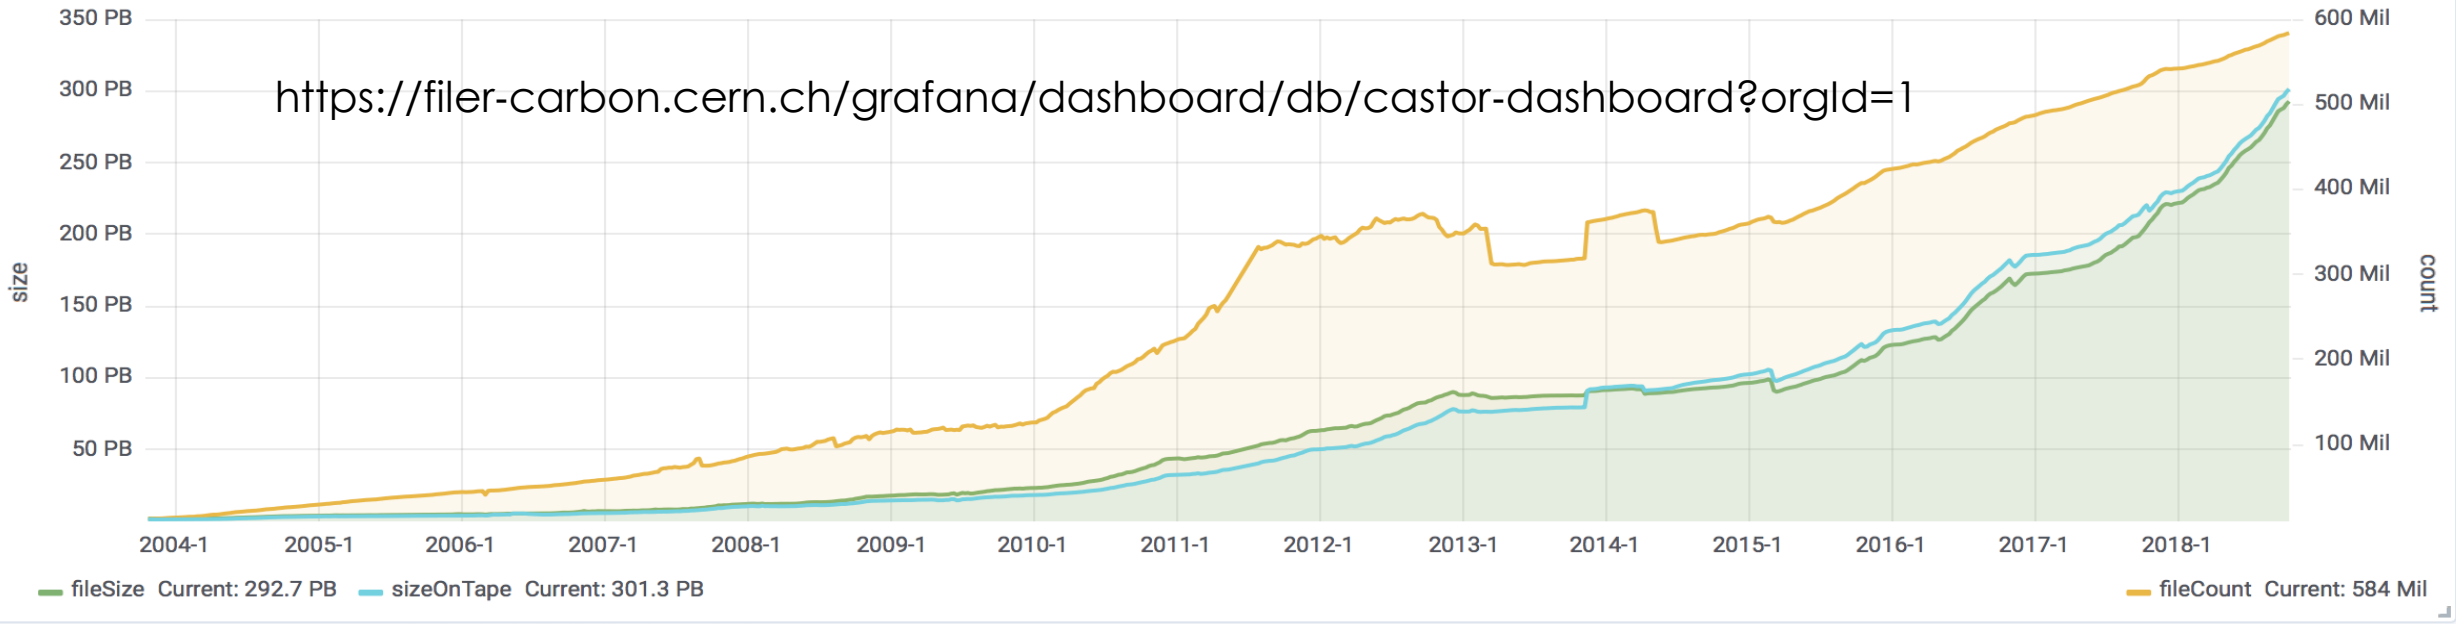

Total data

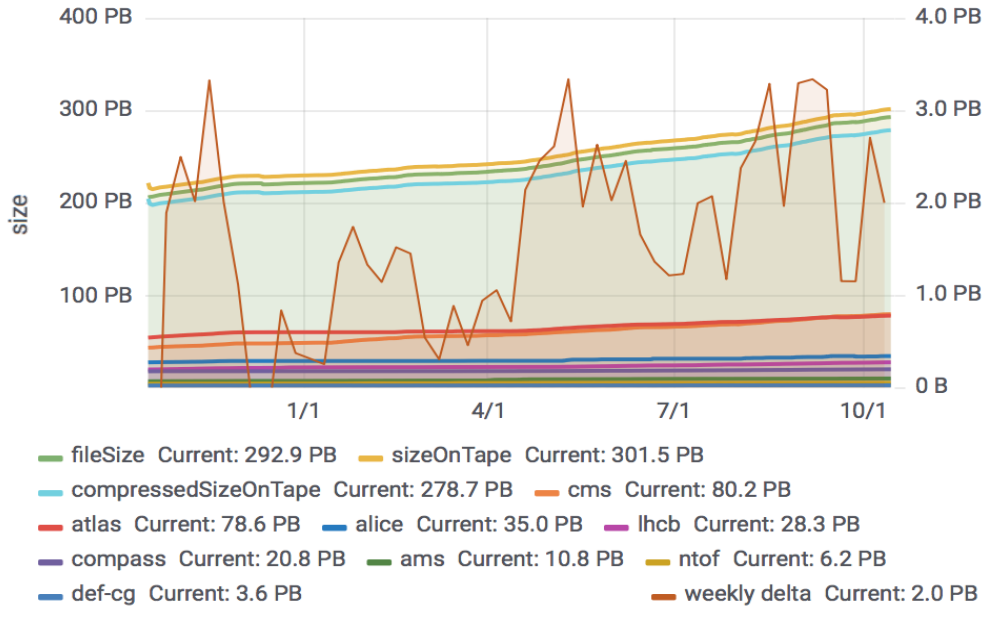

File counters

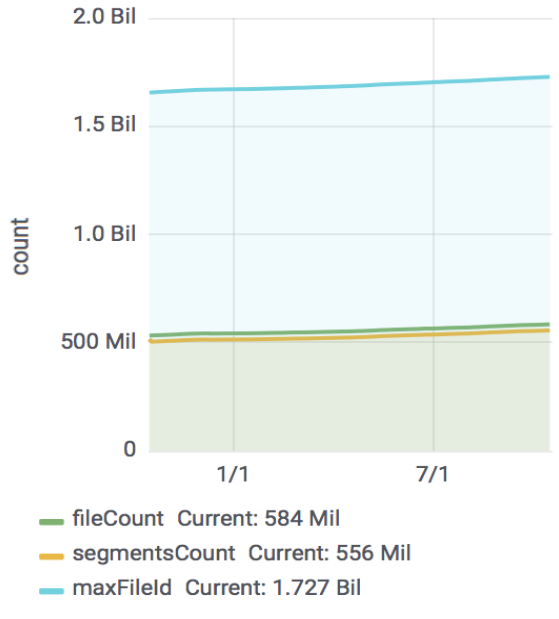

Mean file size

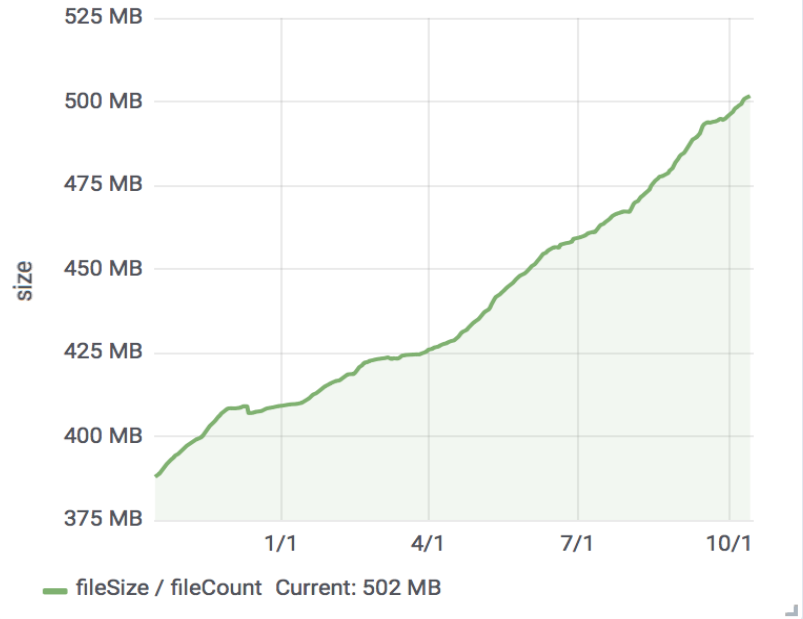

EOS Space Dashboard

Last 1 year Refresh every 5m

Instance All

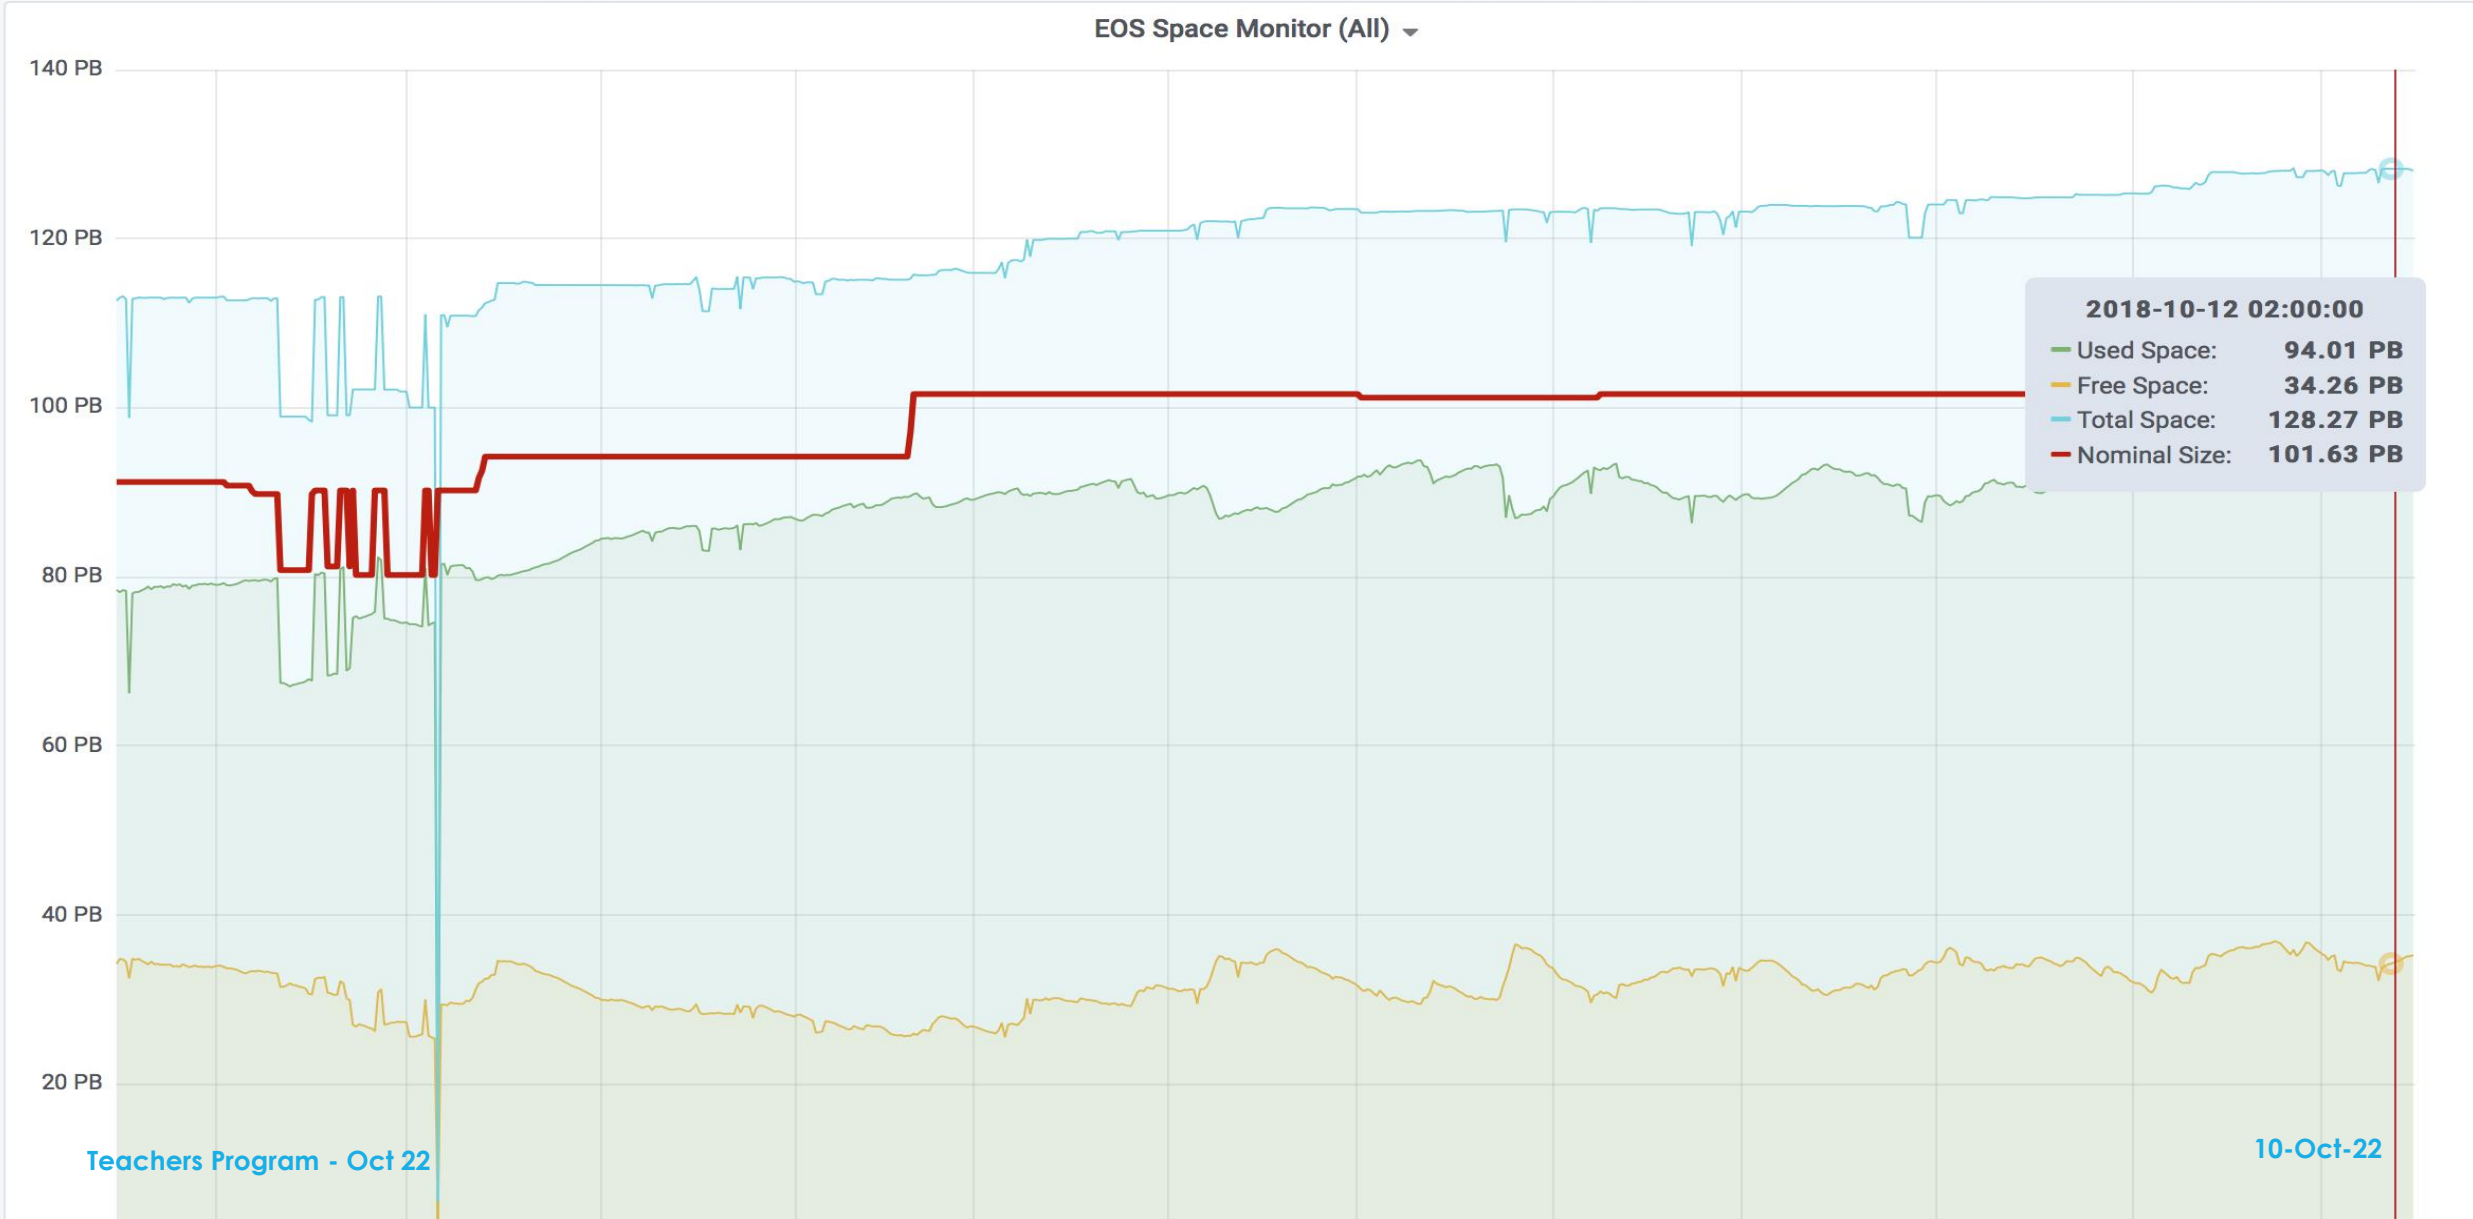

Electricity and cooling

- 2.7 MW consumption (+ ~ 1 MW cooling) from maximum 3.5 MW
  - 480 KW diesel generators
- Protected by UPS
  - Enough to start the diesel generators
  - Enough to shut down non-critical machines*
- Cooling
  - Chilled air via silver ducts enters the false floor and the into the closed server aisles
  - Water-cooled racks in the vault in the basement

WLCG – Worldwide LHC computing grid

- More than 170 data centres in 42 countries with about 800,000 cpu cores
  - CERN provides about 20% of the WLCG resources
  - Allows more than 10,000 physicists to access LHC data
  - >250,000 jobs run concurrently on the Grid
  - Storage is about 400 PB disk and 400 PB of tape globally
  - In 2016, global transfer rates have regularly exceeded 35GB/s
- Key facts and numbers (http://information-technology.web.cern.ch)

www.cern.ch/wlcg
www.cern.ch/wlcg-public

Evolution of the global core processor time delivered by the Worldwide LHC Computing Grid (WLCG)
As seen on the graph, the global central processing unit (CPU) time delivered by WLCG (expressed in billions of HS06 hours per month, HS06 being the HEP-wide benchmark for measuring CPU performance) shows a continual increase. In 2017, WLCG combined the computing resources of about 800 000 computer cores.

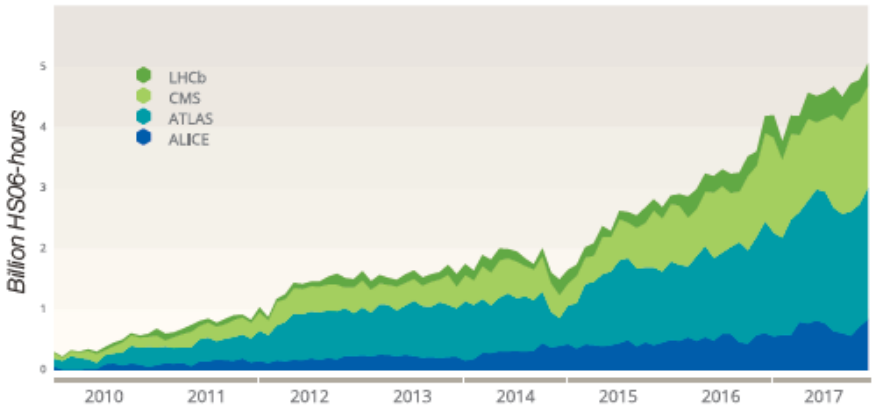

Configuration management at CERN IT

Certification Manager

CERNMegabus

43 000
Puppet managed machines

TEIGI Tool suite

puppet

MCOLLECTIVE

python

PuppetDB

AI-TOOLS

…

riakKV

django

The Puppet cycle

Interactions with the server and the agent

01 **Store manifests into Git**
As a first step, manifests (our config) have to be generated and stored in GitLab.

02 **Register a machine**
A machine will then be created, in a specific hostgroup (eg. webchat/frontend/atlas). It will be registered in Foreman.

03 **Run Puppet**
With the machine ready, the Puppet agent can be executed interactively (or let it run by itself). This will request the catalog (final state) of the machine.

04 **Master asks for hostgroup**
The Puppet master handling the request will ask Foreman for the hostgroup of the machine.

05 **Master asks for manifests**
Once it has the hostgroup, it will obtain the manifests that we defined in GitLab.

06 **Catalog generation**
As a final step, the Puppet master will generate the catalog and return it to the agent, which will apply it to the machine.

Thanks Config team fro the slide

CERNMegabus at CERN IT

A service that provides for **instant communication between services**

CERNMegabus

- The CERNMegabus architecture is based on the **publisher-consumer model and** utilises the **CERN IT messaging infrastructure**
- The publisher and the consumer services comprises of building blocks
  - **configured with Puppet**
  - to use the **CERNMegabus python libraries**
- **Installed on all** Puppet managed **machines** in the **CERN CC**

CERNMegabus architecture

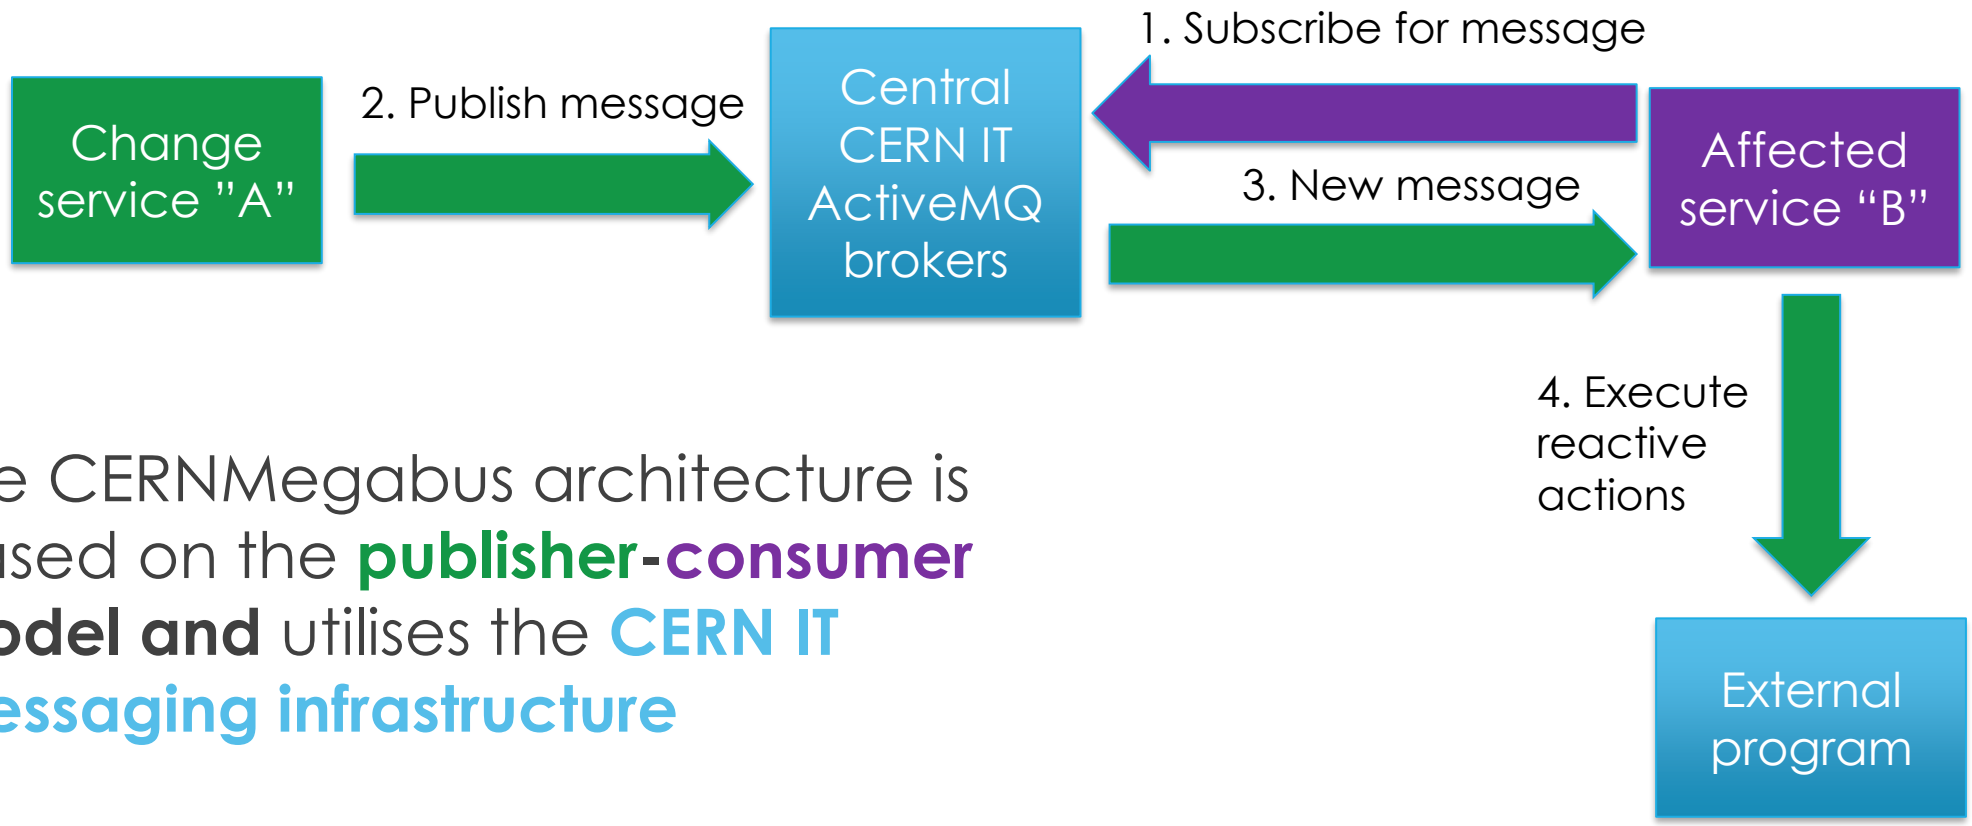

The CERNMegabus architecture is based on the **publisher-consumer model and** utilises the **CERN IT messaging infrastructure**

Already our clients

CASTOR

HAPROXY

BATCH

EOS

CERNMegabus

CLOUD

43 000
Puppet managed machines

TEIGI Tool suite (roger)

DNS Load Balancing

IT Monitoring

CERN Computer Centre Power Cut Management

From roger to EOS/CASTOR/Puppet HAProxy

Set roger state of *castor-lhcb-disknode-X* to disabled

2. Publish message to /topic/roger.**group.**castor**

Central CERN IT ActiveMQ brokers

1. Subscribe for topic /topic/roger.**group.**castor** and hostgroup selector in message

3. Message de-queued

castor-lhcb-headnode-Y

4. Execute `**modifydiskserver Disab castor-lhcb-disknode-X**`

Set castor-lhcb-disknode-X in read-only mode

hostgroup header: `castor-lhcb-diskservers`

{"**new**": {"update_time": "1538633786", "updated_by": "blueuser", "hostname": "***castor-lhcb-disknode-X***", ..., "appstate": "**disabled**"}, "**old**": {"update_time": "1538633774", "updated_by": "somebody", "hostname": "***castor-lhcb-disknode-X*** ", ..., "appstate": "**production**"}

~1 sec

In practice - I

CERN CC Power Cut event

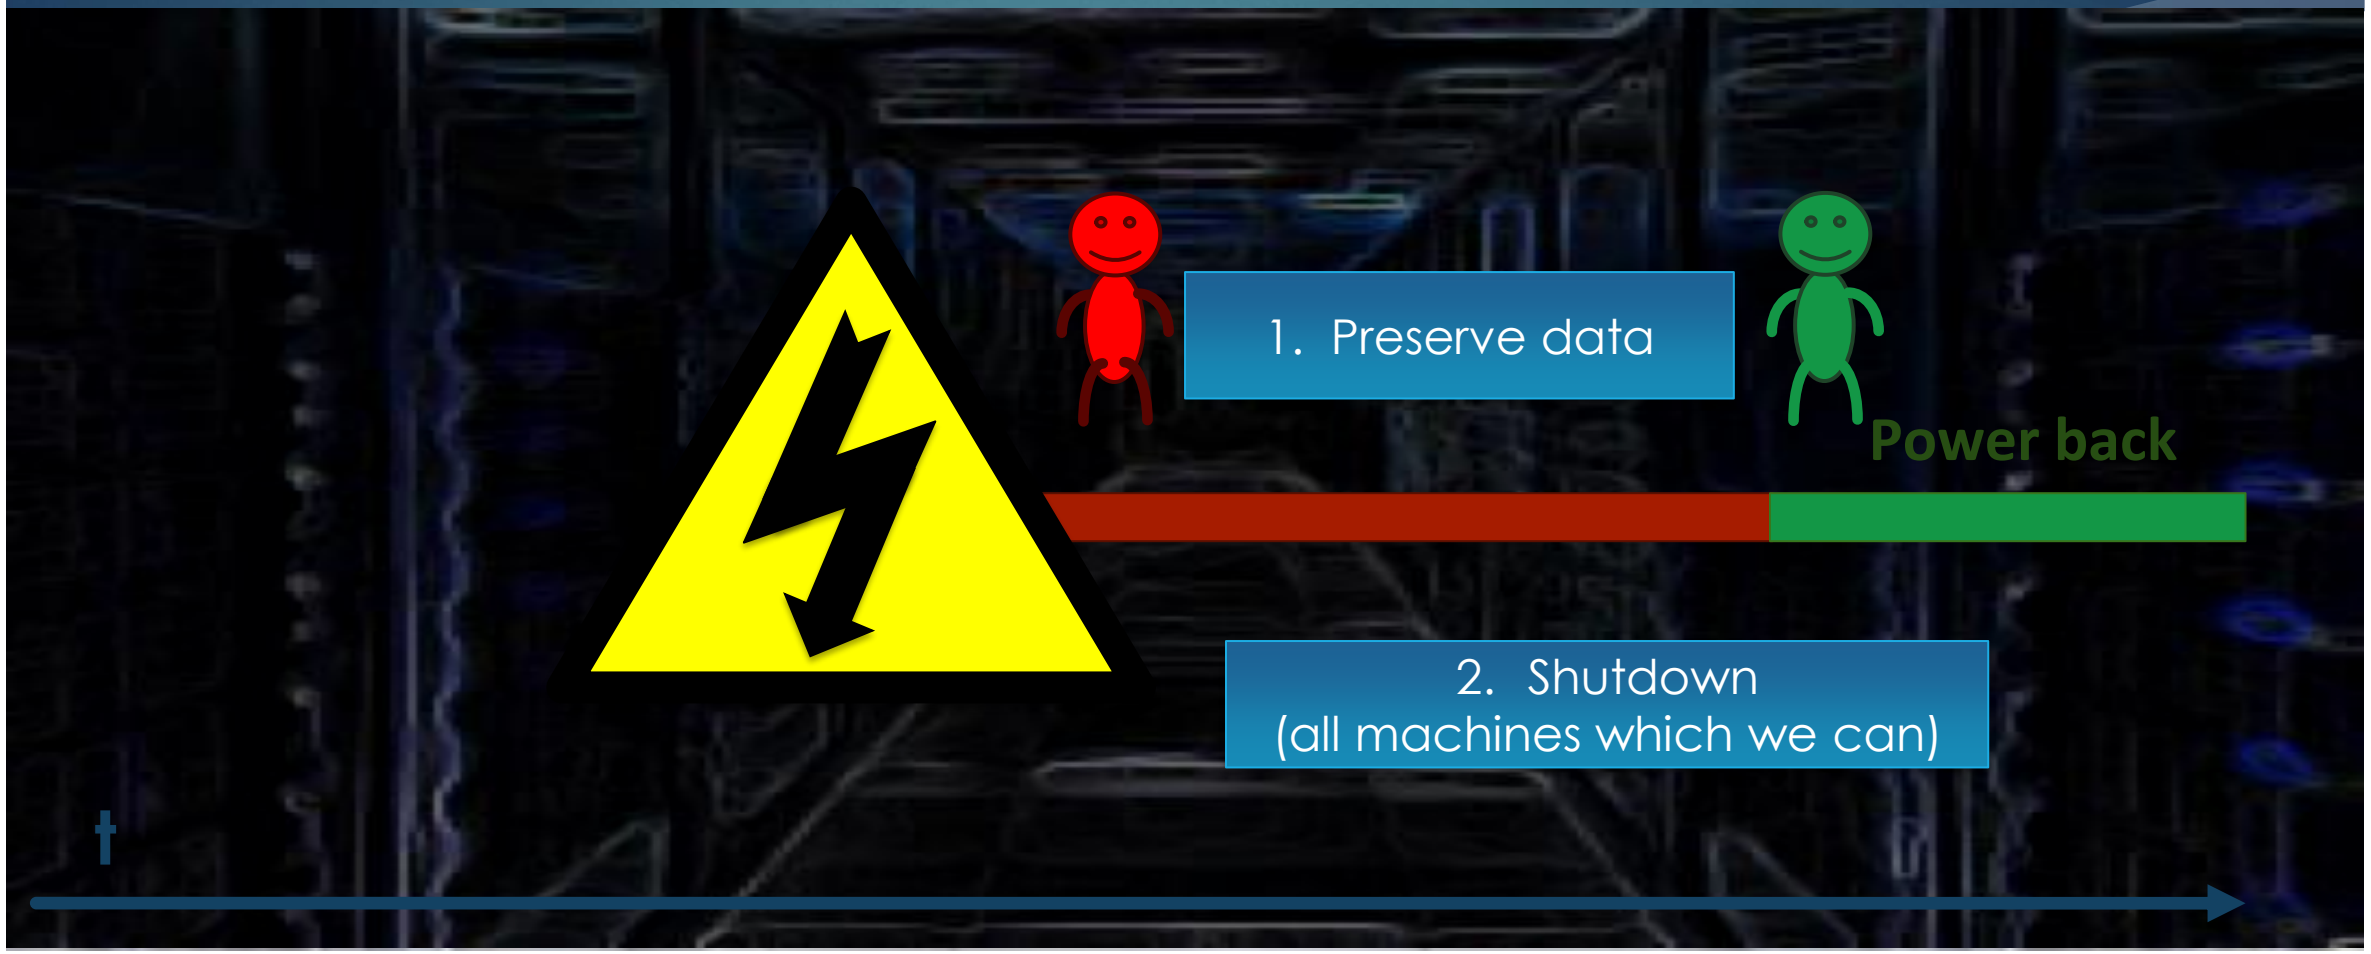

CERN CC Power cut event detection

ccpcoX programmatically detects power cut/power back event

UPS and PLC

UPS STATES

		UPS	Type de UPS	Descriptive	Data Type	MW	Bit	
GROUP 1	Module 1	EBS104*43	Galaxy 6000	Sur batteries	Single	20	0	Act
		EBS104*43	Galaxy 6000	Sur onduleur	Single	20	8	Ina
		EBS104*43	Galaxy 6000	En maintenance	Single	21	0	Act
	Module 2	EBS105*43	Galaxy 6000	Sur batteries	Single	21	8	Act
		EBS105*43	Galaxy 6000	Sur onduleur	Single	22	0	Ina
		EBS105*43	Galaxy 6000	En maintenance	Single	22	8	Act
	Module 3	EBS106*43	Galaxy 6000	Sur batteries	Single	23	0	Act
		EBS106*43	Galaxy 6000	Sur onduleur	Single	23	8	Ina
		EBS106*43	Galaxy 6000	En maintenance	Single	24	0	Act
	Module 4	EBS107*43	Galaxy 6000	Sur batteries	Single	24	8	Act
		EBS107*43	Galaxy 6000	Sur onduleur	Single	25	0	Ina
		EBS107*43	Galaxy 6000	En maintenance	Single	25	8	Act
	N/S	EBS208*43	Galaxy 6000	Danger Bypass	Single	26	0	Ina
		EBS208*43	Galaxy 6000	En maintenance	Single	26	8	Act
GROUP 2	Module 1	EBS204*43	Galaxy 6000	Sur batteries	Single	30	0	Act
		EBS204*43	Galaxy 6000	Sur onduleur	Single	30	8	Ina
		EBS204*43	Galaxy 6000	En maintenance	Single	31	0	Act
	Module 2	EBS205*43	Galaxy 6000	Sur batteries	Single	31	8	Act
		EBS205*43	Galaxy 6000	Sur onduleur	Single	32	0	Ina
		EBS205*43	Galaxy 6000	En maintenance	Single	32	8	Act
	Module 3	EBS206*43	Galaxy 6000	Sur batteries	Single	33	0	Act
		EBS206*43	Galaxy 6000	Sur onduleur	Single	33	8	Ina
		EBS206*43	Galaxy 6000	En maintenance	Single	34	0	Act
	Module 4	EBS207*43	Galaxy 6000	Sur batteries	Single	34	8	Act
		EBS207*43	Galaxy 6000	Sur onduleur	Single	35	0	Ina
		EBS207*43	Galaxy 6000	En maintenance	Single	35	8	Act
	N/S	EBS308*43	Galaxy 6000	Danger Bypass	Single	36	0	Ina
		EBS308*43	Galaxy 6000	En maintenance	Single	36	8	Act
GROUP 3	Module 1	EBS304*43	Galaxy 6000	Sur batteries	Single	40	0	Act
		EBS304*43	Galaxy 6000	Sur onduleur	Single	40	8	Ina
		EBS304*43	Galaxy 6000	En maintenance	Single	41	0	Act
	Module 2	EBS305*43	Galaxy 6000	Sur batteries	Single	41	8	Act
		EBS305*43	Galaxy 6000	Sur onduleur	Single	42	0	Ina
		EBS305*43	Galaxy 6000	En maintenance	Single	42	8	Act
	Module 3	EBS306*43	Galaxy 6000	Sur batteries	Single	43	0	Act
		EBS306*43	Galaxy 6000	Sur onduleur	Single	43	8	Ina
		EBS306*43	Galaxy 6000	En maintenance	Single	44	0	Act
	Module 4	EBS307*43	Galaxy 6000	Sur batteries	Single	44	8	Act
		EBS307*43	Galaxy 6000	Sur onduleur	Single	45	0	Ina
		EBS307*43	Galaxy 6000	En maintenance	Single	45	8	Act
	N/S	EBS709*43	Galaxy 6000	Danger Bypass	Single	46	0	Ina
		EBS709*43	Galaxy 6000	En maintenance	Single	46	8	Act
GROUP 4	Module 1	EBS404*43	Galaxy 7000	Sur batteries	Single	50	0	Act
		EBS404*43	Galaxy 7000	Sur onduleur	Single	50	8	Ina
		EBS404*43	Galaxy 7000	En maintenance	Single	51	0	Act
	Module 2	EBS405*43	Galaxy 7000	Sur batteries	Single	51	8	Act
		EBS405*43	Galaxy 7000	Sur onduleur	Single	52	0	Ina
		EBS405*43	Galaxy 7000	En maintenance	Single	52	8	Act
	Module 3	EBS406*43	Galaxy 7000	Sur batteries	Single	53	0	Act
		EBS406*43	Galaxy 7000	Sur onduleur	Single	53	8	Ina
		EBS406*43	Galaxy 7000	En maintenance	Single	54	0	Act
	Module 4	EBS406*43	Galaxy 7000	En maintenance	Single	54	0	Actif
		EBS407*43	Galaxy 7000	Sur batteries	Single	54	8	Actif
		EBS407*43	Galaxy 7000	Sur onduleur	Single	55	0	Inactif
		EBS407*43	Galaxy 7000	En maintenance	Single	55	8	Actif
	N/S	EBS408*43	Galaxy 7000	Danger Bypass	Single	56	0	Inactif
		EBS408*43	Galaxy 7000	En maintenance	Single	56	8	Actif
GROUP 5	Module 1	EBS704*43	Galaxy 7000	Sur batteries	Single	60	0	Actif
		EBS704*43	Galaxy 7000	Sur onduleur	Single	60	8	Inactif
		EBS704*43	Galaxy 7000	En maintenance	Single	61	0	Actif
	Module 2	EBS705*43	Galaxy 7000	Sur batteries	Single	61	8	Actif
		EBS705*43	Galaxy 7000	Sur onduleur	Single	62	0	Inactif
		EBS705*43	Galaxy 7000	En maintenance	Single	62	8	Actif
	Module 3	EBS706*43	Galaxy 7000	Sur batteries	Single	63	0	Actif
		EBS706*43	Galaxy 7000	Sur onduleur	Single	63	8	Inactif
		EBS706*43	Galaxy 7000	En maintenance	Single	64	0	Actif
	Module 4	EBS707*43	Galaxy 7000	Sur batteries	Single	64	8	Actif
		EBS707*43	Galaxy 7000	Sur onduleur	Single	65	0	Inactif
		EBS707*43	Galaxy 7000	En maintenance	Single	65	8	Actif
	Module 5	EBS708*43	Galaxy 7000	Sur batteries	Single	66	0	Actif
		EBS708*43	Galaxy 7000	Sur onduleur	Single	66	8	Inactif
		EBS708*43	Galaxy 7000	En maintenance	Single	67	0	Actif
	Module 6	EBS720*43	Galaxy 7000	Sur batteries	Single	68	8	Actif
		EBS720*43	Galaxy 7000	Sur onduleur	Single	69	0	Inactif
		EBS720*43	Galaxy 7000	En maintenance	Single	69	8	Actif
	N/S	EBS108*43	Galaxy 7000	Danger Bypass	Single	67	0	Inactif
		EBS108*43	Galaxy 7000	En maintenance	Single	68	8	Actif

EBS406*43	Galaxy 7000	En maintenance	Single	54	0	Actif	Inactif	G4-M3-MAINTENANCE	(54:0) = 1 @2017-12-04 15:37:32
EBS407*43	Galaxy 7000	Sur batteries	Single	54	8	Actif	Inactif	G4-M4-BATTERY	(54:8) = 1 @2017-12-04 15:37:32
EBS407*43	Galaxy 7000	Sur onduleur	Single	55	0	Inactif	Actif	G4-M4-ONDULEUR	(55:0) = 1 @2017-12-04 15:37:32
EBS407*43	Galaxy 7000	En maintenance	Single	55	8	Actif	Inactif	G4-M4-MAINTENANCE	(55:8) = 1 @2017-12-04 15:37:32
EBS408*43	Galaxy 7000	Danger Bypass	Single	56	0	Inactif	Actif	G4-NS-BYPASS	(56:0) = 1 @2017-12-04 15:37:32
EBS408*43	Galaxy 7000	En maintenance	Single	56	8	Actif	Inactif	G4-NS-MAINTENANCE	(56:8) = 1 @2017-12-04 15:37:32
EBS704*43	Galaxy 7000	Sur batteries	Single	60	0	Actif	Inactif	G5-M1-BATTERY	(60:0) = 1 @2017-12-04 15:37:32
EBS704*43	Galaxy 7000	Sur onduleur	Single	60	8	Inactif	Actif	G5-M1-ONDULEUR	(60:8) = 1 @2017-12-04 15:37:32
EBS704*43	Galaxy 7000	En maintenance	Single	61	0	Actif	Inactif	G5-M1-MAINTENANCE	(61:0) = 1 @2017-12-04 15:37:32
EBS705*43	Galaxy 7000	Sur batteries	Single	61	8	Actif	Inactif	G5-M2-BATTERY	(61:8) = 1 @2017-12-04 15:37:32
EBS705*43	Galaxy 7000	Sur onduleur	Single	62	0	Inactif	Actif	G5-M2-ONDULEUR	(62:0) = 1 @2017-12-04 15:37:32
EBS705*43	Galaxy 7000	En maintenance	Single	62	8	Actif	Inactif	G5-M2-MAINTENANCE	(62:8) = 1 @2017-12-04 15:37:32
EBS706*43	Galaxy 7000	Sur batteries	Single	63	0	Actif	Inactif	G5-M3-BATTERY	(63:0) = 1 @2017-12-04 15:37:32
EBS706*43	Galaxy 7000	Sur onduleur	Single	63	8	Inactif	Actif	G5-M3-ONDULEUR	(63:8) = 1 @2017-12-04 15:37:32
EBS706*43	Galaxy 7000	En maintenance	Single	64	0	Actif	Inactif	G5-M3-MAINTENANCE	(64:0) = 1 @2017-12-04 15:37:32
EBS707*43	Galaxy 7000	Sur batteries	Single	64	8	Actif	Inactif	G5-M4-BATTERY	(64:8) = 1 @2017-12-04 15:37:32
EBS707*43	Galaxy 7000	Sur onduleur	Single	65	0	Inactif	Actif	G5-M4-ONDULEUR	(65:0) = 1 @2017-12-04 15:37:32
EBS707*43	Galaxy 7000	En maintenance	Single	65	8	Actif	Inactif	G5-M4-MAINTENANCE	(65:8) = 1 @2017-12-04 15:37:32
EBS708*43	Galaxy 7000	Sur batteries	Single	66	0	Actif	Inactif	G5-M5-BATTERY	(66:0) = 1 @2017-12-04 15:37:32
EBS708*43	Galaxy 7000	Sur onduleur	Single	66	8	Inactif	Actif	G5-M5-ONDULEUR	(66:8) = 1 @2017-12-04 15:37:32
EBS708*43	Galaxy 7000	En maintenance	Single	67	0	Actif	Inactif	G5-M5-MAINTENANCE	(67:0) = 1 @2017-12-04 15:37:32
EBS720*43	Galaxy 7000	Sur batteries	Single	68	8	Actif	Inactif	G5-NS-BYPASS	(67:8) = 1 @2017-12-04 15:37:32
EBS720*43	Galaxy 7000	Sur onduleur	Single	69	0	Inactif	Actif	G5-M6-BATTERY	(68:8) = 1 @2017-12-04 15:37:32
EBS720*43	Galaxy 7000	En maintenance	Single	69	8	Actif	Inactif	G5-NS-MAINTENANCE	(68:0) = 1 @2017-12-04 15:37:32
EBS108*43	Galaxy 7000	Danger Bypass	Single	67	0	Inactif	Actif	G5-M6-ONDULEUR	(69:0) = 1 @2017-12-04 15:37:32
EBS108*43	Galaxy 7000	En maintenance	Single	68	8	Actif	Inactif	G5-M6-MAINTENANCE	(69:8) = 1 @2017-12-04 15:37:32

GALAXY 6000 LOGIC

UPS

	BATTERY	ONDULEUR	MAINTENANCE
MODE NORMAL	●	●	●
MODE MAINTENANCE	●	○	○
MODE BATTERIE	○	●	●
HORS TENSION	●	○	●

NS

	BYPASS	MAINTENANCE
MODE NORMAL	●	●
MODE BYPASS	○	●
MODE MAINTENANCE	○	○

GALAXY 7000 LOGIC

UPS

	BATTERY	ONDULEUR	MAINTENANCE
MODE NORMAL	●	●	●
MODE MAINTENANCE	●	○	○
MODE BATTERIE	○	○	●
HORS TENSION	○	●	○

NS

	BYPASS	MAINTENANCE
MODE NORMAL	●	●
MODE BYPASS	○	●
MODE MAINTENANCE	○	○

Raw data (in case there is some error with pa

Address	Value
20	257
21	257
22	257
23	257
24	257
25	257
26	1
27	0
28	0
29	0
30	257
31	257
32	257
33	257
34	257
35	257
36	1
37	0
38	0
39	0
40	257
41	257
42	257
43	257
44	257
45	257
46	1

CERN CC Power cut event detection algorithm

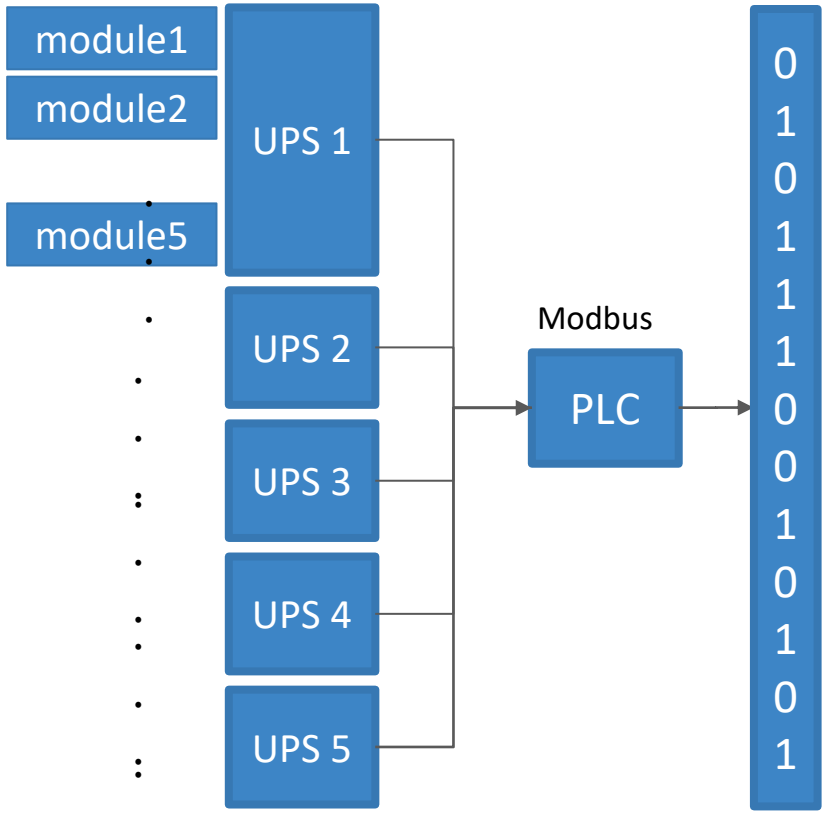

Data collection

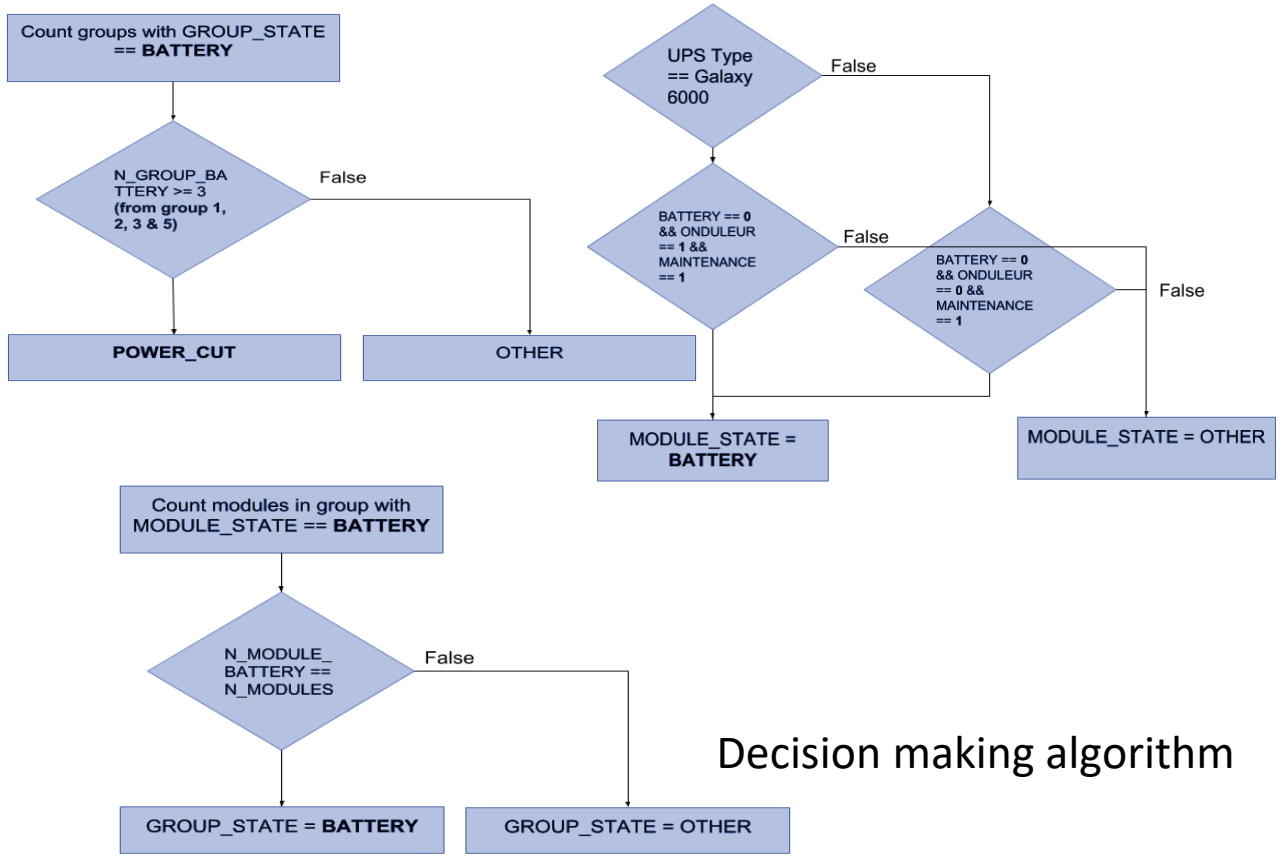

Decision making algorithm

CERN CC Power cut tests

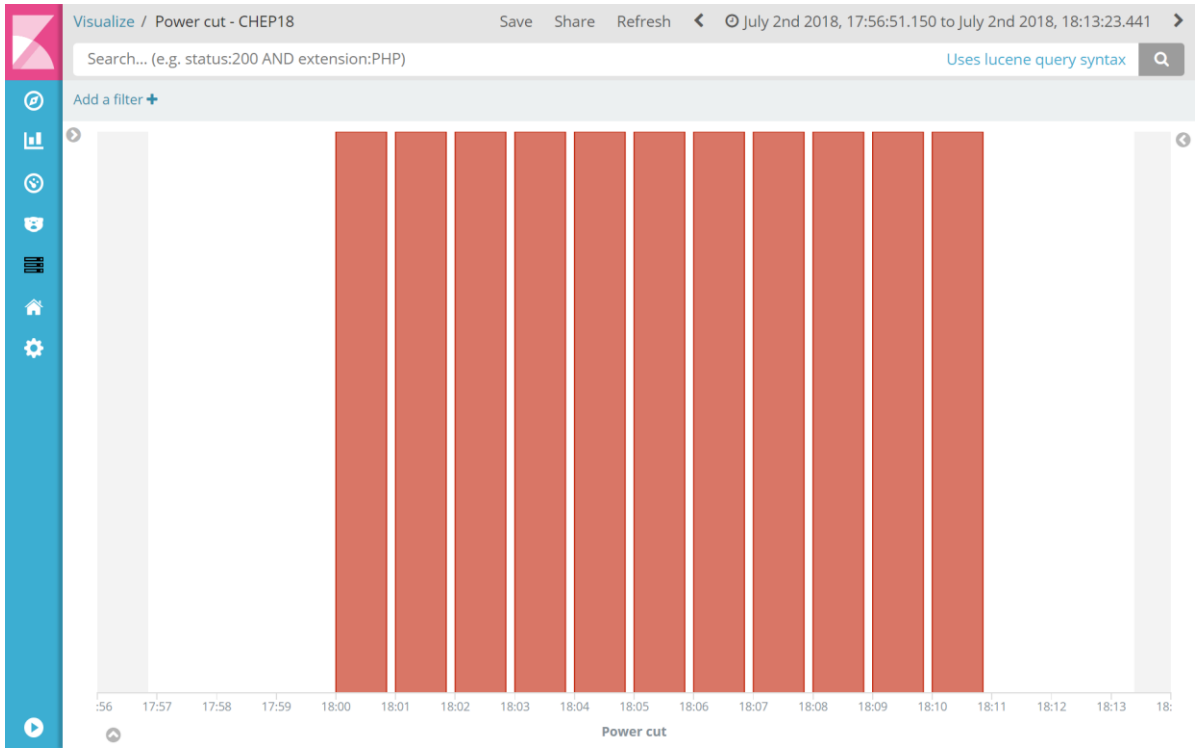

- During mid-annual power cut test on the 2nd of July, 2018
  - Detected power cut
  - Notified the subscribed machines
  - Shutdown the machines, which had been predefined to be shutdown
  - Detected the power back
  - Notified the machines, which h… predefined to wait

Presented at CHEP'18

From CERN CC UPS PLC to CERN CC shutdown

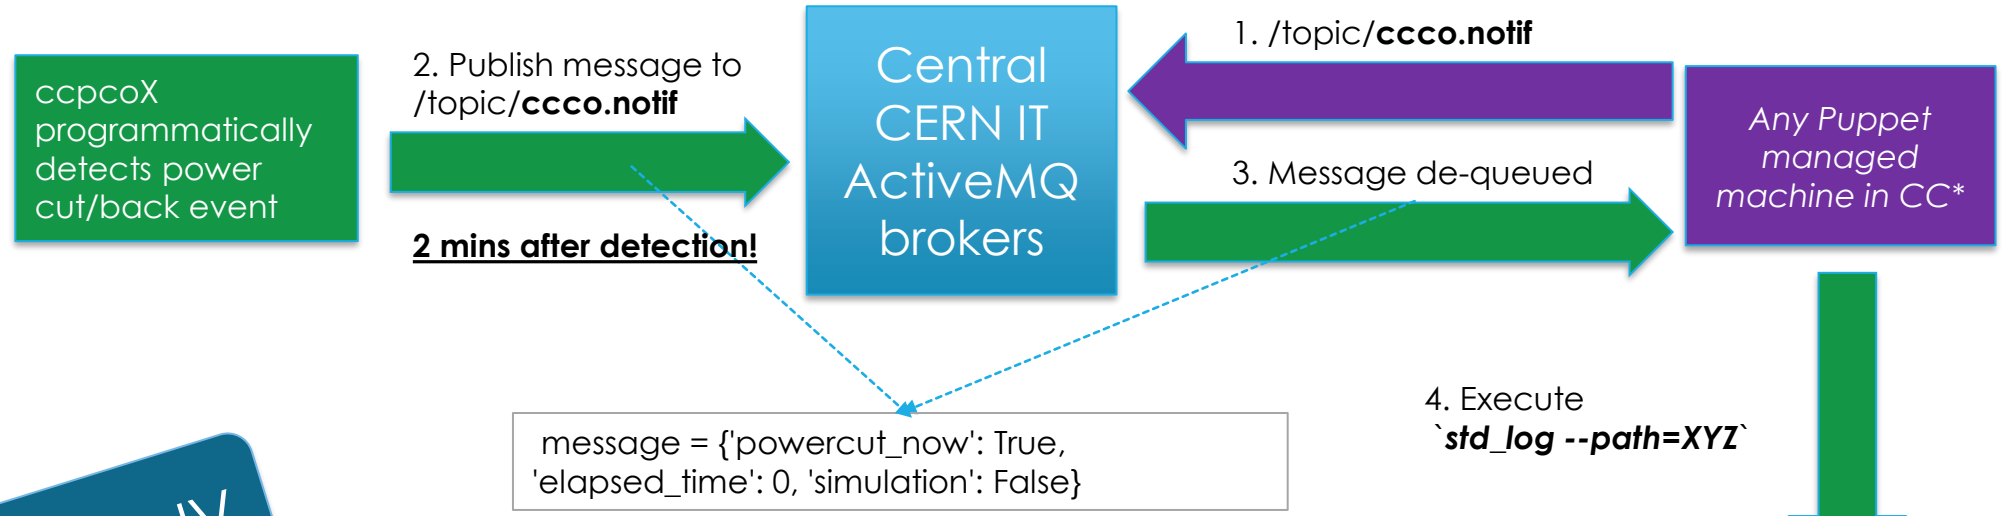

In practice - IV

~2 min*

Thanks

THANKS for listening and ENJOY your visit at CERN

Zhechka