



How to Provide the Run3 and Run4 Tier0 Computing Pledges

Wayne Salter - CERN

Overview

- Some History
- LHCb and ALICE Containers
- Run3 Approach and Experience
- Some More History
- Run4 Approach
- Summary

Some History - I

- End of Wigner contract after an excellent collaboration coming up:
 - Started 1st January 2013
 - Will end 31st December 2019
- Was only seen as a short term fix until we could get our own second DC
- ALICE and LHCb experiments are upgrading their trigger farms for RUN3 (2021)
 - No HW trigger anymore
- They conducted a joint tender for pre-fabricated data centres
 - Winning bid (*) based on containers, each with 24x 48U racks and 500kW equipment load (PUE 1.1)
 - Indirect free air-cooling with adiabatic support

Some History - II

- However, an idea for a Prévessin Computer Centre (PCC) developed in 2016 by Director of Research and Computing
 - A central Data Centre (DC) to meet the needs of IT and ALICE/LHCb for RUN3 and beyond
 - Hence, intended to replace the ALICE/LHCb containers
- Container tender result put on hold during the PCC tendering

Some History - III

- PCC MS launched in December 2016 and Tender in March 2017
- Despite getting interesting (and cost effective replies), the project was blocked by the CERN Management
 - Tried to keep the Tender valid for as long as possible but eventually it was dropped
 - So we needed another solution, as well as, ALICE and LHCb
- For ALICE/LHCb their containers were ordered
- Offer from LHCb for IT to use two of their 6 containers as they won't be using them during RUN3
- But we don't have sufficient budget to fill with new equipment

LHCb Container Installation




LHCb Container Installation



LHCb Container: Electrical Feeds

- No UPS, no diesels
- Single 66kV source
- Two transformers for two independent feeds




Container 5

Container 6



Inside a Container



Approach for RUN3

- Returning equipment from Wigner
 - CPU servers are going to one container
 - Storage servers are going to the CERN DC (and a small amount to the 2nd Network Hub for Disaster Recovery)
 - Equipment not needed will be donated to LHCb (CPU) or left at Wigner for ALICE (storage and CPU)
- New purchases (2019, 2020 and 2021)
 - CPU servers to the 2nd container
 - Storage to be installed in the CERN DC
- Upgrading some water-cooled racks in the Data Centre
 - Changing aging doors
 - Increasing from 80 to 88 racks

Shipment Schedule

- *Shipment 1 – February 2019*
 - 15 pallets (12 CPU servers (150 units), 2 switches, 1 rails, cables, etc.)
- *Shipment 2 – April 2019*
 - 19 pallets (5 CPU servers (54 units), 12 JBODs (72 units), 2 rails, cables, etc.)
- *Shipment 3 – June 2019*
 - 33 pallets (11 CPU servers (129 units), 18 JBODs (144 units), 2 switches, 2 rails, cables, etc.)
- *Shipment 4 – July 2019*
 - 29 pallets (13 CPU servers (150 units), 13 JBODs (104 units), 2 switches, 1 rails, cables, etc.)
- *Shipment 5 – October 2019*
 - 34 pallets (14 CPU servers (174 units), 15 JBODs (120 units), 2 switches, 3 rails, cables, etc.)
- *Shipment 6 – December 2019*
 - 20 pallets (16 CPU servers (183 units), 2 switches, 2 rails, cables, etc.)

Experience So Far - I

- Excellent collaboration with LHCb!!!
- There was an issue with water-tightness between the AHUs and the containers themselves
- The first shipment was logistically very tricky, but the subsequent ones have been relatively smooth
 - Very good support from Wigner!!!
- Containers were not 100% ready for use and so first installations were difficult
 - No ramp (still not)
 - No storage close to the containers
 - Additional containers for storage for equipment and repairs installed later
- There was a leak in one of the AHU due to a design fault
- We have been affected by a scheduled power cut

First Shipment Storage



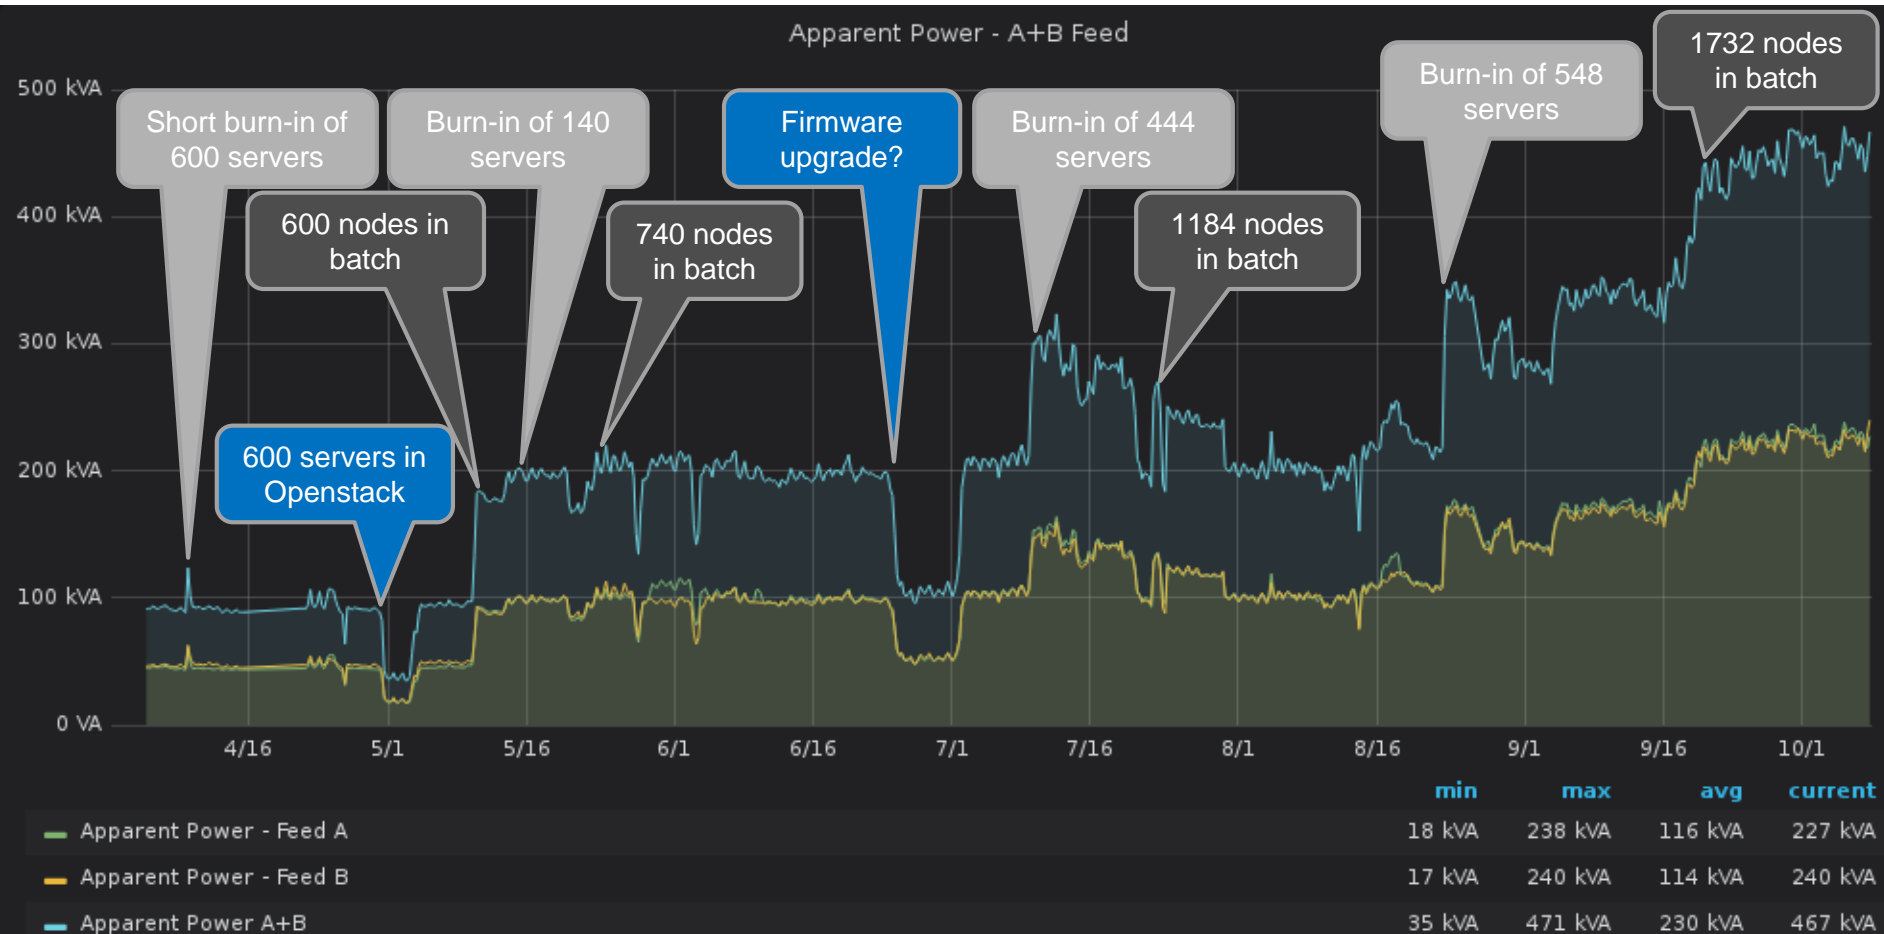
Second Shipment and Storage Container



Experience So Far - II

- Operationally far from ideal
 - Limited space to work
 - Noisy, hot and strong air flow
 - Probably ok for one shot deployment, where the work is completed before the machines are turned on
 - But we are installing in stages!!!!
 - Power strips were pre-install and complicated deployment as there was not enough space to use power tools
 - 2 breakers per phase and it is necessary to balance between them, as well as, between the phases
 - LHCb test equipment was not removed in-time for us to follow our original installation schedule
 - Integration of data into our monitoring system has not been straightforward, but is now done
- Container 6 is now full and we are starting to install new purchases in Container 5

Container 6 Power Usage



More Installations



Storage of Storage Nodes In Current Data Centre



Some More History

- Kept the idea of a new DC alive, but this time, for IT needs only
 - Only cost-effective solution for providing the increased capacity needed for RUN4 and beyond
 - In-time for the end of RUN3
 - Eventually, supported by the CERN Management
 - Included into the Medium Term Plan, which was agreed at the June Finance Committee
 - But foreseen budget has already been cut....

Approach for RUN4

- PCC to be built for the end of RUN3
- Have been discussing internally with technical groups, purchasing and safety
 - Started directly after the FC in June
 - Worked on the MS documents in parallel
- MS went out at the beginning of October
- Tender due out in February 2020
 - Currently working on the Technical Specification
- Adjudication foreseen in December 2020 (Last FC with current management)
- Contract to start in Q1 2021 (unless LS2 extended)

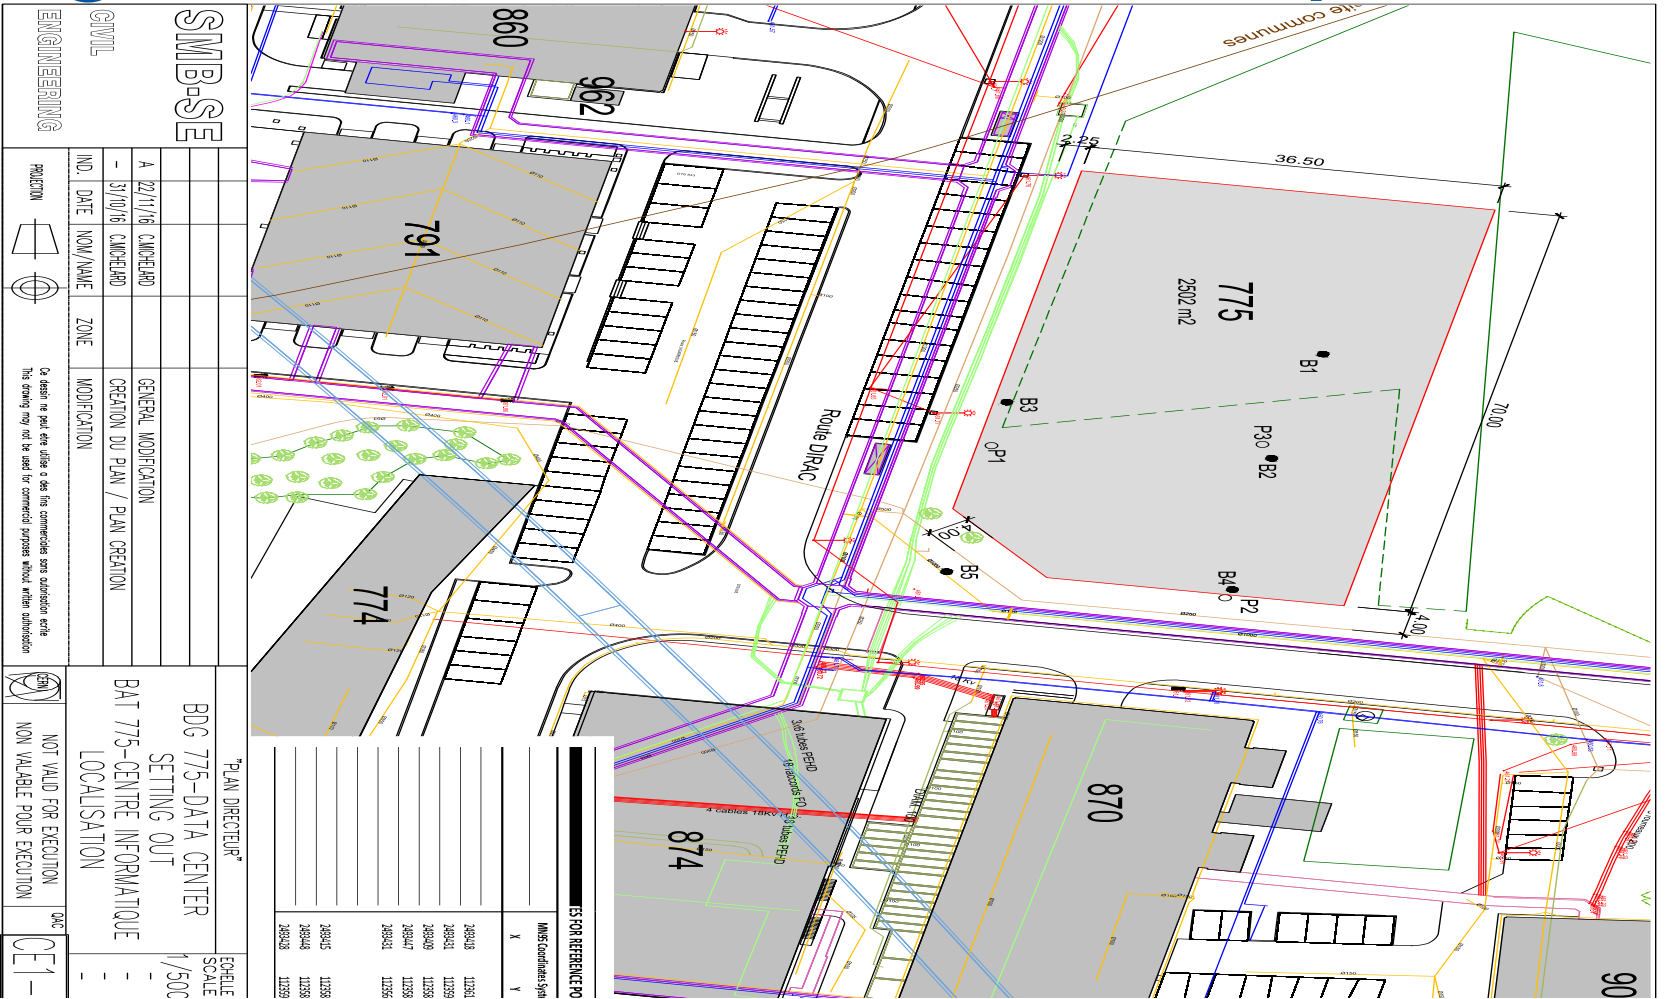
Foreseen Approach and Schedule

- The general approach is for a turnkey Data Centre to be able to take advantage of the expert industry knowledge in this very specialised domain, which is not available at CERN
 - But, HV equipment to be provided by CERN
- Contract would also include the support and maintenance of all technical systems for an initial period of 10 years
 - However, in order to integrate smoothly into the CERN technical infrastructure, involvement of some CERN services will be required.
- Technical groups in the Accelerator and Technical Sector should have the effort available to support the project in 2021-2022
- An optimistic scenario would therefore allow for commissioning of the PCC towards the end of 2022

High Level Functional Specification

- Energy efficient (PUE ≤ 1.1), flexible (supporting varying power densities) and scalable (initially 4MW and expandable in steps to 12MW) DC
- Average Power density of 20kW and maximum of 25kW.
- 2N electrical distribution and N+1 HVAC redundancy
- Including storage and repair rooms as well as rooms for CERN HV equipment, but no offices
 - Toilets required nonetheless
- Only partial UPS coverage (10%-20%)
- ASHRAE A1 temperature environment (allowing excursions up to 32 degrees)
 - Did consider A2 (excursions up to 35 degrees) but decided for A1
- Should provide its own BMS
- Access for delivery trucks and fire brigade
- The building is considered to be a lights out facility, not open to public visits

Agreed location and footprint



PTS FOR REFERENCE PO

	X	Y
	248919	12321
	248921	12329
	248920	12328
	248947	12328
	248921	12326
	248915	12328
	248948	12328
	248922	12329

SMB-SE
ENGINEERING
GWL

BDG 775-DATA CENTER SETTING OUT LOCALISATION
ECHAËLE SCALE 1/500

PLAN DIRECTEUR*
BDG 775-CENTRE INFORMATION LOCALISATION

GENERAL MODIFICATION
CREATION DU PLAN / PLAN CREATION

IND. DATE NOM/NAME ZONE MODIFICATION

A 22/11/16 CAMEHARD
- 31/01/16 CAMEHARD

PROJON

NOT VALID FOR EXECUTION
NON VALABLE POUR EXECUTION

DATE
01-16

* Ce dessin ne peut être utilisé sans la commande sans qualification écrite. This drawing may not be used for contracted purposes without written authorisation.

Summary

- Wigner contract is finishing but no PCC available for RUN3
- But fortunately, LHCb offered us the use of 2 containers for RUN3
- PCC project has been relaunched and has Management Support and financing
- PCC tender next year and should be available by the end of RUN3 and hence provide the required capacity from then on

