

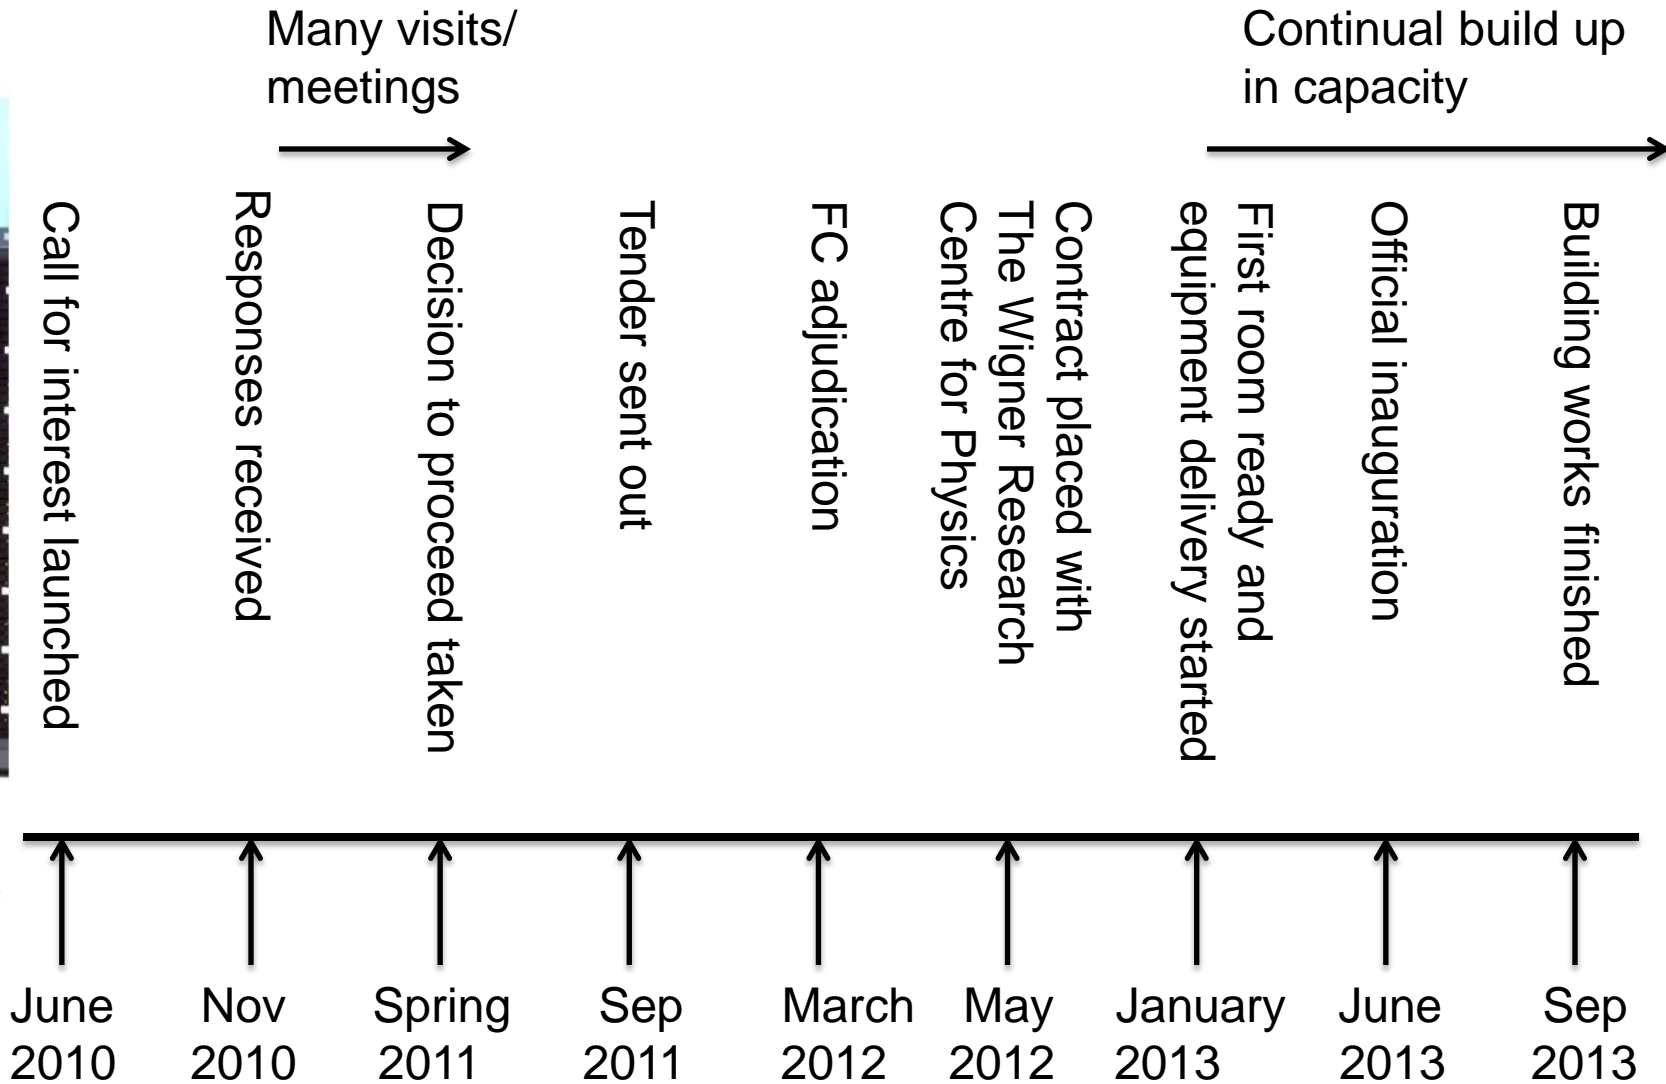
CERN Remote Hosting First Experiences

Wayne Salter

(with input from many colleagues)

HEPiX Autumn Meeting in Lincoln

- Brief History
- Installation Status
- Experience
 - General
 - Commercial
 - Procurement
 - Operations
 - Networking
 - End User Utilisation
- Lessons Learnt
- Conclusions



CF

Brief History in pictures

CERN IT
Department



CERN IT Department
CH-1211 Geneva 23
Switzerland
www.cern.ch/it

HEPiX Autumn Meeting 2014 @ Lincoln - 5



- Two rooms are in operation for CERN with 122 racks used
- 1276 CPU servers – 319 2U quads (25216 cores, 85504 GB RAM, 5904 TB disk)
- 568 external storage units – 4U JBODs each with 24 disks
(52608 TB in total - 1920 TB on 3TB drives and 11712 on 4TB drives)
- Network equipment; 7 high end routers, 43 10GbE and 47 1GbE switches, 1 management router and 107 management switches
- Additional large deliveries expected in December
 - More than doubling CPU capacity and adding 40% more disk storage capacity and requiring use of third room
 - Investigating possibility of having a 3rd 100Gbps link

- On the whole good – generally works well
 - Remote operation and monitoring works well
 - No out of hours support for CERN equipment
- Teams visiting each other was very useful
 - Help given with initial setups
- Over reliance on one person
- Reporting
 - Regular bi-weekly operational telecom
 - Monthly reports (since 2014)
 - Operations and Billing
- Can be time consuming dealing with new requirements, e.g. Russian Tier 1 link

- **Tendering process**
 - Specification as open as possible
 - Adjudication based on a defined ramp up profile, failure rate estimates, and included networking from closest GEANT PoP
- **VAT Exemption**
 - Took many months to sort out and required help from Wigner
- **Insurance split**
 - Discussions still on-going!!
- **Billing**
 - First bill only in 2014 after more than one year of running
 - Detailed spreadsheet as part of monthly operations report

- Detailed instructions to ease reception and installation
 - However, following deliveries is more complex
- Delivery directly to Wigner except for network switches
 - One case of damaged equipment during transport
- Need to provide detailed information in advance on deliveries as well as transport
- Issues with unloading of equipment at Wigner
- Effectively doubled the number of orders to be processed

- Late availability of room for storage and repairs
- Auto-registration and stress testing of machines works well
- Room/rack layout responsibilities ‘unclear’
- Various infrastructure issues
 - Two HV incidents but protected by UPS/diesel
 - Cooling pressure issue causing all chillers to be switched off
 - Leak in cooling pipe
 - Complex new facility not completely understood. Review conducted by TÜV
 - Often slow to get detailed reports

- More difficult than expected to establish good workflows
- Formal procedures and approach being gradually introduced as experience is gained
- Difficult to use full available power
 - Tender estimate of power density does not reflect the reality
- Difficult to verify power consumption figures
- Non-standard setups and debugging of tricky issues are more complicated

- Role of the SysAdmins has not been affected
- Repair Service
 - Runs well
 - Good quality interventions
 - Good response time to SNOW tickets
 - Information flow is more complicated with more parties involved – has not been ideal
 - Data requested not always provided in a timely manner
- Still very limited usage of Wigner for business continuity
 - Lack of second network hub
 - Priority on moving to new critical room at CERN
 - Difficulties in getting allocation of resources for BC

CF

Experience - Networking



- Long discussions on initial network setup in the rooms
- Takes longer to solve simple problems/lot of mail exchange/no out-of-hours support
 - Required changes to operational approach
 - Now giving Wigner access to SPECTRUM monitoring
- Less time for deployment of new equipment (for CERN)
- Availability of 100Gbps links not as expected
 - Long running problems with one of the links (took many months to debug)
 - Over past 100 days; link 1 (99.7%), link 2 (99.96%)
- Broken equipment takes longer to be replaced by manufacturer
 - Try to minimize the number of shipments
 - Shipments must come via CERN

- Complaints of performance of jobs at Wigner
- However
 - Mixture of SLC5/6, Intel/AMD, VM/Bare metal
 - Different type of jobs
 - Locality of data
 - Optimisation of S/W for Intel whilst most CPU server in Wigner were AMD
 - Configuration options, e.g. XROOT TTreeCache
- When comparing like with like only a minimal drop in efficiency
 - However, EOS servers deployed to Wigner
 - Will soon deploy CVMFS service at Wigner
- Investigations are still on-going

- New facility and hence some teething problems as well as one design issue
- Lack of experience on both sides
 - but due to collaborative and flexible approach issues have generally been resolved quickly
- Personal contact is VERY important
 - Help with first installations
 - Teams meeting each other
 - Regular telecoms
- Good communication is important
- Good documentation helps a LOT
 - Still need to improve SLA and other formal arrangements
- Things always take longer than foreseen

- In general everything is running smoothly
- Issues have arisen
 - But in general have been resolved quickly due to flexibility and good relations on both sides
 - VAT and insurances have taken longer due to external parties
- 100Gbps links have not been as stable as expected
- Some questions raised regarding job efficiency
- Full power capacity usage will not be possible due to lower power density than expected
- With experience it should be possible to produce more detailed formal documents next time (....)
- Still waiting to implement more extensive BC
- Contract due to run until end of 2019

The letters 'CF' in a large, white, sans-serif font, positioned in the top left corner of the slide. The background behind the text is a vertical strip of images showing server racks and network equipment.

Thank you for your attention!

CERN IT
Department

- Questions?

