



CF

Computing Facilities

CERN
IT
Department

CERN Computer Facilities Evolution

Wayne Salter
HEPiX May 2011

- Reminder on the current status of CERN Computer Facilities
- Overview of the current issues and anomalies
- Summary and status of the various Evolution Projects
- Closing remarks



- Designed and built in early 1970s
- Fully refurbished around 2000
 - Increased power and improved cooling
- Nominal current capacity: 2.5MW (including 240kW of critical power)
- Extended (usable) capacity: 2.9MW (including 340kW of critical power) forfeiting redundancy for all UPS systems
 - But need to take action in the event of a UPS module failure!
- Small capacity at local Hosting Centre (17 racks and up to 100kW)

- Cooling for critical UPS room insufficient
- Mixing of critical and physics equipment
 - Location and cooling
- No cooling of CC when running on UPS and insufficient stored cooling capacity when on diesel
- Insufficient critical power available
- Approaching the limit of available power for the building
- No redundancy for critical UPS
- Usage of full 'available' 2.9MW implies loss of redundancy for physics UPS
- A/C of CC coupled to adjacent office building
- How to meet CERN's needs in longer-term

- Local hosting
 - Provide additional ‘critical power’
 - Allow some level of business continuity
- On-going improvements of CC
 - Improve efficiency and resilience of cooling system
- Upgrade of current computer centre
 - Increase capacity from 2.9 to 3.5MW
 - Increase ‘critical power’ from 340 to 600kW
 - Address a number of long-term issues
- Remote hosting
 - Address the increasing capacity needs of CERN
 - Increase business continuity coverage

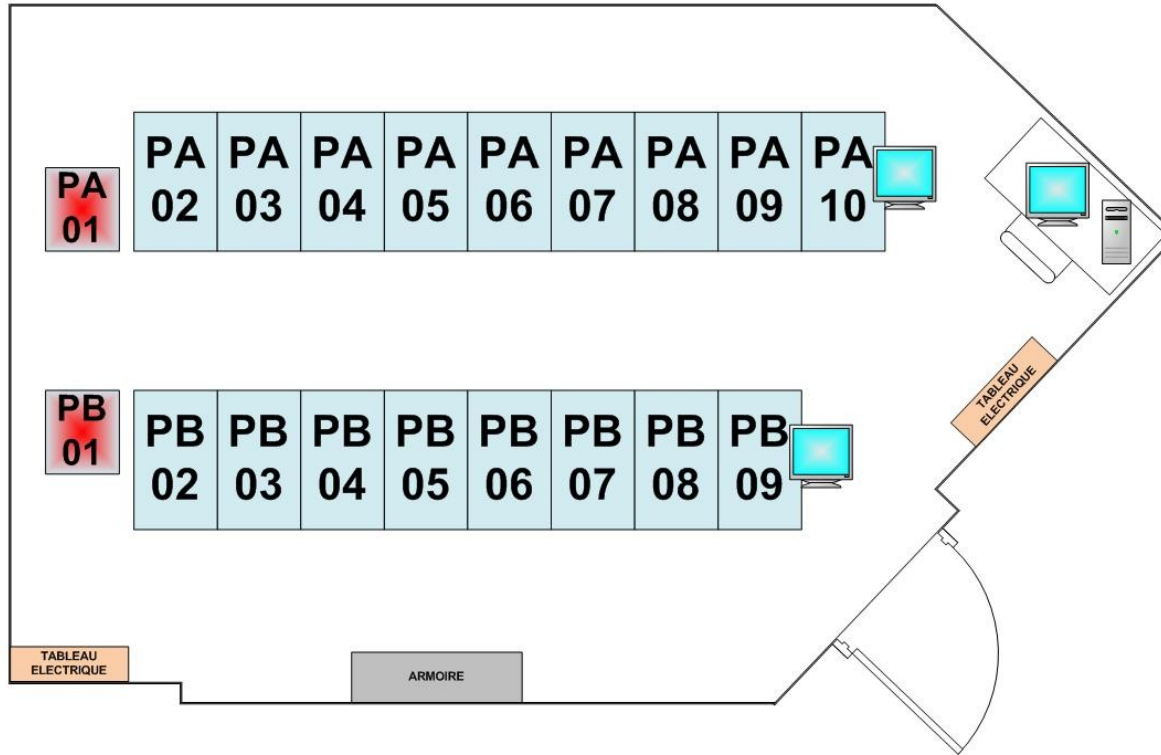
- Reason:
 - Lack of ‘critical power’
 - Provide some level of BC
 - Gain experience with remote hosting
- History
 - Price enquiry for hosting – August 2009
 - Tender for networking – October 2009
 - Planned start – April 2010
 - Actual local hosting contract start – 15th June 2010
 - Connectivity contract start – 1st July 2010
- Contracts for local hosting and network connectivity renewed for a second year


- Agreement for 40m² (actually 36m²) and 100kW (two power feeds)
- 2 dark fibres on diverse paths
- 17 IT racks + 2 network racks
- 209 systems installed
- 14 racks used and 57kW (average power density of 4.1kW/rack)
 - 3 racks remaining and 43kW!
- Interventions at Local Hosting site:
 - Small number of Sys Admin interventions (2-3)*
 - 7-8 Service Manager interventions*
 - 33 Vendor interventions
 - Many installation interventions

* Worrying in terms of 'real' remote hosting!

9994-4

(Update: 30 juin 2010)





 Rack
100x60 cm

 Rack 60x53 cm
Starpoint CS

- Everything took longer than foreseen!
 - Contracts with hosting company and network provider
 - Preparation of CERN room at hosting company
 - Partitioning, power and fibre connections, access card reader, access to InsideEyes, ...
 - Getting equipment into production
- Network connectivity expensive
- Room smaller than expected (36 c.f. 40 m²)
- Problems with ramp
- *However, no significant problems running systems remotely as far as we have seen so far*

- Despite a few teething problems the experience is generally good
 - However, due to proximity not everything done remotely
 - Need to understand reasons and how to avoid
- Still not at full capacity
- Struggling to utilise full available power
 - Low average rack power density ($\sim 4\text{kW}$)
- Good step towards remote Tier0 hosting

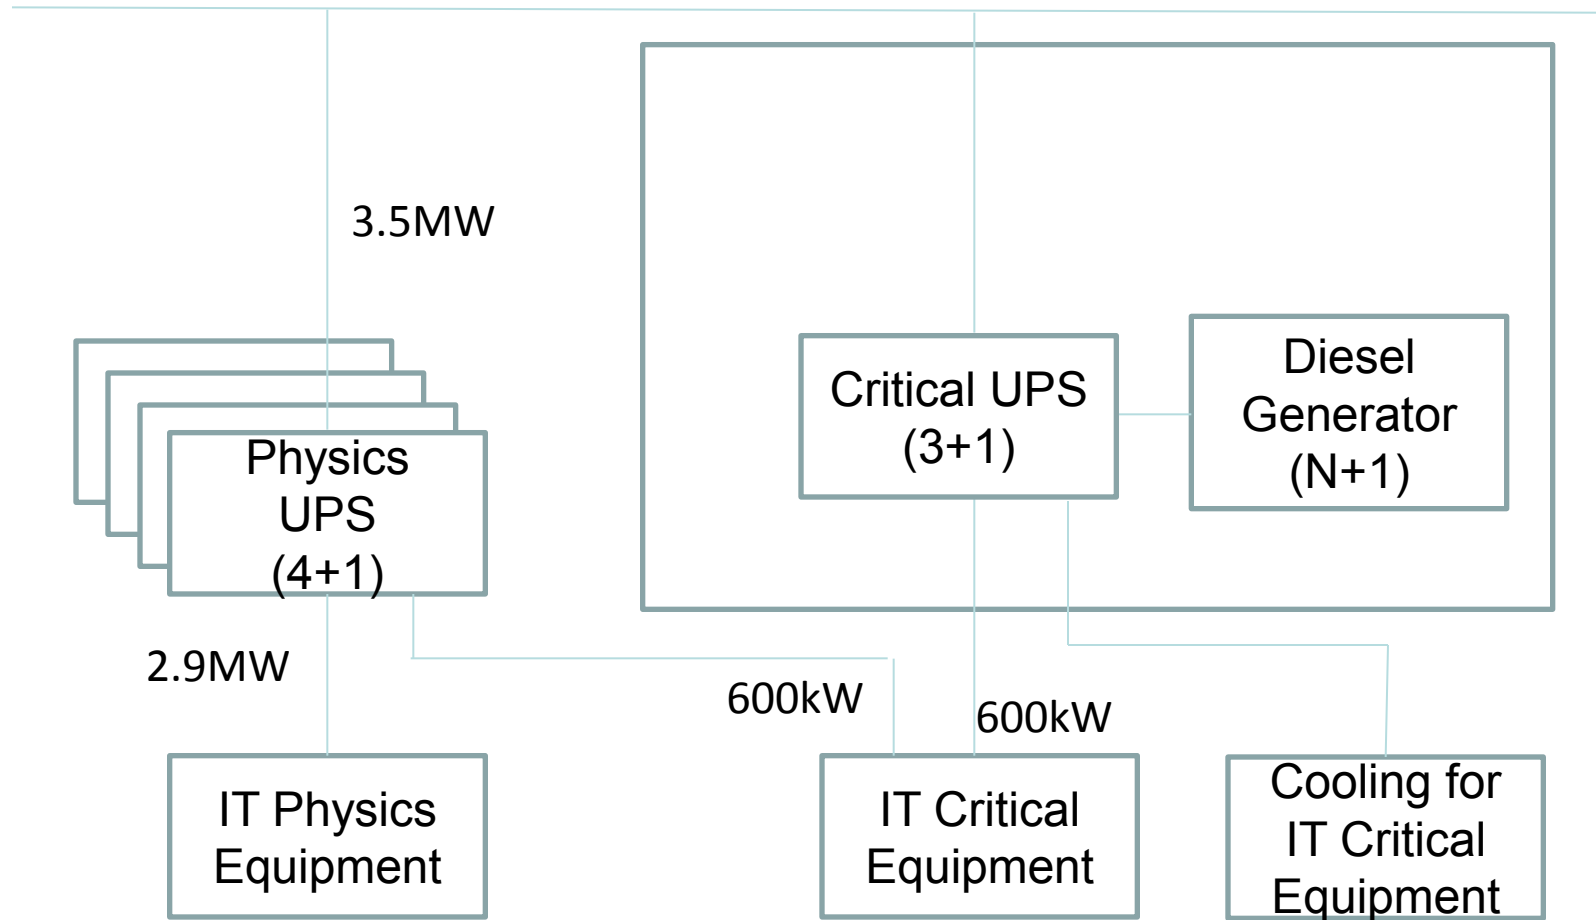


← Wind direction

- Work already done
 - Improved monitoring (allows us to calculate PUE fairly accurately)
 - Correction of air input temperature
 - Reduce intake of hot exhaust air
 - Improve mixing of outside and re-circulating air
 - Full automation of air selection (% of outside vs. re-circulating)
 - Change to anti-freeze protection (delay and better mixing)
 - Above measures predicted to result in 80% saving for chiller power
 - All control system components connected to UPS
- Foreseen potential further improvements
 - Higher server inlet temperature
 - Understand the pressure drop from AHU to servers
 - Optimize air flow (across all aisle)
 - Mixing depending on RH
 - Use variable speed fans for AHU to reduce air flow to only what needed

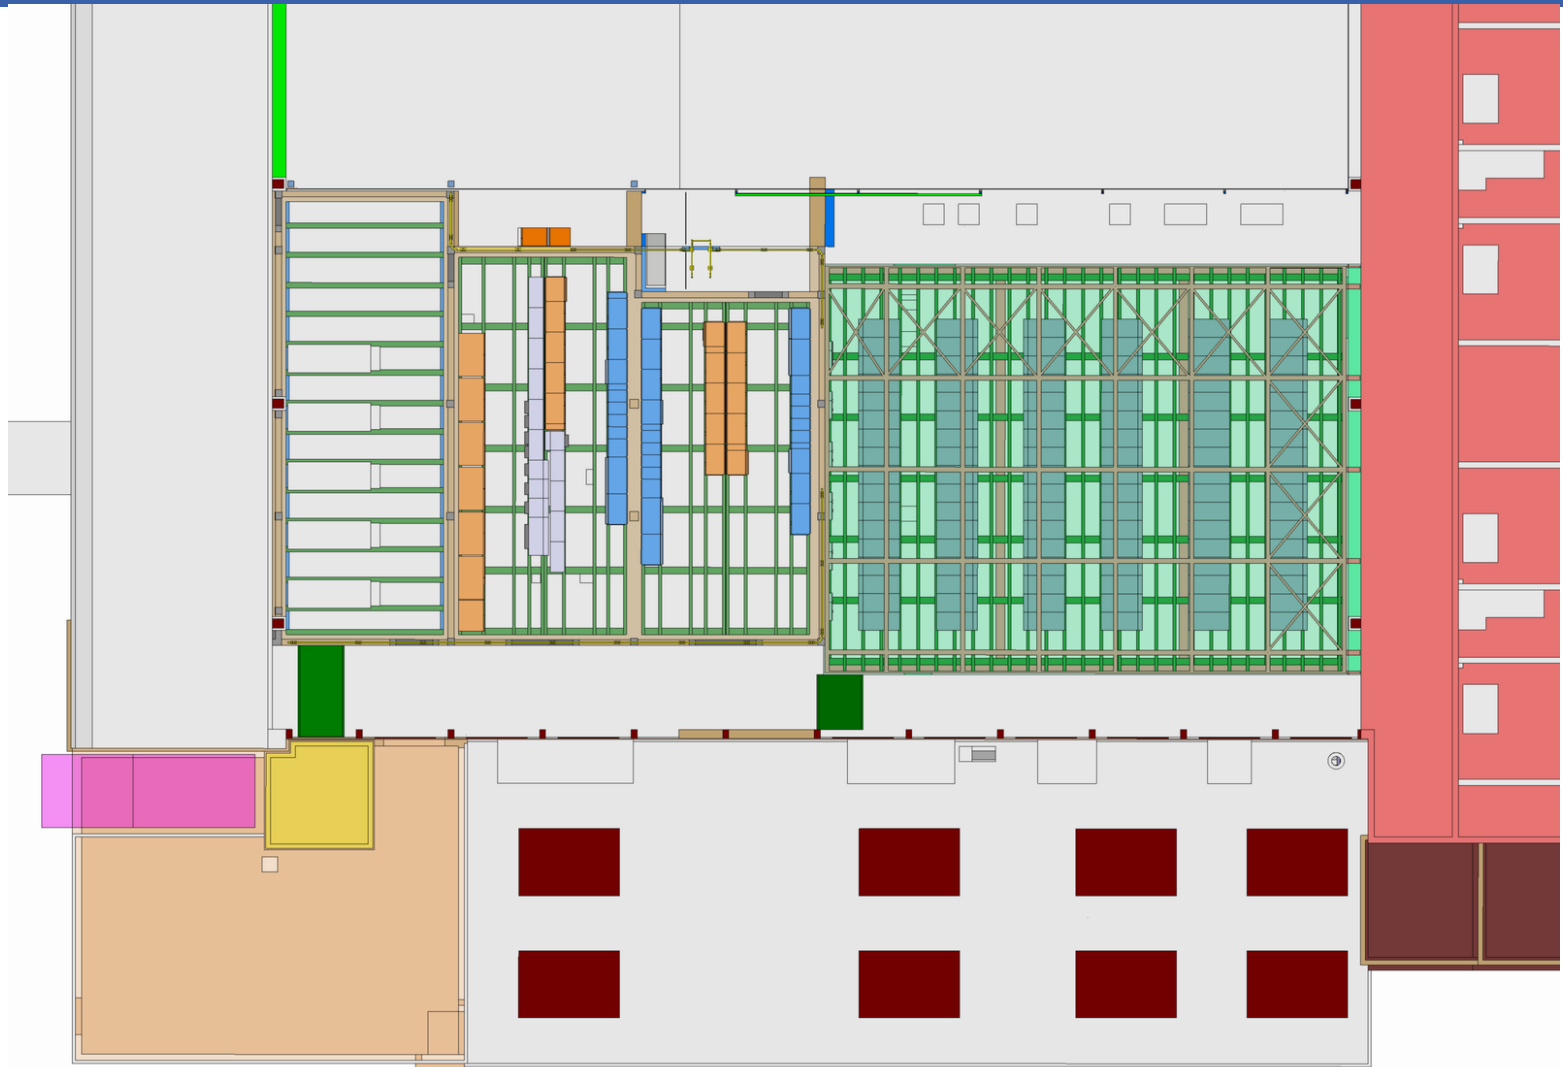
- Solve the cooling issue for the critical UPS room
 - New UPS systems in a different location
- Increase critical capacity to 600kW
- Increase overall power capacity to 3.5MW
- Restore N+1 redundancy for both critical and physics UPS systems
- Secure cooling for critical equipment when running on UPS and extend stored cooling capacity for physics when on diesel
- Decouple the A/C for CC from the adjacent office building

Normal Network





- Barn area of CC building to be converted to house:
 - 3 new electrical rooms
 - An IT room with water cooled racks to house critical equipment (up to 450kW)
 - New ventilation systems for the critical area of the main computer room (and also the critical UPS room in basement)
- New ventilation systems for the Telecoms rooms (CIXP) will be installed in adjacent offices
- New critical UPS systems in basement
- A new partially sunken building for additional chillers and a storage tank for the cooling of the critical areas
- An additional storage tank for extending the stored cooling capacity for physics equipment
- Opportunity to install emergency evacuation stairs

CF





CF






CF





CF



- 'Barn' cleared of IT equipment end of October
- Removal of cabling and ducting finished
- Civil engineering (CE) work just commencing
 - Delayed due to difficulty in freeing offices
- CE works to be completed November 2011
- EL+CV installations Nov/2011-Nov/2012
- Increased physics power available Aug 2012 and increased critical power Nov 2012
- *Project takes a long time and has high cost!*

[Barn Video](#)

- How to provide resources once CERN CC full?
- Studies for a new CC on Prévessin site
 - Four conceptual designs (2008/2009)
 - Lack on site experience
 - Expensive!
 - Lack of support from management
- Interest from Norway to provide a remote hosting facility
 - Initial proposal not deemed suitable
 - Formal offer not convincing
 - Interest from other member states

- Call for interest at FC June 2010
 - How much computing capacity for 4MCHF/year?
 - Is such an approach technically feasible?
 - Is such an approach financially interesting?
 - Deadline end of November 2010
- Response
 - Surprising level of interest – 23+ proposals
 - Wide variation of solutions and capacity offered
 - Many offering > 2MW
 - Assumptions and offers not always clearly understood
 - Wide variation in electricity tariffs (factor of 8!)

- Visits
 - Visit to many sites
 - Others invited to CERN
- Goal
 - To understand better the proposals
 - Clarify CERN's needs
- Benefits
 - See existing installations
 - Triggered us to reconsider some of our ideas/preconceptions
 - Allowed consortia to understand better our needs
 - Collect information for technical specification

- Proposed timeline
 - Official decision on whether to go ahead spring 2011
 - Official letter to be sent explaining procedure
 - Tender during summer 2011 for adjudication end 2011/early 2012
 - Initial installation first half of 2013 to test operational model
 - Gradual build up in capacity in-line with experiment needs

- CERN CC reaching end of capacity
- Further improvements in CC are possible and being implemented in parallel (e.g. monitoring and cooling)
 - Better monitoring capabilities
 - Better efficiency, but
 - We do not get additional computing capacity from this!
- Three options to address providing increased capacity:
 - Local hosting
 - Limited additional capacity
 - CC Upgrade to 3.5 MW
 - Slow and expensive
 - Remote hosting
 - Interesting but introduces new challenges
 - But could allow us to address business continuity properly



ent