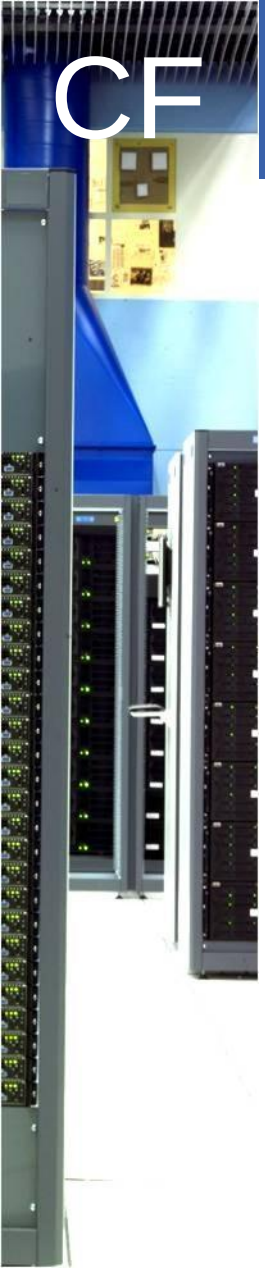

CERN Business Continuity Overview

Wayne Salter

HEPiX April 2012

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

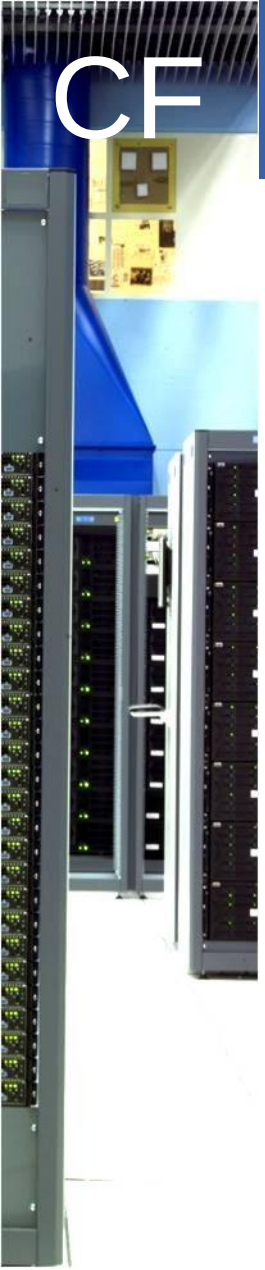

Overview

- What is business continuity?
- What steps can be taken?
- What is the status at CERN?
- What are our current plans?
- Conclusions

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

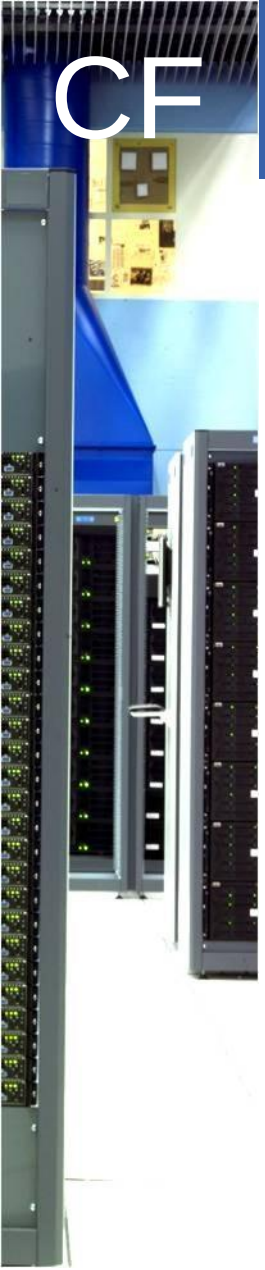

What is business continuity?

- **Business continuity** is the activity performed by an organization to ensure that critical business functions will be available to customers, suppliers, regulators, and other entities that must have access to those functions.
- **Disaster recovery** is the process, policies and procedures related to preparing for recovery or continuation of technology infrastructure critical to an organization after a natural or human-induced disaster. Disaster recovery is a subset of business continuity.

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

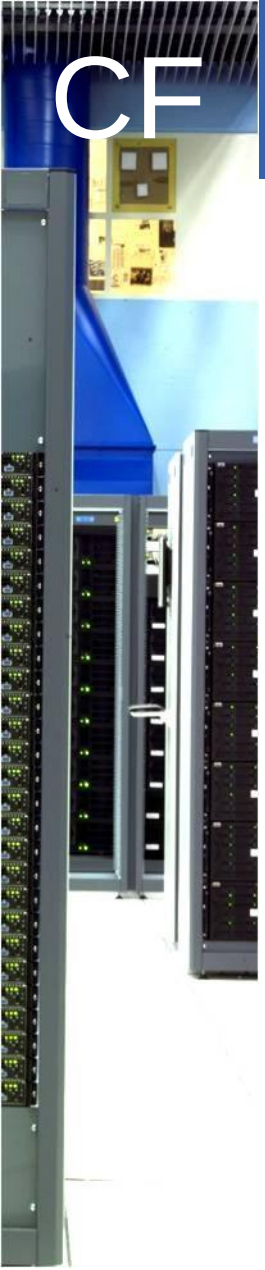

What is business continuity?

CERN IT Department

- **Business continuity** means that services can be provided with an acceptable level of availability (SLA/SLD….).
  - Level of availability will vary between services
- **Disaster recovery** means that services can be restored after a major incident within an acceptable amount of time.
  - Time will vary between services

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

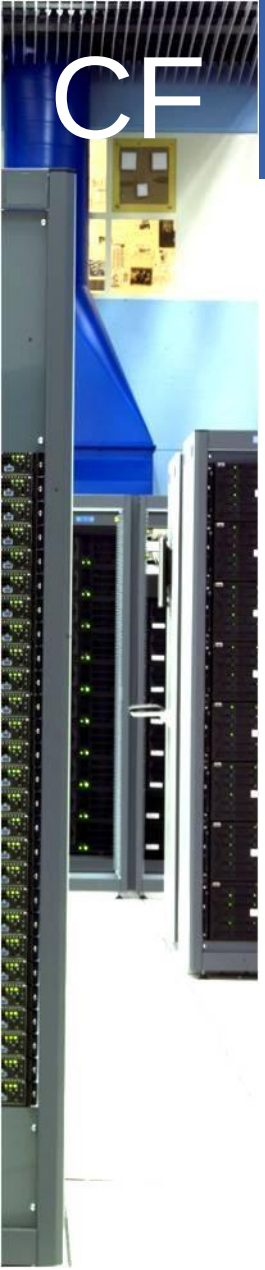

Possible Types of Incident

- Deletion or corruption of a file
- Failure of a disk
- Failure of a server
- Failure of network equipment
- Localised incident over many servers
- Incident effecting a large part of data centre
- Malicious action
- Preventative or corrective maintenance

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

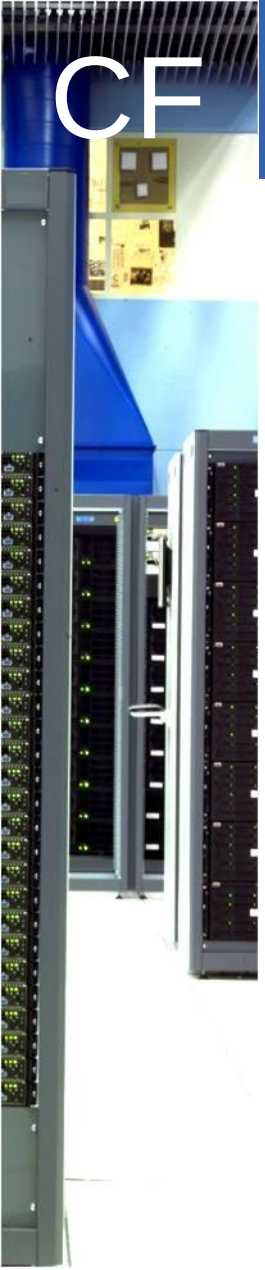

Other Issues to Consider

- When looking at BC for a particular service are all dependencies properly handled?
  - E.g. we have redundant systems in two centres but with common networking
  - How can this be properly verified?
    - Repeated testing!!
- What is actually required in terms of BC?
  - Instant failover?
  - Can one accept a human triggered failover?
  - Can one accept a re-instantiation of service?
  - Is load balancing with loss of capacity acceptable?

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

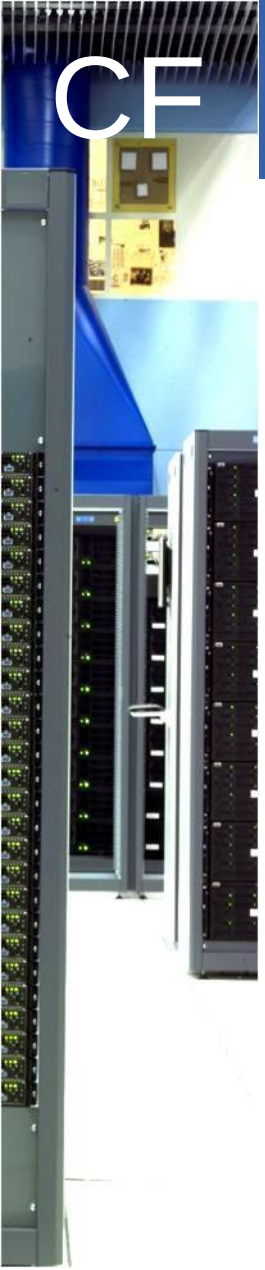

What are typical steps that can be taken?

- Disaster recovery
  - Ensure all software, data, etc. require to provide the services is stored in more than one location
  - The hardware infrastructure can be recreated
    - Try to avoid tight coupling between applications and hardware
  - The applications (with associated data) can be re-established
  - Knowledgeable personnel are available and/or good documentation!

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

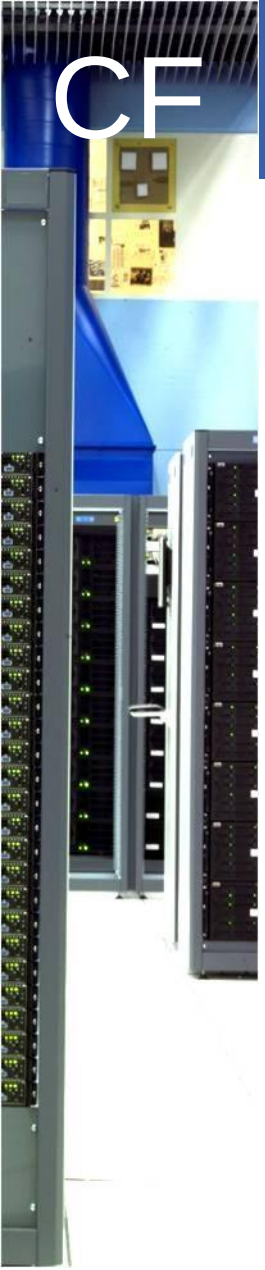

What are typical steps that can be taken?

CERN IT Department

- Business continuity
  - Establish adequate redundancy
  - Where applicable use UPS/Diesel systems (also with appropriate redundancy)
  - Service spread over multiple locations in a centre or across centres
    - Servers, racks, PDUs, power feeds, aisles, rooms,
  - Implement redundant services
  - Ensure the capability to recreate services within a ‘short period of time’
    - Infrastructure, H/W, S/W, personnel, …
  - Good cyber security measures

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

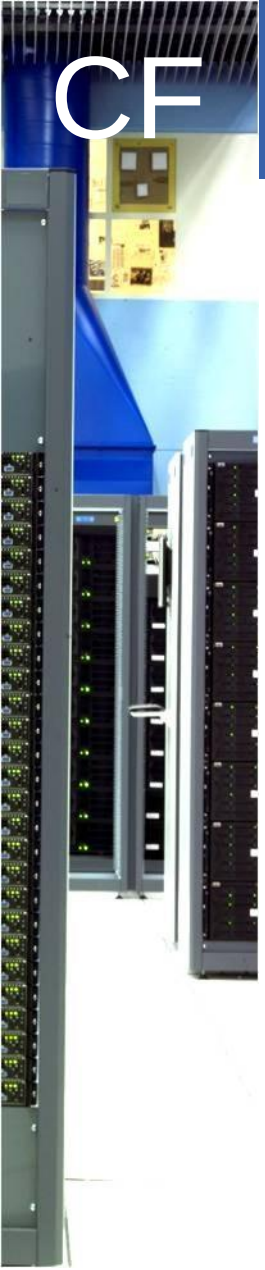

Current Measures @ CERN

- Disaster recovery
  - Critical data backed up to separate building(s)
  - All tools, scripts and configuration data backed up to a separate building
  - Some services running in local hosting site
  - Able to rebuild services on new hardware
  - But where??

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

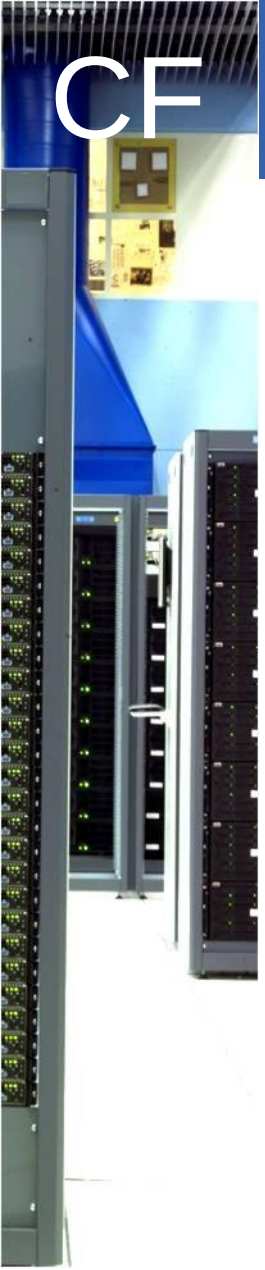

Current Measures @ CERN

- Business continuity
  - Redundancy
    - RAID (H/W or S/W)
    - Multiple power feeds to CC
    - Two independent LV power distribution systems in CC; critical (UPS+diesel), physics (UPS only)
    - Many servers have multiple PSUs (all critical servers)
    - In UPS systems themselves
    - HVAC (chillers, AHU and pumps)
    - Multiple rooms (currently 3 at CERN and soon 4)
  - At service level (including core networking equipment)
  - Distributed star points across the site (but all core equipment in CC)
  - Local hosting site
    - Not primarily for BC
    - Also requires networking in CC

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

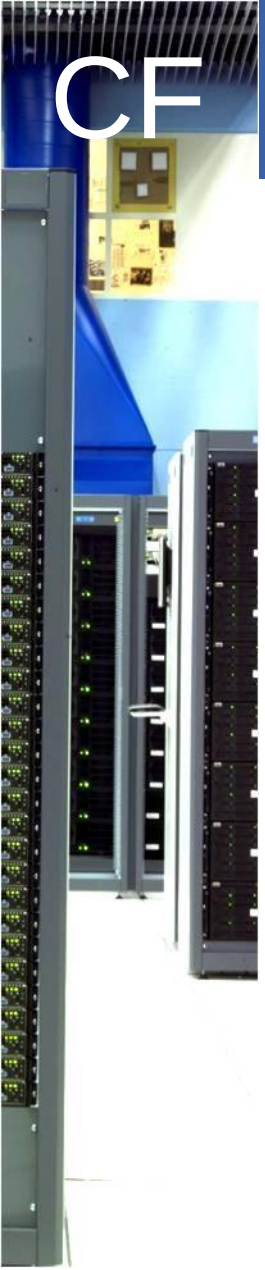

Example 1

- For batch:
  - Many batch nodes
  - Spread across the CC
  - On different PDUs
  - On different UPS systems but only on physics (even machines with multiple PSUs)
  - 2 LSF masters with hot standby and these are on critical/physics power
    - Sharing disk-based state via a PES-managed Netapp appliance (also on critical/physics power)

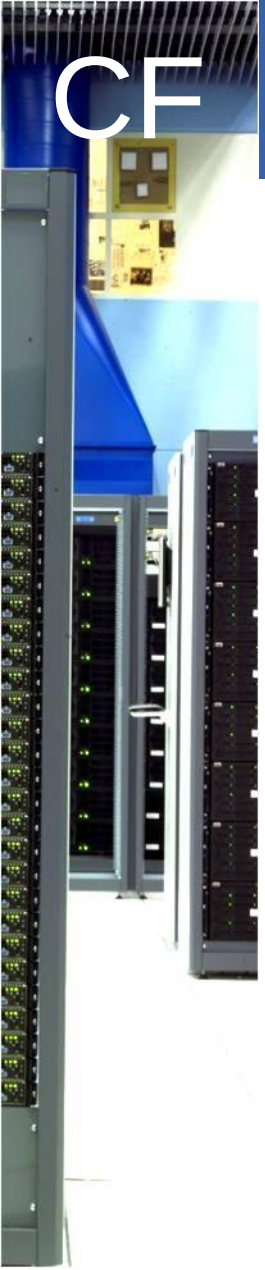

Example 2

- For databases the following approach is taken:
  - Data is backed up to TSM (two locations). Plan to move part of the backup to disk.
  - Data is replicated to a different location (local hosting site) using Data Guard (applying the redo log). Allows queries to be made on replica and hence offload production
  - Storage on critical/physics power but only one out of two servers on critical/physics due to lack of sufficient critical power capacity
  - Replication done asynchronously but continuously (most of the time <3 seconds - data streaming)
    - Risk of losing a few seconds of transactions
    - Looking at doing this synchronously without loading the production system
- Validation of backup, creation of logical copies
  - Reload data from TSM to a clean server and apply all redo logs
  - Make a 'logical copy' to a different location and backup to TSM
    - Allows to go back to a particular point in time independently of Oracle version. e.g. financial data stored for 10 years.
    - Allows to verify all backup and recovery scripts plus establish the time to make a recovery

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

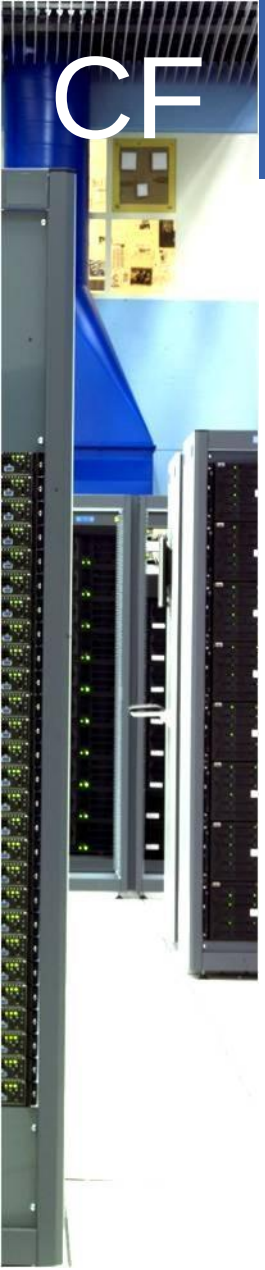

Example 3

- Active Directory
  - Servers at CERN and local hosting site with automatic failover (MS facility)
  - At CERN servers on critical and physics power
  - At local hosting site servers on dual feed
  - All data backed up to TSM (two locations)
- Mail Servers
  - Some mailboxes at CERN and some at local hosting site with mirror at other site
  - MS log shipping used for synchronization
  - Automated failover (MS Exchange feature)
  - At CERN servers on critical and physics power
  - At local hosting site servers on dual feed
  - All data backed up to TSM (two locations)

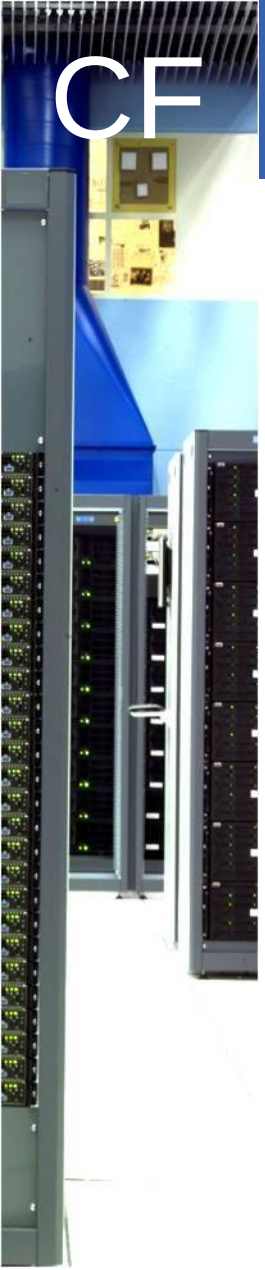

Example 4

- DFS File Servers
  - Home directories split between CERN and local hosting site with backup at other site
  - MS DFS Replication used for synchronization
  - Workspace folders only at CERN but with backup synchronized using homemade script
  - At CERN servers physics power only (due to lack of available critical power)
  - At local hosting site servers on dual feed
  - All data backed up to TSM (two locations)

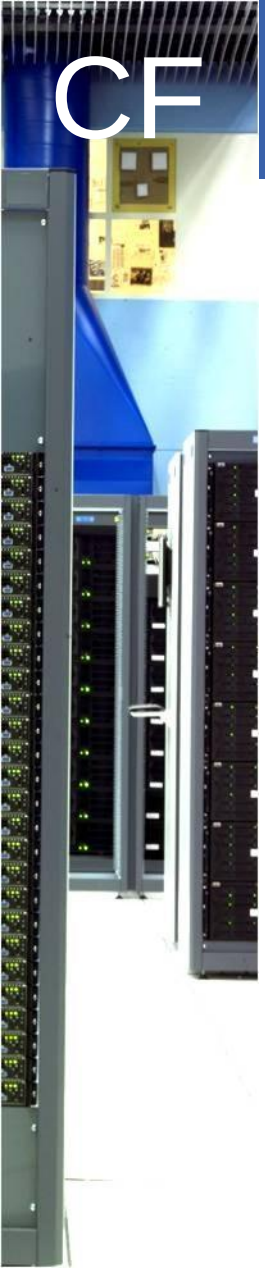

Remaining Major Risks @ CERN

- Extended computer centre power cut
  - No clean separation between critical and physics equipment and their cooling
- Major incident in computer centre building, e.g. fire
  - In an electrical or HVAC room
  - In one of the machine rooms
  - In the telecoms room
- Incident in the fibre room
- Are all dependencies handled correctly?
  - Can only check with testing on a regular basis

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

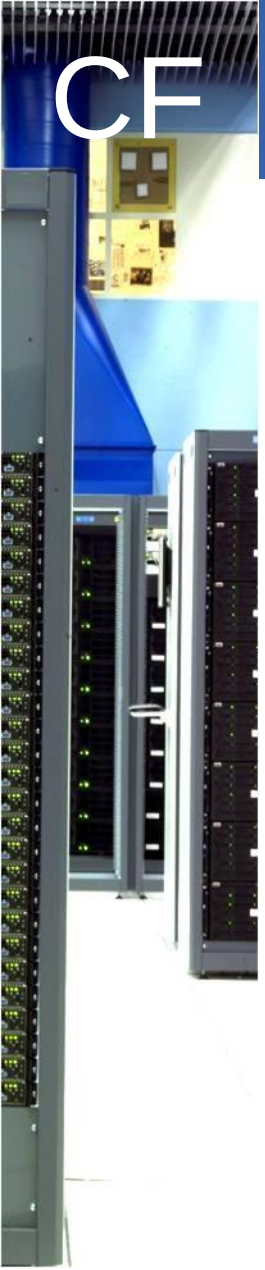

Future Plans

- Securing critical cooling in CC upgrade project
- Remote hosting
  - Duplicate critical services
- First classification of services against three BC options:
  1. Backup
  2. Load balancing
  3. Re-installation
- Second network hub
  - Core networking equipment and connection to distributed star points
  - One of the two links to each of the remote centres
  - Commodity and GEANT Internet connection
  - LHCOPN and LHCONE connections
  - Core networking services, e.g. DNS, DHCP, monitoring, etc.
- Testing
  - Establish test procedures and perform regular tests

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

Possible Network Topology

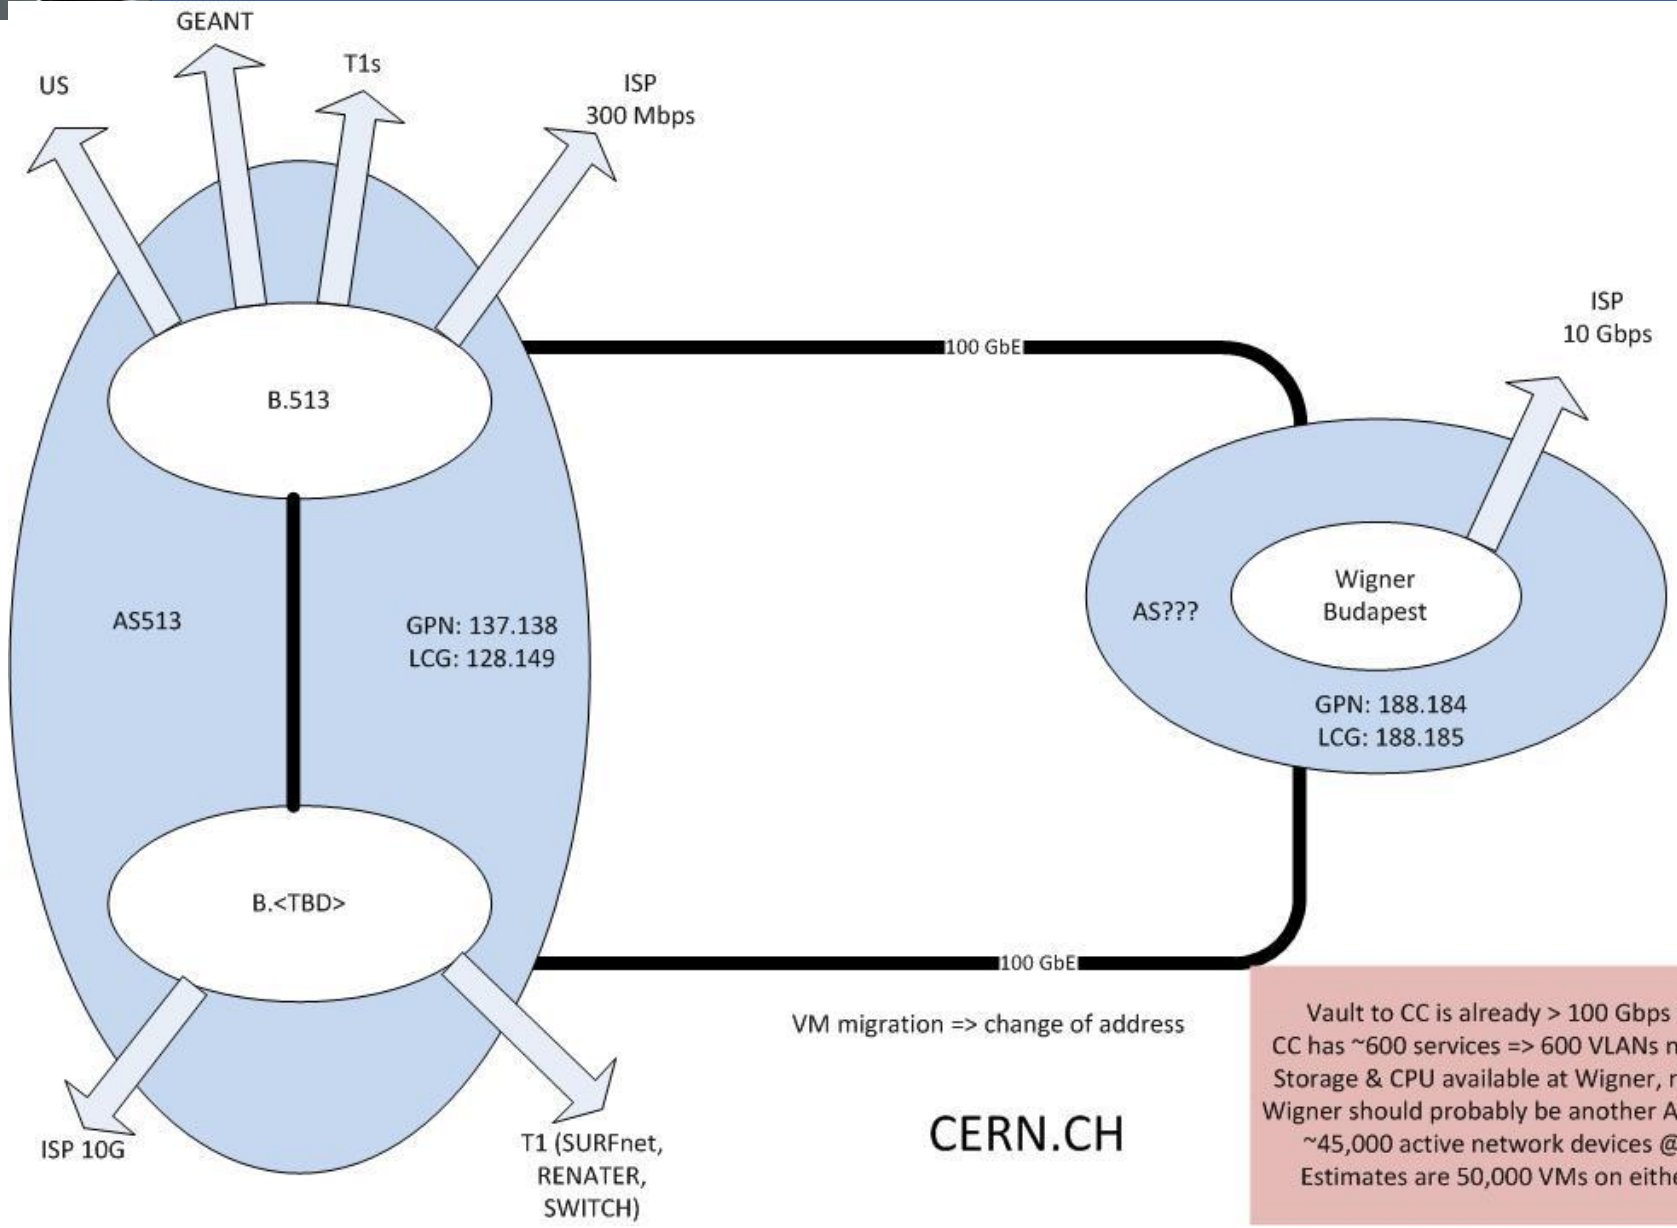

Vault to CC is already > 100 Gbps today
CC has ~600 services => 600 VLANs necessary
Storage & CPU available at Wigner, not tapes
Wigner should probably be another AS number
~45,000 active network devices @CERN
Estimates are 50,000 VMs on either side

Second Network Hub

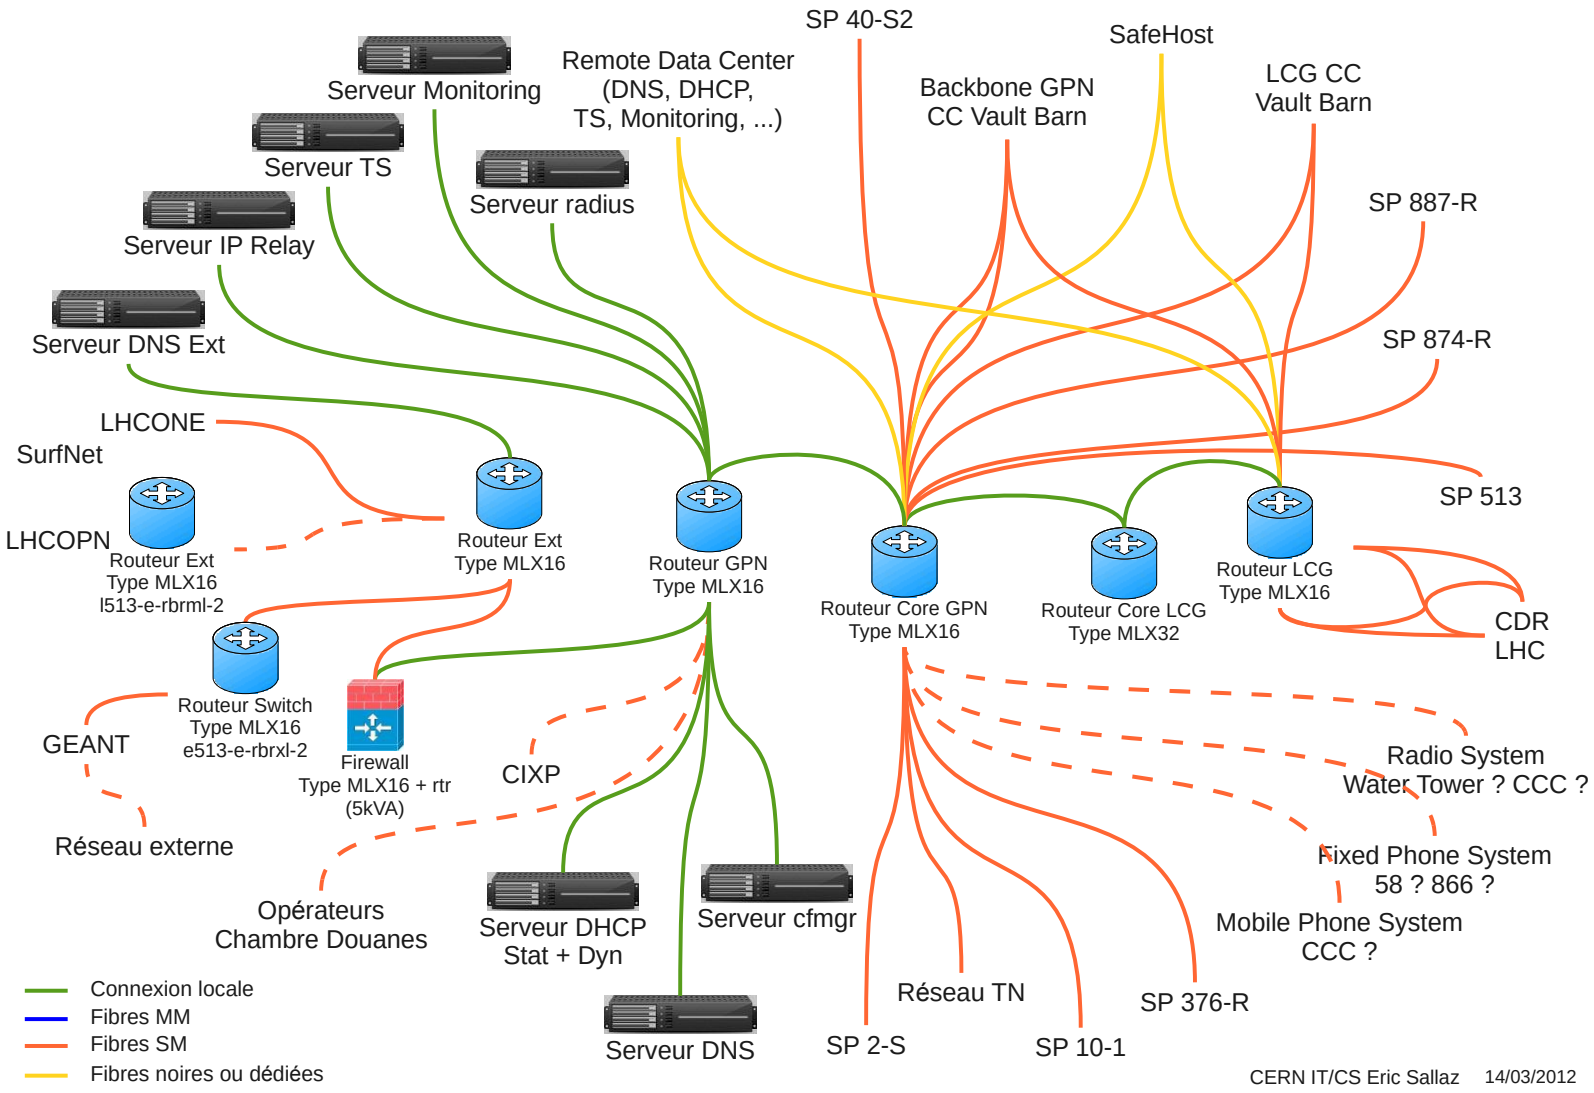

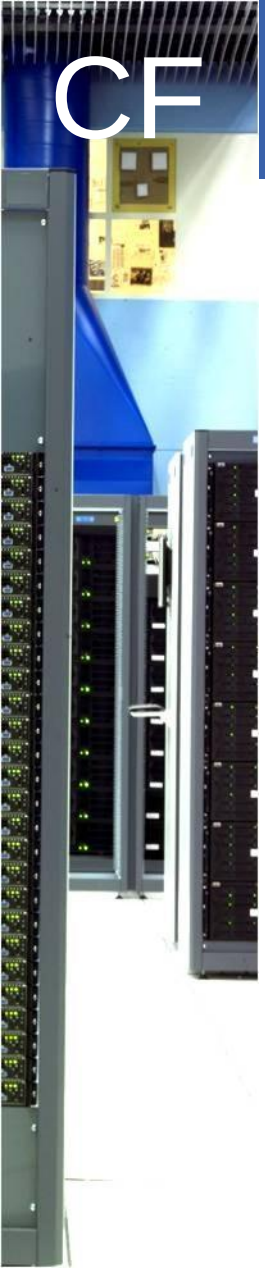

Open Issues

- Full classification of services
  1. Backup
  2. Load balancing
  3. Re-installation
- Latency
  - Expect 20-30ms latency to remote site
  - A priori no known issues but a number of services will need to be tested
- QoS
  - Ideas on implementing a first order ‘QoS’ to avoid network problems in case of peak network loads. e.g. splitting TCP and UDP traffic
  - Needs testing
- Backup strategy at remote site; tape robot, dark disks, cloud storage backup
- Cost

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

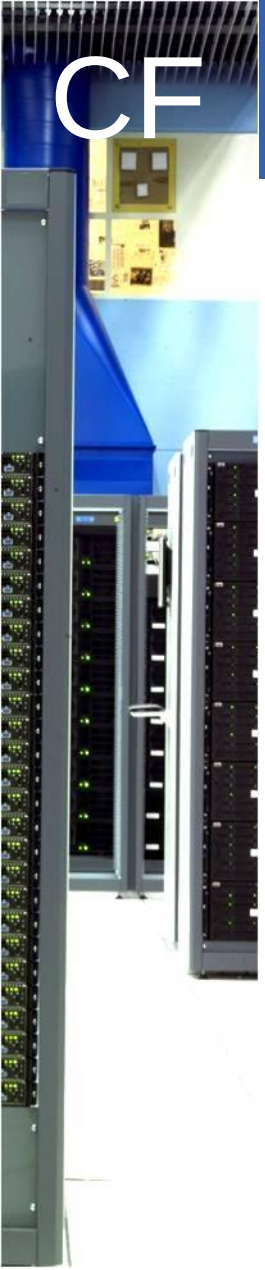

Conclusions

- CERN's approach to business continuity up to now has been largely based on applying a good level of redundancy and backing up data to a separate location
- However, this still leaves the risk of a major incident in the computer centre
- With the new remote centre CERN we will look to implement a more complete business continuity strategy
  - Not only by implementing critical systems in both locations
  - But also by creating a second network hub at CERN

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