

CERN's Business Continuity Working Group

Helge Meinhard / CERN-IT
HEPiX Autumn 2020 in Zoom land
15 October 2020

What Happens to Your Organisation If...

- There is a fire in one of your data centre rooms?
- A floor of your data centre is seriously flooded?
- A tree falls on the main power line supplying your data centre?
- Your organisation is attacked by ransomware?
- An aircraft or lorry crashes into your data centre?
- ...
- A storage cluster crash brings down a huge number of databases, corrupting some of them?
- A block storage service crash brings down a large number of services, corrupting some of their data?
 - And that way, you learn about dependencies you did not know about?
- ...



What Happens to Your IT Department if...

- ... anything like that happens

and

- ... there are several data centre rooms or buildings

or

- ... there are external resources (e.g. clouds) available?

Does your IT department look stupid? Do you?



At CERN...

- ... we have asked ourselves these questions several times
 - In particular when new options became available: The (past) Wigner-Budapest extension to the WLCG Tier-0, the (future) PCC (Prévessin Computer Centre)
 - Even though possibly inspired by these options, the question is more general
- CERN's IT Management decided in 2019 to look at the problem in a more systematic way



Business Continuity Working Group

- April 2020: CERN-IT formed a working group
 - 10 experienced people representing all technical groups, computer security etc.
 - Mandate:
 - Formulate recommendations targeted at increasing resilience of services during disruptions
 - Reduce impacts of disruptions
 - Provide plans for restoring services



Dimensions of BC

- User view: business case
 - How important/critical is a given service?
- Service owner / service manager view: service implementation
 - Dependencies, effort to restore, ...
- Overarching dimension: threats
 - Surely they influence the service view
 - But may impact the user view as well – a major disruption may justify longer time to restore

Business Cases

- Potential classification of services:
 - Emergency services (back within one hour)
 - Essential administrative services (back within a day)
 - Basic working services (back within a few days)
 - Services required by external sites (back within a week)
 - Services required for accelerator running (back within one or two weeks)
 - First level for normal business (back within several weeks)
 - All other services (back within a month or even more)
- Unlike threats and dependencies, need interaction with users (and budget holders, with perhaps different views) in and beyond IT
 - Consider potential constraints by formal agreements, MoUs, ...
- How to ensure that views remain up-to-date with reality?

Service View – Dependencies

- Suggested categorisation:
 - Essential for keeping service running
 - Essential for re-starting service in present form
 - Essential for applying changes to service
 - Important for successfully running service
 - Desirable for successfully running service
- Circular dependencies: Not excluded by principle, but need to remove avoidable ones, and clearly identify and mark the unavoidable ones
- How to ensure that views remain up-to-date with reality?



Service View – Service Granularity

- Say Z depends on Y
- What granularity do we need to consider Y at?
- Can probably be sorted only by input from experts of both Z and Y
- We have the service catalogue used e.g. by the service desk – but granularity clearly too coarse

- Proper understanding important for reaching the right conclusions
- How to ensure that views remain up-to-date with reality?

Threats

- Hardware failures of individual nodes
- Power disruptions at various levels
 - Single-phase or all phases
 - PSU, PDU, room, building, ...
 - What is the right classification?
- Data loss due to intrusions
 - Cybercrime, e.g. ransomware
 - Sabotage
- Network disruptions at various levels
 - Router, switches, room, building, ...
 - What is the right classification?
- Physical impact: cooling, fire, inundation, ...
 - Part / zone of a room, room, building, ...
 - What is the right classification?

How to ensure that views remain up-to-date with reality?



Approach

- All members presented services from their respective groups
- Drilled down in more detail into three services: Indico, Teigi (part of configuration management), DNS
 - Lots of useful insights into which information to collect, and what are potential areas for recommendations

Status

- Template to describe services ready
- About to contact all services
- WG members will provide help wherever needed

- Regular contact with IT Management and with IT service managers

- Agreement that a one-off is of very limited usefulness
 - Group is prepared to launch an iterative process, say once or twice per year



... And Then What With the Data?

- Round of services will deliver large volume of data
- Some tool support is indispensable
 - Storing information
 - Making information available
 - E.g. on dependencies: Show me all services that critically depend on service X; find all circular dependencies
 - Cross-checking information with other sources, finding inconsistencies
 - Providing visualisation wherever adequate

Tool support

- Looked at options
 - Both in terms of free/open-source and commercial solutions
 - Mostly targeted at automatic discovery of dependencies of 'black-box' services
- Currently considering a simple implementation on top of Elasticsearch
 - Data ingress via JSON files – many people know them



Why Bringing This to HEPiX? (1)

- List of services and their dependencies and, to some extent, business cases are CERN-specific
 - Subsets overlap with other sites anyway
- Threats, dependency classifications etc. are definitely not
- Neither is the need for tool support

Why Bringing This to HEPiX? (2)

- Some services to be considered are WLCG- (or otherwise community-) wide
- Potential business continuity implementation may include other sites

Why Bringing This to HEPiX? (3)

- Is your site doing something in similar directions?
 - Is there a potential for collaboration, at least at the level of sharing experience?
- Do you have any experience with or ideas about tools that could help?

- We are all very eager to hear from other sites
 - Ready to consider collaborations

Comments? Questions? Interest?

Contact us now, by mail (see below), or via shared Google document

