Business Continuity/Disaster Recovery/Enterprise Architecture

IT Department Status

Tim Bell
IT BC/DR Lead
5 May 2023

IT strategy 2022-2025

Includes as a provider,

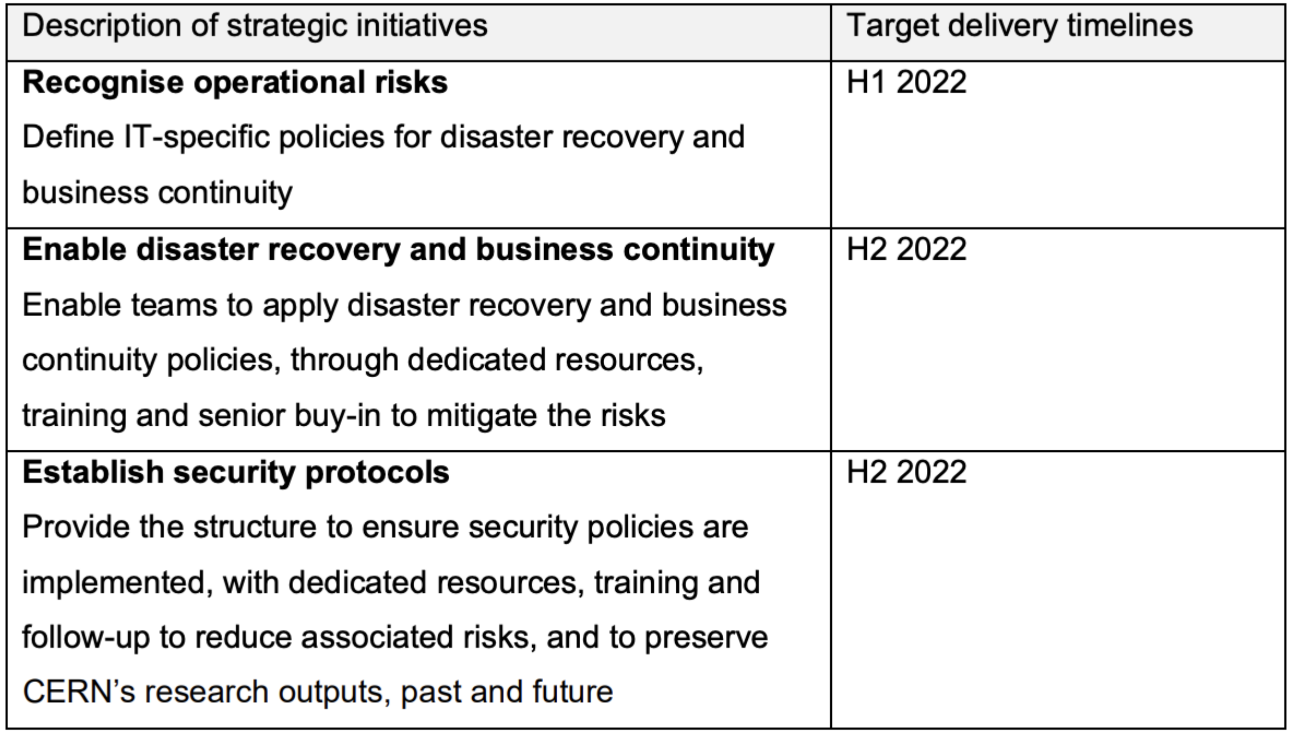

Description of strategic initiatives	Target delivery timelines
Recognise operational risks Define IT-specific policies for disaster recovery and business continuity	H1 2022
Enable disaster recovery and business continuity Enable teams to apply disaster recovery and business continuity policies, through dedicated resources, training and senior buy-in to mitigate the risks	H2 2022
Establish security protocols Provide the structure to ensure security policies are implemented, with dedicated resources, training and follow-up to reduce associated risks, and to preserve CERN's research outputs, past and future	H2 2022

8

Full IT Strategy document

AS-IS External Assessment

Recognise operational risks

Today: IT disaster recovery and business continuity procedures are not adequate. Although failures are limited, the risk is significant to ongoing operations

'We don't have a proper disaster recovery and business continuity plan'

Governance bodies in the project/service lifecycle

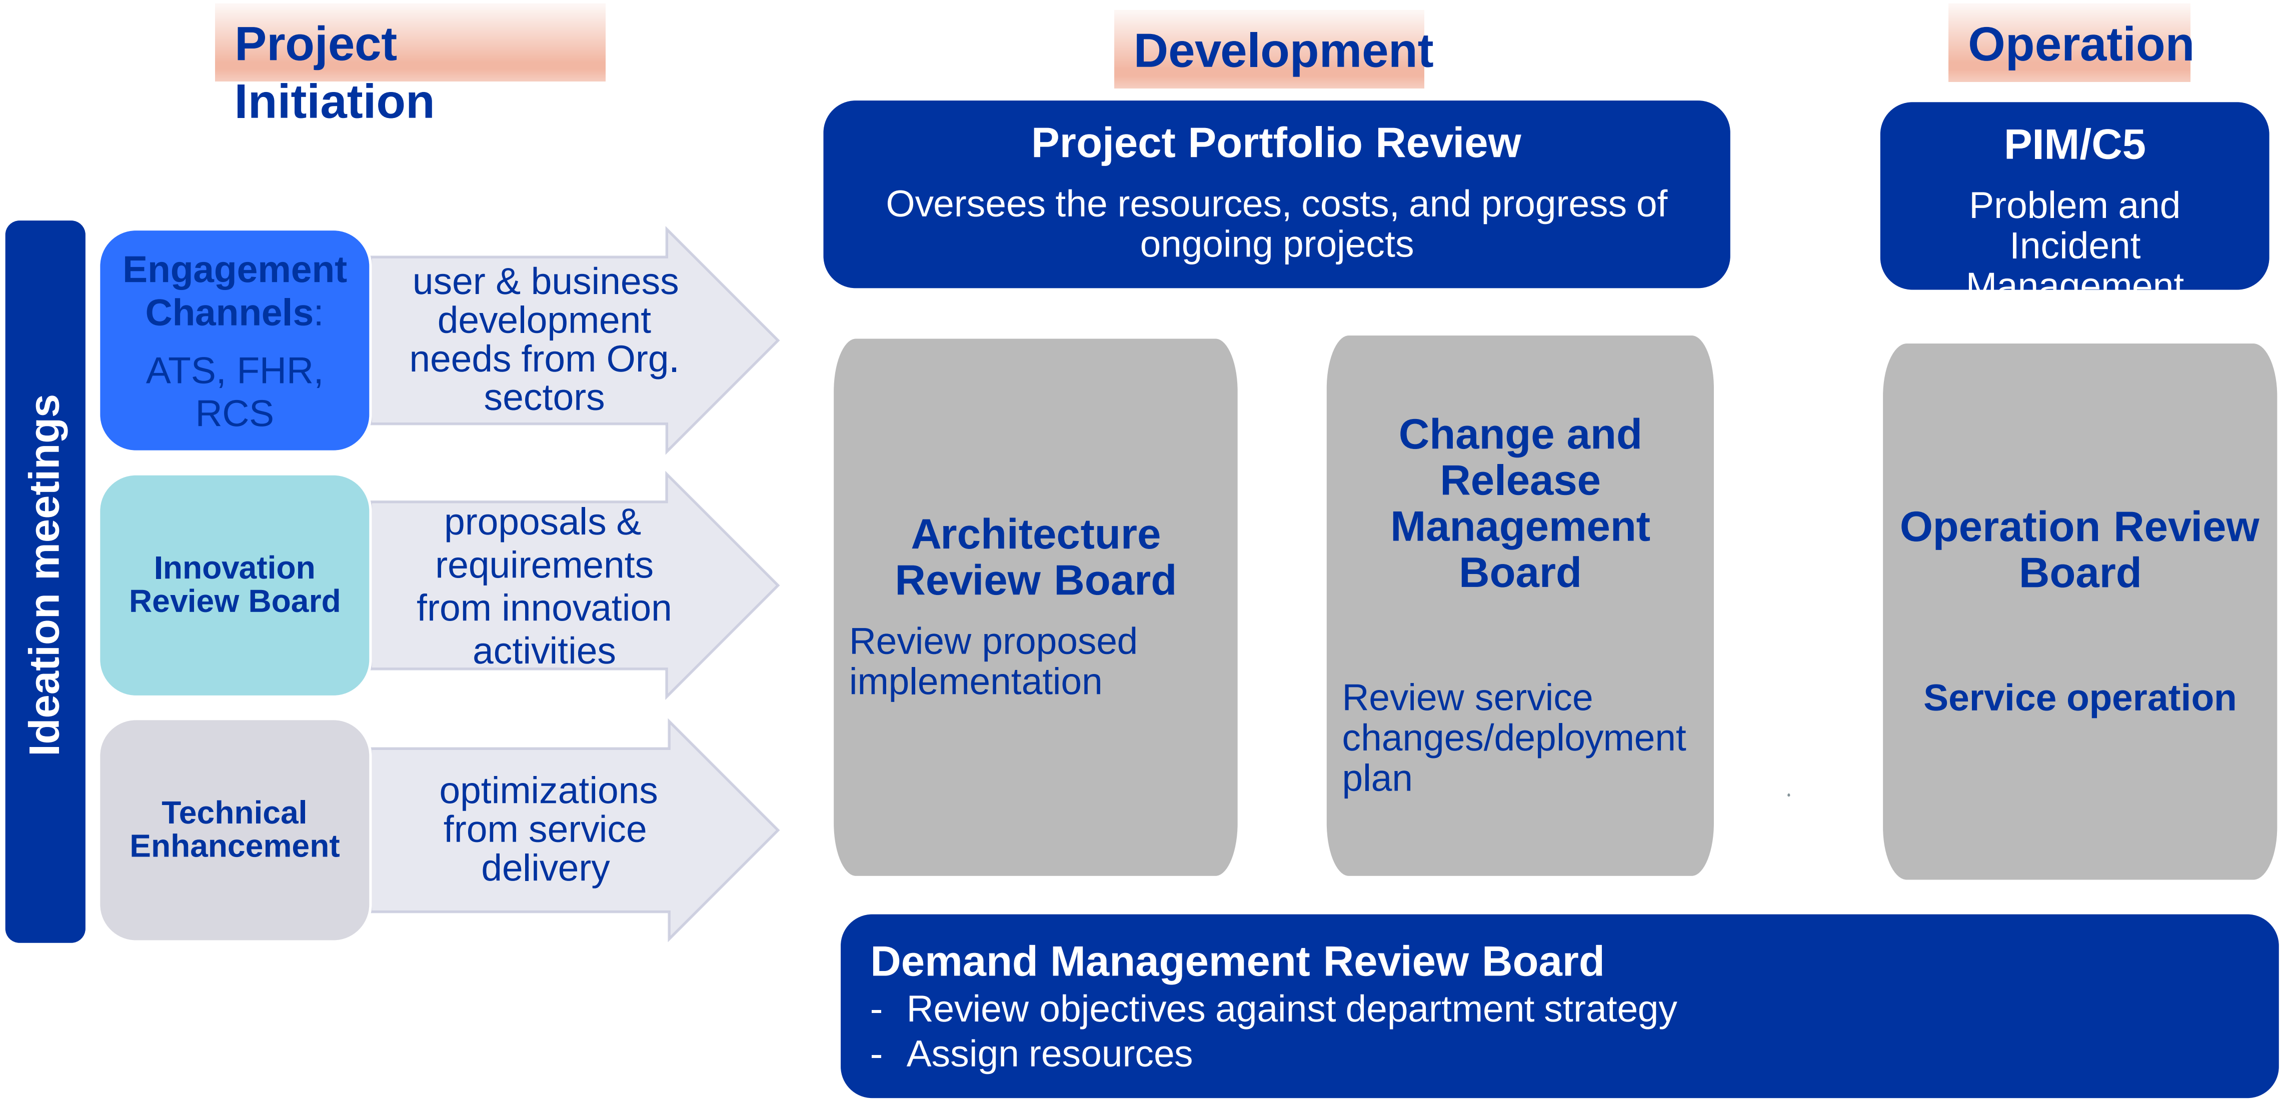

Project Initiation

Development

Operation

Ideation meetings

Engagement Channels: ATS, FHR, RCS

user & business development needs from Org. sectors

Innovation Review Board

proposals & requirements from innovation activities

Technical Enhancement

optimizations from service delivery

Project Portfolio Review

Oversees the resources, costs, and progress of ongoing projects

PIM/C5

Problem and Incident Management

Architecture Review Board

Review proposed implementation

Change and Release Management Board

Review service changes/deployment plan

Operation Review Board

Service operation

Demand Management Review Board

- Review objectives against department strategy
- Assign resources

BC/DR problem statement

- **Risk of significant incident remains**
  - IT has tried several initiatives (2012, 2016, 2019)
- **IT BC/DR project established in 2022 until end 2024 but mismatch between our targets and our current capacity**
  - Estimated 14 person years effort over 2 years, current outlook is around 4 person years spread over 8 people given departures
  - Is budget available? (estimated 1.2M CHF needed for PDC capacity and public cloud)
- **Underlying complex multi-department computing landscape (even more with the experiment workflows)**
  - Each consuming services from others, including external SaaS
  - Service attributes such as responsibilities, recovery times and availabilities are not always clearly understood
  - Personas, processes, systems are often not documented and maintained

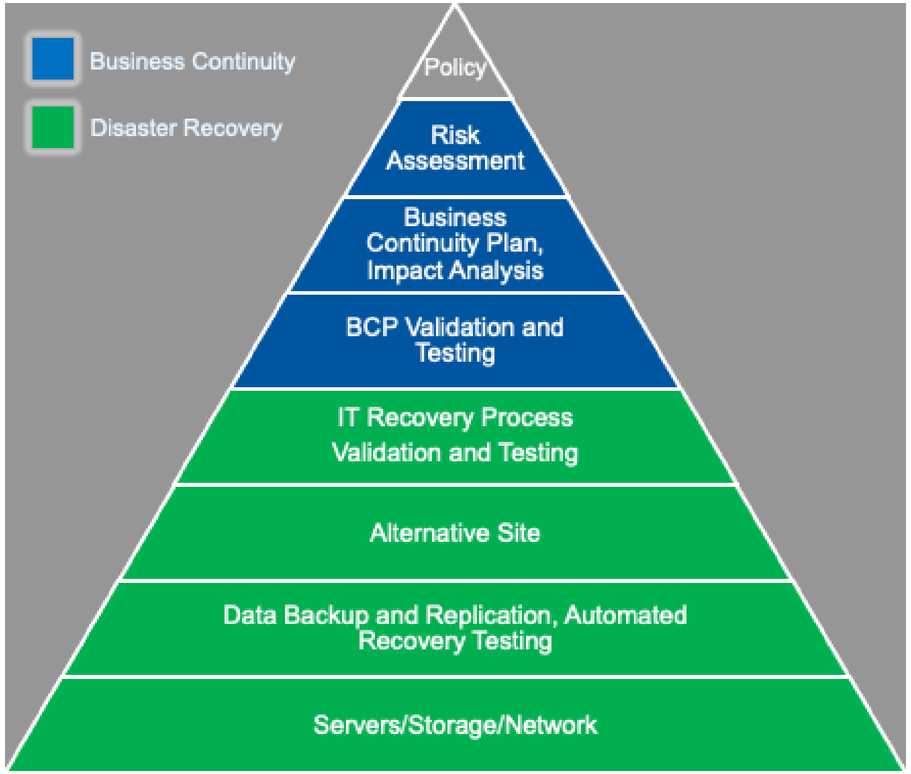
🟦 Business Continuity
🟩 Disaster Recovery
Policy
Risk Assessment
Business Continuity Plan, Impact Analysis
BCP Validation and Testing
IT Recovery Process Validation and Testing
Alternative Site
Data Backup and Replication, Automated Recovery Testing
Servers/Storage/Network

Goal #1: Business Continuity focus

- **Not an IT only problem - need an organisation network to evaluate key processes**
- **Prepare business impact analysis**
  - What happens if a process cannot be executed ?
  - How can the incident be mitigated ?
  - What are the desired recovery objectives (RPO/RTO) ?
  - What is the maximum tolerable downtime (MTD) ?
  - Needs an end-to-end view of how CERN’s critical process are delivered (see #3)

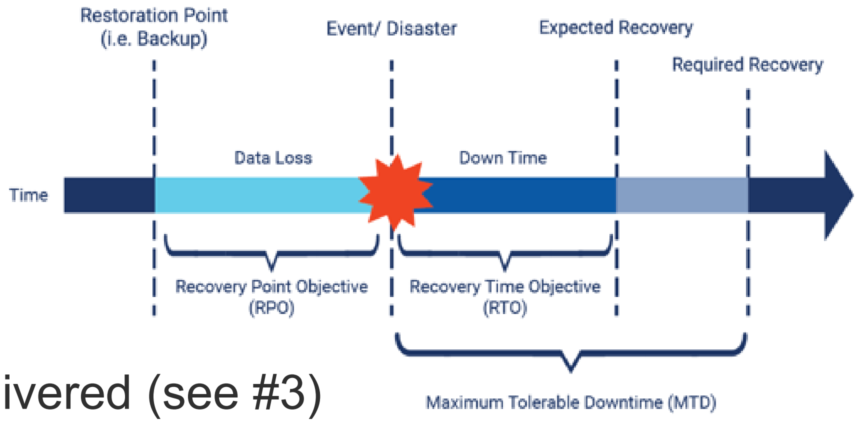

- **Would MTD cut-off reduce effort ?**
  - ~ Only (enviromental|safety|financial|recovery) – mitigations needed for others as significant downtime

Goal #2: Disaster Recovery focus

- **Highest priority is basic building blocks needed by low RTO**
  - Cloud and Storage for services
  - Kubernetes / Openshift
  - Authentication / Authorisation
  - Databases
- **Perform as-is assessment to understand where we are currently (Q2 2023)**
- **Prepare architecture proposals based on PDC availability and external cloud frameworks with the ARB (Q3 2023)**
  - Balance of up-front purchase cost for active-passive compared to cost only if we need it at the expense of higher recovery times
- **Update DR project plan based on available resources and required tasks (Q3 2023)**

Goal #3: Enterprise Architecture (strictly not BC/DR)

- **Understand, document and maintain the multi-department computing landscape**
  - Persona, Processes, Applications, IT Services, IT Functions and the inter-dependencies
  - Service lifecycle roadmap (e.g. pilot, production, maintenance, end of life of service and ensure end user actions aligned)
  - Common vocabulary for computing architectures

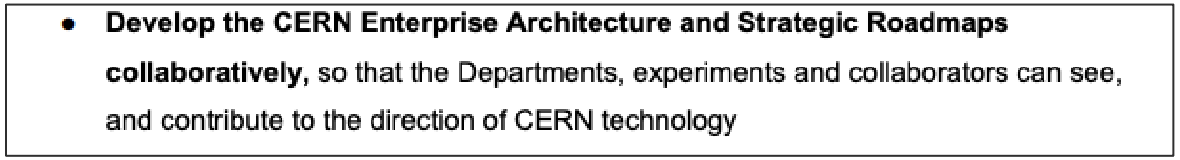

- **Develop the CERN Enterprise Architecture and Strategic Roadmaps collaboratively,** so that the Departments, experiments and collaborators can see, and contribute to the direction of CERN technology

IT Department Strategy 2022-2025

- **How**
  - FHR started 18 months ago and IT input is needed
  - ATS-IT Steering Committee endorsed the EA approach in December 2022
  - Tweak ARB processes and checklists – TOGAF was already selected from the start "What is Architecture?" (LIVE IT)
  - IT plan to be prepared in 2H 2023 and embedded into sector-IT planning, Service Catalog/Levels and BC/DR in 2H 2023 (as per goal defined in the IT Strategy 2022-2025)

- **Who – mirror FHR project structure with experts and ambassadors**
  - Train ARB members as Enterprise Architects,
  - Extended ARB contribute to the EA repository which is aiming to be organisation-wide
  - Service managers will be asked to provide input on current architecture along with updates for ARB and CMRB gates

Mini Business Impact Analysis

- Goal is to map from CERN key processes to recovery objectives of IT service
  - Approach
  - Raw data

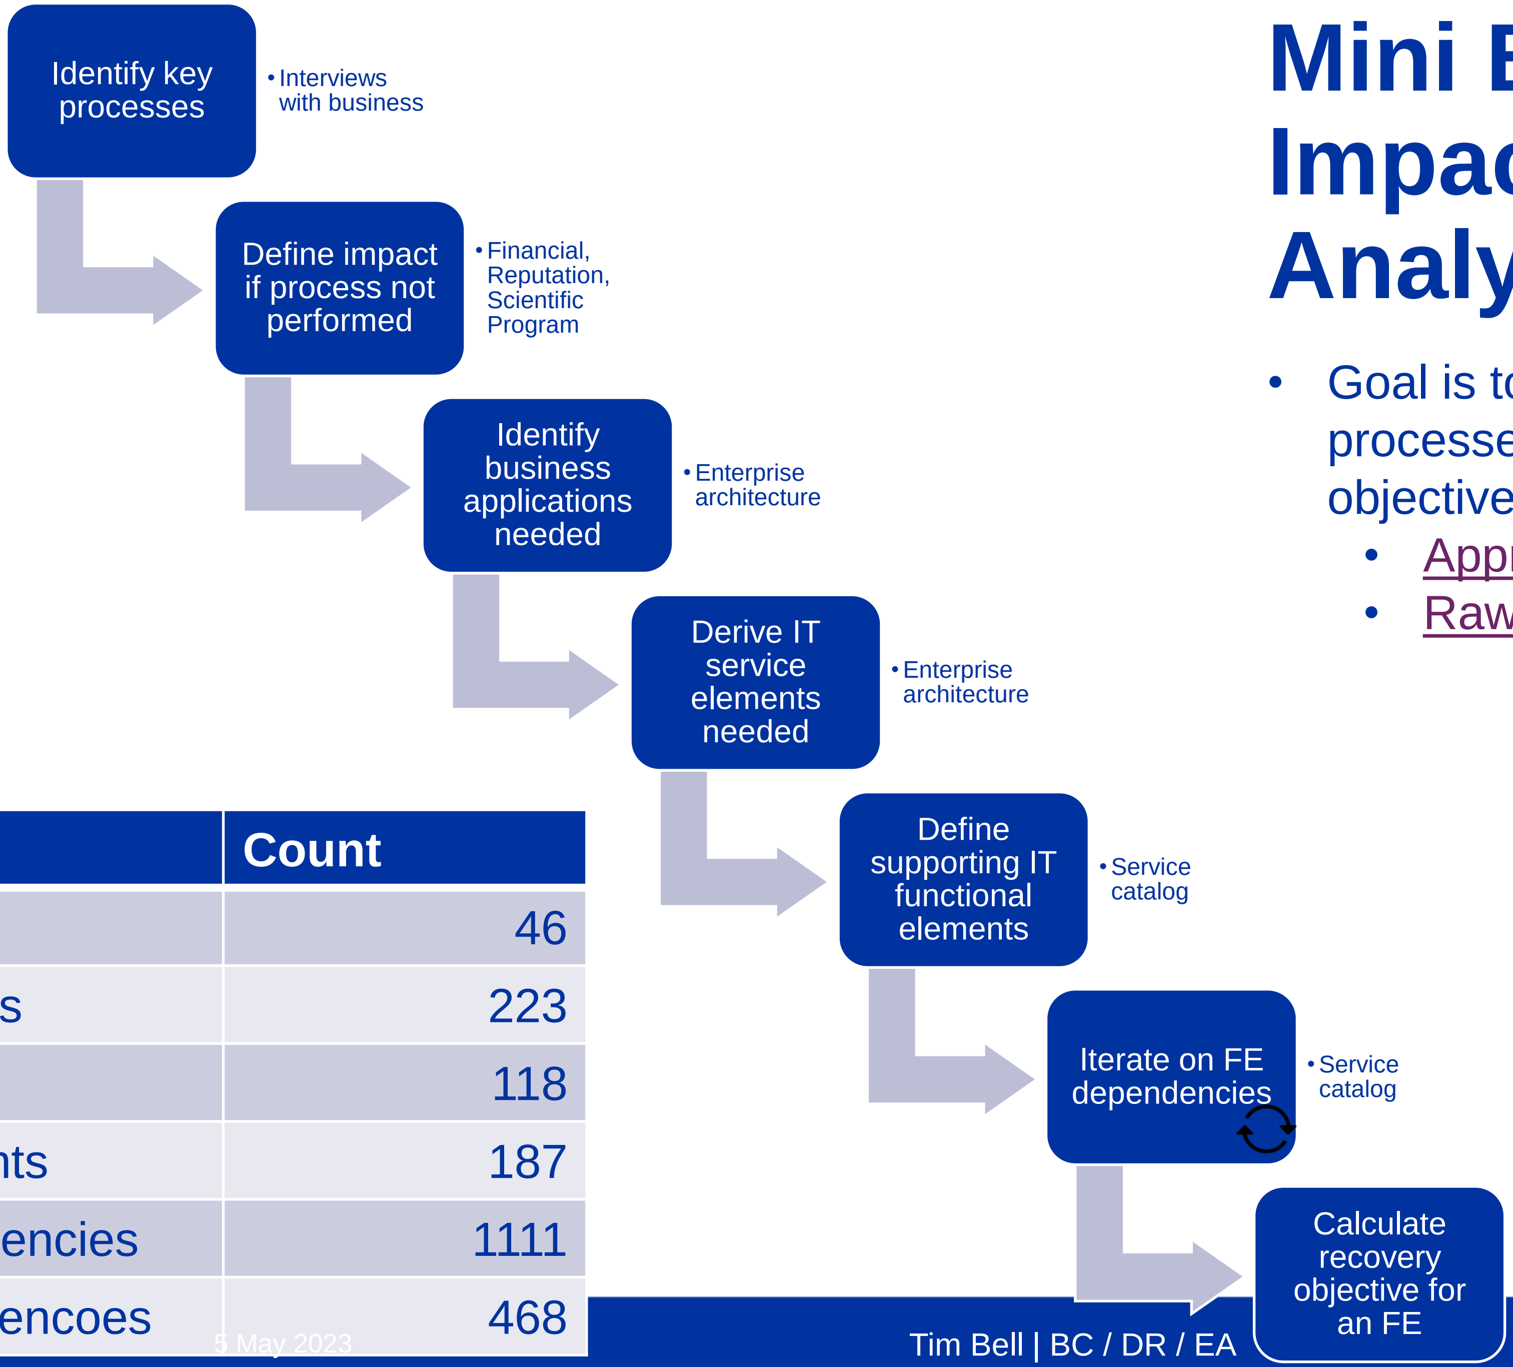

Identify key processes
- Interviews with business

Define impact if process not performed
- Financial, Reputation, Scientific Program

Identify business applications needed
- Enterprise architecture

Derive IT service elements needed
- Enterprise architecture

Define supporting IT functional elements
- Service catalog

Iterate on FE dependencies
- Service catalog

Calculate recovery objective for an FE

Metric	Count
Business Processes	46
Business Applications	223
IT Service Elements	118
IT Functional Elements	187
IT SE on FE Dependencies	1111
IT FE on FE Dependencoes	468

Business Continuity Next Steps

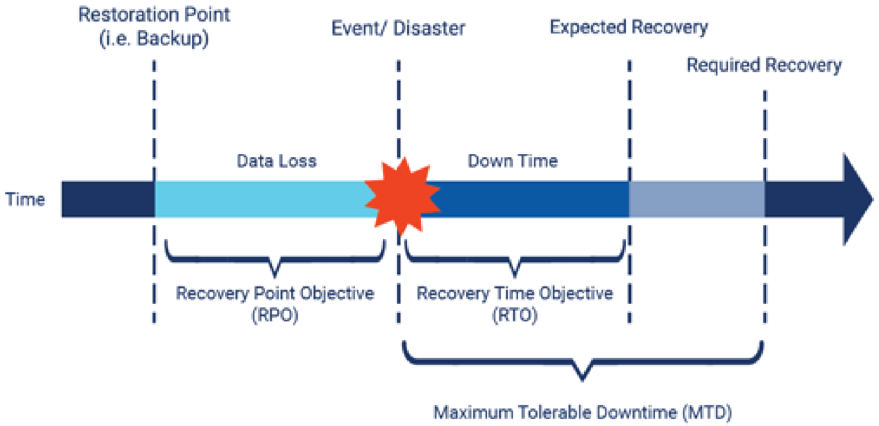

- **FAP-BC have run a procurement for “Impact Analysis on Business of Business Computing Downtime” (DO-33582)**
  - Consultancy to determine the recovery objectives and maximum tolerable downtime for FHR services
  - Consultants will also report on their assessment of CERN’s maturity levels
  - Aim to start in Q2 2023
- **Similar analysis was done in 2016**
- **We can expect a follow up discussion on service levels**
- **This could potentially act as a template for other sectors**

Disaster Recovery

- **Focus of IT efforts in 2023 is establishing a cross-group project team (initially with IT-FA, IT-CD, IT-SD, IT-DA and IT-PW) (DROIT details)**
  - Evaluate plans and tasks in their group's areas of expertise
  - Act as point of contact for their group members to the BC/DR project lead for as-is analysis, service enhancements and test plans
  - Communicate concepts, what's needed and status to their groups
- **Currently performing as-is assessment to identify, for each service (Draft)**
  - Current status of service resilience
  - Functionality to support recovery of other services
  - Upcoming improvements

Disaster Recovery Maturity KPI

- As part of the new operating model, a set of KPIs are being defined to show the current status and track the benefits of the changes going forward.
- Proposed criteria are as follows
  - Data source: An annual self assessment is performed by the disaster recovery operations in IT (DROIT) team. The 2023 report is in preparation here. This operation will be repeated each year in Q1.
  - Collection method : BC/DR lead will perform the CobiT analysis following the annual as-is self assessment. This will produce a value in the range 0-5 based on the average CobiT level for the FEs with a recovery time objective of <= 1 day
  - Detailed criteria in slides #18 (Backups)
  - Reporting: Using spreadsheet at https://cernbox.cern.ch/s/PBAlUp6avpICUcP
  - Current baseline (2023): 0.95 (“Incomplete process”)
- **Note:** it is not a criticism if a service has a lower rating, it reflects previous priorities. The key aspect is the delta as we go forward.

Self Assessment Criteria (CobiT)

Level	CobiT Maturity	CobiT Description	DR Description
0	Incomplete process	The process is not placed or it cannot reach its objective. At this level the process has no objective to achieve. For this reason this level has no attribute.	Limited or no DR capability
1	Performed process	Performed process. The process is in place and achieves its own purpose. This level has only “Process Performance” as process attribute.	Occasional testing and informal estimation of recovery objectives.
2	Repeatable but intuitive	The process is implemented following a series of activities such as planning, monitoring and adjusting activities. The outcomes are established, controlled and maintained. This level has “Performance Management” and “Work Product Management” as process attributes	Resiliency architecture available describing the implementation of the disaster recovery process, RTO/RPOs shared with business, dependencies agreed with associated objectives.
3	Established process	The previous level is now implemented following a defined process that allows the achievement of the process outcomes. This level has “Process Definition” and “Process Deployment” as process attributes.	Success criteria defined. Manually tested confirming correct recovery within RTO/RPO. Test results documented and published.
4	Predictable process	This level implements processes within a defined boundary that allows the achievement of the processes outcomes. This level has “Process Management” and “Process Control” as process attributes.	Process description for how to execute the disaster recovery process documented, qualified staff defined, and DR Test executed, and results validated.
5	Optimising process	This level implements processes in the way that makes it possible to achieve relevant, current and projected business goals. This level has “Process Innovation” and “Process Optimisation” as process attributes.	DR Test part of the standard operational procedures of the service with regular testing of backups and DR steps. RTO/RPO are agreed with the business and documented in the IT service level description.

Using the IT operating model to deliver BC/DR

- **With Architecture, Demand management and Change/Release boards**
  - Define criteria for assessment of operational resilience
  - Best practices defined for backup/restore, resource placement and DR testing
  - Reference industry standard DR architecture patterns for service managers
  - Implement the critical improvements for the highest risks within the IT budget

AWS

Cold	Pilot Light	Warm Standby	Multi Site Active-Active
RTO/RPO : hours / days	RTO/RPO : 10s of minutes	RTO/RPO : minutes	RTO/RPO : seconds
• Less critical systems • Classic restore from backup • Prevision and restore after the event • Cost $	• Data live, minimal capacity • Scale out after the event • Cost $$	• Business critical • Initially running at degraded capacity but usable • Scale to full capacity after the event • Cost $$$	• Minimal downtime • Near zero data loss • Mission critical services • Can be complex and potential production impact • Cost $$$$

Disaster Recovery Next Steps

- **PDC – Building – Infrastructure-as-a-Service**
  - Procurement of the initial Infrastucture-as-a-Service resources for the PDC UPS racks (Q2 order, ~Q4 delivery)
    - Capacity for services with a recovery objective < 1 day based off mini impact analysis and cloud resource report
    - O(1.5M) CHF orders in preparation for fast recovery such as Active-Active or Active-Passive
  - Define architecture for independent cloud region in PDC
    - Test bed in B513 already available to help bootstrap (e.g. burn-in) and test
    - IT Architecture Review Board scheduled in May to review the IaaS proposal
    - Test plan development with DROIT (Q3)
  - Full disconnect test in planning for start / end of LS3
- **Public cloud contracts**
  - For services with a recovery objective > 1 day
  - Low running costs (testing) but could be scaled as needed
  - Some testing with Oracle cloud but will need Public Cloud Framework (in preparation) and budget

Enterprise Architecture Motivations

- **Complex relationships between different CERN departments computing activities are not documented (and often not understood)**

- **Computing evolution is hampered by difficulty of impact analysis and integrating multiple component plans**

- **Result is**
  - “Don’t change anything because you’ll break something”
  - “This service needs to be kept running just in case even though there is a high cost”
  - “Service levels are difficult to define as lower layer dependencies unknown”
  - “There is no common place to plan change and end of life so surprises occur”
  - “Defining downtime impact requires many meetings”

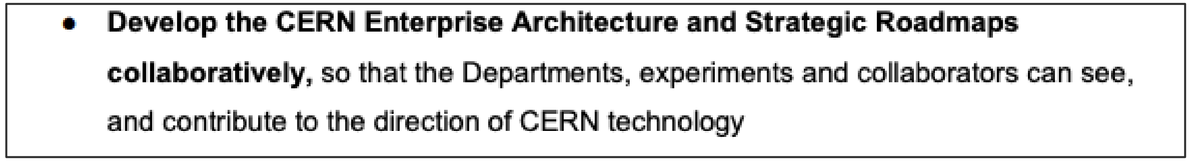

- **Develop the CERN Enterprise Architecture and Strategic Roadmaps collaboratively,** so that the Departments, experiments and collaborators can see, and contribute to the direction of CERN technology

IT Department Strategy 2022-2025

TOGAF : An opengroup Standard for EA

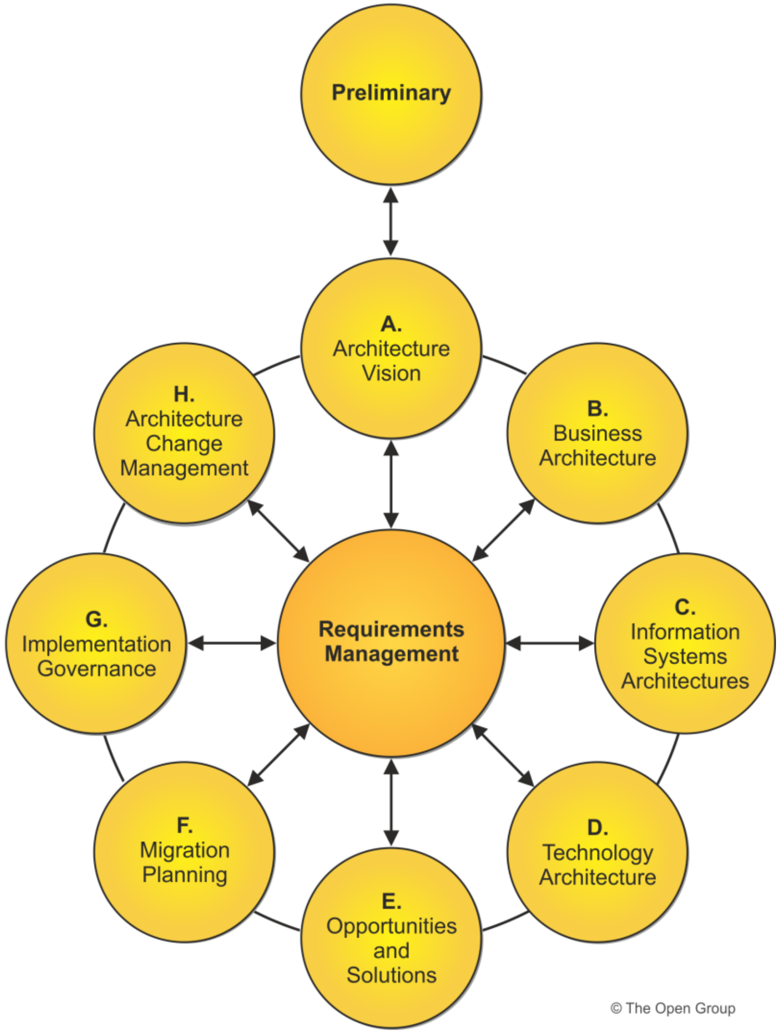

© The Open Group

‘The TOGAF Standard is used by small, medium, and large commercial businesses, as well as government departments, non-government public organizations, and defense agencies

With greatly expanded guidance and how-to material, it enables organizations to operate in an efficient and effective way across a broad range of use-cases, including agile enterprises and Digital Transformation

The TOGAF Standard is designed for the dichotomy of common universal concepts and variable detailed configuration

The structure focuses on what most architects want – more, better, and topical guidance on how to deliver the best Enterprise Architecture that supports their stakeholders and their organization”

Source: https://www.opengroup.org/togaf

TOGAF Technical Reference Model (TRM)

“The objective of the TOGAF TRM is to provide a widely accepted core taxonomy, and an appropriate visual representation of that taxonomy.”

A common ontology for how the organisations computing activities fit together

Source: https://pubs.opengroup.org/togaf-standard/reference-models/trm.html

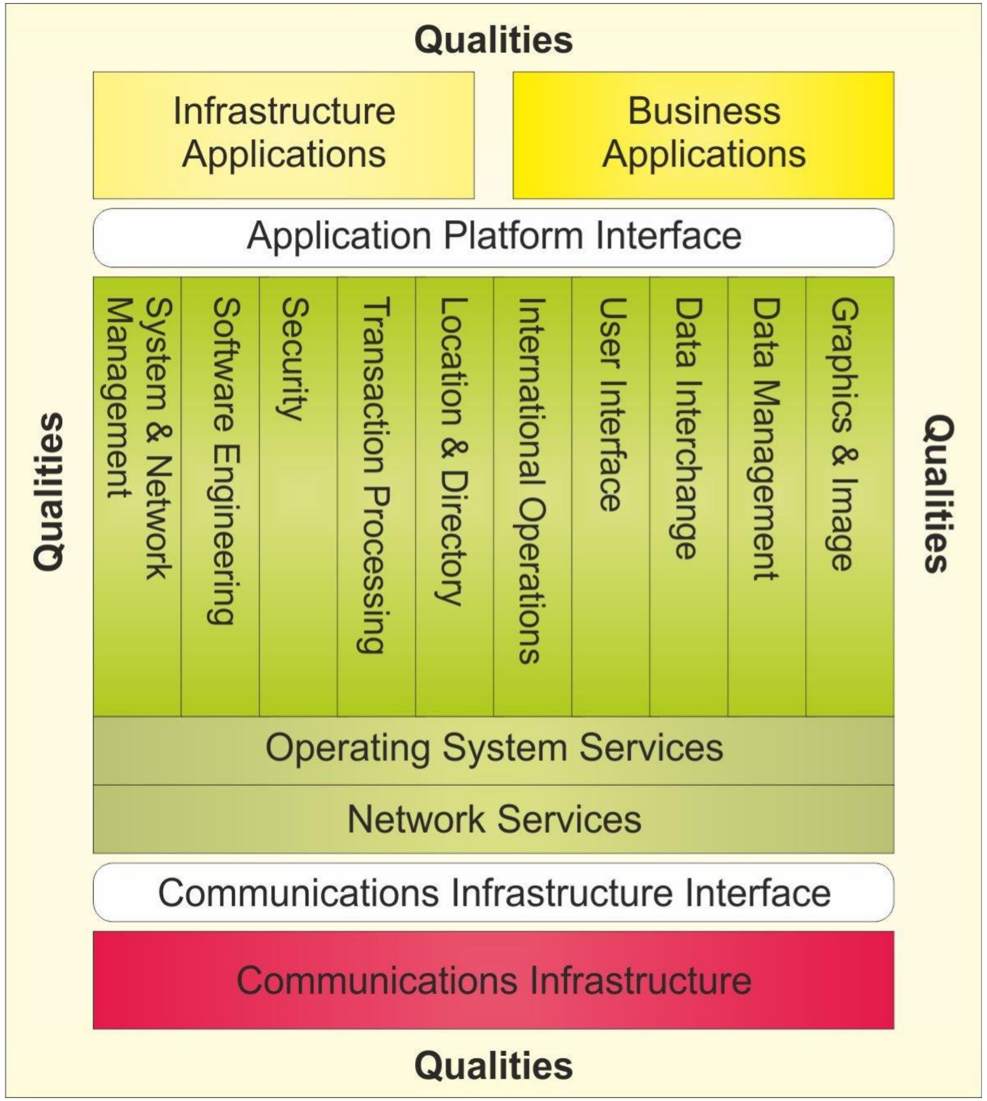

Qualities

Infrastructure Applications

Business Applications

Application Platform Interface

System & Network Management

Software Engineering

Security

Transaction Processing

Location & Directory

International Operations

User Interface

Data Interchange

Data Management

Graphics & Image

Operating System Services

Network Services

Communications Infrastructure Interface

Communications Infrastructure

Qualities

Qualities

Qualities

FHR Analysis : Capabilities

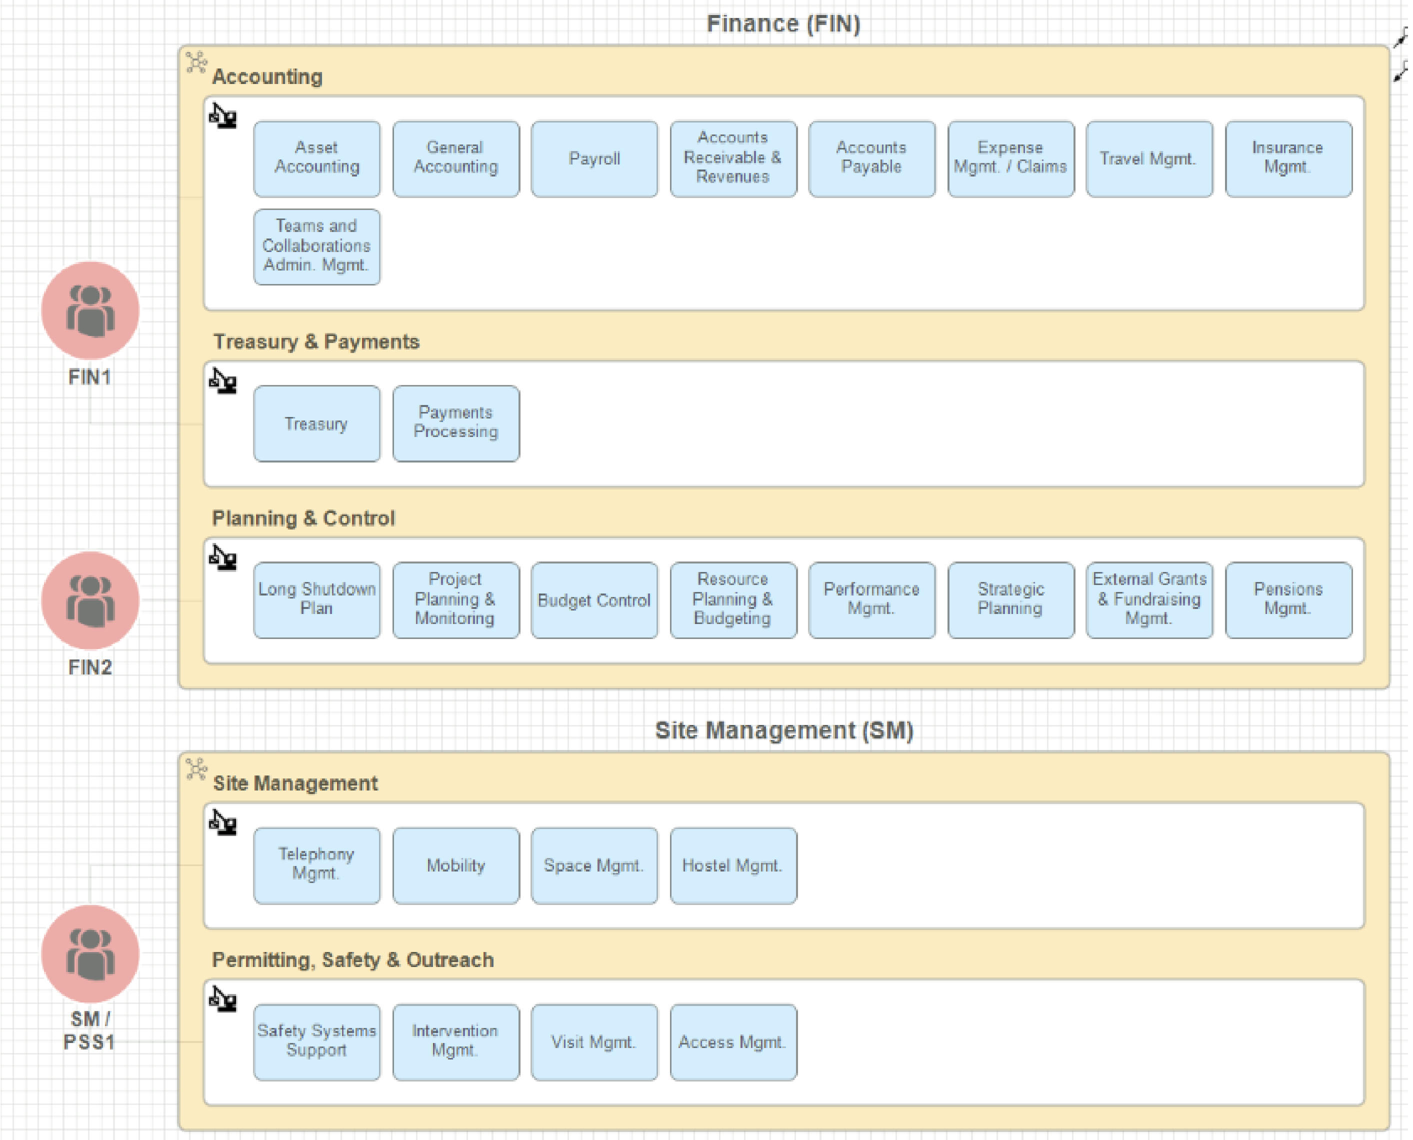

Finance (FIN)

Accounting

- Asset Accounting
- General Accounting
- Payroll
- Accounts Receivable & Revenues
- Accounts Payable
- Expense Mgmt. / Claims
- Travel Mgmt.
- Insurance Mgmt.
- Teams and Collaborations Admin. Mgmt.

Treasury & Payments

- Treasury
- Payments Processing

Planning & Control

- Long Shutdown Plan
- Project Planning & Monitoring
- Budget Control
- Resource Planning & Budgeting
- Performance Mgmt.
- Strategic Planning
- External Grants & Fundraising Mgmt.
- Pensions Mgmt.

FIN1

FIN2

Site Management (SM)

Site Management

- Telephony Mgmt.
- Mobility
- Space Mgmt.
- Hostel Mgmt.

Permitting, Safety & Outreach

- Safety Systems Support
- Intervention Mgmt.
- Visit Mgmt.
- Access Mgmt.

SM / PSS1

FHR Analysis : Sample Process

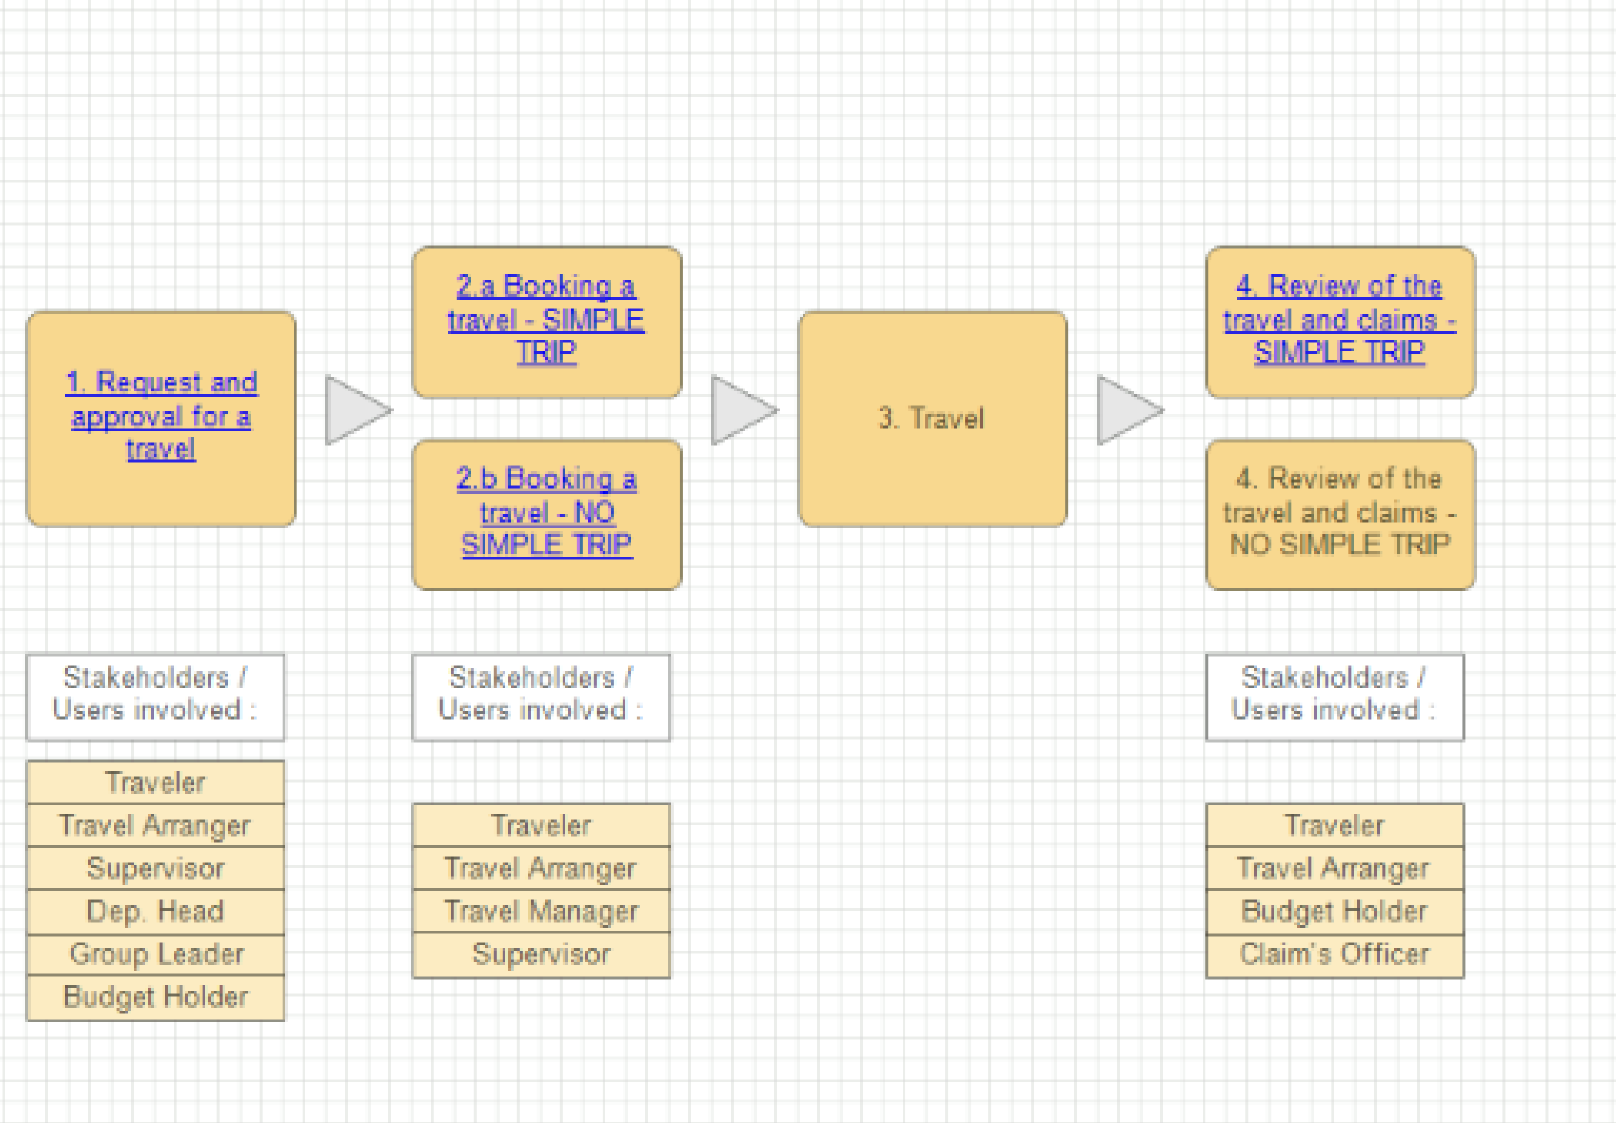

1. Request and approval for a travel

2.a Booking a travel - SIMPLE TRIP

2.b Booking a travel - NO SIMPLE TRIP

3. Travel

4. Review of the travel and claims - SIMPLE TRIP

4. Review of the travel and claims - NO SIMPLE TRIP

Stakeholders / Users involved :
Traveler
Travel Arranger
Supervisor
Dep. Head
Group Leader
Budget Holder

Stakeholders / Users involved :
Traveler
Travel Arranger
Travel Manager
Supervisor

Stakeholders / Users involved :
Traveler
Travel Arranger
Budget Holder
Claim's Officer

FHR Analysis : Process to Application

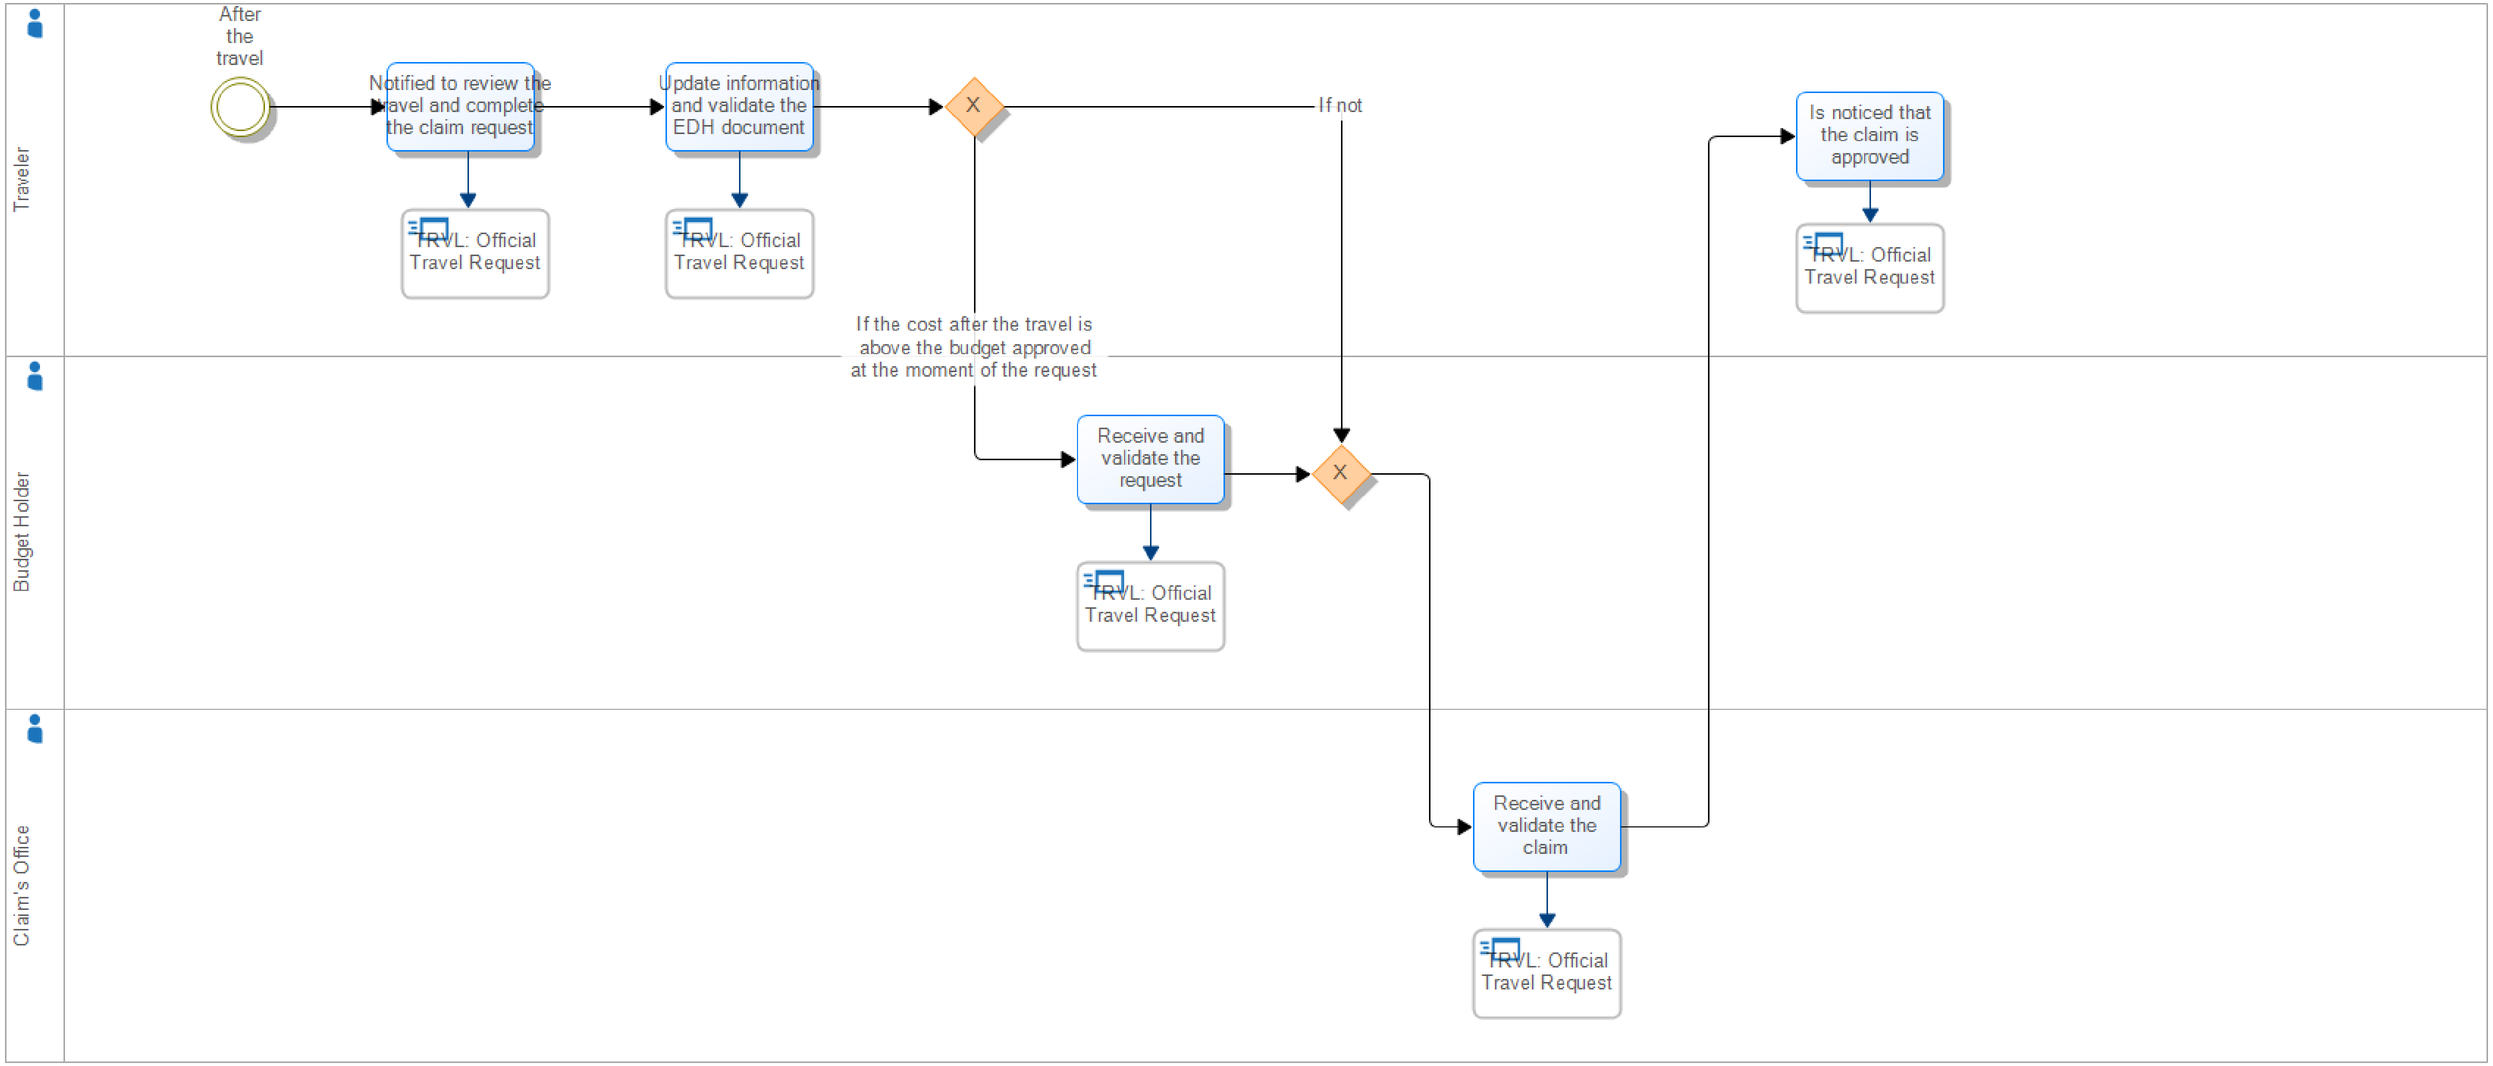
Traveler
Budget Holder
Claim's Office
After the travel
Notified to review the travel and complete the claim request
TRVL: Official Travel Request
Update information and validate the EDH document
TRVL: Official Travel Request
X
If not
If the cost after the travel is above the budget approved at the moment of the request
Receive and validate the request
TRVL: Official Travel Request
X
Receive and validate the claim
TRVL: Official Travel Request
Is noticed that the claim is approved
TRVL: Official Travel Request

Sample Infrastructure Application : acron

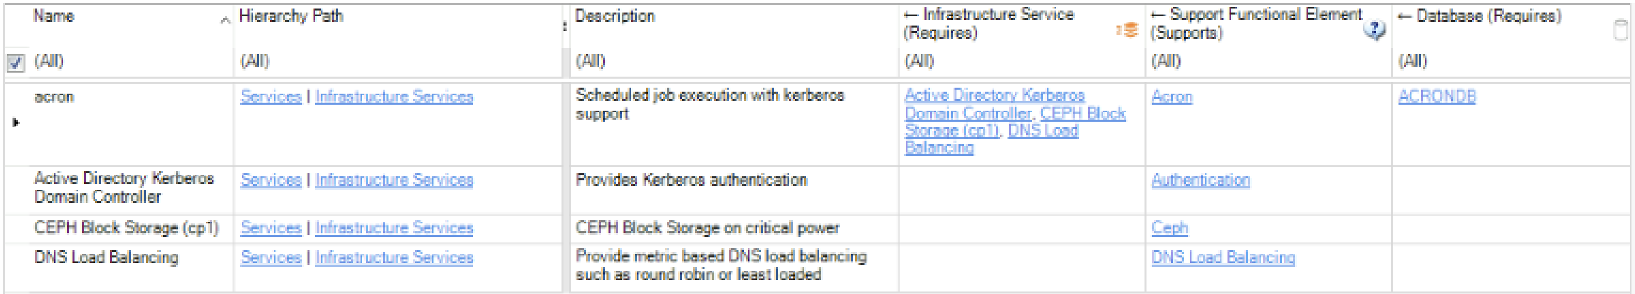

Name	Hierarchy Path	Description	← Infrastructure Service (Requires)	← Support Functional Element (Supports)	← Database (Requires)
☑ (All)	(All)	(All)	(All)	(All)	(All)
acron	Services | Infrastructure Services	Scheduled job execution with kerberos support	Active Directory Kerberos Domain Controller, CEPH Block Storage (cp1), DNS Load Balancing	Acron	ACRONDB
Active Directory Kerberos Domain Controller	Services | Infrastructure Services	Provides Kerberos authentication		Authentication	
CEPH Block Storage (cp1)	Services | Infrastructure Services	CEPH Block Storage on critical power		Ceph	
DNS Load Balancing	Services | Infrastructure Services	Provide metric based DNS load balancing such as round robin or least loaded		DNS Load Balancing	

Catalogs contain properties (e.g. Description) and constraints (with the connection type, e.g. Requires)

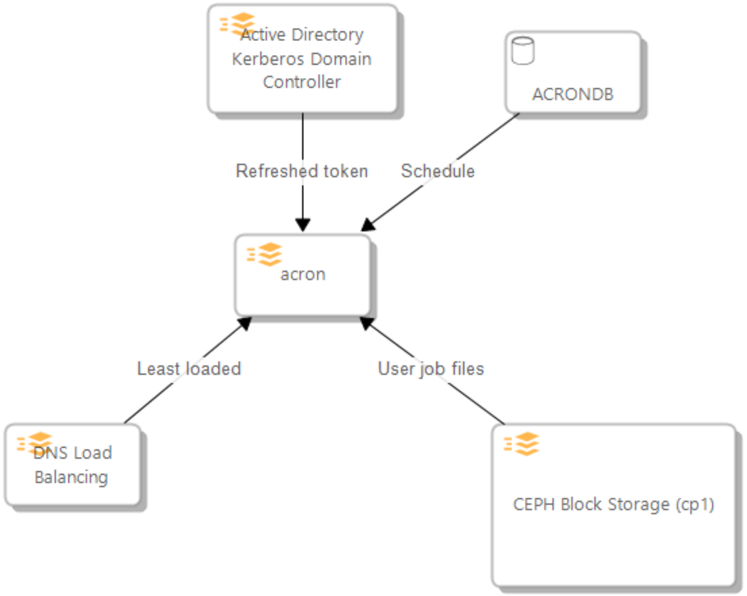

Diagrams

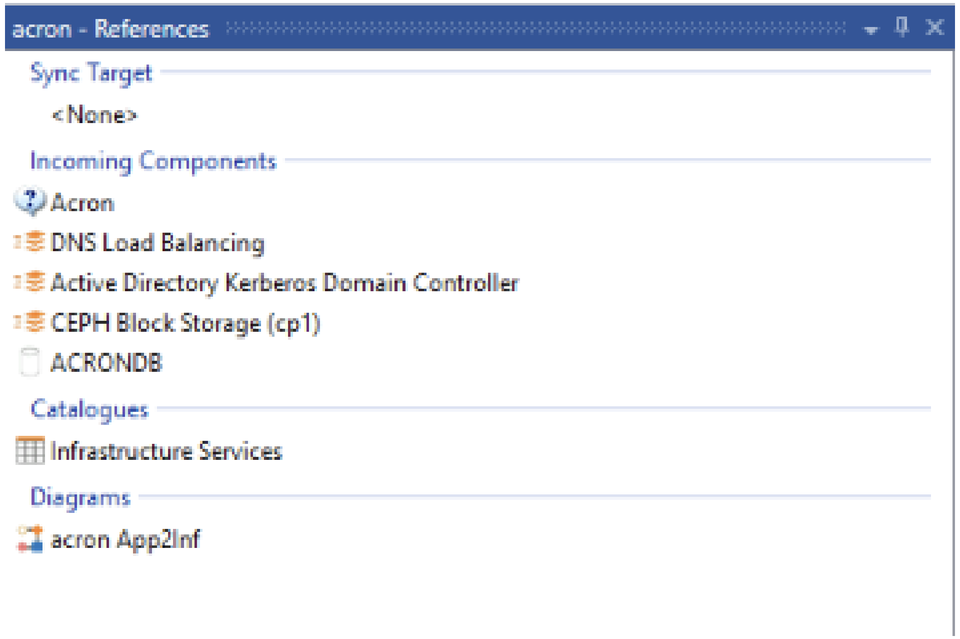

acron - References

Sync Target

<None>

Incoming Components

- Acron
- DNS Load Balancing
- Active Directory Kerberos Domain Controller
- CEPH Block Storage (cp1)
- ACRONDB

Catalogues

- Infrastructure Services

Diagrams

- acron App2Inf

References

Enterprise Architecture : Next Few Months (Q2/3)

- **Working with FHR and ATS on a roadmap**
  - Plan to submit Purpose-Scope-Objective (PSO) initiatives for both FHR-IT and ATS-IT along with the IT DMRB (Draft)
  - Establish a CERN wide forum to share experiences and manage common EA repository
- **An IT Architecture Review Board workplan (Q3)**
  - However, with significant implications on BC/DR e.g. business impact analysis
  - Need to get the organisation computing engineers on board
- **IT Modeling of TOGAF Technical Reference Model (TRM) in Abacus**
  - Map IT services and components to the TOGAF model
  - Define properties / dependencies
  - Investigate diagramming, import processes and built in reports (e.g. availability, performance, …)
- **The SNOW service catalog will initially be used as-is**
  - Many potential evolutions but could lead to pleasing nobody
  - Catalog changes have a significant cost

BC/DR/EA Further Concerns

- **Not all critical services will be resilient and tested until LS3**
  - Cause
    - IT Technical Delivery resource availability
    - Departure of BC/DR team member
    - Related activities needed to improve maturity such as Enterprise Architecture / Service Levels
  - Proposed actions
    - Target PDC first 6 months window for initial DR testing and full disconnect test requested at start/end LS3
    - DMRB allocation of a share of new fellows in Infrastructure-as-a-Service layers

- **Dispersed activities around BC/DR across the organisation**
  - Cause
    - Complex computing services organisation and responsibilities
    - Limited service level agreements between parties
    - No agreement on CERN wide governance structure yet
  - Proposed actions
    - Maintain informal contacts such as EA and power cut preparation team

Incidents : some High Energy Physics examples

Incident	When	Site	Details
Tapes ‘dampened’ due to plumbing mistake	2004	CERN	HEPiX
Tree cuts all power	2005	SLAC	HEPiX
Power outage	2006	CERN	Report
EDH down for 3 days due to RAID failure	2007	CERN	HEPiX
UPS Fire	2009	ASGC/Sineca	Details
20,000 tape files unintentionally deleted	2010	CERN	Details
Site wide power outage	2010	SLAC	HEPiX
Power outage	2014	CERN	Details
Flood of computer centre	2018	INFN	CHEP
Power outage during power test	2021	CERN	ASDF

Additional Material

- **IT BC/DR Project in Indico**
  - https://indico.cern.ch/event/1276940/
- **IT Disaster Recovery Web Site**
  - https://disaster-recovery.web.cern.ch/
- **Draft IT Project Description for BC/DR**
  - https://cernbox.cern.ch/s/ElLIpYoVJpKBjZI

If there are any permission errors, please get in touch

Backups

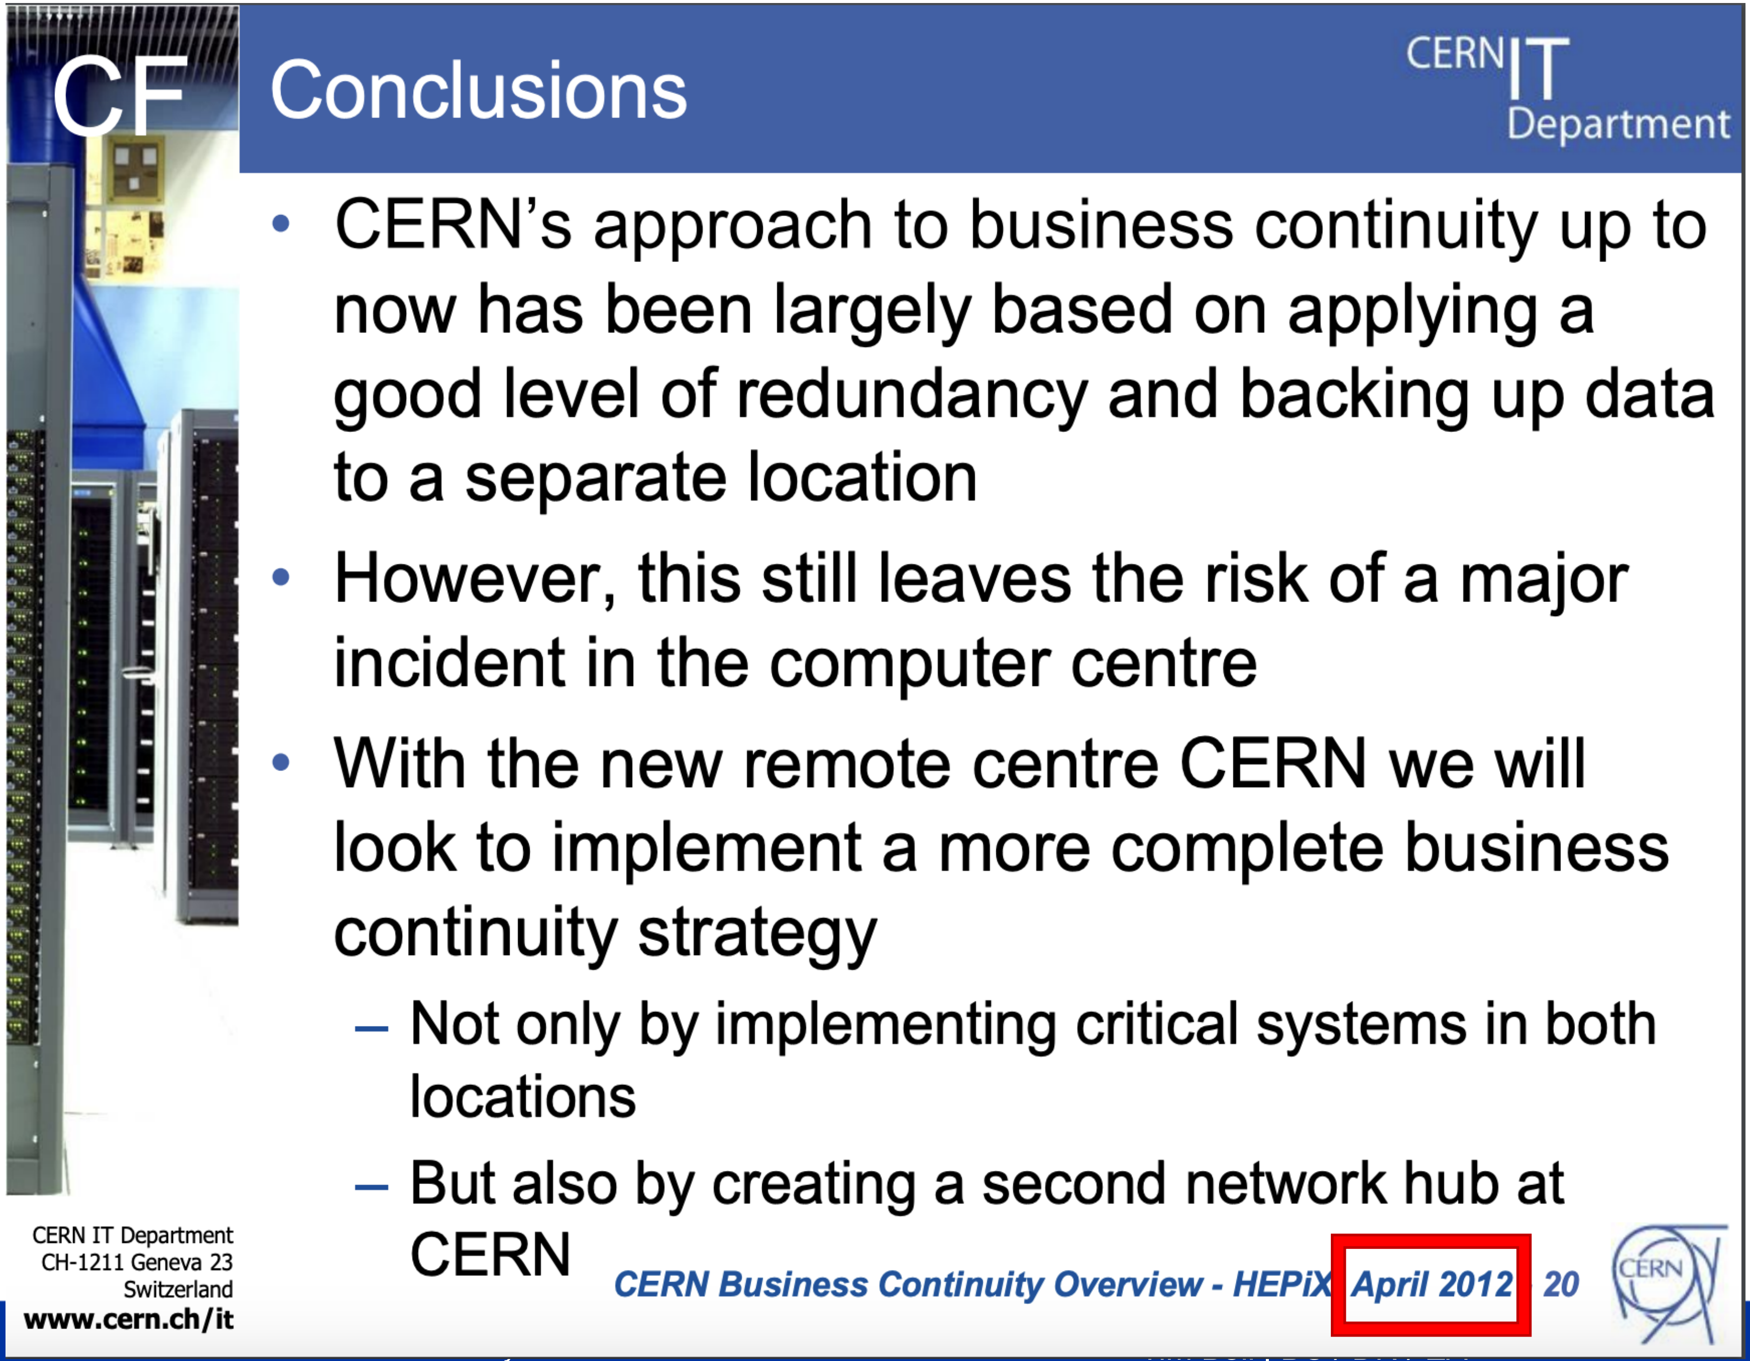

CF

Conclusions

CERN IT Department

- 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

CERN Business Continuity Overview - HEPiX April 2012 20

Wayne Salter
HEPiX 2012

2022 – transformation
2023 – consolidation and evolution

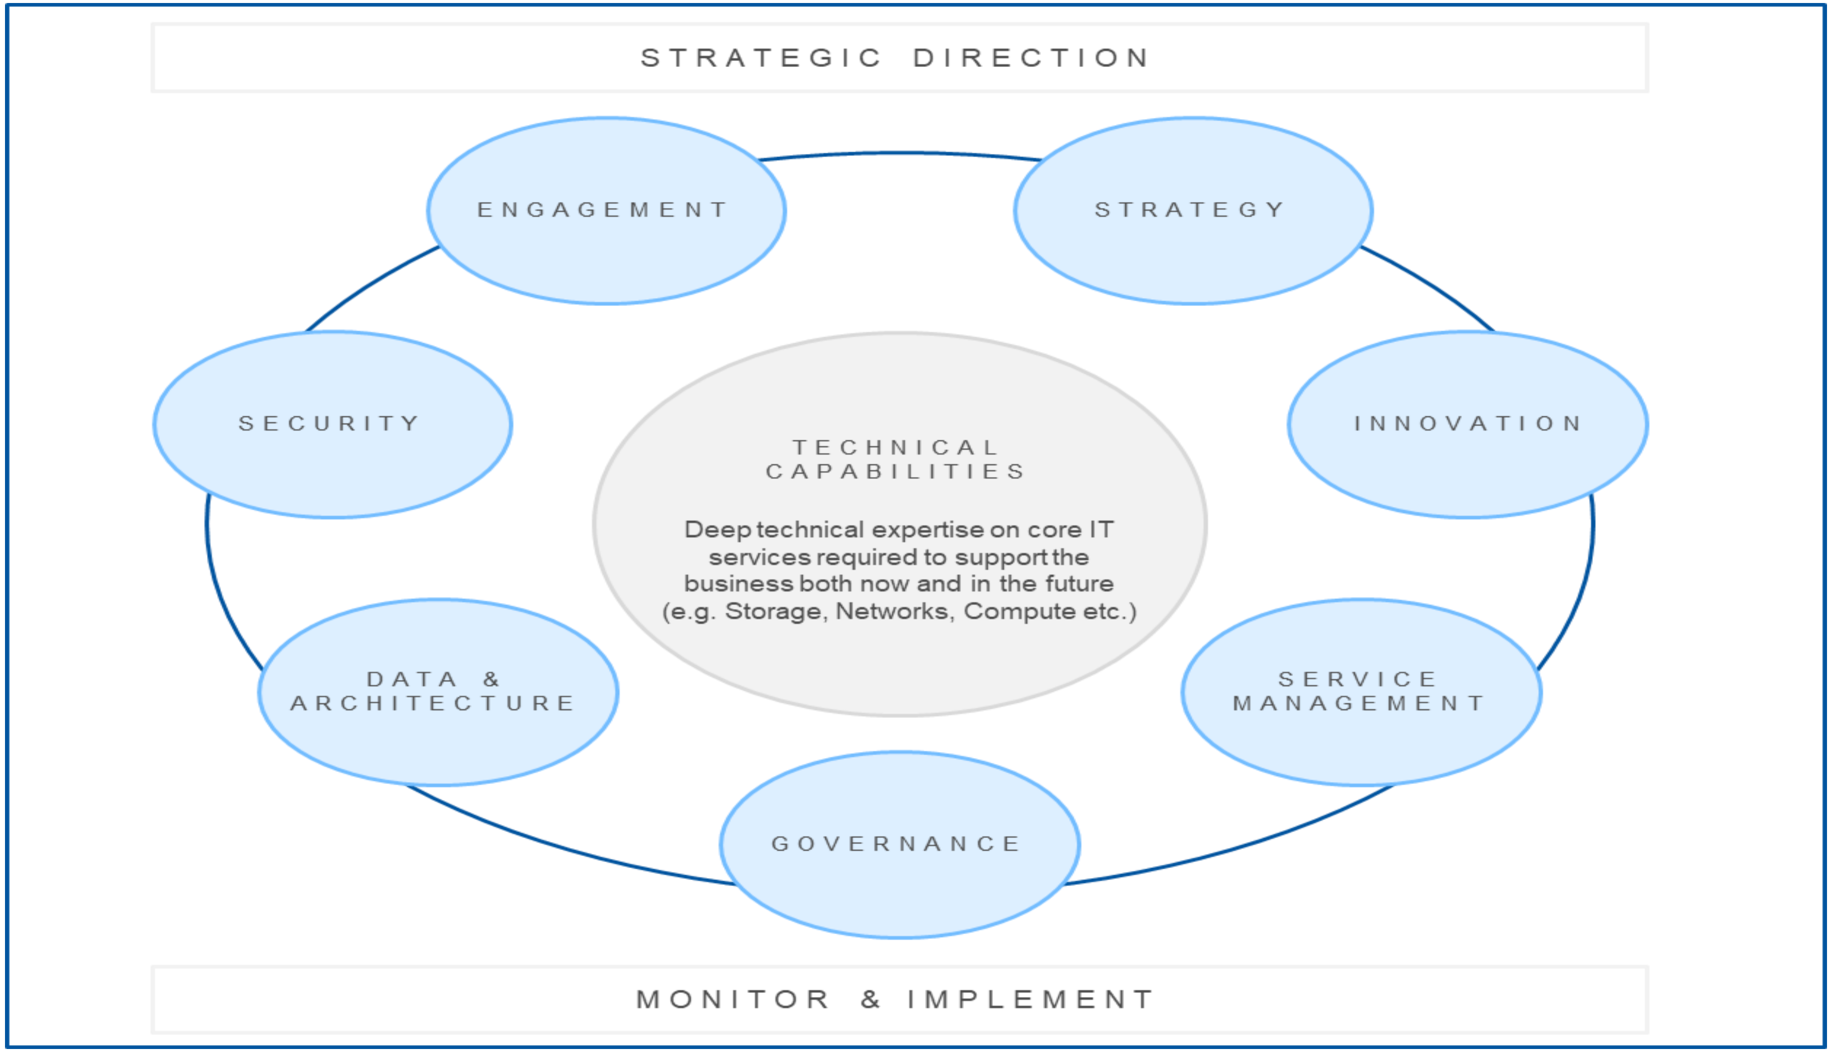

STRATEGIC DIRECTION

ENGAGEMENT

STRATEGY

SECURITY

INNOVATION

TECHNICAL CAPABILITIES

Deep technical expertise on core IT services required to support the business both now and in the future (e.g. Storage, Networks, Compute etc.)

DATA & ARCHITECTURE

SERVICE MANAGEMENT

GOVERNANCE

MONITOR & IMPLEMENT

Disaster Recovery 1st Wave

Need to do a full business impact analysis but..

These are likely to be amongst the first ones

- Authentication (AD, SSO, e-groups, Grappa, …)
- Safety (CSAM, RAMSES, …)
- Access Control (SUSI, ZORA, …)
- SCADA, Cooling/Ventilation + Controls
- Communication tools
  - E-Mail, Chat, Phone, …
- Application servers
  - Foundation, Qualiac, Payroll, IMPACT, HRT, EDH, EDMS, INFOR, …
- Core Web services

Disaster Recovery ‘Easy’ Ones

The (relatively) easy IT ones will also be good to do early

- Mkdocs (!)
- Gitlab, …
- Opendcim, …
- aiadm-homeless
- Puppet et al
- Linuxsoft
- Indico

Summary of Status Based on Live-IT Plans

Area	Milestone	Target	Status
BC	Business Impact Analysis with at least one sector	Q3 2023	On track (Mini)
	Update DMRB with refined resource plan and schedule	Q3 2023	On track
DR	Establish IT cross group co-ordination team	Q1 2023	Done
	Prepare as-is assessment on current capability	Q2 2023	On track (Draft)
	In collaboration with TD, refine implementation to support rapid recovery from incidents for IT critical services, including PCC and public cloud resources as needed	Q4 2023	On track
EA	Evaluate industry standard approach for personas, products, applications, services and ontology with FHR, IT, ATS	Q2 2023	Tight (Draft)
	Prepare project proposal for DMRB	Q3 2023	On track
	Adapt processes with ARB and TD for implementation & maintenance	Q4 2023	Delayed to Q1 2024

home.cern