



ETH Zurich

Speicherumgebung an der ETH Zürich – Lösungen und Herausforderungen



ETH Zurich at a glance



Founded in 1855

- Driving force of industrialisation in Switzerland

ETH Zurich today

- One of the world's leading universities for technology and the natural sciences
- Place of study, research and employment for 30,000 people from over 120 different countries

Reasons for success

- Excellent education
- Groundbreaking fundamental research
- Knowledge transfer that benefits society

Agenda

Short portrait ETH Zürich

ETH IT Services / Facts & Figures / Challenges

Storage Services

Service solutions

Security Measures

Challenges (HW, protocols, tape, SW)

Locations



Campus City Center

Two main locations in Zürich

- *In the city center:* historical main building, built by Gottfried Semper
- *On the Hönggerberg:* Modernern Campus, that exemplifies science, business and the public

Other locations in Switzerland

- *Basel:* Departement of Biosystems
- *Lugano:* Swiss High Performance Computing Center (CSCS)

Other decentralized units

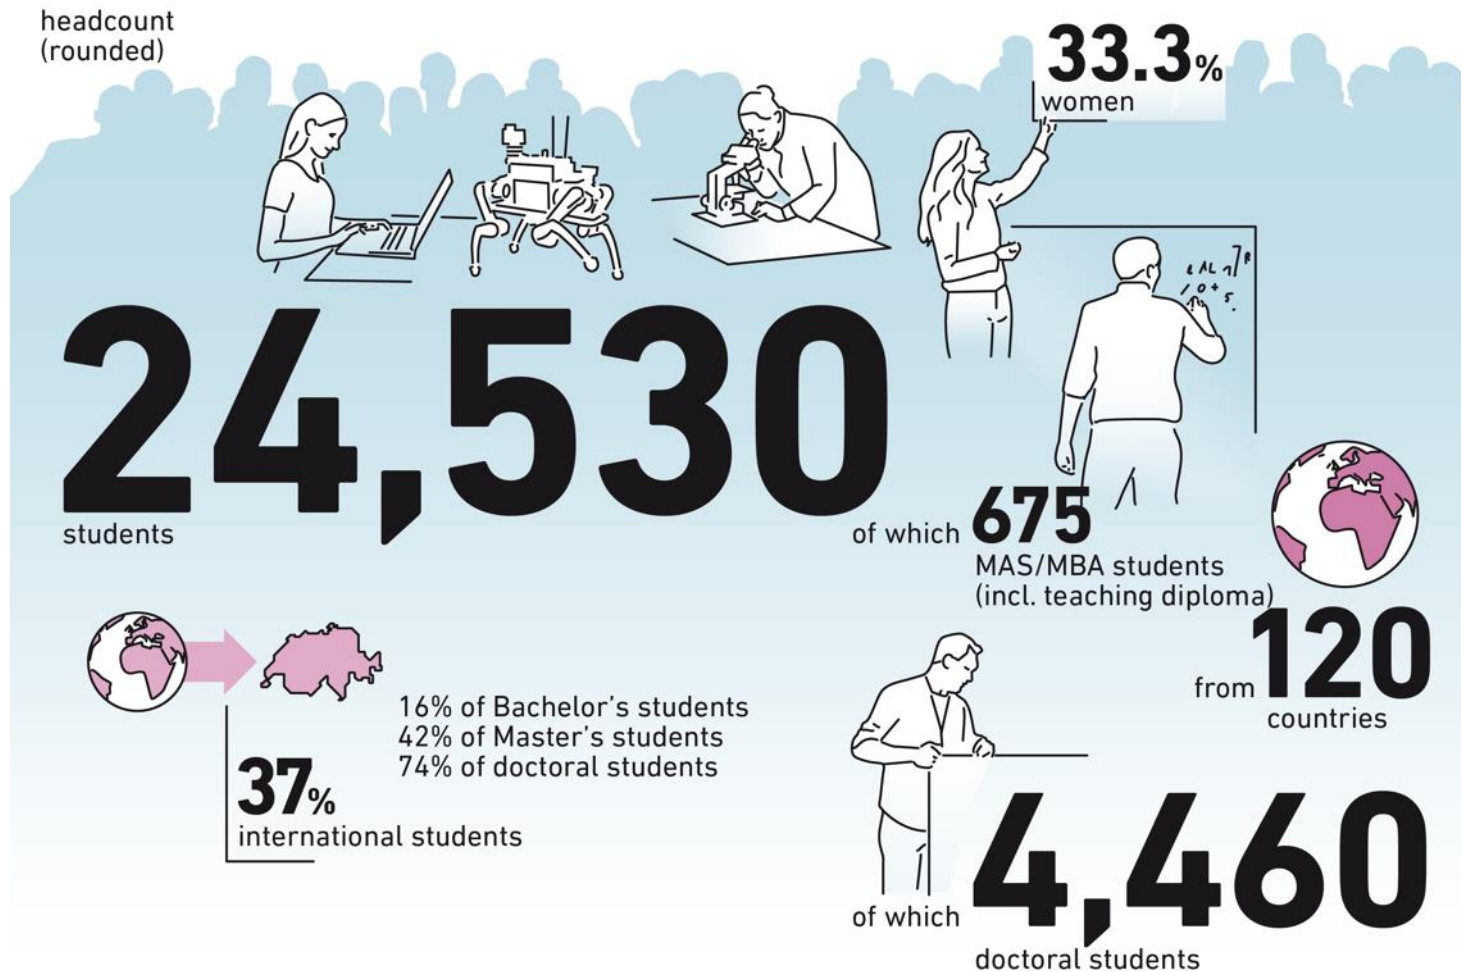
- *Research location in Singapore:* Singapore-ETH Centre (SEC)



Campus Science City, Hönggerberg

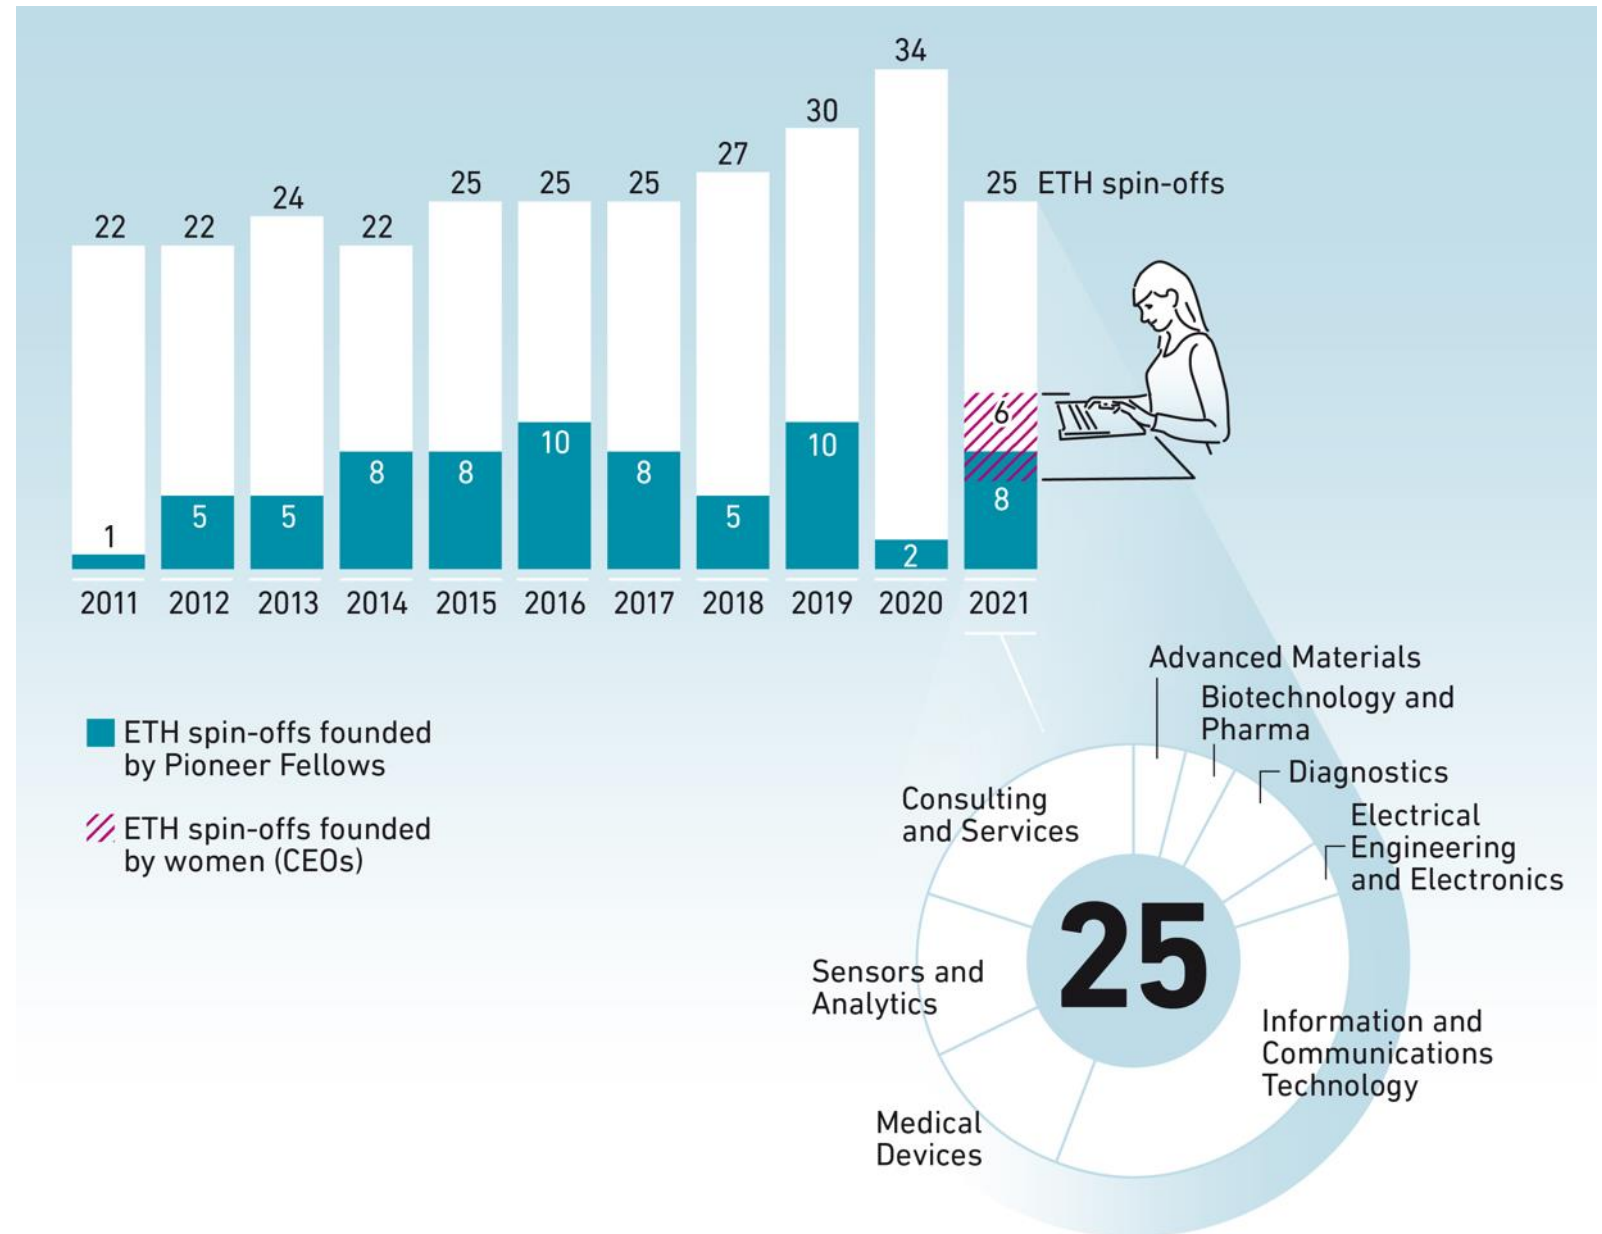


ETH Zurich in figures – students



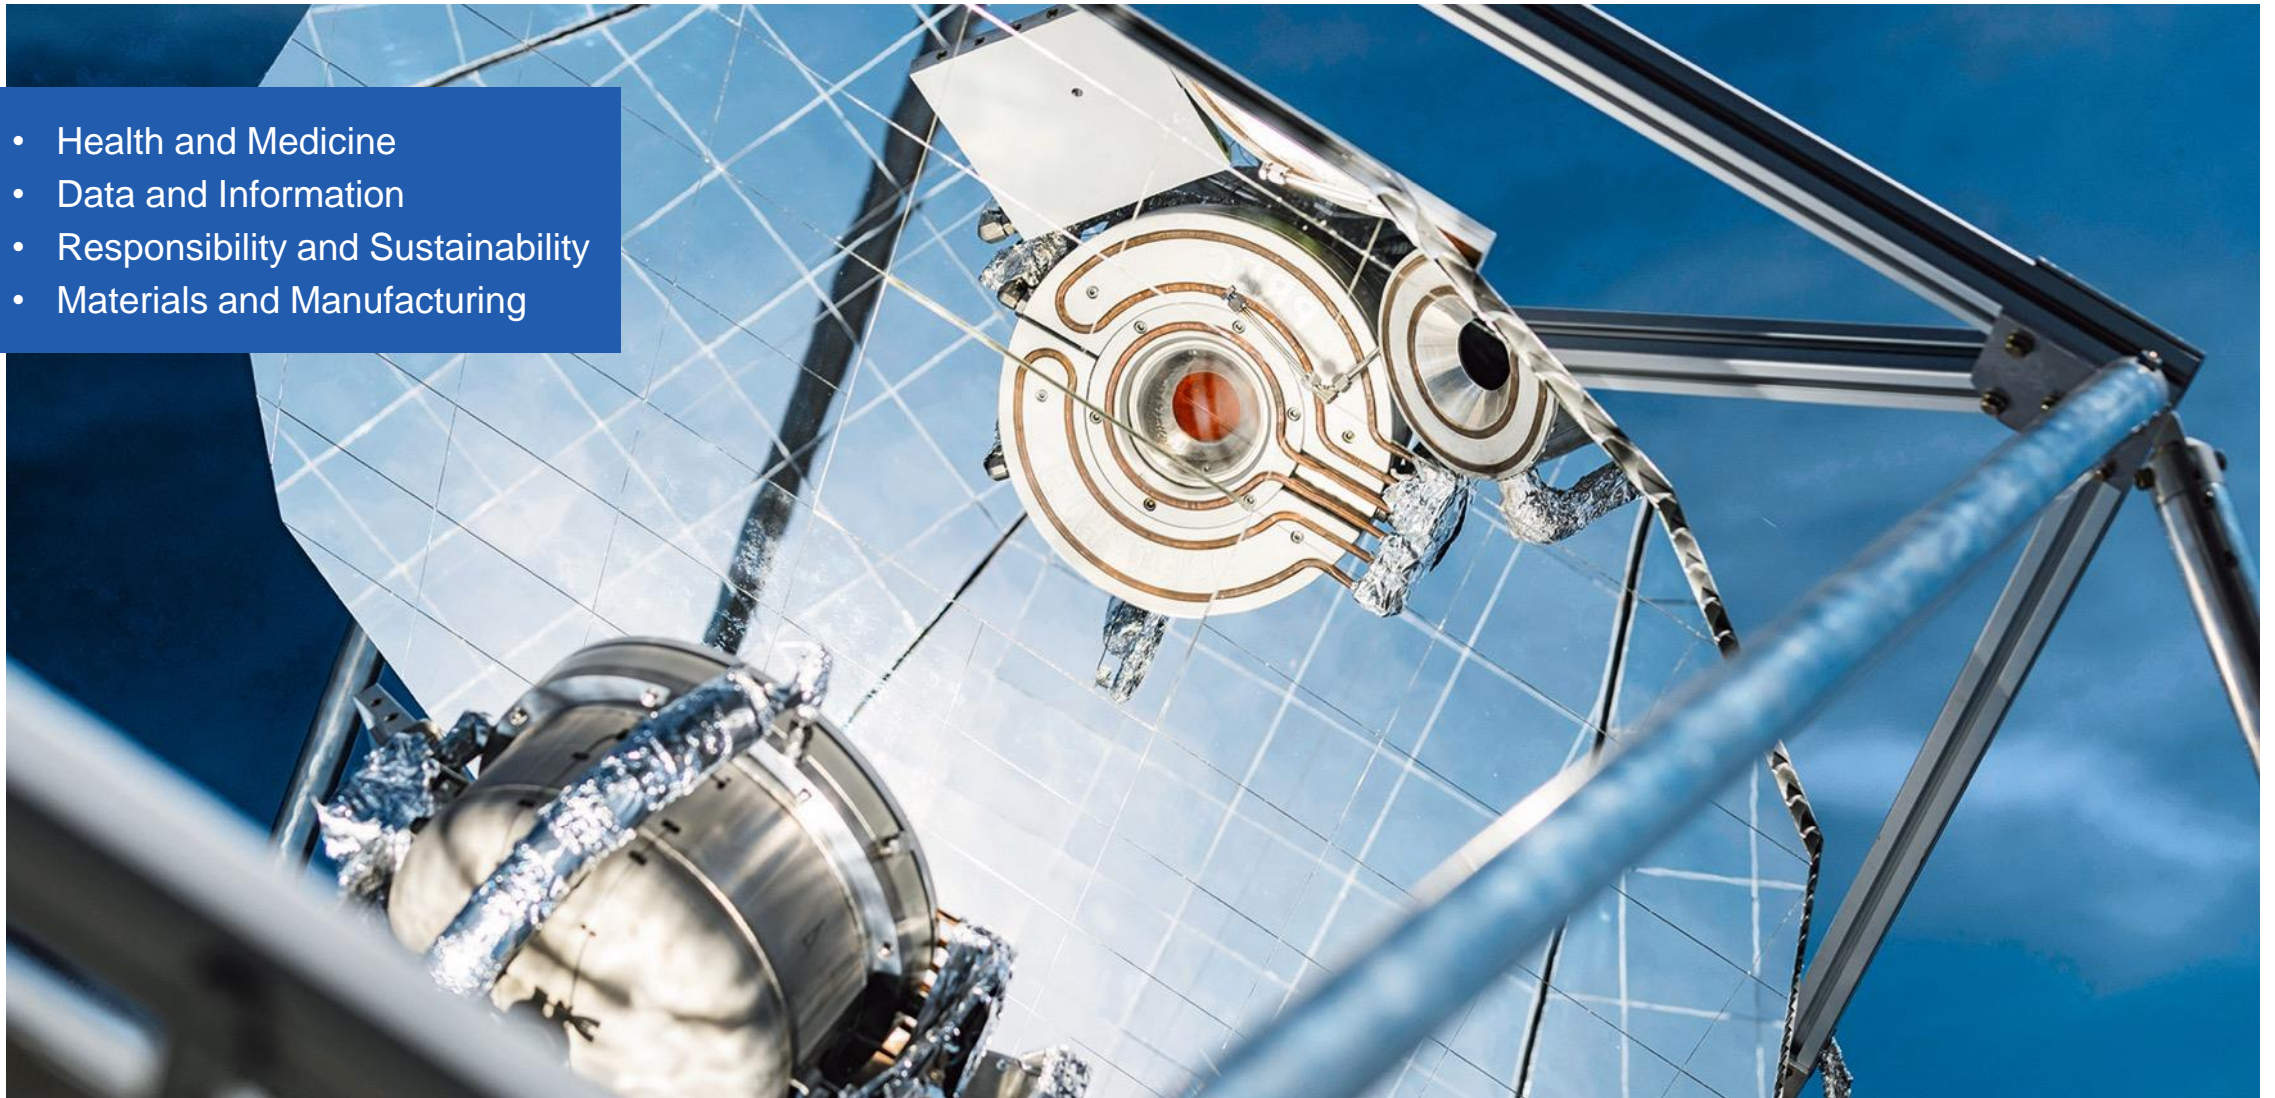
Stimulating the economy

Spin-offs

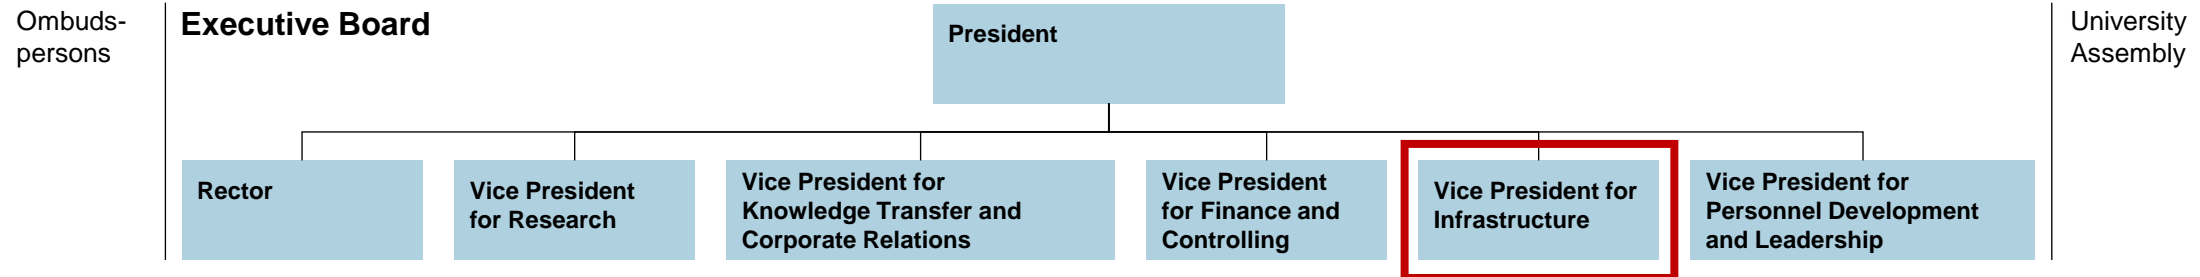


Strategic action areas

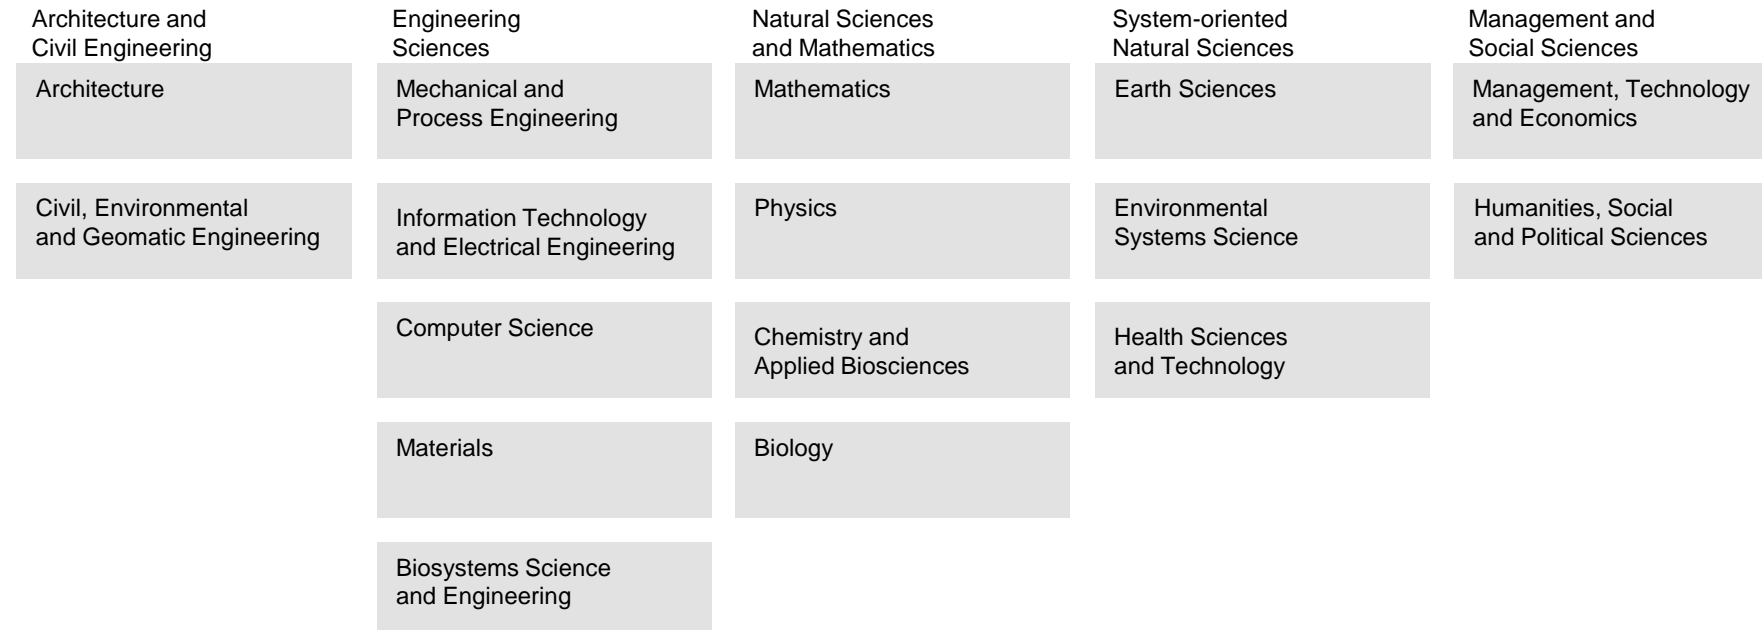
- Health and Medicine
- Data and Information
- Responsibility and Sustainability
- Materials and Manufacturing



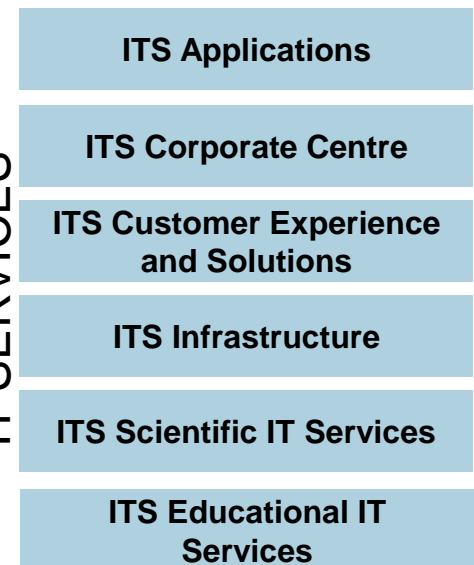
Organisational chart ETH 2024



16 academic departments

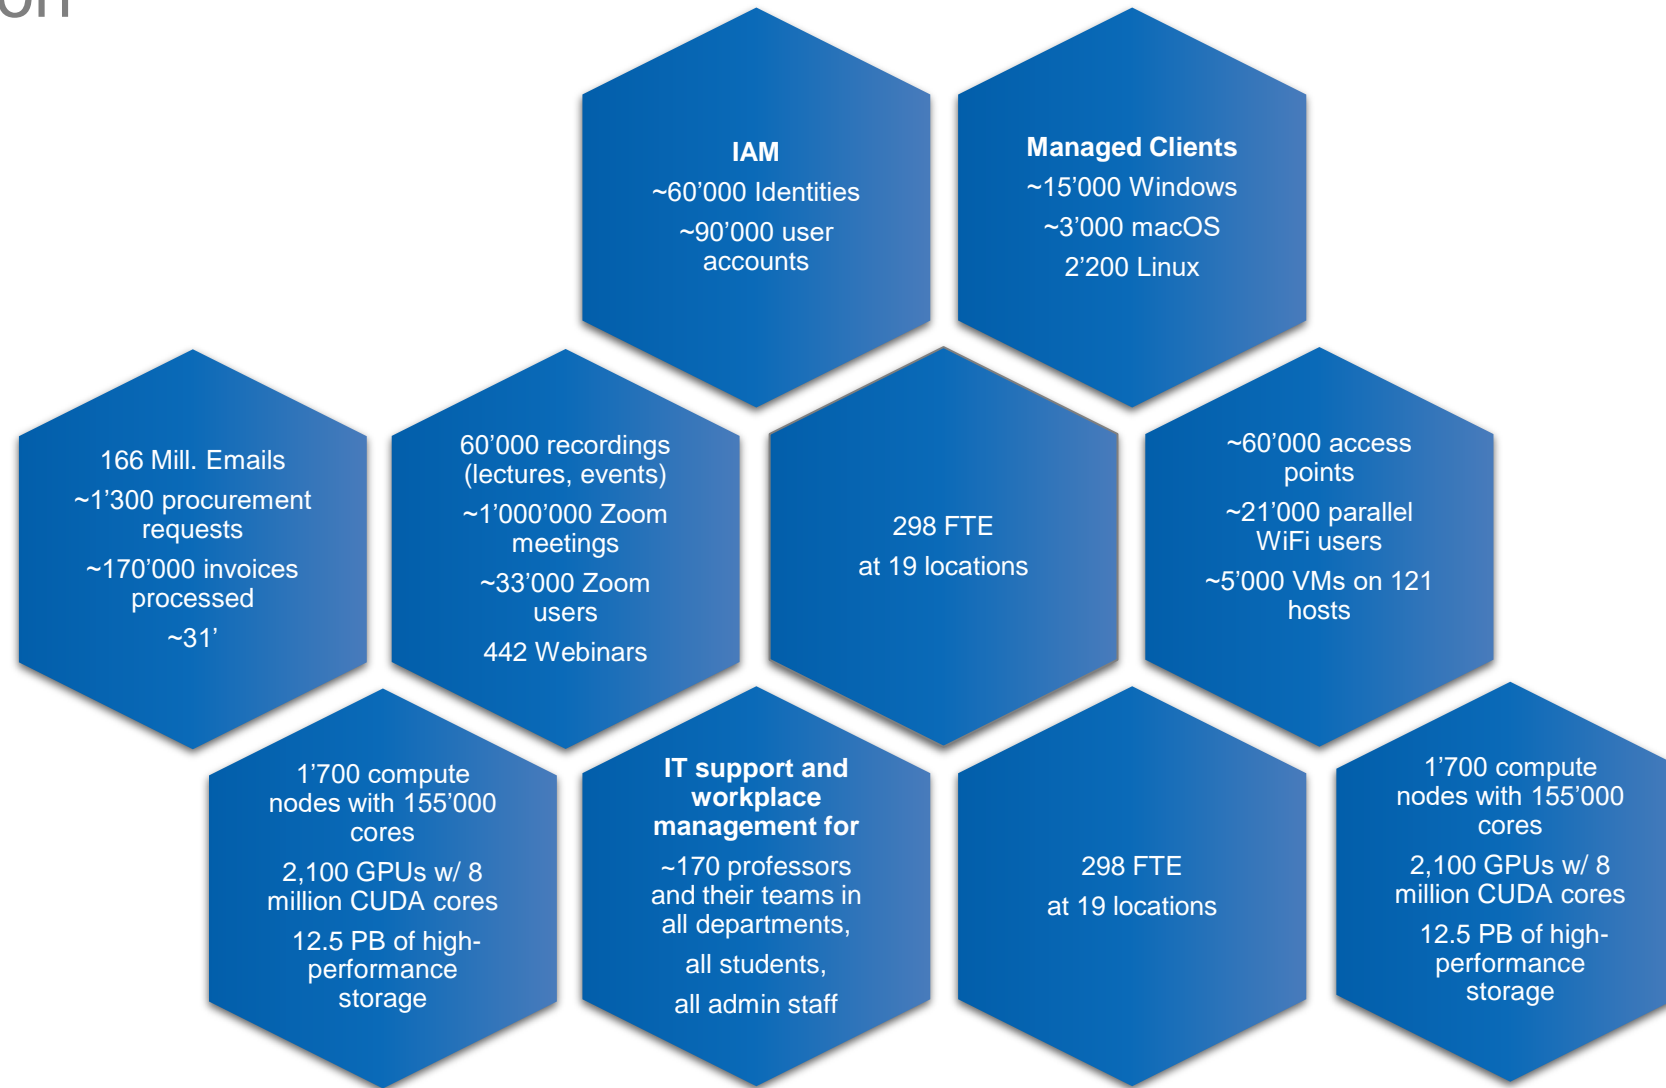


IT SERVICES



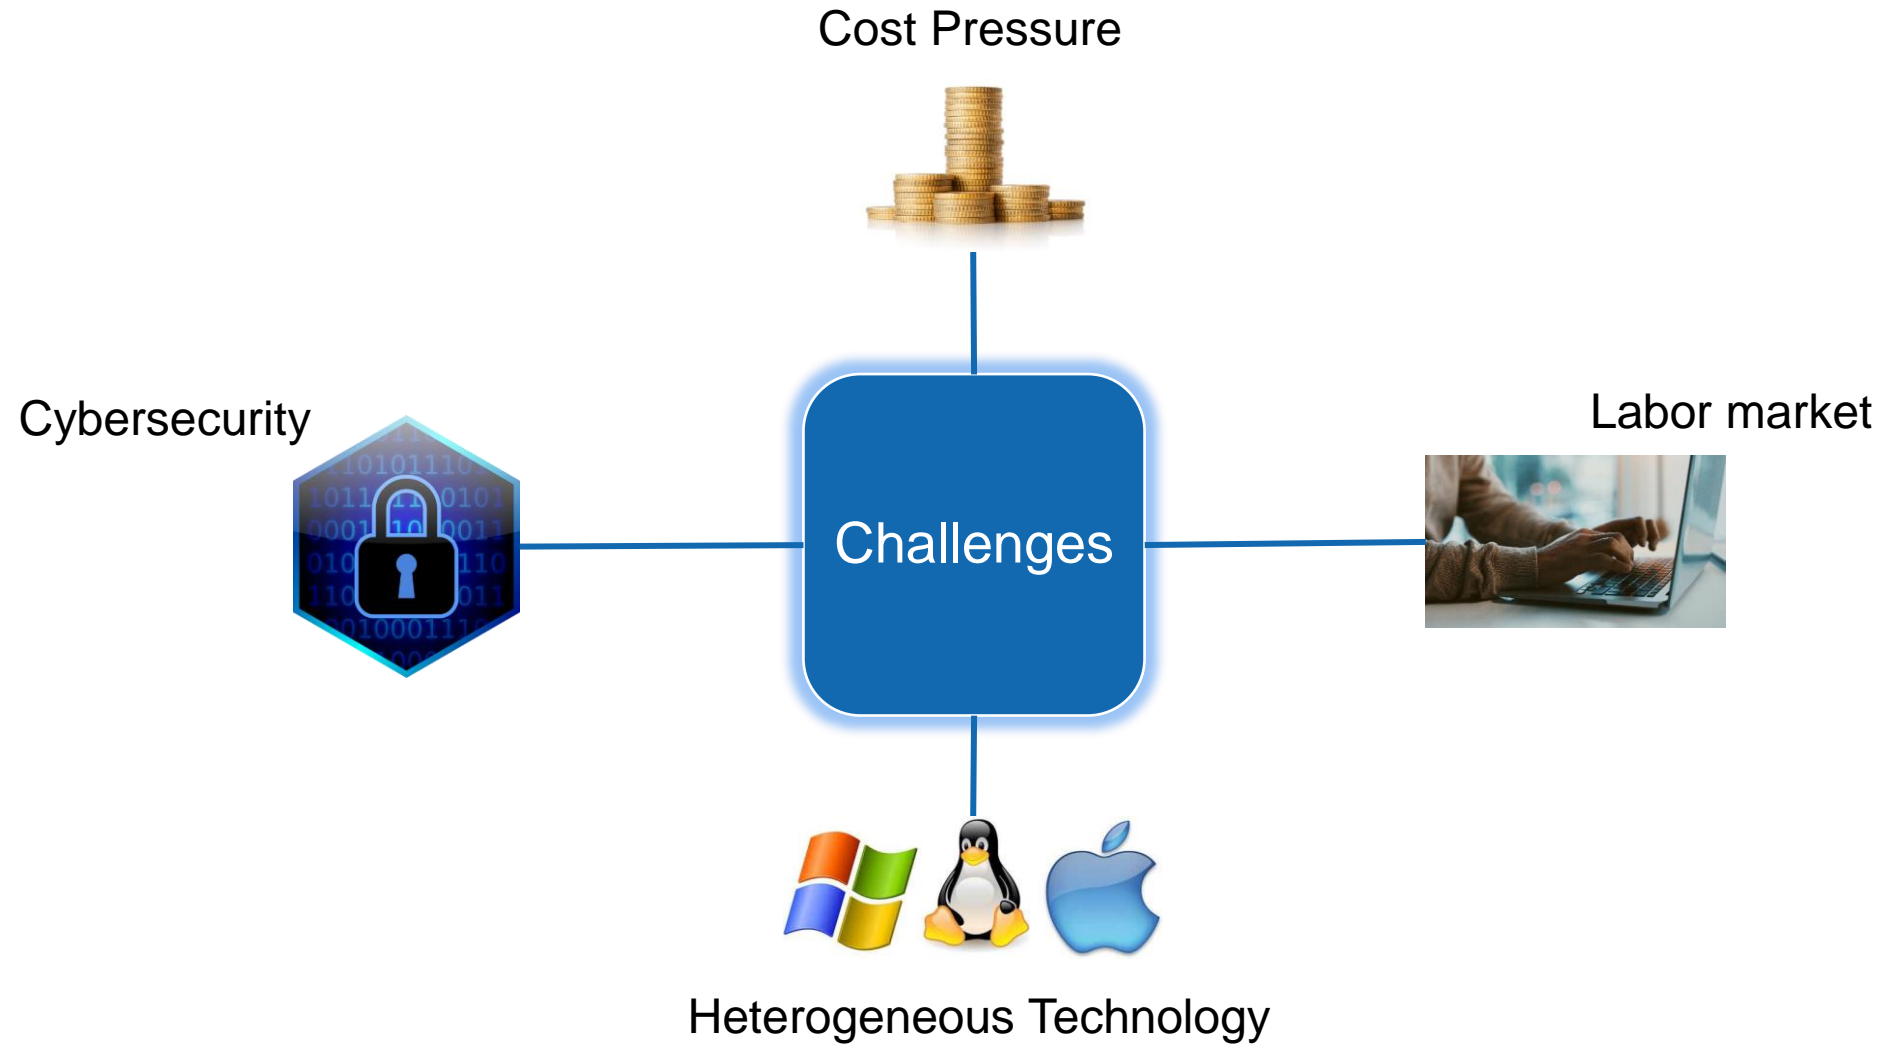
Facts & Figures Selection

As of Dec 2022



Challenges we are facing

Selection



ETH Storage Services Overview

Disk Only

NAS (*IBM Spectrum Scale*)

- High throughput
- Fully managed
- Backed up

CDS

- Cheap
- Not backed up

**SERVER BACKUP
CLIENT BACKUP**

POLYBOX

Tape Only

DBR - Backup & Restore (*Spectrum Protect*)

- Free
- Redundant

LTS - Archive & Retrieve (*StrongLink or SP*)

- Free
- Redundant

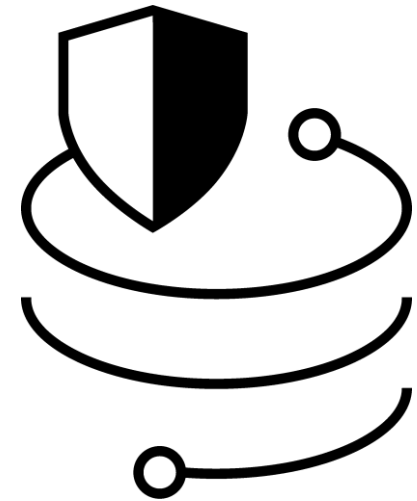
Disk / Tape Hybrid

HSM - Migrate & Recall (*Spectrum Protect*)

- Free
- Could be redundant
- Mostly migrated to LTS

ETH Storage Services Overview – some numbers

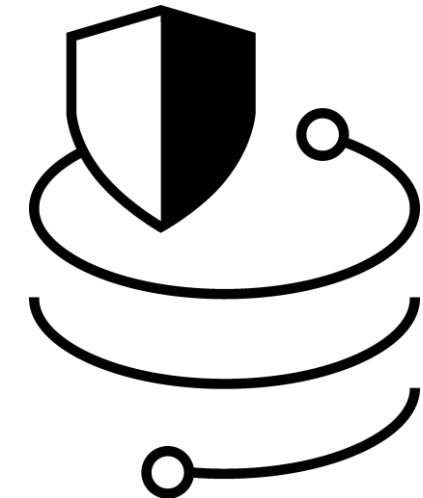
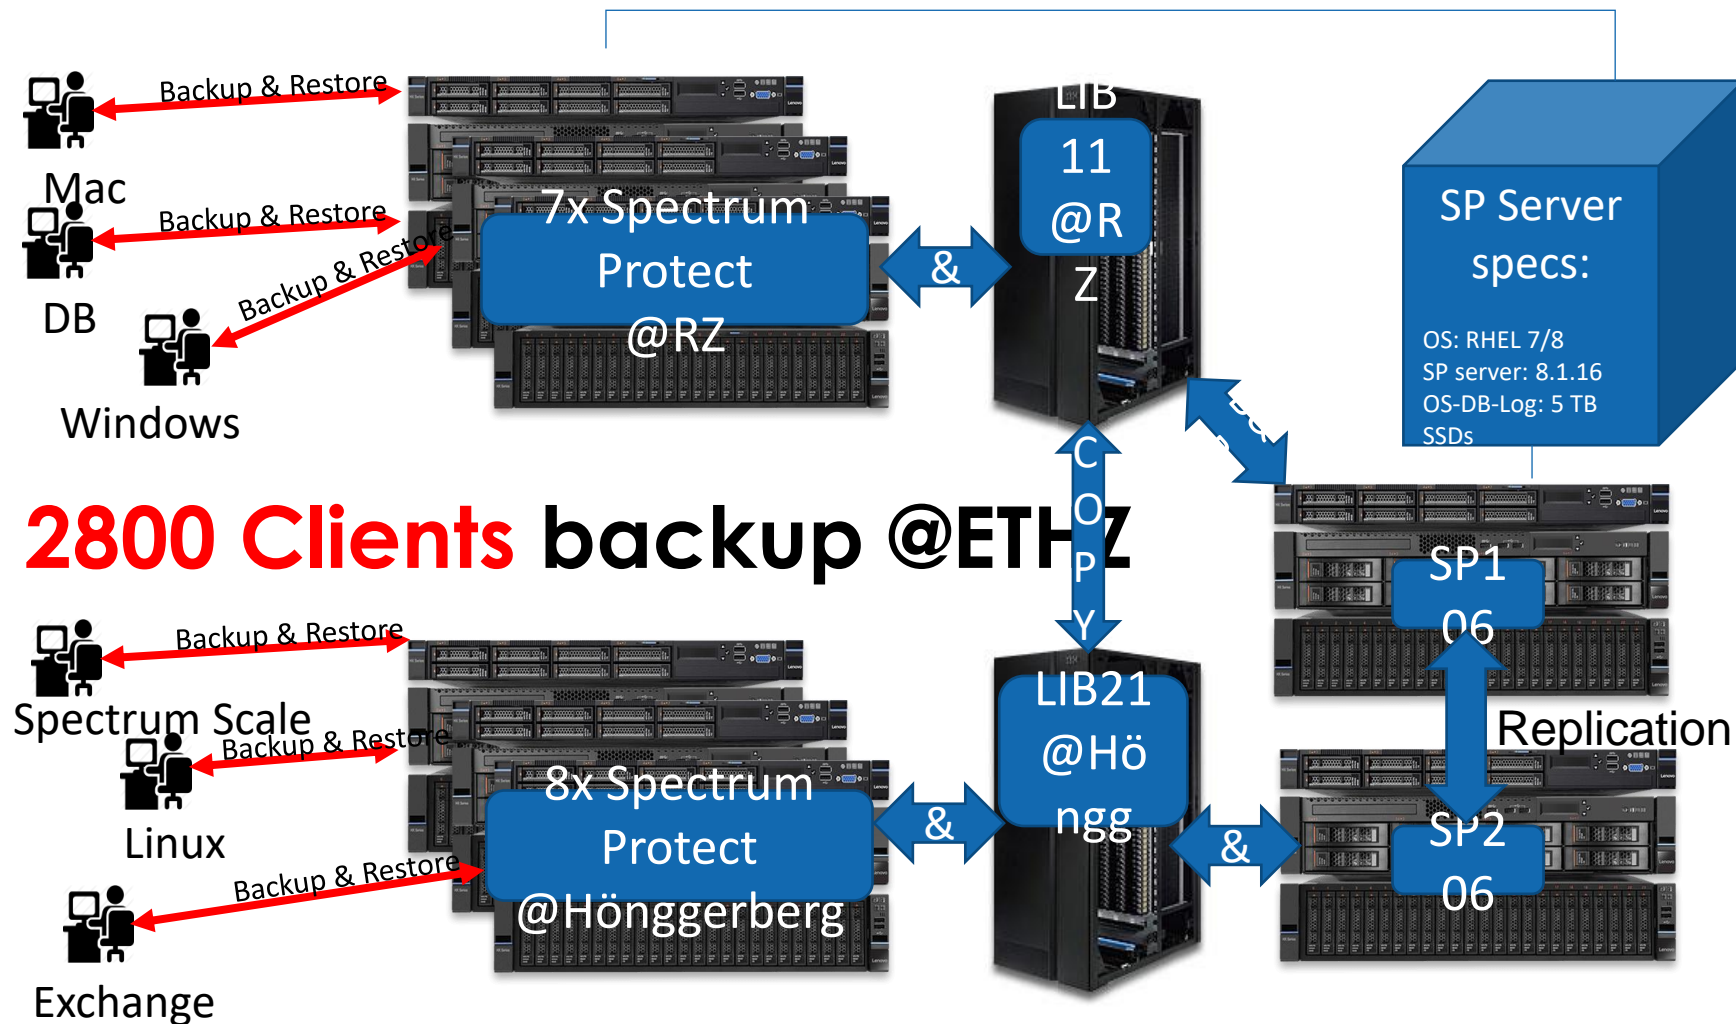
- ~100 PB of tape-only backup (Spectrum Protect)
- ~400 TB of disk-only backup (IBM/Cohesity)
- ~20 PB of disk-only online data (Spectrum Scale/CDS)
- **~24 PB of tape-only archive (StrongLink)**



ETH Storage Vendors

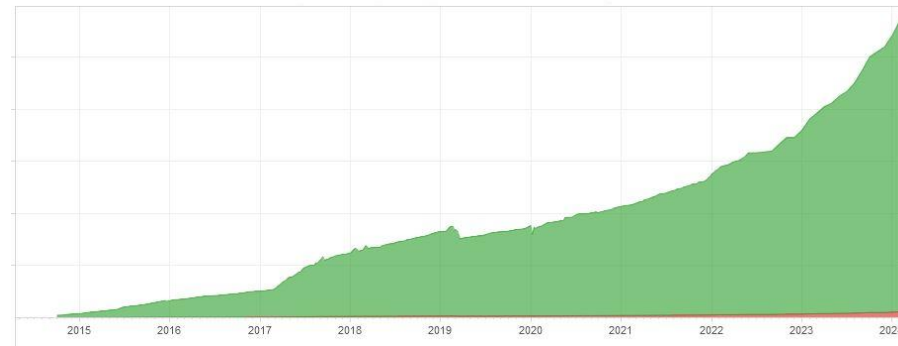


ETH Data Backup and Restore Service Solution

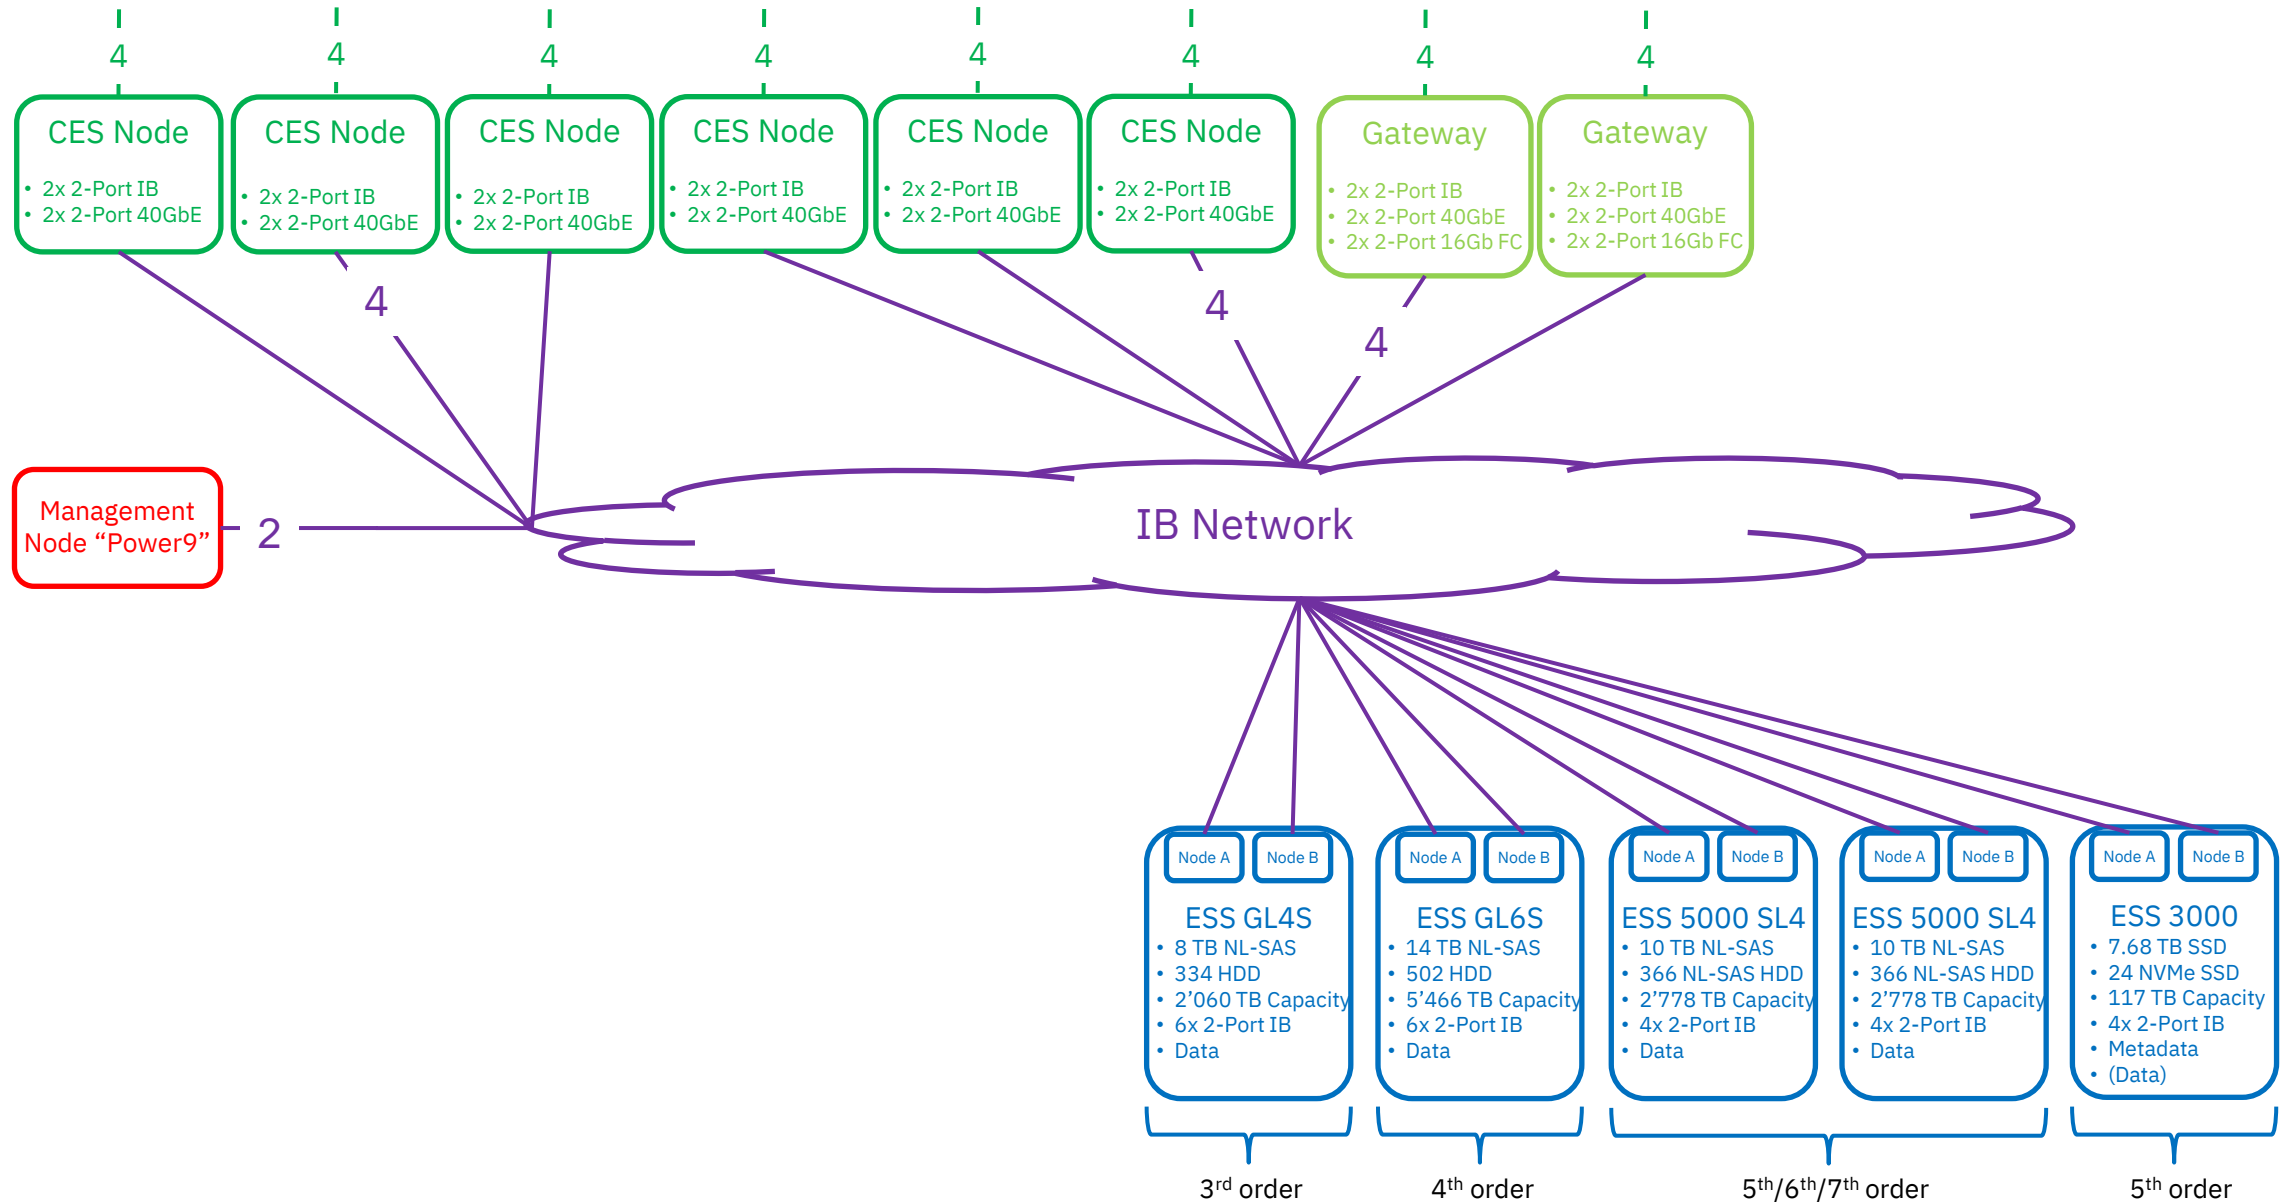


Tape archive numbers (Long Term Storage) – Service Solution

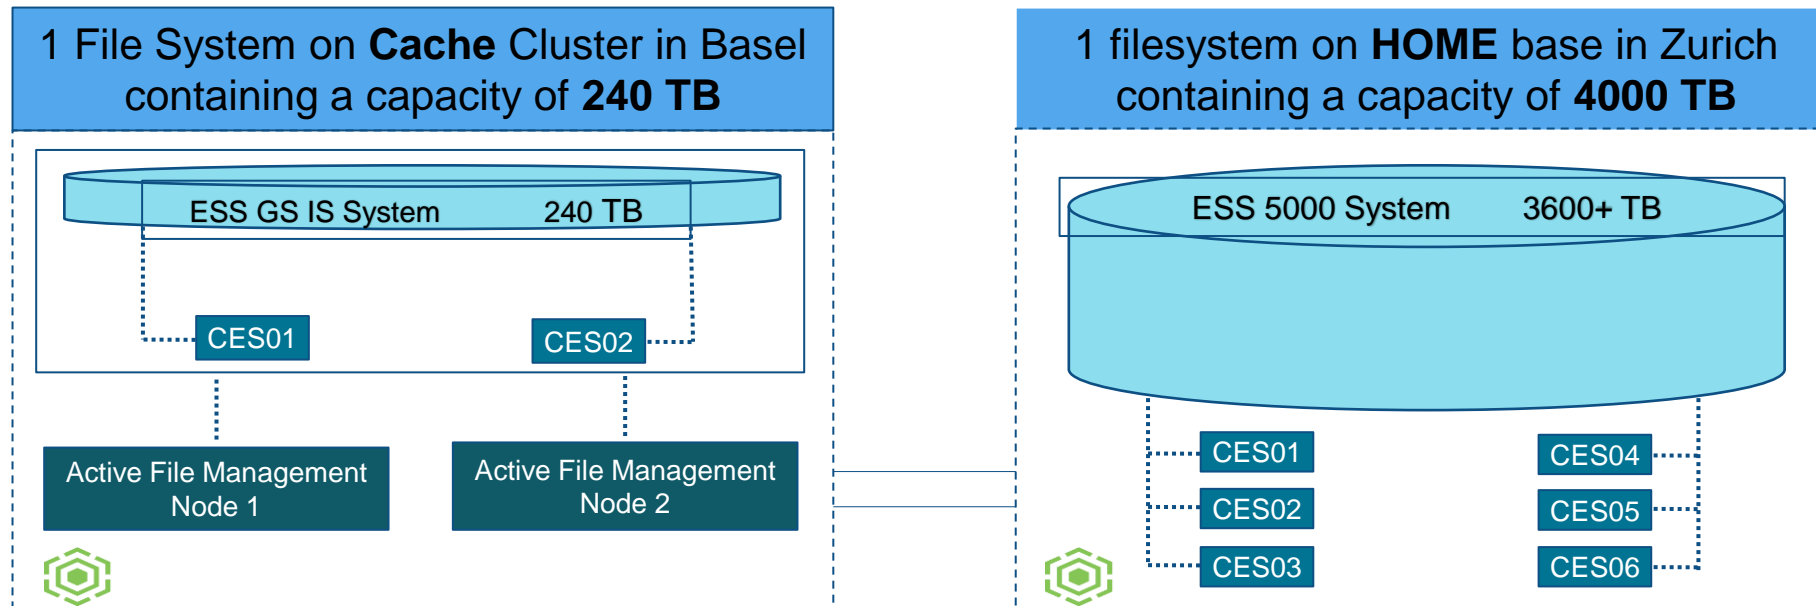
- Around 24 PB on tape (*dual copy*) – 12 PB native
- 2 copies on separate locations of every single file
- 40 shares (NFS & SMB)– one for each department
- ½ million files archived
- File sizes usually between 100 MB to 1.5 TB



Spectrum Scale Service Solution at one of two campus locations



AFM Cache and Home Systems Service Solution



Spanning a distance of 120 KM, 2 AFM nodes enable the transfer of data via a CACHE cluster at 40 GB per second each. To leverage Spectrum Protect servers, Tape-based backup runs against HOME filesets. Six CES docking points for end users enable remote access of the copied data, generating a central storage network with high scalability and at low cost

Security Measures – Tiered Backup

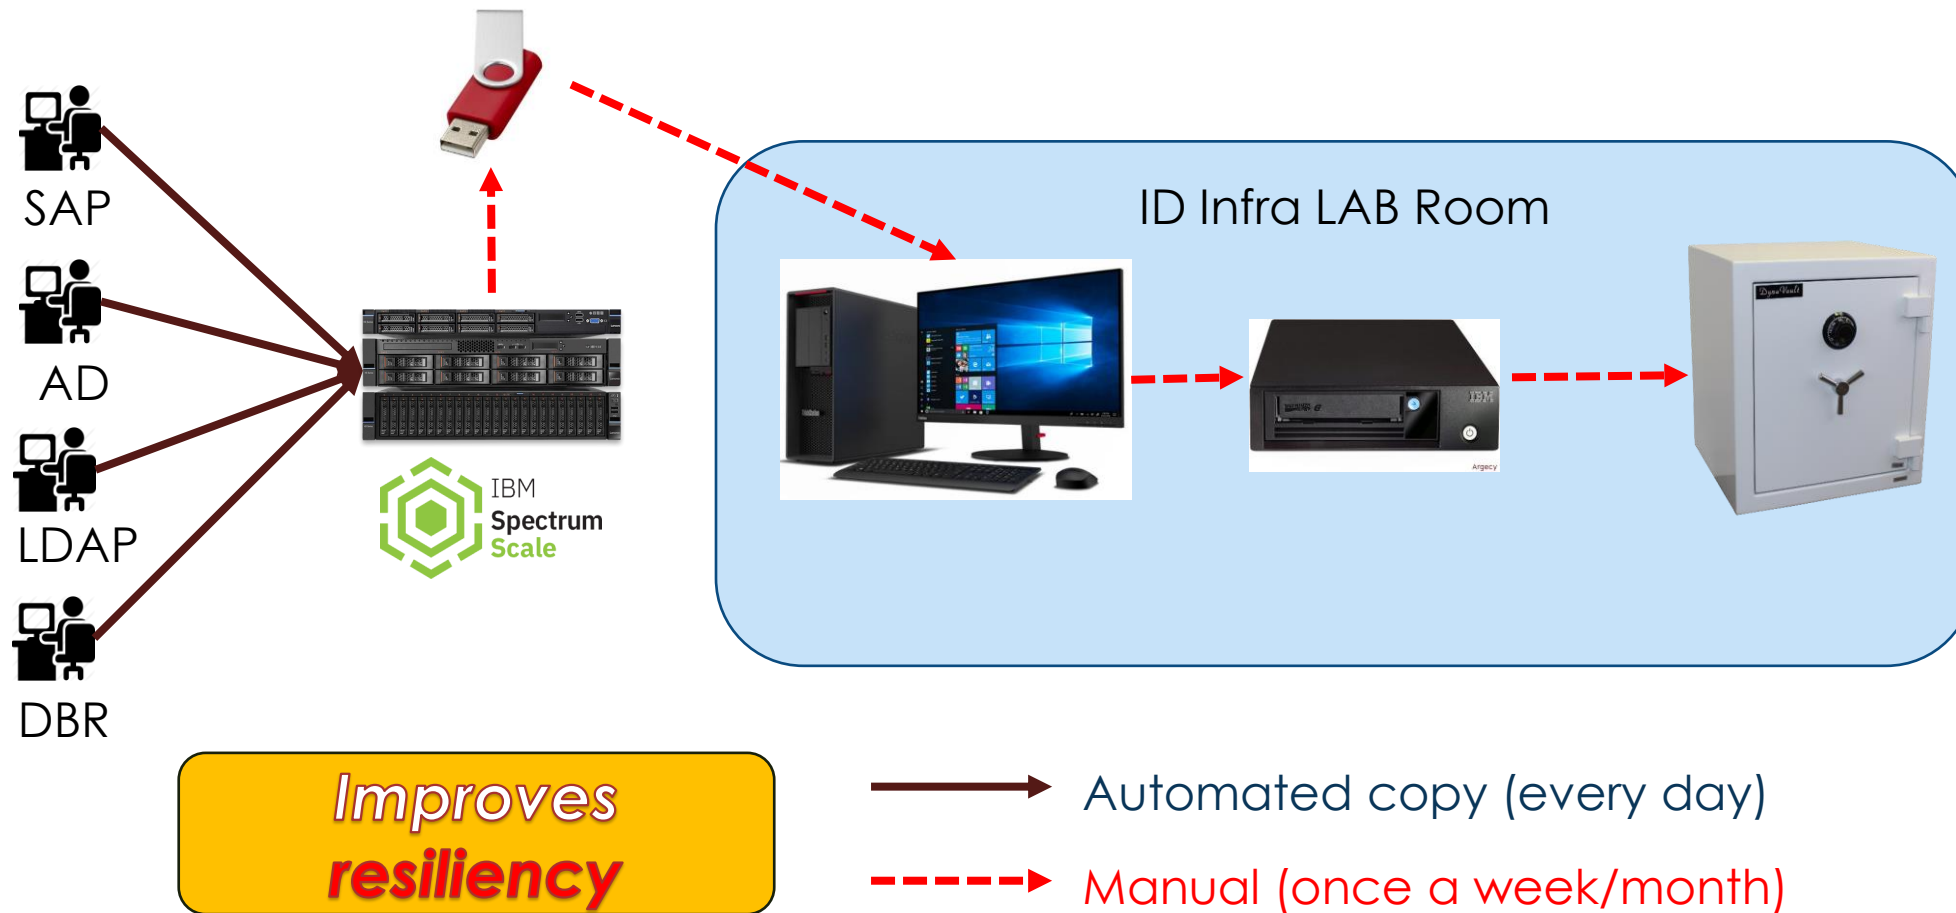
- **Tier 4.** – Common user data with *dual tape copy*
- **Tier 3.** ~ **1 PB** – Important data with *fast restore* required
- **Tier 2.** ~ **200 TB** – Critical data with *logical airgap & mirror*
- **Tier 1.** ~ **20 TB** – Vital data with physical *airgapped mirror*

Can be done **ONLY** if we **first classify data** based on importance, urgency and requirements!

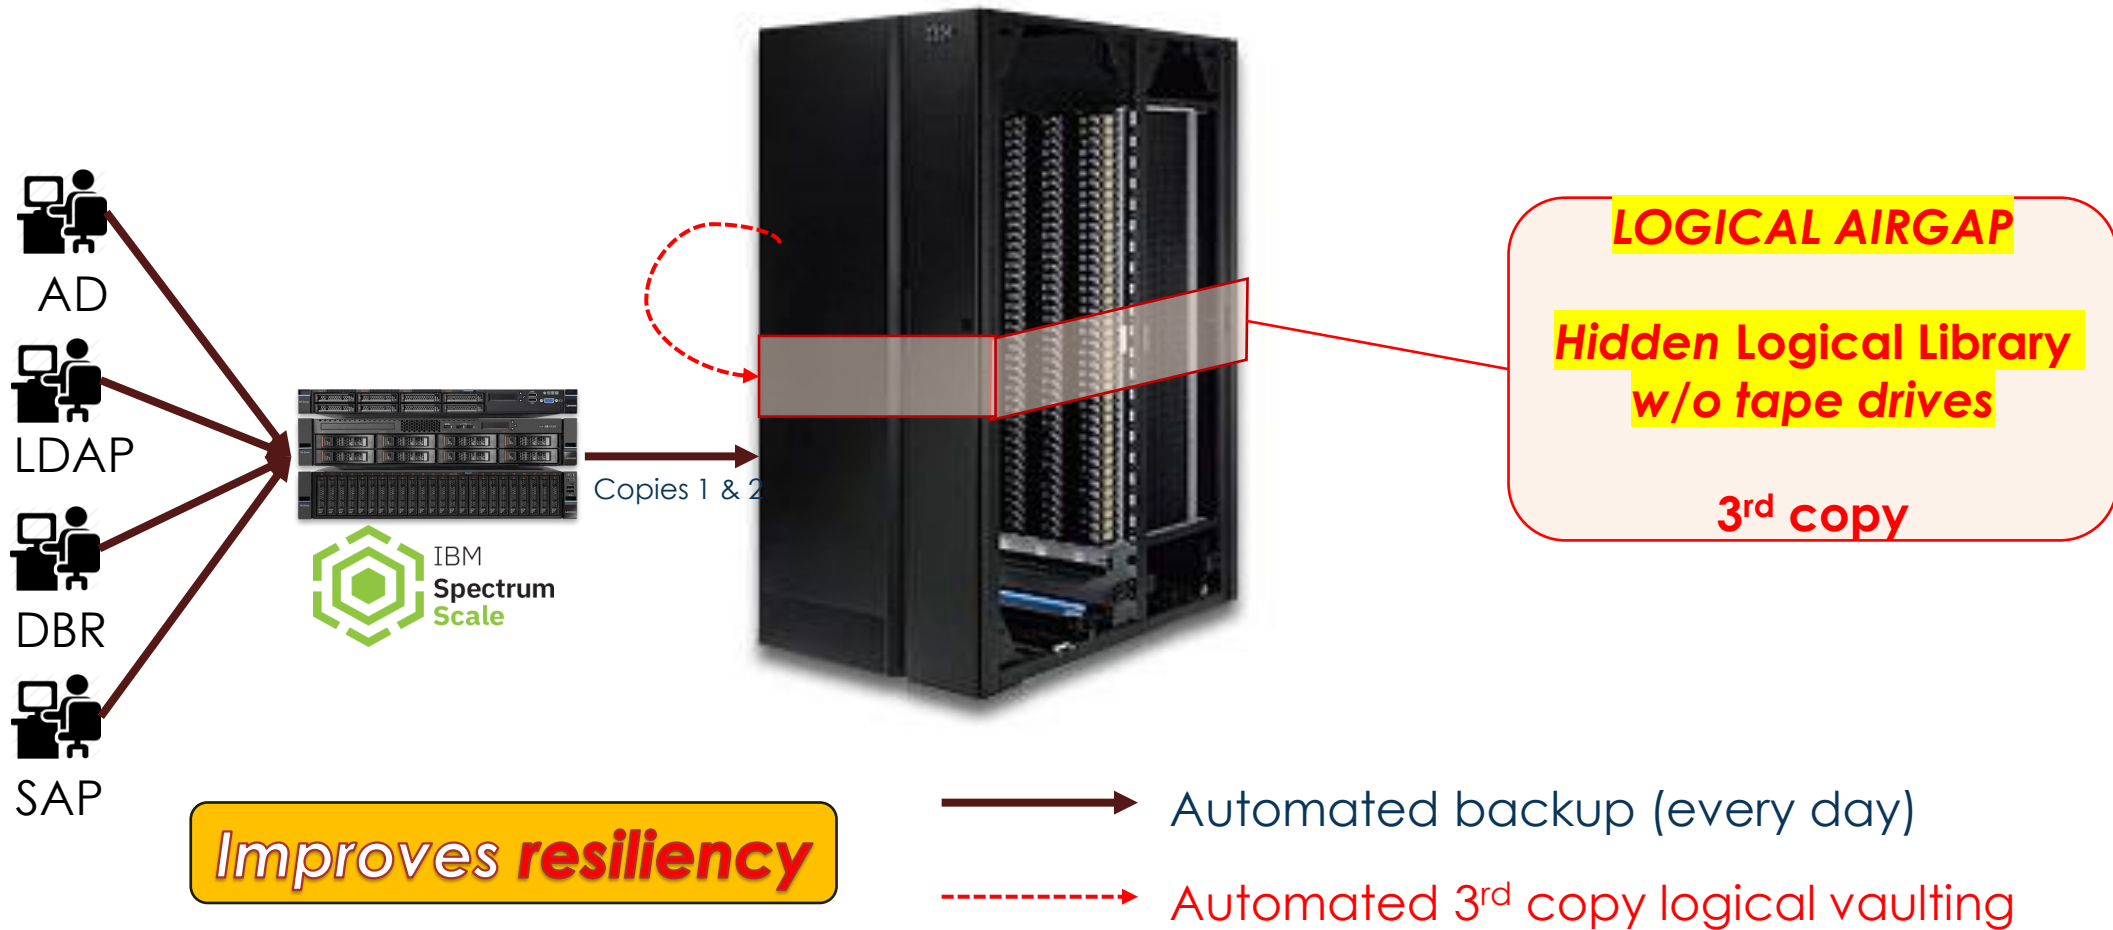


Crown Jewels

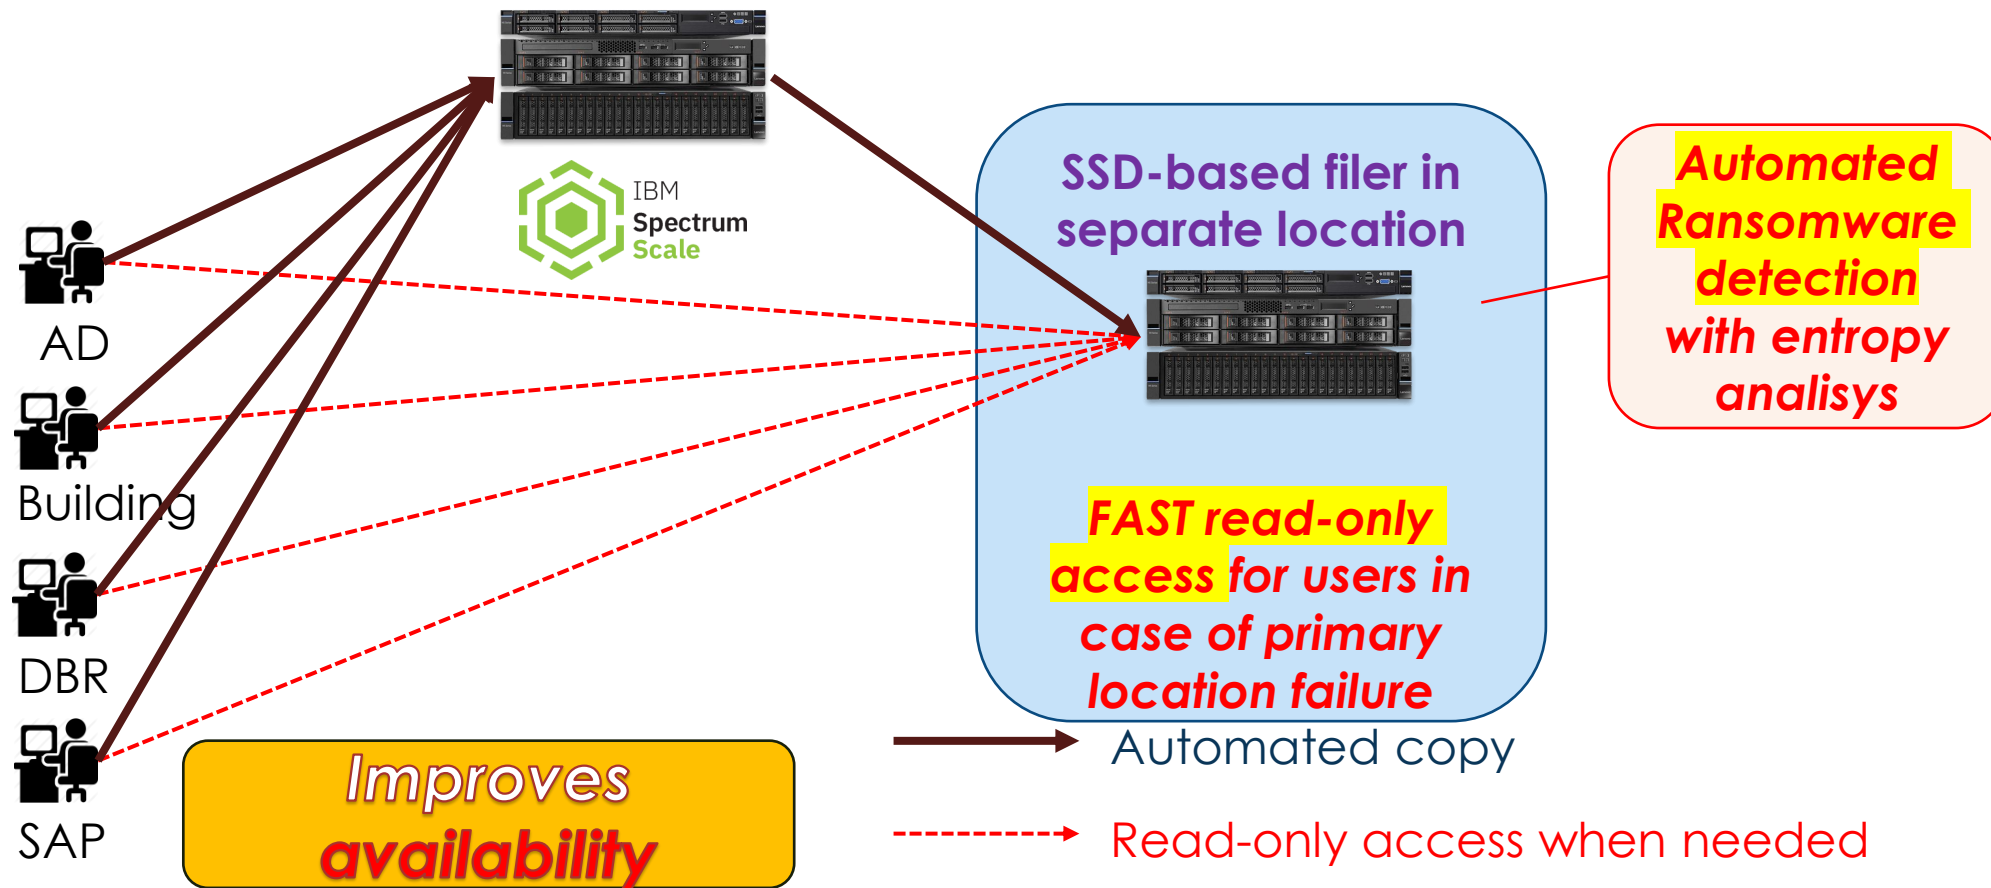
Security Measures – the physical air gap vault



Security Measures – the offline backup copy



Security Measures – the online mirror for Tier 1 and Tier 2



Challenges

- Huge daily growth vs. limited budgets
 - Disk space & network throughput
 - Tape drives & FC throughput
- Protocols (NFS, SMB, S3 all within the SDS layer)
- HW Renewals
- Future of tape
 - will it be tailored for hyper scalers?
- Cyber attacks

Challenges,
challenges
***and more
challenges!***

Answers to your questions ...



Contact:

Johan Verstrepen
johan.verstrepen@id.ethz.ch

Dr. Tilo Steiger
steiger@ethz.ch

Open position: Linux and Spectrum Scale admin