CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

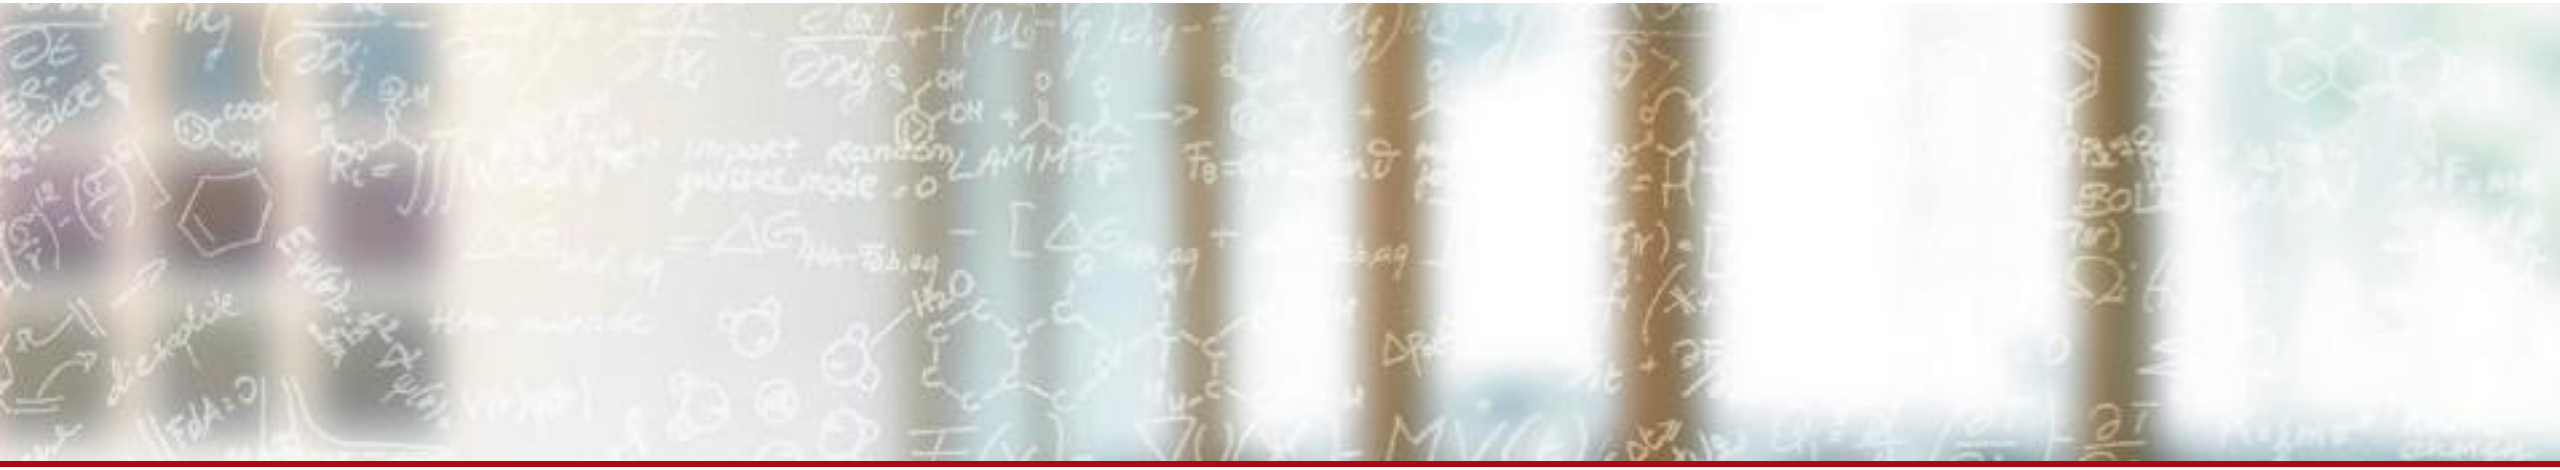

Seamless vCluster migration in CSM

CUG 2024

Miguel Gila, Manuel Sopena Ballesteros, CSCS

May 2024

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

About CSCS

A unit of the Swiss Federal Institute of Technology Zurich (ETH Zurich)

Founded in 1991 in Manno, and moved to Lugano in 2012

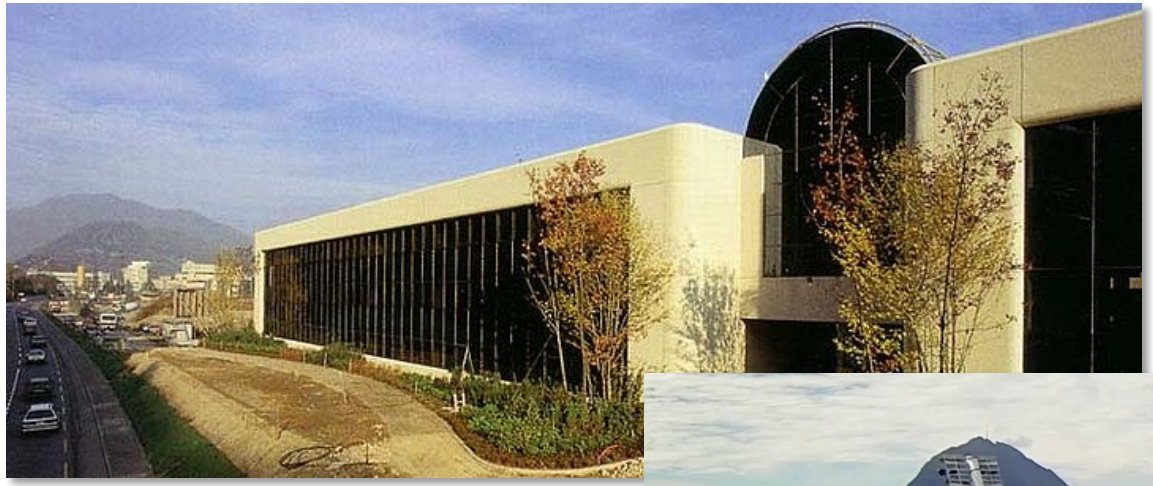

Mission

- We develop and operate a high-performance computing and data research infrastructure that supports world-class science in Switzerland

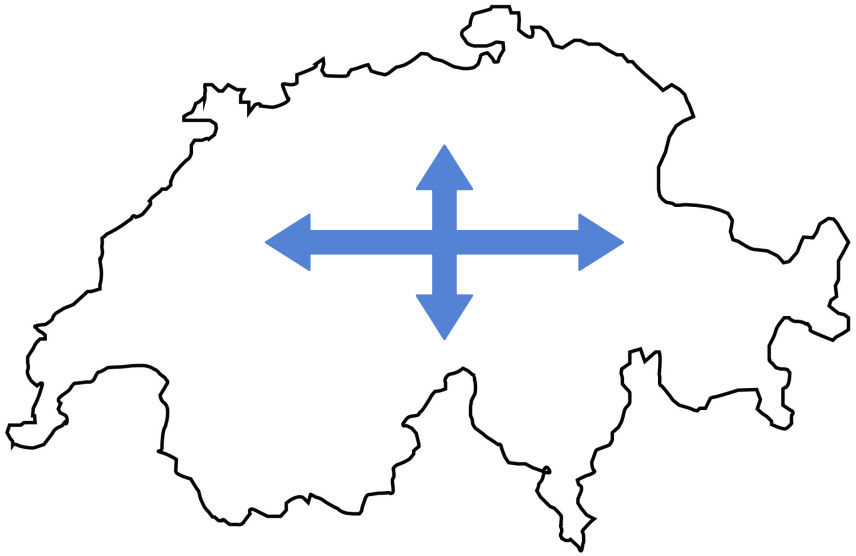

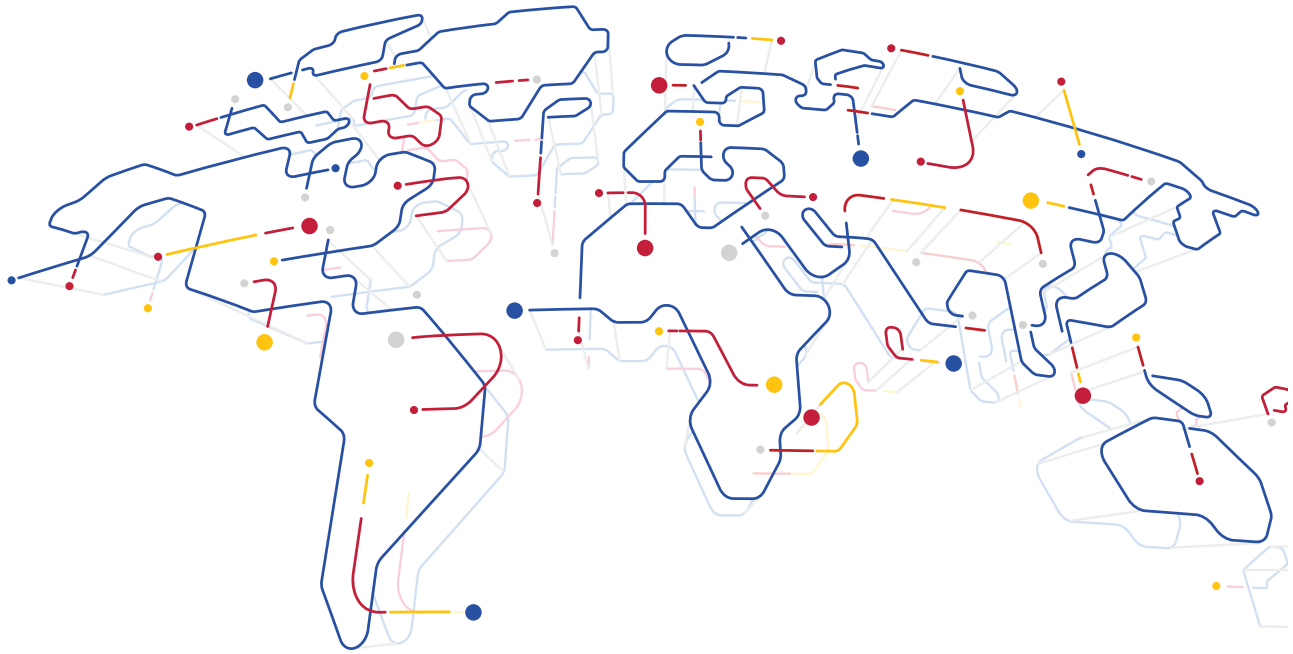

30 years of supercomputing at CSCS

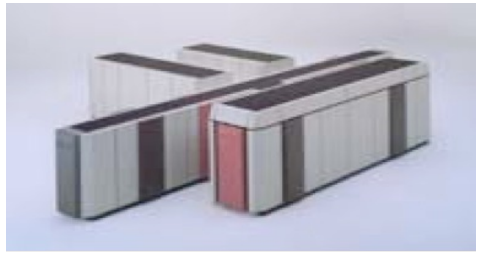

1991 NEC SX3
5.5 GF Adula

1996 NEC SX4
10 GF Gottardo

1999 NEC SX5
64 GF Prometeo

2002 IBM SP4
1.3 TF Venus

2005 Cray XT3
5.8 TF Palu

2006 IBM P5
4.5 TF Blanc

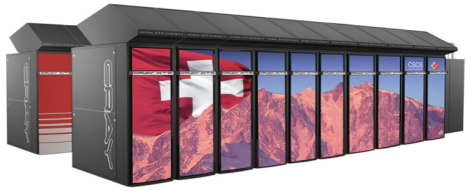

2009-12 Cray XE6 **402 TF**
Monte Rosa

2012-13-16 Cray XC40 / XC50
25 + 2 PF Piz Daint

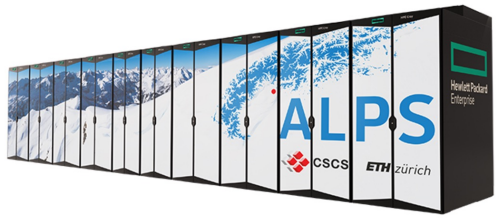

2020 HPE Cray EX
Alps

CSCS is constantly evolving

- Up until very recently, those supercomputers have been operated as single-purpose, vertically integrated entities, almost independent from each other

- In 2016 we started the process of resource consolidation, and in ~2019 the organization was reshaped to adapt to this new operational model

- CSCS went from running many disjoint systems with different scopes, to running a single, bigger infrastructure with more users, and services better tailored for each user community

- Several of those services are critical and must be available 24/7:
  - CSCS operates a supercomputing service for MeteoSwiss

Location of our infrastructure

- As we consolidated services into platforms, and systems into bigger infrastructures, naturally we had to answer the question of: what happens if the big infrastructure suffers a problem?

- We decided to place some parts of the infrastructure on alternative locations:
  - **CSCS @Lugano, Switzerland**
  - EPFL @Lausanne, Switzerland
  - ECMWF @Bologna, Italy

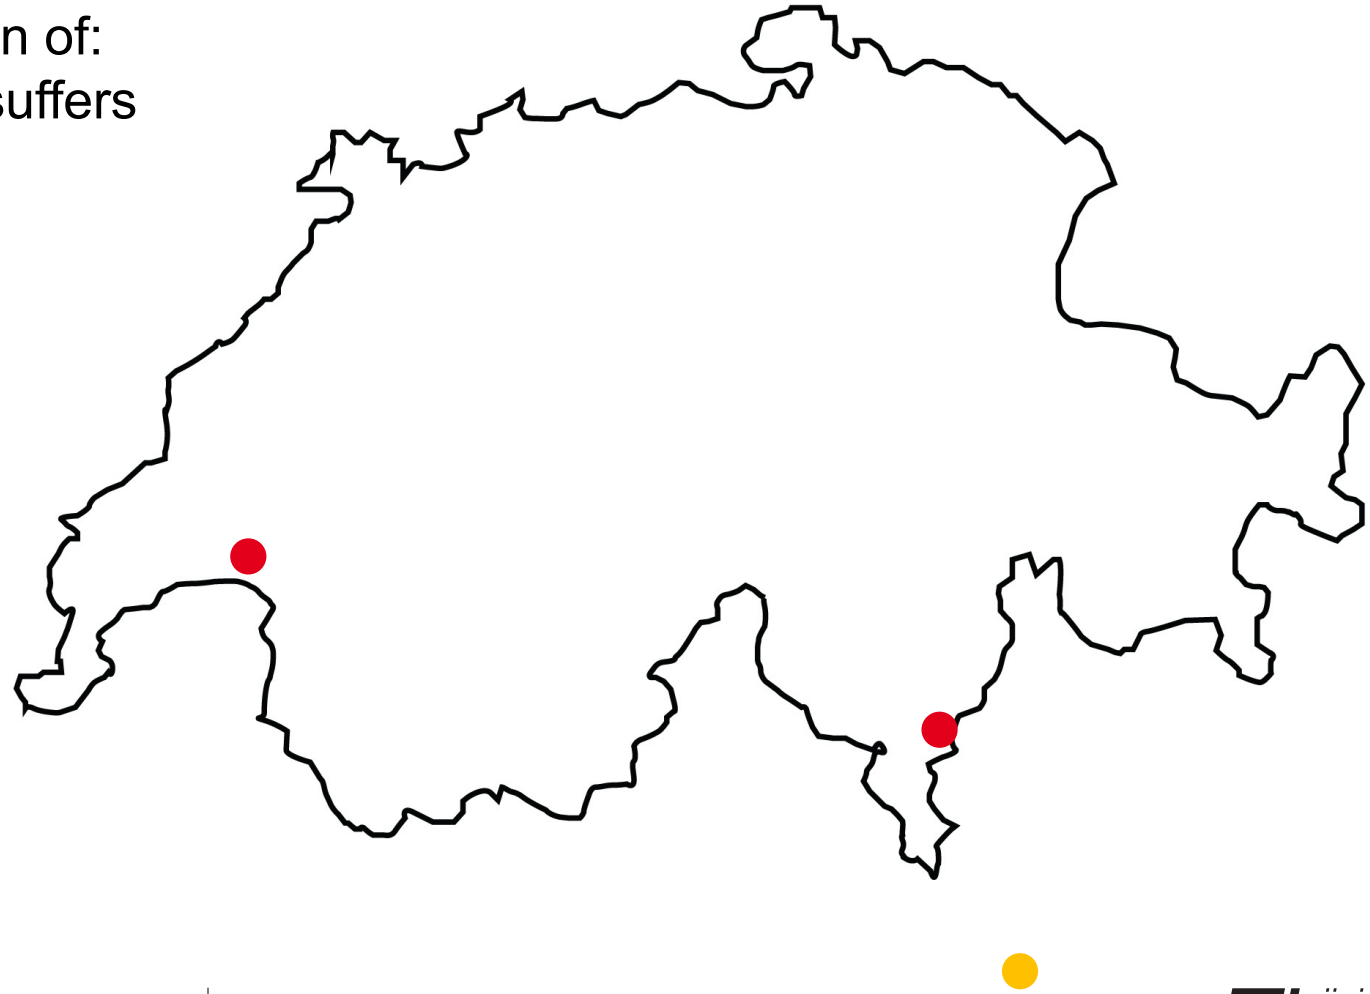

Alps

- Alps is an HPE Cray-EX system
- Phase 0 and Phase 1 are already installed
  - **2020: Phase 0** – AMD Rome CPU nodes
  - **2022: Phase 1** – NVIDIA A100, AMD Mi250x and AMD Milan CPU nodes
- Phase 2
  - **Q1 2024: Phase 2** - NVIDIA Grace-Hopper GH200 nodes

The User Lab scale out Grace-Hopper platform will be the largest **tenant** on Alps

It will replace Daint-GPU with the same number of Grace-Hopper modules (>5000) as there are on Daint-GPU today.

Tenants, Platforms and versatile Clusters (vClusters)

- Infra tenant: institution having access to the management plane
- Platform tenant: institution, project, community that manages the services on a set of platforms
- Platform: a set of vClusters that answer a business/scientific need
- vCluster: a set of services

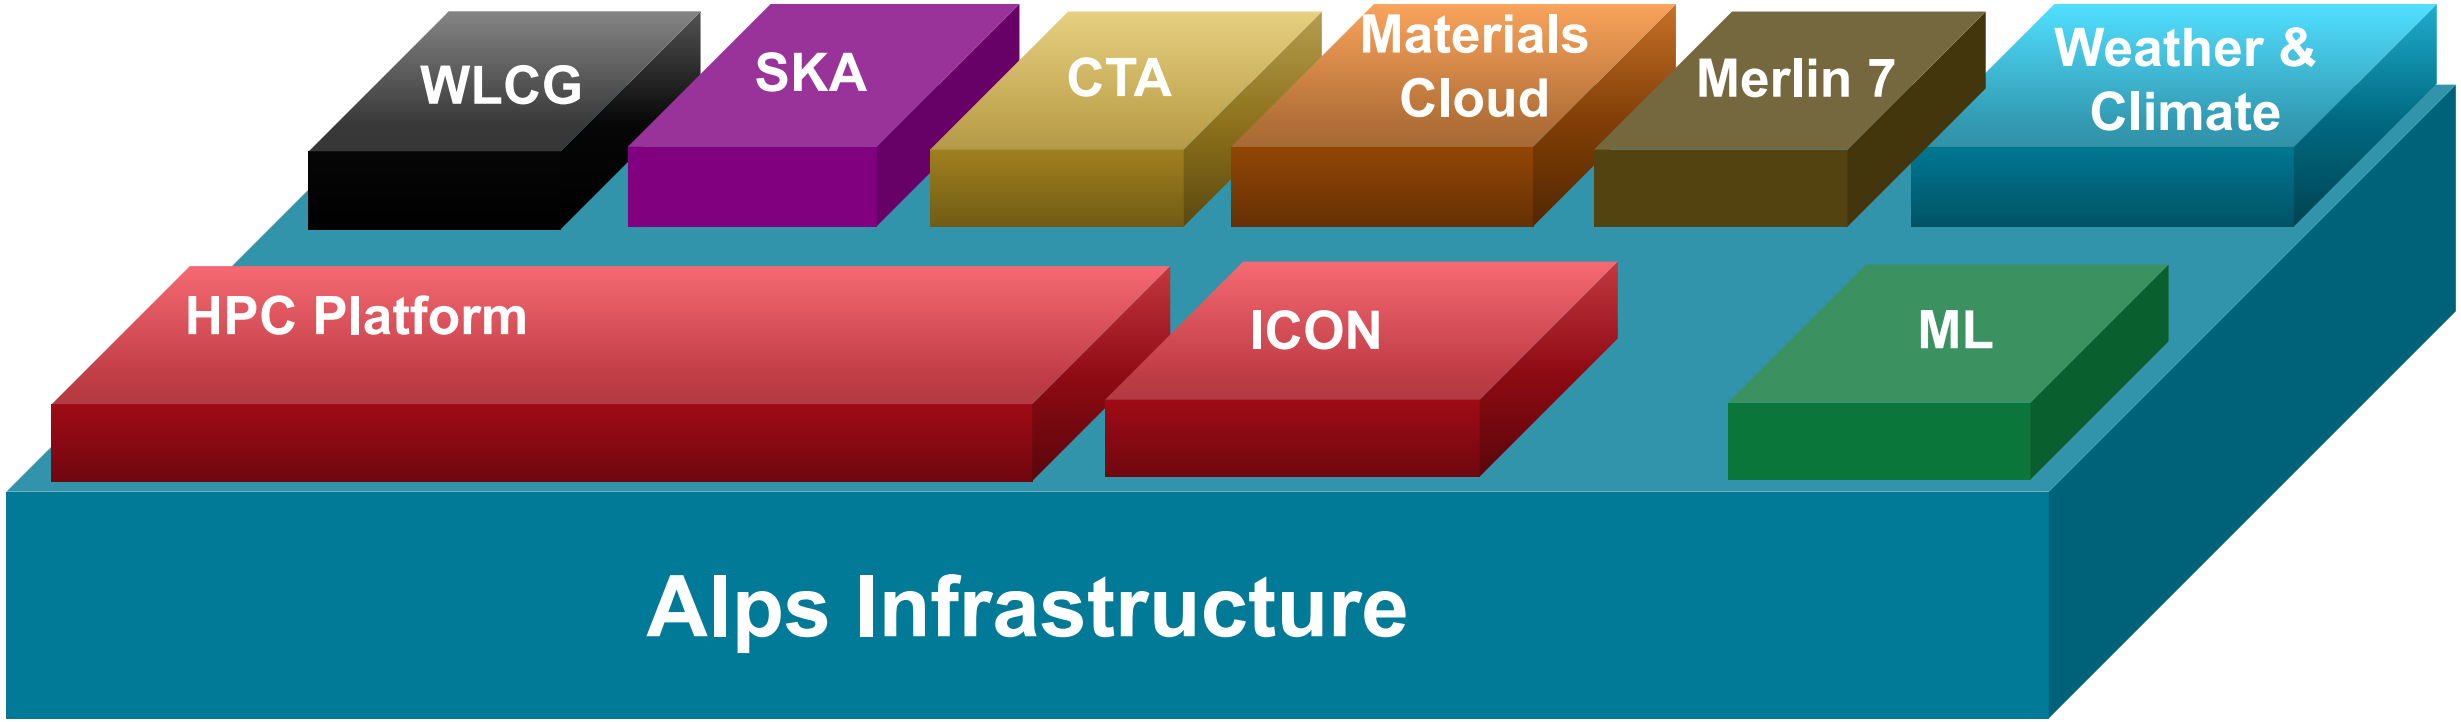

Why do we want to move vClusters?

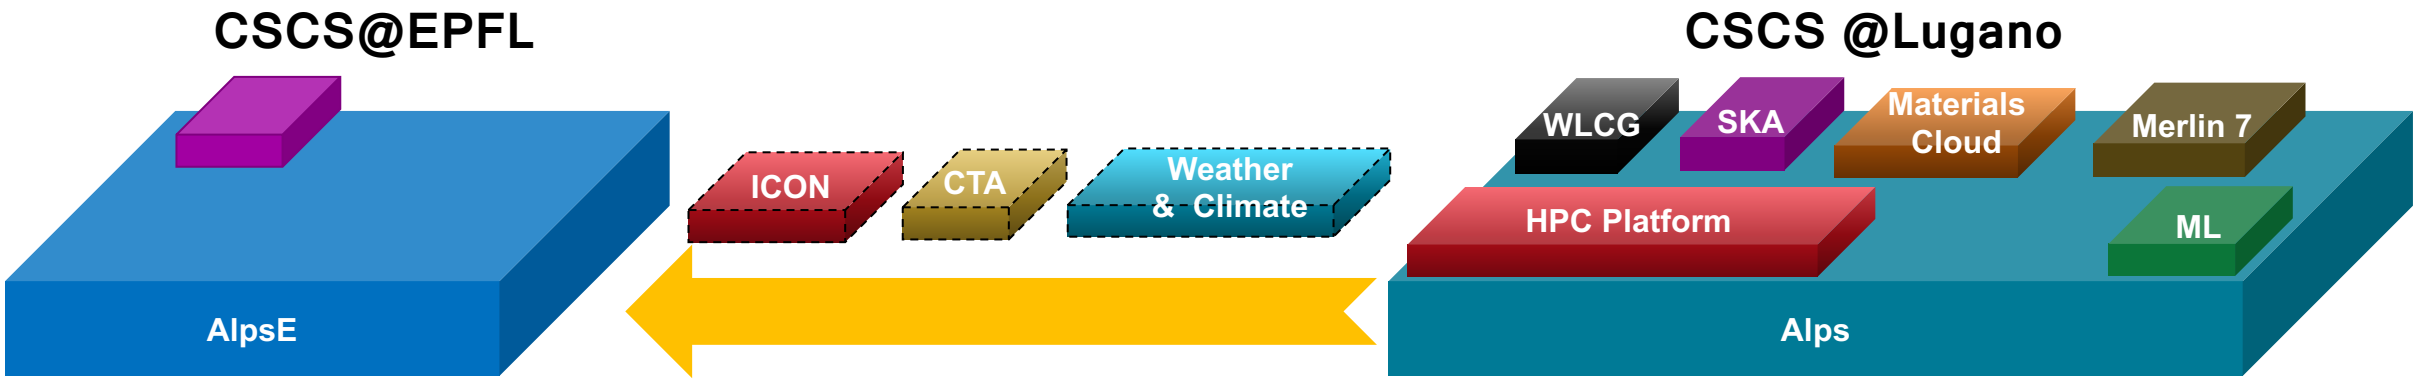

- Operational reasons:
  - Downtimes, upgrades, etc.
- Availability:
  - Some services might need to be available 24/7
  - Hardware constraints with heterogeneous cluster needs

How does this work

Multi-faceted process

- Migrating a cluster, in any cloud-like environment, is a coordinated operation that covers multiple aspects
  - CSM entities (CFS configurations, CFS component boot parameters, etc.)
  - Boot images (IMS images)
  - RPMs (Nexus repository)
  - Network (Capabilities, DNS, virtual IPs, ACLs, etc.)
  - Storage (Filesystems, data replication/sync, etc.)
  - API compatibility (dedicated CSM instances for specific tenants may have a different upgrade cadence)
  - User workflows that can/will be affected by the migration
  - Node selection (different sites have different node configuration)

- Where each one of them need to meet a set of pre-requisites

- And the tool to do the migration, should actually help you doing it

Brief introduction to Mesa/Manta

- Mesa
  - library to interact with CSM APIs built in Rust
  - Extend CSM capabilities
- Manta: CLI application built with Mesa for CSM operators

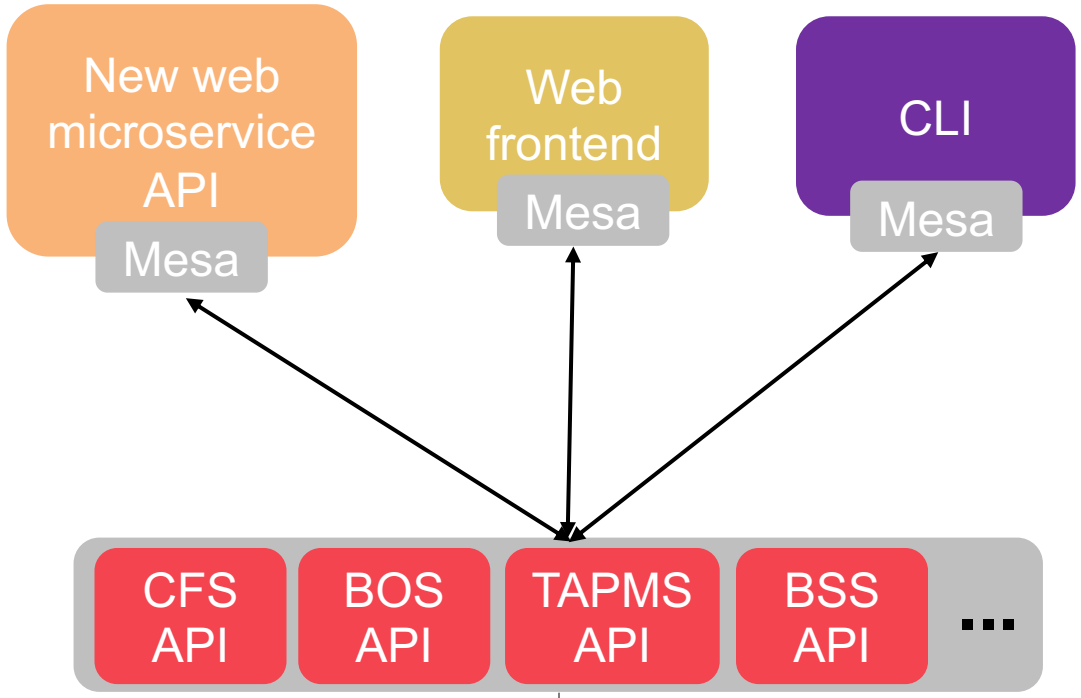

Pre-requisites for CSM

- Configuration management systems need to be consistent
  - If using VCS/Gitea: same history and replicated across sites
  - If using external Git: single location accessible by all sites
- The configuration needs to be abstract enough to work with the differences of each Cray EX
  - xnames, IPs, routes, etc. may be different and need to be configured properly
  - The *initrd* and the images need to know about the CAs
- Network ACLs need to be consistent across sites
- Filesystems and mountpoints need to be accessible across sites

How manta migrate works

- On its simplest form, in CSM a CSCS vCluster consists of
  - HSM group with xnames → Hardware config
  - CFS configuration
  - IMS image

  } Image/OS config
  - BOS session template → Booting config

- Manta does a backup of those 4 entities and dumps them into a bundle
- The bundle can then be restored (or uploaded) into another CSM system that meets the requirements mentioned earlier
- Manta migrate has hooks to run actions before backups and restore operations
  - These hooks are scripts that can apply actions on the non-CSM components of the infrastructure (network, etc.)

Operation

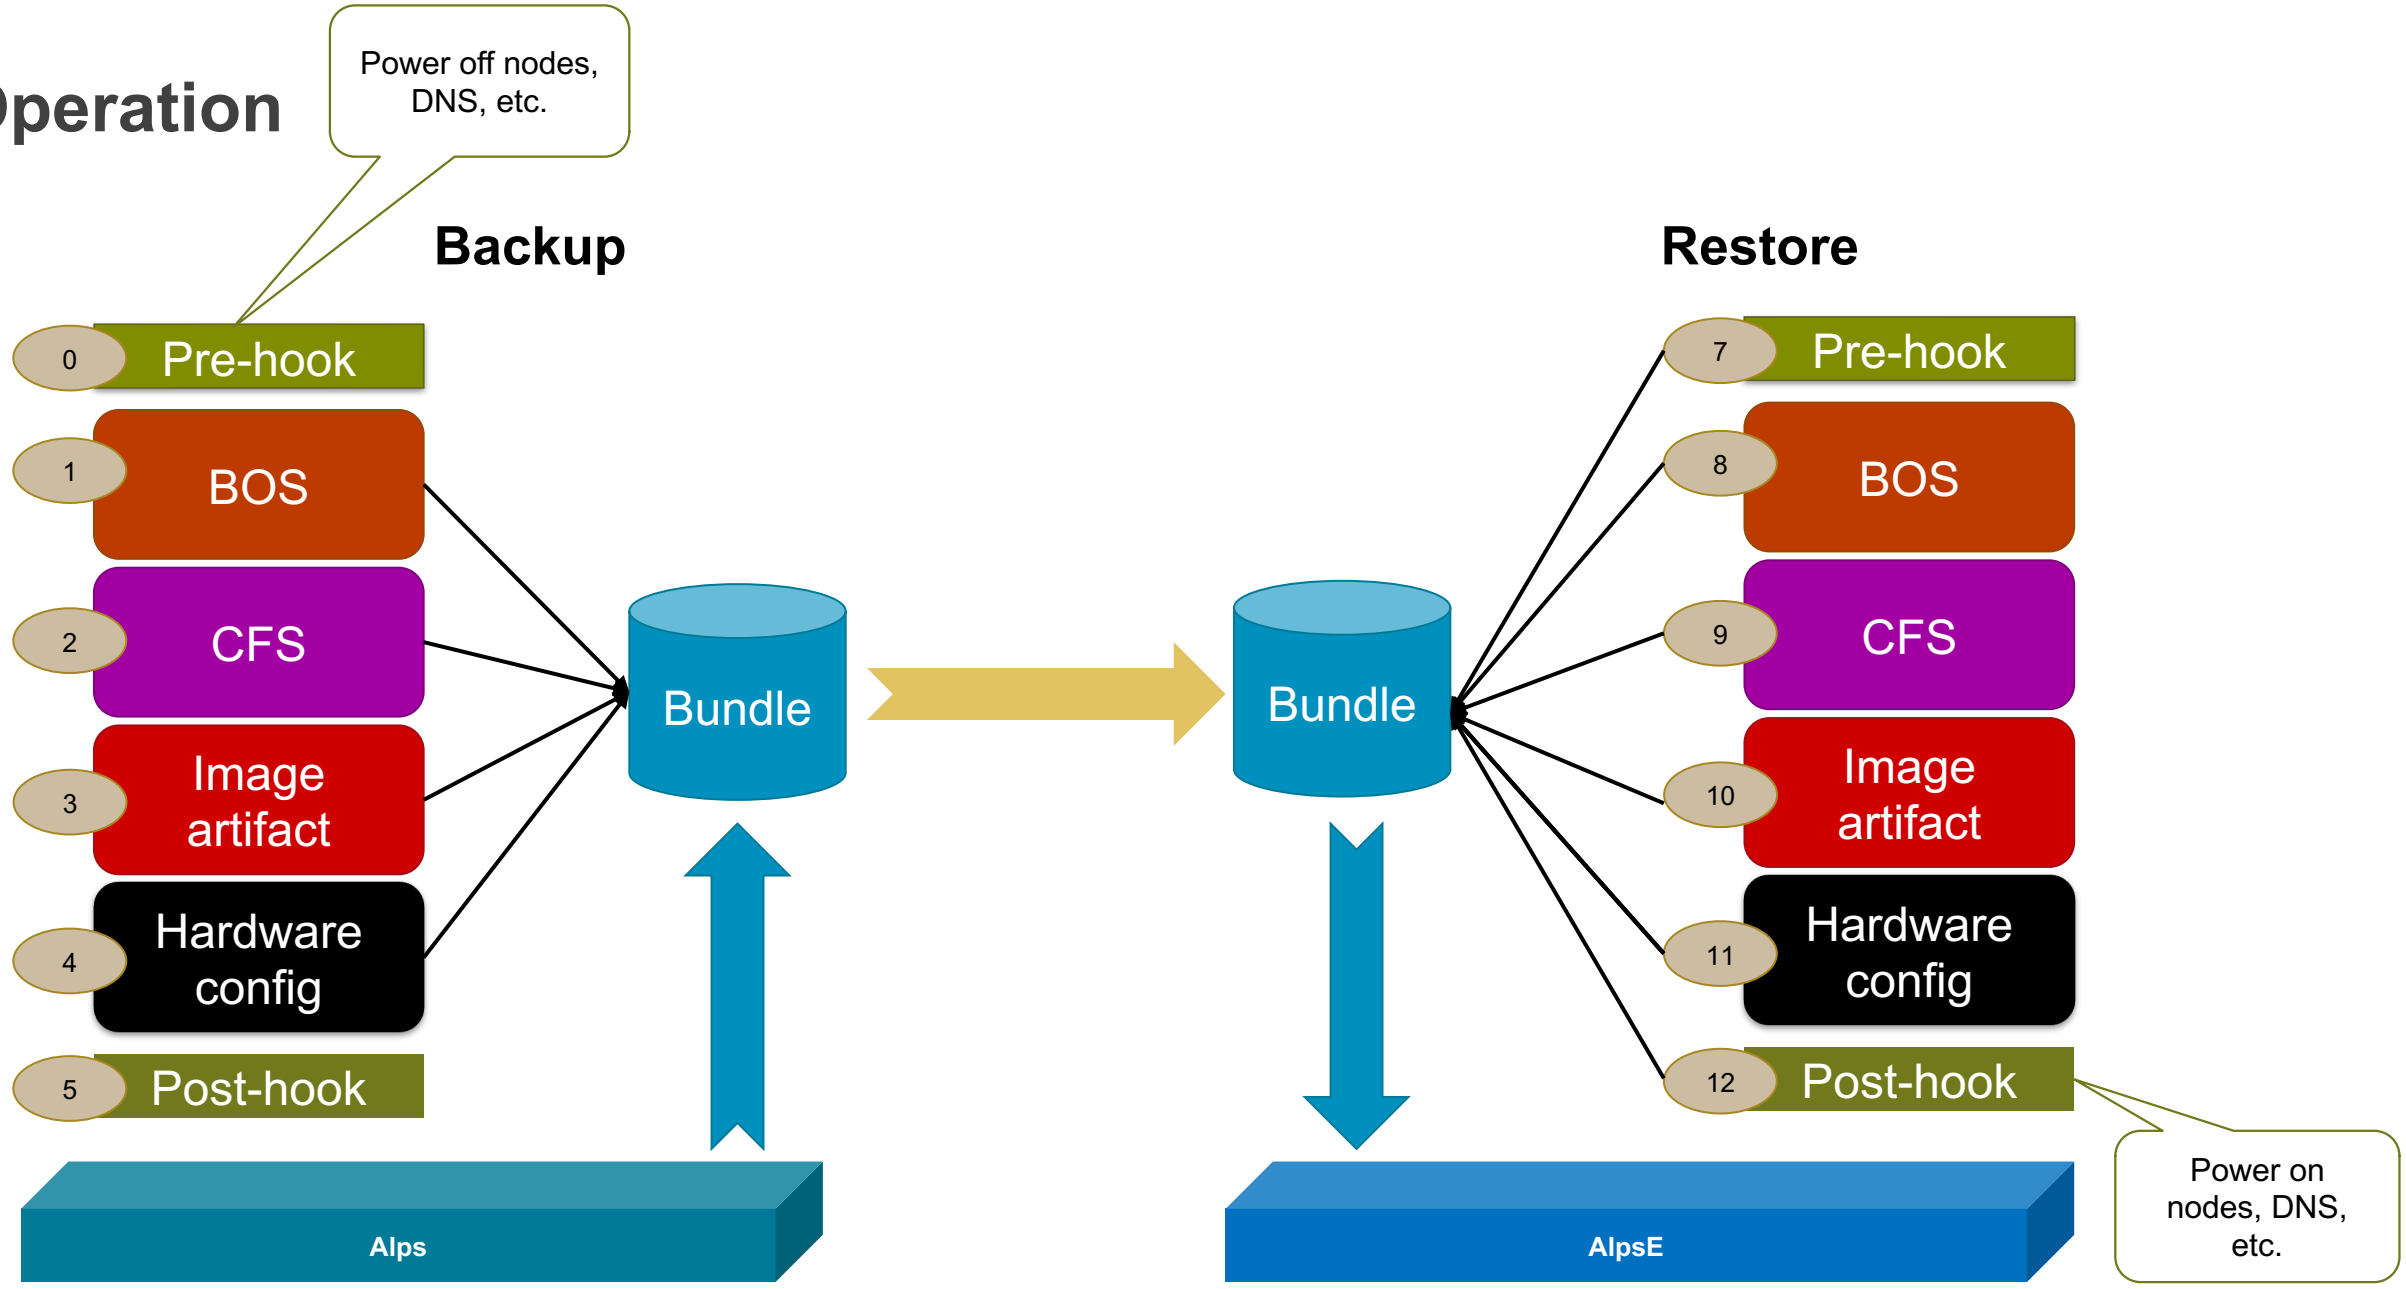
Power off nodes, DNS, etc.
Backup
Restore
0
Pre-hook
1
BOS
2
CFS
3
Image artifact
4
Hardware config
5
Post-hook
Bundle
Bundle
7
Pre-hook
8
BOS
9
CFS
10
Image artifact
11
Hardware config
12
Post-hook
Alps
AlpsE
Power on nodes, DNS, etc.

The network component

- DNS or virtual IPs
  - DNS resolves get updated with every vCluster migration, OR
  - A virtual IP is moved across sites and added to the nodes
- Filesystems that hold data and software stacks must be accessible on all locations, and local scratch can be mounted in the same path
- Data migration/sync is not considered, yet

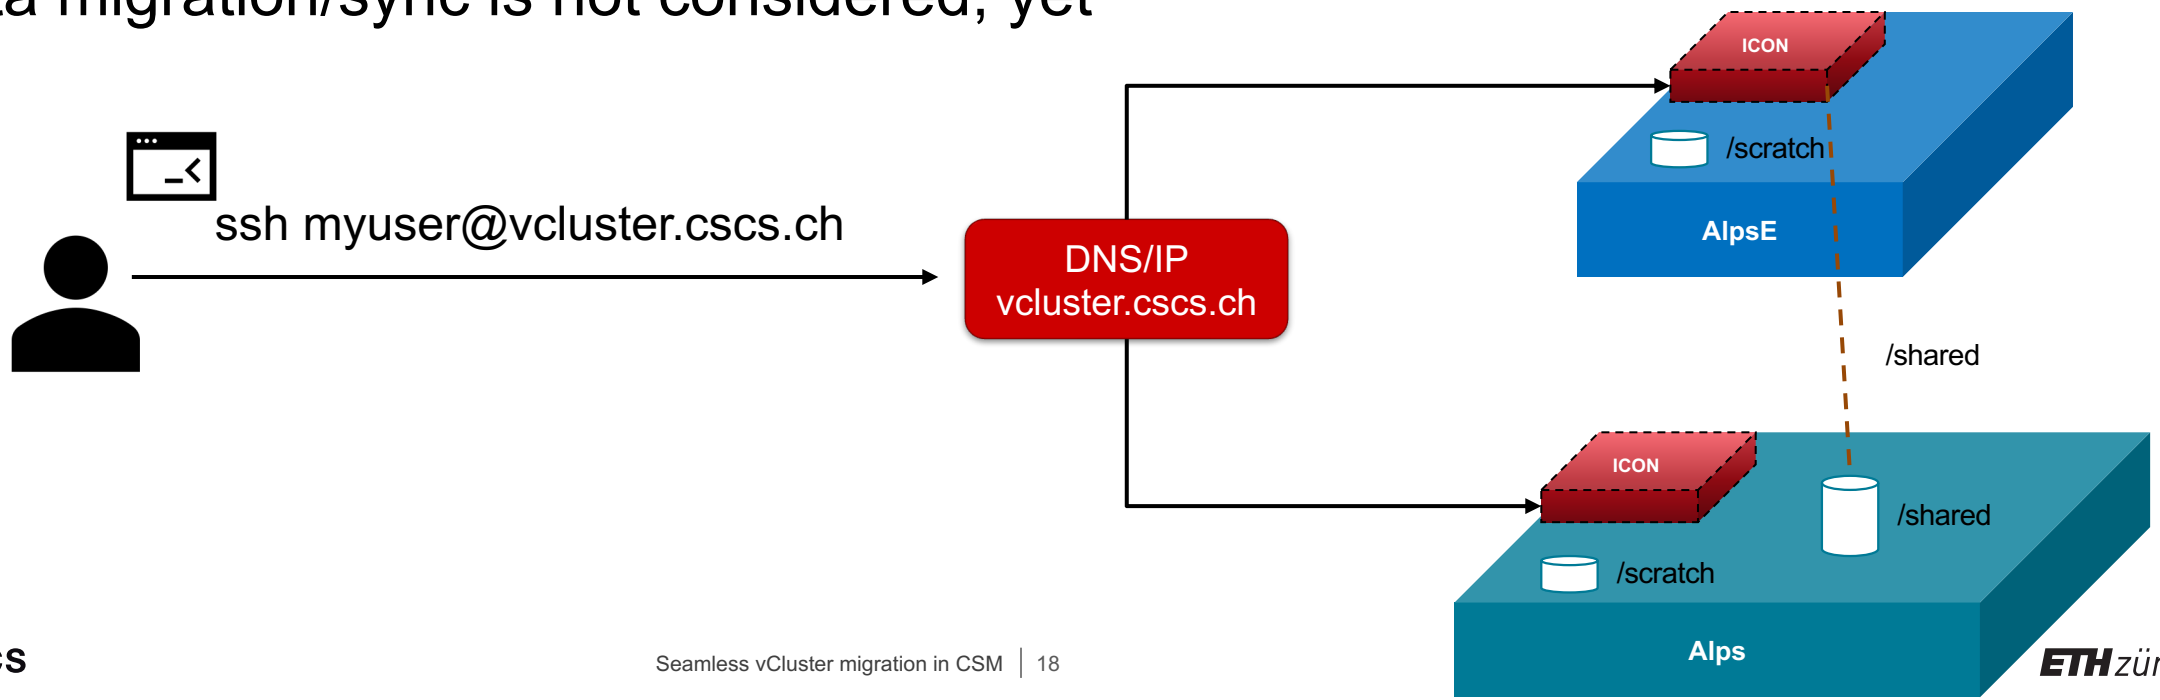

Results and next steps

Moving vClusters: backup

```
$ manta config set site alps
$ manta migrate backup -b gele-cos-template-23.06.0-1 -d $PWD

Migrate backup of the BOS Template: gele-cos-template-23.06.0-1; destination folder: /Users/user/tmp/manta-migrate4

[…] # Removed for cleanliness, complete output at the end of the slide deck

Done, the following image bundle was generated:
    BOS file: /Users/user/tmp/manta-migrate4/gele-cos-template-23.06.0-1.json
    CFS file: /Users/user/tmp/manta-migrate4/gele-cos-config-23.06.0-1.json
    HSM file: /Users/user/tmp/manta-migrate4/gele-cos-template-23.06.0-1-hsm.json
    IMS file: /Users/user/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json
    Image name: gele-cos-23.06.0-1
        file: /Users/user/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/manifest.json
        file: /Users/user/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd
        file: /Users/user/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel
        file: /Users/user/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs
```

Moving vClusters: restore

```
$ manta config set site prealps

$ manta migrate restore --bos-file ./gele-cos-template-23.06.0-1.json --cfs-file ./gele-cos-config-23.06.0-1.json --hsm-file ./gele-cos-template-23.06.0-1-hsm.json --ims-file ./3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json --image-dir ./3e0d418f-7757-4df0-b0f7-0b945495f7bf

Migrate restore of the following image:
    BOS file: ./gele-cos-template-23.06.0-1.json
    CFS file: ./gele-cos-config-23.06.0-1.json
    IMS file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json
    HSM file: ./gele-cos-template-23.06.0-1-hsm.json
    Image name: gele-cos-23.06.0-1
        initrd file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd
        kernel file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel
        rootfs file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs
[…] # Removed for cleanliness, complete output at the end of the slide deck

Uploading BOS sessiontemplate...
There already exists a BOS sessiontemplate with name gele-cos-template-23.06.0-1. It can be replaced, but it's dangerous.
Do you want to overwrite it? yes
Ok, BOS session template gele-cos-template-23.06.0-1 created successfully.

Done, the image bundle, HSM group, CFS configuration and BOS sessiontemplate have been restored.
```

Next steps

- Integrate mesa/manta with the multi-tenancy operator TAPMS
  - Currently manta only backs up HSM group definition, but doesn't consider multi-tenancy
- Add to mesa/manta migration code the ability to use hardware patterns instead of fixed xnames
- Create an API broker make the process agnostic to the infrastructure
- With CSM 1.6 all CFS repos can be external
  - This would make the migration even easier, closer to what a Cloud environment would be
- Port this functionality to be compatible with cloud-commercial environments like Google Cloud or Azure

Moving vClusters: backup, complete output

```
❯ manta config set site alps
❯ manta migrate backup -b gele-cos-template-23.06.0-1 -d $PWD
Migrate backup of the BOS Template: gele-cos-template-23.06.0-1; destination folder: /Users/miguelgi/tmp/manta-migrate4


Downloading BOS session template gele-cos-template-23.06.0-1 to /Users/miguelgi/tmp/manta-migrate4/gele-cos-template-23.06.0-1.json [1/8]
Downloading HSM configuration in bos template gele-cos-template-23.06.0-1 to /Users/miguelgi/tmp/manta-migrate4/gele-cos-template-23.06.0-1-hsm.json [2/8]
Downloading CFS configuration gele-cos-config-23.06.0-1 to /Users/miguelgi/tmp/manta-migrate4/gele-cos-config-23.06.0-1.json [3/8]
Downloading IMS image record 3e0d418f-7757-4df0-b0f7-0b945495f7bf to /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json [4/8]
Downloading image file 3e0d418f-7757-4df0-b0f7-0b945495f7bf/manifest.json (1.51 kB) to /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf/manifest.json [4/8]
[00:00:00] ████████████████████████████████████ (7.61 KiB/s) 1.48 KiB/1.48 KiB  [ETA 0s]                                                                                                Downloading
image file 3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd (51.85 MB) to /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd [5/8]
[00:00:02] ████████████████████████████████████ (20.53 MiB/s) 49.45 MiB/49.45 MiB  [ETA 0s]                                                                                              Downloading
image file 3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel (8.28 MB) to /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel [6/8]
[00:00:00] ████████████████████████████████████ (29.66 MiB/s) 7.90 MiB/7.90 MiB  [ETA 0s]                                                                                               Downloading
image file 3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs (2.54 GB) to /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs [7/8]
[00:01:37] ████████████████████████████████████ (24.92 MiB/s) 2.36 GiB/2.36 GiB  [ETA 0s]
Done, the following image bundle was generated:
    BOS file: /Users/miguelgi/tmp/manta-migrate4/gele-cos-template-23.06.0-1.json
    CFS file: /Users/miguelgi/tmp/manta-migrate4/gele-cos-config-23.06.0-1.json
    HSM file: /Users/miguelgi/tmp/manta-migrate4/gele-cos-template-23.06.0-1-hsm.json
    IMS file: /Users/miguelgi/tmp/manta-migrate4/3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json
    Image name: gele-cos-23.06.0-1
        file: /Users/miguelgi/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/manifest.json
        file: /Users/miguelgi/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd
        file: /Users/miguelgi/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel
        file: /Users/miguelgi/tmp/manta-migrate4//3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs
```

Moving vClusters: restore, complete output

```
❯ manta config set site prealps
❯ manta migrate restore --bos-file ./gele-cos-template-23.06.0-1.json --cfs-file ./gele-cos-config-23.06.0-1.json --hsm-file ./gele-cos-template-23.06.0-1-hsm.json --ims-file ./3e0d418f-
7757-4df0-b0f7-0b945495f7bf-ims.json --image-dir ./3e0d418f-7757-4df0-b0f7-0b945495f7bf
Migrate restore of the following image:
    BOS file: ./gele-cos-template-23.06.0-1.json
    CFS file: ./gele-cos-config-23.06.0-1.json
    IMS file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf-ims.json
    HSM file: ./gele-cos-template-23.06.0-1-hsm.json
    Image name: gele-cos-23.06.0-1
        initrd file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd
        kernel file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel
        rootfs file: ./3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs

Calculating image artifact checksum...
File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd" (51.85 MB)...
[00:00:00] ████████████████████████████████████ (60.89 MiB/s) 49.45 MiB/49.45 MiB  [ETA
0s]                                                                                 File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel" (8.28 MB)...
[00:00:00] ████████████████████████████████████ (63.20 MiB/s) 7.90 MiB/7.90 MiB  [ETA
0s]                                                                                   File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs" (2.54 GB)...
[00:00:40] ████████████████████████████████████ (59.80 MiB/s) 2.36 GiB/2.36 GiB  [ETA 0s]

Registering image with IMS...
Ok, IMS image ID: a64fb212-0ba4-4251-a2cd-18049aa9563f

Uploading image artifacts to s3...
File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/initrd" (51.85 MB) to s3://boot-images/a64fb212-0ba4-4251-a2cd-18049aa9563f/initrd.
[00:00:03] ████████████████████████████████████ (16.04 MiB/s) 49.45 MiB/49.45 MiB  [ETA
0s]                                                                                 File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/kernel" (8.28 MB) to
s3://boot-images/a64fb212-0ba4-4251-a2cd-18049aa9563f/kernel.
[00:00:01] ████████████████████████████████████ (7.84 MiB/s) 7.90 MiB/7.90 MiB  [ETA
0s]                                                                                    File "./3e0d418f-7757-4df0-b0f7-0b945495f7bf/rootfs" (2.54 GB)
to s3://boot-images/a64fb212-0ba4-4251-a2cd-18049aa9563f/rootfs.
[00:02:04] ████████████████████████████████████ (19.37 MiB/s) 2.36 GiB/2.36 GiB  [ETA
0s]                                                                                   File "new-manifest.json" -> s3://boot-images/a64fb212-0ba4-4251-
a2cd-18049aa9563f/manifest.json.
OK

Updating IMS image record with the new location in s3...
Ok

Creating HSM group...
The HSM group gele has been created successfully.
Ok

Uploading CFS configuration...
There already exists a CFS configuration with name gele-cos-config-23.06.0-1. It can be replaced, but it's dangerous as it can trigger automated node reconfiguration.
Do you want to overwrite it? yes
Ok, CFS configuration gele-cos-config-23.06.0-1 created successfully.

Uploading BOS sessiontemplate...
There already exists a BOS sessiontemplate with name gele-cos-template-23.06.0-1. It can be replaced, but it's dangerous.
Do you want to overwrite it? yes
Ok, BOS session template gele-cos-template-23.06.0-1 created successfully.
Done, the image bundle, HSM group, CFS configuration and BOS sessiontemplate have been restored.
```

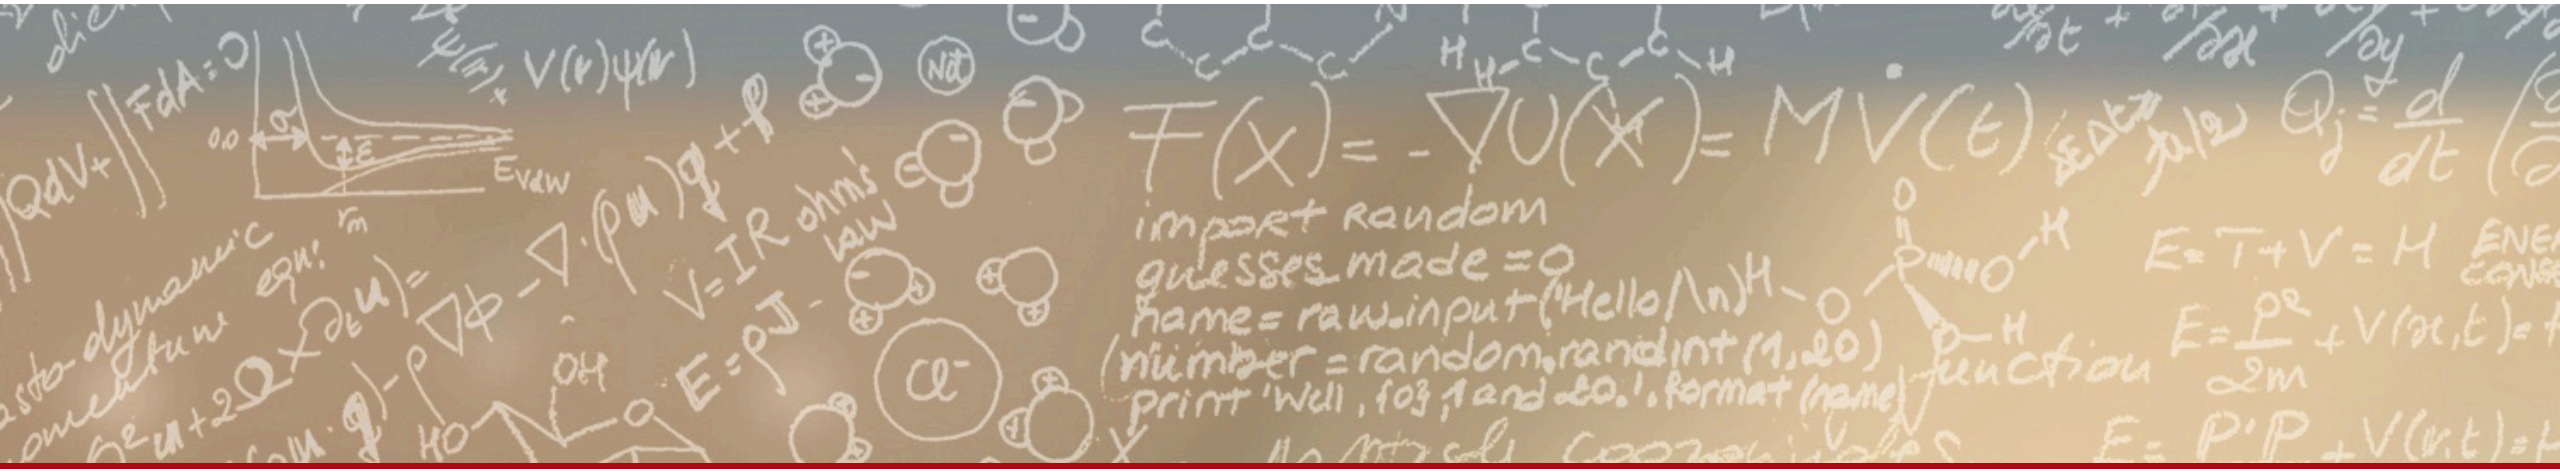

Thank you for your attention.