# CÉCI Technical talks

CÉCI Day 2026


# CÉCI Updates

## New Infrastructure & Evolution

frederic.wautelet@unamur.be

CÉCI Day 2026

# Context

- A infrastructure in transition
- A major evolution of the CÉCI ecosystem

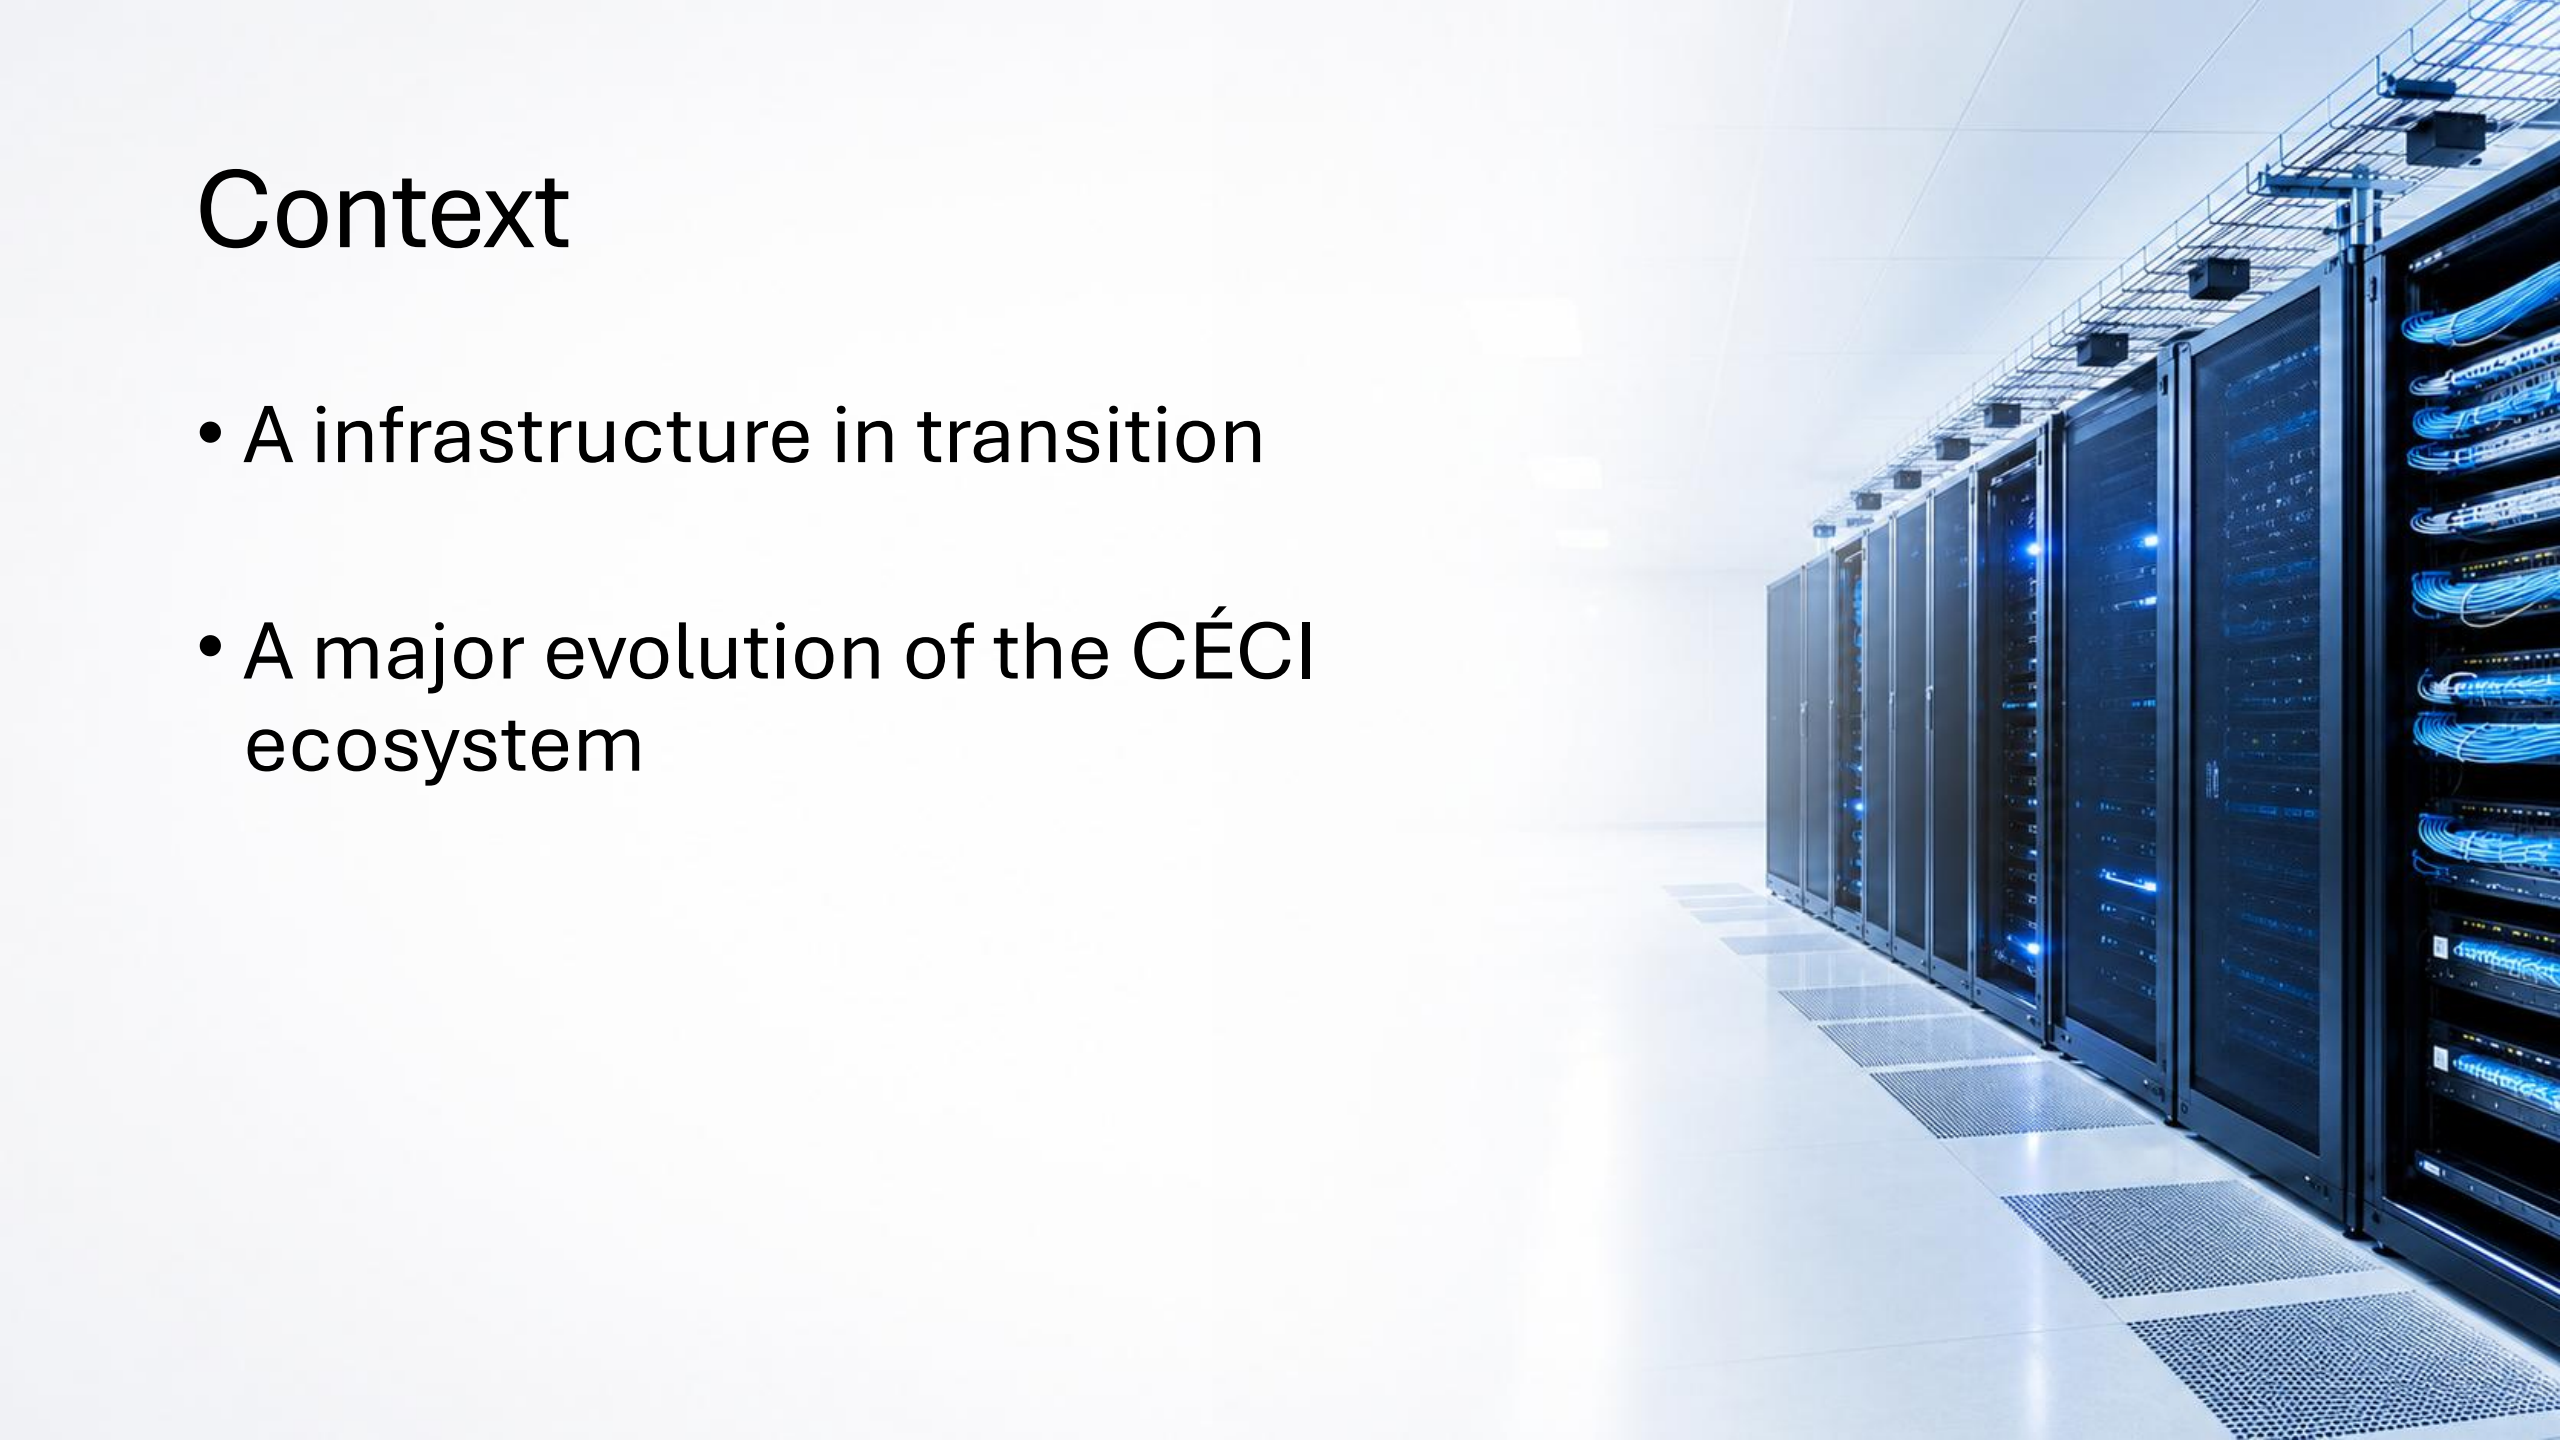

# Philosophy

Different clusters for different needs

# Types of workloads

**HTC / Serial**
independent jobs

**MPI**
parallel distributed jobs

**GPU**
offloaded on accelerators

# CÉCI ecosystem

| Usage | Cluster |
|---|---|
| GPU | Lyra |
| Large MPI | Lemaitre |
| MPI | NIC |
| High-memory | Hercules |
| General purpose | Dragon |

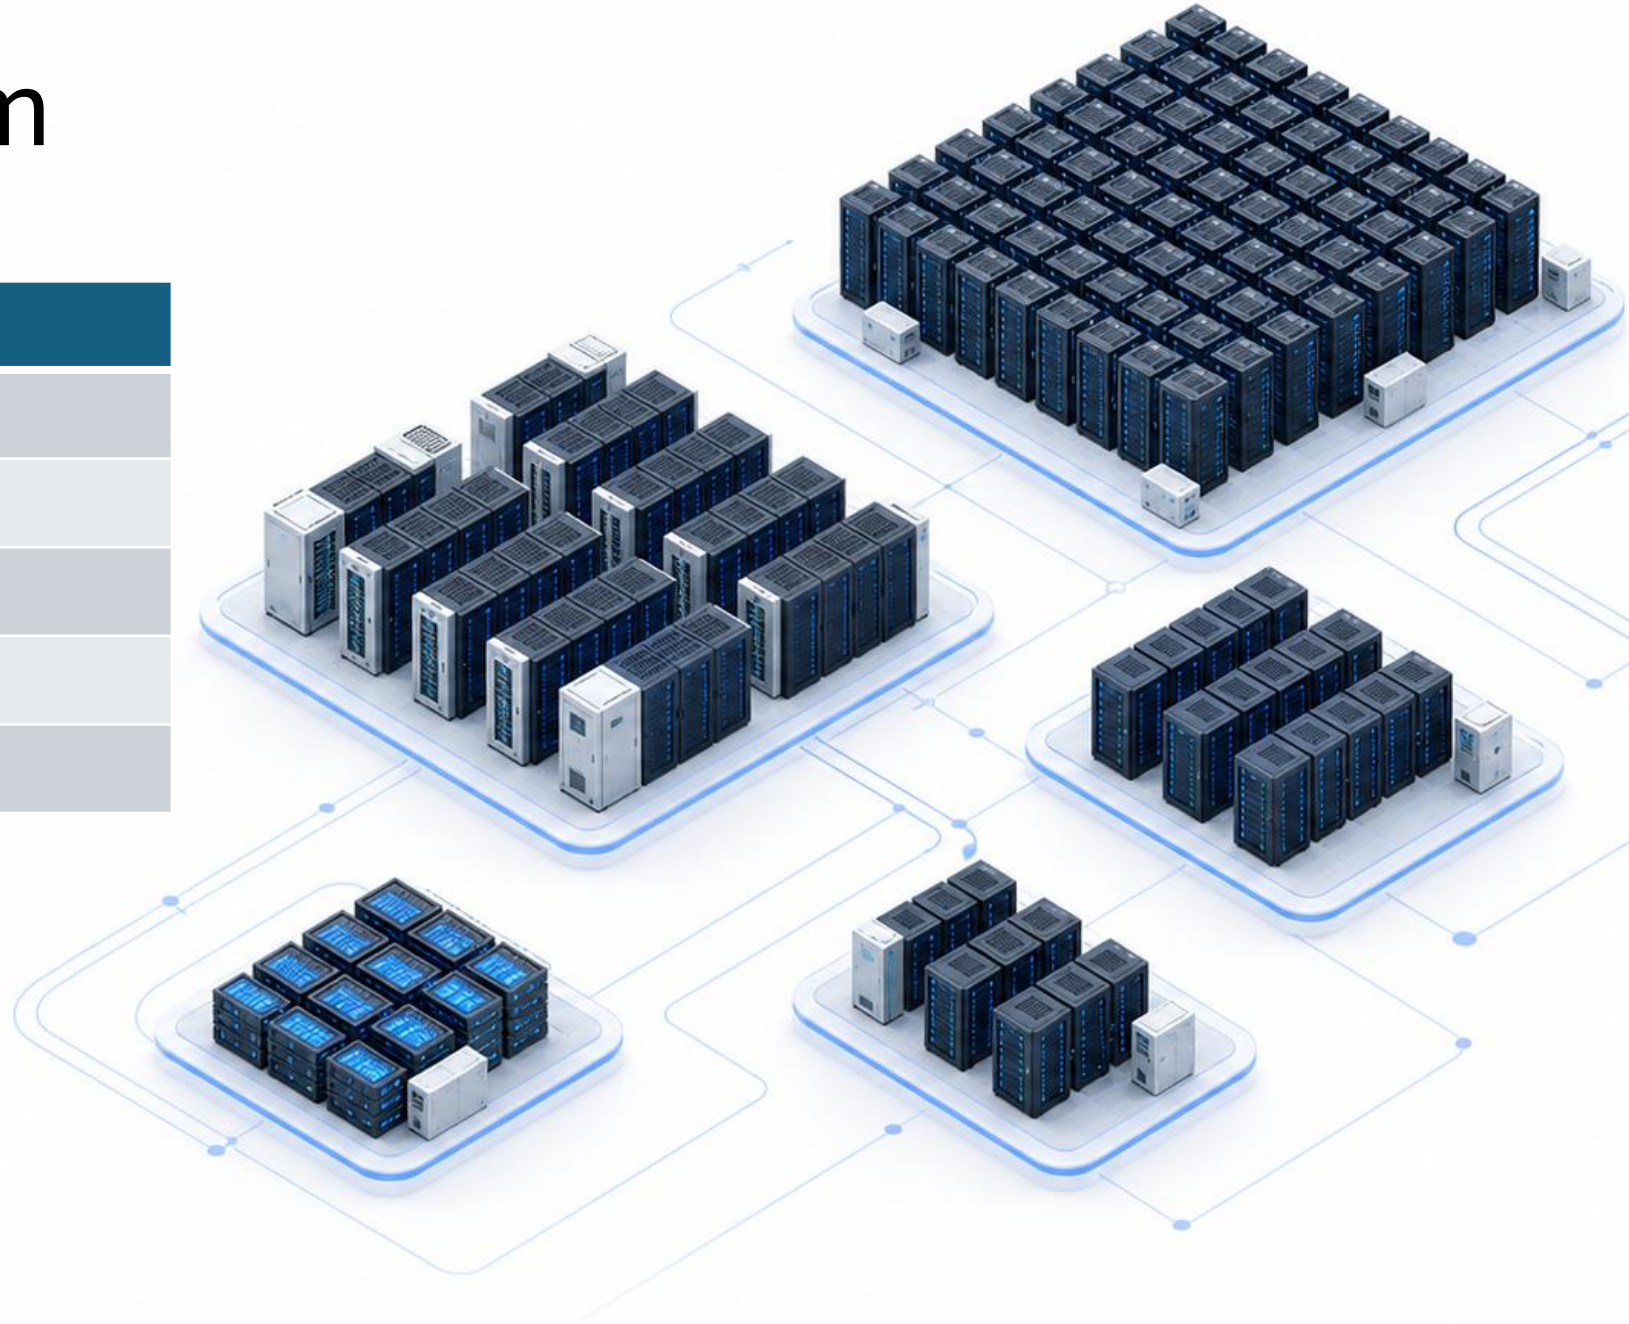

# Why evolve?

- Limitations of the current setup
- Increasing demand for GPU
- More data-intensive workloads
- Aging infrastructure

# Overview

- New CÉCI clusters
- New common storage
- New architectures
- New website

# New cluster generations

- Dragon 3 & Hercules 3
- Successors of existing clusters
- Improved performance
- Modern hardware

# Dragon 3 @ UMons

- Key characteristics
  - Based on latest AMD CPU
  - Nodes with high-performance local scratch
  - High-performance « home » storage
- Transition
  - Deployment **in progress**
  - Dragon 2 will remains available

# Dragon 3: hardware

- 14 compute nodes
  - 2 x AMD Zen5 EPYC 48-Core
  - 768 GB DDR5
  - 7 TB NVME local scratch

- **1344 core total**

- **120 TB storage $HOME**
  - High Performance NVMe SSD

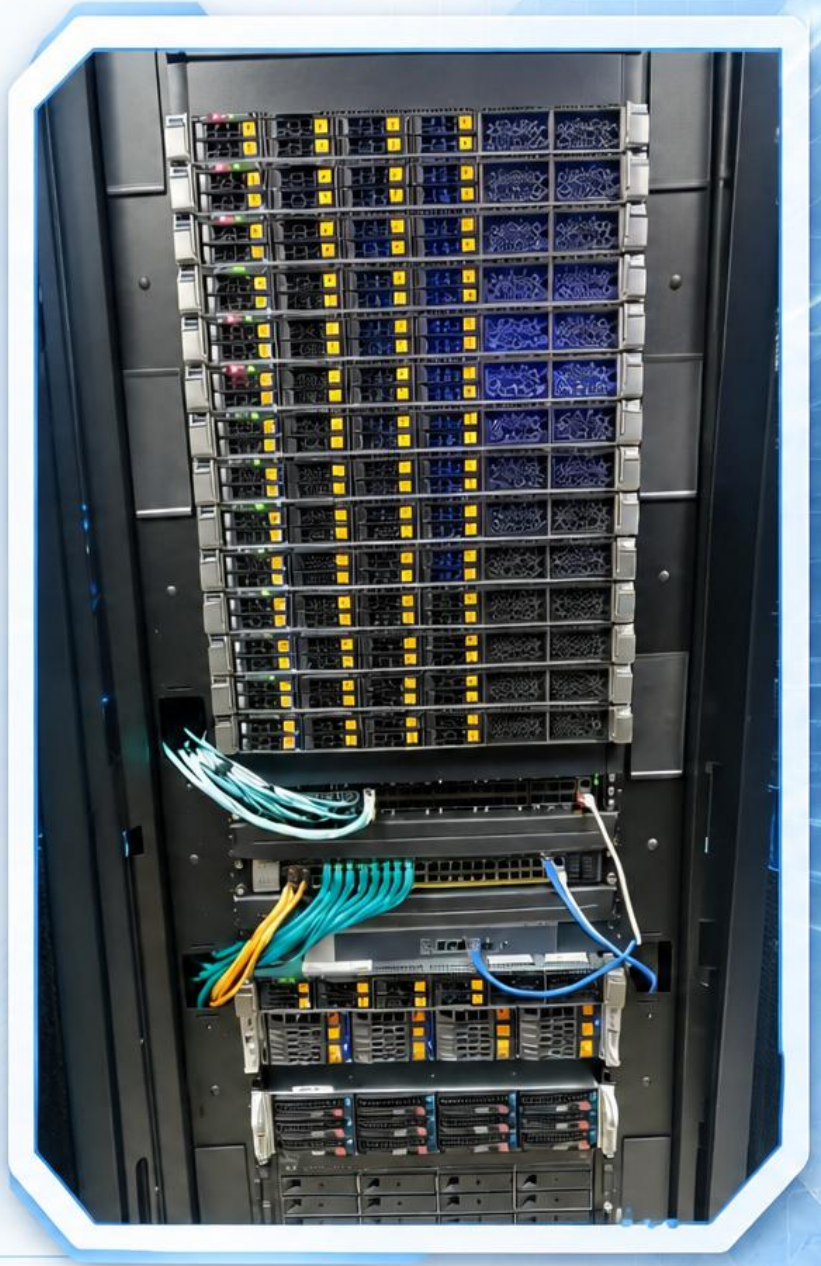

# Hercules 3: Smooth transition

- One hardware infrastructure
  - New nodes already available on Hercules 2
  - Hercules 3 deployed in parallel
- No disruption
  - Nodes gradually migrated from Hercules 2 → Hercules 3
  - Access to both Hercules 2 and 3 during transition
- Hercules 3 ready when fully validated

# Hercules 3: Infrastructure redesign

- System upgrade
  - Rocky Linux 9
  - Need to rebuild all scientific applications
- Focus on resilience & security
  - Reduced downtime during maintenance
  - Faster response to incidents & vulnerabilities

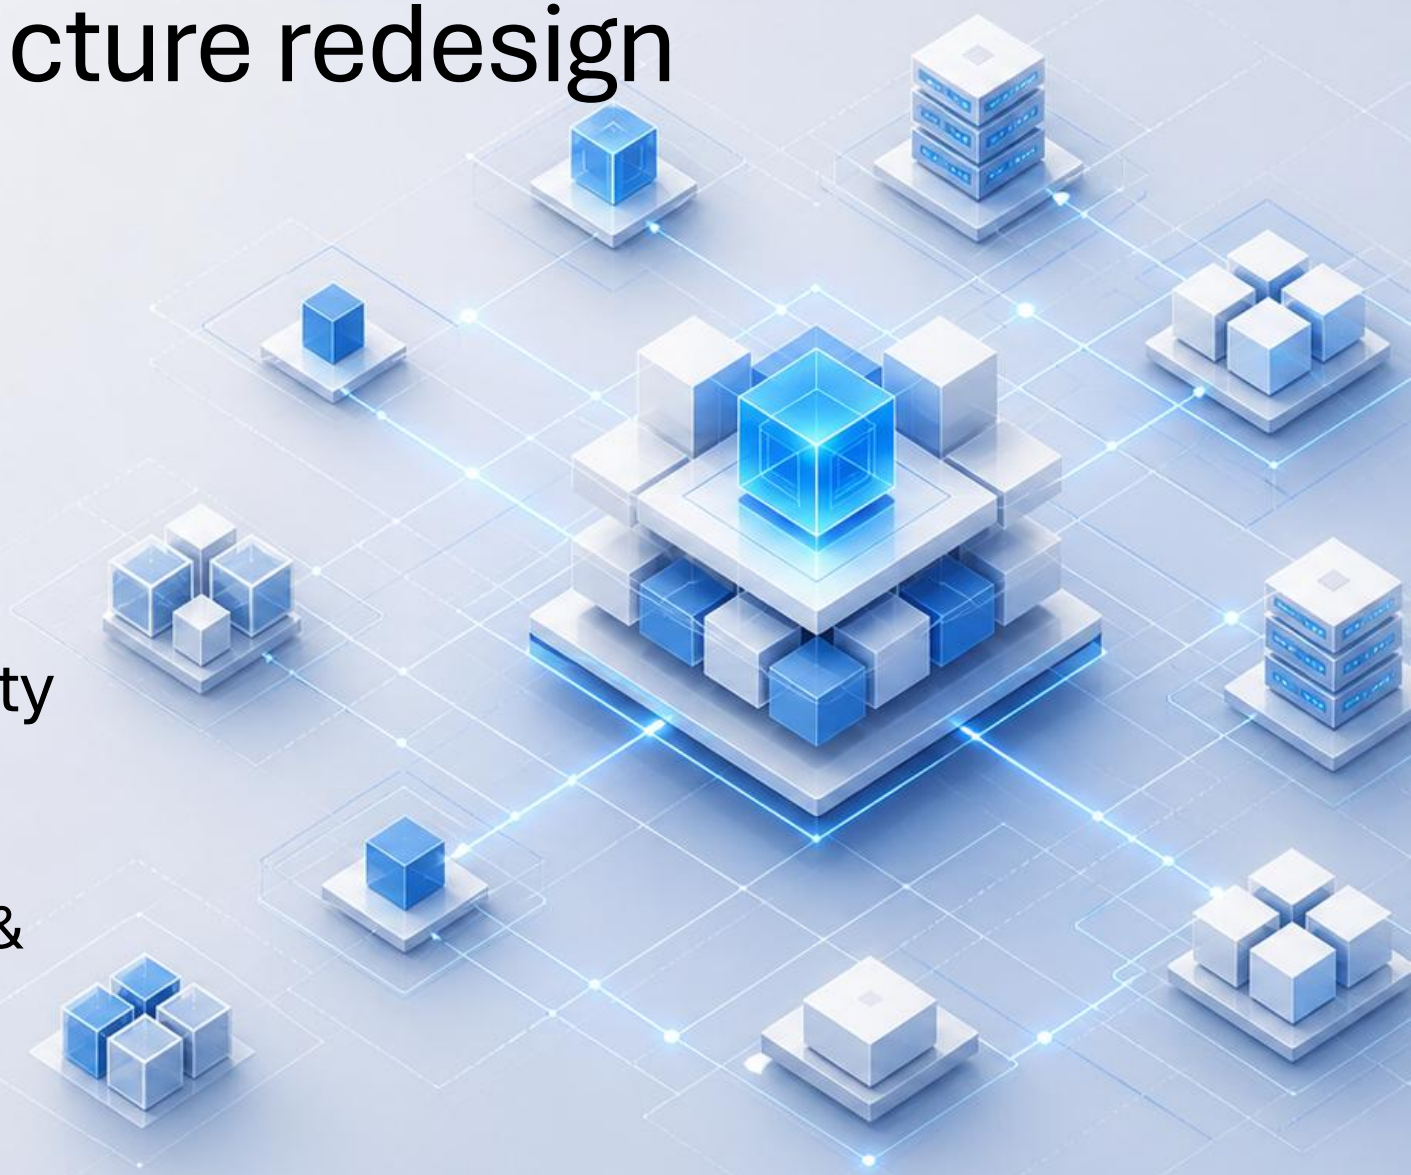

# Hercules 3: Storage

- $HOME
  - 5-nodes Ceph distributed storage
  - 200 TB net
  - **2 TB quota per user**

- $GLOBALSCRATCH ($WORKDIR)
  - **Not available** on Hercules 3
  - But still available on Hercules 2 during transition
  - **Use $LOCALSCRATCH instead**

# Hercules 3: hardware

- 35 nodes from Hercules 2
  - 1216 cores
  - 16 GPUs
- 36 new nodes
  - 26 x 48-core, 384 GB RAM
  - 4 x 48-core, 768 GB RAM
  - 1 x 96-core, **3 TB RAM**

- **71 nodes, 2992 cores total**
- **200 TB storage $HOME**
  - 5-node Ceph

# New clusters

**Dragon 3**

1344 cores

up to 768 GB

serial, shared memory parallel

**Hercules 3**

2992 cores

up to 3 TB

12 GPUs

serial, shared memory parallel, high-memory

# A new CÉCI Common Storage

- A new storage
  - **New distributed storage system available in April 2025**
  - Replaces the old **storage installed in 2015**

- Key idea
  - Based on an Open Source solution: Ceph
  - **Distributed across the 5 universities**

# Impact for users

- More capacity available: 1 PiB
  - Larger datasets can be stored and managed
- Same access from all clusters
  - Unified storage across the CÉCI infrastructure
  - Fastest way to transfer your data from one cluster to another
  - **Available on Lucia**

# CÉCI Home

- To store :
  - your software
  - configuration files
  - small (input) data files or (output) result files

- $CECIHOME
  - 100 GB quota
  - Max 100.000 files

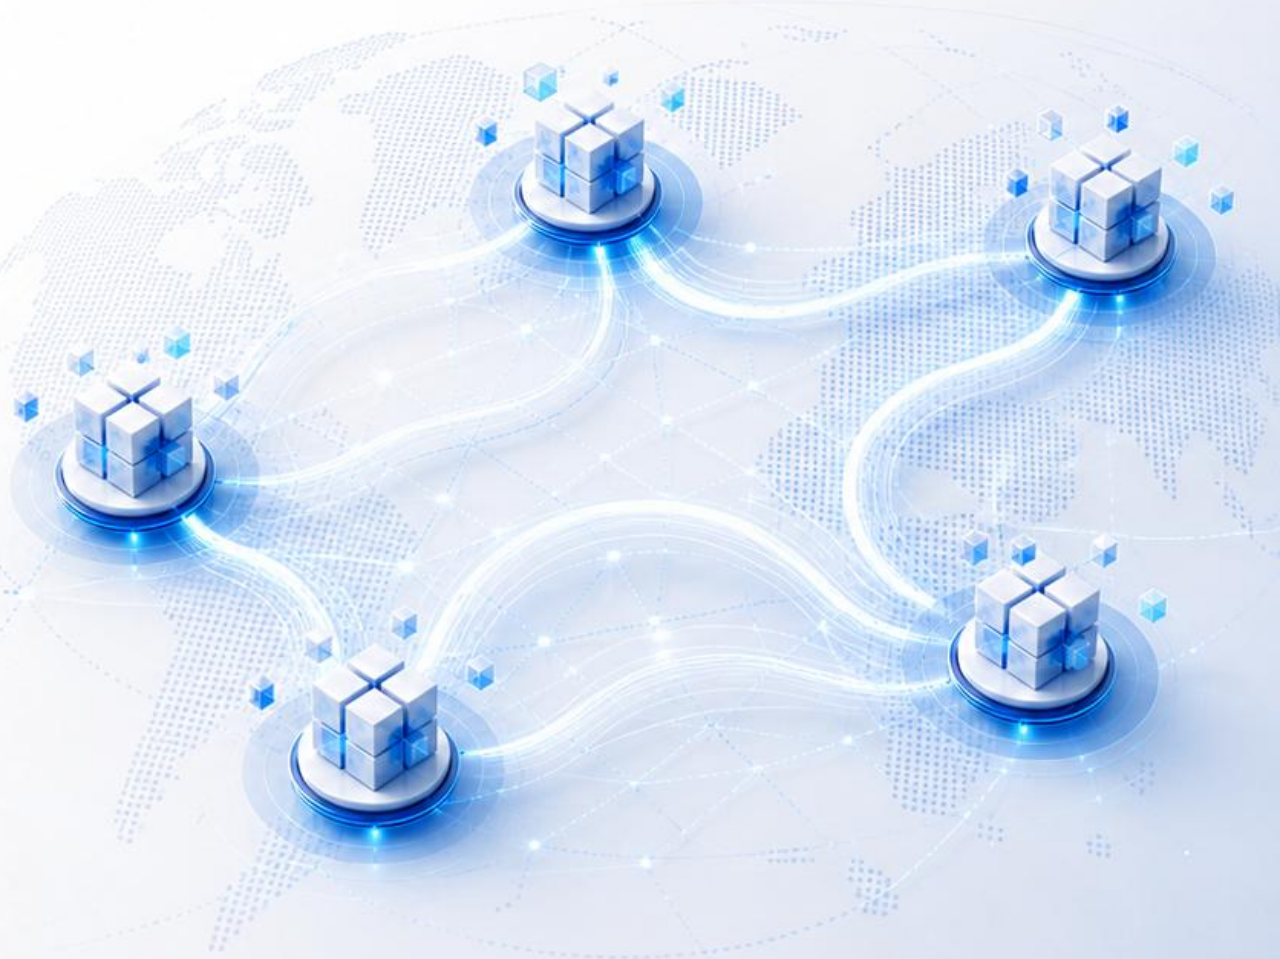

# Common Storage projects

- Allow better collaboration
  - Dedicated to projects
  - Easier sharing within projects and across universities

- /CECI/proj
  - Created upon request

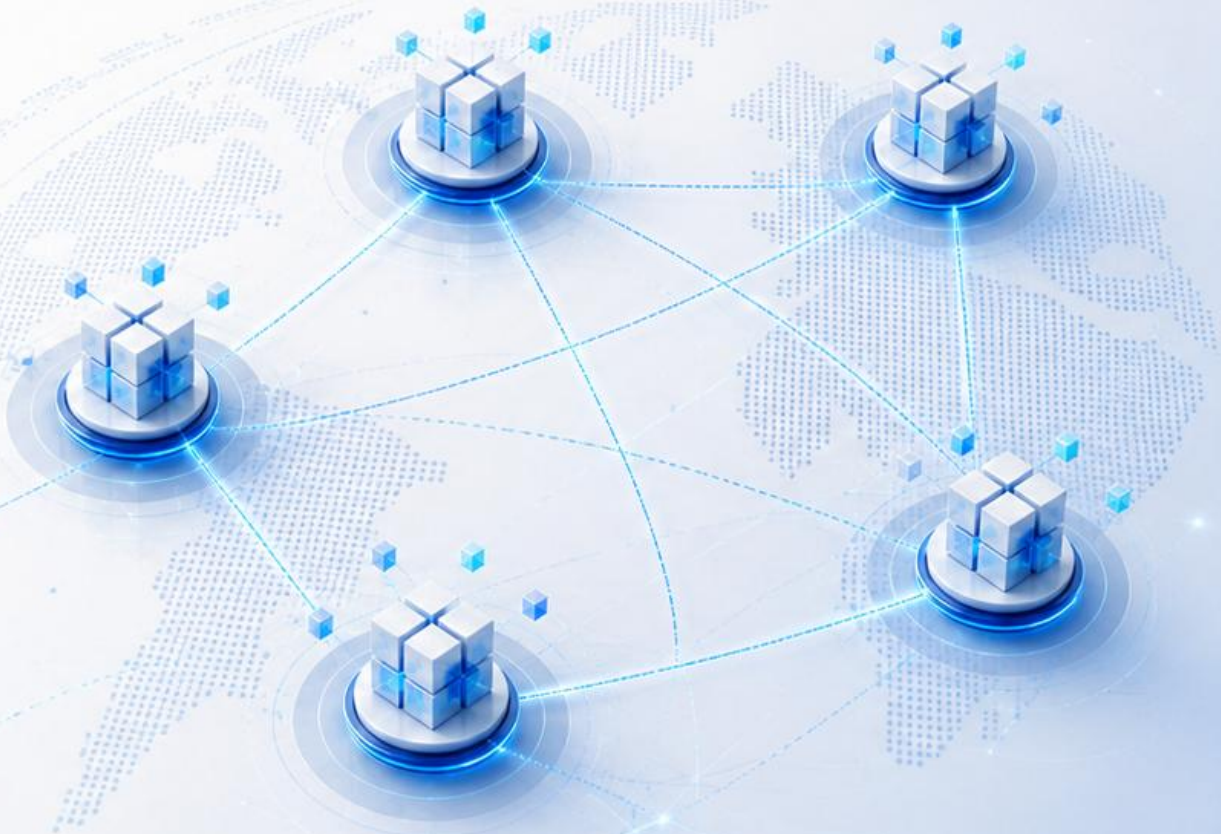

# A new architecture: APU

- Accelerated Processing Unit

- Breaking the CPU–GPU barrier
  - CPU + GPU integrated in a single chip
  - Unified CPU/GPU memory
  - Optimized for double precision (FP64)

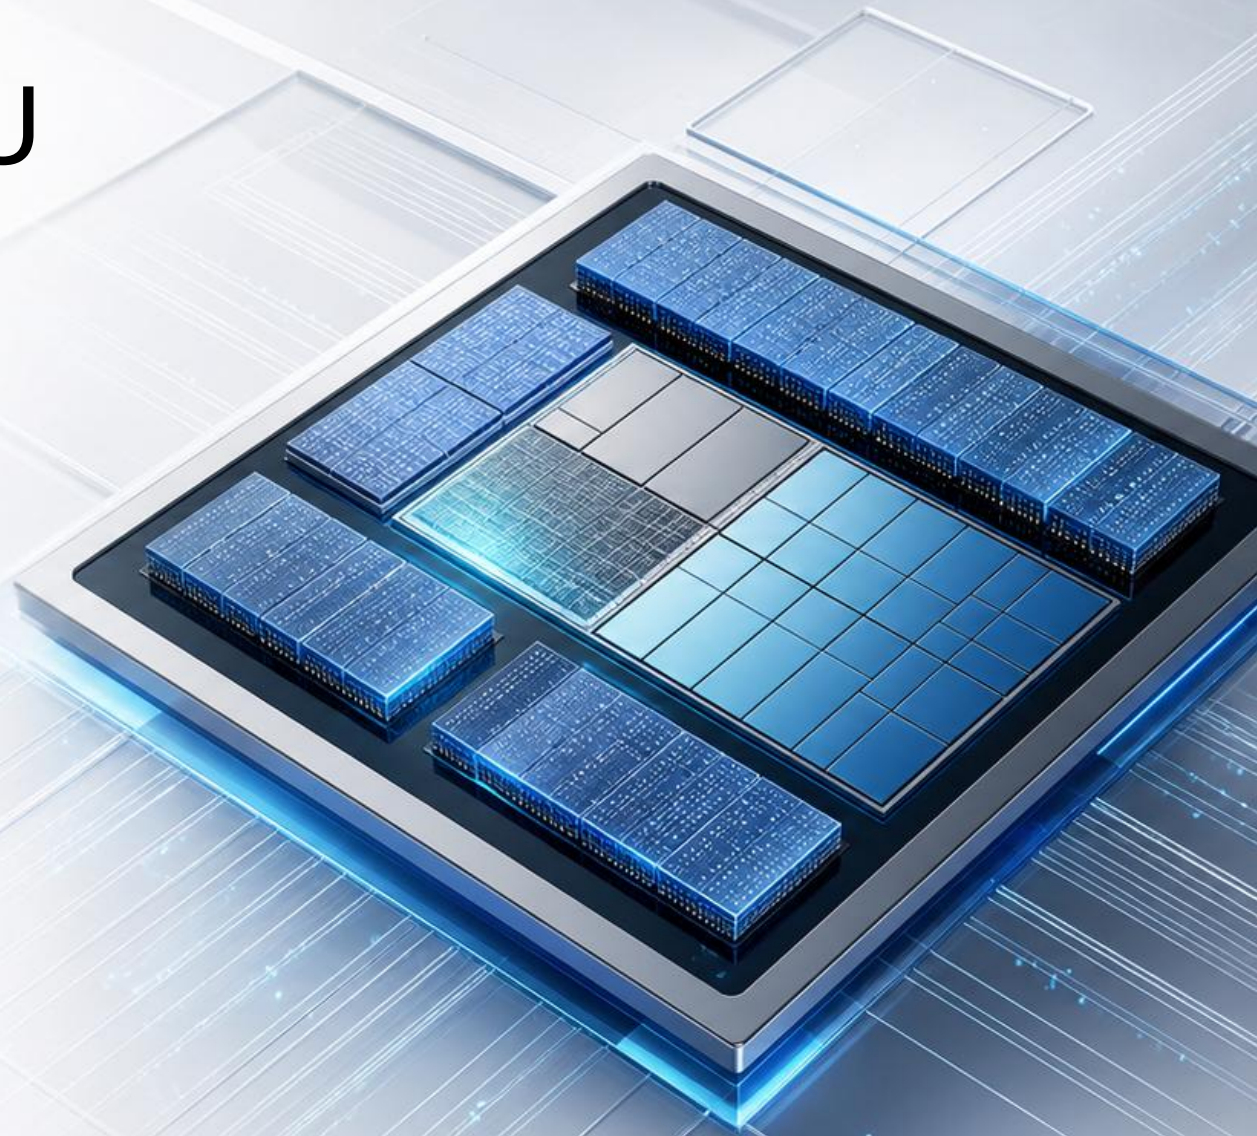

# Why it matters?

- Reduced data movement → faster simulations
- Simpler programming model
- Better energy efficiency
- Designed for next-generation (exascale) systems

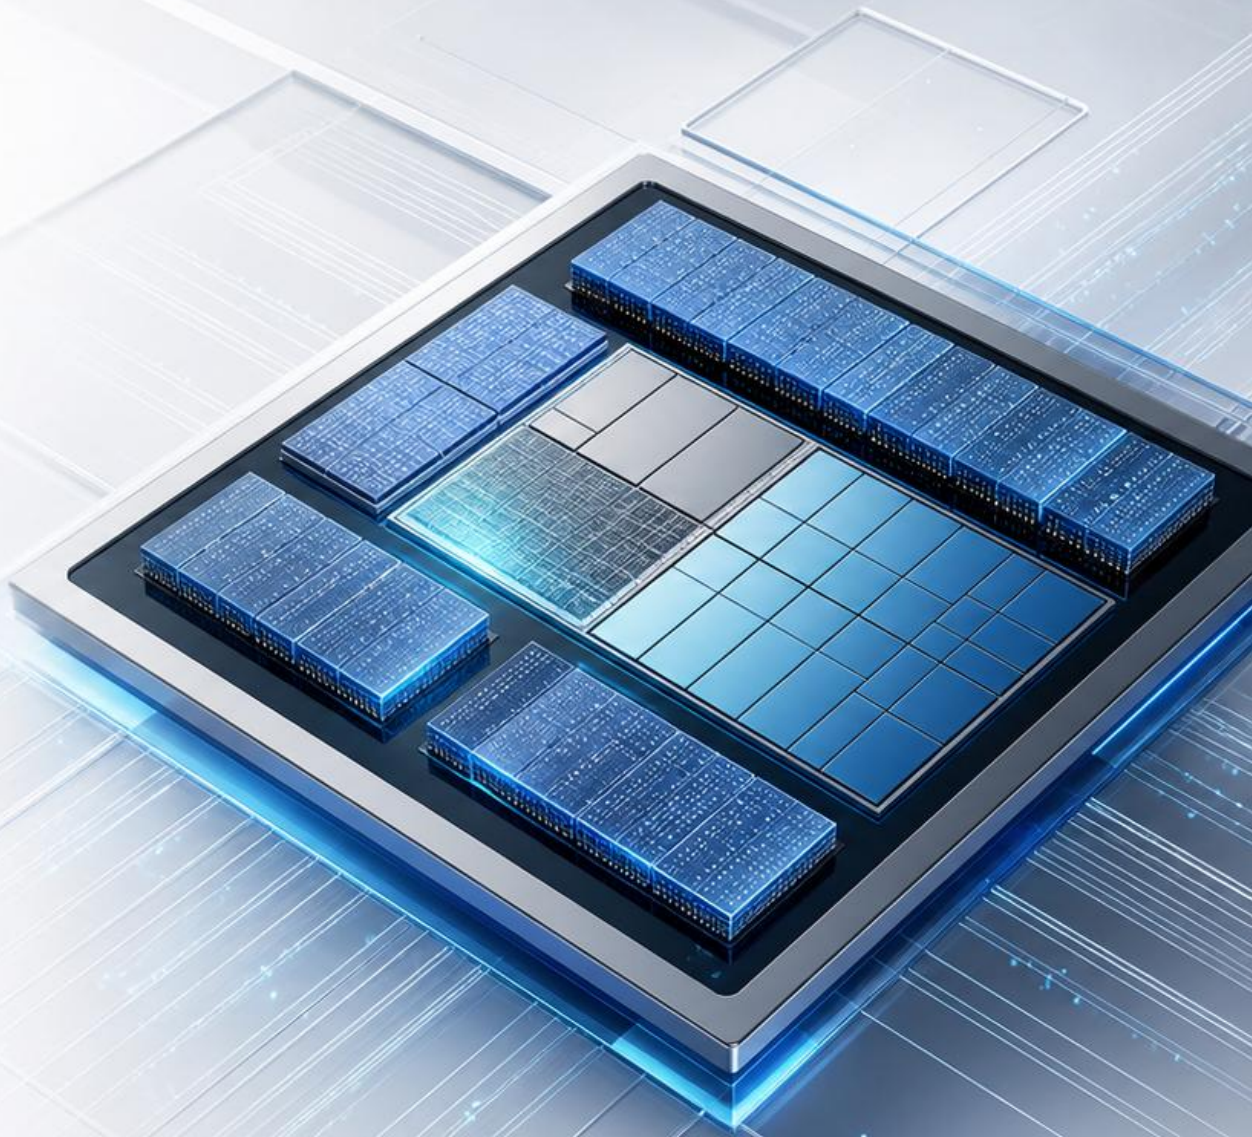

## TRADITIONAL CPU + GPU ARCHITECTURE

Separate compute units with dedicated memory

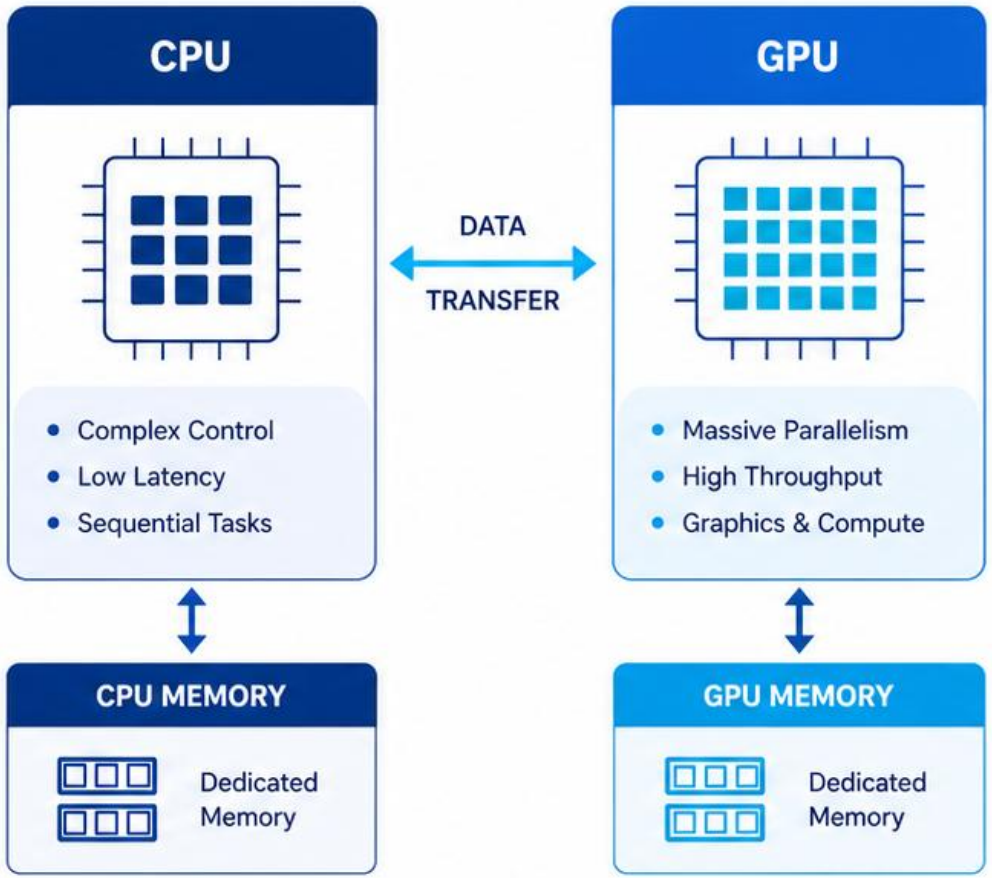


Data must be copied between memories

Higher latency and overhead

Increased power and complexity

VS.

## INTEGRATED APU ARCHITECTURE

Unified compute units with shared memory

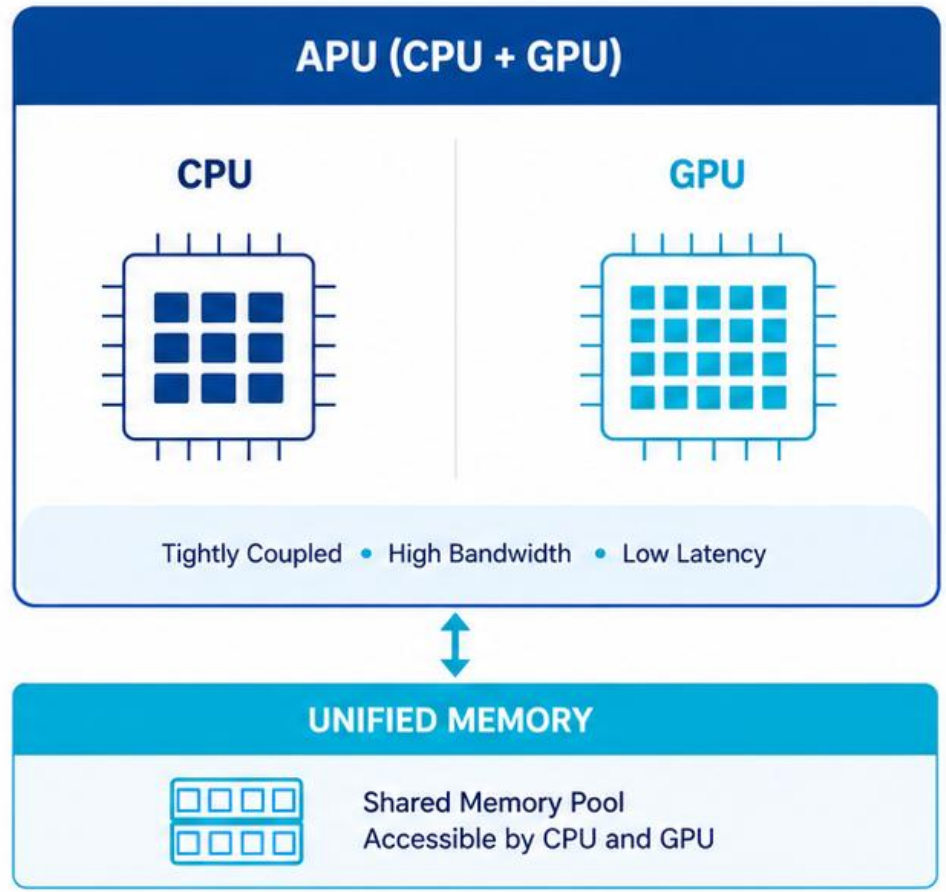


Faster data access and lower latency

Higher performance and efficiency

Lower power and cost

# The New APUs

- 5 nodes with **AMD MI300A APUs**
  - 4 APU units per node
  - **128 GB HBM3 unified memory per APU**
  - CPU: 24 cores per APU
  - **30 TB $LOCALSCRATCH**
- Up to **61 TFLOPS (FP64)** per APU

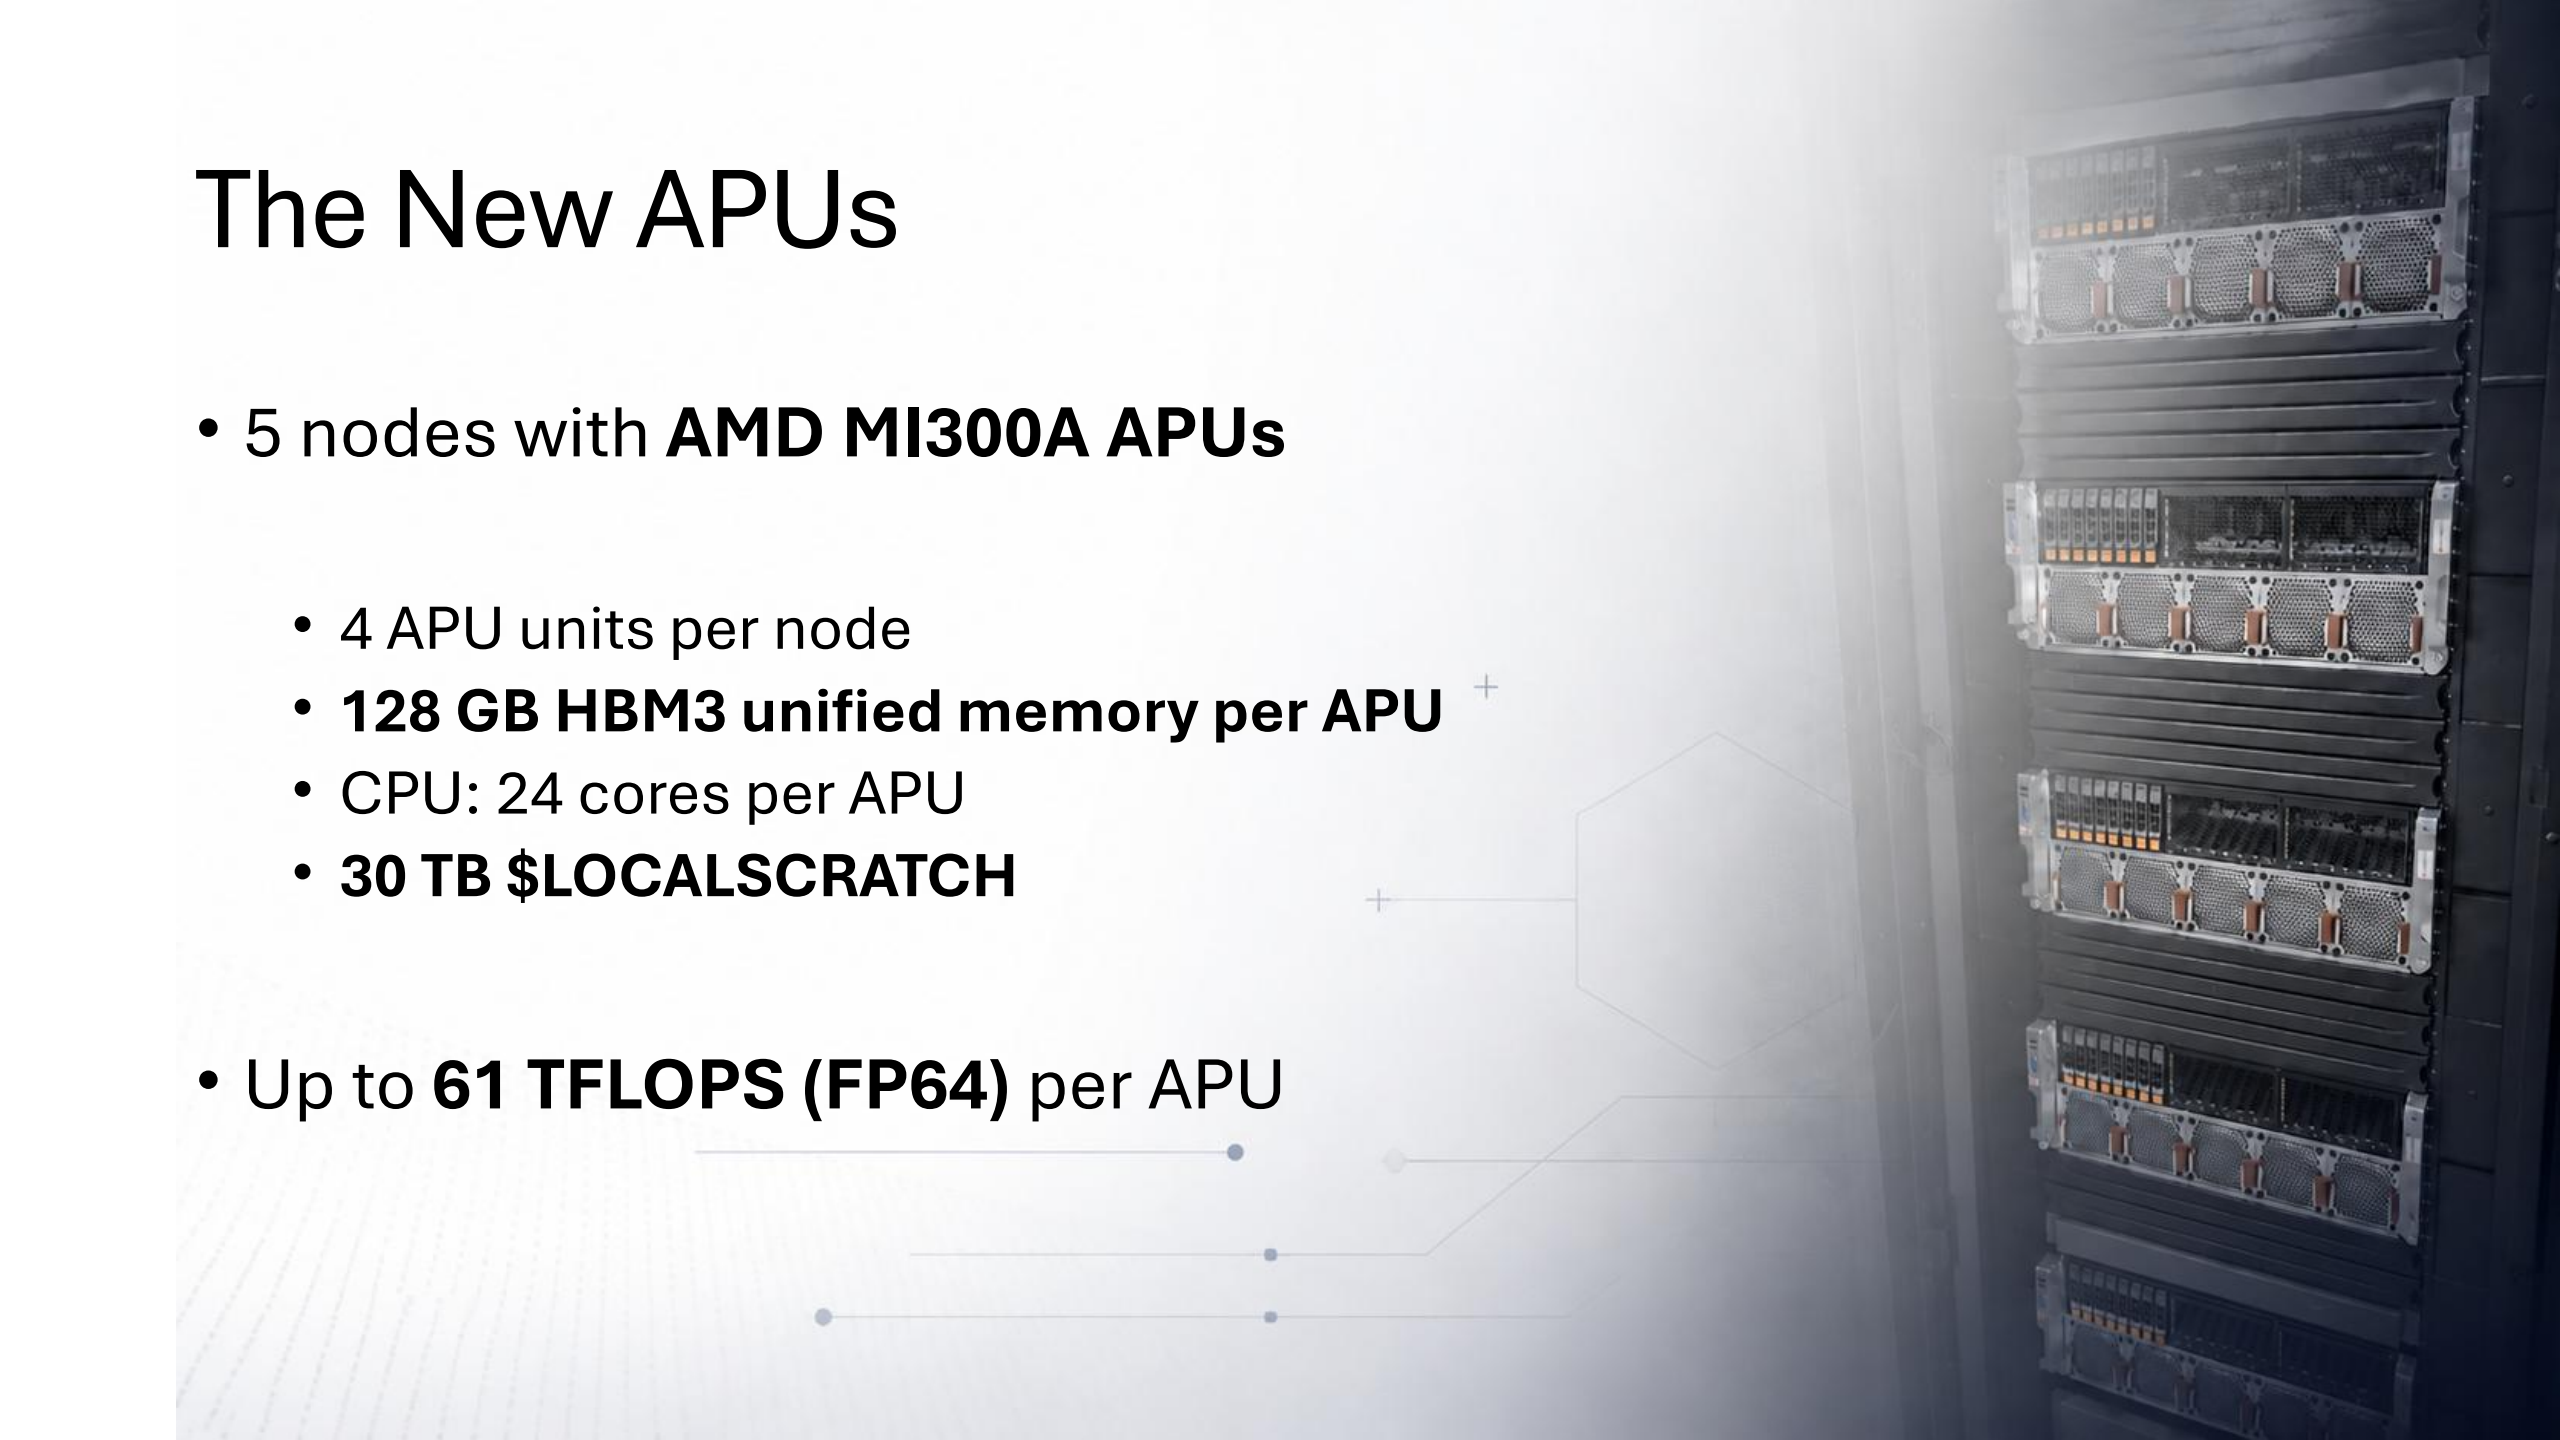

# A new type of workflows

- Designed for tightly coupled CPU–GPU workloads
- Ideal for memory-intensive and large-scale simulations
- Optimized for double precision (FP64)
- **Integrated on Lemaitre 4**
- Expected Q3 2026

# More than new compute nodes

- **Bridge toward EuroHPC systems** (e.g. LUMI)
- Prepare users for future exascale architectures
- Early access to new computing paradigms (APU)

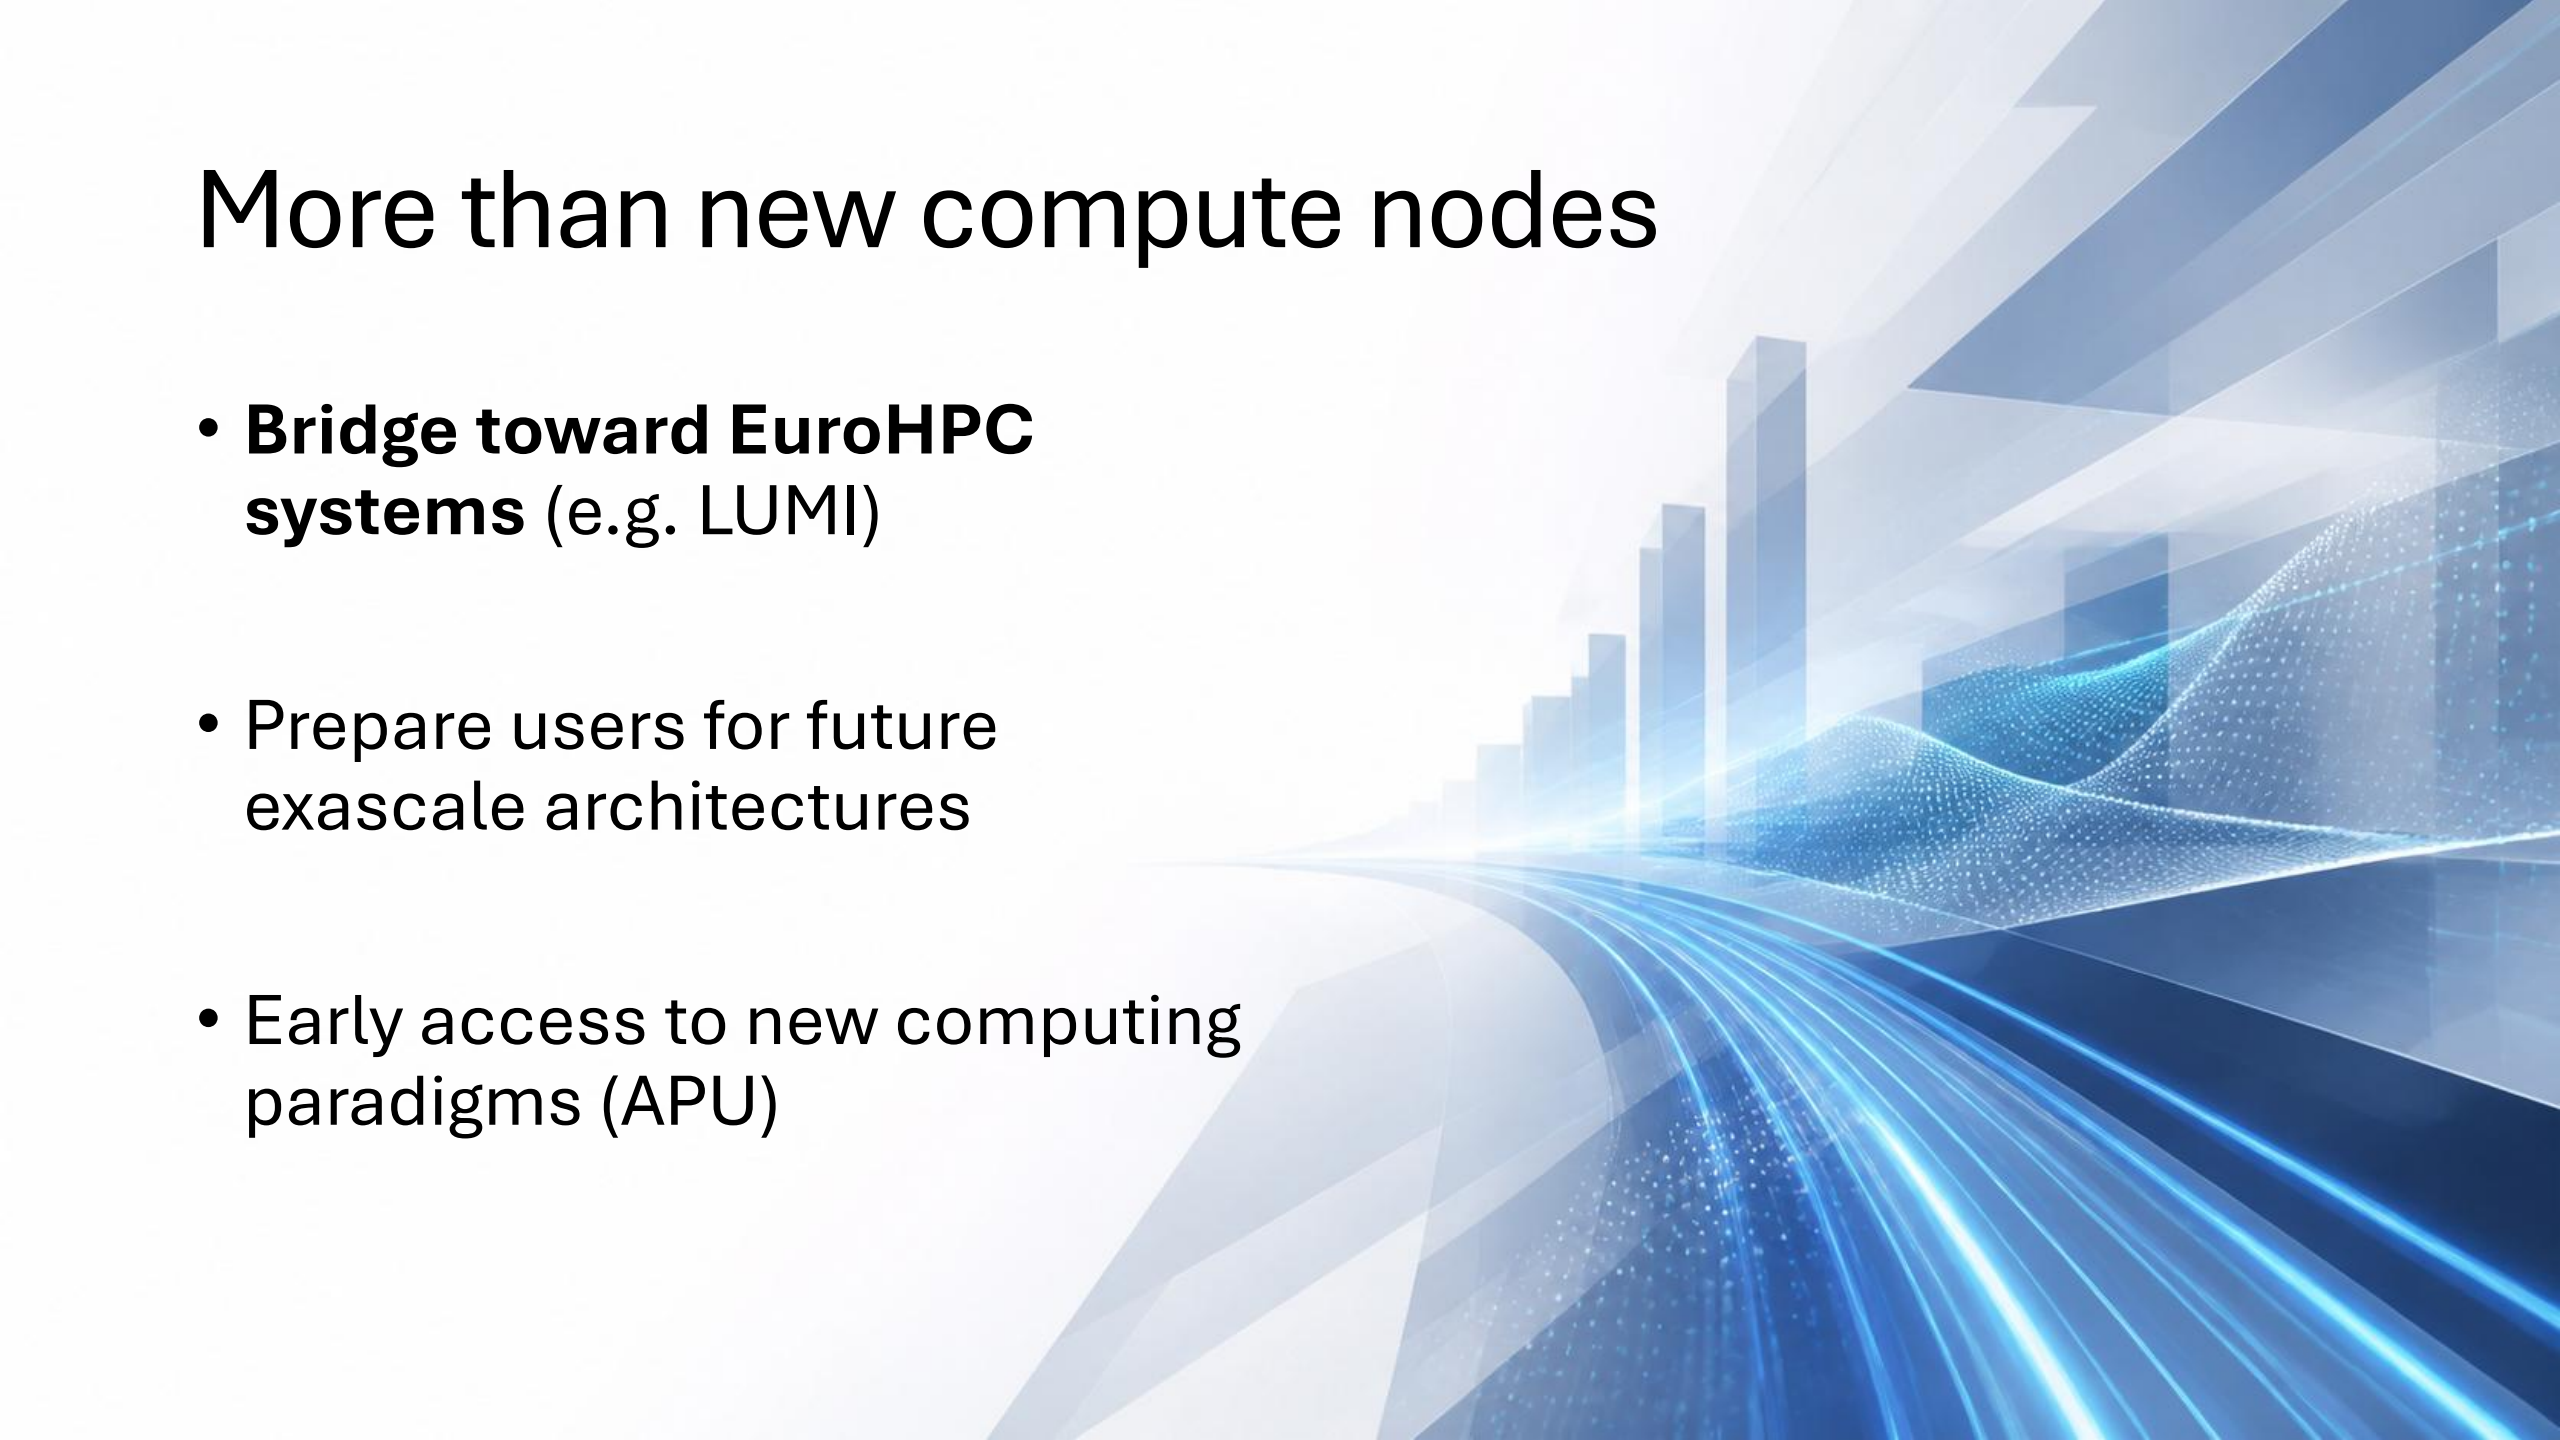

# New websites

## CÉCI homepage

## CÉCI documentation

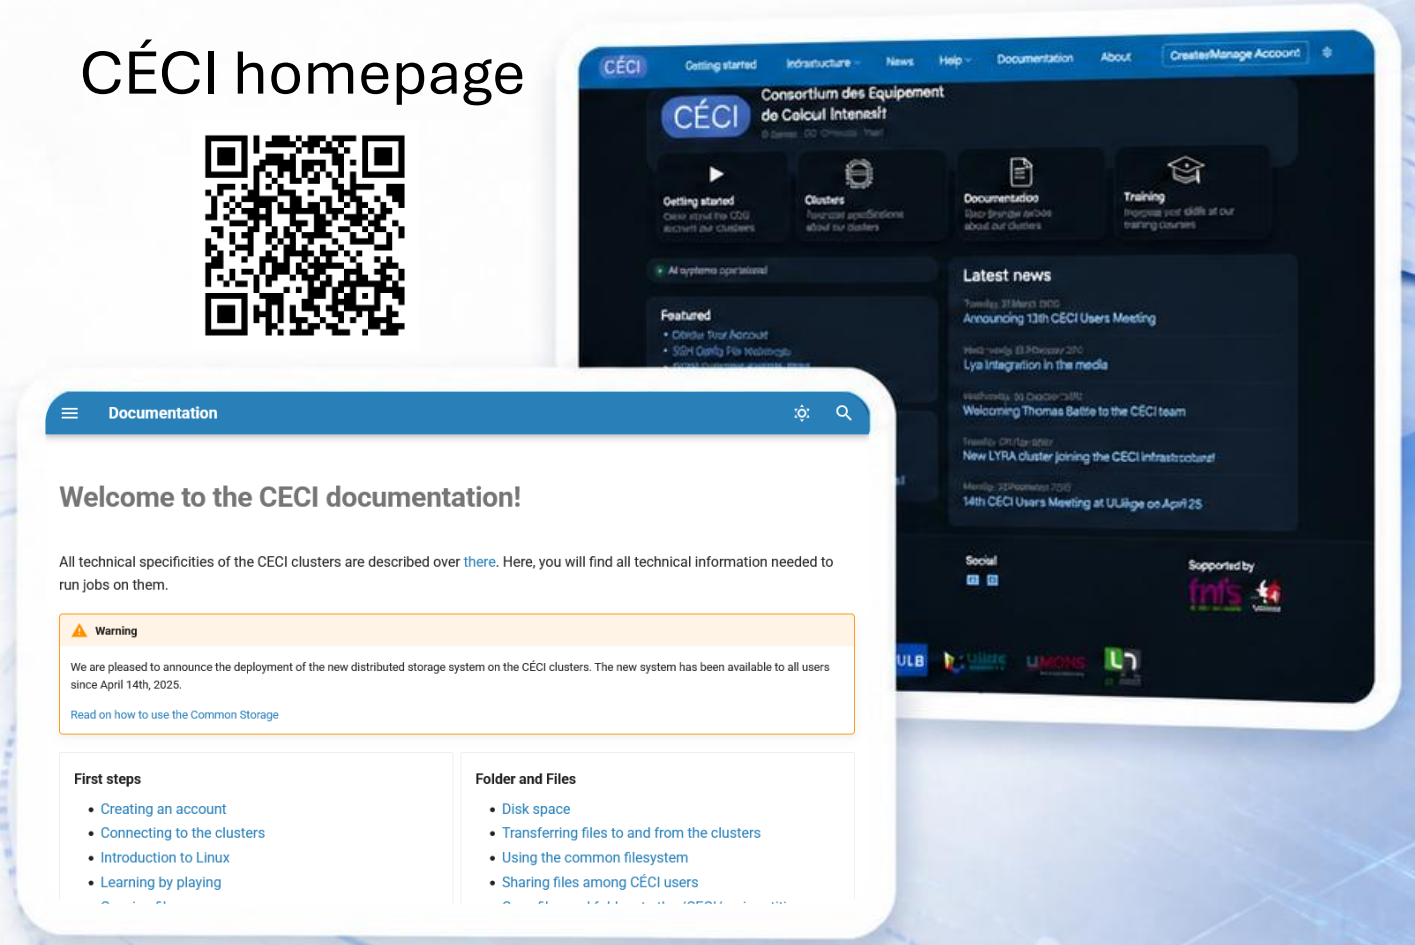

# Conclusion

- New clusters
- New storage
- New architectures
- New websites

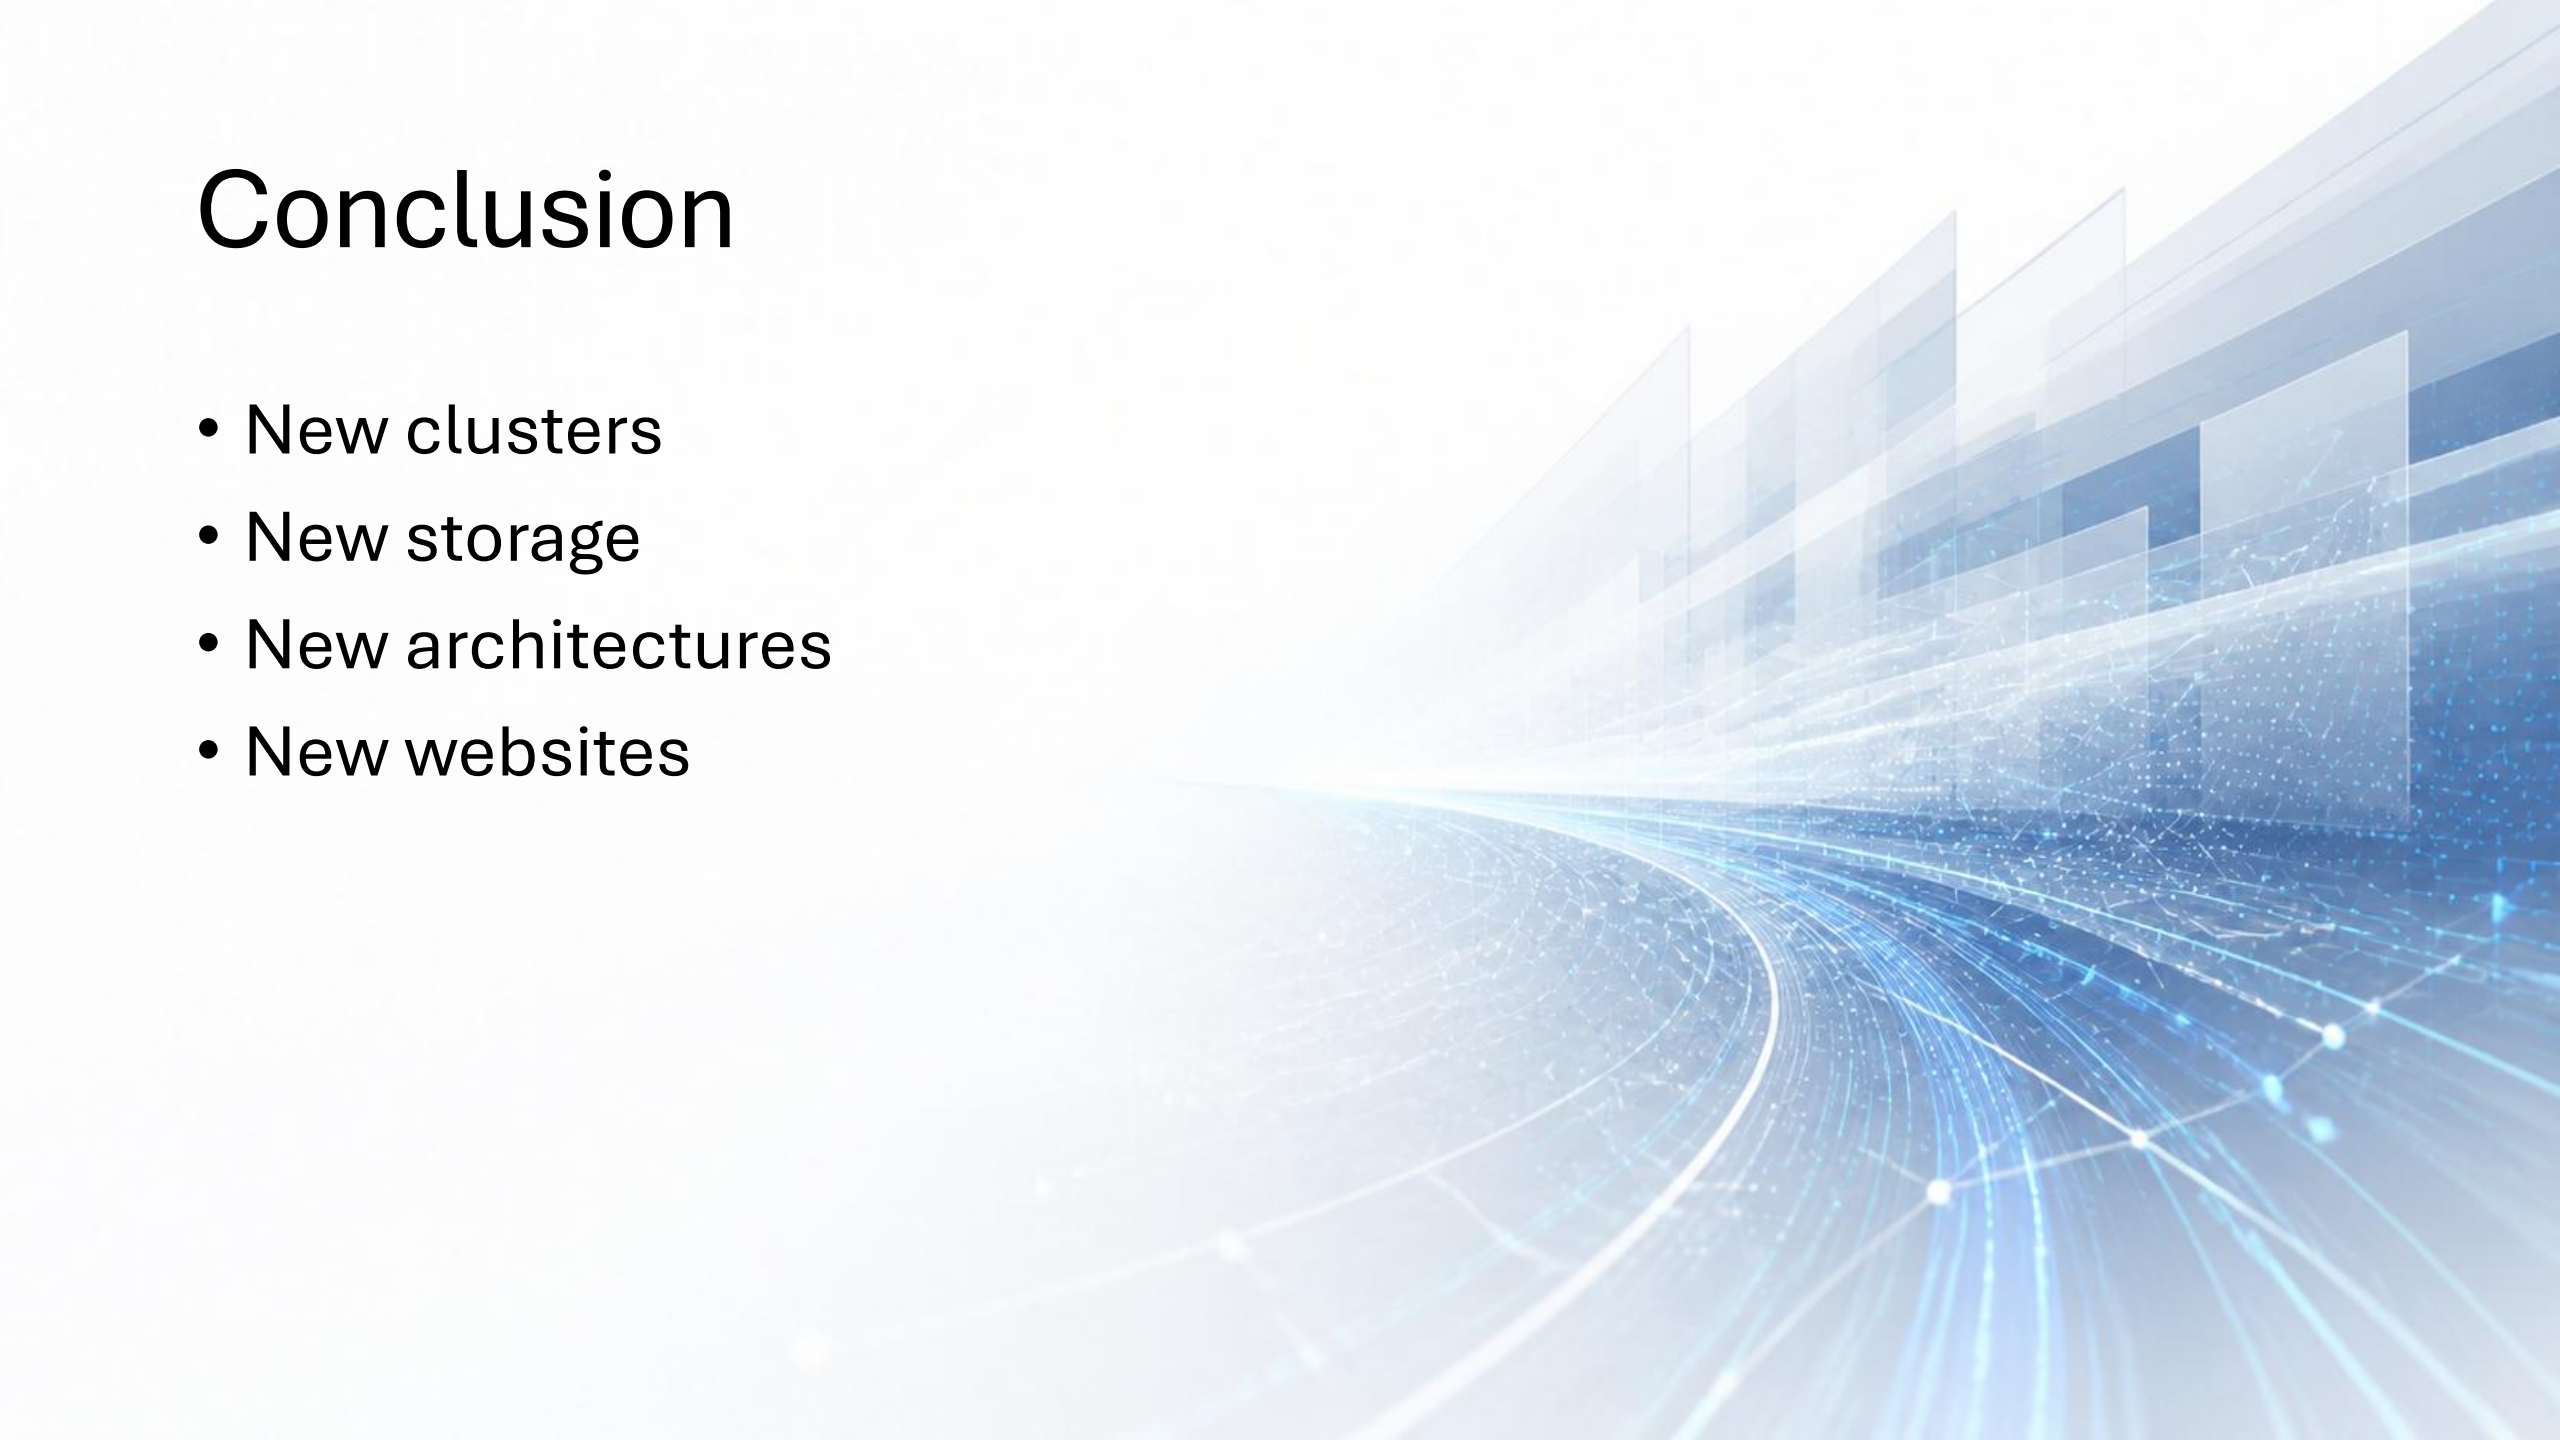

# Thank you for your attention

Wallonie

fnrs LA LIBERTÉ DE CHERCHER

EURO BELGIUM

UCLouvain Université

ULB

LIÈGE université

UMONS

UNIVERSITÉ DE NAMUR