

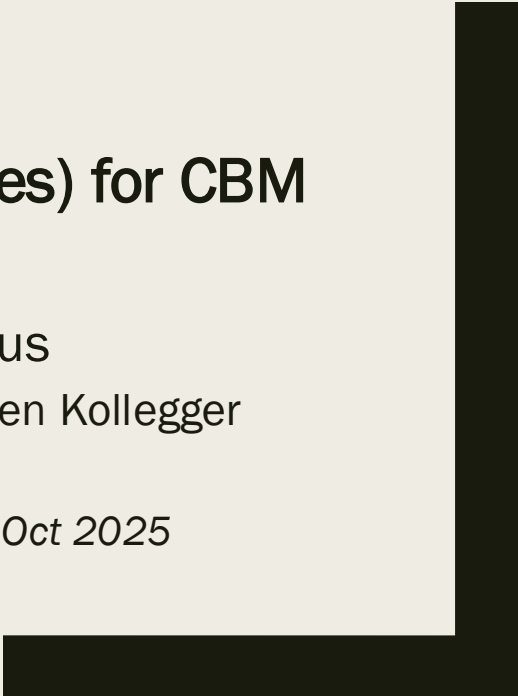


Online (and offline) Processing (and Storage Resources) for CBM

Technical development/Status

Mohammad Al-Turany, Florian Uhlig, Thorsten Kollegger




CBM Collaboration Meeting, Lanzhou, 22. Oct 2025



Status and plans: Cluster

■ At the Moment

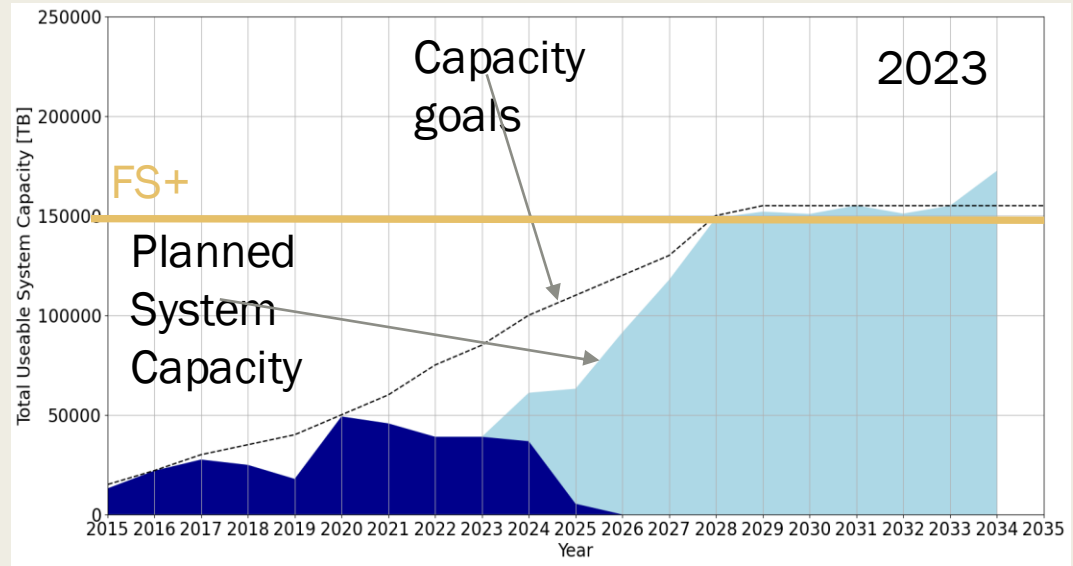

- 70 000 Cores
- 400 AMD MI 100 GPU
- 32 Nvidia H200 GPUs
- HDR InfiniBand



CBM (100 days), NUSTAR (180days), HADES (30 days), and APPA (180 days)

Status and plans: Storage

- At the Moment
 - *About 60 PetaBytes Lustre storage*



Ongoing activities the GSI-IT

- Investigate new Network options
 - *Ongoing POC with Huawei*
- Investigate new storage options/technologies
 - *New “/scratch” system (NVME based) for developers*
 - *Lustre on OceanDisk Hardware*
- Restructure/improve the cluster(s) organization

Huawei POC – Hardware Integration and Evaluation

Huawei provided two OceanDisk 1600 (500 TB) units and four 100 GE switches, including full technical support for this POC.

Objectives:

- Compare **underlay/overlay setups** (setup complexity, performance, monitoring, troubleshooting)
- Integrate hardware into **existing orchestration systems** (full lifecycle management)
- Analyse **Lustre performance over RoCE vs. InfiniBand**

Outcome: Baseline for Ethernet-based data plane operation in FAIR IT environments.

InfiniBand vs RoCE: Advantages and Considerations

Aspect	InfiniBand	RoCE / Ethernet
Latency / Determinism	Excellent (~1–2 μ s), very stable	Good (~2–5 μ s), depends on PFC/ECN tuning
Performance Ecosystem	Mature for MPI, Lustre, HPC	Rapidly improving; still tuning-intensive
Cost & Hardware Availability	Higher, proprietary	Lower, commodity hardware; multi-vendor
Integration & Orchestration	Limited (HPC-specific tools)	Natively supported in cloud/HPC orchestration (Slurm, K8s, OpenStack)
Scalability / Fabric Size	High but IB-specific	Very large via standard Ethernet
Monitoring & Telemetry	Unified IB tools (Subnet Mgr, perf query)	Broader but fragmented ecosystem (sFlow, INT, gNMI)
Vendor Lock-in	Strong (mostly NVIDIA IB)	Weak — multi-vendor ecosystem

InfiniBand vs RoCE: Advantages and Considerations

Aspect	InfiniBand	RoCE / Ethernet
Latency / Determinism	Excellent (~1–2 μ s), very stable	Good (~2–5 μ s), depends on PFC/ECN tuning
Performance Ecosystem	Mature for MPI, Lustre, HPC	Rapidly improving; still tuning-intensive

Pros of RoCE: Unified network fabric, lower cost, cloud-native integration, easier lifecycle automation

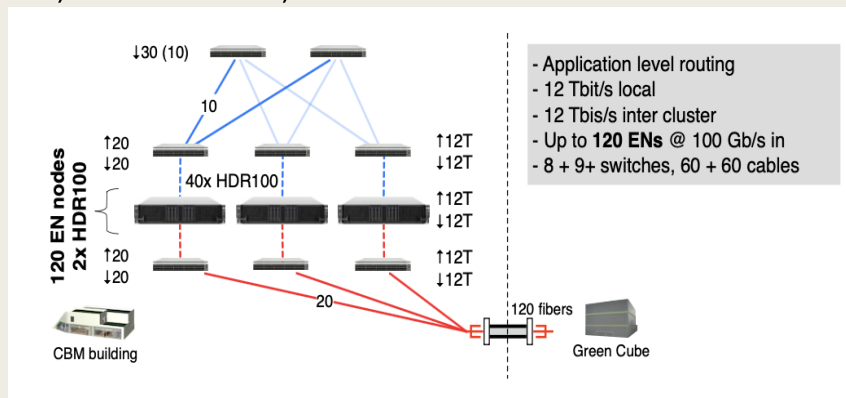
Cons: Higher configuration effort, latency tuning

Monitoring & Telemetry	Unified IB tools (Subnet Mgr, perf query)	Broader but fragmented ecosystem (sFlow, INT, gNMI)
Vendor Lock-in	Strong (mostly NVIDIA IB)	Weak — multi-vendor ecosystem

Potential Impact of RoCE Adoption (?) on CBM Operations

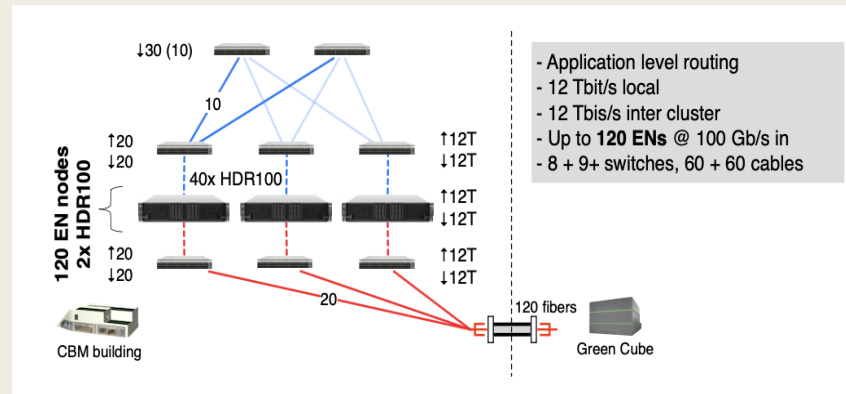
Hardware: No direct impact for CBM but:

- CBM FLES nodes with dual-port InfiniBand HCAs will need to configure one port for RoCE (Ethernet) to communicate with GreenCube.
- Requires driver, firmware, and software validation for RoCE compatibility.



Potential Impact of RoCE Adoption (?) on CBM Operations


- Adaptation of DAQ → storage → processing data flow to hybrid InfiniBand–RoCE topology.
- Technical considerations: Performance parity, QoS, congestion control, and monitoring integration.
- Only the external data link may need to evolve to RoCE.



Potential Impact of RoCE Adoption (?) on CBM Operations

- Adaptation of DAQ → storage → processing data flow to InfiniBand–RoCE topology.
- Technical considerations: Performance, control, and monitoring integration
- Only the external network needs to be converted to RoCE.

[Setup & Integration] → Performance & Analysis → Final Report & Decision




New NVME based filesystem for developers

- Central volatile scratch storage for software development.
 - *Several file servers with fast NVME disks*
 - *Use for compilation, installation, ...*
- Space is limited by quota (per group).
- No backup.
- Read-write-access through submitter nodes
- read-only on batch worker nodes
- So far very positive user feedback

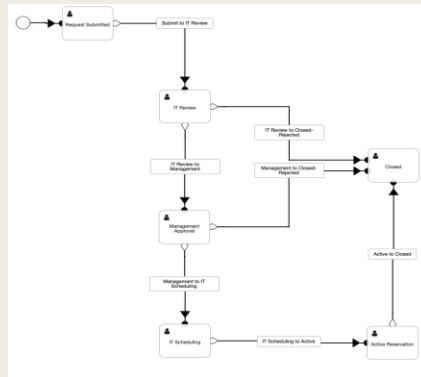
Reservations for online processing

- Up to now we handled this a reservation within our working Slurm environment
- CBM requested per Mail a certain number of Nodes for a certain time
- It worked but with certain constrains and some times problems! e.g:
 - *Submit nodes became unresponsive (due to other users)*
 - *Shared memory segments are not always correctly cleaned up*
 -



Reservations for online processing: Next

- Reservation is handled by *process ticket*
- Dedicated Slurm Controller for Online processing
- Dedicated Submitter nodes
- Full compute nodes allocated for fixed time periods
- Resources isolated from other systems



Process:

Compute Resource Reservation

Title:

mCBM

Request Title

Customer User:

"cslurm" <C.Slurm@gsi.de>

Customer ID:

C.Slurm@gsi.de

Customer

Reservation_Type:

Standard

Reservation Type

Experiment_Name:

CBM

Experiment Name

Number_of_Compute_Nodes:

10

Number of Nodes

Requested_Start_Date:

10 / 05 / 2026

Requested Start Date

Requested_End_Date:

20 / 05 / 2026

Requested End Date

Cost_Center:

1011

Comment:

This is a test

Submit Request

Key Takeaways

- CBM input is essential for the IT system design process.
- CBM will operate within a larger, shared infrastructure — **not as a standalone system.**
- System optimization targets overall FAIR operations, including both **online and offline activities.**
- **CBM-specific** implementations must align with the common system architecture and design principles.