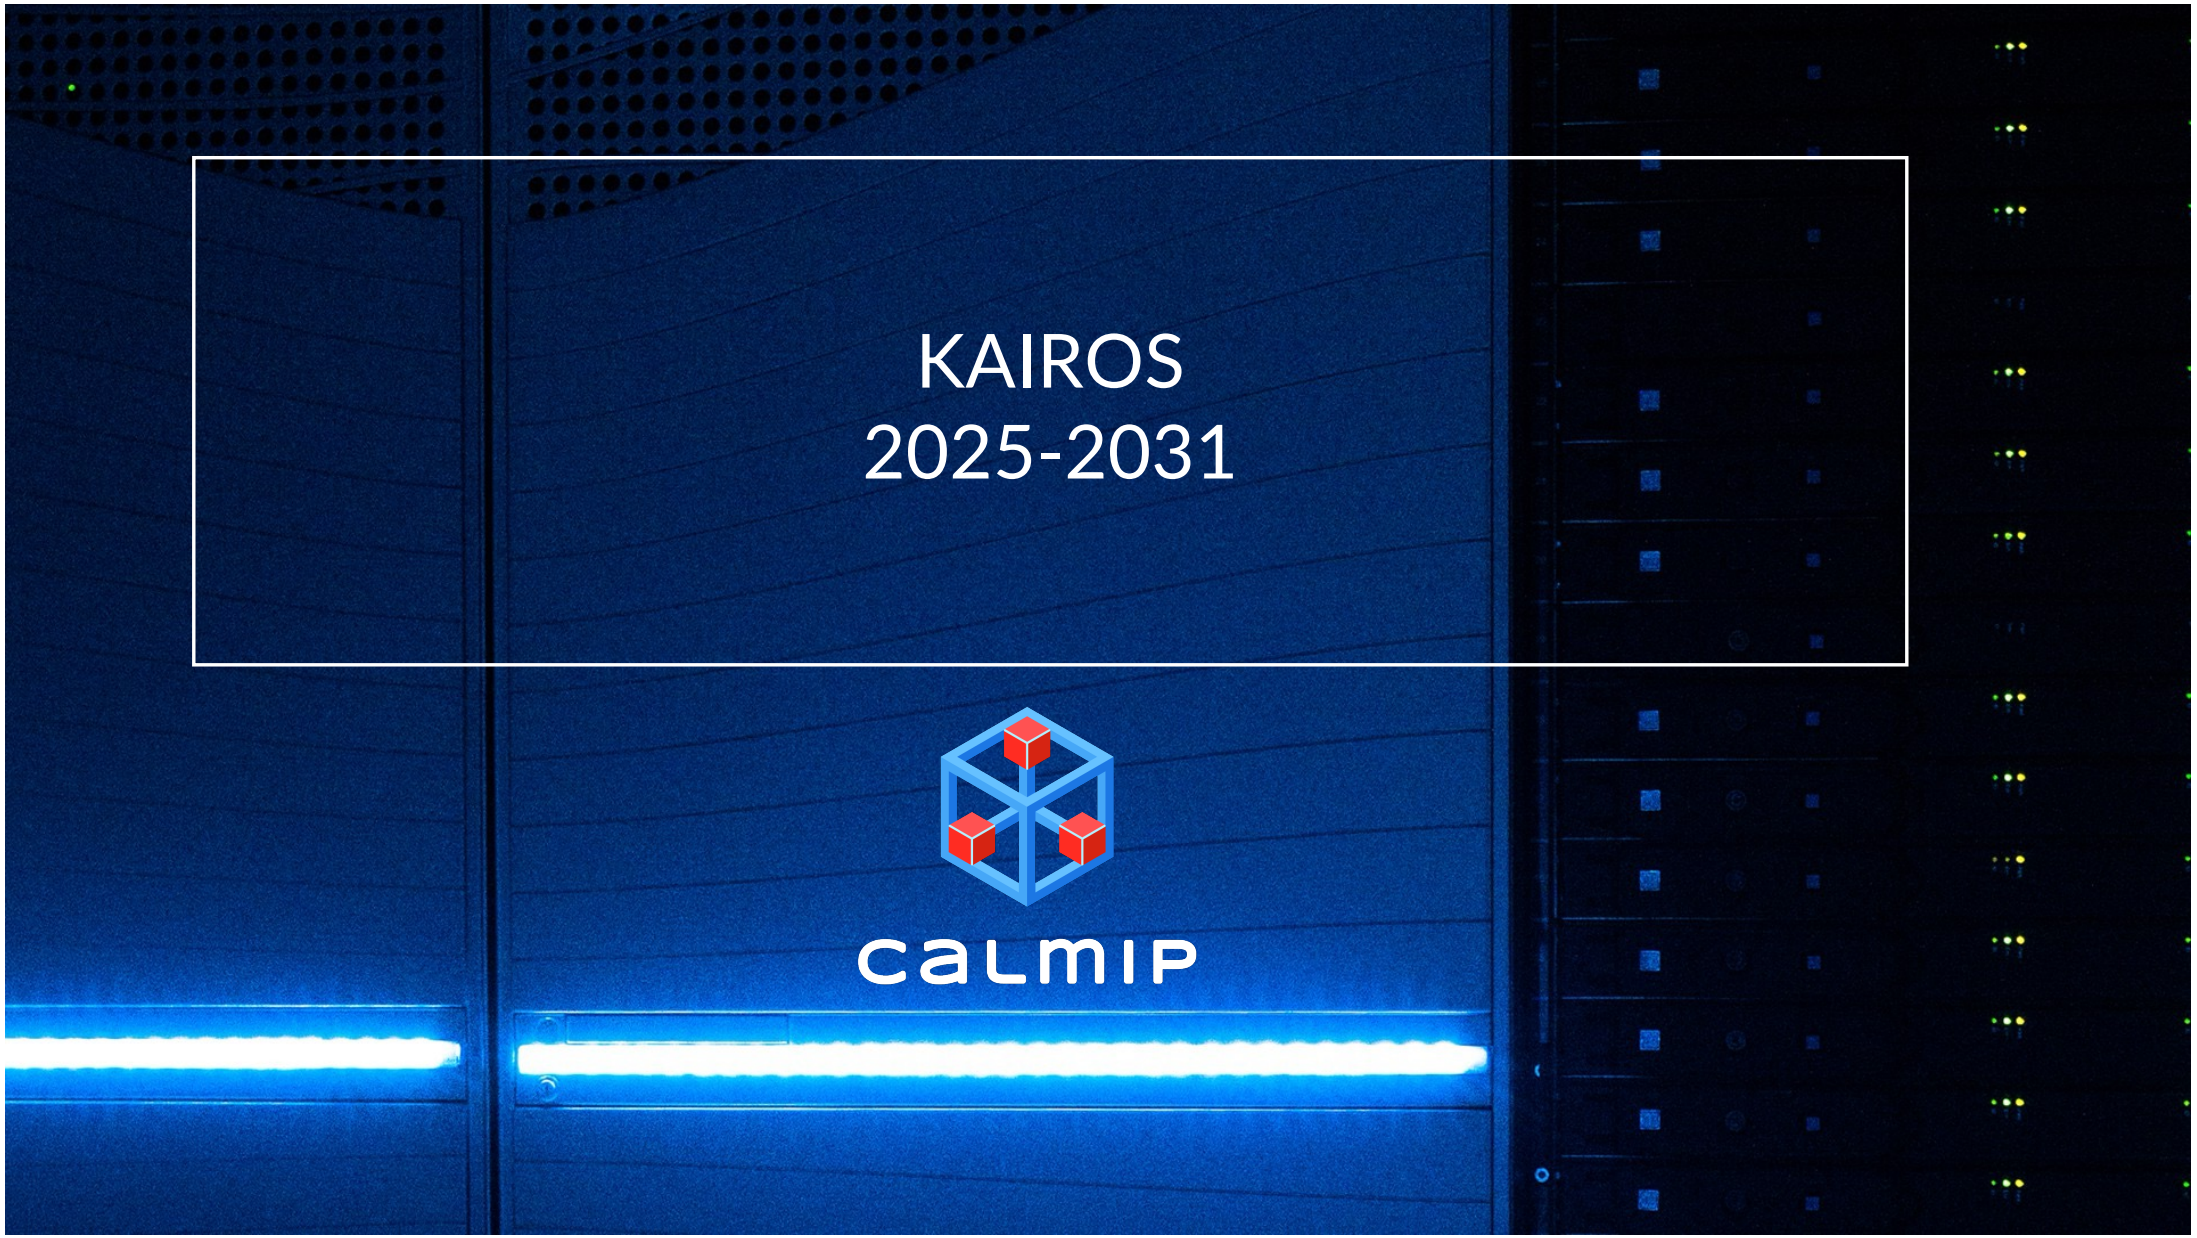
KAIROS
2025-2031
calmip

SUPERCALCULATEUR KAIROS 2025-2031

XH3000 EVIDEN

- 1st@Green500 Nov 2025, St. Louis, SC25
- 24576 cores / 64 GPU / 4.77 Pflop/s peak
- 186 kW HPL (consommation énergétique Calcul stable depuis 2014)
- Partition Accélérée (GPU) : 3.05 PF Rmax (High Performance Linpack)
- Réseau rapide Infiniband NDR (200 Go/s)
- 5 années de Maintenance
- 100 Jours Accompagnement
- Refroidissement à cœur – eau chaude
- Installation : Automne 2025
- Production : début 2026

Crédit photo CNRS

CALMIP DATA

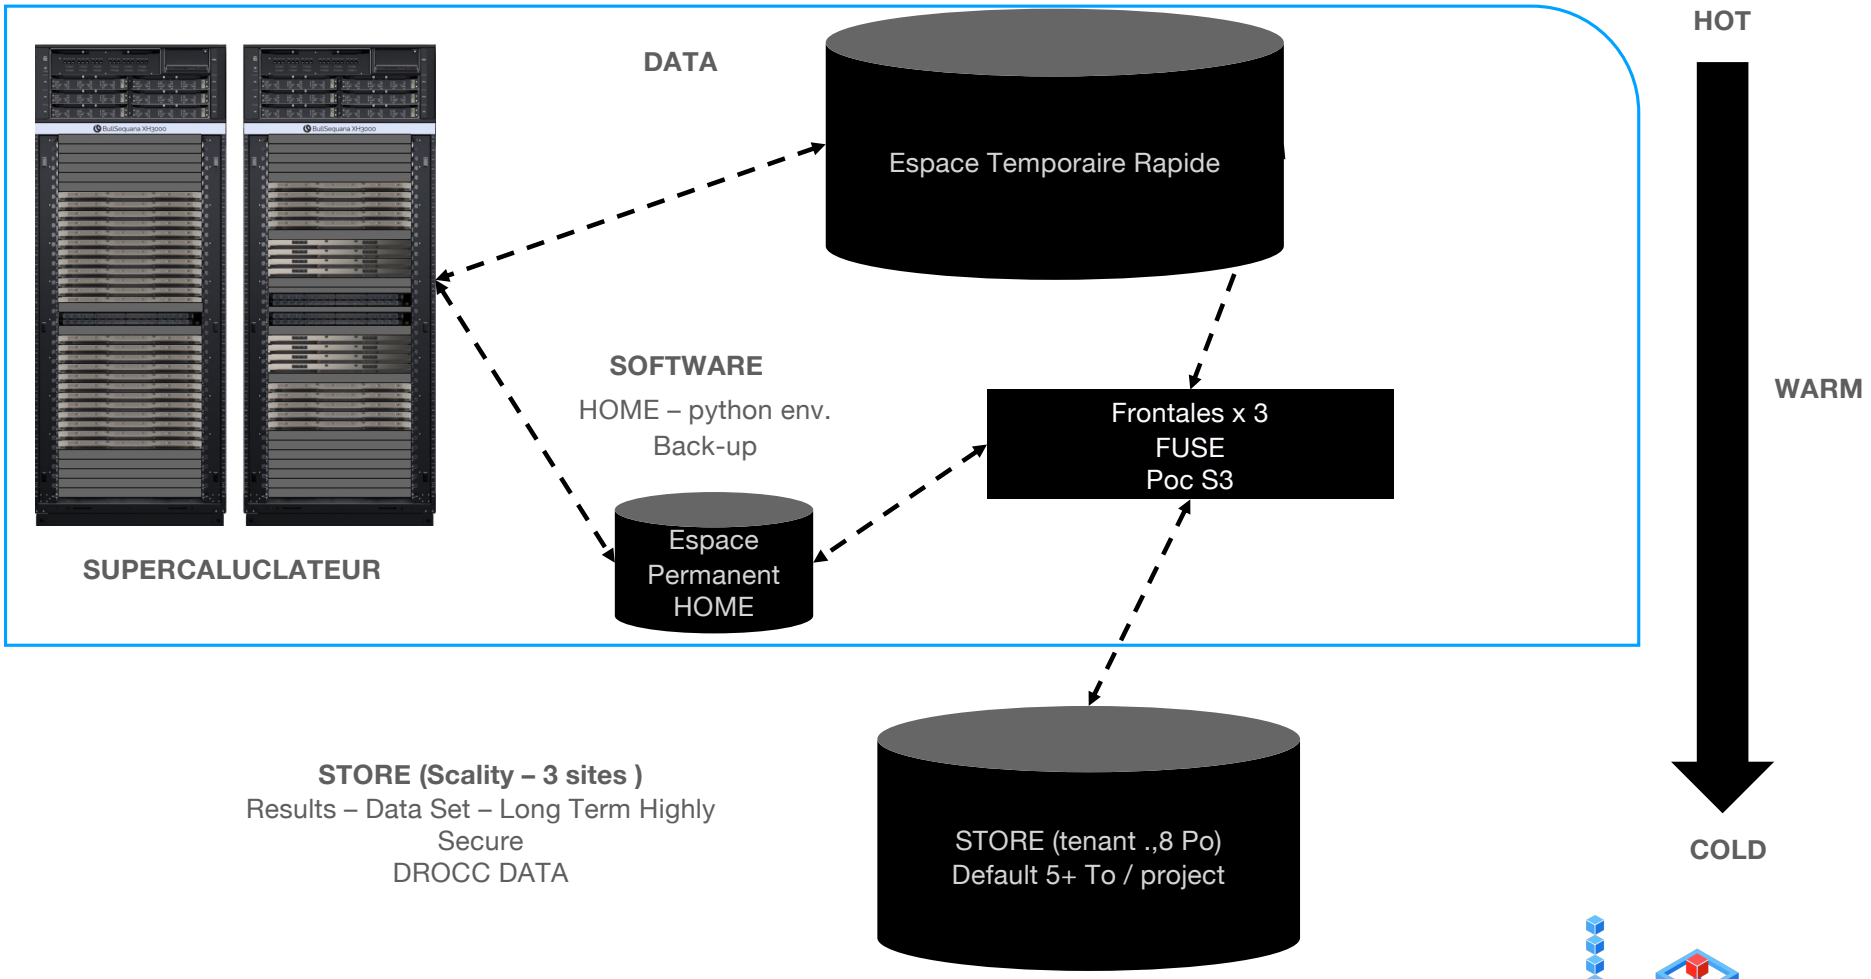
DATA
Espace Temporaire Rapide
SOFTWARE
HOME – python env.
Back-up
Frontales x 3
FUSE
Poc S3
Espace
Permanent
HOME
SUPERCALUCLATEUR
HOT
WARM
COLD
STORE (Scality – 3 sites)
Results – Data Set – Long Term Highly
Secure
DROCC DATA
STORE (tenant .,8 Po)
Default 5+ To / project

SUPERCOMPUTER KAIROS : PARTITIONS

Partition Accelerated (GRACE-HOPPER)

3.2 PF Peak – 3.05 PF Rmax (High Performance Linpack)

- GRACE-HOPPER node : 4 x [72 cores ARM – GPU H200] – 864 Go Unified memory – Nvlink
- 16 nodes GRACE-HOPPER – 4608 cores ARM (Nvidia) – 64 GPU H200-96 Go HBM3 (Nvidia)

Partition Massivaly Parallel processing 1.35 PF Peak

1.35 PF Peak

- MPP node : 2 x Intel granit 96 cores 2,1 Ghz – MRDIMM - 8,8 GB/s/core – 768 GB
- 102 compute node – 19584 cores x86 « Granit » (Intel)

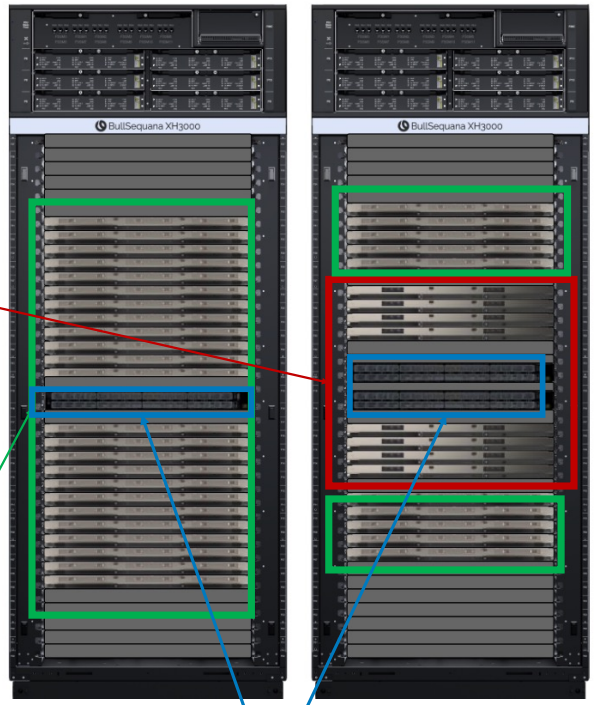

Infiniband NDR 200GB/s

CALMIP

SUPERCALCULATEUR KAIROS : PARTITIONS

XH3000 EVIDEN (refroidissement à cœur –eau chaude)

- Partition Accélérée (Acc.) : 3.2 PF Peak – 3.03 PF Rmax (High Performance Linpack)
- 16 nœuds GRACE-HOPPER – 4608 cores ARM (Nvidia) – 64 GPU H200-96 Go HBM3 (Nvidia)
- 1 nœud GRACE-HOPPER : 4 x [72 cores ARM – GPU H200] – 864 Go Unified memory - Nvlink
- Partition Massively Parallel Processing (MPP) : 1.5 PF Peak
- 102 Nœuds de calcul – 19584 cores x86 (Intel)
- 1 nœud MPP : 2 x processeurs Intel granit 96 cores 2,1 Ghz – MRDIMM - 8,8 Go/s/cœur – 768 go
- Réseau Rapide InfiniBand NDR 200 Go/s

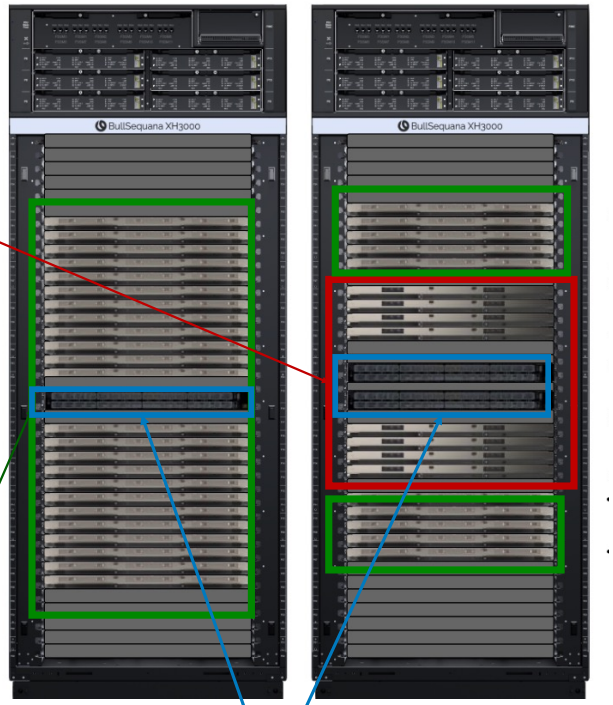

CALMIP

SUPERCALCULATEUR KAIROS: SERVICE

SERVICE (2 Rack air – Hébergement DROCC TierIII+)

- 3 nœuds de logins
- 3x[2xGranit rapid 96c – 384 Go]
- 2 nœuds de visualisation
- VirtualGL – TurboVNC
- 2 x [2 x Granit-rapid 96c – 768 Go – Nvidia L40 48Go]
- 2 nœuds large mémoire (2 x 1.5 To RAM)
- Espace NFS permanent 80 To
- Espace Temporaire Rapide : 2.3 Po IBM Storage Scale
- 400 To NVMe
- 290 Go/s lecture
- 145 Go/s écriture

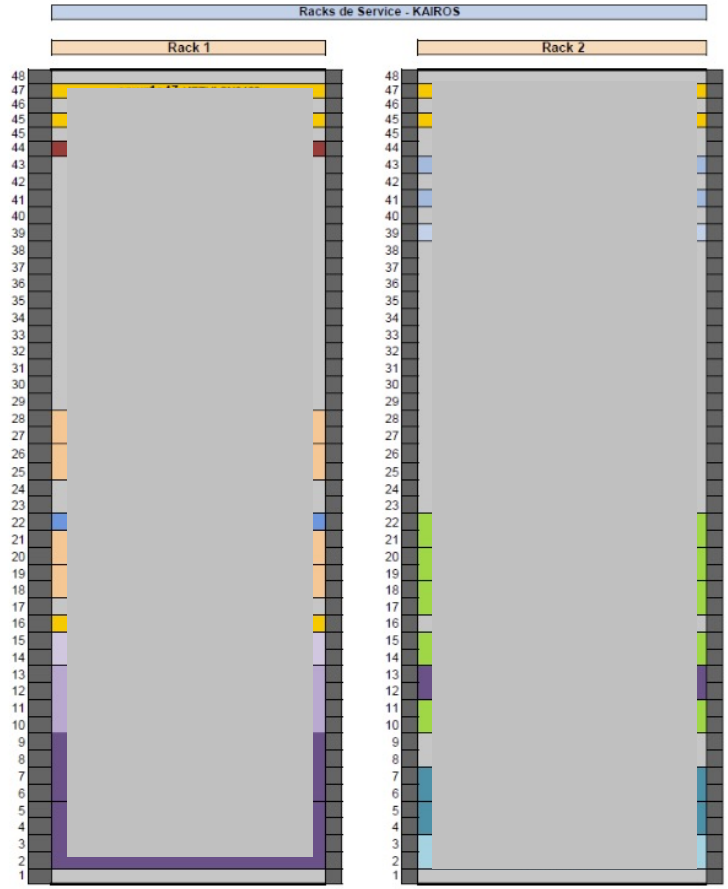

KAIROS: ENVIRONNEMENT

- X86
- Intel oneAPI HPC Toolkit (+Intel MPI, Vtune, Advisor, ITAC) + Support Intel + training
- Support norme OpenMP5.0 ; 5.1, 5.2 partiellement ; quelques spécifications OpenMP 6.0 (TR12)
- GNU Compiler
- Arm + GPU
- GNU compiler
- ARM Suite (compiler, Performance Libraries)
- NVPHC (compilers, cuda, math libs, OpenMPI, NCCL, Nsight, …) DCGM ?
- Support norme OpenACC 2.6, quelques fonctionnalités de la norme OpenACC 2.7
- Linaro Forge (DDT+MAP) + training
- 128 process + 16 GPU ARM & x86
- ARGOS (EVIDEN) + training
- Energy monitoring, job
- Profiling, etc ….
- Eviden OpenMPI (OpenMPI 4.x, OpenMPI 5.x, IB, BXI, support norme MPI3.1)

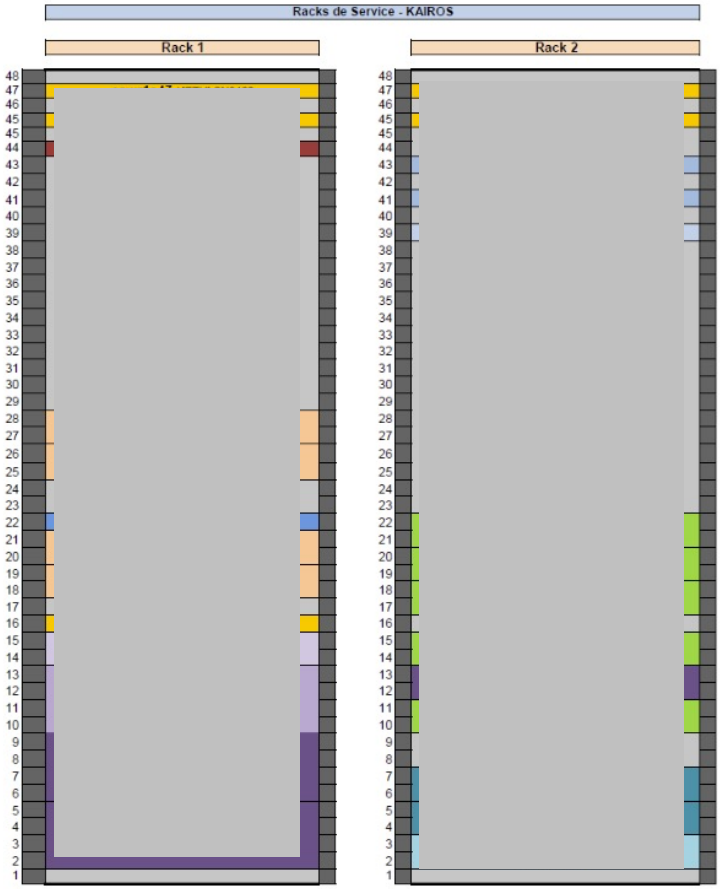

CALMIP

KAIROS (PARTITION ACC.) : GRACE-HOPPER

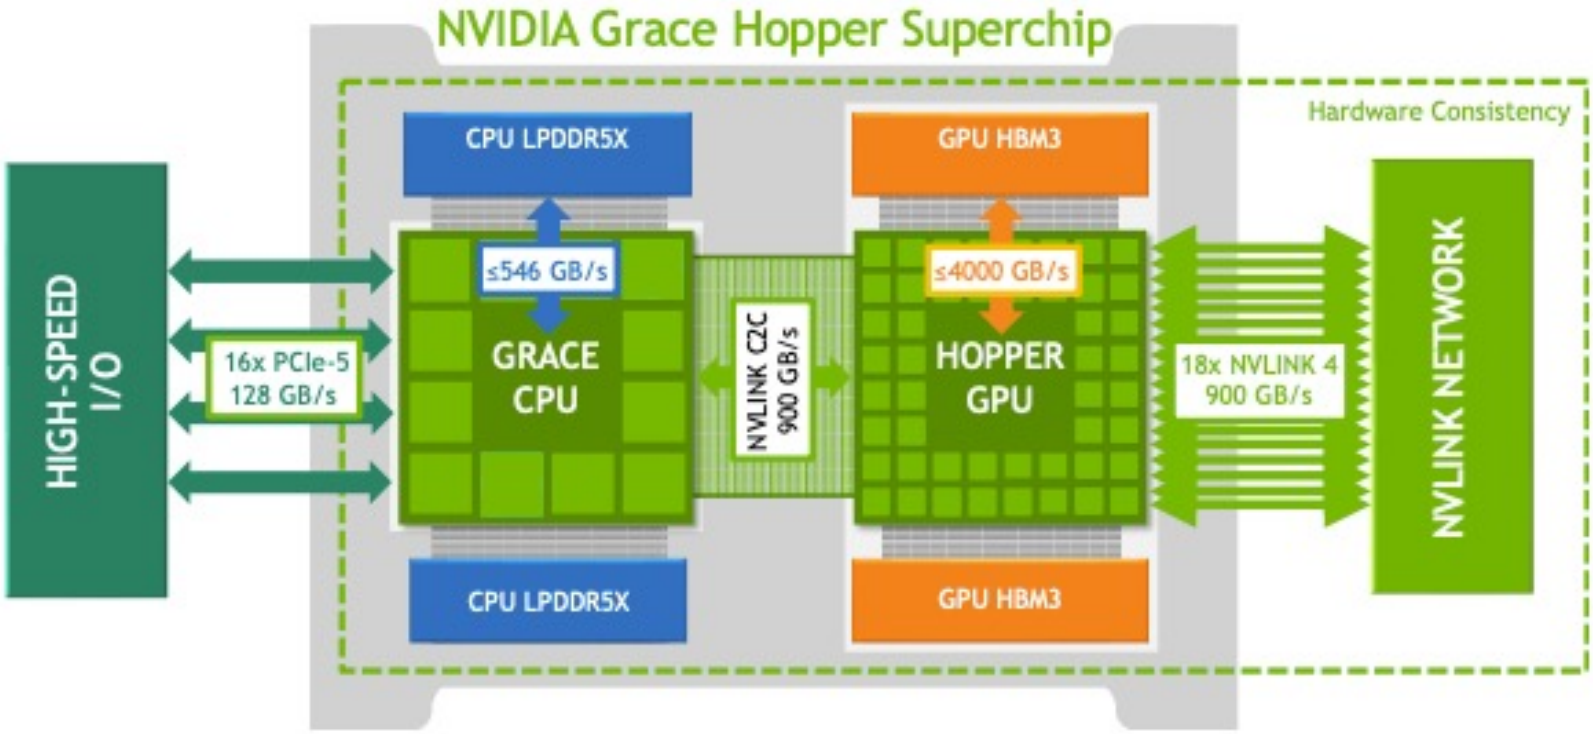

680 WATT TDP
120 Go LPDDR
96 Go HBM3 –STREAM 3.75 TB/s

CALMIP

KAIROS (PARTITION ACC.) : NŒUD GRACE-HOPPER

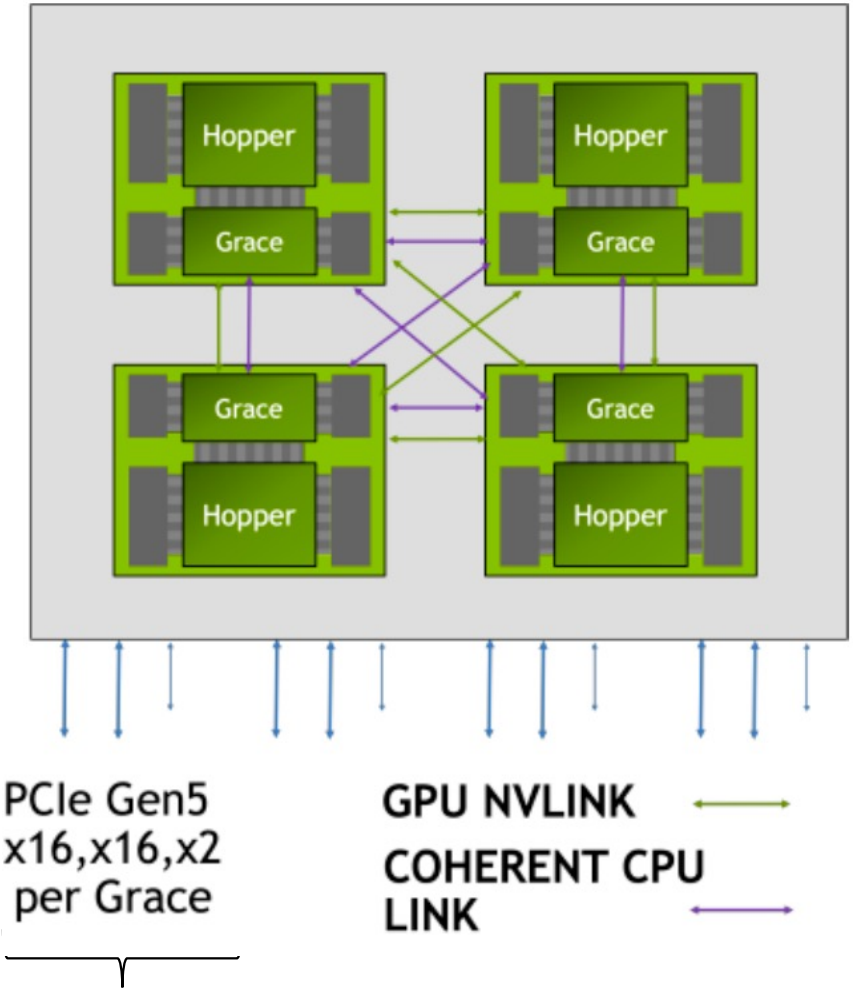

4x [72 cores ARM – GPU H200]
864 Go Unified memory
200 TF (GPU) – 14 TF (CPU)

4 X link IB NDR (200 Go/s)

CALMIP

KAIROS (PARTITION MPP): NŒUD GRANIT-RAPID

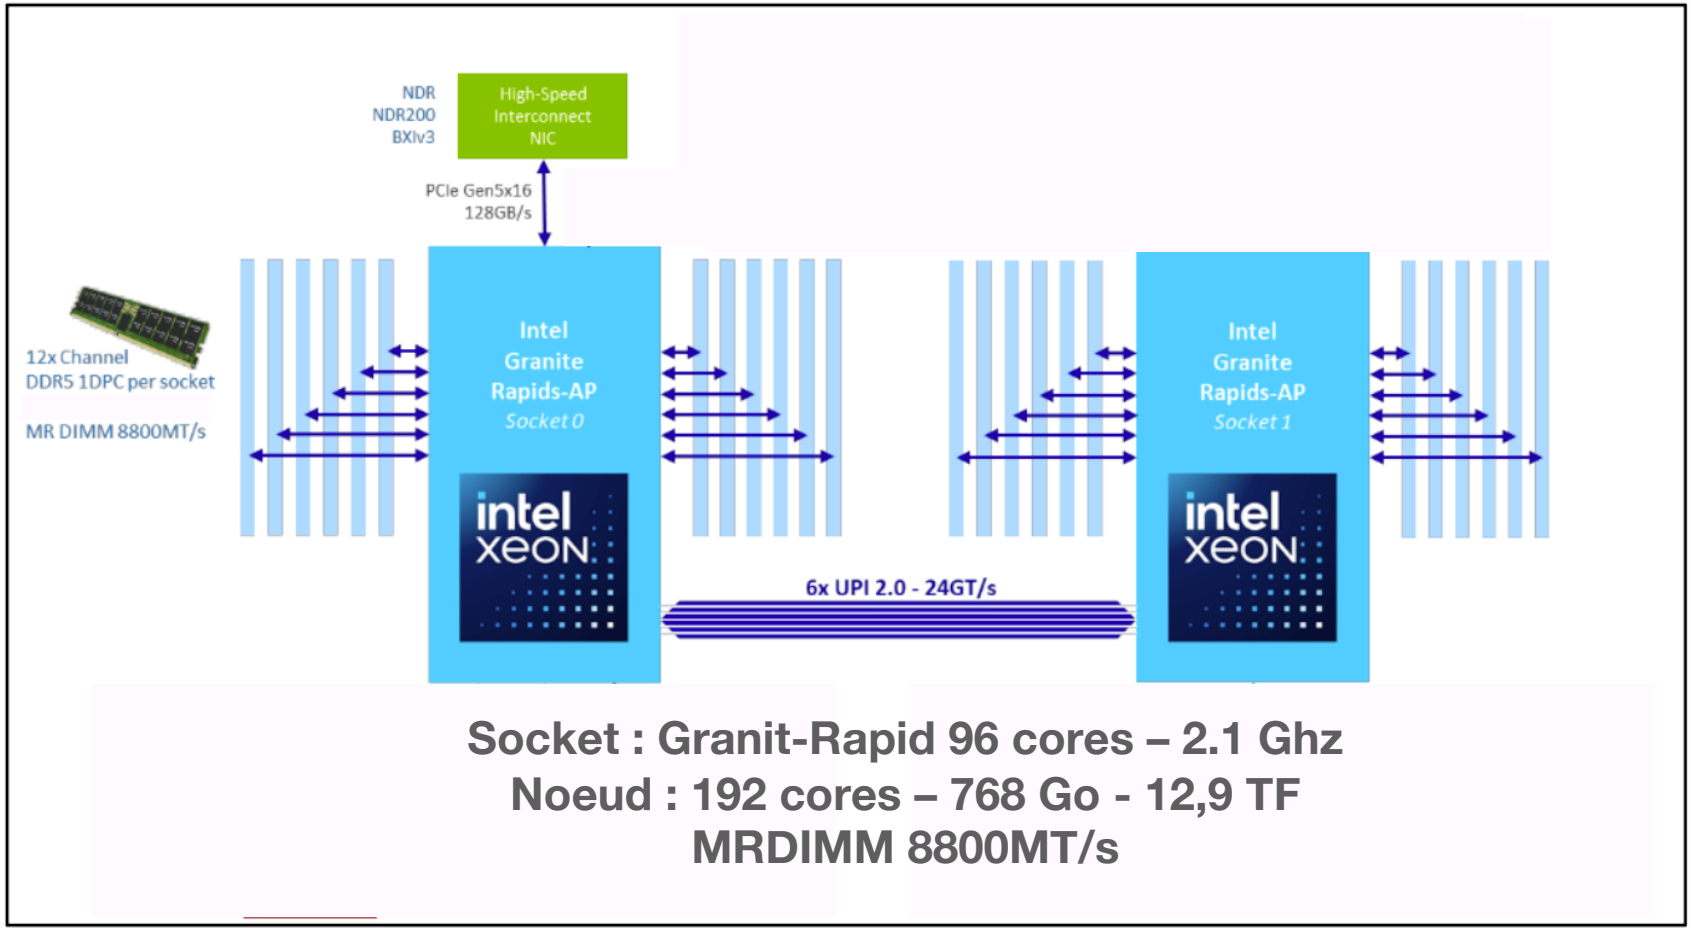

Figure 1- BullSequana XH3150 (TINO): architecture noeud

CALMIP