Funcionamiento y usos de una Supercomputadora

Marcos Mazzini | CPA conicet / CCAD (UNC)

Disclaimer:

- El centro de cómputo de Clementina XXI está en proceso de conformación.
- El equipo todavía está en fase de validación por parte Lenovo.
- Mis opiniones son mías, no representan una posición oficial.
- Los datos son datos 🙂

CUÁNDO ES HPC?

High Performance Computing

- Realizar **miles de billones** de operaciones en punto flotante por segundo **FLOPS** (>1 PetaFLOP actualmente)
- **Almacenar** Petabytes de datos con acceso y dsponibilidad continua y rendimiento cercano a un disco local.
- **Paralelizar** las tareas entre todos los procesadores/ GPUs comunicando en tiempo real los cambios de estado eficientemente en términos de **escalabilidad**.
- Opcionalmente, **visualizar** una representación tridimensional de los resultados obtenidos a partir de los cálculos efectuados

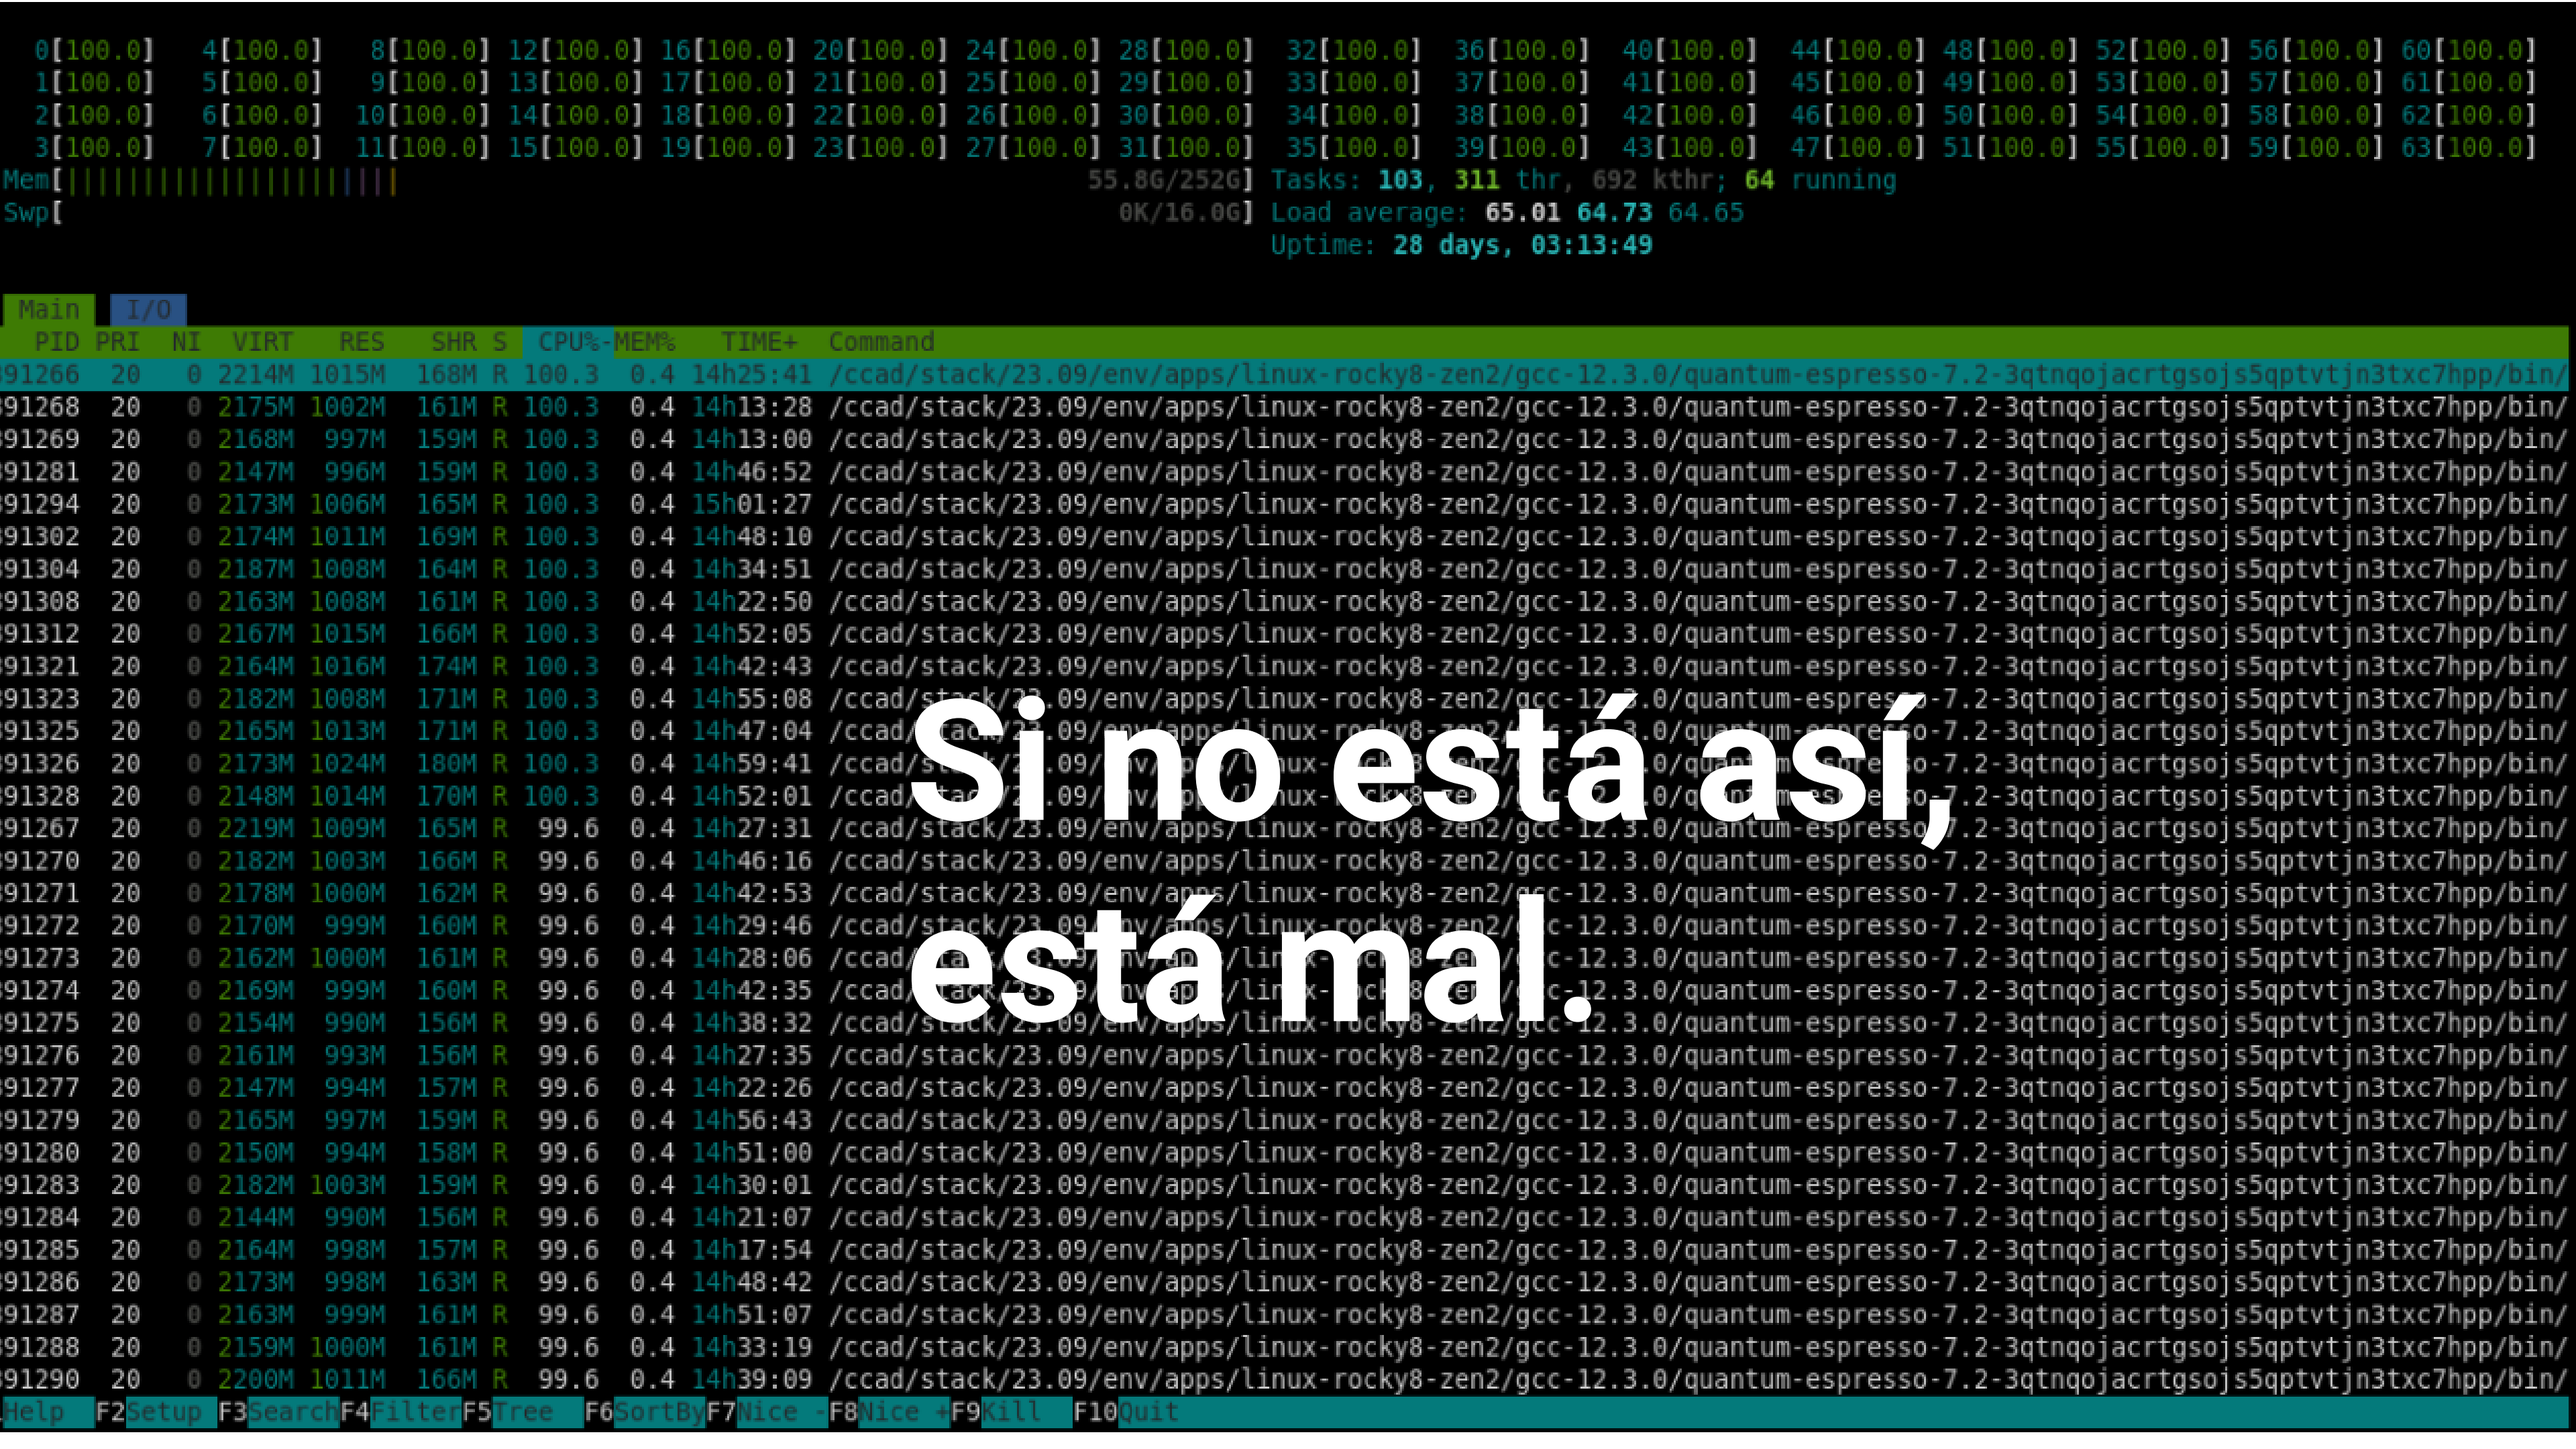
0[100.0] 4[100.0] 8[100.0] 12[100.0] 16[100.0] 20[100.0] 24[100.0] 28[100.0] 32[100.0] 36[100.0] 40[100.0] 44[100.0] 48[100.0] 52[100.0] 56[100.0] 60[100.0]
1[100.0] 5[100.0] 9[100.0] 13[100.0] 17[100.0] 21[100.0] 25[100.0] 29[100.0] 33[100.0] 37[100.0] 41[100.0] 45[100.0] 49[100.0] 53[100.0] 57[100.0] 61[100.0]
2[100.0] 6[100.0] 10[100.0] 14[100.0] 18[100.0] 22[100.0] 26[100.0] 30[100.0] 34[100.0] 38[100.0] 42[100.0] 46[100.0] 50[100.0] 54[100.0] 58[100.0] 62[100.0]
3[100.0] 7[100.0] 11[100.0] 15[100.0] 19[100.0] 23[100.0] 27[100.0] 31[100.0] 35[100.0] 39[100.0] 43[100.0] 47[100.0] 51[100.0] 55[100.0] 59[100.0] 63[100.0]
Mem[55.8G/252G] Tasks: 103, 311 thr, 692 kthr; 64 running
Swp[0K/16.0G] Load average: 65.01 64.73 64.65
Uptime: 28 days, 03:13:49
Main I/O
PID PRI NI VIRT RES SHR S CPU%- MEM% TIME+ Command
91266 20 0 2214M 1015M 168M R 100.3 0.4 14h25:41 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91268 20 0 2175M 1002M 161M R 100.3 0.4 14h13:28 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91269 20 0 2168M 997M 159M R 100.3 0.4 14h13:00 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91281 20 0 2147M 996M 159M R 100.3 0.4 14h46:52 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91294 20 0 2173M 1006M 165M R 100.3 0.4 15h01:27 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91302 20 0 2174M 1011M 169M R 100.3 0.4 14h48:10 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91304 20 0 2187M 1008M 164M R 100.3 0.4 14h34:51 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91308 20 0 2163M 1008M 161M R 100.3 0.4 14h22:50 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91312 20 0 2167M 1015M 166M R 100.3 0.4 14h52:05 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91321 20 0 2164M 1016M 174M R 100.3 0.4 14h42:43 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91323 20 0 2182M 1008M 171M R 100.3 0.4 14h55:08 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91325 20 0 2165M 1013M 171M R 100.3 0.4 14h47:04 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91326 20 0 2173M 1024M 180M R 100.3 0.4 14h59:41 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91328 20 0 2148M 1014M 170M R 100.3 0.4 14h52:01 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91267 20 0 2219M 1009M 165M R 99.6 0.4 14h27:31 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91270 20 0 2182M 1003M 166M R 99.6 0.4 14h46:16 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91271 20 0 2178M 1000M 162M R 99.6 0.4 14h42:53 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91272 20 0 2170M 999M 160M R 99.6 0.4 14h29:46 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91273 20 0 2162M 1000M 161M R 99.6 0.4 14h28:06 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91274 20 0 2169M 999M 160M R 99.6 0.4 14h42:35 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91275 20 0 2154M 990M 156M R 99.6 0.4 14h38:32 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91276 20 0 2161M 993M 156M R 99.6 0.4 14h27:35 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91277 20 0 2147M 994M 157M R 99.6 0.4 14h22:26 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91279 20 0 2165M 997M 159M R 99.6 0.4 14h56:43 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91280 20 0 2150M 994M 158M R 99.6 0.4 14h51:00 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91283 20 0 2182M 1003M 159M R 99.6 0.4 14h30:01 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91284 20 0 2144M 990M 156M R 99.6 0.4 14h21:07 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91285 20 0 2164M 998M 157M R 99.6 0.4 14h17:54 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91286 20 0 2173M 998M 163M R 99.6 0.4 14h48:42 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91287 20 0 2163M 999M 161M R 99.6 0.4 14h51:07 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91288 20 0 2159M 1000M 161M R 99.6 0.4 14h33:19 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
91290 20 0 2200M 1011M 166M R 99.6 0.4 14h39:09 /ccad/stack/23.09/env/apps/linux-rocky8-zen2/gcc-12.3.0/quantum-espresso-7.2-3qtnqojacrtgsojs5qptvtjn3txc7hpp/bin/
F1Help F2Setup F3SearchF4FilterF5Tree F6SortByF7Nice -F8Nice +F9Kill F10Quit
Si no está así, está mal.

También las GPUs.

INFRAESTRUCTURA

Datacenter

Refrigeracion

Aire / Agua / Inmersión

Alimentación Eléctrica

Consumo eléctrico

Storage

Filesystem Paralelo

Interconexión

Infiniband
Slingshot
Ethernet + RDMA

Refrigeración Direct Water Cooling

99% heat capture at 40°C

CPUs | GPUs | DIMMs | Fuentes

CLEMENTINA XXI

74 NODOS GPU + 6 NODOS CPU

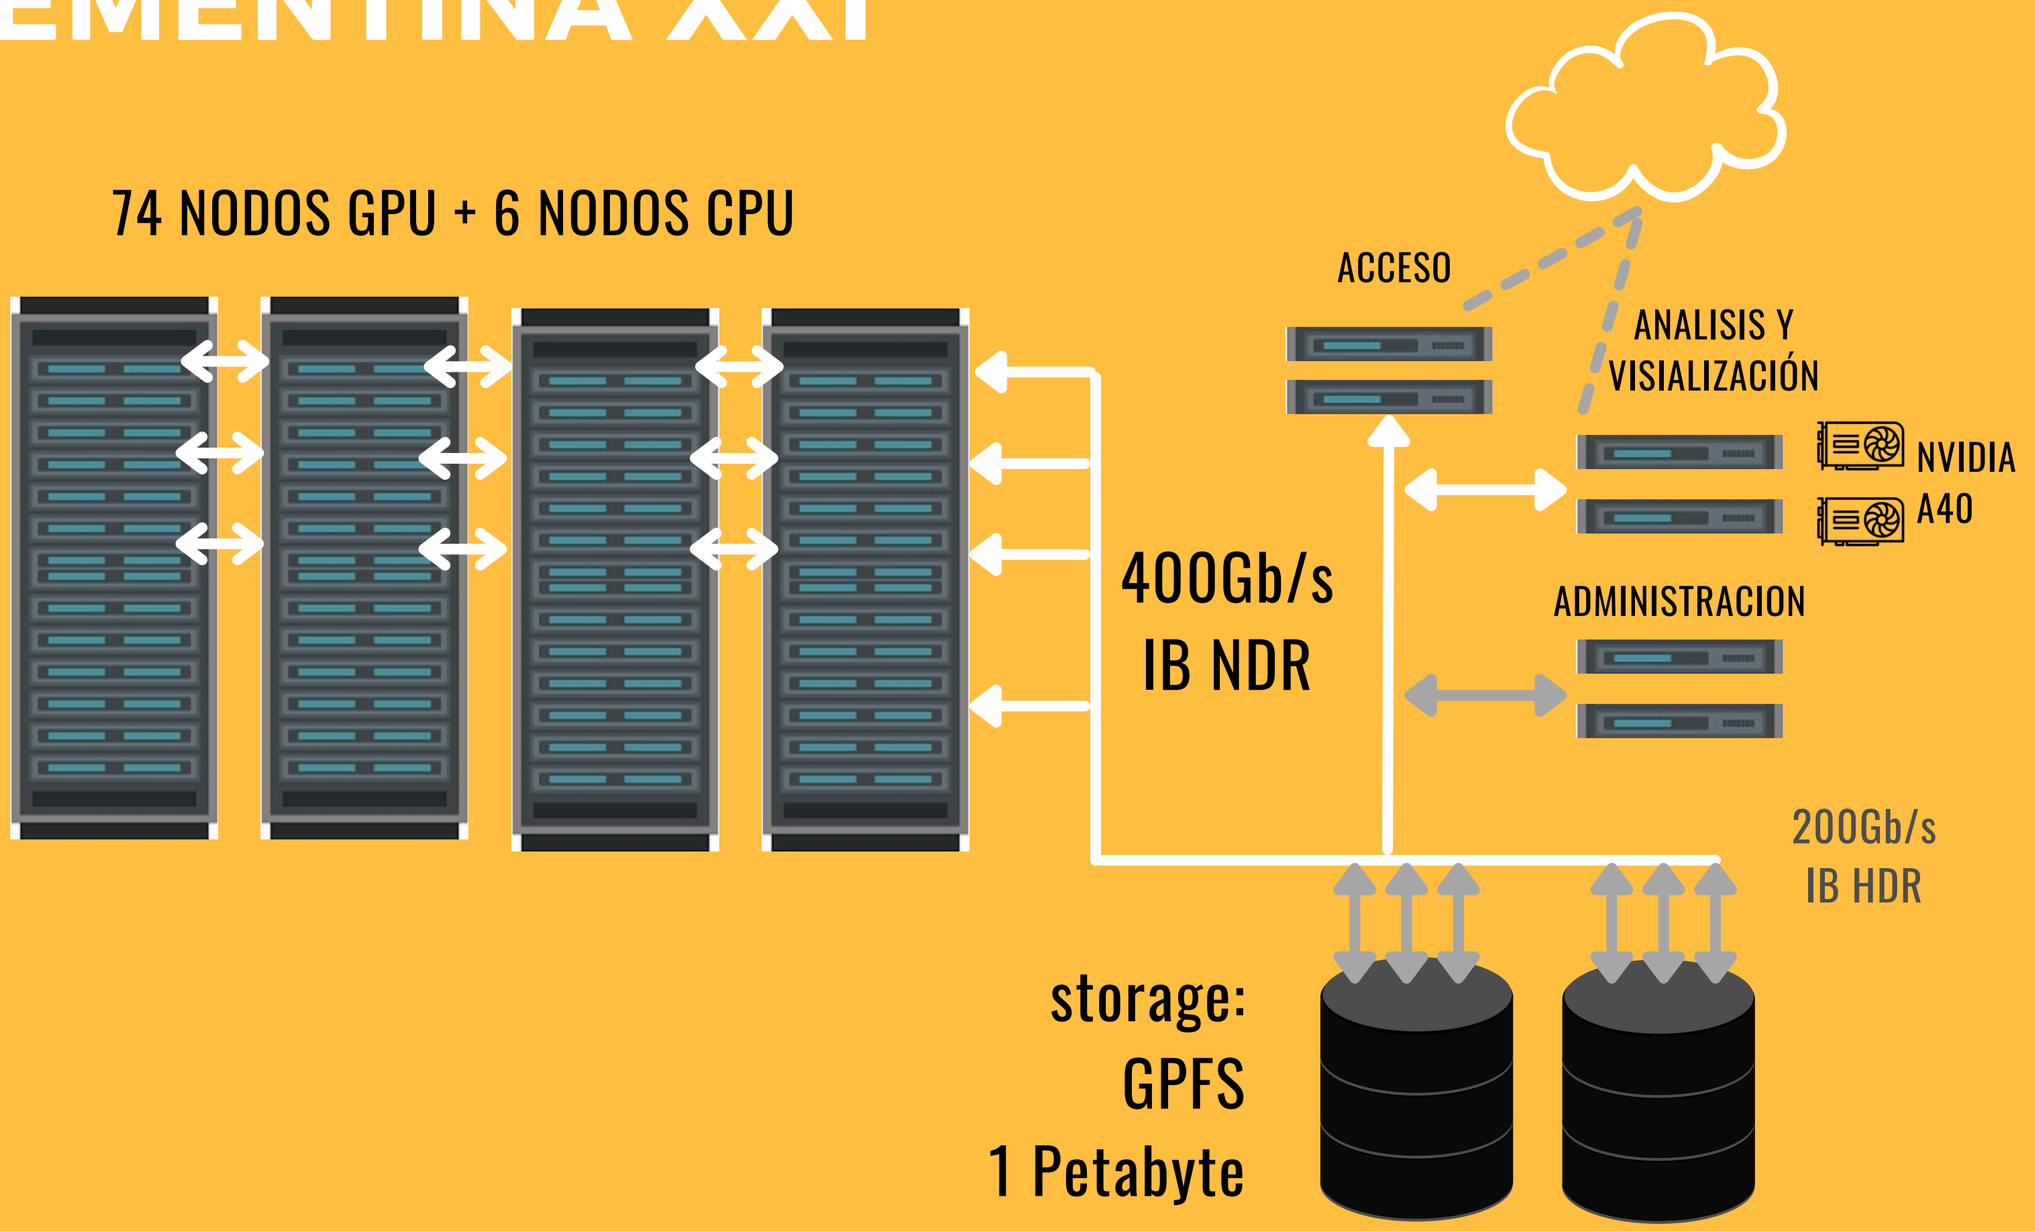

NODO

RED INFINIBAND
CONNECTX-7 NDR
QSFP400

4 X

INTEL GPU MAX 1550
(PONTE VECCHIO)
128 GB POR GPU, 600W TDP

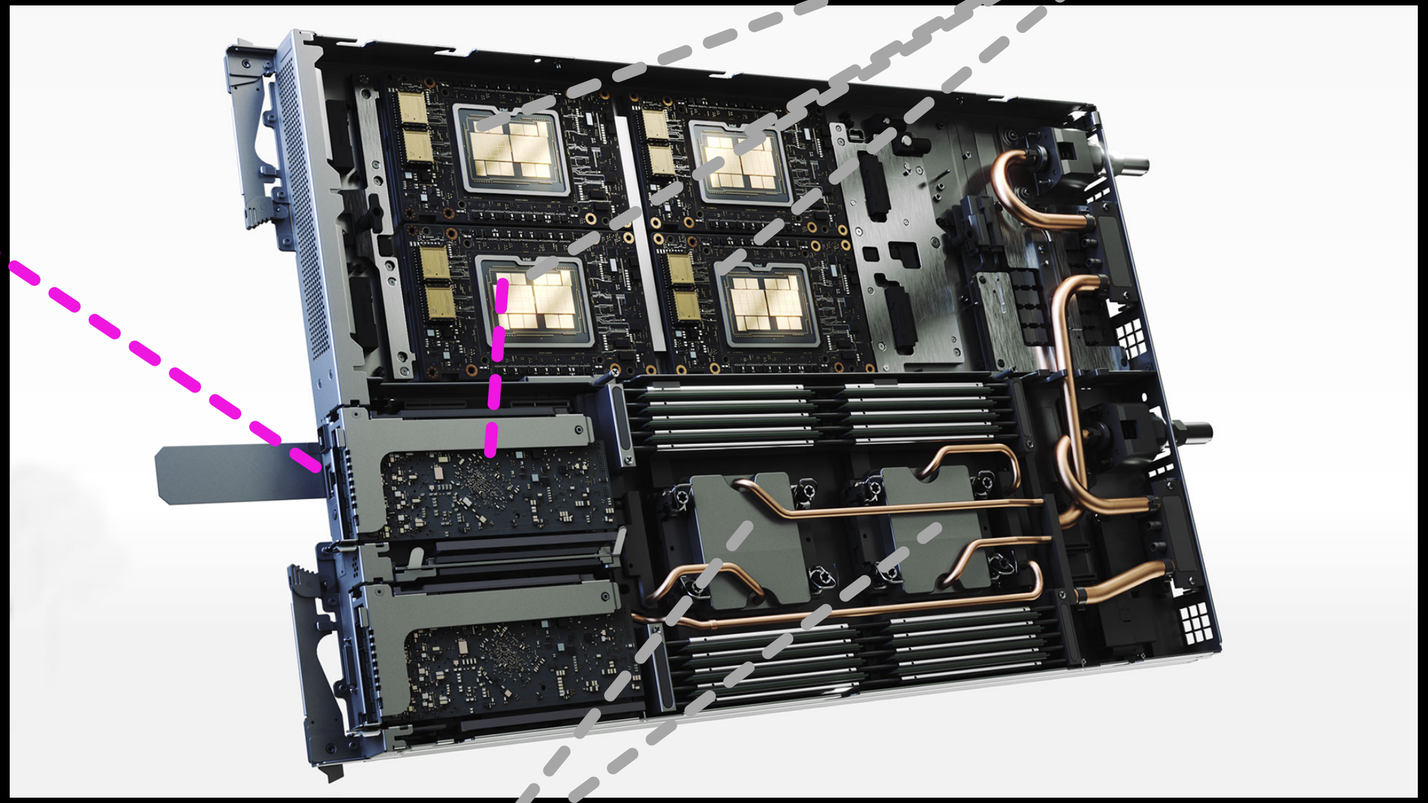

2 X

INTEL XEON MAX 9462 @ 2,7GHZ (SAPHIRE RAPIDS)
64 GB HBM + 256 GB RAM DDR5 4800 MT/S

Intel GPU MAX 1550 (PVC)

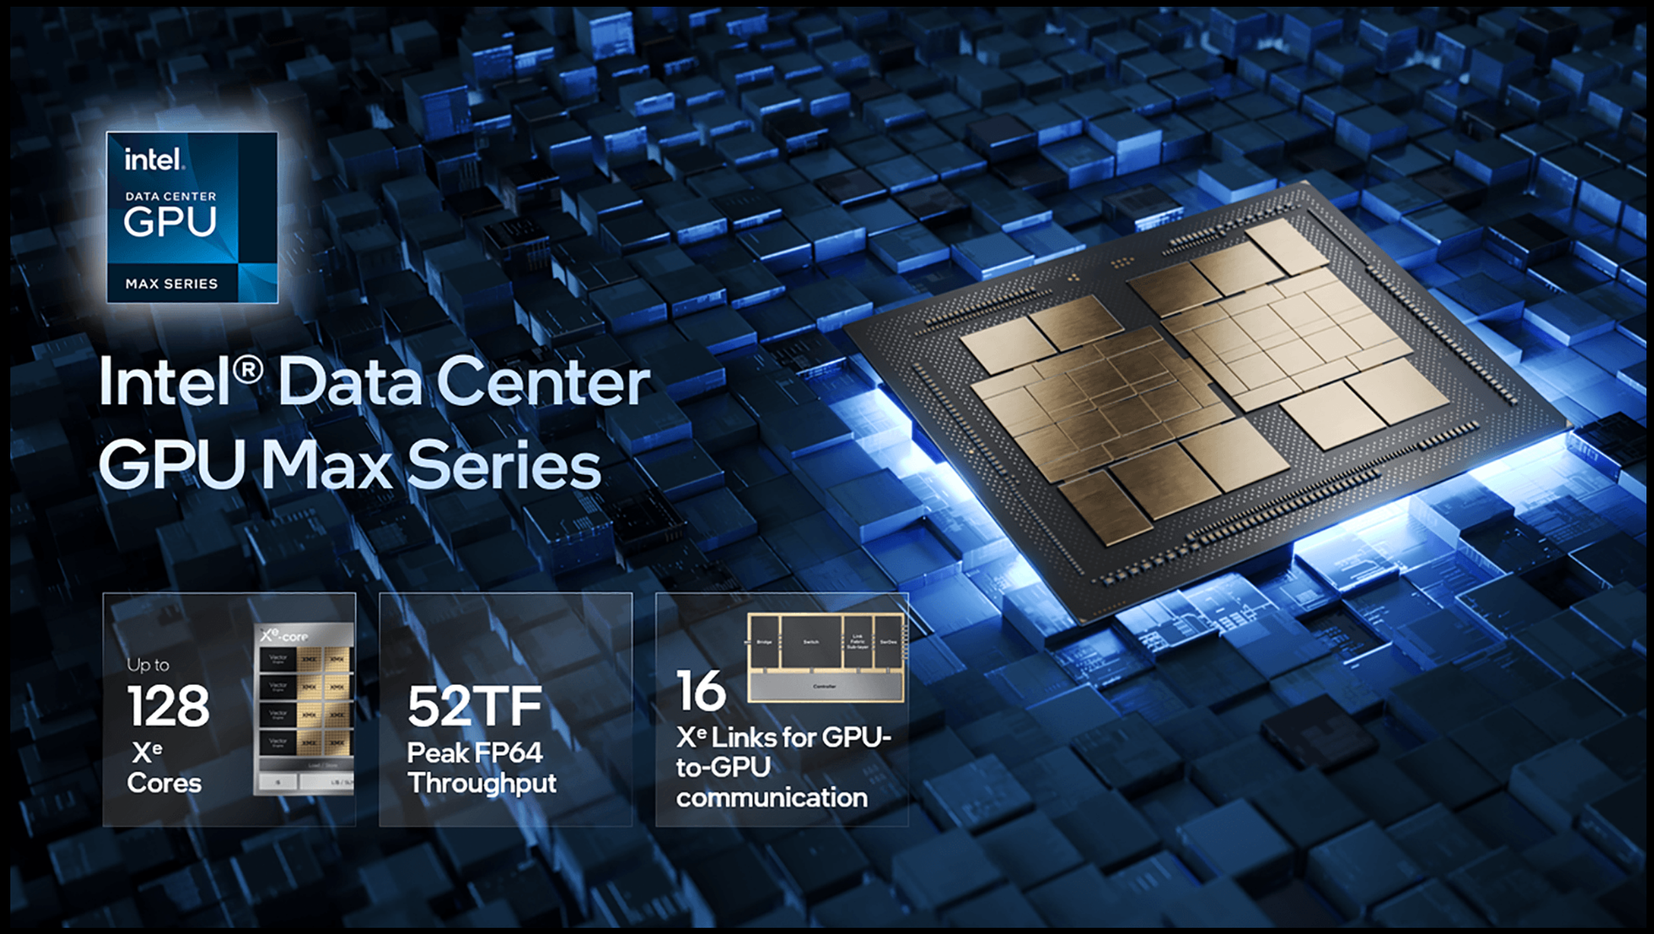

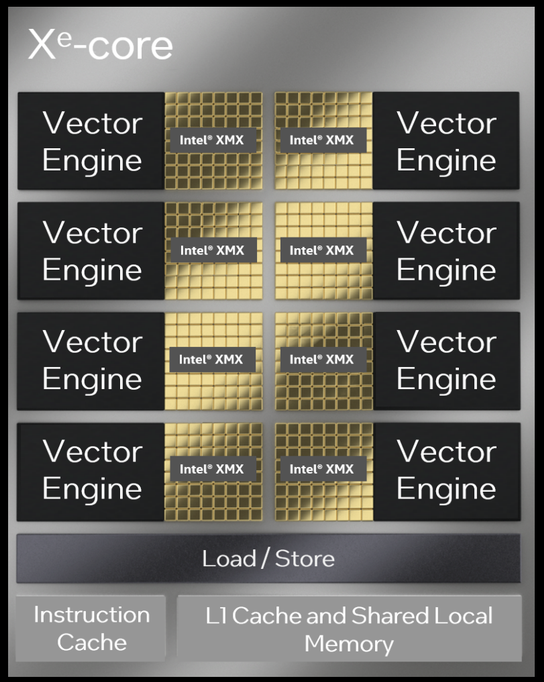

128 Xe cores

128GB HBM2E
PCI-EXPRESS 5.0 X16
TDP/TBP 600 W

4 X

oneAPI

SYCL™

Intel Xeon Scalable 9462 (SPR)

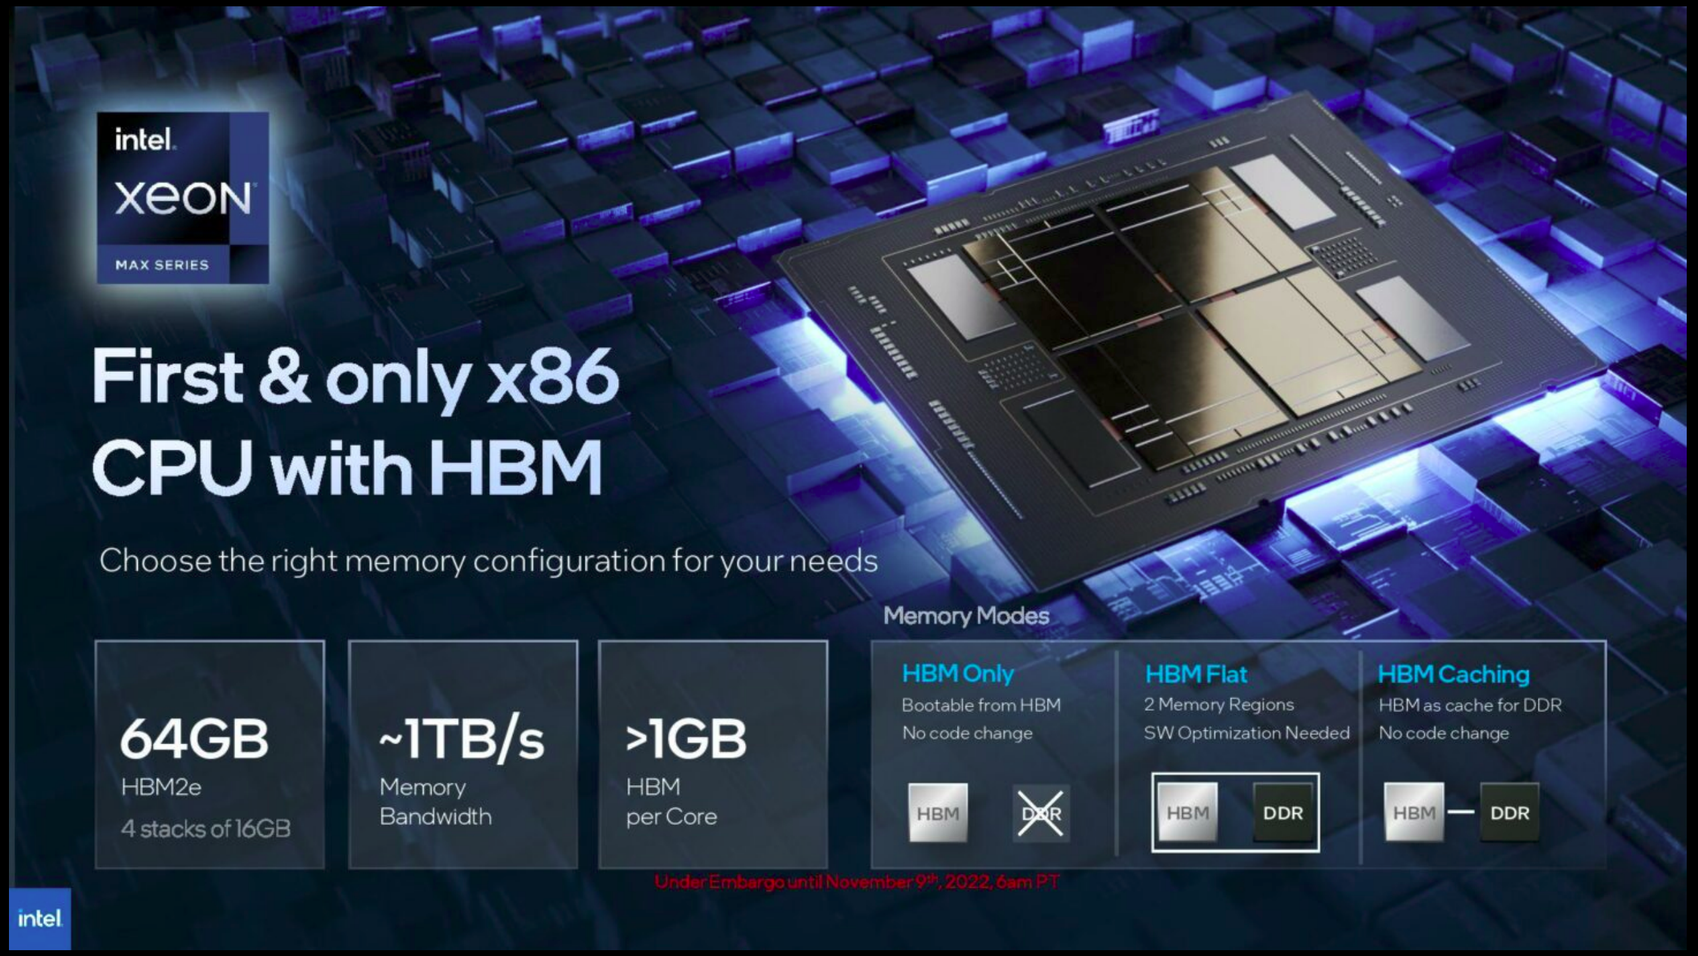

32 CORES GOLDEN COVE
64 GB HBM
256 GB DE RAM (8 MODULOS)
TDP/TBP 350 W

2 X

Cómo se usa?

EL USO ES NO INTERACTIVO

01

Conexion SSH

Consola remota modo texto

02

Sistema de colas

Scripts no intetractivos con:

- Software propio
- Aplicaciones específicas

03

Análisis de resulatados

Visualización y transferencia de archivos

Top500: FLOPS

- **Rpeak:**
  - **Máxima performance pico teórica**

- **Rmax:**
  - **Máxima performance efectiva**
    - **Mejor resultado ejecutando HPL**

OPs en FP64 (double de C)

Top 5

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	**Frontier** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,699,904	1,194.00	1,679.82	22,703
2	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	**LUMI** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	**Leonardo** - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,824,768	238.70	304.47	7,404
5	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096

Top500: llegaremos?

74 nodos
x 4 GPUs
x 52 TFLOPS peak
= 15.329 TFLOPS

15.329 TFLOPS x 0,50%
+ 400 TFLOPS CPU
= ~8,06 PFLOPS

79	**LANTA** - HPE Cray EX235n, AMD EPYC 7713 64C 2GHz, NVIDIA A100 40GB, Slingshot-11, HPE NSTDA Supercomputer Center (ThaiSC) Thailand	87,296	8.15	13.77	310
80	**TSUBAME3.0** - SGI ICE XA, IP139-SXM2, Xeon E5-2680v4 14C 2.4GHz, Intel Omni-Path, NVIDIA Tesla P100 SXM2, HPE GSIC Center, Tokyo Institute of Technology Japan	135,828	8.12	12.13	792
81	**Roxy** - Apollo 2000, Xeon Gold 6248 20C 2.5GHz, Infiniband EDR, HPE Government United States	174,720	8.12	13.98	
82	**HoreKa-Green** - ThinkSystem SD650 V2, Xeon Platinum 8368 38C 2.4GHz, NVIDIA A100 40GB, Mellanox HDR Infiniband, Lenovo Karlsruher Institut für Technologie (KIT) Germany	90,932	8.03	15.01	359
83	**Narwhal** - HPE Cray EX, AMD EPYC 7H12 64C 2.6GHz, Slingshot-10, HPE Navy DSRC United States	276,480	8.02	11.50	
84	**Plasma Simulator** - SX-Aurora TSUBASA A412-8, Vector Engine Type10AE 8C 1.58GHz, Infiniband HDR 200, NEC National Institute for Fusion Science (NIFS) Japan	34,560	7.89	10.51	1,550

Recurso
+ Modelado numérico
+ Optimización
= Valor agregado + ROI

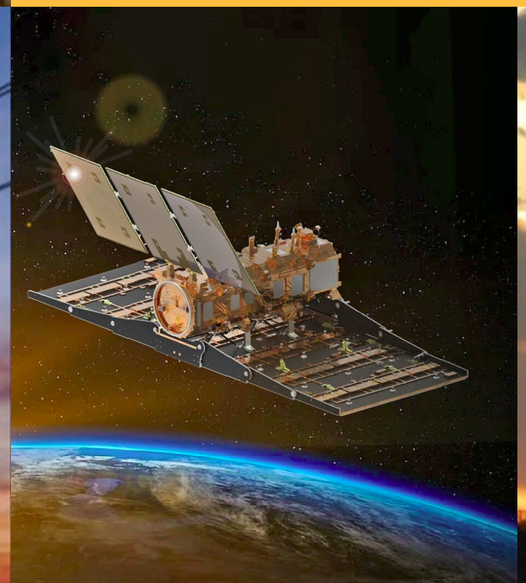

Qué esperar en los próximos años?

- Tener equipos ES soberanía.
- Los equipos pasan, la gente queda.
- El HPC va a seguir creciendo.
- Demanda global de personal capacitado.
- Academia e Industria.

Más de 60 años para recuperar potencia de cálculo competitiva en Argentina.

Clementina ... clusters ... Clementina XXI ...

Preguntas?

Marcos Mazzini | CPA conicet / CCAD (UNC)