



*Breaking silos architecting Multi-Protocol Data Services for Modern HPC & AI
at University of Pisa*



Locarno, April 2026 - HPC-AI Advisory Council Swiss Conference

Roberto Dognini, VAST Data
Maurizio Davini, University of Pisa


HPC & AI Together and ... AI moving to Production



**HPC Users
are still there**



**AI Models are
Improving Rapidly**



**Inference requires
High Performance
and reliable data
services**

Deep Agents AIQ

Access Control

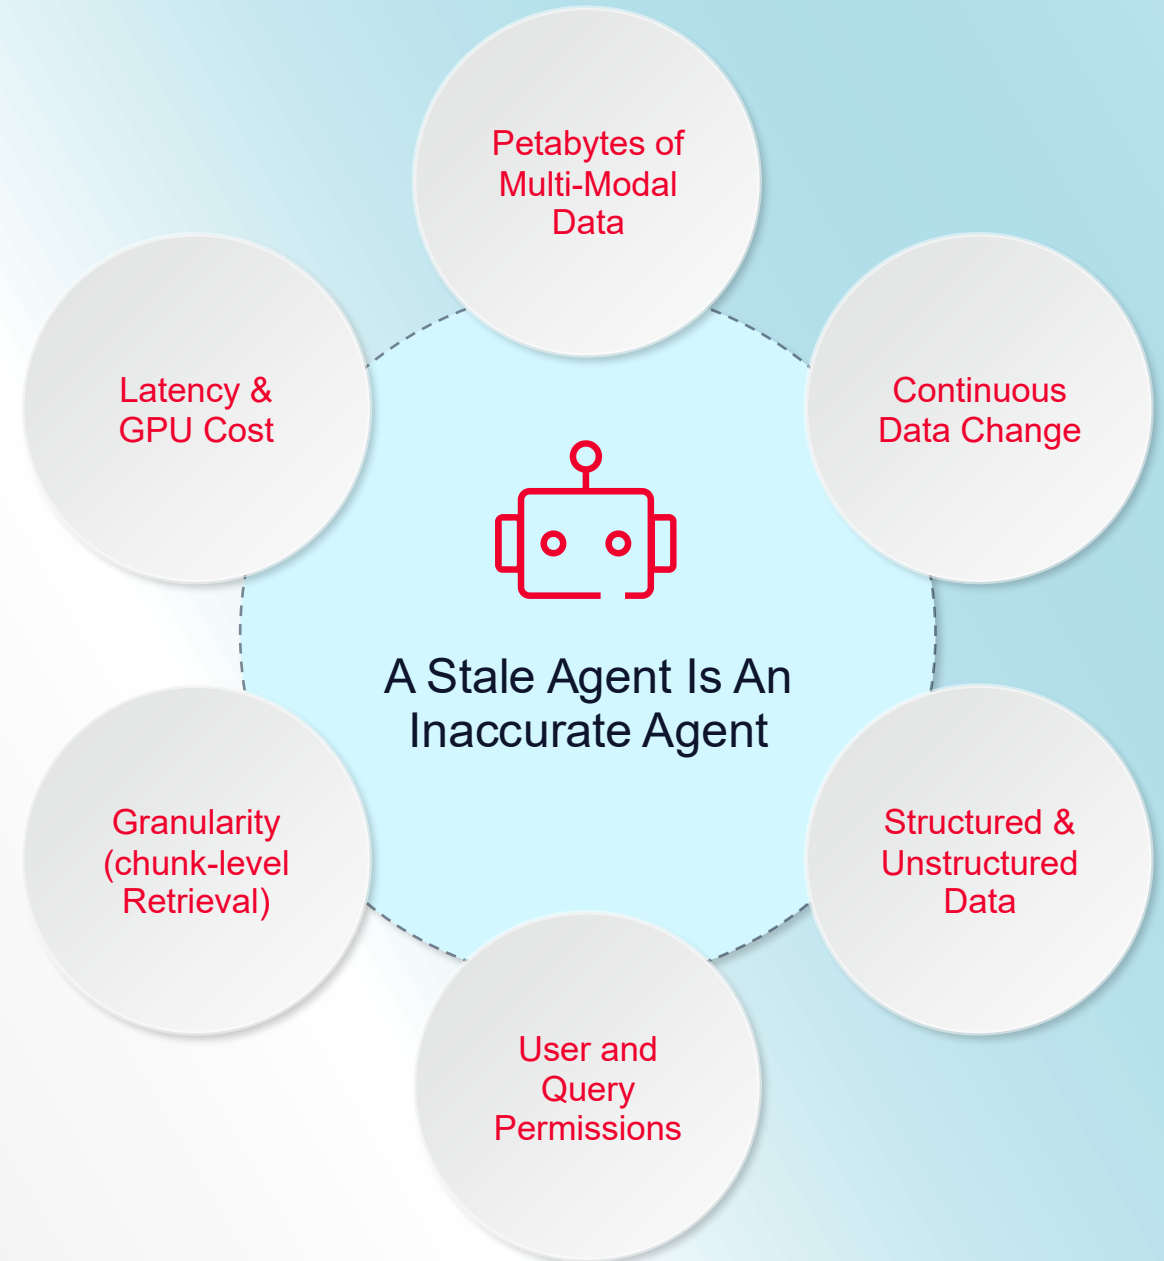
Access to Real-Time Data

Ability to Act with Tools

**Agents are
approaching**


AI Models Are Not AI's Limiting Success Factor

Inference at scale is a data systems problem



The Challenge With Scaling Vectors


Fast and Expensive



1T Vectors | 768 Dimensions = \$40M of HW

VS.

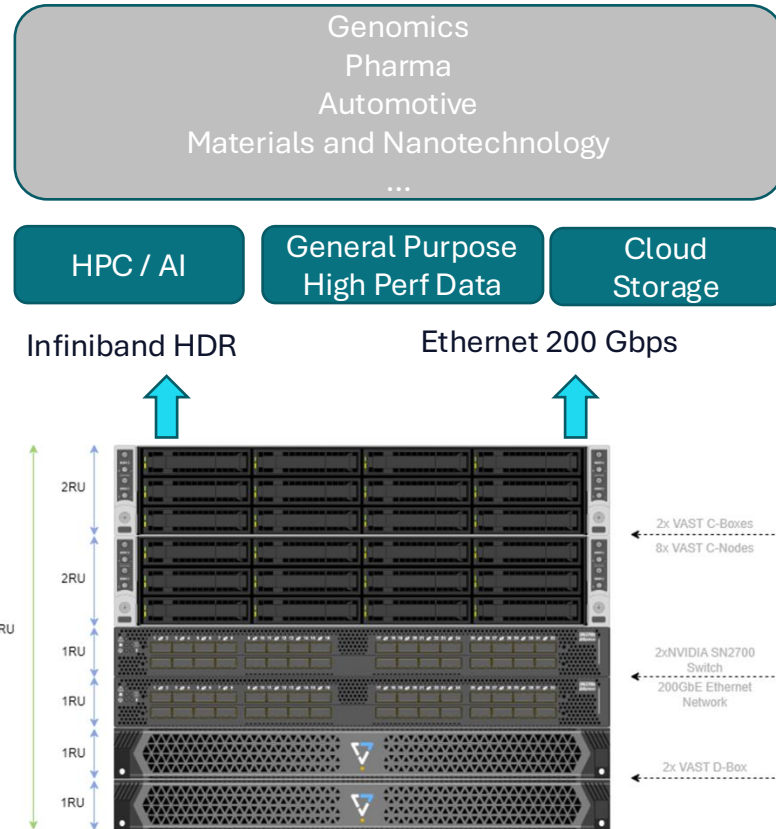
Slow and Cheap



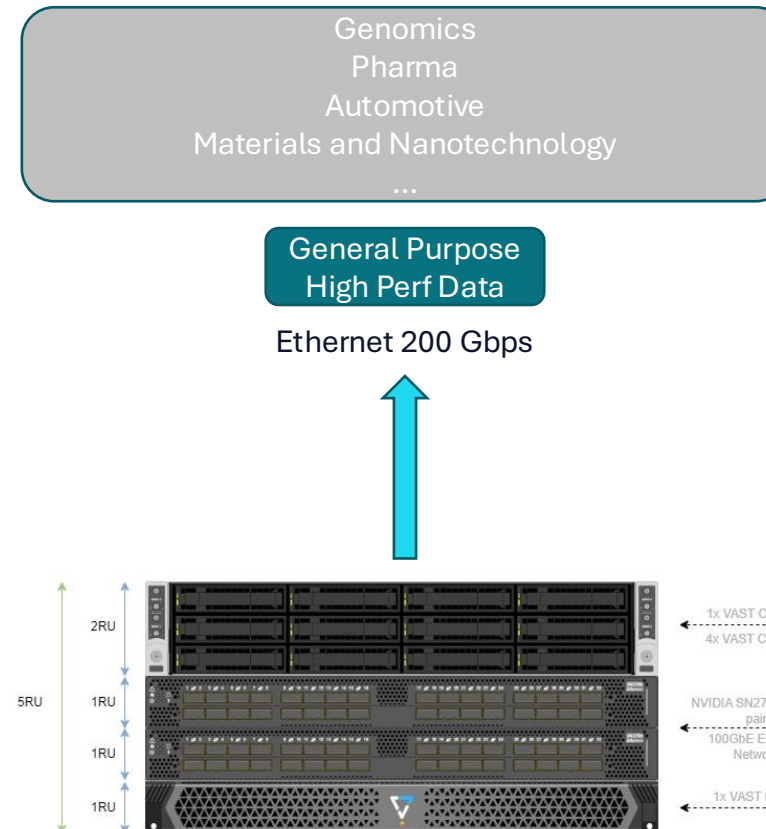
Wasted GPU Resources

From HPC Scratch to High Performance Multiprotocol Reliable Data Service

Production SITE



DR SITE



552.7 TB

Available HW Usable Capacity

690.9 TB

Available HW Effective Capacity
(with 1.25: 1 Data Reduction)

100 GB/s

Read Bandwidth

28 GB/s

Write Bandwidth Max

550.000

Random File Operations (4k)

238.8 TB

Available HW Usable Capacity

298.5 TB

Available HW Effective Capacity
(with 1.25: 1 Data Reduction)

50 GB/s

Read Bandwidth

14 GB/s

Write Bandwidth Max

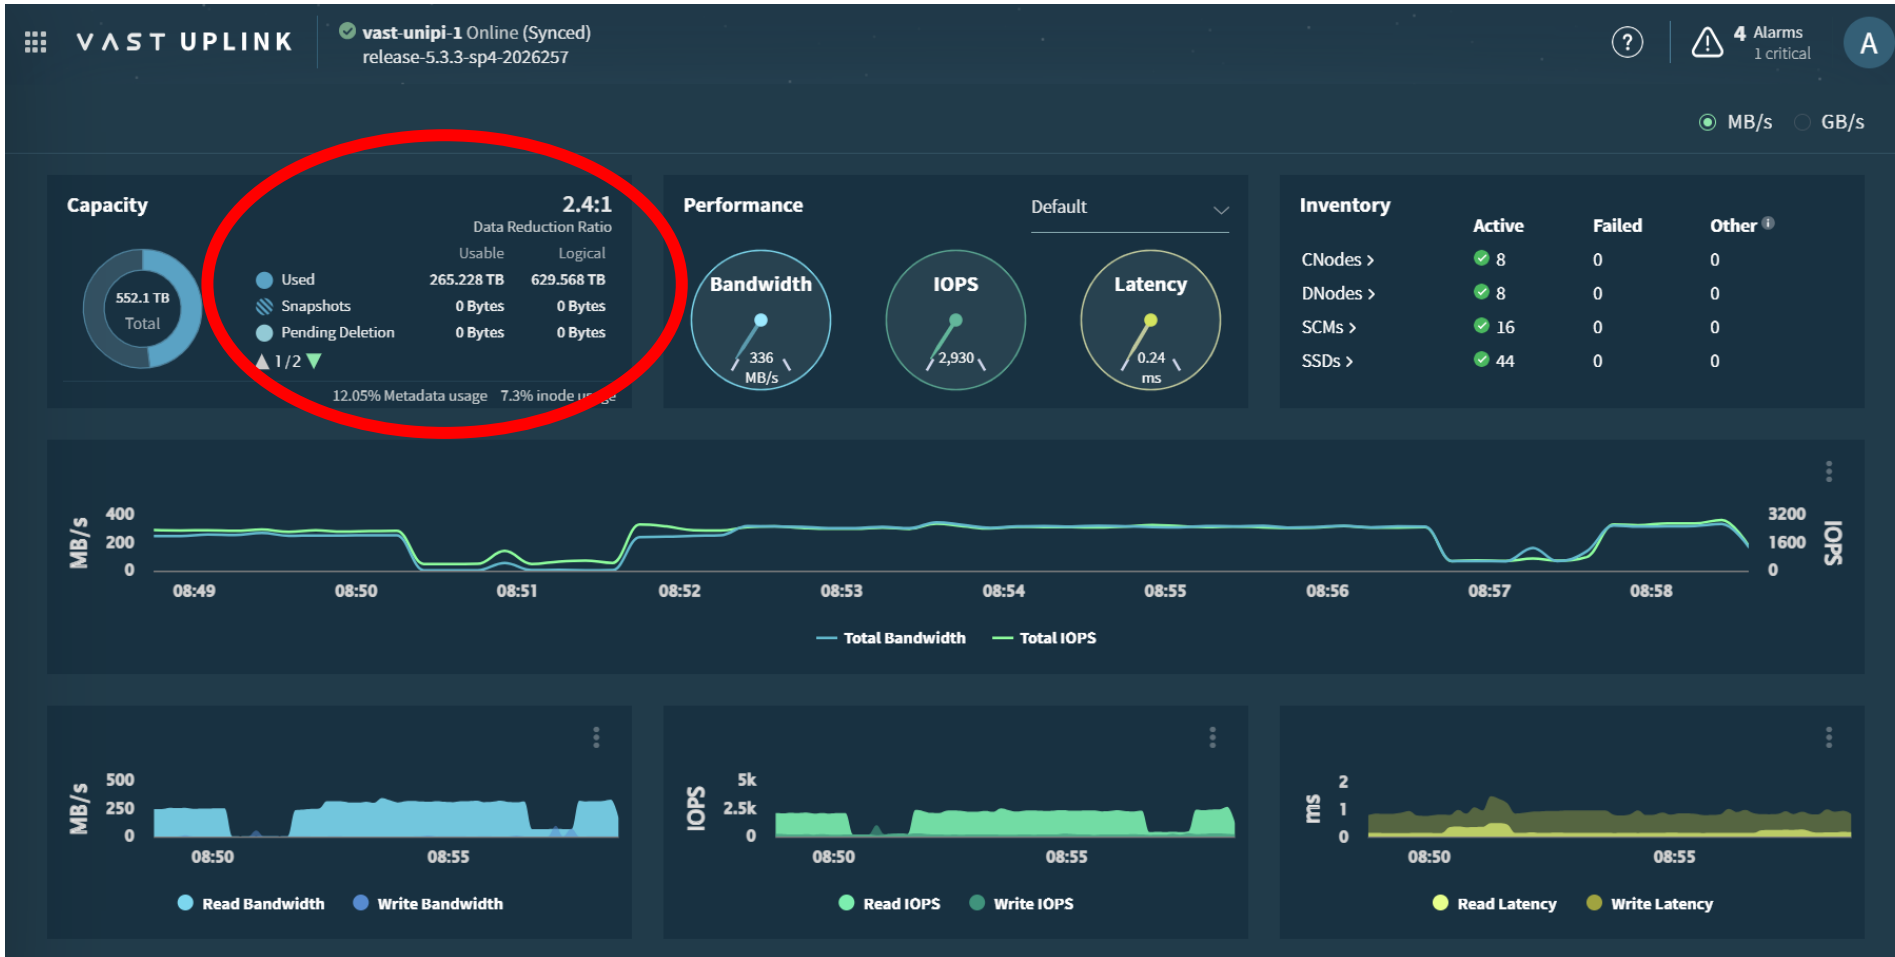
330.000

Random File Operations (4k)

Disaster Recovery - Async Replica - RPO 15min

High performance NFS, S3, SMB, CSI and COSI with Multitenancy - QoS - Auditing - Integrated Snapshot data protection - SQL Metadata Catalog

An interesting positive surprise after a few months...




552.7 TB

Available HW Usable Capacity

690.9 TB

Available HW Effective Capacity
(with 1.25: 1 Data Reduction)



552.7 TB

Available HW Usable Capacity

1324 TB 😊

Available HW Effective Capacity
(with 2.4: 1 Data Reduction)

Conclusion

- Vast Cluster is the central storage platform for Campus Network
- Mixed Use : backplane for cloud storage, HPC/AI storage con HPC cluster and AI systems
- Storage backplane connections to Genomics systems and simulations
- Storage backplane for satellite images acquisition

- Key storage for Google Cloud Storage re-migration

Work in progress – Additional Services we will offer to our users

- VAST Data integrated Columnar Database for high performance Advanced Analytics functionalities
- VAST Data extreme high scalable and easy to use integrated Vector DB



UNIVERSITÀ
DI PISA



VAST

Thank you