

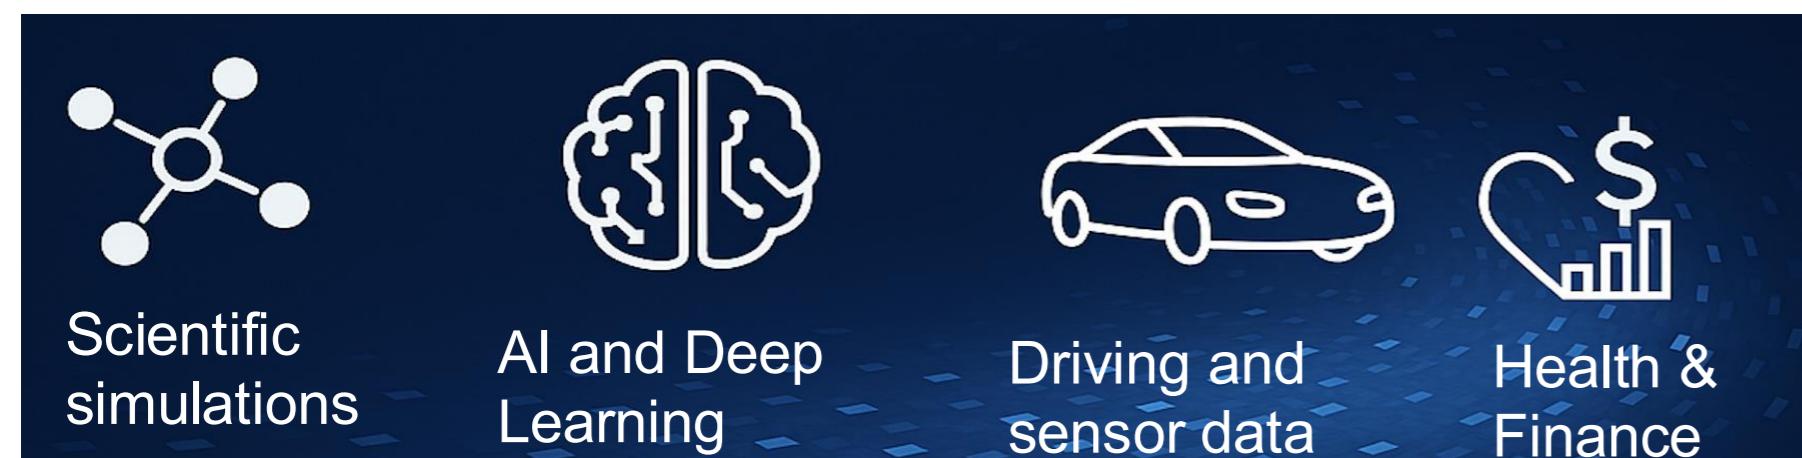
Accelerating AI Workloads with GPUs and RDMA-Enabled S3 and NFS Data Access

Madhu Thorat

Anandhu K

IBM

1. AI and HPC applications generate massive amounts of data !

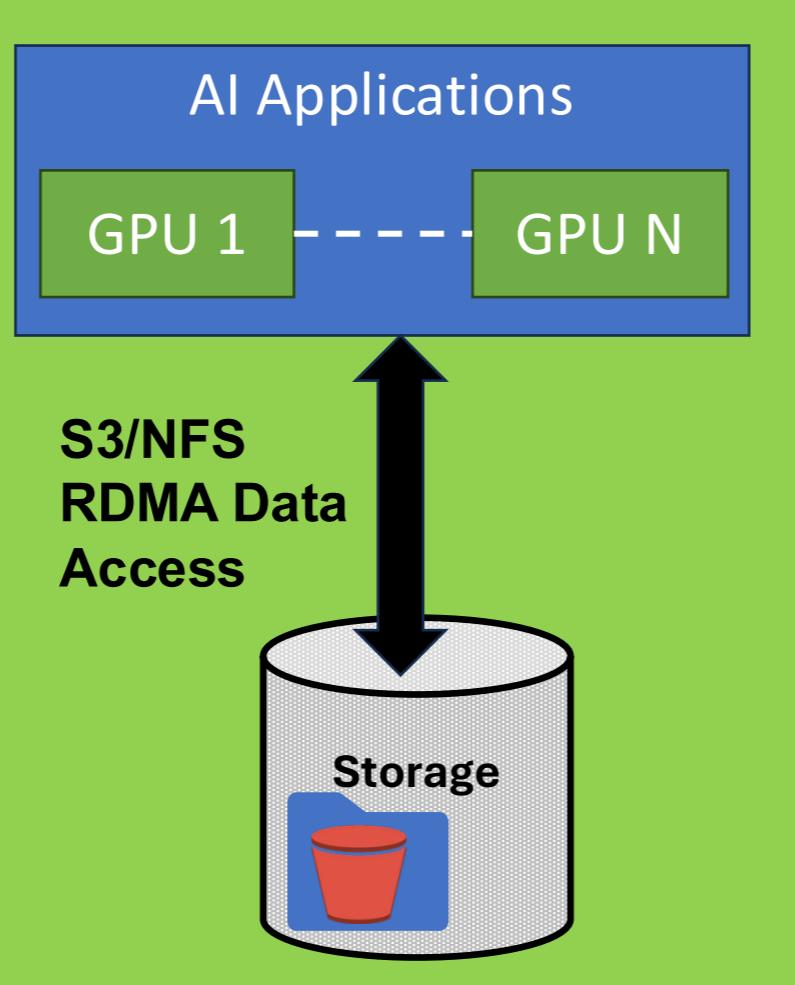


Need & Benefits of RDMA access for High-Performance Storage

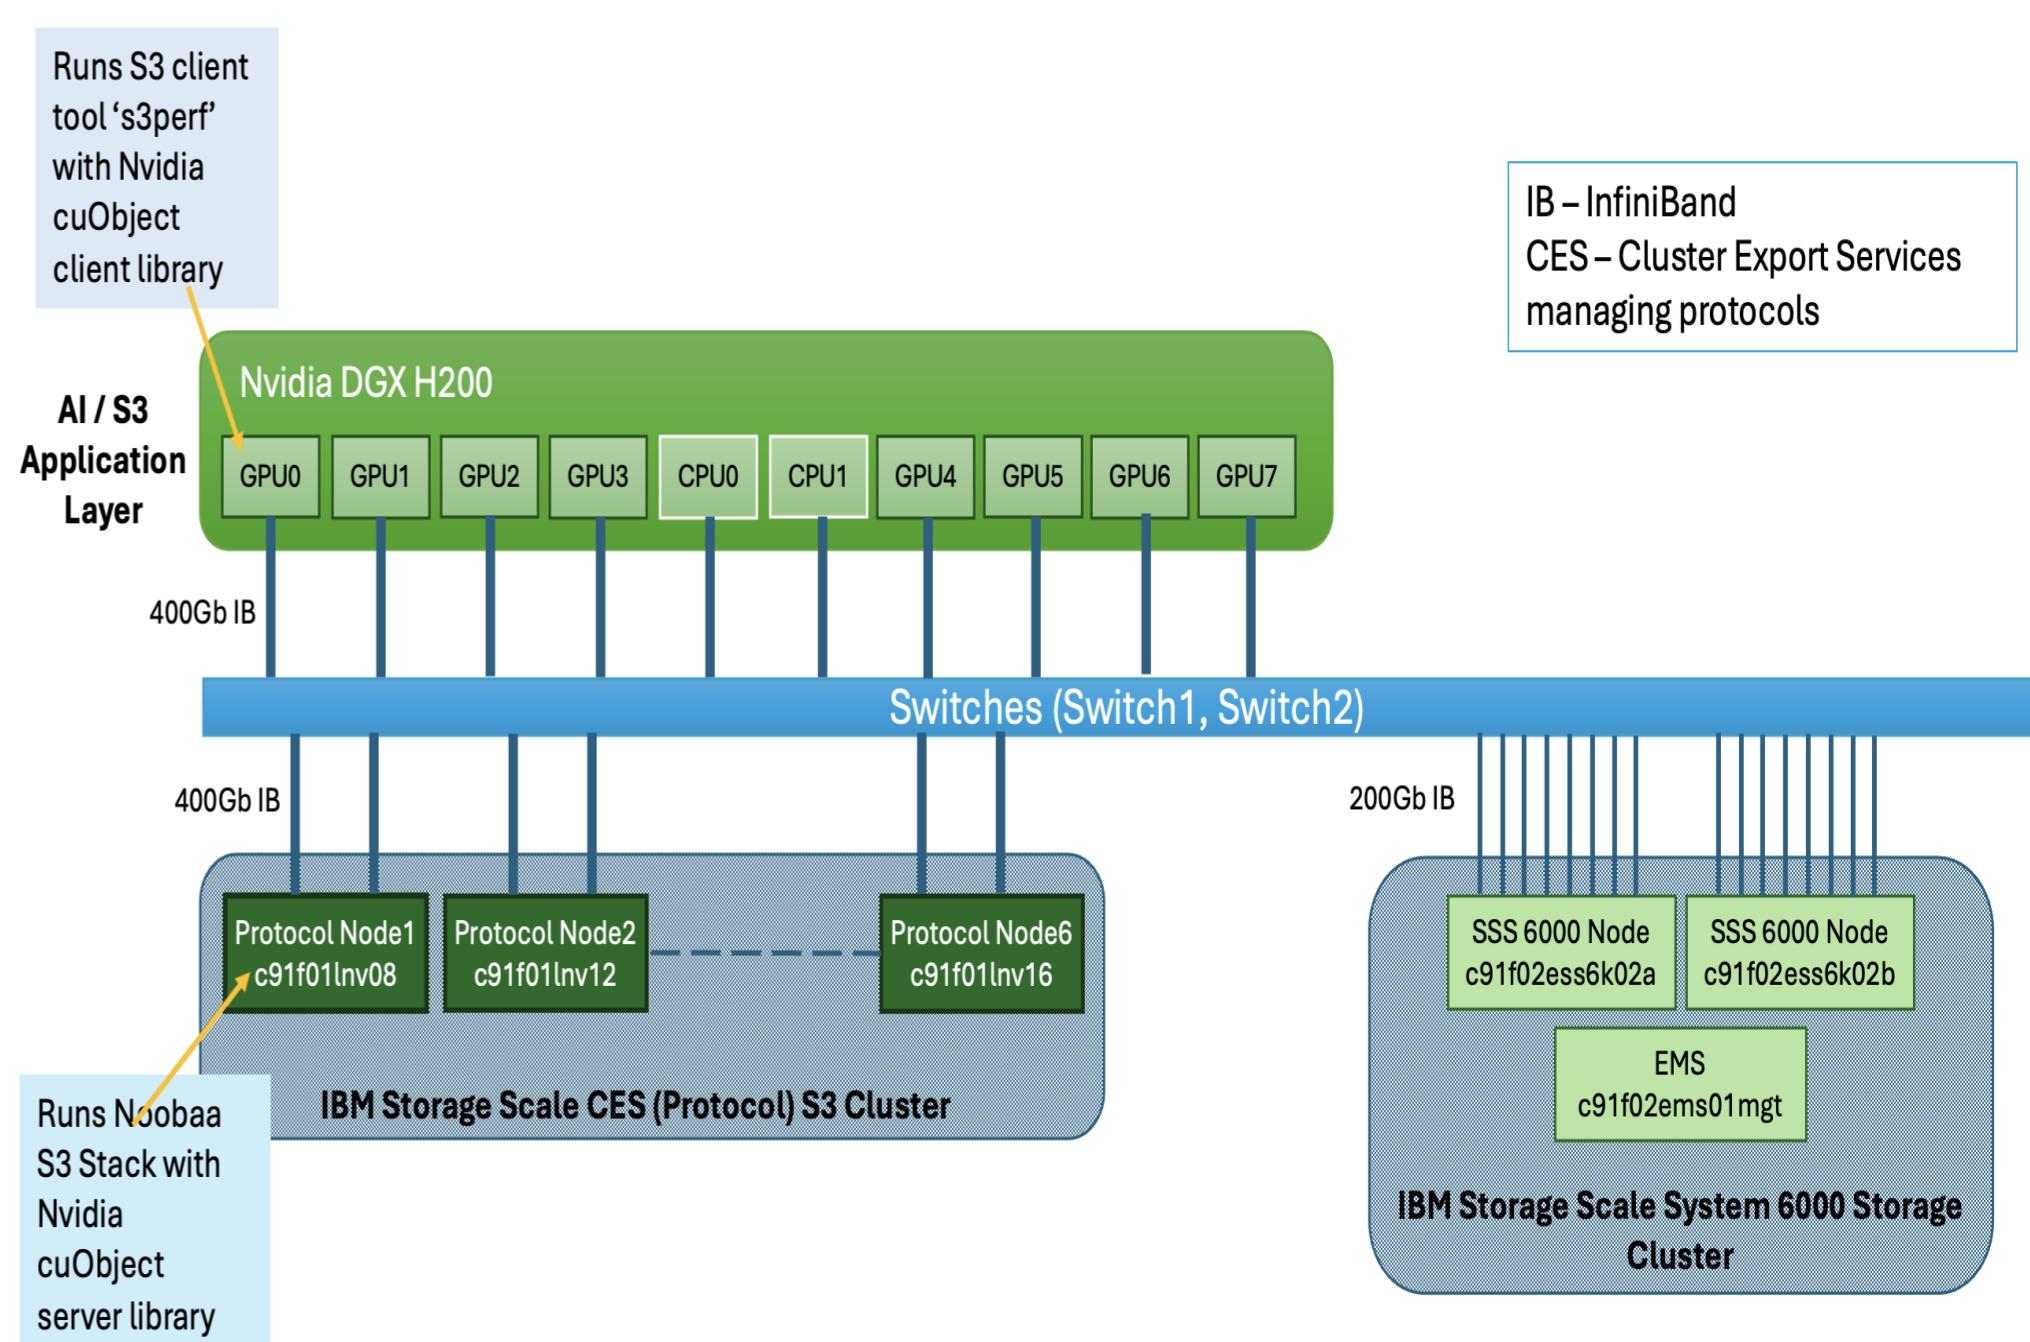
AI applications processing uses GPUs and requires fast data access through protocols like S3, NFS to access backend storage.

Benefits of RDMA Data Access to Storage

- Reduced CPU resource overhead by enabling memory to memory data movement and zero-copy transfers.
- Higher throughput with lower latency.
- Control path and data path can be separated which helps boosts the AI data access performance.

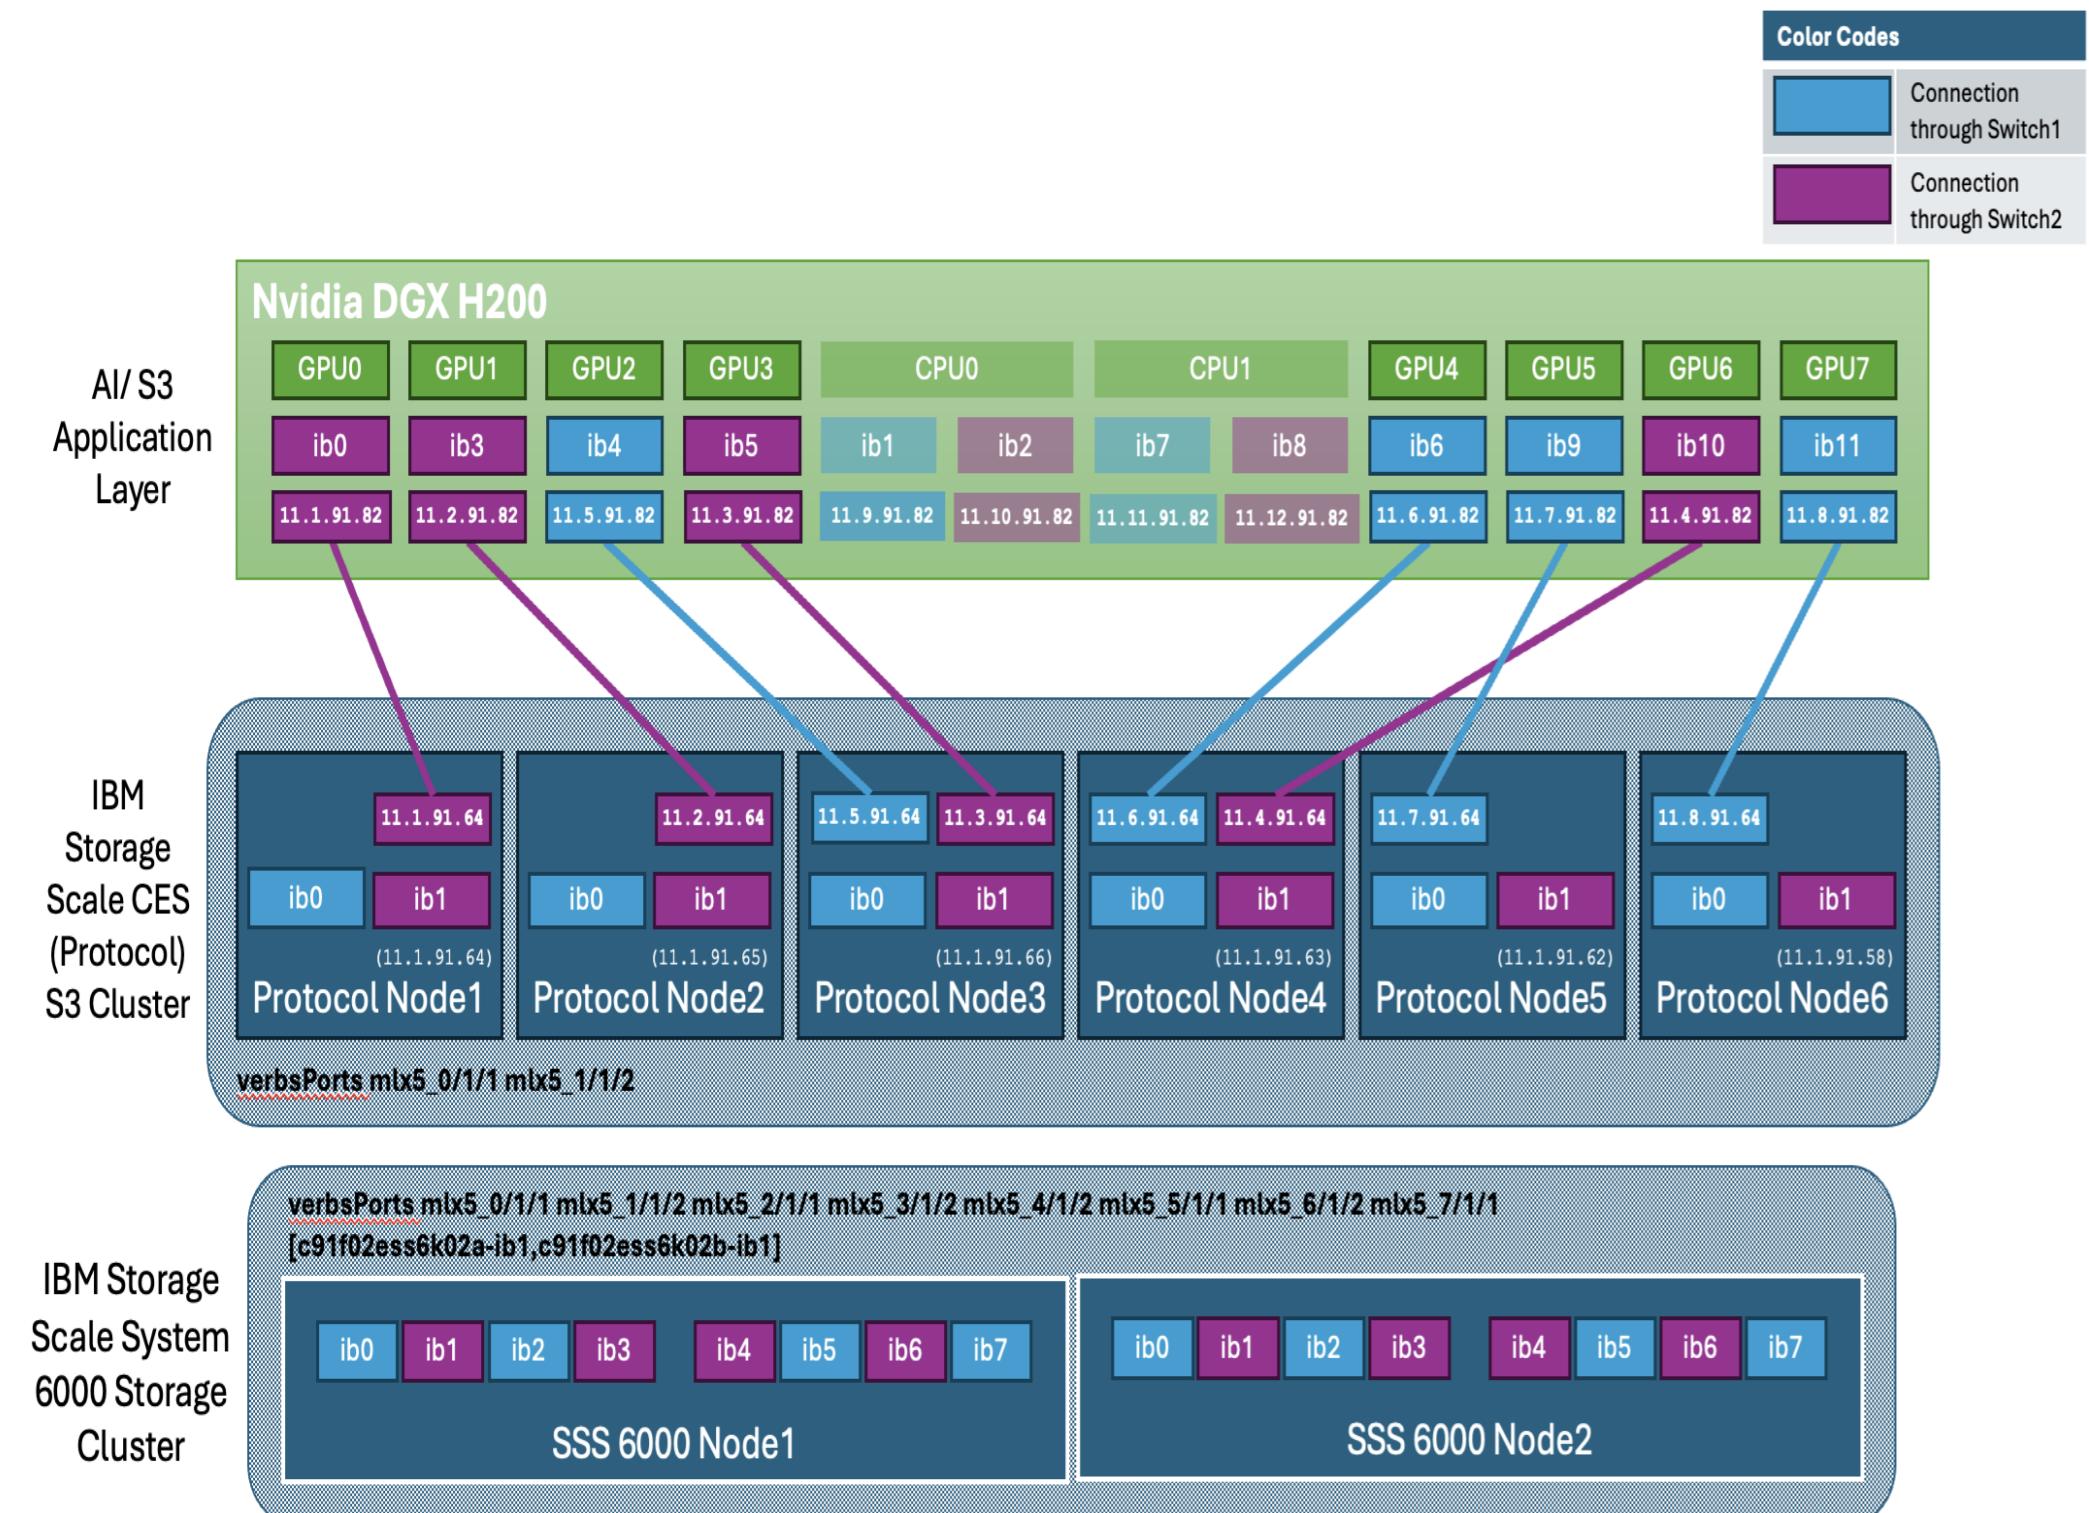


2. Method: Testbed - Lab Infrastructure Setup for measuring performance for RDMA-enabled S3 protocol data access

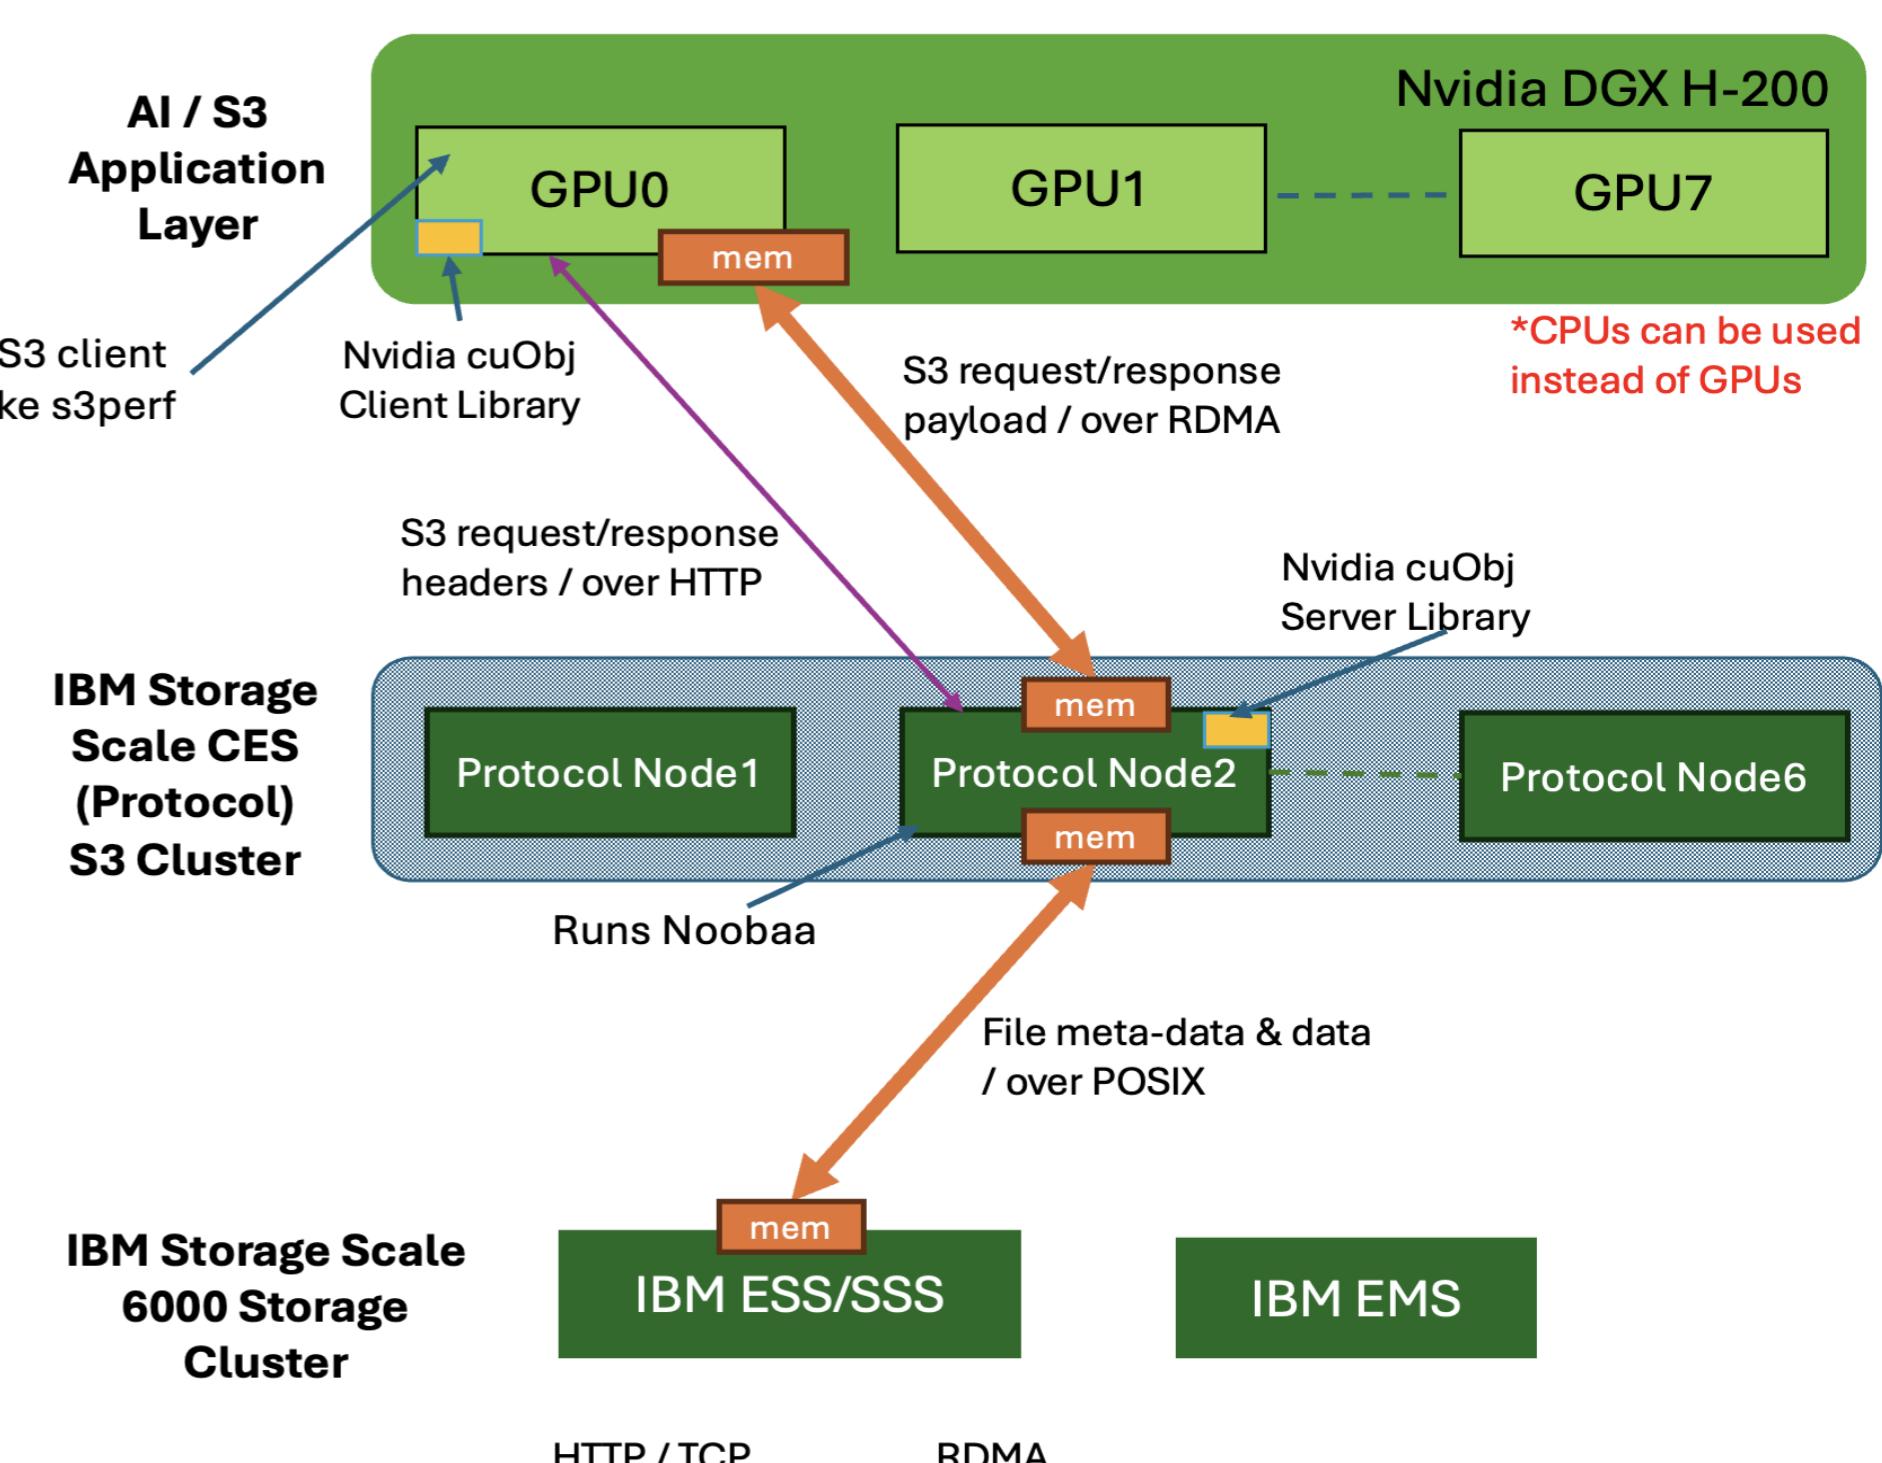


* Only IB connections are shown, but TCP/IP connections would be used for cluster management and for handling S3 request/response headers

3. Method: Testbed - RDMA Network Topology



4. Performance Benchmark Reference Data Flow Architecture for RDMA-enabled S3 protocol data access



5. Results: Performance Benchmark Results for S3 over RDMA with GPUs and IBM Storage Scale System 6000

Object Size	READ Bandwidth	WRITE Bandwidth
1 GB	~297 GB/s	~153 GB/s
128 MB	~308 GB/s	~149 GB/s
64 MB	~308 GB/s	~140 GB/s

Small Objects:

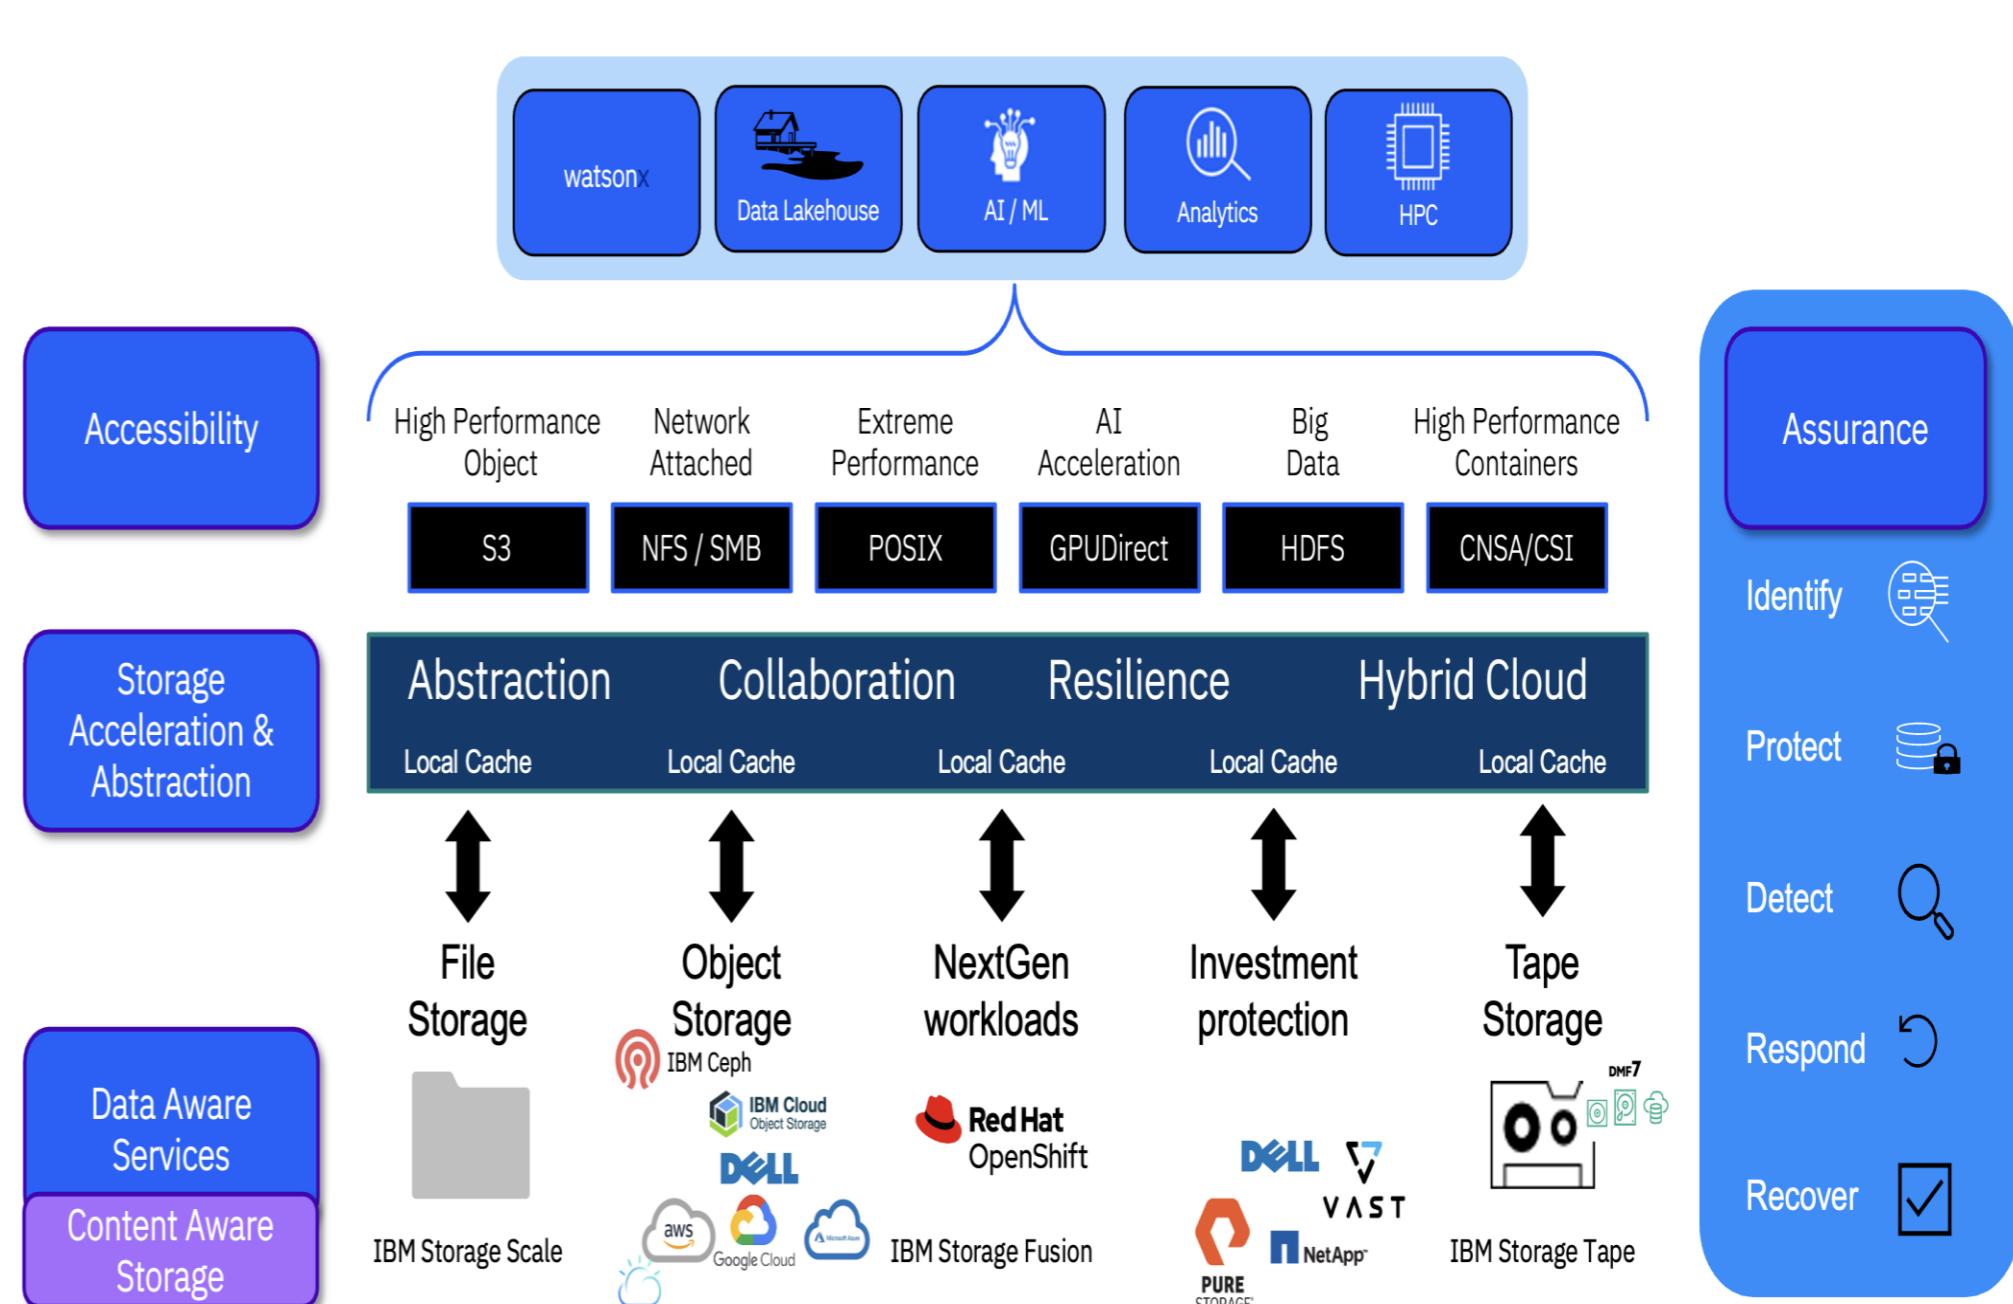
Object Size	READ Ops/sec	WRITE Ops/sec
16 MB	~17K ops/s	~4.6K ops/s
1 MB	~23.5K ops/s	~4.6K ops/s

* Performance was measured with single IBM Storage Scale 6000 cluster with 3 nodes.

Performance with POSIX for 1GB files: ~320 GB/s read, ~155 GB/s write

Conclusion: Despite protocol overhead, S3 over RDMA achieved 90%+ of POSIX performance, showing its suitability for high-throughput AI workloads which prefer S3 access. Similarly, NFS over RDMA can provide high-performance data access for AI workloads

6. Global Data Platform to provide Multi-site, Large Scale Enterprise Infrastructure for AI applications based on IBM Storage Scale supporting GDS based access



References

References:

- [1] IBM Storage Scale Document <https://www.ibm.com/docs/en/storage-scale>
- [2] IBM Storage Scale S3 Support Document <https://www.ibm.com/docs/en/storage-scale/6.0.0?topic=reference-s3-protocol>