EMC In HPC – The Journey So Far, The Road Ahead – With Dell!

Percy Tzelnic, SVP, EMC Fellow

EMC Investment In HPC R&D And Advanced Technology

- Investment in HPC is challenging for all vendors:
  - **Supercomputing** – accelerating needs for performance per watt, per dollar, and per application, vs. extensive investment needs in R&D, with lumpy revenues, in a slow growing market
  - **Commercial HPC** and large-scale data analytics markets (faster growing, more predictable) – also need large investments in technology, justified through a clear vision of how technology will trickle down into enterprise
- EMC has not been traditionally a major vendor in either domain, yet…
- …Guided by Tucci and Gelsinger, EMC decided in 2010 to make a substantial investment into core technologies that enable the exascale I/O path
  - Joint development with the DOE Labs, in particular Los Alamos
  - R&D funded via the Federal FastForward I/O
  - Work with Partners – tech companies, such as Intel, Micron, Cray, Mellanox, Penguin Computing, Bull, etc.
  - Work within Open Source – LANL PLFS, OpenSFS / Lustre, U of Clemson / OrangeFS / Omnibond, etc.
  - Work with Universities and Research – UTexas (TACC), Purdue, UCSD, UCSC, Michigan, etc.

Focus Of EMC Research And Advanced Development

- Pioneering efforts on Burst Buffers
  - Supplied the first Burst Buffers to the Buffy testbed for Trinity, at Los Alamos
- aBBa – in situ analysis,
  - Active processing in burst buffer (via function shipping)
  - Simulation steering
- IOD (FastForward Burst Buffer), with Intel, Cray and HDF Group
- Co-design with Cray of the burst buffer for Trinity
- 2 TIERS™
  - Evolution of aBBa technology
  - Into a broader market (including enterprise and hyperscale)
  - Novel I/O Stack architecture based on Stack Disaggregation
- VMware advanced development of virtual RDMA, integration with GPGPUs, and more (vHPC = Virtual HPC!)

Not this:

But this:

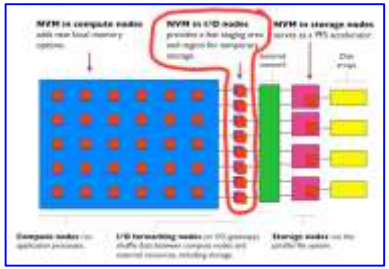

FastForward Architecture

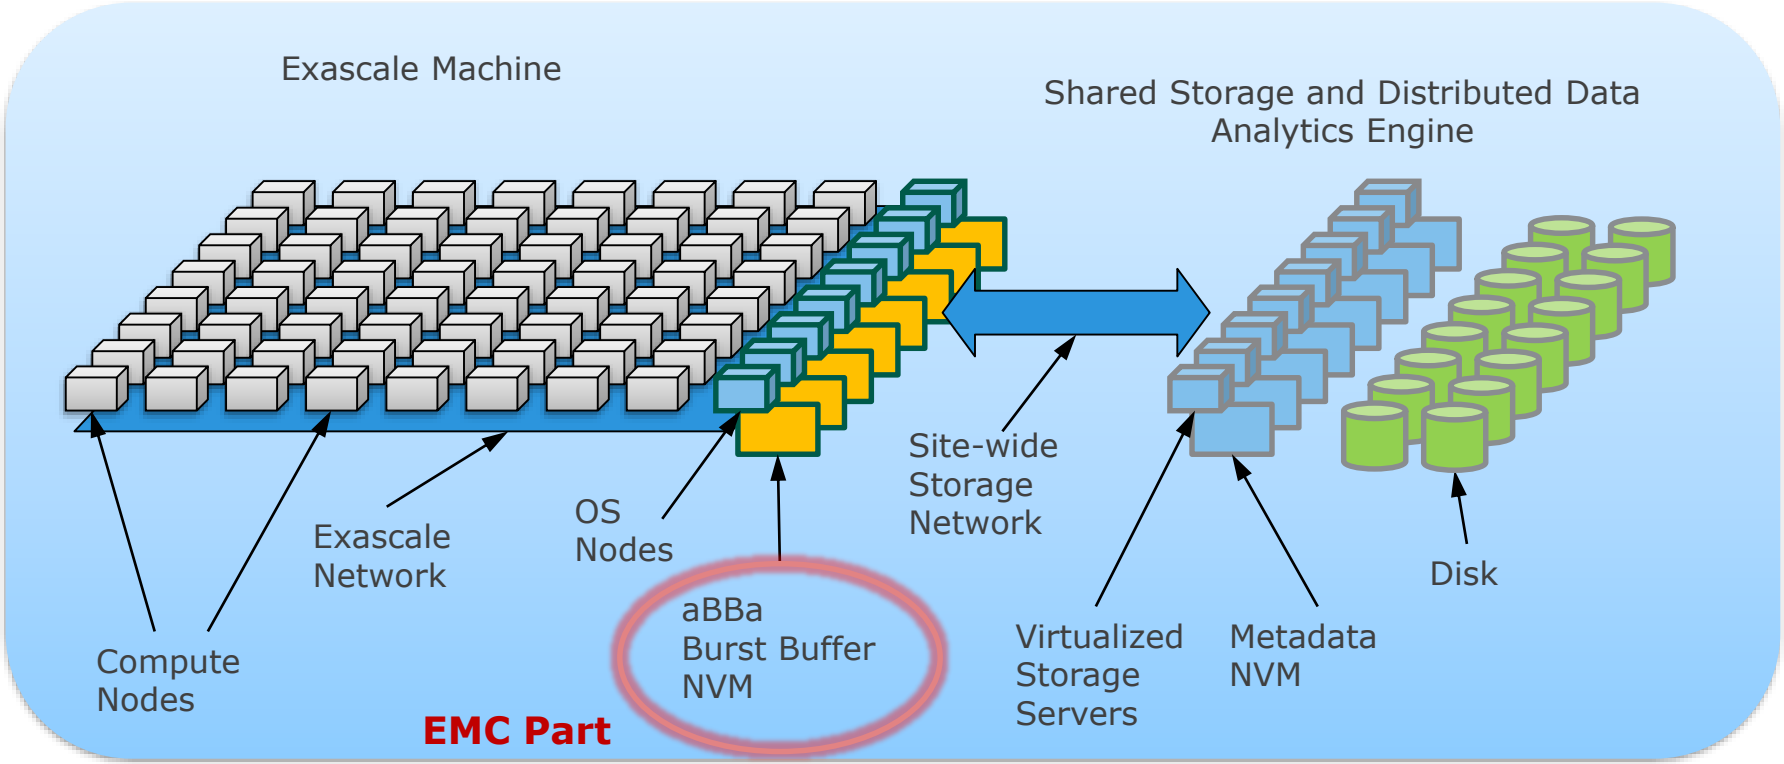

What Is The Bottom Line For EMC, So Far?

- The most notable successes in the HPC markets have been in verticals where people buy storage independently
  - Significant penetration by Isilon in the Life Sciences, O&G and Manufacturing, as well as Education /Research, and Federal
  - Significant progress with DSSD and ScaleIO as platforms for Parallel File Systems (Lustre and GPFS)
  - Incipient presence with ECS, as object store for campaign archival

- (Differently) notable: aBBa / Burst Buffer was not productized at EMC, despite an excellent suitable platform, DSSD
  - The Burst Buffer market is driven by server vendors!
  - Since Burst Buffer is for the most in the public domain, some variants of it are now part of other vendors offerings, those with a stronger server / system presence
  - The work on Parallel File Systems (Lustre, GPFS, OrangeFS, StorNext) has been monetized by EMC in several markets, as solution offerings based on EMC storage (VNX, DSSD, ScaleIO)

EMC HPC Assets

EMC OTS Technologies For HPC

- **Isilon**
  - Secondary storage, active archival, analytics target
  - With CloudPool, hierarchical storage to the cloud
  - Capacity Tier with 2 TIERS™ software
- **DSSD**
  - High performance, HA, high density store for PFS
  - Burst Buffer with aBBa software
  - Unique integration with OrangeFS
  - Fast Tier with 2 TIERS™ software
- **ScaleIO**
  - SDS store for PFS
  - HPC in the Cloud solution
  - Unique integration with OrangeFS
  - Fast Tier with 2 TIERS™ software
- **ECS**
  - Cold archival object store
  - With GrauData PDM HSM, hierarchical from PFS
  - Capacity Tier with 2 TIERS™ software

- **VMware** – vHPC, vRDMA, HPC in the Cloud
- **Virtustream** – HPC in the Cloud

EMC Technology Assets And Connections

- aBBa + Lustre-aBBa
- Storage Platforms
- Lustre expertise and partnership with Intel
- OrangeFS expertise and partnership with Omnibond
- OEM agreement with Penguin Computing
- Collaboration with Bull
- CRADA (collaborative research and adv. dev. agreement) with LANL
- Strong relationships with the LLNL, SNL, ORNL, LBNL, ANL, PNNL
- Strong relationship with TACC (DSSD for Lustre and Spectrum Scale)
- Strong relationship with GrauData (PDM HSM)

What Is The Future In HPC With EMC Now Part Of Dell Technologies?

- At EMC, *the motivation for significant investment* in HPC R&D and Advanced Technology has been:
  - The creation of a clear and strong vision of how these technologies will trickle down into the HPC commercial markets, and into big data analytics in enterprise 
  - The establishment of reputation as an I/O and Storage technology leader in the HPC Community 
  - The campaign for finding powerful server and system vendors as partners, based on the strengths of EMC storage and technology 

We see great synergy between the products and technologies that the two companies, united as Dell Technologies, can bring to the HPC market

Dell Technologies HPC Vision to Democratize and Advance HPC

Dell Technologies will help more people make **more innovations and discoveries** than any other HPC systems vendor in the world, via an **innovative, cost-effective portfolio of solutions** that integrate **Dell Technologies and partner innovations with community standards**.