



## Comments on a specific HDS versus EMC decision

ABC,

Here are some comments for sharing with your client, in prep for a concall Thursday (if still needed), under a couple of headings:

### **Vendor advantages of no real consequence:**

- HDS systems can scale to greater storage capacity, but nobody really uses this—other issues arise long before the architecture scale is exhausted.
- EMC has better VMware integration/features, but nobody uses this to any great extent yet—it is still “just” a SAN, and that’s all we WANT it to be (not a server, not a cloud, not an application engine).
- EMC is selling ‘integration not isolation’ (like Oracle is doing), but we still really WANT storage ‘isolation’ so we can manage it with existing skill sets and for tuning/fault-isolation.
- HDS claims that their various systems can all be managed through a single GUI, but EMC has a similar ‘single pane of glass’ offering.
- HDS can virtualize third-party storage systems and add features to them, but nobody does this to any great extent because of the expensive vendor maintenance costs on the post-warranty third party systems.
- HDS boasts of supporting 2.5” drives (versus EMC), but this advantage is negated by their lack of support for 15K drives of greater than 146GB (in these specific configurations).
- EMC boasts of having 128 cores versus HDS’s 32, but HDS has tons of ASICs which offload the data cores.
- EMC boasts of being able to apply multiple configuration changes simultaneously (versus HDS’ one), but we generally only apply one at a time anyway—regardless of how ‘safe’ it looked in QA/TestDev.

### **(Possible) Vendor advantages of (possible) consequence:**

- HDS can scale architectural ‘dimensions’ (that is, cache, storage, processors, ports) independently, whereas the EMC requires whole ‘engines’ to be added. This affords some short-term scalability advantage for HDS, but we cannot scale one dimension very far before the other ‘components’ have to catch up with it. [Think of replacing an external 100GB harddrive attached to a vintage 1990 PC, with a 2TB SATA drive, in an IDE toaster—you wouldn’t get the full advantage of the SATA drive because the other components just couldn’t keep up with it.]
- EMC has a smaller chunk-size in FAST VP Sub-LUN tiering than HDS (7.6KB versus 42MB). That means faster, more granular data tiering, but also means more metadata and processing overhead. This could be of significant advantage in cases of radical application usage volatility (like applications being swapped in and out of memory too frequently), but might not be of use in high-demand applications—simply because we might not even PUT those critical application databases INTO an auto-tiering pool. We might force them into tier 0 or 1, to make performance very predictable.

- ❑ HDS has partitioning within the global cache, allowing QoS to be implemented for certain application sets. This is a possible ‘predictability’ advantage for application performance control—but only for a narrow range of apps.
- ❑ HDS’s cache is a global cache (like the DMX systems were), versus the distributed caches of the VMAX. This is a possible speed advantage of HDS, but it is offset somewhat by the DMA-structure of the VMAX (that is, the various engines can access the local cache of other engines via direct memory access—using ASICs—versus normal bus-hampered access protocols).

**Vendor differences that probably DO matter in this case:**

- ❑ The new HDS VPS models do not support Fibre Channel in-ports, perhaps rendering some HBA investments useless. [I am presuming that the VSP is under consideration, and not the USP V – which DOES support FC, but not the 2.5” SAS drives.]
- ❑ Although I think EMC overstates their DR advantages routinely, it is still agreed that their SRDF and other replication solutions are the best in the market.
- ❑ HDS does not support 15K drive sizes greater than 146GB, causing the physical/energy footprint to be larger than EMC’s. And it increases the number of ‘moving parts’ that can break—for equivalent storage capacity figures. But—looking at the future-- if the 15K drive tier is ‘going away’ (in favor of two-tier systems of SSD’s and 10K SATA/NL-SAS drives) then this might become unimportant within 18months. But currently, it IS an issue with HDS systems—a limitation.
- ❑ EMC’s experience within your client (due to current installed base/footprint) is a major advantage—existing skill sets, documentation, relationships, connections, etc. These are real benefits of staying within the EMC product family (at least for the VMAX—the VNX is less familiar and represents a learning investment, but less than that for HDS though)

All in all, this would make EMC to be the front-runner (assuming the numbers and contract terms are reasonable).

**From a vendor-management standpoint, I should comment that:**

1. Your client should press EMC to develop 2.5” drive capabilities and to give them preferential pre-launch discounts for in-term upgrades (just pick a number for “% of drives to be replaced for free” within 18 months or so).
2. Your client should emphasize that they will start testing HDS on non-critical applications within 12 months from implementation of the new EMC gear, and that this means that many of the incumbent-advantages will be gone by the NEXT time the refresh time comes around. Remind EMC that they have to DEMONSTRATE additional savings and values on a bi-annual basis (and at every upgrade and new feature launch) to solidify/continue their ‘insider’ status.
3. They should inform EMC that they are NOT interested in their vision of integration, federation, the cloud, running server applications on their SANS—but that they are ONLY interested in storage performance, reliability, and cost-effectiveness. EMC should be reminded that THIS is what will help them keep the business, and not the fancy visions of little or no relevance to high-performance data center operations.

We can discuss any of these on the call, if need be—I hope this helps them some. They are right to keep the storage plant on a 3-year refresh, for reasons of maintenance and innovation.

Thanks—Glenn Miller, VP Strategic Advisory Services [Aug 2011]