

Manage Rising Disk Prices with Storage Virtualization

WebTech Q&A Session – January 25, 2012

1. Where does the virtualization functionality reside?

Solutions from Hitachi Data Systems offer virtualization from the storage controller. These extend the mature controller-based storage services Hitachi developed to support both internal and external heterogeneous storage.

2. Will the solution make my current infrastructure more complicated?

Instead of needing extra hardware layers (management servers, application blades, new switches and separate management software), Hitachi virtualization is easily deployed by enabling an existing software key within the Hitachi Virtual Storage Platform.

3. Can the solution be implemented in smaller phases? How easy is it to roll back?

Virtualization technology in the Virtual Storage Platform can be turned on or off, and rolled out in big implementations or small pilot projects. Organizations also have the option to choose whether to use the system as just another storage system or fully activate its existing virtualization functionality. The system can also be deployed as just another storage system and then used to easily "discover" any externally attached heterogeneous storage volumes, which can be assigned as virtual volumes.

4. How reliable is the tiered storage model currently supported by the virtualization solution?

Hitachi Data Systems believes in supporting lower cost, high-capacity tiers of storage as externally attached but physically separated storage frames from those of the high-end, Tier 1 storage. This not only offers greater customer buying choice and the option to reuse legacy systems as lower level storage tiers, but also eliminates the "single point of failure" risk to both tiers of data. This physical separation of tiers also allows any rebuild efforts ultimately required on high-capacity disks to be more easily isolated with no impact on the performance and availability of any surrounding application data. Since Hitachi virtualization technology allows data to be seamlessly migrated in the background during such rebuilds, access to the impacted tier's data is also assured even during this type of maintenance activity.

5. How do you ensure extra virtualization elements won't decrease system reliability or availability?

Reliability and availability levels for the Virtual Storage Platform do not rely on clustering technology. Instead, these are based on the already mature and proven reliability of the underlying Hitachi storage controller, which has already been field-proven for reliability in large enterprise storage environments. In fact, Hitachi Data Systems offers a 100% data availability warranty.

6. How easy is it to migrate or replicate data between systems or back out once you begin the migration process?

Any-to-any replication, seamless data migration and centralized data management are part of Hitachi virtualization functionality. Hitachi controller-based virtualization (in the Virtual Storage Platform) has been carefully integrated with Hitachi Universal Volume Manager and other Hitachi software. This results in a single set of management tools for performing snapshot, migration, archiving, backup and replication of data across heterogeneous storage systems, even from mainframe systems to lower-cost open systems.

7. Is zero page reclaim disruptive? Is there any perceptible overhead if it's not disruptive?

It certainly could never lead to data loss or corruption in any way. Of course there is some processing activity but it will be minimal. There is I/O too, to read the pages to detect zeros, and after reclaim there may be some rebalance as well. That is intentionally kept to a minimal rate. Even on a small pool it is only a small percent of the total pool capacity; on large pools it is fractions of a percent. We always design such features to operate without impacting production. If you run ZPR at a time when the pool is overloaded, then ZPR will just get aborted.

8. Is zero page reclaim available on enterprise-class Hitachi Data Systems storage only?

Zero page reclaim is available on Hitachi Data Systems enterprise (VSP) and mid-range (AMS) storage. You need to know of a difference in implementation of Hitachi Dynamic Provisioning (HDP) that has an impact here. AMS manages HDP data in 32KB pages inside 1GB chunks. A chunk is dedicated to a DP-VOL. ZPR on AMS will recover pages within a chunk (so they can immediately be re-used in the same DP-VOL for new data). ZPR will also return a chunk to the pool for reuse in other DP-VOL provided the chunk is now completely free. The effect is that ZPR is a little less effective on AMS.

9. Is it a feature within the storage management software suite?

Yes. You can initiate ZPR within Hitachi Command Suite using CLI or GUI: On the Resources tab, select Storage Systems in the navigation tree. Expand the tree for the target storage system that includes the appropriate DP pool. On the DP Vols tab, select one or more volumes and click Reclaim Zero Pages. Click Submit.

10. In order to use HDT on my external volumes, do I have to migrate the data to the thin pool, or will the system automatically thin the data on the external array?

External volumes can be incorporated into HDT and HDP in two different ways:

- A. The external data volumes can be migrated into an already-established HDT/HDP pool. This could be built from internal drives or external ones. You would use Hitachi Tiered Storage Manager to achieve this. Most customers enable the ZPR function at this step, so zero storage is reclaimed automatically.
Also, as pages are relocated between tiers in HDT, an automatic ZPR is performed “for free.” Pages that are relocated down-tier are less-used; this may well be because they are empty.
- B. External volumes can also be used as pool volumes in HDT or HDP. These would be formatted to zero during creation of the tier. If they contained user data, you would migrate this data to a pool first as above. Containing data would be very rare, as the volumes used for an external pool would normally be created from parity groups already free on the external array. We discourage “just throwing in any spare storage” to a pool. With HDP, it will result in a poorly performing pool. With HDT, the poorly implemented tier will tend to slow the whole pool. Pools with external storage should be designed following best practices. (See January HDT Webtech sessions for more description.)

11. What kind of capacity efficiency technologies are available for file data in network attached storage?

The Hitachi NAS Platform supports several capacity efficiency technologies:

- A. File-based dynamic “thin” provisioning helps grow capacity on demand to ensure that capacity is not pre-allocated with potential for unused space
- B. Intelligent file tiering and migration automatically places the right data on the right-cost storage for the right balance of price and performance. Hitachi NAS Platform, powered by BlueArc® has a built-in policy manager that automatically migrates inactive data externally to a Hitachi Content Platform with cheaper and higher density drives, to help reclaim expensive primary storage capacity
- C. Pointer-based snapshots and file cloning (BlueArcJetClone) enables capacity efficient replication and data protection without full copies
- D. Object-based replication (BlueArc JetMirror) enables efficient replication of data over WAN to remote locations, thus saving capacity from making full replicas