


Reduce the Cost and Complexity of Distributed IT Environments

WebTech Q&A Session, November 3, 2010

- 1. Is the deletion of the data at the Hitachi Data Ingestor (HDI) based on access frequency time?**
No. Deletion of data from HDI is based on a capacity “high water” mark of 50% of the total 4TB of usable storage on HDI. When data is deleted, it is replaced by a 4KB stub file. The stub file is transparent to applications and provides a link to the content replicated to HCP, so the content is always available for recall. Only content already replicated to HCP is stubbed and subsequently deleted from HDI. So, at any given time, there will be up to 2TB of original content stored in the HDI file system, and any content above 2TB will be represented by stub files.
- 2. Can the movement of data from the HDI to Hitachi Content Platform (HCP) be scheduled?**
Yes. There is a default schedule for the movement of data from HDI to HCP, and that schedule may be adjusted. HDS can provide guidance for adjusting the schedule based on customer workloads and performance requirements.
- 3. How is HDI different from NAS?**
This solution is intended for customers looking to reduce their storage footprint and administrative overhead at the edge. For example, many distributed organizations, such as retail, banking, Insurance and healthcare, do not have the IT staff or expertise at their distributed locations to properly manage, protect and backup the storage residing at the edge. Although HDI uses industry standard NAS protocols (CIFS and NFS) for access by applications and end-users, HDI does not function as a conventional NAS system. First, HDI is designed as a purpose built ingestor to HCP and functions as a caching device offering NFS/CIFS ingestion and performance acceleration. Second, while other NAS platforms may offer some level of migration or tiering functionality, HDI has automated policies that migrate all resident files to HCP. Since all of the data is replicated to HCP, HDI can recover all of its files in the event of a failure. This makes HDI a nearly bottomless, backup-free solution, reducing management overhead and freeing IT from constant capacity planning. Hitachi Data Ingestor has a limited amount of useable, local storage, which cannot be expanded, and requires the Hitachi Content Platform as its replication target to expand its capacity.
- 4. Can you explain in more detail how autonomic refresh works?**
For any autonomic tech refresh, HCP already supports multiple back ends. This new feature automates the retirement or upgrading of older technologies and platforms. Point HCP to the new system or capacity and it will start writing to the new array. Meanwhile movement of existing data from the old array to the new and read access to the old array will both stop once the system has been completely migrated. This really reduces the migration complexity, delays “fork lift” upgrades,



and minimizes downtime. The platform can remain for decades and adapt to newer, cheaper, better technologies as they become available.

5. Why do you call this backup free?

HDI replicates all of its data, metadata, and configuration settings to HCP. HDI can rebuild from HCP. HCP is backup free because of DPL (Dynamic Data Protection Levels), metadata protection across multiple raid groups and RAID levels, active data integrity checking, replication, failover, etc.