

WebTech: Designing Storage Tiers

Webinar Questions and Answers

1. Most large customers have worked with this for several years- pooled storage. Are we able to push this down to SMB?

Advanced tools like Tiered Storage Manager are not available on Hitachi Adaptable Modular Storage and Workgroup Modular Storage systems, but other than that everything talked about in this session applies. You can do dynamic, nondisruptive data migration on the Adaptable Modular Storage and Workgroup Modular Storage platforms using the Modular Volume Migration feature of the Hitachi Resource Manager software.

2. When you say classify data, are you implying use of an off the shelf product? Any products you can share?

No, what was meant was a manual classification of data, coupled with tools to help create tiers and manage migrations. There are many consulting companies that specialize in helping companies with data classification, but I'm not aware of any off the shelf products.

3. Does Hitachi have a product on managing offline storage (tape)?

Hitachi offers a complete server-based backup solution, the Hitachi Data Protection Suite. In addition we offer an advance reporting tool, HiCommand[®] Backup Services Manager, which is used to report on third party backup products such as NetBackup and Tivoli Storage Manager for compliance purposes or to improve disk and backup device usage.

4. Will your Hitachi Universal Storage Platform architecture work with server virtualization technologies like VMware?

Absolutely, Hitachi Storage Systems are fully cross certified with VMware server environments. Natural synergies exist across the board from the single view of all storage and replication offloading features that Hitachi Data Systems storage systems provide to VMware servers.

5. Is it possible to schedule migrations within Hitachi Tiered Storage? For instance, move data set 'X' once it reaches a certain age?

HiCommand Tiered Storage Manager can be used to create and schedule background migrations and offers a complete CLI for automation purposes. Scripts could be either be written by the customer or as a Services offering to recognize application data sets and automatically control movement due to aging. Hitachi Data Systems calls this 'Policy-driven Management'.

6. Does the Tiered Storage Manager handle other storage in addition to Hitachi – that is, EMC, NetApp, etc.?

Yes, only if it is virtualized behind a Universal Storage Platform.

7. When referring to guaranteed speed per tier, should IOPS be used as the standard singular metric? If so, how do you measure the application needs and determine the disk configuration that can support it?

Depending on the application's performance requirements it may be more important to optimize bandwidth over IOPS or not. You probably also want to consider the ratio of random versus sequential operations. Using a tool such as Performance Monitor of Tuning Manager you can profile the application's storage behavior.

8. The administrative benefits are obvious, what type of performance degradation should be expected by virtualizing storage tier? That is, what performance cost is introduced by virtualizing storage via the Universal Storage Platform or Universal Storage PlatformV?

There is some degradation, but given the substantial extra bandwidth in the SAN and the ability to utilize the faster cache and algorithms on the Universal Storage Platform not much. Exactly how much depends on both the particular hardware combinations involved and the application behavior.

9. Will they change HiCommand Tiered Storage Manager to allow migrating a migration group to multiple tiers?

No. A single migration group can only be migrated from one tier to another. If you have reason to migrate application data to more than one tier, use more than one migration group.

10. When data is moved from one tier to the next what does it look like on the 1st tier. Are there stub files or is it a move?

It's an actual move, no stub files are involved.

11. Are there plans to support tiered storage manager on top of Hitachi Dynamic Provisioning?

Yes. Tiered Storage Manager will support migration of Dynamic Provisioning virtual volumes in the 5.9 release due out in the first calendar quarter of 2008.

12. I've been told Tuning Manager metrics can be used to automatically move data between storage tiers - can you explain how this integration works/is defined?

Tiered Storage Manager can use dynamic input from Tuning Manager in defining tiers, but there is no 'automatic' ability to move data unless you script it using the CLI.

13. Are there plans to support tiered storage manager on Adaptable Modular Storage systems?

Not at this time.

14. Is there a recommended backup policy for tiers other than first?

No one size fits all. Differing requirements for backup policies are one of the things driving differentiation between tiers.

15. What is the migration path for the Storage Domain Controller or upgrading from the Universal Storage Platform to the Universal Storage Platform V?

It is a Hitachi Data Systems Global Solution Services offering. Depending on the circumstance, they will use one of a number of tools to accomplish the move.

16. Is there a performance hit during migrations? If so, how significant?

There is some performance hit during migrations, how much depends on a number of factors. Generally the performance hit is low, very similar to the performance hit experienced with replication products such as ShadowImage®, however if it is an issue migrations can be queued for execution during a slow period.

17. Are the performance characteristics criteria that Tiered Storage Manager gets from Tuning Manager also available from EMC or IBM storage as well?

No, Tiered Storage Manager can only collect performance metrics from Tuning Manager which in turn only supports Hitachi storage systems.

18. Are there any success stories for automatic policy-based storage migrations?

Currently, we have one published reference story of a Swedish customer.

19. Will be there any future integration of Tiered Storage Manager with file awareness, for migrating files from tier to tier?

This is an area Hitachi Data Systems is moving to, but we cannot make any product announcements at this time.

20. Another aspect of tiers is that low-end apps are more likely to use IP SAN and higher end apps to use FC SAN. Is there any way to move back and forth between these? Or to have automated tiers must everything be on one protocol?

Tiers can be on any supported protocol.

21. Is there a process or equation that can be used to determine IOPS for different disk configurations such as RAID level or RPM speed, in order to assign those disk configurations to tiers.

As a rule of thumb for cache miss workloads, you should use 180 IOPS (small block) random per 15k FC/SAS disk, 120 IOPS for a 10k FC/SAS disk, and 75 IOPS for a SATA disk. Also consider the physical IOPS write penalty by RAID type: RAID-10 = 2, RAID-5 = 4, and RAID-6 = 6. Add two more to each if using SATA drives rather than FC (automatic read after write function for SATA). There is no Read penalty – all RAID types have 1 physical IOP per read. It is important to determine your read-write ratio as well as the random-sequential mix before making any performance estimates. For example, with 15k FC disks in RAID-5 7d+1p, in a cache miss situation, for 100% reads expect about 1440 random IOPS, and for 100% write expect about 360 IOPS. You can see much higher disk performance if the requests keep the disk heads in a narrow region on the disks in a RAID Group (locality of reference) or if some degree of cache hits are present (20% is typical on Open Systems). Higher performance on writes is also likely if a lot of writes to the same region are occurring and a single full – stripe write destage to disk is possible.



22. We stripe LDEVs at the OS level, would we still be able to leverage Tiered Storage Manager?

Only if you manually kept track of all of the physical LDEVs used as stripes in the OS level LDEV and migrated them as a unit.