

Hitachi Dynamic Tiering WebTech Series- Session 1 of 3 Q&A Sessions – July 13, 2011

1. Is there a separate management interface for Hitachi Dynamic Tiering (HDT) or is it integrated into Hitachi Command Suite?

HDT shares all the interfaces of other provisioning styles on Hitachi Virtual Storage Platform (VSP). HDT is an evolution from HDP, which in turn is an evolution of classic provisioning so commonality is leveraged. This also has allowed us to remain largely compatible with all provisioning tools. Hitachi Command Suite, Storage Navigator 2 and the new raidcom scripting method are all available.

2. Is HDT set on lun and raid group levels? Can we set exclusions to some raid groups or luns?

Whether you are using HDT or not depends on the type of pool. A pool is either HDP or HDT; if your DP-VOL is from an HDT pool, its pages will come from the pool's tiers. Tier management will be automatic and based on the pool strategy set. Currently, the only exclusions are:

- a. Do not put it in the HDT pool. Sometimes an HDP pool or even classic LDEVs will be the best option.
- b. You are allowed to "lock" a DP-VOL into its current location in the tiers.
- c. We are considering alternate methods of "fixing" specific data into a tier. However, if this type of strategy is adopted, there becomes a point where the data is no longer "dynamically tiered." This would become an issue because HDT will not add value and its operation would be restricted for other users.

3. What is the impact of monitoring and relocating more frequently than the default values?

You have 2 possible questions here: 1) Default automatic policy or 2) manual versus automatic policies.

- a. There is no default automatic policy. When configuring, you will have to make a choice among the strategies available.
- b. If you choose manual operation, you can create newer policies. There are some operational "minimums" that need to be considered in scripting. Immediately starting and stopping monitoring without waiting for all state changes will cause the policies to be rejected.

In both cases there are a number of factors:

1. Monitoring will happen continuously unless manually disabled, but it is an internal operation and intended to have trivial impact.
2. As mentioned in the presentation, frequent relocation cycles could mean that cycles are not completed. The cycle will just be cancelled and a new one started with newer monitor data. Relocation is not an essential operation and can be seen as a strategic improvement.
 - a. We always consider that supporting the production workload is more important than any management activity. This is a Hitachi core value. Relocation therefore follows "limit" rules similar to other products such as Initial copy, ShadowImage®, COW, ZPR, rebalance, and migration.
 - b. The relocation rate is intentionally limited to minimize the probability of any impact.
 - c. If the pool or a tier is overloaded, relocation will not occur.
3. If cycles are continuously being aborted, then there is probably a combination of 2 issues. The 1st issue is too little storage in one or more tiers (upper probably). The 2nd issue may be a flat distribution. If the distribution is totally flat, it may not be suitable for HDT.

4. If relocation cancelation occurs near migration, it should be ignored or possibly disabled until the pool is under the production load only.
4. How do we keep track of what type of storage an app is using since HDT may move them around?
The command **raidcom get ldev** will indicate the tier usage of each DP-VOL. Or you can use DP-VOL tier properties to see tier usage (immediate) and a graph with performance data for the next cycle and weighted.
5. What is the suggestion from HDS to do chargeback with HDT on?
There are 2 possibilities:
 - a. We suggest that you charge all users an average price based on the pool average cost. This will be unfair to some customers, but should eventually average out.
 - b. Another option we suggest would be to obtain the tier residence of all DP-VOLs in each pool and correlate this with owner (i.e. **raidcom get ldev -ldev_list dp_volume -pool 0**)
6. If the promotion and demotions occur too often, will continuous monitoring cause an impact on the performance of the array ?
See answer 3 for more information. Relocation is a lower priority. Even if continuous relocation was indicated, production should not be affected.
7. Does HDT work with LUSE volumes?
No, LUSE is an obsolete technology and is not supported by HDP or HDT. You can have 60TB volumes with HDT and HDP on the VSP, which removes most needs for LUSE. Unfortunately, volumes over 4TB may not be replicated.
8. Will there be any indicators to show when relocation starts and ends?
Yes, please refer to the audit log in section 3.
9. If a LUN is eligible for movement and during moving, there is an increase in activity on the disk, will HDT abort the movement or continue?
LUNs are never eligible for movement but pages in a LUN are. The limit is on a pool and tier basis not per LUN.
10. Will Hitachi Data Systems publish a whitepaper (best practice) when using HDT and ASM together, not just ASM on Hitachi Dynamic Provisioning?
It is currently under consideration.
11. Is there a best practice for each tier? (i.e. 10% T0 30% T1 etc.)
This is covered in part 2. There are rules of thumb, but it is easy to find contra-examples.
12. What VSP mcode code level was this referenced on?
HDT base features are V01 (GA) but a couple of features are V02: subscription, continuous monitor mode and external disks . Raidcom has evolved with new features as each internal update is made.
13. How do we model locality of reference for different applications?
Locality of reference basically means the spikiness of the tier properties curve. Internally, we have started to build up a database of common "characteristics," which include messaging, OLTP, file and etc. These are available to HDS staff as the HDT performance modeler. More information is included in part 2.
14. Slide 35 and 36 show Tier Cap Utilization LT 100 %. Is this a real case? Aren't Tier-1 in mature pools always at close to 100% ?
There will always be the 2% and often the 8% reserves. Once P% nears 60%, you will probably see used

percentage drop from 100% in order to prevent overloading the tier.

15. Will the HDT page work well in Oracle ASM environment?

We believe so and have seen good results in lab and field. At first, you might think that the server side striping would lead to flat curves, but we have good evidence that it does not. For example, the locality on REDO is high. You will find some tablespaces busier than others. Some tables as well as parts of tables and indexes will show up on the top. Obviously, query design may have an impact. But what is good for Oracle is also suitable for HDT. In fact, Oracle spends most of their time designing algorithms to improve locality and reduce true randomness. As I said, storage would be unusable without locality. Database would also be useless without locality (record/table/transaction etc.).

16. Is it possible to lock the state of the HDT pool for specific workload period and then apply it when needed?

You can lock DP-VOL individually or you change to manual mode and stop relocate or monitor as appropriate. You cannot however “store” a profile, only the last monitor period is usable.

17. How is over-provisioning tracked? Where is the original LDEV created (i.e. tier 2)? Are the tiers where the sub LUNs get relocated to?

Please refer to parts 2 and 3. HDT does not consider over-provisioning. It involves pool capacity monitoring, which HDP HDT provide. LDEVs are not created in tiers, but pages are when they are used. The subscription limit is “on creation”. The pool limits are “on page add”.

18. Is 60% utilization of tier measured in terms of performance or capacity?

It’s a measure of performance; utilization% compares actual IOPS with a nominal IOPS rate based on the type and number of drives in the tier.

19. Will Hitachi Device Manager report on how much of a LUN is in each tier?

For Device Manager, see details about DP-VOL. below. It is also trended in Tuning Manager. HCS 7.1 also adds several of HDT capabilities. See part 3 for more information.

Hardware Tier Information				
H/W Tier No.	Drive Type	Drive Speed (RPM)	RAID Level	Used Capacity
1	SSD	0	RAID5(3D+1P)	924 MB
2	SAS	15000	RAID5(3D+1P)	52.58 GB
3	SAS	10000	RAID5(3D+1P)	34.98 GB

20. We have used Hitachi Tuning Manager to dynamically move LUNs based off the event to different HDP pools. Is this based on policies or solely on hardware with no human interaction?

What you describe is similar to the GSS Policy Based Management described in the beginning of the presentation. This is a strategic technology. HDT is more tactical and is solely based on the strategies you set. No intervention is required or possible. You could combine the two, but only at the pool level (hot pool and cold pool) not at the tier level. There is no move to LUN from tier 0 command currently because HDT is not designed to work this way.

21. What is the max capacity per hour that is typical for relocates? Is it possible to check HDT pool performance and utilization at tier level?

Please reference the slides in presentations 2 and 3 for information on HDT pool performance and utilization at tier level.