

# Hitachi Dynamic Tiering: Overview, Managing and Best Practices

## WebTech Q&A Session, November 6 and 13, 2013

- 1. What is a typical tier distribution (from a percentage perspective) in an HDT pool? In other words, how much disk should be part of Tier 1 relative to Tier 2 or Tier 3?**

Application profiling and an understanding of capacity and performance growth will determine the final tier distribution percentage numbers. In many 3-tier pools, the top tier is likely to be smaller than the other tiers, and the third tier is likely to be the largest tier. However, proper sizing needs to be done to ensure that no tier is undersized, which can have a large impact on the overall performance of the pool.
- 2. Can you monitor part of the day only (say, 8:00 a.m. to 6:00 p.m.) and relocate based only on this time period?**

Yes. One feature of the mobility (aka Hitachi Tiered Storage Manager) license is the Hitachi Dynamic Tiering (HDT) monitoring schedule template, whereby you can monitor for only a specified period. Relocation will be based only against the activity of this monitoring period.
- 3. Do you recommend extensively pinning volumes using tiering policies?**

Extensively pinning volumes using tiering policies negates the value proposition of HDT. For some application volumes or perhaps for a short period of time, pinning a volume may be appropriate. However if pinning a large number of volumes for an extended period of time, you are in effect emulating how Hitachi Dynamic Provisioning works, and using only a single class of disk geometry for those applications.
- 4. Is there a priority between promotion and demotion?**

Yes. Promotion of pages has a higher priority than demotion. In general, pages will be promoted, and once all pages eligible for promotion have been relocated, demotion will occur. Further, relocation priority is dependent on (in order):

  1. The presence of data-placement profiles.
  2. The use of tier policies.
  3. Normal relocation activities.
- 5. With solid-state drives (SSD), 15K, 10K and 7.2K, how would a 3-tier system work?**

In a 3-tier system, you could choose any of the 4 types of disk geometry listed here. However, the faster geometry will always be ranked higher over slower speed disk when the pool is made up of internal disk only. For example, you could have a tier order of SSD / 10K and 7.2K, or a tier order of

15K / 10K / 7.2K. Regardless of its disk speed, if external storage is part of the HDT pool it will always have a lower tier ranking than internal storage.

**6. What is the page size for data movement between tiers?**

42MB pages.

**7. Can monitoring be specific to LUNs or it is for an entire array?**

You can “disable” monitoring of specific volumes by using pinning policies. Otherwise, monitoring is set at the pool level.

**8. Where are the tier properties: Hitachi Command Suite or Hitachi Storage Navigator?**

The view tier properties graph is available in Hitachi Device Manager 7.5 or later, Hitachi Command Suite Mobility 7.5 or later, and in Storage Navigator when you are at the pool level (select More Actions, lower right-hand corner).

**9. What tool is used to gather this graph and information?**

The view tier properties graph.

**10. Looking at the tier property table graph, what criteria should I use to decide which of the 3 tiers I should upgrade to improve or resolve performance problems?**

The performance utilization and/or the capacity utilization figures are a useful guide for determining which tier to expand to remediate performance problems. If the capacity is fully used, then pages that need a high-performing disk may be forced to reside on a poor-performing disk (there is no room to move these pages to the better disk). If the performance utilization is greater than 60%, the tier is not able to generate the I/O activity required of the pages on that tier. Either the pool or the tier is under-configured to deliver the performance requirements. Or as we often see with the third tier, the upper tiers are completely used and highly active pages are residing on the lowest tier, unable to be promoted to the tiers that can deliver the performance requirements.

**Among customers with whom you are familiar, what single configuration error is most common?**

Incorrect sizing of tiers, generally the result of not doing proper application profiling.

**11. If the HDT buffer space for a new page assignment is changed after HDT has already been implemented, what are the ramifications of making such a change? For example, what if the Tier 3 assignment was initially set to 0%, then later on changed to 8%?**

As a default, new page assignment is set at 8%; if set to 0%, then pages from that tier will not be assigned during a new write request, and this could impact performance. Those pages will have to wait for the relocation cycle to be placed to a different tier. As a general rule, buffer settings should not be changed.

**12. Can we expand virtual volume on the fly non-disruptively?**

Yes, using Hitachi Command Suite or Storage Navigator.

**On what timelines would a relocation take place if I have to move a block from lowest to highest tier? Does the block go to the highest tier directly or does it go up the hierarchy, waiting for the next rebalance?**

Promotion occurs during the relocation cycle that follows a monitoring period. If required, a page can be promoted from the bottom to the top tier in one relocation process.

- 13. Did you say you that you can convert Hitachi Dynamic Provisioning (HDP) pool to HTP pool or the other way around?**  
HDP to HDT.
- 14. Can notifications be set up to go out if the performance utilization of a pool is too high?**  
Yes, you can set alarms in Hitachi Tuning Manager related to pool performance, but the criteria to set up this alarm would be different than that measured by HDT. Consult an HDS solutions consultant for more information about this.
- 15. Can static LUNs be restricted to a low tier?**  
Yes.
- 16. What does "performance utilization" measure?**  
Performance utilization is only an approximation of tier utilization. It is based on assumptions of read/write ratios (50-50). It does not factor in RAID I/O (parity I/O), or I/O activity related to items such as rebalancing, pool shrinkage, and similar factors.
- 17. Will there be a difference using Flash Memory Drives (FMD) as Tier 1 and an external storage as Tier 2, compared to having the second tier internal in the Hitachi Virtual Storage Platform?**  
External storage is always going to have a lower rank than internal storage, even if the disk geometry of the external storage is faster than that of the internal storage.
- 18. Wouldn't it be important to monitor for shorter periods and to move the data quickly after the monitoring cycle to promote the hot data when it requires performance?**  
It depends on the application's I/O activity. Shorter monitoring periods will allow for earlier movement of pages if need be.
- 19. Any issues using HDT for VMware data stores for guest operating systems? What about Oracle database files?**  
HDS has a number of recommended configurations for Oracle and VMware environments. A local solutions consultant can help you build out an HDT architecture that is best suited for your applications and hypervisor technology.
- 20. With shared tenancy, reporting on HDT would be very important. So if I have 10 clients accessing a single HDT pool, how can I charge customers based on their utilization?**  
There are a number of reports from a number of tools that can show you the utilization of a pool volume. Your local solutions consultant can facilitate a meeting to show you how can use these tools to meet this objective.
- 21. Assuming the relocation procedure depends on the storage processor utilization, how frequently can the relocation be scheduled in a day, and what is the time gap between relocation jobs?**  
Relocation cycles are defined using Hitachi Command Suite and/or Storage Navigator, and relocation follows a monitoring period (which can occur at ½-hour, 1-hour, 4-hour, 8-hour and 24-hour intervals). Processor utilization will determine the pace at which relocation will occur.

- 22. How long should the relocation take? When is it a concern that it is taking too long?**  
The standard pace is 3 TBs per day for relocation, although it can be higher or lower than this. If relocation is taking a long time, it is likely that the array processor utilization rates are high.
- 23. What are best practices for monitoring HDT and quantifying its actual cost effectiveness in a particular customer implementation?**  
Arrange a discussion with a local storage economics expert, who can help you understand this.
- 24. How does Dynamic Tiering monitoring and relocation impact host reads versus host writes?**  
Monitoring and relocation does not impact host I/O activity.
- 25. Where can I find the best documentation of tiering policies and how page movement happens in different situations?**  
Contact your local solutions consultant.
- 26. How long is the data pattern studied before promotion and demotion is initiated? How long between analysis of data, or is it continuous?**  
Monitoring is for a defined period of time. Monitoring of data is always happening; the time period used and mode type will determine when and where pages are relocated.
- 27. What is the difference between HDT for Hitachi Unified Storage 100 and HDT for Hitachi Virtual Storage Platform or HUS VM?**  
Conceptually, they are the same but there are differences in that the underlying technology, HDP, is slightly different on the HUS 100 than on VSP or HUS VM. Contact your local solutions consultant for a full explanation.
- 28. Is there any impact on rebalance if one uses overprovisioning?**  
No impact.
- 29. Can you review the part about using different RAID types in same tier?**  
Intermixing RAID configurations is supported, meaning that on the same tier you could have a combination of RAID5 (x+y) and RAID6 (x+y). However, this is not recommended.
- 30. What are your thoughts on spanning HDT across internal and external (virtualized) storage?**  
This is fully supported, but having proper sizing and performance requirements is highly recommended.
- 31. How could someone compare and contrast this solution with EMC, IBM, and other tiering technologies?**  
EMC = FAST2; IBM = EasyTier. There are other players. Contact your local solutions consultant for more information on the differences between HDT and offerings from other vendors.
- 32. A long monitoring cycle will not promote pages until the monitoring cycle is complete. At the end of the monitoring cycle, the pages may look like they were hot but 24 hours later it becomes clear that the pages didn't really need to be promoted. Is what they really needed was to be promoted 23 hours earlier when the data required higher disk performance?**  
Set your monitoring cycle to ½ hour or 1 hour.

# Hitachi Dynamic Tiering: An In-Depth Look at Managing HDT and Best Practices

## WebTech Q&A Session, November 13, 2013

### 1. Why can't I run performance utilization at 100% on all 3 tiers?

If a tier's performance approaches 60% utilization HDT aims to move I/O down a tier. The 60% sustained I/O figure is used to accommodate peaks and prevent queuing. When performance is greater than 60%, the tier range is increased. This reduces the tier's utilized capacity, meaning that all of the tier's capacity will not be used – but this is better than overloading the tier with too much IOPH. If all tiers are over 60% (i.e., >60% / >60% / >60%), the pool is over-utilized and tier ranges are increased again.

### 2. If I move a server into HDT, how would I know when to move a second server into the pool? Can I mix my redo/mirror logs with the data pool so that I can take advantage of flash, or do I need to separate redo/mirror logs into a separate pool contention?

Application profiling would be a consideration here. Speak to an HDS solutions consultant.

### 3. Why can't you redefine tier boundary, to lessen SATA load?

Redefining the tier boundary will not lessen the load on a tier of SATA; the buffer zone is a small percentage of the size of each tier.

### 4. Is it ok to turn on HDT even though you don't have multiple tiers at this time?

Yes, but it requires a HDT license if you want to make a HDP pool an HDT pool. If you want to monitor a pool to show HDT-like characteristics on the VSP or HUS VM (no license required), speak to your local solutions consultant about System Option Mode 93(SOM937). It is available to enable HDT monitoring for an HDP-Pool without Dynamic Tiering license.

### 5. Can we do reporting through Command Suite?

Yes. All of the reports shown in the presentation were from Hitachi Tuning Manager.

### 6. Are all of the reports and graphs canned or do you need to build them? Is there any version level required for Tuning Manager?

The reports referenced in the WebTech session come with the Hitachi Tuning Manager release 7.6 and are available on the HDS Community Website, and can also be provided to you by your local solutions consultant. The Agent for RAID needs to support DataModelVersion 8.4, so the Agent for RAID from at least Tuning Manager 7.4 should be used.

- 
- 7. What version of Hitachi Command Suite supports HDT? Do you have many customers utilizing external storage with HDT enabled and are there any special considerations we should think about in those scenarios?**

HCS 7.0 onwards supports HDT. Yes, we have many customers using external storage. There are many considerations that you should be aware of (performance, availability, RAID configurations, etc.) that require closer examination. Speak with your local solutions consultant to setup a session to discuss using external storage.

- 8. Can you break the tier performance graphs down to individual DP-VOL to show what pages sit on which tier and compare this over a sequence of relocation cycles to show the difference per DP-VOL that is being moved between tiers?**

Yes.

- 9. Are those heat maps available for HDT on HUS VM?**

Yes.