

# Increase Savings with File Systems, Thin Provisioning and Storage Reclamation

## WebTech Q&A Session: January 20, 2010

**1. Is there any work being done to implement ongoing reclamation in any file systems besides VxFS?**

The approach being taken today is to have file system vendors review the T10 draft and cooperate with Hitachi to develop an appropriate reclamation scheme. The Write Same approach is agnostic so any file system vendor can develop to it. In the industry every thin provisioning supplier does storage mapping differently (for instance different granular storage allotment sizes) so file system vendors will need to cooperate with the storage vendors to maximize their efforts.

**2. What about security of this handshaking process? Is there any chance for valid data to be deleted inadvertently?**

The approach using the SCSI Write Same command has a built in level of “security” comparable to a standard SCSI Write command. In other words you do not expect that there is a security issue concerning a Write command being directed into a storage address that is inappropriate. The Write Same SCSI command would have the same built-in safe guards.

The file system’s purpose is to reserve storage addresses (LBAs and LBA ranges) for the use by a file. The file system will perform the same fundamental operation to reserve (dedicate) the storage addresses that are the reclamation target of the Write Same command during the execution of the command.

**3. If you have an Adaptable Modular Storage 2500 virtualized behind a Hitachi Universal Storage Platform® VM do you still need to script for reclamation as in the Adaptable Modular Storage, or will it be automatic as in the Universal Storage Platform VM?**

Yes, a stand-alone Adaptable Modular Storage system can be thin provisioned, or if it is virtualized behind a Universal Storage Platform V or VM it can also be thin provisioned. In this latter case you can use thin provisioning on the Adaptable Modular Storage system behind the Universal Storage Platform (not recommended) in which case the Universal Storage Platform will see and treat them as normal volumes and the write same process won’t work. Or you can use the thin provisioning capability of the Universal Storage Platform, in which case it will work smoothly and transparently.


**4. Is Hitachi currently working with any other OS vendor native file systems in the area of TP reclamation?**

Yes.

**5. Can virtualized storage say Hitachi Adaptable Modular Storage be thin provisioned behind a Universal Storage Platform?**

Yes, however the ability to use reclamation changes.

Consider the hybrid configuration of the following model:



Host LUN -> USPV eLUN -> UVM -> virtualized LUN -> AMS DPVOL -> AMS HDP Pool.

The host based thin provisioning reclamation will not operate in the above model. The reason for this is that the host does not 'see' a thin provisioned LUN therefore the host will not initiate a reclamation operation.

- 6. What happens to the view from the Universal Storage Platform of the virtualized LUN, I assume it still looks the same size, etc., however the reclamation of the storage gets passed through the Universal Storage Platform down to the Adaptable Modular Storage to make it available to the thin provisioned pool on the Adaptable Modular Storage**

The LUN does not change its logical size due to any reclamation process. The host using the LUN will not see any changes.

The reclamation of the virtualized storage is not possible (please see previous question and answer)