Software Product Description

HSM - Hierarchical Storage Manager

Description:

HSM is an add-in to the VMS I/O system which provides VMS with the ability to have transparent file migration between active and near line storage in one or more steps. When files are migrated ("shelved") from normal disk storage to backing storage, a marking is left on them which is automatically read so that when the file is opened, it is automatically retrieved by HSM from nearline storage. Thus, a user or program need not be aware at all that such shelving occurred and no operator intervention is needed to perform the "unshelving" operation. This is distinguished from operation where a user must first request an archived file be reloaded, which requires detailed advance knowledge of such needs. The files appear to have been on disk all the time, but in fact the online disk space is conserved. It is also distinct from modes of operation where a file's location visibly changes. These, too, require that programs be told where the new site is, which can be awkward. HSM provides total transparency of file migration, invisible to programs and users apart from small delays where files must be unshelved.

In addition, HSM provides two unique "soft link" abilities which complement unshelving, and manages volume space.

The basic capabilities of HSM are these:

* Files can be shelved (by space-making script or by command) and unshelved automatically from nearline storage when they are opened. The process opening the files then sees a successful open with no side effects. Shelved files can be stored in compressed form if this is desired, and can be stored in any desired location. Storage of shelved files on tapes or the like can be done also.

* Files can be "soft linked" to other files, even across disks. This mode of access can be used for a sort of permanent shelving on another volume by truncating the original file to zero blocks. The soft link operates extremely fast and causes the file in question to be opened in its new location, with the channel restored on close so that again a program observes no change, but the file is accessed transparently at its new site. Where the new site is a read/write device this can be most effective.

* Files can be "soft linked" in a "readonly" mode to another file. In this mode, suitable for read-only backing storage, whenever a file is opened for read-only access, it is transparently and instantly opened on its linked site, wherever that may be on nearline storage. When such a file is opened for any kind of writing, however, it is treated as a shelved file and is unshelved and replaced on normal disk before the open is done. Thus any read/write access will find the file in a suitable location for its open to succeed, transparently. (Notice that soft linked storage must be on disks and must not be in compressed form. The decompression would induce an unwanted delay in access.)

* Disk space can be managed. Whenever an extend or create (or inswap) would not have adequate space on disk, HSM starts a "make-space" script which is tailored by a fullscreen utility to match site policy. This policy can select files based on access time, size, name, or characteristics for space making.

HSM tailoring can be done either for the entire system or for any number of disks at a time. Thus it is possible to have a multi-tier migration strategy with each tier managed by a separate HSM server, so that files may migrate toward slower storage but still be retrievable transparently no matter how far down the hierarchy they have migrated. Features of HSM can be separately enabled or disabled as well on a per server basis should this be desired.

File migration can also be handled by a simple command which runs a menu driven selection of how the file is to migrate and in which mode. Provision exists to regenerate file markings in case they are lost or to audit the markings. Also, it is possible to specify exempt images which are not subject to file unshelving, or to set a process temporarily as exempt from unshelving or softlinks, so that operations which must view the disk without triggering unshelving can be easily run.

HSM can be run in a mode where there is essentially no overhead at all imposed (just a few instructions added along some paths and no disk access) for any files except those which need softlinks or possible unshelving. There is no limit to how many files may be so marked on a disk. A fullscreen setup script allows one to select the HSM run modes. Even if HSM is forced to examine all files for its markings, the overhead imposes no added disk access and costs only a tiny added time (typically a percent or two) in open intensive applications. In addition, HSM can be turned off or back on at any convenient point should this be desired.

Support:

HSM runs on VAX VMS 5.5 or greater or AXP VMS 6.1 or greater. The same facilities exist across all systems. HSM must be installed on each cluster node of a VMScluster where it is to be used but imposes no restrictions on types of disk it works for. HSM will work with any file structure used by VMS, so long as a disk class device is used to hold it. It is specifically NOT limited to use with ODS-2 disks. Where HSM is run on only some nodes of a VMScluster, its features will not be usable on the other nodes, but no ill effects will occur.