Evolution of Data Tiering

From Hierarchical Storage Management to High Performance MicroTiering

Data tiering has established itself as an integral part of virtual storage in the enterprise storage market, especially where terabytes and petabytes of data are involved. Tiering automatically places data on the most optimal storage *tier* based on usage patterns without the need for human intervention. This technology has become essential as the volume of data grows at annual rates of up to 100% in many situations.

A good example of where data tiering can be essential is in web search or business information retrieval where there is no advance warning of data traffic surges.

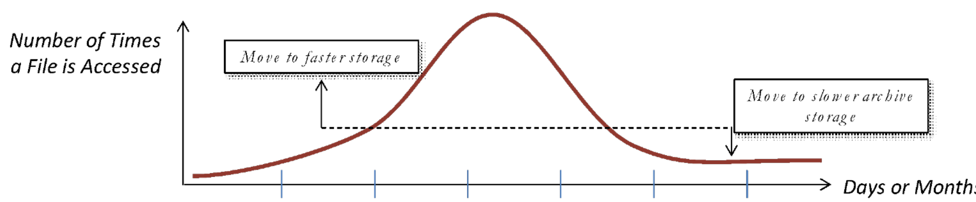

Example Tiering Based on File Activity Levels

In poorly designed storage systems, if data is not on the right performance tier when needed, any detrimental access delays may cause system overload, unsatisfied customers or missed opportunities especially if in financial markets. Similar situations are now arising in next generation business analytics environments which are increasingly merging data from multiple sources including the web and require very fast response times to data requests.

Improved response times resulting from data tiering are the primary motivation for tiering. Customers and/or applications experience much faster response times and improved experience on ultra-fast storage versus slower, larger bulk storage. Modern automated data tiering ensure that important data (or ***hot data***) is located on the highest performance storage media at very short notice, then later moved back to the slower and usually much cheaper archive storage when no longer needed (i.e. becomes ***cold data***).

The economic value of data tiering is straight forward. Storage with the fastest response times are substantially more expensive on a per terabyte basis than regular hard drives, often in the range of 100s of times more. Virtual storage combined with data tiering facilitates a mixed pool of fast expensive with slow in-expensive storage to provide the optimum performance-cost tradeoffs.

Evolution of Data Tiering

HSM (Hierarchical Storage Management) is one of the earliest forms of automated data tiering and was implemented as a means of moving data up to then expensive disk drives from tape backup archives when needed and returned back to archives once the usage had dropped off. HSM was implemented mostly as an application layer solution or built into the application itself and was ***file aware*** i.e. it typically looked at file names, extensions and file metadata along with that file's usage patterns to determine how to handle that file on a periodic basis. For example, every month, the HSM application would examine file age and usage and backup specific file types e.g. spreadsheet, financial database or legal files, to an archive, but ignore all temporary data or non-critical data files. The most common use of HSM in large storage systems was with enterprise class Fibre Channel or SCSI hard disk drives as the primary tier, less expensive storage such as low cost SCSI or IDE disks as possibly an intermediate storage tier and usually a tape library tier for archiving.

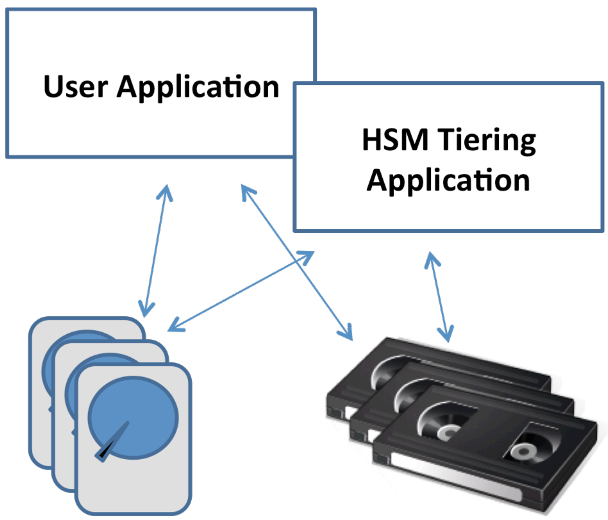

HSM (Pre-Virtual Storage)

ILM (Information Lifecycle Management) introduced a more comprehensive approach to tiering that included data classification engines for sorting data by importance and started to introduce the concept of making the data more fluidly move ***up and down the tiers*** based on activity levels. ILM also introduced the concept of lifecycle management with its aim to manage data throughout the whole lifecycle from creation, usage and eventual deletion. More sophisticated policies were introduced that allowed users or system managers to specify a criterion as to when and where data should be moved and most importantly, if it can be simply moved or deleted.

As file sizes grew however, it become more difficult to managed data at file level only, especially as in some applications, several large databases of several gigabytes at a time could be briefly and frequently touched for a quick or simple lookup. With no restrictions, file-only tiering would result in the whole file being pulled up to fast storage tiers and ultimately could result in a significantly degraded response time as too much data was being moved around unnecessarily.

Sub-File Tiering

Hence the concept of sub-file tiering was born i.e. allowing ***sections*** of a file that were being frequently accessed to be ***promoted*** (i.e. moved up to a faster tier) rather than the whole file. This dramatically reduced overall storage fabric and network traffic. However, as it became impractical for HSM or ILM applications to cater to every application and file format and maintain file consistency, many applications that used such large files started to implement the data tiering function within the application itself. Overall, complexity started to grow and storage management systems were having to track a significant number of variables at file and sub-file level, as well as deal with the specifics of the particular operating system they were part of.

Virtual Storage and Block-Level Tiering

As the use of industry standard storage area networks (SANs) increased, more intelligent storage devices were deployed as part of a broader effort to virtualize storage i.e. abstract storage to the point where users and IT managers no longer needed to be aware of specifics of the storage they were running on. Virtualization introduced a set of new management tools allowing IT managers to more easily add, change, move and provision storage. Features such as thin provisioning, block level tiering, remote replication and snapshots were now very possible and most importantly, implemented without the increased use of any system level or application software as was the case with earlier HSM and ILM systems. These newer systems often operated in concert with some higher level software combined with many new block level features.

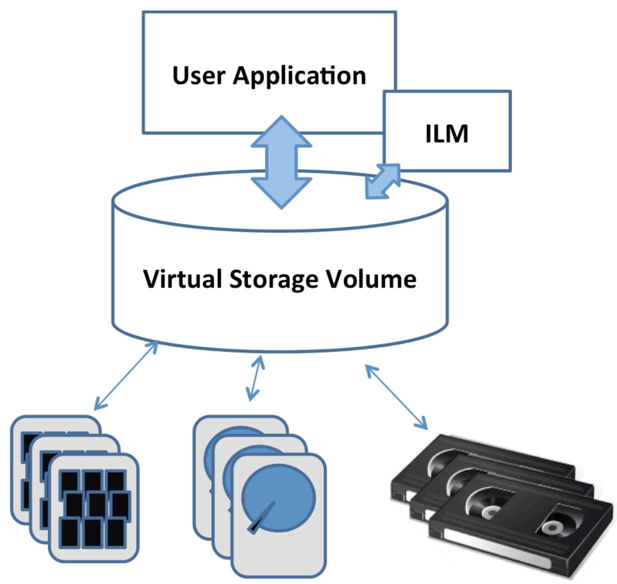

ILM and Virtual Storage

Block level virtual storage is a method of manipulating very low level storage I/O requests before they reach the physical storage, allowing the storage request to be remapped from a virtual device that the application or host computer sees to the actual requests that eventually hit the underlying storage. The significance of this was that storage could now be dynamically shrunk, grown, older storage retired and replaced with newer faster storage while the host computer system(s) continued to use the storage. In an Internet enabled age where 24/7 availability of information is becoming common place, storage virtualization has become a necessary way to implement storage. Most systems increasingly adopted the use of SANs in order to manage the scalability of this new generation of virtual storage devices, combined with the new trend of virtualizing the server itself using hypervisor technology.

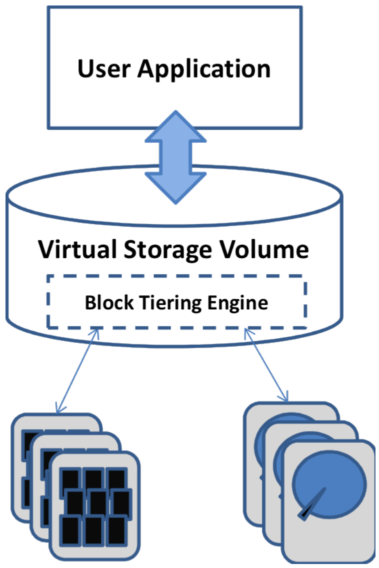

Virtual Storage and Integral Block Level Tiering

With the introduction of virtual storage technology, the ability to tier at data block level is now possible i.e. at the storage layers instead of the file or applications layers. Blocks are defined as one or more of the underlying physical devices blocks or sectors, which in mainstream storage devices is typically 512 or 4K bytes. Tiering usually groups these into virtual pages of 'N' sectors in order to reduce the number of mapping and statistics variables required to maintain the block tiering environment. These pages can range in size from a few 10s of Kilobytes up to several Megabytes.

One of the significant benefits of block level tiering is that sub-file tiering occurs without file or application helper software, thus is totally transparent to the application and its operating system. Furthermore, block-level tiering can be implemented either in software in the driver layers of the computer operating system or in the hardware storage IO layers in an embedded processor. Prior to the Enmotus Virtual Storage Processor introduction, block-level tiering has been limited to software driver layers only and thus limited to a handful of operating systems. Hardware accelerated solutions are capable of operating in any operating system and have the added benefit of being able to operate at sub-hypervisor level in virtual server or workstation environments.

Flash Storage Tiers

Solid state disk drives (SSDs), and specifically PCIe based SSDs, significantly disrupted the legacy virtual storage model. While SAN based solid state appliances exist and have a strong place in the storage world, PCIe SSDs have demonstrated huge performance gains in excess of 1 million IOs per second when placed directly on the primary computer's local bus. However, moving the SSD into the server as a standalone storage island has fragmented the traditional virtual storage tiering and caching architecture, making many of the existing storage management tools obsolete in the process.

The benefits of flash based storage devices have demonstrated to be huge in today's information driven environments and have improved data response times by as much as 10,000 times over a single equivalent hard drive for example. Furthermore, the power savings by eliminating several racks of short stroked high performance disk drives has resulted in energy savings of up to 1000% for the fast tier of the storage due to the sheer number of physical hard drives they can replace.

The problem is that these new devices have created a new storage island that is decoupled from the SAN virtual storage pool, so the use of host server based applications or caching layers have to be re-introduced into the server software layers to "re-integrate" this new form of storage into the storage hierarchy. Furthermore, the pricing level of these new flash devices is targeted at multiple SAN disk array unit replacement costing several 10s of thousands of dollars each, hence are still somewhat out of the reach of mainstream computing. So what is being done to service the mainstream server markets that don't have petabytes of storage to manage?

Block Level MicroTiering

As SSDs become an integral part of servers, there is a growing need for ultra-high performance data tiering that can transparently move data in real time between fast and slow storage. Furthermore, as big data architectures start to proliferate into enterprise computer-storage networks, we find an increased reliance on a networked cluster of virtualized and distributed storage-servers rather than the centralized server-SAN architectures of today. Combined with advances low cost hard drives, today's servers can easily accommodate 64TB to 128TB of capacity in a single rack mount chassis

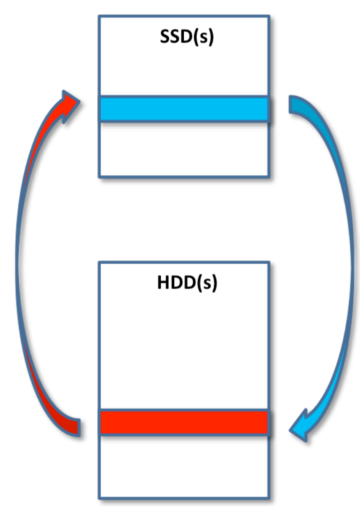

Enmotus MicroTiering

This creates an exciting opportunity to create a new class of distributed tiering architecture. One example is the DAS virtualization architecture offered by Enmotus. Their product implements a high performance form of tiering referred to as MicroTiering which is a local "in-the-box" approach to tiering that is optimized around ultra-high performance PCIe, SAS or SATA attached SSDs. MicroTiering virtualizes and tiers these devices with direct attached hard drives that are also industry standard SAS or SATA attached. A novel patent-pending mechanism is used to move data between tiers in real time.

The MicroTiering solution utilizes an intelligent PCIe based data processing engine with hardware accelerated MicroTiering, and works with most mainstream operating system or hypervisor environments. This includes highly distributed mapping system architectures such as Hadoop used in big data web search environments. The benefit of embedded MicroTiering and virtualization is its ability to fully automate and "drop-in" to existing configured servers or bare metal server systems due to OS and hypervisor agnostic abilities. Further, embedded MicroTiering allows deployment of SSDs into skinny-hypervisor environments requiring only a lightweight driver to support all of the tiering capabilities.

Conclusion

Data tiering has evolved over the years and become a mainstream and integral part of any storage system. There will always be fast expensive and slow storage, and information continues to be created more than deleted, so expect many forms of data tiering to appear to address the issue of information availability and help reduce operating and system costs associated with maintaining that data.