Cost Optimized Storage Performance

How Much Flash Do You Need?

Abstract:
Flash storage has become a necessity for applications requiring high IOPS and low latency, but not all data requires this level of performance, especially as it ages. The key is identifying what data requires performance and what data does not, and with this information, being able to automatically assign the appropriate storage resources to this data.

Enmotus Virtual SSD – Cost Optimized Storage Performance

The question isn't if you need flash; clearly that answer is YES. The real question is "**How much flash do you need?**" Your "Hot" or "Active" data needs to be stored on flash so that your critical applications are guaranteed the best possible IOPS and latency. Most data is accessed infrequently within a relatively short period of time after creation. It doesn't make sense to continue to store this data on expensive flash storage. At the same time, it is critical that this data be easily accessible when needed.

The optimal solution, from a cost and performance point of view, has your active data on fast **primary storage** and your cold data on cost effective capacity centric **secondary storage**, and continuously adapts to changing access patterns. The key is being able to identify the actual working set and then apply the proper storage resources in order to optimize performance while containing costs. The challenge is that manually identifying hot and cold data and managing it is extremely difficult.

Enmotus' Virtual SSD Software (VSSD) uniquely addresses this challenge. The software analyzes storage workloads in real time and dynamically provisions flash storage to the workloads that require it without adding any overhead. Unlike caching architectures, the flash media is utilized as primary storage for both your read and write active workloads, while your cold data is stored cost effectively on capacity efficient secondary storage media. The secondary media can either be traditional spinning hard drives, resulting in a very cost effective solution, or it can also be lower cost bulk flash storage, resulting in a best of breed all flash solution.

The result is an **Intelligent Data Management** solution that automatically places your active data on Primary Storage and moves the cold data to Secondary Storage. This is all done seamlessly without any user intervention. To the operating system, the primary and secondary storage looks like one single volume, so all data is always immediately accessible.

The Enmotus Advantage

The VSSD software features three key technological advances: Virtualization, Analytics and Automation (Optimization). This provides unparalleled value to datacenters by balancing application requirements with the underlying storage resources, while minimizing management resources.

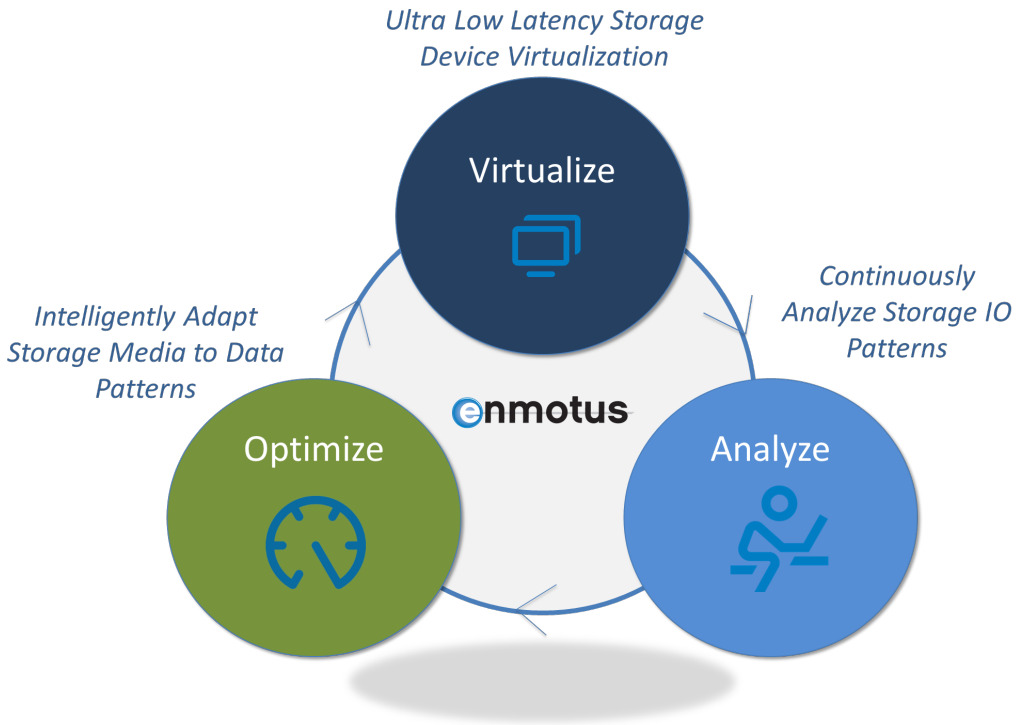

Virtualization:
Enmotus' technology blends fast storage media with capacity storage media into a single virtual volume with almost no overhead, resulting in storage that has performance characteristics of the fast media but capacity and cost points closer to that of the secondary storage media. A key differentiator is the flexibility provided by the ability to work with any block storage devices, be they local to the machine or remotely connected via a SAN. For example, cost conscious virtual drives could be comprised of SAS/SATA fast media and HDD for capacity. Likewise, performance driven requirements could be comprised of NVMe flash for the fast media blended with SAS/SATA SSDs for the capacity media. This configuration provides the benefits of NVMe performance but at competitive all flash cost points. Likewise, ultra high performance requirements could use NVDIMMs or 3D XPoint for the performance media blended with whatever capacity media meets the requirements.

Analytics:

Data bases, as well as most applications, typically only access a small part of the data set, meaning the IO does not have a wide range across the entire volume. The challenge is identifying the size of the working data set, as well as its location within the volume. Enmotus' Device Analytics works at two levels: embedded in the kernel IO path directly collecting raw information about the storage IO transaction size and frequency, and user level tools to visually identify how much of your data active, thereby eliminating guesswork. This allows you to accurately size your flash purchases to meet the needs of your applications. Easy to use visual monitoring tools show IO activity in real time so you can validate that flash has been provisioned to your active applications.

Automation:

The statistics collected from the device analytics engine are utilized to dynamically provision the flash resources to those applications that require it in real time. The Enmotus technology re-maps the volume continuously without user intervention. If access patterns change, the volume is automatically remapped to maintain optimal performance. Advanced pinning capabilities allow you to lock files to either the performance or the capacity tier as needed. These tools allow you to allocate performance media to batch processes that are executed on a predefined schedule.

Upgrade to flash performance

Dell and Enmotus have partnered to bring the benefits of Storage Automation and Analytics bundled with Dell's PowerEdge Servers. Dell PowerEdge servers are designed to deliver in demanding mission critical applications such as CRM, ERP, OLTP and large databases. Your choice of performance storage media, such as NVMe flash, blended capacity storage media, along with Enmotus Virtual SSD Software provides you the optimal balance of performance, storage utilization and cost.

The Enmotus Virtual SSD enables flash equivalent performance but at significantly lower cost points. In the majority of applications, the actual data working set is just a fraction of the entire volume. A properly sized flash performance tier will provide enough capacity to keep all of the active data set on the flash, while the majority of the data is kept on cost effective media.

The following graph compares the performance of 3 different configurations: an all HDD configuration, an all-flash configuration, and a virtual SSD configuration. As shown in the graph, the Virtual SSD demonstrates all flash equivalent performance[1]. The key differentiation being that that the majority of the capacity of the virtual SSD is comprised of hard drives. Cost savings for a capacity equivalent Virtual SSD vs. All Flash can be as much as 5X or more.

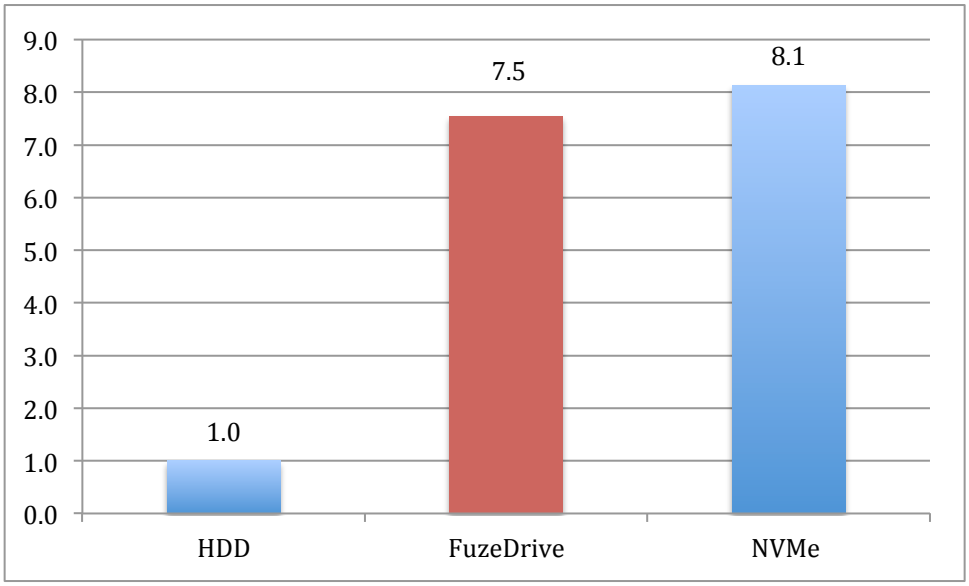

Figure 1. This chart illustrates that the Enmotus Virtual SSD performance is over 7X greater than a hard drive only, solution, but more importantly, it offers SSD equivalent performance at a substantially lower cost point.

[1] Testing was done using a Dell R730XD PowerEdge server, Microsoft SQL and HammerDB, an open source load testing and benchmarking tool. The HDD configuration used 15K 2.5" drives. The All Flash configuration used Dell Express Flash NVMe SSDs, and the Virtual SSD configuration used a combination of the 2.5" hard drives as well as the Express Flash NVMe SSD.

Flexibility – Buy What You Need:

Our advance analytics identify the size of your active data, allowing you to size the performance tier to your specific needs. Enmotus Virtual SSD provides the flexibility to configure storage to your specific requirements, which include performance, capacity and cost. Since it works with any block storage devices, you can mix and match the performance and capacity storage of your choice to create a configuration that meets your specific needs. VSSD supports any ratio of primary to secondary storage, so you can easily size each storage tier to meet your individual needs, providing enough flexibility to satisfy the need for any enterprise application. As needs change in the future, add more performance storage while keeping the system online.

The following 2 charts are examples of potential cost saving when using the Enmotus Virtual SSD. Costs may vary depending on what class and capacity of SSD and HDD drives are used. Costs in the examples are only for the price of the storage media.

The high capacity example shows a price delta greater than $46K between an all-flash solution and the Virtual SSD, a difference of 5X. Given that most data is rarely accessed, it provides much greater value. It doesn't make business sense to store this data on your most expensive storage, but it does make good sense to know that it will be promoted automatically to fast storage when needed.

High Capacity Example[2]

Configuration	Drives	Total Capacity	Total Cost	$ Savings
All Flash	46x960GB SSD	44.160 TB	$58,374	NA
Virtual SSD	2x960GB SSD 22x2TB HDD	44.960 TB	$11,690	$46,414

The low capacity example might satisfy the needs of a typically small to medium enterprise application. Again the savings are quite substantial.

Low Capacity Example[2]

Configuration	Drives	Total Capacity	Total Cost	$ Savings
All Flash	10x960GB SSD	9.6 TB	$12,690	NA
Virtual SSD	2x960GB SSD 5x2TB HDD	10.96TB	$4,618	$8,342
Virtual SSD	2x400GB SSD 5x2TB HDD	10.4TB	$3,732	$8,958

[2] Virtual SSD assumes 2 mirrored SSDs for redundancy. Drive price for 960GB SSD = $1,269. Drive Price for 400GB SSD = $826. Drive price for 2TB HDD = $416

Conclusion:

Dell PowerEdge servers with the Enmotus Virtual SSD software provide best in class performance, cost and manageability. The Virtual SSD software enables flash performance without breaking the bank. Its flexibility allows you to configure your storage to meet the needs of your specific application and cost points. Full automation means you know your storage is continuously optimized so you can focus your energy on more important tasks.