

Virident Storage Class Memory *Performance Without Compromise*

The Application Performance Problem

Delivering application performance in a cost effective manner is a challenge for today's enterprise IT organizations. Most applications suffer from having to wait longer for data from the legacy storage infrastructure, even as the computing horsepower increases over time with the advent of faster processors and more cores in today's datacenter servers. Traditional storage devices that rely on mechanical spinning media cannot keep up with the performance and throughput of today's CPU architectures, leaving applications and processors sitting idle waiting for data, driving down utilization levels in servers.

Traditional Approaches Are Inefficient

There are several traditional methods that are leveraged in today's enterprises to address application performance problems, but all of them deliver less than satisfactory results.

Adding Storage

Storage throughput and performance can be increased by adding more HDD spindles to expensive storage arrays and by reducing the amount of data tied to a given spindle. This is called 'short stroking'. Besides the increased direct cost of this approach, the required additional floor space, power, and management overhead increase the cost even further.

Adding Memory

Adding more DRAM to a server is another approach to the problem. Moving datasets to fast memory reduces the need to access slower storage media to retrieve application data. But the issues associated with this approach are the exponential cost of memory as the capacity grows, and the non-persistent nature of DRAM

which requires data sets to be re-staged when a server comes back online.

Both the traditional approaches described above lead to server and storage architectures that are out of balance, where applications are still not operating at optimal performance levels and the costs are not acceptable.

Flash-based Approaches

Flash memory has been leveraged during recent years to address the application performance problems discussed above. Flash memory offers many benefits, including a silicon-based media that delivers data access at very low latencies and is non-volatile such that it does not lose its data when power is lost. It combines many of the benefits of DRAM and magnetic media in a single product. This is well understood in the market today.

However, most vendors have taken a 'bolt on' approach to leveraging this technology for time-to-market reasons, which is sub-optimal in terms of delivering the full potential of flash for data access bandwidth and latency. Also, as users are finding out, performance of flash storage over extended periods, with real workloads and at full capacity, can be very different from their datasheet numbers. Two such popular but sub-optimal approaches are described below.

Flash Packaged as Disk Drives

Many products simply put flash media in a traditional drive form factor, and connect it to the server using standard storage interconnects such as SATA. Such implementations do not fully utilize the inherent performance of flash media in terms of bandwidth and access latency.

VIRIDENT



Flash 'Drives' on a PCIe Card

Another recent approach is to put 'naked drives' behind a storage controller, aggregated on a single PCIe card. This also incurs the overhead of having to deal with storage interfaces, which are optimized for spinning media rather than flash, which is more similar to memory. In addition, these products leverage external processors in order to assist with the heavy amount of processing required on the card. These processors run very slow as compared to the host CPUs, and therefore decrease performance. These architectures are also more error prone given the number of discrete devices that have to all work together, usually from different vendors. The disparate devices that must interoperate properly may include storage controllers, RAID controllers, and 3rd party flash controllers.

Finally, flash storage delivered as standalone disk drives or aggregated on a PCIe card is inefficient and more costly to manage. In both approaches, you are still presented with separate drives. To have a single pool of capacity, you have to use 3rd party software RAID solutions to create a single host volume, or choose to protect against failures using software RAID protection. This adds additional overhead and reduces performance.

The Virident Solution – FlashMAX SCM with vFAS

Virident FlashMAX™ Storage Class Memory (SCM)

Virident's unique approach to solving the application performance problems described above incorporates specialized software and hardware that combine to implement Storage Class Memory (SCM) within the server. Virident's SCM architecture has been designed to tightly integrate different kinds of flash media, hardware and software to deliver memory-class performance with storage-class capacity and persistence.

Virident's FlashMAX SCM devices and associated software deliver performance without compromise, along with HDD-like capacity in a very compact, universal

form factor. FlashMAX devices range in capacity from 300 GB to 1.4 TB, in both SLC and MLC types of flash media. They have been designed from the ground up to fully exploit today's computer architectures with many, fast CPU cores and the PCI Express interconnect bus to deliver maximum application performance. The accompanying software allows applications to access data stored in the SCM in the most efficient manner possible, far surpassing traditional storage protocols and interconnects. In addition, FlashMAX devices deliver performance over the entire lifetime of the device, including when the device is at near full or full capacity.

Virident vFAS

vFAS stands for Virident Flash-management with Adaptive Scheduling, which is the Virident software technology that delivers the most efficient access to flash media for applications. In addition to providing optimized access for peak performance, vFAS also includes many sophisticated techniques for ensuring that applications get a steady, sustained stream of data at all times. vFAS virtualizes the underlying flash media to present a standard block device interface to applications, without leveraging inefficient storage protocols or interconnects resulting in unprecedented gains in application performance without any change to the application.

Memory-like performance and Hard Disk Drive-like capacities

For application compatibility reasons, it is important to maintain the interfaces associated with the disk drive world if applications are going to seamlessly take advantage of the benefits of flash-based storage. However, very few comparable solutions have figured out how to do this effectively without incurring the burden of traditional storage interconnects like SATA or SAS, which places a performance penalty on the solution as described above. Virident's FlashMAX SCM devices present a traditional block storage volume to the host so that applications can easily access it without realizing that it is a different type of media. But that is where the similarity to a disk drive, and disk drive interfaces ends.



vFAS has been designed to treat flash media much more like an extension of memory, while maintaining a traditional block storage interface for applications. All of this is done without leveraging storage protocols, storage controllers, or storage interconnects. The result is access latencies under 20us, which is closer to DRAM performance than storage. vFAS allows applications to enjoy the benefit of a new memory tier due to this optimized access to the media, where the Flash performs as a Storage Class Memory (SCM) device rather than a traditional storage device. The results are clearly demonstrated by the ability of Virident's FlashMAX with vFAS to deliver over 1.4 Million IOPS in a single low-profile card – the only product in the market capable of doing so.

The two diagrams below outline the differences between the traditional 'bolt-on' approaches to leveraging Flash media in servers as compared to the Virident vFAS SCM approach.

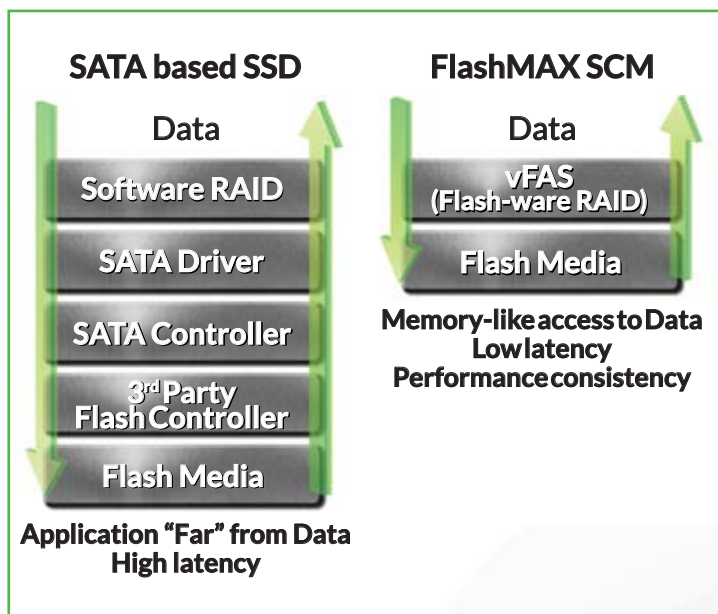


Figure 1. 'Bolt-On' approach to leveraging Flash (on the left), which can be done on a PCIe card or using a storage slot inside a server, leverage a lot of layers and associated context switches, which combine to keep the application 'far away' from the data in performance terms. Virident's vFAS (on the right) provides optimized access, allowing applications to access flash more like memory than storage, all while presenting a traditional block device to the system/application for application compatibility reasons.

Unconditional Performance

Many flash-based storage products on the market today deliver good performance in one dimension or another, but not across all. In addition, they will show significant performance dips when background flash maintenance tasks run as well as dramatically decreased performance when capacity utilization is high.

Virident's FlashMAX with vFAS delivers consistent performance across all application workloads, as well as when the device is fully utilized. FlashMAX with vFAS delivers application performance whether it is peak small block read performance, where a single low profile card can deliver over 1.4 million IOPS, or sustained mixed read/write performance when the drive is nearly full.

vFAS has sophisticated scheduling algorithms that allow flash management tasks like garbage collection and grooming, to occur without impacting application performance. In addition, the FlashMAX card leverages multiple hardware components that provide for a wide 'fan out' of operations, allowing hundreds of operations to be carried out in parallel. The results of this architecture are realized when the device is nearly full, where all competing products will see significant performance declines. Virident FlashMAX with vFAS delivers 2x higher performance than comparable flash storage solutions when measured at near full capacity.

Besides delivering maximum sustained performance to applications, this architecture also allows Virident to deliver datasheet performance numbers at the advertised capacity, whereas comparable solutions must run in 'performance mode' which is about 50% of the advertised capacity.

Simplified Management

100% of the capacity available on a FlashMAX card is available as a single host volume on the server without having to leverage 3rd party software RAID products to stripe across multiple drives. Many competing solutions do not offer this capability. With FlashMAX, you can have a single volume presented to the operating system up to the formatted capacity, which is currently a maximum of 1.4 TB, depending on the FlashMAX product used.



Flash-aware RAID

Virident FlashMAX adapters with vFAS support flash-aware RAID for enhanced reliability and data availability. This feature allows the discrete Flash components included on the card to be isolated as separate slices to provide 7+1 RAID protection, implemented in a flash optimized manner. The flash modules on a card are spread across RAID groups, allowing for multiple failures to occur without disrupting an application's access to the data. This delivers continuity of operations in the event that there are failures with the flash media itself.

Global and Local Wear Leveling

Virident FlashMAX with vFAS offers global wear leveling to maximize the lifetime of the flash media. Also, data is relocated to other parts of flash that are less-used whenever needed. Most of the 'bolt-on' approaches that leverage 3rd party flash controllers only do very local wear leveling, where the flash managed by the individual controller (several are used on each card) is wear leveled, but wear leveling cannot occur across the entire device. This makes these products susceptible to 'hot spots' and premature wear. FlashMAX's sophisticated wear leveling delivers maximum lifetime of the flash media.

Conclusion

Today's IT organizations are looking for solutions to application performance problems in order to remain competitive, reduce capital expenditures and reduce management and operating costs. In order to address this challenge, data center architectures are being re-designed to take advantage a new class of media to accelerate application performance. A new tier of storage is being deployed within servers as a result.

Virident's FlashMAX SCM with vFAS offer the best price-performance on the market today, due to its uncompromised performance across diverse workloads and over the entire lifecycle of the device. In addition, FlashMAX offers the highest density products on the market today by delivering up to 1.4 TB in a low profile device.

Why compromise when architecting new application servers for performance by leveraging flash-based products that do not deliver value in all dimensions? With Virident's FlashMAX SCM with vFAS, you no longer have to compromise. Today's IT organizations can now realize unprecedented gains in application performance, along with unconditional performance over the entire lifecycle of the device. These benefits are revolutionizing the way leading IT organizations are thinking about their future server and storage architecture in order to maximize application performance.



VIRIDENT

(408) 503-0100
www.virident.com