

STORAGE SWITZERLAND

VMWARE SERVER SIDE CACHING - FLASH SSD OR DRAM



George Crump, Lead Analyst

Server virtualization has a well-documented problem; it creates an I/O blender that can bring storage networks and storage systems to their knees. The I/O blender vs. the desire to design increasingly dense virtual machine (VM) architectures has created a new market of solutions that leverage server side flash SSD and caching software to offload much of the storage I/O from the storage network and the storage system. Flash-based server side caches have their challenges though and it may be time to investigate DRAM based server side cache alternatives.

Comparing Flash Based Caching to DRAM-Based Caching

There are a variety of problems that server side flash SSD caching solutions need to overcome. These challenges are why so many caching vendors have struggled with initial adoption.

Installation and Testing

The first step in any IT project is to implement the potential solution and then to evaluate it. This is especially important in cache solutions since every I/O profile is different. The IT Planner needs to see exactly what the performance improvement will be in THEIR environment.

The good news is an increasing number of caching solutions have a free trial download which makes acquiring the software easy. Almost all of these solutions, however, require flash SSD to be installed in order to work. Flash can't be downloaded, it has to be ordered and installed. This means paying for the hardware upfront, scheduling downtime and installing the SSD device.

A cache that leverages DRAM, like the [Infinio Accelerator](#), can utilize the DRAM that is already in the server. Based on Storage Switzerland's analysis, it is not uncommon for the average VMware Host to have less than 50% of DRAM actually be in use. Consequently, in many cases, no additional hardware needs to be procured and installed to immediately start deriving the benefits of a local server cache.

This also makes production use case testing safer. As stated above, I/O profiles are typically very different. While testing for stability is important, the only way to truly know the impact of a caching resource is to apply it to production and have it work with real-world workloads. Furthermore, not requiring the pre-purchase of SSD and scheduling downtime to install SSD capacity, is an even bigger advantage on production systems.

Once the impact of cache technology on production workloads is verified with DRAM based capacity, then flash can be added to the cache if it is needed.

Reliability

Another well-documented challenge that flash-only caching solutions need to overcome is the rapid decay rate of flash cells as they are used. Even though most server-side cache solutions are read-only, they still have to be written to first in order to get data loaded into the cache. If the cache is relatively small, then the cache turnover rate (new data replacing old data) can be fairly high. That turnover rate not only leads to poor performance, it also leads to high flash failure.

The concern about cache turnover rate leads IT planners to design flash caching areas that are larger than they actually need just so that data will be refreshed less frequently in the caching area. Doing so increases performance and reliability of the flash but it also significantly raises the cost of the solution.

DRAM has no write concerns, it does not wear out over time and data is written to DRAM as fast as it can be read. This means that DRAM caches can be sized exactly to the size of the active data set, making the impact of a high cache turnover rate negligible. Data in a DRAM cache performs better and re-loading of cached data is faster. While DRAM is more expensive than flash, the total deployed cost can turn out to be less expensive than a

flash-based alternative; particularly when underutilized server DRAM can be leveraged as a cache resource.

Efficiency

An increasing number of server-side caching technologies are able to optimize the amount of storage memory they use. This is typically achieved by some form of deduplication or compression that eliminates redundant data. When data stored on a slower medium, like a shared HDD or SSD system, is being promoted to the server-side cache, the data in the cache needs to be scanned and analyzed for potential duplicates. In both flash and DRAM use cases, this process takes fractions of seconds. Since DRAM is faster than flash, however, it can perform its efficiency processes much more quickly, which actually allows for a more detailed data analysis and lower cache latency.

Share-ability

Another important aspect of server side caching in a VMware environment is how those caches handle VM mobility. What happens to a cache when a VM is migrated from one host to another? Most server side solutions are resident in each single server, and they don't pool their caching resources. If a VM is migrated from one host to another, the caching software must account for this in some way.

The most common method for server side caching solutions is evicting the cache for that VM on the original server and then rebuilding it on the receiving server. This first technique assumes that the receiving server has the caching software and hardware already installed. Secondly, it also means that the newly migrated VM has to suffer HDD-like performance until the cache is rebuilt.

Another option is to create a shared cache by pooling capacity from the various hosts. This solves the migration problem and adds a layer of redundancy in case of a cache failure. It also allows for the use of smaller amounts of storage memory per physical host since these resources can now be pooled.

A shared cache pool does add a layer of latency because data has to be distributed across a private network segment. Here again, DRAM helps alleviate latency issues thanks to its high performance capabilities.

Volatility

While the advantages of DRAM caching have been well outlined in this article, there is one disadvantage, and it is the chief reason why many vendors choose flash as the primary caching area - volatility. If DRAM simply loses power, all the data that is stored in it is gone. Flash does not have this problem; if power is lost, data remains in flash storage.

DRAM caching solutions overcome this problem in several ways. First, they count on the very high probability that most servers in a data center, especially those that are hosting dozens of VMs, have a battery backup. Second, most of these caching solutions are read only, meaning that if there is a failure, all the data is already safely stored on hard disk storage.

Finally, caching solutions like those offered by Infinio, thanks to their global pooling of memory resources, can offer some level of redundancy within their cache pool. This not only helps to protect against power failure but also against failure of a server within the cluster or a memory module within the shared pool.

Conclusion

While flash-based storage solutions garner a lot of attention, thanks to their solid performance and high capacity, DRAM offers higher performance that is right-sized for typical server cache use cases. As importantly, DRAM is already prevalent in servers and can be immediately leveraged without resorting to additional capital expenditures or requiring a lengthy maintenance installation window.

About Storage Switzerland

Storage Switzerland is an analyst firm focused on the virtualization and storage marketplaces. For more information please visit our web site: <http://www.storage-switzerland.com>

Copyright © 2013 Storage Switzerland, Inc. - All rights reserved