

Server-Side Cache

Introduction


I/O makes or breaks the system, so storage performance is always a big deal, especially in virtualized environments, where VMs are hungry for IOPS. As it's extremely expensive to implement an all-Flash, even more so – all-RAM storage, both considered to be overkill. Thus, a combination of slower spindle, faster Flash and much faster RAM tiers is typically used in the industry.

Problem

This multi-tier approach brings up a few issues. Faster and more reliable memory **costs more**. High-performance or non-volatile memory usage significantly raises the resulting price of the system. Talking about flash, SLC flash memory is very expensive, while cheaper but still costly MLC flash has shorter lifetime due to up to 10 times lower write/erase cycles limit.

VM workload is dominated by random I/O, which is difficult to predict. Insufficient cache memory would mean low cache hit ratio. The idea to get more cache memory stumbles upon high **price**. Additionally, certain processes, like starting moved VMs on destination node after vMotion, Live Migration or XenMotion, cause **drops in performance**, because the VMs are started in a “cold state”. This means that the VM is started without its data in cache, causing a dramatic loss of performance.


Using conventional RAM instead of non-volatile memory, renders the cache prone to errors, because data in **cache will be lost** in case of hardware malfunction or power outage. On the other hand, utilizing power-independent memory raises the problem of high price, referenced above.



All Flash Storage

500.000 IOPS

=




High-Performance storage is very expensive

Solution

Price of the system is lower, because StarWind uses commodity hardware, like conventional RAM, SATA SSD, MLC FLASH or so for caching. Expensive flash memory will have better lifetime because of possible SSD usage in either write-through or write-back mode, while StarWind uses conventional RAM as a write buffer and Level 1 cache to adsorb writes. This approach turns Flash into Level 2 cache, reducing the amount of write cycles going through it, prolonging its life. Additionally, space reduction technologies, namely in-line deduplication and compression, lower the actual amount of data to be written, allowing flash memory to last even longer, cutting the bills by requiring rare replacements.

StarWind raises **performance** by having bigger cache for the same money. Commodity inexpensive hardware is used – MLC flash instead of costly SLC flash, so more memory may be bought to meet workload requirements. In addition, starting VM's after migration doesn't affect the performance, as all moved VMs are started in "hot state", because caches are kept coherent – synchronized between nodes. This means that the destination node holds the required data in cache, so the VM starts without loss in performance.

Reliability is kept at maximum – StarWind basically mirrors data in cache between multiple nodes, creating distributed cache. This way, even when the power goes out, all data is safe, because redundant replicas are stored on all nodes. Besides, cache blocks are digitally signed, negating the possibility of silent data corruption – bit rot. Additionally, space reduction technologies lengthen the life of flash memory, lowering the risk of failure.



StarWind creates High-Performance storage with commodity hardware affordable

Conclusion

StarWind uses commodity inexpensive – DRAM, SATA, MLC and high-performance – PCIe, DIMM Flash hardware providing maximum possible performance and reliability without breaking the bank.

In 2016, Gartner named StarWind “Cool Vendor for Compute Platforms”.

Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.