

White Paper

Server-Side Cache

Introduction


Contemporary market demands result in mission-critical applications having to process unprecedented amounts of data, generating astounding numbers of I/O. With this in mind, the storage subsystem has to be able to provide more speed for virtual machines (VMs) for them to handle all I/O requests correctly. NVMe storage technology, which is quite pricy, was designed to help achieve that. However, NVMe hardware itself lacks the necessary mechanisms to fully deliver its benefits, particularly if used as shared storage. Hence, there's a need for a system that can combine NVMe and regular SSD in a unified cost-efficient storage pool.

Problem

In essence, an all-SSD storage infrastructure can easily satisfy the predominant majority of workloads. Nonetheless, you'll need to consider your chosen NAND type. For example, due to their limited erase cycles, the currently popular MLC and QLC flash storage would probably be not the best fit for an IT environment with frequent, many daily rewrites.

All-NVMe storage has the power to cater to even the most I/O demanding applications. Unfortunately, its entire potential is left unused in usual, daily usage. Considering the higher price of NVMe, equipping your infrastructure with it and not utilizing all its capabilities is simply a waste of money, yes, but also of the brilliant technology itself.

That being said, building all-NVMe or all-SLC-SSD storage systems will sum up to serious costs, which, frankly, may not be the way to go for every company.



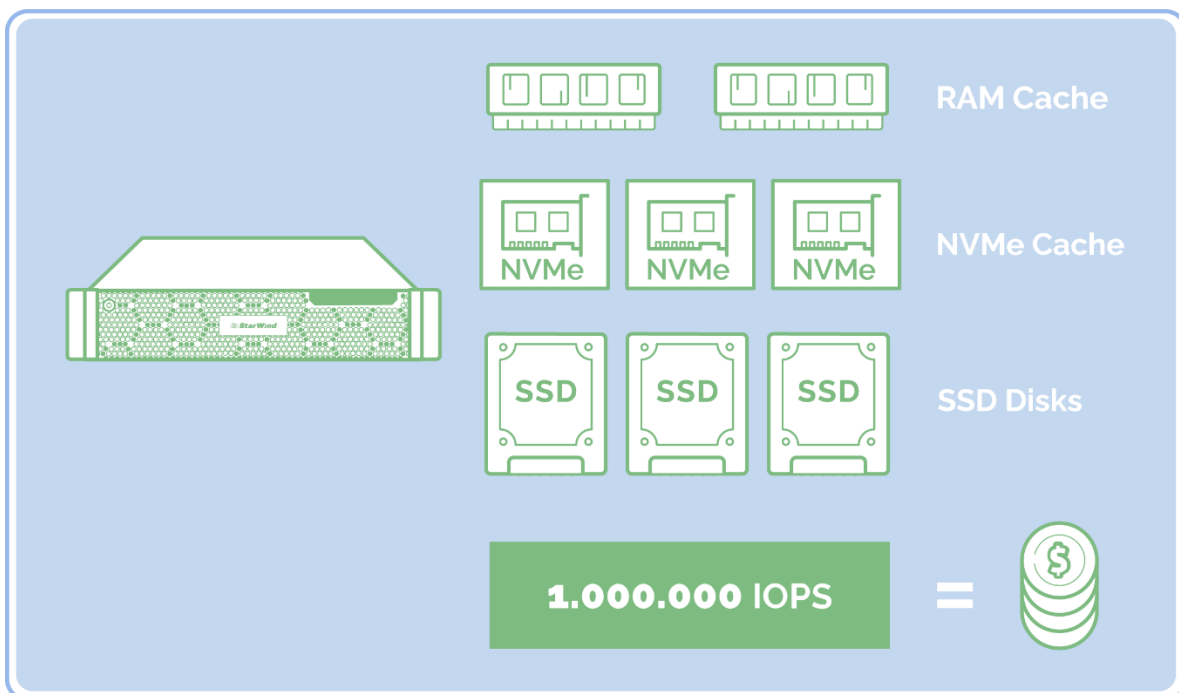
High-performance storage "boasts" a high price tag

Solution

StarWind has designed a way to combine redundant NVMe drives, such as Intel Optane™, and inexpensive MLC and QLC flash into a reasonably priced solution that's just as powerful and reliable as all-NVMe and SLC systems. While StarWind NVMe caching handles all the writes, MLC and QLC are effectively used as the main storage capacity wasting fewer erase cycles.

Consequently, StarWind armors your IOPS-hungry applications with all the available NVMe performance without you having to actually buy all-NVMe storage. Such a configuration enables your storage system to work smart: none of its performance potential is wasted or gone unused and all the necessary power is always there to process any required workloads. Apart from this, in a scenario where adding NVMe flash as caching layer is not an option, StarWind can easily and effectively alter RAM resources into L1 cache to secure faster writes for designated applications.

All of these options are available while securing seamless and solid fault tolerance: StarWind "mirrors" data in cache between hypervisor nodes thereby creating distributed cache. As a result, in case any failure or outage occurs, all data is safe and kept integral thanks to identical copies being stored on all nodes in the StarWind shared storage pool.



StarWind creates high-performance storage while using affordable commodity hardware

Conclusion

StarWind lets nothing go to waste. Our technology efficaciously combines the fastest flash, NVMe, with affordable SSDs in a unified shared storage pool. We further enhance it by repurposing a certain part of it into L1 cache that will be dedicated to your most demanding applications. That way, all of your IT resources are smartly utilized to the best of their potential without you having to pour substantial investment into building cutting-edge hardware-based storage infrastructure.

Contacts



Phone number

+1 617 449 77 17
+44 2037 691 857
+34 629 03 07 17



Fax

+1 617 507 58 45



Voice Mail

+1 866 790 26 46



Customer Support Portal

www.starwind.com/support

Support Forum

www.starwind.com/forums

Sales

sales@starwind.com

General Information

info@starwind.com

To learn more, visit www.starwind.com