

Announcing Zero Move Tiering for FlashBlade



Summary

Zero Move Tiering for FlashBlade delivers two different performance classes of storage from the same array. This innovative approach to data tiering eliminates the costly burden of managing multiple systems for the same results.

[bws_pdfprint display='pdf,print']



We're excited to announce the release of Zero Move Tiering (ZMT) for [FlashBlade®](#)—an innovative and simplified approach to delivering two different performance classes of storage from the same array. ZMT provides the TCO economic benefit of data tiering while eliminating the complicated and costly overhead needed to manage multiple systems for the same results.

Data Tiering Challenges

It's been widely accepted for quite some time that workload data has generally fallen into two access classifications: "hot" data that is frequently accessed and "cold" data that is leveraged less often. And, while ratios may vary, it's safe to assume most data centers operate with at least 10%-30% of their data requiring high-performance delivery, while the other 70%-90% can be accessed at a more moderate rate. This is why data tiering was created—it provides the ability to align stored data to repositories that best match the performance need.

Many legacy vendors offer tiering as an option for their solutions because the practice enables better overall TCO for a storage ecosystem. Unfortunately, legacy data tiering solutions come up short on an important part of that TCO. Their approach has involved moving data between separate physical tiers to solve the challenge, which lowers the initial capital expenditures (CAPEX) required to deliver necessary space. However, this also introduces the unexpected operational cost burdens of managing separate control planes for two separate physical repositories for hot and cold data, paying for licensing to make it work, and introducing delays in the data retrieval process from network or array latencies due to data movement. This burden, if not accounted for, will quickly reduce or even negate the original TCO (CAPEX) benefits of a tiering solution.

How Is Pure Storage's Zero Move Tiering for FlashBlade Different?

Our approach to data tiering dramatically lowers the operational burden and complications of managing different classes of storage while maintaining the CAPEX savings. ZMT is native to Purity//FB working **continuously** behind the scenes to balance and deliver workload data from higher- and lower-performing storage classes, **all from the same namespace** on the FlashBlade array **with zero movement between tiers**. At launch, file systems can be manually pinned to hot and cold tiers to be delivered based on their assigned performance class. Policy-based tiering will soon follow in a future update. Our approach increases data tiering efficiency and avoids the need for admins to devise complex data movement policies over the network between separate storage endpoints to ensure a more predictable performance balance between hot and cold data sets. ZMT provides differentiated SLAs at the file or object level and avoids any latency due to reassembly of blocks.

ZMT also eliminates complicated node-based data movement licenses, making its deployment and management an integral Purity//FB feature. All Purity features with ZMT are available, such as snapshots, replication, [SafeMode™](#), WORM immutability, and data-at-rest encryption. Finally, ZMT ensures data is accessible from the same namespace using the same endpoints, which offers the added benefit of keeping the storage footprint small to reduce power and cooling.

FlashBlade® Zero Move Tiering
Single Cluster and Namespace

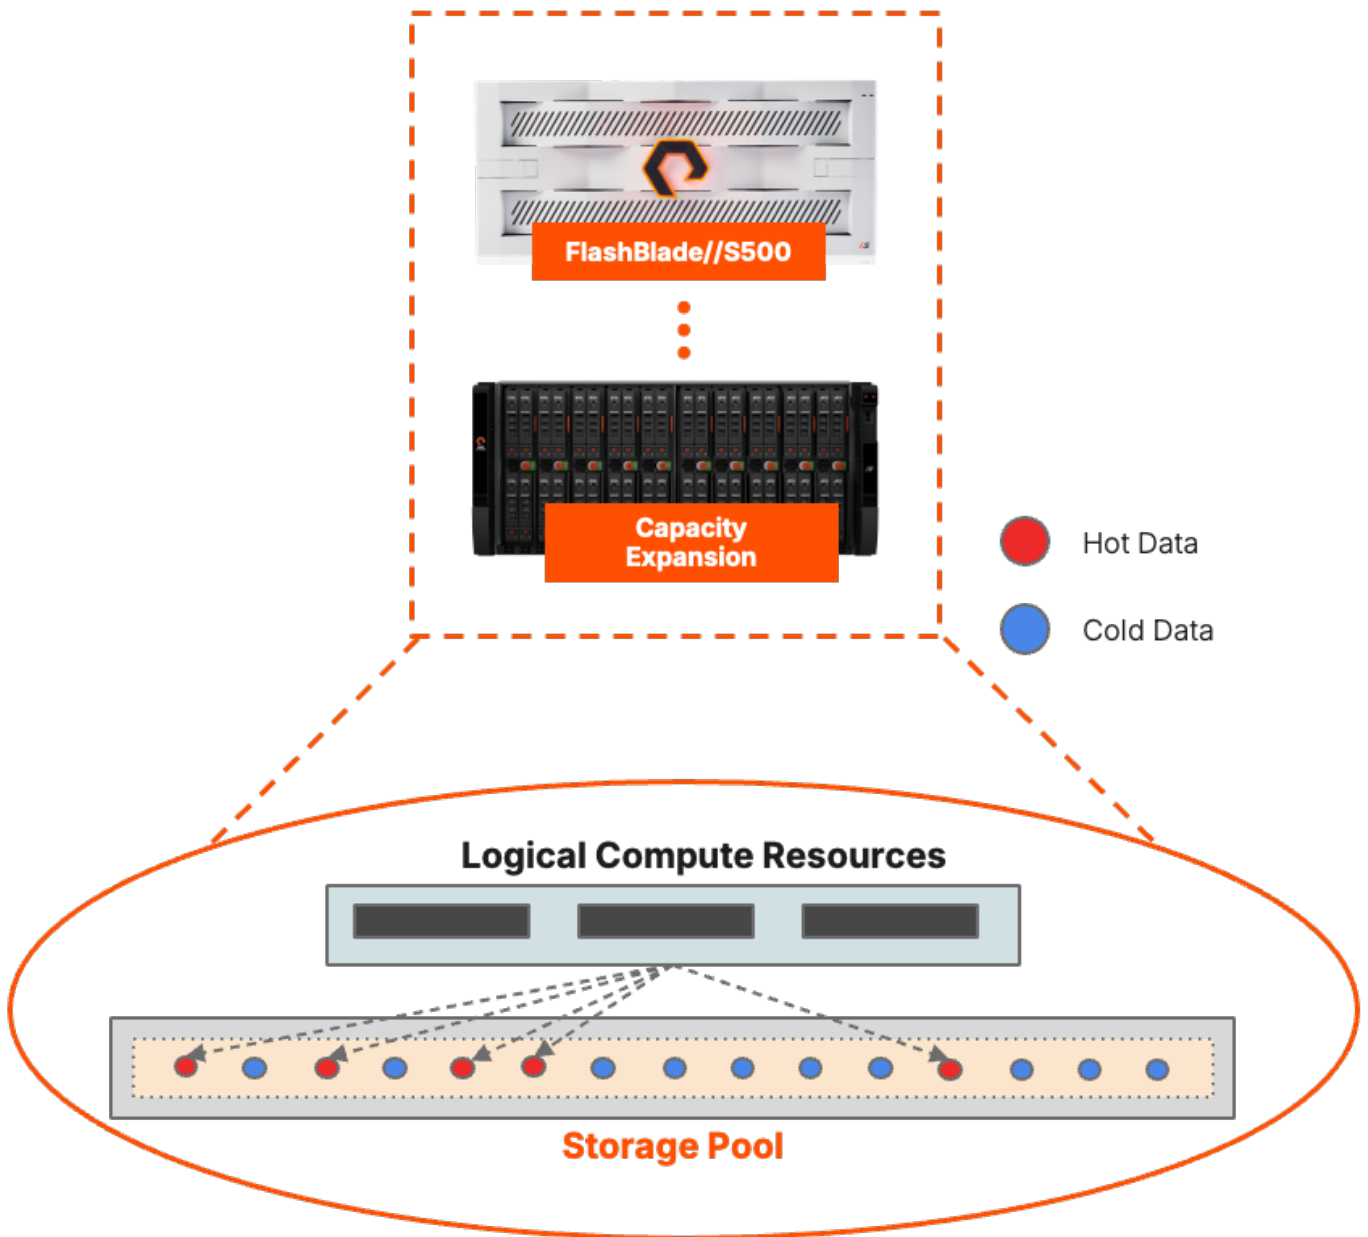


Figure 1: Logical diagram of Zero Move Tiering for FlashBlade.

Compelling Use Cases for ZMT

ZMT will provide instantaneous value to workloads that rely on high-throughput storage for one function but can leverage a lower-performance class for related functions. For example, AI training workloads will have better performance balance from active data models being serviced with a faster SLA than where previous models would be stored for occasional reference. [EDA workflows](#) can also benefit from ZMT by leveraging the hot tier for design and modeling scratch space, while lower-performance storage houses tertiary data that supports the process. Finally, data analytics workloads are great candidates because they leverage hot tiers for processing queries but can use less performant storage for the rest of the data set.

Get to Know More about Zero Move Tiering for FlashBlade

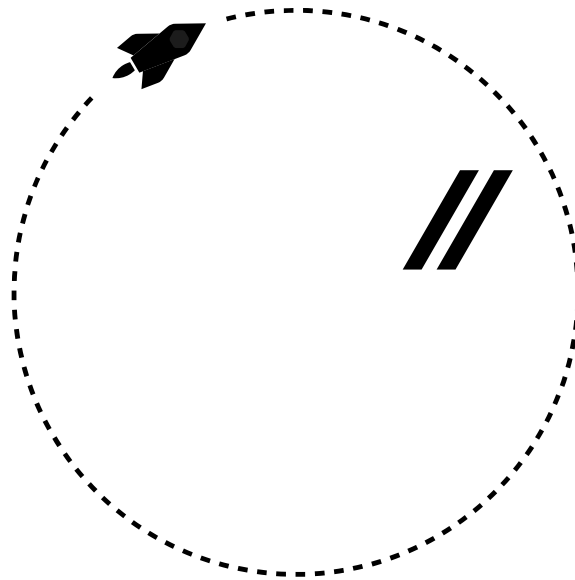
ZMT will be available on launch for the highest-performing workloads that also need large cold data sets to be within the same namespace for quick access. It will be offered as an expansion to FlashBlade//S500 and will allow pinning and transitioning of whole filesystems. Bucket transition for object storage as well as granular performance differentiation for individual files and objects will be offered later.

Contact your Pure Storage account team today to learn more about how ZMT can make costly and complicated legacy tiering solutions a thing of the past.

Pure//Launch

Register for our October 30 webinar to learn more about the next era of storage.

[Register Now](#)



WRITTEN BY:

[Don Poorman](#)

Pure//Launch

Join the webinar on October 30 to learn more about this and other innovations from Pure Storage.

Attend the Webinar

