

Accelerate applications dynamically

HPE 3PAR Adaptive Flash Cache

Use flash capacity to adaptively extend system cache, accelerate workloads, and improve service levels.

Flash-optimized for hybrid arrays

HPE 3PAR Adaptive Flash Cache is included as part of the HPE 3PAR Operating System Suite version 3.2.1 and later, and is supported on all HPE 3PAR StoreServ Storage arrays that have a blend of solid-state drives (SSDs) and hard disk drives (HDDs).

Benefits

- **Reduced latency**
Improves performance for random read-intensive workloads
- **Increased throughput**
Accelerates write throughput in mixed workloads
- **Responds dynamically**
Provides smart and adaptive data placement based on application and workload demands
- **Supports multi-tenancy**
Enables HPE 3PAR Adaptive Flash Cache across the entire system or select particular workloads to accelerate—the choice is yours

Serving more with less

A major goal for both application consolidation and multitenancy is the opportunity to drive down the overall cost of serving multiple I/O-intensive workloads. One of the most effective methods for accelerating read-intensive applications is increasing system cache. However, DRAM cache is expensive, and your system only supports a finite amount of it. HPE 3PAR Adaptive Flash Cache¹ allows you to extend your system's cache virtually—by leveraging flash capacity from your array's SSD tier—to accelerate performance dynamically and flexibly by just leveraging Flash.

HPE 3PAR Adaptive Flash Cache² is a software feature that can as much as half the application read response times by significantly reducing latency to accelerate application performance and improve service levels.

Accelerating read-intensive workloads

Random read-intensive workloads are associated with many common applications and environments, including server virtualization deployments, online transaction processing (OLTP) databases, and Web servers. Due to the system's Mesh-Active design, system-wide striping of data across resources, and enhanced caching algorithms, HPE 3PAR StoreServ Storage is able to deliver extremely high performance while relying on a minimal amount of DRAM cache.

In addition to these built-in read acceleration mechanisms, another way to augment application performance is to place data for read-intensive applications onto your array's high-performance SSD tier. Yet in some environments, this may not be sufficient. For example, what if you want to accelerate read-intensive workloads across the system and not just for a particular application?

HPE 3PAR Adaptive Flash Cache provides you with an additional tool to accelerate your read-intensive workloads by allowing you to use SSD capacity to extend system DRAM cache virtually. Unlike migrating I/O intensive applications to SSDs, this solution gives you the ability to accelerate read-intensive workloads across your array—not just those applications whose data resides on your system's SSD tier. In I/O-intensive environments, this can make all the difference when it comes to your ability to ensure application quality of service (QoS) and meet service-level agreements (SLAs), particularly in multitenant environments.

With this feature, you can extend system cache by up to 8 TB per controller node pair (depending on your array model) while maintaining the flexibility to enable this feature system-wide or only for selected applications or workloads. If you choose, you can maintain full control over allocation of system resources, QoS, and placement of application data onto the system's SSD tier, so you don't have to compromise any of the multitenancy benefits offered by your array.

¹ Adaptive Flash Cache is supported on HPE 3PAR StoreServ 8200, 8400 and 20800 Storage systems.

² Up to 32 TB of Flash Cache support is available on HPE 3PAR StoreServ 20800 Storage system.

Turning flash into cache

Just as storage controllers are limited for DRAM cache they can support, spinning media is limited in the performance it can deliver. However, by using SSDs to extend system cache, it is possible to leverage the low latency and high bandwidth offered by flash-based media while taking advantage of the larger and more economical capacities offered by HDDs.

HPE 3PAR Adaptive Flash Cache gives you the ability to extend your array's on-board DRAM cache by using a portion of the system's SSD capacity. Extending DRAM cache increases read cache hit percentage and accelerates host I/O, decreasing backend requests to the HDD tier. As fewer random read requests hit the HDD tier, more DRAM resources will be available for writes.

Redirecting a high percentage of host I/O from DRAM cache to SSD capacity reduces latency as well as backend load, resulting in increased application performance with less reliance on high-performance memory. Memory is therefore freed up to service other requests while host I/O is accelerated.

With HPE 3PAR Adaptive Cache enabled, the system satisfies read requests by using DRAM cache up to a threshold of 90 percent utilization, at which point small blocks of random read data are automatically destaged to the virtualized cache extension created from SSD capacity. Similarly depending on the I/O request small block random write requests will be de-staged from DRAM to Adaptive Flash Cache.

³ HPE 3PAR Adaptive Optimization Software is an optional software product sold as part of the HPE 3PAR Data Optimization Software Suite.



Sign up for updates

★ Rate this document

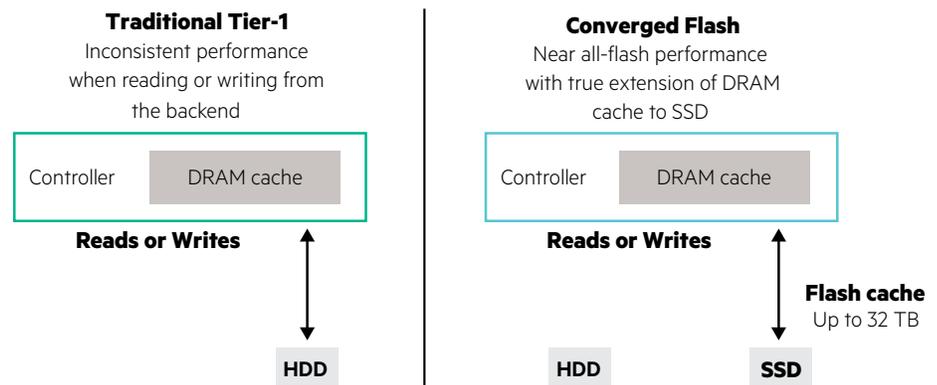


Figure 1: Accelerate Performance with HPE 3PAR Adaptive Flash Cache

Getting the most from your flash investments

One of the most robust benefits of HPE 3PAR Adaptive Flash Cache is the feature's compatibility with HPE 3PAR Adaptive Optimization software.³

On its own, HPE 3PAR Adaptive Optimization is a powerful tool for increasing IOPS and reducing latency with minimal investment. Adaptive Optimization operates on the principle of I/O access density, which is based on the observation that for most applications, a very small amount of the total user data receives most of the I/O accesses. By configuring only the very active data blocks on your SSD tier and the rest of your application data onto more economical SAS or nearline SAS storage, HPE 3PAR Adaptive Optimization can deliver higher performance at a lower cost, particularly in cases where more than 90 percent of the IOPS being served are relying on as little as five percent of your system's allocated capacity.

HPE 3PAR Adaptive Optimization Software dynamically moves data between your array's storage tiers based on analysis of data access patterns, including both reads and writes. As already described, HPE 3PAR Adaptive Flash Cache uses SSD capacity to extend on-board DRAM cache in real time to

accelerate random read requests. Depending on your configuration, these two processes can share the same SSD capacity tier, thereby enabling you to get the most out of your flash investment.

Refer to the [Adaptive Optimization for HPE 3PAR StoreServ Storage white paper](#) to learn more about Adaptive Optimization and this combined solution.

See for yourself

With the current version of the HPE 3PAR OS, you can start using HPE 3PAR Adaptive Flash Cache today to accelerate applications non-disruptively, without physically extending system cache or purchasing additional hardware or software. Starting with as few as four SSDs, you can improve performance and achieve predictably lower latency.

To see for yourself how this solution can benefit your particular environment, take advantage of the simulation mode built into HPE 3PAR Adaptive Flash Cache. This simulation mode analyzes your specific workloads and can give you an accurate picture of how HPE 3PAR Adaptive Flash Cache can benefit your specific application workload.

Learn more at hp.com/go/3parstoreserv