Technical white paper

HP SmartCache technology

Table of contents

Abstract

As an IT professional, you need cost-effective ways to increase the performance of applications. However, as CPU core counts continue to grow, and addressable memory space increases, storage can become a bottleneck for virtualized and data-intensive workloads.

HP ProLiant Gen8 servers address application performance problems by including dynamic workload acceleration—an architectural innovation that employs HP SmartCache technology to accelerate access to multi-terabytes of data in a direct-attached storage (DAS) environment.

This paper describes how HP SmartCache technology converges solid-state drive (SSD) and hard-disk drive (HDD) technologies with intelligent control to offer high performance and high capacity without incurring the cost of an all-SSD configuration.

Introduction

Enterprise storage requirements are growing rapidly in terms of performance and capacity. While most applications need a balance of those attributes, some applications do not. For instance, applications that process many read and write requests to the storage subsystem need greater performance. Applications that manage archival data that grows continuously need greater capacity.

Two storage technologies commonly address application needs. The SSD, with very low latency, is the answer for performance sensitive applications. The HDD, with long-term reliability and economical cost per GB, meets large-capacity needs. Figure 1 illustrates the suitability of SSDs and HDDs for meeting the needs of various applications.

- High-frequency trading applications require very low latency. An all-SSD solution provides the performance this environment requires.
- Online transaction processing (OLTP) and analytics push the boundaries of both performance and capacity growth. A mix of SSDs and HDDs best meets those needs.
- Archival and smartphone applications require unstructured data growth; this need is best addressed with high-capacity HDDs.

Figure 1. Different applications have varying storage needs.

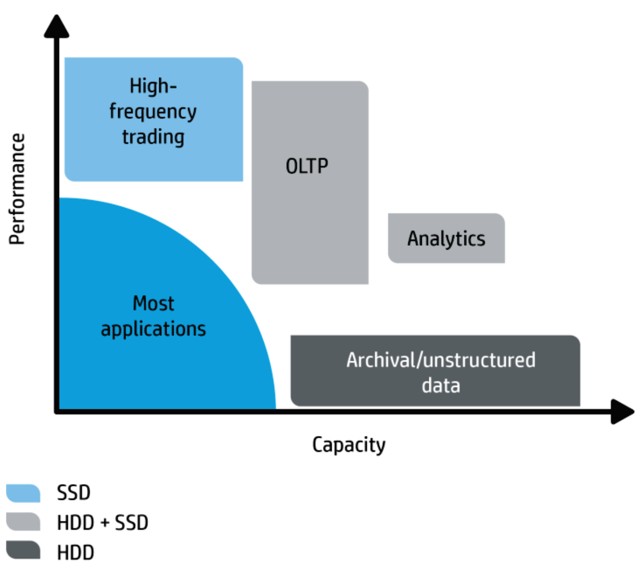

Comparing storage technology performance

SSDs provide significant performance advantages, as compared to traditional HDDs. Figure 2 illustrates the relative performance in IOPS and latency for a high-performance HDD, a mainstream SSD, and an HP IO Accelerator. (The HP IO Accelerator is an interface card containing flash memory and a custom controller that significantly improves data throughput.)

As shown in Figure 2, for each vertical bar, the top number indicates the IOPS for 100 percent random reads; the bottom number indicates the IOPS for 100 percent random writes. The chart illustrates how the mainstream SSD and HP IO Accelerator vastly outperform the HDD, due to the IO Accelerator's ability to read and write data much faster than the HDD.

Figure 2. IOPS and latency performance of storage hardware.

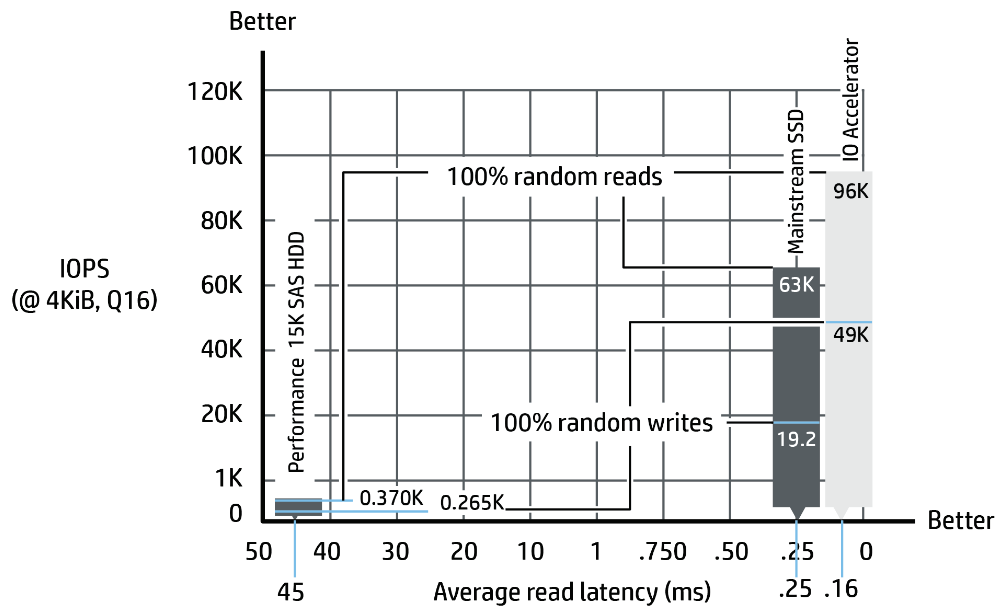

Benchmark Iometer 2006

SSD technology is the foundation for both the HP IO Accelerator and SATA SSD; however, their performance capabilities differ. The underlying SSD implementation in the HP IO Accelerator provides faster read and write performance:

- The HP IO Accelerator's read performance is roughly 35 percent higher than that of the mainstream SATA SSD.
- The HP IO Accelerator's write performance is roughly 2.5X of the mainstream SATA SSD.

HP IO Accelerators and SSDs obviously address performance requirements, but the cost might be prohibitive for applications that require high capacity and scalable performance.

What about hybrid drives?

Hybrid drives combine a small SSD with traditional HDD technology in an integrated package. These drives are a good option in the consumer market. In the server space, which is characterized by numerous HDDs, hybrid drives lack the flexibility to scale SSD technology differently, based on application needs. If an application references data from only one drive out of dozens in a hybrid drive configuration, performance will suffer because of the small SSD capacity on the single drive. In addition, the unused SSD capacity on the other hybrid disks becomes a stranded resource. With a larger SSD dedicated for caching, performance scales more effectively as capacity or the number of HDDs increases.

Why caching?

You may be wondering if SSDs are superior in performance, why not use HP IO Accelerators and SSDs in all server configurations. The answer is simple—you need to weigh the requirements for capacity and performance against the increased costs for those abilities. As the cost of SSDs is trending downward, SSD technology is making inroads into the enterprise server market to address the low-latency requirements for certain applications. However, even with the performance benefit and lower cost, SSD technology will not replace HDD technology in the near future because HDDs remain an economical and durable (long-term) storage medium. Also, SSDs support a finite number of write cycles before beginning to degrade. Even so, SSD durability is improving and the capacity is increasing to where SSDs are a more viable solution for enterprise applications.

Using a high-performing all-SSD configuration might be feasible for applications with modest capacity requirements, but it is difficult to ignore HDDs for their high capacity and very economical price—particularly in enterprise systems that require thousands of storage devices.

To ease your storage decisions, HP SmartCache technology provides an incremental, economical way to introduce SSD technology into your environment. HP SmartCache provides SSD performance benefits, while simultaneously preserving your investments in previously purchased hard disk assets.

How does caching work?

Intelligent caching places data in a lower-latency device so responses to future requests for that data can occur much faster. If an application requests data that is in the cache (called a “cache hit”), then the lower-latency device can service that transaction. Otherwise, a “cache miss” occurs and the data must be accessed from the original, slower device. As more cache hits occur, overall performance improves.

Caching technology has been proven in many areas of computing, due to the nature of how applications access data.

- Processors use faster memory caches to speed up instruction execution as they fetch data from memory DIMMs.
- HDDs contain caches (also known as buffers) that allow data to be queued to and from the magnetic media.
- Host operating systems utilize caches to improve application performance.

SSDs bring caching to storage systems to provide a cost-effective performance boost, especially to server configurations.

HP SmartCache architecture

HP SmartCache combines different technologies and device types to close the cost/performance gap. In the HP SmartCache architecture, a copy of the data resides on the HDD, as well as on the lower-latency SSD used for caching. The basic HP SmartCache architecture includes the following three elements:

- Bulk storage—any supported HDD or SSD that can be attached to an HP Smart Array controller.
- Accelerator—a lower-latency SSD device that caches data; the capacity of the accelerator is less than that of the bulk storage device.
- Metadata—information held in a relatively small storage area of the Flash-Backed Write Cache (FBWC) memory that maps to the location of information residing on the accelerator and bulk storage devices.

The flexible HP SmartCache architecture supports numerous device types for bulk storage, accelerator, and metadata. The HP SmartCache architecture can adapt to accommodate the evolution of storage devices.

HP SmartCache for DAS

The direct-attached HP SmartCache solution includes the three elements of the HP SmartCache architecture; HDDs serving as bulk storage, SSDs as the accelerator, and Flash-Backed Write Cache (FBWC)[1] memory for storing the metadata (see Figure 3). For this implementation, the SmartCache control layer resides in the firmware of the onboard Smart Array Controller of the HP ProLiant Gen8 Server, below the operating system and driver. This architecture allows caching for devices connected to a single array controller.

Figure 3. HP SmartCache can be applied to DAS.

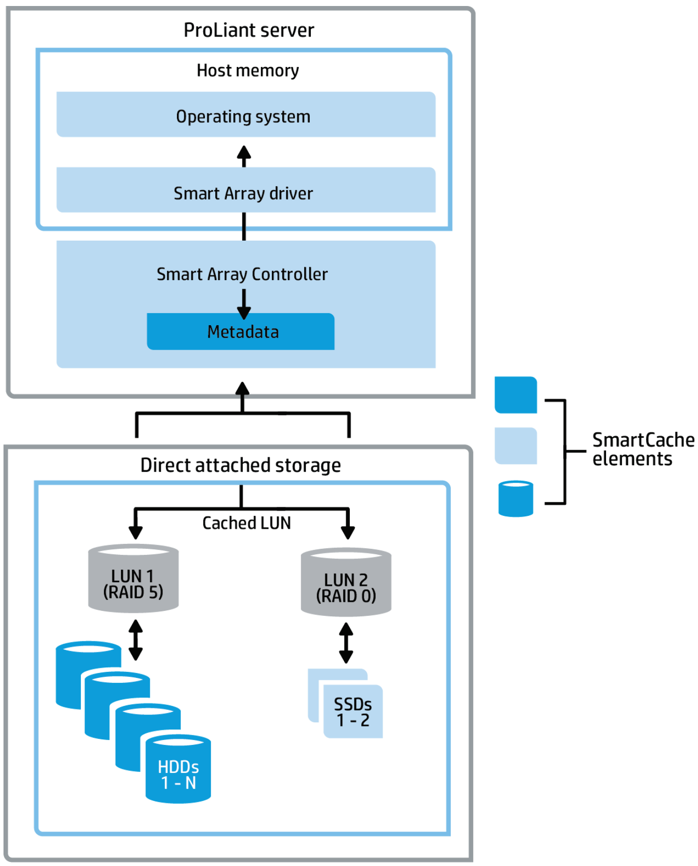

HP SmartCache offers flexibility in creating logical disk volumes from HDDs.

- The bulk storage volume may be any RAID configuration supported by the Smart Array Controller.
- Each bulk storage volume can have its own cache volume or none at all.
- Cache volumes can be created and assigned dynamically without adversely impacting applications running on the server.
- Cache volumes may be only RAID 0.

Only SSDs can be used for cache volumes, and a cache volume can only be assigned to a single logical disk volume. The HP SmartCache solution consumes a portion of the FBWC memory module on the Smart Array Controller for metadata. To ensure sufficient storage for accelerator and metadata, we recommend:

- 1 or 2 GB of FBWC memory, or
- Roughly 1 GB of metadata space for every terabyte of cache volume space

When using the HP SmartCache solution for direct-attached storage, legacy cache is still present and operational; it utilizes the remaining space in the FBWC memory. When using HP SmartCache, the Smart Array DDR cache will be set to 100 percent write operation which enables write-back functionality in DDR cache. This allows SSDs configured as HP SmartCache volumes to provide a much larger read cache.

[1] FWBC is an integrated persistent memory device used for caching operations conducted within the Smart Array Controller's firmware. Writes are posted to the FBWC, acknowledged immediately, and later completed to the actual HDDs for improved latency of write operations.

Since the HP SmartCache solution is implemented within the Smart Array firmware, it can be used on any operating system supported by HP ProLiant servers. The HP SmartCache solution operates transparently, with no dependencies on operating system, device driver software, file system type, or applications.

HP SmartCache supports write-through caching. When an application writes data to the disk, the Smart Array Controller writes the data to the HDDs; if the data is also in the cache volume, the data is written to the SSD.

Management and analytics

To configure HP SmartCache, use the HP Smart Array configuration utilities: HP Smart Storage Administrator (HP SSA), HP SSA command-line (SSA-CLI), HP Array Configuration Utility (ACU), or HP ACU command-line version (ACU-CLI). These utilities incorporate analytics to provide more visibility into the disk storage subsystem. This way, you can see what is happening within HP SmartCache, as well as view detailed information on both the physical and logical disks. In addition, integration with HP Insight Control comes standard, providing information and alerts about HP SmartCache status.

Future capabilities will allow you to profile your applications, and then model the HP SmartCache performance through a simulator. This will allow you to make various accelerator substitutions, according to capacity and performance capabilities.

Case study: OLTP workload

Data integrity is the most important characteristic for a disk storage subsystem. HP SmartCache technology provides a reliable, enterprise-class solution to improve data integrity. Enterprises are most interested in HP SmartCache because of its performance (dependent on the application) and price-performance characteristics. HP SmartCache improves application performance when:

- There is more read traffic than write traffic.
- Data is read repeatedly, which is common in many applications.

For our OLTP workload case study, we ran a Transaction Processing Performance Council type E (TpsE)-like[2] benchmark on an HP ProLiant DL380 Server with 48 GB of memory, an HP Smart Array P421 Controller, and a 2-GB FBWC module configured as 100 percent write cache.

Figure 4. The configuration used for the TpsE-like case study.

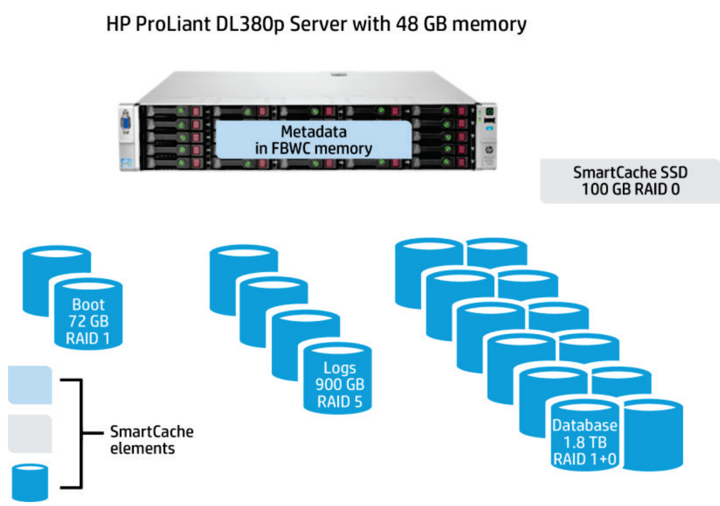

[2] Modified TPCe workload as defined in the presentation titled, “Using Solid State Drives As a Mid-Tier Cache in Enterprise Database OLTP Applications,” by Badriddine Khessib, Microsoft® for TPC Technology Conference 2010—Singapore (tpc.org/tpctc2010/TPCTC2010-12.pdf).

The Smart Array P421 Controller disk configuration included the following:

- Boot volume—2 x 72 GB 10K 6G SAS drives configured in RAID 1
- Database logs—4 x 300 GB 10K 6G SAS drives configured in RAID 5
- Database—12 x 300 GB 10K 6G SAS drives configured in RAID 1+0
- HP SmartCache—1 x 100 GB 6G MLC SSD drive configured in RAID 0. The HP SmartCache was configured in write-through mode and assigned to cache the database volume.

The database size was roughly 1.6 TB during the two-hour test run. Figure 5 shows the performance of the TpsE-like workload with HP SmartCache (grey) and without HP SmartCache (blue). TpsE performance of the TpsE-like workload was significantly better when using HP SmartCache. The SSD used as a cache was less than seven percent the capacity of the database. The cache-hit ratio of HP SmartCache was 84 percent, meaning that roughly five out of every six disk reads to the database volume were serviced by the SSDs vs. the HDDs. A smaller amount of HP SmartCache would have yielded a smaller cache-hit percentage and lower performance gains.

Figure 5. TpsE-like workload performance with and without HP SmartCache.

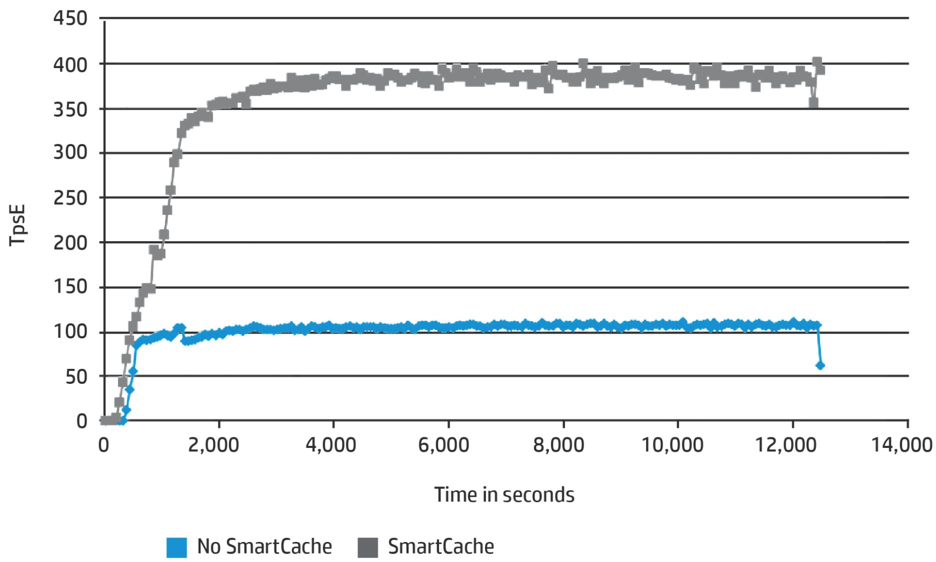

Table 1 compares the database logical volume I/O statistics without and with HP SmartCache. With the addition of HP SmartCache, performance results change dramatically, simply by adding a 100 GB SSD.

Table 1. TpsE-like workload statistics

I/O analysis	Without HP SmartCache		With HP SmartCache	
Accesses	32,767,015		159,239,037 (almost 5X more)	
Reads	29,114,040		140,159,417 (almost 5X more)	
Writes	3,652,975		19,079,620 (5X more)	
Percent reads	89		88	
Percent writes	11		12	
Average IOPS, read	4,382		19,464 (more than 4X better)	
Average IOPS, write	550		2,650 (almost 5X better)	
Duration (seconds)	6,644		7,201	
I/O distribution (reads)	**Count**	**Avg. latency (ms)**	**Count**	**Avg. latency (ms)**
0% at 4 k	0	0.00	0	0.00
97.6% at 8 k	28,430,311	27.97	140,159,417	6.23
0% at 16 k	71	28.21	0	0.00
<1% at 32 k	1,457	21.54	0	0.00
2.3% at 64 k	682,201	18.79	0	0.00
0% at 128 k	0	0.00	0	0.00
0% at 256 k	0	0.00	0	0.00
100% (total)	29,114,040	27.75	140,159,417	6.23
I/O distribution (writes)	**Count**	**Avg. latency (ms)**	**Count**	**Avg. latency (ms)**
<1% at 4 k	93	10.18	98	45.39
82.9% at 8 k	3,028,218	0.13	18,015,518	89.35
16% at 16 k	585,670	0.15	1,031,087	89.27
<1% at 32 k	18,322	0.22	18,745	89.90
<1% at 64 k	13,607	0.29	11,419	90.06
<1% at 128 k	5,856	0.36	2,326	90.01
<1% at 256 k	1,209	0.46	427	90.10
100% (total)	3,652,975	0.13	19,079,620	89.35

Taking a closer look at the data from Table 1, we see the transaction volume through the database logical drive increased significantly with HP SmartCache. As shown in Figure 6, the disk subsystem without SmartCache was able to process approximately 29 million reads and 3.6 million writes (blue). With HP SmartCache, the disk subsystem performance increased significantly to approximately 140 million reads and 19 million writes (gray) within roughly the same period. With HP SmartCache, the disk subsystem was able to process approximately four times as many requests than without HP SmartCache.

Figure 6. Transaction volume through database volume.

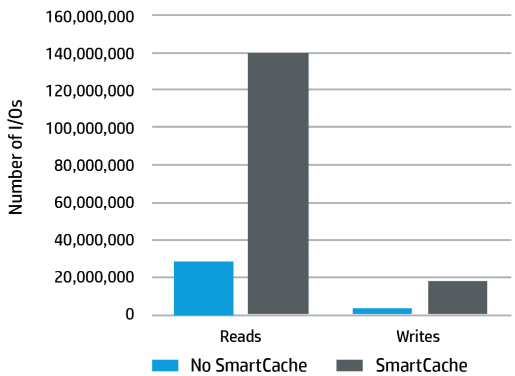

Looking at the data another way, we see the differences in the number of IOPS with and without HP SmartCache (Figure 7). Without HP SmartCache, the disk subsystem was able to process approximately 5,000 requests per second, combining both read and write traffic. With HP SmartCache, the disk subsystem was able to process approximately 22,000 requests per second, with both reads and writes. These figures show the disk subsystem with HP SmartCache is able to process four times as much data than the system without HP SmartCache.

Figure 7. HP SmartCache supports increased IOPS.

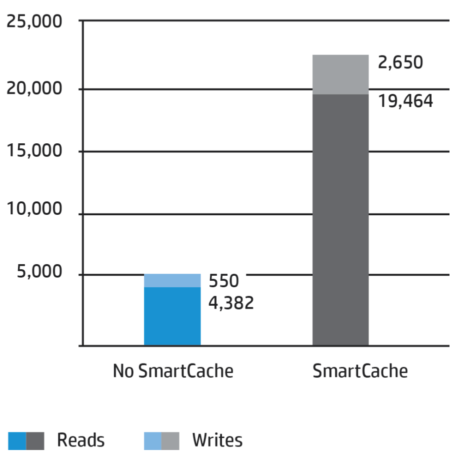

Furthermore, without HP SmartCache, we can see that the average read latency per request was almost 28 milliseconds (ms), and the average write latency was practically non-existent at 0.13 ms (Figure 8). This workload highlights the benefit of the FBWC of the Smart Array Controller. The FBWC of the Smart Array Controller completes the write operation by saving the data within the FBWC memory, which is much quicker than writing the data to the HDDs.

Once data is saved in the FBWC, the Smart Array Controller ensures the data is written to the HDDs. The primary function of the FBWC is to protect the data in the FBWC in case of unplanned power loss to the server. From an application workload perspective, the write load of the application on the Smart Array Controller is not high enough nor sustained long enough to fill the capacity of the FBWC, so the write I/O response time is excellent.

Figure 8. Average latency per I/O.

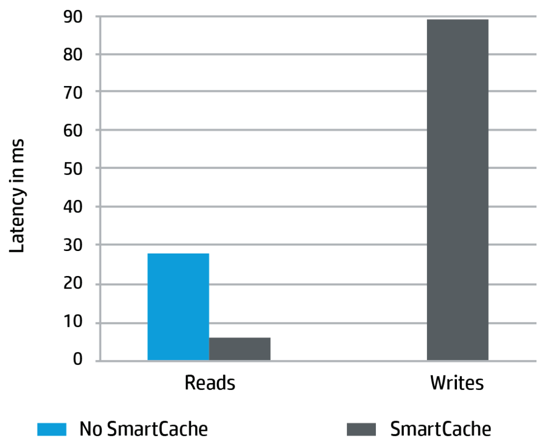

Figure 9 also shows that read latency drops from almost 28 to 6.23 ms per request when HP SmartCache is enabled. This behavior is expected as more reads are satisfied by the accelerator device or SSD, which further demonstrates the value of HP SmartCache. What might not be expected is the increase in write latency when HP SmartCache is enabled, which will be explained shortly.

The distribution of latencies per transaction for 8 KiB read traffic (Figure 9) shows the benefits of HP SmartCache. With HP SmartCache enabled, 95 percent of the read requests completed in less than 10 ms . Without HP SmartCache, the distribution of latencies is distributed across the spectrum of response times because there is more dependency on the rotation access time of the HDDs. This rotational latency results in a near 28 ms average per read request.

Figure 9. 8 KiB read latency distribution.

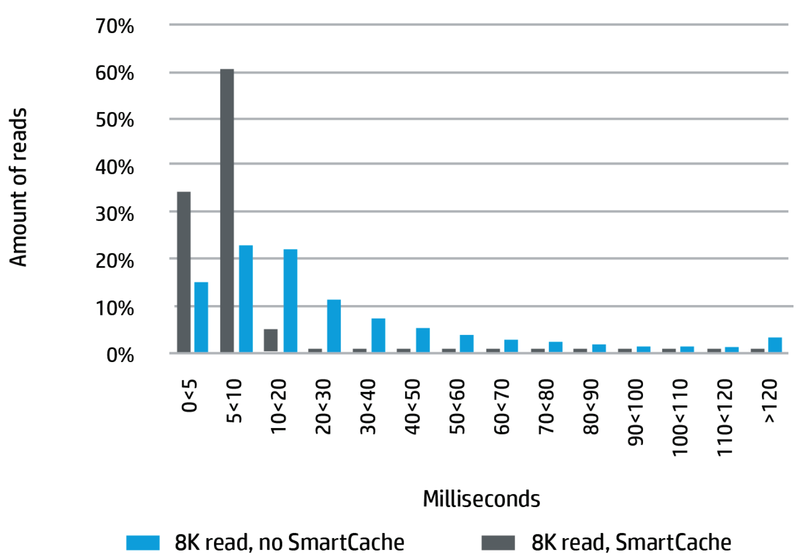

Write latency jumped significantly from an average of 0.13 ms without HP SmartCache to 89.35 ms with HP SmartCache (Figure 10). Without HP SmartCache, the write requests were completed with low latency by the FBWC. With HP SmartCache, the increase in read requests serviced by the faster SSD results in a workload shift on the HDDs from mostly reads to mostly writes.

Also, note that using HP SmartCache results in the volume of writes from the application exceeding the FBWC capacity. When the FBWC memory is full, the write latencies increase because write operations cannot complete until data is written to the HDDs. This explains the increase in write latencies when using HP SmartCache. However, the benefit of substantially increased read performance far outweighs the increase in write latency, resulting in better overall performance to the application.

Figure 10. 8 KiB write latency distribution.

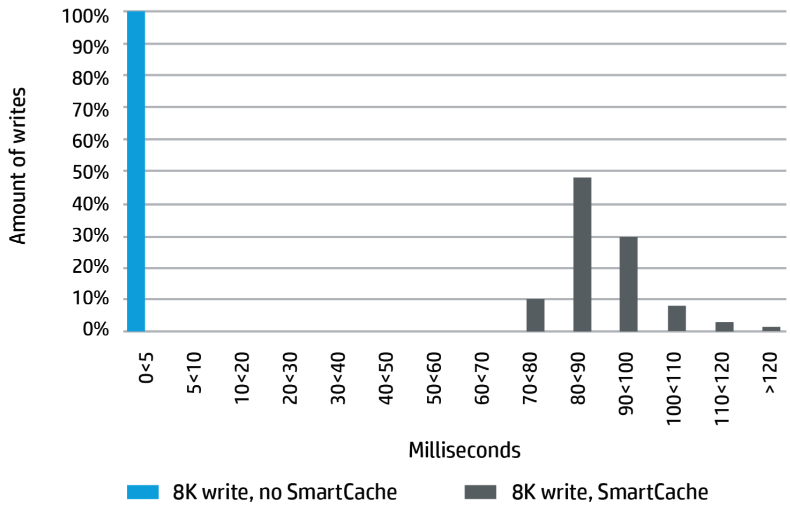

Conclusion

While other applications such as databases, analytics, and virtual desktops can benefit from acceleration techniques, TpsE-like workloads present the biggest challenge—and HP SmartCache successfully overcomes it. The characteristics of TpsE—more read traffic (88 percent) than write traffic (12 percent) and repetitive data requests—allow HP SmartCache to enhance performance, while also significantly reducing storage costs.

Other types of workloads might not benefit from HP SmartCache, including those with the following characteristics:

- Non-repetitive I/O—Applications that do not access data multiple times over a period of time. Examples include routine backups, because they only access every file once; and disk benchmark utilities, because their disk access patterns are typically random.
- Write-heavy I/O—Applications that perform predominantly write traffic vs. read traffic, such as database logging volumes, because the majority of traffic to these volumes is writes to the HDDs.

HP SmartCache is the cost-effective, scalable, high-performance solution for accelerating data access for multi-TB data sets for DAS environments.

Resources and additional links

HP SmartCache
hp.com/go/smartcache

HP SSA
hp.com/go/hpssa

HP Dynamic Workload Acceleration
h20195.www2.hp.com/v2/GetPDF.aspx/4AA3-9649ENW.pdf

HP PCIe IO Accelerator for ProLiant servers
h18006.www1.hp.com/products/servers/proliantstorage/solid-state/benefits.html?jumpid=reg_r1002_usen_c-001_title_r0007

Learn more at
hp.com/go/smartcache

4AA5-1706ENW, April 2014