

A photograph of two men walking through a server room. The man on the left is wearing a blue shirt and grey trousers, pointing towards a server rack. The man on the right is wearing a red shirt and dark trousers. The server racks are black with yellow handles and are filled with server components. The room has a white ceiling with exposed pipes and lights.

Brochure

Adaptive Data Reduction

HPE 3PAR StoreServ Storage with Zero Detect, Deduplication, Compression, and Data Packing



Hewlett Packard
Enterprise

Adaptive Data Reduction technologies and benefits:

Zero Detect

- First-level, preventative data reduction technology.
- Autonomic and automatic—you get all the benefits by doing nothing.
- Unique to HPE 3PAR StoreServ Storage and built into the HPE 3PAR ASIC.
- Hardware-embedded, so there is no impact to the system.
- Can be combined with other data reduction technologies for additive benefit.

Deduplication

- Second-level, preventative data reduction technology.
- Unique HPE 3PAR implementation features deduplication engines built into the HPE 3PAR ASIC.
- Selective nature allows per-volume granularity when enabling deduplication.
- Hardware-assisted, so there is minimal impact to system operations.
- Can be combined with other data reduction technologies for additive benefit.

Compression

- Complementary, always-inline data reduction technology.
- Unique HPE 3PAR implementation identifies incompressible data and stores it in its uncompressed state, saving CPU resources and reducing latency.
- Can be enabled on a per-volume basis in the management console for greater system efficiency.
- Increases endurance of solid-state media by reducing written data.
- Can be combined with Zero Detect and deduplication for maximum cost reduction benefit.

Data Packing

- Advanced compaction technology unique to HPE 3PAR StoreServ.
- Always inline, works to virtually eliminate garbage collection typically associated with data reduction.
- Enabled automatically when data reduction is enabled.
- Increases system efficiency, performance and further improves endurance by writing at flash-native block sizes.

Reduce storage costs with Adaptive Data Reduction for total system efficiency

Adaptive Data Reduction is a collection of technologies that come standard with HPE StoreServ Storage which are designed to reduce your data footprint. When used alone or in combination, these technologies help you get the most out of your system's flash capacity and reduce your total cost of storage while improving flash media endurance. When paired with the industry's largest flash drives,¹ Adaptive Data Reduction technologies change the economics of flash, making it possible to realize performance and cost benefits for every application and service.

This brochure is your guide to understanding Adaptive Data Reduction technologies for HPE 3PAR StoreServ Storage arrays, their unique advantages, and how to combine them for additive benefit.

Zero Detect

Zero Detect is a data reduction technology unique to HPE 3PAR StoreServ Storage driven by one of the many dedicated engines built into the HPE 3PAR ASIC. Designed to reduce the cost of storage by identifying and removing repeated data from incoming data streams in real-time, this hardware-accelerated capability reduces the amount of capacity required to store data without impacting performance because operations take place at the hardware layer, not in software.

During normal operations, hosts often write an extended string of zeros to a storage array as part of a write stream. Zero Detect is the feature within the HPE 3PAR ASIC that examines all incoming write streams, identifies extended strings of zeros, and removes them—preventing unnecessary data from ever being written to storage. As a result, the duplicated data never consumes capacity on the array.

¹ HPE 3PAR StoreServ Storage supports 15.36 TB SSDs, the largest enterprise SSD available on the market based on publicly available information available as of date of publication.

How does the HPE 3PAR ASIC make a difference?

The HPE 3PAR ASIC contains dedicated, memory-speed engines for many functions, including RAID parity calculations, cluster communications, Zero Detect and more.

In fact, the HPE 3PAR ASIC is so fast and so efficient, it would take four dedicated 16-core CPUs to achieve performance on par with just the dedicated deduplication engines within a four-node, 800-series HPE 3PAR array—a collection of CPUs that would require 20X more power than the HPE 3PAR ASIC.

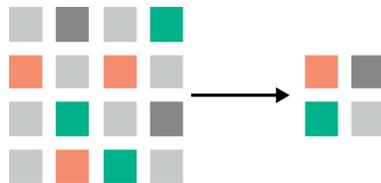
This is one of the reasons that Hewlett Packard Enterprise is able to offer such high-performance deduplication yet maintain extremely low, all-flash array power and cooling requirements.

Since Zero Detect is performed within the HPE 3PAR ASIC, not only do all operations take place inline and at wire-speed, but they consume no CPU cycles so do not impact system operations.

Zero Detect is considered the very first level of data reduction. In addition, because it is completely autonomic and hardware-embedded, Zero Detect works independently from other data reduction technologies, meaning savings can be made on all data and in combination with other data reduction technologies.

Deduplication

Like Zero Detect, deduplication on the HPE 3PAR StoreServ array uses the HPE 3PAR ASIC and is designed to reduce the amount of capacity needed to store data by minimizing the amount of data actually being written to storage.



However, unlike Zero Detect, the system is not looking for strings of zeros but data that is more complex and duplicative of data that has already been written to storage. To achieve this, the system assigns a unique “fingerprint” to write requests as they are processed by the array. The system saves these “fingerprints” for future reference so that when it receives a request to write the same data more than once, it is able to identify these subsequent requests as duplicative by cross-referencing their “fingerprints” against those it has already captured. A match reveals that the incoming request contains duplicative data, at which point the system performs a detailed verification check and then discards the duplicative data rather than writing it to storage, instead recording a “pointer” to the original data.

Deduplication on HPE 3PAR StoreServ Storage leverages the HPE 3PAR ASIC to perform tasks that are typically CPU-intensive. Each HPE 3PAR ASIC features three dedicated deduplication engines, each of which is capable of creating data “fingerprints” at memory speed. In addition, the HPE 3PAR ASIC is essential to delivering enterprise-class deduplication by ensuring that data sharing the same “fingerprint” identical at the bit-level, thus preventing corruption.

Futureproofing your flash investment

HPE 3PAR Data Packing is optimized for flash but also for next-generation solid-state technologies, including storage-class memory (SCM).

Next-generation memory technologies have different behavior from today's NAND flash—offering significantly higher performance and lower latency—but they also have much higher endurance and different write patterns.

Data Packing is already optimized for these future technologies by using an update-in-place algorithm instead of an allocate-on-write approach like many “built for flash” platforms. This, combined with flash-friendly page sizes, aligns well with future technology introductions.

This enables HPE 3PAR to adopt these next-gen memory technologies and gain the performance benefits they provide as soon as they become available, without needing to alter the architecture of the system.

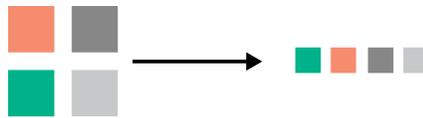
Like Zero Detect, deduplication on HPE 3PAR StoreServ Storage also takes place inline, meaning that it requires no post-process data operations that would result in duplicate data being written to storage and then scrubbed. This makes the HPE 3PAR implementation extremely efficient from both a performance and capacity perspective. When paired with Express Indexing—an HPE 3PAR technology that significantly accelerates data fingerprint lookup and comparison—deduplication on HPE 3PAR offers substantial cost, storage footprint, power, and cooling reductions. It also improves flash media endurance by reducing data writes.

Deduplication on HPE 3PAR arrays works extremely well over a wide range of data types, offering significant reductions in the amount of flash capacity required to store virtual server and virtual desktop workloads in particular. The HPE 3PAR implementation of deduplication on HPE 3PAR is also highly compatible with environments that store multiple copies of data—for example, for testing, development, and user acceptance testing (UAT). This alone can dramatically reduce both the cost of flash and the data center footprint, making flash immediately affordable for a variety of mainstream applications.

Deduplication on HPE 3PAR works independently from other data reduction technologies, meaning savings can be made in combination with Zero Detect and compression.

Compression

While Zero Detect and deduplication both reduce the amount of flash required to store data by eliminating unnecessary data, compression works by looking inside data streams for opportunities to reduce the overall size of the data set.



As with other data reduction technologies, the HPE 3PAR ASIC plays a key, although indirect, role in compression. Instead of using dedicated engines within the ASIC to perform compression, the HPE 3PAR ASIC is used to offload other resource-intensive operations from the CPUs, thus freeing them up to perform compression operations. In addition, the system spreads these compression operations broadly across multiple CPU cores in order to expedite data compression. Furthermore, a unique HPE 3PAR technology called Express Scan improves compression efficiency by identifying incompressible streams of data and storing them in their native formats instead of wasting CPU cycles by attempting to compress data that is incompressible.

Like all HPE 3PAR data reduction operations, compression runs inline for optimal efficiency. Since there are no post-process tasks to manage, and uncompressed data is never stored on the array, system resources are not required for cleanup or post-processing. This is not only important to increase the endurance of flash, but to ensure consistent performance by not needing to invoke resource-intensive post-process tasks.

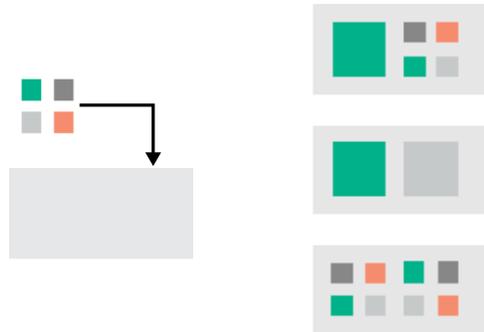
Compression complements Zero Detect and deduplication by adding a third method for reducing the amount of flash required to store a given amount of data. The combination of these three technologies is key to minimizing the cost of flash and making it an economical choice for nearly any application.

Optimize your IT investment strategy with new ways to acquire, pay for and use technology, in lock-step with your business and transformation goals.

hpe.com/solutions/hpefinancialservices

Data Packing

Increasing storage efficiency through deduplication and compression reduces the cost of flash by increasing the amount of data that can be stored on a given amount of capacity. But maintaining that efficiency system-wide over time is a much bigger challenge. To overcome this challenge, HPE 3PAR StoreServ Storage uses a unique software feature called Data Packing that combines data reduction and flash efficiency technologies to maintain peak capacity efficiency over time.



To understand how Data Packing works, it is important to understand that once data has been deduplicated and compressed, the result is a set of odd-sized “pages” in cache that are very inefficient to write to flash. Data Packing takes these random sized pages and packs them into small, fixed-size pages. This allows the system to attain a higher total system efficiency as compared to other all-flash platforms. The uniform, “packed” pages are set to a flash-native size, resulting in excellent efficiency but also excellent performance, as the resulting reads and writes to and from flash are performed at their internal page size. This also improves endurance as written data doesn’t cross multiple internal pages, resulting in very efficient use of flash pages.

Data Packing also packs together pages with good data locality, ensuring that excessive amounts of garbage aren’t created, in sharp contrast to many other implementations that require post-process garbage collection to tidy up the large amounts of garbage created by data overwritten by hosts. Reducing the need for resource-intensive garbage collection tasks has a positive impact on overall system performance. In addition, the use of Data Packing allows HPE 3PAR arrays to offer 100% inline data reduction with absolutely no post-processing, which is important to enterprise 24x7 environments that cannot accommodate “quiet” times for housekeeping tasks.

Due to the incredible efficiency gained through Data Packing, HPE 3PAR StoreServ arrays offer one of the highest raw-to-effective ratios among major all-flash arrays while maintaining high levels of performance.²

² Based on HPE analysis of publicly available data at the time of publication.

Selectively applying data reduction

Data reduction technologies on HPE 3PAR StoreServ Storage are designed to make flash more affordable. However, not all data reduction technologies make sense for all workloads. Good examples are compressed media formats and databases that apply compression or encryption before data is written to the storage system. In these scenarios, employing compression technologies likely provides little or no additional savings, as the efficiencies have already been attained at the application layer. In this case, it can be beneficial to disable compression at the volume level to free up array resources that can then be used for other operations. In addition, all data reduction technologies—regardless of implementation—can impact latency. If an application requires the highest levels of performance and the lowest latency, disabling data reduction can improve performance for that application while allowing the system to continue offering savings for other applications.

To address this, HPE 3PAR StoreServ Storage arrays offer the flexibility to selectively apply data reduction, allowing you to choose, on a per-volume basis, which features are enabled. Applications that demand the highest levels of performance and the lowest latency levels can then meet service level agreements (SLAs) while other workloads gain the storage efficiencies that make flash affordable for all of your mainstream applications.

Learn more at
[**hpe.com/3par**](http://hpe.com/3par)



Sign up for updates
