Brochure

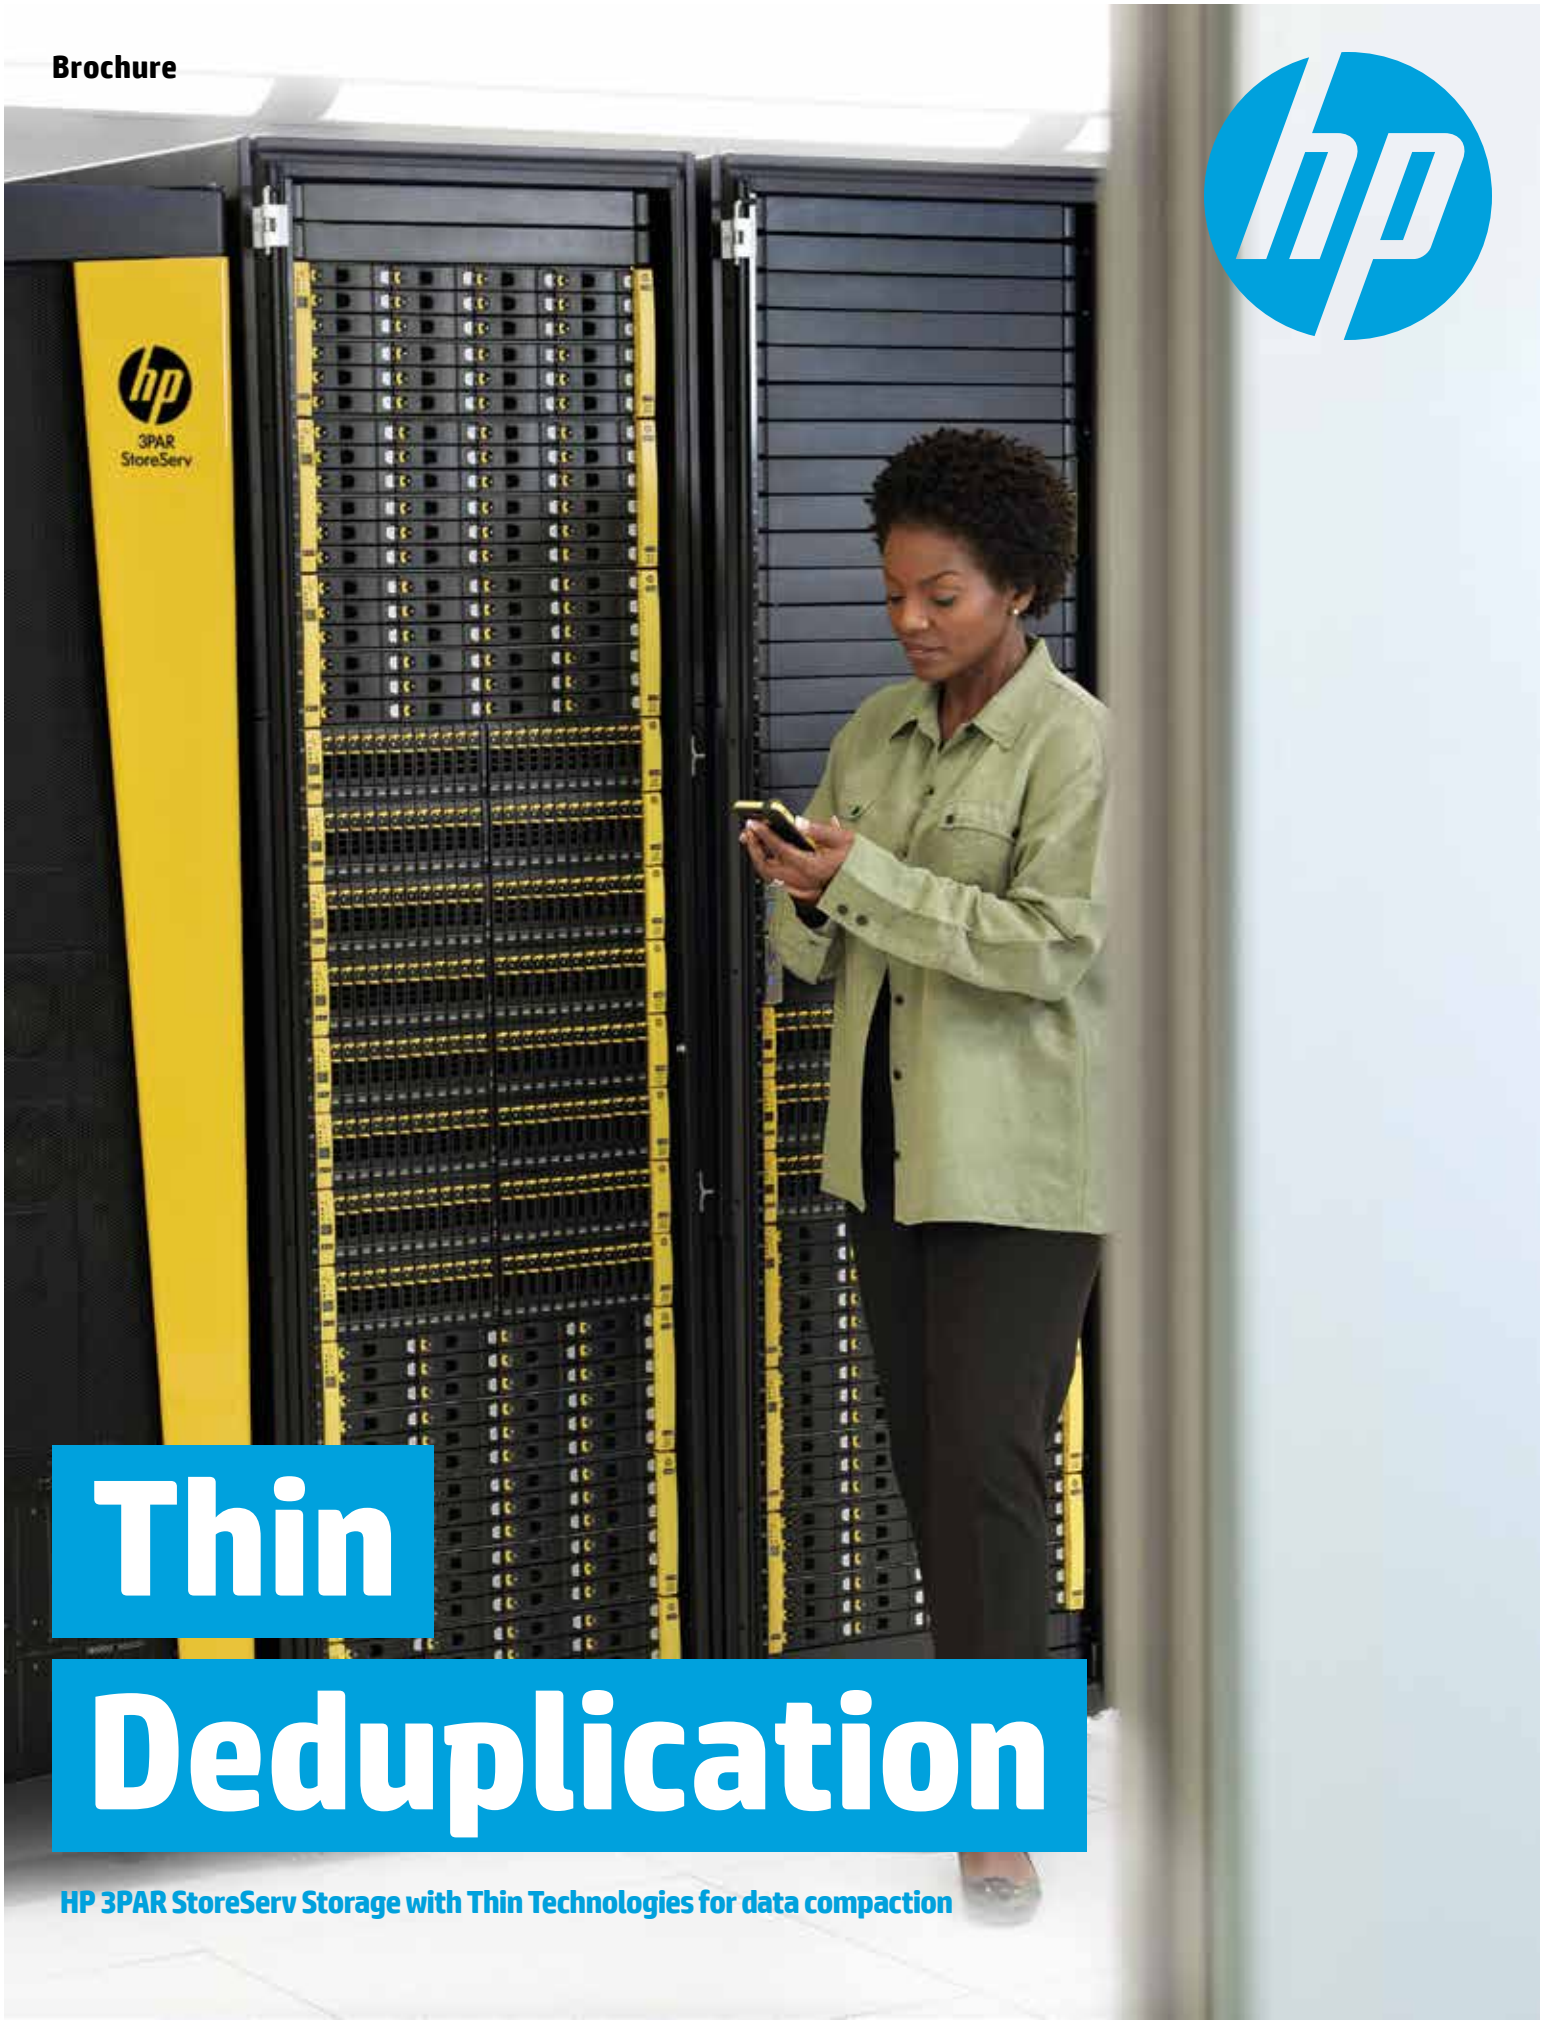


# Thin Deduplication

HP 3PAR StoreServ Storage with Thin Technologies for data compaction

# Use data compaction to start thin, get thin, and stay thin

### HP 3PAR StoreServ capacity efficiency benefits

**Start thin:**

- Just-in-time, reservationless thin provisioning eliminates pre-allocation and pooling
- HP 3PAR Virtual Copy software enables reservationless snapshots that only consume capacity for changed data
- HP 3PAR Thin Clones[1] software instantly creates non-duplicative VM clones for Microsoft® Hyper-V and VMware ESXi
- System-wide striping simultaneously allows high disk utilization and performance
- RAID 50/60 uses granular, 1 GB chunklets instead of disks for increased capacity efficiency

**Get thin:**

- HP 3PAR Thin Deduplication[3] software delivers inline, block-level deduplication without performance or capacity inefficiency tradeoffs
- Built-in, zero-detection mechanism drives efficient inline zero block deduplication at the hardware layer
- HP 3PAR Thin Conversion software leverages zero-detection to drive inline, hardware-accelerated "fat-to-thin" volume conversions

**Stay thin:**

- HP 3PAR Thin Persistence software reclaims allocated but unused space associated with deleted data within volumes
- Inline deduplication increases capacity efficiency, protects flash performance, and extends flash media life span
- HP 3PAR Thin Copy Reclamation reclaims unneeded space from snapshots and remote copies

[1, 2, 3] Available in a future release. Supported only on HP 3PAR StoreServ 7450 systems.

More energy efficient, more compact, and more predictable than spinning drives, flash-based media is the future of high performance storage. Solid state drive (SSD) technology is both mature and capable of giving you incredible performance in a superbly efficient package. However, there is still a significant cost difference between SSDs and spinning media—meaning that, now more than ever, efficiency mechanisms like thin provisioning, block-level deduplication, and space reclamation are essential to getting the most out of SSD capacity.

The introduction of HP 3PAR Thin Deduplication and Thin Clones software[2] demonstrates how HP continues to set the gold standard for hardware-accelerated thin technologies. These technologies drive up capacity efficiency and extend flash media life span to make flash more cost-efficient. Now you can start thin, get thin, and stay thin with an enterprise-class solution that lets you extend your flash storage investments without compromising performance or resiliency.

## Why deduplication for primary storage?

With the increasing use of flash, deduplication for primary storage arrays has become critical. The cost differential between SSDs and hard disk drives (HDDs) requires compaction technologies like thin provisioning and deduplication to make flash-based media more cost-efficient. Widespread deployment of server virtualization is also driving the demand for primary storage deduplication. The potential benefits of deduplication correlate directly with data redundancy. Virtual machine (VM) images and client virtualization environments with hosted virtual desktops are both characterized by a high degree of data redundancy, meaning that these are two use cases where primary deduplication fits well.

The main issue that primary deduplication typically faces is that, particularly in virtualized environments, primary storage arrays are subjected to unpredictable performance demands that can require simultaneously low latency and high throughput. The impact of deduplication on performance is determined by various parameters such as whether deduplication takes place inline or as a background process and the level of granularity used for deduplication operations. This means that deduplicating data at a fine level of granularity while simultaneously delivering space savings improvements generally requires a lot of CPU processing power and memory—more than most primary storage arrays have to spare. This can force tradeoffs that restrict the overall efficiency of primary deduplication.

Designed for mission-critical environments, HP 3PAR StoreServ Storage offers the only solution in the industry that uses built-in, silicon-based mechanisms and a unique Express Indexing feature to protect flash performance while delivering extremely efficient, extremely granular block-level deduplication. Unlike other approaches, HP 3PAR Thin Deduplication software performs a full check on all data before marking it as duplicated, which is essential to ensuring data integrity for mission-critical environments. To understand how this approach is unique and why it's superior, it is important to first understand how HP 3PAR Thin Technologies drive

data compaction through features that provide a broad range of capacity efficiency benefits and leverage hardware acceleration to preserve high performance and ensure the resiliency demanded of Tier-1 storage.

## HP 3PAR Thin Technologies for data compaction

### The key enabling mechanisms for HP 3PAR Thin Deduplication

- HP 3PAR ASICs with built-in data signature generation and bit-bit compare on match
- Hardware offload engines for identifying duplicated data
- Fast lookup tables that store location pointers to accelerate data access
- Tri-layer addresses translation mechanism akin to Virtual Memory lookup tables

Compaction technologies such as thin provisioning and thin reclamation offer efficiency benefits for primary storage that can significantly reduce both capital and operational costs with spinning media and SSDs. However, thin technologies can vary widely in how they are implemented, and this can greatly impact the ability to reduce capacity requirements and extend SSD life span without forcing performance trade-offs. Not only is HP 3PAR StoreServ Storage viewed as the industry's thin technology leader, but third-party testing and competitive analysis confirm that HP 3PAR StoreServ offers the most comprehensive and efficient thin technologies among the major enterprise storage platforms.[4]

Unlike competitive offerings, the "thin" mechanisms unique to HP 3PAR StoreServ Storage allow the platform to offer a range of thin technologies that do not require pre-planning or up-front space reservations.[5] These technologies deliver capacity utilization rates of 80 percent or greater, and reduce capacity consumption by 50 percent or more.[6] In addition, HP 3PAR Thin Technologies protect SSD performance and extend flash-based media life span while ensuring resiliency.

HP 3PAR Thin Technologies including HP 3PAR Thin Provisioning, Thin Conversion, Thin Deduplication, Thin Persistence, and Thin Copy Reclamation achieve data compaction through leveraging built-in hardware capabilities and Express Indexing. Inline, real-time HP 3PAR Thin Conversion only available with HP 3PAR StoreServ Storage provides simple data mobility across clusters or generations of storage systems without downtime, completely changing IT refresh cycles. HP 3PAR Thin Persistence and Thin Copy Reclamation apply these unique mechanisms to ensure that capacity remains thin by reclaiming allocated but unused space at a granular level. Now, HP 3PAR Thin Deduplication and related HP 3PAR Thin Clones software take thin efficiency to the next level when combined with the all-flash HP 3PAR StoreServ 7450 Storage array.

HP 3PAR Thin Deduplication software is able to deduplicate data in-line with a high degree of granularity to provide capacity efficiency that is superior to other approaches without monopolizing CPU resources or compromising data integrity.

## Enterprise-class deduplication

HP 3PAR StoreServ Storage employs purpose-built HP 3PAR ASICs at the heart of each controller node that feature efficient, silicon-based mechanisms to drive inline deduplication.

This implementation relies on the HP 3PAR ASICs to generate and assign signatures to each unique incoming write request. HP 3PAR Express Indexing, a mechanism that accelerates data signature comparison, is used for ultra-fast detection of duplicate write requests in order to preventing duplicate data from being written.

When a new I/O request comes in, HP 3PAR Express Indexing performs instant lookups using metadata tables in order to compare the signatures of the incoming request to signatures of data already stored in the array. When a match is found, HP 3PAR Express Indexing flags the duplicate request and prevents it from being written to the back end. Instead, a pointer is added to the metadata table to reference the existing data blocks. To ensure data integrity, HP 3PAR Thin Deduplication software relies on the controller node ASICs to perform bit-to-bit comparison before any new write update is marked as a duplicate.

With HP 3PAR Thin Deduplication software, the CPU-intensive jobs of calculating signatures for incoming data and verifying reads are offloaded to the ASICs, freeing up processor cycles to deliver advanced data services and service I/O requests. This hardware-assisted approach enables inline deduplication that carries multiple benefits, including increased capacity efficiency, flash performance protection, and flash media lifespan extension.

Without the purpose-built HP 3PAR ASICs and HP 3PAR Express Indexing, other storage architectures lack the processing power to simultaneously drive ultra-fast inline deduplication and the high performance levels demanded by flash-based media. HP 3PAR Thin Deduplication software enables these functions to take place without contention, without sacrificing performance, and while concurrently delivering advanced data services such as replication, federated data mobility, and quality of service level enforcements.

[4, 5] HP Thin Technologies: A Competitive Comparison, Edison Group 2012. h20195.www2.hp.com/V2/GetDocument.aspx?docname=4AA4-4079ENW&cc=us&lc=en

[6] Requires the use of HP 3PAR Thin Conversion Software and HP 3PAR Thin Provisioning Software. For details, refer to the Get Thin Guarantee Terms and Conditions. For more information: hp.com/storage/getthin

**65%**

The average capacity savings achieved by HP 3PAR StoreServ Storage customers today using HP 3PAR Thin Technologies, based on installed base reports.

**Figure 1.** ASIC-based signature generation for inline deduplication

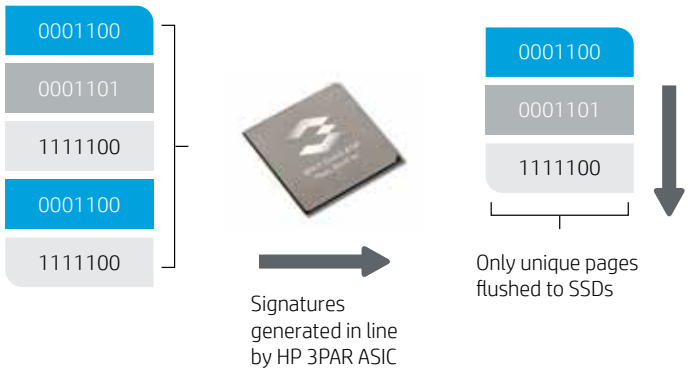


The HP 3PAR StoreServ 7450 Storage with HP 3PAR Thin Deduplication software is the only solution in the industry that uses silicon-based signature generation, allowing inline deduplication to take place while protecting flash performance levels.

## HP 3PAR Thin Clones for non-duplicative VM cloning

An extension of HP 3PAR Thin Deduplication for server virtualization environments, HP 3PAR Thin Clones software enables the creation of non-duplicative VM clones with Microsoft Hyper-V and VMware ESXi.

**Figure 2.** Non-duplicative VM clones

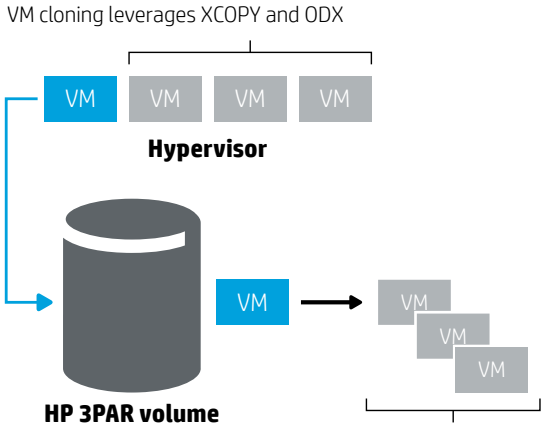


- Clones are created on-the-fly without pre-allocating any storage
- New data is deduplicated inline

These VM clones are created instantly by leveraging copy offload for VMware vStorage APIs for Array Integration (VAAI) and Microsoft Offloaded Data Transfer (ODX) technology without increasing capacity consumption on the HP 3PAR StoreServ Storage system. HP 3PAR Thin Clones software leverages Thin Deduplication to update the metadata table without copying data, relying on inline deduplication technology to reduce capacity footprint as new write request come in.

**Learn more at**
**hp.com/go/StoreServ**

Customize your IT lifecycle management, from acquisition of new IT, management of existing assets, and removal of unneeded equipment.
**hp.com/go/hpfinancialservices**

Developing solutions for major social and environmental challenges
**hp.com/hpinfo/globalcitizenship**

**Sign up for updates**
**hp.com/go/getupdated**

Share with colleagues

Rate this document



4AA5-2367ENW, June 2014