

# Improving Storage efficiency of Microsoft SQL Server databases

## Enhancing SQL Server management with HPE 3PAR StoreServ Adaptive Data Reduction

### HPE 3PAR StoreServ Adaptive Data Reduction technologies

#### Zero Detect

- First-level, preventative data reduction technology
- Autonomic and automatic—you get all the benefits by doing nothing

#### Deduplication

- Second-level, preventative data reduction technology
- Unique HPE 3PAR implementation features deduplication engines built into the HPE 3PAR ASIC

#### Compression

- Complementary, compaction-based data reduction technology
- HPE 3PAR ASIC performs compression, deduplication, leaving the controller CPU less burdened to perform other tasks

#### Data Packing

- Complementary, advanced compaction technology unique to HPE 3PAR StoreServ
- Supports efficiency without compromising performance even when deduplication and compression are both enabled

The exponential growth of business-critical data residing on databases makes it extremely important to ensure the performance, lossless data consistency, and reduction in database sprawl. In short, efficient database management is critical.

What technologies do you employ for efficient management of your SQL Server database and storage? If you already employ SQL Server compression capabilities, exploring storage-layer data reduction technologies can help you to further reduce data footprint for efficient database management.

### What is HPE 3PAR StoreServ Adaptive Data Reduction technology?

HPE 3PAR Adaptive Data Reduction is a collection of capacity-efficiency technologies that come as a standard with HPE StoreServ Storage. Adaptive Data Reduction technologies include Zero Detect, Data Packing, compression, and deduplication. These technologies enhance SQL Server compression, enabling you to get the most out of your system's flash capacity, which in turn reduces your total cost of data storage.

### Cost reduction

Zero Detect is a hardware data reduction technology unique to HPE 3PAR StoreServ built into its ASIC chips. It helps reduce the cost of storage by identifying and removing duplicated zero data from incoming data streams in real time. This results in reduced database size on disk, and ultimately, cost reduction. Compression works by looking inside data streams for opportunities to reduce the overall size of the data set through compaction.

SQL Server customers using compression could see dramatic data reduction savings. As seen in figure 1, Hewlett Packard Enterprise internal testing demonstrates expected compression ratios are 2:1, which drives significant footprint reduction. You can maximize your Microsoft® SQL Server investment combined with 2:1 compression, and lower the cost per GB.

### Expected compression results

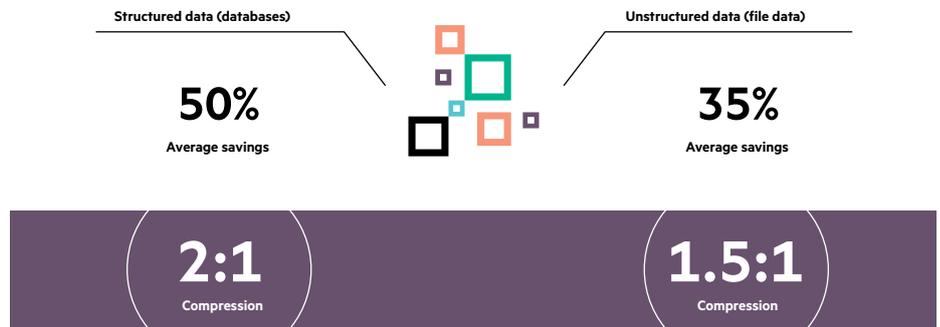


Figure 1. HPE 3PAR StoreServ expected compression results

### CPU efficiency

Express Scan is a technology designed to identify incompressible streams of data and store them in their native format. It prevents wastage of CPU cycles in attempting to compress incompressible data—adding a new dimension to efficiency and improving system performance.

### Total system efficiency

Data Packing allows HPE 3PAR systems to maintain their efficiency by packing deduplicated, compressed data together into a fixed, flash-native page that allows HPE 3PAR to gain a higher total system efficiency than other all-flash platforms. This also reduces the read and write loads on the flash media, and in turn, increases the endurance of flash.

### Data integrity

HPE 3PAR data reduction is lossless, so the original data can be reconstructed from the compressed data. Consider various use cases for lossless compression. Whether it is the need to reconstruct the data in financial and healthcare industries or simply text reconstruction, it is paramount to have the ability to reconstruct the original data and yield needed information for later enhancement and analysis.

### Flexibility for optimal solution

HPE 3PAR arrays allow each of the data reduction technologies to be applied independently and with volume level granularity, providing flexibility for individual workloads with different characteristics. This allows the storage administrator to select which technologies are enabled on a per-volume basis. This in turn allows applications that demand high levels of performance to achieve the service level agreements (SLAs) determined by the business, while gaining the storage efficiencies required to reduce the cost of flash on the remaining data sets.

### Summary

By employing data reduction technologies companies can improve the overall daily performance and capacity of their SQL Server databases and improve ROI for existing hardware, maintenance, and future hardware requisitions. With HPE 3PAR Adaptive Data Reduction capabilities, you can now maximize your Microsoft SQL Server investment combined with 2:1 compression, and lower the cost per GB without compromising data consistency. HPE 3PAR Adaptive Data Reduction allows flexibility and ability to select which technologies are enabled on a per-volume basis for complete flexibility.

For more information on all Data Reduction capabilities offered with HPE 3PAR StoreServ, consult the [Adaptive Data Reduction Brochure](#).

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



Sign up for updates