StoreVault™
A NetApp Division

# Trading Capacity for Data Protection –

## A Guide to Capacity Overhead on a StoreVault Storage System

How capacity is calculated

What to expect

Benefits of redundancy

## introduction

So you just bought a new storage array and the reported capacity is significantly less than the raw capacity you just paid for. Why is that? This paper takes a straightforward look at the causes of differences between raw and usable capacity. Some of the differences are obvious; for example, the losses due to RAID are easily understood. But other differences are less well known. Many of these losses are applicable to all manufacturers, but some elements that are unique to NetApp provide a higher level of data integrity, reliability, and performance.

All storage devices that offer redundancy also trade capacity to provide varying levels of protection, and this tradeoff is usually a compromise with performance. A StoreVault uses capacity tradeoffs through a number of technologies to provide improved reliability for your data and improved system availability. RAID 4 consumes one drive to provide the ability to recover data in the event of a disk drive failure. RAID-DP consumes an additional drive, but provides outstanding protection by enabling recovery from two concurrent disk drive failures. Adding a global hot spare allows the system to recover faster, often before a drive actually fails, which is a pretty good insurance policy for the cost of an extra drive!

For this paper, the capacity losses are divided into two groups: drive-level losses, which are common in every disk drive; and system level losses, which account for what happens after the array is built and depend on the system configuration.

## drive-level capacity losses

### Bytes to Bytes

Disk drive manufacturers have long described capacity in a decimal style: a megabyte is defined as 1,000,000 bytes and a gigabyte is defined as 1,000,000,000 bytes. This is industry-standard practice, and for convenience and consistency, systems manufacturers also use this definition to describe the raw capacity of arrays. But as any system administrator knows, this is not the usable capacity as reported by the operating system.

Usable capacity is defined in the way the file systems consume it: 1KB is 1,024 bytes, and 1GB is 1,073,741,824 bytes, so it's necessary to correct the raw capacity by translating from "disk-drive speak" to "file-system speak."

### Block Checksum

NetApp uses a data protection scheme that provides assurance for data integrity. This scheme involves using every ninth block as a checksum for the preceding eight data blocks. One way to think of this is as "in-line RAID.": The checksum block can be used to check for inconsistencies in the previous eight data blocks. If an inconsistency is discovered, then the system can recreate the bad block or blocks from the parity drive and the other drives in the RAID group. This prevents a full RAID rebuild when a bad block is detected. If multiple checksum errors are found, then a Rapid RAID Rebuild can be kicked off in the background. This process allows a graceful drive replacement without a performance impact, instead of the usual hit taken with a RAID 5 rebuild.

Block Checksum consumes 1/9 of the available capacity per disk drive to provide this very secure reliability. Block checksum provides a similar level of data protection on SATA drives as that found on enterprise-class storage with Fibre Channel drives, at a much lower cost and with no performance penalty.

### RAID

There is a small (approximately 5%) capacity overhead associated with efficiently operating RAID at the disk level. This overhead allows low-level disk-based management tasks to be performed, such as labeling and record maintenance.

### Right Sizing

Not all disk drives have precisely the same number of data blocks, so to ensure that the system manufacturer has flexibility and to allow blocks to be flagged as not used on individual drives, a lowest common denominator approach is used. Right sizing consumes a few gigabytes per drive, but it provides additional spare capacity on the drives to enable small disk defects to be isolated without compromising the usable space and without unnecessarily generating a failed disk and a RAID rebuild.

### Result of Drive-Level Overhead

These combined overhead values (block check sum, RAID, and right sizing) total just above 17% of the published disk size. So, a 500GB drive, for example, will have 413.2GB available space in the StoreVault system.

## system-level capacity losses

### RAID 4

As with any other single parity scheme, RAID 4 consumes one drive of capacity in an array. This one drive contains all of the parity information for all of the data drives in the array and ensures that your data is safe in the event of a disk drive failure. If a drive fails, then that drive is replaced (or a hot spare is used) and rebuilt using the information on the other drives.

In a 4-drive RAID 4 array, one drive is used for parity, so 75% of the total capacity is available. In a 12-drive array, one drive is used for parity, so 92% of the total capacity is available. It is clear that the larger the number of spindles in an array, the larger the usable percentage capacity.

RAID 4 has an advantage over RAID 5 in that additional drives can be added to the array to expand actaual data storage capacity at any time without disruption. To add capacity to a RAID 5 array, the entire array must be copied to scratch space, drives added, the array rebuilt, and finally the data must be migrated back to the array—a time-consuming, tedious, disruptive, and potentially disastrous procedure.

### RAID-DP™

Unique to NetApp, RAID-DP is a double parity scheme that provides protection for your data even in the event of two concurrent drive failures. Two drives are used for parity information, so RAID-DP consumes two drives of capacity in an array. A RAID 4 array can be converted to RAID-DP at any time with the addition of an extra drive and a couple of mouse clicks.

RAID-DP is like RAID 4 in that additional drives can be added to the array to expand capacity at any time without disruption. This is very different from RAID 5 and RAID 6 where the entire array must be rebuilt to add capacity.

In a 5-drive RAID-DP array, two drives are used for parity so 60% of the total capacity is available. In a 12-drive array, two drives are used for parity so 83% of the total capacity is available. As with RAID 4, the larger the number of spindles in an array, the larger the usable percentage capacity. Incidentally the performance also gets better with a higher number of drives.

### Hot Spare

Hot spares are used in an array to provide an immediate target for the contents of a drive that is to be replaced. The failing drive may have actually failed; or more likely it will have been detected as showing signs of potential future failure and a Rapid RAID Rebuild initiated. Hot spares of course are not used for data while they are maintained as spares, so their capacity should be regarded as overhead when considering usable capacity in a system. One or more hot spares can be used with RAID 4 and RAID-DP to provide faster recovery.

### WAFL® Overhead

All file systems consume some overhead. The NetApp WAFL (Write Anywhere File Layout) system reserves 10% of capacity in order to very efficiently lay out the data for fast retrieval and minimal file fragmentation, trading a little capacity for a significant improvement in performance and reduced system management.

### Snapshot™ Reserve

NetApp Snapshot technology enables instantaneous backup and almost instant restore by maintaining images of the file system at points in time. Best practice, established over many years, suggests that 20% is a good reserve for Snapshot copies, and this is the default setting on all StoreVault products. Although this 20% is "reserved," it is actually guidance for administrators. The actual Snapshot consumption of disk space depends on the environment and types of applications. The administrator of a StoreVault storage system can adjust the amount of reserve space for Snapshots on each of the individual storage areas within the system.

## real capacity yields for StoreVault

The following tables show exactly what to expect with a StoreVault, depending on the configuration and the number of drives in the array. The configurations are RAID 4, RAID 4 with hot spare, RAID-DP, and RAID-DP with hot spare.

"Total Raw GB" is the capacity indicated by the number of drives; for example, 6 x 500GB drives gives a total raw capacity of 3,000GB, or 3TB.

"Data Raw GB" is the capacity indicated by the total number of drives less the parity and spare drives.

"Usable GB" is the storage capacity indicated by StoreVault Manager. This is the capacity available for your data after subtracting all the losses detailed previously without a Snapshot Retrieve.

"Usable %" is "Usable GB" as a percentage of "Total Raw GB".

## Usable Capacity with RAID 4 and No Hot Spare

| Total Drives | Data Drives | RAID 4 and No Hot Spare | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Using 250GB Drives | | | | Using 500GB Drives | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 4 | 3 | 1000 | 750 | 557 | 56% | 2000 | 1500 | 1116 | 56% |
| 5 | 4 | 1250 | 1000 | 743 | 59% | 2500 | 2000 | 1487 | 59% |
| 6 | 5 | 1500 | 1250 | 929 | 62% | 3000 | 2500 | 1859 | 62% |
| 7 | 6 | 1750 | 1500 | 1115 | 64% | 3500 | 3000 | 2231 | 64% |
| 8 | 7 | 2000 | 1750 | 1300 | 65% | 4000 | 3500 | 2603 | 65% |
| 9* | 8 | 2250 | 2000 | 1486 | 66% | 4500 | 4000 | 2975 | 66% |
| 10* | 9 | 2500 | 2250 | 1672 | 67% | 5000 | 4500 | 3347 | 67% |
| 11* | 10 | 2750 | 2500 | 1858 | 68% | 5500 | 5000 | 3719 | 68% |
| 12* | 11 | 3000 | 2750 | 2044 | 68% | 6000 | 5500 | 4091 | 68% |
| | | Using 750GB Drives ** | | | | Using 1TB Drives** | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 4 | 3 | 3000 | 2250 | 1672 | 56% | 4000 | 3000 | 2231 | 56% |
| 5 | 4 | 3750 | 3000 | 2229 | 59% | 5000 | 4000 | 2975 | 60% |
| 6 | 5 | 4500 | 3750 | 2787 | 62% | 6000 | 5000 | 3719 | 62% |
| 7 | 6 | 5250 | 4500 | 3344 | 64% | 7000 | 6000 | 4463 | 64% |
| 8 | 7 | 6000 | 5250 | 3901 | 65% | 8000 | 7000 | 5206 | 65% |
| 9* | 8 | 6750 | 6000 | 4459 | 66% | 9000 | 8000 | 5950 | 66% |
| 10* | 9 | 7500 | 6750 | 5016 | 67% | 10000 | 9000 | 6694 | 67% |
| 11* | 10 | 8250 | 7500 | 5573 | 68% | 11000 | 10000 | 7438 | 68% |
| 12* | 11 | 9000 | 8250 | 6131 | 68% | 12000 | 11000 | 8181 | 68% |

* StoreVault S500 and S550 only
** StoreVault S300 and S550 only

## Usable Capacity with RAID-DP and No Hot Spare or RAID 4 Plus Hot Spare

| Total Drives | Data Drives | RAID-DP And No Hot Spare or RAID4 And One Hot Spare | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Using 250GB Drives | | | | Using 500GB Drives | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 5 | 3 | 1250 | 750 | 557 | 45% | 2500 | 1500 | 1116 | 45% |
| 6 | 4 | 1500 | 1000 | 743 | 50% | 3000 | 2000 | 1487 | 50% |
| 7 | 5 | 1750 | 1250 | 929 | 53% | 3500 | 2500 | 1859 | 53% |
| 8 | 6 | 2000 | 1500 | 1115 | 56% | 4000 | 3000 | 2231 | 56% |
| 9* | 7 | 2250 | 1750 | 1300 | 58% | 4500 | 3500 | 2603 | 58% |
| 10* | 8 | 2500 | 2000 | 1486 | 59% | 5000 | 4000 | 2975 | 59% |
| 11* | 9 | 2750 | 2250 | 1672 | 61% | 5500 | 4500 | 3347 | 61% |
| 12* | 10 | 3000 | 2500 | 1858 | 62% | 6000 | 5000 | 3719 | 62% |
| | | Using 750GB Drives ** | | | | Using 1TB Drives** | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 5 | 3 | 3750 | 2250 | 1672 | 45% | 5000 | 3000 | 2231 | 45% |
| 6 | 4 | 4500 | 3000 | 2229 | 50% | 6000 | 4000 | 2975 | 50% |
| 7 | 5 | 5250 | 3750 | 2787 | 53% | 7000 | 5000 | 3719 | 53% |
| 8 | 6 | 6000 | 4500 | 3344 | 56% | 8000 | 6000 | 4463 | 56% |
| 9* | 7 | 6750 | 5250 | 3901 | 58% | 9000 | 7000 | 5206 | 58% |
| 10* | 8 | 7500 | 6000 | 4459 | 59% | 10000 | 8000 | 5950 | 60% |
| 11* | 9 | 8250 | 6750 | 5016 | 61% | 11000 | 9000 | 6694 | 61% |
| 12* | 10 | 9000 | 7500 | 5573 | 62% | 12000 | 10000 | 7438 | 62% |

* StoreVault S500 and S550 only
** StoreVault S300 and S550 only

## Usable Capacity with RAID-DP Plus Hot Spare

| Total Drives | Data Drives | RAID-DP and One Hot Spare | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | Using 250GB Drives | | | | Using 500GB Drives | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 6 | 3 | 1500 | 750 | 557 | 37% | 3000 | 1500 | 1116 | 37% |
| 7 | 4 | 1750 | 1000 | 743 | 42% | 3500 | 2000 | 1487 | 42% |
| 8 | 5 | 2000 | 1250 | 929 | 46% | 4000 | 2500 | 1859 | 46% |
| 9* | 6 | 2250 | 1500 | 1115 | 50% | 4500 | 3000 | 2231 | 50% |
| 10* | 7 | 2500 | 1750 | 1300 | 52% | 5000 | 3500 | 2603 | 52% |
| 11* | 8 | 2750 | 2000 | 1486 | 54% | 5500 | 4000 | 2975 | 54% |
| 12* | 9 | 3000 | 2250 | 1672 | 56% | 6000 | 4500 | 3347 | 56% |
| | | Using 750GB Drives ** | | | | Using 1TB Drives** | | | |
| | | Total Raw GB | Data Raw GB | Usable GB | Usable % | Total Raw GB | Data Raw GB | Usable GB | Usable % |
| 6 | 3 | 4500 | 2250 | 1672 | 37% | 6000 | 3000 | 2231 | 37% |
| 7 | 4 | 5250 | 3000 | 2229 | 42% | 7000 | 4000 | 2975 | 43% |
| 8 | 5 | 6000 | 3750 | 2787 | 46% | 8000 | 5000 | 3719 | 46% |
| 9* | 6 | 6750 | 4500 | 3344 | 50% | 9000 | 6000 | 4463 | 50% |
| 10* | 7 | 7500 | 5250 | 3901 | 52% | 10000 | 7000 | 5206 | 52% |
| 11* | 8 | 8250 | 6000 | 4459 | 54% | 11000 | 8000 | 5950 | 54% |
| 12* | 9 | 9000 | 6750 | 5016 | 56% | 12000 | 9000 | 6694 | 56% |

* StoreVault S500 and S550 only

** StoreVault S300 and S550 only

## summary

Capacity overhead is a fact of life with storage systems. However, as demonstrated in this paper, the benefits of performance, reliability, and data integrity derived from these losses are very worthwhile. As the number of disk drives increases, the percentage of capacity usable for data increases. It is worth noting that the performance of the array also increases with increased spindle count.

A note on best practice: NetApp strongly recommends that customers consider implementing RAID-DP and a hot spare. This provides maximum protection for your data and maximum likelihood of avoiding a disruptive rebuild. There is a price to pay for that additional redundancy, but it seems a small price compared to the benefits and additional protection provided.

The aim of the StoreVault product line is to combine simplicity with robust data protection. This paper presents the actual usable capacities in a simple format and explains the reasons for the differences. When you compare usable storage capacities, make sure that you are getting the value of additional protection for your data.

For more information, visit www.storevault.com.

NetApp®

Network Appliance, Inc.
495 East Java Drive
Sunnyvale, CA 94089

About Network Appliance

Network Appliance is a world leader in network storage solutions for today's data-intensive world. Since its inception in 1992, Network Appliance has delivered technology, product, and partner firsts that simplify data management. Information about Network Appliance solutions and services is available at **www.netapp.com.**

For more information on StoreVault by NetApp, go to **www.storevault.com**



10-193-111-3 01/2008