



Highlights

- Accelerate applications while helping lower storage costs
 - Increase effective storage capacity without sacrificing flash performance
 - Gain the benefits of inline data deduplication, compression and thin provisioning in one easily integrated system
-

Accelerate storage savings and efficiency

All-flash IBM FlashSystem and Permabit SANblox data reduction make fast storage more affordable.

Business data volumes are essentially doubling every two years, yet spending on IT worldwide is flat or even declining in some sectors.¹ The collision of these trends places extraordinary pressure on data center managers to wring every penny of cost savings and efficiency possible from their data storage solutions. Data reduction strategies that incorporate both deduplication and compression increase effective storage capacity, often enabling enterprises to increase storage density, shrink storage footprint and thus reduce overall IT costs.

Flash storage provides many benefits, including much higher performance than other storage media, greater reliability and lower operational expenses. Data reduction optimizes flash capacity, making flash even more cost-effective, but unless the reduction is done efficiently, it can impact storage performance, latency and resource utilization. IBM has been working closely with data-reduction technology industry-leader Permabit Technology Corporation to design, test and validate solution architectures based on IBM® FlashSystem® all-flash storage arrays and the Permabit Albireo SANblox™ high-performance data reduction appliance. Together, IBM FlashSystem and Permabit SANblox offer a high-performance, high-efficiency low-cost storage solution with enterprise-grade performance, availability and ease of installation.



Driving down the cost of data storage

Two different but related technologies—data deduplication and compression—have been developed to help storage systems handle larger volumes of data at significantly lower cost. When deployed separately, each technology can dramatically increase effective storage capacity for certain types of application workloads. When deployed together, they provide significantly greater benefits across a wide spectrum of workloads and business environments.

Data deduplication

Deduplication works across files to reduce storage needs by eliminating redundant data. Only one unique instance of the data is actually retained on the storage media. Redundant data is replaced with a pointer to the initial data instance. For example, a single server in a virtual server farm may contain a hundred copies of an operating system (OS) and various applications, each copy potentially taking up 100 GB of storage or more and thus all the redundant copies requiring 10 TB of extra storage capacity. With data deduplication, only one instance of the OS and each application is stored; each subsequent instance is then referenced back to the one saved copy, reducing 10 TB of storage demand to 100 GB.

Data deduplication offers other benefits, as well. Lower storage space requirements can lead to both capital expense savings and reduced operational costs, thanks to the storage media that is not powered or cooled. Deduplication also reduces the amount of data that must be sent across a wide area network (WAN) for

remote backups, replication and disaster recovery. In actual practice, data deduplication is best used in conjunction with data compression; together, these techniques are very effective at optimizing the use of storage space as well as its associated costs and overhead.

As with data compression, data deduplication must be run inline to enable the most effective use of storage system resources while delivering immediate storage savings. A second critical consideration with deduplication is granularity. In general, the smaller the deduplication block size, the more efficient the deduplication becomes. A block size of 4 KB is used by all modern operating systems. Systems that deduplicate across larger block sizes typically see from 20 to more than 40 percent less savings than solutions utilizing 4 KB blocks.

Data compression

Where deduplication works across files, data compression works within files to reduce the number of bits needed to represent data. Compressing data lowers storage capacity requirements, just as deduplication does, and also speeds file transfers across networks, thus helping decrease costs for storage hardware and network bandwidth.

Compression is performed by applying formulas or algorithms that replace longer bit strings with shorter ones, using a dictionary for the conversion. Compression can also be achieved by inserting a reference or pointer into a string of bits that the

compression engine has already seen. With the technology available today, compression rates of 5:1 and higher can be achieved, depending on the type of data and application workload. Data compression works best in traditional database and online transaction processing (OLTP) environments, as well as with many data warehouse and analytics workloads.

Originally, data reduction technologies such as compression were developed to reduce the capacity requirements and costs of archived information. This meant that the technologies were applied when their use was convenient and out of the data path. But an increasing number of enterprises are looking to leverage the advantages of data reduction to deliver immediate storage savings for active production data sets. Even for these use cases, many compression solutions still do their work “post-process” and do not perform well in random input/output (I/O)-intensive primary storage environments. The best data compression solutions, however, perform compression inline in real time and provide good results for all I/O patterns with minimal impacts to storage latency and performance.

With deduplication working across files and compression working within files, these become complementary solutions, yielding synergistic benefits and more efficient use of processing resources.

Finally, with all data reduction strategies, system performance and the efficient use of resources are crucial factors. Many data reduction implementations are just not designed to perform under primary workloads, which are typically characterized by heavy random I/O. These solutions generally revert to a post-process approach, consuming storage for data caching and increasing storage costs. This heavy I/O often puts so much load on IT systems that it impacts the performance of data management operations such as snapshot and replication activities. In addition, some data reduction solutions require that data reduction always be enabled—even with data sets and workloads that will not necessarily benefit (for example, where data is already compressed at the application level). This requirement can further impact the economics of storage.

Permabit, a data reduction market leader

Permabit is a pioneer in the development of primary storage data reduction technologies such as data deduplication and compression. Innovative products from Permabit enable leading storage vendors to help their customers cut effective cost and increase effective capacity, reduce time to market and gain competitive advantage. Just as server virtualization revolutionized the economics of compute, Permabit data reduction technologies are transforming the economics of storage.

SANblox

SANblox from Permabit is a high-performance data reduction appliance for Fibre Channel storage area network (SAN)-attached primary storage arrays. SANblox delivers inline data deduplication and compression to flash, hybrid or disk-based storage environments.

The industry-leading capabilities of SANblox enable SAN-attached systems such as IBM FlashSystem arrays to deliver higher effective storage capacity at lower cost. SANblox incorporates Permabit data deduplication and Permabit HIOPS Compression™ software into a 2U (high availability) appliance capable of delivering 6:1 data reduction rates or greater in primary storage environments.

Permabit technologies are combined with Red Hat Enterprise Linux, Red Hat Enterprise Linux High Availability, and Emulex Fibre Channel software to ensure compatibility, ease of management and comprehensive support for Fibre Channel SANs. SANblox is able to deliver predictable, high-performance data reduction with minimal impact to back-end storage operations because it operates as a separate physical appliance. And SANblox functionality can be selectively applied to data sets that will benefit from data reduction, while remaining out of the data path of workloads that will see no benefit from compression or deduplication.

Unlike gateway products that perform write-back caching to mask poor performance, SANblox commits data to storage before a write is acknowledged. Writes are not acknowledged

before being persistently stored, and both data and metadata are maintained in the enterprise storage device. This improves data safety because once a write completes, data is immediately under the protection of the array.

SANblox can benefit a wide variety of data sets and applications. It provides immediate benefit for applications that store large amounts of redundant data such as virtual server farms and virtual desktop infrastructures (VDIs). It is also beneficial for software that produces compressible data such as traditional databases, e-commerce/OLTP, data warehouses, data analytics and big-data applications, including those that process and store system logs and sensor data. These applications are the backbone of the finance, healthcare, education, retail, manufacturing and energy industries, among many others.

IBM FlashSystem market-leading all-flash solutions

IBM FlashSystem is a family of all-flash storage arrays engineered to address the most demanding enterprise performance, reliability, comprehensive feature set and cost requirements. The family currently includes IBM FlashSystem 900, the work-horse base model designed to provide extreme performance, low latency and data protection for business-critical application acceleration. IBM FlashSystem V9000 adds the industry-leading storage services and virtualization capabilities of IBM Spectrum Virtualize™ to create a complete storage solution for active data.



Figure 1. IBM FlashSystem 900

Powered by IBM FlashCore™ technology, IBM FlashSystem 900 delivers the extreme performance, enterprise reliability and operational efficiencies required to provide enterprise-grade targeted application acceleration. For storage environments already equipped with robust storage management capabilities such as those provided by IBM Spectrum Virtualize or IBM Spectrum Scale™, the ultra-low latency, market-leading performance and ease of implementation offered by IBM FlashSystem 900 make it an ideal choice to transform a business-critical application into an engine of business growth and competitive advantage.



Figure 2. IBM FlashSystem V9000

IBM FlashSystem V9000 offers the advantages of software-defined storage at the speed of flash. These all-flash storage systems deliver the full capabilities of the hardware-accelerated architecture, IBM MicroLatency® modules and advanced flash management provided by IBM FlashCore technology. These capabilities are coupled with a rich set of storage services and virtualization features, including IBM Real-time Compression™, dynamic tiering, thin provisioning, snapshots, cloning, replication, data copy services and high-availability configurations. Thanks to the downward spiral of flash acquisition costs combined with the extraordinary operational cost savings offered by flash, many enterprises are now deploying IBM FlashSystem V9000 to great advantage as their primary storage solution across their entire active data set.

Integrating data reduction and storage performance

As a cornerstone technology for next-generation storage, data reduction makes it possible to deliver high-performance flash at the cost of spinning disk. As a result, storage arrays with data reduction are stealing market share from products that lack it. The message coming from the marketplace is clear—next-generation storage arrays must have inline data reduction technology to compete.

Solution architecture

SANblox delivers fine-grained (4 KB block) inline deduplication and compression for a best-in-class combination of efficiency and performance. An easy to use, web-based interface allows administrators to aggregate IBM FlashSystem high-performance flash into a pool of deduplicated and compressed storage. This storage pool can be carved up and presented as SANblox logical unit number (LUN) storage volumes for use by applications. When data is written to SANblox, Permabit software monitors the write requests that come down from the SANblox LUNs and uses deduplication technology to identify duplicates. HIOPS Compression then compresses any unique blocks before they are written to the storage device. From then on, the software keeps track of stored data and manages it in a way that is seamless to the user. Data reduction through block-level deduplication and compression increases the amount of usable IBM FlashSystem storage, helping lower costs by increasing the effective capacity of the all-flash storage array.

SANblox is configured in high-availability pairs to provide transparent failover. Data reliability is assured because all data is immediately written to back-end block storage with no write

caching. A suite of management tools allows administrators to easily provision and monitor SANblox through command-line tools, a web-based user interface or SMTP-based event monitoring.

Confirming compatibility and performance

Permabit and IBM engaged in a joint effort to provide up-to-date performance and data reduction results with SANblox and IBM FlashSystem storage. The data collected enables comparisons between the combined IBM FlashSystem/Permabit SANblox solution and other high-performance storage and data reduction solutions.

Test plan

The tests used synthetic data generators such as Flexible IO (fio). The tests were executed using the following read/write characteristics:

- 100 percent sequential write, 128 KB block size
- Random read/write mixes
- 100, 70, 50, 30 and 0 percent reads
- 4 KB, 8 KB, 16 KB, 32 KB and 64 KB block sizes
- 100 percent sequential reads

With two data patterns:

- Unique data only
- Data with 3:1 compression and 3:1 deduplication

Test results

SANblox achieved the following performance with IBM FlashSystem 900 as a target:

- Minimum latency of 0.4 ms with 0.7 ms latency on 4 KB reads and writes, demonstrating only a 200 to 500 μ s latency overhead through the SANblox appliance
- 1:1 linear performance scaling confirmed by adding a second pair of SANblox nodes, all using the single IBM FlashSystem 900 target
- 350,000 peak IOPS at 4 KB, with the expected ability to scale to 700,000 IOPS (4 KB reads) with a single IBM FlashSystem 900 array and two pairs of SANblox appliances
- 50 to 80 percent load with deduplication maintained latencies averaging 1 ms
- Failover in 20 to 45 seconds

With IBM FlashSystem V9000 as the target, SANblox demonstrated similar performance, plus:

- Linear performance scaling when two SANblox high-availability pairs were deployed with a single IBM FlashSystem V9000 array
- Slight improvement in latency due to IBM FlashSystem V9000 controller cache

Market leaders deliver

No one disagrees that flash storage is much faster than conventional systems, but historically, flash-based storage adoption rates have suffered due to a perception that it is more costly per unit of capacity. Even though evaluations of overall storage costs are becoming much more nuanced, and flash prices have steadily fallen, many enterprises still don't fully recognize the

value of much higher performance versus higher acquisition costs. Adding data reduction capabilities to the mix through solutions such as SANblox changes this equation significantly. Suddenly flash cost per capacity becomes very competitive with conventional arrays, while the performance, throughput and ultra-low latency benefits of IBM FlashSystem storage make the joint solution extremely attractive.

And with SANblox providing data reduction capabilities that can be implemented per storage volume, enterprises can target the data sets and workloads that offer the most data reduction benefits—even while maintaining the flexibility to apply the full benefits offered by market-leading all-flash arrays such as IBM FlashSystem to business-critical applications and workloads that require the highest performance available.

Low cost, high performance, broad workload ranges, mature technologies and great flexibility. These are the hallmarks of the solutions you expect from innovators and market leaders such as IBM and Permabit.

For more information

To learn more about IBM FlashSystem A9000, please contact your IBM representative or IBM Business Partner, or visit the following website: ibm.com/systems/storage/flash

Additionally, IBM Global Financing provides numerous payment options to help you acquire the technology you need to grow your business. We provide full lifecycle management of IT products and services, from acquisition to disposition. For more information, visit: ibm.com/financing



© Copyright IBM Corporation 2016

IBM Systems
Route 100
Somers, NY 10589

Produced in the United States of America
February 2016

IBM, the IBM logo, ibm.com, IBM FlashSystem, IBM FlashCore, MicroLatency, IBM Spectrum Virtualize, Real-time Compression, and IBM Spectrum Scale are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml

Permabit Albireo SANblox and HIOPS Compression are trademarks of Permabit.

This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

It is the user’s responsibility to evaluate and verify the operation of any other products or programs with IBM products and programs.

THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.

Permabit Albireo SANblox and HIOPS Compression are not IBM products or offerings. Permabit Albireo SANblox and HIOPS Compression are sold or licensed, as the case may be, to users under Permabit’s terms and conditions, which are provided with the product or offering. Availability, and any and all warranties, services and support for Permabit Albireo SANblox and HIOPS Compression is the direct responsibility of, and is provided directly to users by, Permabit.

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

¹ Phil Goodwin and Ashish Nadkarni, “The Copy Data Problem: Analysis Update,” *IDC*, February 2015.
<http://www.idc.com/getdoc.jsp?containerId=254354>



Please Recycle
