



Megan Gilge
Karen Orlando

Flash or SSD: Why and When to Use IBM FlashSystem

Overview

This IBM® Redpaper™ publication explains how to select an IBM FlashSystem™ or solid-state drive (SSD) solution. It describes why and when to use FlashSystem products, and reviews total cost of ownership (TCO), economics, performance, scalability, power, and cooling. Read this guide for information about selecting the correct solution (SSD or flash technology).

This paper also reviews examples of FlashSystem storage bundled with solutions such as hybrid arrays and IBM System Storage® SAN Volume Controller (SVC). Furthermore, it compares FlashSystem storage to other storage such as PCIe Express (PCIe) products and hard disk drives (HDDs).

Business value of FlashSystem products

FlashSystem storage systems deliver high performance, efficiency, and reliability for shared enterprise storage environments, which can help you address performance issues with your most important applications and infrastructure. FlashSystem products also provide capability for big data and cloud environments. These storage systems can either complement or replace traditional hard disk arrays in many business-critical applications:

- ▶ Online transaction processing (OLTP)
- ▶ Business intelligence (BI)
- ▶ Online analytical processing (OLAP)
- ▶ Virtual desktop infrastructures (VDI)
- ▶ High-performance computing (HPC)
- ▶ Content delivery solutions (such as cloud storage and video-on-demand)

As standard shared primary data storage devices, FlashSystem storage delivers performance exponentially beyond that of most traditional arrays, even those arrays that incorporate SSDs or other flash technology. These storage systems can also be used as the top tier of storage

alongside traditional arrays in tiered storage architectures such as the IBM Easy Tier® functionality available in the System Storage SVC storage virtualization platform.

FlashSystem storage helps you achieve a competitive advantage in a world of 24x7 continuous and complex operations:

- ▶ FlashSystem products are designed to deliver the lowest latency and the highest IOPS in the market today at an economical price.
- ▶ FlashSystem products can be used as Tier 0 storage for solutions with SVC.
- ▶ FlashSystem products are designed to be small and fast. You can use SVC to add functionality such as real-time compression and mirroring.
- ▶ FlashSystem products are especially useful when cooling, low power consumption, and small footprint are key to accelerating business-critical applications.

Storage technologies, latency, and performance

Figure 1 shows the typical points that add to latency for various types of storage technologies.

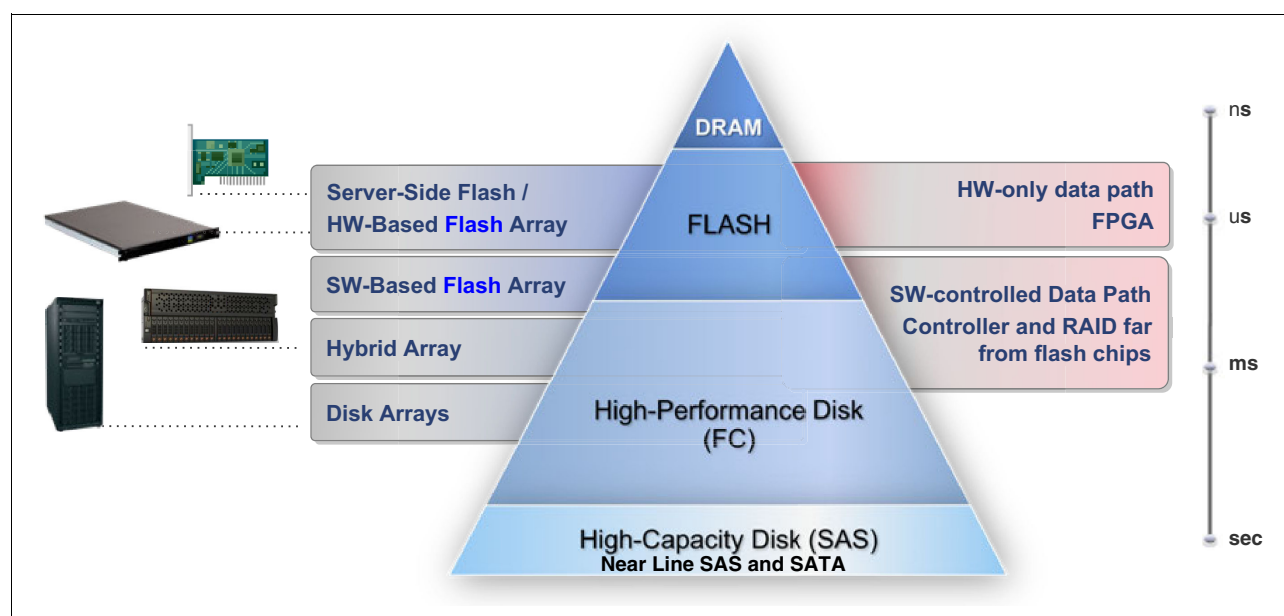


Figure 1 Latency seen in storage technologies

DRAM has the lowest latency, but is the most expensive. In-server direct attach flash memory and externally shared flash memory (storage area network, or SAN) with a hardware-only data path are less expensive than DRAM, and are designed for low latency. FlashSystem storage uses hardware-based flash arrays with a hardware-only data path.

Hardware-based flash arrays provide maximum performance and lowest latency with optimized field-programmable gate array (FPGA) hardware-only data path, custom hardware design, massive parallelism handling data, and 25 us write latency into flash media for FlashSystem storage.

For software-based flash arrays, data flows through software, regardless of the array's internal bus connection to flash memory. The connections can be PCIe, SSD, or proprietary. Because there is additional time required for software processing, these solutions have much lower performance and latency than flash technology.

Hybrid arrays are non-deterministic performance solutions. They are tiered architectures that usually implement flash technology via SSD or PCIe cards. These arrays were originally designed for disk-only solutions, with little to no design specification to maximize the benefits of flash memory.

High-capacity disk Serial Attached SCSI (SAS) has lower performance and cost, with Near Line SAS and SATA disks having the lowest performance and cost. Note that the lower end storage such as Near Line SAS and SATA can be more prone to hardware failure than higher-end HDDs.

If you compare one SSD with one HDD, or one SSD RAID with one HDD RAID, an SSD device always performs faster than an HDD device, and a flash volume performs faster than an SSD. This is especially true for OLTP, which is critical for SAP and IBM DB2®.

If you consider price per TB (terabyte), FlashSystem storage is more expensive than SSD and HDD. However, if speed is important, flash configurations are cheaper. When evaluating a storage solution, always consider the total cost of ownership and the return on your investment. Storage solutions often affect the entire environment. For example, if you use FlashSystem storage to meet requirements that cannot be met with scale-out storage, you might also be able to eliminate costly hardware and software licenses. When you simplify your storage environment, you can reduce the resources and expense required to administer it. Figure 2 depicts how FlashSystem storage can reduce your overall business expenses.

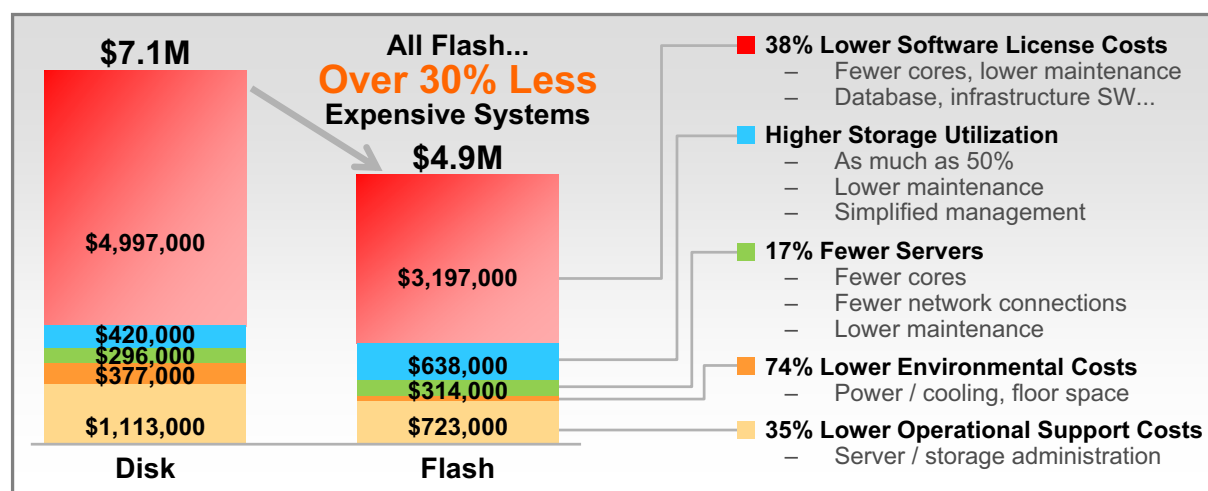


Figure 2 FlashSystem economics

Also consider how faster storage can improve application performance and reduce costs. Faster storage can decrease transaction time, improve the rate of transactions, and increase revenue. FlashSystem storage products can also help you reduce requirements for energy, cooling, and floor space.

As shown in Figure 2, clients report over a 30% cost savings by using FlashSystem storage as compared to disk. They also experience up to 38% overall lower software license costs due to fewer cores and lower software maintenance, and note as much as a 50% improvement in storage utilization, lower maintenance, and ease of management. With an improvement of 17% fewer servers, environments have fewer cores and network connections, and therefore lower maintenance costs. Lower operational support costs achieved by using less power, cooling, and required floor space result in up to a 35% savings.

The difference between SSD and FlashSystem storage

Flash memory is a storage technology. It is available in multiple forms, such as an IBM FlashSystem product or an SSD product with a hard disk form factor. FlashSystem products have extreme performance because they have a hardware-only data path, but other more traditional types of SSD storage often require software processing.

For example, an SSD device with a hard disk form factor has flash memory that is put into a carrier. This carrier is inserted into an array like a hard disk drive. The speed of storage access is limited by the following technology, which cannot keep up with flash technology, and adds latency:

- ▶ Array controllers and software layers
- ▶ SAS controllers and shared bus
- ▶ Tiering and shared data path
- ▶ Form factor enclosure

Top reasons to use IBM FlashSystem storage

IBM FlashSystem products are fast and efficient. Lower power consumption and fast, reliable operation helps increase revenue (return on investment, or ROI) and saves more money (TCO). You can use IBM FlashSystem products to improve performance and efficiency without rearchitecting applications or storage infrastructure. With FlashSystem products, you can achieve these benefits:

- ▶ Supplement your existing infrastructure
- ▶ Assign FlashSystem storage IOP-hungry, latency-sensitive, business-critical workloads
- ▶ Decrease overall response times
- ▶ Increase efficiency and use across the IT stack
- ▶ Improve storage performance

Extreme performance

Tip: If you have 5 or more SSDs, it can be more cost-effective to use FlashSystem storage. Also, the latency that is caused by a software data path and application processing can be greatly reduced by FlashSystem products that have a hardware-only data path and low latency.

FlashSystem products enable businesses to unleash the power of performance, scale, and insight to drive services and products to market faster.

FlashSystem products increase application performance as much as 10x faster than other storage solutions. When compared to equivalent disk systems, IBM flash memory solutions deliver capacity in a single 1U rack, and are 19 times more cost efficient in dollars.

These solutions include the latest in industry-standard, solid-state flash memory technology, including enterprise multi-level cell (eMLC) flash technology and single-level cell (SLC) flash technology. Data is moved through the system as quickly as possible, with no bottlenecks. No cheaper, slower MLC technology is used by these products.

Capacity varies in the range of 1 - 10 TB for FlashSystem 710 and 810 models and 6 - 24 TB for FlashSystem 720 and 820 models.

IBM MicroLatency

FlashSystem products deliver extremely fast response time to accelerate critical applications. MicroLatency (that is, roughly 100-microsecond access time) enables faster decision making by facilitating an extreme-performance data path to accelerate critical applications and help you achieve a true market-based competitive advantage.

DRAM on each module helps enable fast writes at 25 microseconds. Purpose-driven, highly parallel design maximizes host CPU efficiency and productivity.

MacroEfficiency

FlashSystem products can help you consolidate hardware and software, increase deployment speed, reduce work for IT staff, and provide power and cooling savings:

- ▶ FlashSystem storage can help you achieve business benefits with the following components:
 - A 1U form factor, which has a minimal footprint for optimum ROI.
 - Two dual-port 8 GB Fibre Channel controllers or dual-port 40 GB QDR InfiniBand controllers.
 - 350 watt or less power draw.
 - Hot-swap flash modules to enable uninterrupted operations.
 - You can place up to a petabyte (PB) of FlashSystem storage in a single rack, on a single floor tile.
- ▶ FlashSystem storage systems offer the following energy-efficiency features to save energy, reduce operational costs, increase energy availability, and contribute to a green environment:
 - Energy-efficient flash components help lower operational costs.
 - FlashSystem storage systems offer one of industry's best IOPS per watt ratio to maximize energy savings.
 - FlashSystem storage systems use hexagonal ventilation holes, a part of IBM Calibrated Vecteded Cooling™ technology. Hexagonal holes can be grouped more densely than round holes, providing more efficient airflow through the system.

Enterprise reliability

FlashSystem products have durable and reliable designs that use enterprise-class flash technology and patented data-protection technology.

FlashSystem storage uses flash solid-state storage technology. FlashSystem devices are designed for cost-effective, high-storage performance that is used to accelerate critical business applications. FlashSystem devices feature patented Variable Stripe RAID (VSR™), 2D flash RAID, Active Spare support, error correction code (ECC) at the chip level, and other reliability technologies.

Two-dimensional flash RAID eliminates single points of failure and provides enhanced system-level reliability. VSR technology helps reduce business interruptions and prevent chip failures to enhance the two-dimensional protection mechanism. It also maintains performance capacity levels.

Hot-swappable flash modules and redundant components with built-in battery backup help boost data availability and IT productivity. An available integrated spare flash card limits downtime.

SVC products with FlashSystem storage

You can use IBM Storage System SVC products to add advanced storage functionality to the extreme performance of FlashSystem storage. SVC products have a maximum of 100 µs of latency and add the following features:

- ▶ Thin provisioning to allocate storage “just-in-time”
- ▶ Improved utilization to harvest all SAN capacity
- ▶ Disaster avoidance with location-proof data availability
- ▶ Easy Tier for storage efficiency
- ▶ IBM FlashCopy® for point-in-time copies
- ▶ Mirroring and copy services for data replication and protection
- ▶ Real-time compression to place up to five times more data in the same physical space

You can order FlashSystem storage individually, or an enhanced version of FlashSystem storage bundled with the SVC system, which adds the advantages of the just-mentioned functionality and storage hypervising.

For information about how to order FlashSystem storage or the FlashSystem solution bundling features, see the New IBM FlashSystem 720 and FlashSystem 820 high-performance flash memory systems Announcement A13-0197 at the following site:

http://www.ibm.com/common/ssi/ShowDoc.wss?docURL=/common/ssi/rep_ca/7/649/ENUSA13-0197/index.html&lang=en&request_locale=en#h2-prodnox

SVC compatibility

The SVC product supports a wide range of host operating systems, server platforms, Fibre Channel (FC) SAN switches, and storage controllers. An interoperability matrix for each SVC version is available at the following website:

<http://www.ibm.com/systems/storage/software/virtualization/svc/interop.html>

Server-side flash memory and PCIe cards

Server-side directly attached PCIe products typically have less latency than disk arrays. Each PCIe slot is connected to the system board using a dedicated data path or bus, which provides extremely fast performance that is critical for high-performance applications. They can also be ideal for caching and acceleration.

Although PCIe storage provides better latency for a single host than SAN-attached SSD storage with a disk form factor, and has a slightly lower latency than SAN-attached FlashSystem storage, it must be managed differently. Some PCIe adapters require software and CPU cycles for RAID striping and management, but most disk form factor SSDs can be managed by the operating system and controlled as drives. Remember that RAID performance is limited by the speed of the slowest device, so it is important to consider the speed of all storage devices. For best performance, it might also be necessary to use caching software that is optimized for PCIe.

It is also important to consider the capacity of the hosting system when you evaluate a server-side flash PCIe solution. PCIe products do not require the availability of drive bays on the host system. However, the host system must have adequate resources and processing power to manage the data, and must have enough available slots to support the required number of PCIe adapters. Some adapters might require more than one available slot. Depending on performance requirements, required availability, scalability, and advanced functionality, it might be better to select another solution such as a FlashSystem product or other external SSDs such as IBM DS8870 storage. Figure 3 summarizes key differentiators between server-side flash PCIe cards and FlashSystem storage.

	Server-Side Flash and PCIe Cards	FlashSystem Storage
Sharing	<ul style="list-style-type: none"> PCIe card benefits just one server. 	<ul style="list-style-type: none"> External Flash arrays can be shared by many servers in a SAN.
Use of Server resources	<ul style="list-style-type: none"> PCIe cards can use up to 20%-30% of a server's CPU, and use the server's memory for Flash management as well. 	<ul style="list-style-type: none"> FlashSystem arrays use no host system resources.
Availability	<ul style="list-style-type: none"> RAID often can only be set up by installing more than 1 card. No HA in PCIe cards. 	<ul style="list-style-type: none"> FlashSystem 720/820 are HA and offer the above levels of protection, plus system level 2D Flash RAID. FlashSystem 810/820 use eMLC technology which is 10x more reliable than MLC.
Scalability	<ul style="list-style-type: none"> PCIe cards do not scale. 	<ul style="list-style-type: none"> FlashSystem arrays can easily scale in a SAN environment.
Advanced Functionality	<ul style="list-style-type: none"> No ability to do snapshots, replication to other cards or arrays, etc. 	<ul style="list-style-type: none"> Together with IBM SVC, FlashSystem arrays can offer the advanced functionality of state of the art disk arrays.

Figure 3 Server-side flash PCIe cards and FlashSystem storage considerations

Also consider the configuration of the adapter when you select a solution. For example, some PCIe adapters are essentially fast flash memory that is inserted into a slot. However, some modular products include RAID controllers and flash memory modules that can be inserted into the controller. This configuration is more prone to mechanical failure, but can provide benefits such as RAID mirroring for modules on the same controller, the ability to expand storage, and reduction in software and system resources that are required to manage the controller.

PCIe products are available for both IBM System x® and IBM POWER® system software.

For System x products, IBM High IOPS PCIe adapters provide up to 640 GB of SLC storage or 1.2 TB of MLC storage. See the *IBM High IOPS SSD PCIe Adapters*, TIPS0729 IBM Redbooks® Product Guide for more information:

<http://www.redbooks.ibm.com/abstracts/tips0729.html?Open>

For POWER system software, a PCIe RAID adapter provides support for 1, 2, or 4 modules of SSD storage. Each module has 177 GB of eMLC flash memory. Modules can be mirrored at the adapter level. For more information, see Announcement 110-150, dated August 17, 2010, *PCIe SSD can reduce I/O-bound batch times and improve interactive or query response time*:

http://www.ibm.com/common/ssi/rep_ca/0/897/ENUS110-150/ENUS110-150.PDF

SSD solutions

In many ways solid-state drives (SSDs) are similar to hard disks. They are typically connected with a SATA interface and they store files just as any other drive does. The difference is that, instead of storing data on one or more spinning platters, an SSD writes and reads data to and from nonvolatile flash memory. Because they do not contain moving parts, solid-state drives deliver much better performance than HDDs.

The fastest hard disk drives can read and write data at 200 - 400 MB per second, and access data in a few milliseconds. In comparison, the fastest solid-state drives can achieve a much higher rate of transfer that essentially saturates the SATA interface, and their typical access times are a fraction of a single millisecond (ms). Because they have no moving parts, another SSD advantage is durability. Solid-state drives are not as susceptible to damage or degraded performance as HDDs. Until the availability of flash memory with a hardware-only data path, traditional SSDs and PCIe solutions were the fastest storage options available.

SSDs offer a number of potential benefits over hard disk drives:

- ▶ Up to 100 times the throughput and 10 times lower response time than 15,000 RPM spinning disks
- ▶ Lower power consumption
- ▶ Less heat generation
- ▶ Lower acoustical noise

Many IBM storage systems support SSDs:

- ▶ IBM System Storage DS8000®
- ▶ IBM Flex System™ V7000 Storage Node
- ▶ IBM XIV® Storage System Gen3
- ▶ IBM System Storage DS3500
- ▶ IBM Storage Subsystem DCS3700
- ▶ IBM System Storage DS5000

Use SSDs to optimize performance. When there are more than five SSDs, FlashSystem is a more cost- and performance-effective choice. Because SSDs are composed of flash memory that is put inside a carrier and inserted into an array like a hard disk drive, there is added latency caused by controllers, software layers, and shared data paths.

HDD solutions

Although traditional hard disk drives (HDDs) are far more advanced and higher-performing than their predecessors, on many levels the basic underlying technology remains unchanged. All HDDs consist of quickly rotating magnetic platters that are paired with read/write heads that travel over the platters' surfaces to retrieve or record data. HDD solutions are cheaper per terabyte than SSD and FlashSystem solutions, and hardware-based encryption is available. However, they are more prone to failure because of moving parts. HDDs are the most common performance bottlenecks on any system. Why? Simply because unlike RAM and CPU, hard disk drives have mechanical moving parts. This means that every time you need to open, save, modify a file, or do anything else, the disk needs to spin.

The technology is mature, reliable, and relatively inexpensive compared with other storage options. Hard disk drives are also available in relatively high capacities. Today's fastest hard disks can read and write data at a rate of 200 - 400 MB per second with access times less

than a few milliseconds. However, this speed is still significantly slower than even some of the most affordable solid-state drives and FlashSystem products.

Use HDD when you need a large amount of storage and are not as concerned about achieving peak system performance.

Hybrid solutions

Hybrid solutions typically consist of a tiered approach that incorporates HDD capacity with SSD and flash speeds. Some of the advantages of hybrid storage products include reduced cost, improved capacity, and manageability. Though HDD remains a popular drive option, most products in the IBM storage product portfolio also offer SSD capability. IBM System Storage DS8870 and XIV Storage System are offered as *hybrid configurations*, and use solid-state drives (SSDs) with hard disk drives (HDDs).

Easy Tier capability facilitates the HDD and SSD hybrid configuration usage on the DS8000 family of products such as DS8870, Storwize V7000, and SVC. IBM FlashSystem products also comprise part of the IBM all-flash technology and hybrid (disk and flash technology) solutions.

Hybrid arrays provide functionality and consolidation in addition to tiering and data movement. However, SSD competes with other disks for resources. When low latency and fast access are essential, FlashSystem products provide the best performance for Easy Tier. DS8870 systems can have higher latency than solutions with FlashSystem and SVC products, but provide enhanced support for encryption, high availability, and disaster recovery. It is also important to consider whether your environment requires IBM PowerHA® or IBM Geographically Dispersed Parallel Sysplex™ (IBM GDPS®). Consider using the DS8870 high-end disk storage systems solution to meet the stringent demands of your enterprise storage environment.

Easy Tier helps enable more effective storage consolidation by taking the guesswork out of deploying solid-state drives. Easy Tier can automatically and dynamically move the appropriate data to the appropriate drive tier in the system, based on ongoing performance monitoring. Easy Tier performs data relocation within SSD and HDD storage pools (also referred to as a *hybrid solution*) to achieve the optimum performance, by moving hottest extents to SSD, and cooler extents to HDD. It is done automatically, dynamically, and transparently to hosts, that is, without disruption to applications. This feature significantly improves the overall storage cost performance and simplifies the performance tuning and management.

Hot data and cold data

When designing a storage solution, consider the following types of data:

- ▶ Hot data is frequently accessed data. Migrating it from HDD to SSD or to FlashSystem storage can provide significant performance improvement.
- ▶ Cold data is less frequently accessed data or data that is not accessed at all. Migrating it from HDD to SSD might provide little or no benefit.

Easy Tier assumes that SSDs do not benefit much from sequential workloads and that nearline SAS disks are good candidates for data that is primarily accessed sequentially. Tier 2 features the lowest disks (Near-Line SAS), tier 1 features the next performance level (SAS), and tier 0 features the fastest disks (SSD/flash memory). DS8870 systems typically have higher latency than solutions with FlashSystem and SVC products, but provide enhanced support for encryption, high availability, and disaster recovery. It is also important to consider whether your environment requires PowerHA or GDPS support.

SVC storage grouping

The SVC product helps administrators control storage growth more effectively by moving low-activity or inactive data into a hierarchy of lower-cost storage. Administrators can free disk space on higher-value storage for more important, active data. It is achieved by creating various groups of storage, corresponding to underlying storage with various characteristics, for example, speed and reliability. With SVC software, you can better match the cost of the storage used to the value of data placed on it.

Figure 4 illustrates the Easy Tier solution running on SVC software. Using this solution, you can combine FlashSystem storage with SSDs and HDDs with manual FlashSystem data placement.

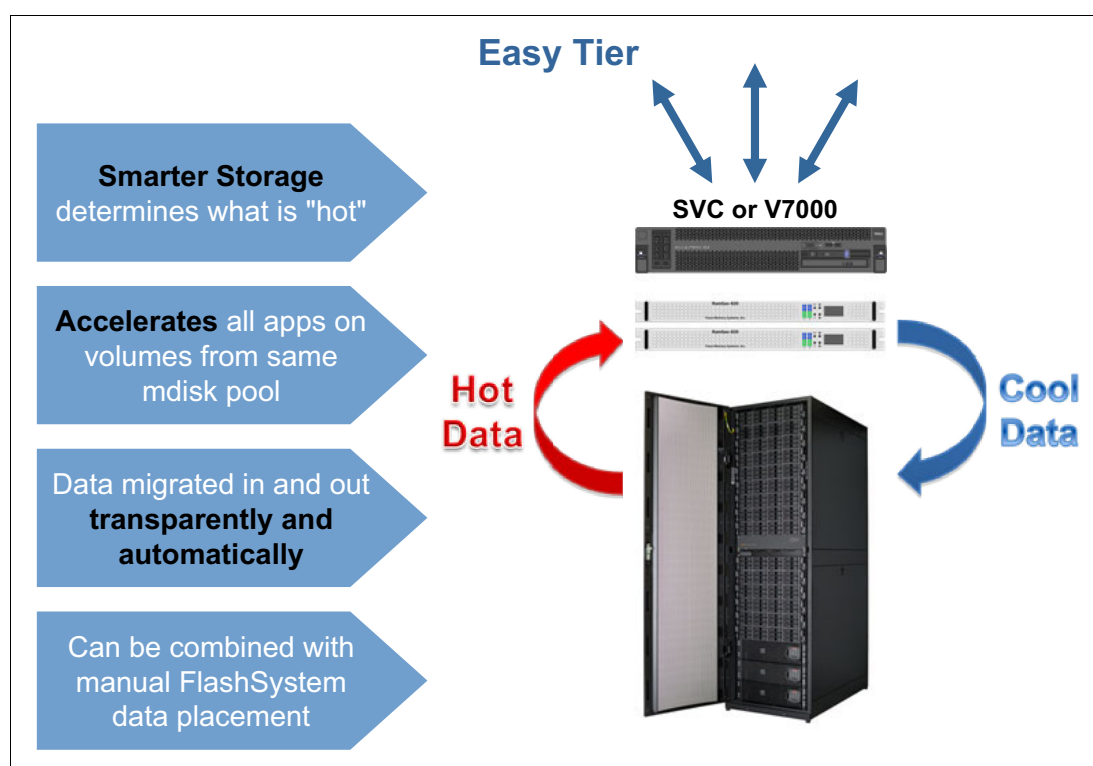


Figure 4 Easy Tier on SVC

The workload or I/O characteristics of any installation is unique, even though there are general or similar patterns. Many installations do more random I/Os during the day when online transactions are the main activity. After the main online period finishes, the batch period starts, during which time the I/Os become more sequential in nature. For example, with typical financial institutions, the market open period is a critical time because of the flood of transactions that come in during a relatively short period of time. Month-end, quarter-end, and year-end periods are generally times when I/O transactions might double or even more, compared to normal days.

Shifting workload and I/O patterns make it nearly impossible to micromanage placement of data at the extent level. No function is available to provide manual placement of data by extent, but this placement can be automatically managed by Easy Tier.

System x or PureFlex Systems

FlashSystem products can be used with System x servers and other x86 architecture servers with various standard FC interfaces. Host interfaces must be capable of at least 8 Gbps operation for maximum performance.

New components and capabilities help you go beyond blades. With a broad range of x86 compute nodes, the new Flex System V7000 storage node, enhanced networking capabilities and improved system management capabilities, you can upgrade your existing blade infrastructure or build your own system and make your IT simpler, more flexible, more open, and more efficient.

When extreme performance is important, FlashSystem storage products typically provide better performance for multiple servers or extensive data access than other solutions such as server-based PCIe storage. FlashSystem storage has latency of less than a microsecond, but latency for PCIe solutions is typically in the range of milliseconds. Implementing FlashSystem storage can help reduce costs for servers and licensing, because fewer systems are required in the environment. Its fast access time and low latency can reduce response time and the amount of time required to process jobs.

FlashSystem products require minimal reconfiguration of your storage environment. You can move your most I/O-intensive applications to FlashSystem storage, and improve performance of your entire application environment. In a tiered environment, FlashSystem storage can also be used as tier 0 storage.

IBM Intelligent Cluster

IBM Intelligent Cluster™ solutions are built, tested, delivered, and installed by IBM, and supported as a single end-to-end solution instead of being treated as hundreds of individual components. IBM provides single point-of-contact, solution-level support that includes both IBM and third-party components to deliver maximum system availability throughout the life of the system, so you can spend less time maintaining systems and more time delivering results.

The IBM Platform High Performance Computing (HPC) comprehensive management solution lets you quickly put the system to work with minimal effort and skill. Platform HPC optimizes resources automatically to maximize usage, and the easy-to-use web-based tools make application deployment and ongoing management simple.

High-performance flash memory systems for Intelligent Clusters were announced 04/30/2013 and include rack mount FlashSystem storage products. In this solution for enterprise deployment, you can take advantage of IBM FlashSystem 720 and FlashSystem 820 storage. VSR technology included in IBM flash memory capability enhances system resiliency without sacrificing performance or usable capacity.

Note: See Announcement 113-067, dated April 30, 2013, *New IBM FlashSystem portfolio of high-performance flash memory systems for Intelligent Clusters* for details:

http://www-01.ibm.com/common/ssi/rep_ca/7/897/ENUS113-067/ENUS113-067.PDF

Flash memory systems for Intelligent Clusters

In addition to VSR data protection, FlashSystem 720 and FlashSystem 820 storage systems incorporate other reliability features, including:

- ▶ ECC to provide bit-level reconstruction of data from flash chips
- ▶ Checksums and data integrity fields to protect all internal data transfers within the system
- ▶ Over-provisioning to enhance write endurance and decrease write amplification

Use FlashSystem in Intelligent Cluster for High Performance Computing environments to maximize your storage and reduce latency (see Figure 5).

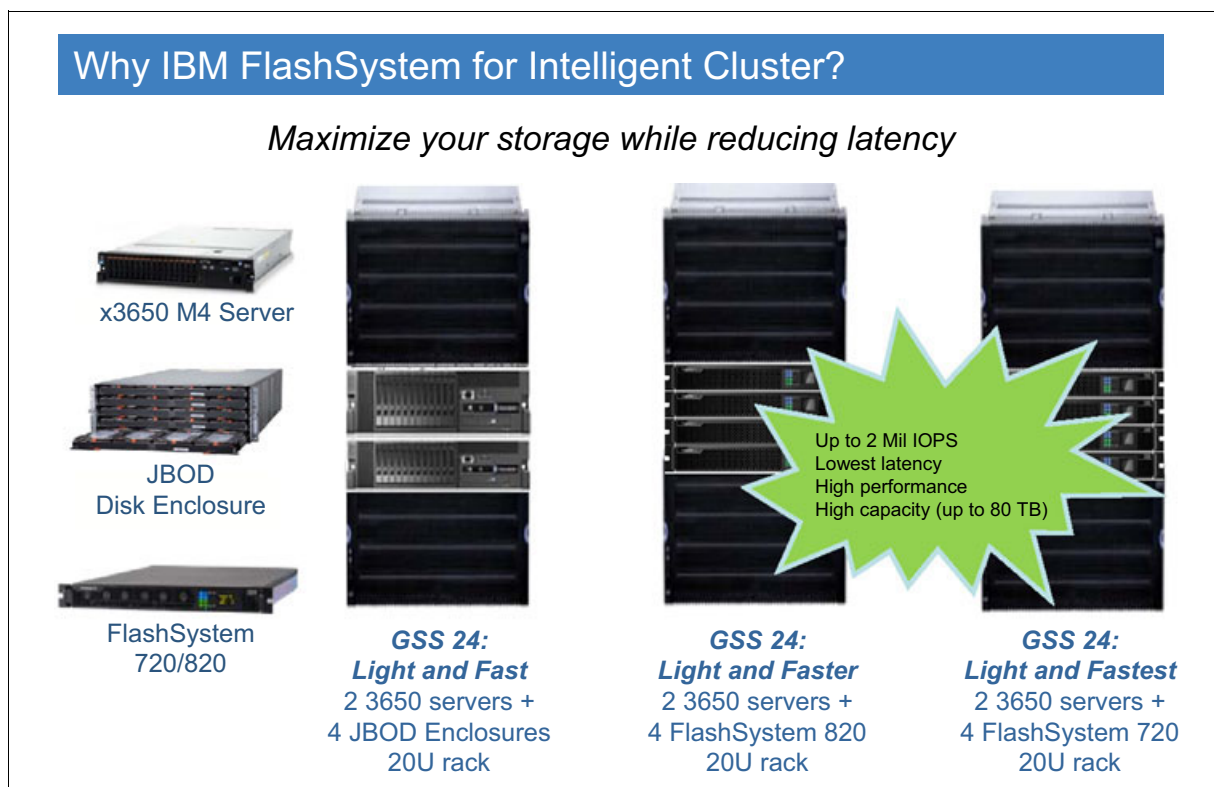


Figure 5 Intelligent Cluster and FlashSystem storage

For additional FlashSystem in Intelligent Cluster information see:

<http://www-03.ibm.com/systems/x/hardware/largescale/cluster/index.html>

Use cases

This section describes use cases for integrating IBM FlashSystem products into an IT environment.

FlashSystem products can be ideal for OLTP and OLAP applications. OLTP handles many short transactions in quick, random bursts. Use FlashSystem products to accelerate OLTP for a variety of business needs:

- ▶ Analytics
- ▶ Financial applications
- ▶ Gaming
- ▶ Real-time billing

- ▶ Trading
- ▶ Real-time monitoring
- ▶ Query acceleration for databases such as DB2 and Oracle

OLAP handles larger chunks of data, but must retrieve data from many sources and respond to complex queries quickly. You can use OLAP for a similarly diverse set of business needs:

- ▶ Business intelligence
- ▶ Batch processing
- ▶ ERP systems
- ▶ Reporting
- ▶ Massive data feeds

Also consider FlashSystem storage for SAP and video applications.

Online analytical processing (OLAP)

You commonly use data warehouses are commonly used with OLAP workloads in decision support systems, such as financial analysis. OLAP queries are typically complex, and they process large volumes of data from multiple sources. Accurate, real-time operational data is critical for analytics. Delays can significantly increase your business and financial risks.

Many servers have adequate RAM and processor power to process massive amounts of data (frequently referred to as *big data*). However, the I/O that is required for reading data from storage for processing in the OLAP database server can frequently reduce performance. Delays come primarily from batch data loads and performance issues due to handling heavy, complex queries that use I/O resources.

IBM FlashSystem storage solutions help address these challenges in the following ways:

- ▶ Dramatically boosting the performance of OLAP workloads with distributed scale-out architecture, providing almost linear and virtually unlimited performance and capacity scalability
- ▶ Significantly improving response time for better and more timely decision making

An OLAP solution with IBM flash memory systems consists of a set of components:

- ▶ Database servers (IBM System x or IBM Power Systems™) to run data management software such as DB2, Microsoft SQL Server, or Oracle Database
- ▶ Flash memory systems (IBM FlashSystem 720 or 820) to host the entire data set or partitioned subsets of data
- ▶ A private network (such as 10 Gb Ethernet or QDR/FDR InfiniBand) that is used to provide high-speed connectivity across database servers in a cluster
- ▶ A storage area network (SAN) that is used to provide connectivity across database servers and storage systems

Using FlashSystem products in SAP landscapes

Figure 6 shows an example of how to integrate a FlashSystem into your SAP landscapes.

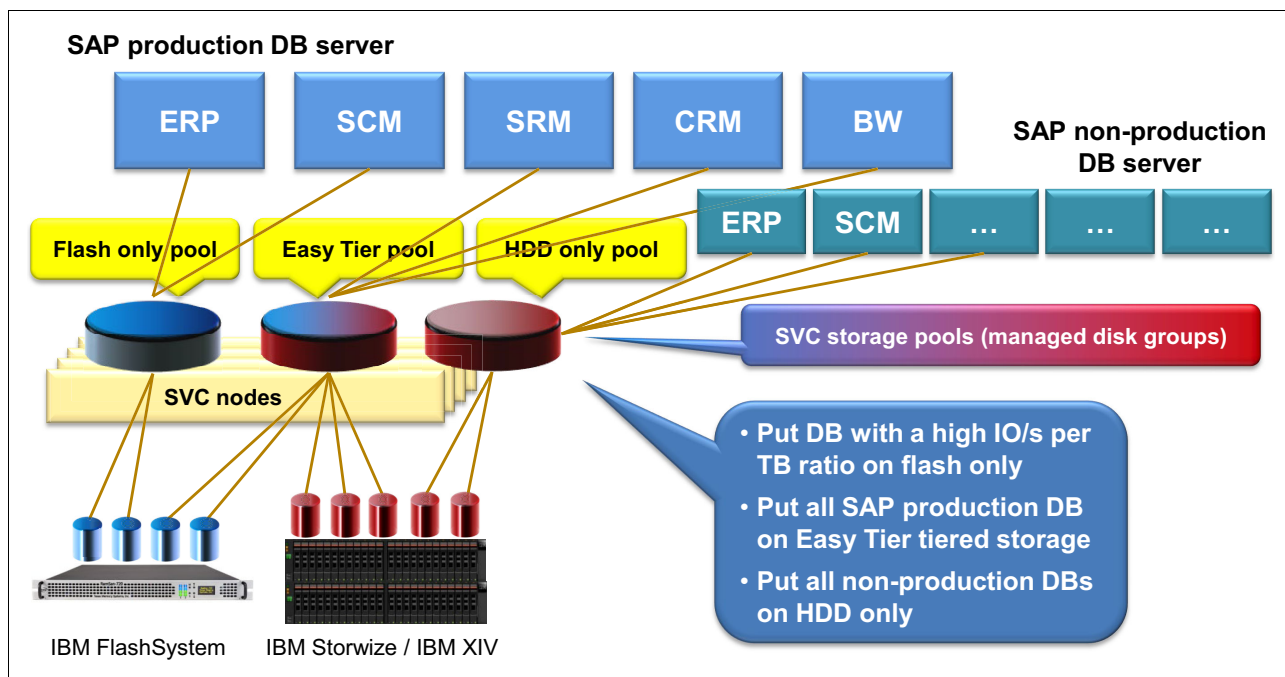


Figure 6 Integrating an IBM FlashSystem into SAP landscapes

In this example, volumes (logical unit numbers, or LUNs) are created on the storage system. You put these volumes into the storage pools of the SVC Managed Disk Group as Managed Disks (Mdisks). There are three pools:

- ▶ A flash-memory-only pool
- ▶ An Easy Tier pool
- ▶ An HDD-only pool

You also put the SAP databases into the pools. For example, place a high performance production database into the flash-memory-only pool. Place all other production databases into Easy Tier. You might also include quality assurance (QA) systems. Put any remaining databases, such as Development and Test, into the HDD-only pool.

Using FlashSystem products to improve video production

Suppose that you work for a large digital content provider that needs to provide fast video delivery to hundreds of partners. To make the video content compatible with the variety of platforms, every piece of content must be transcoded. Video transcoding allows content that is created in one format to be playable on a different format. Rapid transcoding and transmission are vital to your business. Your current disk-based storage performs I/O operations too slowly to keep up with the transcoders. Slow storage affects both the quality of transcoding and the delivery of content.

You investigate several options for a new, nondisruptive storage solution. SSD solutions do not provide the required performance. You select a FlashSystem 710 product, which has 5 TB of SLC flash memory that delivers 400,000 IOPS read and write performance.

The FlashSystem 710 storage is integrated into the existing 8 GB Fibre Channel SAN infrastructure with no changes to code or workflow. The time that is required for transcoding

jobs drops by as much as 70 percent. Your new solution reduces storage bottlenecks, which improves the speed and efficiency of other file operations by as much as 40%.

Supported platforms

For supported hardware and platforms for the IBM FlashSystem portfolio of products refer to the following link:

<http://www.ibm.com/systems/support/storage/ssic>

Related information

For more information, see the following documents:

- ▶ IBM FlashSystem family product page
<http://www.ibm.com/storage/flash>
- ▶ IBM Redbooks Solution and Product Guides for IBM FlashSystem family
<http://www.redbooks.ibm.com/redbooks.nsf/searchsite?SearchView&query=flashss>
- ▶ IBM Redbooks Product Guides for IBM Flex System
<http://www.redbooks.ibm.com/Redbooks.nsf/portals/puresystems?Open&page=pgbycat>
- ▶ IBM Flex System Product page
http://www.ibm.com/systems/pureflex/flex_overview.html
- ▶ IBM Flex System Enterprise Chassis Product Guide, TIPS0863
<http://www.redbooks.ibm.com/abstracts/tips0863.html>
- ▶ IBM Disk Storage Systems Product page
<http://www.ibm.com/systems/storage/disk/index.html?LNK=browse>
- ▶ DS8000: Introducing Solid State Drives IBM Redpaper publication
<http://www.redbooks.ibm.com/abstracts/redp4522.html?Open>
- ▶ IBM System Storage Solutions Handbook IBM Redbooks publication
<http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/sg245250.html?Open>
- ▶ IBM Support Portal
<http://ibm.com/support/entry/portal/>
- ▶ IBM System Storage Interoperation Center (SSIC)
<http://www.ibm.com/systems/support/storage/ssic>
- ▶ IBM System x Solid state storage adapters product page
<http://www.ibm.com/systems/x/options/storage/solidstate/adapters.html>

Authors

This paper was produced by a team of specialists from around the world working at the International Technical Support Organization, Poughkeepsie Center.

Megan Gilge is a Technical Writer in the IBM International Technical Support Organization. Before joining the ITSO, she was an Information Developer in the IBM Semiconductor Solutions and User Technologies areas. Megan holds a degree in Liberal Arts from Michigan Technological University.

Karen Orlando is a Project Leader at the International Technical Support Organization, Tucson Arizona Center. Karen has over 25 years in the IT industry, with extensive experience in open systems, and information and software development of IBM hardware and software storage. She holds a degree in Business Information Systems from the University of Phoenix and is Project Management Professional (PMP)-certified since 2005.

Special thanks for his contributions to this project to:

Levi B Norman
IBM Systems & Technology Group, Storage Platform

Now you can become a published author, too!

Here's an opportunity to spotlight your skills, grow your career, and become a published author—all at the same time! Join an ITSO residency project and help write a book in your area of expertise, while honing your experience using leading-edge technologies. Your efforts will help to increase product acceptance and customer satisfaction, as you expand your network of technical contacts and relationships. Residencies run from two to six weeks in length, and you can participate either in person or as a remote resident working from your home base.

Find out more about the residency program, browse the residency index, and apply online at:

ibm.com/redbooks/residencies.html

Stay connected to IBM Redbooks

- ▶ Find us on Facebook:
<http://www.facebook.com/IBMRedbooks>
- ▶ Follow us on Twitter:
<http://twitter.com/ibmredbooks>
- ▶ Look for us on LinkedIn:
<http://www.linkedin.com/groups?home=&gid=2130806>
- ▶ Explore new Redbooks publications, residencies, and workshops with the IBM Redbooks weekly newsletter:
<https://www.redbooks.ibm.com/Redbooks.nsf/subscribe?OpenForm>
- ▶ Stay current on recent Redbooks publications with RSS Feeds:
<http://www.redbooks.ibm.com/rss.html>

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing, IBM Corporation, North Castle Drive, Armonk, NY 10504-1785 U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

© Copyright International Business Machines Corporation 2013. All rights reserved.

Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

This document REDP-5020-00 was created or updated on June 26, 2013.



Send us your comments in one of the following ways:

- Use the online **Contact us** review Redbooks form found at:
<http://ibm.com/redbooks>
- Send your comments in an email to:
http://redbooks@us.ibm.com
- Mail your comments to:
IBM Corporation, International Technical Support Organization
Dept. HYTD Mail Station P099
2455 South Road
Poughkeepsie, NY 12601-5400 U.S.A.




Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. These and other IBM trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at <http://www.ibm.com/legal/copytrade.shtml>

The following terms are trademarks of the International Business Machines Corporation in the United States, other countries, or both:

Calibrated Vectored Cooling™
DB2®
DS8000®
Easy Tier®
FlashCopy®
FlashSystem™
GDPS®
Geographically Dispersed Parallel
Sysplex™

IBM®
IBM FlashSystem™
IBM Flex System™
Intelligent Cluster™
Parallel Sysplex®
POWER®
Power Systems™
PowerHA®
Redbooks®

Redpaper™
Redbooks (logo) ®
Storwize®
System Storage®
System x®
XIV®

The following terms are trademarks of other companies:

VSR, and the Texas Memory Systems logo are trademarks or registered trademarks of Texas Memory Systems, an IBM Company.

Microsoft, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.