Realtime publishers

The Shortcut Guide to Balancing Storage Costs and Performance with Hybrid Storage

sponsored by

FUSION-iO

Dan Sullivan

Copyright Statement

Chapter 1: Advantages of Hybrid Storage

The past several years have witnessed rapid changes in the way IT services are delivered. Virtualization is commonly used to improve the efficiency of both compute and storage services, and the advent of software-defined networking is bringing the benefits of virtualization to network services. A new generation of tools for analyzing large volumes of data are becoming established as common parts of enterprise IT infrastructure. Employees are using their own tablets, smartphones, and laptops for work-related activities, driving the adoption of virtual desktop infrastructure and related services. Although these changes are pushing the established boundaries of different parts of IT operations, they share a common characteristic: These changes place increasing demands on IT service providers to increase levels or application and storage performance or at least maintain existing levels of performance in the face of increasing demand (see Figure 1.1).

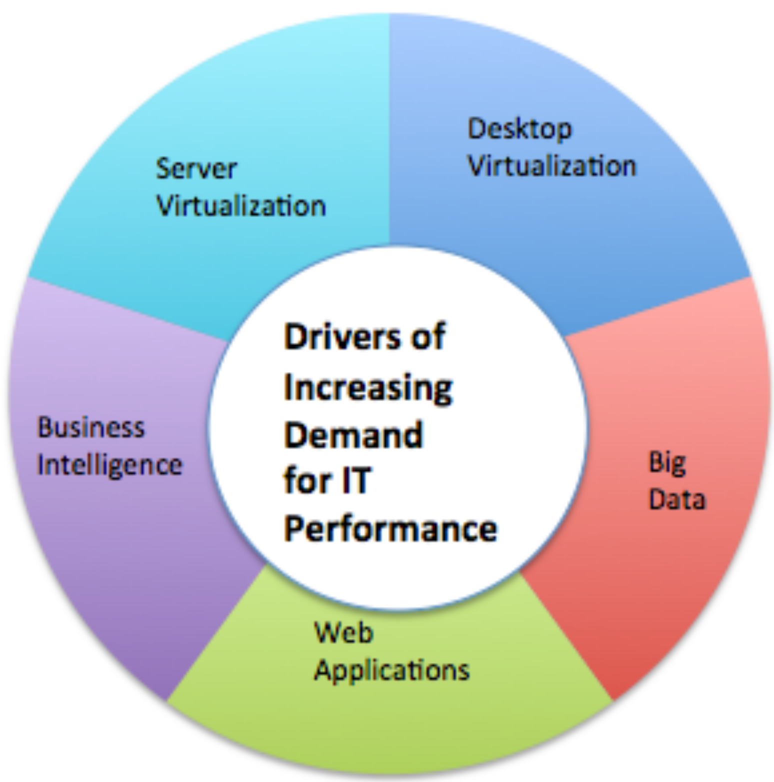

Figure 1.1: Multiple factors are increasing demands for application and service performance.

Although IT professionals are focused on meeting service level requirements with well-designed software and scalable, reliable architectures, these professionals are constrained by limited resources. Reasonably talented IT professionals with ample budgets can design multi-tiered systems to address a range of business requirements while meeting performance demands. But what if budgets are constrained? Tradeoffs have to be made and design decisions about algorithms and architectures can have a significant impact on the performance of mission-critical applications.

Overview of Flash Deployment in Hybrid Storage Systems

The Shortcut Guide to Best Practices for Flash Deployment in Hybrid Storage is designed to help systems architects, IT managers, and directors understand their options when it comes to deploying flash storage technology in their environment. This guide is divided into four chapters, each of which addresses a relevant aspect of flash technology and hybrid storage.

In this first chapter, the focus is on the advantages of hybrid storage and the need to balance the benefits of flash performance with the additional costs relative to disk storage. The chapter also addresses requirements for effective and efficient hybrid storage deployments.

The second chapter considers the role of hybrid flash and disk storage in server virtualization and desktop virtualization. It also addresses the benefits of manageability with regards to quality of service commitments.

In the third chapter, the focus shifts to database applications and the implications of disk-based storage for database performance as well as the challenges of maintaining consistent levels of performance with storage Quality of Service (QoS) features, such as guaranteed IOPS and latency service level. It also covers the benefits of flash for maintaining consistent I/O performance and the ability to scale Input/Output Operations Per Second (IOPS) as needed.

The guide concludes with a discussion of best practices for deploying hybrid flash in the enterprise. Topics include analyzing workloads, architecture issues, real-time QoS controls, data management practices, and planning for hybrid storage.

Balancing Cost and Performance: Advantages of Flash and Disk Drives

IT professionals and business users that depend on IT technology have been fortunate. The computing resources available to us have grown at exponential rates. Moore's Law, for example, conjectures that the number of transistors on an integrated circuit double about every 2 years, and David House of Intel predicted that the performance of integrated circuits would double every 18 months. Unfortunately, in spite of this type of sustained advancement in chip technology, the same cannot be said about the performance of disk drives.

The volume of data that can be stored on hard disk drives has certainly grown. Just several decades ago, disk storage was measured in hundreds of kilobytes, and multi-terabyte disks are common today (see Figure 1.2). Disk technology is keeping pace with the demand to store increasing volumes of data, but the speed of I/O operations on those devices is not growing at the same rate.

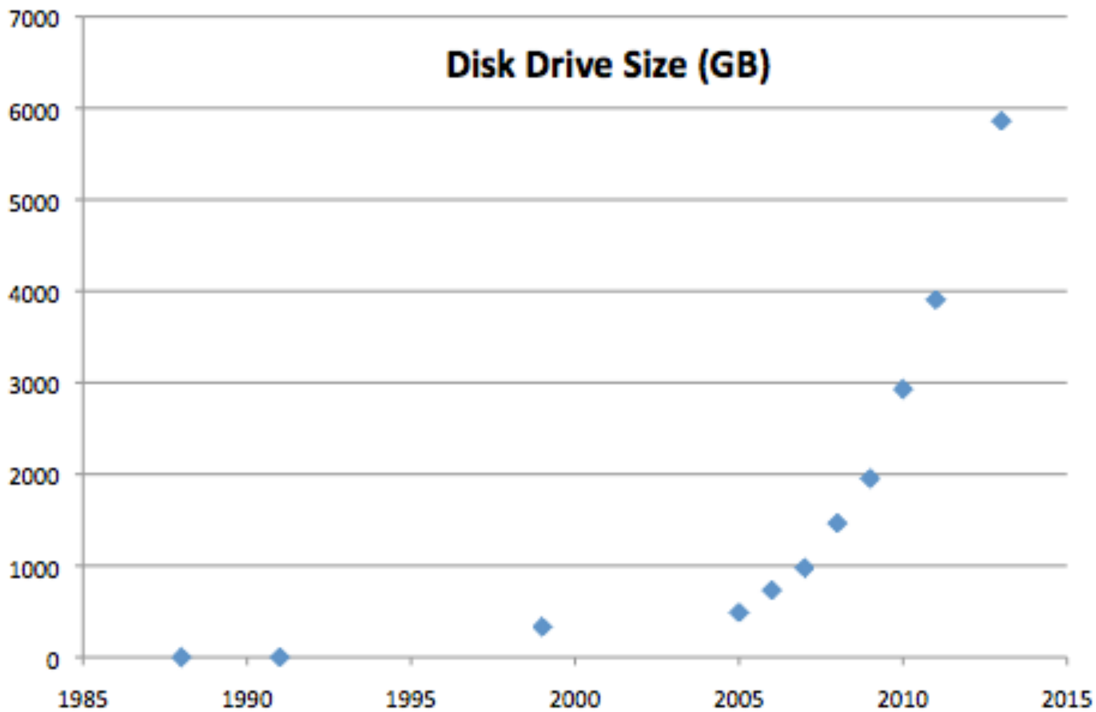

Figure 1.2: The capacity of disk drives has increased exponentially in recent years, enabling advances in business intelligence, big data analytics, and the expansion of business-critical applications.

Although the storage capacity of hard disk drives has significantly increased, the performance of those drives has not kept pace with improvements in CPU performance (see Figure 1.3).

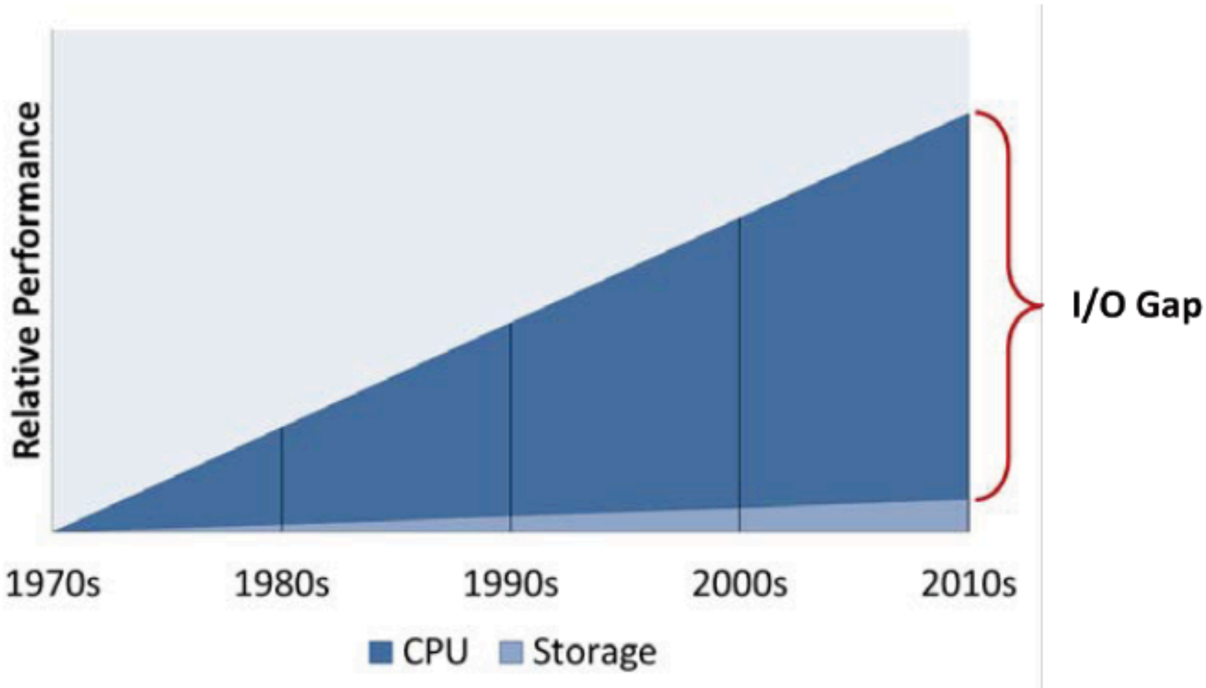

Figure 1.3: The growth in CPU performance has far outpaced the performance of storage devices over the past several decades. As a result, there is a gap in I/O performance relative to CPU performance.

The relative disparities in performance between computing and storage resources can adversely affect application performance. You might find, for example, that your CPU utilization is low because your I/O-bound application is limited by the speed of hard disk drives and network latencies. A logical course of action in this case is to virtualize the server to improve CPU utilization. However, doing so might, in fact, exacerbate the problems with I/O performance because more processes will be contending for limited I/O resources.

To maximize the benefit of advanced CPUs, you need to architect systems at a level that accounts for the strengths and weaknesses of all major components, including CPU, storage, and networking. It is here that the advantages of flash storage become clear.

Advantages of Flash Storage

Flash storage devices have a number of advantages over hard disk drives. Flash devices have no moving parts and are therefore less susceptible to mechanical breakdown. Although reliability and lack of noise are important advantages, perhaps the most significant advantage of flash storage is performance. Flash storage can sustain hundreds of thousands of IOPS, compared with hundreds of IOPS from hard disk drives, making flash storage ideal for data-intensive operations, such as those involving databases.

The disadvantage of flash is, not surprisingly, the cost. Flash storage can cost significantly more per gigabyte than hard disk drives cost.

Advantages of Hard Disk Drives

Hard disk drives, as noted earlier, are providing increasingly large amounts of storage capacity. When storage use cases call for large volumes of storage but with limited needs to low-latency response times, disk storage has a clear advantage. Archival storage, for example, is appropriate for hard disk drive storage. Applications that execute high volumes of sequential writes, such as logging, are appropriate applications for hard disk drives.

Mixed Workload Environments

It is not unusual to have a mixed workload environment with a combination of mission-critical, business-critical, and non-critical applications. Mission-critical applications might include email, databases, and customer-facing Web applications. When these systems are down or functioning poorly, there is a direct impact on productivity, customers' abilities to interact with the business, and ultimately the bottom line. As a result, mission-critical applications should maintain the highest levels of QoS even if it adversely impacts non-mission-critical applications (see Figure 1.4).

Business-critical applications are important to the operations of an organization but are somewhat resilient to variations in performance. For example, business intelligence reports might normally be distributed to managers by 9am every day. If for some reason the reports are delayed an hour or two, it might be an inconvenience and delay some management operations, but the adverse impact is less than if a mission-critical application were down. Similarly, application performance can be somewhat inconsistent without severely adversely impacting operations.

Non-mission-critical applications have the greatest tolerance for variation in performance. Consider a case in which mixed workloads are running on several virtual machines.

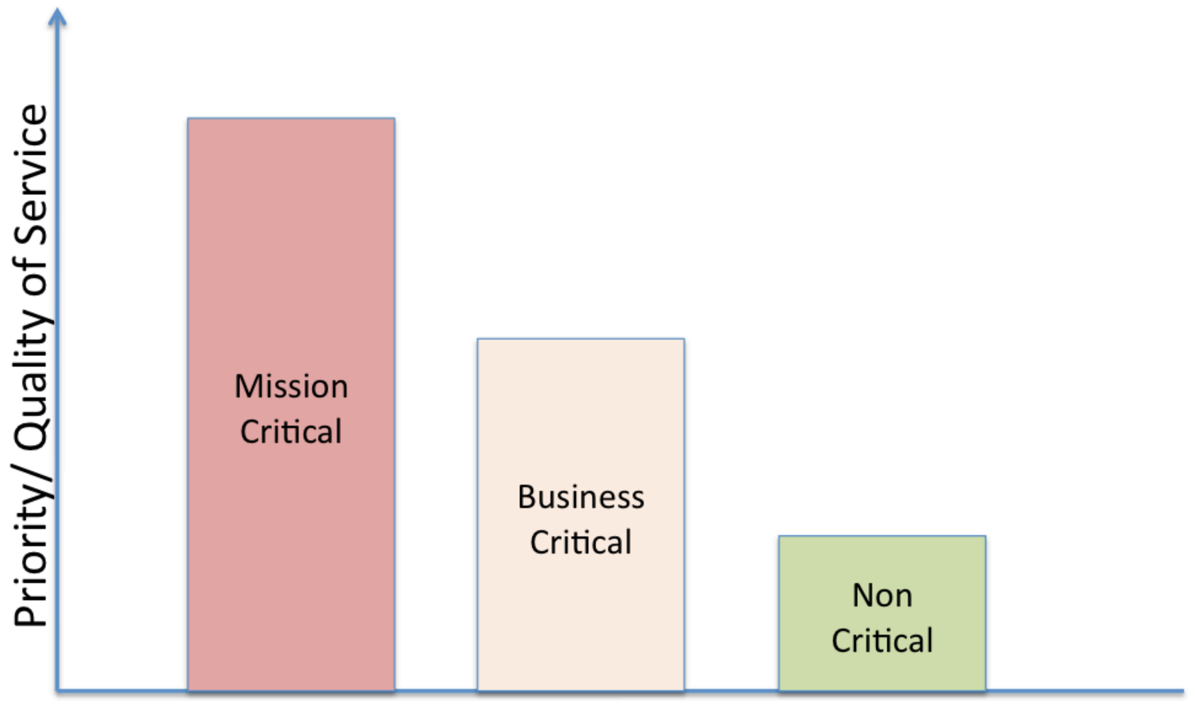

Figure 1.4: Mixed workloads typically consist of applications with varying levels of importance to an organization. One of the key considerations when designing an architecture for mixed workloads is understanding and meeting QoS requirements.

To meet the needs of mission-critical applications at affordable costs, a hybrid storage system based on a combination of flash storage and hard disk drives is called for.

Requirements for Effective and Efficient Hybrid Storage

In particular, an effective and efficient hybrid storage deployment will entail three key characteristics:

- Optimum use of flash
- Cost effective serial-attached SCSI (SAS) disk drives
- QoS performance management

It is important to note that all three characteristics are essential to maximizing the value of a hybrid storage system. Addressing one or two of these qualities might lead to marginal improvements in performance or savings in storage costs, but applications will not realize the full potential benefit of a well-designed hybrid storage system.

Fast Flash Implementations

The physics of flash devices determine the upper bounds on the potential performance of the devices, but other implementation choices can significantly alter the performance profile of a flash drive:

- The form factor used to house the flash drive
- The location of the flash drive relative to the CPU
- The use of application-specific interfaces

The choices you make with regards to flash implementations help to determine the latency of read and write operations to flash devices. For example, if flash is installed in a storage array, read and write operations are dependent on the network. Network latencies can significantly impact the speed at which I/O operations complete.

The protocols that are used to access a flash device also influence the speed of I/O operations. When read and write operations have to be mapped through multiple layers of protocols, the speed of the flash device is adversely affected.

Also, consider the reliability and endurance of flash. Over time, NAND chips on flash devices can fail. Wear-leveling algorithms that equalize the use of NAND chips can increase the useful life of a flash device. In addition, flash designs that employ redundancy at the chip and block level while isolating failed chips can avoid disruption of operations due to chip failures.

Cost-Effective SAS Drives

SAS drives offer favorable cost-per-gigabyte ratios. When an organization has to store and manage large volumes of data in highly reliable and redundant ways, SAS drives are a logical choice.

From an architectural perspective, the goal of maximizing performance while minimizing costs becomes a balancing act between the amount of more costly but higher performance flash drives and less costly but poorer performing SAS disk drives (see Figure 1.5).

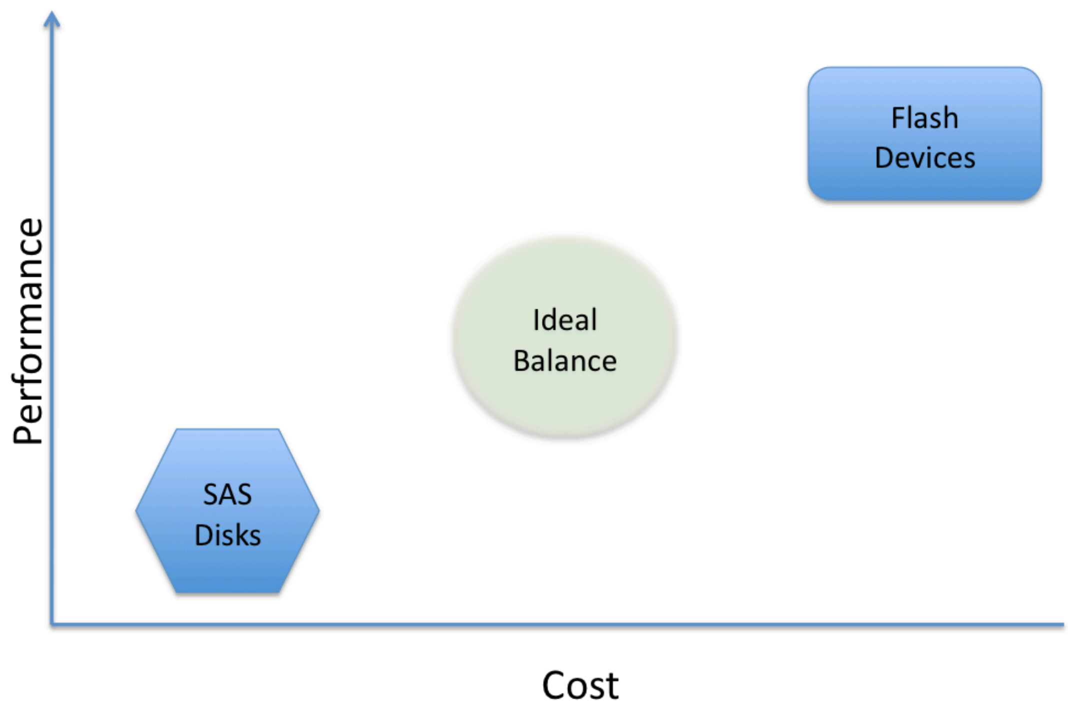

Figure 1.5: The ideal hybrid deployment leverages the cost effectiveness of SAS disks with the performance advantages of flash devices.

QoS Performance Management

The third element of efficient and effective hybrid storage solutions is the ability to allocate resources according to QoS requirements. Mixed workloads are common. Some workloads, especially mission-critical applications, should have priority over less important applications.

In virtualized environments, business-critical applications may run on the same physical server as non-critical applications run. This setup can make sense when the goal is to optimize CPU utilization. A business-critical application may only require 60 to 70 percent of CPU availability of a server, leaving 30 to 40 percent for other applications. During peak demand periods, however, the business-critical application should take precedence over the non-critical applications.

Hybrid storage solutions that implement policy-based resource allocation can be tailored to effectively allocate storage resources according to business value and simply by a general resource allocation algorithm.

Hybrid storage solutions combine flash storage with hard disk drive storage to take advantage of the benefits of both. There are several ways to deploy hybrid storage in a multi-tiered environment, and each approach has advantages and disadvantages.

Options for Deploying Flash Storage

Flash storage is a technology based on integrated circuits and lends itself to more deployment strategies than are found with hard disk drives. Several options for deploying flash storage exist:

- Solid state drive (SSD) flash
- Peripheral Component Interconnect Express (PCIe) flash
- Local to server flash used as cache
- All flash array
- Hybrid array

Each method has its advantages but not all are suitable for enterprise storage environments. As the following sections describe each method, it is important to keep in mind the over-arching need to maintain high I/O performance while controlling costs. Some solutions favor performance, such as all-flash array, while others have favorable cost characteristics, such as SSD in a server. The most common objective IT staff face is finding an optimal balance between cost and performance.

SSD Flash

SSD flash combines flash storage devices with hard disk drive form factors. From the outside, SSD look similar to hard disk drives; for example, SSD can have 2.5" form factors with the same interface as a hard disk drive. The SSD flash also supports the same kind of block I/O protocols that hard disk drives support. This setup allows SSD to be used wherever a standard hard drive could be used.

An advantage of this approach is that there are virtually no changes to hardware or low-level I/O software required to use SSD. This benefit can be appealing, especially for power users who want their laptops to boot quickly or analysts working with modest data sets on their workstations. Existing hard drives can be swapped out and replaced with faster SSD for data storage or new SSD can be installed to store the operating system (OS) to speed boot operations. This solution is not optimal for server storage, however (see Figure 1.6).

The overhead of controller bottlenecks, disk interfaces, and RAID firmware can all adversely impact performance of the flash drive. In addition, when RAID 5 or 6 is used, some of the flash storage is used for parity or mirroring. This setup reduces the total amount of storage available to applications and increases the average cost of usable gigabytes of the flash drive.

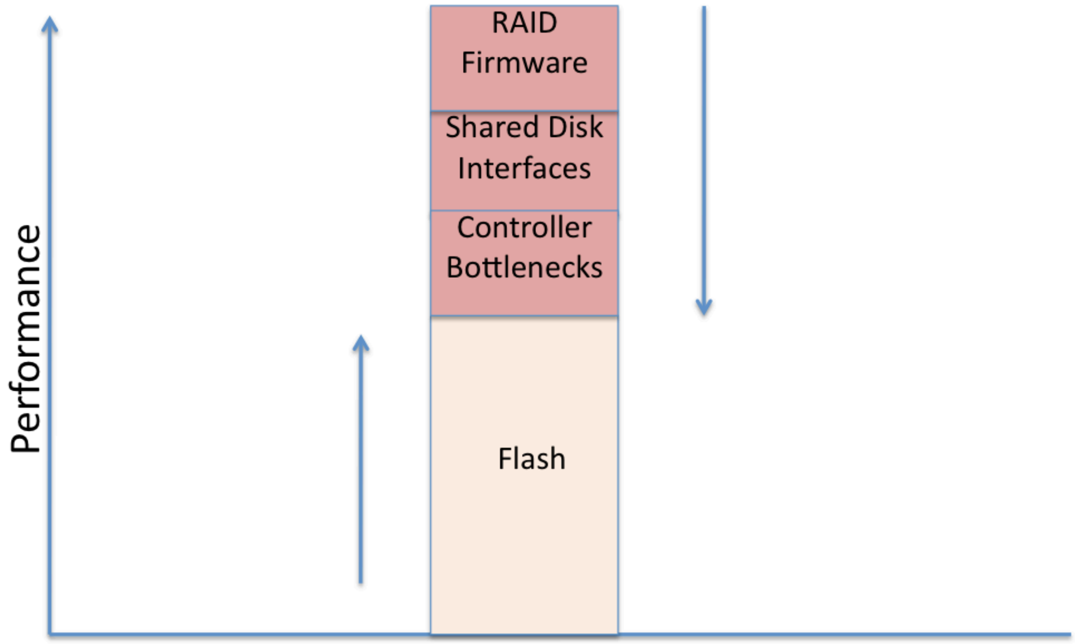

Figure 1.6: Flash circuitry increases performance relative to disk drive technology; however, disk drive protocols and interfaces hinder overall performance. The ability to plug-in flash drives to hard drive infrastructure comes at a substantial performance cost.

PCIe Flash

PCIe high-speed serial bus is a standard for data interchange based on point-to-point topology. Data between devices on a PCIe bus can be transmitted across multiple lanes, enabling high throughput.

A clear advantage of a PCIe-based implementation is that data can be transferred to and from the CPU at faster speeds than if the data were transmitted over the network to a shared storage device, which Figure 1.7 depicts. Similarly, PCIe in the server outperforms SSD in the server by avoiding unnecessary controllers, embedded processors, and storage protocols.

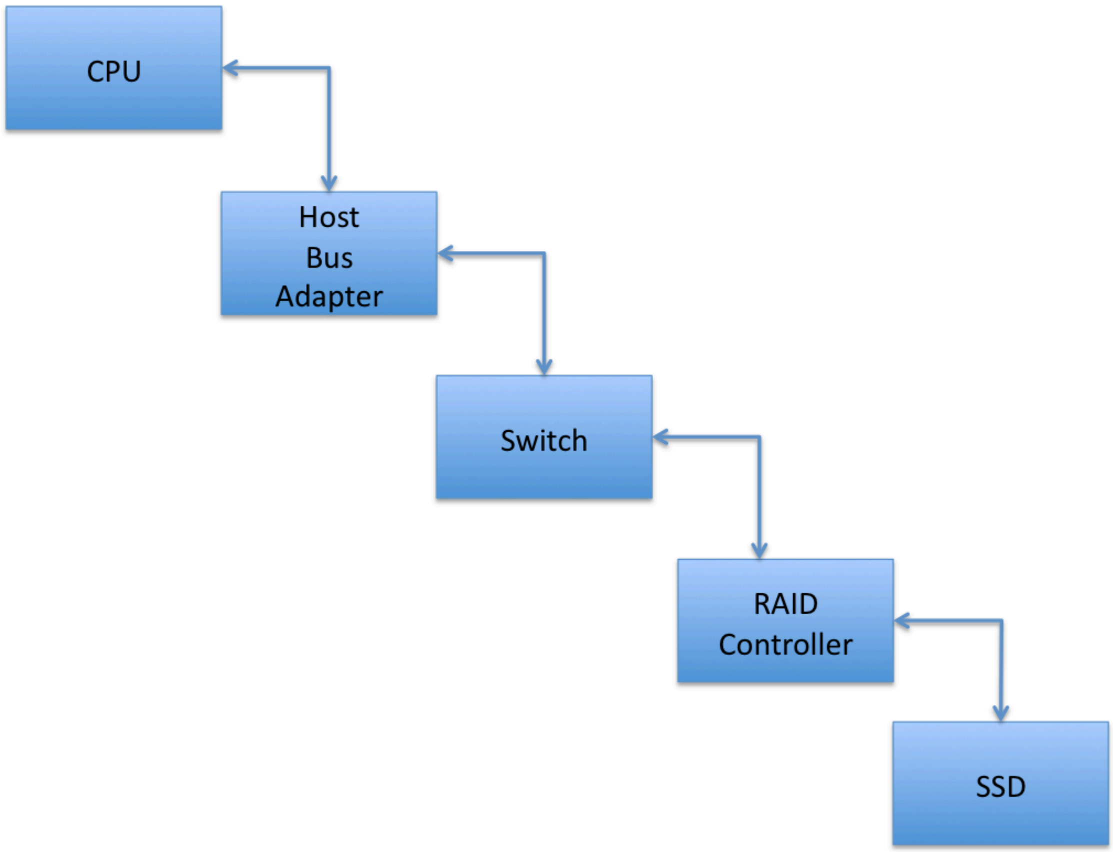

Figure 1.7: Moving data back and forth from a CPU to an SSD entails long latencies due to the number of components that must be traversed to fetch or store data.

PCIe-based flash storage entails a much simpler data path, as Figure 1.8 shows.

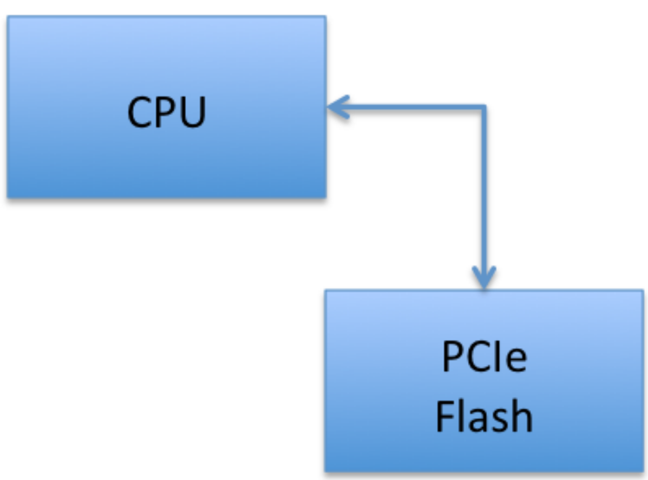

Figure 1.8: PCIe-based flash storage improves latency over SSD-based implementations.

Local to Server as Cache

A third method for deploying flash drives is to install the flash device local to a server as a cache or direct attached storage (DAS). Using flash as a cache local to a server can improve the speed of I/O operations. Reads from flash are faster than are reads from hard drives, so data-intensive operations, such as big data analytics and high-utilization virtualized servers, can avoid some of the latency associated with reading from hard disk drives.

If high availability is a concern when using flash and DAS, then pairs of flash drives can be installed on a server with the second drive used as a redundant copy. As flash memory is consumed, data can be written to less expensive disk drives in the background without adversely impacting application performance.

In highly virtualized environments, there may be many servers with virtual machines migrated across the set of available servers. Unless all servers are configured with local flash drives for caching or as DAS, users might find inconsistent performance across their virtual machines. Virtualization platforms may further restrict the use of local storage. In order to take advantage of some advanced virtualization features, such as virtual machine migration, the virtual machines are required to use shared memory. In such cases, the advantages of flash devices can be realized by using an all-flash array.

All-Flash Array

As the name implies, an all-flash array uses the solid state technology for all storage. There are no hard drives, which means no mechanical parts, more reliability, less noise, and higher performance. An all-flash storage array provides the best performance possible compared with hard disk drive arrays. The drawback, of course, is cost.

The cost of flash drives has decreased in the past several years and a number of vendors are using deduplication to drive down the cost per gigabyte of all-flash arrays. There was a time in which hard drive prices were climbing (due primarily to a disruption in the supply chain caused by a natural disaster), but those prices have leveled off. For the near future, hard drives will likely continue to maintain a cost advantage over flash drives.

There might be some business cases that would justify an all-flash array. For example, real-time currency or equities trading might involve such costly transactions that delays in executing such trades could warrant the lowest transaction latencies possible. In such a scenario, an all-flash array would likely be one of several high-performance components in the application infrastructure.

Ideally, you want the performance benefits of an all-flash array but with a price point approaching that of a hard disk drive storage array.

Hybrid Storage Arrays

Hybrid storage arrays are today's best option for balancing cost and performance. Flash devices provide performance levels needed by demanding applications such as database servers, virtual desktop infrastructures, and business intelligence applications. Hard disk drives provide large volume storage options while maintaining reasonable cost-per-gigabyte ratios.

There are multiple ways to configure hybrid storage arrays. One could, for example, simply replace some hard drives in an array with SSD. This chapter has already discussed the limitations of that approach. Although there is some performance improvement, this setup does not realize the full potential of flash performance.

To optimize the benefits of hybrid storage, the hybrid device needs to be controlled by software that leverages the inherent strengths of each type of storage device. Flash drives should be used as the point of first read and write. Doing so allows I/O operations to complete to a point that a logical transaction can be considered committed. If such data was only stored in flash drives, then you would need large-volume flash drives to accommodate typical enterprise applications. This setup is difficult to justify, in most cases. As Figure 1.9 illustrates, a better option is to write data from flash drive to hard disk drive after a transaction is complete.

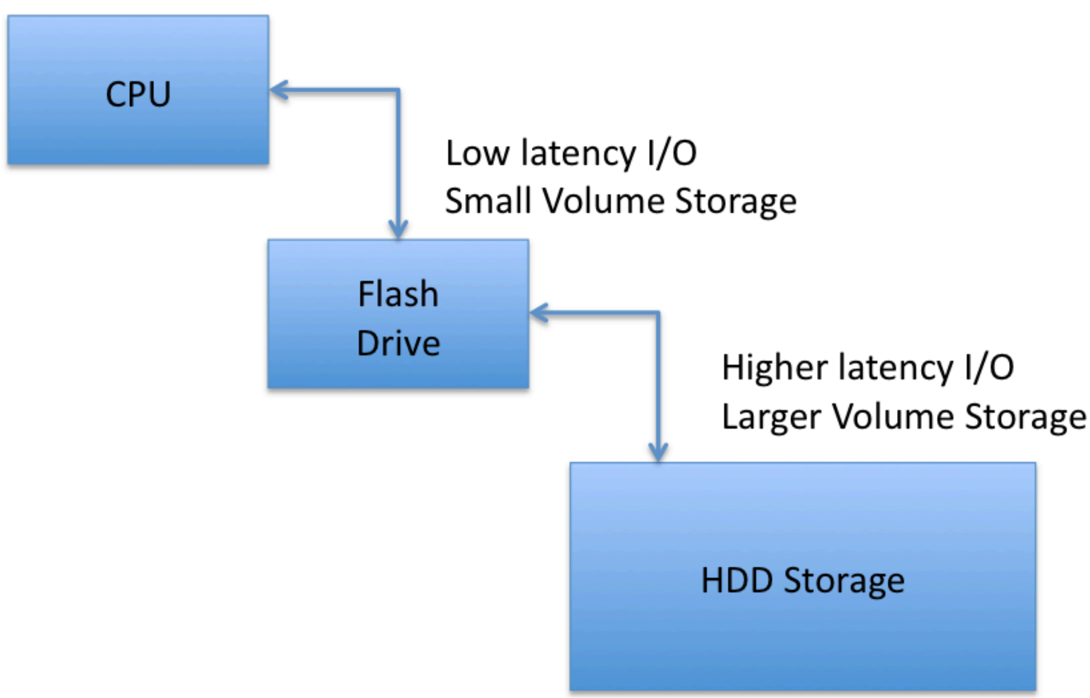

Figure 1.9: Using a combination of flash drives and hard disk drive storage allows for a balanced storage solution that keeps costs down while achieving performance objectives.

Multiple Uses of Flash Devices

Flash drives can be used through a storage infrastructure. Servers can use local flash in the form of PCIe flash drives for low-latency data transfers between the CPU and flash storage. Storage arrays can use combinations of flash devices and hard disk drives controlled by software to optimize performance while maintaining reliability and availability.

The next chapters of this guide will move from a general discussion of the role of hybrid storage to specific application areas such as database servers and virtual desktop infrastructure.