

# Complements of StorageIO

This chapter download from the book “Cloud and Virtual Data Storage Networking” (CRC Press) by noted IT industry veteran and Server StorageIO founder Greg Schulz is complements of The Server and StorageIO Group (StorageIO). Learn more about the techniques, trends, technologies and products covered in this book by visiting [storageio.com](http://storageio.com) and [storageioblog.com](http://storageioblog.com) and register for events and other promotions. Follow us on twitter @storageio or on Google+ among other social media venues.

**Brouwer** *Storage Consultancy*  
Seminar for Storage Professionals with Greg Schulz  
May 7th, 8th, 9th 2012 in Nijkerk Holland  
[Click here to learn more](#)



Visit [storageio.com/events](http://storageio.com/events) to see upcoming seminars and activities

Cloud and Virtual Data Storage Networking has been added to the Intel Recommended Reading List (IRRL) for Developers. Click on the image below to learn more about the IRRL.



The Recommended Reading List is a valuable resource for technical professionals who want to thoroughly explore topics such as software threading, wireless technologies, power management, and more. Dozens of industry technologists, corporate fellows, and engineers have helped by suggesting books and reviewing the list.

Learn more about Cloud and Virtual Data Storage Networking (CRC Press) by visiting [storageio.com/books](http://storageio.com/books)

To become a chapter download sponsor or to discuss your other project opportunities contact us at [info@storageio.com](mailto:info@storageio.com).

#1 VM Backup

Veeam Backup & Replication for

**Hyper-V**

is HERE!

[Learn More](#)

**VEEAM**

StorageIOblog.com site sponsor

# Chapter 2

---

# Cloud, Virtualization, and Data Storage Networking Fundamentals

---

*The more space you have, the more it seems to get filled up.*

– Greg Schulz

## In This Chapter

- Storage (hardware, software, and management tools)
- Block, file, direct attached, networked, and cloud storage
- Input/output, networking, and related convergence topics
- Public and private cloud products and services
- Virtualization (applications, desktop, server, storage, and networking)

This chapter provides a primer and overview of major IT resource components as well as how information is supported by them. Key themes and buzzwords discussed include block, file, object storage, and data sharing, along with public and private cloud products and services. Additional themes and buzzards include file systems, objects, Direct Attached Storage (DAS), Network Attached Storage (NAS), storage area networks (SANs), and virtualization.

Chapter 1 provided the big picture of why there is a need for clouds, storage, and networking, along with virtualization, to address business and IT challenges. For those already familiar with storage, input/output (I/O) networking, virtualization, and cloud fundamentals, feel free to skip over or just skim this chapter.

## 2.1. Getting Started

Data storage is taken for granted by many people while not being really understood, particularly when there is no more room to save files or photos. Then storage becomes frustrating, if not a nightmare, when you cannot find the file or document that you need. Even worse is after a disaster (fire, flood, hurricane, virus or data corruption, theft or accidental deletion), when you realize what should have been preserved was not adequately protected. Cost is also a concern, whether it is how much you have to pay to park your files, videos, or other data somewhere, or to have it backed up or to purchase more disk storage devices.

As was pointed out in Chapter 1, there is no such thing as a data or information recession. As a society we have become addicted to information-related services at home, at work, as well as when in transit. If you do not believe me, try this simple test: See how long you can go without checking your email, texting, using your cell phone or PDA, looking at a website, listening to satellite radio or watching TV (regular, HD, or IP), accessing your bank account, or shopping on-line. Even this book depended on lots of data storage resources for backups, copies, revisions, artwork, and other items.

Many resources are needed to support information services, including applications and management software. Also necessary are I/O and networking for connectivity between servers and data storage, infrastructure resource management (IRM) tasks, processes, procedures, and best practices. These items apply whether your information services are being deployed or accessed from a traditional IT, virtualized, or kept in a cloud environment.

## 2.2. Server and Storage I/O Fundamentals

Servers, also known as computers, are important in a discussion about cloud, virtualization, and data storage networking in that they have multiple functions. The most common—and obvious—function is that servers run the applications or programs that deliver information services. These programs are also responsible for generating I/O data and networking activity. Another role that servers play in cloud and virtualized data centers is that of functioning as data or storage appliances performing tasks that in some cases were previously done by purpose-built storage systems.

Servers vary in physical size, cost, performance, availability, capacity, and energy consumption, and they have specific features for different target markets or applications. Packaging also varies across different types of servers, ranging from small hand-held portable digital assistants (PDAs) to large-frame or full cabinet-sized mainframe

servers. Another form of server packaging is a virtual server, where a hypervisor, such as Microsoft Hyper-V, VMware vSphere, or Citrix Xen, among others, is used to create virtual machines (VMs) from physical machines (PMs). Cloud-based compute or server resources can also leverage or require a VM.

Computers or servers are targeted for different markets, including small office/home office (SOHO), small/medium business (SMB), small/medium enterprise (SME), and ultra-large-scale or extreme scaling, including high-performance computing (HPC). Servers are also positioned for different price bands and deployment scenarios.

General categories of servers and computers include:

- Laptops, desktops, and workstations
- Small floor-standing towers or rack-mounted 1U and 2U servers
- Medium-sized floor-standing towers or larger rack-mounted servers
- Blade centers and blade systems
- Large-sized floor-standing servers, including mainframes
- Specialized fault-tolerant, rugged, and embedded processing or real-time servers
- Physical and virtual along with cloud-based servers

Servers have different names—email server, database server, application server, Web server, video or file server, network server, security server, backup server, or storage server, depending on their use. In the examples just given, what defines the type of server is the type of software being used to deliver a type of service. This can lead to confusion when looking at servers, because a server may be able to support different types of workloads, thus it should be considered a server, storage, part of a network, or an application platform. Sometimes the term “appliance” will be used for a server; this is indicative of the type of service the combined hardware and software solution are providing.

Although technically not a type of server, some manufacturers use the term “tin-wrapped” software in an attempt to not be classified as an appliance, server, or hardware vendor but still wanting their software to be positioned more as a turnkey solution. The idea is to avoid being perceived as a software-only solution that requires integration with hardware. These systems usually use off-the-shelf, commercially available general-purpose servers with the vendor’s software technology preintegrated and installed, ready for use. Thus, tin-wrapped software is a turnkey software solution with some “tin,” or hardware, wrapped around it.

A variation of the tin-wrapped software model is the software-wrapped appliance or a virtual appliance. Under this model, vendors use a virtual machine to host their software on the same physical server or appliance that is being used for other functions. For example, a database vendor or virtual tape library software vendor may install its solution into separate VMs on a physical server, with applications running in other VMs or partitions. This approach works in terms of consolidating underutilized servers, but caution should be exercised to avoid overconsolidating and oversubscribing available physical hardware resources, particularly for time-sensitive applications. Keep in mind that cloud, virtual, and tin-wrapped servers or software still need physical compute, memory, I/O, networking, and storage resources.

### 2.2.1. Server and I/O Architectures

Generally speaking, servers (see Figure 2.1), regardless of specific vendor implementation, have a common architecture. That architecture consists of a central processing unit (CPU) or processor, memory, internal busses or communication chips, and I/O ports for communicating with the outside world via networks or storage devices. Computers need to perform I/O to various devices, and at the heart of many I/O and networking connectivity solutions is the Peripheral Component Interconnect (PCI) industry-standard interface.

PCI is a standard that specifies the chipsets that are used to communicate between CPUs and memory with the outside world of I/O and networking device peripherals. Figure 2.1 shows an example of a PCI implementation including various components such as bridges, adapter slots, and adapter types. PCIe leverages multiple serial unidirectional point-to-point links, known as lanes, compared to traditional PCI, which uses a parallel bus-based design.

The most current version of PCI, as defined by the PCI Special Interest Group (PCISIG), is PCI Express (PCIe). Backwards compatibility exists by bridging previous generations, including PCIx and PCI, off of a native PCIe bus or, in the past, bridging a PCIe bus to a PCIx native implementation. Examples of PCI, PCIx, and PCIe adapters include Ethernet, Fibre Channel, Fibre Channel over Ethernet (FCoE), InfiniBand Architecture (IBA), SAS, SATA, Universal Serial Bus (USB), and 1394 Firewire. There are also many specialized devices such as analog-to-digital data acquisition, video surveillance, medical monitoring, and other data acquisition or metering (e.g., data collection) devices.

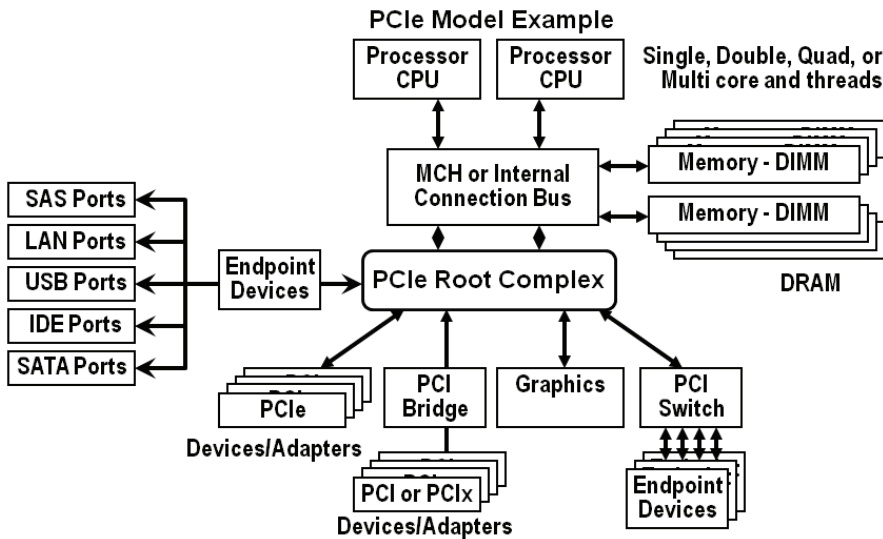
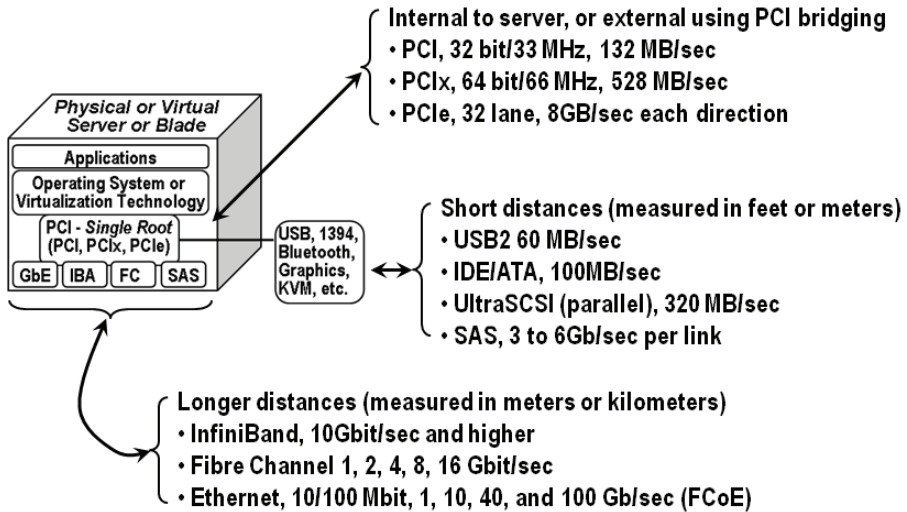


Figure 2.1 Generic computer or server hardware architecture. (Source: Greg Schulz, *The Green and Virtual Data Center*, CRC Press, Boca Raton, FL, 2009.)



**Figure 2.2** General computer and I/O connectivity model. (Source: Greg Schulz, *The Green and Virtual Data Center*, CRC Press, Boca Raton, FL, 2009.)

While the specific components and the number of components will vary depending on the server, in general, servers have one of more of the following:

- Compute or CPU chips or sockets
- One or more cores per CPU socket or chips capable of single or multithreading
- Internal communication and I/O buses for connecting components
- Some main memory, commonly dynamic random access memory (DRAM)
- Optional sockets for expansion memory, extra CPUs, and I/O expansion slots
- Attachment for keyboards, video, and monitors (KVM)
- I/O connectivity for attachment of peripherals including networks and storage
- I/O networking connectivity ports and expansion slots such as PCIe
- Optional internal disk storage and expansion slots for external storage
- Power supplies and cooling fans

Figure 2.2 shows a generic computer and I/O connectivity model, which will vary depending on specific vendor packaging and market focus. For example, some computers have more and faster processors (CPUs) and cores along with larger amounts of main memory as well as extensive connectivity or expansion options. Other computers or servers are physically smaller, lower in price and with fewer resources (CPU, memory, and I/O expansion capabilities), targeted at different needs.

In Figure 2.2, the component closest to the main processor has the fastest I/O connectivity; however, it will also be the most expensive, distance limited, and require special components. Moving farther away from the main processor, I/O still remains

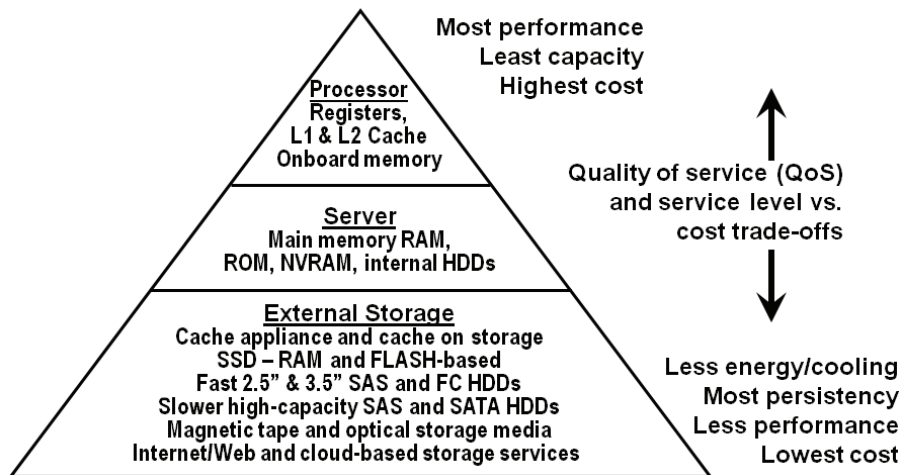
fast with distance—measured in feet or meters instead of inches—but is more flexible and cost-effective.

In general, the faster a processor or server, the more prone to a performance impact it will be when having to wait for slower I/O operations. Fast servers need lower latency and better-performing I/O connectivity and networks. Better performing means lower latency, more input/output operations per second (IOPS), as well as improved bandwidth to meet various application profiles and types of operations.

### 2.2.2. Storage Hierarchy

The storage hierarchy extends from memory inside servers out to external shared storage, including virtual and cloud accessible resources. Too often, discussions separate or even distance the relationship between server memory and data storage—after all, one is considered to be a server topic and the other a disk discussion. However, the two are very much interrelated and thus benefit as well as impact each other. Servers need I/O networking to communicate with other servers, with users of information services, and with local, remote, or cloud storage resources.

In Figure 2.3, an example of the storage or memory hierarchy is shown, ranging from fast processor core or L1 (level 1) and L2 (level 2) on-board processor memory to slow, low-cost, high-capacity removable storage. At the top of the pyramid is the fastest, lowest-latency, most expensive memory or storage, which is also less able to be shared without overhead with other processors or servers. At the bottom of the pyramid is the lowest-cost storage, with the highest capacity while being portable and sharable.



**Figure 2.3** Memory and storage pyramid. (Source: Greg Schulz, *The Green and Virtual Data Center*, CRC Press, Boca Raton, FL, 2009.)

The importance of main or processor (server) memory and external storage is that virtual machines need memory to exist when active and a place on disk to reside when not in memory. Keep in mind that a virtual machine is a computer whose components are emulated via data structures stored and accessed via memory. The more VMs there are, the more memory is required—and not just more, but faster memory is also important.

Another demand driver for increased memory capabilities on servers are applications such as database, video rendering, and business analytics, modeling, and simulations, for which large amounts of data are kept in memory for speed of access. As a result, more memory is being installed into denser footprints in servers along with more sockets for processors. An example is the IBM Power7-based Power 750, with 512 GB of DRAM and 32 cores (4 sockets, each with 8 cores).

Why are there all these different types of storage? The answer, with all technology points set aside, comes down to economics. There is a price to pay for performance. For some applications this can be considered the cost of doing business, or even a business enabler, to buy time when time is money and using a low-cost storage could have a corresponding impact on performance.

Storage can be tiered, with the applicable memory or storage technology applied to the task at hand. At a lower cost (and slower) than RAM-based memory, disk storage, along with NVRAM and FLASH-based memory devices, are also persistent. As noted in Figure 2.3, cloud is listed as a tier of storage to be used as a tool to complement other technologies. In other words, it is important from both cost and performance perspectives to use the right tool for the task at hand to enable smarter, more intelligent, and effective information services delivery.

### **2.2.3. From Bits to Bytes**

Storage is an extension of memory, and cache is a convergence of memory and external media. Digital data at a low level is stored as 1's and 0's, binary bits indicating on or off, implemented using different physical techniques depending on the physical media (disk, tape, optical, solid-state memory). The bits are generally grouped into bytes (1 byte = 8 bits) and subsequently organized into larger groups for different purposes.

The gibibyte is an international system of units (SI) measure (see Table 2.1) used for data storage and networking in base 2 format, where 1 GiBi  $2^{30}$  bytes or 1,073,741,824 bytes. Another common unit of measure used for data storage and servers as well as their operating systems is the gigabyte (GB) in base 10 (decimal) format. A GB, also sometimes shown as a GByte, represents  $10^9$  or 1,000,000,000 bytes.

Computer memory is typically represented in base 2, with disk storage often being shown in both base 2 and base 10. For example, the 7200-RPM Seagate Momentus XT Hybrid Hard Disk Drive (HHDD) that I used in my laptop for writing this book is advertised as 500 GB. The HHDD is a traditional 2.5-in. hard disk drive (HDD) that also has an integrated 4-GB flash solid-state device (SSD) and 32 MB of DRAM. Before any operating system, RAID (redundant array of independent disks), or other formatting and overhead, the HHDD presents 500,107,862,016 bytes based on 976,773,168

(512-byte) sectors or 500 GB. However, a common question is what happened to the missing 36,763,049,984 bytes of storage capacity (e.g.,  $500 \times 2^{30}$  base 2).

The Seagate Momentus XT (ST95005620AS) 500-GB HHDD guarantees 976,773,168 (512-byte) sectors. Seagate is using the standard of 1 GB = 1 billion bytes (See Table 2.1). Note that accessible storage capacity can vary depending on operating environment and upper-level formatting (e.g., operating system, RAID, controllers, snapshot, or other overhead).

The common mistake or assumption is that a 500-GB disk drive has 1,073,741,824 ( $2^{30}$ )  $\times$  500 or 536,870,912,000 bytes of accessible capacity (before overhead). However, in the example above, the disk drive presents only 500,107,862,016 bytes, leaving some to wonder where the other 36,763,049,984 bytes went to. The answer is that they did not go anywhere because they were never there, depending on what numbering base you were using or assuming. This has led to most vendors including in their packaging along with documentation what the actually accessible capacity is before environment overhead is subtracted. As an example, in a Windows environment, after formatting overhead, the empty disk shows as having a capacity of 465.74 GB.

The importance of the above is to understand that if you need a specific amount of data storage capacity, get what you expect and need. This means understanding the various ways and locations of where as well as how storage capacity is measured. Factor in overhead of controllers, RAID, spares, operating and file system, and volume managers, along with data protection such as snapshots or other reserved space, as part of storage capacity.

Table 2.1 shows data standard units of measures in both base 2 and base 10. Some of the capacities shown in Table 2.1 may seem unimaginably large today. However, keep in mind that 10 years ago a 9-GB disk drive spinning at 7200 (7.2K) revolutions per minute (RPM) was considered to be large-capacity and fast. By comparison, in late 2010, fast, energy-efficient SAS and Fibre Channel 15.5K-RPM 600-GB disk drives were shipping along with high-capacity 7.2K SAS and SATA 2-TB drives and some 3-TB consumer drives, with even larger drives soon to be appearing on the market. We will talk more about hard disk drives (HDDs), solid-state devices (SSDs), and other related technologies as well as trends in later chapters.

**Table 2.1 Storage Counting  
Numbering and Units of Measures**

		Base 2			Base 10
kibi	ki	$2^{10}$	kilo	k, K	$10^3$
mebi	Mi	$2^{20}$	mega	M	$10^6$
gibi	Gi	$2^{30}$	giga	G	$10^9$
tebi	Ti	$2^{40}$	tera	TB	$10^{12}$
pebi	Pi	$2^{50}$	peta	P	$10^{15}$
exbi	Ei	$2^{60}$	exa	E	$10^{18}$
zebi	Zi	$2^{70}$	zetta	Z	$10^{21}$
yobi	Yi	$2^{80}$	yotta	Y	$10^{24}$

### 2.2.4. Disk Storage Fundamentals

Figure 2.3 shows the basics of storage hierarchy, core or primary memory to external dedicated and shared storage. Storage can be dedicated internal Direct Attached Storage (DAS) or external shared DAS in addition to being networked and shared on a local or remote or cloud basis.

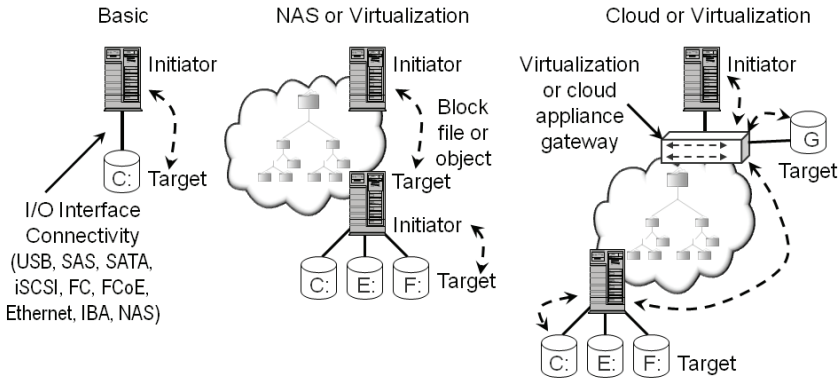


Figure 2.4 Initiator and target examples.

### 2.2.5. Initiators and Targets

A fundamental storage and I/O networking concept is that of the initiator (client or source) and target (server or destination) as shown in Figure 2.4. There is an initiator and a target in all types of storage and access mechanisms across physical, virtual, and cloud technologies. The topologies, underlying implementations, and specific functionality will vary with vendor-specific products.

Servers or other initiators initiate I/O requests (reads, writes, and status inquiries) of targets that then respond to the requests. Initiators are either taught (configured) targets or discover them at boot or start-up. These targets can be block, file, object, or some other service-based destination providing and responding to storage I/O requests. For example, a server with an initiator identifier or address makes a request to a block storage device using the SCSI command set (e.g., SAS, iSCSI, Fibre Channel, FCoE, or SRP on InfiniBand), where the target is a SCSI logical unit (LUN).

Building on the previous example, the same server can initiate I/O activity such as a file read or write request using NFS or CIFS over TCP/IP on an Ethernet-based network to a NAS target. In the case of a file request, the initiator has a mount point that is used to direct read and write requests to the target which responds to the requests. Servers also initiate I/O requests to cloud targets or destinations. While the protocols can differ for block, file, and object or specific application programming interface (API), the basic functionality of initiators and target exists even if referred to by different nomenclature.

While initiators are typically servers, they can also be storage systems, appliances, or gateways that function both as a target and an initiator. A common example of a storage system acting as both a target and an initiator is local or remote replication. In this scenario, a server sends data to be written to the target device that in turn initiates a write or update operating to another target storage system.

Another example of a target also being an initiator is a virtualization appliance or a cloud point-of-presence (cPoP) access gateway. These devices that are targets then initiate data to be copied to another physical, virtual, or cloud device.

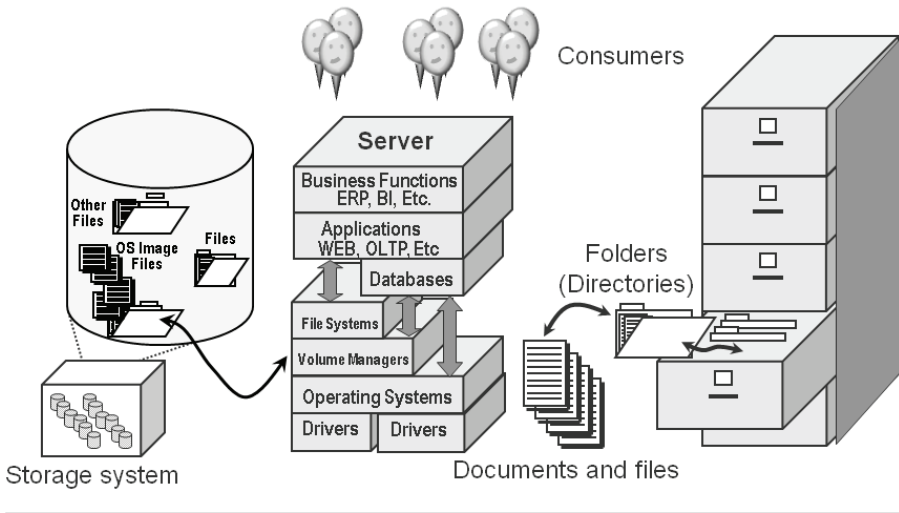


Figure 2.5 How data and information is stored.

### 2.2.6. How Data Is Written to and Read from a Storage Device

In Figure 2.5, an application creates a file and then saves it—for example, a Word document created and saved to disk. The application—Word in this example—works with the underlying operating system or file system to ensure that the data is safely written to the appropriate location on the specific storage system or disk drive. The operating system or file system is responsible for working with applications to maintain directories or folders where files are stored.

The operating system, file system, or database shown in Figure 2.5 is responsible for mapping the file system where the folder or directories are located to a specific disk drive or LUN or a volume on a storage system. A storage system that receives data from a server via file access—as in the case of a NAS device or file server—is, in turn, responsible for mapping and tracking where blocks of data are written to on specific storage devices.

Data is accessed on the disk storage device (Figure 2.6) by a physical and a logical address, sometimes known as a physical block number (PBN) and a logical block number (LBN). The file system or an application performing direct (raw) I/O keeps track of what storage is mapped to which logical blocks on what storage volumes. Within

the storage controller and disk drive, a mapping table is maintained to associate logical blocks with physical block locations on the disk or other medium such as tape.

When data is written to disk, regardless of whether it is an object, file, Web database, or video, the lowest common denominator is a block of storage (Figure 2.6). Blocks of storage have been traditionally organized into 512 bytes, which aligned with memory page sizes. While 512-byte blocks and memory page sizes are still common, given larger-capacity disk drives as well as larger storage systems, 4-K (e.g.,  $8 \times 512$  bytes) block sizes are appearing.

Larger block sizes enable more data to be managed or kept track of in the same footprint by requiring fewer pointers or directory entries. For example, using a 4-K block size, eight times the amount of data can be kept track of by operating systems or storage controllers in the same footprint. Another benefit is that with data access patterns changing along with larger I/O operations, 4 K makes for more efficient operations than the equivalent  $8 \times 512$  byte operations for the same amount of data to be moved.

At another detailed layer, the disk drive or flash solid-state device also handles bad block vectoring or replacement transparently to the storage controller or operating system. Note that this form or level of bad block repair is independent of upper-level data protection and availability features, including RAID, backup/restore, replication, snapshots, or continuous data protection (CDP), among others.

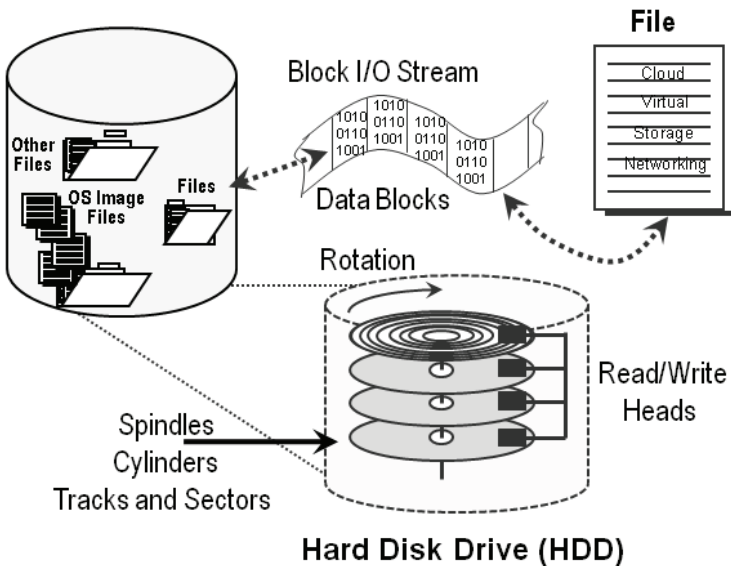


Figure 2.6 Hard disk drive storage organization.

### 2.2.7. Storage Sharing vs. Data Sharing

Storage and data sharing may sound like the same thing, and the phrases are often used interchangeably; however, they are quite different and not at all interchangeable. Sharing

*storage* means being able to have a disk drive or storage system accessible by two or more initiators or host computer servers. By being shared, only portions of the disk device or storage system is accessible to specific servers or initiators, as seen in Figure 2.7.

For example, the C: E: and F: storage devices or volumes are only accessible to the servers that own them. On the right side of Figure 2.7, each server is shown with its own dedicated storage. In the middle of Figure 2.7 a shared storage device is shown, with each server having its own LUN, volume, or partition. Data sharing is shown with different servers having access to the D: volume, where they can access various documents, objects, or VM files with applicable security authorization.

With shared storage, different servers can initiate I/O activity to the portion of storage to which they have access, which might be a partition, logical drive or volume, or LUN. For high-availability (HA) clustering, shared storage can be accessed by multiple servers running software that maintains data integrity and coherent access to the device.

Shared *data*, on the other hand, involves multiple servers being able to read or write to the same file via file serving and sharing software. File serving or sharing software is typically found in most operating systems as well as within NAS that support common protocols including Network File System (NFS) and Common Internet File System (CIFS).

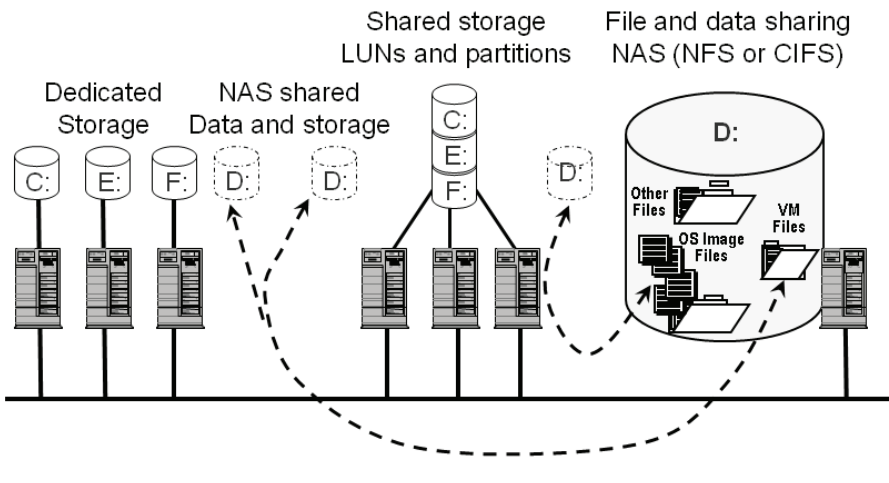


Figure 2.7 Storage and data sharing.

### 2.2.8. Different Types of Storage: Not All Data Storage Is the Same

There are many different types of storage (Figure 2.8) for different application requirements and various usage scenarios. Some storage is performance-oriented for bandwidth (throughput), measured in megabytes or gigabytes per second, or in terms of response time (latency), or as the number of I/O operations per second (IOPS). Storage can be

optimized or targeted for on-line active and primary usage, near-line for idle or inactive data, or off-line, where the focus can be high capacity at low cost.

On-line or active applications are those for which data is being worked with, read, and written, such as file systems, home directories, databases, and email. Near-line or applications with idle data include reference material or repositories, backups, and archives.

Some general categories of storage include:

- Shared or dedicated, internal or external to a server/computer
- Local, remote or cloud, block, file, or object
- On-line active or high-performance primary
- Inactive or idle, near-line or off-line

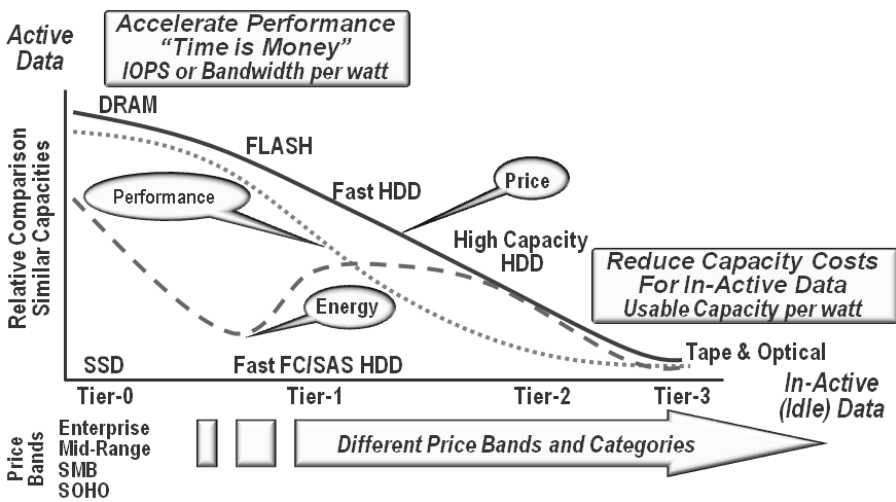


Figure 2.8 Types and tiers of storage media.

### 2.2.8.1. Structured, Unstructured, and Meta Data

Some applications store and access data in highly structured databases (e.g., databases or application-specific organization) such as IBM DB2/UDB, Intersystem’s Caché, Microsoft SQLserver, MySQL, Oracle 11g, and SAP Sybase, among others. Other applications have their own predefined stores or pools where storage is allocated either as a LUN or as a volume in the case of block access. In the case of block-based access, the application, database, or file system works with applications to know what files to access. In other scenarios, applications access a file by name in a folder or directory that is part of a mount point or share in a file system, either locally or on a remote file server or NAS device. Another means of data access is via an applications program interface (API), where a client requests information from a server via a defined mechanism.

In the case of a database, applications only need to know the schema or how the database is organized to make queries with tools including Structured Query Language

(SQL), while the database handles read or writing of data either in a block or file system mode. For block mode the database is assigned LUNs or storage space where it creates its files or datasets that it manages. If using a file system, the database leverages the underlying file system to handle some of the storage management tasks.

Structured data has defined attributes, making searching or other functions relatively easy. This structure, however, can make adding or changing the organization more complex or costly. As a result, there is the growing category of unstructured data, also known as file-accessed data. The value proposition of unstructured data is that there is no formal organization other than files stored in a folder or directory in a file system. File names can vary depending on the specific file system or operating system environment, with attributes—including creation and modification dates, ownership, and extension—that can usually indicate what application it is associated with along with security and file size.

Some file systems and files can support additional meta data (data about the data) or properties. The flexibility of unstructured data causes challenges when it comes to being able to search or determine what the files contain. For example, with a database, the schema or organization makes it relatively easy to search and determine what is stored. However, with unstructured data, additional meta data needs to be discovered via tools including eDiscovery (search) and classification tools. Additional meta data that can be discovered and stored in a meta database or repository includes information about the contents of files, along with dependencies on other information or applications. The use of structured or unstructured data depends on preferences, the performance of the specific file system or storage being used, desired flexibility, and other criteria.

In general, storage is accessed locally, remotely, or via a cloud using:

- Application Programming Interface (API)
- Block-based access of disk partitions, LUNs, or volumes
- File-based using local or networked file systems
- Object-based access

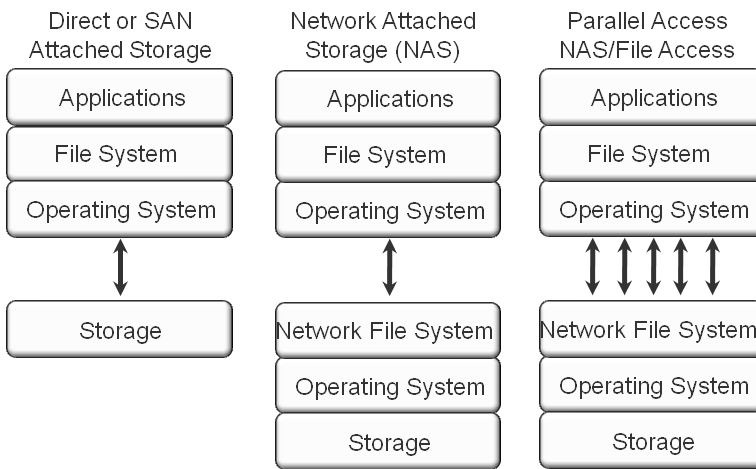
In early generations of computers, back when dinosaurs roamed the earth (maybe not quite that far back), in an era before file systems and volume managers, programmers (aka the user) had to know physically where files or data were saved and to be read. Fast forward a few decades and a few computer generations, and saving a file has become fairly transparent and relatively easy. Granted, you still need to know some form of an address or directory or folder or share where the data is located, but you do not have to worry about knowing what starting and stop locations on the disk or storage system to access.

### 2.2.8.2. Block Storage Access

Block-based data access is the lowest level of access and the fundamental building block for all storage. This means that block-based data access is relevant for cloud and virtualized storage as well as storage networks. Regardless of whether applications on

servers (initiators) request data via an API, object, file, or via a block-based request from a file system, database, email, or other means, it ultimately gets resolved and handled at the lowest layers as blocks of data read from or written to disks, solid-state disks, or tape-based storage.

For those who see or access data via a file system, database, document management system, SharePoint, email, or some other application, the block-based details have been abstracted for you. That abstraction occurs at many levels, beginning at the disk drive, the storage system, or controller, perhaps implementing RAID to which it is attached as a target, as well as via additional layers including virtualized storage, device drivers and file systems, and volume managers, databases, and applications.



**Figure 2.9** File access examples.

### 2.2.8.3. Files Access, File Systems, and Objects

File-based data access of unstructured data (Figure 2.9) is seeing rapid growth due to ease of use and flexibility for traditional environments as well as for virtual and cloud data infrastructures. File-based data access is simplified by abstracting the underlying block-based components, enabling information to be accessed via filename.

File system software provides the abstraction of file-based access on a local or remote basis. Instead of knowing where the physical data is located on a disk or storage system, meaningful file names are used, along with directory structures or folders, for organization purposes. With file sharing, the client or initiator makes requests to the filer (target or destination) to process I/O requests on a file basis. The filer of the file system software presents (serves) data from the storage device to other host servers via the network using a file-sharing protocol while maintaining data coherency.

NAS filers can have dedicated HDD or SSD-based storage, external third-party storage, or a combination of both. NAS systems that leverage or support attachment

of third-party storage from other vendors are known as gateways, NAS heads, routers, or virtual filers, among other marketing names. For example, NetApp has vFilers that support attachment of their own storage as well as systems from HDS, IBM, and many others. The advantage of using a NAS gateway such as those from Dell, BlueArc, EMC, HP, IBM, NetApp, Oracle, and others is the ability to reuse and repurpose existing storage systems for investment protection or purchasing flexibility.

NAS systems with integrated or dedicated storage across different market segments include BlueArc, Cisco, and Data robotics, Dell, EMC, Fujitsu, HP, IBM, Iomega, NetApp, Overland, Oracle, and Seagate. NAS software, in addition to propriety solutions from those mentioned and others, includes that from Microsoft (Windows Storage Server) and Oracle (ZFS), as well as others.

Some examples of data or information access protocols are shown in Table 2.2; others include HTTP (HyperText Transfer Protocol), FTP (File Transfer Protocol), WebDAV, Bit Torrent, REST, SOAP, eXtensible Access Method (XAM), and Digital Imaging and Communications in Medicine (DICOM), among others. Table 2.3 shows examples of software and solutions, including a mix of general-purpose, specialized, parallel, and clustered file systems.

Note that products shown in the Product column of Table 2.3 are software-based and may or may not be currently available as a software-only solution. Some of the vendors, particularly those who have acquired the software from its developer, have chosen to make it available only as a bundled preconfigured solution or via original equipment manufacturers (OEM) to other solution partners. Check specific vendors' websites and supported configuration for additional details or limitations. Other file systems and clustered file systems software or bundled solutions include, among others, SGI XFS and CXFS, Panasas PanFS, and Symantec/Veritas file system and clustered file system.

**Table 2.2 Common File Access Protocols**

Acronym	Protocol Name	Comment
AFP	Apple File Protocol	Apple file serving and sharing
CIFS	Common Internet File System	Microsoft Windows file serving and sharing
NFS	Network File System	File sharing for Unix, Linux, Windows and others
pNFS	Parallel NFS	Part of NFS standard supporting parallel file access suited for reading or writing large sequential files

#### 2.2.8.4. Object and API Storage Access

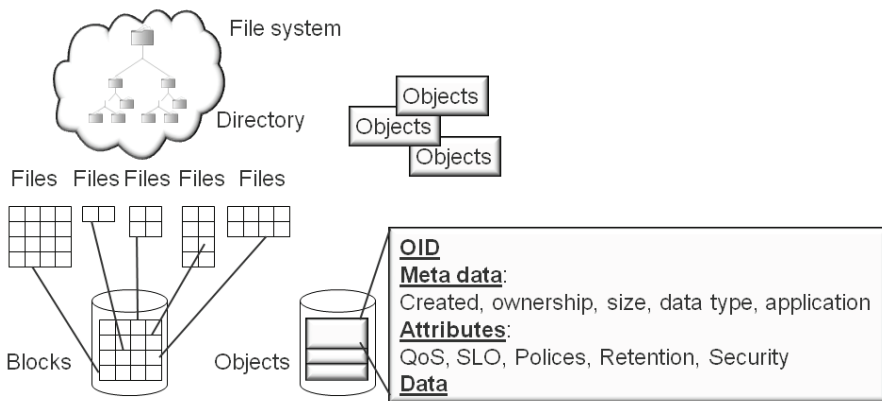
Object-based storage, or content-addressable storage (CAS), is continuous building on the block and file storage access models. As previously discussed, file systems store files in a hierarchy directory structure mapped onto an underlying block storage device. Instead

**Table 2.3 Common File System Software Stacks or Products**

Product	Vendor	Comment
Exanet	Dell	Software bundled with Dell hardware as a scale-out or bulk NAS
GPFS	IBM	Software bundled with IBM hardware (SONAS) for scale-out NAS
IBRIX	HP	Software bundled with HP hardware as scale-out or bulk NAS
Lustre	Oracle	Scale-out parallel file system software
Polyserve	HP	Software bundled with HP hardware as scale-out or bulk NAS
SAM QFS	Oracle	Software integrated with Storage Archive Manager for hierarchical storage management
SFS	Symantec	Scale-out file system software sold by Symantec and others
WSS	Microsoft	Windows Storage Servers is used in many entry-level NAS products
ZFS	Oracle	Oracle bundled solution as 7000 series storage or by original equipment manufacturers

of a storage system reading and writing blocks of data as required by a file system, object-based storage works with objects. With block- and file-based storage, applications coordinate with file systems where and how data is read and written from a storage device with the data in blocks and with little meta data attached to the stored information.

In the case of object storage, instead of as a group of blocks of data, data is stored as an object that contains meta data as well as the information being stored. The object (Figure 2.10) is defined by an application or some other entity and organized in such



**Figure 2.10** Block-and-file vs. object-based storage access.

a way that information about the data (meta data) is also attached to the data being stored, independent of the file system, database, or other organizational mechanisms.

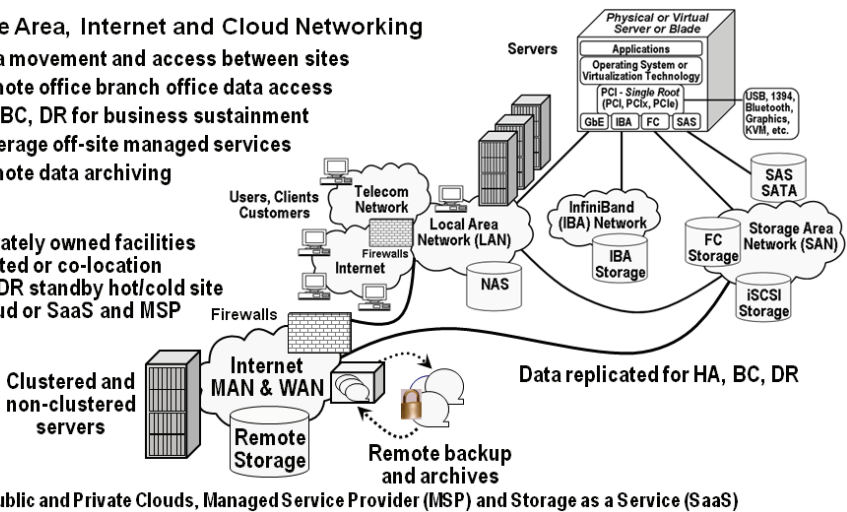
Instead of a server or file system retaining the information about the data in a file system directory and inode (directory entry), CAS stores and retrieves information using unique key identifiers that are derived from the contents of the data being stored. If the data being stored is changed, the identifier also is changed. While there are variations in how CAS is implemented, the common approach is to add a level of abstraction as well as preserve data uniqueness for compliance and other retention applications.

Vendors with CAS and object-based storage solutions include Amazon, Clever-safe, Dell, EMC, HDS, HP, IBM, NetApp, Oracle, Panasas, and Scality. An emerging International Committee for Information Technology (INCITS) and an ANSI specification, the T10 Object Storage Device (OSD), is also evolving to formalize object storage and associated technologies. It is also worth mentioning that the ANSI T10 group (www.t10.org) is responsible for the SCSI command set found in open-system block storage solutions. Another, related group is ANSI T11, which focuses on Fibre Channel matters.

**Wide Area, Internet and Cloud Networking**

- Data movement and access between sites
- Remote office branch office data access
- HA, BC, DR for business sustainment
- Leverage off-site managed services
- Remote data archiving

- Privately owned facilities
- Hosted or co-location
- BC/DR standby hot/cold site
- Cloud or SaaS and MSP



**Figure 2.11** I/O and data storage networks. (Source: Greg Schulz, *The Green and Virtual Data Center*, CRC Press, Boca Raton, FL, 2009.)

**2.3. I/O Connectivity and Networking Fundamentals**

There are many different types of I/O and networking protocols, interfaces, and transport media, as shown in Figure 2.11. While networks and storage I/O interfaces support different aspects of computing, they both support moving information between computing and I/O interfaces. Over time, storage I/O interfaces have become specialized

to support the needs and characteristics of moving data between servers and storage as well as between storage devices.

Local area networks (LANs) and wide area networks (WANs) are used for:

- Accessing and moving data to or from public/private clouds
- Data movement, staging, sharing, and distribution
- Storage access and file or data sharing (NAS)
- High-availability clustering and workload balancing
- Backup/restore for business continuance and disaster recovery
- Web and other client access, including PDAs, terminals, etc.
- Voice and video applications including Voice-over-IP (VoIP)

The term *networked storage* is often assumed to mean network attached storage (NAS) as opposed to a storage area network (SAN). In a general context, networked storage can mean storage that is accessed via some form of I/O network.

SAN and NAS are both part of storage networking. SAN is associated with Fibre Channel block-based access and NAS with LAN NFS or CIFS (SMB) file-based access. Each has its place to address different business needs. A SAN can also provide high-speed backup of NAS filers using Fibre Channel to access shared tape devices. Similar to the benefits of host servers, NAS filers also benefit from storage and backup sharing for resiliency. With a SAN, unless concurrent access software such as HACMP, Quantum Stornext, or some other clustered shared access tool is used, the LUN, device, or volume is owned by a single operating system or virtual machine guest at a time. In the case of NAS, data can be accessed by multiple servers, as the owner of the data, which provides concurrency as well as integrity, is the file server.

Some storage networking benefits include:

- Physically removing storage from servers
- Improved server resiliency and clusters
- Diskless servers using shared resources
- Storage and data sharing and consolidation
- Improved backup and recovery resource sharing
- Improved distance, capacity, and performance
- Simplified management via consolidated resources
- Lower total cost of ownership (TCO) from resource sharing

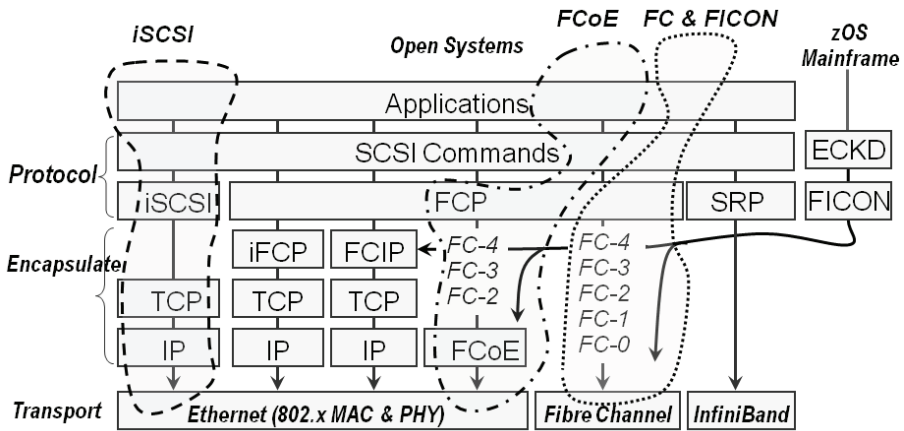
A storage network can be as simple as a point-to-point connection between one or more servers attached to and sharing one or more storage devices including disk and tape. A storage network can also be as complex as multiple subnets (segments or regions) spanning local, metropolitan, and global sites and using multiple topologies and technologies.

There are many different types of servers optimized for various applications, performance, capacity, and price points (e.g., tiered servers), with various tiers of storage. Similar to storage tiers, there also different tiers of I/O connectivity and networking (Figure 2.11). Storage and I/O interconnects (Figure 2.11) have also evolved from various vendors'

proprietary interfaces and protocols to industry-standard Fibre Channel, InfiniBand, Serial Attached SCSI (SAS), and Serial ATA (SATA) as well as Ethernet-based storage.

With the exception of IBM legacy mainframes that utilize count key data (CKD) or extended count key data (ECKD) protocols, open systems-based computers, networking, and storage devices have standardized on the SCSI command set for block-based I/O. Physical traditional parallel SCSI cabling has given way to serial-based connectivity for block storage access. For example, SAS, iSCSI (TCP/IP on Ethernet), Fibre Channel, and SRP (InfiniBand) all rely on the SCSI command set mapped onto different transports.

Looking at Figure 2.12, one might ask why so many different networks and transports are needed: Why not just move everything to Ethernet and TCP/IP? The simple answer is that the different interfaces and transports are used to meet different needs, enabling the most applicable tool or technology to be used for the task at hand. An ongoing trend, however, has been toward convergence of protocols, transports, and networking cabling, as shown in Figure 2.12. In fact, the number of protocols has essentially converged to one for open-systems block I/O using the SCSI command set [SAS, Fibre Channel, and Fibre Channel over Ethernet (FCoE), iSCSI or SRP on InfiniBand (native SCSI on InfiniBand vs. iSCSI on IP on InfiniBand)].



**Figure 2.12** Data center I/O protocols, interfaces, and transports. (Source: Greg Schulz, *The Green and Virtual Data Center*, CRC Press, Boca Raton, FL, 2009.)

## 2.4. IT Clouds

There are many different types of clouds (public, private, and hybrid), with different types of functionalities and service personalities (e.g., storage of objects, backup, archive, generator file storage space, application-specific using various APIs or interfaces). There are cloud services to which you move your applications or whose software you use as a service, those to which you move your data and use cloud applications that

reside in the same or a different location to access that data, and those in which your data is accessed via a cloud gateway, router, cloud point of presence (cpop), software, or other agent. Then there are products that are located at your site that enable cloudlike operations or management, and some that can be considered private clouds.

Many cloud definitions are based on a particular product or product focus area where more than one vendor aligns around common themes. For example, several vendors use the public cloud services model in which you can have your own private or shared space using a specific API such as REST, DICOM, SOAP, or others. Some solutions or products are designed for building services that, in turn, can be sold or provided to others. There is no one right cloud approach; rather, there are various approaches to align with your specific needs and preferences.

Cloud functionality varies by public, private, or hybrid cloud as well as for a fee or free, depending on the specific model or how it is deployed. Products can be turnkey off the shelf, custom, or a combination, including hardware, software, and services. Functionality can vary based on cost, service-level agreement, or type of service or product. Some services are based on shared infrastructure, whereas others are dedicated or isolated as well as having the ability to specify in what geographic region the data is parked for regulatory compliance, encryption for security, import/export capabilities for large volumes of data, along with audit and management tools.

Clouds can be a product, technology, or service as well as a management paradigm. They can leverage various technologies, including storage (DAS, SAN, and NAS, along with disk, SSD, and tape), servers, file systems along with various types of networking protocols or access methods, as well as associated management tools, metrics, and best practices.

Services provided by public or private clouds include application-specific software (Salesforce.com, ADP/payroll, concur expense reporting tools), archiving and data preservation, backup and restore, business continuance/disaster recovery, business analytics and simulation, compute capabilities, database and data warehousing, document sharing (Google Docs), email, collaboration and messaging, file sharing or hosting, along with object storage, office functions (word processing, spreadsheet, calendaring), photo, video, audio storage, presentations, slide content sharing, SharePoint and document management, video surveillance and security, and virtual machines.

From the xSP nomenclature, where “x” is replaced with various letters representing different themes—such as “I” for Internet or Infrastructure—“x as a Service” is now being used in a similar manner, as “xaaS,” for example, Archive as a Service (AaaS), Application as a Service (AaaS), Backup as a Service (BaaS), Desktop as a Service (DaaS), Disk as a Service (DaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS), Software as a Service (SaaS), and Storage as a Service (SaaS), among many others.

Some consider the only XaaS categories to be AaaS, PaaS, and IaaS, and that everything else must be included under those umbrellas. The reason for this is that those are the models that line up with the product or service they are selling or supporting, so it makes sense to keep conversations focused around those themes. In those instances it makes sense, but, realistically, there are other categories for the more broadly focused aspects of public and private cloud products, services, and management paradigms.

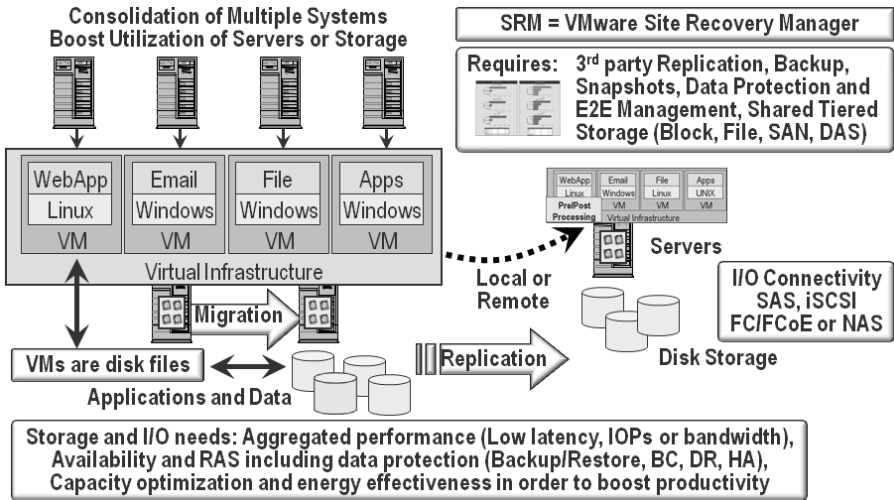


Figure 2.13 Various forms of virtualization, storage, and I/O.

## 2.5. Virtualization: Servers, Storage, and Networking

There are many facets of virtualization (Figure 2.13). Aggregation has become well known and a popular approach to consolidate underutilized IT resources including servers, storage, and networks. The benefits of consolidation include improved efficiency by eliminating underutilized servers or storage to free up electrical power, cooling requirements, floor space, and management activity, or to reuse and repurpose servers that have been made surplus to enable growth or support new application service capabilities.

Figure 2.13 shows two examples of virtualization being used, with consolidation on the left side and transparency for emulation and abstraction to support scaling on the right. On the consolidation side, the operating systems and applications of multiple underutilized physical servers are shown being consolidated onto a single or, for redundancy, multiple servers in a virtual environment with a separate virtual machine emulating a physical machine. In this example, each of the operating systems and applications that were previously running on their own dedicated server now run on a virtual server to boost utilization and reduce the number of physical servers needed.

For applications and data that do not lend themselves to consolidation, a different form of virtualization is to enable transparency of physical resources to support interoperability and coexistence between new and existing software tools, servers, storage, and networking technologies, for example, enabling new, more energy-efficient servers or storage with improved performance to coexist with existing resources and applications.

Another facet of virtualization transparency is to enable new technologies to be moved into and out of running or active production environments to facilitate technology upgrades and replacements. Virtualization can also be used to adjust physical resources to changing application demands such as seasonal planned or unplanned

workload increases. Transparency via virtualization also enables routine planned and unplanned maintenance functions to be performed on IT resources without disrupting applications and users of IT services.

## 2.6. Virtualization and Storage Services

Various storage virtualization services are implemented in different locations to support various tasks. In Figure 2.14 are shown examples of pooling or aggregation for both block- and file-based storage, virtual tape libraries for coexistence and interoperability with existing IT hardware and software resources, global or virtual file systems, transparent data migration of data for technology upgrades and maintenance, and support for high availability (HA) and business continuance/disaster recovery (BC/DR).

One of the most commonly talked about forms of storage virtualization is aggregation and pooling solutions. Aggregation and pooling for consolidation of LUNs, file systems, and volume pooling, and associated management, are intended to increase capacity utilization and investment protection, including supporting heterogeneous data management across different tiers, categories, and price bands of storage from various vendors. Given the focus on consolidation of storage and other IT resources along with continued technology maturity, more aggregation and pooling solutions can be expected to be deployed as storage virtualization matures.

While aggregation and pooling are growing in popularity in terms of deployment, most current storage virtualization solutions are forms of abstraction. Abstraction and technology transparency include device emulation, interoperability, coexistence, backward compatibility, transition to new technology with transparent data movement and migration, support for HA and BC/DR, data replication or mirroring (local and remote), snapshots, backup, and data archiving.

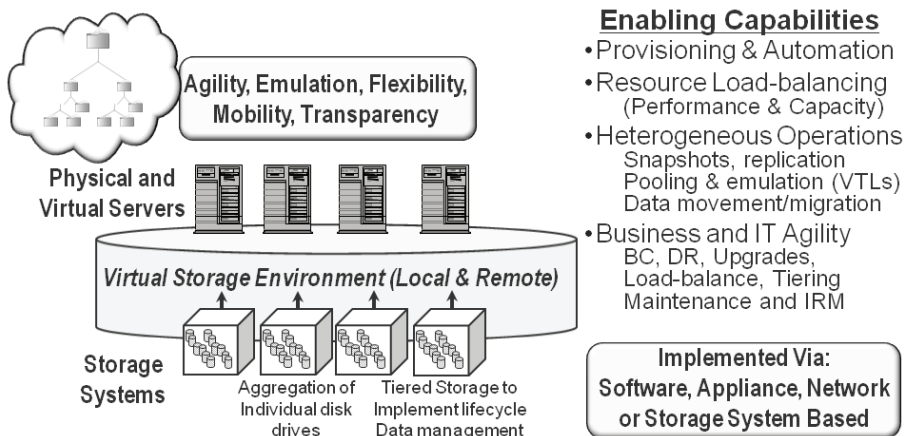


Figure 2.14 The many forms of storage virtualization.

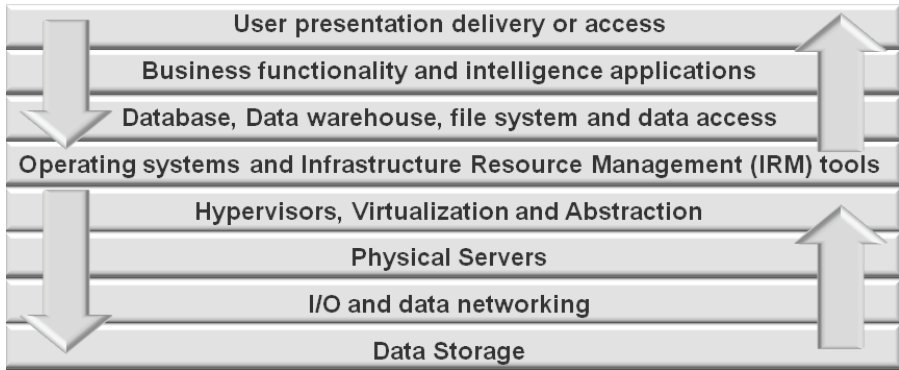


Figure 2.15 Data infrastructure stack and relationships.

## 2.7. Data and Storage Access

Figure 2.15 brings things together in terms of the topics covered in this chapter. In general, different layers and protocols interact to support information services delivery. Looking at Figure 2.15, storage is at the bottom and the tenants of data infrastructure components are farther up the stack. Data infrastructures needed to support information factories involve more than just disk or storage devices on a local or remote basis. The delivery of information services also relies on infrastructure resource management (IRM) tools, file systems, databases, virtualization, cloud technologies, and other tools and technologies.

The items or layers in Figure 2.15 can be consolidated or increased for additional detail, including middleware and other component items. To help put things into perspective, let’s take a moment to look at where servers and storage have been, currently are at, as well as where they are going. Figure 2.16 shows from left to right how servers

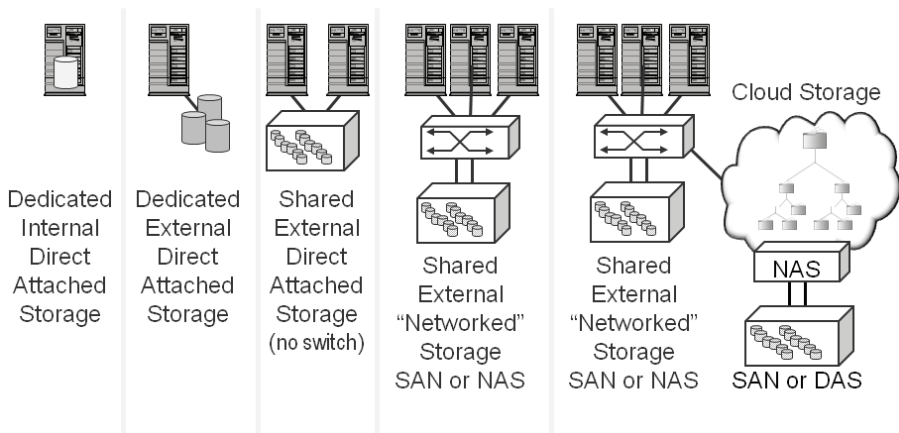


Figure 2.16 The server and storage I/O continuum.

and storage have evolved from being closely coupled as well as propriety to unbundled, open, and interoperable. Also shown is the evolution from dedicated to shared Direct Attached Storage (DAS) to networked along with local and remote, physical, virtual, and cloud.

### **2.7.1. Direct Attached Storage (DAS)**

Figure 2.16 shows different storage access scenarios, including dedicated internal DAS, dedicated external DAS, shared external DAS, shared external networked (SAN or NAS) storage, and cloud-accessible storage. DAS is also called point-to-point storage, in that a server attaches directly to storage systems adapter ports using iSCSI, Fibre Channel, or SAS without a switch. It is important to keep in mind that DAS does not have to mean dedicated internal storage; it can also mean external shared direct accessible storage using SAS, iSCSI, or Fibre Channel.

While a general industry trend is toward increased use of networked storage, that is, either block SAN using iSCSI, Fibre Channel, or FCoE as well as NAS using NFS and CIFS, there is still a strong if not growing use of DAS for deploying virtualization and clouds.

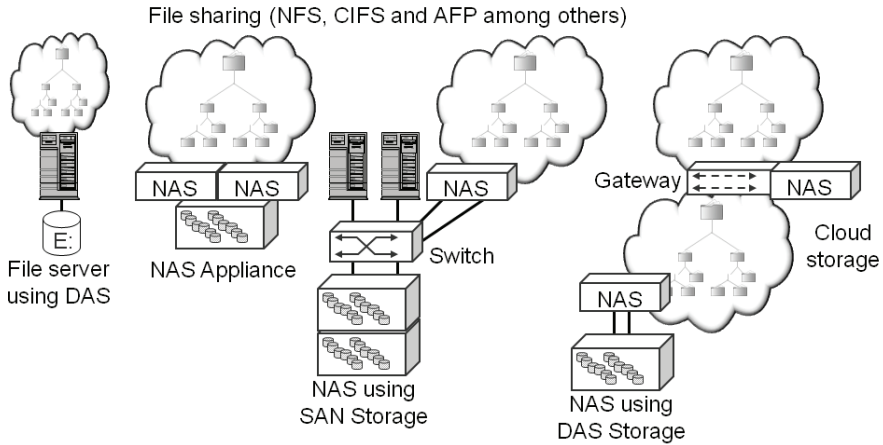
For example, storage and data services, including backup/restore, data protection, and archive solutions may present their functionality to clients or initiators (e.g., servers) via SAN block or NAS; however, the back-end storage may in fact be DAS-based. Another example is a cloud-based storage solution that presents iSCSI LUNs or virtual tape, HTTP, FTP, WebDav, or NAS NFS or CIFS access to clients, while the underlying storage may be DAS SAS (Figure 2.16).

### **2.7.2. Networked Storage: Network Attached Storage (NAS)**

Four examples of NAS are shown in Figure 2.17. Shown first, on the left, is a server sharing internal or external storage using NFS, AFP, or Windows CIFS, among other software. The next example is a high-availability NAS appliance that supports various file- and data-sharing protocols such as NFS and CIFS along with integrated storage. In the middle of Figure 2.17 is shown a SAN with a NAS appliance without integrated storage that accesses shared storage. At the far right of Figure 2.17 is a hybrid showing a NAS system that also has access to a cloud point of presence (cPOP), gateway, or appliance for accessing cloud data. The cloud-based storage in the lower right-hand corner of Figure 2.17 is a NAS server with external DAS storage such as a shared SAS RAID storage system.

### **2.7.3. Networked Storage: Storage Area Network (SAN)**

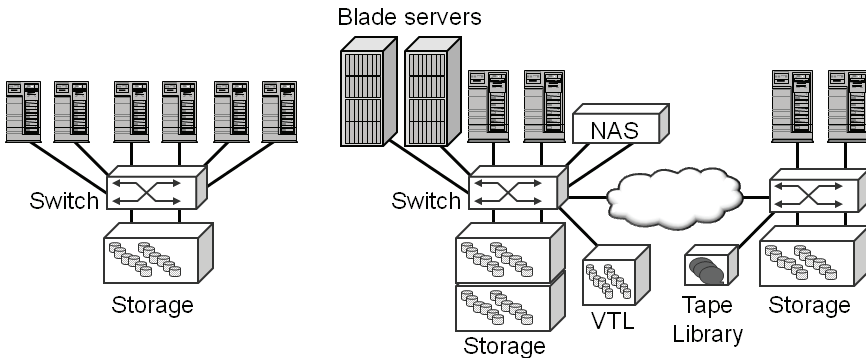
Storage area network (SAN) examples are shown in Figure 2.18, with a small or simple configuration on the left and a more complex variation on the right. On the left side of



**Figure 2.17** Network Attached Storage (NAS) examples.

Figure 2.18, multiple servers attach to a SAN switch, which in turn attaches to one or more storage systems. Not shown would be a high-availability configuration in which a pair of switches is used to connect servers and storage via redundant paths.

The SAN interface or protocols that can be used include shared SAS with a switch, iSCSI using Ethernet, Fibre Channel, and Fibre Channel over Ethernet or InfiniBand. The example on the right side of Figure 2.18 shows multiple blade servers along with traditional servers and a NAS gateway appliance attached to a SAN switch or director (large switch). Also attached to the switch are multiple storage systems that support deduplication, virtual tape libraries (VTLs), NAS, and other functionalities for backup/restore along with archiving. The example on the right in Figure 2.18 could be a pair of switches in a single location, across a campus, or across a metropolitan or wide area basis.



**Figure 2.18** Storage area network (SAN) example.

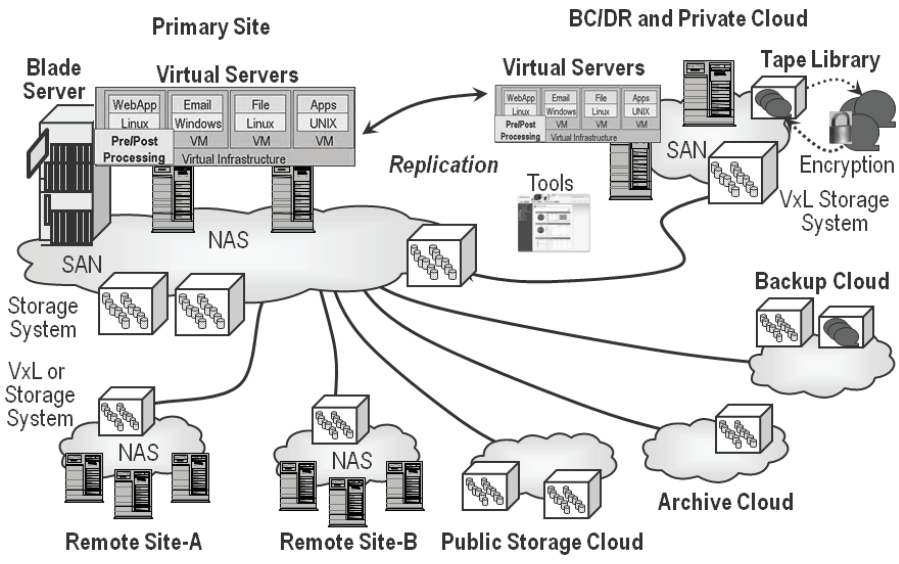


Figure 2.19 Cloud storage example.

### 2.7.4. Networked Storage: Public and Private Clouds

Figure 2.19 shows various cloud products and services supporting public and private capabilities combined with DAS, SAN, and NAS storage access for local as well as remote locations. Also shown are physical machines (PM) and virtual machine (VM) servers, primary on-line storage, and technologies for data protection. To support data protection, replication between locations along with backup/restore using disk to disk (D2D) as well as tape for long-term archiving or data preservation is also provided.

## 2.8. Common Questions

*Do virtual servers need virtual storage?* While virtual servers can benefit from features found in many virtual storage systems, generally speaking, virtual servers do not need or require virtual storage. However, virtual servers do need access to shared storage such as external SAS, iSCSI, Fibre Channel, FCoE, or NAS.

*Do clouds need cloud storage?* The answer depends on your definition of a cloud. For example, private clouds can use traditional storage products combined with additional management tools and best practices. Private clouds can also leverage cloud-specific solutions as well as external, third-party-provided public clouds. Public clouds often include storage as part of their solution offering or partner with cloud storage services providers.

*Doesn't DAS mean dedicated internal storage?* DAS is often mistaken for meaning dedicated internal attached storage inside a server. While this is true, DAS can also

refer to external shared storage that is directly attached to one or more servers without using switches over interfaces including SAS, iSCSI, and Fibre Channel. Another name for external shared DAS is point-to-point storage, such as where iSCSI storage is connected directly to a server via its Ethernet ports without a switch.

## 2.9. Chapter Summary

Not all data and information is the same in terms of frequency of access and retention, yet typically it is all treated the same. While the cost per unit of storage is decreasing, the amount of storage that can be managed per person is not scaling at a proportional rate, resulting in a storage management efficiency gap. Information data can be stored in different formats using various interfaces to access it.

General action items include:

- Fast servers need fast storage and networks.
- You cannot have software without hardware, and hardware needs software.
- Clouds and virtual environments continue to rely on physical resources.
- There are many types of storage and access for different purposes.

Unless you are in a hurry to keep reading to get through this book, now is a good time to take a break, relax, and think about what we have covered so far, before moving on to the next section. Also visit my website at [www.storageio.com](http://www.storageio.com) and my blog at [www.storageioblog.com](http://www.storageioblog.com), where you will find additional details, discussions, and related content to the material discussed here and in subsequent chapters.

# Complements of StorageIO

This chapter download from the book “Cloud and Virtual Data Storage Networking” (CRC Press) by noted IT industry veteran and Server StorageIO founder Greg Schulz is complements of The Server and StorageIO Group (StorageIO). Learn more about the techniques, trends, technologies and products covered in this book by visiting [storageio.com](http://storageio.com) and [storageioblog.com](http://storageioblog.com) and register for events and other promotions. Follow us on twitter @storageio or on Google+ among other social media venues.

**Brouwer** *Storage Consultancy*  
Seminar for Storage Professionals with Greg Schulz  
May 7th, 8th, 9th 2012 in Nijkerk Holland  
[Click here to learn more](#)



Visit [storageio.com/events](http://storageio.com/events) to see upcoming seminars and activities

Cloud and Virtual Data Storage Networking has been added to the Intel Recommended Reading List (IRRL) for Developers. Click on the image below to learn more about the IRRL.



The Recommended Reading List is a valuable resource for technical professionals who want to thoroughly explore topics such as software threading, wireless technologies, power management, and more. Dozens of industry technologists, corporate fellows, and engineers have helped by suggesting books and reviewing the list.

Learn more about Cloud and Virtual Data Storage Networking (CRC Press) by visiting [storageio.com/books](http://storageio.com/books)

To become a chapter download sponsor or to discuss your other project opportunities contact us at [info@storageio.com](mailto:info@storageio.com).

#1 VM Backup

Veeam Backup & Replication for

**Hyper-V**

is HERE!

[Learn More](#)

**VEEAM**

StorageIOblog.com site sponsor