

1

Introduction

The purpose of this chapter is to convey the basic idea underlying this book. To this end we will first describe conventional server-centric IT architecture and sketch out its limitations (Section 1.1). We will then introduce the alternative approach of storage-centric IT architecture (Section 1.2), explaining its advantages using the case study ‘Replacing a Server with Storage Networks’ (Section 1.3). Finally, we explain the structure of the entire book and discuss which subjects are not covered (Section 1.4).

1.1 SERVER-CENTRIC IT ARCHITECTURE AND ITS LIMITATIONS

In conventional IT architectures, storage devices are normally only connected to a single server (Figure 1.1). To increase fault tolerance, storage devices are sometimes connected to two servers, with only one server actually able to use the storage device at any one time. In both cases, the storage device exists only in relation to the server to which it is connected. Other servers cannot directly access the data; they always have to go through the server that is connected to the storage device. This conventional IT architecture is therefore called server-centric IT architecture. In this approach, servers and storage devices are generally connected together by SCSI cables.

As mentioned above, in conventional server-centric IT architecture storage devices exist only in relation to the one or two servers to which they are connected. The failure of both of these computers would make it impossible to access this data. Most companies find

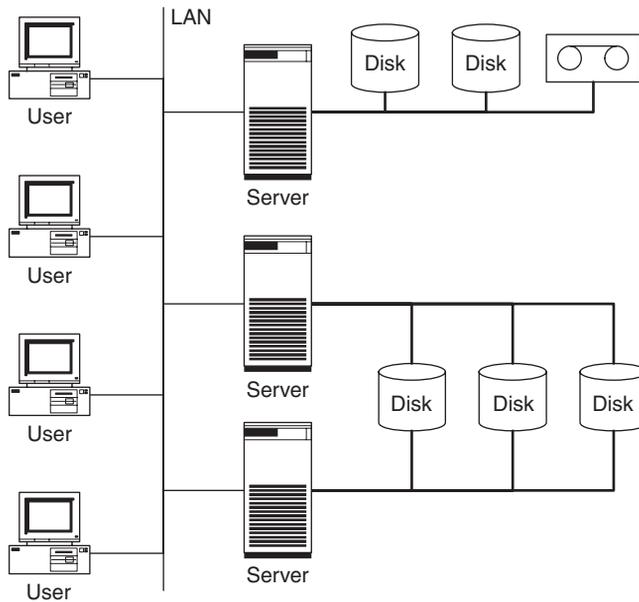


Figure 1.1 In a server-centric IT architecture storage devices exist only in relation to servers.

this unacceptable: at least some of the company data (for example, patient files, websites) must be available around the clock.

Although the storage density of hard disks and tapes is increasing all the time due to ongoing technical development, the need for installed storage is increasing even faster. Consequently, it is necessary to connect ever more storage devices to a computer. This throws up the problem that each computer can accommodate only a limited number of I/O cards (for example, SCSI cards). Furthermore, the length of SCSI cables is limited to a maximum of 25 m. This means that the storage capacity that can be connected to a computer using conventional technologies is limited. Conventional technologies are therefore no longer sufficient to satisfy the growing demand for storage capacity.

In server-centric IT environments the storage device is statically assigned to the computer to which it is connected. In general, a computer cannot access storage devices that are connected to a different computer. This means that if a computer requires more storage space than is connected to it, it is no help whatsoever that another computer still has attached storage space, which is not currently used (Figure 1.2).

Last, but not least, storage devices are often scattered throughout an entire building or branch. Sometimes this is because new computers are set up all over the campus without any great consideration and then upgraded repeatedly. Alternatively, computers may be consciously set up where the user accesses the data in order to reduce LAN data traffic. The result is that the storage devices are distributed throughout many rooms, which are

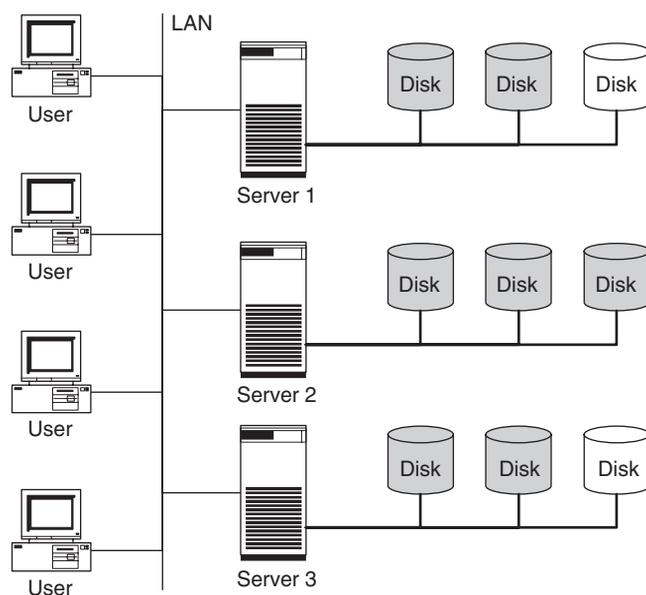


Figure 1.2 The storage capacity on server 2 is full. It cannot make use of the fact that there is still storage space free on server 1 and server 3.

neither protected against unauthorised access nor sufficiently air-conditioned. This may sound over the top, but many system administrators could write a book about replacing defective hard disks that are scattered all over the country.

1.2 STORAGE-CENTRIC IT ARCHITECTURE AND ITS ADVANTAGES

Storage networks can solve the problems of server-centric IT architecture that we have just discussed. Furthermore, storage networks open up new possibilities for data management. The idea behind storage networks is that the SCSI cable is replaced by a network that is installed in addition to the existing LAN and is primarily used for data exchange between computers and storage devices (Figure 1.3).

In contrast to server-centric IT architecture, in storage networks storage devices exist completely independently of any computer. Several servers can access the same storage device directly over the storage network without another server having to be involved. Storage devices are thus placed at the centre of the IT architecture; servers, on the other

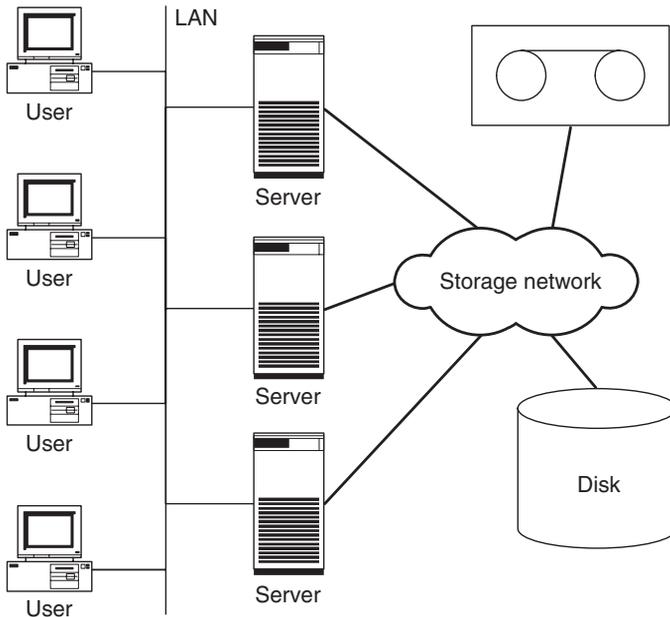


Figure 1.3 In storage-centric IT architecture the SCSI cables are replaced by a network. Storage devices now exist independently of a server.

hand, become an appendage of the storage devices that ‘just process data’. IT architectures with storage networks are therefore known as storage-centric IT architectures.

When a storage network is introduced, the storage devices are usually also consolidated. This involves replacing the many small hard disks attached to the computers with a large disk subsystem. Disk subsystems currently (in the year 2009) have a maximum storage capacity of up to a petabyte. The storage network permits all computers to access the disk subsystem and share it. Free storage capacity can thus be flexibly assigned to the computer that needs it at the time. In the same manner, many small tape libraries can be replaced by one big one.

More and more companies are converting their IT systems to a storage-centric IT architecture. It has now become a permanent component of large data centres and the IT systems of large companies. In our experience, more and more medium-sized companies and public institutions are now considering storage networks. Even today, most storage capacity is no longer fitted into the case of a server (internal storage device), but has its own case (external storage device).

1.3 CASE STUDY: REPLACING A SERVER WITH STORAGE NETWORKS

In the following we will illustrate some advantages of storage-centric IT architecture using a case study: in a production environment an application server is no longer powerful enough. The ageing computer must be replaced by a higher-performance device. Whereas such a measure can be very complicated in a conventional, server-centric IT architecture, it can be carried out very elegantly in a storage network.

1. Before the exchange, the old computer is connected to a storage device via the storage network, which it uses partially (Figure 1.4 shows stages 1, 2 and 3).
2. First, the necessary application software is installed on the new computer. The new computer is then set up at the location at which it will ultimately stand. With storage networks it is possible to set up the computer and storage device several kilometres apart.

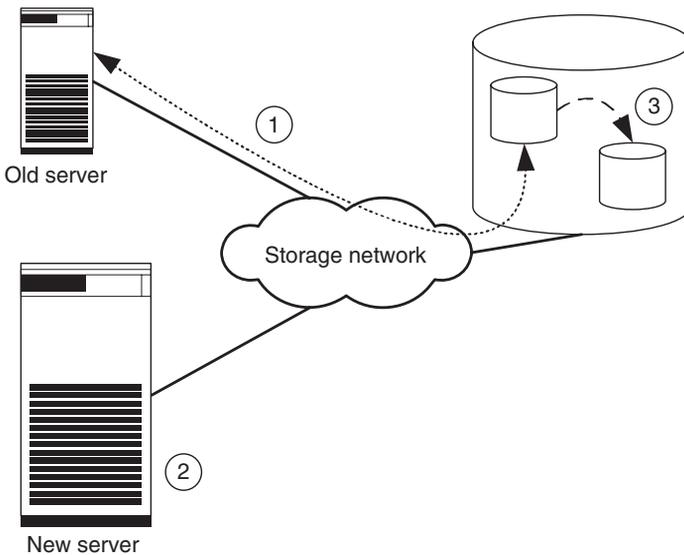


Figure 1.4 The old server is connected to a storage device via a storage network (1). The new server is assembled and connected to the storage network (2). To generate test data the production data is copied within the storage device (3).

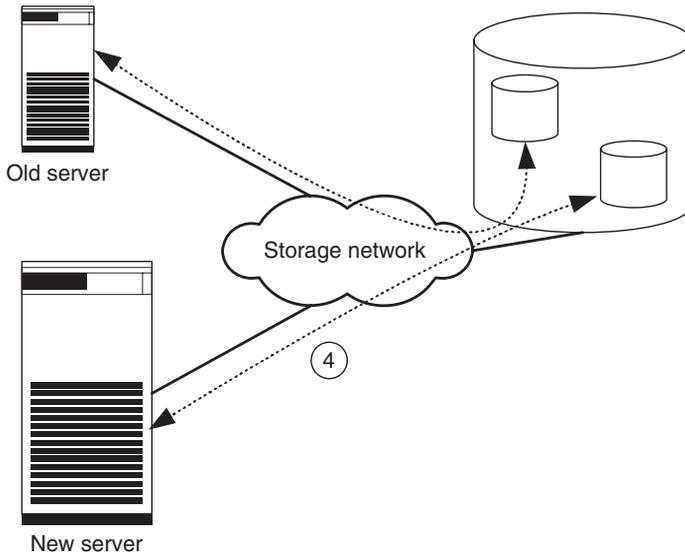


Figure 1.5 Old server and new server share the storage system. The new server is intensively tested using the copied production data (4).

3. Next, the production data for generating test data within the disk subsystem is copied. Modern storage systems can (practically) copy even terabyte-sized data files within seconds. This function is called instant copy and is explained in more detail in Chapter 2. To copy data it is often necessary to shut down the applications, so that the copied data is in a consistent state. Consistency is necessary to permit the application to resume operation with the data. Some applications are also capable of keeping a consistent state on the disk during operation (online backup mode of database systems, snapshots of file systems).
4. Then the copied data is assigned to the new computer and the new computer is tested intensively (Figure 1.5). If the storage system is placed under such an extreme load by the tests that its performance is no longer sufficient for the actual application, the data must first be transferred to a second storage system by means of remote mirroring. Remote mirroring is also explained in more detail in Chapter 2.
5. After successful testing, both computers are shut down and the production data assigned to the new server. The assignment of the production data to the new server also takes just a few seconds (Figure 1.6 shows steps 5 and 6).
6. Finally, the new server is restarted with the production data.

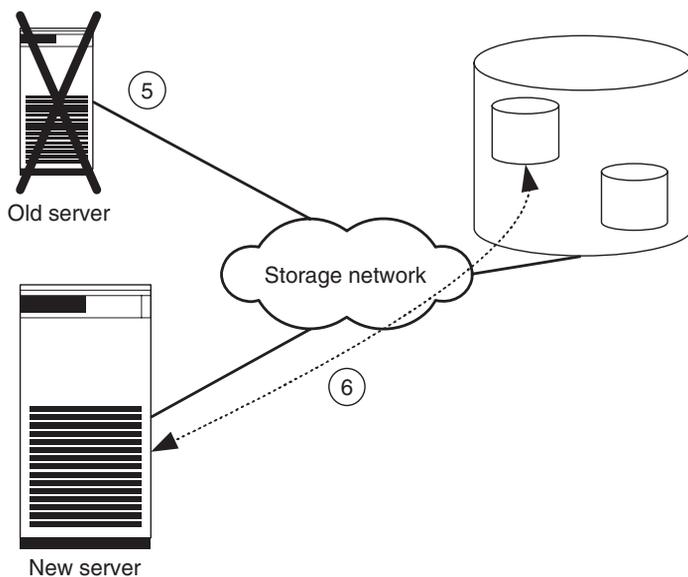


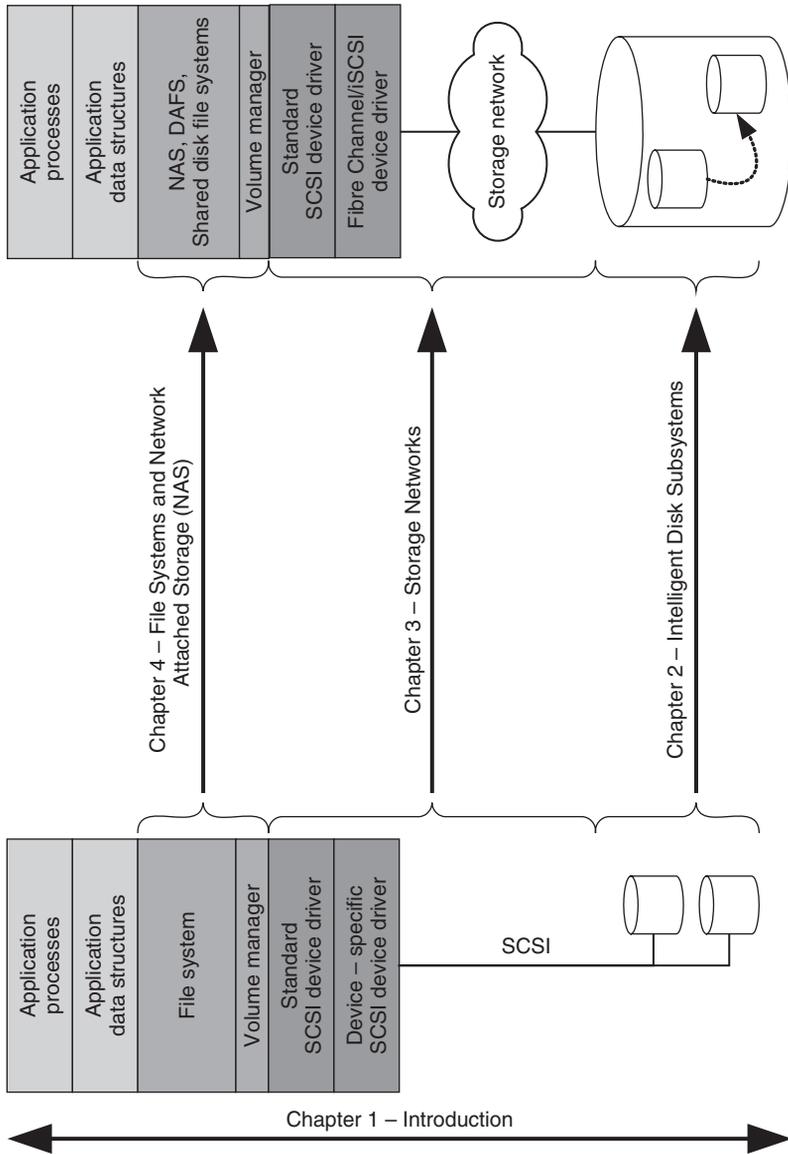
Figure 1.6 Finally, the old server is powered down (5) and the new server is started up with the production data (6).

1.4 THE STRUCTURE OF THE BOOK

One objective of this book is to illustrate the benefits of storage networks. In order to provide an introduction to this subject, this chapter has presented a few fundamental problems of conventional server-centric IT architecture and concluded by mentioning a few advantages of storage-centric IT architecture based upon the upgrade of an application server. The remaining chapters deal with the concepts and techniques that have already been sketched out and discuss further case studies in detail. The book is structured around the path from the storage device to the application (Figure 1.7).

In modern IT systems, data is normally stored on hard disks and tapes. It is more economical to procure and manage a few large storage systems than several small ones. This means that the individual disk drives are being replaced by disk subsystems. In contrast to a file server, an intelligent disk subsystem can be visualised as a hard disk server; other servers can use these hard disks that are exported via the storage network just as they can use locally connected disk drives. Chapter 2 shows what modern disk subsystems can do in addition to the instant copy and remote mirroring functions mentioned above.

- Chapter 12 – The SNIA Shared Storage Model
- Chapter 11 – Removable Media Management
- Chapter 10 – Management of Storage Networks
- Chapter 9 – Business Continuity
- Chapter 8 – Archiving
- Chapter 7 – Network Backup
- Chapter 6 – Application of Storage Networks
- Chapter 5 – Storage Virtualisation



The hardware of tapes and tape libraries changes only slightly as a result of the transition to storage networks, so we only touch upon this subject in the book. In Section 6.2.2 we will discuss the sharing of large tape libraries by several servers and access to these over a storage network and Chapter 11 will present the management of removable media including – among other removable media – tapes and tape libraries.

Fibre Channel has established itself as a technology with which storage networks can be efficiently realised for both open systems (Unix, Windows, Novell Netware, MacOS, OS/400) and mainframes. Where Fibre Channel introduces a new transmission technology, its alternative Internet SCSI (iSCSI) is based upon the proven TCP/IP and Gigabit Ethernet. InfiniBand and Fibre Channel over Ethernet (FCoE) are two additional approaches to consolidate all data traffic (storage, cluster) onto a single transmission technology. All these technologies are subject of Chapter 3.

File systems are of interest in this book for two reasons. First, pre-configured file servers, also known as Network Attached Storage (NAS), have established themselves as an important building block for current IT systems. Storage networks can also be realised using NAS servers. In contrast to the block-oriented data traffic of Fibre Channel, iSCSI and FCoE in this approach whole files or file fragments are transferred.

So-called shared-disk file systems represent the other interesting development in the field of file systems. In shared-disk file systems, several computers can access the same data area in an intelligent disk subsystem over the storage network. The performance of shared-disk file systems is currently significantly better than those of Network File System (NFS), Common Internet File System (CIFS), AppleTalk or the above-mentioned NAS servers. Examples of problems are discussed on the basis of shared-disk file systems that must also be solved in the same manner for comparable applications such as parallel databases. Chapter 4 deals with Network Attached Storage (NAS) and shared-disk file systems.

The first four chapters of the book discuss fundamental components and technologies with regard to storage networks. As storage networks have become more widespread, it has become clear that the implementation of a storage network alone is not sufficient to make efficient use of the resources of ever growing storage networks. Chapter 5 sketches

Figure 1.7 The book is divided into two main parts. The first part discusses the fundamental techniques that underlie storage networks. In particular, these apply to intelligent disk subsystems (Chapter 2), block-oriented storage networks (Chapter 3) and file systems (Chapter 4). We also outline how virtualisation can manage storage more efficiently (Chapter 5). The second part of the book discusses the application of these new technologies. In particular, we discuss standard applications such as storage pooling and clustering (Chapter 6), backup (Chapter 7), archiving (Chapter 8) and business continuity (Chapter 9). These chapters show how storage networks help to develop IT systems that are more flexible, fault-tolerant and powerful than traditional systems. We then discuss the management of storage networks (Chapter 10) and removable media (Chapter 11). Finally, the SNIA Shared Storage Model is presented (Chapter 12).

out the difficulties associated with the use of storage networks and it introduces storage virtualisation – an approach that aims to reduce the total cost of ownership (TCO) for accessing and managing huge amounts of data. It further discusses possible locations for the realisation of storage virtualisation and discusses various alternative approaches to storage virtualisation such as virtualisation on block level and virtualisation on file level or symmetric and asymmetric storage virtualisation.

The first chapters introduce a whole range of new technologies. In Chapter 6 we turn our attention to the application of these new techniques. This chapter uses many case studies to show how storage networks help in the design of IT systems that are more flexible and more fault-tolerant than conventional server-centric IT systems.

Data protection (Backup) is a central application in every IT system. Using network backup systems it is possible to back up heterogeneous IT environments with several thousands of computers largely automatically. Chapter 7 explains the fundamentals of network backup and shows how these new techniques help to back up data even more efficiently. Once again, this clarifies the limitations of server-centric IT architecture and the benefits of the storage-centric IT architecture.

Digital archiving is another important application in storage networks. The law requires that more and more data is kept for years, decades and even longer under strictly regulated conditions. For example, none of the archived data is allowed to be changed or deleted prior to the expiration of the retention times. Due to the long retention times and technical progress, data is required to be copied periodically to new storage media or systems. Chapter 8 discusses the fundamental requirements for digital archiving and presents a number of techniques and solutions that are based on them.

Continuous access to business-critical data and applications, even in a crisis situation, is essential for a company's ability to exist. This does not only apply to those areas one thinks of automatically in this context, such as stock broking, air traffic control, patient data, or Internet companies like Amazon and eBay. An increasing number of smaller and medium-sized companies are now delivering their products to customers worldwide or are tightly integrated into the production processes of larger companies, such as automobile manufacturers, through just-in-time production and contractually agreed delivery times. Chapter 9 introduces the area of business continuity with special consideration of storage networks and discusses different techniques along with possible solutions.

Storage networks are complex systems made up of numerous individual components. As one of the first steps in the management of storage networks it is necessary to understand the current state. This calls for tools that help to answer such questions as 'Which server occupies how much space on which disk subsystem?', 'Which servers are connected to my storage network at all?', 'Which hardware components are in use and how great is the load upon the network?'. In this connection the monitoring of the storage network with regard to faults and performance and capacity bottlenecks of file systems is also important. The second step relates to the automation of the management of storage networks: important subjects are rule-based error handling and the automatic allocation of free storage capacity. Chapter 10 deals in detail with the management of storage networks and in this connection also discusses standards such as Simple Network Management Protocol

(SNMP), Common Information Model/Web-based Enterprise Management (CIM/WBEM) and Storage Management Initiative Specification (SMI-S).

Removable media represent a central component of the storage architecture of large data centres. Storage networks allow several servers, and thus several different applications, to share media and libraries. Therefore, the management of removable media in storage networks is becoming increasingly important. Chapter 11 deals with the requirements of removable media management and it introduces the IEEE 1244 Standard for Removable Media Management.

Storage networks are a complex subject area. There is still a lack of unified terminology, with different manufacturers using the same term to refer to different features and, conversely, describing the same feature using different terms. As a result, it is often unclear what kind of a product is being offered by a manufacturer and which functions a customer can ultimately expect from this product. It is thus difficult for the customer to compare the products of the individual manufacturers and to work out the differences between the alternatives on offer. For this reason, the Technical Council of the Storage Networking Industry Association (SNIA) has introduced the so-called Shared Storage Model in 2001 in order to unify the terminology and descriptive models used by the storage network industry. We introduce this model in Chapter 12.

What doesn't this book cover?

In order to define the content it is also important to know which subjects are not covered:

- *Specific products*
Product lifecycles are too short for specific products to be discussed in a book. Products change, concepts do not.
- *Economic aspects*
This book primarily deals with the technical aspects of storage networks. It discusses concepts and approaches to solutions. Prices change very frequently, concepts do not.
- *Excessively technical details*
The book is an introduction to storage networks. It does not deal with the details necessary for the development of components for storage networks. The communication of the overall picture is more important to us.
- *The planning and implementation of storage networks*
Planning and implementation require knowledge of specific products, but products change very frequently. Planning and implementation require a great deal of experience. This book, on the other hand, is designed as an introduction. Inexperienced readers should consult experts when introducing a storage network. Furthermore, a specific implementation must always take into account the specific environment in question. It is precisely this that this book cannot do.

