

""UGTXGT/EGPVTR 'K'CTEJ KGEVWTG'CPF 'KUNIKO KCVKQP U'
 ""

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

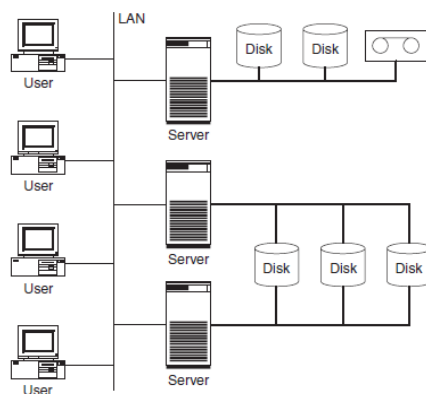


Figure 1.1 In a server-centric IT architecture storage devices exist only in relation to servers.

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 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.

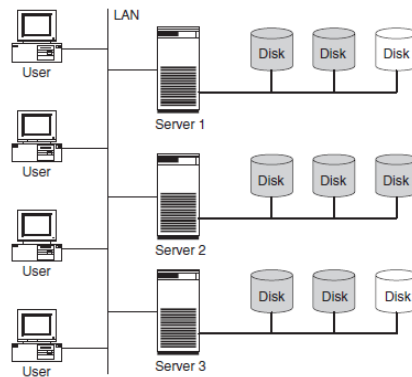


Figure 1.2 The storage capacity on server 2 is full. It cannot make use of the fact that there is still storage space free on server 1 and server 3.