SYMMETRIC STORAGE VIRTUALISATION IN THE NETWORK

The symmetric and asymmetric virtualisation models are representatives of storage virtualization in the network. In both approaches it is possible to perform virtualisation both on block and on file level. In both models the virtualisation entity that undertakes the separation between physical and logical storage is placed in the storage network in the
form of a specialised server or a device. This holds all the meta-information needed for the virtualisation. The virtualisation entity is therefore also called the metadata controller.
Its duties also include the management of storage resources and the control of all storage functions that are offered in addition to virtualisation. Symmetric and asymmetric virtualisation differ primarily with regard to their distribution of data and control flow. Data flow is the transfer of the application data between the servers and storage devices. The control flow consists of all metadata and control information necessary for virtualisation between virtualisation entity and storage devices and servers. In symmetric storage virtualisation the data flow and the control flow travel down the same path. By contrast, in asymmetric virtualisation the data flow is separated from the control flow.

6.7.1 Symmetric storage virtualisation

In symmetric storage virtualisation the data and control flow go down the same path (Figure 5.15). This means that the abstraction from physical to logical storage necessary for virtualisation must take place within the data flow. As a result, the metadata controller
is positioned precisely in the data flow between server and storage devices, which is why symmetric virtualisation is also called in-band virtualization.

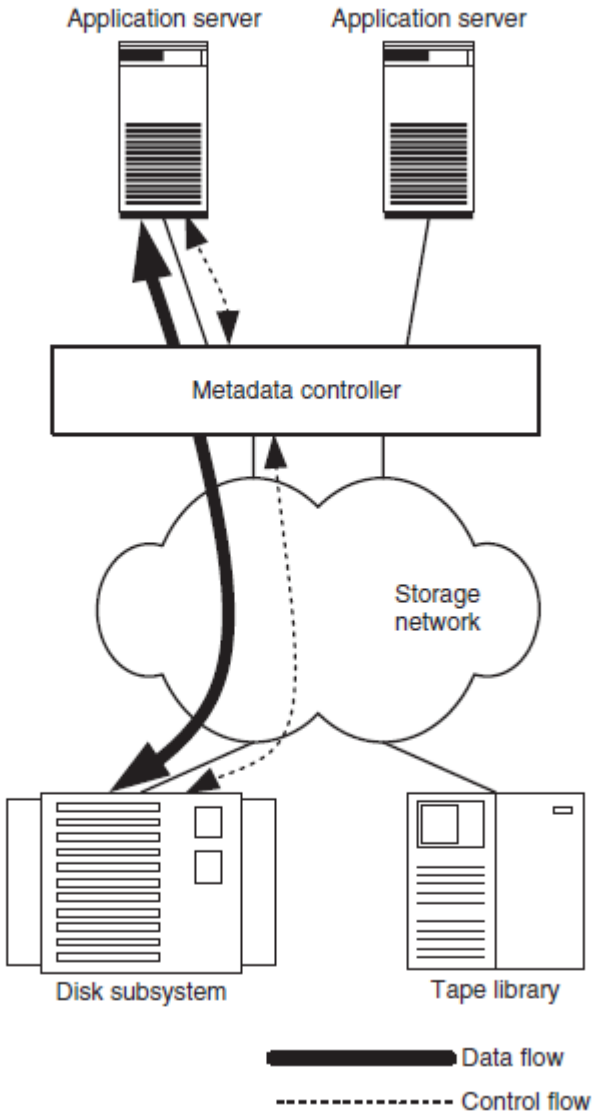

Figure 6.15 In symmetric virtualisation, data and control flow travel down the same path. The abstraction from physical to logical storage takes place within the data stream. In addition to the control of the virtualisation, all data between servers and storage devices now flow through the metadata controller. To this end virtualisation is logically structured in two layers: the layer for the management of the logical volumes and the data access layer (Figure 5.16):

1. The volume management layer is responsible for the management and configuration of the storage devices that can be accessed directly or via a storage network and it provides the aggregation of these resources into logical disks. 2. The

data access layer makes the logical drives available for access either on block or file level, depending upon what degree of abstraction is required. These logical drives can thus be made available to the application servers by means of appropriate protocols. In the case of virtualisation on block level, this occurs in the form of a virtual disk and in the case of virtualisation on file level it takes place in the form of a file system.

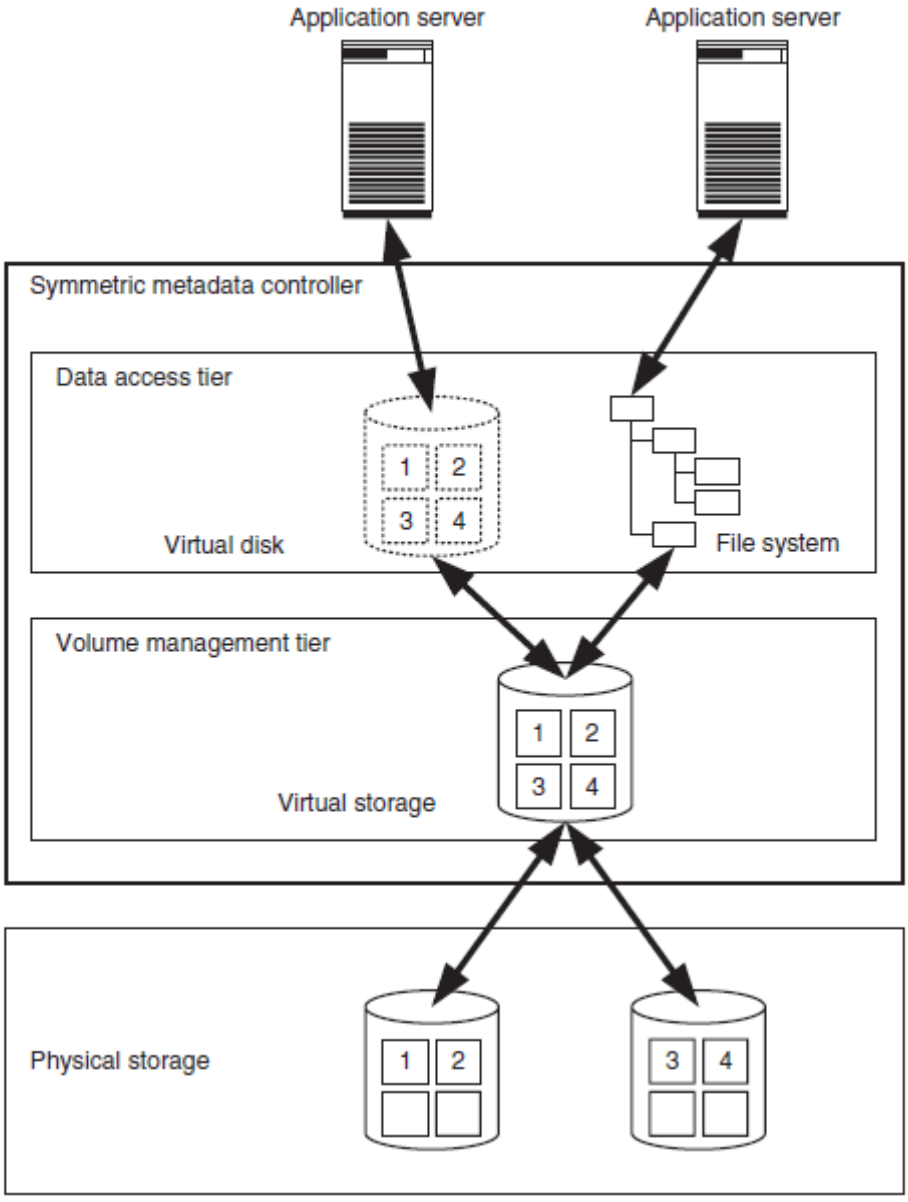

Figure 6.16 In symmetric virtualisation the metadata controller consists of a data access layer and a volume management layer.

In symmetric virtualisation all data flow through the metadata controller, which means that this represents a potential bottleneck. To increase performance, therefore, the metadata controller is upgraded by the addition of a cache. With the use of caching and symmetric virtualisation it is even possible to improve the performance of an existing storage network as long as exclusively write-intensive applications are not used.

A further issue is fault-tolerance. A single metadata controller represents a single point of failure. The use of cluster technology (Section 6.3.2) makes it possible to remove the single point of failure by using several metadata controllers in parallel. In addition, a corresponding load distribution provides a performance increase. However, a configuration failure or a software failure of that cluster can lead to data loss on all virtualised resources.

In the case of a network-based virtualisation spanning several servers and storage devices, this can halt the activity of a complete data centre (Section 6.3.4).

Thus the advantages of symmetric virtualisation are evident:

• The application servers can easily be provided with data access both on block and file level, regardless of the underlying physical storage devices.

• The administrator has complete control over which storage resources are available to which servers at a central point. This increases security and eases the administration.

• Assuming that the appropriate protocols are supported, symmetric virtualisation does not place any limit on specific operating system platforms. It can thus also be used in heterogeneous environments.

• The performance of existing storage networks can be improved by the use of caching and clustering in the metadata controllers.

• The use of a metadata controller means that techniques such as snapshots or mirroring can be implemented in a simple manner, since they control the storage

access directly. They can also be used on storage devices such as JBODs or simple RAID arrays that do not provide to these techniques themselves.

The disadvantages of a symmetric virtualisation are:

• Each individual metadata controller must be administered. If several metadata controllers are used in a cluster arrangement, then the administration is relatively complex and time-consuming particularly due to the cross-computer data access layer. This disadvantage can, however, be reduced by the use of a central administration console for the metadata controller.

• Several controllers plus cluster technology are indispensable to guarantee the fault-tolerance of data access.

• As an additional element in the data path, the controller can lead to performance problems, which makes the use of caching or load distribution over several controllers indispensable.

• It can sometimes be difficult to move the data between storage devices if this is managed by different metadata controllers.

Source : http://elearningatria.files.wordpress.com/2013/10/cse-viii-storage-area-networks-06cs833-notes.pdf