Understanding Storage Virtualization of Infortrend ESVA

White paper

Abstract

This white paper introduces different ways of implementing storage virtualization and illustrates how the virtualization technology embedded in Infortrend ESVA storage systems benefits users with optimized efficiency of storage utilization.

Storage Virtualization Overview

Storage virtualization is a technology to create logical abstraction that isolates applications from physical storage systems. It allows administrators to gather, organize, allocate and manage a heterogeneous collection of storage resources without regard to the applications using it. The main benefits of storage virtualization are improved resource utilization and simplified management, which make the technology especially attractive in the challenging economy.

Below we will introduce the three basic approaches to implement storage virtualization: host-based, network-based and storage-based.

Host-based Virtualization

Host-based virtualization runs logical volume management software on the host operating system to intercept I/O requests, perform metadata lookup and redirect the requests to storage devices. All storage resources accessible to the host can be integrated into the virtualized architecture. Requiring no hardware infrastructure change, the virtualization can often be easily deployed at affordable costs.

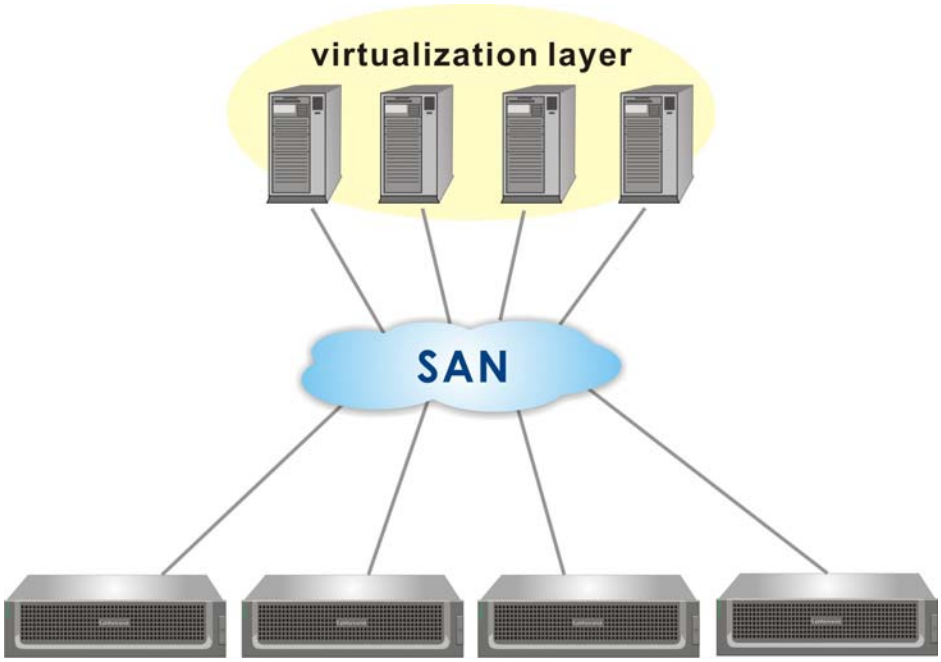

Figure 1. Host-based Virtualization

To deploy host-based virtualization, users have to make sure that their host operating system is compatible with the virtualization software. Since the software is installed on each host machine, it introduces extra server overhead. Moreover, due to the

exclusive storage ownership, if a host crashes or is offline, the storage assets it owns will also become inaccessible. To solve the problem of single point of failure, users need additional cluster volume management software to coordinate between each server's view of external storage and in turn enable more granular storage sharing.

Network-based Virtualization

Network-based virtualization runs software on dedicated appliances or intelligent switches. They are both positioned between host and storage to perform metadata mapping and I/O redirection.

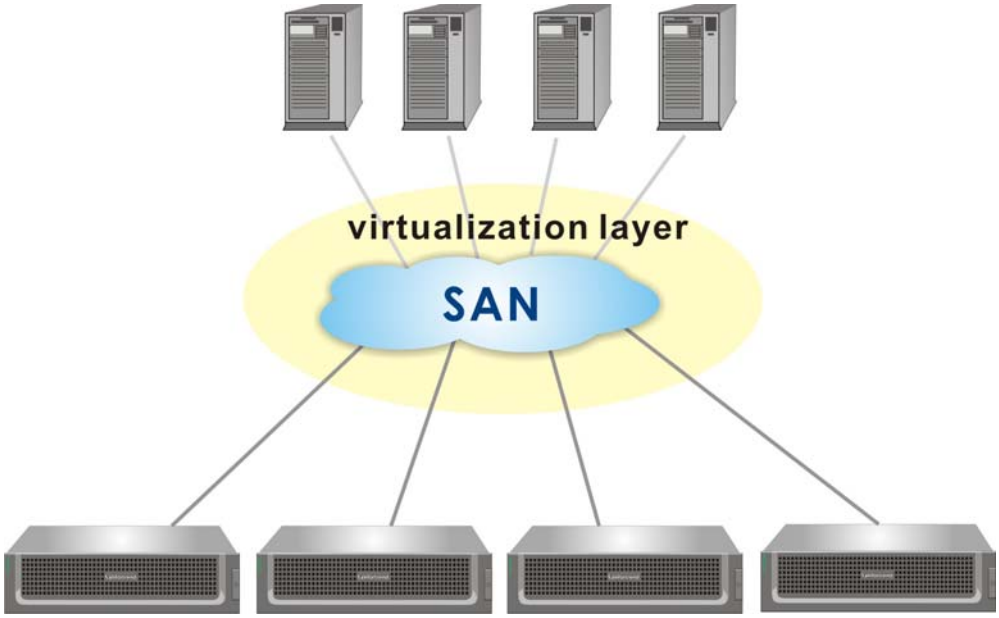

Figure 2. Network-based Virtualization

In-Band vs. Out-of-Band

For appliance-based solutions, there are two kinds of implementations: in-band and out-of-band. In-band virtualization devices are located in the data path between host and storage devices. Hosts perform I/Os to the in-band appliance and the appliance in turn performs I/Os to the storage devices. All I/O requests have to pass through the appliance and the host never directly interacts with storage devices.

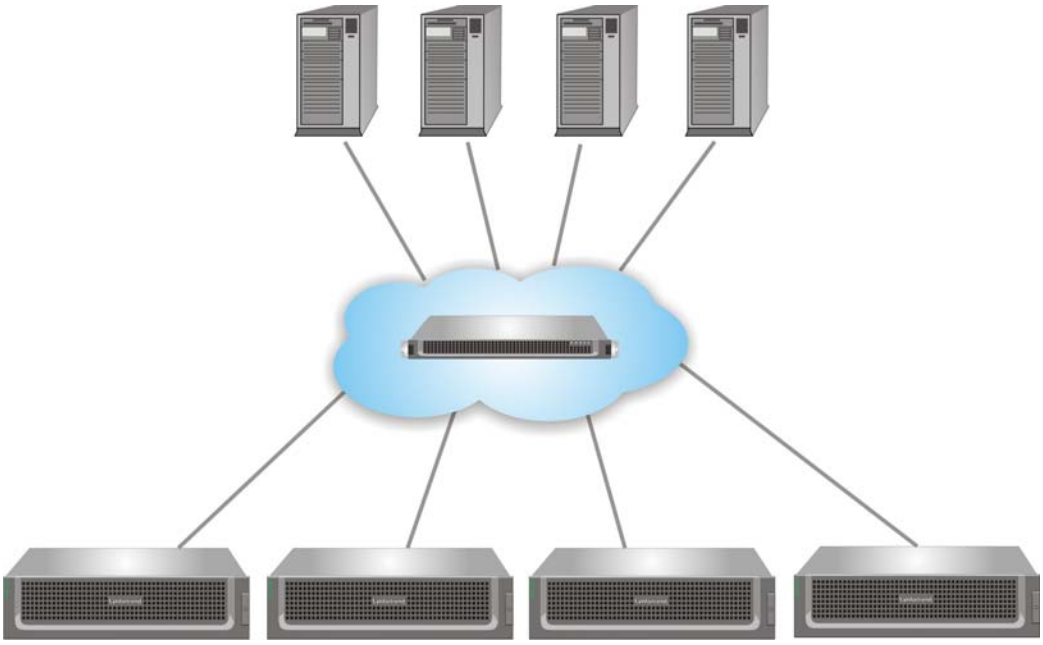

Figure 3. In-band Virtualization

Instead of directly handling I/O requests, out-of-band virtualization devices only perform the metadata mapping function. Users need to install additional software in the host so that the host knows to first request the location of the actual data from the appliance before sending I/Os.

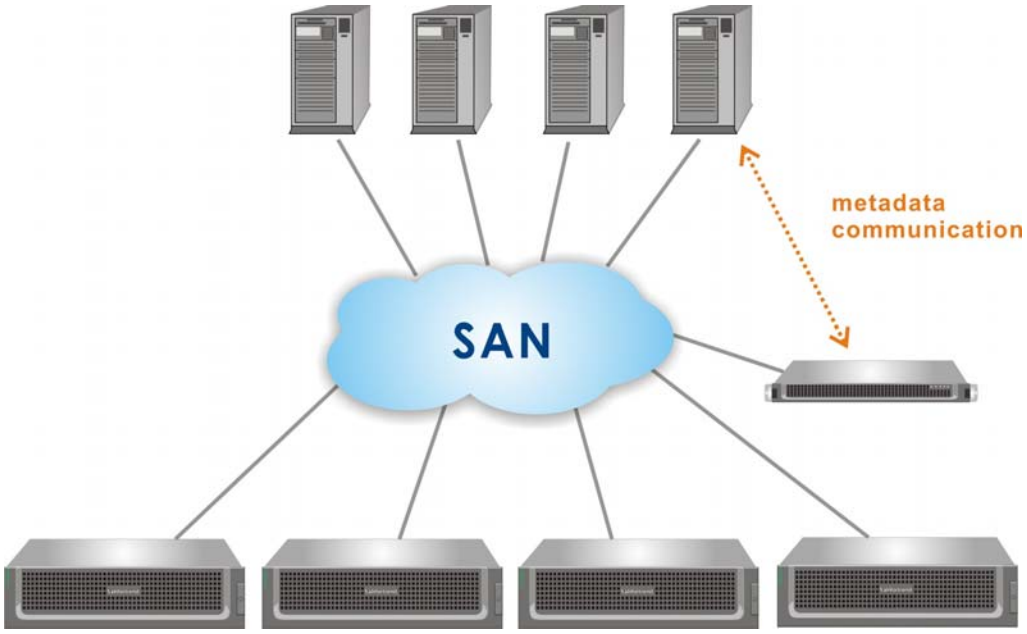

Figure 4. Out-of-band Virtualization

With network-based storage virtualization, users can realize true heterogeneous

virtualization and simplify management with a single interface for all virtualized storage. If users adopt in-band appliance-based solutions, they can further improve performance through data caching. However, the implementation of network-based virtualization is limited by the vendor support matrix of the virtualization device. Moreover, out-of-band solutions require specific host-based software and in-band solutions may degrade performance and constitute single point of failure in the data path.

Storage-based Virtualization

Storage-based virtualization embeds the technology in the storage array. Although this way of virtualization is applicable only across storage devices from the same vendor, it can bring users many substantial benefits.

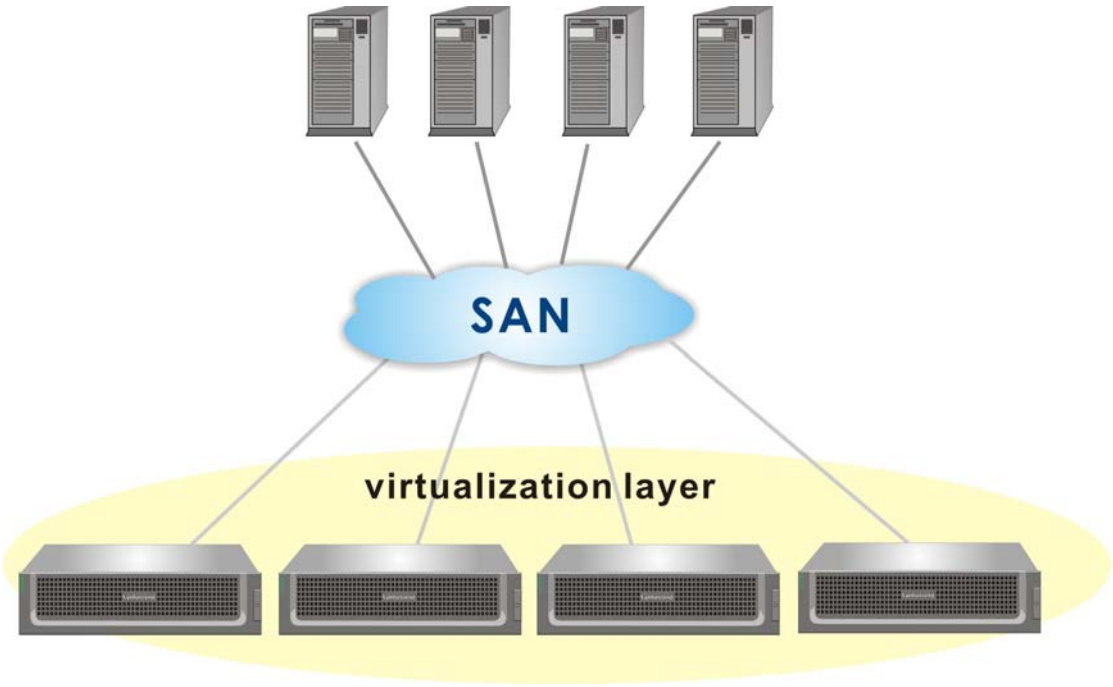

Figure 5. Storage-based Virtualization

Please see the following section to understand how Infortrend ESVA benefits users based on this technology.

Infortrend ESVA

ESVA stands for Enterprise Scalable Virtualized Architecture, Infortrend's storage solution line aimed at mid-range enterprises. Based on storage virtualization technology, ESVA helps users increase the efficiency of storage utilization with features such as thin provisioning and scale-out expansion. To facilitate data backup and

enable high availability deployments, ESVA also enriches the virtualized foundation with data protection features such as snapshot, volume copy, and remote replication.

Architecture

Shown below is an ESVA virtual pool in its most basic form. Disk drives are separately managed by dual controllers in an ESVA unit. One disk drive is automatically reserved as an enclosure hot spare. These drives form two logical drives and become woven into a virtual pool.

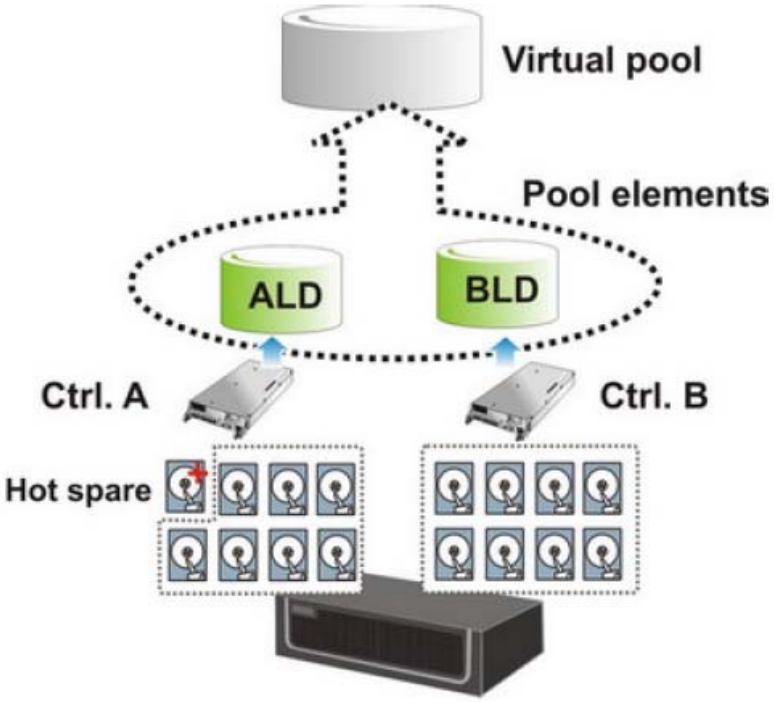

Figure 6: A Basic Virtual Pool

A virtual pool can span across multiple systems. The pooled capacity expands either horizontally by adding more ESVA units or vertically by attaching expansion enclosures. While more raw disks are added, they can be added into an existing pool, or used to create new pools.

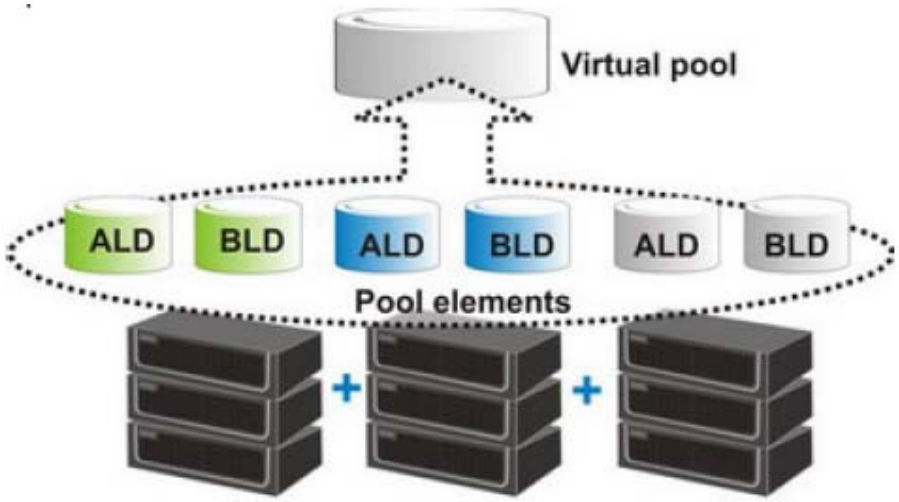

Figure 7: An Extended Virtual Pool

Dynamic Capacity Allocation

The idea of thin provisioning is to allocate a large virtual capacity for an application server regardless of the actual physical capacity available. The technology allows users to flexibly allocate the storage space consolidated through virtualization as virtual volumes to each application. If an application is filling a volume to full, the volume can be easily expanded by drawing space from the virtual pool. If the virtual pool itself runs out of storage capacity, it can be expanded by appending expansion enclosures or another ESVA system. The “just-in-time capacity” essentially eliminates allocated but unused storage and greatly simplifies storage provisioning tasks, reducing administration costs.

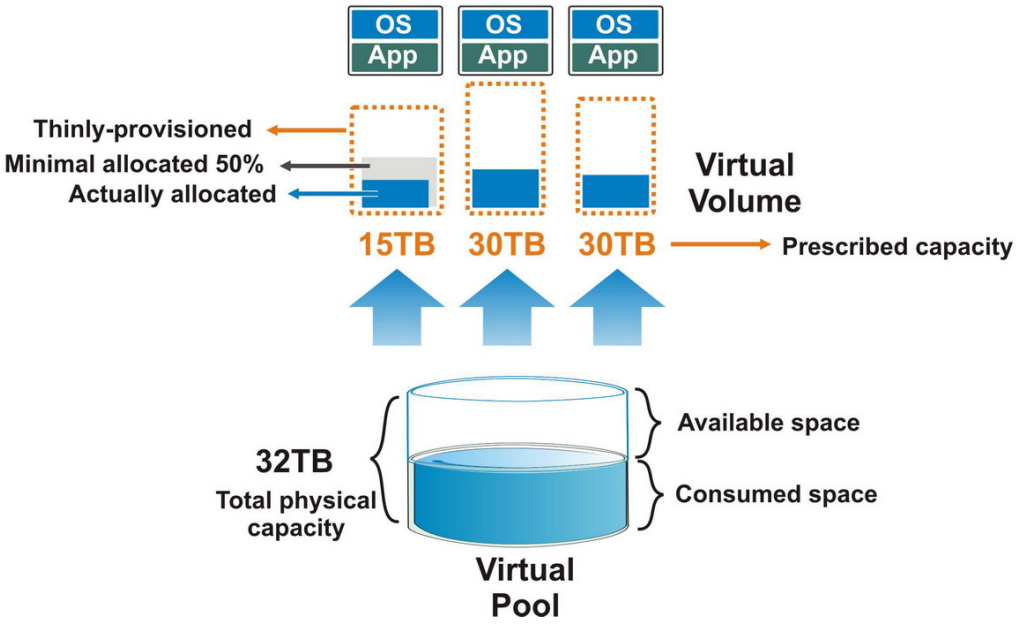

Figure 8. Thin provisioning

Scalable Performance

In addition to capacity efficiency, combining resources into virtual pools also allows scale-out performance. ESVA's load-balancing technology optimizes system performance without any manual intervention. When host issues a write request to a storage pool, its composing data blocks are automatically distributed to storage systems in a balanced way. As shown in **Figure 7**, the six data blocks composing the write request are evenly distributed on the three member storage systems, with each member receiving two blocks.

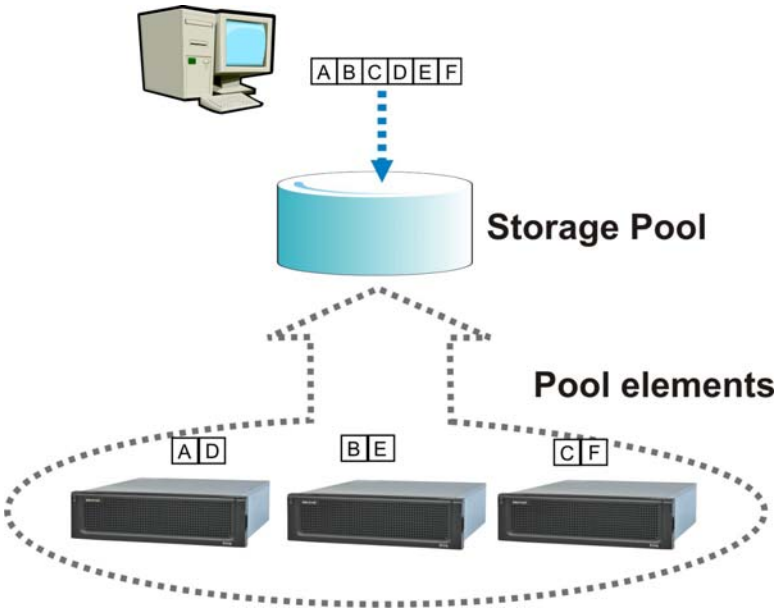

Figure 9. Distributed Load Balancing

Conclusion

As storage-based virtualization is independent from host platforms and easy to deploy with no additional hardware or software, it is an ideal choice to help companies improve storage utilization efficiency at minimized extra costs. Designed for mid-range enterprises, Infortrend ESVA enriches its virtualized architecture with thin provisioning, scale-out technology and data protection to bring users optimal resource efficiency, management convenience, operational resilience and increased availability at reduced costs. The power to effectively and promptly respond to changing IT challenges will translate into enduring competitive edges in the business world.