Using ESVA to Optimize Returns on Investment

White paper

Abstract

This white paper explains how ESVA helps users optimize returns of achievement by achieving maximized capacity utilization, fully-utilized performance, prioritized volume accessibility, and cost-effective, as-needed scaling with its advanced technologies.

The Infortrend ESVA Series is a leading-edge storage solution designed for mid-range enterprise Fibre Channel or iSCSI SAN. At affordable price, it meets mission-critical storage demands for performance, scalability and reliability with advanced hardware design and comprehensive data services. On the innovative Enterprise Scalable Virtualized Architecture, various features, including storage virtualization, thin provisioning, distributed load balancing, automatic data migration, prioritized volume accessibility, and array-based snapshot and replication, are consolidated to bring users three main benefits: optimize returns on investment, simplify storage infrastructure and maximize application productivity. In this document, we will illustrate in details how ESVA technologies help users optimize returns on investment in a business environment.

Maximized Capacity Utilization

Low storage utilization has always been a prominent problem in the IT environment. With direct-attached storage (DAS), the utilization is about 30%. Even if DAS is consolidated into a storage area network (SAN), the utilization is still often no more than 50%. One of the main reasons accounting for low utilization is that the free capacity dispersedly located on different physical systems can't be consolidated for uses.

As shown in **Figure 1**, in SAN configuration, there is 5TB of free space left on the two storage systems.

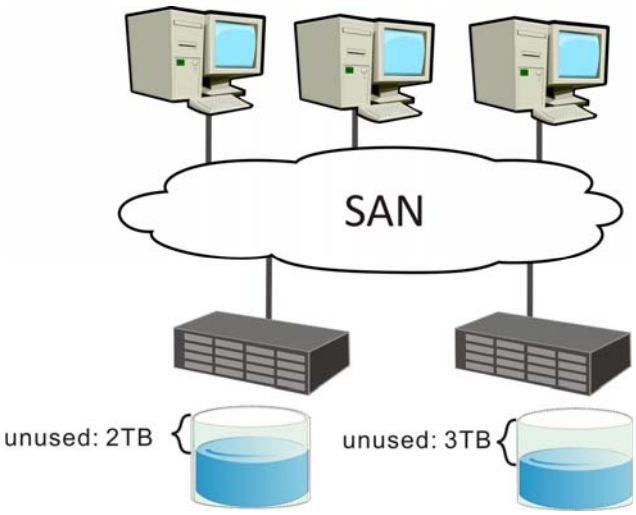

Figure 1. Dispersed Unused Space

However, if users want to introduce a new application requiring 4TB in, the physical boundary makes them unable to consolidate the dispersedly located capacities into a data volume for uses. Users will therefore need to add additional capacity for the new application, and this will reduce utilization. To solve the problem, users need the help of

storage virtualization technology.

Storage virtualization is the basis of ESVA. It is a technology to aggregate resources on different physical storage systems into a single “storage pool.” It can eliminate the problem of dispersed space because the consolidation it enables overcomes physical boundaries. As shown in **Figure 2**, in the virtualized environment, the consolidated 5TB of free space can be freely leveraged by the newly-introduced application.

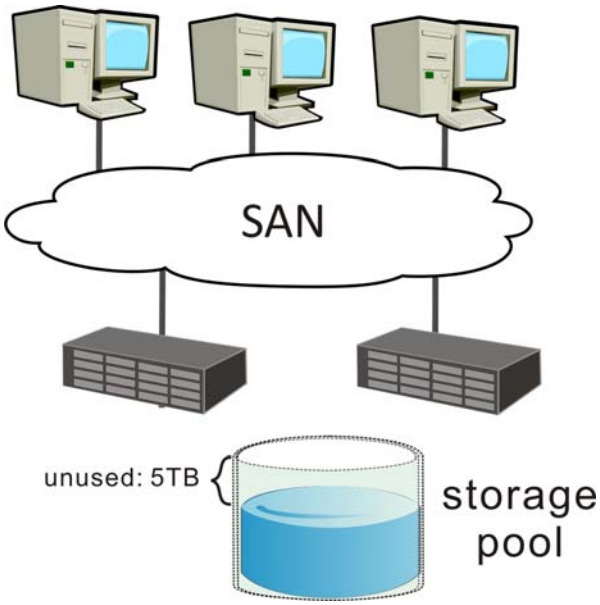

Figure 2. Consolidated Unused Space

Another reason accounting for low storage utilization is the traditional way of provisioning. In the traditional way of provisioning, for an application to leverage a certain amount of capacity, IT managers have to first create a data volume of the amount and then allocate it to the host. The act of allocation means privileging a certain host to use a certain amount of capacity. Since expanding data volumes causes downtime, IT managers tend to create data volumes of the amount much larger than immediately required to delay expansion hassles. That is why this way of provisioning is often called thick provisioning. As shown in **Figure 3**, host 1 currently has only 5TB of data but is assigned a data volume of 20TB. And host 2 currently has 3TB of data but is assigned a volume of 10TB.

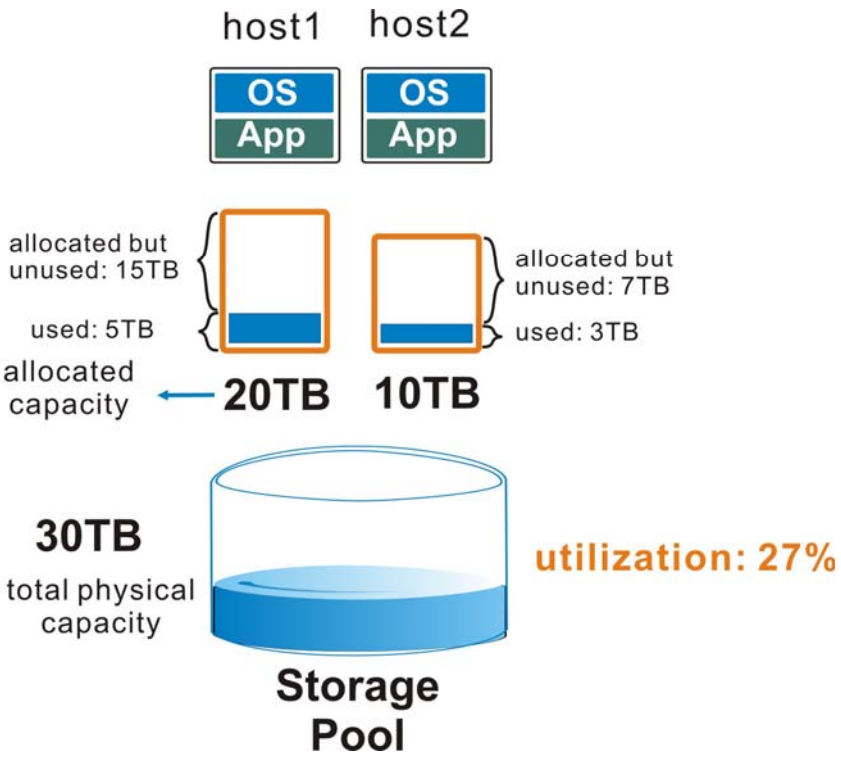

Figure 3. Traditional Provisioning

The 22TB of free space is wasted because the allocation prevents other applications from using it. These large and underutilized data volumes lead to the storage utilization as low as 27%. They also inflict much wasted investment in space, power and cooling expenses. Infortrend ESVA Series delivers thin provisioning capability to solve the utilization problem.

Thin provisioning is a technology to allocate just-in-time free space to applications when data are written. It largely improves utilization because it eliminates allocated but unused space. In the figure below, we can see a storage pool of 10TB.

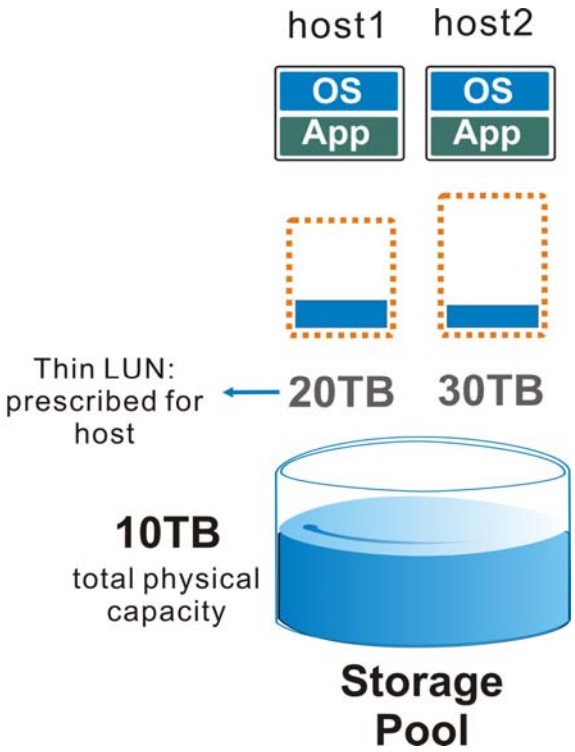

Figure 4. Thin Provisioning

With thin provisioning, users can “bluff” capacity to applications by prescribing 20TB of space to host 1 and 30TB to host 2. The prescription does not fix any amount of space for the applications. The pooled capacity is dynamically allocated when data writes occur. Seeing the total capacity is using up, users can expand the pool on-the-fly and the added capacity will be automatically leveraged by applications. The addition of capacity is completely transparent to the online service.

By implementing thin provisioning capability on virtualized architecture, Infortrend makes ESVA a solution to help users make the most efficient utilization of storage capacity. Besides minimizing the wasted expenditure on underutilized data volumes, ESVA also greatly profits users by delaying storage acquisition for them to leverage storage price declines.

Fully-utilized Performance

Besides capacity, ESVA also ensures users to make the most of their storage computing power. In the traditional IT environment, since applications produce data in different I/O rates, the uneven data distribution causes poor performance utilization. While some channels are overloaded and become performance-dragging bottlenecks, some are underutilized. To optimize overall performance, administrators have to regularly re-allot workloads among hardware resources concerning required capacity,

performance and service levels of applications. This task takes much time, effort and downtime to complete, and what's worse, only if workloads vary, the optimized state would be easily broken. To restore the state, the whole process has to be gone through all over again.

In ESVA's scale-out architecture, optimized performance can be easily achieved without any manual intervention. When host issues a write request to a storage pool, its composing data blocks would be distributed to element storage systems in a balanced way. As shown in **Figure 5**, the six data blocks composing the write request are evenly distributed on the three member storage systems, with each member receiving two blocks.

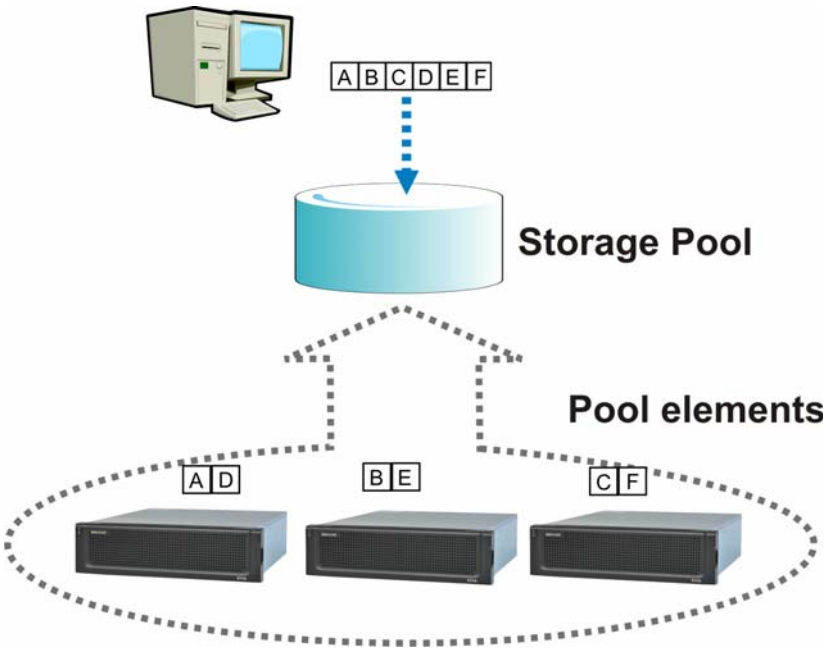

Figure 5. Distributed Load Balancing

The automatic and continuous load-balancing enables fully-utilized performance to boost productivity. Even if users add or remove subsystems so that the configuration changes, the optimized state can still be maintained since the existing data would be migrated to ensure balanced workloads. In **Figure 6**, originally there are three data blocks on each of the two member systems. When a new member is introduced in, the data blocks are migrated to make all members equally share the workloads, with each carrying two blocks.

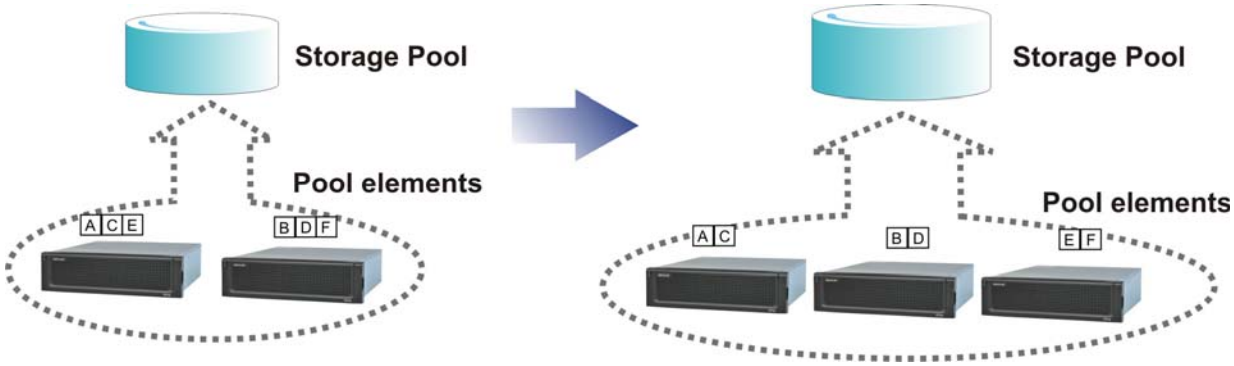

Figure 6. Automatic Data Migration

With ESVA, users can be assured that their storage is always working in its fullest power to drive the best business profits for them.

Prioritized Volume Accessibility

In default configuration of ESVA's high-performance architecture, storage processing power is distributed among applications according to their workloads. Based on the principle that the read/write requests are processed in the order they come into the pool, the most data-intensive application can get the most storage power for its service. As shown in **Figure 7**, the backup server delivers the highest I/O rates and therefore consumes the most resources.

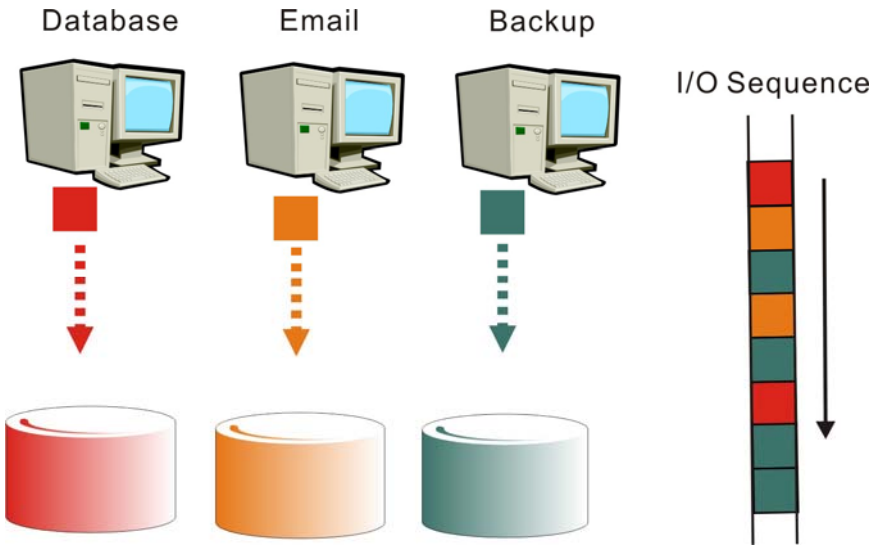

Figure 7. Un-prioritized Read/Write Requests

This way of computing power allocation causes problem because in some business environment, the most data-intensive application is not necessarily the most business-critical one. For example, in the situation shown in **Figure 7**, when the

backup job is intensively done, the performance of database application would be seriously degraded and may not be able to achieve the ideal service level. To make sure that applications are served by proper amount of resources according to their importance, Infortrend equips ESVA with a volume accessibility prioritizing mechanism. By setting priority (high, middle or low) for respective data volumes, users can make the read/write requests processed in their priority order, instead of in their coming order. In **Figure 8** we can see that all database I/Os are put ahead in the I/O sequence because the volume for database service is set with High priority.

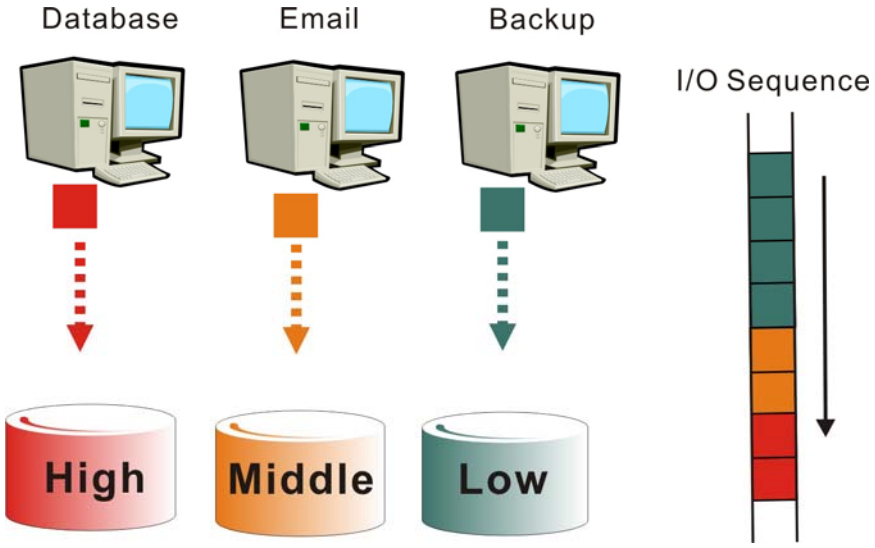

Figure 7. Prioritized Read/Write Requests

With this prioritizing mechanism, the computing resources can be allocated in the most efficient way to meet individual performance needs of all applications.

Cost-effective, As-needed Scaling

In traditional scale-up architecture, storage scaling is achieved by “replacing.” When the performance or capacity limitation of a single storage system is reached, users have to buy a new system to replace the old one. This kind of scaling is not cot-effective in two senses. First, the current investment can not be protected. After the replacement is done, the old system will typically be discarded or relegated to secondary applications. It can’t be used to enhancing performance or providing capacity for the primary application. The second reason of poor cost-efficiency is that users are forced to spend much upfront investment on what they do not think necessary now because their upgrade choices are limited by what vendors offer. For example, even if users think they currently need only two times better performance, they may have to buy a new system delivering 5 to 6 times better performance

because it is the only available next-level upgrade choice. What's worse, the new system often comes with poorer price/performance than the old system. ESVA's scale-out architecture offers an ideal solution to these problems.

With ESVA, storage scaling is achieved by "adding." As shown in the figure below, expanding a storage pool is as simple as connecting a new system and adding it to the pool. Thanks to the technology of distributed load balancing illustrated above, by the simple act of adding, performance can be linearly increased along with capacity.

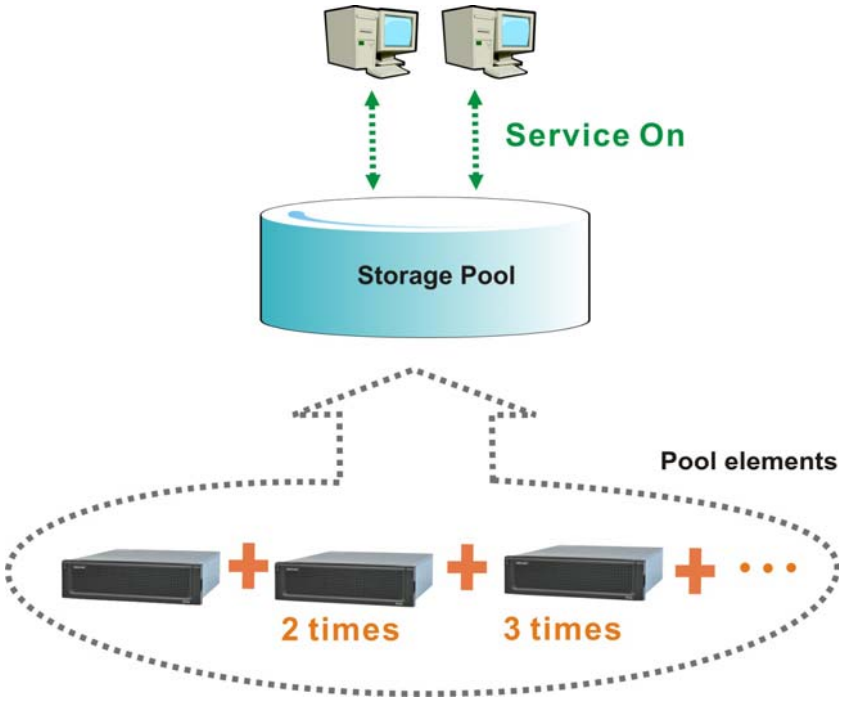

Figure 8. Scaling for Linearly Growing Capacity and Performance

Adding one more system means two times better performance while adding two more means three times better. The flexible and incremental way of scaling ESVA enables not only minimizes wasted investment on what users do not currently need but protects current investment and offers excellent price/performance.

Conclusion

Making every penny of investment drive the most business profits is a shared goal and also a challenging task of all companies. Achieving good returns on investment (ROI) depends on efficient utilization of resources. ESVA's advanced technologies, including

storage virtualization, thin provisioning, distributed load balancing and prioritized volume accessibility, ensures that the consolidated capacity and performance are amply allocated to meet application requirements without causing wastes. When more resources are required, storage can also be easily scaled while protecting current investment and offering ideal price/performance. During the hard economic time, ESVA helps users satisfy data demands with the least possible storage investment.