Plenty of alternatives, plenty of confusion about storage for VMS

By Justin Schorn

Recently there has been a lot of confusion about what storage technology to use on multi-server video projects. Small, single server deployments are straightforward because local storage is normally a reasonable option and easy to deploy. The storage strategy for larger projects is more complex and dictated by many factors: the size of the project, the scope and scale of the existing infrastructure, and of course, the IT budget. The critical decision is choosing between a storage area network (SAN) and network attached storage (NAS). The factors affecting that decision have changed, as both SAN and NAS technology have evolved and virtualization has become a major factor. Network infrastructure may also be a decisive factor; especially if Fibre Channel fabric is already in place or if 10 Gig Ethernet is already in the plan.

Let's begin by clarifying the difference between a SAN and NAS. Understanding the differences and similarities will go a long way towards understanding where each is useful and appropriate.

A SAN is a block storage device. Any device that exposes its storage externally as a block device falls into this category such as an external hard drive or Direct Attach Storage (DAS). We call it an external hard drive when we attach it to a desktop. We call it a DAS when we attach it directly to a server. We call it a SAN when we add some form of networking, generally a switch and a fibre-based cabling system, between the device and the server that is consuming the storage. Common protocols for communicating with block storage include iSCSI and Fibre Channel. In the end, a computer attaching to a block storage device will always see the storage presented as a disk drive.

A NAS is a file storage device. This means that it exposes its storage as a network filesystem. So any computer attaching to this storage does not see a disk drive but instead sees a filesystem. Users and servers attach to the NAS primarily using TCP/IP over Ethernet, and the NAS has its own IP address. Common protocols for communicating with file storage devices include NFS, SMB / CIFS and AFP.

What separates block storage and file storage is the type of interface that they present to the outside world. Both types have the option of providing extended features beneath the "demarcation point" before they hand off the storage to the outside. Both may (or may not) provide RAID, logical volume management, monitoring, etc.

While both potentially can sit on a network and allow multiple computers to attach to them, only a file storage device has the ability to arbitrate that access. This is very important and cannot be glossed over. Block storage appears as a disk drive. Only one server can write or change the data at once. A file storage device, on the other hand, has natural arbitration as the NAS itself handles the communications for access to the file system.

This distinction is what makes a NAS attractive to VMS. Conventionally, a SAN was used to provide the needed performance on large scale VMS projects. With a SAN using fibre channel instead of network protocols for data transfer (required by a NAS), it reduces latency and improves performance. However using block storage adds complexity since there is a need to statically map each server to a defined RAID target, or LUN in SCSI terms. Thus, the larger the system, the more mappings will exist. In times when outages are experienced (such as a VMS server failing or network disruption), the need to statically map the LUNS from one VMS server to another occurs. This results in the loss of access to recorded video while the LUNS are re-mapped to the new server.

Conversely, using a NAS storage element, no re-mapping of video storage is necessary since it can provide a single volume that all VMS servers can simultaneously access on demand. Think of this as a big C: drive but accessible via an IP address. Some NAS providers have this single volume functionality across multiple nodes, which is referred to in the industry as Scale-Out NAS. This flexibility makes it easy to leverage advanced VMS features such as automatic server failover, load balancing and scalability (since any server can get access to any recorded video). A Scale-Out NAS system can provide access to the same data from any node and, thus, can scale linearly in capacity and performance. For example, an Isilon NAS cluster can provide up to 40PB in a single volume, making failover and automatic load balancing easy. All servers simply record to a single name space and no re-mappings of storage location per camera is ever required. Though data transfer rates may still be slightly higher for SANs, the gap is closing and the overall value proposition of a NAS is stronger than a SAN in almost all scenarios.

If a single enterprise class NAS is out of your budget, using a group of smaller mid-market NAS devices can be considered such as Iomega® StorCenter™ Network Storage devices. Aimetis Symphony 6.12 offers the ability to automatically record video redundantly to more than one NAS or use recording failover functionality to record to a redundant NAS. In this case, each NAS operates independently but any server can gain access to any recorded video regardless of which NAS was used, since storage locations are nothing more than network shares. While a NAS can be used with all Symphony redundancy configurations, a SAN can still be used if each LUN is only being accessed by one server in the farm.

To be fair, there are ways of sharing storage from a block device, but this is not handled by the block storage device itself. Block storage devices are made "shareable" by using a clustered file system. This setup can be complex, making the attractiveness of this approach less appealing as the technology improvements in both Ethernet and file storage devices have closed the performance gap dramatically.

There are situations when block storage is the best choice. In single server deployments such as banking or retail, virtual machines (VM) are becoming increasingly common, with the VMS server running as a VM to reduce the number of physical servers IT needs to manage. Server virtualization offers advanced features such as high availability, or the ability to move a running VM from one host to another. Block storage is very common for virtual servers and works well, but with all server virtualization products now supporting the NFS protocol, NAS devices can provide a worthy, cost-effective alternative to fiber channel SANs for shared storage between VMs.

Finally, it is common today for storage devices to include both block storage and file storage from the same device. The EMC VNX Series is an example of a unified storage system. With unified storage, how the device is configured determines whether it behaves as block storage (SAN) or file storage (NAS). This drives home the point that this is purely a protocol or interface distinction, not one of size, capability, reliability, performance and features. Ultimately, both block or file storage can be used on your video project, but do not prematurely overlook file storage as the value proposition has significantly improved for projects of any size.

Justin Schorn is the Vice President, Product Management for Aimetis, Waterloo, ON., and an expert in IP technology and video management software.

aimetis