Differentiating A Video Server: What Sets Omneon Apart

Whitepaper

OMNEON VIDEONETWORKS
www.omneon.com

Differentiating a Media Server: What Sets Omneon Apart

By Paul Turner
Vice President, Product Management
Omneon Video Networks

The use of video servers in multichannel television operations has accelerated dramatically in recent years. The move to file-based storage and management of media has increased accordingly. Facility architectures based on video servers, along with file-based media management, offer a number of clear advantages over traditional tape-based operations. Those advantages include improved workflow efficiencies, decreased operating expenses, and more programming options. In fact, most, if not all, new facilities being built today will be all digital and tapeless, and many existing facilities, rushing to keep up, are in the process of transitioning their current operations in the same way.

Many server designs can support the playout of multiple channels of video, and most manufacturers understand the standard issues of storing and playing that video. The basic flaw in server designs has been a concept flaw that comes from clinging too closely to traditional broadcast technologies. The problem is that the designers are still thinking just in terms of video and not in terms of digital files. The fact is, with digital, broadcast has crossed over into the realm of information technology (IT) and we are at a point where products need to address and take advantage of these two traditions. Most servers today have remained focused nearly exclusively on the video tradition and have ignored many of the advances that have been developed as best practices of IT for managing storage and networking infrastructures.

Taking advantage of technologies and techniques perfected in the IT world, Omneon has designed a media server system that is extremely robust, fault tolerant (resilient), and easy to repair or expand. At the same time, an Omneon system enables LAN and WAN connectivity to extend the capabilities of a server beyond those of simple playback. It's not just about playing video; it's about building a platform capable of becoming the central shared storage component of a complete facility-wide, network-based solution.

Reliability/Ease of Maintenance – Direct Cost of Ownership Benefits

The importance of establishing a server infrastructure with 24x7 reliability is obvious to anyone in the broadcast industry. Having the system go down is not an option. Redundant systems are mandatory, including power supplies, fans, and RAID protected storage. And while redundancy is built into the costs of a system initially, there is an additional hidden cost of ownership which must be considered: the cost of repair, particularly the opportunity cost.

Components do fail, and the point of redundancy at various levels is to provide the system with the ability to survive such circumstances. Given that components will fail, it is important to examine the issues involved in replacing them.

Consider the case where one channel output in a system fails.Usually, systems are designed with some number of "spare" outputs to ensure that in the case of an I/O failure, the program material can simply be re-routed to another output, maintaining uninterrupted operations. Some time later, the failed I/O must be replaced. Many systems fall short at this point. Because most servers are built around the philosophy of incorporating the I/O into the main chassis, the I/Os are rendered non-“hot-swappable.” In some cases, the main chassis must be removed entirely from the rack in order to swap out an I/O. Since chassis-based I/O boards are not hot swappable, at the very least a reset of the main chassis is required. So, now, in order to replace the failed I/O, playback on all of the other channels must be shut down in order to re-initialize the replaced I/O. In some cases, the system power may need to be cycled, as add-in boards cannot be removed or inserted with the power on.

Let’s look at the same scenario on an Omneon media server. The I/Os are self contained entities, and are indeed hot swappable, even having their own self contained power supplies. With an Omneon server, the failed I/O can simply be removed and a new one put in its place. All other channels remain in place and fully functional. On power up of the new I/O device, nothing else is disturbed (the rest of the system remains on, even on-air, throughout the procedure), the channel is automatically recognized and initialized, and video starts to play out of that channel immediately. There is no disturbance whatsoever to any of the other channels in the system.

A similar case exists for the file system controllers within an Omneon system. In the event of a failure, the user simply takes out the failed controller and puts in a new controller. On power up, the controller interrogates the disk subsystem, is given the information about the existing file system, and is autoconfigured and operational in a matter of seconds. All other server vendors, without exception, use embedded file system controllers, with the same issue as detailed in the I/O discussion. Recognizing the inherent limitation in file system controllers, companies have looked at other techniques, one of which is the concept of “RAIDing” data across multiple file controllers, much as you RAID data across disks. In this way, the failure of a controller does leave the system operational — but it introduces new flaws into the system. With this approach, each controller is responsible for its own set of drives. Therefore, when a controller fails, the associated drives are taken out of the file system. Now, any recordings that are made while the controller is down are physically striped across fewer drive sets. Upon repair of the failed controller, its associated drives are returned to the “pool.” Unfortunately, now the entire file system has to be rebuilt, as some files are now striped across fewer drives than others. This process takes so long that one manufacturer goes so far as to recommend that customers do not record files while a controller (or I/O, given the above discussion) is out of operation. If a customer’s operation is a round-the-clock operation with continual ingest, this approach is clearly not acceptable. The Omneon architecture maintains the full file system at all times, negating this issue. Once again, the Omneon architecture is the more sophisticated, resulting in reduced maintenance considerations.

Another important consideration with regard to reliability is the redundancy of the server’s operating system (OS). An OS that resides on a single drive is clear cause for concern. Drives do fail, and the loss of the OS will result in a complete server failure. Many manufacturers, therefore, go to significant lengths to offer redundancy of the main system OS. Omneon takes a different approach: there is no system disk in an Omneon server — the OS is held in non-volatile memory on the file system controller. This has two benefits — (a) no possibility of mechanical failure, and (b) much faster boot time. The only other potential point of failure would then be corruption of the NV RAM. And for that possibility, Omneon systems keep a copy of the OS in a special location on the RAID protected, high-speed disk system. In the event of NV RAM corruption, the controller simply goes out to the drives and reloads the system.

Of course, all of this redundancy is for naught unless the user is made aware that a fault condition exists. A power supply can fail, and the system will most likely stay up, running on a back-up power supply. But this back-up power supply system has no redundancy, and will not be protected until the faulty supply is replaced. It is important, then, that an automatic scheme should notify the engineering department that maintenance is required. SNMP is often used, and some servers, such as those from Omneon, go so far as to send an email or page an engineer to inform them when the system is running in an unprotected state.

Customer support

Many issues are relatively easy to diagnose (at least to the component level), but there are some that present symptoms that cannot be interpreted by on-site customer maintenance staff. In these cases, the availability of support from the manufacturer is of the utmost importance. Omneon customer support goes above and beyond, providing not only fast response but proactive remote system diagnostics as well. The company maintains a NOC (network operations center) in its US and EAME offices from which trained service personnel can remotely monitor a system, diagnose problems via a network connection, recommend corrective action, and if necessary issue replacement parts or update software. The remote diagnostic capabilities are very comprehensive, allowing support personnel to diagnose all the way down to LIP (loop insertion protocol) issues in a Fibre Channel cable.

Omneon's commitment to customer service doesn't stop there. Operators in the NOC can also offer pre-emptive service capabilities, monitoring the operational status of a customer's system and warning onsite staff of an impending component failure before it happens. An example of such a situation would be a disk drive that is developing a significant number of bad sectors, indicating that it is likely to fail soon. Omneon's remote monitoring service can identify the questionable drive, notify the customer, and then onsite staff can replace the component in advance of the actual failure.

Sophisticated File System

As mentioned earlier, protection of storage is critical to the system's reliability. Manufacturers have many options when choosing a RAID scheme, and each solution offers unique benefits, but each also will uniquely affect the storage capacity that is actually available for media. In most cases, manufacturers tend to choose RAID3 or RAID5 protection schemes. Both technologies employ the concept of data drives plus parity drives for redundancy purposes. The difference between these two schemes is simply whether there is a dedicated drive for parity, or parity is striped across the disks along with the data.

In light of the fact that the broadcast industry is moving away from a tape based infrastructure to file based storage and management, Omneon recognized that simply using a fixed size RAID stripe for storage was not the best solution technically nor the most cost effective. Instead, Omneon designed an advanced file system to manage the storage, dynamically matching file sizes with RAID topologies. This maximizes both performance and storage capacity, regardless of physical file size. An example helps to illustrate the point.

Most playout servers use a RAID stripe that is configured for storage of large files. The general rule is to make the smallest “chunk” that can be written to disk as large as possible. By doing so, you can read more data per read operation, simplifying buffering issues further down the line. This is reasonable, as media files are generally very large. A consequence of this decision is that small files can take up substantially more storage space than the size of the file would indicate. The crux of the issue here is that given a large “chunk” size, a small file is “striped” across many disks, but only a fraction of each chunk is actually filled with real data. The rest remains unused, and the resulting file takes up much more space than is necessary. This issue is illustrated in figure 1. The chunk size is set by the file system for the most efficient storage of large files, and is regularly 64 KB or larger. If we assume a chunk size of 64 KB, and a small file of say 64 KB, then in the RAID system shown, the original file would be split into 4 chunks (i.e. 16 KB of data), plus another 16 KB chunk of redundancy data. The system can only write in chunks of 64 KB to each drive, so it “zero fills” each chunk of the file to make it 64 KB, and then writes the chunks to each disk. So the original 64 KB file now takes up 320 KB of storage.

An Omneon file system dynamically allocates the most efficient stripe methodology in order to use a more appropriate striping system for each file size, thus maximizing return on storage investment. Intelligent use of the correct RAID technology is the only way of maximizing storage efficiency for varying file sizes. (By the way, there’s absolutely no reason why you can’t have more than one RAID striping configuration on an array – you just have to manage them correctly.) This issue is of increasing importance as servers start to embrace the idea of standard networking topologies, with the resulting increase in the number of smaller files stored on the server.

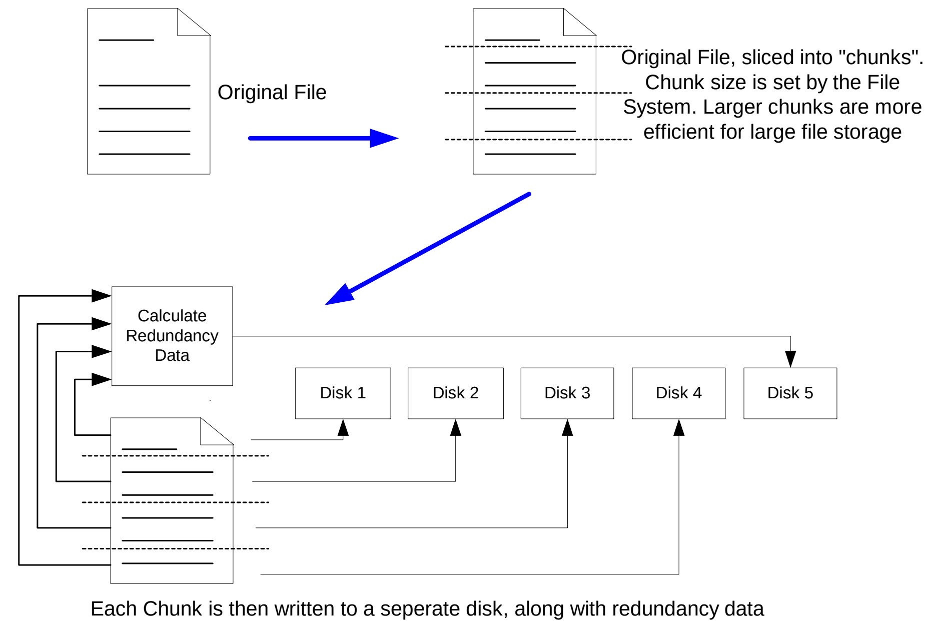

Fig 1. Typical RAID3 "workflow"

In addition to the sophisticated control of the disk subsystem, it should be noted that all Omneon servers support industry standard access protocols (such as CIFS) providing built in SAMBA capability for clients to view content within the file system. Simply put, this means that an Omneon server will appear on the network as a mountable drive. Thus NLE systems, graphic systems – in fact any system on the network – can connect directly to the Omneon system and mount it as external NAS storage without the need for 3rd party gateways. This dramatically reduces overall system cost, while maximizing flexibility of operations. For the Macintosh community, Omneon servers can also be mounted under AFP 3.0 (Apple File Protocol) simultaneously with the SAMBA mount. Finally, as a "lowest common denominator," Omneon servers offer full FTP capabilities, acting as an FTP server in addition to their other duties.

Clearly, the advantages of a file system which has been designed to offer all of the speed and flexibility of access demanded by broadcast applications while simultaneously offering full Gigibit Ethernet based network connectivity, allow Omneon servers to function in a multitude of roles, all at the same time.

External Control

Another area in which Omneon servers demonstrate architectural leadership is in the control system. In and of itself, a playout server can be considered to be a "bit bucket" – video goes in, and some time later, video comes out. External devices control both the ingest and playout of media, and this requires the use of both standard and expanded control protocols.

The lowest common denominators in the world of control protocols are the ubiquitous VDCP (Video Disk Control Protocol), originally developed by Louth (now Harris), and the Sony BVW protocol. In almost all cases, servers respond to some subset of the full gamut of commands offered by these protocols, usually limited by the technical capabilities of the server. Access to server control protocols, is either via RS-422 (point-to-point) connectivity or over the network via Ethernet. Often referred to as a control API (Application Programming Interface), it is here that a server manufacturer will expose the most intimate levels of control access to the system, and it is here that the maximum functionality can be eked out of the server.

While Omneon servers do respond to the above control protocols, they also offer a sophisticated network based control API. Several manufacturers of automation systems use this API to control Omneon servers, as it offers features not available via the simplistic automation protocols detailed above. For example, VDCP offers “look ahead” of only 1 clip – in other words, the automation system can only pre-load the server one clip ahead. This can be very inefficient, particularly when a number of small clips must be played “back to back,” as the control protocol may now limit the minimum clip length which may be played. Omneon servers have an internal “timeline” offering the ability to pre-cache a significant number of clips. The Omneon control API exposes this timeline to the automation system, removing the bottleneck of an old control mechanism.
The use of the Omneon control API fundamentally changes the level of sophistication and flexibility available in controlling a server, enhancing the capabilities of the overall system.

System Expandability

Customers will usually want to specify a server of sufficient size to allow for expected future needs, not just the immediate requirements. This is not always possible, and estimates of future needs are just that – estimates. Thus ease of post-install expansion is an important issue. The Omneon architecture allows seamless incremental expansion, all the way from a one or two channel system to a 150 channel server with significant amounts of IP bandwidth, all without the need to duplicate material on multiple servers (a frequent scenario in other servers), and without the associated media management nightmare.

The modular nature of the Omneon architecture is the key to its expansion flexibility. The architecture’s benefits were detailed earlier in regard to the ease of maintenance of the system, but the same advantages come into play when we consider expandability. For example, if you need another 3 output channels, you simply add them to the existing system while it is running. If you need extra bandwidth, simply add another file system controller. All Omneon servers are naturally SAN capable, requiring no external components to enable multiple servers to access the same material in the existing storage subsystem.

As technology and the market changes, the preferred video format and storage compression technologies may also change. Recall the early days of playout servers, when video I/O was often baseband video. Today it is becoming more and more common in a number of applications to use DVB/ASI as the I/O format of choice. Once again, in an Omneon system, you simply add a DVB/ASI MediaPort (the I/O module in Omneon systems). The file system is already configured to store “video as data,” so no other changes are required. As new formats emerge in the marketplace, Omneon adds format capability through the addition of new I/O devices to its product line. Each of these new I/O devices is plug-and-play within any existing Omneon system, allowing customers to easily add new format support without any disruptive system upgrade.

In comparison, for other server vendors, adding channels of I/O is definitely not something that can happen while the system is running, and will sooner or later require the addition of another server chassis, or even an entirely separate system. Once a system boundary is crossed, the customer has no choice but to duplicate material, as their systems can only communicate via a gateway, removing the ability to easily share material without the addition of dedicated SAN access technology. The ability of an Omneon system to simply add components means that it is no longer necessary to bear the cost of a large architecture from day one in anticipation of expansion. I/Os can be added as needed, without the need for extra controllers or storage. Additional bandwidth can be added without the need for additional I/O or storage. Extra redundancy can be added, without the need to add storage.

The final piece in expanding a system is increasing its storage capacity. Yet again, Omneon’s sophisticated file system comes to the fore, enabling intelligent scaling of the storage subsystem, without the need for re-striping. Simply connect additional storage to the system, and it will be used as soon as the initial storage is filled up.

The term that Omneon uses for this expansion capability is “smart scalability” — because the systems can be independently expanded in terms of I/O, bandwidth, redundancy or storage, intelligently and without penalty or disruption.

From this discussion, it is clear that Omneon servers have a number of significant competitive advantages based on their superior architecture, reliability, ease of expansion, and networking capabilities. Designed for the broadcast market, but embracing IT standards and best practices, Omneon servers are the choice for those building for the future.