How to Buy a Video Server

Introduction

The adoption of video servers in television operations worldwide is a well established fact. That adoption has also accelerated the industry's move to file-based storage and management of media. Video servers are the core of facility architectures and enable file-based media management, which offers media companies the potential for significant workflow improvements, including reduced operating expenses and expanded programming options.

Many server designs can support the playout of multiple channels of video, and most manufacturers understand the standard issues of storing and playing out that video. File-based workflows add another dimension to the traditional video server, which now must also embrace IT (information technology) best practices for the management of storage and networking infrastructures along with the interchange of files between systems. With all of these changes in the infrastructure, broadcasters often struggle when trying to weigh the various attributes of the video servers available to them.

As they try to choose the best platform to suit their needs, broadcasters should remember that their decision is not simply about recording and playing out video; it's about choosing a platform that is both robust enough and sophisticated enough to serve as the central media storage component of a complete, facility-wide, network-based solution, while also ensuring that this solution can grow and adapt cost-effectively as market requirements change.

In addressing the purchase of a video server, this article will focus both on the main attributes a buyer should consider when selecting a system and on the questions that every customer should ask of a prospective media server vendor.

Reliability

The importance of establishing a server infrastructure with 24/7 reliability is obvious to anyone in the broadcast industry. Downtime as a result of server failure is not an option. Redundant systems, including power supplies, fans, and RAID-protected storage, are mandatory. While redundancy is built into the initial system costs, the broadcaster should factor in the additional hidden costs of ownership, including system repair and, particularly, the opportunity cost of a failed component.

RAID Storage

Built-in redundancy is vital to every mission-critical storage system, and this redundancy most typically takes the form of RAID protection. RAID stands for "redundant array of independent disks." In RAID-based systems, information is stored on multiple disks, with additional data being written to yet more disks, called "parity disks," in order to offer error correction. The exact method by which this is done is beyond the scope of this article, but the basic operation is simple. If a drive fails, RAID storage uses remaining data in conjunction with data on the parity disks to reconstruct the data that was on that failed drive. This technique means that the server can continue to play back video even in the event of a complete drive failure. The most common RAID schemes deployed in modern video servers are RAID3 and RAID5. (The distinction between these two schemes is also beyond the scope of this article, but data on them is readily available on the Internet).

Key RAID Capabilities

Broadcasters should, of course, validate that any server they are considering does indeed have RAID protection for the storage, but it is just as important that they validate operation with regard to alarming and rebuild time.

RAID generally protects against a single drive failure in any array, which means that a broadcaster can't leave a failed drive in a RAID system for any length of time. Until the failed drive in the array is replaced, the server system lacks redundancy. Thus, if another drive in that array fails, the data on the array is lost permanently. As a result, it is extremely important that the server let the broadcaster know immediately that a drive has failed, so that corrective action can be taken promptly.

Many servers will provide a visual warning via their user interface (UI) indicating that the array has experienced a disk failure. In many cases, this model is insufficient, as users do not look at the server's UI with any frequency, concentrating instead on the UI of the control or automation system that drives the server. Broadcasters therefore should demand that any server they evaluate have a broader method of alarming error conditions. Sophisticated media server systems will provide alarms via e-mail or SNMP, both of which provide proactive announcements of a failure to the operations and engineering staff so that corrective measures such as drive replacement can be initiated. This is vital to reliable operation of the system.

RAID systems, when properly administered, can "self heal" by writing data onto the new (and empty) drive that replaces the failed unit. Through its error correcting mechanism, the RAID system can write data that exactly matches the data that once resided on the failed drive. This process is neither instantaneous, nor without impact. Broadcasters should look for servers that offer the shortest rebuild time on a failed disk, remembering that an array remains compromised until the failed drive is rebuilt and that another drive failure during the rebuild could render the data on the array useless.

Having a "hot spare" in the array is a convenient way to ensure that the array always has a spare drive ready in case one of the data drives fails. In this scenario, the RAID system will instantly (and automatically) start the rebuild process, putting the recovered data on the hot spare drive. This approach reduces the period of vulnerability, as the system doesn't have to wait on engineering staff to replace the failed drive before it can start the rebuild.

The broadcaster also should ask prospective vendors about the impact of a rebuild on the operation of the server — whether the rebuild operation will use up some of the bandwidth of the storage system, and if this reduction in available bandwidth will limit the normal operation of the system. A reduction in available IP bandwidth or in the number of I/O channels is an example of the compromises often required by less well-engineered servers.

Ease of Maintenance

In real-world operation, components do fail. Acknowledging this, broadcasters should be aware that the speed at which functional components can be replaced has a direct impact on the operational cost of the server. Critical to the analysis of any server, then, is an examination of how easily major components can be replaced when they fail. The impact of component failure and replacement on the rest of the server also is an important consideration.

If the entire server must be shut down during replacement of a failed I/O, the ongoing cost of ownership will include the lost opportunity cost of having to take down perfectly functional channels in order to fix the failed channel. The vendor should be able to provide information about how long a system shutdown would take, how the shutdown should be performed, and what issues can arise if the shutdown process is performed incorrectly. Downtime has a real cost, and it's one that must be considered along with the initial cost of the server. In many cases, broadcasters will opt to over-provision a playout server with more I/O channels than are required so that in the event of an I/O failure, the channel may be played out of one of the "spare" I/Os. While there is, of course, an up-front cost to this approach, the failed I/O will ultimately need to be replaced as described above, and potential shut down costs incurred.

Expandability

Virtually every broadcaster will eventually want to expand its media server. Most commonly, this is in the form of acquiring additional storage or attaching more I/Os to the existing system, and more recently, it has become common to wish to increase the available IP bandwidth of the system. Once again, the purchaser should examine the ability of a media server to expand, in what increments, and at what cost. Questions for the vendor should address whether additional channels come with additional storage, even if the broadcaster doesn't want to increase the storage at this point, and further, if adding storage will require that the server be taken offline for an extended period of time so that the system can rewrite the existing media to its newly expanded storage. These are real "cost of ownership" questions that should be asked and answered before the initial purchase decision is made. After all, the average operation life of a server is five to seven years, and expansion is almost guaranteed to occur during that period.

Flexibility

Media servers are no longer playback "islands" standing alone within the facility architecture. Today they interact significantly with other systems in the workflow and the efficiency of that workflow depends largely on the server's ability to integrate seamlessly with those other systems, particularly over IP connections. In addition to supplying ingest/playout functions, the media server also may act as storage for nonlinear editing (NLE) systems. Ideally, this will allow for "edit-in-place" operations, in which the NLEs do not have their own local storage (and therefore their own copies of the media, which must be managed), but instead pull material frame by frame from the server over Ethernet. For this model to be viable, the media server must be flexible enough to ingest, store, and playout material in the multiple formats these external processes may use. If the media server lacks this capacity, transcoding steps will be required, and these steps inevitably slow down the flow of material through the workflow, adding to costs in terms of initial capital and ongoing operations. A flexible media server will minimize or even eliminate such system costs.

Edit-in-place functionality also requires that the server be able to accommodate industry-standard wrappers such as QuickTime® and MXF. The use of proprietary wrappers for material interchange invariably escalates system costs as custom rewrapping and transcoding applications must be employed. If the media server supports standard wrappers and formats, then any rewraps or transcodes required by the workflow can be accommodated by non-proprietary products, again offering the greatest freedom of choice in designing the system or when modifying the system to accommodate future market changes.

Support

Media servers are mission critical devices. They sit at the center of the workflow, and they are expected to operate flawlessly in the 24/7 environment typified by broadcast television. In any complex workflow, issues are going to arise, and it is vital that a broadcaster be able to rely on support from its vendors. Without a doubt, the best approach in evaluating the quality of this support for any given vendor is to ask for contact information at existing deployments and contact the user directly to validate the level of support actually being provided. Another direct approach in determining quality and availability of service is simply to call the manufacturer's support line late on a weekend and see if a support engineer is available for immediate response. Both of these techniques will allow the broadcaster to observe the level of support it is likely to receive as a customer.

Conclusion

Many factors influence a purchase decision, and broadcasters ultimately must decide which factors and features are the most important to them in their environment. When looking to purchase a media server, there are some basic requirements the broadcaster must seek to satisfy above all others. The highest priority always must be reliability. The purchase of a media server with sophisticated built-in capabilities offers no benefit if the server isn't operational. The loss of even one I/O can have significant consequences if the server has not been designed to withstand such a failure. The highest priority within in this consideration is the resiliency of the data stored on the server, and the management of the storage (alarms, etc.) in all scenarios.

Given the presence of those characteristics, the ability of the server to grow (both in capacity and number of channels supported) in a cost-effective manner as needs change must weigh heavily in the analysis of cost of ownership over the life of the server. Support for a wide range of industry-standard formats and wrappers over both real-time and IP connections is of growing importance, and this capability will continue to grow in importance as broadcasters worldwide continue to adopt file-based workflows.

All of these facets of media server functionality are important, but equally important is the level of confidence the broadcaster has in the quality and consistency of support that will be provided by the vendor, no matter what the time, no matter what the day. At the end of the day, the successful broadcaster's business depends on a reliable 24/7, 365-day-a-year system.

More information on buying a video server can be found at www.omneon.com.