

WHITEPAPER

Considerations When Choosing A Media Server

Matt Allard, *Director of Marketing, Servers & Digital Production*

October 2011

A media server is a much more complex system than one might first think. When choosing such a system for a facility, there are a wide number of issues that should be carefully considered.

IT Optimization

Today's best media servers leverage the core technologies of the general IT industry, and then optimize and enhance those technologies specifically for the real-time, high-availability requirements of television broadcast and video production. The best suppliers build systems that use standard file systems, storage, and networking. These suppliers provide added value by optimizing and tuning IT technologies for media applications in ways IT vendors can't or won't. The implementation of the latest technology produces more cost-effective solutions.

A media server should incorporate a variety of the latest technologies. These include 8 Gb Fibre Channel backbones with 15/7.2K SAS drives in RAID-5 or RAID-6 topology. Storage components alone will not provide a satisfactory solution, they must also be integrated with a complete infrastructure. An IT server component must be included to act as a data bridge and to manage dynamic and deterministic bandwidth requirements. For optimum throughput, a 64-bit operating system and file system are desirable. A way to provide scalability and cost-effectiveness is to adapt the iSCSI protocol over 10 Gb/1 Gb Ethernet to provide real-time and deterministic performance for media movement.

Technologies and infrastructure should be combined to manage bandwidth. Bandwidth is a primary factor in determining how reliably a system will perform and what capabilities it can provide. A server system should be architected to manage the multiple levels of bandwidth that may be needed for deterministic operations. This can be provided by a shared file system that includes Quality of Service mechanisms.

Three levels of managed bandwidth should be provided:

- The first level with highest priority is for real-time media where clients are designed to never exceed stated bandwidths. This real-time performance is a service level not needed by most regular IT systems.
- At the other end of the spectrum is managed shared bandwidth for non real-time operations such as file transfers.

- In between these two levels there should also be a reserved bandwidth level for time-critical production activities, such as editing, where each individual client gets allocated bandwidth that it cannot exceed.

All three levels should be available simultaneously. On many systems there is only one level of bandwidth management and it cannot be relied on to perform all tasks reliably and consistently. A proper media server system has built-in redundancy, buffering, and multiple levels of processors for guaranteed throughput. The supplier of such a system characterizes the storage performance for various production uses, and measures bandwidth performance by individual LUNs, servers, and controllers.

Another unique aspect of media servers is file management. Unlike files found in most regular IT server systems, media server files are large... very large. The system, including storage controllers and file system, must be specially tuned for such files.

In a standard IT server system there are a large number of small, simple files with rapid reads/writes to and from storage. It is not necessary to read during the write process; delays and latencies can be tolerated.

In media, delays and latencies are not acceptable as they translate to black frames and audio dropouts. In media, working with large files means storage has many simultaneous reads during the write process.

For proper media file management, the server system should provide specialized buffering and pre-fetching, enhanced on-the-fly error correction, packaging of multi-track essence inside media file containers, and optimization of the application layer to read/write efficiently and deterministically. When assessing a server system for media production, this is an area that should be examined very critically. Many systems that claim they are suitable for media do not offer these optimizations, and consequently fall short of delivering the constant performance necessary for content delivery.

Platform Flexibility

A media server should be versatile so it can deliver different solutions suited to a variety of unique requirements. It needs to scale to any dimension of performance and redundancy. Expandability needs to scale along multiple, independent vectors. A system may need to increase media channels, storage capacity, file transfer bandwidth, or operational bandwidth for production needs such as off-speed play and editing. A server for broadcast and video production must be equally adept as both a media server and a file server.

Because there are no second tries when delivering media data, redundancy on multiple levels is required. The system must provide redundant data paths between media clients, media servers, and storage—a design that inherently has “no single point of failure.” In terms of implementation, this means specific practices, such as redundant RAID controllers, redundant network interface controllers connected via multiple servers, and multiple switches should be available. On a device level, there should be redundant power supplies, redundant cooling, and redundant storage drives.

For media interfacing, a standard PC with some processor cards is simply not reliable or robust enough. A purpose-built, appliance-type device for broadcast and production that incorporates selected IT technologies should be used to provide a highly available play and record service for 24/7 environments.

With this approach, a variety of benefits can be attained. The device can be ruggedized for use even in mobile environments. Modular components can be easily accessed and serviced. A dedicated device can utilize specialized embedded and real-time operating systems along with multi-level general purpose and customized processors to offer unique and complex features with reliable performance. Some of these features include multi-purpose channels that can manage audio/video recording, audio/video playout, recording of super slow-motion, recording and playout of combined left-eye/right-eye 3D, recording and playout of combined video/key, multiple compression formats, and resolution up/down/crossconversion.

All this can be augmented with the creation and management of low-resolution proxy media as well. At the same time, this type of device incorporates all the advantages of a standard platform for storage, interface displays, graphics, and network connectivity for simple integration and lower costs.

An integral part of any media server system is storage. Broadcast and production environments have diverse requirements, so a server system must offer a range of storage solutions.

The simplest implementation is serial attached SCSI (SAS) drives that are located internally to the media client. To simply increase capacity, external RAID chassis with serial attached SCSI drives can be connected to the media client through a host bus adapter.

For the highest performance and flexibility, a media server should support a Storage Area Network (SAN). A SAN is a dedicated storage network that provides access to consolidated storage. The networked storage devices are accessible to the servers such that they appear as if they are locally attached to the operating system, allowing direct read/write requests to the storage disks. The server component acts as a data bridge between the Ethernet and Fibre Channel connectivity. It also hosts a shared file system to provide a file abstraction that applications and networks can interact with. A SAN provides the scalability, connectivity, bandwidth, and storage to permit a range of production scenarios.

As a server platform provides the essential infrastructure for file-based production, there are a number of features that should be provided to make interoperability easy and efficient. Files should be stored as elementary streams on disk for direct editing and metadata operations. There should be import and export of media in parallel with other operations to networked storage, removable media, editors, and archive. Media files should support standard video compression formats such as MPEG-2, MPEG-4, DV, and AVC-Intra. These files should be available for sharing in common wrappers such as MXF and QuickTime. Other file types such as WAV for audio, QuickTime 32 for graphics, and MPEG transport/program stream for video should also be compatible. Files should be transferable using standard FTP and CIFS protocol, as well as via the direct iSCSI connections. To help automate the process of moving files, there should be watch folder services to bring files on and off the server platform.

As content owners want to re-purpose material, creating, saving, and exchanging metadata must be associated and managed with media files by the server system using an open standard such as XML.

Managing the server system and its capabilities, as well as keeping operating costs and total cost of ownership low should be provided for. To simplify how systems are used and maintained, there should be an integrated toolset and utilities to assist in deploying, maintaining, configuring, and upgrading the system.

Application Integration

A server system is only as useful and interesting as the applications that it can be used with. To have complete solutions, a server supplier must also enable applications that people need to use. There should be applications that permit people to use a server system for workflows in live production, news, entertainment, and playout. A server system supplier needs to create applications to enable innovative and affordable solutions. Some of these applications include ingest, replay, clip store, editing, channel control, and playout. Application services should be provided to offer proxy operations, common edit decision lists/playlists, common metadata keywords and markers, and content management. Having these applications available show that the supplier understands the details of how complete solutions need to work across the entire broadcast infrastructure.

User experience and feedback are what should drive application design. Applications should focus on reducing learning curves and making usage simple. Ease of use should also apply to upgrades and maintenance. All applications should have common user interfaces and terminology. They should share common file format support, proxy, metadata, and edit decision list/playlist information.

No media server supplier can produce all the applications that may be needed by production users. Therefore, it is necessary to have an API available so additional applications can be developed by third-party companies and even end users. The API should support common programming environments that do not require proprietary languages. A traditional capability for an API is for media channel control. With the transition to file-based production, the API must also incorporate content management capabilities. The API must provide direct access to media assets. There should be access through network services for system-wide assets with the ability to initiate file transfer through FTP and CIFS protocols. All new functionality should be openly supported through the API. To create sophisticated application frameworks, there should be an API that supports a service oriented architecture (SOA) abstraction layer to provide services for content management, proxy, ingest, playout, search, transfer, and metadata.

With a well executed API, a media server supplier should have a staff of development support engineers to foster a diverse community of active developers. They will help to enable the supplier to develop a broad array of supporting applications.

Testing and Validation

For a media server to provide all these capabilities and integrations, an extensive amount of research and development has to occur long before any solution is delivered. While careful software development and hardware design at the engineering level are expected, there are some other critical technical processes that need to be incorporated as well.

With the goal of having tight integration of a variety of standard IT components and sub-systems, rigorous validation has to be performed.

Some of the technology evaluation criteria that should be included as part of validation:

- Does a certain technology deliver on the promise of its features and specifications?
- Can the technology be effectively utilized for its intended application?
- Can the technology function well as a building block for the overall system?
- Will the development cycle to integrate the technology be predictable?
- Will the technology scale as expected and operate under various loads?
- What will be the operational behavior in error conditions?

Some of the evaluation criteria for storage validation should include:

- Benchmarking performance over multiple days during rebuilds, and while running various server input output configurations
- Price versus performance versus latency
- Bounded latency with augmented storage sub-systems that respond within a certain time limit
- Hot-swapping storage controllers under a full load, live firmware updates under a full load

Some of the evaluation criteria for the file system validation should include:

- With a shared SAN, how is access managed and controlled for concurrent access of real-time and non real-time applications?
- Automatic management of the file system profile so defragmentation is not required even under 24/7 operation

Some of the evaluation criteria for network data movement validation should include:

- Management of the network data stacks and fabric so that transfers are effectively lossless
- Characterization, configuration and management of device initiators and targets, network interface cards, and network switches so they operate as one clean lossless connection from end-to-end
- Analysis of network switch internal architecture, port setup, and microcode versions

Testing and Validation (cont.)

Once a design has been completed and software created, a multi-stage component/system validation and testing process must be completed. This process should be implemented on a long-term basis and be highly scalable. For on-going support and new releases, the functionality and behavior of each technology piece, and the overall system, must have extensive regression testing.

As technology advances, there will be subsequent generations of components. As they are incorporated, the system must be retested with new versions of storage drives, RAID controllers, file system, data servers, network interface cards, host bus adaptors, and switches to ensure the system behaves as specified.

To put the necessary effort into perspective here are some examples of what typically should be done with every generation of components to guarantee the performance of the system.

Storage RAID:

Six months with two engineers at the media server system provider and three engineers at the supplier to validate.

Three months for firmware release, regression testing, QA, configuration tool analysis.

Shared Storage Drive Series:

Four to six months with two engineers to develop firmware, validate, and test.

Part of the validation is for extended multiple-year serviceability so that storage can be expanded with future generations of drives and have the drives, drive groups, rebuilds all work within the expected performance levels.

Internal Drive Series:

Four to six months with two engineers to develop firmware, validate, and test

Performance analysis and validation, failure analysis, special mode page handling for media usage, integration of management and serviceability tools.

As with shared drives, not every drive family or drive supplier will meet the necessary standards.

Media Networking:

After years of initial development, nine months with three engineers to develop the latest 10 Gigabit-based media networking tools, performance benchmarking, failures analysis, supplier qualifications, and integration.

Six months of release testing, scaling, and configuration tool analysis

Services and Support

After a media server has been selected, the interaction with the supplier needs to enter a new phase. Now the supplier must deploy service expertise to effectively design and implement a system that meets your business needs, and support that system efficiently over its useful life. Evaluating the supplier's capabilities in these areas is critical for realizing the benefits of the new technology, minimizing risk, and controlling total cost of ownership.

Effective system design is essential if you are to realize all of the benefits of media server technology. In this step, the supplier's system architects must engage in a discovery process to gather detailed technical requirements, and then translate those requirements into a system design. An experienced system architect will

ask basic questions about your requirements, such as preferred compression and bit rates, projected amounts of storage, and numbers of ingest and playout channels. A true consultative approach from best-in-class suppliers should also explore potential future expansion of the system, interface points with other systems in the enterprise, long-term data continuity planning, and total system fault-tolerance. These considerations should influence the design process as much as the simple bits, bytes, and channels data, which often take precedence outside of a consultative engagement.

If an effective system design is the vision, then the reality is built during on-site system implementation. Here, the media server supplier's team of project managers, field service engineers, and trainers

combine to commission new systems in real-world environments. The system must be configured and tested for optimum performance, and then the users must be trained on operation and maintenance procedures.

Every media server supplier should have these basic capabilities. For systems of greater complexity and size, such basic implementation capabilities must be orchestrated by a detailed project management methodology to control project cost, timing, and risk. In any supplier evaluation process, each supplier should be asked to explain their project management methodology, including the statement of work, project scheduling, supplier/vendor communication, system documentation, and issue tracking.

Services and Support (cont.)

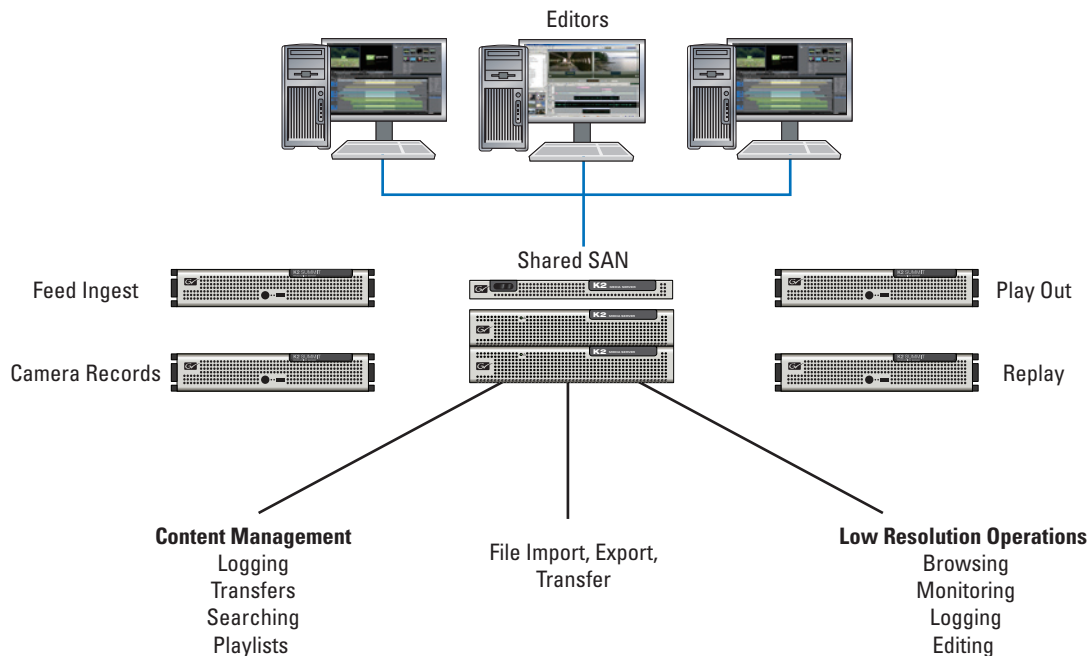
A properly designed and implemented media server system can provide years of operation as the heart of various media delivery infrastructures. However, all systems inevitably require technical support to troubleshoot failures, keep the system current with supported software releases, replace defective hardware components, and generally maintain system uptime. Potential suppliers must demonstrate their capabilities and investment in these core areas to ensure that capital investments are protected. Beyond core technical support capabilities, media server suppliers with advanced customer support infrastructure should be capable of providing high-availability services such as:

- 24/7 technical phone support
- Remote system diagnosis
- Access to continuing software releases and associated installation services
- Advance exchange hardware replacement with next business day delivery
- Field service engineering

Media server suppliers who can offer such high-availability services are able to dramatically minimize any system downtime while ensuring the profitable flow of media assets to the overall enterprise, on-air playout, and online operations. In addition, best-in-class customer support organizations are able to offer high-touch services such as dedicated technical account

management, 24/7 remote system monitoring, and even outsourced engineering support.

Often these services are offered in the context of a system support agreement with prescribed response time parameters. Not only do support agreements obligate the system supplier to a specific level of performance, they also definitively control cost of ownership. A full evaluation of potential media server suppliers should include a review of their customer support capabilities, high-availability services, and system support agreements.



GLOBAL SERVICES



Grass Valley Global Services specializes in the defining of, deployment of, and support of today's dynamic file-based workflows, based on Grass Valley and third-party solutions. With Grass Valley Global Services, you can achieve your operational goals in the most efficient and cost-effective way possible with a partner you can trust.

www.grassvalley.com/support

Define: We help you to define your business and technology requirements and then design solutions to meet them.

Deploy: Our professional service organization, backed up with proven project management methodologies, can take you from design through deployment, commissioning, and training.

Support: We offer a complete SLA portfolio to keep your systems running and help plan for your long-term maintenance needs.

Join the Conversation at **GrassValleyLive** on Facebook, Twitter, and YouTube.

