

FlashNet



content archive & storage management
for media & entertainment

Your content, where you need it

Content storage and life cycle management is becoming an increasingly important consideration for the modern broadcaster. Material goes through many stages during its life-cycle - ingest, edit, manipulation, playout and archive - and the requirement to have sequences and clips instantly available to all of these areas has become paramount, even after the files have been transferred to long-term storage. Making sure that data is in the correct location, whether it be online for playout or offline for archive, is vital to smooth, interruption-free operation.

SGL FlashNet fulfils this need. It tracks and manages your data throughout its storage life-cycle, working seamlessly with automation, asset management, NLE and video server solutions in order that data is always available when required. Using SGL FlashNet, a central storage repository for the entire installation becomes a reality. The archive becomes an extension of the main storage capacity, available to all departments transparently. Whether the archive consists of storage network devices, disk arrays or tape libraries (or a combination of these), it simply appears as an additional storage area to all applications.

SGL FlashNet takes control of archive and content management painlessly and without effort. When requested by the asset management or automation applications, FlashNet moves data to near-line disk or removable media (such as tape or optical disk). FlashNet's database tracks the location of all material using Unique Identifiers, and makes these available to the asset management or automation system. When a clip is required, the requesting application passes the UID it requires to FlashNet, and the clip or sequence is immediately restored to online storage.

SGL FlashNet also supports Partial File Restore for a number of formats. PFR allows the user to request only a certain portion of an archived clip to be restored from the archive, rather than having to restore an entire clip and then edit down to the required start and end times. The 'new' clip is allocated its own UID and is treated as a separate asset in the database.

Resilience. Scalability. Power.

To develop the best possible archive management solution, SGL established detailed discussions and pilot schemes with major broadcast vendors and customers to ascertain the needs of the modern broadcast facility, in particular the requirements of the broadcast archive. The result of these discussions is a new architecture for SGL's content management software, which moves away from the more traditional single server model to provide broadcasters with greater flexibility, expansion, capacity and power in the archive. The clustered architecture at the heart of FlashNet has been designed with these needs in mind. A Microsoft SQL database drives an infinitely-scalable cluster of servers, transferring data into and out of the archive at the highest possible speeds. Each server runs identical software, which means each can perform any task, and each is connected via a switch to every archive device (e.g. tape drive or disk archive) so that the archive always operates at maximum capacity.

This evolution means that FlashNet can operate across a cluster of any size. For the smallest operations, all three FlashNet modules can be installed on a single physical server. Larger organisations, which typically require

How it all works

SGL FlashNet's architecture is based around clustering. Each installation comprises 3 base modules: the database, the resource allocation module and I/O modules. These modules are installed on physical servers, known as nodes. Each node is identical, running the same software as all others, and each node runs the I/O module. This means that every node in the FlashNet cluster is attached to the storage devices (disk/tape library etc) and is capable of performing data transfers. FlashNet's highly scalable architecture means that a cluster can comprise a single physical machine running all modules, or many nodes for high levels of resilience and data throughput. A clustered architecture with a central database gives SGL FlashNet unrivalled resilience. Component failover is automatic*; when a module fails another node in the cluster automatically takes over operation of that module.

The database

The SQL database is the brain of the FlashNet cluster. It holds all information about the archive, including the location of material, the job queue and even the health and status of cluster components. As each job is sent from the requesting application it is passed to the Job Queue table in the database. The Resource

Allocation Module (which may be running on a different node) analyses the Job Queue table and allocates each job to a node for I/O transfer. The database keeps a constant picture of the entire archive, and load balances job distribution based on the status of nodes in the cluster. For high availability and redundancy considerations, the database server can be installed across a Microsoft® cluster which provides automatic failover in the event that the primary database server fails.

The Resource Allocation Module

The Resource Allocation Module (RAM) runs as a service on all nodes. Although present on all nodes, it is only ever active on one node at a time. The RAM analyses the Job Queue in the FlashNet database every five seconds and distributes the jobs in the queue to nodes that currently have I/O resources available. Clusters containing two or more nodes provide high RAM availability; in the event that the currently active RAM fails, the RAM on another node in the cluster becomes active and begins processing the job queue. This failover is instantaneous, and requires no user intervention.

I/O modules

Each node in the cluster is capable of performing data transfers. This functionality is known as the node's I/O module. All data transfer into and out of the archive takes place via a FlashNet cluster node. The amount of nodes in the cluster defines the throughput capacity; in simple terms, more nodes equate to greater transfer rates. Each node is connected to the archive devices (tape library, optical library and/or disk cache) via one or more fibre channel switches. This allows all nodes to communicate with all drives and devices. This architecture means that any node can initiate transfers to any free archive device; individual nodes are not connected to individual devices. The database also records the usage of all components in the archive, (cont.) the archive

FlashNet



content archive & storage management
for media & entertainment

and uses this data to automatically load balance across the cluster, right down to the level of individual HBA ports to maintain component health to the highest possible level. The automatic failover capabilities in FlashNet also enable active upgrade and maintenance procedures. When cluster software has to be updated, or maintenance is scheduled, individual servers can be shut down safely while the rest of the archive remains operational.

Scalability

A FlashNet cluster is scalable to a virtually infinite degree, increasing the speed and capacity of data throughput as the number of nodes in the cluster grows. Smaller operations may require a lower-cost archive, consisting of, for example, a FlashNet cluster installed on a single physical server; this provides a compact yet highly efficient archive system, whereby all 3 modules run alongside each other on the same machine. Larger facilities traditionally require greater speed and redundancy in the archive; using identical FlashNet software a cluster of 20 or more nodes can be assembled, concurrently driving 40 high speed tape drives at speeds of up to 80MB/s each and writing to several terabytes of disk cache. Whatever the size of the initial cluster, as data transfer requirements increase the cluster grows with the simple addition of further nodes, each of which widens the access channel into and out of the archive. More nodes mean that your content gets where you need it faster.

Power

FlashNet's power comes from the nodes that drive the archive devices. Each node is configured identically throughout the cluster, typically installed with at least one dual-port fibre channel card; these connect via one or more fibre channel switches to the archive devices (e.g. tape drives). This configuration allows each port to communicate with every drive - no single fibre connection connects to a specific drive, which allows any available node to transfer data to or from any available drive. This keeps the archive operating at maximum speed at all times. Expanding the cluster by adding more nodes increases the amount of connections into the archive (as long as there are sufficient archive devices to be driven); more nodes mean higher data transfer speeds.

Resilience

High availability is a key theme of FlashNet's clustered architecture. All nodes are in constant communication with the central database, and each module on each node reports its status every few seconds with a 'heartbeat'. If a node fails, or is taken off-line for maintenance, its workload is automatically redistributed throughout the cluster, keeping the flow of data to and from the archive consistent and uninterrupted. In standard configurations the main database is also fully redundant, using the latest clustering and mirroring techniques and software to ensure that content is available 24x7.

In addition, SGL FlashNet's unique system of storing index information on the media to which data is written means that in the event of ultimate system failure the database can be rebuilt from the individual media.

The database also records the usage of all components in the archive, and uses this data to automatically load balance across the cluster, right down to the level of individual HBA ports to maintain component health to the highest possible level. The automatic failover capabilities in FlashNet also enable active

upgrade and maintenance procedures. When cluster software has to be updated, or maintenance is scheduled, individual servers can be shut down safely while the rest of the archive remains operational.

StorageManager

The optional StorageManager module takes content storage management a step further. This module automatically manages the data within the archive, moving and copying content across media types, storage areas and even geographical locations to provide the most complete data management system available to the broadcast industry. StorageManager also defragments the archive, automatically moving material and repurposing media to make the archive as cost-efficient as possible. StorageManager consists of two sub-modules, Defrag and LifeCycle. Each acts based on rules defined by the user. When data or volumes meet the criteria specified in the rules set, StorageManager instigates FlashNet restore and archive processes to move, copy or delete data as required, and to relabel any media that is eligible for defrag. The spare media is then added back into FlashNet's media pool, streamlining media usage. StorageManager works transparently to all automation and asset management systems. Data is always available regardless of its location within the archive.

LifeCycle

Storing clips based on their content is a highly effective means of keeping storage costs to a minimum, and ensures that local and remote storage is handled in the most efficient manner possible. LifeCycle automates the movement of material through different levels of storage in the archive, moving, copying and deleting clips as defined in user-defined LifeCycle rules. Material is automatically moved based on various attributes: for example, using LifeCycle material archived to nearline disk for fast-access storage can be automatically moved to longer-term tape archive once it has been inactive for a defined period, e.g. 30 days. Automatic data movement based on age and usage means that material is always stored on the most suitable medium and in the most appropriate location. Whilst movements are in progress the clip is always available to asset management and automation. There is no 'down time'.

Defrag

Defrag helps broadcasters cut expensive media costs by automatically scanning the archive for volumes from which the majority of material has been deleted. As users delete clips from FlashNet's UID database media becomes fragmented, containing a mixture of useful and deleted data. Defrag works on a set of user-defined rules that stipulate how much non-useful data a volume should contain before it is defragged. When a volume reaches this level, Defrag instigates FlashNet restore and archive processes to move the 'good' data to a disk cache, and from there to a different media group. When all 'good' data has been moved, the original volume no longer contains useful data and is re-formatted, to be used by FlashNet in standard operations. All Defrag operations take place automatically, and can be set to use only a certain amount of resources (e.g. tape drives) so as to impact on daily operations as little as possible. SGL FlashNet is compatible with all of the solutions provided by major vendors in the broadcast industry. It has been tested with all of the leading asset management, video server and automation systems to provide the greatest level of third-party compatibility and integration of any broadcast archive solution.

