



# Quantel White Paper

## **Genetic Engineering: Teamworking infrastructure for post and DI**



Edition 1.2 March 2009  
030 – 23 – 4096 – 45

## About this White Paper

All film post production, broadcast channel branding, TV commercials or programme making is a team effort. A group of people with different skills use a number of systems with different tools.

Fifteen years ago that involved film or videotape. That meant physically copying media at each stage of the process. It simply wasn't possible to share media, so there were many practical problems and bottlenecks when working as a team.

Today, file based technology has largely replaced tape and film during Post production – but we're still a long way yet from true teamworking

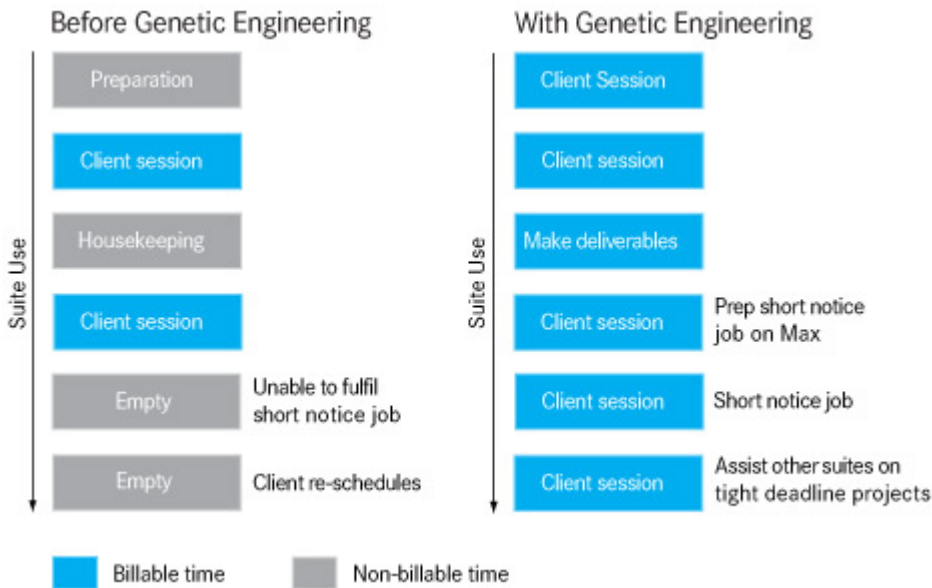
Genetic Engineering is an award-winning open teamworking infrastructure from Quantel which overcomes many current technical limitations and enables a truly collaborative Post and DI environment to be built without boosting management overheads.

This paper describes the capabilities of the new infrastructure and details how it increases efficiency in the multi-vendor real-world Post, DI or Broadcast Creative Services pipeline.

## Teamworking in Post and DI

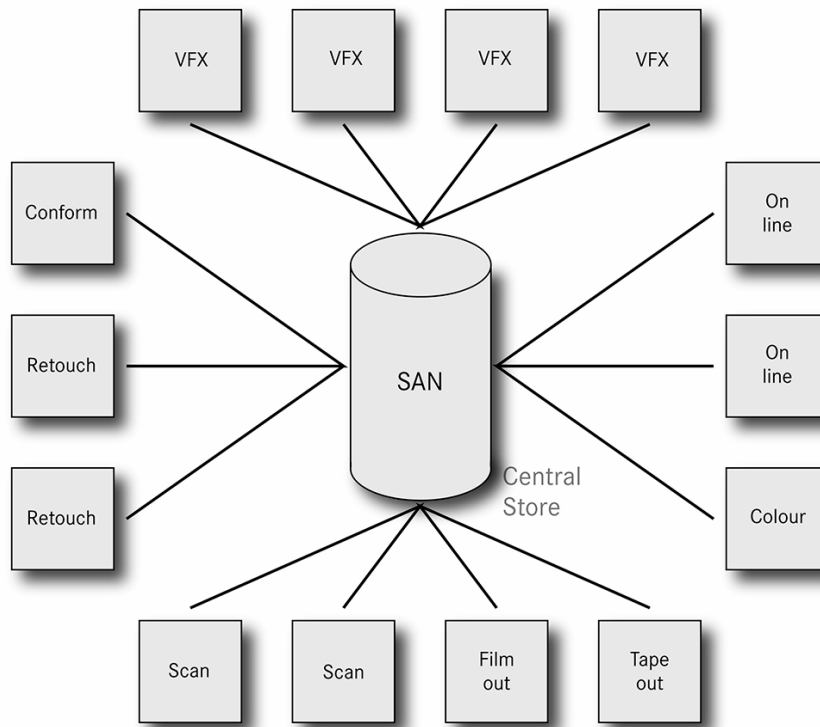
Teamworking is an essential part of a competitive post facility, broadcast creative services department or digital film facility today. Different people with different skills working on different equipment combine their talents to deliver great results with maximum creativity, on-time and in-budget for their clients. There are many processes involved, for example ingest, conform, VFX, grading, audio, versioning, and deliverables, which occur in a complex parallel and serial pipeline. Especially at high resolution, e.g. HD, 2K, 4K and Steroscopic 3D, moving media between applications is inefficient and costly in both time and disk space.

In practice this means that staff, equipment and rooms are not used to maximum efficiency. A typical room utilisation may look like the left hand side of this picture:



The 'billable' time in a suite is limited, with much unbillable time spent loading or unloading media to and from the machine. A short notice booking may not be practical there is insufficient time to load material. A client may suddenly reschedule a job and the suite is left with no booking. Meanwhile a different suite may be very busy and in need of help. So, what is needed is a 'joined up' workflow, shown on the right hand side, which allows best use to be made of equipment, rooms and staff.

Before Genetic Engineering, these multiple processes and inefficiencies drove many facilities to look at a way to share data between the different applications and people involved using a SAN.



A SAN model for teamworking

The idea of enabling the various applications to share a common pool of storage using a SAN is a good one. No time is wasted when transferring between applications and as a consequence there is considerable operational flexibility about what happens and when. However, while SANs bring many benefits, there are many real-world issues that arise:

- SAN performance may not be adequate to support enough real time clients without imposing unrealistic limits such as leaving 50% of disk space free for internal mechanisms to guarantee performance.
- SAN performance under fault conditions, e.g. rebuilding failed disks may be an automatic process but it can impact the number of real-time concurrent users
- Providing multiple streams of 2K and 4K remains a challenge in some systems. All the above make providing a mission critical shared environment almost impossible so local storage is often introduced where absolute reliability is required, i.e. client review, ingest and playout. Of course this adds inefficiency, increases costs and produces more copies of media to be managed.
- Applications may copy data to reorder it (edit it), this quickly leads to not only inefficient use of disk space but also adds a management overhead to manage all the copies of data.
- Different applications may support different file formats, resolutions, colour spaces or even bit-depths. These application-imposed limits mean media must be converted before it can be shared meaning new file conversion applications and more copies of data to be stored.
- Incompatible metadata between different applications and the inability to roundtrip dark metadata mean flexibility is lost when work moves between systems and external asset management needs to be introduced to track what happened where. This becomes a big overhead once the usual creative changes part way through the process start to occur.

These issues are well understood by manufacturers of high-end edit systems and are the reason why close-coupled direct attach storage is still widely used in practice with SAN based solutions.

*(Continued)*

Even where shared storage is offered it is closely policed storage that is not open to all in order that performance parameters can be guaranteed. This of course limits its usefulness and it tends to fit in the above diagram as just another lump of storage with media needing to move between it and other disks and applications in the facility with all the inevitable time, storage and management overheads.

### **An ideal shared storage solution**

An ideal sharing infrastructure for post and DI would suffer from none of the above drawbacks. An ideal infrastructure would:

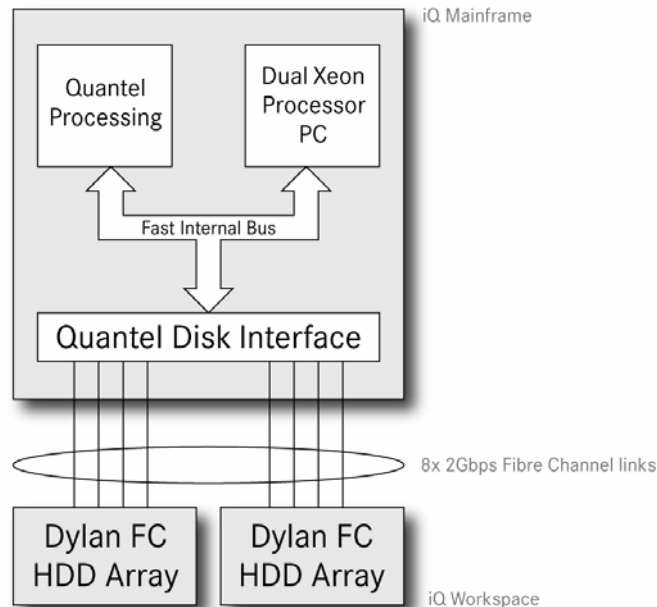
1. Provide deterministic performance for specified clients even at 4K when required without unnecessary overheads (no leave 50% free requirement)
2. Provide guaranteed performance under severe fault conditions
3. Never copy media for re-ordering  
guaranteed fast access allows live 'play' of edits
4. Provide high bandwidth open 'files and folders' access to media in the pool for clients not requiring deterministic access
5. Solve application file format, resolution, colour space and bit-depth issues without requiring multiple copies of media
6. Keep track of media relationships  
e.g. know which four layers in a composite relate through to the final flattened result.
7. Provide facilities for storing and managing both open and dark metadata
8. Scale to any extent

The above specification presents a tough technical challenge, yet even addressing only most of the above points would be a large step forward for teamworking in Post, Broadcast Creative Services and DI. At a stroke such an infrastructure would lower costs both directly in savings in disk space and indirectly in time and management effort. Genetic Engineering already addresses points one to six and seven and eight are future development targets.

Genetic Engineering can be seen either as a complementary technology to a SAN (especially in supporting an existing SAN) or as an alternative approach.

## Quantel Workspace

The foundation of Genetic Engineering can be found in the Dylan workspace used in today's Quantel Post and DI systems such as Pablo, iQ and eQ. This direct attach storage provides fault-tolerant, deterministic, high speed access to media, does not copy to re-order media and equally importantly includes powerful media management that tracks media relationships down to frame level.



iQ workspace high-bandwidth direct attach storage

**High Speed Access** The above system can support continual playout of real time 4K, i.e. over one GByte/second (GB/s) and also error-free dual 2K stereoscopic playback. The performance is maintained even with the disk over 90% full and irrespective of where the media is physically located on the disk media. The disk never requires defragmenting so all the bandwidth is available to the user at all times.

**Fault tolerance** In addition to dual power supplies the storage provides auto-rebuild of hot swapped disks without compromising performance. Where required dual parity protection can be implemented to protect against the failure of a second disk in an array while another is rebuilding. Again this does not impact performance.

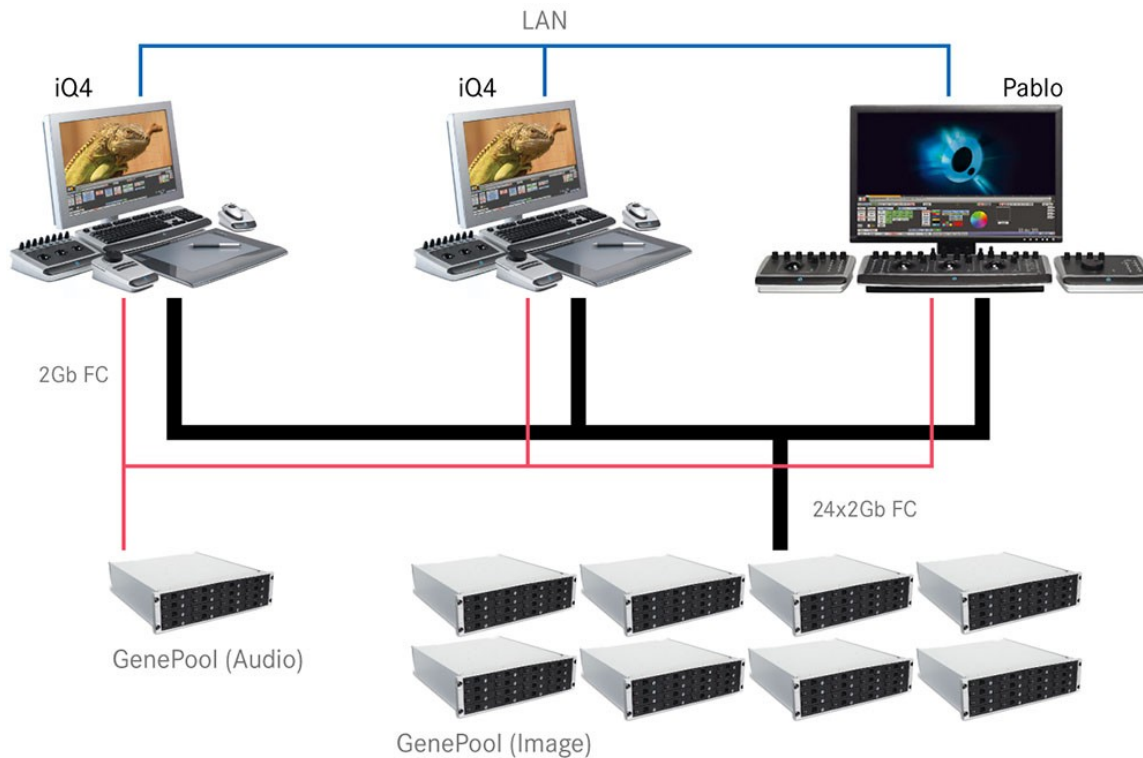
**No copy editing** The disk is capable of playing out media in any order in real-time. It does not need to copy or move media in order to play out edited selections of shots. Multiple edited clips can refer to the same media. This technique makes storage incredibly efficient as the media only exists once on the disk irrespective of how many times it is used.

**Frame-based media management** To implement no-copy editing requires a strong media management mechanism or inefficiencies are easily reintroduced. The media management tracks each usage of every frame, this ensures that long rushes are not locked from deletion simply because a few frames are used elsewhere in an edit and also that each user can safely delete unused material as the system will not delete frames if they are being used in other clips. Both the above mechanisms are completely transparent to the user, the storage manages the process without any user input allowing the user more creative (billable) time and to spend less (non-billable) time on media and disk management.

The new infrastructure extends these benefits not only to multiple Quantel systems sharing the same disk pool but also to 3<sup>rd</sup> party systems accessing the shared storage pool. The paper will consider these two aspects of sharing separately.

## Genetic Engineering: Quantel to Quantel Sharing

Genetic Engineering infrastructure allows multiple Quantel systems to share the same 'GenePool' workspace. This shared workspace replaces the direct attach storage. Each system accessing the pool retains exactly the same capabilities as it would have with direct attach storage – there are no compromises that have to be made to benefit from the shared infrastructure.

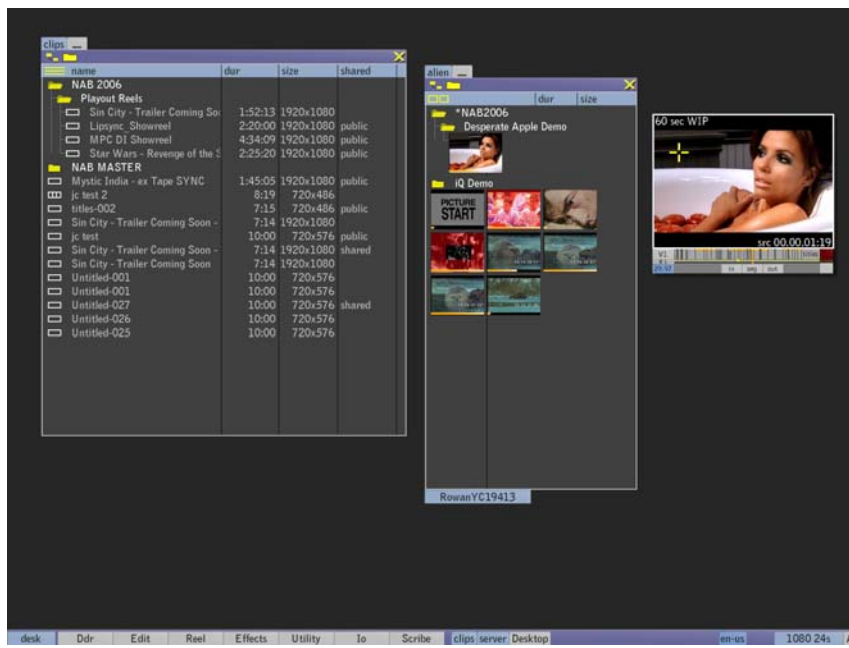


Quantel to Quantel teamworking infrastructure using GenePool shared Workspace

In order to increase reliability and minimise system components there is no storage controller required, each mainframe manages its own storage just as it does with direct attach workspace. Audio is stored in a separate disk array connected by fibre channel to each mainframe. New mechanisms are provided to ensure each system knows the current state of the free space so there is no risk of two systems attempting to write to the same free space. This inter-system communication takes place over the LAN whenever space needs to be allocated. This method also automatically provides fault tolerance should a mainframe become temporarily unavailable.

Due to the fibre channel topology each mainframe has access to the same disk bandwidth as it would have with direct attach storage so its performance remains unchanged. In the above diagram each system can reliably and deterministically play 4K (or could play for example Stereo3D) without any impact on the systems' operation. So for example one system can be doing a client-attended grading session, another can be performing quality control and the third playing out a 4K stream as HDRGB to HDCAMSR for a high quality deliverable. There is only ever one copy of any particular media and each system can, in theory, have independent access to it. In practice there is an access control system implemented to allow users to select which systems have access to individual clips in their libraries.

Each system can see into the others' libraries and clips can be browsed and sorted just as in the local library.



Accessing a clip on another Quantel system is a simple drag and drop operation

Dragging a clip from a remote library onto your desktop causes a new virtual copy of the clip to be created in your system, no new media is created and the frame-based media management kicks-in and increases the usage count of every original frame in the clip. Everything remains live, all the edits, the tail frames and all the history are intact. Resolution Co-existence still applies so, for example, a 2K reel graded on one system can be played for SD DVD from another without creating any new media or waiting for anything to render.

This instant sharing opens up some new efficient workflows for Quantel users. There is considerable flexibility to handle last minute schedule and editorial changes from clients as any project can be available to work on in any suite at a moments notice. This flexibility is simply not available if you have to wait for several Terabytes (TB) of data to be moved! Material can be loaded and conformed in a 'backroom' environment before being instantly available for grading in a client suite. An assist product with video and data i/o - Max - allows customers to build the most cost effective environment. Max means that file and video deliverables can take place away from the client suite increasing billable time.


Existing infrastructure connections and mechanisms, for example to the facility SAN and NAS, are unaffected by this new infrastructure. However any media imported into the workspace will now be available to all the workstations on the pool as required.

The Genetic Engineering infrastructure can support up to three Quantel workstations (e.g. one Pablo, one Max and one iQ or any combination of these three) Current high-end workspace configurations can vary from 16 to 64 hours of 2K, i.e. from around 20TB to 80TB. It is expected that more workspace will be able to be supported in future. Third party SANs can be directly connected to Quantel workstations or using a 'remote i/o' function through 'Sam' (described overleaf) via fibre channel for super-fast completely background file transfers, all controlled from the creative workstations.

Existing eQ, iQ and Pablo systems can be upgraded to take advantage of the benefits of shared workspace and existing Dylan FC workspace reallocated to a GenePool. Three different shared workspace configurations are available, one supporting up to HDRGB with eQ, Pablo HD and Max HD, one supporting up to 2K for DI with iQ2K, Pablo2K and Max 2K and one supporting up to 4K for DI with iQ4K, Pablo4K and Max 4K.

## Files and folders open access

If the GenePool infrastructure only provided Quantel to Quantel sharing it would be a useful development for the post and DI pipeline but this is only half the story. As has already been stated Post, DI and Broadcast Creative Services today are about teamworking and the GenePool infrastructure integrates any system capable of working with media in standard file systems. A shared access product, Sam, provides open read/write access to media in the GenePool.




Windows view of Media presented via Sam using CIFS

Using CIFS (Common Internet Filing System) third party applications, running Linux (2.6 Kernel) or Windows XP/Vista or Mac OSX can gain read/write access to media in the shared pool. There is no modification or special API required, any application that can access media over a network will be able to access the shared pool 'out of the box'. The Sam CIFS server supports multiple simultaneous clients each accessing the media at the same time without blocking. Depending on configuration and the network Sam can deliver several hundred MB/s to support 3<sup>rd</sup> party applications. Of course any file i/o activity does not impact on the performance of the Quantel mainframes connected to the shared pool.

**Common Internet Filing System (CIFS)** is a cross platform protocol used for representing disks across networks. CIFS has been around for many years but it is currently undergoing something of a renaissance in IT circles as it fits in well with new architectures such as File Area Networks. Industry giants such as Microsoft (for Windows Vista) and IBM (for Linux) are investing heavily in CIFS. The use of a protocol with such wide industry support ensures fast access to the latest advances made in the IT world.

The combination of open access and creative high end workstations dramatically simplifies the workflow in many post and DI pipelines. For example a DI facility can now scan direct to the shared workspace, it can dust bust on media in the shared workspace and finally output back to film from the shared workspace (GenePool). Sam provides enough bandwidth for film scanners, dust busters and many other kinds of systems to work simultaneously on the GenePool even though they may both demand 70 -100MB/s data bandwidth.



Typical DI facility with Genetic Engineering

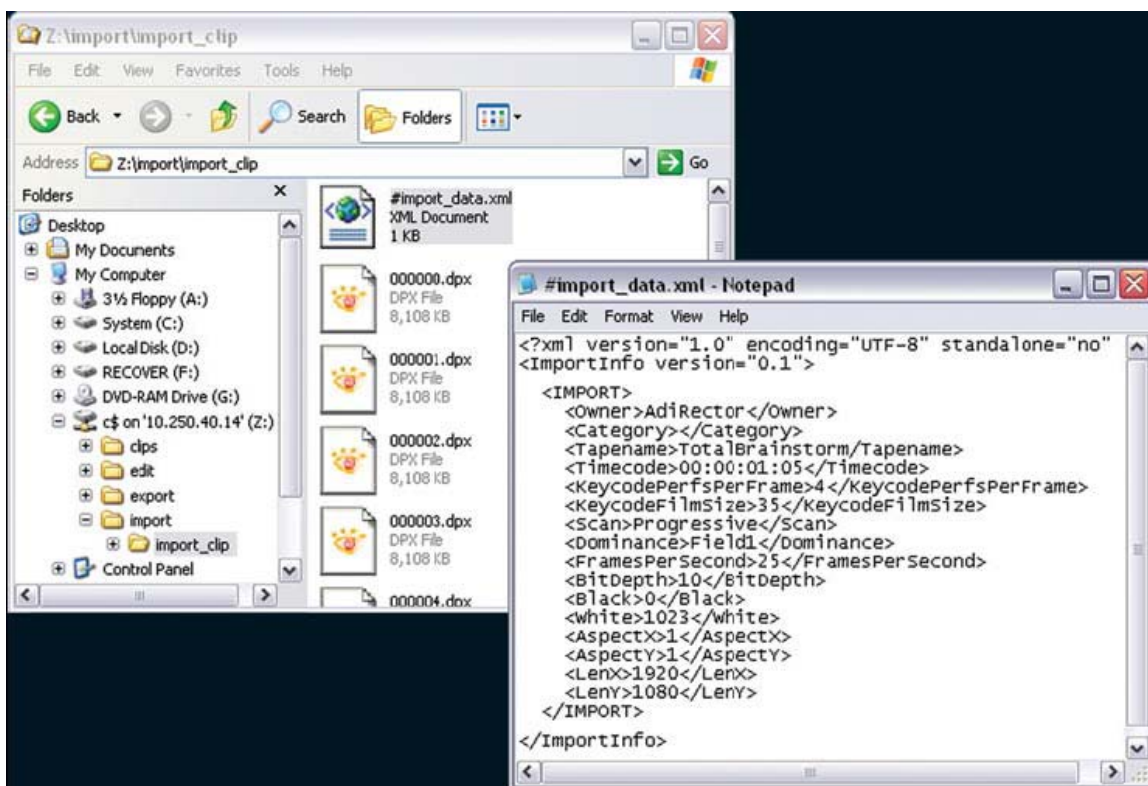
There are headline savings in infrastructure costs but also significant saving in time and management effort tracking media between different storage subsystems.

## Opening up the metadata

In the typical DI facility, Post house or Broadcast Creative Services department, there is a significant amount of management effort involved in tracking and ensuring the essential metadata needed to conform, process and display media correctly is in place and is accurate. Metadata is usually transported in file headers, making it far from human readable and non-too easy to access and modify.

When importing media Quantel systems read the metadata that is present and allow the operator to verify and modify it via the UI but this is generally something that is best not done in a creative suite. The new infrastructure provides a novel mechanism for handling metadata allowing it to be both human-readable and easily modified.

As a clip is written into the Sam import folder, any metadata in the file header is read and virtualised in an editable XML file. Opening and reading the XML file allows the metadata to be verified, editing the XML file allows the metadata to be modified.



Sam presents metadata as XML so it can easily be modified

On import this gives access to such metadata fundamentals as keycode, tape name and time code, technical parameters such as pixel aspect ratio and black and white values and descriptive text fields such as owner and category.

For reading a clip a similar mechanism allows metadata to be modified, for example to assign a specific time code to the clip. When files are read from Sam, the metadata is placed into the appropriate slots in the file headers. The mechanism is extended for reads to cover how the clip is virtualised, see the following section. Allowing the metadata to be handled in this non-file format specific way simplifies the task of any external system that needs to work with metadata.

## Virtualisation and Distribute

A further aspect of Sam is to virtualize media, initially to different resolutions and a common RGB colorspace but in the future to different file formats and bit depths. For example a clip that contains media originally from Digital Betacam (SD 4:2:2) from film (2K 4:4:4) and from .r3D Red camera files, can be presented to the outside world as if it was all HD RGB in dpx, Tiff or non compressed Quicktime files. Sam uses purpose-built hardware to handle the computationally intensive conversions on-the-fly as the media is read by the 3<sup>rd</sup> party system.



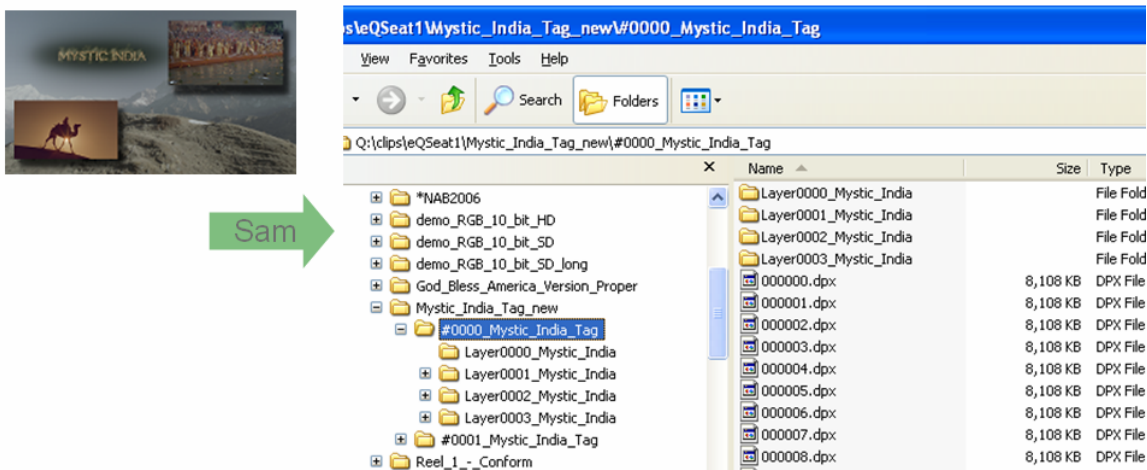
Purpose-built hardware enables on-the-fly media virtualisation

As always no new media is created on disk during these operations. Each application simply has access to the media in the format it needs it when it needs it without having to explicitly manage file conversions. How the clip is virtualised is controlled by modifying an associated XML file.

A 'distribute' function on each workstation then allows media to be re-formatted and made available for third party systems without any rendering – for example an HD clip can be made available in SD to a desktop application that cannot play back HD.


## Integrated Media Management

The files and folders view is not simply presenting a view of flattened files but it also exposes media relationships.



Sam exposes media relationships

The above four layer composite is presented not only as a flattened file but also as a series of ingredients. This allows 3<sup>rd</sup> party systems to not only access the media itself but also to access the history of how the media was created, all without any external media management system. This is a truly powerful environment for VFX or Broadcast Creative services.



HD VFX Facility using Genetic Engineering

A VFX facility can use multiple desktop VFX systems each accessing media in the GenePool. Media does not have to be transferred to the desktop systems they can work direct on the media in the shared pool, this greatly simplifies media management and also makes it straightforward to see an overall view of the current state of the work in progress as all the media is one location and can be instantly conformed for viewing. Full history remains available and files are never flattened making changes much easier than today.

### Genetic Engineering: A New Teamworking Infrastructure for Post and DI

The new infrastructure combines for the first time a high-performance HD, 2K, 4K and Stereoscopic 3D capable teamworking environment with open access for 3<sup>rd</sup> party systems. The benefits are enormous.

- Guaranteed multi-stream HD, 2K, 4K and Stereo3D for Post and DI
- Integrate Quantel and third party software applications for Broadcast Creative Services
- Combine tools and talent on collaborative VFX
- Performance maintained even under extreme conditions
- No time wasted defragmenting the shared storage to maintain performance
- New efficient workflows for teamworking and scheduling flexibility
- Large (up to 80TB) workspace holds multiple projects simultaneously
- No copy editing maximises storage efficiency
- Less requirement for multiple storage systems simplifies facility management
- Open file presentation allows different toolsets to be easily combined on the same project
- Integrated media management automatically tracks media simplifying external systems
- Integrated media management exposes media relationships enabling more efficient teamworking
- Virtualisation enables easier teamworking between different applications
- Virtualisation enables much more efficient disk usage and simplifies media management
- Increased flexibility to handle late changes

## Acknowledgements

The author thanks Simon Rogers and James Cain in Quantel R&D for all their help in preparing this paper. They are responsible for the original thinking that makes the new infrastructure such an exciting development and also for its implementation that is benefiting post, DI and broadcast facilities worldwide.

Please contact [steve.owen@quantel.com](mailto:steve.owen@quantel.com) (or [mark.horton@quantel.com](mailto:mark.horton@quantel.com) ) if you have any questions on this white paper.