High Performance Collaboration
stone® shared White Paper

TABLE OF CONTENTS

HIGH PERFORMANCE COLLABORATION

With computer performance and storage capacity continuing to increase at a rapid pace, a digital revolution is taking place in the film and television post-production industry. Fuelled by the dominance of computer software in post-production, studios are dramatically changing their configuration to data-centric workflows: the use of pure data formats as the asset to host and work on within their walls. This is having a phenomenal impact on the process of creating content, shifting the facility from an analog infrastructure with digital 'pools' to a full digital pipeline where massive advances in creativity and productivity are possible.

Core to these new digital infrastructures are new solutions that provide centralised access to data enabling multiple users to collaborate more effectively. In digital media environments these centralised data storage systems are required to provide a high-level of performance in order to provide real-time access to data. Discreet's **stone**® **shared** storage area network (SAN) is one such solution.

Although no two digital media productions ever have exactly the same infrastructure requirements, it is possible to draw some general conclusions as to the type of workflows required for different types of production. In this white paper we look at typical workflows for digital intermediate applications including scanning, grading and printing; post-production applications including editing, visual effects, compositing, 3D, paint and rotoscoping; and broadcast applications including episodic programming and news.

A digital revolution is taking place: studios are dramatically changing their configuration to data-centric workflows where massive advances in creativity and productivity are possible.

Each workflow requires a different amount of shared storage capacity and input/output (I/O) bandwidth. Based on the media format being used, the type of collaboration required and the number of real-time systems that need access to the SAN, it is possible to determine the appropriate SAN configuration that will serve the needs of the facility.

An example of shared storage use in a post-production facility

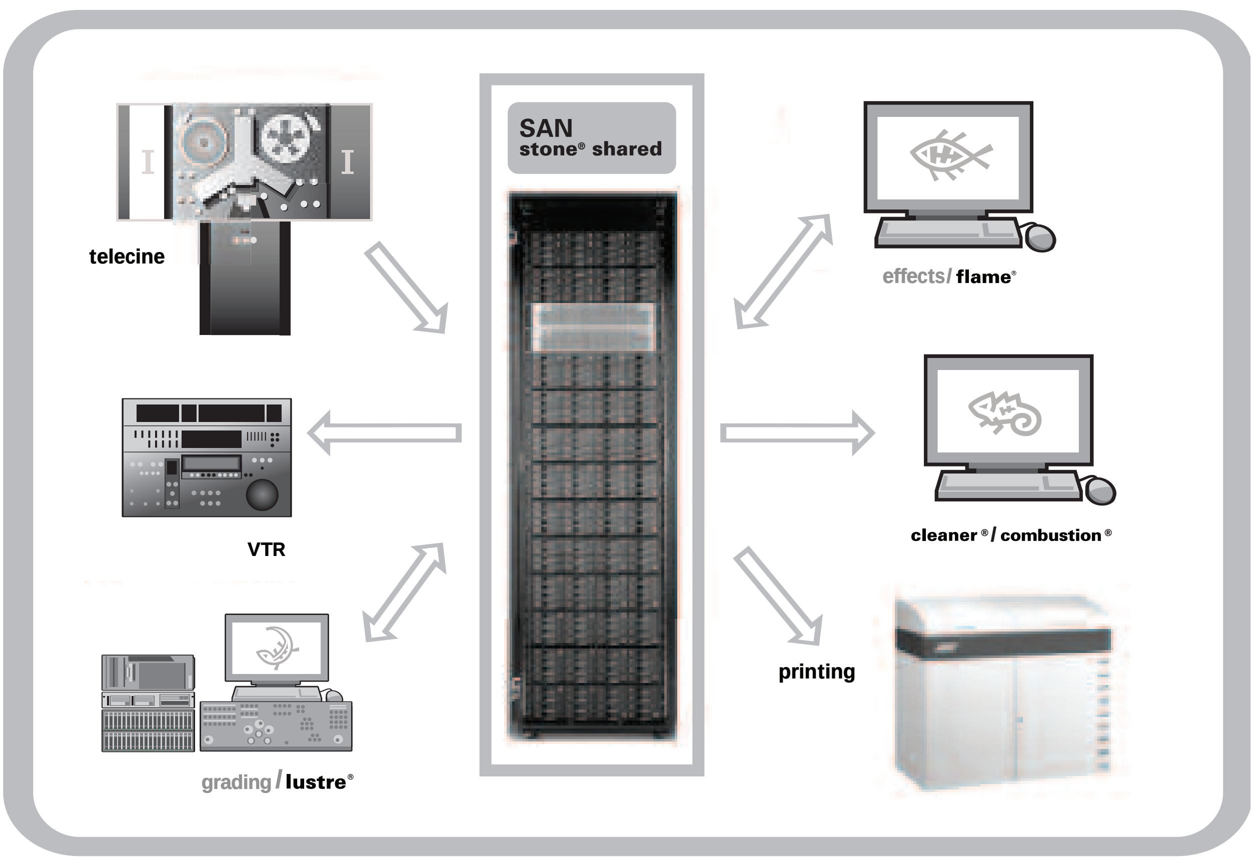

When configuring a SAN it is important to take into consideration every application's exact usage of that SAN. Some applications only need to write to the SAN (e.g. scanners), others only need to read from the SAN (e.g. printers), while others may need to both read and write to the SAN (e.g. a grading system).

Real-time applications that need to simultaneously read and write to a SAN require at least double the bandwidth of applications that only ever need to read, or write, to the SAN at any given time. Applications like VTRs, VCRs and DDRs read and write to the SAN but never simultaneously. When used they are either reading from the SAN or writing to the SAN and so behave like printers and scanners.

Also, before a SAN configuration can be established, the media formats it will be required to support must be well understood. All formats are not equal and some formats will put significantly more stress on resources than others. The following table describes the typical bandwidth requirements for digital media applications.

Typical Bandwidth Requirements for Digital Media Applications

FORMAT	RESOLUTION	BYTES/FRAME	MB/S	GB/HOUR
CCIR601-525 8bit (NTSC)	720x486x8@30	8398080	30.03	105.59
CCIR601-525 10bit (NTSC)	720x486x10@30	11197440	40.05	140.78
CCIR601-625 8bit (PAL)	720x576x8@25	9953280	29.66	104.28
CCIR601-625 10bit (PAL)	720x576x10@25	13271040	39.55	139.05
720p24 8bit	1280x720x8@24	22118400	63.28	222.47
720p24 10bit	1280x720x10@24	29491200	84.38	296.63
720p50 8bit	1280x720x8@50	22118400	131.84	463.49
720p50 10bit	1280x720x10@50	29491200	175.78	617.98
720p60 8bit	1280x720x8@60	22118400	158.20	556.18
720p60 10bit	1280x720x10@60	29491200	210.94	741.58
1080p24 8bit	1920x1080x8@24	49766400	142.38	500.56
1080p24 10bit	1920x1080x10@24	66355200	189.84	667.42
1080p25/i50 8bit	1920x1080x8@25	49766400	148.32	521.42
1080p25/i50 10bit	1920x1080x10@25	66355200	197.75	695.23
1080p30/i60 8bit	1920x1080x8@30	49766400	177.98	625.71
1080p30/i60 10bit	1920x1080x10@30	66355200	237.30	834.27
2K 10bit	2048x1556x10@24	101974016	291.75	1025.68
4K 10bit	4096x3112x10@24	407896064	1167.00	4102.73

Based on a sustainable bandwidth of 1000 MB/second[1], the number of playback streams and real-time systems that can be supported by **stone shared** configuration is as follows:[2]

Playback streams & real-time systems supported by stone shared

FORMAT	MB/S	THEORETICAL # OF RT STREAMS[1]	THEORETICAL # OF DUAL STREAMS[3]	DISCREET SYSTEMS[4]
CCIR601-525 8bit (NTSC)	30.03	33	11	5
CCIR601-525 10bit (NTSC)	40.05	25	8	5
CCIR601-625 8bit (PAL)	29.66	33	11	5
CCIR601-625 10bit (PAL)	39.55	25	8	5
720p24 8bit	63.28	15	5	5
720p24 10bit	84.38	11	3	5
720p50 8bit	131.84	7	2	3
720p50 10bit	175.78	5	1	2
720p60 8bit	158.20	6	2	3
720p60 10bit	210.94	4	1	2
1080p24 8bit	142.38	7	2	3
1080p24 10bit	189.84	5	1	2
1080p25/i50 8bit	148.32	6	2	3
1080p25/i50 10bit	197.75	5	1	2
1080p30/i60 8bit	177.98	5	1	2
1080p30/i60 10bit	237.30	4	1	2
2K 10bit	291.75	3	1	3
4K 10bit	1167.00	1	0	1

[1] A **stone shared** standard configuration offers 5-, 10 or 21-Terabytes of storage with 1000 MB/seconds of I/O. However in order to ensure peak performance it is recommended that a SAN be configured to have sufficient bandwidth headroom to avoid performance conflicts. SAN performance decreases once storage capacity reaches > 50%.
[2] These figures are approximate and meant as a guideline only.
[3] Dual stream implies that you can play back 2 streams and write 1 per system.
[4] Discreet supports a maximum of 5 systems in standard configurations. Professional services are required for non-standard (larger) configurations.

discreet T 514.393.1616 1.800.869.3504 F 514.393.0110 10 RUE DUKE. MONTRÉAL (QUÉBEC) CANADA H3C 2L7

Digital Intermediate (Feature Film Production)

A digital intermediate (DI) workflow will typically comprise systems for film scanning, editing, conforming, quality control, dust-busting and film repair, colour grading, digital effects (to replace opticals) and printing. Depending on whether the facility is primarily a DI facility or also involved in the creation of special effects other systems may include 3D animation, compositing, paint, rotoscoping and other visual effects applications.

Typical DI workflow

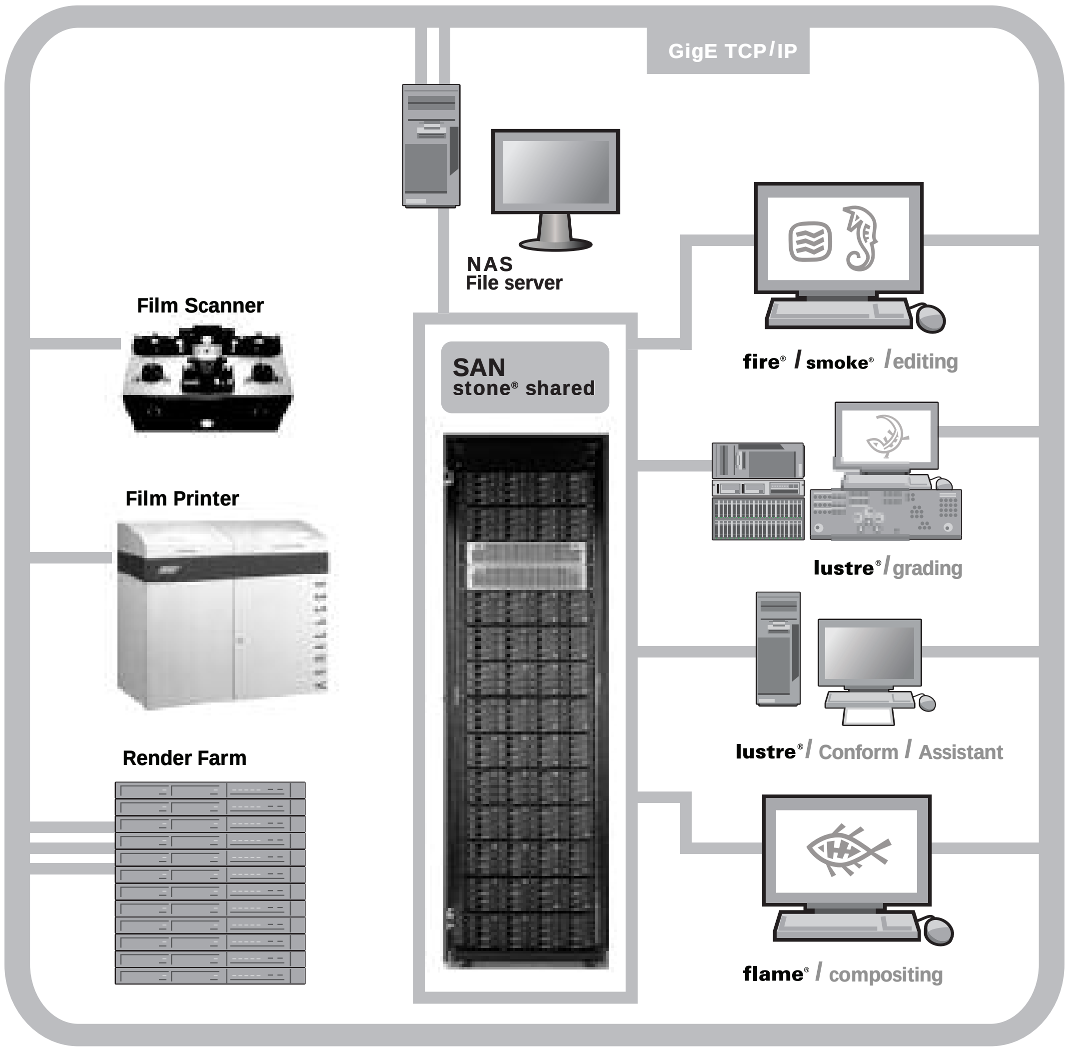

In a typical DI workflow only applications requiring real-time access to data need to be connected directly to **stone shared** via Fibre-Channel. However, many DI applications cannot process data in real-time and therefore do not require direct access: High-quality film scanners, laser printers and render farms are all examples of DI applications that process data at rates that are significantly slower than real-time.

A film scanner operating at one second per frame needs less then 15 MB/s bandwidth to record 2K data to the storage device. Render farms may take seconds, minutes or even hours to render a frame, depending on the complexity of the render. Any delay in reading or writing a frame to the storage device is easily buffered and will not have a critical impact on overall application performance. As a result such applications do not require high bandwidth for proper operation and standard 100baseT or GigE networks are generally sufficient. In these cases the application will connect to **stone shared** as a NAS device via a file-server. This allows more Fibre-Channel bandwidth to be reserved for real-time applications.

The File Server has a direct Fibre-Channel connection to **stone® shared** and can be configured with multiple network ports bonded together to create higher bandwidth pipes if required. For facilities with a large number of non real-time seats (e.g. with large 3D CGI departments) there may be a lot of network traffic to and from the file-server. This traffic is controlled via Network Load Balancing (NLB) Software. The NLB software ensures that bandwidth aggregation requirements of the file-server are met and dynamically balances similar network interface devices into a single IP address. The traffic is handled automatically by the software for greater ease of use to the client.

Real-time applications such as the interactive creative seats required for screening, colour grading or high-speed compositing, are given a direct Fibre-Channel connection to **stone shared**. Each of these systems requires a minimum level of interactivity and real-time capabilities that can be guaranteed via SGI's Guaranteed Rate I/O (GRIO) feature. However only a certain number of systems can be guaranteed real-time access to media depending on the bandwidth capacity of the SAN.

In the typical DI workflow diagram, 2K 10-bit log DPX film scans are transported over the gigE network from the scanner to **stone shared** and written onto the SAN by the file-server. Once a film reel has been scanned, it is available to the **lustre®** grading system, **flame®** compositing system and **fire®** editing system in real-time through direct Fibre-Channel connections and CXFS™ GRIO. At three streams of 2K a standard **stone shared** configuration is close to maximum capacity. Adding more real-time applications will require SAN customisation.

Real-time applications such as the interactive creative seats required for screening, colour grading or high-speed compositing, are given a direct Fibre-Channel connection to stone shared.*

The scanned film can be graded by the **lustre** grading system in real-time, rendered grades are stored back on the SAN and can be accessed (along with the original frames) by the **fire** and **flame** systems. A **fire** editor or **flame** artist simply browses **stone shared**, selects the clips they need and Soft Imports them. Soft Import creates a virtual reference clip the Discreet application so that at any time the editor/artist can simply click on the clip to play the selected frames from **stone shared**. If an effect is processed on the clip **fire** and **flame** will automatically search for the frames referenced in the clip from **stone shared**. The final renders can be published back to **stone shared** so that other applications can access them.

The film scans and rendered files on **stone shared** are available to any system on the house network via the file-server. If film restoration is an integral part of the Digital Intermediate workflow it may require direct connection to **stone shared** if the restoration system is real-time. Otherwise, non-interactive or automatic restoration applications may not need a high bandwidth connection.

Some DI facilities also provide DVD encoding services. This service will not typically require a high speed connection, since the actual encoding is often not supervised, and mostly a rendering process. Access to the data will be available via the file-server.

stone shared is designed to provide the capability and bandwidth required for today's DI workflows. The large amounts of data and collaborative nature of DI, make **stone shared** a necessary item for any DI facility. **stone shared** is open, and works well with all required components. Its easy integration means little downtime and its simplicity means that IT departments can easily manage it. **stone shared** is a competitively priced solution that uses standard components known and trusted within the industry and by IT professionals. This Discreet branded solution offers the one-stop-shop performance, quality and support advantages that all of Discreet's solutions offer.

[3] Dual stream = 2 streams read and one stream write per system
[4] Discreet only supports a maximum of 5 systems on standard configurations. Larger SANs must be custom configured. Discreet systems currently only require dual stream read performance for HD and single stream for 2K+.
* Discreet **inferno**, **flame** and **flint** applications require **stone® direct**

stone shared is open, and works well with all required components. Its easy integration means little downtime and its simplicity means that IT departments can easily manage it.

Advertising (Commercials, Music Videos, and Effects)

An advertising/post-production facility workflow will typically comprise of systems that can include telecine, colour correction, online conforming and finishing, graphics, effects & compositing, editing, 3d animation, audio sweetening, DVD authoring and dubbing to name but a few. The requirements of a typical post workflow differ from a digital intermediate infrastructure only in that the post house does not typically have film scanners and printers and does not have to typically assemble the huge amounts of data and footage required for film production. However each post-house is different, and each may have elements of DI for enhanced workflow or because of market demand.

Many post-houses base their workflow around a few high-end online suites. These may typically be **flame** or **smoke** suites. These are used for client supervised sessions and are the suites where the final product (commercial, music video, TV program) is assembled and finished. For simple projects only one operator and suite may be required. For more complex projects several artists/editors may be required to work on the project over a fairly long period of time. In this case the ability to share material becomes very useful[5]. **stone shared** provides a reliable, high-performance centralised storage repository for such projects.

Typical Advertising / Post Production workflow

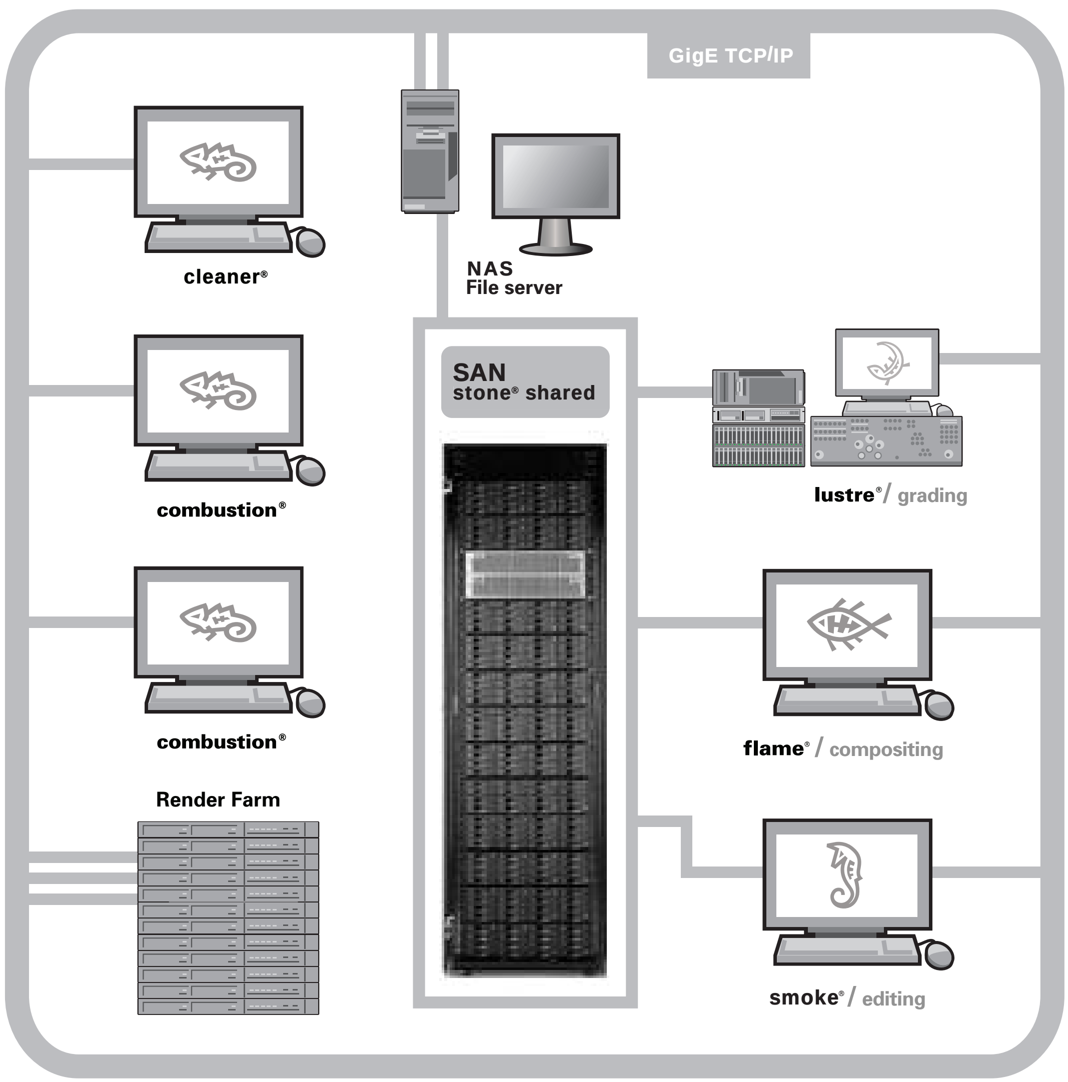

discreet T 514.393.1616 1.800.869.3504 F 514.393.0110 10 RUE DUKE. MONTRÉAL (QUÉBEC) CANADA H3C 2L7

Most online suites also require the support of a sometimes significant number of graphics workstations creating 3D animations, 2D graphics, text, and other elements for the project. Even though the graphics stations often have high processing performance, they typically do not require sustained real-time access to the online media and a direct Fibre-Channel connection to **stone shared**. This is because the CG work is not typically client supervised and does not require the interaction of the online suite. Since direct real-time access to a SAN is always expensive compared to network access via the file-server, the latter is often the better solution for connecting these stations to **stone shared**. Often the data rates from the CG systems will be much smaller than those of online editing, effects and grading systems.

Television (SD/HD Episodics, News)

Episodic Facilities

Television HD/SD episodic programming workflows often consist of a full pipeline for ingest, offline, grading, graphics and effects, compositing, conform, encoding, and final delivery. This workflow is tightly integrated, and all pieces must work well together to deliver on the stringent needs of television broadcasters in terms of rapid turn-around and throughput.

Videotape workflows are not only non-collaborative, they can reduce quality and usually limit the amount of creative work that can be done per episode. Making last minute changes is also harder. stone shared can help greatly increase the amount of time that can be spent making the program look, feel and sell better.

High-quality, prime-time TV shows shot on film are typically scanned in from a telecine with a first-light pass to the SAN. At the same time, a low-res offline version is created. The offline EDL can be used by **lustre**, and **fire** to collaboratively assemble the episode. Common elements are made available on **stone shared**, and can be easily and instantly integrated. Since all the data is online, any modifications to titles, captions, or the likewise can be made quickly and easily. Any shots requiring composition or FX are directly available within the **flame**, and CG elements created by the CG department can be easily shared via the SAN. Also those elements can remain available for future episodes, or as special items for DVD authoring.

The final graded and assembled program is laid out to videotape, and delivered for broadcast. A DVD version can also be produced at the same time, reducing cost and offering greater flexibility for deliverables. All data is resolution independent, not requiring down-conversion, before final layout to tape.

This provides the TV producer with greater flexibility, and reduces the amount of tape-based workflow. Episodic workflows can ideally benefit from **stone shared**. Today, most are still based on videotape workflows. Videotape workflows are not only non-collaborative, they can reduce quality and usually limit the amount of creative work that can be done per episode. Making last minute changes is also harder. **stone shared** can help greatly increase the amount of time that can be spent making the program look, feel and sell better.

With more and more episodics being shot in HD and the emergence of high quality HD formats like HDCAM SR, **stone shared** also provides the ideal shared storage solution for high-quality collaborative HD work.

[5] **Note:** current versions of IFF/FS/B do not support video I/O directly to **stone shared** so stone direct is the main storage point for video I/O. In datacine environments 4:4:4 scans via dual-link to **stone shared** provide a powerful means of centralised collaboration as well as an easy means of integrating advanced digital colour grading into the workflow via **lustre**.

Typical Television Episodic workflow

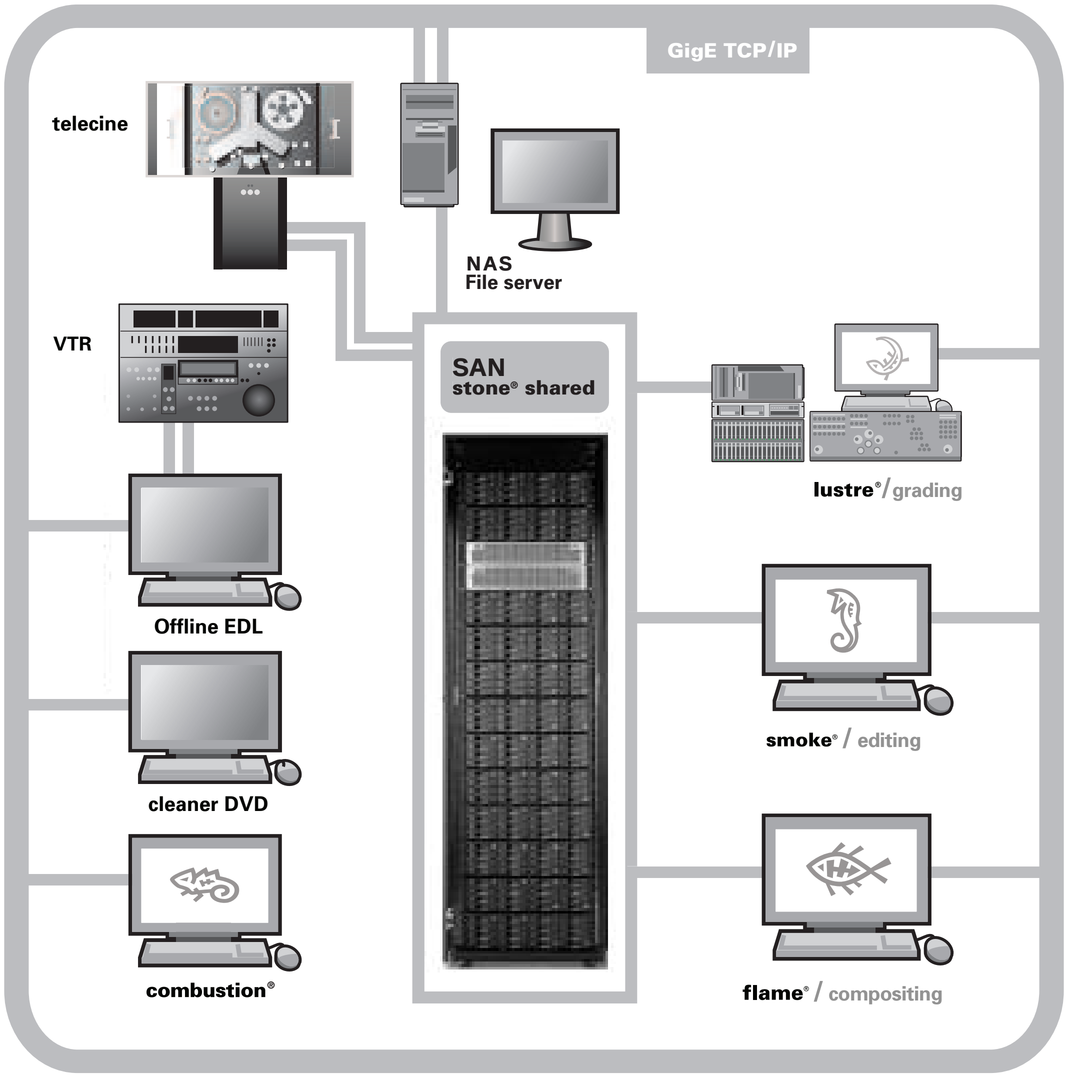

NEWS Facilities

News facilities also have extremely fast paced workflows. Many news shots come in via satellite feeds or videotape and require rapid assembly and quick integration of graphics and archive footage to create the days newsreels.

The GFX and effects departments of newsrooms must create all the GFX used for the day's show, as well as create all the promos, opens, and specials. A news facility's greatest requirement is the immediate availability of data. All data is ingested, and placed on the SAN for anyone's use. The ready availability of archival footage is also necessary, as it needs to be easily integrated. A large SAN can allow for a lot of available data. Most news facilities have large amounts of accessible data. **stone shared** fits this requirement.

Newsroom workflows are not always ideally suited to Discreet's systems applications as a high degree of this work is done with compressed media. Typically each GFX bay, and edit bay work as separate islands, or are connected to a specialised video server. Although **stone shared** is capable of supporting these workflows Discreet does not recommend its branded solution for these applications. Clients interested in shared storage can purchase directly from SGI®. Newsroom environments that are running a few Discreet systems for their graphics and promotions departments may find great value in an open SAN solution like **stone shared** to provide collaboration not only between Discreet systems but also with the rest of the newsroom.

Typical News workflow

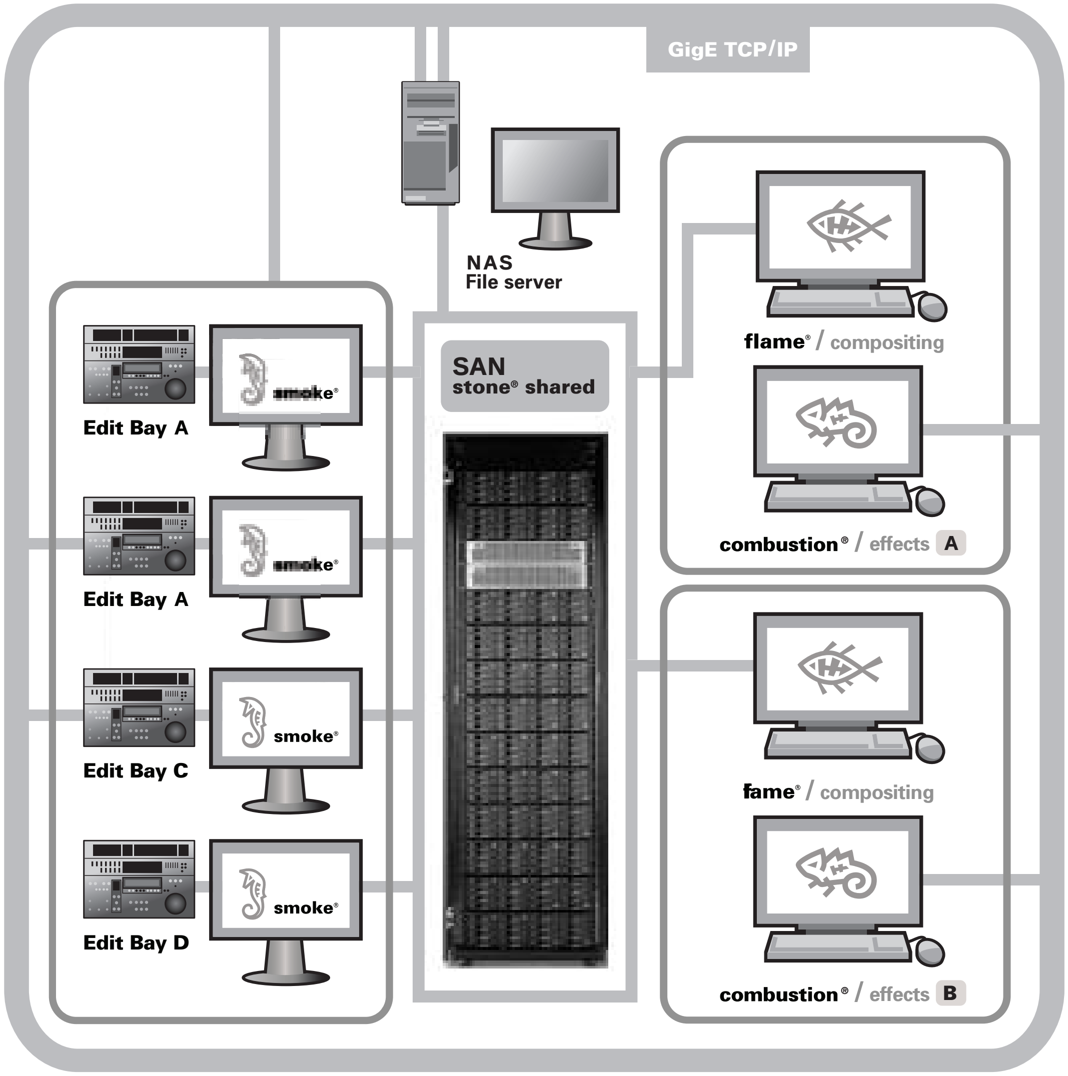

INTEROPERABILITY AND DATA SHARING

Market Drivers

Why are shared storage solutions and application interoperability of increasing importance to digital media creators such as film studios, post-production companies and broadcasters?

As digital technology continues to advance in leaps and bounds, it enables ever-greater levels of creativity and invention. The top tier of innovative content creators continuously push new digital technologies to their limits and natural competition forces the others to keep pace with both the innovators and the ever-increasing expectations of their clients. The down side is that digital media production is becoming increasingly complex.

As production complexity increases media creation companies are looking harder and harder at improving efficiency so that they can better manage increasing production requirements while ensuring that costs and schedules remain on target. As a result there is an increasing need for solutions that can help speed the production process without compromising either quality or the creativity– after all quality and creativity play a critical part in determining the eventual value of the media asset being produced – and in consequence the value (and income) of the company providing the services.

As a result forward thinking digital media creators are looking for solutions that will

1- support the highest quality raw material required for digital production – RGB data,
2- enable facility-wide (and not vendor-specific) collaboration, and
3- be able to scale to meet future facility needs.

With complexity increasing, it is impossible for any one vendor to supply all the tools required in the post-production process. Therefore, any infrastructure solution that is designed to meet the needs of only a single application or manufacturer will not be able to deliver maximum efficiency. It is for this reason that Discreet is committed to open, non-proprietary, infrastructure solutions as embodied by **stone shared** and the open access capabilities of its advanced systems.

Media and Meta-data

There are two kinds of interoperability that can help increase productivity within digital post-production: interoperability of media and interoperability of meta-data. Media is the raw material that all post-production is based on and meta-data is the value that is added in that it describes the creative decisions that have been made and applied to the media.

For media to be efficient it should be easily shared and accessible. The most interoperable media formats are file formats that are public domain or have been standardised by organisations such as SMPTE, DCI/ASC, etc. Their properties are well known, fully documented and publicly available and they can therefore be supported by every manufacturer. Standard file formats include DPX, TIFF and MPEG. Non-compressed formats like DPX and TIFF are the most optimal formats for post-production purposes as they allow extensive manipulation without quality degradation. Since modern post-production is almost entirely done on computer systems, RGB formats offer significantly better quality than legacy digital video (YCbCr) formats. The worst media formats for collaborative post are compressed or sub-sampled (4:2:2) proprietary formats as they cannot be shared between applications and they eventually end up degrading the quality of the production.

Many of the high-end, high profile systems and software manufacturers have created, and continue to create their own proprietary formats, limiting interoperability.

Meta-data sharing is more complex since meta-data is often very specific to the application being used. Creative applications are defined by the tools they offer for making changes to media. But since these tools are generally the intellectual property of the manufacturer and their sole means of generating revenue, there is very little opportunity for standardisation. As a result only the simplest, most common meta-data will ever be standardised and fully interoperable. Today this is mainly limited to basic edit decisions (length and order of clips, handles, simple transitions, track layering etc). Otherwise interoperability is obtained by individual agreements between different manufacturers to collaborate or through the use of application specific APIs (application programming interfaces).

Ideally all software would simply be interoperable, and all data shared. Unfortunately this is not the reality. Many of the high-end, high profile systems and software manufacturers have created, and continue to create their own proprietary formats. Each manufacturer puts the interests of their applications first, and since those interests differ, deciding on formats that are universally interoperable is difficult. However in the past years there has been a lot of focus on just that. AAF, or MPEG-4, and XML are examples of formats coming together to develop interoperable data.

From EDLs to XML, DPX to FBX, OpenGL to OpenML and OMF to AAF, Discreet is actively participating in researching, implementing and defining open standards and technologies as well as gathering customer feedback on their use and value to develop solutions that are as open and as interoperable as possible given cur-rent constraints and limitations. **stone**® **shared** and open access to Discreet's advanced systems (**inferno**®, **flame**®, **flint**®, **fire**®, **smoke**® and **backdraft**®) are an integral part of this initiative.

Note: The workflow shown in the above diagram has a single point of failure in the Fibre-Channel Switch. This limitation can be removed by configuring the Fibre-Channel mesh with complete redundant paths, though this will increase costs.

Open Access in Discreet's Solutions

Discreet's advanced systems applications interoperate with a SAN environment in many different ways. Whether it is read/write, soft-import/publish, normal import/export, or archiving, advanced systems can utilise a SAN environment to greatly improve efficiency and workflow. Several new tools have been created to help take advantage of **stone shared**.

lustre is designed to work in a completely open environment working with standard formats like DPX stored on open file-systems. **lustre** can read and write data directly to and from **stone shared** which can scale to support multiple streams of 2K data in real-time to several **lustre** master stations.

Discreet's open access initiative will provide the same capabilities to its other advanced systems, **inferno**, **flame**, **flint**, **fire**, **smoke** and **backdraft**. Phase 1 of open access with **inferno** 6, **flame** 9, **flint** 9, **fire** 6.5, **smoke** 6.5 and **backdraft** 6.5 provides the ability to read and publish to standard file systems including CXFS and **stone shared**. Phase 2 will provide read and write capabilities. However, guaranteed real-time video I/O requires a stone direct storage array with the **stone** file-system.

Soft-Import and Publish

Soft Import and Publish allow Discreet's editing and effects systems to work directly with standard image files, such as DPX, RGB, TGA, and TIFF for both source and final media collaboration thereby avoiding unnecessary duplication of material. Discreet systems can read the media directly from a shared storage device, such as a SAN, without having to first copy it to the local stone file-system (as would be the case with standard Imports).

Every time a Soft Clip is played or processed the application immediately fetches the original file from the storage device. At the same time, other systems and applications can also access the media providing collaboration.

The ability to read media directly from a standard file system is termed Soft Import. Soft import will create a "Soft" Clip in the Library (The Discreet application's media database) that points to the external media. Every time a Soft Clip is played or processed the application immediately fetches the original file from the storage device. At the same time, other systems and applications can also access the media providing collaboration. If another application changes the media (overwriting the old file but keeping the name), the Discreet application will load the new file. In this way other applications can dynamically update the media and the Discreet system will always use the latest results.

All Discreet systems apart from **lustre** will render all intermediate and final results to the **stone direct** storage in Phase 1. This can reduce the load on the SAN, since only original sources are ever accessed from it. In this way the SAN does not need to store intermediate files and caches that are generally only relevant to the person working on the Discreet system. Final results can be published back to the shared storage device for further collaboration with other applications.

Publish is different from an export in the sense that the published clip has an entry in the Library and when played or processed the exported clip will behave exactly like a Soft Import.

Any media that is soft-imported remains soft imported even after conforming, or doing I/O. There is an option to import (i.e. copy) the media referenced by a Soft Clip to the stone file-system (stonifise) at any time from the library. Media that has already been rendered to the **stone** file-system will not be re-rendered, making the system efficient in its data usage.

Please refer to the Infrastructure Components document for more details on Soft-Import/Publish.
Please refer to each product's User Guide for details of supported import and export formats.

OPEN ACCESS WORKFLOW

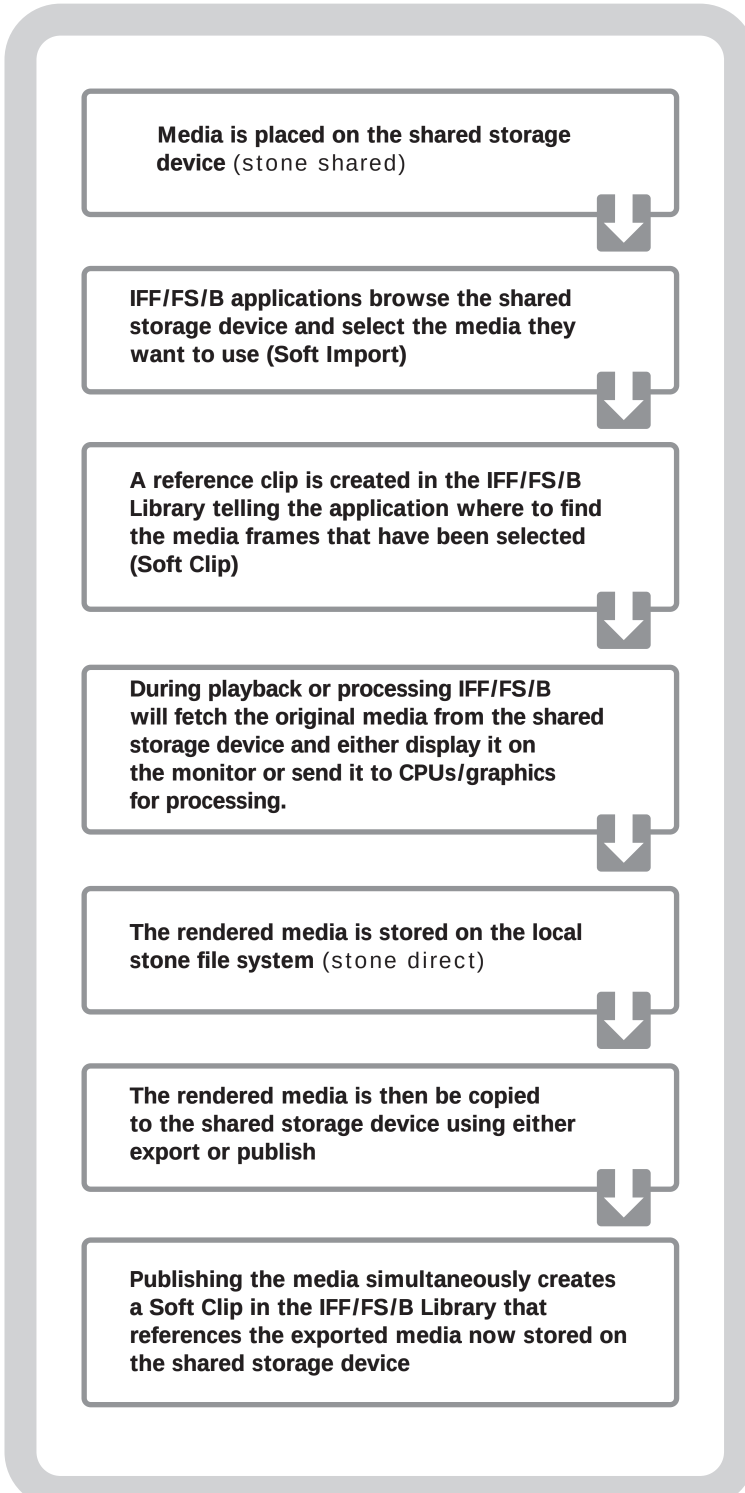

Asset Management

In typical shared storage and collaborative workflows, many of the hardware components and software applications that make up the workflow are of third party origin and not from Discreet. For this reason, media stored on the SAN will typically be of a standard image file type such as DPX which provides easy interoperability, and collaboration with other applications. However, these standard formats are, by definition, not highly optimsed for specific applications (as opposed to the raw RGB format Discreet applications use on the stone file-system). This adds performance and operational overheads which facilities must take into account when working with such solutions.

Standard formats of third party origin add performance and operational overheads which facilities must take into account.

An Asset Management system will have more tools focused on the management, cataloging, removal, searching, and archiving of the image files from the SAN. Discreet applications can only manage the data that is within its own library. Discreet systems applications include their own proprietary asset management for themselves (Clip Library). It just isn't accessible to the outside world, and doesn't manage all the media on the SAN.

This generally requires specialised Asset Management Systems designed to manage the assets of the entire facility not those of specific applications. There are many asset management systems available from different vendors, with different options depending on the facilities needs. Using open access workflows allows third party asset management systems to integrate well with Discreet systems applications, **lustre** and **stone shared**.

stone shared OVERVIEW

stone shared is Discreet's high-performance, large capacity Storage Area Network (SAN) solution designed to provide high aggregate performance and centralised data storage to multiple high-definition workstations.

The **stone shared** solution consists of robust, high performance data storage hardware combined with a clustered file-system designed for multi-user real-time concurrent access to high-resolution data from a single storage device.

stone shared disk arrays are based on Fibre-Channel disk technology, the industry standard in high performance, combining high read/writer speeds, low latency, RAID data integrity, and high bandwidth for large-scale connectivity. Direct Fibre-Channel connection to the file-system provides increased bandwidth and efficiency when transferring data between the hosts and the storage.

stone shared features SGI's CXFS clustered file-system allowing multiple computer workstations to share the same physical storage device and manages all file locking and file-system metadata. The CXFS file-system provides similar functionality and has the same attributes as a standard XFS file-system but can be accessed directly from multiple operating systems including IRIX, Linux1 and Windows XP.

stone shared is designed to provide the data integrity, security, and the performance you expect from a Discreet storage product.

Note: Sustained real-time playback and I/O is not guaranteed in current implementations of **stone shared**. Performance will vary depending on how many workstations are simultaneously accessing the storage device and the resolution of the media they are accessing, the I/O performance, characteristics of the application and its native operating system.

Benefits

- **stone shared** allows instant data sharing without network mounts.
  CXFS supports all major operating systems including IRIX, Windows, Linux[6] and Mac OS X[7].
- Extremely high performance aggregate bandwidth
- Full RAID data protection guarantees data integrity

stone shared is an open SAN solution with a non-proprietary file system designed to integrate, and collaborate well with third party applications, as well as Discreet systems and software applications. Using a SAN can enhance the clients' workflow by reducing the amount of copied content, and creating a more collaborative environment. The SAN offers huge scalability in terms of bandwidth, and storage capacity. It also offers a file-server solution for non-direct access availability to all systems on the network fabric.

stone shared is an open SAN solution with a non-proprietary file system designed to integrate, and collaborate well with third party applications, as well as Discreet systems and software applications.

Since the Discreet **stone shared** is an "open" SAN architecture, it follows the IT world of storage, and administration. Unlike proprietary SAN solutions such as Avid Unity, **stone shared** hardware and software is based on standard components that are well understood by IT professionals and well documented. A wealth of detailed technical information is available for reference. This makes it easier for IT professionals to integrate, utilise, and maintain the SAN.

Connectivity - How it all works together

Connectivity is how the various SAN hardware components interface with each other. Typically a SAN uses a Fibre-Channel fabric to transport data between the storage devices and the host systems. Fibre-Channel is used because it offers the dual benefit of high speeds, and the ability to send data long distances without losing performance or reliability.

stone shared also utilises Gigabit Ethernet (GigE) connectivity to transfer file-system meta-data and status information between the various servers that manage the SAN. The GigE connectivity is only used to manage the SAN. It is not used to transfer application data. All application (user) data is transferred over Fiber Channel. Together the Fibre-Channel drive interface and the GigE meta-data and file-server connectivity combine to create a more powerful and efficient SAN. **stone shared** includes various maintenance applications, such as disk utilities, bandwidth reservation and handling, and administration utilities, designed to maximise the performance and capabilities of the SAN. To do so, these applications have been optimised to use the resources of the meta-data servers and to communicate with their host components using GigE thereby minimising any risk of their impacting the performance of the host systems.

[6] Intel 32-Bit Architecture is scheduled for early 2005. Guaranteed Rate I/O IRIX/Windows-only
[7] Mac OS X will be supported at a later date.

Facility-wide Discreet infrastructure

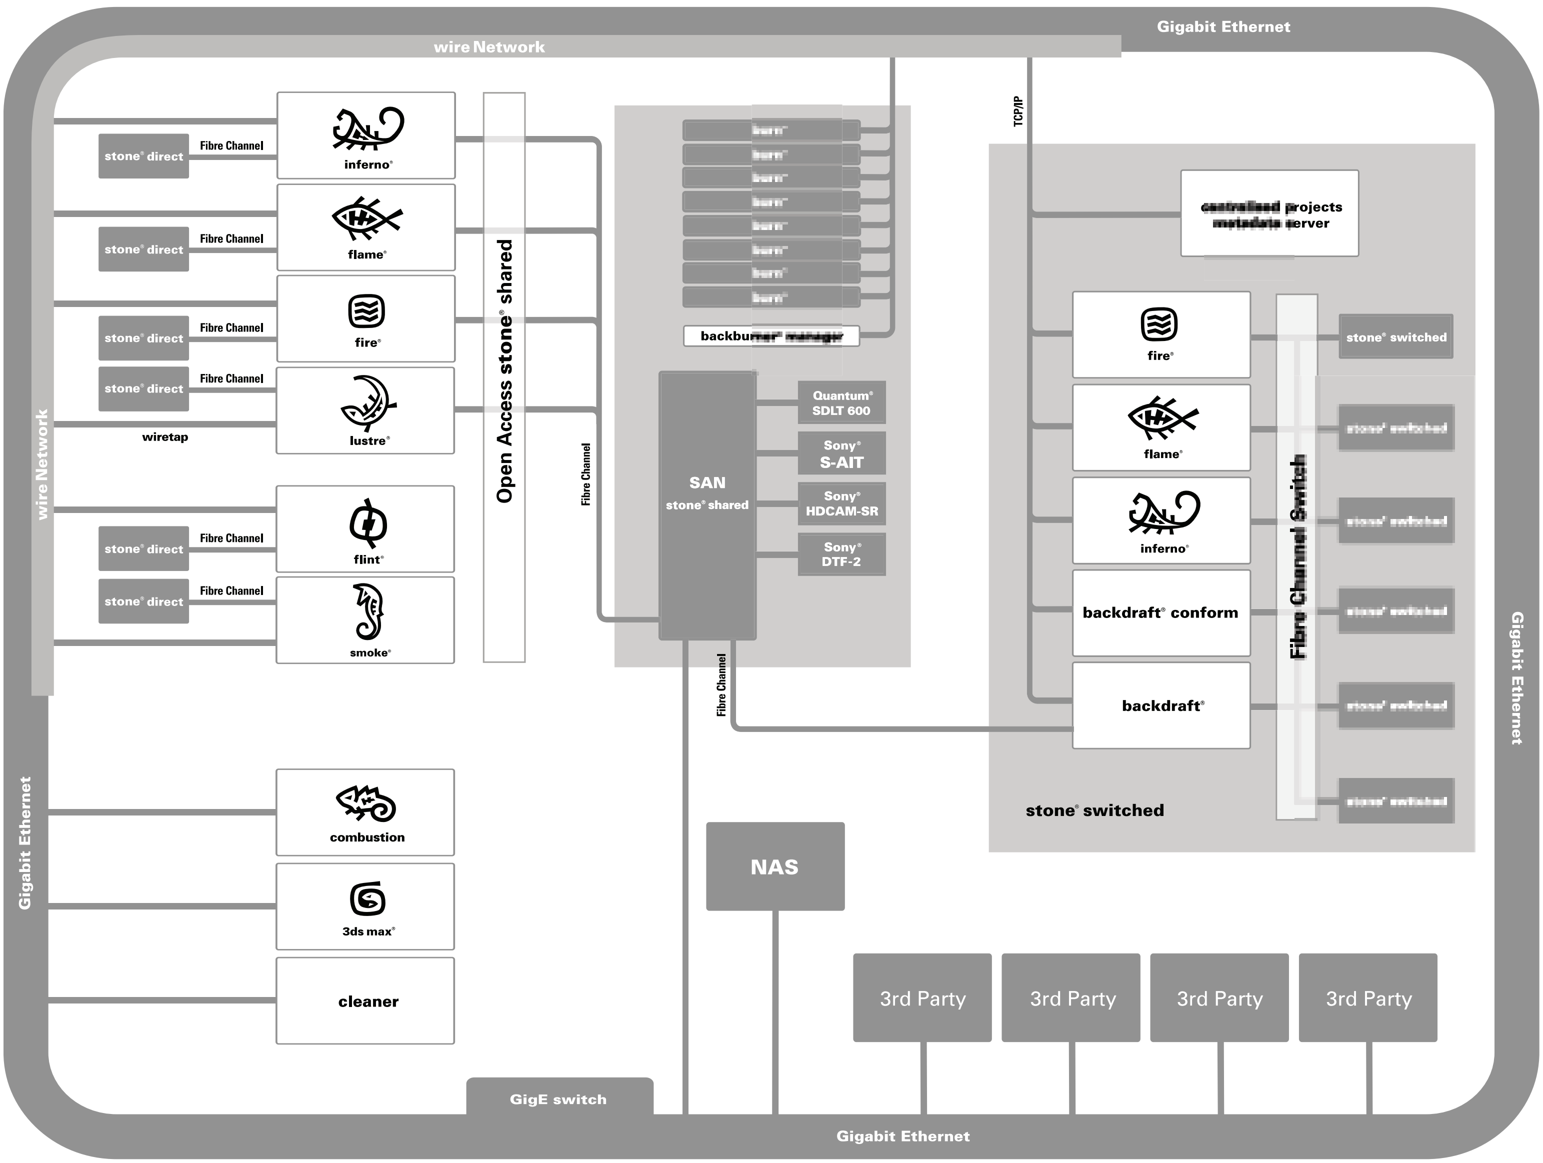

stone shared TECHNOLOGY

stone shared is made up of 3 major components, the hardware that physically stores and manipulates the data, the SAN management software that tells the hardware what to do with the data and the shared file-system that allows applications to access and share the data. These are described below.

stone shared consists exclusively of high-performance, high-quality hardware components designed to store, transfer and protect large banks of data as efficiently and effectively as possible. The hardware includes disk drives, drive enclosures, Fibre-Channel switches, RAID controllers and, at the heart of it all, the SSH8500 appliance based on DataDirect Networks (DDN) Silicon Storage Appliance (S2A) intelligent network infrastructure device (see below for more details).

stone shared is available in standard configurations of 5, 10 and 21 Terabytes and is fairly simple to deploy so that on-site installment times are significantly reduced. **stone shared** is scaleable up to 130 Terabytes through Discreet Professional Services.

Unlike proprietary SAN solutions designed to support only the applications of a single vendor, **stone shared** is designed to enable multi-vendor environments[8] to access and back up centralised data in an easy, cost effective and reliable manner.

All this means that the **stone shared** hardware is designed for best performance, simplicity, and reliability, giving maximum performance for price.

stone shared Hardware Infrastructure

stone shared hardware consists of the following hardware components

- SSH8500 appliance based on DataDirect Networks Silicon Storage (S2A) technology
- SGI Origin 350 file-server[9]
- SGI Origin 350 meta-data server[3]
- Fibre-Channel storage units (JBOD drives and enclosures)
- Fibre-Channel switches and cabling**
- Adapters

** The Fibre-Channel switches and cables are commonly referred to as the SAN “fabric”.

The unique hardware component of **stone shared** is the SSH8500 appliance which has consolidated and fully integrated SAN functionality into one powerful, easy-to-manage,

stone® shared hardware components

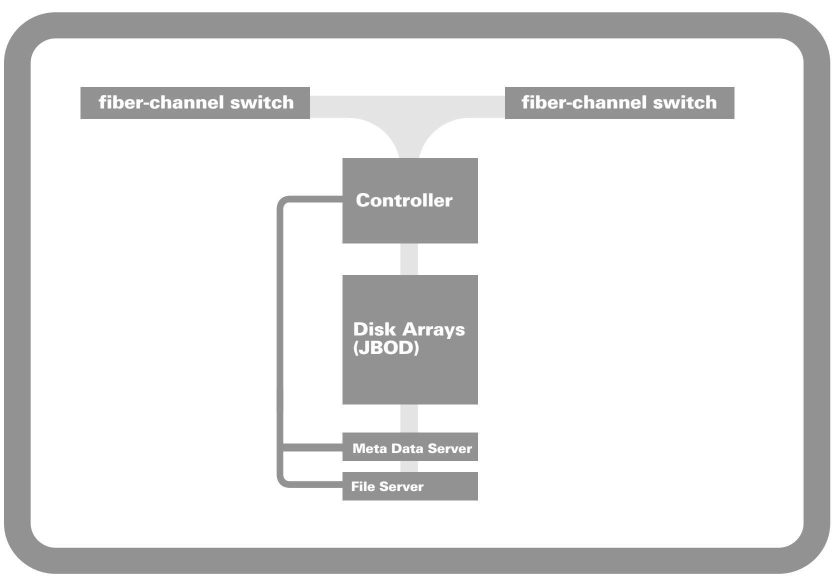

[8] Discreet does not support or guarantee the performance of other vendor's applications
[9] We always use two meta-data servers to ensure high availability. The primary meta-data server can be used for file serving

stone shared **removes the complexity, congestion, latency and contention of switched fabric SANs. This exceptional combination delivers dramatically lower total cost of ownership and improved return on IT investments.**

intelligent device. The SSH offers a fully pipelined, parallel processing architecture designed to deliver a highly-reliable SAN storage solution that is both easy to install and manage and easy to scale in capacity and performance.

The unique zero-bottleneck architecture was developed by DataDirect and removes the complexity, congestion, latency and contention of switched fabric SANs. This exceptional combination delivers dramatically lower total cost of ownership and improved return on IT investments.

Traditional SAN technologies

Today, most SAN technologies rely on complex and often inefficient hardware and Fibre-Channel Switch handling. This causes congested Fibre-Channel fabrics, added latencies and port contention as illustrated below. The end result is that the SAN components cannot be used to their maximum efficiency resulting in lowered performance which in digital media applications often means less real-time streams, less clients or reduced data quality through the use of image compression – as is the case with compressed HD SANs.

Traditional SAN: complex & inefficient

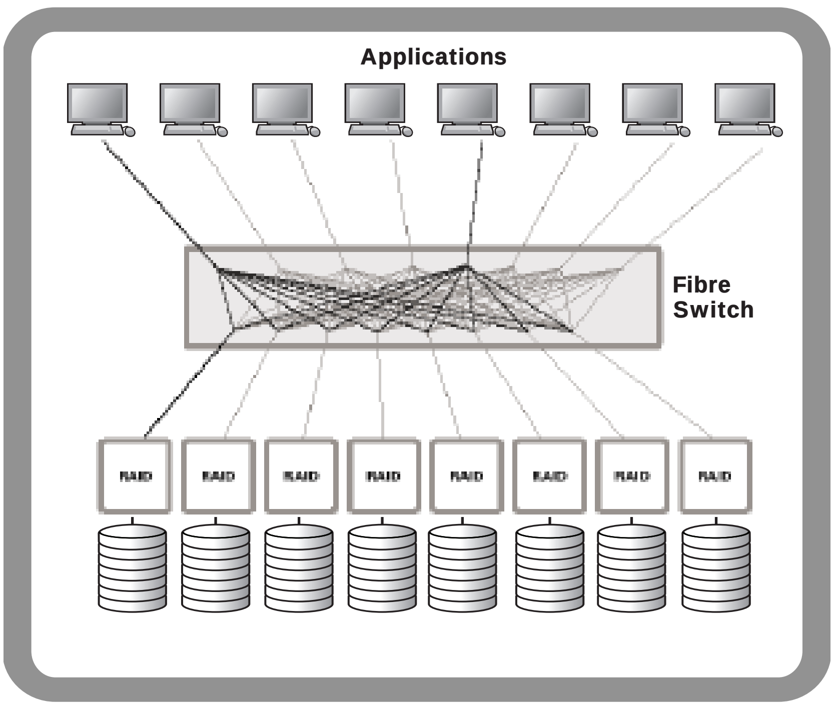

Each host in this traditional SAN configuration shown above is required to perform its own disk striping, which is very costly to the host CPU load. It also adds a high overhead in switch congestion, since each host must span the entire switch. The result is a congested Fibre-Channel fabric, with high switching latencies.

discreet T 514.393.1616 1.800.869.3504 F 514.393.0110 10 RUE DUKE. MONTRÉAL (QUÉBEC) CANADA H3C 2L7

stone shared **SSH8500 SAN Technology**

stone shared is designed to relieve the congestion and contention stresses that are inherent to traditional SANs, while simultaneously reducing the complexity and number of components required to reach a given level of SAN performance.

stone shared is a high-availability SAN with mirrored meta-data for failover, and low latency due to its reduced traffic congestion workflow. The **stone shared** solution uses host parallelism to reduce switching latencies and reduce port contention – as a result it is the ideal solution for high quality, non-compressed high definition workflows. Based on the unique design of the DataDirect Networks Silicon Storage Appliance (S2A), the **stone shared** SSH8500 uses an advanced pipelined, parallel processing architecture, caching, directRAID, ASICs (Application Specific Integrated Circuits) and system and file management technologies.

The stone shared SSH8500 uses an advanced pipelined, parallel processing architecture, caching, directRAID, ASICs and system and file management technologies.

Using zero-latency High Parallel Port Technology (HPPT™), each SSH8500 can orchestrate a coherent flow of data throughout the SAN, from users to storage, at up to 800MB/second (200MB/s per port) per singlet or 1600 MB/second per couplet. This is accomplished through virtualised host and storage connections, a DMA-speed shared data access space, advanced network-optimised directRAID™ data protection and security—all acting in harmony with sophisticated Fibre Channel storage management intelligence embedded within the SSH8500.

The SSH8500 integrates these technologies into a single plug and play device providing simple, centralised and secure data and network management. With its modular design, the SSH8500 can be “coupled” together forming data access redundancy while maintaining fully pipelined, parallel bandwidth to the same disk storage. This modular architecture ensures high data availability and uptime along with application performance. The system provides full bandwidth to all host ports simultaneously, without host striping with its PowerLUN™ technology. The SSH8500’s PowerLUN and host-port parallelism technology removes the need for host-based striping making **stone shared** the ideal storage back-end infrastructure for large clusters, heterogeneous shared and parallel file-system implementations.

The **stone shared** hardware RAID controller is optimised with intelligent latency management, look ahead caching, and distribution and is designed to reduce contention and congestion within the system. Also should stone stared experience a disk failure and require repair, the system has been designed to continue to operate with no performance loss.

stone® shared: optimised & efficient

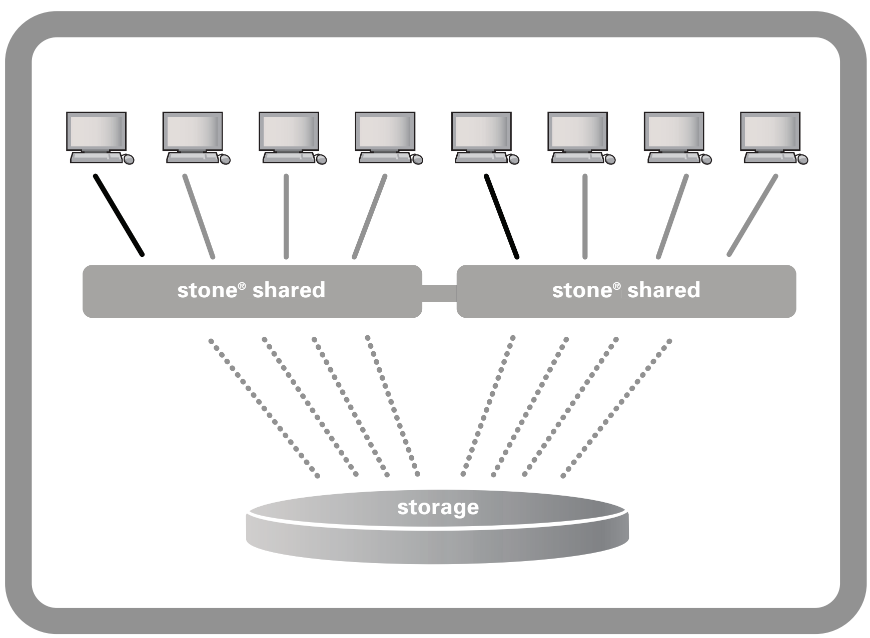

Note: The workflow shown in the above diagram has a single point of failure in the Fibre-Channel Switch. This limitation can be removed by configuring the Fibre-Channel mesh with complete redundant paths, though this will increase costs.

stone shared **SSH8500 Features**

- Single 2U S2A8500 with Four 2Gb/s Ports or Dual 4U with Eight 2Gb/s Ports
- 20 Back end Fibre-Channel Loops Supporting Fibre and/or SATA Disks.
- Up to 1120 Disk Drives, 512 Direct Host logins, and 8192 LUNs supported
- 5TB to 130TB with Fibre Channel Disks, 20TB to 250TB with SATA disks
- Sustained Performance with sequential large block up to 1.5 GBytes/sec
- Full Fibre-Channel Duplex Performance on every port
- Real time Any to Any Virtualisation
- No performance penalties in degraded mode operation, Very fast rebuild rate
- PowerLUN™ 1 Gbyte/sec+ individual LUNs without host-based striping
- Up to 20GB of Cache, LUN-in-Cache Solid State Disk functionality
- SNMP, GUI and Telnet support
- Optional directMirror™ multi-clone mirroring capability

stone shared **Management Software**

stone shared hardware consists of several software components that monitor and manage the state and performance of the SAN

- DataDirect Networks' hardware and system monitoring
- DataDirect Networks' RAID data protection
- SGI file-system's device and system management
- SGI file-system's client and device configuration
- SGI file-system's performance monitoring and bandwidth reservation
- SGI file-system's system security

stone shared **CXFS™ Shared File-System**

stone shared features SGI InfiniteStorage CXFS shared file-system. CXFS is a powerful cross-platform file-system capable of delivering superior performance while scaling massively to meet the most demanding needs of digital media production. Key characteristics of the SGI CXFS shared file-system include:

- High performance: capable of sustaining extremely high data rates (read/write)
- Multi-platform support for IRIX, Windows, Linux and OSX
- Mediation and management of concurrent shared file access
- High workload: multiple processes on multiple hosts directly reading/writing the same file
- Optional guaranteed Rate I/O (GRIO) : guarantees required bandwidths to the user/application
- Highly scalable: up to 18 million TB (theoretical limit)
- Instant, no-copy data sharing among all major file systems
- Reduces the management associated with replicated data
- Simplified RAID administration

Note: Many of the same tools and utilities available for XFS are also available for CXFS.

The CXFS shared file-system offers the highest performance and reliability of any shared file-system. It offers Discreet clients a powerful means of centrally storing their data and sharing it amongst both Discreet and 3rd party applications. Although not fully supported by the **stone shared** hardware CXFS has been designed to scale up to a massive 18 million terabytes of storage and to be able to access files that are as large as 9 million terabytes per file. This design makes CXFS an incredibly powerful solution for digital media applications.

Bandwidth Reservation and SGI CXFS GRIO v.2

Post-production companies are working at increasingly higher resolutions and with ever larger datasets which in turn places greater demands on the bandwidth required to service even a single creative workstation. At the same time, increasing project complexity is requiring more and more people to collaborate on projects concurrently and share the same data. These two trends are placing increasingly heavy loads on shared storage architectures. To address these trends shared storage solutions need to be both scalable and manageable. A key component to managing the resources of a shared storage system is bandwidth reservation.

Bandwidth reservation is the ability to reserve pre-defined amounts of I/O bandwidth for specific systems or applications in order to improve their performance, particularly for systems that may require real-time data-streaming for client supervised sessions. In a shared storage environment situations will arise where multiple applications all ask for access to the same data at the same time. In such cases, the shared storage device will attempt to divide its bandwidth resources between each request and – depending on how many requests there are – each application may well end up receiving insufficient bandwidth to accomplish the task required such as streaming HD media in real-time. This is very different to direct access storage solutions such as stone direct and stone switched which provide the host system with fixed data rates with no interruptions from other systems.

Bandwith Reservation

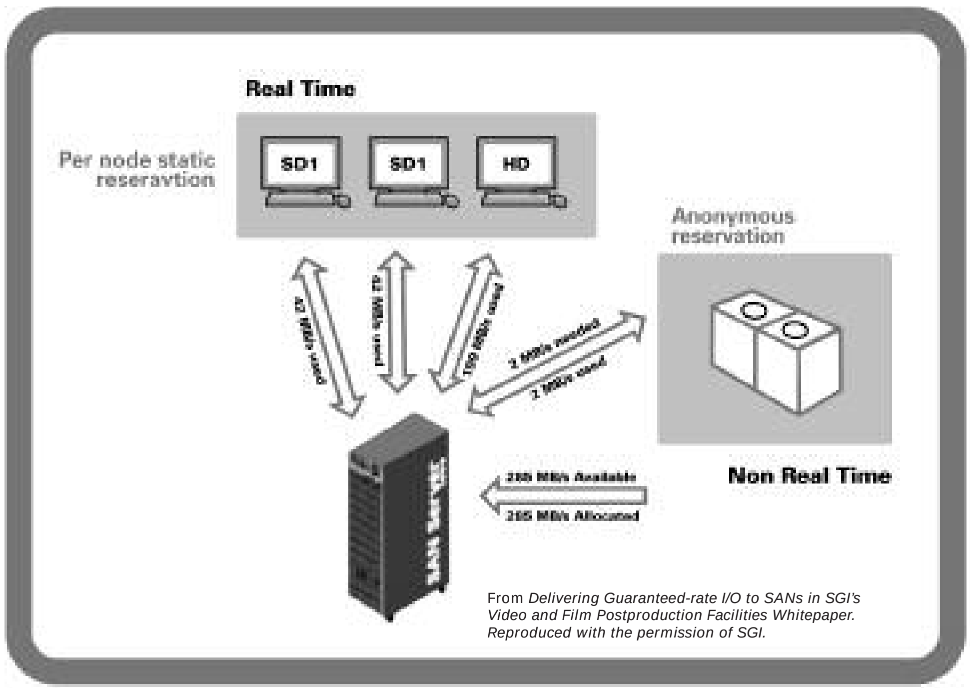

From *Delivering Guaranteed-rate I/O to SANs in SGI's Video and Film Postproduction Facilities Whitepaper. Reproduced with the permission of SGI.*

Non critical applications can also be given priorities to their claim on any remaining available bandwidth after the high-priority applications have been served.

The SGI InfiniteStorage shared file-system CXFS has been specifically designed to deliver guaranteed performance to mission-critical applications ranging from satellite data transmission to on-air broadcasting. To accomplish this CXFS uses a system termed Guaranteed Rate I/O or GRIO, which guarantees that specific required levels of bandwidth will be made available to any pre-defined high-priority applications on the CXFS SAN, enabling them to carry out essential tasks without the risk of interruption from other systems or applications on the SAN.

GRIO can guarantee a dynamic or static stream of data to high-priority applications for any specified period of time. The data-rates required to serve these applications will be deducted from the overall available bandwidth and
the remainder divided among the demands of non-critical systems and applications. Non critical applications can also be given priorities to their claim on any remaining available bandwidth after the high-priority applications have been served. GRIO's central service will throttle their access to available bandwidth (increasing or decreasing it) based upon fluctuations in the total reaming SAN bandwidth available. This capability scales with the SAN environment and does not restrict configuration flexibility.

The ability to assign bandwidth to non-guaranteed applications is called anonymous stream creation. Anonymous stream creation happens when any system requests I/O, but is not registered as a Guaranteed I/O stream. The Anonymous stream can NOT affect any guaranteed I/O stream. For any system to have guaranteed disk I/O, it must have a registered stream with GRIO. Any system not GRIO aware has a lower priority than GRIO aware systems. Systems can be GRIO aware by using the GRIO management software, or via an API in which the Application can make direct calls to GRIO.

Current versions of Discreet systems (**inferno** 6, **flame** 9, **flint** 9, **fire** 6.5, **smoke** 6.5, backdraft 6.5 and **lustre** 2.5) utilise the node level implementation of GRIO. All systems directly connected to the CXFS will require node level registration. The node level registration allows the system administrator to allocate bandwidth to each system on the SAN.

SAN administrators should be aware that there are many applications and utilities associated with the host systems that are not GRIO aware, and that may impact the GRIO service affecting real-time performance. For this reason Discreet does not guarantee real-time on **stone shared** though real-time is possible if the SAN is managed correctly by the facility. The SAN administrator will need to manage and control the use of applications or utilities that are not GRIO aware and that may affect performance. Some examples include the use of the ‘cp’ command to perform a background copy on a system that is trying to play media in real-time since ‘cp’ is unaware of GRIO. Other commands such as detailed directory listings (ls –l) and remove (rm) are metadata intensive, and since the metadata is not managed by GRIO they too may affect the performance of the SAN, as may disk de-fragmentation utilities that are local disk intensive. It is critical that system administrators monitor the possible impact during operations on the SAN.

Note: See SGI White Paper Delivering Guaranteed-rate I/O to SANs in Video and Film Postproduction Facilities for more info on GRIO.