

CommandSoft™

work together

Workgroup Editing SAN White Paper

Work Together — Sharing at the Speed of Light

Editing Stations, The Fabric, The Shared Storage Pool

FibreJet — Stations, Volumes, Projects, and the Database

Sharing at the Speed of Light

Workflow — Multiple-Readers, Write-Access, Mounting

Table of Contents

EDITING STATIONS, THE FABRIC, THE SHARED STORAGE POOL	2
Editing Stations	4
The Fabric	4
The Shared Storage Pool	5
FIBREJET – STATIONS, VOLUMES, PROJECTS, AND THE DATABASE	5
WORKFLOW – MULTIPLE-READERS, WRITE-ACCESS, MOUNTING	6

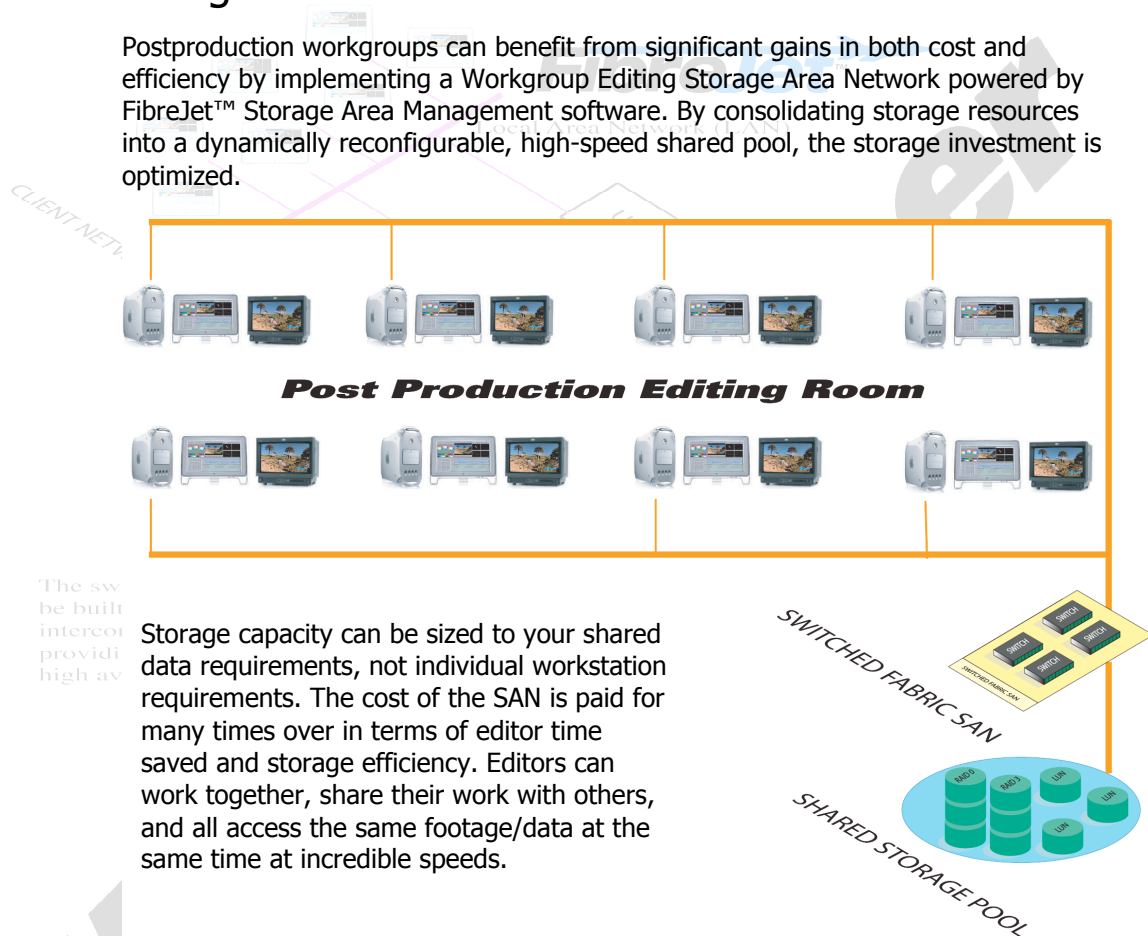
The switched Fabric mesh can be built of a large group of interconnected switches providing complete 99.999% high availability in the paths.

CommandSoft
work together
Sharing at the Speed of Light

CommandSoft Workgroup Editing SAN White Paper

Work Together — Sharing at the Speed of Light

Postproduction workgroups can benefit from significant gains in both cost and efficiency by implementing a Workgroup Editing Storage Area Network powered by FibreJet™ Storage Area Management software. By consolidating storage resources into a dynamically reconfigurable, high-speed shared pool, the storage investment is optimized.



Storage capacity can be sized to your shared data requirements, not individual workstation requirements. The cost of the SAN is paid for many times over in terms of editor time saved and storage efficiency. Editors can work together, share their work with others, and all access the same footage/data at the same time at incredible speeds.

Editing Stations, The Fabric, The Shared Storage Pool

Although this paper discusses Fibre Channel as the network, the theory of operation and software is equally applicable to other network technologies, including iSCSI, SCSI-3, SSA, etc. The Workgroup Editing Storage Area Network (SAN) powered by FibreJet™ Storage Area Management (SAM) software consists of an arrangement of hardware in a network. A group of editing stations is connected through fibre optic cabling to a network fabric, made up of Fibre Channel switches. Then, a pool of Fibre Channel enabled storage is attached to the fabric. These elements are discussed in the following sections.

FibreJet™ software is critical to the workgroup editing solution because it dynamically manages file system traffic so that data is safe and doesn't get clobbered when everyone tries to write to the storage. It does this transparently, dynamically, and doesn't require any administration in the daily use and reconfiguration of the network. Additionally no editing station restarts are required to make changes. Without FibreJet™, this would not happen, as discussed below.

Once everything is physically connected, all the editing stations would be able to see all the storage, as if it were directly attached to a single editing station. When a workgroup of multiple editing stations is actually involved, this is a disastrous situation because all the editing stations would be able to write to the storage, uncoordinated, and in a non-cache-coherent manner. This results in data corruption of the underlying file systems as all the traffic would clobber the storage when the editing stations try to write to the same locations at the same time.

Some SAN systems approach this problem by implementing a method called "shared-nothing". In this approach, the physical network connects everything, but access is restricted either by software, or hardware such as "switch" zoning, such that each editing station would be allowed to see only that portion of storage it was granted access. The following summarizes why this approach is not appropriate for a workgroup-editing environment:

- It defeats the ability to actually share the storage, at the same time, with others. Because of this, the other reasons are moot, but they are relevant because they illustrate the other advantages this system has over traditional Storage Area Management (SAM) solutions.
- Changes are not dynamic, and often require restarting the editing workstation.
- It is a management nightmare, requiring trained, full-time administration, to make changes.
- The granularity of control doesn't match real world use of storage, at the individual file system level, but is limited to an entire storage device, or storage device logical unit (LU) per access permission.

So the main issue with this approach is that it defeats the main benefit of a workgroup SAN, which is the ability to actually share the storage, at the same time, with others. Another equally important issue with this approach is management, and flexibility. With this approach, in almost all circumstances today, changing the configuration, access permissions, or zones actually requires the editing station to do something before the changes will be seen. This inevitably involves restarting the editing station to reflect the changes! Even worse, while the changes are being made some editing stations involved might need to shutdown and thus stop working! Also, simply making the changes is complex and requires a trained administrator to manage the storage system, its allocations, and its permissions full-time. Lastly, these approaches greatly suffer because the granularity of control is often an entire storage device, and at best a single logical unit (LU) of the storage device, as opposed to a file system.

These are some of the reasons FibreJet™ software is the right fit for the workgroup editing SAN solution. It eliminates all the problems as illustrated by the other approaches.


Editing Stations

Editing stations are enabled for the workgroup editing SAN by the addition of one or more Fibre Channel PCI cards and coresponding fibre optic cabling to attach to the network fabric. Then a USB hardware key, configured to enable the purchased software features, is attached to the editing station, and the various FibreJet™ software is loaded from the CD.



The Shared Storage Pool

All the Fibre Channel storage devices are attached to the fabric, preferably using fibre optic, as opposed to copper cabling, completing the SAN.

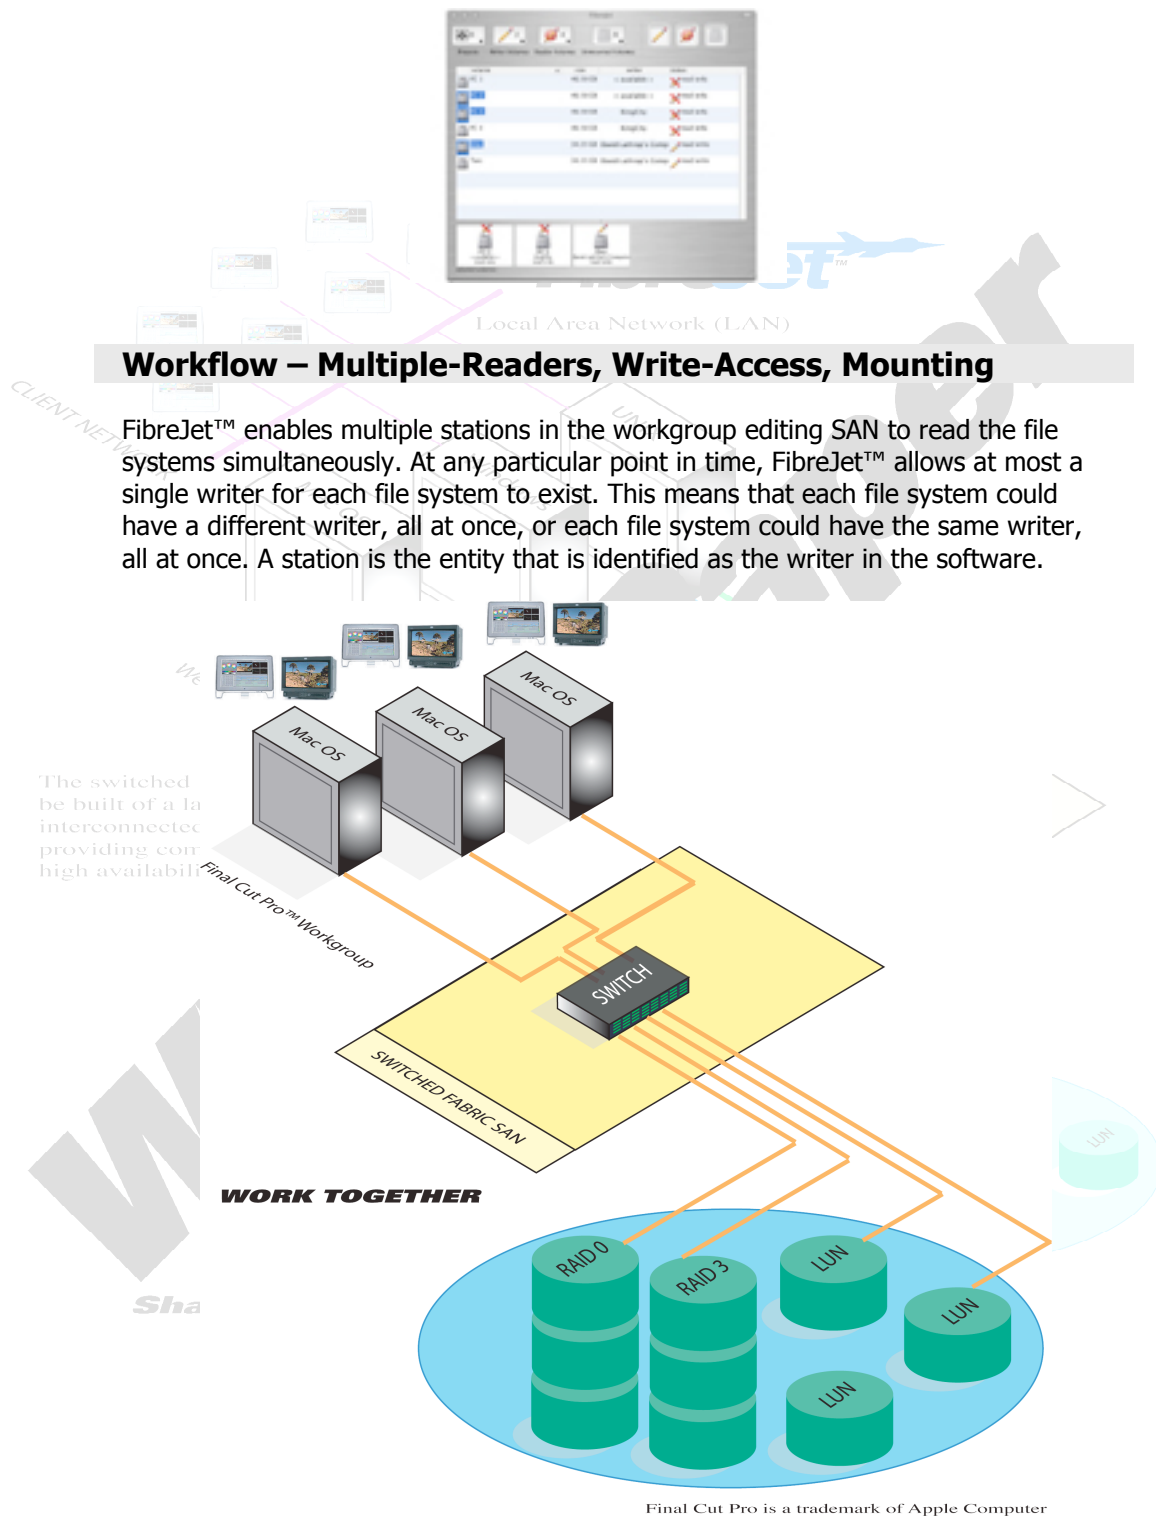


FibreJet – Stations, Volumes, Projects, and the Database

Relevant objects in the workgroup editing SAN are represented in the software:

- Stations represent the editing stations. They are automatically discovered and managed transparently. They show up in the software under the computer's name that was assigned when the machine was set up. However, they are uniquely identified by their primary Ethernet MAC (media access control) address.
- Volumes represent the file systems discovered on the storage. They are an abstract logical entity. For example, a file system may actually exist across several physical storages devices, in the case of RAID-0 (striping) environment. Alternatively, a file system may actually exist only on a small partition (section) of a logical unit (LU) of a storage device. There is a direct one-to-one mapping between volumes and the named file systems you see mounted on your computer.
- Projects represent named and password-protected groups of volumes. If a station does not know the password for a project, it will be unable to see any of the volumes associated with that project. A station may access multiple projects at any time.
- The Database is a special area of the SAN storage chosen when the system is originally setup. This raw-disk partition database contains a cross-platform database implementation that is where the software keeps track of all the stations, volumes, projects, their state, and reconfiguration traffic in the operating network.

The main business of the software is to dynamically control and reconfigure access to volumes. A Volume can have several states, including mounted, unmounted, read-only, read-write, inaccessible.



If a user of a station wishes to place new output to a piece of storage they must first decide on which volume the new output will be placed. The output could be any kind of data, including newly digitized footage, rendered footage, animation, special effect graphics etc...

If the user doesn't already have write-access to the volume in question (as identified in the software), the user simply requests write access using the software. If no one else has write access to that volume, then the user is automatically granted write access, and is identified as the writer in the database for other users to see. If requesting write access from another user, the user may type a message to be sent to the other user as part of the request. If the request is successful, the requesting user is again identified as the writer.

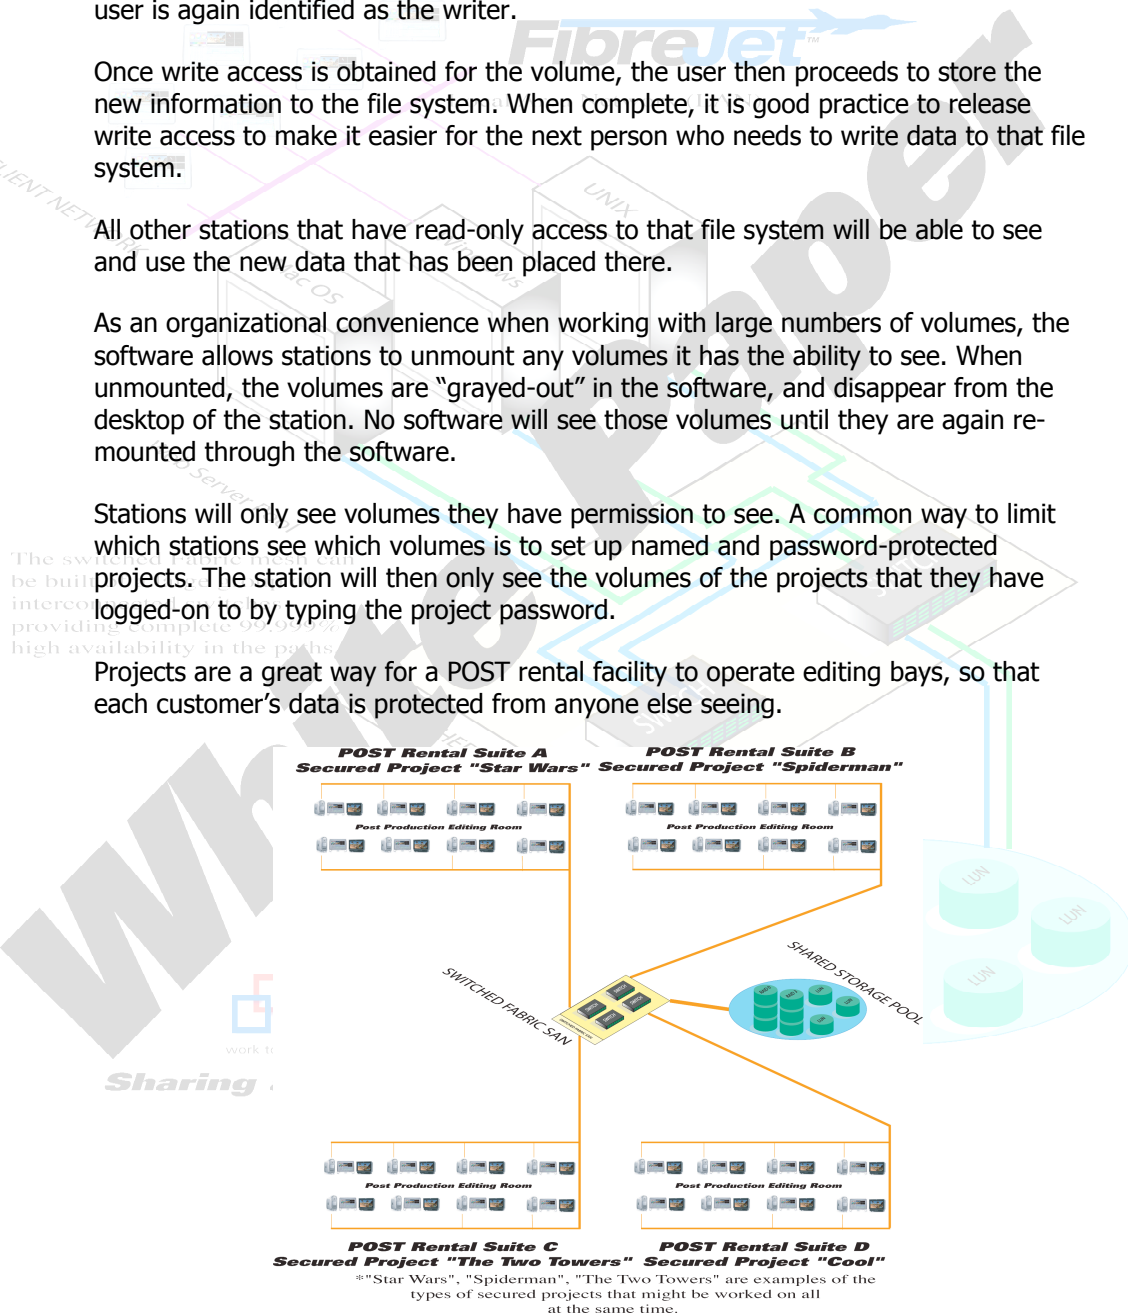
Once write access is obtained for the volume, the user then proceeds to store the new information to the file system. When complete, it is good practice to release write access to make it easier for the next person who needs to write data to that file system.


All other stations that have read-only access to that file system will be able to see and use the new data that has been placed there.

As an organizational convenience when working with large numbers of volumes, the software allows stations to unmount any volumes it has the ability to see. When unmounted, the volumes are "grayed-out" in the software, and disappear from the desktop of the station. No software will see those volumes until they are again re-mounted through the software.

Stations will only see volumes they have permission to see. A common way to limit which stations see which volumes is to set up named and password-protected projects. The station will then only see the volumes of the projects that they have logged-on to by typing the project password.

Projects are a great way for a POST rental facility to operate editing bays, so that each customer's data is protected from anyone else seeing.





Index

access permission, 3

administration, 3

administrator, 4

allocations, 4

cabling, 4, 6

consolidating storage, 2

copper, 6

data is safe, 3

Database, 6

editing station restarts, 3

Editing Stations, 4

Fabric, 4

Fibre Channel, 4

fibre optic, 4, 6

FibreJet™, 3, 4, 7

file system, 4

file system traffic, 3

file systems, 6

flexibility, 3

granularity of control, 3, 4

groups of volumes, 6

inaccessible, 7

iSCSI, 2

logical unit, 4

management, 3

mounted, 7

Mounting, 7

multiple editing stations, 3

multiple projects, 6

Multiple-Readers, 7

network congestion, 4

output, 8

partition, 6

password-protected, 6

PCI cards, 4

permissions, 4

ports, 4

Postproduction, 2

project password, 9

Projects, 6, 9

RAID-0, 6

read-only, 7

read-write, 7

reconfiguration, 3

Redundancy, 4

requests write access, 8

restarting, 3

SCSI-3, 2

share the storage, 3

Shared Storage Pool, 6

shared-nothing, 3

shutdown, 4

SSA, 2

states, 7

Stations, 6, 9

Storage Area Management, 2

Storage Area Network, 2

storage device, 3, 4

storage device logical unit, 3

striping, 6

switches, 4

Trunking, 4

unmounted, 7

USB hardware key, 4

Volumes, 6, 9

Work Together, 2

Workflow, 7

workgroup-editing environment, 3

write access, 8


write to the storage, 3



Write-Access, 7

writer, 7

zones, 4

zoning, 3





work together

Release 2003.01.01

CommandSoft™, Inc.

610 Anacapa Street
Santa Barbara, CA 93101
U.S.A.

Phone: 805-730-7772
Fax: 805-730-7779
Email: info@commandsoft.com
Website: <http://www.commandsoft.com>