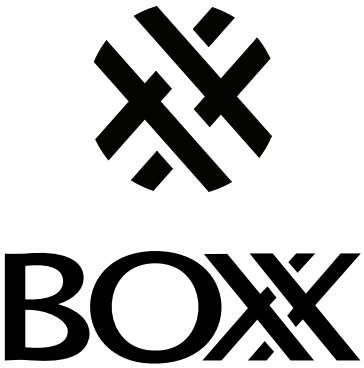

WHITE PAPER

DATA STORAGE:
Making the Right Choice for Your Digital Media Pipeline

A BOXXLabs White Paper from BOXX Technologies, Inc.

Contents

WHAT YOU'LL LEARN IN THIS WHITE PAPER

Nobody has to tell you how important data storage has become for today's digital media industry. The high-resolution assets being created for today's visual effects and motion media design pipelines require much more disc storage capabilities than even just a few years ago. With typical projects taking up multiple terabytes of high-speed storage, choosing the right solution for your data storage needs is a crucial strategy for growing your business.

At BOXX Technologies, we understand this challenge because we specialize in developing high-performance computing platforms for digital media artists around the world. In the same way that we've become experts in the needs of visual effects (VFX) professionals, we have also developed a clear understanding of the data storage requirements for the digital media pipeline.

There is no single data storage solution that's right for every digital media agency. Choosing a solution depends on understanding specific project needs, technology strategies, budget requirements and business goals. That's why we have provided a detailed description and assessment of all major options — both pros and cons — along with each option's appropriate target applications.

After reading this white paper, you will be able to make a more informed decision for developing a data storage solution that fits the way you want to serve your clients, support your pipeline and grow your agency.

HISTORICAL OVERVIEW - THE STORAGE CHALLENGE

For years, solving data storage problems for digital media applications has been an expensive and challenging exercise. Individual artists as well as large facilities have struggled to manage their media storage needs while constantly running into (and trying to push past) the inherent limitations of the technology available at the time.

For instance, just a few short years ago, all non-linear digital video editing systems (NLE's) required expensive, bulky and noisy external disc arrays –often costing more than the PC and editing software combined. In these "early days" of digital video production, the efficiency of digital editing and simple effects compositing was often hampered due to limitations in disc I/O performance and storage capacity. Even very large disc arrays (several hundred gigabytes at the time) that cost thousands of dollars could only provide enough space for two or three hours of digitized and compressed NTSC (television resolution) video footage. It was very rare for all the reels of raw footage for a given project to be digitized and available online at any one time. The work needed to be produced in segments due to the inherent limitations in drive space.

For effects compositing, the workflow slowed to a crawl as multiple layers of full motion video needed to be manipulated in near real-time. Artists lost precious seconds — not to mention their patience — waiting for the application to catch up with their mouse. The slow-down was often caused by the application pausing while the hard drives could send the appropriate frames of video to the screen.

Until recently, high-definition video editing and cinematic visual effects work always required the highest–performing disc storage systems that money could buy. These applications pushed storage capacities to the very edge, and only the largest facilities could afford the massive expense systems. For most other facilities, working in these higher resolutions was at best impractical or usually just not feasible.

For computer animation, the story was largely the same. Using a powerful enough workstation and graphics subsystem, most 3D scenes could be modeled, textured and animated with relative ease. As is still the case today, the process of creating and rendering complex animations was not greatly impeded by the performance or capacity of the workstation's hard drive(s).

Unfortunately, in order to actually view the final, rendered animation sequences, the rendered frames often needed to be moved over the network to a digital disk recorder or "DDR".

DDR's for professional use are essentially stand-alone, networked video "appliances" equipped with multiple-disc RAID arrays and enough disc I/O bandwidth to play-out multiple, individual frames (as real-time "video") to an attached video monitor or tape deck. Different from today's consumer disk recorders found attached to the living room TV, these DDRs were expensive ($4K to $30K), and the workflow often required moving large amounts of data across a slow network before the footage can be seen. This process was often painful and slow. While these devices could easily support the playback of standard NTSC resolution, the ability to play back uncompressed high-definition animations in real-time could be found on only the most high-end and expensive DDRs – or not at all.

For digital film production (both editing and effects), hard disc performance limitations have always been a major source of concern and expense. The bandwidth and storage capacity requirements for working with digitized film are enormous. In fact, until quite recently, performance requirements for many digital film production pipelines were simply unattainable. There were no disc storage systems fast enough to allow the digitized frames to be played back or manipulated in real-time.

PROS AND CONS OF CONSUMER DISC DRIVE TECHNOLOGY

The good news is that advances in the last few generations of consumer hard disc technologies have smoothed the road dramatically for many digital media applications. The cost and technology challenges, at least for some applications, have been largely minimized due to increased speed, greater capacity, and easier implementation. These advances have allowed many digital artists to be more productive without upgrading their workstations from the stock configuration or investing in high-performance, third-party storage solutions.

Digital media applications in this category typically include corporate video, web/interactive media, CAD/CAM, and architectural design. Consumer-level disc technology is generally sufficient for many of these purposes, in both speed and capacity. Standard graphics workstations are usually installed with SATA disc drives, and these are normally sufficient for the average user.

A single consumer-level disc drive can easily store 30 to 70 hours of digital video (DV) while providing

enough disc I/O performance to allow for quick and efficient editing in any number of software applications. Even some multilayer compositing work , while still limited because of disc I/O bandwidth, is now quite commonly performed using stock, configurations with a single, internal SATA drive. Over the past few years, we have begun to see even high-end digital video editing and effects tools for commercial television and other media are making use of advances in SATA drive technology.

However, we need to keep in mind that the storage requirements for editing and compositing uncompressed high-definition and digital film formats still require more performance than standard, consumer-level disc drives can deliver on their own. For example, some digital artists use a combination of SATA disc drives and professional RAID controller hardware to provide the required levels of performance, capacity and data redundancy. This combination of consumer and professional hardware can meet technical requirements while helping to keep costs to a minimum.

The same combination of consumer and professional-level technology can be used in working with computer animation. Playing back rendered, uncompressed animation sequences in NTSC resolution requires about 30MB/s, easily done on any single SATA disc drive. This allows animators to work on practically any PC (including a laptop) and be able to view their finished animations in smooth, full motion. On the other hand, viewing animations in uncompressed high-definition is still beyond the performance capability of any single disc drive. Here again, the combination of multiple, inexpensive SATA drives with RAID controller hardware can enable digital artists to view high-definition animations on practically any PC without the use of a DDR or other expensive hardware.

In short, a number of today's digital media applications can be supported by off-the-shelf, consumer SATA drives or consumer-professional "hybrid" solutions. However, as we'll see in the next section, solving storage and connectivity issues for high-end digital media is another issue entirely.

MASSIVE BANDWIDTH REQUIREMENTS FOR VFX AND DI

At the top of the digital media production "food chain" are a number of applications used to create high-definition content for national television and VFX content for films.

The visual effects produced by these applications are world class. However, these applications also require truly vast amounts of raw storage capacity, data throughput and redundancy - far more than can

be achieved using standard off-the-shelf hardware or any combination of consumer-level and professional technology.

Let's look at a few numbers to better understand what "vast amounts" really means.

"2K" digital film footage (typically 2048x1536 pixels) requires about 13MB to 20MB per frame. At 24 frames per second, a minimum of 20GB of storage is required for 60 seconds of footage. That is about 1.2 terabytes for one hour - a huge amount of data. Now add the requirement for backup redundancy and multiple copies of the footage being used simultaneously in different areas of the post-production facility and the enormous scale of the storage requirements starts to become clear.

For "4K" (typically 4096 x 3112 pixels) digital film pipelines, the amount of disc storage needed for 60 seconds can be up to 70GB. That translates to 1.2GB/second or about 4.2 terabytes for 60 minutes of storage. To put this into perspective, 1.2GB/s is roughly the entire capacity of a standard DVD every 3.9 seconds, or 7 minutes of DV footage every second. Achieving disc I/O bandwidth in this performance regime is no small feat, and it easily pushes the limits of today's most leading edge storage technologies such as fibre channel and SAS (serial-attached SCSI).

In the following sections of this white paper, we will take you through a step-by-step examination of technologies now used in high-end digital media production.

EVALUATING THE TARGET APPLICATION

First, the target application(s) must be evaluated and the specific storage requirements analyzed. The criteria used in making disc storage-related purchase decisions can vary widely, and must be evaluated against all the available storage options.

- Bandwidth – What is the requirement for sustained data throughput (I/O bandwidth)?
- Storage Capacity- How many minutes (or frames) of media need to be accessible and "on-line"?
- Redundancy- How important is fault tolerance to the target application or pipeline?
- Price- Are there economic limitations that need to be considered?

As in most decisions, some desired features must be given up in preference for others.

In addition, keep in mind that different applications within the digital media production pipeline have uniquely different requirements. In general, the key decision-making criteria used in determining appropriate storage solutions can be represented by the following chart, using the indicated relative levels of importance:

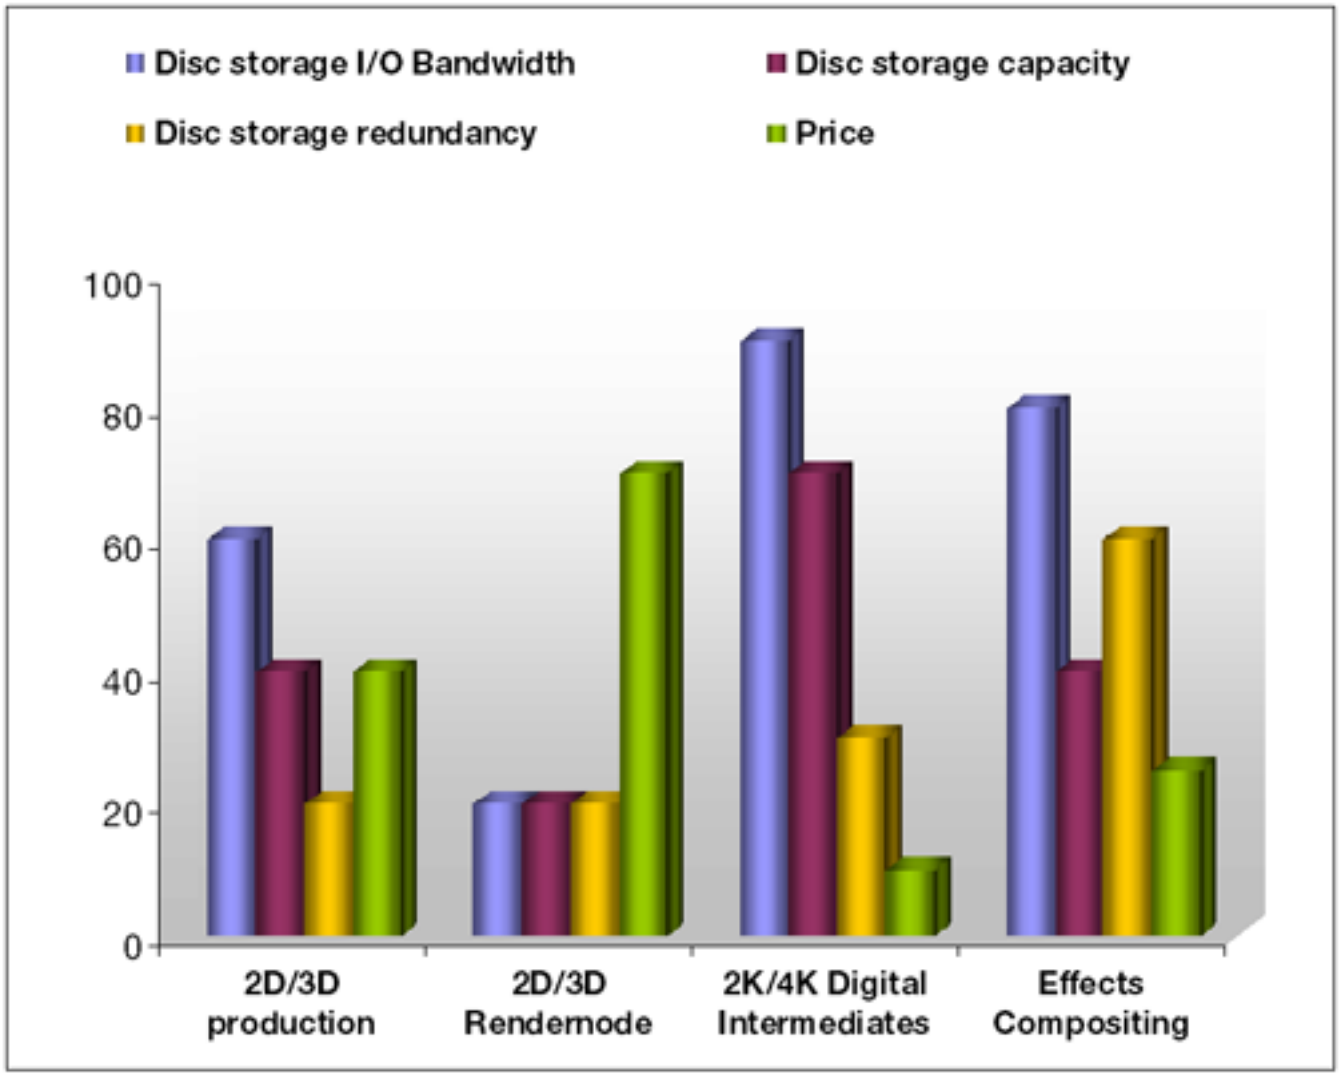

As the chart shows, the criteria for different applications within the digital media production pipeline can be quite varied. For certain systems, price is the most important factor and disc performance is not as important. For other applications, overall disc I/O bandwidth is crucial, and cost is hardly a factor. For still others, redundancy and capacity are paramount – and cost and speed factors are secondary.

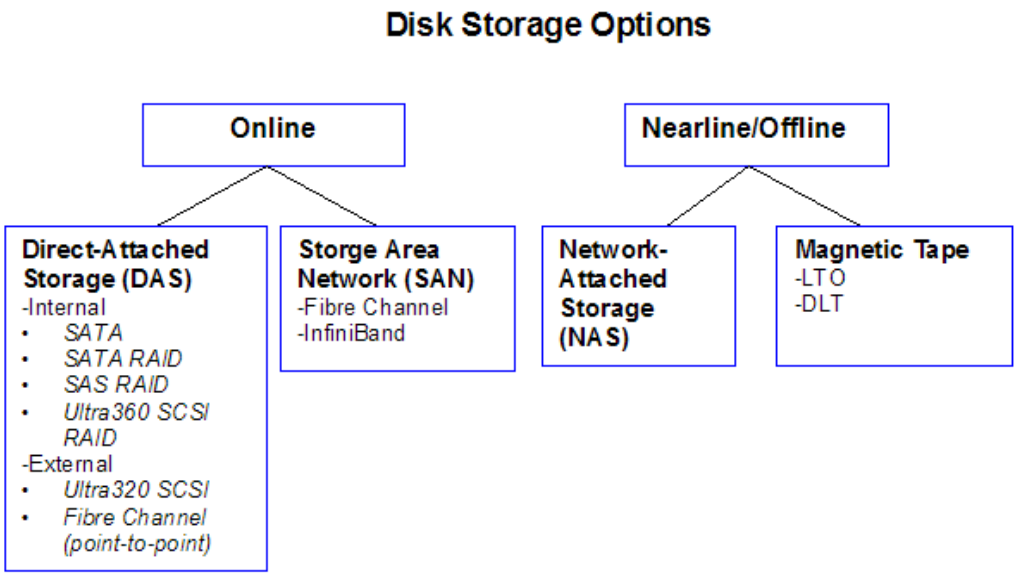

KNOW YOUR STORAGE OPTIONS

There are many different classifications of data storage devices and methodologies. Below is an overview of the major storage technologies currently utilized in the digital media industry.

ONLINE STORAGE

Online storage is the term for any storage system that is immediately available to the host machine/host application at full performance specifications. Online storage may be internal, external or remote. Within this classification, there are two options - Direct-Attached Storage (DAS) and Storage Area Networks (SAN).

DIRECT-ATTACHED STORAGE (DAS) is any disc storage system which is either internally or externally attached to the host machine (PC, workstation, etc).

Internal DAS solutions include some of the most common types of disc storage for digital media production, and can come in several different forms:

- **SATA** – Provides consumer-level performance using un-enhanced, single SATA drives. This is often the most practical and cost-effective configuration for storing the operating system and applications. For most print and web media production, and some DV editing applications, a single SATA drive can also be sufficient for storing and working with the actual media as well. SATA is ideal for render nodes. Maximum throughput is about 40MB/s, with single drive capacities upwards of one terabyte.
  PROS – Inexpensive, easy to install/upgrade, relatively large capacity.
  CONS – Relatively low throughput, not highly scalable, not as reliable as enterprise-class solutions, marginal stability at upper performance limits.
  TARGET APPLICATIONS – Web/print media, editorial style DV, 3D modeling/animation, render nodes, license/render servers.

- **SATA RAID** - Combines inexpensive consumer SATA drives with high-performance RAID controller hardware to achieve high bandwidth at relatively low cost. The 3DBOXX 8300 RTX with 15 internal SATA drives can achieve capacities up to 8 terabytes and up to 700MB/s throughput speeds in a RAID5 configuration.
  PROS – Inexpensive, very good performance, high storage capacity, wide variety of solutions and vendors
  CONS – Not as reliable as enterprise-class solutions, marginal stability for some apps.
  TARGET APPLICATIONS – DV/HDV/HD editing and effects, uncompressed 2K playback (digital dailies, etc), renderfarm framestores, dataset-driven design (BIM, scientific visualization).

- **Ultra360 SCSI RAID** – Provides an extremely reliable and time-tested solution for high-performance applications. Although SCSI technology is aging, it is still considered to be a valid platform for some applications.
  PROS – Maximum reliability, very good performance, proven technology.
  CONS – Aging technology which will soon be phased-out, limited drive sizes.
  TARGET APPLICATIONS – Any applications where legacy SCSI controllers are still being utilized.

- **SAS RAID** – Combines SAS drives and high-performance RAID controller hardware. This is a relatively new technology and is positioned to replace Ultra360 SCSI where maximum speed, maximum performance, scalability and reliability are needed.
  PROS – maximum performance, very good scalability, very good stability
  CONS – expensive, small drive sizes, not many vendors yet
  TARGET APPLICATIONS – 2K/4K DI pipelines, high-end cinematics and visual effects compositing for television and film, large-scale renderfarm storage, studio-class digital asset storage, etc.

EXTERNAL DAS STORAGE SOLUTIONS are typically used to overcome the space limitations inherent with "internal" storage (enclosed in the same chassis as the PC itself). Many digital media applications require the storage performance and capacity that can only be achieved when many (sometimes many dozens of) individual disc volumes are combined into a single, managed system.

External storage enclosures provide power to the individual drives as well as a sturdy, rigid mounting structure. Some external systems include their own disc RAID controller, and can perform all disc management and administration operations independently of the host PC. In other configurations, the external disc array depends on a RAID controller installed as a PCI card(s) *inside* a host PC. There are several cable connector and data transport technologies found in use on external disc storage solutions. External hard disc arrays can achieve extremely high data rates (over 1GB/s) and capacities upwards of hundreds of Terabytes.

- **Ultra320 SCSI** – Provides a reliable, time-tested hard drive technology and storage connectivity platform. Ultra320 SCSI has a maximum throughput of 320MB/s per channel. Multiple channels/cards can be installed into a host PC to achieve very high data rates. External Ultra360 SCSI connections are limited to 10 feet.
  PROS – Maximum reliability, very good performance, proven technology
  CONS– Aging technology/will soon be phased-out, limited drive sizes, limited range
  TARGET APPLICATIONS – Any applications where legacy SCSI controllers are still being utilized

- **FIBRE CHANNEL (POINT-TO-POINT)** – A very high-speed data interface that provides arguably the highest performance and most robust data storage technology available. Fibre channel is also the most expensive storage platform. Typically, an external fibre channel array will include its own controller and can be used to connect directly to a host PC over optical or copper cable up to several hundred feet away. A Fibre Channel array may actually use SAS or SATA drives internally, and is almost always configured as a true

RAID system., Fibre Channel arrays are ideal for mission-critical or ultra-high performance applications where ultimate performance is a minimum requirement. Fibre Channel data rates fall between 2Gb/s (copper or optical) and 4Gb/s (optical only) per channel with virtually unlimited scalability. Throughputs of over 1GB/s are easily attained on this plat form.

PROS – Maximum reliability, ultimate performance and stability, proven technology, ultimate scalability, long cable length

CONS– Expensive

TARGET APPLICATIONS – High-end resolution visual effects for 2K/4K DI, mission-critical, facility-wide media/asset storage, etc.

STORAGE AREA NETWORK (SAN) is typically used for facility-wide, shared storage applications. Commonly, SAN systems are set up to allow multiple users ("clients") to share a single, large pool of data (media) over a managed, dedicated, high-speed network. SAN network topologies are either "switched" or multipath "switched fabric" topologies which use one of two competing connectivity interfaces: fibre channel or InfiniBand.

Both Fibre Channel and InfiniBand SAN solutions typically require several dedicated components: a "server" PC running a SAN management application, a switch (or multiple switches in the case of a switched fabric), and of course a managed RAID-configured disc storage array. These RAID arrays can be built using SCSI, SAS or Fibre Channel drives.

Aside from obvious implications for enhanced collaboration on large projects, SAN topologies are often an appealing option for ergonomic, thermal or acoustic reasons. Large-scale storage arrays are often noisy and can generate considerable heat. SAN technology allows the "client" workstation(s) to be physically located away from the disc storage array, while maintaining high-speed connectivity to "online" media.

SAN systems are also typically the most robust in terms of reliability and fault tolerance. They can be configured to allow continued multi-user access and nominal operation even with simultaneous failed drives and failed SAN management server. SAN storage capacities are virtually limitless, and aggregate data rates can exceed 2GB/s

PROS – Ultimate flexibility, very high throughput and stability, ultimate scalability, multi-user capability

CONS– Expensive, complex to design/deploy

TARGET APPLICATIONS – High-end, high-resolution visual effects for 2K/4K DI, mission-critical facility-wide media/asset storage, etc.

NEARLINE/OFFLINE STORAGE

NETWORK-ATTACHED STORAGE (NAS) is data storage technology which can be connected directly to a computer network to provide centralized data access and storage over ethernet. Using traditional LAN protocols such as TCP/IP over Ethernet, NAS enables additional storage to be quickly added by plugging it into a network hub or switch. NAS removes the responsibility of file serving from other PC's on the network. NAS is often considered "nearline" storage because data stored on a NAS device can be easily transferred via the network to DAS or SAN systems and thereby made available for use at speeds high enough for the creative process.

PROS – inexpensive and scalable, easy to implement

CONS– relatively slow, bandwidth limited to the speed of the IP network, non-realtime

TARGET APPLICATIONS – temporary storage of shared assets, can also be used as a framestore for distributed rendering

MAGNETIC TAPE ARCHIVES: Digital Linear Tape (DLT), Linear Tape-Open (LTO) are various formats of magnetic tape and drive systems used for temporary to long-term data storage and archiving. They all have varying capacities and speeds but are effectively the same in terms of implementation and usage. Tape cartridge capacities range from 40GB to 600GB of data storage and feature data transfer speeds ranging from 5MB/s up to over 100MB/s.

PROS – more economical than disks for archival data, but this is changing since disk capacities are increasing enormously. Tape cartridges are shock/heat resistant.

CONS– can be expensive to implement. Sequential format: locating a specific record requires reading every record in front of it or searching for markers that identify predefined partitions. Tapes stored for long periods must be periodically recopied or the tightly coiled magnetic surfaces may contaminate each other.

TARGET APPLICATIONS – archival of raw media as well as VFX assets/renders and final "master" of finished media

CHOOSE A VENDOR WHO UNDERSTANDS YOUR BUSINESS

As a digital media artist, you have a rare combination of creative talent, technical proficiency and industry know-how — and you should expect nothing less from your technology vendors, whether for data storage or for computing platforms.

At BOXX, we understand your business because it's *our* business as well. We're recognized specialists in digital media, helping artists like you to create worldclass images, effects, animation and architectural visualizations — on time, on target and on budget.

We know that you don't get paid to wait. You're paid to produce. And when time is money, you need technology solutions geared specifically to supporting an efficient, smoothly-running digital media pipeline.

That's why BOXX has worked hard over the years developing an industry-specific range of products, backed by technical expertise, support, and strategic business relationships with software and hardware vendors. Our Innovative Integration process leverages the best of standards-based technology to produce high-performance workstations, render nodes, and custom solutions. We maintain close engineering relationships with all major graphics vendors to make sure that our products leverage these new technologies for customers. BOXX Labs thoroughly tests and optimizes all our platforms with VFX applications to find the sweet spot that allows our users to improve workflow. BOXX also runs advanced R&D projects that push the limits of parallelized graphics processing for rendering in particular.

SO WHAT DOES THIS MEAN TO YOU AS A DIGITAL ARTIST?

It means that with BOXX, you have a vendor who can talk your talk and understands the technology you need in order to do the work you love. BOXX can give you practical, expert advice about how best to run the digital media applications you use every day. BOXX knows VFX.

Feel free to give us a call at 1.877.877.BOXX (2699) or visit our site at www.boxxtech.com. We'll be happy to tell you more about the unique, proven advantages of BOXX solutions.

APPENDIX: DISC STORAGE TERMS

DISC ARRAY – a set of multiple, identical hard disc volumes, typically mounted and interconnected within a single chassis. May be “internal” or “external” to a workstation or server PC, and can consist of any type of disc(s) or connection specification.
E-SATA – (External SATA) – A connection specification (part of SATA II spec) providing up to 300MB/s external connection between disc and controller up to 2 meters.
FAULT TOLERANCE – The capacity for a device or system to recover from a hardware or software failure. Also see “RAID”
FIBRE CHANNEL – A serial communication protocol that provides high-speed, low-latency data transmission between workstations and storage servers, and between storage controllers and individual disc volumes. Fibre channel devices can be connected via either copper or fiber optic cables and provide data bandwidth of 2Gbps to 4Gbps per channel. Fiber optic connectivity between fibre channel devices has become the standard as optical GBICs for switches have dropped dramatically in price over the past few years and speeds for optical connections have increased.
FIBRE CHANNEL DISC – Typically a SCSI disc with a copper, fibre channel connection to the disc controller.
GBIC – (Gigabit Interface Converter) – A generic term referring to a connection interface (a physical plug-in port) between fibre channel storage controllers, switches and client machines. Can provide a connection for either copper or optical fibre channel cables.
HBA – Host Bus Adapter – A fibre channel connection (usually an add-in card) for PC’s which allows connectivity to a fibre channel storage device or SAN. Can be fitted with either optical or copper GBICs.
JBOD – (Just a Bunch of Discs) –Often been used refer to ANY disc array, but it is more formally used to refer to any non-redundant (non-RAID) disc array.
RAID – (Redundant Array of Inexpensive Discs) – a set of specifications that define various methodologies for configuring a multiple disc array. There are many different RAID “levels” providing varying combinations of data I/O performance, fault tolerance and scalability.
RAID0 – (pronounced “RAID zero”) –more accurately known as a “stripe set”, RAID0 is NOT actually a true RAID configuration. RAID0 provides for maximized data throughput (I/O bandwidth) but provides no fault tolerance. Data being stored to the array is broken into pieces and “striped” across all discs simultaneously. In this configuration, each disc in the array stores separate pieces of the data. If a single drive fails, the logical volume will fail and any existing data will be lost.
RAID5 – a common RAID configuration used frequently in the digital media production space. RAID5 provides both very good data throughput performance AND a moderate level of fault tolerance. If a single drive fails, the disc array can continue to function (although in a degraded, slower capacity).

APPENDIX: DISC STORAGE TERMS

RAID50 – a common RAID configuration used extensively for digital media production which combines multiple RAID5 logical volumes across a RAID0 stripeset. - Provides expanded data throughput performance over RAID5.

SAN – (STORAGE AREA NETWORK) - A generic term used to refer to a dedicated, high-speed network connecting multiple "client" workstations to a single, shared storage array.

SAS - (SERIAL ATTACHED SCSI) – Next-generation SCSI technology. This combines traditional SCSI disc discs and fits them with a serial connection interface and provides greater flexibility, performance, reliability, and connectivity than the current parallel SCSI mainstream interface. The SAS specification allows up to 3Gbps throughput per channel and virtually unlimited scaling via the use of SAS expanders

SATA- (SERIAL ATA) The successor to parallel ATA for data storage; a point-to-point connection that delivers full throughput to each storage device by allowing a single controller to manage multiple ports. Serial ATA technology starts at transfer rates of 1.5Gb/s with a roadmap up to 3.0Gb/s and ultimately up to 6.0Gb/s; enables affordable disk-based storage with thinner cabling than Parallel ATA.

SCSI (SMALL COMPUTER SYSTEM INTERFACE) - An aging, high-performance hard disc and interconnect specification. The current spec is Ultra320 SCSI, which supports up to 14 hard discs and a maximum of 320MB/s per channel.

STORAGE CONTROLLER – A device (typically a card installed inside a PC) designed to control the activity and data flow to-and-from a hard disc or hard disc array. It can provide control for ATA, SCSI, SATA and SAS disc technologies.

STRIPE SET – see RAID50.

SWITCH –A networking device that performs transparent bridging between network nodes and other devices at up-to the native speed of the hardware technologies being used. A switch can connect Ethernet, Fibre Channel or other types of packet switched network segments together to form a network.

SWITCHED FABRIC - A high-performance computer network topology with many devices connecting with each over redundant paths via network switches, used in some storage area networks and other high-speed interconnects, including Fibre Channel and InfiniBand. Typically, the more devices or "nodes" that are added, the faster the aggregate bandwidth of the network fabric.

VOLUME (OR DISC VOLUME) – A single disc storage device. There are two kinds of disc volumes – "physical volumes" and "logical volumes." A physical volume is just that — a single physical hard disk unit. A "logical volume" is seen by the host device or operating system as a single disc, but it may actually be comprised of many individual physical disc volumes. Most RAID controllers combine multiple physical discs into a "logical volume" which is seen by the operating system (and by the PC user) as a single, large disk drive.

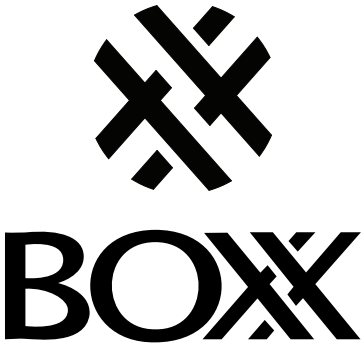
BOXX