

# The Ever Growing Challenges of Data Storage

Electronic data storage needs continue to grow. As your client's organization produces more information in electronic format, storage space is becoming increasingly important. Managing data storage for performance, integrity, and scalability is the next summit in Information Technology management and planning.

It wasn't long ago that having a single volume in the terabyte size range was rare for extremely large organizations. With the advent of IDE RAID and SATA RAID capabilities, large storage systems are within reach of medium to small businesses.

Let's put things into perspective – how much space is 1TB?

Number of Bytes	What that relates to
1 Byte	One character (letter or number)
1KB (Kilobyte) 1000 bytes	3 or 4 typed manuscript style pages
1MB (Megabyte) 1,000,000 bytes	Average size of a novel (300-400pgs); 1 diskette
1GB (Gigabyte) 1,000,000,000 bytes	Approximately 20 sets of encyclopedias
1TB (Terabyte) 1,000,000,000,000 bytes	A small library (approx. 5,000 books)

To get the best performance and reliability from any storage space, strategic storage planning is essential. This month's technical article will review the importance of the file system and planning considerations.

## The File System's Role

The file system's role is a layer above the storage device(s) itself. The file system manages the individual allocation units of the volume and provides hierarchical organization for the files. Managing the allocation units of the files requires algorithms that will know where to write file data and have a method of verifying that the data was written correctly.

Hierarchical organization is the logical formation of directories and underlying structures. For instance, a storage volume that has millions of files on it will have specific data that describes the directory or folder structure of where these files belong. This directory or folder structure has integrity checks and balances to ensure that the indices reliably point to the user data.

Today's file systems track more than just the name of the file or directory structure. Additional information called Metadata is also stored. Metadata is data about data. Essentially, the file system is saving more details about your files and is storing this along with attributes of the file. Some file systems record only the minimum of metadata (file name, size, time and date, start address), while other file systems record more information (file name, size, multiple time and dates, security details such as Read/Write/Execute/Delete privileges).

Some file systems are designed for specific hardware and storage media. For instance, the file systems used for CD-ROMs are quite different than those for floppy diskettes. Forcing these file systems on other media may be possible, but not practical. So while specific storage media, such as CD-ROM, DVD-ROM, magnetic-optical disks, and tape, have unique file systems, hard disk and hard disk storage systems can work with many different file systems.

Understanding these extra features of file systems will help in choosing the best one for the needs of the volume.

## **File System Considerations**

During server planning, more time and research is spent on hardware, data space requirements, and application specifications than on how the data will be stored. The file system can become a low priority during the planning stages of a file or data server because the file system is inherent to the operating system. Sometimes it is assumed that this is best fit. However, your storage requirements may call for a more robust method of data organization on the hard disk(s). Investigate whether the operating system you are planning to use allows the other file systems to be used.

If you have a choice of file systems, here are some requirements to consider:

- Volume Size
- Estimated number of files on the volume
- Estimated size of files on the volume
- Shared volume requirements
- Backup Requirements

### **Volume Size**

Volume size is an important place to start for planning. However, this is only the start since strategic planning involves scalability—can it grow as the need arises without interruption of service to the users? The axiom of filling free space is all too true for data volumes. It is not uncommon to add a terabyte of storage and in six months it's already half full.

Two terabytes (2TB) has become the initial hurdle for many file systems. This limit starts with the SCSI command set being limited to 32-bit logical block addressing. Therefore, a single SCSI LUN using 512 byte block size cannot access over 2TB. File systems that have been used on these systems have been 'adjusted' to handle extremely large volumes. However, volumes that are nearing the 2TB limit may be stressing the limits of the file system.

### **Estimated Number of Files on the Volume**

The next item to plan for is the number of files that could potentially be stored on the volume. Earlier we discussed Metadata and how the file system uses this to describe the files that are stored. This means there is going to be a certain amount of volume space used by the file system just to manage the files that are there.

File systems that are not built for excessively large directories will slow down applications that access them. This can adversely affect users that have thousands of files on a volume that has millions of files.

### **Estimated Size of Files on the Volume**

The next consideration is the sizes of the files that will be on the volume. Organizations that are running large database servers usually have the need to be able to pre-allocate very large files in the gigabyte range of sizes. The file system and operating system need to be able to handle this level of input and output. For these types of enterprises' systems, expectations are high for performance and integrity. Will the file system be able to handle those extremely large files?

### **Shared Volume Requirements**

There are mixed environments in many organizations today. Some organizations may have three or four different platforms of computer systems; from mainframe systems to 64-bit Sun machines, from Apple desktops to Intel based machines. Some of these systems may share storage space. Will the volume support mixed data types? Additionally, will the operating system that manages the file system allow for different types of data streams to be accessed simultaneously?

### **Backup Requirements**

Large volumes present a challenge for backup procedures. Due to the amount of data, restorations can take days. There are some file systems that have 'Snap-shot' technology incorporated into the backup software. This technology saves critical file system metadata. This, along with incremental file backups, is part of entire system scheme of data archiving.

These considerations should be matched with hardware specifications to get the best performance, integrity, and growth capability.

## **Ontrack's Recovery Capabilities for Large Volumes**

Despite the best planning, failures do happen. Ontrack monitors the technological advancements in the storage industry. This includes research and development in new hardware and file systems. Ontrack's software development staff continues to provide proprietary recovery software for use in all our data recovery labs worldwide.

The recovery tools that Ontrack's engineers use have been updated to meet the 2TB barrier. For the past year Ontrack has seen an increase in large, multi-terabyte size volumes. Due to Ontrack's dedication to provide complete recovery solutions, we have been successful in these types of recoveries.

In many cases, Ontrack's Remote Data Recovery<sup>®</sup> has become the standard recovery process for these large volumes because shipping in the drives is impractical. Even if a RAID configurations is lost or if one drive has failed, Remote Data Recovery can retrieve the original data.

As mentioned previously, terabyte volumes are becoming more common. If your client calls you because of errors or problems accessing large terabyte volumes, call Ontrack for assistance. Your sales representative and a qualified engineer will discuss with you and your client all of the options available to get the volume accessible in the quickest manner.

Your client's storage needs will continue to grow and partnering with Ontrack for recovery services means that we'll be there if there is ever a data disaster.

## **References**

Microsoft Volume Recommendations

[http://www.microsoft.com/resources/documentation/WindowsServ/2003/all/deployguide/en-us/Default.asp?url=/resources/documentation/WindowsServ/2003/all/deployguide/en-us/sdcbc\\_sto\\_cokp.asp](http://www.microsoft.com/resources/documentation/WindowsServ/2003/all/deployguide/en-us/Default.asp?url=/resources/documentation/WindowsServ/2003/all/deployguide/en-us/sdcbc_sto_cokp.asp)

Linux Files Systems Overview

<http://linux.org.mt/article/filesystems>

File System IO Benchmark Tester

<http://www.iozone.org>

SCSI – Protocol Changes

SAS Protocol/Architecture – SBC-3 breaks 2TB barrier—

[http://www.scsita.org/aboutscsi/sas/tutorials/SAS\\_SCSI\\_upper\\_layers\\_public.pdf](http://www.scsita.org/aboutscsi/sas/tutorials/SAS_SCSI_upper_layers_public.pdf)

Data limits – From Infotrend, a RAID manufacturer

<http://www.infortrend.com/document/htm/FAQ-max.htm>