

## Introduction: The Need for Centralized Backup/Restore

With critical data dispersed across networks, the need to provide a complete solution for enterprise backup/recovery has become a recurring topic in today's computing environment. Successful retrieval of important data is critical not only to meeting an organization's current needs, but also to measuring that site's ongoing stability. A study by Freeman & Associates noted that sites unable to recover critical data within 10 days of a disastrous loss typically are bought out or sold within one year of the data loss. According to the same study, since half of the disastrous data losses result solely from human error, no network is beyond the risk of a technical disaster. Thus, prudent organizations, by necessity, must focus on the most effective and secure ways to recover from such losses.

In addition to being able to recover from disastrous data loss, an enterprise backup/recovery system also should boost productivity at a site and facilitate its daily operation. Today's Information Technology (IT) administrator spends increasing amounts of time tracking data: from chasing backup tapes to restoring various increments of data, to making room on servers for more data. Yet while the demand for IT service is increasing, the IT administrator is being asked to do more with fewer resources. Companies no longer "staff up" data centers. Instead, the emphasis is on decentralizing the computing resources and reducing the level of central staffing.

This trend toward reduced staffing runs counter to the importance of data security as well as to its cost. A recent NCSA survey predicts that the cost of rebuilding lost data will increase from \$2 million in 1996 to more than \$10 million by the turn of the century.

Fortunately, with today's technology, it is possible to design a secure backup/restore system that backs up, restores, and archives data from a central location, across a distributed network, whether the central repository is in the next room or across the country.

To appreciate the importance of such a system, one must review the methods that until recently were employed for backup. Indeed, in some sites, these methods are still implemented today. In addition, one must consider the best ways to secure decentralized data. This white paper will show how BBars addresses such issues, and how the BBars approach has revolutionized backup/restore technologies.

## The Traditional Backup Process

Traditional backups use a “Media Rotation Scheme.” The administrator inserts a tape to a drive on each server, waits for the tape to back up the drive’s data, labels the tape, and then stores the tape in a safe location. Each day, the administrator repeats the procedure, using a fresh set of tapes. The tapes then are rotated out of storage when their retention periods expire. Unfortunately, this traditional method of backup relies too heavily on the length of time the media are kept on file, and on the amount of human intervention involved.

For example, because media are rotated out depending on the media’s creation date, the rotation scheme is non-discriminant to the data on any particular tape. As a result, all of the files on a tape rotate at the same time: when the tape is rotated.

For critical files, such a rotation could be disastrous. If that critical file is removed from the server for any reason (disk cleanup, someone deeming the file unnecessary, or intentional deletion by a disgruntled employee), that file is no longer available for backup. And once that file’s backup media are rotated out via the tape rotation scheme, the file itself is gone, never to be recovered. Even if the file resides on three tapes (using a GFS rotation: Grandfather, Father, and Son), all tapes will eventually be rotated out and the critical file will be gone.

The risks of such manual intervention are exceedingly high. Manual intervention causes half of all disastrous data loss. Quite simply, the more manual involvement, the greater the chance of disaster. It is implausible to suggest that humans will not make mistakes, especially in today’s environment where one must accomplish so much in a limited time frame. In the backup process, these human lapses mean that tapes will be mislabeled, misplaced, or rotated out by accident.

## The BBars Difference

Given the increasing demands for reliable data backups, today's IT administrators are being forced to reconsider the traditional backup scheme. To be acceptable in today's computing environment, a backup/restore system must qualify as a comprehensive, secure enterprise solution. As a result, it must meet certain requirements:

- Media-independent file retentions
- Central data storage
- Simple recovery
- Flexible security
- Comprehensive reporting
- Scalability and compatibility
- Ease of use

In addition, today's backup/restore system should be able to accomplish tasks that older systems were never asked to perform. Among these tasks is the management of the media library itself. With awesome amounts of data being stored, a backup library has the potential to grow to colossal levels. Thus, the backup solution selected for a site should include management options for the output it creates.

BBars has been designed to handle all of these requirements. It provides centralized storage, independent retentions, flexible security, and a host of other features sorely needed in today's backup/restore environment. Further, BBars has been designed to keep media inventory to a minimum by enabling the site to utilize each tape to its fullest capacity.

The following pages briefly describe how BBars resolves each issue to implement an enterprise-wide solution for today's backup/restore needs.

## *Media-Independent, File-Based Retentions*

The length of time a file is retained in the backup system should be determined by the administrator, not by the constraints of the backup system. Unlike the archaic method of backup previously discussed, BBars provides a “file-based” approach to retention. The administrator determines how long each file will be retained, as well as how many versions of each file will be retained in backup. For companies that must store 401K data or other legally required files, this feature is particularly critical.

The administrator also has the option of excluding files from backup. These file specifications are definable at any path level, from the server/workstation level all the way down to the individual file level. In addition, these file specifications are applied automatically, even if the user is accessing the workstation/server at the time of backup.

## *Central Data Storage*

Automated, centralized backup and storage has always been the ultimate “safe” solution for an organization’s security. But many backup/restore specialists have been trained to think that such centralized backup is no longer possible. Given today’s decentralized computing environment, and the restrictions on data transfer due to bandwidth constraints, these gurus believe that a centralized backup is not feasible.

Fortunately, this conventional wisdom is simply not so. B&L Associates, Inc., as specialists in the backup field for more than 25 years, has made centralized backup a reality. To handle these demands, we use three procedures: Modeling, Take Two, and Full Creation.

**Modeling** uses a *logical backup* rather than a physical backup. Logical backups reduce the physical demands on a network. As an example, assume that file ABC is resident on Server1 and Server2, both of which are backed up daily at 3 a.m. and 5 a.m., respectively. BBars will physically back up file ABC on Server1 at 3 a.m., and enter the necessary information in the BBars database. At 5 a.m., when BBars starts to back up file ABC on Server2, it recognizes this file (by parsing the file itself, rather than by its name, date, and time stamp) as being identical to the file ABC from Server1. Thus, another physical backup of the file is not necessary. As a result, rather than physically backing up

ABC again, which would increase network traffic, BBars inserts a pointer in its database to remind itself that file ABC is resident on both Server1 and Server2. Since the comparison of files uses internal data (rather than relying solely on a file name or date/time stamp), BBars ensures the integrity of its “logical” backups. The odds of any logical file being somehow different from its related physical file is  $2^{64}$ , or 4,000,000,000<sup>2</sup>.

In the **Take Two** process, two backup copies are created from the same *physical* backup. This allows one copy to be stored onsite and one copy offsite. Where once an administrator would have to create two backup cycles for two copies of identical data, BBars now eliminates the problem: One physical copy generates two storage copies.

The **Full Creation** process allows BBars to create a complete or full backup from the various incremental backups it has previously stored. Studies show that, on a daily basis, only 3-5% of a disk’s contents change. Thus, each time a full backup is created, it most likely is duplicating much of its previous work. Because BBars knows exactly what each server/workstation looked like at the time of its last backup (or on any given day, for that matter), there is no need for such duplication. Full backups never have to be taken across the network unless specifically requested. If the need for full backup over the network diminishes, the backup process traffic can be reduced by 90%.

### *Simple Recovery Processes*

File backups do not, by themselves, guarantee data protection. Restoring lost or misplaced files from the backup is equally critical. Backup/restore systems should not only be in place for disastrous data loss, but also for simple file and directory recovery as part of the normal work routine. Its unique backup mechanism and comprehensive file structures enable BBars to track every file version and its storage location automatically. As a result, recovery has never been easier.

The BBars administrator may request any BBars backup version for retrieval. The request may target any version of a file, or multiple files from a path, or all the files in a directory. Requests also may be qualified by date/time. They may be limited to certain files, such as “add only files not resident” or “add only archived files.” Or, if security permits, the request may restore files from one workstation/server to another workstation/server with no additional effort.

Recovery from a disastrous data loss is equally impressive. If BBars is notified that a disastrous data loss has occurred, one click of the mouse enables BBars to automatically specify the tapes needed for restoring. If there is an optional tape library attached, BBars can start loading these tapes automatically. Servers can be restored with no manual intervention whatsoever. In addition, BBars can be directed to use only its offsite versions for restoration.

### *Flexible Security*

Security is an essential part of any backup system. If security is not designed properly, manipulations within the system, intentional or unintentional, can destroy a site's backup/restore process. Yet, with a central backup/restore system, administrators also may desire some flexibility. Certain responsibilities may need to be assigned to local administrators or to the users themselves.

With BBars, comprehensive security features allow the system administrator to control all access, from deciding who can use the system to which workstations and features they may operate. Access is controlled at all levels: every BBars menu, every function under every menu, and every report generated from the system. The same user is also allowed to have different user profiles at different workstations.

BBars includes an efficient setup for any security scheme. There is no need to repeat specifications for each individual user, for example. Once a profile is defined specifying the type of access permitted, it is simply applied by name to users who are allowed such access. In addition, like workstations/servers, users can be grouped together if they inherit the same specifications. This eliminates the hassle of reapplying changes to every single user, because the change is inherited throughout the specified grouping.

### *Comprehensive Reporting*

Regardless of size, it is critical that administrators know which backups were completed, which are pending, and which were aborted. It does the administrator no good to spend time trying to restore a file only to discover that the file, for whatever reason, was never actually backed up.

With BBars, an activity report is on the administrator's desk every

morning, online or on paper, tailor-made to the administrator's preferences. The report can be customized to show a certain time frame, to show any or all workstations/servers or groups, to show all details or only a subset of detail. It can also show information on pending, completed, or aborted backups and restores, or show every file of every station that was backed up or restored. The amount of detail is up to the administrator. Moreover, BBars can provide 20 other reports covering any and all aspects of its operations.

### ***Scalability and Compatibility***

One byproduct of the trend in decentralized computing is the proliferation of workstations and servers. Thus, a fundamental requirement for any type of system is its scalability. For a backup/restore system, scalability is a prerequisite for service in an enterprise. It must grow with the enterprise, but avoid the "growing pains."

BBars meets this requirement. It is completely scalable. Adding a new workstation or server to BBars is simple. Just install a small agent piece on the workstation/server, define that station/server to BBars, and the new machine is now ready to be backed up or restored. Users never have to worry about outgrowing BBars.

BBars also supports numerous operating systems, including Windows, Windows95, Windows NT, OS/2, Novell NetWare 3.1 and 4.1, SVR4, Unisys UNIX, AIX, Linux, Ultrix, SCO Unixware, Solaris, and SunOS.

### ***Ease of Use***

Because BBars is completely automated in every aspect, once it is installed and configured to your site, BBars can operate without any manual intervention, if desired. However, should manual intervention be needed, BBars can be administered from any Windows (3.1, 95, or NT) workstation defined to the BBars system. If security levels permit, backup and restore can be implemented to any workstation/server from any workstation/server. Controlling and moving data has never been easier.

## *Tape Library Management Aids*

Tape handling is a necessity with most backup/restore systems. With some, it is a significant problem. Since new tapes are created every night for hundreds or thousands of machines, storing the media becomes a difficult task. BBars addresses this increasing problem with an application called TapeMerger. TapeMerger combines tapes whose capacity falls below a certain threshold set by the administrator. For example, if threshold capacity is 30% full, TapeMerger automatically merges the tapes falling below 30% to one tape. Merging not only alleviates the problem of tape overflow in the library, it also transfers data to new media, which prevents tapes from sitting idle for long periods of time. For the end user, this is a great benefit.

## **Summary**

Today, as distributed, open systems reach mission-critical mass, organizations are searching for ways to store and control decentralized data. The optimum solution is an automated, centralized backup/restore operation. For this optimum solution to be effective, however, certain requirements must be met. Backups must be timely, flexible, and media-independent. They must be secure. Restoration must be simple and effective. The system must be scalable and comprehensive.

BBars provides such a system.

Although market hearsay contends that centralized backup cannot be timely, the BBars technology demonstrates the fallacy of such beliefs. BBars takes backup and restoration to new levels. With its centralized data storage technology, its thorough reporting, flexible security, ease of use, and features designed specifically for enterprise backup/restore, BBars is the correct solution to an increasing problem.

## About B&L Associates, Inc.

Since 1975, B&L Associates has been providing operations automation tools for information systems. With a user base spanning 19 countries, our results have been impressive. Because of our partnership approach, encompassing implementation to customization, we have achieved a superior degree of customer satisfaction among our user base. To obtain more information about B&L or our firm, call 1-800-OK-BANDL or email us at [info@bandl.com](mailto:info@bandl.com).