Backup You Can Trust℠

The BRU Advantage

The following information is offered to detail the key differences between the TOLIS Group, Inc. BRU™ data protection products and those products based on the traditional tar and cpio tape formats.

The Reality of Backup: Simply stated, a backup must be accurate in order to successfully recover data. Though true and simple in concept, successful execution is a rather complicated process. Depending on the system topology (local or network), there are enough potential points of failure in the data path wherein data can be corrupted to mandate measures be taken to protect it. Data is pulled off the disk farm onto the bus to bring it into memory, it comes out of memory onto the bus again, and travels to the HBA (Host Bus Adapter), across the SCSI cable, and into the input buffers of the archive device (tape, most usually) and any other devices on the SCSI bus (Figure 1.). The belief that filesystem data ultimately entering the archive device cannot/has not been corrupted is a dangerous one. Equally dangerous as well, is the belief that the archive device's ECC (Error Correction Code) is sufficient to provide protection. Garbage in, garbage out.

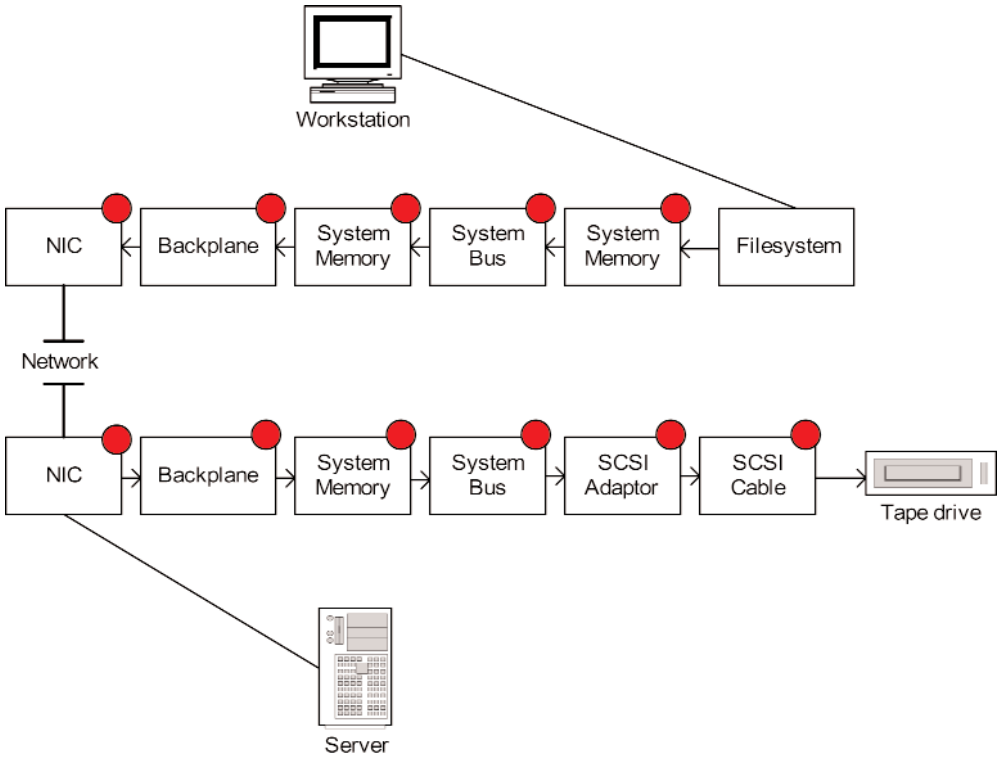

Figure 1. Areas wherein data can be corrupted before archiving

BRU's sole purpose is to deliver reliable backup and recovery, while simplifying that which has become a complex task in today's computing environment — protecting one's data. BRU accomplishes its purpose, and the unique tape format implemented in all BRU products forms the basis of its leadership reliability.

Depending on the system topology (local or network), there are enough potential points of failure in the data path wherein data can be corrupted to mandate measures be taken to protect it.

Assuring Back Up Reliability: BRU incorporates mechanisms that provide accountability of each and every bit of data to be backed-up. Without the checks and balances of this accountability, the validity of data stored for future access cannot be assured. Taking measures to protect data is necessary to assure one's peace-of-mind.

BRU applies 32-bit checksum technology to both the metadata and file data coming off the filesystem and places the checksum into the header of each buffer block. These buffer blocks are then assembled into blocks written to tape. The administrator has the ability to define the block size to fine-tune the backup system for optimal performance. The checksum value in each buffer block is the critical factor used during the verification operation to validate the accuracy of the backup.

Solutions based on the `cpio` format implement checksumming on the metadata only. The actual file data is not calculated into the checksum value; therefore corrupt file data cannot be detected. Checksumming is not implemented in products based on the tar format (see `tar` verification below).

The "server" component of BRU Server, a BRU backup solution for large heterogeneous system networks, is implemented on either a low cost Linux system (x86, PPC) or on a Mac OS X System. These ultra-reliable operating systems complement the ultra-reliable BRU technology, and the pair serves as the foundation for a rock-solid backup system (Figure 2.). Additional support of other UNIXes will follow, check TOLIS' website for newly supported platforms.

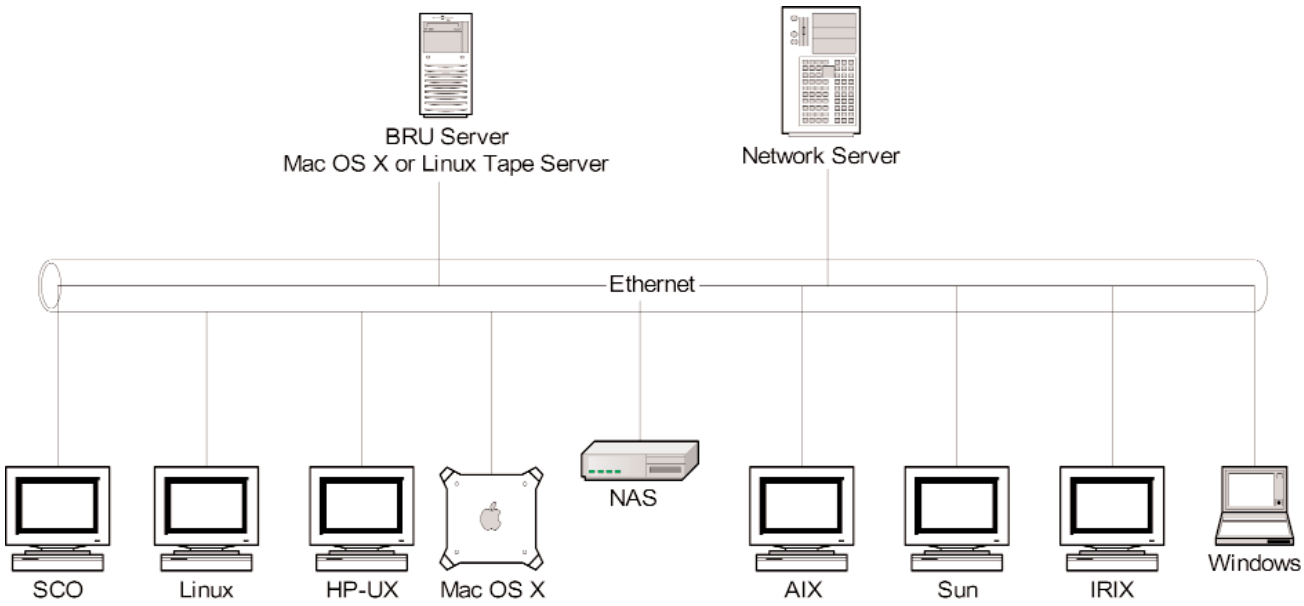

Figure 2.0 BRU Server typical network backup system

Verifying the Backup's Accuracy: Backups need to be verified for accuracy, otherwise there is no assurance the data can be recovered correctly. Experiencing an aborted restore is absolutely the wrong time to find out the backup was no good! All BRU products offer verification choice and flexibility. BRU backups can be verified on-line immediately following the backup, off-line, or even on an inexpensive or spare system elsewhere, regardless of platform or OS. This flexibility is possible since the checksum is written to tape and provides all of the authentication information required to verify the backup. Each backup can be fully verified without the need for continued access to the original data. Additionally, verifying BRU tapes off-line shortens the on-line backup window, and an administrator can even delegate the task of verifying backups to a junior person. There's even backward compatibility between BRU Server and BRU tapes.

When BRU tapes are verified, the checksum value is compared to the metadata and file data to assure that what was expected to go on tape actually made it. If an invalid signature within any buffer block is detected, an error message is issued, pointing to the specific location of that bad buffer block (Figure 3.).

```
warning - archive read error at block 342
```

Figure 3.0 BRU error reporting

This detailed reporting negates searching through voluminous reports to identify where the backup failed, and thus any restore, could or would fail. Note the contrast in messaging reported by `tar` and `cpio` (Figure 4.).

```
tape I/O read error
```

Figure 4.0 tar & cpio error reporting

Solutions based on the `cpio` tape format verify only the metadata and cannot detect corrupted file data. This scenario provides a false sense of security since there is still no assurance that all data will be recovered accurately. Solutions based on the tar tape format necessitate a "differences mode" (bit-by-bit comparison of disk and tape contents) verification be conducted on-line, immediately following the actual backup while the system is in a quiescent state. This verification operation virtually doubles the backup window. Should the system be accessed during the verification stage, an administrator will have to search through a plethora of error messages to determine if someone did an "`ls`" on a particular directory changing the atime (access time), or someone deleted a file, or find out the backup was actually bad.

Without the checks and balances of accountability, the validity of data stored for future access cannot be assured.

Recovering Your Data: Protecting data is all about recovery. If it cannot be successfully recovered, why back it up? The BRU tape format was born to deliver absolute assurance that data can be accurately recovered. Rather than checksumming across an entire file, checksumming each buffer block allows provides BRU the ability to recover a majority of a file should bits within it become corrupted following archiving onto tape. During a restore, if BRU detects a bad buffer block, an error message identifying the location is issued (Figure 3) and the tape is advanced until a good header is located. When the data is validated, the restore continues. If not, BRU advances the tape again, and again, until it detects good data and then continues the restore. Recovering data quickly is an issue because most usually it is needed immediately. In addition to QFA (Quick File Access), BRU's architecture supports faster wholesale data recovery, addressed in the performance section of this paper.

Back Up Performance: Following reliability considerations, the throughput performance of a backup system is important — both how quickly the backup can be accomplished, and how quickly can the data be restored. On the backup side, if the system must be placed in a quiescent state during the backup and/or verification operation, then it is not not available to support the business.

Mechanisms that shorten the backup and recovery window, make administering the backup system easier, and assure backup system availability are of high consideration and value. Each of these BRU characteristics will be addressed.

Experiencing an aborted restore is absolutely the wrong time to find out the backup was no good! ... Protecting data is all about recovery. If it cannot be successfully recovered, why back it up?

Although the BRU tape format is the most robust in terms of delivering reliability, it incurs no throughput performance penalty. BRU can stream the highest performance tape drive available.

Backup system throughput will always be impeded by the slowest link. Slow client systems, slow network capability, sub-optimal tape systems, or worse, backup software that cannot deliver the required performance can all impede timely backup execution. Assuming your system is appropriately configured, the focus should then be on the backup software.

As mentioned, BRU technology incurs very low overhead and can stream the most advanced tape drive. In single system environments, BRU's throughput capability is not the bottleneck to backup execution time. When the system topology is more complex, such as in a client/server architecture involving many systems, the architecture and capabilities of the backup software become key.

In the large networked system environment, BRU Server can run simultaneous data streams from multiple clients to multiple tape devices (Figure 5.). This means backups can be executed in up to 75% less time than other approaches, depending, of course, on the system topology. Simultaneous writing also allows data from any particular client to be captured contiguously on a tape, or tapes, depending on the client filesystem size. Should a particular client system need to be recovered, the restore is quite straightforward and efficient.

Other solutions provide a multiplexing capability, wherein multiple data streams from different clients on the network are channeled into a single tape device. Multiplexing, though it sounds good, negatively impacts several issues: namely backup throughput and data recoverability. By funneling all client data to a single tape device, backup times can be greatly lengthened. In effect, the software serves as the bottleneck to performance. In multiplexing backup schemes, software algorithms poll clients intermittently for their data. Since a single stream of data is written onto tape, files from any particular server can be spread across many tapes when numerous clients are being backed-up.

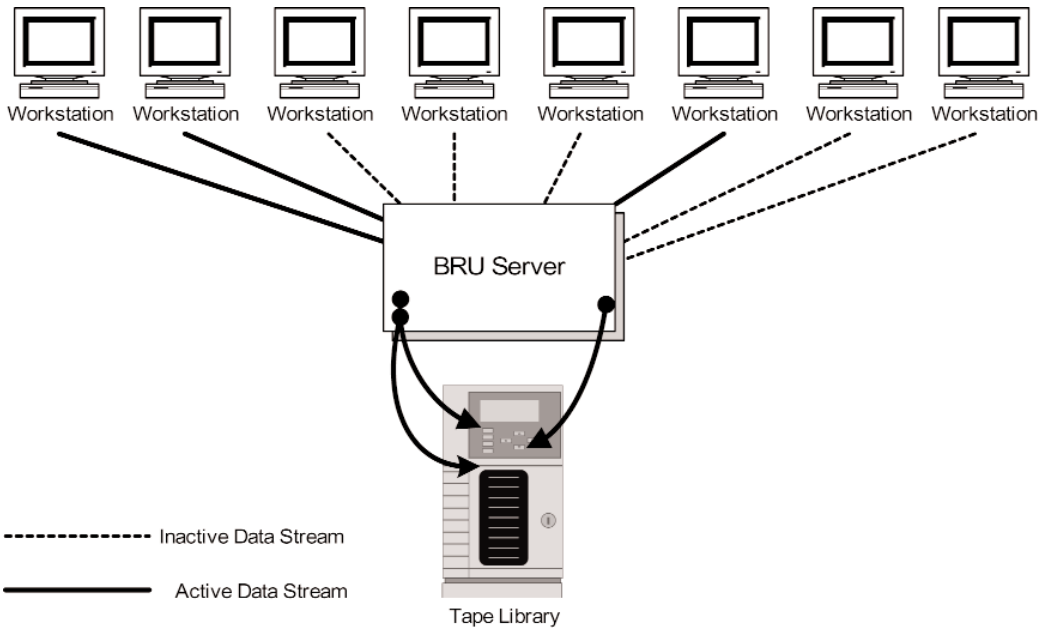

Figure 5.0 BRU Server simultaneous data streams

This means recovering data for that particular client would be complicated since its data has been spread across multiple tapes, requiring multiple tape loads and searches.

To further shrink the backup window, BRU Server provides robust disk staging support (D2D), whereby data is backed up to interim disk prior to committing it to tape. This allows multiple clients to perform their backups simultaneously. BRU Server's disk staging capability is fully configurable, and is fully verifiable like a backup to tape.

Flexibility: A system/network administrator's job can be a difficult one. Each aspect of BRU's development has been audited to assure ease-of-use. BRU-Server's automated configuration of client systems and backup devices alleviates the typical cumbersome manual effort required to define the backup system components.

When unwanted/unexpected events occur during a backup or restore operation, it's key that reporting be detailed and informative to guide the administrator. BRU's real world reporting allows administrators to pinpoint what occurred. For instance, when backing up active systems, BRU will report that a particular "file grew while archiving" if someone added data to it as it was actively being backed-up. Conversely, if someone deletes information out of a file while it is being backed-up, BRU reports "file truncated while archiving." In both cases, BRU will always archive the file even though things changed.

Mechanisms that shorten the backup and recovery window, make administering the backup system easier, and assure backup system availability are of high consideration and value.

Additionally, BRU Server handles the details of tapes within libraries by allowing the segregation of media into "destinations."

... backups can be executed in up to 75% less time than other approaches, depending, of course, on the system topology.

Backup systems are intended to assure the availability of archived data should it be needed. But what backs up the backup system? After all, the availability of the backup system is the natural precursor to being able to recover data.

Unlike other backup solutions, BRU Server writes the catalogs to tape as well as the tape server's hard disk drive (HDD). Should the tape server's HDD fail (a weak link in backup systems), simply replace the drive, reload BRU Server, and select "import tapes" from the BRU Server Console. BRU Server automatically reads the catalogs on each tape and replaces the catalog onto the tape server HDD. Whereas the time to back up your data may have taken hours, the process to restore the catalogs takes minutes.

When using any other backup software application, the process to recreate the catalogs for placement on the tape server HDD is one of protracted manual endeavor.

Enhanced backup system availability is also provided by BRU Server's Drive Autobalancing™ capability. Drive Autobalancing™ insures no drive is used more often than any other within a multi-drive environment. Tape drives do fail, and this capability optimizes the overall uptime of a tape library subsystem.

In an emergency situation, data can even be taken off BRU Server tapes with any UNIX/Linux system. Simply download and install a version of BRU specific to your platform from TOLIS' Web site and read the tapes. The full function demos support data restores without the requirement to register. Backward compatibility to read BRU Server tapes extends back to BRU 6.1.

Depending on the size of the files being backed-up, it is estimated BRU will consume anywhere from 8-12% of additional tape capacity to contain the checksum values and additional control metadata.

When unwanted/unexpected events occur during a backup or restore operation, it's key that reporting be detailed and informative to guide the administrator.

Seemingly inefficient at first glance, the additional tape being used to contain failsafe mechanisms is a small, cost-effective investment to assure your data is safely protected. Without the fullest protection, your data is at risk.

Conclusion: Data protection is the TOLIS Group's sole business, and the company's knowledge and expertise is evident in the products. BRU solutions are ultra-reliable, easy to implement and use, are highly intuitive, and no expensive classes are needed to understand how to use them. The documentation, written by people with system and network administration experience, is fully descriptive; containing the useful information you need to know. Should technical support be needed, TOLIS engineers answer your questions.

BRU products span the computing spectrum, from the datacenter to the home, and the argument to implement them is compelling. In a time of high cost and questionable return, BRU delivers outstanding value.

Backup systems are intended to assure the availability of archived data should it be needed ... [and] without the fullest protection, your data is at risk.

For full details on the robustness of the BRU products, please visit: http://www.tolisgroup.com