



Backup You Can TrustSM

Reliable Verification vs. Tape Drive Read After Write and Hardware ECC

At TOLIS Group, we strongly believe that verified backups are the only backups that can be trusted when disaster strikes and data needs to be restored. Over the 18 years that BRU has been providing reliable backup and restore, we've heard numerous horror stories about worthless backups caused by bad backups or bad backup practices. A frightening trend in the backup industry is the dependence on a mechanism in the tape drive hardware called "Read After Write" verification and drive-based "Hardware ECC" (error correction code). Several vendors are using these tape drive features as a way of lessening the importance of a data verification pass or as the explanation as to why their application is limited to supporting certain tape drives. While "Read After Write" and "Hardware ECC" are helpful features, don't be fooled into thinking they provide foolproof data protection.

While these features may seem new or high-tech, they have actually been available in the tape drive industry for many years. If these features provided the level of assurance that many software vendors are now touting, why are we still witness to unrecoverable backups? Both drive based operations were available in early quarter inch (QIC) drives from Archive, Cipher, Wangtek, and Tandberg Data as far back as the QIC-11 standard (circa 1985). They were implemented to enable tape drives to automatically recognize and correct data errors caused by media failure — such as areas where the data was not correctly written.

There is only one way to truly validate that the data you think was backed up actually made it onto the tape - an application controlled verification pass.

On the next page, we'll examine these two processes further:

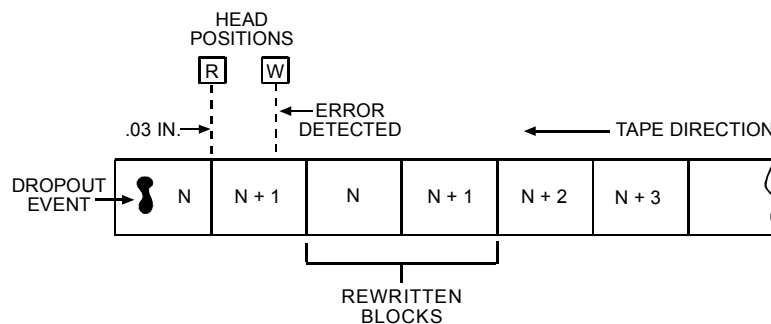


Figure 1 — Read After Write Logic

The "Read After Write" process uses a specially designed head assembly. The "head" that you can see in a tape drive is actually made up of multiple magnetic elements - the "actual" heads. These actual heads are positioned within the head assembly to allow the write head to completely write one or more blocks of data before the read head examines the previously data written - thus the name, "Read After Write." As figure 1 illustrates, in the occurrence of an error in data block n , the drive completes the writing of data block $n+1$ and then rewrites data blocks n and $n+1$ to ensure continuity in the data stream. This rewrite can occur up to 16 or more times consecutively before the drive decides it has reached a hard error on the tape and completely aborts the write process.

The second operation designed to reduce the potential for data errors is "Hardware ECC," or Error Correction Code. "Hardware ECC" is a process that stores the information redundantly so that bad data can be repaired. For most modern tape devices, the combination of "Read After Write" and "Hardware ECC" allow legitimate vendor claims of 1 in 10²¹ bits of unrecognized bad data - once the data is inside of the drive.

While these processes are very good at catching errors in the data stream caused by problems with your tape media or the drive's hardware, they do nothing to verify that the data stream coming into the drive's input buffers from the data bus is anything more

than garbage. And, we all know that "garbage in" means "garbage out". These errors can occur in system memory, on the motherboard, in the drive cabling, and many other areas of a system that would never be seen by the tape drive's methods

At TOLIS Group, we strongly believe that verified backups are the only backups that can be trusted when disaster strikes and data needs to be restored ... [Read after Write and ECC] do nothing to verify that the data stream coming into the drive's input buffers from the data bus is anything more than garbage.

There is only one way to truly validate that the data you think was backed up actually made it onto the tape - an application controlled verification pass. There are two generally accepted methods for verifying the backed up data - a bit-by-bit comparison, or a checksum-based re-read of the tape data.

Our BRU and BRU Server products provide both of these mechanisms for verification. Our preferred process is BRU's Autoscan™ or AnyTime Verify™ mechanisms. Since BRU generates 32 bit checksums immediately upon reading each 2KB of data from your filesystems, we are able to re-read the tape during a post backup pass and validate that every 2KB of data is what we originally read from the filesystem; not just on a file by file basis, but for each 2KB of data within each file. By not involving the filesystem in this verify pass, BRU effectively reduces your backup window requirement by 50%.

Also, since you can use BRU's AnyTime Verify mechanism as its name suggests - any time - and since the process only involves the checksum values for the backed up data, you can even verify your backups on another system running a different operating system.

If your backup window is not a problem, BRU also supports the more recognized bit-by-bit comparison verification process. This option re-reads your backup media and compares it, bit-by-bit, against the same file(s) on your system. The drawback to this type of verification is the filesystems must be kept quiet and unused during both the backup and verification passes.

Remember, tapes drives have offered "Read After Write" verification and "Hardware ECC" capabilities for many years. But, people still end up with unrecoverable backups. Don't let this happen to you! Make certain that your backup application provides true error detection during the entire backup process.