



Backup You Can TrustSM

BRU's Unique Availability Characteristics

Preface: The process of backing up data is not an end result. Rather, it's a means to assure the access to archived information when it's most needed — usually at a critical juncture.

The availability to off-line data is directly attributable to the accuracy of the backup, which is directly related to characteristics embedded into the design of a backup tool's architecture. The presence or lack of these characteristics may not be immediately evident to someone evaluating backup solutions, but they certainly have a profound impact on performance. Additionally, the backup operation is intrusive and takes time away from on-line computing when the system is supporting the business. Mechanisms employed to minimize the time to back up data are of high value, as are those mechanisms that help assure the uptime/availability of the backup system.

The process of backing up data is not an end result. Rather, it's a means to assure the access to archived information when it's most needed.

The characteristics of BRU™ technology that deliver high availability to archived data and backup system availability are closely intertwined and presented below. Most characteristics are common across the BRU and BRU Server product sets. It will be noted when a characteristic is solely associated with a particular product set.

Design Approach: Structure is defined by its foundation - the more firm the foundation, the more viable the structure. The "BRU engine" was developed from a cornerstone mandate to successfully recover data. This mandate has not only formed the philosophical foundation of BRU, it has defined the unique BRU data format that has proven itself for 18+ years. The backup operation was in turn designed to support the established recoverability guidelines. Support for superfluous "bells and whistles" or unreliable data transmission protocols that fall outside of the BRU backup and recovery design parameters are not incorporated into the products.

It is these philosophical differences that set BRU technology apart from the competition in design, implementation and execution.

Is the Backup Good? Is the backup good? A simple question, but one not easily nor efficiently answered using the vast majority of backup tools. The availability of archived information, and the ability to recover it accurately, necessitates a good backup. BRU technology assures the accuracy of the backup by applying unique verification techniques.

Each and every bit of data archived is accounted for beyond the non-existent or limiting bit-level

comparisons used by some tools. BRU's comprehensive auditing capabilities add an enhanced dimension to the verification of backups.

BRU calculates checksums and includes these checksums with the data written to the archive media. These checksums calculate not only on the actual data, but the metadata as well, in contrast to those approaches that checksum only the metadata. It is important to note that alterations to the metadata can result in an inaccurate backup and the lost recoverability of data. BRU's unique format allows the backup to be fully authenticated via a background operation, allowing the resumption of normal system activity immediately following the backup operation. This means faster resumption of normal system duties.

Some backup tools sample random files to verify the backup. If the sampled files are accurately restored, the entire backup is deemed good. This strategy is seriously flawed, and provides a false sense of security that the complete backup is good. When using BRU, the question of a backup's accuracy is easily and most accurately answered. No guesswork is involved.

Is the backup good? A simple question, but one not easily nor efficiently answered ... When using BRU, the question of a backup's accuracy is easily and most accurately answered.

Minimizing the Backup Window: Bit-level comparisons that verify the backup double the time required to backup the system since an amount of time equal to the backup time must be spent to perform the verification. To ease the administrative effort of this operation, the system should be kept in a quiesced state so values don't change. If the system is not kept in a quiesced state, the administrator must then deal with potentially voluminous error messages to determine their origin — whether someone actually changed data, or simply read it, or the data was indeed incorrectly written to the archive.

As mentioned above, BRU writes the checksum (data & metadata) to the archive media. This technique allows the verification of the backup after the system is placed back into on-line, operational status. Since the BRU format is operating system transparent, AnyTimeVerify™ allows verification of the backup on the primary system, on a different system, or even on a different operating system platform. Operation is straightforward and can even be done by junior personnel. Using AnyTimeVerify™ reduces the verified backup window by 50%. BRU's AUTOSCAN™ feature can be used if a bit-level comparison is desired.

From a throughput perspective, BRU is a "small" application, and incurs very little overhead. Because of this low overhead, backups are completed at the speed of the system's capability to further accelerate backup completion time. BRU is also "tunable" to set buffer size parameters to match the optimum performance characteristics of the tape devices being implemented. BRU also supports double buffering, a technique to increase throughput by creating "parent" and "child" processes to simultaneously manage disk and archive device I/O to increase performance. (Please note: not all system configurations support double buffering)

BRU Server for client/server-based heterogeneous system networks accelerates the backup process significantly by writing multiple client data simultaneously to multiple tape devices for the ultimate in backup throughput. This performance advantage can deliver throughput up to 75% faster than traditional approaches.

Furthermore, when hardware data compression isn't an option, BRU's highly efficient compression algorithms can be used to further reduce the backup window. The compression is done in memory (no disk space is required), and a special error-detection format is implemented. To realize heightened efficiency, BRU will not attempt to compress files that have already been compressed.

The backup time efficiency realized will be dependent upon the type of system data, since not all data can be effectively compressed (i.e. .MP3 and .jpeg files).

Backup System Availability: When the backup system isn't available, archived data simply cannot be recovered — unless BRU is being used (explained under "Availability of Off-line Data").

BRU implements mechanisms to help assure overall backup system availability such as monitoring tools that will report dirty tape drives, bad tapes, and bad tape drives to allow for pro-active remedy before data must be recovered. In smaller system environments, BRU's tape directory capability allows the preview of a tape's contents, thereby eliminating time-consuming reads that drain system resources. In larger, client/server-based network topologies, BRU Server implements a number of unique features to assure availability of the backup system. These features follow.

A backup system is only as capable as its weakest element. BRU Server is designed to run on either an ultra-reliable Linux tape server attached to the network, or a Mac OS X system. The proven stability and uptime characteristics of Linux and Mac OS X alleviate operating system compatibility and reliability as a backup system issue and failure point.

BRU Server is the only client/server-based backup solution that writes the catalogs to tape as well as to the tape server hard disk drive (HDD). This feature provides yet another high level of backup system availability. Should the HDD fail, recovery of the catalogs is a straightforward and efficient process. This is important since the catalogs sync the tapes to the application so data can be located, recognized, and recovered. Should the HDD fail while using other products, the rebuilding of the catalogs is a difficult process involving significant time, effort, and lack of archived data availability assuming the data can be rebuilt.

In computing environments that implement tape autoloaders/libraries configured with multiple drives, BRU Server's AutoBalancing™ capability spreads the duty cycle across all of the drives. Assuring balanced access across the drives helps to alleviate premature failure which bolsters tape subsystem availability.

Availability of Off-Line Data: As previously mentioned, the backup must be accurate in order to successfully recover the data. The backup's accuracy is the prime factor, but how the data was placed on the archive media affects its recoverability as well. The BRU format calculates and writes the checksums to the header of buffer blocks, and not across an entire file. This unique technique is key when corrupted data is encountered during a restore.

A backup system is only as capable as its weakest element ... how the data was placed on the archive media affects its recoverability as well.

A number of events following the backup can corrupt data on the archive media, and these events do happen. Traditional backup approaches abort the restore when bad data is detected. This might be acceptable if it happens to be the last bit to be recovered, since all but the last bit will have been recovered. When it's not the "last bit," a significant amount of data can be irretrievably lost. When BRU detects bad data, it attempts to re-read the data multiple times. If the data cannot be read, BRU logs the location detail, advances to the next buffer block and reads the media. The process continues until good data is read and the recovery then continues to completion. Of all the backup tools available, BRU returns the most data.

When the computing topology is a network of multiple systems, the pathway the data takes on its course to be laid on tape has a direct effect on its availability. BRU Server simultaneously backs up multiple clients to multiple tape devices. Think of this process as a multi-lane highway with all lanes open and the data from client systems traveling the lanes. This is important because the data from any particular client is streamed contiguously onto tape, spanning across tapes when necessary. Information recovery is therefore straightforward and efficient. The recovery of a single file is

simplified because the data is localized on the media. A complete restore of the client's data is also time-efficient because a mount, search, and read of multiple tapes is not required.

In contrast, other client/server-based backup tools do not provide the throughput of BRU Server. They intermittently poll clients for their data and constrict the flow of "traffic" to a "single highway lane" writing to a single tape device. Pumping data into a single stream results in the data of any particular client being spread across many tape volumes even though that client's data could have been contained on a single tape cartridge.

Restoring all of a particular client system's data is slower and inefficient because multiple volumes must be mounted and the data subsequently pulled off them.

Additionally, if a tape is lost (this too happens), the data of many client systems will be lost since that data is dispersed among many tape volumes. Additionally, BRU's multiple simultaneous stream technique can deliver a 75% greater throughput than other tools.

Pumping data into a single stream results in the data of any particular client being spread across many tape volumes even though that client's data could have been contained on a single tape cartridge.

Should a catastrophic event render a backup system implementing BRU Server totally incapable of operation, data can still be readily recovered. BRU Server tapes are backward

compatible, and can be read with any version of BRU back to v6.1 (17.0 is currently shipping). BRU is available for download from the TOLIS web site, and supports all common UNIX variants. Simply download BRU onto a different or spare system and data can be easily pulled off the BRU Server tapes. This is yet another example of BRU flexibility and availability characteristics.

Conclusion: All BRU products do one thing, and do it very well. They backup data accurately, and successfully recover that information better than any other data protection tool available. Care has been taken in BRU's design to assure users that their archived data is safely protected. TOLIS' "Backup You Can Trust" service mark mandates this level of performance.

All BRU products do one thing, and do it very well...they restore information.