

Data integrity with LiveVault



LiveVault data protection

Increasingly, today's businesses recognize the value and convenience of using online backup to protect their server data. For any enterprise considering an online backup service, it is critical to ensure the integrity of data, both during backup and restore, as well as during longer retentions. With this safeguard in place, you can recover your data when disaster strikes. LiveVault® offers on-disk retention for up to seven years and ensures that protected data is error-free during backup, transmission, storage, and restoration.

Data in transit

For ongoing backups and restores, LiveVault performs data integrity checks on all packets of data sent between protected servers, vaults, and Turbo Restore Appliances using the Secure Sockets Layer (SSL) protocol on top of TCP/IP. These checks ensure that any alterations, additions, or deletions of data during transmission—either accidental or malicious—are detected and corrected by retransmitting the data.

In addition, LiveVault transmits the file ID and block number information redundantly and validates that the data blocks arrive in the expected order, without gaps, and that the header information is consistent. When messages arrive at the vault, they are written to disk.

The LiveVault vaulting software does no processing on the backed-up data. It arrives compressed and encrypted, and is stored this way. LiveVault vaulting software executes on standard Windows® 2008 R2 servers with RAID 5 disk systems, and the software uses standard Microsoft® file system I/O. Source server applications, such as SQL Server® and Exchange, also rely on the integrity of the Microsoft file system.

Data in vaults

Unlike most applications, LiveVault takes special steps to validate the integrity of the data on disk. LiveVault vaults employ proprietary technology to index, catalog, and organize the data blocks. This technology enables LiveVault to run four types of consistency checks:



1. On every write operation, LiveVault validates the header information to ensure that blocks are written in the intended order, and that the header information is consistent.
2. Once a day, a block consistency checker runs automatically. This validates that all the blocks are present that the vault needs to restore the latest backups. The block consistency check operates against LiveVault's block index files, but it does not actually read the encrypted data blocks from disk.
3. Every two weeks, a full catalog consistency checker runs automatically. This validates that, if requested, every file could be completely restored, but it does not actually read the encrypted data blocks from disk.
4. Once a month, a data consistency checker runs automatically. This simulates a full restore of the latest version of every file. It reads the encrypted data blocks from disk to flush out any lurking disk errors.

In addition, the "rotation" software helps ensure the integrity of data at rest. As time passes, various backup versions reach the expiration of their retention periods, but some of the data blocks—the ones that have been static—are still needed and are "moved forward" into a new baseline by the rotation software.

This process involves copying the blocks and provides an additional check that disk blocks are accessible, and that the header information is correct. Unlike tapes—which age on a shelf until you attempt a restore—LiveVault reads and rewrites data throughout its retention period, regularly testing the integrity of the storage systems and software.

Mirrored data centers

For high availability, LiveVault maintains all backup data on two independent vaults at separate geographic locations. Backup data that arrives at the primary vault is then sent to the secondary vault, where it is revalidated and stored in the same manner as on the primary vault. If there should ever be an unrecoverable error on one of the vaults, the other vault provides continuity of ongoing backups and gives you the ability to restore historical versions. The vaults continuously cross-check each other to see that each has the same set of backup versions, and they cross-replicate between each other as necessary. When a vault comes back online after repair, it automatically acquires data that it is missing from the other vault.

Learn more at

www.livevault.com