

Virtual Data Recovery

White Paper Discussion on Exploiting the Virtual Tape System For Enhanced Vaulting & Disaster Recovery

Disaster Recovery for Virtual Tape Data

There are several processes in use today for creating and storing *dual copies* of critical data for long-term vaulting and Disaster Recovery purposes. This paper is intended to discuss the most common processes, which involve backups created on tape. It is the paper's intention to primarily focus on these processes where a virtual tape system has been implemented.

Tape D/R Processes When Not Utilizing the Virtual Tape System.

Creating dual copies on site. The most common process for creating dual copies is for the application or a separate backup job to create backups on tape which will then be shipped to a separate storage location. This process historically requires a number of concurrent native tape drives which depend on the amount and frequency of your backup creation needs. The process is straightforward and can easily be audited. Although simplistic, the targeted media is often dramatically under utilized due to the simple nature of the applications focus to back up its own critical datasets which results in higher costs associated with vaulting data because more slots are required and the media is under utilized.

Furthermore, the site must always ensure enough concurrent tape drives are available which has been an accepted cost until recently. If an application needs backup data for recovery, the data must be brought back from a vaulted location and at that point, the integrity of the D/R process is compromised. Local backup copies can be created to alleviate data integrity issues. The cost for this additional process would be to double the media, tape drive hours, floor space and processing time.

Creating dual copies electronically to remote locations. Another dual copy process is to create the copies directly to tape devices residing at a remote location. This process eliminates the potential D/R integrity issue discussed above and can reduce the amount of local administration costs. However, total costs are higher due to the expense of additional remote administration costs and tape devices for remote storage as well as communication hardware, software, and transmission expenses.

Tape D/R Processes When Utilizing the Virtual Tape System.

Using an Import/Export process. Virtual Import/Export solutions offer customers the capability to temporarily utilize the Virtual Tape System for critical data but not without these process drawbacks:

- The user must identify the virtual volsers to be exported.
- Selected Virtual volsers are then migrated to a physical container volume and when the container volume is removed from the virtual tape system, the Tape Management Catalog has no record of those Volumes and therefore the vault rotation of those Volumes can not be managed correctly by the Tape Management System.
- Stacking tape files with different vault rotations can compromise the D/R process.
- The time it takes to copy an exported Volume to other media can be unacceptable. There are no automated functions for JCL creation and job management without a VTL at the DR site.
- Does not exploit the Virtual Tape System for continuous & seamless D/R purposes.
- Places heavy duty on the VTL associated with Dataset recalls on the back end and limited tape drives which can impede daily production jobs
- Creates a cartridge in a proprietary format that can only be read by the VTL

Using a Peer to Peer process. This process involves adding ample virtual devices and supporting hardware/software for virtual data duplication. The dual copies are automatic and do not require the user to specify critical data. If the additional virtual equipment is located in a disaster recovery location, the dual copies will not need to be transported and the entire process can be somewhat transparent to the user. Consideration has to be given to the additional costs of such a solution which can include more virtual storage devices, hardware & software communication, and transmission costs than expected since this solution creates

dual copies of all virtual data, in lieu of user selected critical data. In addition, potential slow-downs can be caused by inadequate bandwidth, or component failures and force the local VTL to enter slow-down mode.

Using tape migration software to move critical data sets out of the VTS.

This solution requires an automated tape copying process which can identify critical data sets within the VTS, build JCL for the conversion jobs, provide or access a data set migration utility to move the data to a native Volser which will be stored off-site, and automatically update the TMS and MVS catalogue.

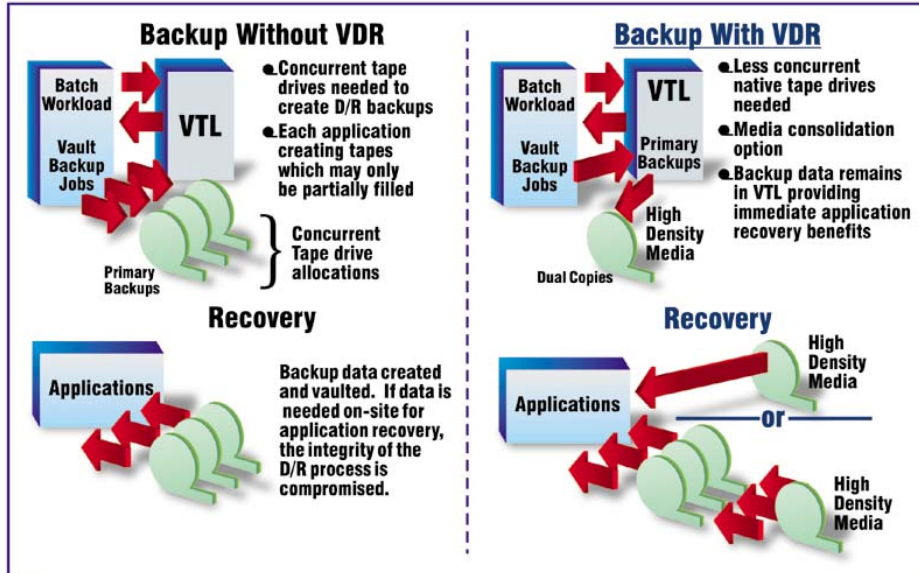
The benefits are:

- Critical application data can be created in the VTS thereby further exploiting new virtual technology
- Lower cost solution for creating backups
- Backups can be stored on higher density media, thereby reducing overall vaulted tape storage and potential transportation costs.

There are some potentially negative issues the user should consider with this process also:

- The data is permanently removed from the VTS, thereby eliminating virtual benefits from a subset of tape data.
- If an application needs backup data for local recovery, the data must be brought back from a vaulted location and at that point integrity of the D/R process can be compromised.
- This will most likely impact other applications since the data is now stacked.
- Stacking files with different vault rotations can cause rotations to occur earlier than defined in the vault pattern dataset.

Virtual Data Recovery (VDR)



Virtual Data Recovery was created to fully exploit virtual tape system technology with *dual copy* creation for VTL Data Availability and/or Disaster Recovery. The following section describes the continuous process of vaulting critical tape data while providing a cost effective approach which fully supports any existing vault process.

Simplistically, the S/390 user defines to VDR what critical tape data should be duplicated. VDR will maintain a *dual copy* table that contains the name of the original data set (generic wild carded filters) to be copied and the new name of the target data set. The user must then decide the frequency required for creating duplex copies of business critical data. For example, if the current disaster recovery policy incorporates a daily movement of tapes from the data center to the remote vault, then the most pragmatic timing for VDR processing is following the completion of the overnight batch schedule.

Having identified the storage management policy, the job Scheduler then introduces a very simple batch process to duplex the specified datasets. The VDR function then utilizes the base data movement function to duplex the required tape files and volumes. There is significant intelligence applied to this processing, as the VDR function cross references the Tape Management Subsystem (TMS) inventory to determine whether a copy operation is required or not. The VDR function builds a list of data sets to copy, and passes this request to the base copy engine. VDR will continuously track the relationship between the primary backup and the *dual copy*.

VDR can be applied to most current D/R processes without modifying existing vaulting rules. One common D/R process is to create a backup copy of critical application data for storage off site or to protected locations. While creating these backups, concurrent native tape drives are needed to create all the backups necessary to meet a certain processing window. The backups are then usually shipped to the vault location. Many times datasets return from the vault offsite location before their specified time for various reasons. If the VTL is exploited for D/R purposes, the primary backup is created in the VTL and VDR creates a dual copy that is then sent to the vault offsite location.

The *dual copies* can be consolidated on high density media during creation. In addition, only one native tape drive is needed to create the *dual copies* which can drastically reduce the number of concurrent native tape drives needed to create backups since the primary backups were targeted to the virtual device. The end result of the VDR process is a cataloged *dual copy* of the primary backup data to be located off-site. The Virtual Tape System will continue to store the primary backup allowing instant access to all applications which need an immediate recovery of a file for any reason. This often allows the user to eliminate the extra on-site copies retained on physical media and results in the dramatic reduction in storage racks and storage space for these cartridges.

Recovery

If the S/390 customer encounters a data center outage, then they can invoke the VDR recovery function at the disaster recovery site to rename the *dual copy back* to the primary backup. Obviously this particular VDR function only applies to tape files and volumes that were copied using VDR processing.

Recovery of *dual copy* data can be handled in several ways, depending on the Disaster Recovery needs. One option is to recover all *dual copy* data from existing media. The advantage for this *in place recovery* is time and storage. There would be no data movement necessary. The VDR function would provide all the dataset file processing and associated cataloguing necessary to ensure the *dual copy* is renamed back to the original name and switched to the primary dataset. Then the application can begin the standard recovery process directly from the VDR cartridges without any data movement. VDR also allows the user to address concurrent dataset requests during the recovery by allowing the co-location function to intelligently organize and stack the dual copies.

Another recovery option is to recover to any targeted media. This option allows the user to copy the *dual copies* to selected media instead of recovering on the existing media, thereby minimizing the potential processing conflicts caused by multiple datasets stored on the same tape volume being requested by concurrent application recoveries. The cost for this flexibility is time and storage. *As with the in place recovery*, the VDR function would provide all the dataset file processing and associated cataloguing necessary to ensure the *dual copy* is renamed back to the original and is switched to the primary backup dataset.

A Hybrid of the two processes described above can also be used at the DR site.

- First using **copy recovery** to restore some number of days into either a VTL or native media at the DR site.
- And then a second process using **the backup media as source** to recover the rest of the past week or month.

Co location also allows the user to prioritise their application recovery to support the Business Continuity Planning requirements to insure rapid recovery.

In summary, VDR provides the VTL customer with the ability to provide increased levels of resilience for business critical data.

Benefits

VDR reduces the cost of current tape vaulting methodologies and increases the integrity of the D/R process in the following ways:

- VDR is a selective data process, thereby reducing the amount of tape media required.
- VDR creates & maintains a cataloged copy and allows the virtual device to retain a copy thereby ensuring a copy on-site in addition to a secured location. This maintains or improves the integrity of the vaulting process if backup data is needed at the data center.
- Reduces native drives necessary to create daily vault data by copying directly to the virtual device.
- Reduces backup media requirements by copying to higher density media and providing recovery flexibility.
- The data center does not need to retrieve the vaulted cartridges until all files are eligible to return, thereby maintaining the integrity of the D/R process.
- If the vaulting process involves more than one vault, VDR assures the appropriate backup data is in the correct vault.

Conclusion

Virtual tape devices should play a greater role in the disaster recovery process. The virtual tape device utilizing VDR can provide increased data integrity within the D/R framework, faster resolution to application recovery, and reduce the requirement for native tape drives. Having overcome concerns regarding tape capacity optimization and performance, S/390 customers can now focus on data resilience challenges. The key to exploiting the Virtual Tape System so that it can be properly used within your vaulting process is the Virtual Data Recovery solution described above.

The Virtual Data Recovery solution provides the required functionality to resolve data duplication shortcomings associated with other tape migration software. Additionally the VDR function can circumvent any issues that might be associated with virtual tape servers, where volumes may need to be exported to a physical cartridge for transport to a remote vault. Thus the VTL customer can utilise VDR to duplicate and consolidate their business critical data, onto the minimum number of Tape Management System controlled media.

In conclusion, the VDR function can ensure Virtual Tape System benefits are maximized by using it as a component of D/R while maintaining or even improving the daily vaulting processes.