

White Paper on DASD Backup Automation

Introduction

As a component of the Disaster Recovery process (or Business Continuance), most Data Centers make backup copies of their DASD volumes to facilitate the restoration process of replicating their operating environment at a given synchronization point at a Disaster Recovery Center. The Data Center usually starts with a standard OS/390 floor system provided by the hot site vendor and begins the restoration process to completely restore all Data Center DASD.

Some companies may lay down a mini-system and begin the rebuilding process after an IPL of their mini-system, then through other processes using their tailored operating system, catalogs, application software and other application data would be reloaded on initialized DASD volumes. At that point, the data center would re-IPL and then complete the application recovery process or begin testing.

Using Tape for Backing up DASD Volumes

Over the years, Data Centers have been able to utilize relatively low cost tape, 3420/3480/3490, (from 50 cents to 4 dollars) to backup DASD volumes. It might take anywhere from 2 to 10 tapes to backup one 3390 model 3. As tape capacities have grown and the cost of tape cartridges have also risen (\$90.00 to \$150.00), Data Centers have been confronted with the need to better utilize the media by implementing a method of stacking multiple DASD volumes onto a single cartridge.

The Build Process

In many Data Centers, the process of adding additional volumes is controlled by a Storage Administrator that is responsible for allocating and de-allocating DASD resources to the various customers, departments or application pools within the organization. When the Storage Administrator has completed the allocation, it is often that person's responsibility to notify a systems programmer to add the newly created volume to the backup process, and maybe another person may need to update the restore process and maybe another person that updates the vaulting rules. Because we are human there is always some degree of "human error". These modifications/changes to procedures may be overlooked or substantially delayed due to workloads, and communication errors. The end result, not all DASD volumes are backed up when they should be backed up, resulting in a potential failure at disaster recovery during testing or an actual disaster recovery. In a manually controlled system there is not an easy way to audit the processes. One would have to wonder, are all critical volumes backed up and the restores built for each of those volumes?

In most OS/390 shops today, the Data Center tape operations staff uses DASD backup utilities such as DFDSS or FDR. These utilities require the Storage Administrator or Systems Programmer to create specific JCL to identify each critical DASD volume that is to be backed up, including the control cards directing the utility on how the backup should be performed. Many Storage Administrators have written specific JCL statements to make backup copies of a DASD volume onto a 3480 or 3490 cartridge. These cartridges would then be picked, packed and shipped to the archive or disaster recovery vault until it was required for disaster recovery. In addition, many Data Centers implemented strict change control procedures to identify all hardware, software and policy changes in the environment to insure that the appropriate DASD volumes are being backed up. Change Control Procedures insure that volume substitutions were properly addressed, making necessary changes to account for changes in the environment such as; new DASD volumes, changes in allocations or assignments, new DASD and Tape device types. In addition to writing the backup JCL, the Storage Administrator may also be required to write restore JCL taking into account the environment at the Disaster Recovery site, such as quantity and type of devices available for restoration, both Tape devices and DASD Models.

Execution Processes

The time required to complete full volume backups can be substantial, sometimes requiring over 24 hours. Many variables impact this window, such as the number and utilization of the DASD Volumes, the number and type of channels and the number and type of tape drives. Because of this lengthy process, many Data Centers are forced to only run full volume backups over weekend. The result of only running full volumes weekly requires complex procedures to complete a mid-week recovery.

More often than not companies require more than a just a weekend backup to achieve Business Continuity. What's really required to stay in business after a long term Data Center outage? This makes it necessary to create multiple synchronization points during the week for application recovery. As such, the Storage Administrator has been tasked with the requirement of also setting up or customizing incremental and/or differential backups. Because Incremental and Differential backups are more efficient than full volume backups, this will allow the business to create more synchronization points during the week.

The incremental or differential backup completes in less time and uses less tape, which requires a smaller number of backup tapes for shipment to the vault. The backup utilities allow the Storage Administrator to create full volume DASD backups once a week, and then make copies of only the datasets that have changed since the last full volume was taken. This method is more complex both in the Backup Process and the Recovery Process.

Some Data Centers have implemented the newer Advanced Copy Functions available with products like TimeFinder or FlashCopy. These newer advances allow the data center to basically mirror each DASD volume and complete a full volume backup of the mirrored volume daily. This simplifies and speeds up the Recovery Process. At the same time these Advanced Copy Functions will add some complexity to the Backup Process.

New Technology for High Density Tapes

Although the cost per megabyte has dropped on a cartridge, as the density, reliability, and the speed of tapes has increased, the cost per tape has also risen dramatically. It has become even more imperative that the Data Center make every effort in the effective use of tape media to reduce the overall cost of the tapes being shipped offsite. This will also reduce the storage and vaulting cost charged by the vault provider.

This need has resulted in company resources to create homegrown solutions in an effort to help stack DASD backups onto the newer higher-density media. All of the historical issues of matching DASD volume capacities and tape capacities need constant monitoring to insure that backups are completed correctly and tapes are being efficiently utilized. The Change Control procedures are still required to address changes in the environment for both Tape and DASD technology.

These Homegrown solutions are often created using a combination of TSO utilities, SAS, REXX, and some assembly language programming that requires maintenance by highly paid individuals. Each time there is a change to the environment such as releases of the backup tools, changes in operating systems, and physical media changes, the Homegrown utility may also require changes. These solutions are often partial or incomplete and often require extensive maintenance and monitoring by the individuals that created them. In many cases, the author of the utility is promoted, moves on, retires, is reassigned, or just no longer available to keep the maintenance current. Someone then needs to be trained to understand the inter-workings of the Homegrown utility and take over the maintenance, and because this utility is written using several system utilities more than one person may be required to maintain this system. The short comings of a Homegrown utility and the maintenance requirements are often discovered at the time of a disaster recovery. Furthermore, the disaster recovery platform is often different than originally expected so the user must make significant JCL changes to allow the DASD restoration to occur.

Automation/Solution

The Data Center needs a solution that provides an automated method which identifies volumes for backup based on "Pre-defined Policies" these policies would have inclusion and exclusion capability. The Storage Administrator should also have the ability to organize the backup process, which volumes are backed up and when, what volumes are stacked together or separated. These decisions may be based on DASD volume ownership, customer or applications. Backups could be scheduled in waves to create multiple sync points.

The Storage Administrator should have the choice of maximizing the utilization of the tape media or a number of DASD volumes per tape. The Storage Administrator should also be able to direct the utility to detect the actual used space of the DASD volume or the allocated space based on Data Center policies. The solution should be significantly adaptive in order to self adjust to meet the utilization goals of tape media usage, the quantity of DASD devices and their respective utilization.

The Storage Administrator should be able to incorporate the accumulated knowledge of the current processes, incorporate those into a systematized process that will automatically audit for changes, incorporate those changes when building the Backup JCL using vendor supported Backup tools like DFDSS and FDR. In addition, the Storage Administrator should have a tool that pre-builds the restore JCL at the primary site that matches the DASD configuration of the DR Site. Building the Restore JCL at the host site will act as an audit of the DR DASD configuration, if you have not contracted enough or the correct DASD models an exception will be reported. In addition, the reporting should identify all the DASD backup tapes that will be required at the DR site during recovery. This tool should also allow for changes to be made at the DR Site to support unexpected configuration changes with minimal human intervention to insure proper and complete DASD restoration.

The Storage Administrator should have a tool that would identify all the DASD backup tapes that will be required at the DR site during recovery. In addition the tool should provide comprehensive reporting on the backup and restore processes.

Today's costs of conducting a DR test can be quite costly. Many hours are expended in preparation for the test through out the company, not just the data center staff. The absence of a DASD volume during a DR test can result in a diminished test, a very long delay or a complete outright failure! In either case, the delay and the need to quickly reschedule a test can be extremely time consuming, costly and an additional drain on human resources.

DASD Backup Supervisor from OpenTech Systems meets all of these needs and will greatly increase the probability of successfully restoring all required DASD volumes at the DR Site.

- **An Audit Function -**
 - Monitors all DASD volumes and reports any changes, additions or deletions
 - Builds a VOLUME File that contain key points of control
- **A Backup Generation Function -**
 - Automates the creation of Backup JCL that conforms to Data Center Policies
 - Allows Storage Administrators many choices in the control of what gets backed up when and how
- **A Restore Generation Function -**
 - Automates the creation of Restore JCL at the Data Center for the DR Hot Site
 - Audits the DR sites DASD configuration
 - The ability to re-create the Restore JCL at the DR Hot Site if necessary to adapt any last minute changes or surprises
 - Generation of Tape Pull Listing

Summary

In the Data Center, Managers are being asked to provide more capability, with fewer resources, and maintain the high degree of reliability that mainframe Data Centers have delivered historically. With cost-conscience management and the high visibility of Business Continuity, Data Center managers can ill afford to have a failed Disaster Recovery test caused by homegrown tools that are not current with today's technology and environment. With DASD Backup Supervisor the Data Center Management can expect to have more control over the use of resources such as Tape Drives, Cartridge usage, and time necessary to conduct the backup process as well as reducing the time necessary to restore the DASD volumes at the Disaster Recovery site allowing more time to test the critical applications.