

JPL's Move from StorageTek Redwood to 9940 drives/media

Cris Windoffer
Jet Propulsion Laboratory
Pasadena, California

Introduction

StorageTek has end-of-lifed the Redwood Tape drive so JPL needed to find a replacement tape drive system for our Cray SV1 DMF data. We chose the StorageTek 9940 tape drive as the replacement. It is supported on the Cray SV1 and has similar drive access, capacity and bandwidth. It has much better reliability (and cheaper maintenance) than the Redwood drives. No user changes would be necessary to access their DMF data.

We are in the process of moving our DMF storage from StorageTek Redwoods to 9940 drives/media. We thought a site talk of our experience/tips might be of interest.

The Problem

JPL has long used DMF on its Crays and partnered with StorageTek for the hardware and media. With StorageTek's announcement of end-of-life we wanted to maintain our Cray DMF data and move to a newer tape technology. JPL has eight terabytes of data on the Cray SV1. This data has both primary and backup copies. We would like to move this data (again making two copies) during user production time with minimal user impact and maximum efficiency.

Hardware

Our existing hardware is three StorageTek Redwood drives in a Powderhorn silo. StorageTek recommended their newer tape drive, the 9940. The 9940 is supported by the Powderhorn silo and ACSLS v6.0.1. We consulted with Cray to see if they would support the new 9940 tape drives. Their support team tested and approved the 9940s with SCSI connectivity (fibre channel is not supported). So JPL purchased three 9940 tape drives and new tape media.

For the transition from Redwood to 9940s we configured the hardware with two Redwood drives and two 9940 drives. This would allow us to copy data from one Redwood drive while writing to two 9940 drives, and leaves one Redwood drive for continued user access.

Configure/Test Tape Devices

JPL configured the new tape devices per Cray's specifications. On the 9940 front-panel menu, we set each drive to 3590 emulation. In the TEXT_TAPECONFIG file we used the "type=VTAPE" parameter. We also defined a new group type of 9940. We then stopped and started the tape daemon and it saw the new drives. To use the new drives in testing, first we had to change our UDB limits adding user privilege on the second tape type. Also we labeled the

tapes with the “tplabel -g 9940” command. At this phase the TEXT_TAPECONFIG file kept our default group device as Redwood so all references to the new devices included the “-g 9940” parameter.

We moved on to testing tape read and write access using DUMP and RESTORE and TAR and UNTAR scripts. All testing was successful and we were ready to move on to trying DMF with our tape drives.

Configure/Test DMF

In order to test the new drives with DMF, we defined a 9940 device and two tape MSP pools in the DMF_CONFIG file. Cray specified the new 9940 device should emulate a 3490 with DMF v2.5.5.2. In the DMF_CONFIG file we set “DEVICE_TYPE 3490”. During testing we created two new pools of tapes and made them third and fourth priority over our two user pools (“MSP_NAMES redwd1 redwd2 silo3 silo4”). We used the standard UDB parameter “archmed:1:” to allow the test users to read and write to these two new pools of tapes. We then created the empty pool databases with the DMVOLADM commands and added tapes to those pools.

After successful DMF testing, we changed our DMF configuration to write all new user production data to the new tape pools by making them priority one and two (“MSP_NAMES silo3 silo4 redwd1 redwd2”). Old user data was still available in read mode from the two Redwood pools.

Moving The Old Redwood Data to New 9940 Media

Cray recommended we use the DMMOVE command to move Redwood data to 9940 tapes. However, they counseled the DMMOVE command could bottleneck all tape drives. After consulting with Cray we determined to limit DMMOVE to one Redwood tape at a time. So JPL wrote scripts to locate the files on a per tape basis and use a DMSELECT input format for those files. We did this by dumping the databases with DMCATADM “dump all” command and matched that to the DMSELECT “-m redwd1” output. We then created one file per tape that contained a list of files contained on the tape. After starting the move process from Redwood to 9940 we found the moves were going too slowly - months to complete? We made a number of improvements to speed the process.

1. We added LDCACHE to the DMF database file systems. This greatly sped up the job on tapes that had thousands of files (Note: Cray does not recommend this. A crash could corrupt the DMF databases.).
2. We gave priority to the move processes with the RENICE command.
3. To further speed up the move we changed from hard deleting the old Redwood data after successful move, to just leaving the Redwood data in the DMF databases. This greatly sped up the database access by cutting the number of database transactions in half.

4. We next cut back on regular DMF and DUMP CRON jobs to minimize system tape activity (and turn off DMF sparse and hard-delete scripts).
5. Watch out for log and journal files growing too large and filling file systems.
6. Borrow and daisy chain additional tape drives?
7. Automate tape error handling.

Ongoing Problems

We have had a number of on-going problems including dirty tapes, dirty tape drives, hung channels and processes. Senior operations staff and system staff worked together to create interrupt procedures to allow tape drive cleaning and maintenance, kill and restart procedures to clear hung processes or channels as well as normal gigaring reset procedures.

Summary

This task has now gone into a production and monitoring mode and operations support is able to efficiently keep the tape drives busy. We expect our eight terabytes of data to be successfully copied by the time I'm giving this paper, down from 100 days to 65 days, to 35 days.

The Supercomputer used in this investigation was provided by funding from JPL Institutional Computing and Information Services and the NASA Offices of Earth Science, Aeronautics, and Space Science.