



LIBRARY OF  
CONGRESS

# Fixity File Checking

Designing Storage Architectures

September 2012

Scott Rife



LIBRARY OF  
CONGRESS

# Problem

**“This is an Archive. We can’t afford to lose anything!”**

- Our customers are custodians to the history of the United States and do not want to consider the potential loss of content that is certain to happen at some point

## **Solutions**

- At least 2 copies of everything digital
- Test and monitor for the failures
- Refresh the damaged copy from the good copy
- This process must be as automated as possible
- Someday we’re going to lose something
  - What’s that likelihood?
  - What costs are reasonable to reduce that?



LIBRARY OF  
CONGRESS

## Word on the street

**USC Shoah (8 PB collection of videos from holocaust survivors) reported that they were seeing read errors on 1 out of every 5 T10Kb tapes they were staging for rewrite**

- We had seen no unrecoverable errors in staging
- We store a *sha1* digest for each file in an Oracle db
- Wrote some *perl* code to stage our oldest tapes and found no errors in 600 TB of content staged
- Clearly there are other issues contributing to these errors



LIBRARY OF  
CONGRESS

# Solution

## **Archive Integrity Checker**

- LOC contracted with a developer to write code to systematically check all the content and keep track of the status of each file
- Code runs via cron and checks content by date of ingest and last check
- Files that are to be staged are sorted by tape and then sorted by location on tape and staged in that order
- Oracle and IBM have alternatives to our homegrown solution



LIBRARY OF  
CONGRESS

# Oracle solutions

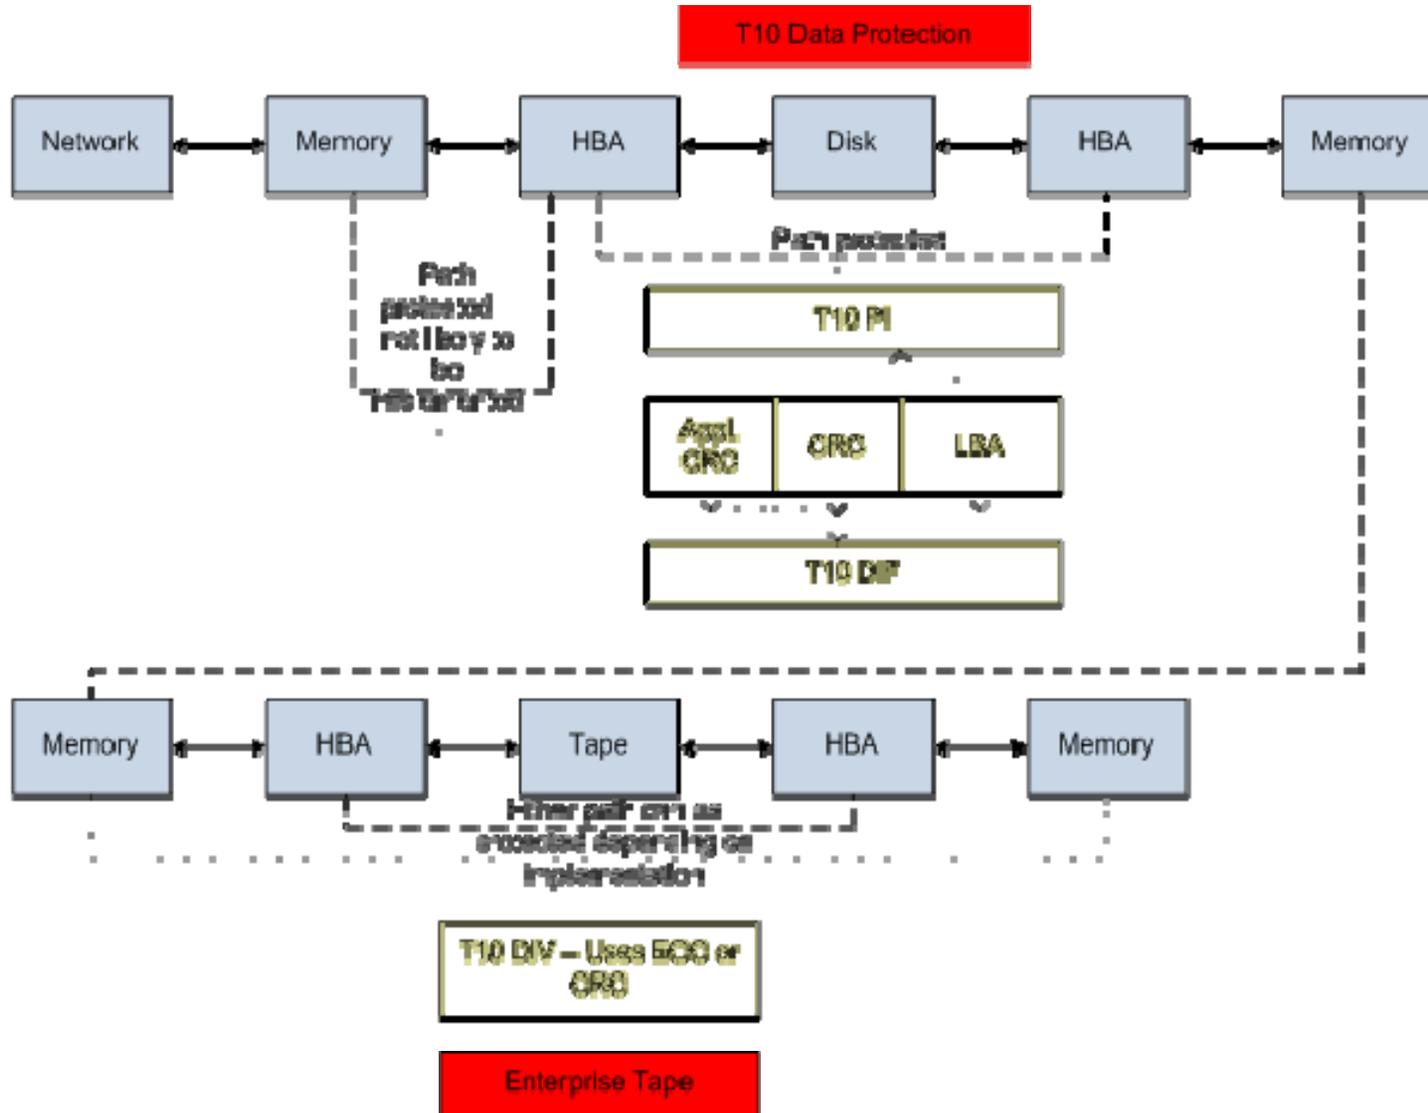
## Verification flag (ssum)

- Enabled via software for a collection of files (filesystem, directory, file).
- Generates a 32 bit running checksum for file and stores in SAM's meta data (viewable via *sls -D*).
- File staged back to disk from tape and verified
  - Can be time consuming and slow down archiving
  - We turned it off
- Can be turned on to generate only and used when files are staged back from tape later.



LIBRARY OF  
CONGRESS

# Datapath integrity field





LIBRARY OF  
CONGRESS

# Oracle solutions - continued

## **T10-DIV**

- A variation of the T10-DIF (Data Integrity Field) that adds a CRC to each data frame from the tape kernel driver to the tape drive
- Verifies the FC path to to tape drive and verifies each write to tape by reading afterward
- Requires Solaris 11 and SAM 5.3
- Can be used to verify the tape content without staging back to disk
- Oracle is working on a tape to tape migration that will use this information to validate the content read from tape during the migration



LIBRARY OF  
CONGRESS

# IBM HPSS solutions

- End to end check sums for data content integrity
  - User choice (ie *SHA1*, *md5*, *etc*)
- Tape to tape migration
  - Updates meta data with new tape location
  - Verifies using check sums
- IBM just released an LSI drive array that supports T10-DIF



LIBRARY OF  
CONGRESS

# Future

- Packard Campus will be using T10-DIV by the end of the CY 2012 on our T10KC tapes.
- Verification flag will not be used
- Oracle TA (Tape Analytics) will be used to monitor our tape drives/libraries
  - Drive code upgrade to 7.70 or 8.01
    - need ACSLS 8.x for the 8.X drive code (not available for Solaris 11)
  - Purchase HBT card for library (additional memory)