



Reliability and Recovery: How Can We Make Better Storage Devices?

Tom Coughlin


Coughlin Associates

www.tomcoughlin.com

Outline

- Modern Storage Devices
- What causes failure and data – all bits must die!
- Avoiding the loss of data
- Conclusions

Modern Storage Devices

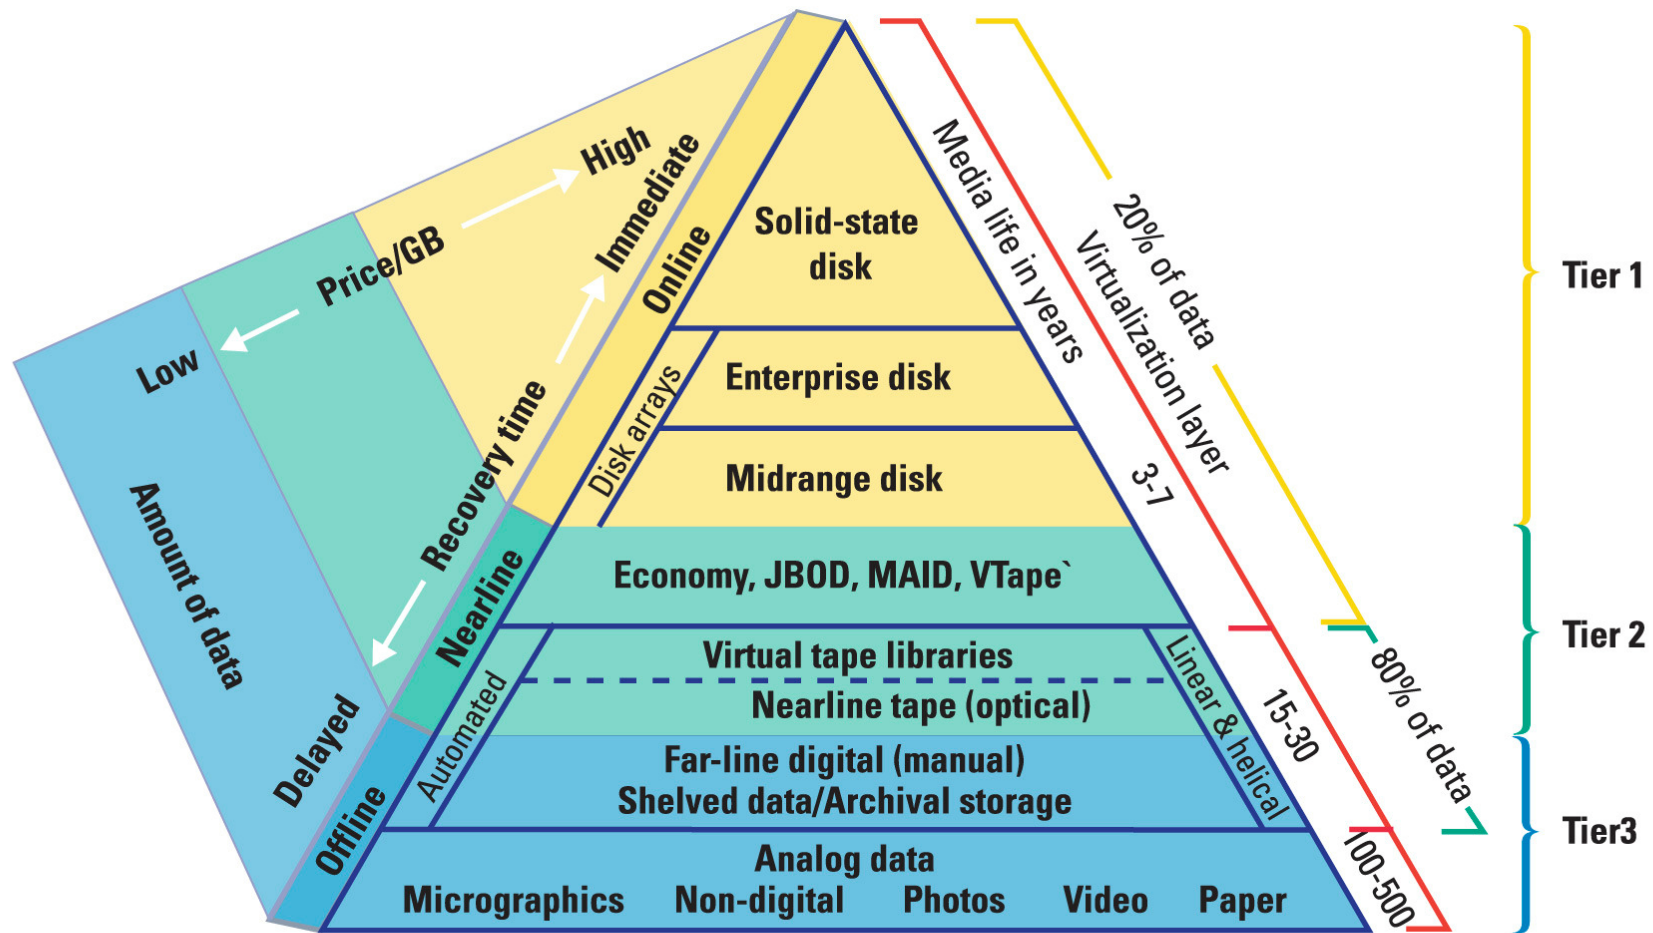


Common Storage Devices

- Hard disk drives
 - Performance drives (SAS and FC)
 - Capacity drives (SATA)
- Magnetic data tape
 - LTO 4, 5
 - Other formats
- Optical Discs
 - CD, DVD, Blu-ray
- Flash Memory
 - Card formats
 - SSDs



Traditional Enterprise Storage Tiers




Source: Horizon Information Strategies

A close-up photograph of a white marble statue of a winged figure, possibly Nike, with one arm raised. The statue is set against a background of light-colored stone columns. The text "All Bits Must Die!" is superimposed in a bold, black, cursive font over the center of the image.

***All Bits Must
Die!***

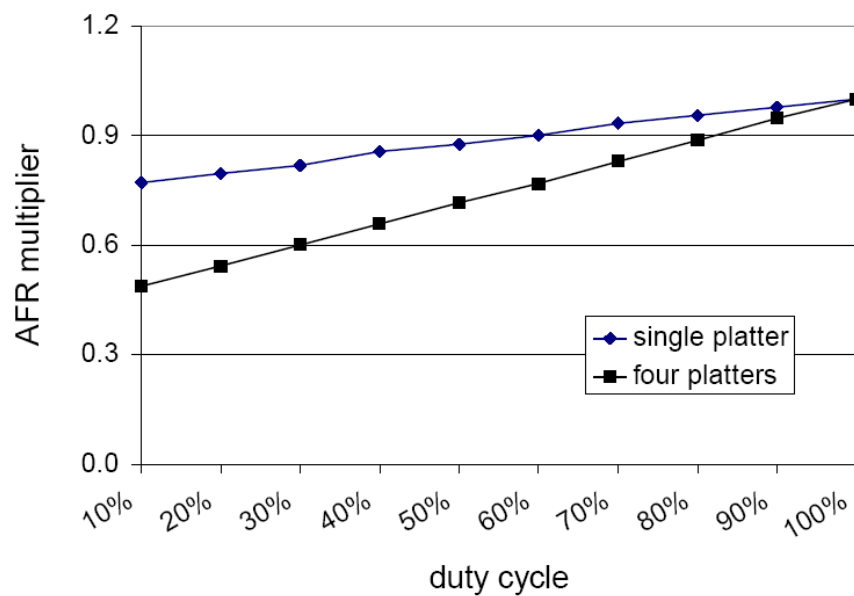


Failure Bathtub Curve



Some Causes of Storage Device Failure


- Hard Disk Drives
 - Head crashes from many causes
 - Stiction and failure of mechanical components
 - Firmware problems causing malfunctions
 - Shock and vibration
 - Power supply circuitry problems
 - Off-track and other issues leading to loss of servo data on tracks
- Flash Memory Storage
 - Wear of storage cells during erase/write steps
 - Firmware problems causing malfunctions
 - Power supply issues
 - Read disturbs and other operational sources of failure
 - Cosmic ray and other radiation damage
- Optical Discs
 - Scratches, pits and delamination of media
 - Firmware problems causing malfunctions
 - Power supply issues
 - Warping and deformation of media

HDD Reliability Factors



“Estimating Drive Reliability in Desktop Computers and Consumer Electronics Systems”, G. Cole, Technology Paper TP-338.1, Seagate Technology, Nov. 2000

Showing Combination of Various Failure Mechanisms Produces Overall Failure Rate




$h(t) = f(t)/R(t)$, $f(t)$ is the probability density function and $R(t)$ is reliability as a function of time

Reliability Analysis of Disk Drive Failure Mechanisms, S. Shah & J. G. Elerath,

When do Storage Devices Really Fail?


Typical HDD Failure Test Results

- In some cases the integrator or the manufacturer may not be testing the drive in the same way that it failed in (such as vibration induced failure or data loss in a storage array)




"From Server Class Disk Drives: How Reliable are They?", J. G. Elerath and S. Shah, Proc. Annual Reliability & Maintainability Symposium., January 2004

Causes of Storage Device Data Loss



- Data can be lost on storage devices even through the device itself may still be operable
- For instance some fault in OS or a premature system shut down may cause error in the file system making it difficult to reassemble data files
- Firmware faults in a storage device or other devices in the system can cause file corruption
- forensic data recovery techniques may be needed to get the data back

AVOIDING THE LOSS OF DATA




Improving Reliability or Living with Failure Prone Storage Devices

- Use SMART attributes to detect drive degradation and need to back-up or transfer data to new device
- What are the implications of content management approaches to drive life
 - E.g. frequent erasure of recorded data could shorten life of drive
- Back-up of device data
 - External drives (DAS)
 - Home entertainment network backup to NAS
 - Backup to cloud (remote) storage service

External Storage and Backup



Clickfree Backup



iVDR external storage



HP Media Server, V2



Any storage device becomes networked storage

Survival Storage

(External SSD Torture Test)

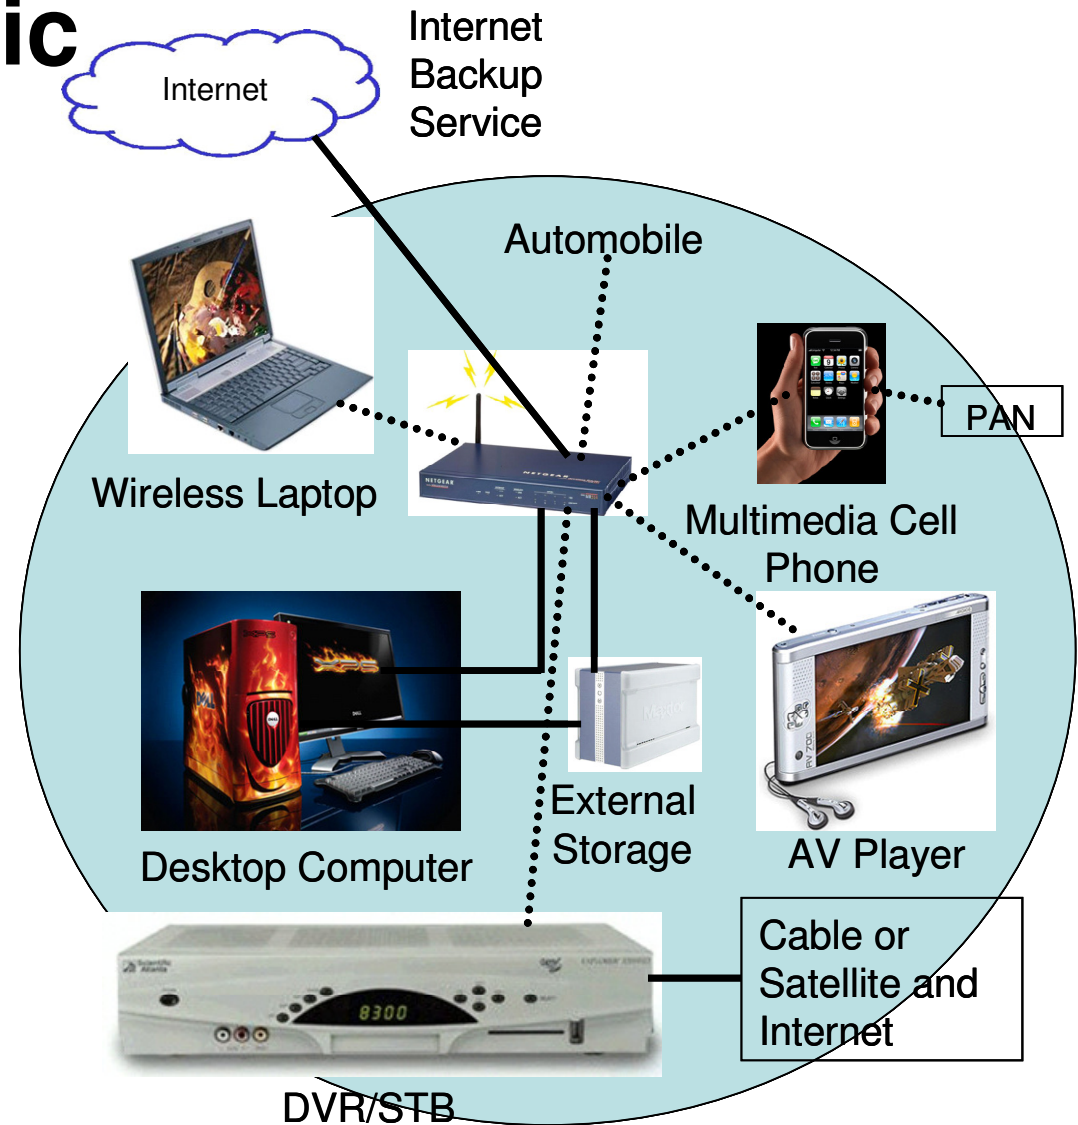


**ioSafe Demo-
Lition**

January 8, 2010

Home Storage Utility—Automatic Data Backup

- Everything is connected
- Content and storage is shared—many copies storage is cheap
- Content is managed, indexed and automatically backed up



Data Recovery Value and Cost




Figure 1. A graphical representation of the relationship between the value of lost data to the consumer and the cost of the recovery.


"Data Recovery from Solid State Drives", Scott Holewiniski, Gillware, Inc., 2010

© 2010 Coughlin Associates

Data Recovery from Storage Devices

(Discussion with a Data Recovery Company, Gillware, Inc.)

- The average cost of recovering data from a HDD is about \$694
- The average cost of recovering data from a SSD or other flash memory device is about \$2,850



SSDs are more expensive to recover data from because:

- They are newer on the market and the data recovery techniques developed over many years for HDDs are not always applicable for SSDs and other flash memory devices
- Also imaging of the device content, a common practice in data recovery, is difficult to do with multi-chip flash memory storage devices

Easier SSD Data Recovery

There are six ways by which SSD manufacturers can assist data recovery partners, helping to improve SSD data recovery success rates and reducing costs:

1. Provide technical details of the FTL
2. Supply documentation of vendor-specific ATA commands for firmware and FTL manipulation
3. Allow access to the appropriate cipher in the presence of hardware-level encryption
4. Provide information about the ECC implementation
5. Grant access to data sheets for SSD controllers and non-volatile memory
6. Supply controller emulation tools
7. Co-develop with data recovery company engineers improved systems for obtaining memory chip reads

“Data Recovery from Solid State Drives”, Scott Holewiniski, Gillware, Inc., 2010

Making Storage Devices Easier to Recover Data From

- Full device encryption is becoming common in many enterprises. Some sort of industry partnerships between storage device makers and data recovery companies would help to make it easier to recover self encrypted data.
- HDDs have a vendor-specific mechanism for manipulating device firmware over the ATA interface and sometimes through other means, such as an undocumented RS-232 or JTAG connection.
- SSDs need a similar way to access them,
 - certain firmware corruptions could potentially be repaired and the device restored to a functional state.
 - the ability to obtain the raw contents of the memory chip would be incredibly valuable.
 - In either situation the need to desolder and individually read each memory chip is eliminated.



Conclusions



- People need more storage, everywhere and always!
- We have more data on more things using more types of storage devices than ever before, and more chances to lose that data
- By understanding the causes of storage device failure and data loss we can take actions to reduce the odds of failure and data loss
- Backup and data protection will help prevent data loss even with the failure of one or more individual storage devices.
- Changes in the design of storage devices could make data recovery easier and less expensive when it is necessary

