

Disk Failure Investigations at the Internet Archive

Thomas Schwarz³, Mary Baker², Steven Bassi³, Bruce Baumgart¹, Wayne Flagg¹, Catherine van Ingen⁴, Kobus Joste⁵, Mark Manasse⁴, Mehul Shah²

¹ Internet Archive, San Francisco, CA

²HP Labs, Palo Alto, CA

³Santa Clara University, Santa Clara, CA

⁴Microsoft Research, BARC, San Francisco, CA

⁵Bixdata LLC, New York, NY




Purpose

- Better quantitative understanding why and how disks fail.
- Current numbers by disk manufacturers:
 - Focus not historical.
 - Determine quality of current products.
 - Determine budget for warranty funds.
 - Use artificially accelerated tests.
 - Do not address silent data corruption (bit-rot)
 - Large number of returned disks show no problem, but customer problem was presumably real.





Endeavour

- Distill data from already collected data by the Internet Archive.
- Data collection should be planned
 **Archive Observatory**
- Larger scale bit-rot experiments
 - Using donated IA surplus equipment.



Internet Archive

- Nonprofit that safeguards digital and digitalizable data
 - WayBack Machine:
 - Internet snapshots
 - Grows at 25 TB/Month
 - Stored in ~100MB archive files (ARC)
 - Other media: films, music, ...
 - Locations in Alexandria, Amsterdam, San Francisco,
 - Stores data on desktop ATA disks
 - Four disk pizza-box-form-factor



Historical Data Set

- MD5 values of each ARC file
- SMART statistics
- Kernel logs
- Replacement notebook entries



Chains

- Measurements have gaps / periods without data.
 - Caused by:
 - Failure
 - Temporary removal
 - Loss of measurements
 - Extraordinary states of system
- Notion of chains:
 - Contiguous sets of measurements of single object.
- Developed database scheme in preparation for the Archive Observatory.



Results

- Bit-rot data
 - Less than feared, but still substantial
 - Analyzed 112,865,205 MD5 for 4,717,158 are available.
 - Analyzed only 1,496,572 MD5 for 36,169
 - Observed some transient changes
 - Of 10 disks with a permanent change of MD5, 2 were ultimately attributed to disk failure.



Results

- S.M.A.R.T. data
 - Observed 136 “oscillating” chains
 - Attributed 94 to spurious indication of healthy state.
 - Remaining 42 chains represent 27 (out of 4978) disks.
 - 8 of these disks have left the population prematurely.
 - In all cases, SMART gave ample warning of failure.
 - Next step: Correlating SMART and bitrot data.



Results

- Disk Failure Rates
 - Archive **does not see** disk infant mortality
 - 2 incidences of unrelated batch failure
 - Attributed by manufacturer to shipping or handling mistreatment.
 - Disk failures now at 2.00%, but were as high as 6%
 - Combined failure rate (motherboard, cpu, memory, ...) 2.54%



Results

- Bit failure experiment
 - 2PB of data
 - So far, no genuine incident has been discovered
 - But many MD5 misreadings.
 - <http://www.us.archive.org/ao/>
 - <http://www.us.archive.org/ao/wip2.php>
 - <http://www.us.archive.org/ao/wip.php>



Lessons Learned

- Understanding of disk failure is becoming more important.
- Instrumentation and analysis is much more difficult than thought.
- Need
 - Other sites to instrument their observatory
 - More controlled experiments