
OPEN
Compute Project
OCP U.S. SUMMIT 2017
Santa Clara, CA

SEAGATE

Re-evaluating HDD Error Correction

Bit Error Rates, UER and Cloud Hard Drives

David Burks

Seagate Technology

OPEN HARDWARE. OPEN SOFTWARE. OPEN FUTURE.

OPEN Compute Project


!
CAUTION

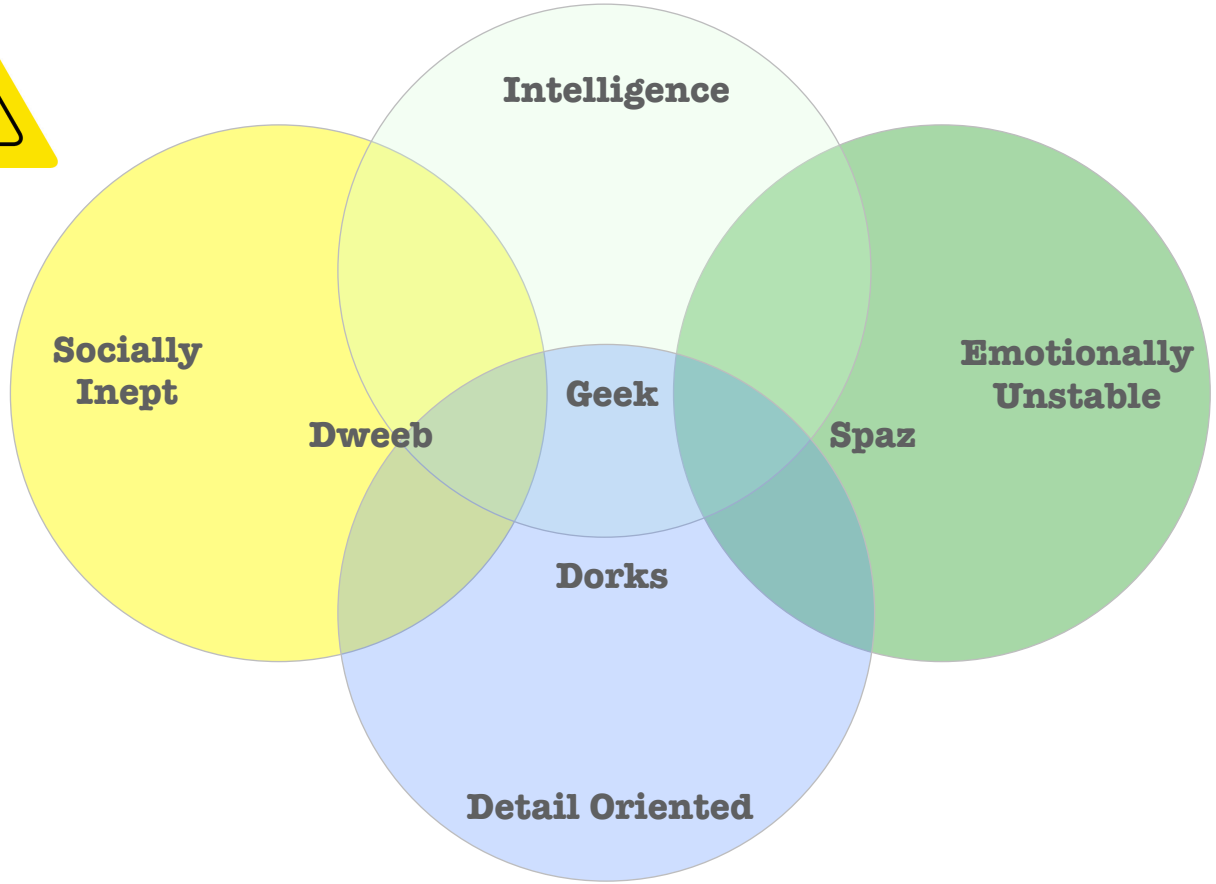
Intelligence
Socially Inept
Emotionally Unstable
Dweeb
Geek
Spaz
Dorks
Detail Oriented

A Short History of HDDs & Error Correction

The 50’s

IBM RMAC 305

Error Correction Capabilities

What’s Error Correction?

Our reliability spec insures that 70% of our drives power up straight out of the box!

The 60’s

IBM 1311

Error Correction Capabilities

Be really careful with it!

- “Precaution should be exercised to prevent foreign particles from entering the 1311 disk pack…”
- “Use of cleaning implements that raise dust such as brooms and feather dusters should be avoided.”

The 70’s

IBM 3340

Error Correction Capabilities

We’ll tell what might work…

3340/3344 Recovery Action Table (continued)

Action	Explanation
8	a. Perform error correction function. b. Examine bit 7 of the file mask. If this bit is off, go to step c. If this bit is on, return to user with indication that data has been corrected. (User is operating in PCI fetch mode and must supply restart recovery action.) *Note: Only applies with OS/360.* c. Increment the seek argument by one. Cylinder bytes and the high-order head byte are obtained from the user. The low-order head byte is obtained from bits 3 through 7 of sense byte 6. d. Construct Restart CCW 2. e. Complete the interrupted operation and continue the user’s chain (if appropriate) by executing the following command chain. Seek (argument from step c)† Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC* –8 Restart CCW 2 TIC (channel status word) *Note: If the modified seek argument is not within the user’s extent, then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must perform Action 2.*

The 70’s

IBM 3340

Error Correction Capabilities

We’ll tell what might work…

3340/3344 Recovery Action Table (continued)

Action	Explanation
8	a. Perform error correction function. b. Examine bit 7 of the file mask. If this bit is off, go to step c. If this bit is on, return to user with indication that data has been corrected. (User is operating in PCI fetch mode and must supply restart recovery action.) *Note: Only applies with OS/360.* c. Increment the seek argument by one. Cylinder bytes and the high-order head byte are obtained from the user. The low-order head byte is obtained from bits 3 through 7 of sense byte 6. d. Construct Restart CCW 2. e. Complete the interrupted operation and continue the user’s chain (if appropriate) by executing the following command chain. Seek (argument from step c)† Set File Mask (same as original) Set Sector (argument 0) Search ID Equal (record 1) TIC* –8 Restart CCW 2 TIC (channel status word) *Note: If the modified seek argument is not within the user’s extent, then IOS must supply the correct seek argument before issuing the seek. If that is impossible, then IOS must perform Action 2.*

The 80’s

Seagate ST 506

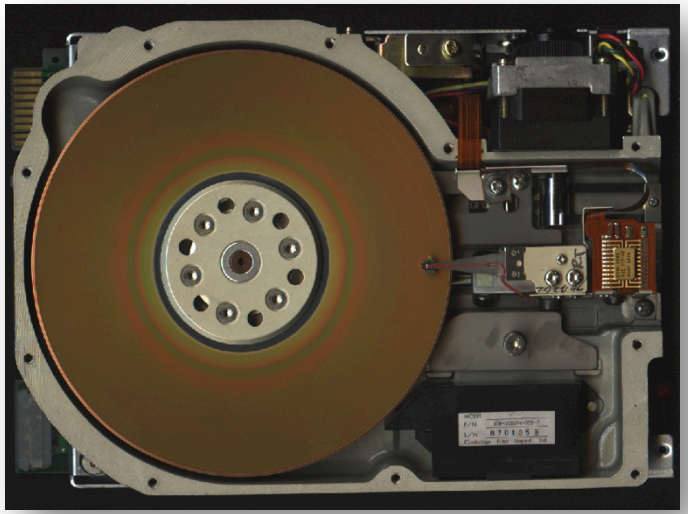

Error Correction Capabilities

Published UER!

1 bit 100,000,000,000 will be uncorrectable
(10 e12)

The 90’s

Seagate ST 4051

Error Correction Capabilities

UER Improvements…

1 bit 10,000,000,000,000 will be uncorrectable
(10 e14)

The 21 Century

Seagate 15K Enterprise

Error Correction Capabilities

The era of modern ECC!

1 bit 1,000,000,000,000,000 will be uncorrectable
(10 e16)

And the Era of Cloud Computing….

What is a Cloud HDD?

They don't yet exist!

But if they did, would they need the awesome error correction features of modern HDDS?

Replication

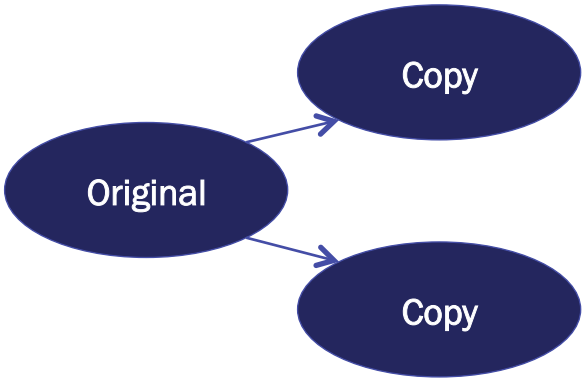

Erasure Coding

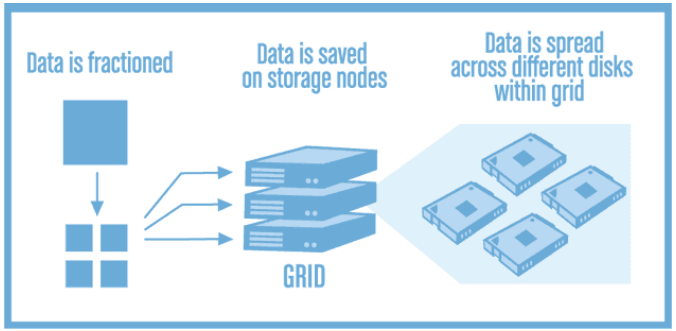

Error Correction and Tail Latency

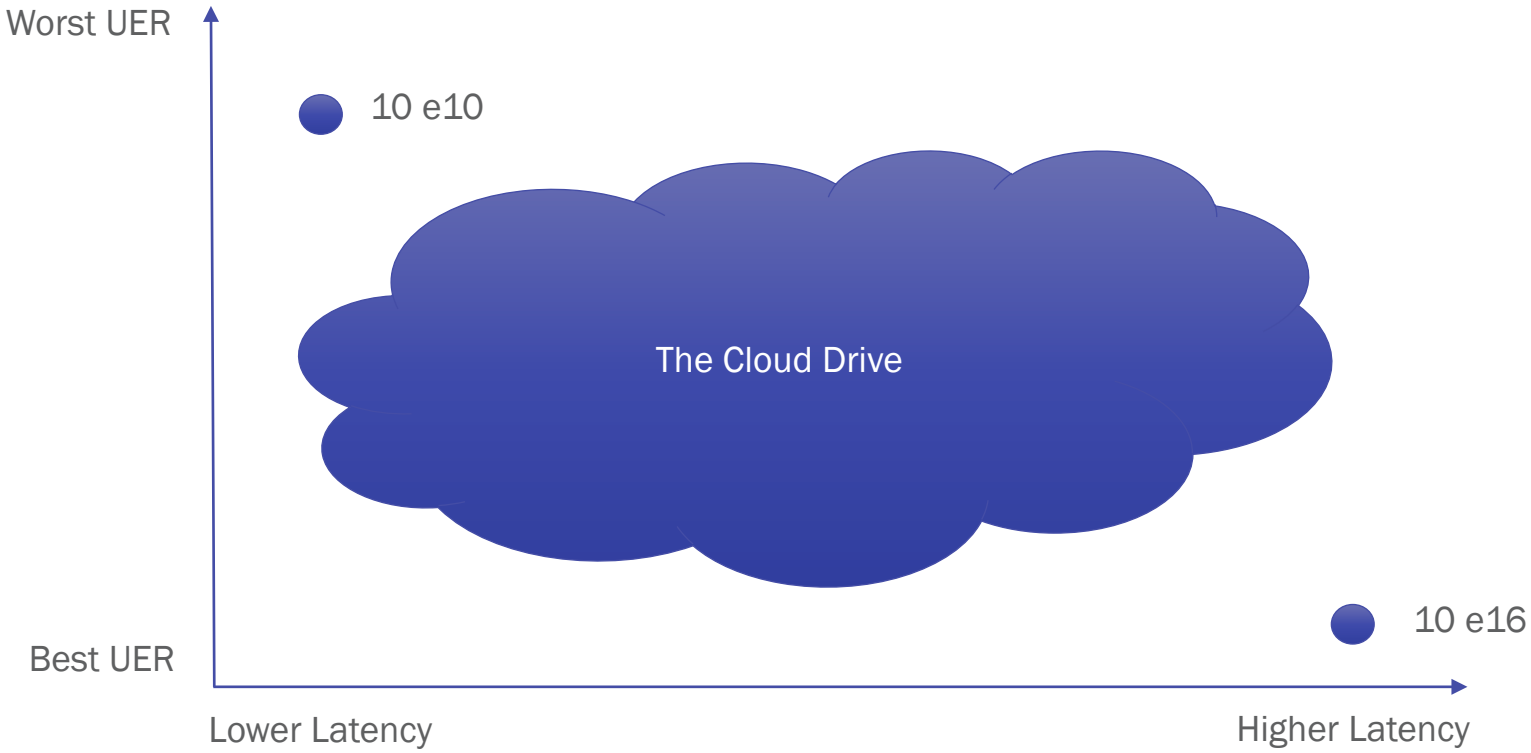

Back to Future

- What if we took UER back to 1980?
- What might that achieve for a Cloud HDD?

NEAR TERM

- ✓ Better command completion times (CCT)
- ✓ Lower tail latency & improved SLA (better performance)

LONG TERM

- ✓ Improved HDD yields
- ✓ Better TTM for new capacities
- ✓ Reduced costs for the HDD ecosystem driving lower $/TB

A Classic “Chicken & Egg” Challenge

- Error rate assumptions are “baked” into many, many aspects of HDD design & manufacturing.
  - ✓ Head & media engineering & manufacturing criteria
  - ✓ ASIC design
  - ✓ Servo-mechanical capabilities
  - ✓ Firmware
  - ✓ Test processes

This opportunity demands industry collaboration!

Key Questions to Pursue

- Is there a “sweet spot” of reduced UER common the the CSP community?
- How would this best be determined?
- Can OCP be leveraged to organize, standardize and operationalize opportunities like these for the benefit of the booming Cloud market place?

QUESTIONS, COMMENTS & RIDICULE

Visit our Booth B7

Learn about the industry's broadest catalog of storage devices

ENTERPRISE & CLOUD

PCIe SSD

SAS SFF SSD

SATA SFF SSD

15K SFF HDD

10K SFF HDD

7.2K SFF HDD

Performance 3.5 HDD

Traditional Nearline

Content Nearline

Kinetic HDD

Content SMR HDD

CLIENT (DESKTOP & MOBILE)

Desktop HDD / SSHD

Laptop HDD / SSHD

Client SSD

Ultra Mobile HDD / SSHD

Ultra Thin HDD

NAS

Enterprise NAS

NAS HDD

SURVEILLANCE

Traditional Nearline

Surveillance HDD

VIDEO & MEDIA

Video 3.5 HDD

Video 2.5 HDD

Presenter Information:

DAVID BURKS

Director
Product Line Management

David,P.Burks@seagate.com

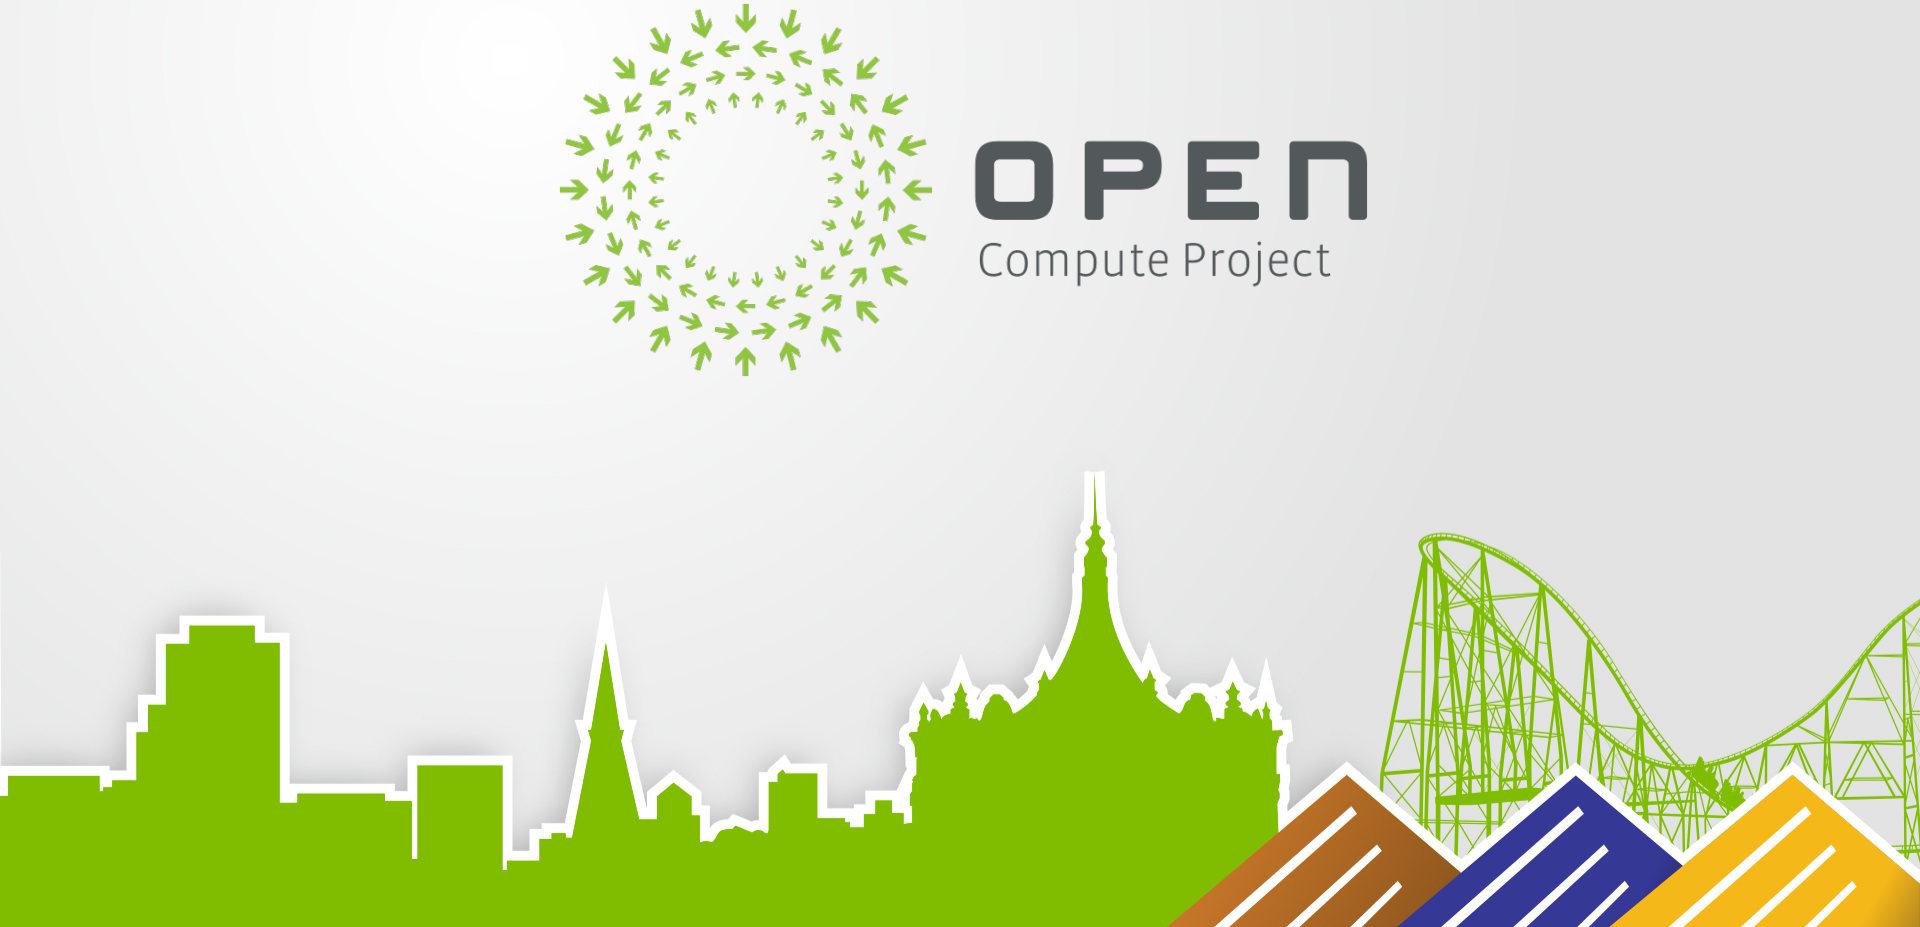
OPEN
Compute Project