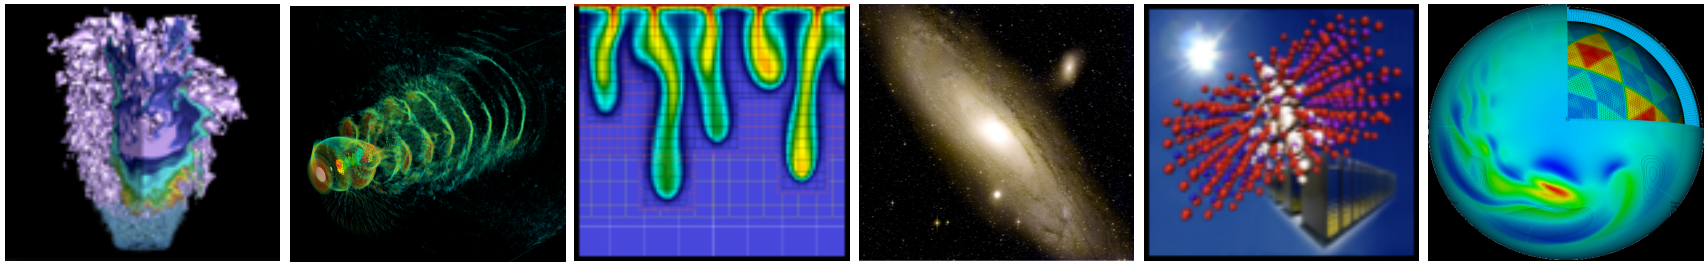

MIR Analysis

Damian Hazen, Jason Hick,
Wayne Hurlbert, Mike Welcome

April 19, 2011

U.S. DEPARTMENT OF ENERGY | Office of Science

NERSC National Energy Research Scientific Computing Center

BERKELEY LAB Lawrence Berkeley National Laboratory

Understanding Our Goals

- **Conduct collaborative research on tape technologies**
- **Provide feedback to Oracle tape engineering about information that customers need in order to improve "storage planning" and "tape system operations"**
- **Find out if the MIR contains this information**
- **If so, encourage Oracle to make the information available for customers**
  - Work with partners (Crossroads, Quotium) that already have expertise in the area of tape system/drive analytics
  - Integrate with existing products (SLC, ACSLS, …?) to prevent having yet another tool. For example, something we can enable or disable at will, that will then collect information and store it somewhere we already have today upon each mount/dismount.
  - Provide superior analytics for real-time assessment of tape operations (problem diagnosis) or storage planning (decision making)

Two Main Uses for MIR

- **MIR Assisted Search: Locating bad data, positioning around bad data, seek optimized to particular tape layout.**
  - Data is laid out in tracks and wraps (serpentine)
  - Positioning information would enable the application to know the optimal way to access a set of files on tape
  - Or enable a tape recovery utility to read until error and position around the error
- **Problem analysis: Identifying problem tapes, categorization of problems, breadth of problems, trending over time, proactive response**
  - Population analysis: Given a population of tapes with their MIR information, can I determine any trends that are important to operation of the system:
    - What are the top x% of tapes with permanent read (or write) errors?
    - What are the top x% of tapes with soft read (or write) errors?
    - What tapes have been mounted in a certain drive or written at a certain microcode level?
  - Mount analysis: Given what the stats looked like before and now, should action be taken?
    - Files that couldn't be read or written
    - Per mount transfer statistics (bandwidth, amount of data, cartridge)

Part of Media Information Record

- **Looked at T10KA/B information which provides:**
  - Lifetime Statistics
    - Serial number of tape
    - Meters of tape
    - Number of mounts
    - Number of permanent errors
  - Per Mount Statistics
    - Meters of tape
    - Read/write bytes
    - Number of soft errors
    - Time spent positioning
    - Location of error (need location MIR info to make sense of it)
    - Drive error occurred on
    - Microcode level of drive

The Process

- **The tape must be mounted in the drive. Unfortunately, current mount's MIR information not sync'ed to tape until dismount.**
- **However, made attempt to correlate FSC information with specific mounts.**
  - Knowing you can do this means you could correlate client (transfers)/host device information with tape drive/cartridge information. That would be very useful!
- **Added SCSI commands to extract the MIR information from the tape into our HPSS mover code before dismount.**
- **Collected this information into a database for each tape dismount since July 2010 (over 400,000 MIR mount stats, ~10,000 per week).**
- **Used this database to produce and analyze the statistics (our environment has 92 T10KB drives and about 20,000 T10KB cartridges).**
- **The key is in knowing how to interpret the MIR records and having the FSC dictionary.**

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Preempting Failures

- **From the analysis we've conducted, we don't think this is possible, but we intend to conduct more research in this area.**
- **There are indicators of problems**
  - For instance, error correction
    - There are several different parameters that provide the degree of error correction.
    - Problem is, our statistics show that they aren't absolute indicators of unreadable tapes/data. And they don't appear to change gradually.
    - As technology shrinks, degree of error correction is increasing and expected to increase.

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read and write errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Permanent Write Errors per Volume

Since we have a history of fault symptom codes, we can trend them over time. Also have number of blocks so we can identify the volumes with the greatest amount of errors.

```
[hpssdb@flanders bin] ./mir_errors.pl -o write

Mounts resulting in permanent write errors:

VOLUME   DATE       TIME     WRITEBYTES   READBYTES    FSC  BLOCK
ED000000 12/16/2010 15:02:09 128          62           3627 0
ED000100 12/16/2010 15:02:02 128          62           3627 0
ED000100 12/16/2010 22:48:51 11059657472  24793996     3773 9
ED000100 12/17/2010 00:58:51 16285949124  24793996     393B 234683
ED000100 12/17/2010 15:43:53 4141281384   41769360     3773 6
ED000100 12/17/2010 16:28:53 7372263864   24793996     3773 10
ED000100 12/17/2010 16:28:53 7372263864   24793996     48C5 314545
ED000100 12/17/2010 16:58:53 5658387628   24793996     3773 9
ED000100 12/17/2010 23:28:53 609054704    41769360     3773 10
ED000100 12/18/2010 06:28:54 65792028     24793996     37F6 522547
ED000100 12/20/2010 08:14:00 793644880    24793996     3773 215
ED000100 12/20/2010 14:39:01 43553460     24793996     37F6 531652
ED000100 12/20/2010 20:04:02 998174280    24793996     3773 21
ED000100 12/20/2010 21:14:02 2394213876   33281678     3773 10
ED000100 12/20/2010 21:14:02 2394213876   33281678     37F6 579154
ED000100 12/20/2010 22:24:02 3152878092   24793996     3773 58
ED000100 12/20/2010 22:59:02 1264061924   24793996     3773 27
ED000100 12/21/2010 02:09:03 1088093156   24793996     3773 15
ED000100 12/21/2010 02:39:03 1520575316   24793996     3773 215
```

Plotting Count of FSC Write Errors each Day

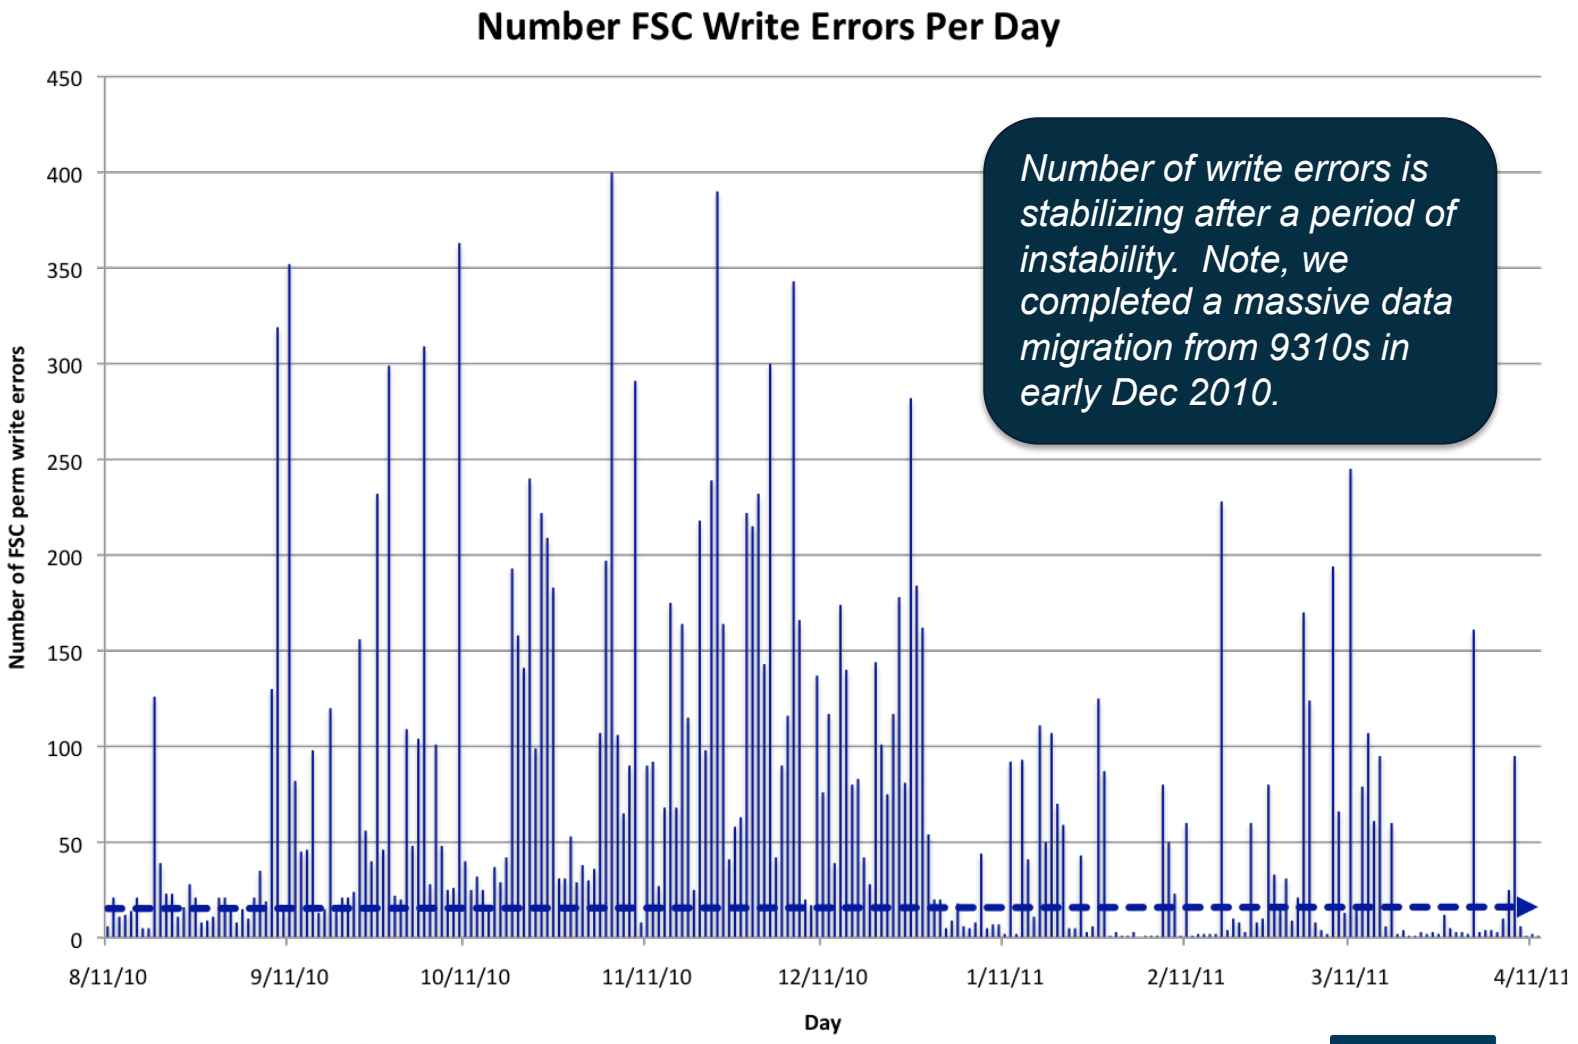

Plotting Count of FSC Read Errors each Day

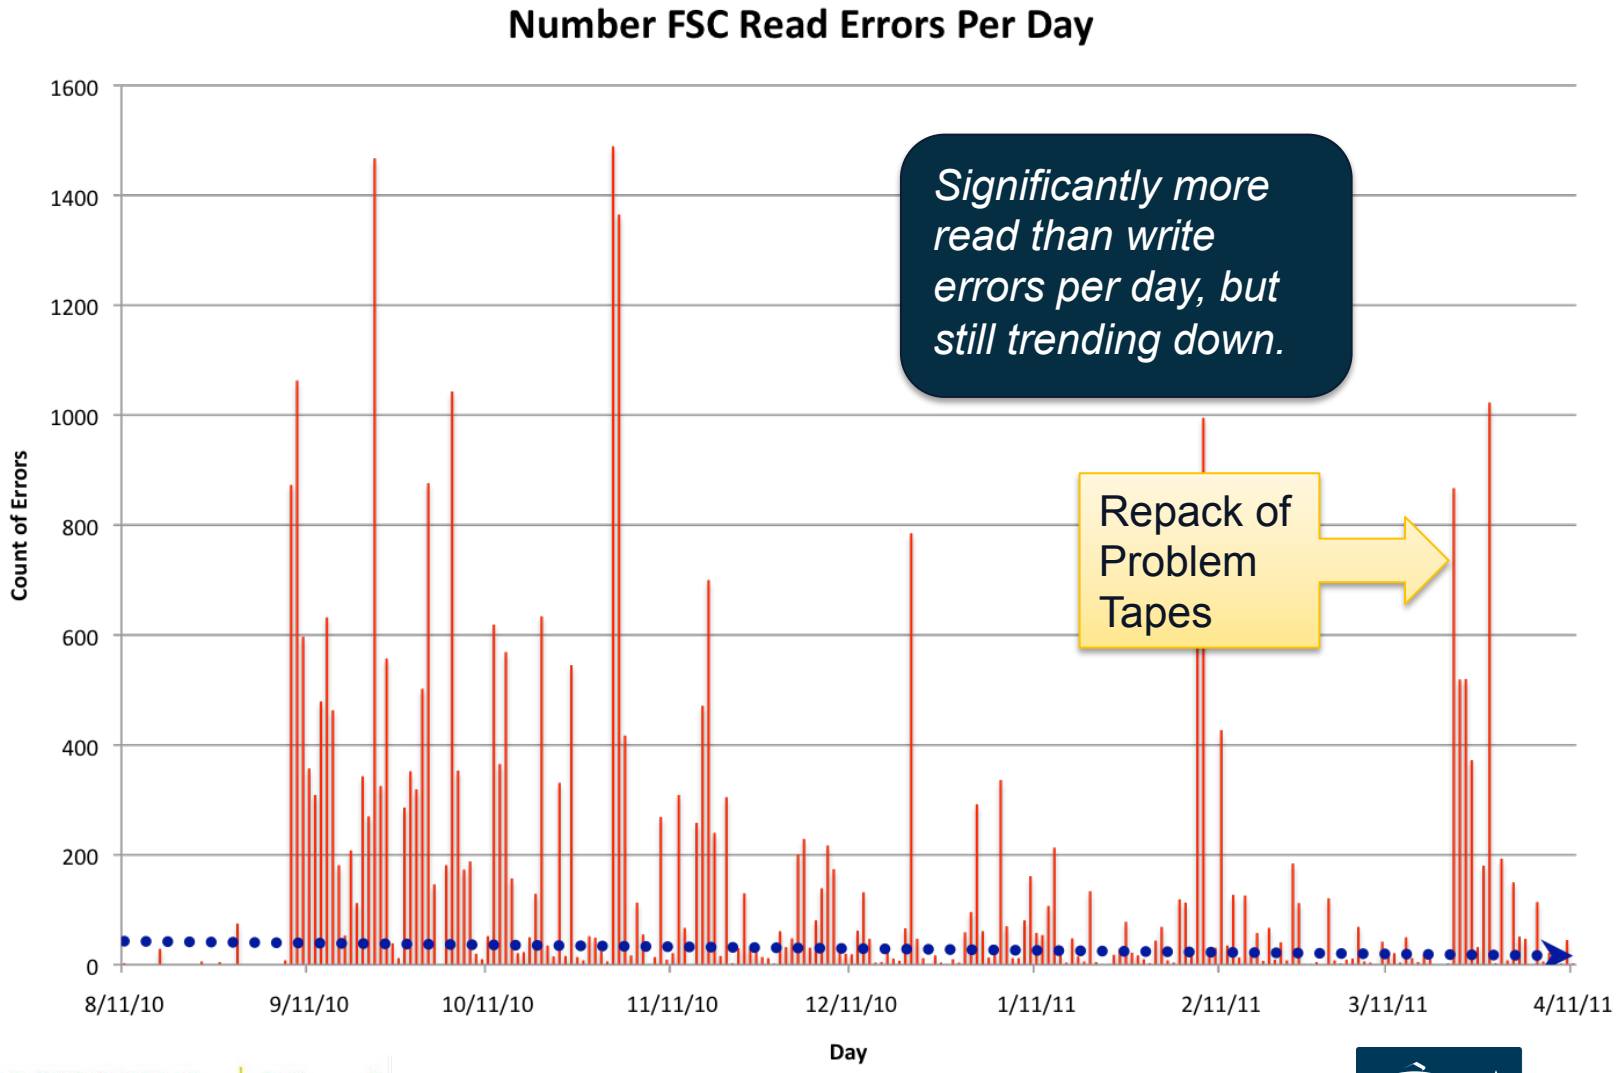

Plotting Count of FSC Errors by Code

Log scale

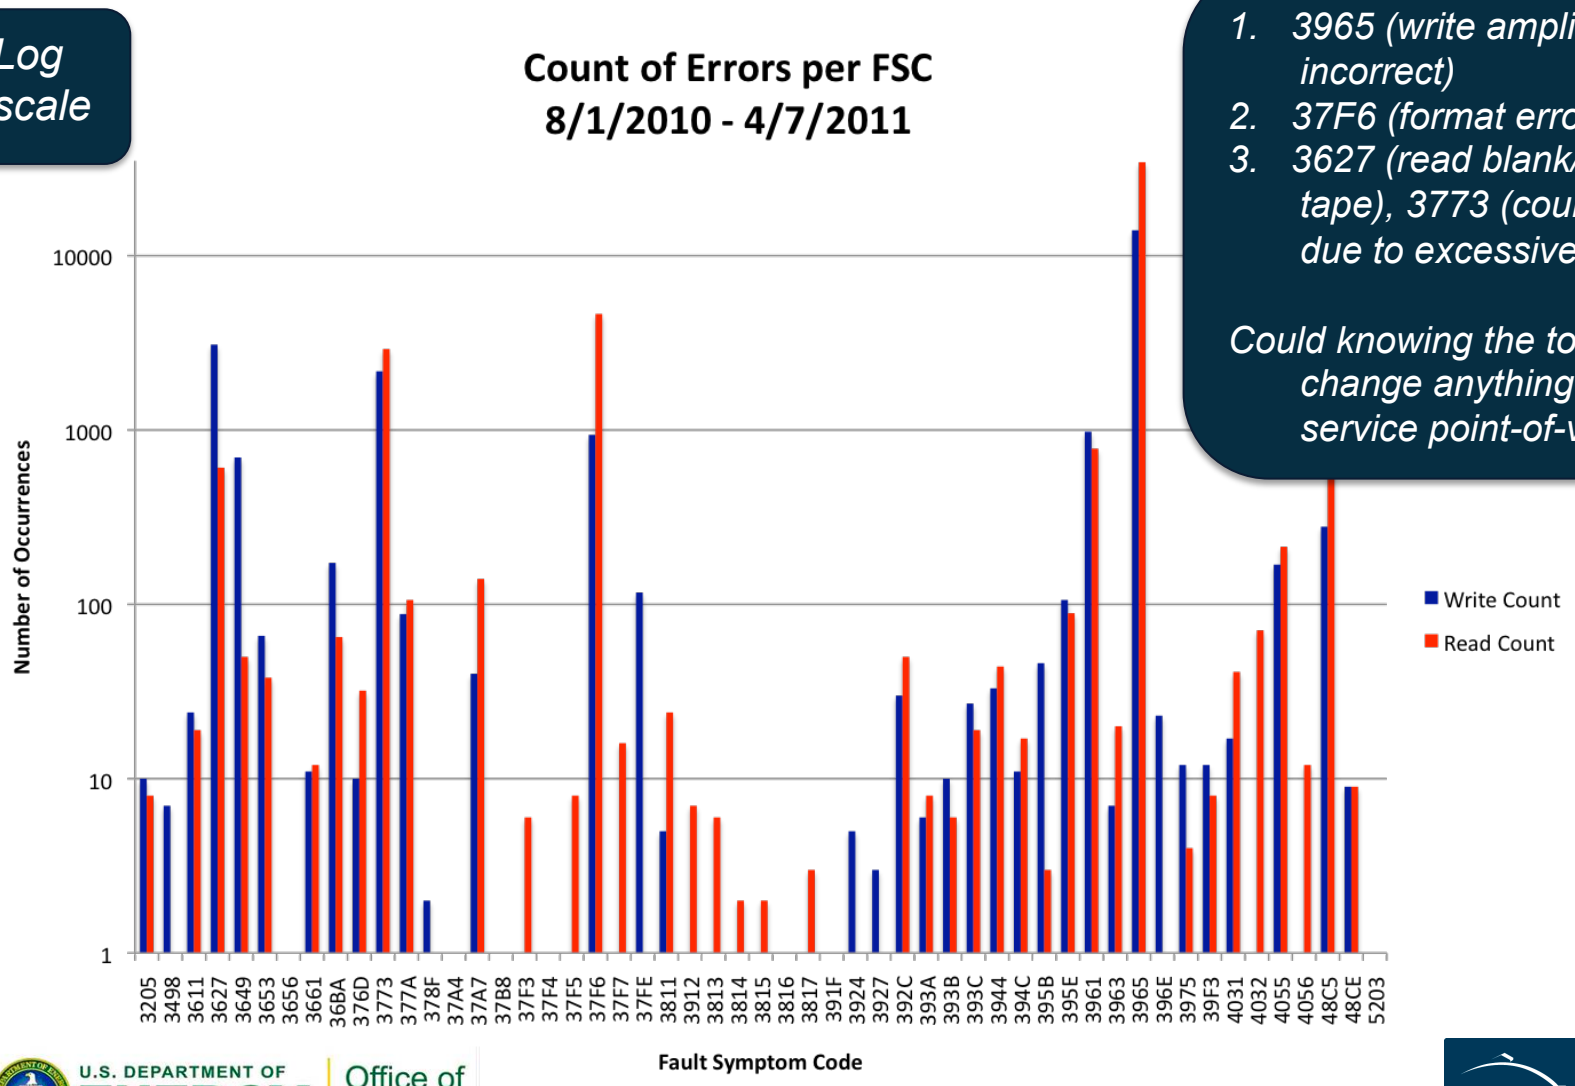

1. *3965 (write amplitude incorrect)*
2. *37F6 (format error on read)*
3. *3627 (read blank/VOLSAFE tape), 3773 (could not read due to excessive errors)*

Could knowing the top few change anything from a service point-of-view?

So why is that useful?

- The trends aid a site in knowing that both the SE and admins/operators are making positive changes to the tape subsystem (e.g. policy on replacing drives is effective or not, stable microcode levels for that site, …)
- Enables comparing across sites to understand site specific vs. systemic problems.
- Prioritize problem resolution (e.g. which errors are affecting us the most?, which are most serious?)

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Permanent Read Errors per Volume

flanders.nersc.gov — ssh — 49×22

```
[hpssdb@flanders bin] ./mir_volume.pl -o read

Drives volumes had perm read errors on:

VOLUME
        DRIVESN BLOCK

EM000000
        572004003872 1294556
        572004003872 1302917
        572004003872 1942861
        572004003872 4456623
        572004003872 4459186
EM000200
        572004004231 1063150
        572004004231 1065730
        572004004231 1077758
        572004004231 1084474
        572004004231 1090378
        572004004231 5801869
EM000400
        572004003815 1184692
```

Volumes, the drives they failed to be read on, and the block number of blocks for data that couldn't be read for that mount & drive.

Plotting Count of Bad Blocks on Reads

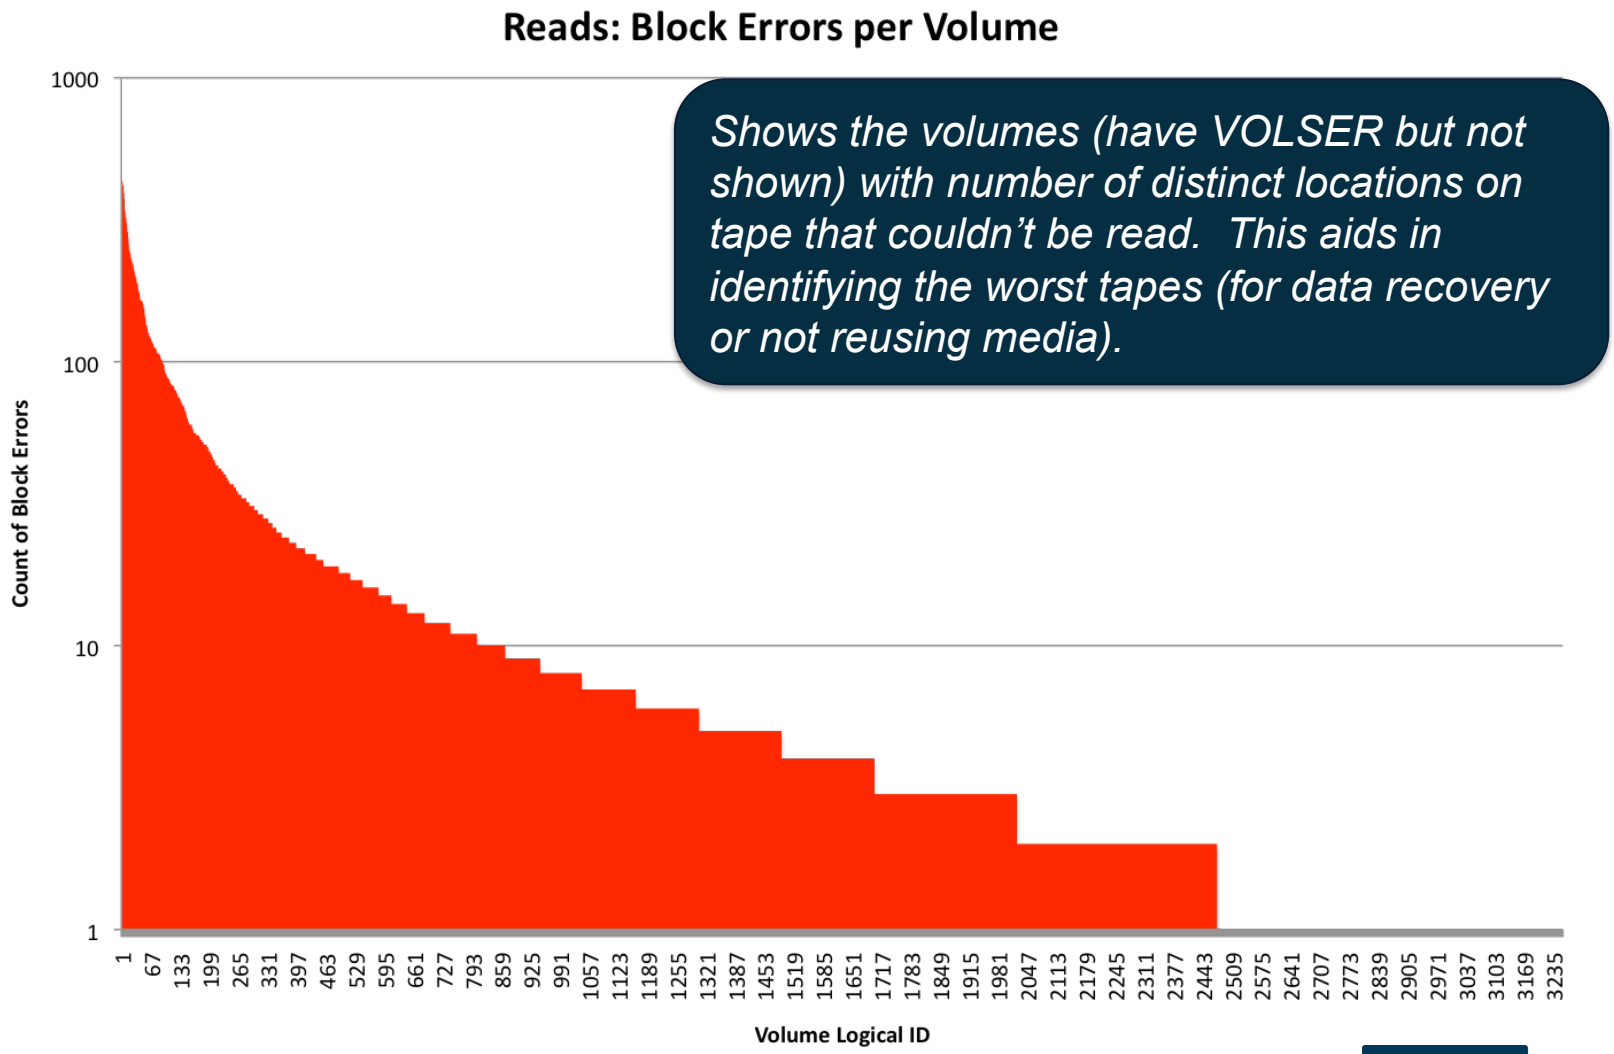

Shows the volumes (have VOLSER but not shown) with number of distinct locations on tape that couldn't be read. This aids in identifying the worst tapes (for data recovery or not reusing media).

Plotting Count of Bad Blocks on Reads

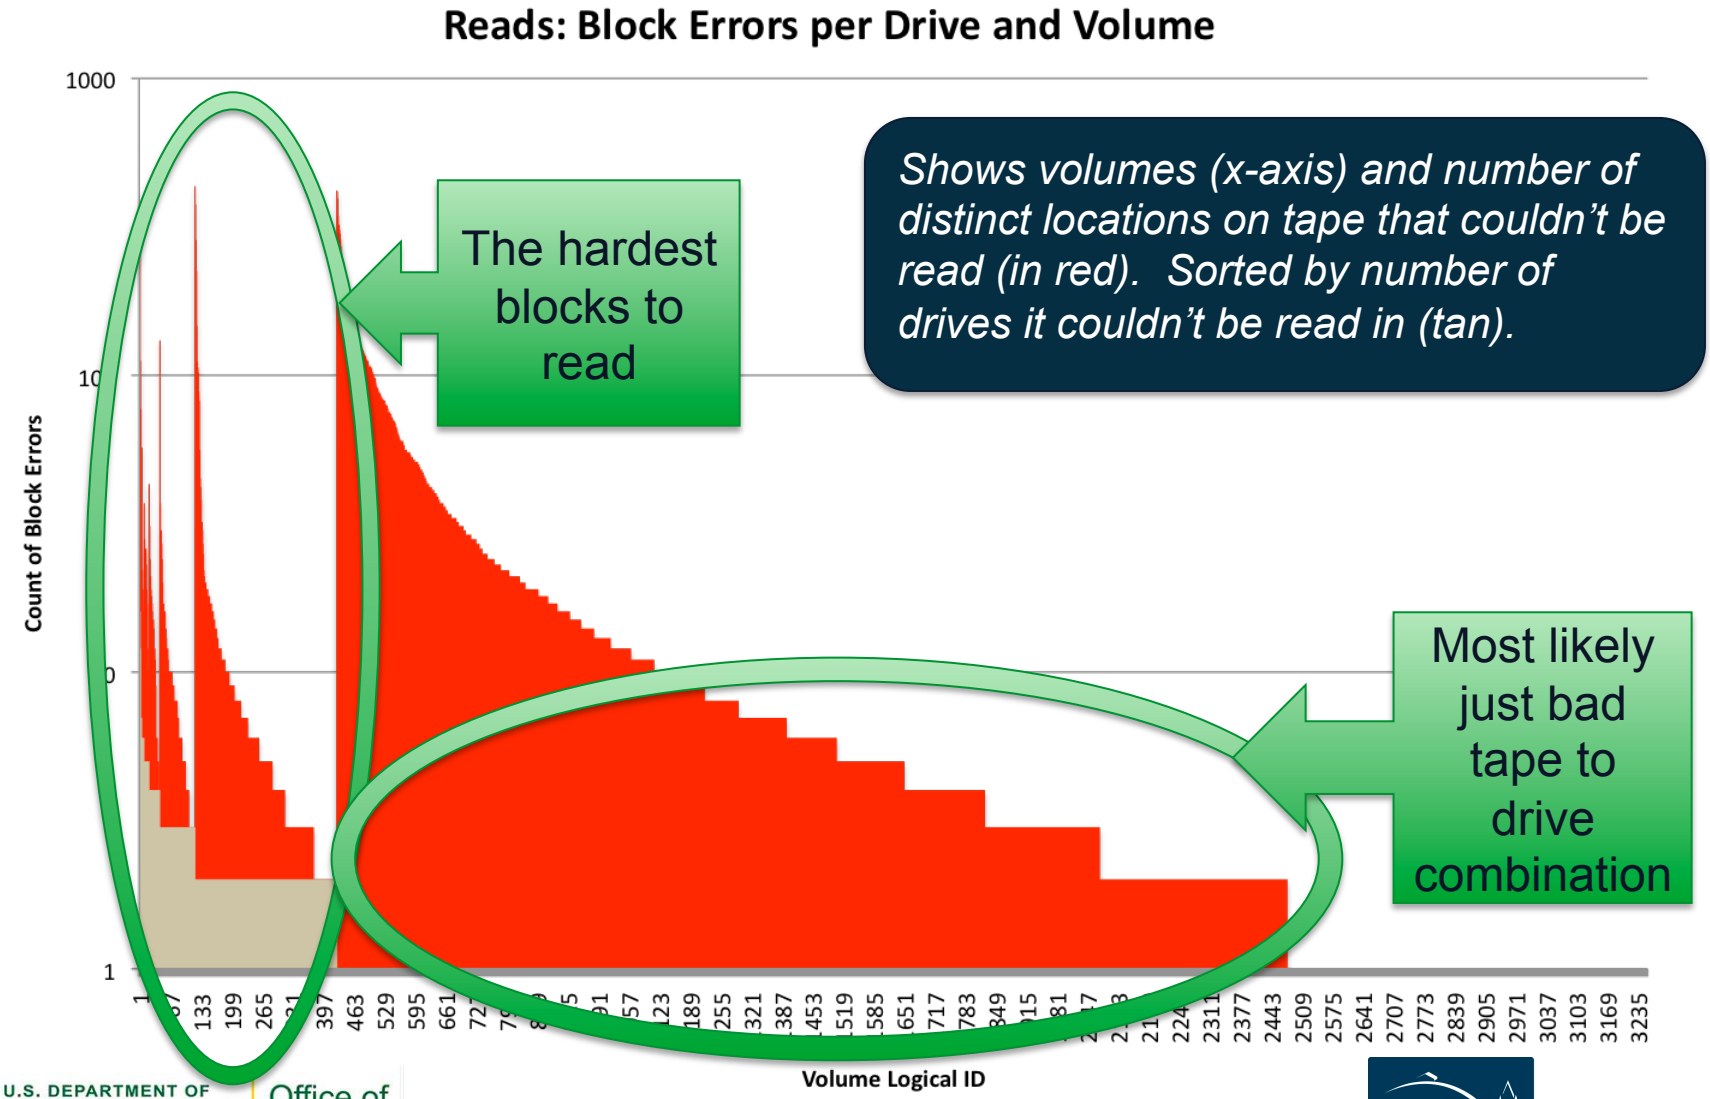

Plotting Count of Bad Blocks on Writes

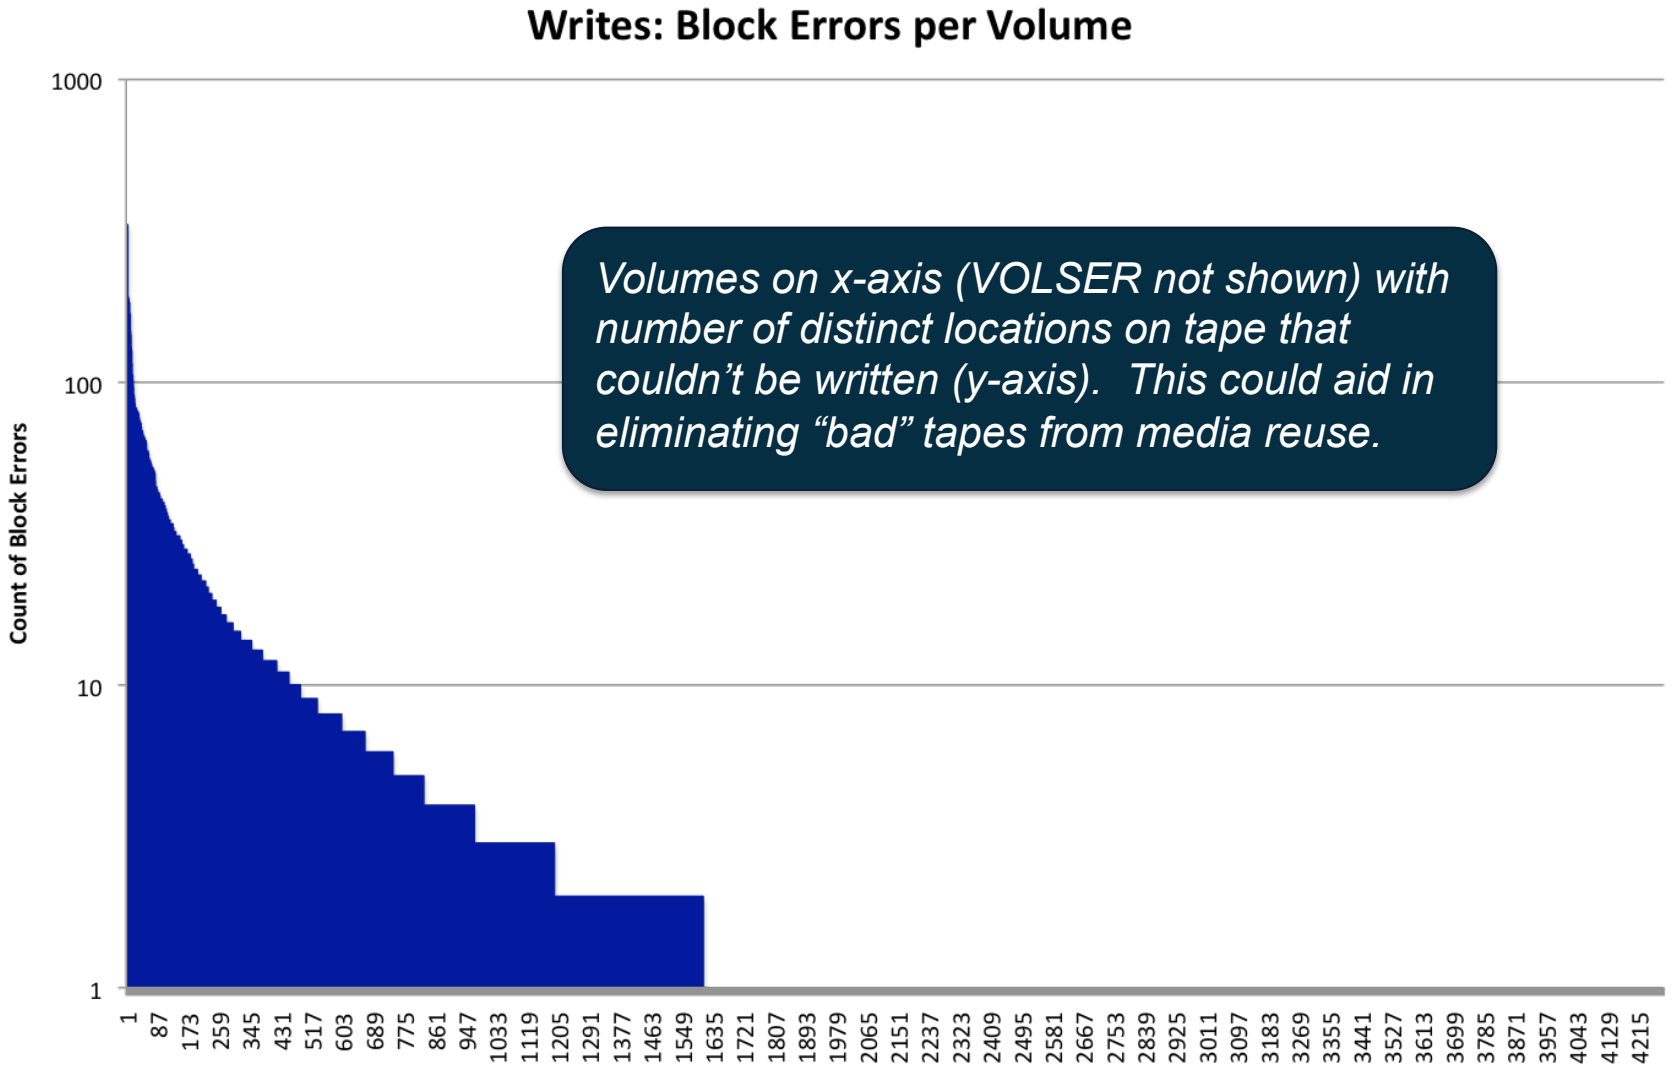

Volumes on x-axis (VOLSER not shown) with number of distinct locations on tape that couldn't be written (y-axis). This could aid in eliminating "bad" tapes from media reuse.

Plotting Count of Bad Blocks on Writes

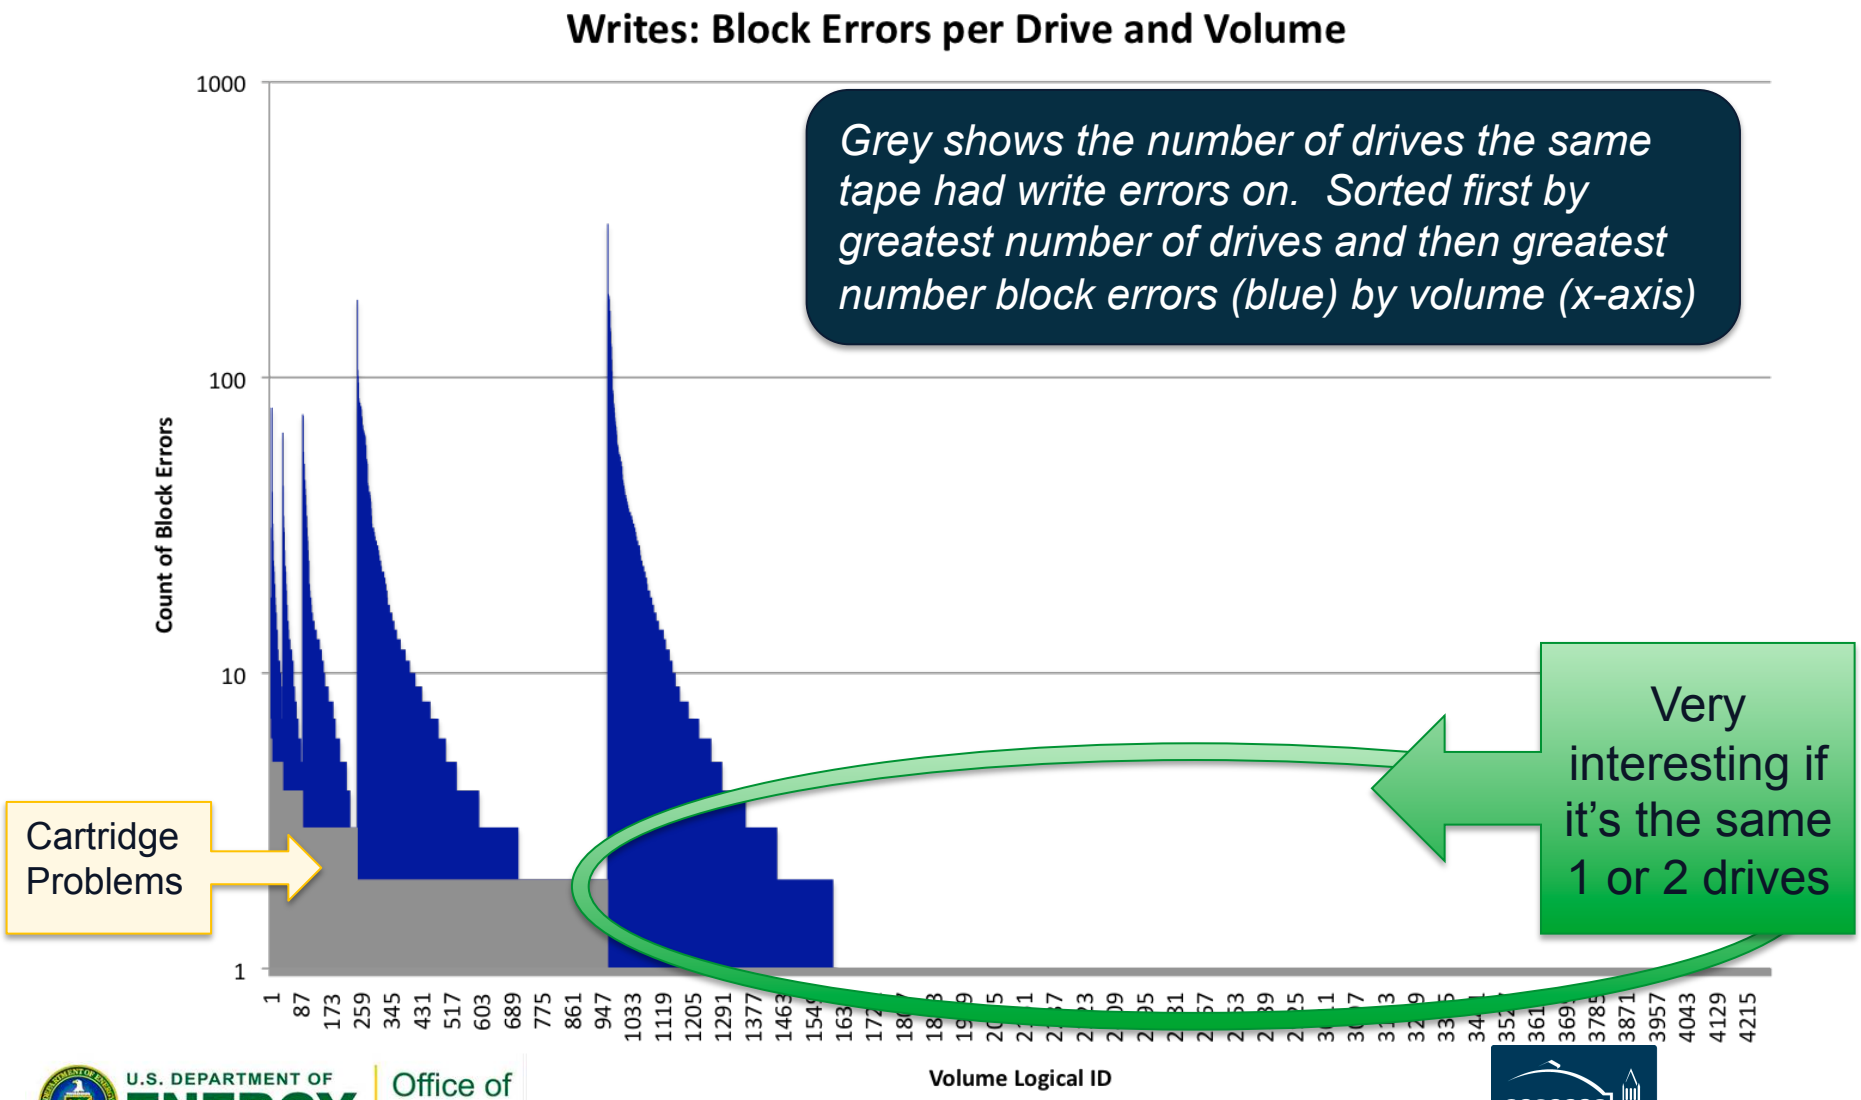

Recap for Volume/Drive Error Statistics

- Aid in determining media vs. drive issues, at least mismatches (i.e. this cart has trouble in this drive)
- Can identify the “worst” X number of tapes to avoid reusing them
- Can identify the tapes with the most blocks that can’t be read for data recovery or avoiding media reuse
- Can easily trend these over time to get an idea of whether tape subsystem is getting more or less error prone

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Looking at Errors after Grouping Volumes into Lots

flanders.nersc.gov — ssh — 64×25

```
[hpssdb@flanders bin] ./mir_lot.pl -i 1000

Volumes and their error statistics:

VOLUME VOLUMESN NUMERRORS LOT
ED000100 00000051007518431630 72 6
ED000200 00000051007518535130 155 6
ED000300 00000051007516562530 1025 6
ED000400 00000051007519401930 48 6
ED000800 00000051007501425030 2 6
ED001100 00000051007516461730 39 6
ED004900 00000071014805022834 2 10
ED005000 00000071014705023934 8 10
ED005200 00000071014705024134 4 10
ED006100 00000071011206002234 26 10
ED006500 00000071011204003234 2 9
EM001800 00000050820510150430 8 2
EM004800 00000050820509253830 1 2
EM014300 00000050821823085230 2 3
EM014800 00000050821801335130 100 3
EM018500 00000050820407383930 7 2
EM025200 00000050820300531530 2 2
EM025400 00000050820300560730 4 2
EM026800 00000050818819580430 2 2
EM027000 00000050822011071130 5 4
```

Volumes, their cartridge serial number, the total number of permanent read/write errors, and the lot grouping number.

It would be even more useful if we knew how to understand the volume serial numbers

- *Do they go up over time?*
- *Are tapes in a box sequential?*
- *Are tapes in an order sequential?*

Plotting Errors by Lot

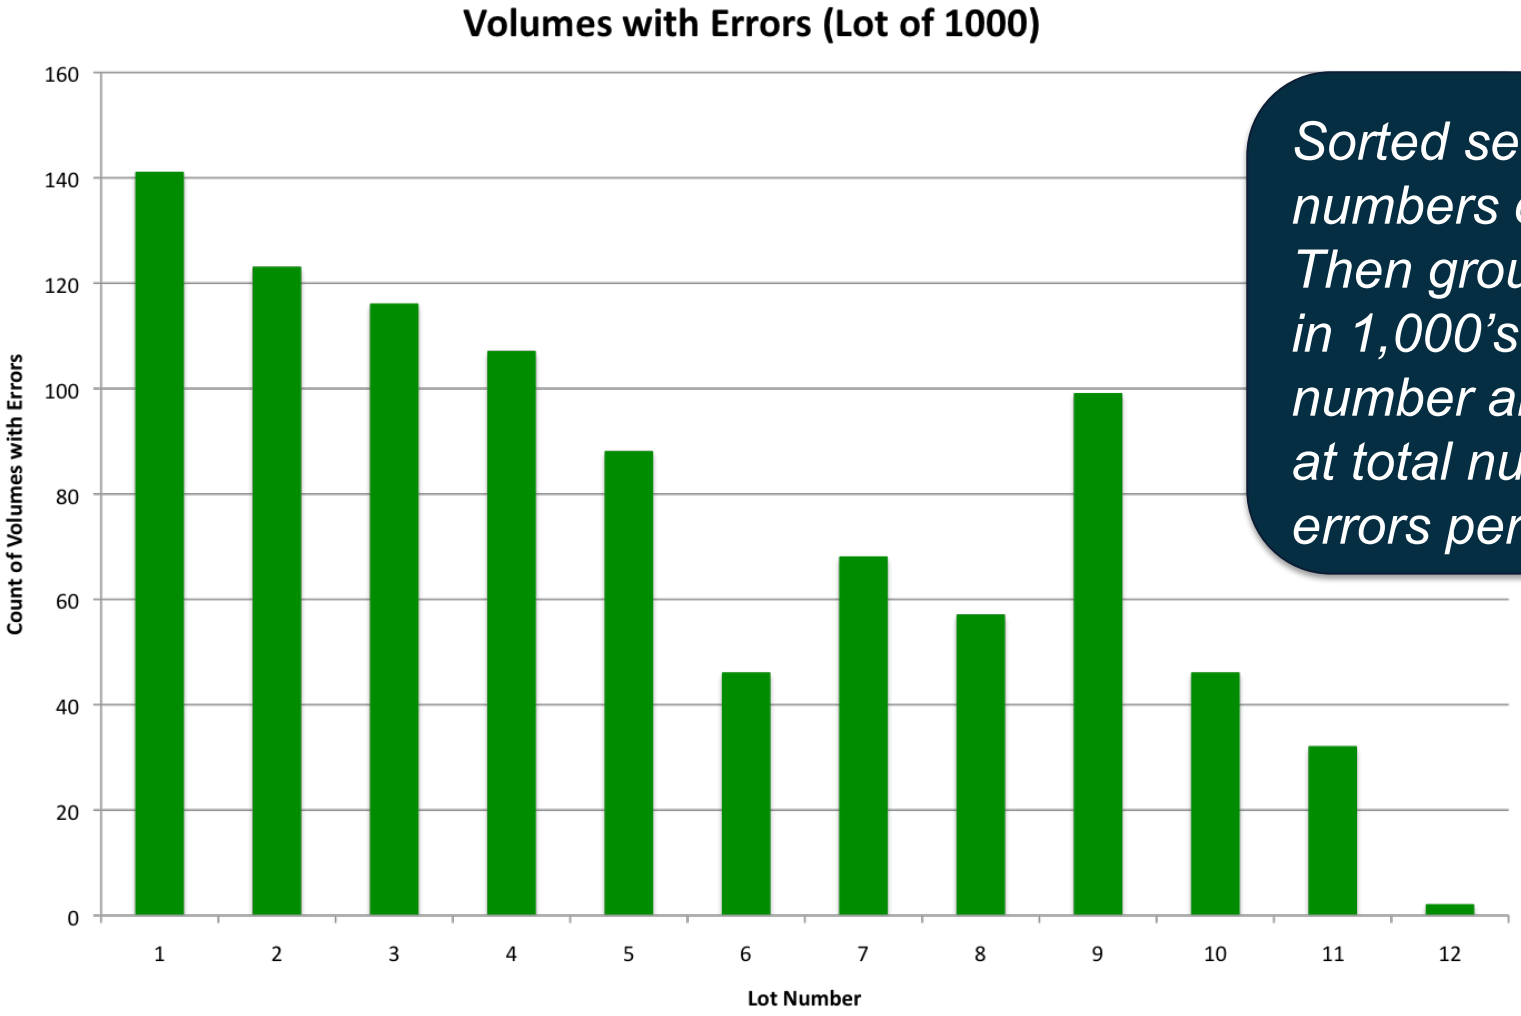

Sorted serial numbers of all tapes. Then grouped them in 1,000's by serial number and looked at total number of errors per volume

Plotting Errors by Lot

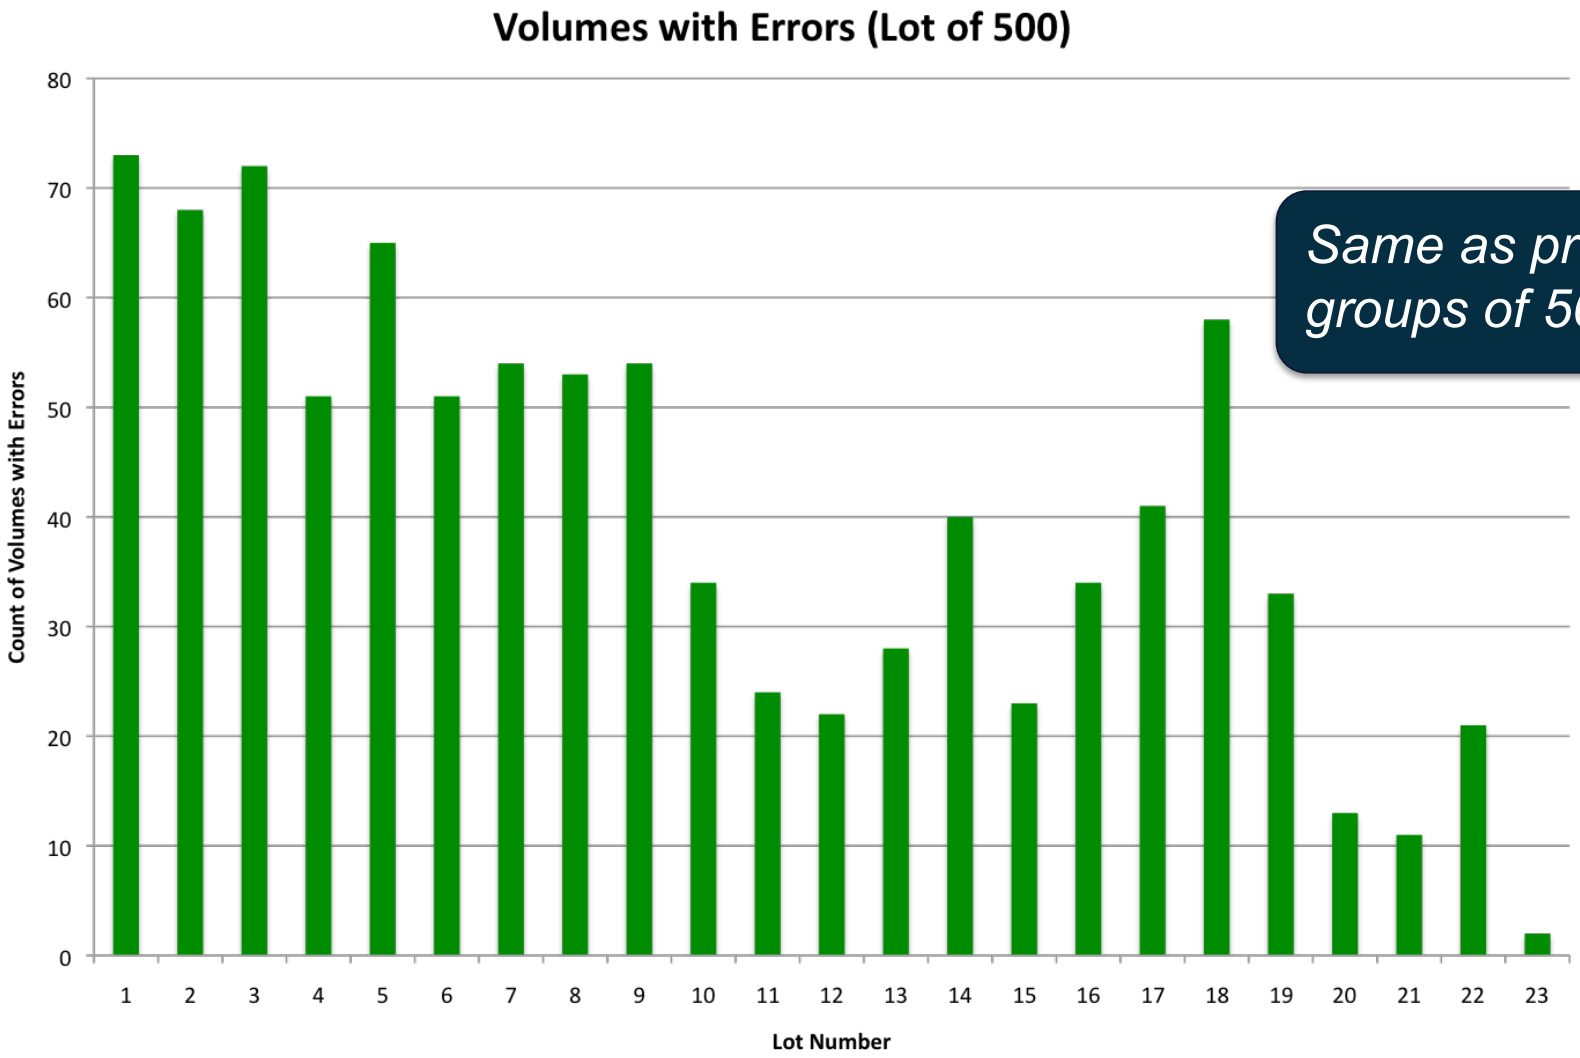

Same as previous but groups of 500 tapes

Plotting Errors by Lot

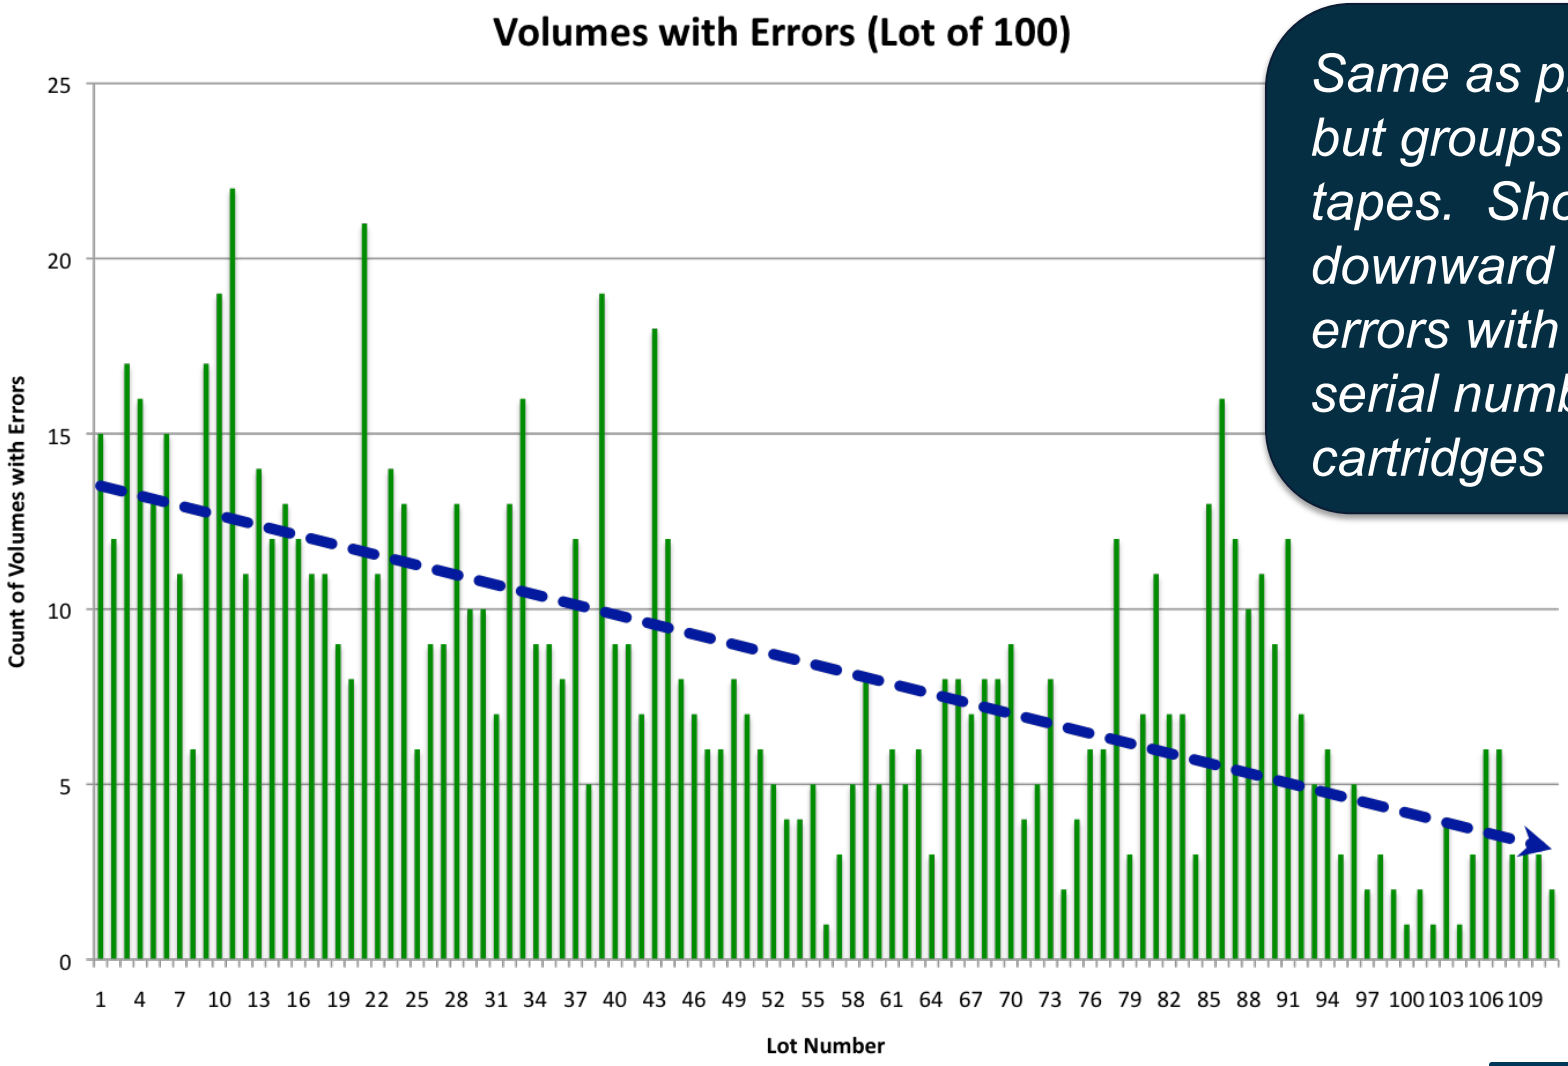

Same as previous but groups of 100 tapes. Showing downward trend in errors with higher serial numbered cartridges

Plotting Errors by Lot

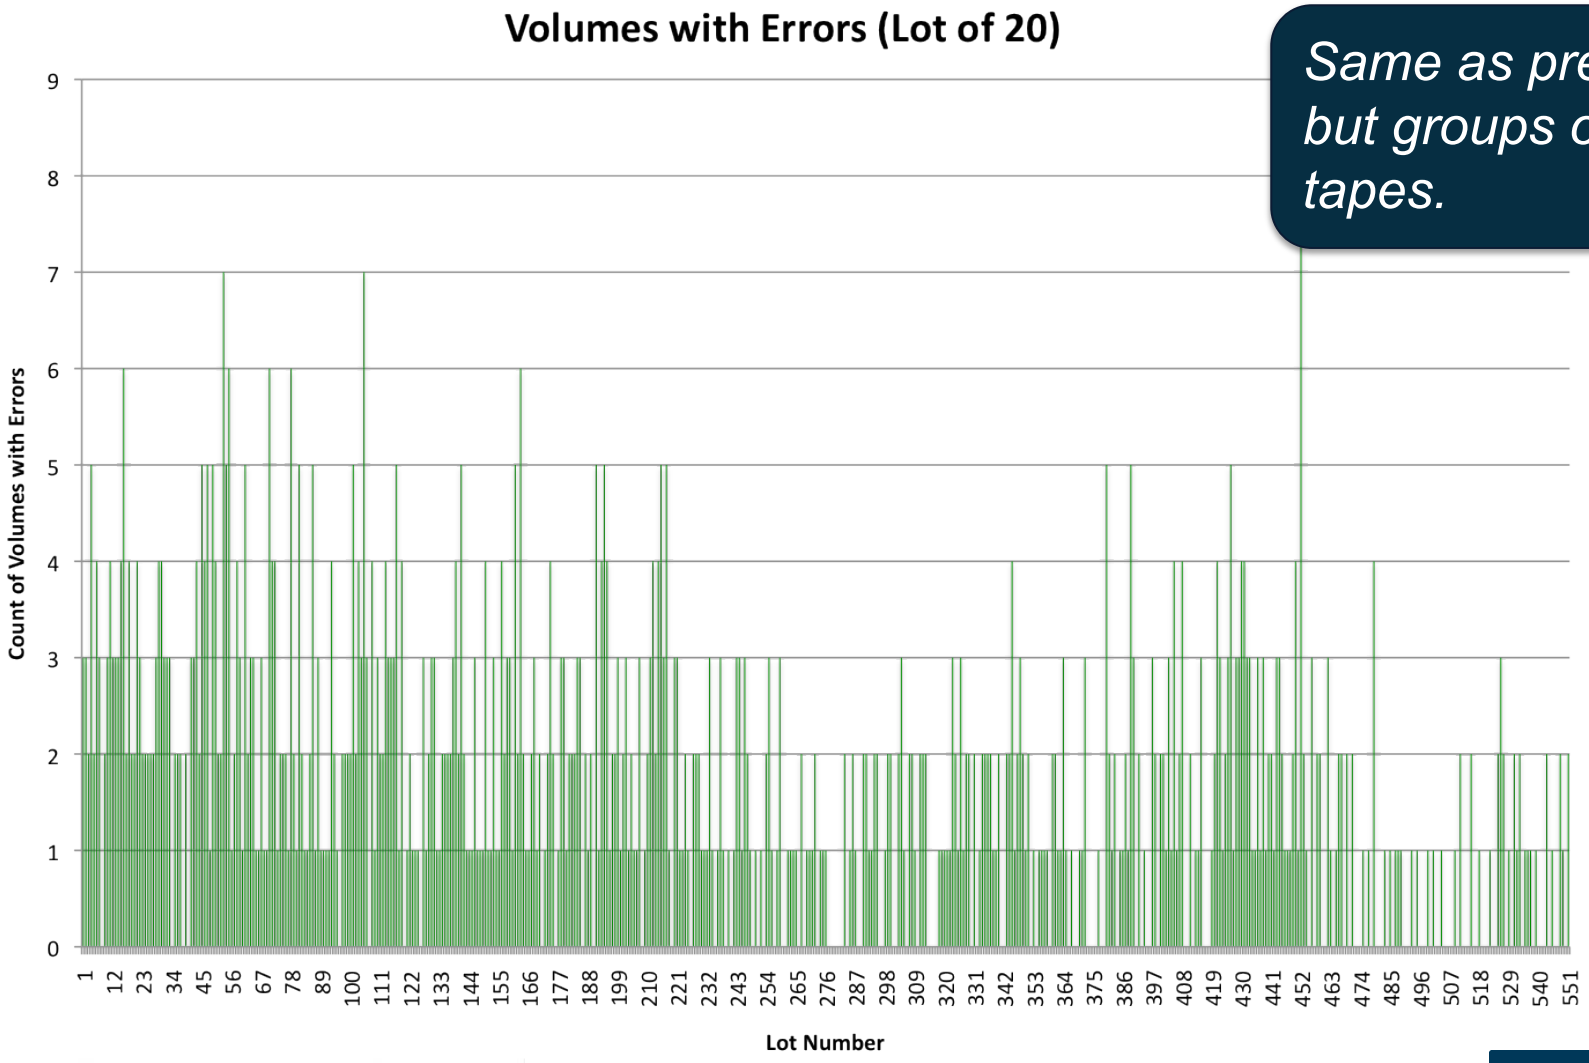

Same as previous but groups of 20 tapes.

Improving our operations

- **We don't know a file can't be read until someone tries to read it.**
  - We have the ability to repack problem files (or volumes) to new tapes, but don't take action until a user identifies a file they need. Very reactive.
  - Even if we can't predict a problem, it would be helpful to quickly generate a list of volumes and files per volume that had read errors.
  - A list of tape with drive combinations that resulted in errors would be helpful when trying to read the data.
- **In reusing media, we currently have no way of identifying the spectrum of good to bad media.**
  - If you could quickly produce a sorted list of volumes by different criteria (# read errors, # soft errors, # of meters, approximate age) that would be helpful in identifying which not to reuse.
- **We don't have an easy way of identifying tape lot problems that may exist.**
  - Would be useful to produce statistics on sequences of cartridge SNs.
- **We don't have an easy way to identify drives that write volumes that can't be read by other drives.**
  - Some tapes could only be read/written in specific drives. Could produce a matrix showing which volume or drive is statistically problematic.

Drive to Tape Stats

- We have previously shown that we do know which tapes worked or failed in which drives
- We haven't produced a matrix to show good/bad combinations of tapes and drives
- But we provided graphs of tape failures with drive mount history that aids in determining this

Summary

- **It is useful to analyze available MIR data for trends to aid in storage planning and operation of the system.**
- **An ideal solution would be to collect this information automatically (configurable per drive would be best)**
  - Enable users/admins to make comments for certain dates (tag significant events e.g. microcode updates, known operational issues)
  - Enable users/admins to not collect stats if using a drive/tape that is known to be bad (e.g. repack, data recovery)
- **To date, the following MIR fields prove most valuable to us:**
  - FSC code information (number, type of FSC code – read/write, date, drive it occurred on, microcode level)
  - Information about cartridges (age by cartridge SN, lot/grouping by cartridge SN, number of errors, drives cart could be read in/was written on, number of "bad" blocks, location of "bad" data)
  - Degree of error correction by cartridge
  - Statistically speaking whether the drive or the cartridge is likely at issue
- **We didn't highlight it here, but can also:**
  - List of volumes with their read block errors mapped to file names
  - List of errors per drive per microcode level

Future Directions

- We need greater collaboration with Oracle to make further progress
- Syncing of MIR data with specific command to get current mount stats and MAS with FSC code info
- Refine the separation of read statistics from write statistics (important to what information you care about)
- Easier extraction of important information from MIR
- MAS locations of blocks with errors to pathname of file