Hard Drive Stressor Series

Alan Nagl

Director of Technical Services and Sales

HDSTOR

Table of Contents

Introduction to the Series

In this series of articles, the goal is to educate the reader on hard drive classes and their unique features. To provide enough comprehension of the design features incorporated into today's drive choices that the reader will be able to make informed choices for their system design, ultimately, ensuring a more reliable system.

Throughout, the series will be examining the use of hard drives in the real world and the known stressors that occur. Some of the stress mechanisms are universal, such as heat, while others are unique to the specific use cases.

For each stressor, we will go into some significant detail and examine *why* these occur, as well as to outline their effect on the hard drives.

Once the stressor has been outlined, we will discuss the state of the art for hard drives and what features have been incorporated to combat these particular stressors, paying specific attention to which drives are best suited to deal with each type of stress.

At the conclusion of the series, we will tie what we have learned together and hopefully, give the reader confidence in their choice of drive models.

Stress: Heat

Heat and hard drives is an interesting subject, though certainly not new. Drives are actually quite robust against heat, and once you understand the effects of heat on the drive, their capabilities are really pretty impressive.

The reason drives are sensitive to heat is due to the fundamentals of their operation: magnetics.

Drives use magnetics to store digital data, and magnetic materials change properties over temperature.

The recording heads, the magnetic media layer, the voice coil actuator for head positioning, and the spindle motor, all rely on magnetic materials. Each of these materials are a bit different and each have their own unique dimensions and tolerances.

Stress Detail: Heat

Some key examples:

- When the media is cold, the magnetic recording layer's resistance to polarity change (called coercivity) increases, making it more difficult to influence polarity change, which would be required to "write" to the disc. This is also true in the opposite. When the media is hot, it's coercivity becomes "soft" and it is easily influenced to change magnetic polarity. If this were taken to the extreme, the magnetic layer could conceivably become so low in coercivity, that the magnetic bits would simply lose their will to maintain their polarity. This would translate to data erasure and total failure.
- Permanent magnets will lose magnetic strength (measured in Gauss) when they become warm, and gain it when they get cooler. When the magnets used inside the Spindle Motor get warm, they will lose some of their magnetic Gauss. The drive could potentially begin to slow down it's Revolutions per Minute (RPM) if this were not compensated for.
- The same is true for the magnets used to create a permanent magnetic field for the head actuating voice coil motor. With loss of Gauss, the seek speed will slow. In the extreme, this could mean the servo positioning system would not be able to accurately calibrate head movement.

The Expected Effect of Heat

Drive makers know the effect of heat on their product, because they test for it. Some companies will publish temperature de-rating curves, that will outline the effect.

Drive designers typically make the following assumptions:

- Mean environmental ambient temperature: 25ºC
- Mean operating temperature for the population: 35ºC (case temp)
- A stable material property temperature range of about 70ºC

So, this is the basis for the bulk of design considerations used in the specifications for Operating Temperatures for a hard drive:

Mean Operating Temp: 35º

Plus/Minus 50% of Stable Temp Range: +/- 35º

Result:

Operating Temp Range = 0 – 70ºC

Heat in Surveillance

Heat is a factor for all hard drives, and indeed for all devices at some point.

So, why is this a concern for Surveillance specifically? Easy: heat build-up.

With “most” types of systems, their usage model allows periods of higher and lower workload, depending upon the interactions with a user.

For Surveillance, the camera is expected to catch all the action, and the drive is responsible for scribing multiple cameras’ views, *all the time.* With no significant opportunity to cool off, drives in Surveillance run warm, period.

This has a negative effect on the reliability of the population, and remember, it’s not good for the other components either.

The 24 hr Heat Profile

- Both temperature extremes and temperature swings are stressful for devices that rely on magnetics
- For Surveillance, we typically only need to worry about one temperature related stress, long term heat
- Notice that even in a very well cooled DVR, the temperature relief is minimal, and the heat buildup is relentless

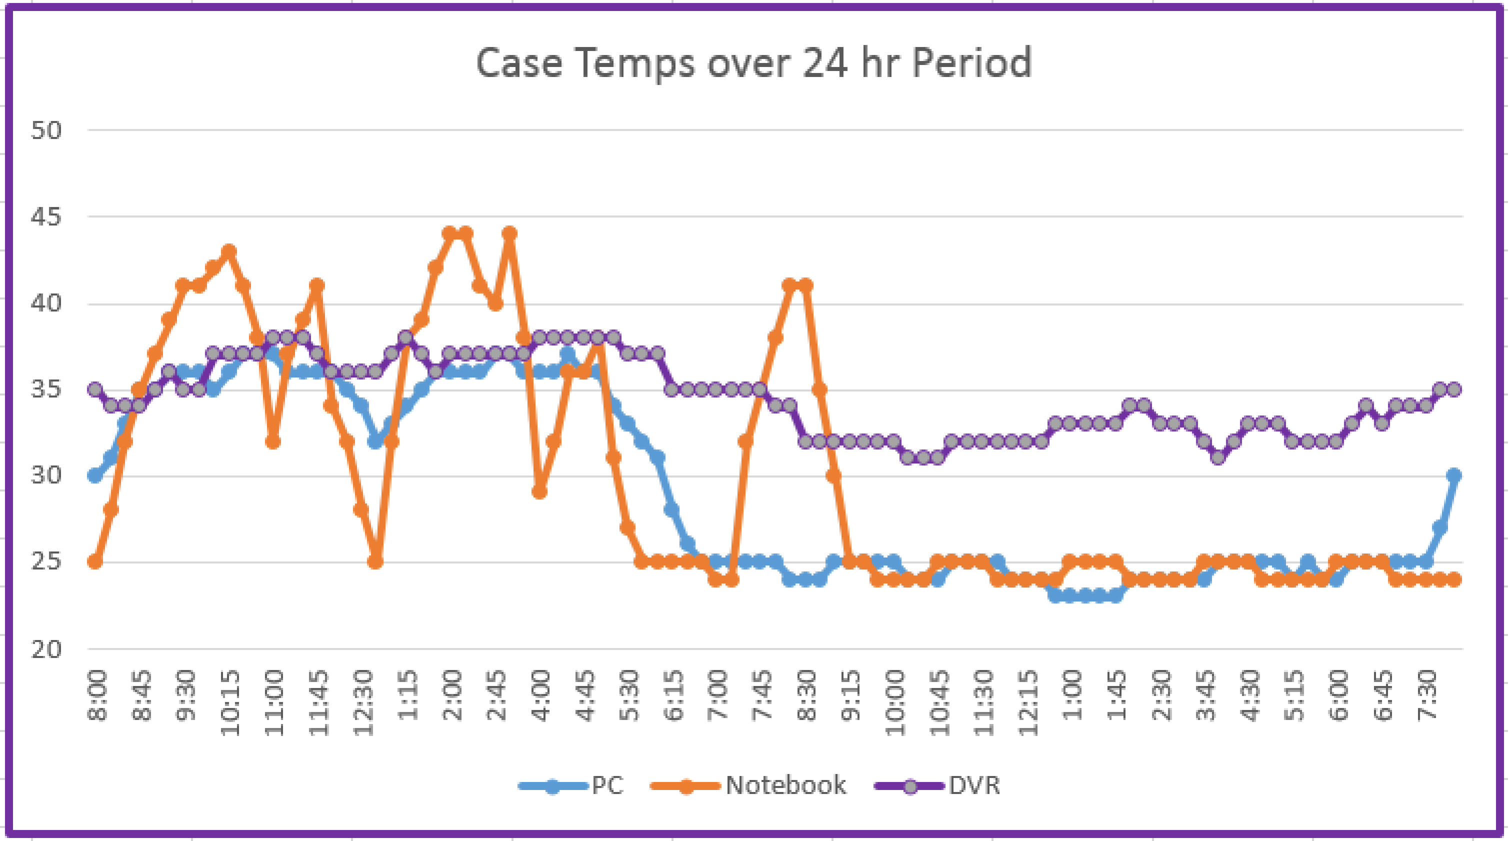

Heat Induced Reliability Derating

- This is routinely tested in drive Design Centers
- The key is to remember: Constant heat is the issue, not short term heat

As outlined in this graph, the Annualized Failure Rate (AFR) of drives in hot systems can easily be expected to double

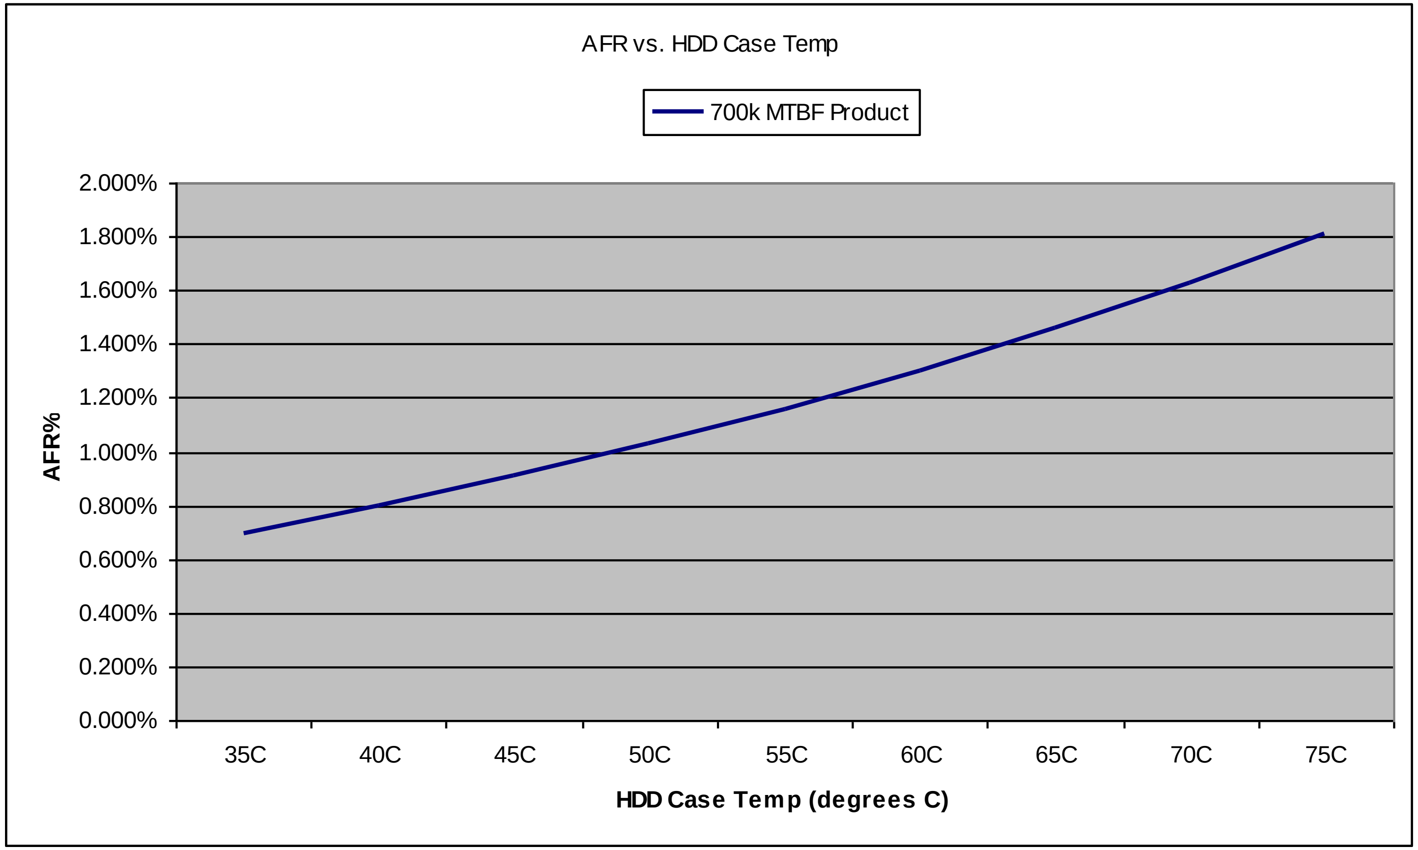

Heat: What to Do About It

- Heat is not unique to any one usage, but the potential for long term exposure to heat is a reality for certain use cases
- Keep this in mind when designing systems:
  - Ventilation schemes matter more for us
  - Chassis density is a pitfall for compounding the issue
  - If you plan to run 24/7, fans are not optional
  - Test your design for HDD temperature rise in still air, which is worst case
  - Ask your hard drive provider for help with a Reliability Projection
  - Use the test data and projection to establish expectations for failure rates

Stress: Rotational Vibration

Vibration of any type is a potential stressor to any device. The key to managing the stress levels induced by vibration is in knowing the device's susceptibility. This is typically defined by three parameters:

- Orientation of the energy (or Axis)
- Amplitude –the strength, or intensity of the energy
- Frequency –the speed at which the energy oscillates

Hard drives are not particularly susceptible to linear vibration, such as the X, Y and Z axis of traditional laboratory vibration tolerance testing.

This is, in part, due to the mass of the Basecasting, and the careful balancing of the Head Stack Assembly (HSA)

The other reason, is because ***linear vibration doesn't really exist!***

Say What?!

- It’s true. In nature, all movements are nonlinear.
- Even in laboratory conditions, with extreme measures taken, man fails to create perfectly linear vibration.
  - Put an accelerometer on the cross plane of any vibration test, or any vibrating surface, and you’ll find out of plane motion.
- This is caused by “asymmetry”, or “the inability to achieve true symmetry”.
- All materials will respond to energy inputs in their own unique way, based upon their stiffness, spring rate, dimensions, etc.
- Since no two pieces of matter are exactly identical, no two parts will ever react exactly the same. Fix these two parts together, and they can only move in one possible pattern, rotationally.

Sources of Rotational Vibration

- All vibrations, within the system or externally applied to the system, become rotational vibration
- Within the system, we have vibration sources:
  - Fans
  - DVD ROM/Writer
  - Adjacent Drives


Emitted vibration

- The hard drive itself is the largest contributor to disruptive energy
  - Sir Isaac Newton taught us that for every action, there is an opposite and equal reaction
  - Every time a drive executes a seek, it applies a torque to the base, which in turn, is applied to the chassis

Rotational Vibration Effects

Using a Fixed Shaft motor, and adding an additional screw, through the top cover, restricts the motion of the gyroscopic precession

The Disc Pack within the drive, is, in effect, a gyroscope. As such, it is susceptible to all of the traits of gyroscopes.

This illustration outline the basics of "Precession", the response a gyroscope will have to an external force

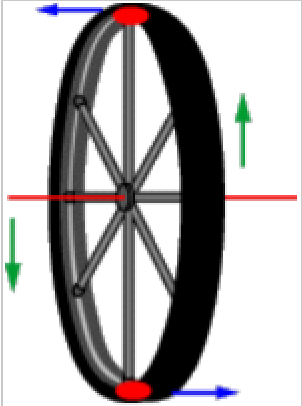

As forces are applied to the axle, the two points identified will attempt to move in the indicated directions.

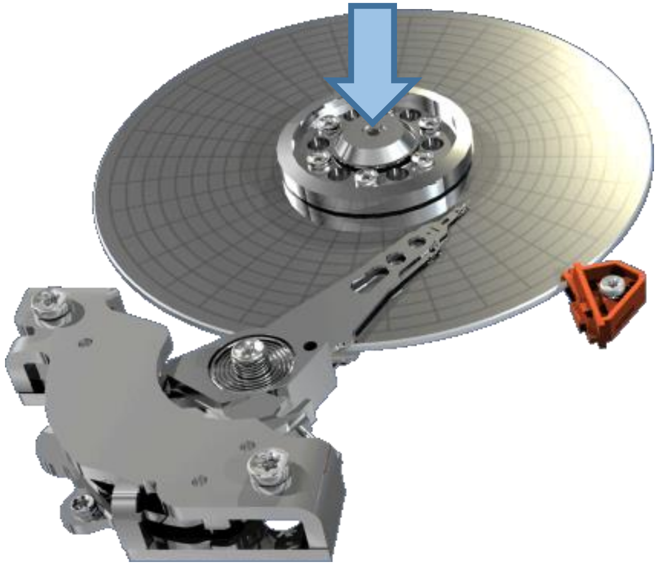

Rotational Vibrations will induce "Precession" of the Spindle Motor, if both ends are not fixed.

Stress: Power On Hours (POH)

- Running 24/7 is tough on all equipment
- Standard “Desktop” drives are designed around 40 hrs/week usage
  - 5400 RPM drives consume ~5.5 Watts
  - 7200 RPM drives consume ~8.0 Watts
- Powering a device off (or Stand By) allows it to cool
  - When devices are asked to run 24/7, they never cool
- Constant data transfer increases probability of error, and accumulates heat
  - Both of these are known Reliability degraders

Design for Power On Hours

- Multiple attributes can be enhanced to resist these stresses
  - Spindle Motors
  - Recording Subsystem
  - ASIC design
- Tuning can enhance drive life by anticipating workload type
- Unique screening during manufacture can “weed-out” weaker players
  - Factory CERT testing has been conditioned to emulate specific use cases
  - Drives incapable of tolerating these uses are screened out and get reprocessed for easier applications

Stress: Workload

- A measure of work, as serviced by the drive
- Typically defined by Volume of Data over Time (ie, TB/year)
- More data means:
  - Higher power consumed
  - More heat produced
  - Accelerated wear
  - Greater statistical probability of error

Design for Workload

- How drives design for higher workloads:
  - More robust Reader/Writer Element
    - Allows for more data transferred at higher data rates, extends drive life
  - More robust "Read Channel Preamp"
    - The ASIC that processes raw data into the controller is built with beefier silicon and tighter tolerances
  - Tougher Spindle bearings
    - Hydrodynamic bearing surfaces are enlarged, motor windings are encased in stiffer materials
- Be prepared for high workloads by choosing components that are rated and tested for it. There is no shortcut

Stress: RAID File System

- There are *many* implementations of RAID
- RAID 0 and 1 present no elevated stress, all other types, **do**
- RAID can be controlled through software, or hardware
  - The hard drive does not know (or care) which.
  - Both represent a dictated set of data storage instructions the drive complies with
- The primary goals for RAID are:
  - Enhanced write performance through multiplication of spindles/heads
  - Enhanced data security through redundancy and parity

Stress: RAID File System

- Data security is a highly desirable trait in many system types
  - RAID, with parity, can provide this (RAID 5 is most common)
  - System designers simply need to be aware of the effects on the components within this type of system, and choose accordingly
- The greater the data security, the more “Parity” is required
- The more Parity employed, the longer the RAID rebuild time
- While RAID can prevent data loss, it comes at a price…
  - It’s much tougher on drives
  - Requires more robust, more expensive drives
  - Sacrifices format efficiency, raising costs

Example: Streaming Video

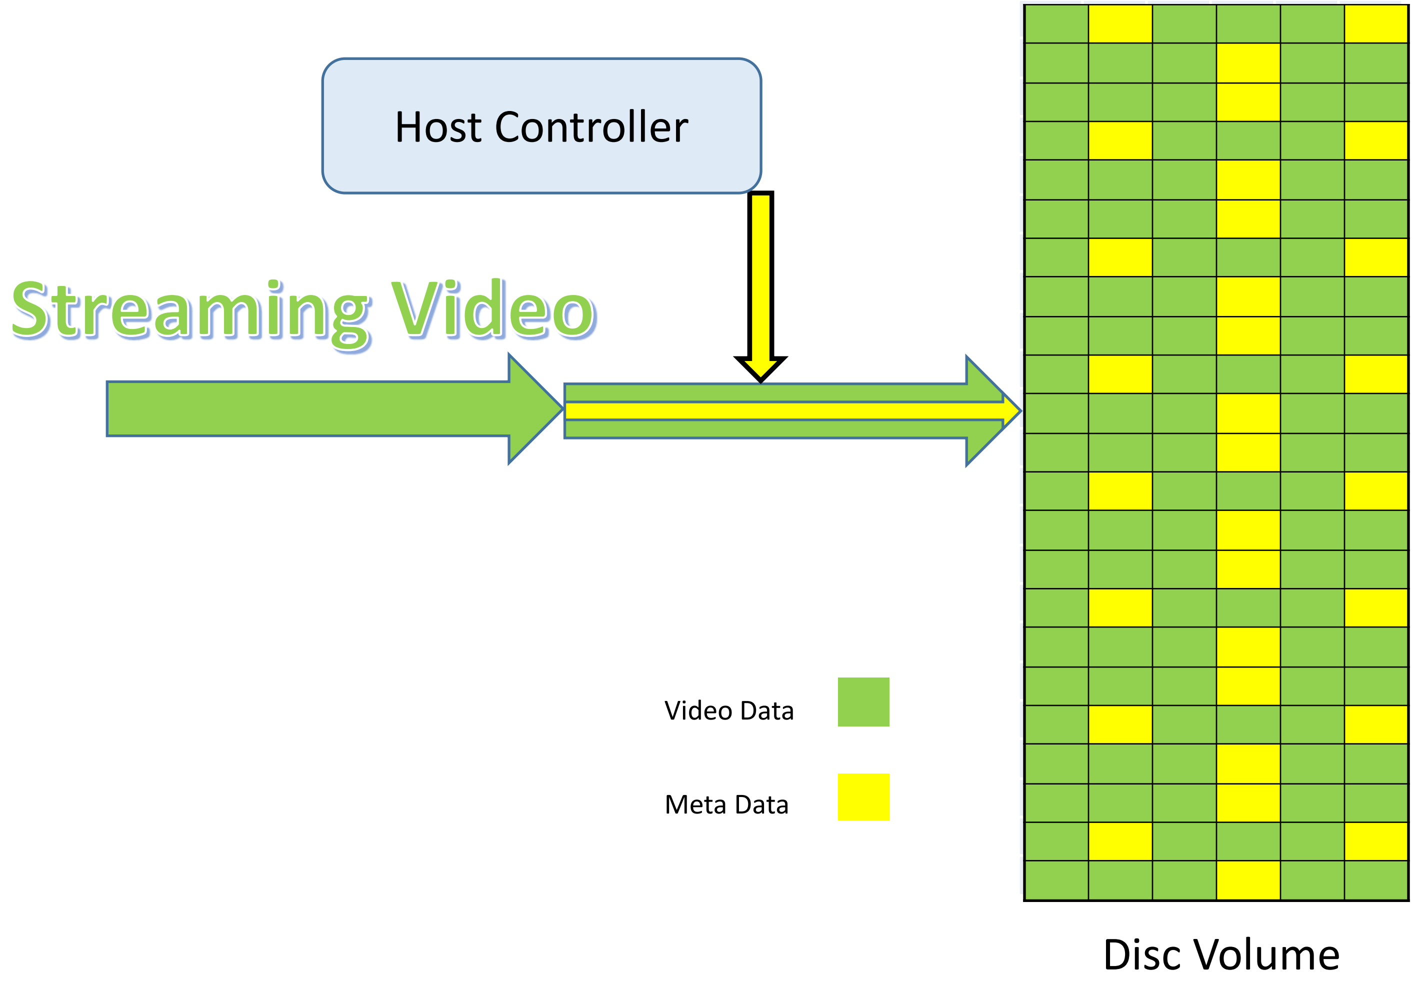

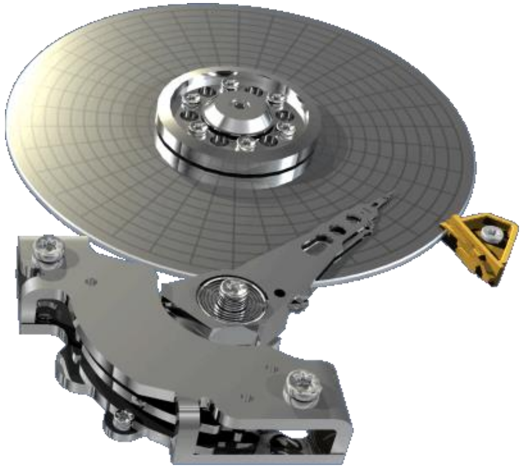

Streaming File System Means
Very Low Movement

Detail: Streaming File System

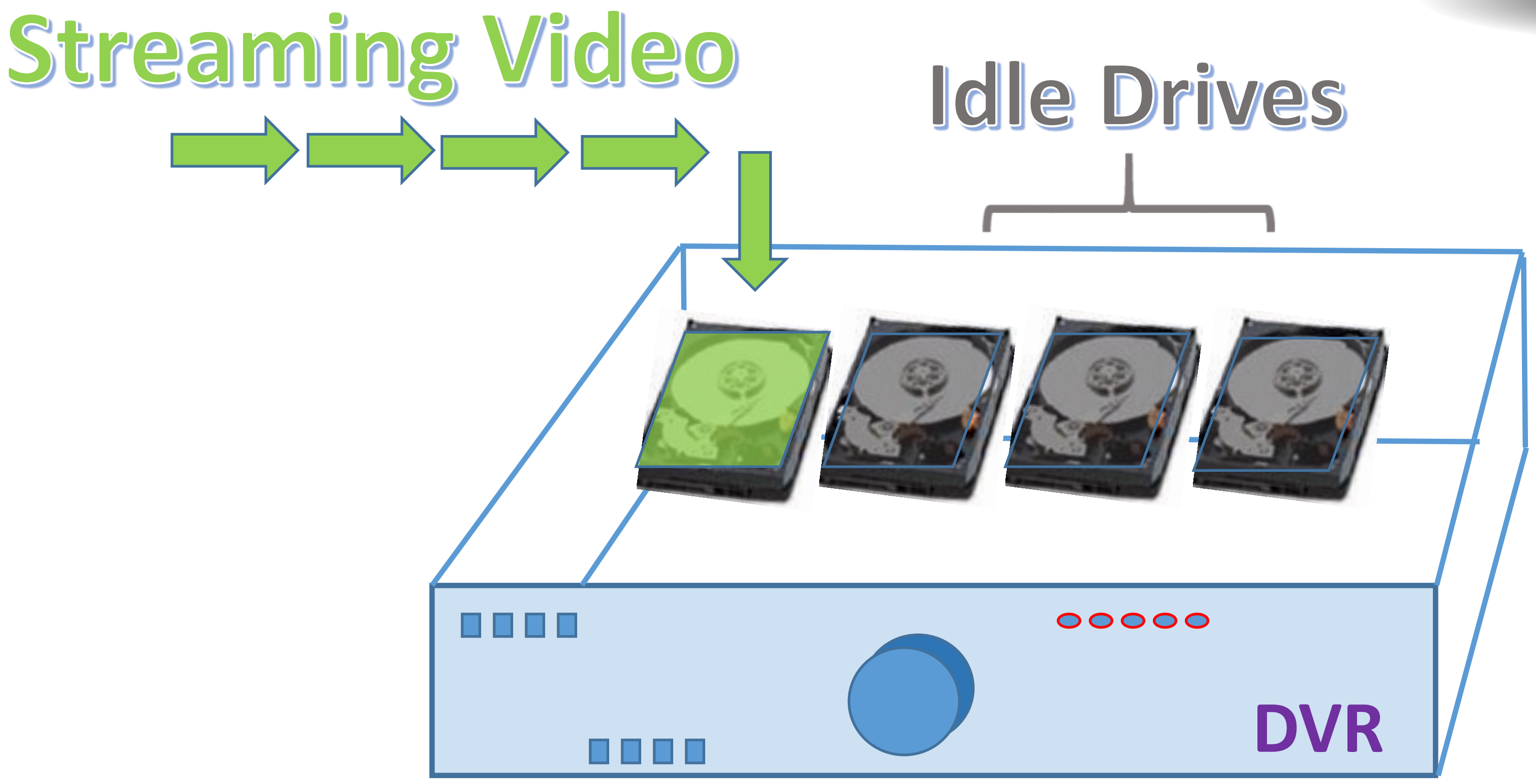

Example: RAID File System

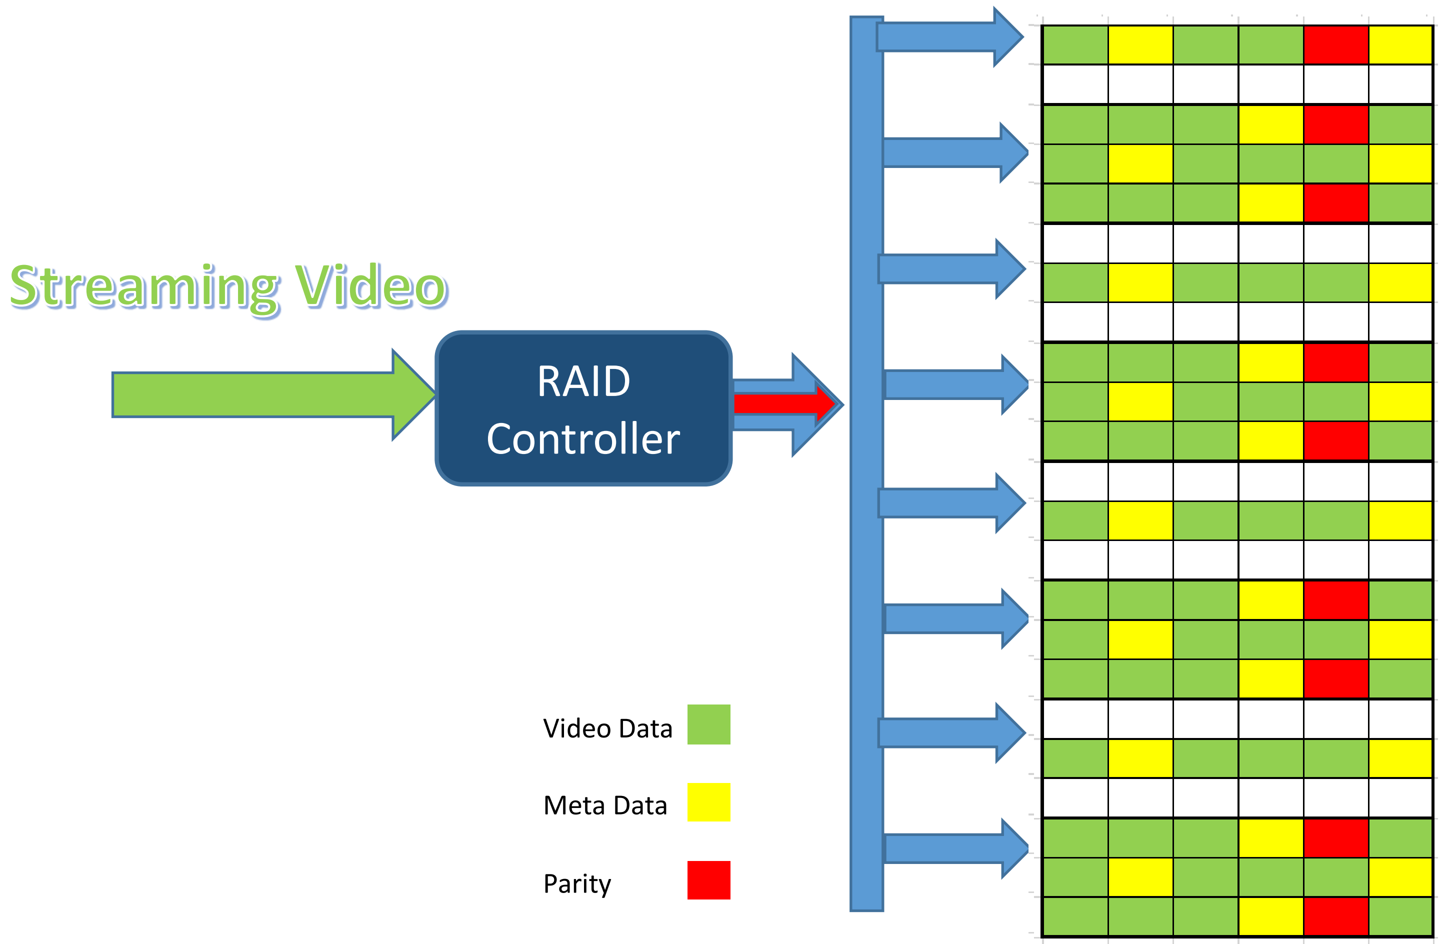

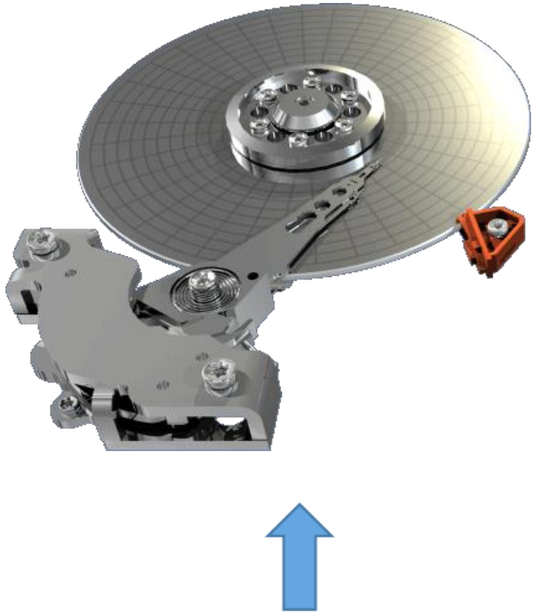

RAID means:

High Activity, High Energy

Stress: RAID File System

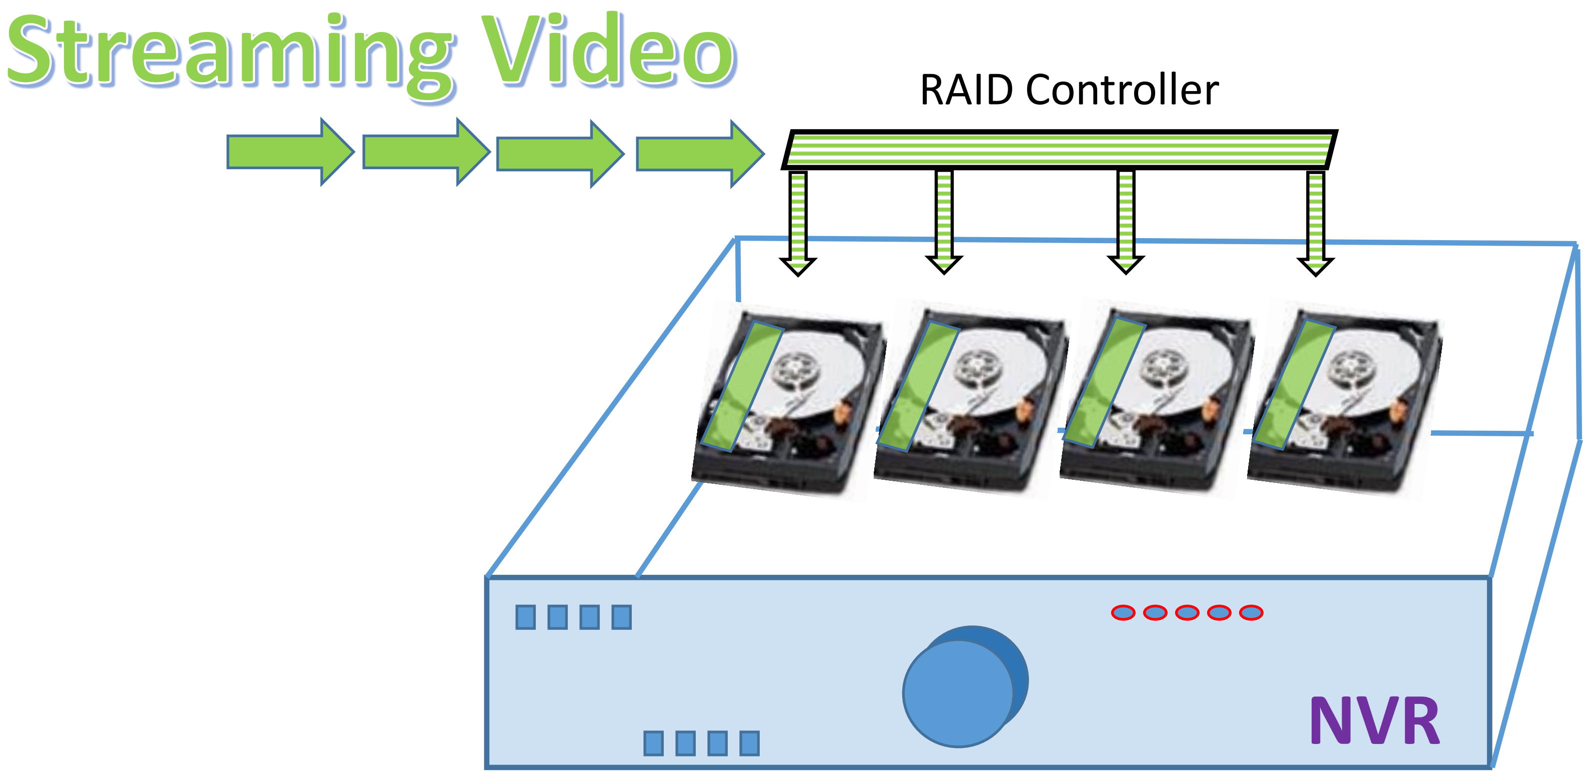

Design for RAID File Systems

- The higher activity means lots of vibrations within the chassis
  - Seek activity “self-excites” each drive
  - Adjacent drives “excite” each other
  - High speed fans contribute even more energy
- This energy is typically referred to as “RV”, or Rotational Vibration
- Enterprise class drives are designed to tolerate this
  - Mechanics are enhanced via:
    - ***Fixed Shaft Motor***
    - ***RV Feed-Forward Sensors***
      - Faster processor and more complex “Off-Track” prediction algorithm
      - RV Feed-Forward sensors measure repeatable rotational energy

Stress: Population Density

- In a Rack-Mounted Chassis, component density rises…
- More components within the system means:
  - Increased opportunity for chassis-borne vibrations:
    - Driving “Offtrack” instances that standard Servo can’t respond to
  - Shorter material spans:
    - Driving frequency responses higher
  - Increased power consumption:
    - More heat production
    - Increasing fan counts
    - Raising fan speeds

Design For Population Density

- The enhancements required for RAID, apply to Chassis Density
    - Faster Servo Response
    - RV Feed-Forward
- Enhanced robustness for key components
    - Fixed Shaft Spindle Motors
    - Heads and Media
    - ASICs
- Enterprise class drives typically include all these features

Stressors Conclusions

- In this series, we have explored multiple use cases and system design features that present challenges to hard drives.
- As these stresses are added up, the challenge becomes insurmountable for the “common” drive, and an Enterprise drive is required.
- The real challenge is in knowing which stressors your system exhibits, in which combinations and for what durations. Then, translating this understanding to an appropriate hard drive choice.
- While it is not necessary to use Enterprise class drives for every system to become reliable, it is virtually impossible to get reliability out of drives that were not designed for certain stress types.

Conclusion

- Hopefully you'll find yourself having greater confidence in your drive selection abilities following this course.
- Many systems have some combination of stressors, but not all of them, making selection potentially confusing.
- These are the cases that really require laboratory testing in order to determine the required component selection.
- HDSTOR can help you to design this test plan, and advise the best selection, based upon your test results.
- Thanks for taking time to participate with HDSTOR's educational program.

The Author

Alan Nagl

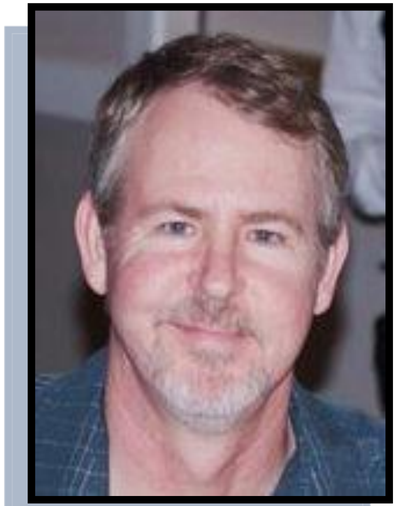

- With over 30 years of HDD industry experience, Alan brings a well rounded and practical approach to drive selection, integration, and system design
- Alan has extensive experience in:
  - R&D Engineering
  - Manufacturing
  - Reliability Engineering
  - System Design

MiniScribe	-8 years
Fujitsu	-3 years
Seagate	-17 years
Xyratex	-3 years