STORAGE DEVELOPER CONFERENCE
SDC 21
BY Developers FOR Developers
Virtual Conference
September 28-29, 2021

Introduction to HDD Field Accessible Reliability Metrics

Using disc drive metrics for predictive analytics

Paul Burnett, Seagate Technology
Matt Shumway, Seagate Technology

Data On Our Lives

Speaker Photo Will Be Placed Here

- **Intent: Provide a data structure that enables Fleet Management and Failure Prediction with visibility for open access to device metrics.**
  - FARM log provides access to metrics that captures health, drive usage and environmental parameters.
  - Ongoing FARM log collection and monitoring enables and supports failure prediction (recommend daily).
  - FARM is designed with ease of use in mind and all metrics are fully accessible by all customers.

Why FARM

Speaker Photo Will Be Placed Here

- Why use FARM when there are SMART monitoring metrics?
  - FARM provides a lot more parameters that SMART.
    - Host needs access to device workload, internal metrics to do better analysis and device management at scale.
  - These can be monitored to build history and do analytics and gain insight to device usage and its overall health.
  - Provides access to metrics that were only available using proprietary logs and tools.
  - Pulling the Log pull will not impact IOPS throughput or induce any latencies to the host.
  - Allows maintainability, testability and parse-ability.

FARM Basics

- SATA FARM Log Size is 96KB and the log structure is based on pages similar to ATA Device Statistics (32, 512-byte blocks per page).
  - FARM Log Structure has five pages implemented.
- SAS FARM Log Size is ~9KB and the log structure is based on T10 standard parameter codes.
  - Follows similar FARM Log Structure as SATA but still follows same basic principle for security concerns on content visibility.
- Includes metrics at Device and Head level.
- Statistical groups combined into single pages.
  - General Drive Information, Workload statistics, Error statistics, Environmental statistics and Reliability statistics.
  - There are ~170 different metrics available in the log for SATA and ~140 for SAS.
- The parameters in the log are have valid and supported bits for each parameter.

Speaker Photo Will Be Placed Here

FARM Basics - General

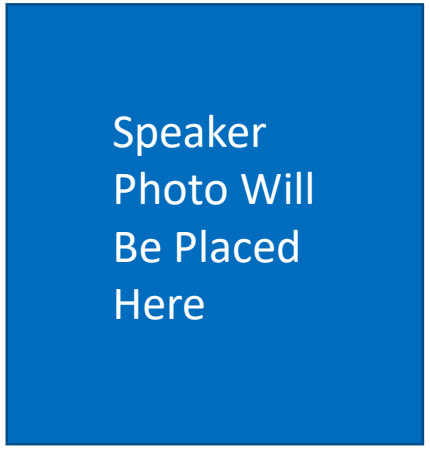

- Enhancements to FARM Metrics is reviewed on an ongoing basis to improve drive health detection and usage for failure prediction.
- FARM uses a vendor unique SATA log that can be accessed by using Read Log Extended Command (Address: 0xA6).
- FARM uses a vendor unique SAS log that can be accessed by using Log Sense Command (Log Page 3Dh, Sub Page 03).

FARM Key Metrics

What are the key metrics within the FARM log for monitoring (47).

- Link to full list of metrics

Speaker Photo Will Be Placed Here

Page 1(General)	Page 3 (Error Statistics)	Page 5 (Reliability Statistics) - by head / zone /radii
Power on Hours	No of Unrecoverable Read Errors	Host Selftest Disc Slip (uinches)
Power cycle Count	No of Unrecoverable Write Errors	Host Selftest BER Zone 0
Head Load Events	No of Reallocated Sectors	Number of Resident G-List Entries
Hardware Reset Count	No of Reallocated Candidate Sectors	Number of Pending Defect Entries
Page 4 (Environmental Statistics)	No of Read Recovery Attempts	BackGnd Scan Write Refresh Count
Curr Temp in Celsius	Uncorrectable Errors (Attr 187) *	Magneto Resistive Head Resistance
Highest Temp in Celsius	CTO Count Total*	Magneto Resistive Head Resistance 2nd Reader
Lowest Temp in Celsius	CTO Count Total >5s*	BackGnd Scan Ought to scans count
Page 2 (Workload)	CTO Count Total >7.5s*	BackGnd Scan Need to scans count
Total Num of Read Commands	Total Flash LED (Assert) Events	BackGnd ScanWrite Fault Scans
Total Num of Write Commands	**Page 5 (Reliability Statistics)**	Write Power On Hours in (sec)
Total Num of Random Read Commands	Raw Error Rate*	Current Head Selftest @ of codewords at iteration level, avg
Total Num of Random Write Commands	Raw Error Rate Norm*	Current Head Selftest amplitude, avg
Total Num of Other Commands	Raw Error Rate Worst*	Current Head Selftest asymmetry, avg
Logical Sectors Written	Seek Error Rate*	Current Head Selftest trimmed mean bits in error.
Logical Sectors Read	Seek Error Rate Norm*	Current Head Selftest iterations to converge
* = SATA only	Seek Error Rate Worst*	Applied fly height clearance delta (0.001 Angstrom)

Risk Prediction using FARM

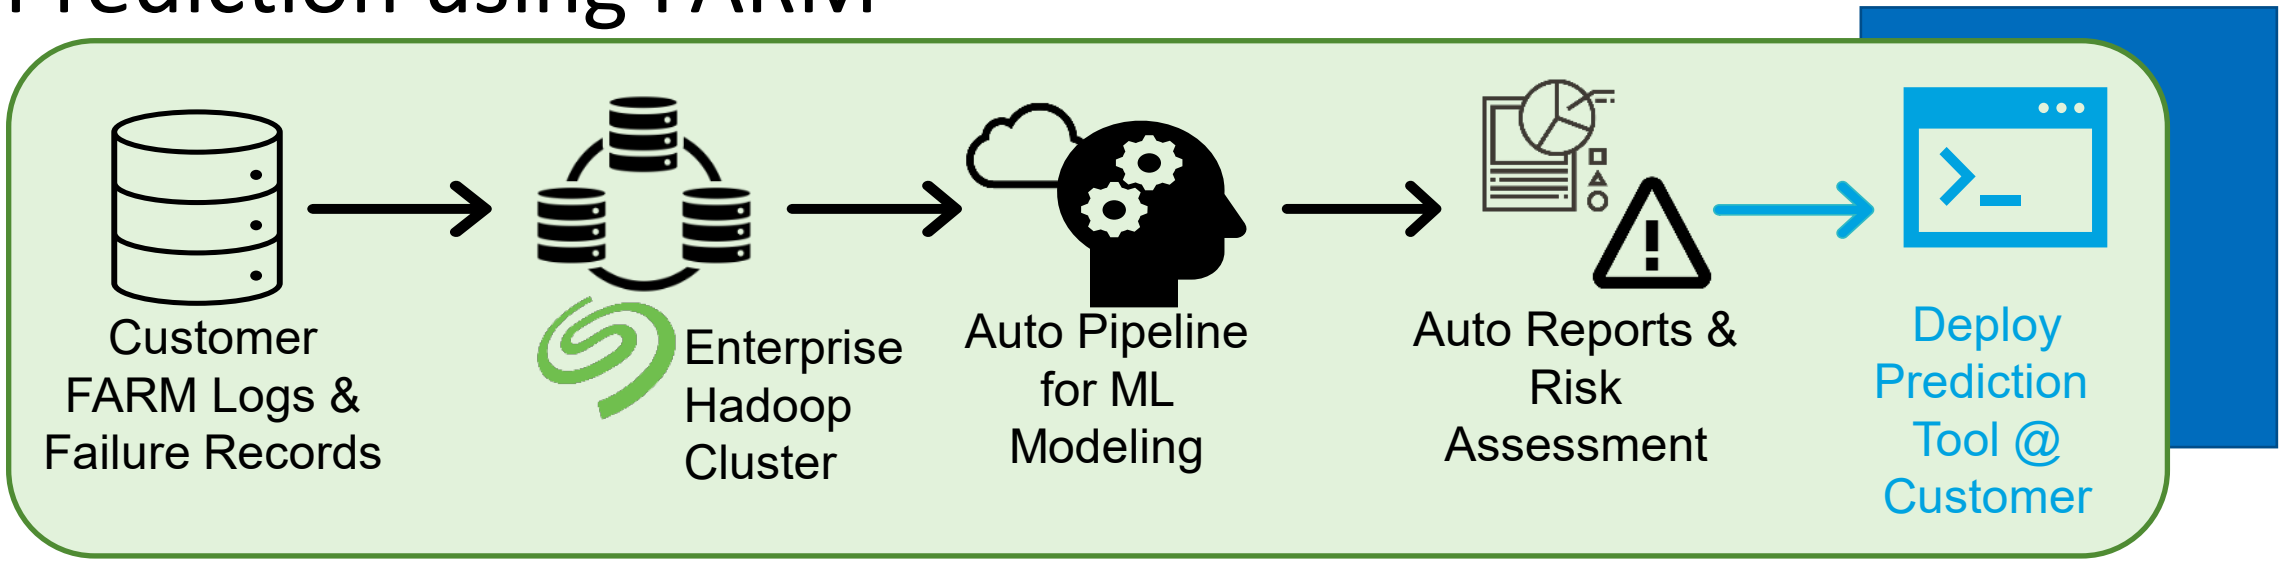

- The big question: Can we predict potential upcoming drive failures in the customer fleet?
- FARM benefits: considerably more data available than SMART attributes.
- Engagements with several Cloud customers, some with regular risk prediction reporting and with deploying the prediction tool at their sites.
- Our Method:
  - Store daily feed of FARM logs from customer fleet + past failure records
  - Train an ML model using past failure data to predict future failures based on current log values
    - Customized model specific to customer environment, product, etc.
  - Send software package to customers to monitor their systems

Example of ML Results

Speaker Photo Will Be Placed Here

- Daily Customer data fed into Seagate systems
  - FARM log data from ~100K's to 1M's
  - Failure instances reported to Seagate
- Train classification model to predict future drive failure within 7 days.
- How well does the model perform? Evaluating the test data typically shows good results:
  - ~0.16% False Positive Rate
  - 36% of the failures are captured with a 7-day prediction window
- A low false positive rate indicates a strong understanding of what a "good drive" is.
- Evaluating the model over a longer prediction window increases the failure capture rate at a cost false positive rate rate.

Sample Prediction Results from Test Data

	Predicted Healthy	Predicted to Fail
Healthy	612477	970
Failed	1234	694

How is the data helping customers?

Speaker Photo Will Be Placed Here

China IT News3.0 | News list | Search | Find | Categories | Post news | RSS subscribe | Mobile Ap

Home > News content

Use big data modeling to do real-time physical examination for hard disk. Seagate and Tencent cloud operate in this way

via:驱动之家 time:2020/10/12 12:52:11 readed:173

With the expansion of data center scale, the number of hard disks used to store data is growing. In the cloud computing platform, millions of hard disks are used as data storage carriers. The importance of efficient, reliable and stable hard disk storage system for cloud services is self-evident.

Because hard disk is the most important storage carrier in cloud computing system, the health management of hard disk is the fundamental of cloud service robustness and reliability. As a result, Seagate and the Tencent cloud server operation team began deploying FARM technology in Tencent's cloud data center in 2018, allowing real-time access to each hard disk of the cloud platform and data cente

Why can farm predict the health status of hard disk? How is this technology applied? Today we will systematically trace the source for you to see what kind of mysterious black

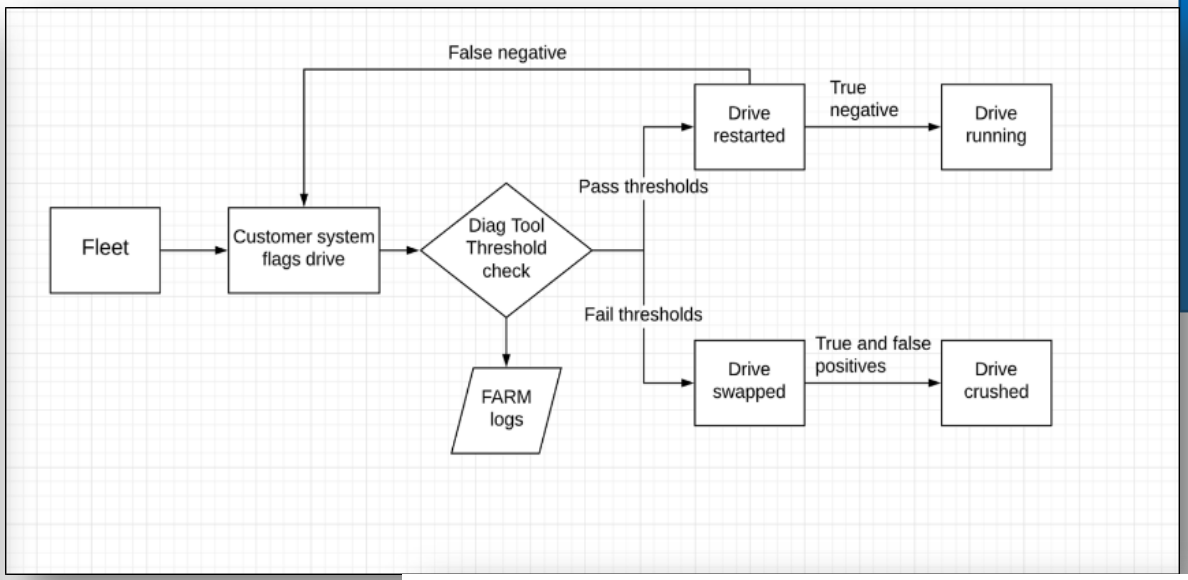

[Reference Link]

Google Cloud

Blog | Latest Stories | What's New | Product News | Solutions & Technologies | Topics

AI & MACHINE LEARNING

Google Cloud and Seagate: Transforming hard-disk drive maintenance with predictive ML

Nitin Aggarwal
Technical Program Manager

Data centers may be in the midst of a flash revolution, but managing hard disk drives (HDDs) is still paramount. According to IDC, stored data will increase 17.8% by 2024 with

Industry Standard?

Should FARM be an industry standard?

- Pro:
  - Highly desirable for common tools for customers.
- Con:
  - Limits supplier to provide most up to date parameters.
    - Processing through industry standards is time consuming and can change implementations which causes additional changes.
    - Not all suppliers use equivalent metrics depending on internal design.

Speaker Photo Will Be Placed Here

SSD Support

Is FARM in Solid State Drives?

- Currently FARM is not in SSD drives.
- Internal to Seagate, metrics have been defined.

Speaker Photo Will Be Placed Here

FARM Tools

FARM Tools and Accessibility

GitHub Support

Can Seagate Tools can be made available to customers for FARM log collection?

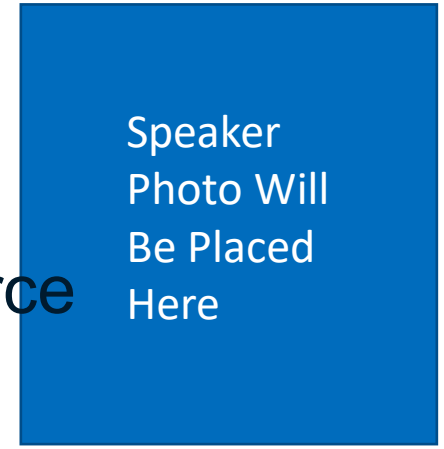

- Seagate Open SeaChest Tools set - Free and Open Source Software (FOSS).
  - openSeaChest Utilities
    - Supports cross hardware and software platform command line tools
    - Use **openSeaChest_Logs** application for Pulling FARM.
  - opensea-api – C based API and libraries
  - openSeaChest Log Parser
    - Supports parsing of FARM to JSON, TEXT & CSV format.
  - opensea-parser – C++ based API and library to parse FARM
- Seagate Prebuild binaries
  - **SeaDragon_LogUtils** will pull FARM and **SeaDragon_LogParser** will parse.
- Open Source packages like **sg3_utils** through bash scripting.

GitHub Support

Can Seagate Tools can be made available to customers for FARM log collection?

- Seagate github page:
  - https://github.com/seagate
- openSeaChest Utilities
  - https://github.com/Seagate/openSeaChest
- opensea-api – C based API & libraries
  - https://github.com/Seagate/opensea-api
- openSeaChest Log Parser
  - https://github.com/Seagate/openSeaChest_LogParser
- opensea-parser – C++ based API & libararies to parse FARM
  - https://github.com/Seagate/opensea-parser

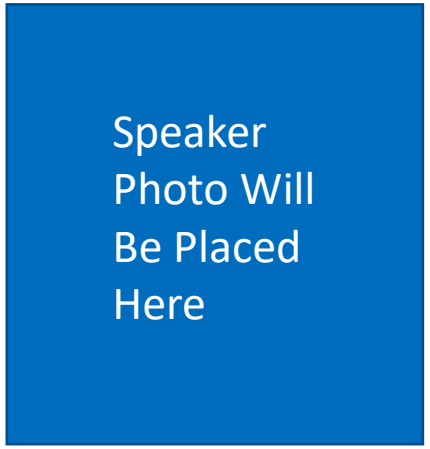

Please take a moment to rate this session.

Your feedback is important to us.

FARM

Log Structure for reference

FARM Metrics

The total list of FARM Metrics categorized

Speaker Photo Will Be Placed Here

Page 0 (Header)	Page 1(General)	Page 2 (Workload)	Page 3 (Error Statistics)	Page 4 (Environmental Statistics)	Page 5 (Reliability Statistics)	Page 5 (Reliability Statistics) - by head / zone /radii
Signature	Serial Num	Total Num of Read Commands	No of Unrecoverable Read Errors	Curr Temp in Celsius	**Time of Last IDD test**	**Host Selftest Disc Slip (uinches)**
Log Major Rev	WWN	Total Num of Write Commands	No of Unrecoverable Write Errors	Highest Temp in Celsius	**Sub Cmd of Last IDD Test**	**Host Selftest BER Zone 0**
Log Minor Rev	Interface	Total Num of Random Read Commands	No of Reallocated Sectors	Lowest Temp in Celsius	No of Glist Reclamations	**BackGnd Scan Write Refresh Count**
Pages Supported	Drive Model Number	Total Num of Random Write Commands	No of Read Recovery Attempts	Avg Short Term temp*	Servo Status	**Delta VGA Skip Write Detect**
Log Size in Bytes	48bit Capacity	Total Num of Other Commands	No of Mechanical Start Failures	Avg Long Term temp*	**No of Alts before IDD**	**Running Avg VGA Skip Write Detect**
Max Hds Supported	Physical Sector Size	Logical Sectors Written	No of Reallocated Candidate Sectors	Highest Avg Short Term Temp*	**No of Alts after IDD**	**Filtered VGA Skip Write Detect**
	Logical Sector Size	Logical Sectors Read	No of ASR Events*	Lowest Avg Short Term Temp*	**No of Resident Glist before IDD**	**Skip Write Detect Threshold Exceeded Count**
	Device Buffer Size	Dither events during current power cycle	No of CRC Errors*	Highest Avg Long Term Temp*	**No of Resident Glist after IDD**	**AC Feedforward Sine compensation**
	Number of Heads	Times dither was held off during random workloads	Spin Retry Cnt Raw*	Lowest Avg Long Term Temp*	Number of RAW Operations	**AC Feedforward Cosine compensation**
	Device Form Factor	Times dither was held off during sequential workloads	Spin Retry Cnt Norm*	Time in over Temp*	Raw Error Rate*	**Micro-Piezo Transducer Calibration compensation**
	Rotational Rate	Read commands from 0-3.125@ of LBA space @	Spin Retry Cnt Worst ever*	Time in under Temp*	Raw Error Rate Norm*	Magneto Resistive Head Resistance
	Firmware Revision	Read commands from 3.125-25@ of LBA space @	CTO Count Total*	Specified Max Op Temp	Raw Error Rate Worst*	Magneto Resistive Head Resistance 2nd Reader
	ATA Security State*	Read commands from 25-50@ of LBA space @	CTO Count Total >5s*	Specified Min Op Temp	Seek Error Rate*	**Servo Timing Mark Detects @**
	ATA Features Supported*	Read commands from 50-100@ of LBA space @	CTO Count Total >7.5s*	Over Limit Shock Events Count*	Seek Error Rate Norm*	**Servo Velocity Observer @**
	ATA Features Enabled*	Write commands from 0-3.125@ of LBA space @	Uncorrectable Errors (Attr 187) *	High Fly Write Count*	Seek Error Rate Worst*	**Servo Velocity Observer No Timing Mark Detects @**
	Power on Hours	Write commands from 3.125-25@ of LBA space @	IOEDC Errors*	Current Relative Humidity (0.1@ RH)	High Priority Unload Events* (Emergency retract)	**Current Head Selftest @ of codewords at iteration level, avg**
	Spindle Power on Hours*	Write commands from 25-50@ of LBA space @	**Total Flash LED (Assert) Events**	Humidity Mix Ratio	**Micro Actuator LockOut @**	**Current Head Selftest amplitude, avg**
	Head Flight Hours*	Write commands from 50-100@ of LBA space @	**Index of last entry in Flash LED Info array below**	**Current Motor Power scalar @**	**Number of disc slip recalibrations performed**	**Current Head Selftest asymmetry, avg**
	Head Load Events	Read Commands of transfer length <=16KB space @	**Timestamp (us) of last 8 Flash LED (assert) Events**	Current 12V input (mV) @	Helium Pressure Threshold Trip (1 – trip 0 – no trip)	Number of Resident G-List Entries
	Power cycle Count	Read Commands of transfer length (16KB – 512KB] @	**Power Cycle of the last 8 Flash LED (assert) Events**	Minimum 12V input (mV) @	**RV Absolute Mean @**	Number of Pending Defect Entries
	Hardware Reset Count	Read Commands of transfer length (512KB – 2MB] @	FRU code if smart trip from most recent SMART Frame +	Maximum 12V input (mV) @	**Max RV Absolute Mean @**	**BackGnd Scan Ought to scans count**
	SpinUp time (ms)	Read Commands of transfer length > 2MB @	**Info on the last 8 Read Recoveries, wrapping array**	Current 5V input in mV @	Idle Time @	**BackGnd Scan Need to scans count**
	NVC Status on Power on	Write Commands of transfer length <=16KB @	**Super Parity Recoveries**	Minimum 5V input (mV) @	**Number of Scrub List Entries Before IDD Scan**	**BackGnd ScanWrite Fault Scans**
	NVC Available time after Save User Data to Media	Write Commands of transfer length (16KB – 512KB] @	**Reallocation Reason**	Maximum 5V input (mV) @	**Number of Scrub List Entries After IDD Scan**	**Write Power On Hours in (sec)**
	Power On Timestamp (ms) for time restricted data-Min	Write Commands of transfer length (512KB – 2MB] @	Unrecoverable Read errors due to Error Recovery Control		Number of BackGnd Scans Performed	**BackGnd Scan Write Count Threshold**
	Power On Timestamp (ms) for time restricted data -Max	Write Commands of transfer length > 2MB @			Number of LBAs Corrected by ISP	**Current Head Selftest trimmed mean bits in error.**
	Time to ready of the last power cycle	Count of Queue Depth =1 at 30s intervals @	Legend		**Number of Valid Parity Sectors**	**Current Head Selftest iterations to converge**
	Time drive is held in staggered spin	Count of Queue Depth =2 at 30s intervals @	LUN Based Parameter		**Number of LBAs Corrected by Parity Sector**	**Applied fly height clearance delta (0.001 Angstrom)**
	Reman Depop Status	Count of Queue Depth 3-4 at 30s intervals @	By Head Based Parameter			**Cumulative Lifetime Unrecoverable Read Repeating**
		Count of Queue Depth 5-8 at 30s intervals @	by Head by zone or radii parameter			**Cumulative Lifetime Unrecoverable Read Unique**
		Count of Queue Depth 9-16 at 30s intervals @	**Bold - Extremely Seagate Specific**			
		Count of Queue Depth 17-32 at 30s intervals @	@ = cover a specified time range of data			
		Count of Queue Depth 33-64 at 30s intervals @	+ not supported by SATA			
		Count of Queue Depth >64 at 30s intervals @	* not supported by SAS			

Log Structure-SATA

- **Highlights**
  - Log, all pages, and values are self-describing
  - All log pages are the same length
  - Page 0 is reserved as a log descriptor
    - Includes information such as:
      - Unique signature
      - Defines pages supported
      - Number of 512-byte blocks per page
  - Statistical groups combined into single page
    - Examples include:
      - General Drive Information
      - Workload Statistics
      - Error Statistics
      - Environmental Statistics
      - Reliability Statistics

Speaker Photo Will Be Placed Here

Log Page 0 Directory Page
Log Page 1 General Drive Information
Log Page 2 Workload Statistics
Log Page 3 Error Statistics
Log Page 4 Environmental Statistics
Log Page 5 Reliability Statistics

Log Page Structure-SAS

- **Highlights**
  - Each log page has a header
    - Includes log page identification
  - Each parameter has a self-describing status byte
    - Examples include:
      - Parameter Validity Bit
      - Parameter Supported Bit
  - New parameters added to the end of the log page

Speaker Photo Will Be Placed

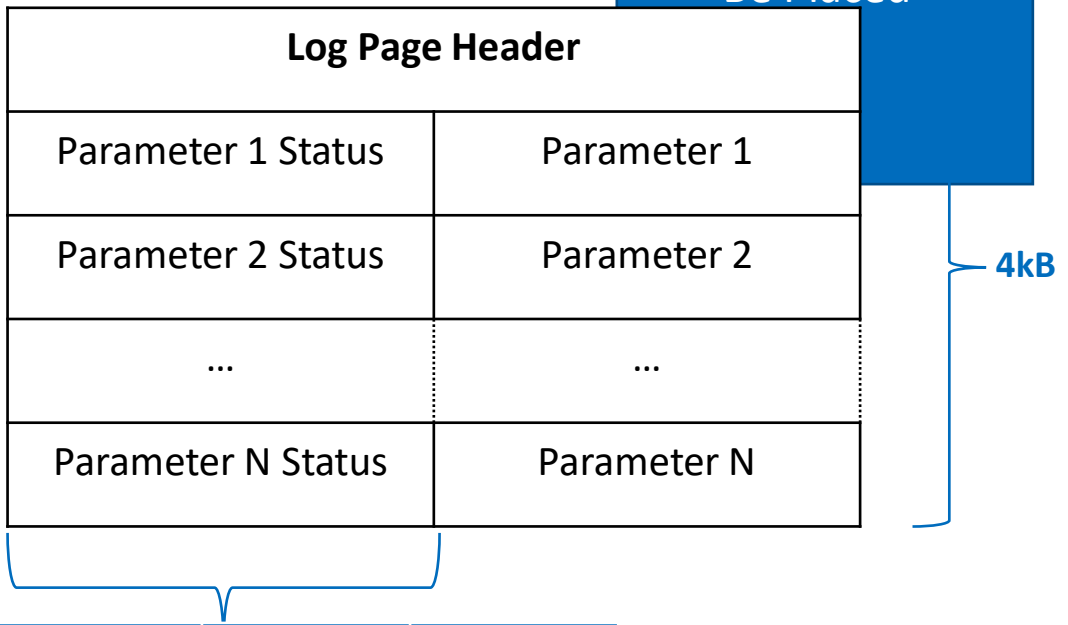

Bit	7	6	5	4	3	2	1	0
Example Status	Field is Supported	Field is Valid	Other Status	Other Status	Other Status	Other Status	Other Status	Other Status

Links to Internet for Seagate and Cloud for FARM

Speaker Photo Will Be Placed Here

- Google Cloud and Seagate: Transforming hard-disk drive maintenance with predictive ML
- Use big data modeling to do real-time physical examination for hard disk. Seagate and Tencent cloud operate in this way

