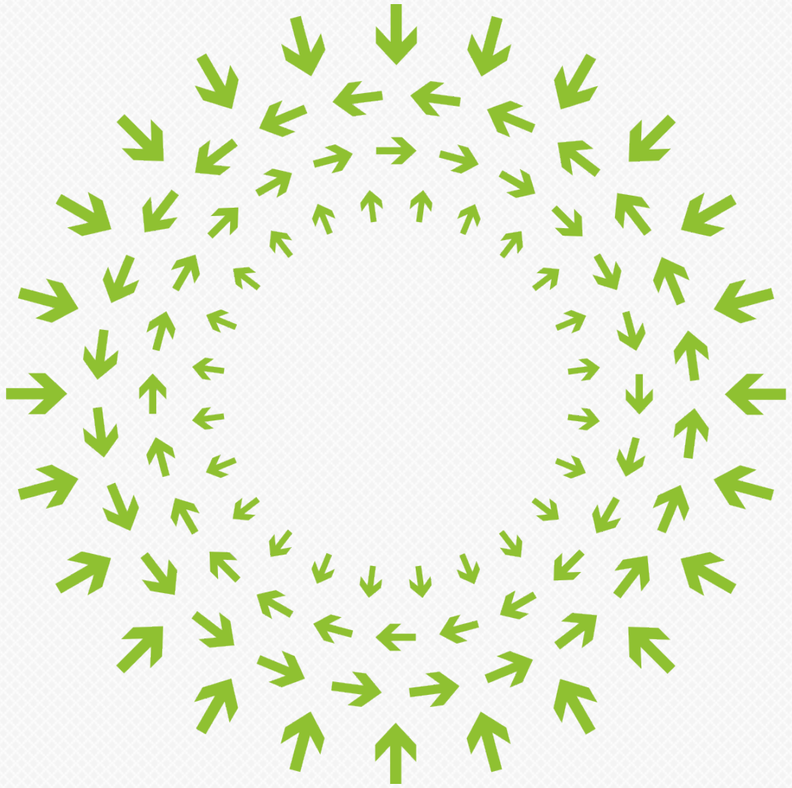

OPEN Compute Summit

January 28–29, 2014San Jose

OPEN Compute Summit

Archive & Cold Storage Device Trends

Jun Liu
Seagate Technology
Cloud Architect

Josh Tinker
Seagate Technology
Enterprise Sr Product Manager

Customer Driven Cold Storage Devices

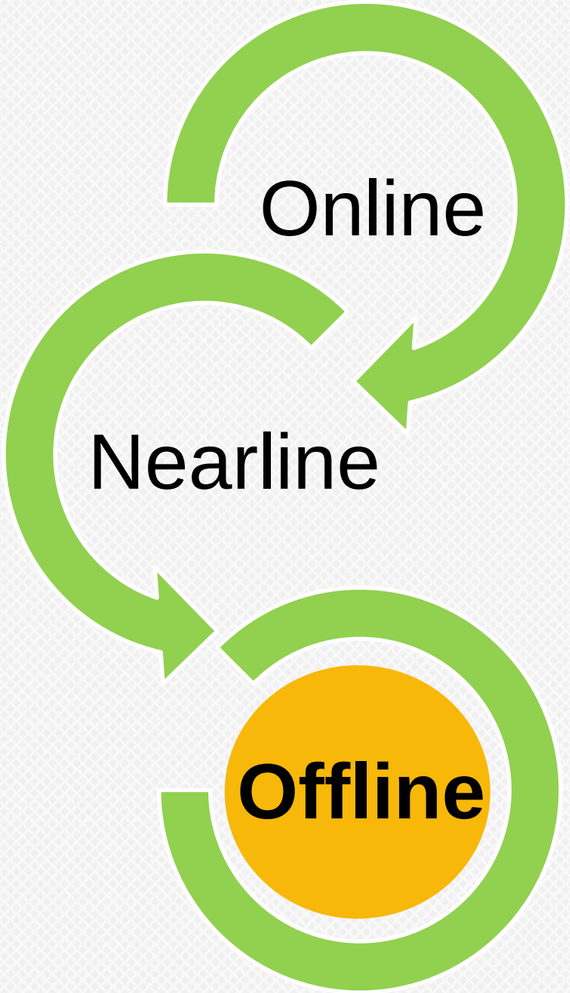

- Market Demands
  - Cloud service customers want to keep their content forever
- Regulatory requirements
  - Health care records, e-Discovery, Financial data, Government data
- Storage is inexpensive and data management is complex and takes too much time
- Enables new cloud services for CSPs to offer…

Key Requirements

Key Requirements

- Lowest $/GB
- Highest Capacity
- Most efficient Energy Consumption Profile
- Suitable in High Density Storage Environment

Key Requirements - Solutions

- Lowest $/GB
- Highest Capacity

- Most efficient Energy Consumption Profile → **Power Management**
- Suitable in High Density Storage Environment → **Drive Health Management**

How do we lower $/GB?

Increase Capacity

$$\frac{\$}{GB}$$

Best option

- Near term?

SMR: Shingled Magnetic Recording

- Today's most cost effective way to increase capacity

OPEN
Compute Summit

What is SMR?

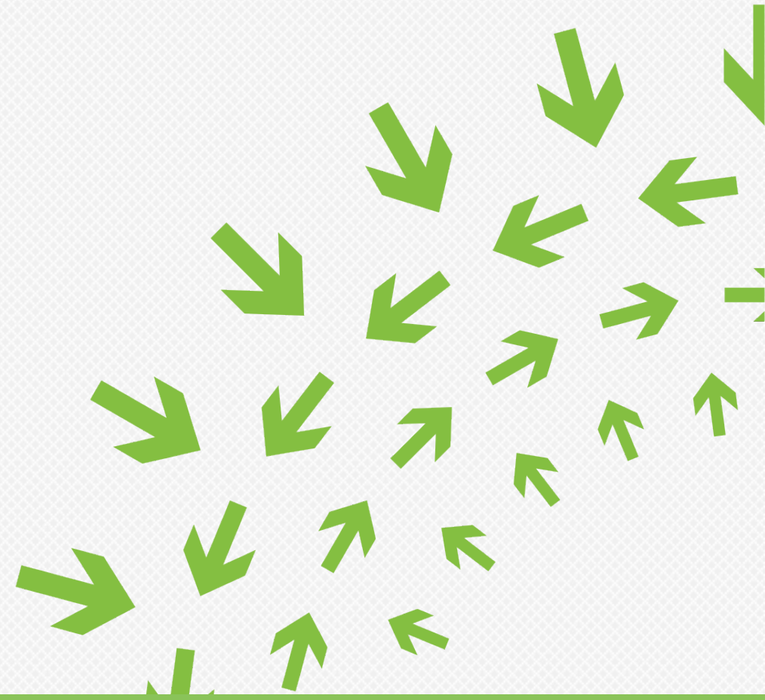

Why SMR?

Conventional, Non-SMR Writes

- Conventional writer width defines TPI
- Conventional reader is much narrower than writer

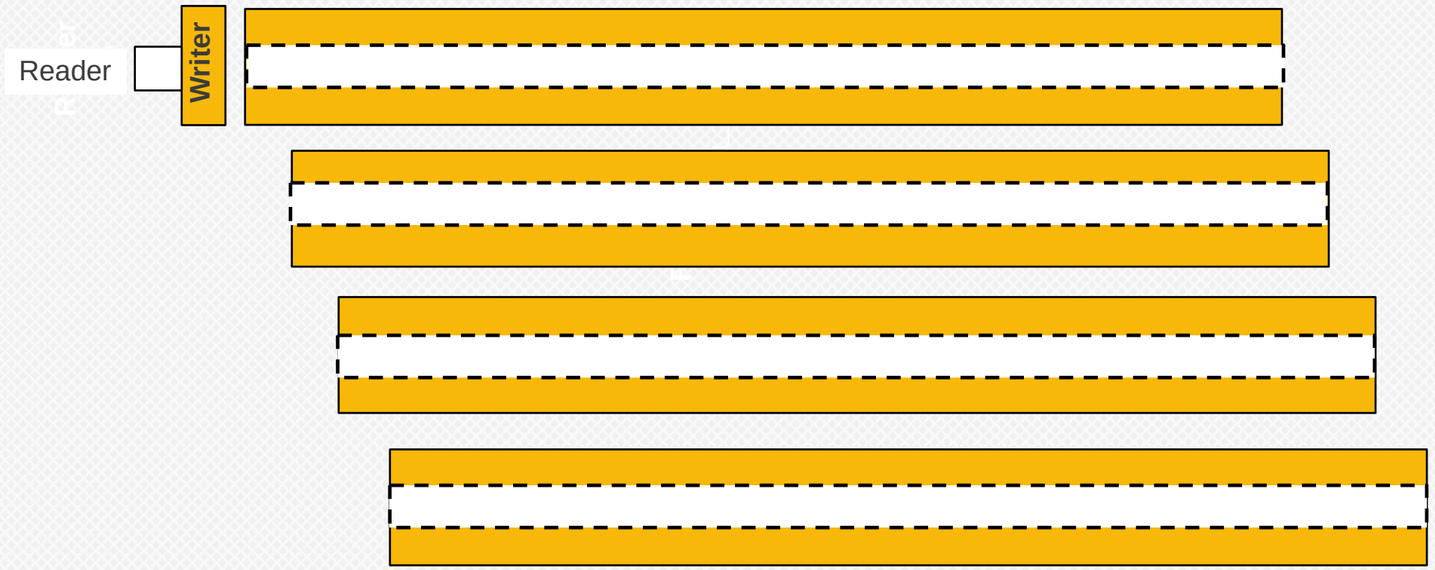

SMR squeezes tracks, enabling more TB

SMR Writes

- Data is written wide and the write to the next adjacent track trims the previous track
  - Data is written in progressing track order
  - The last track of a band is not trimmed
  - No random in-place writes
- Readable tracks are narrower than originally written track

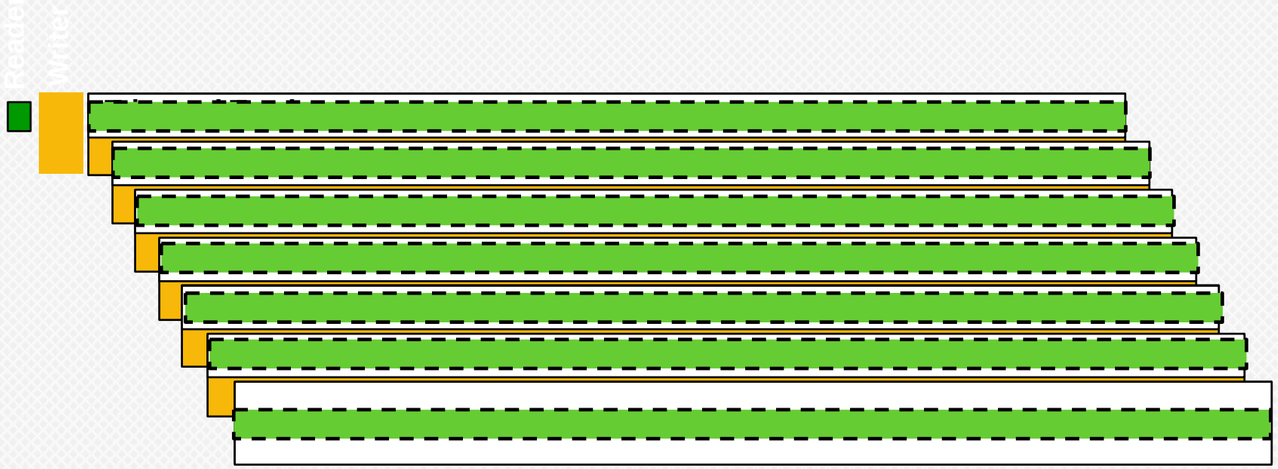

What about Random Writes?

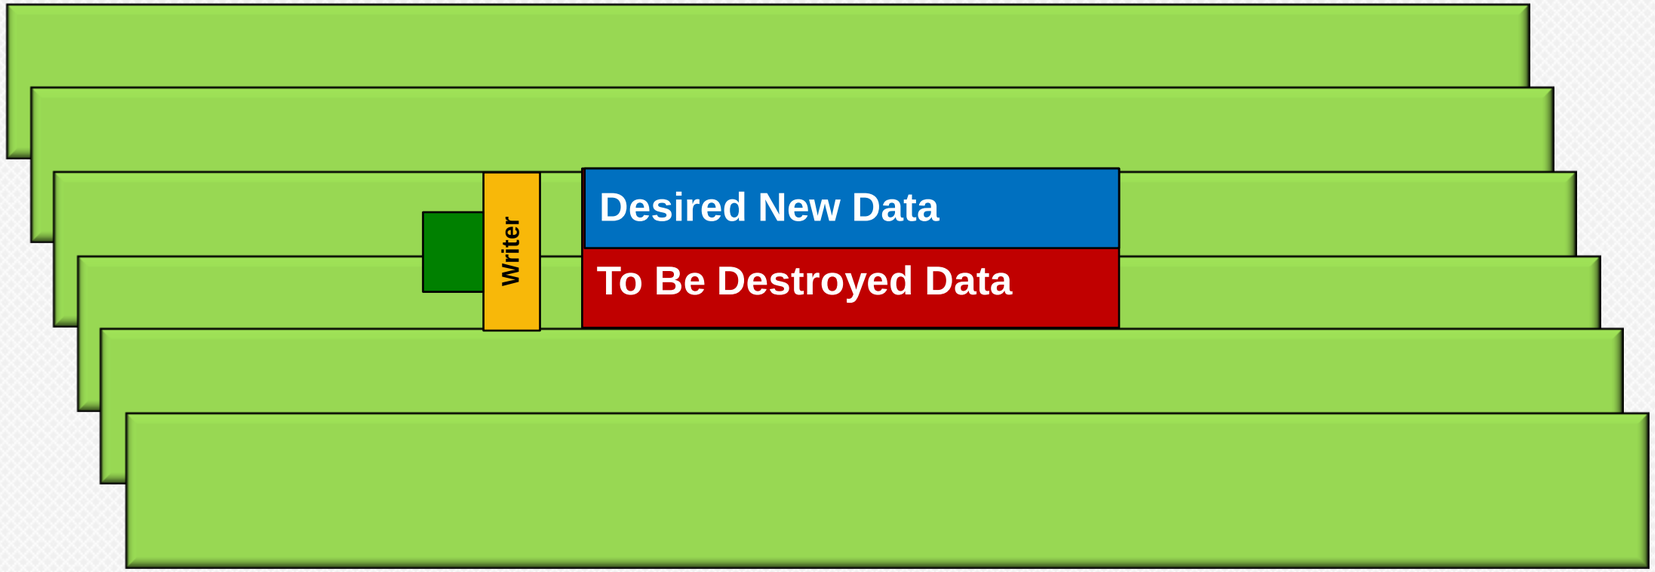

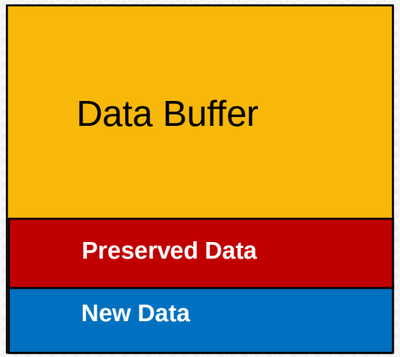

* New Data is in Buffer

* Preserve neighboring data (read) before writing New Data

What About Random Writes?

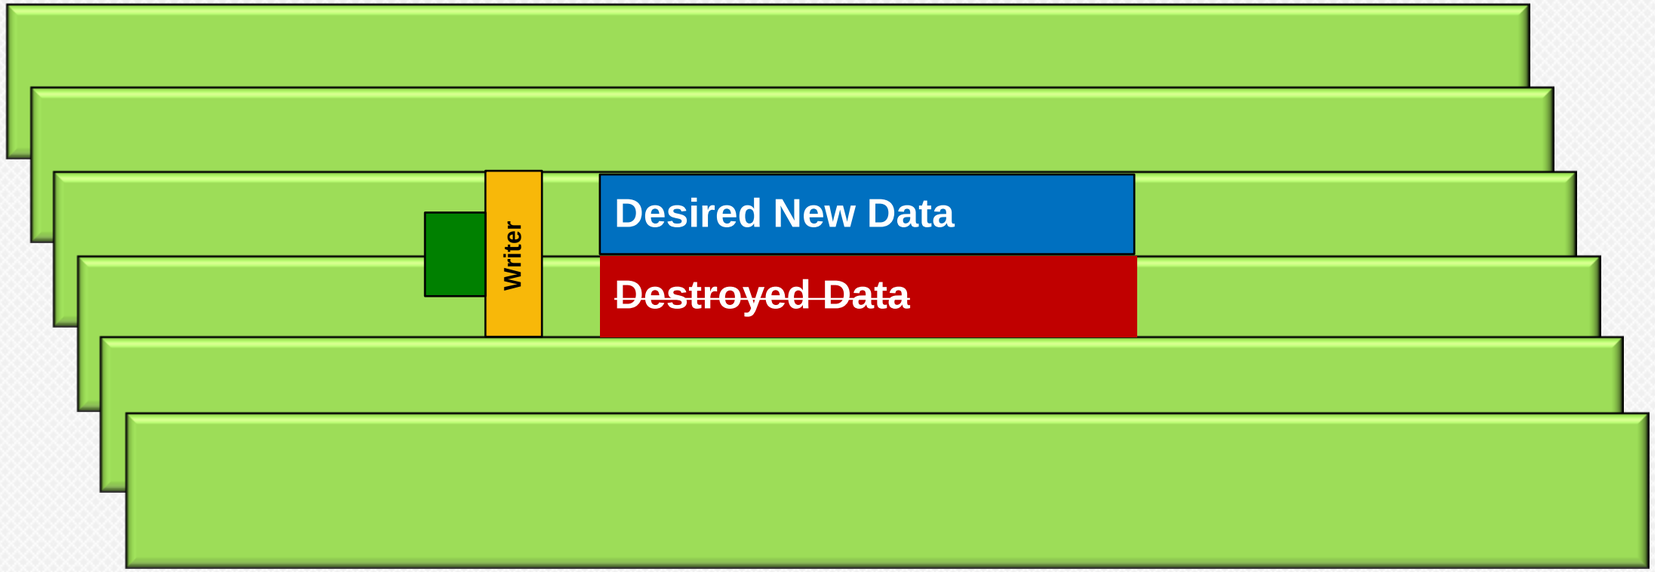

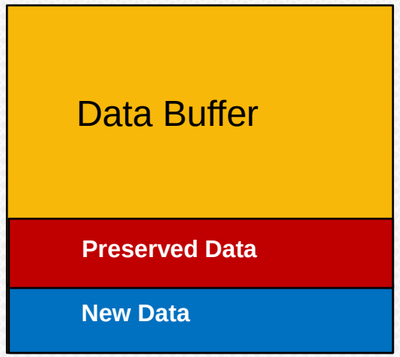

* New Data is written

* Neighboring data is destroyed

Grouping tracks into Bands

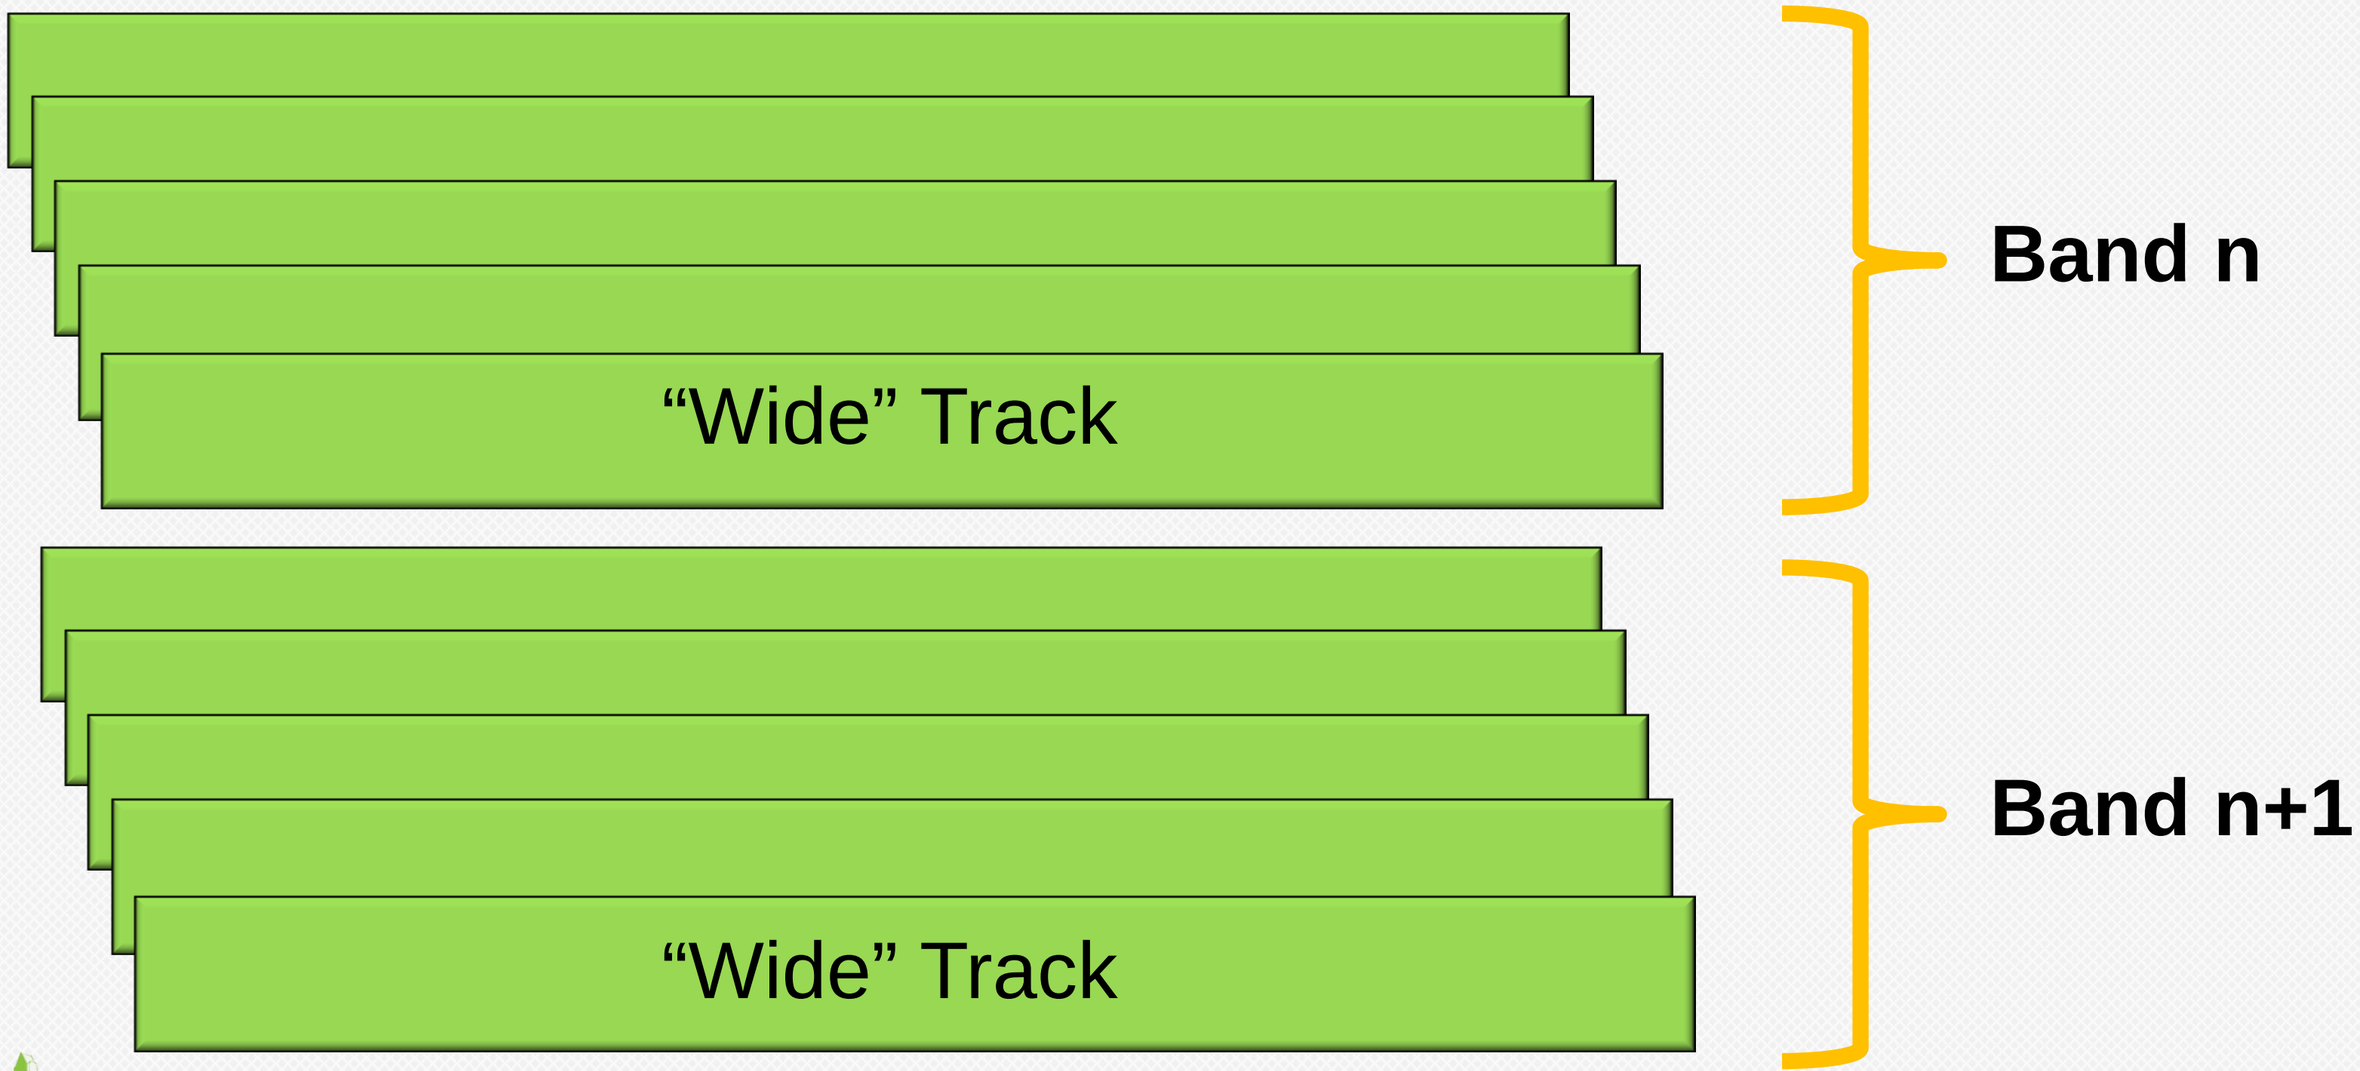

What other strategies do we have to handle random performance?

SMR Architectures

New Communication Method

Standard progressing in T10 (SAS) & T13 (SATA)

New rules for writing

Host Aware SMR

- Host can optimize write behavior
- What if new write rules are not 100% followed?
  - Drive Managed performance
- Backward Compatible

Restricted SMR

- Host must optimize write behavior
- What if new write rules are not 100% followed?
  - Device rejects the request
- Not backward compatible – new device type

Drive Managed SMR HDD

No Standard required
Drive manages all requests
Fully backward compatible

OPEN
Compute Summit

Power Management

Power Management

- Enables storage devices to operate in lower power states with faster recovery latencies
- Firmware Features enable hosts to manage or control power states
  - Automated Power Management (APM)
  - PowerChoice™

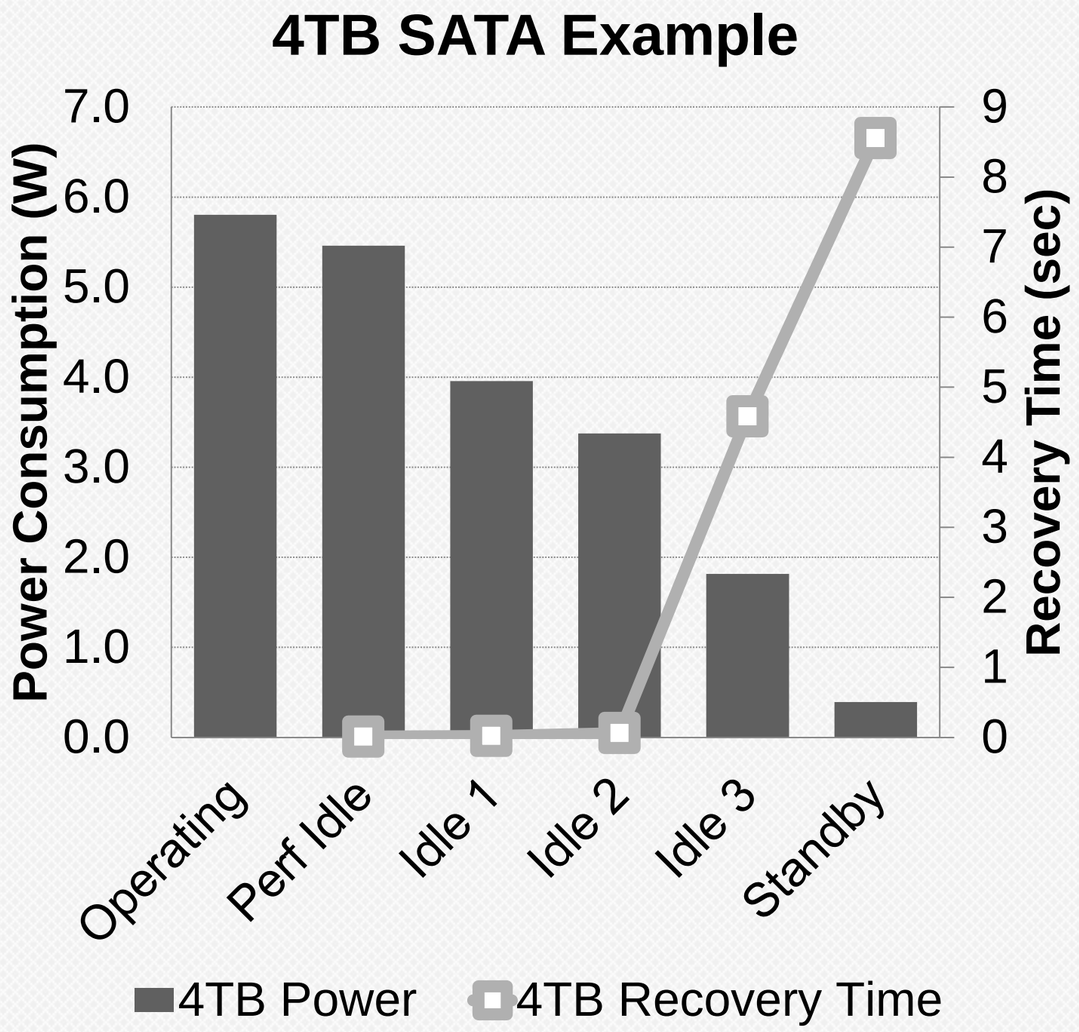

Drive Heath Management …

- In-Field Drive Diagnostics (IDD)

What is In-Field Drive Diagnostics (IDD)?

- Drives conduct self diagnostics in-situ – No Drive Removal!
  - Firmware Feature
  - Host initiates IDD

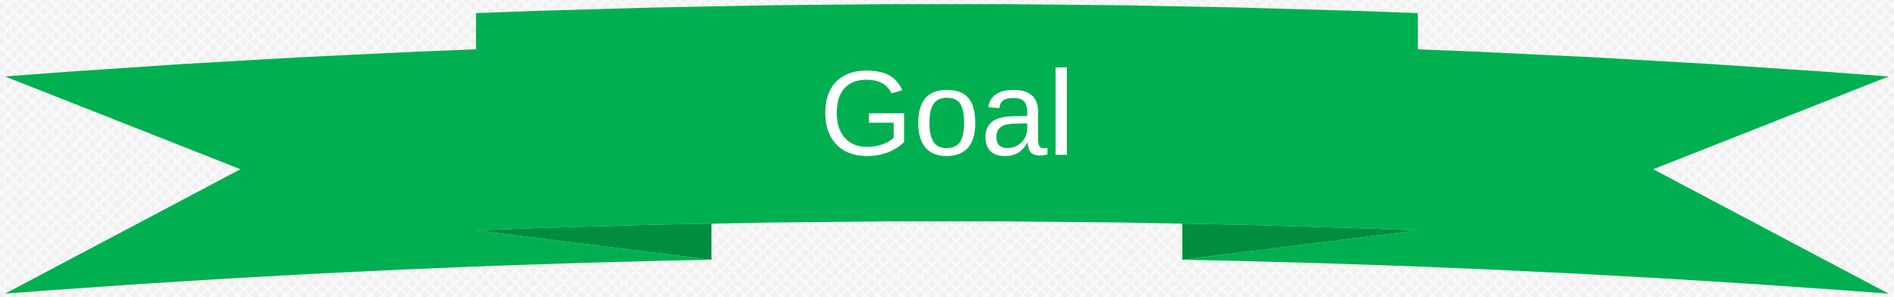

Reduce operating cost

by reducing unnecessary, costly, and time-consuming drive removal

In-Drive Diagnostics Details

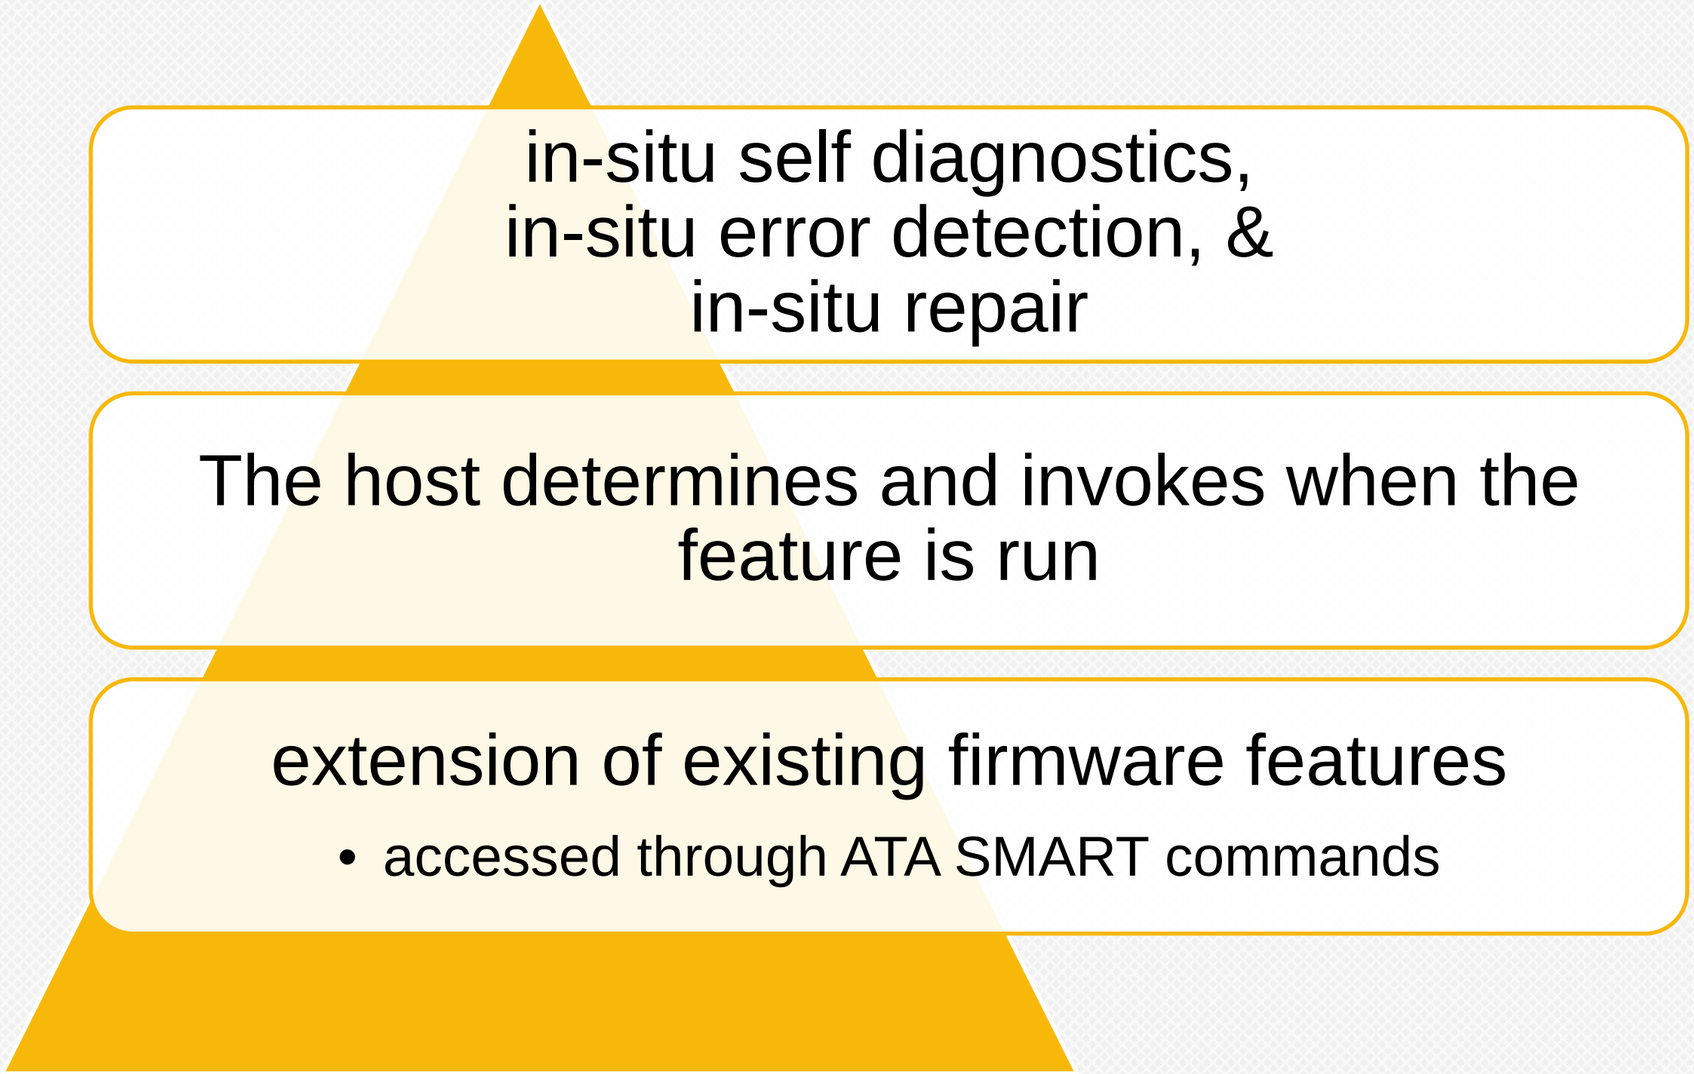

Value to Customers

IDD

1) Reduces TCO for large scale-out data centers
2) Aligns with hyper-scale data center automation philosophy and architecture
3) Reduces costs associated with drive maintenance – HDD checks, HDD pulls and replacement, HDD returns, and human errors related to drive maintenance
4) Improves operational efficiencies – load balancing, minimizes data migration and resources needed for migrating data

Call to Action for Cold Storage Devices

Engage with your HDD provider on SMR

- Participate in changes to standards and file systems
- T10 (SAS) & T13 (SATA)

Lower the operational costs of Cold Storage

- Lower energy consumption → Power Management features
- Better operational efficiency → In-Drive-Diagnostics for drive health

OPEN
Compute Summit

Thank You

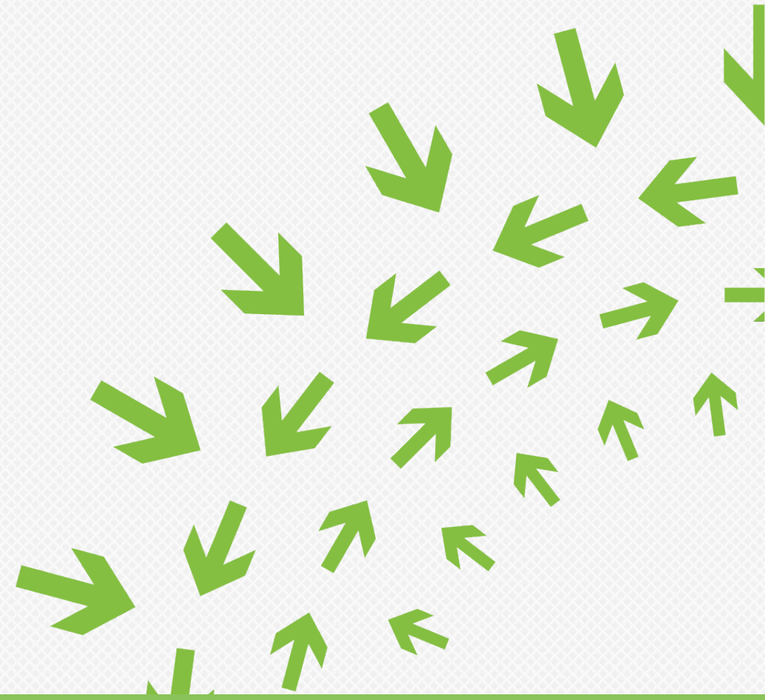

Jun Liu
Cloud Architect
Jun.a.liu@seagate.com

Josh Tinker
Sr. Product Manager
joshua.b.tinker@seagate.com