HPC Archive Storage

Trends, Technologies, and Integrated Solutions

April Alstrin, PhD
Director, Enterprise Tape Development
February 11, 2016

ORACLE

Program Agenda

1. Oracle StorageTek Strategy
2. StorageTek Libraries Update
3. StorageTek T10000D Update
4. Future Technology Update

Disk vs Tape

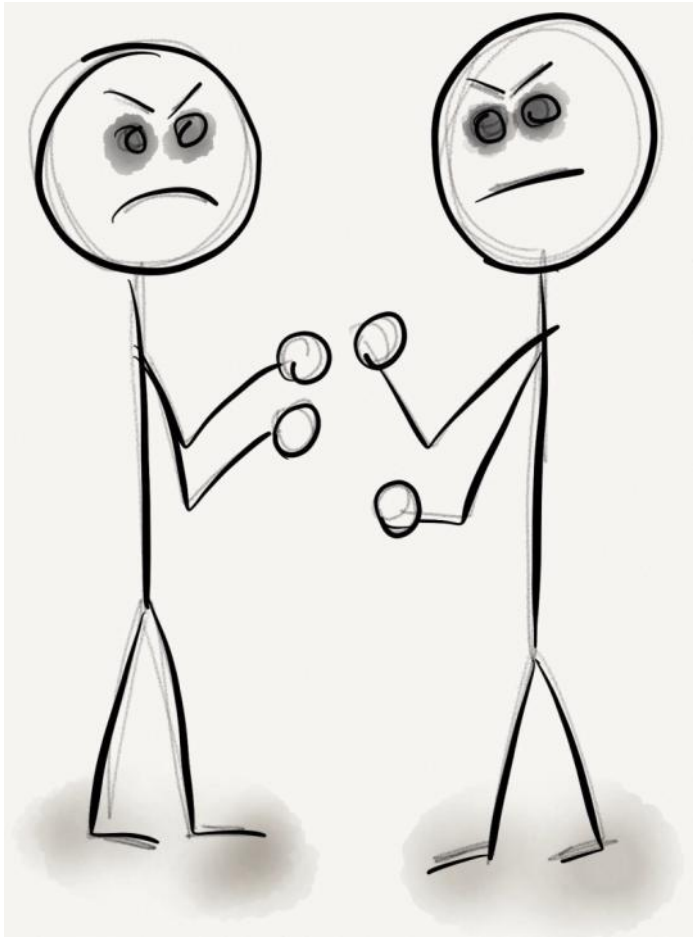

The strengths of disk and tape can be combined to deliver

lower risk of data loss,

faster performance,

and **lower TCO**

than either technology can offer on its own

Data Growth and Storage Requirements

35,000 Exabytes in 2020

Information is growing at 50% a year *

80% of data is never used after 90 days

Cost to power and manage storage is increasing

800 Exabytes in 2009

Source: IDC Digital Universe Study, sponsored by EMC

Breakout of storage growth by storage media type

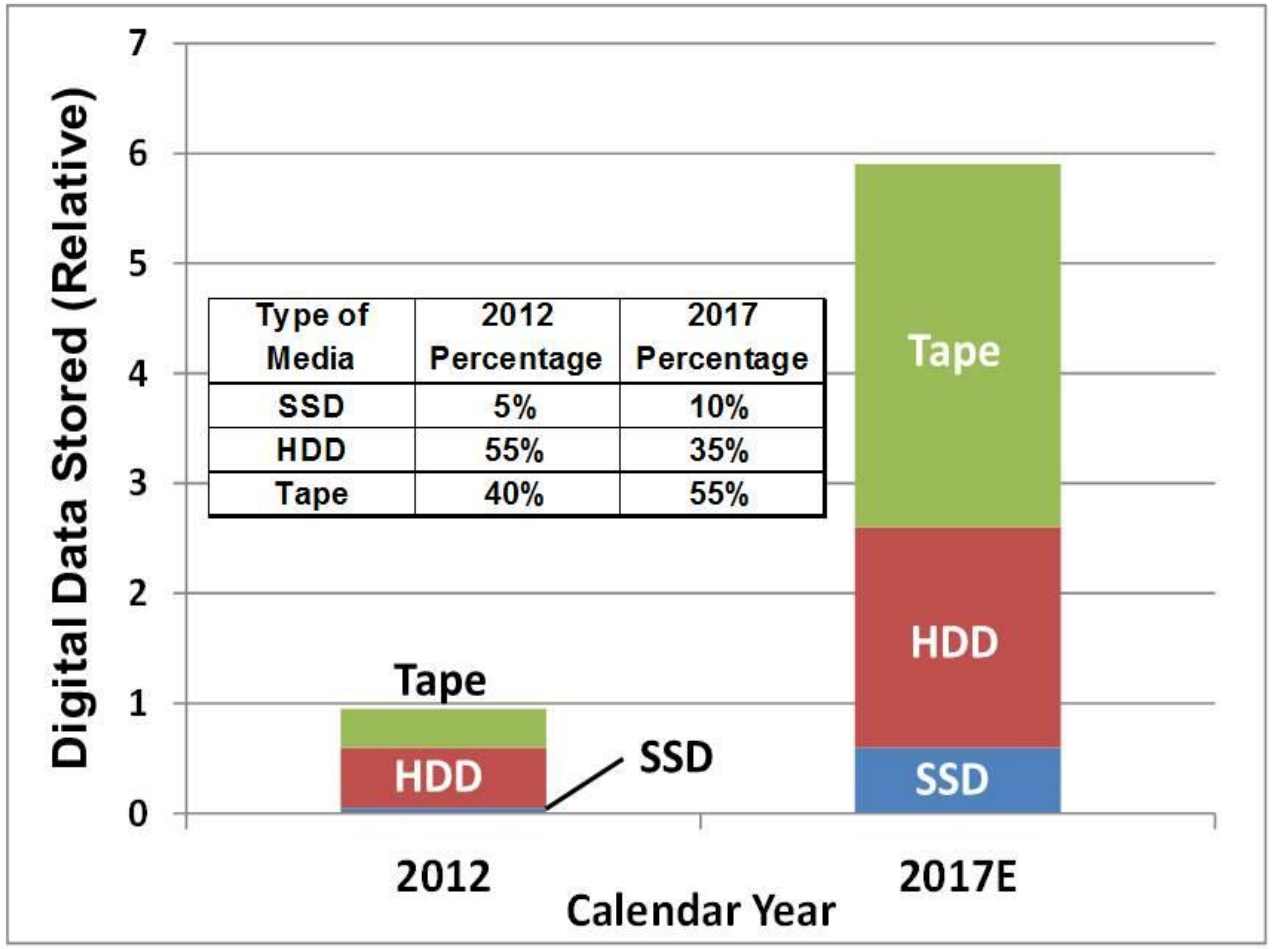

F. Moore, "The Era of Colossal Content Arrives," Horison.

Tape Usage is Predominant and Growing

Data Protection - Tape is Standard

78%

of Enterprises use both Disk and Tape in Data Protection for a Tiered Solution[1]

Archive - Fastest Growing Use Case

52% Tape CAGR

Majority of data stored on tape [2]

1 Gartner Group. Referenced in http://www.cio.in/sponsored-white-paper/gartner-reveals-best-practices-data-backup
2 Horison, Inc. The Era of Colossal Content Arrives

Storage Device Price Trends and Predictions

Areal Density Trends Drive Lower Cost per GB

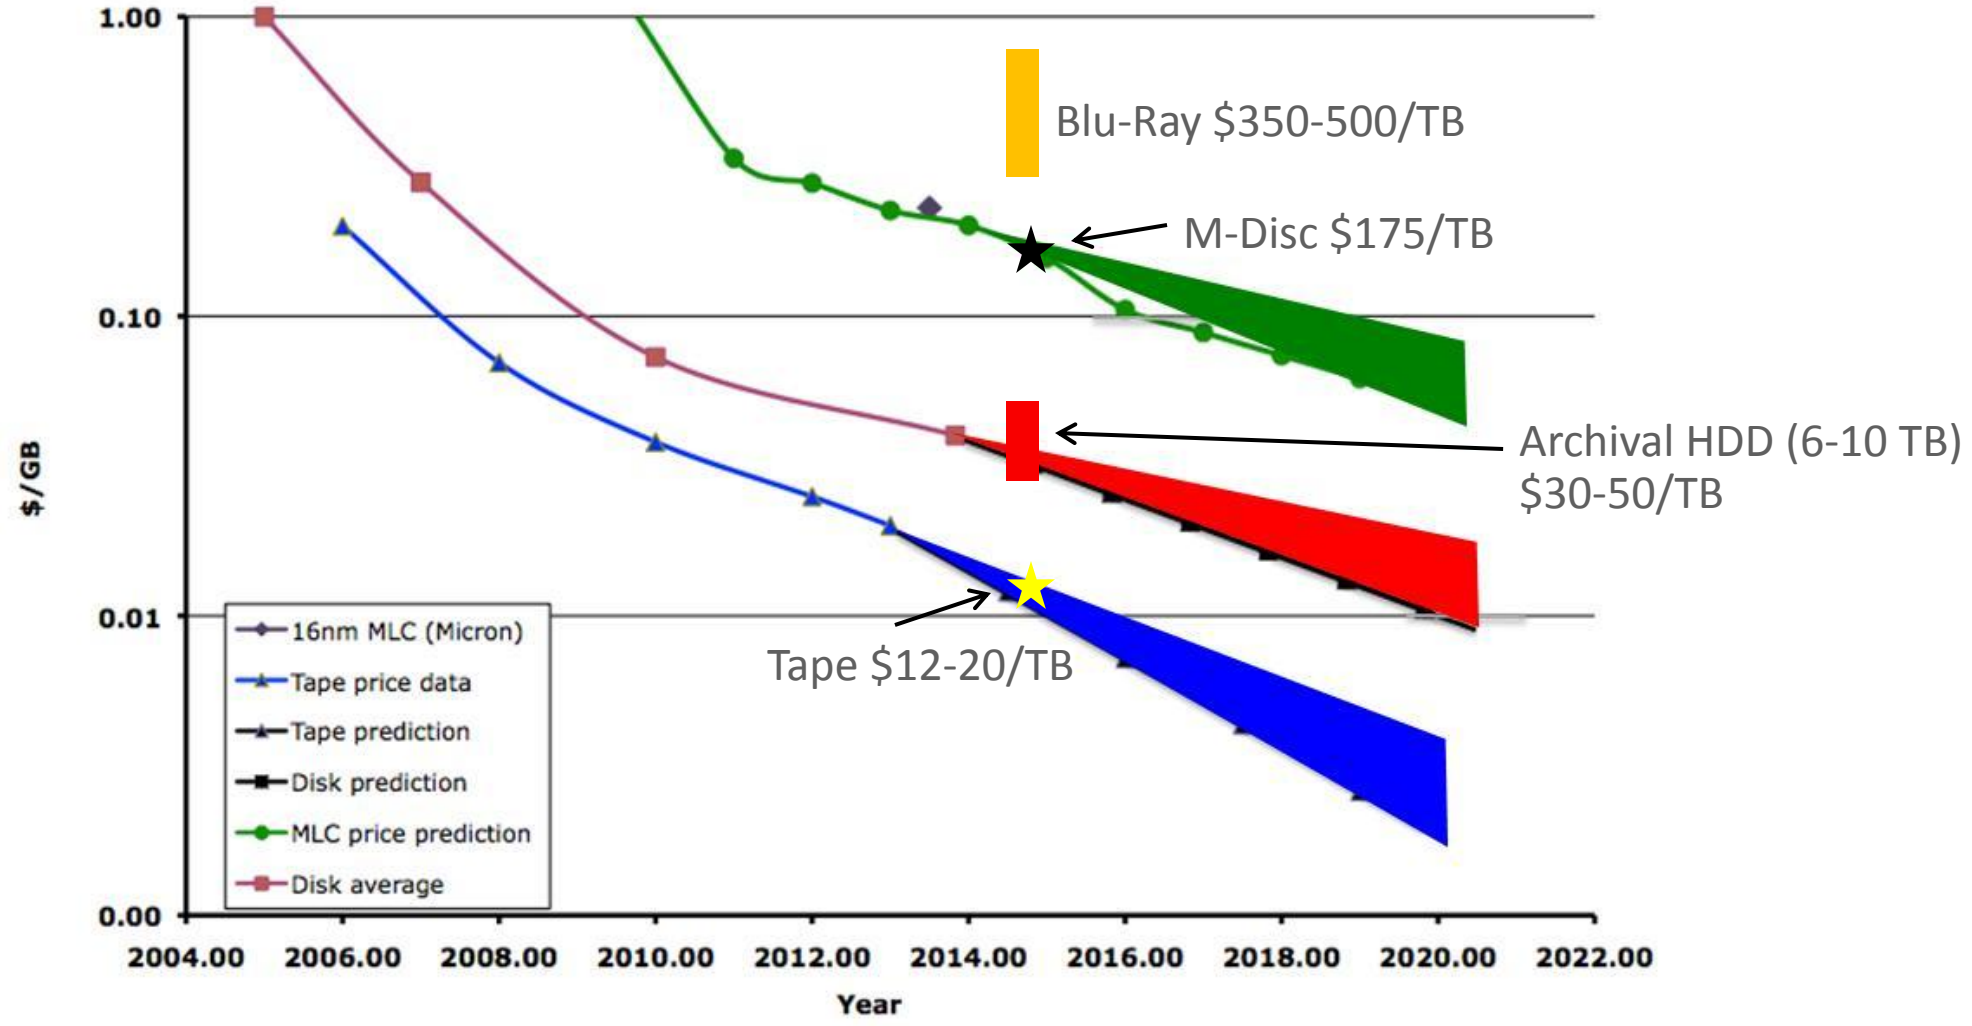

Advantages Continue to Propel Investment in Tape

Lowest cost per TB

Lowest data security risk

Lowest energy use

Lowest migration frequency

Oracle's Complete Storage Portfolio

Engineered Efficiency to Lower Costs and Improve Performance

Engineered Systems

Exadata | Exalogic | Big Data Appliance | SPARC SuperCluster

NAS Storage

ZS Series

SAN Storage

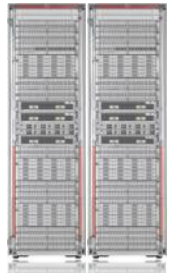

FS Series

Archive Storage

LTO | T10000 | SL150 | SL3000 | SL8500 | VSM

Cloud Storage

Deployment Options: Private, Public, Hybrid

Services: IaaS, PaaS, SaaS

Consumption Options: Build, Manage, Subscribe

Storage Software

Storage Management: OISP, OEM, ASM, StorageTek Tape Analytics, CAM, ACSLS, ELS

Automated Tiering: Partitions, Hybrid Storage Pools, DIVArchive, HSM, VSM

Data Reduction: 11g ACO, ADO, HCC, RMAN, ZFS Storage Appliance Dedup/Comp

Data Protection: Data Guard, RMAN, OSB, ZFS Storage Appliance Snap/Rep, MaxRep

Security/Encryption: ASO, Oracle Key Manager, Disk/Tape Encryption

Oracle's Broomfield, Colorado Campus

- Archive Storage R&D
  - 225,000 ft^2 facility
  - 300+ staff dedicated to tape
- Oracle IT
- Cloud
- Recording Head Operations
  - only vendor with captive wafer fab
  - integration with T10000 engineering
  - #1 market share (56%)

Oracle StorageTek Tape Strategy

- Increase pace of innovation
  - Increased Oracle investment
  - Leveraging Oracle software and hardware for tape products
- Focus on technology for heterogeneous environments
  - Reduce $/TB and lower management costs
  - Increase performance and reliability
- Address key growth areas
  - High Performance Computing
  - Digital Archiving
  - Cloud

Oracle Acquisition of Front Porch Digital DIVA Solutions

ORACLE®

DIVA

#1 Content Storage Management Software

Leader in securely migrating, managing and delivering media content on-site and in the cloud.

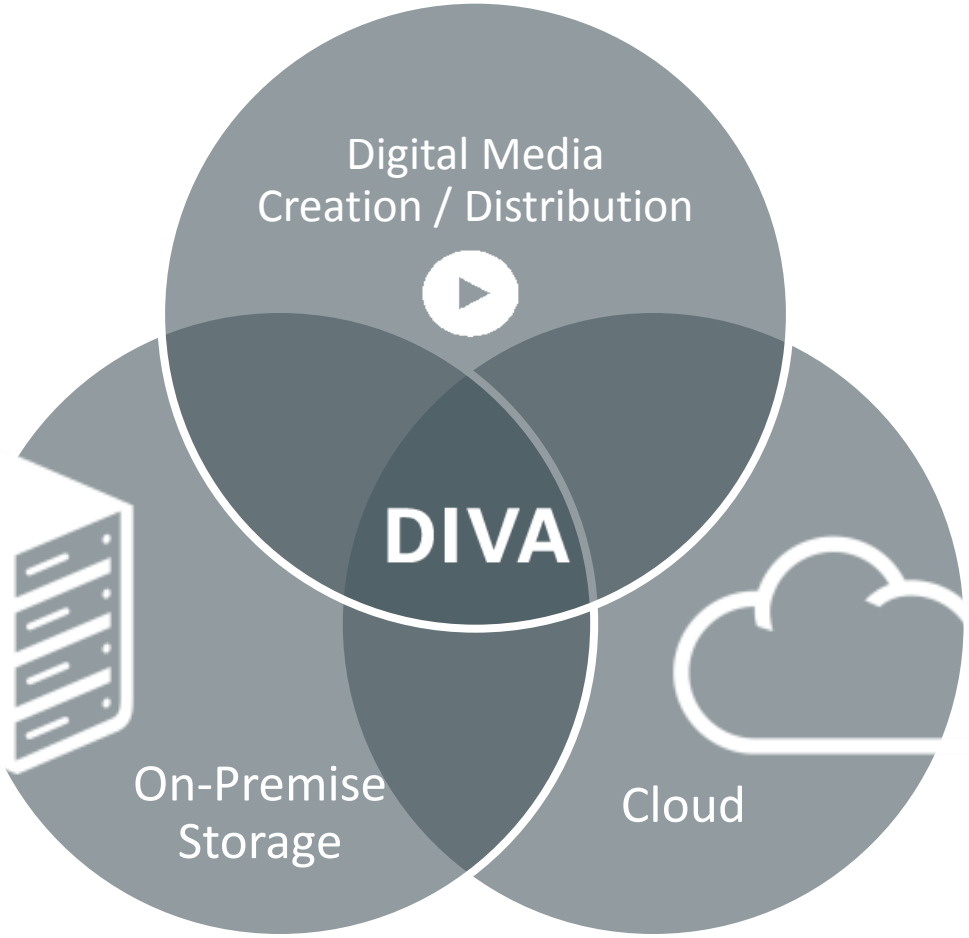

New! Oracle Storage Cloud Archive Service

Key Features

- On-demand capacity, scales to petabytes
- Multiple redundant copies of data for the highest availability
- Supports customer encrypted data
- Automatic data integrity checks for durability
- Industry standard RESTful APIs - Swift

New! Oracle Storage Cloud Archive Service

Benefits

- Ideal for large data sets of fixed content, e.g. email archives, digital video, scientific data, preservation data, etc.
- Absolutely the lowest cost solution for long term data retention
- Simplified data center provisioning and capacity planning
- Right-size on-premise assets
- Integrated with Oracle and third-party backup, archive and preservation software (check for availability)

New! Oracle Storage Cloud Archive Service

Amazon Glacier

$0.007 - $0.013*

$/GB/Mo.

Google Nearline

$0.011

$/GB/Mo.

ORACLE

$0.001

$/GB/Mo.

* Region dependent

Another Archive Tier for HPC and Scientific Data

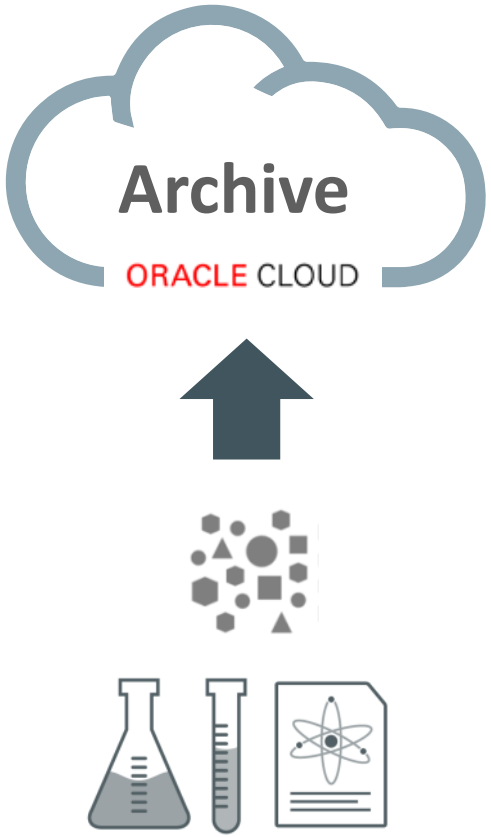

Key Benefits

- Lowest TCO for long term retention of HPC and scientific data
- Never run out of storage with petabytes of capacity on-demand
- Move dormant data to ultra low cost cloud
- Infinite retention for future validation or data re-use
- Durable storage with high degrees of data integrity validation

Program Agenda

1. Oracle StorageTek Strategy
2. StorageTek Libraries Update
3. StorageTek T10000D Update
4. Future Technology Update

StorageTek SL8500 Modular Library System

- Manage Growth
  - Non-Disruptively scale from 1,450 to 100,000 slots and 640 drives
  - Supports up to 850PB
- Reduce Costs
  - Consolidate storage with advanced partitioning and Any Cartridge Any Slot technology
  - Save 33% on TCO with T10000D
- Minimize Risk
  - Only vendor with redundant and hot-swappable robotics, electronics and pass-through ports

StorageTek SL3000 Modular Library System

- Manage Growth
  - Scale at your own pace and pay only for what you need with Capacity on Demand
  - Highest density - up to 592TB/sq ft
  - 200 to 5,925 slots supporting up to 50PB
  - Up to 56 drives for 48TB/hr throughput
- Reduce Costs
  - Consolidate storage with advanced partitioning and Any Cartridge Any Slot technology
- Minimize Risk
  - Only vendor with redundant and hot-swappable robotics and electronics

New! SL8500 Bulk CAP

- Provides independent CAP at each rail
- 4x improvement in CAP capacity
- Up to 5x faster load/unload performance
- Assign CAPs to Partitions
- Available as upgrade for existing SL8500

New! SL8500 Bulk CAP

CAP access door

CAP control panel

SL8500 CAP Comparison

	Rotational CAPs	New Bulk CAP
Total Capacity	78 for two CAPs installed	288 for both sides all four rails
Library CAP count	Up to 2	8
Capacity per CAP	39 slots per CAP	36 slots per CAP
Rail Orientation	Single CAP spans 3 rails	2 CAPs per Rail
Serviceability	Non-disruptive with Safety Door	Non-disruptive with Safety Door

Program Agenda

1. Oracle StorageTek Strategy
2. StorageTek Libraries Update
3. StorageTek T10000D Update
4. Future Technology Update

LTO Roadmap

StorageTek LTO-7 Tape Drives

- Up to 6 TB native capacity
- 300 MB/sec data rate (uncompressed)

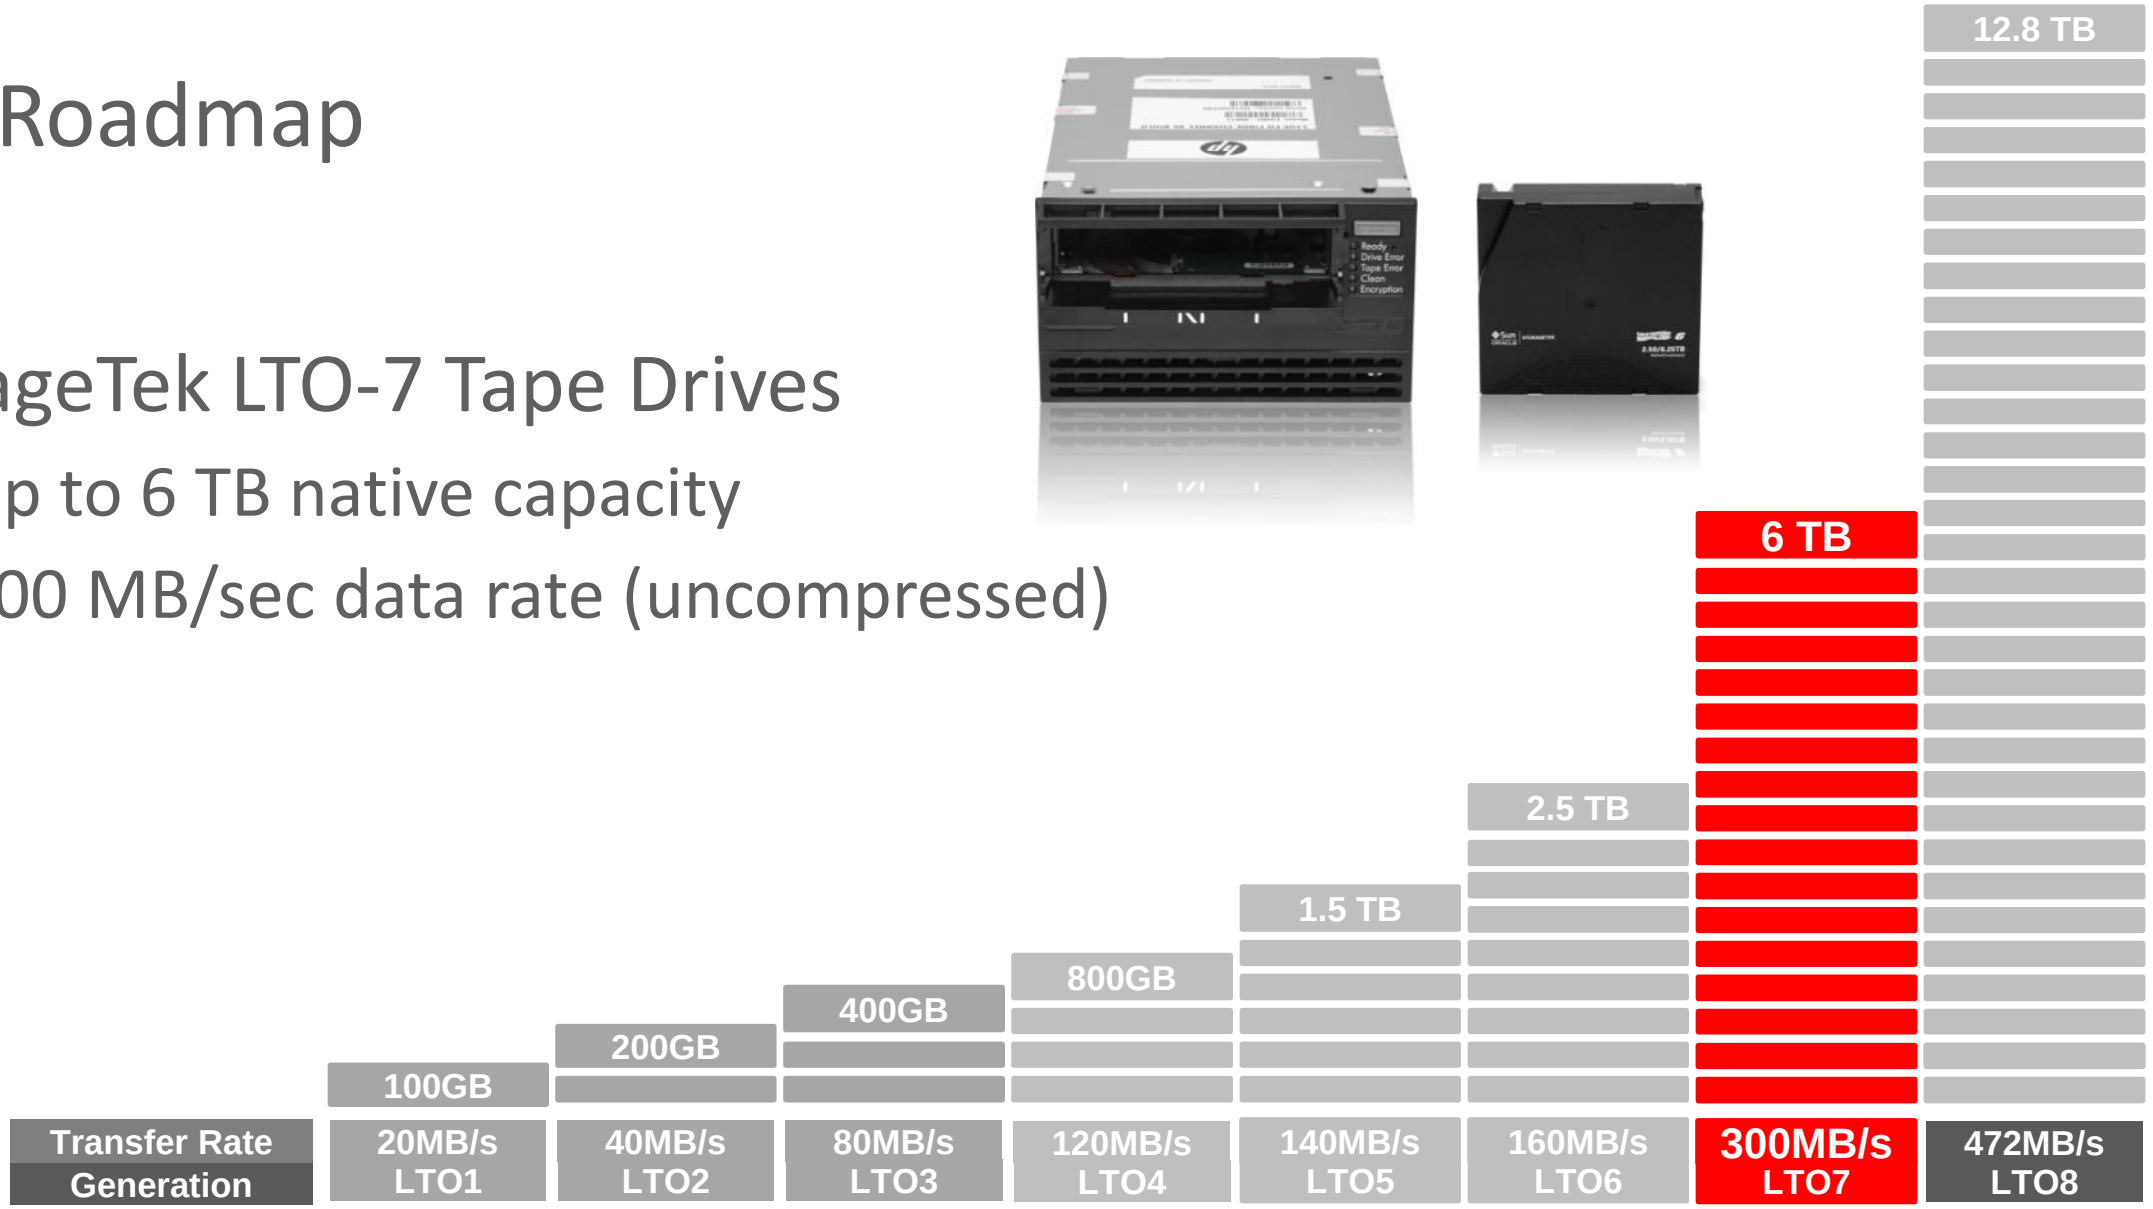

StorageTek T10000D Tape Drive

8.5 TB Capacity

252 MB/sec native data rate
up to 800 MB/sec compressed

Interface flexibility

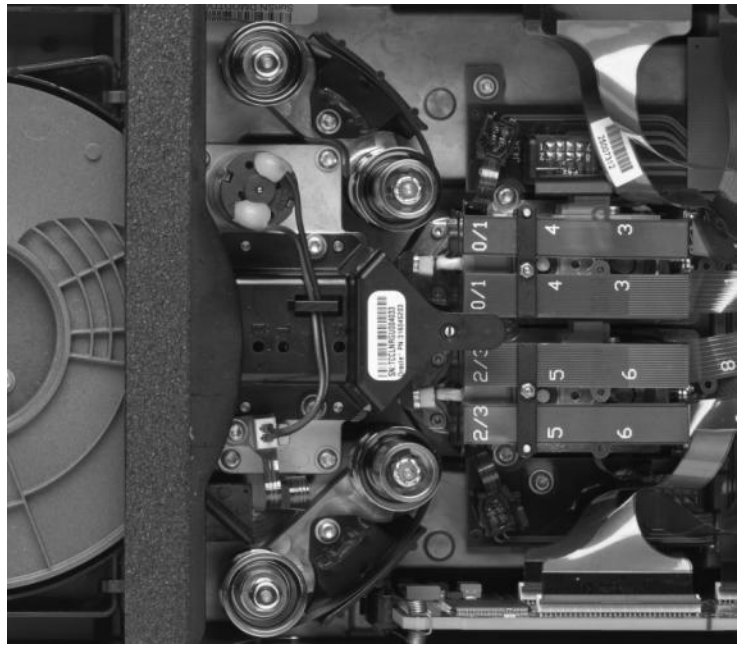

High stability archive media

Designed for 24x7 environments

Increases capacity of T10000C-written cartridges by 55%

Reads T10000A, T10000B, and T10000C-written cartridges

T10000D Feature Set

- File Handling
- Partitioning
- Data Integrity Checking
- In-Drive Media Validation
- **New!** Read Access Ordering
- **New!** Drive-to-Drive Migration

StorageTek Data Integrity Validation

Discover Corrupted Records **Before** They Are Written to Tape

Write Step 1:	Write Step 2:	StorageTek Data Integrity Validation	Write Step 3:
Server pulls file from disk, and **hardware** generates 32-bit CRC over each record within file*	Server **appends** CRC to each record and sends to T10000D	T10000D performs CRC comparison, alerts server of corrupted record	Server re-sends only the failing record

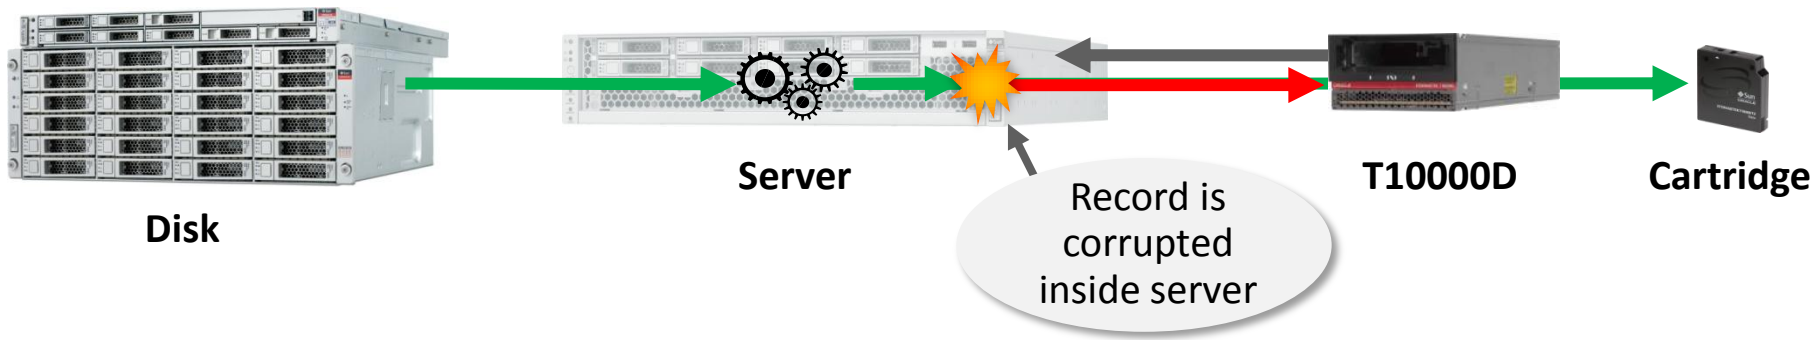

StorageTek T10000 Media Validation

Easily Validate the Integrity of Your Digital Assets with T10000D

Step 1:	Step 2:	Step 3:	Step 4:
Application sends standard 16 byte verify command to T10000D	T10000D reads records and checks DIV-CRCs (if requested)	T10000D sends validation status to server	T10000D sends drive stats to STA 2.0 during unload

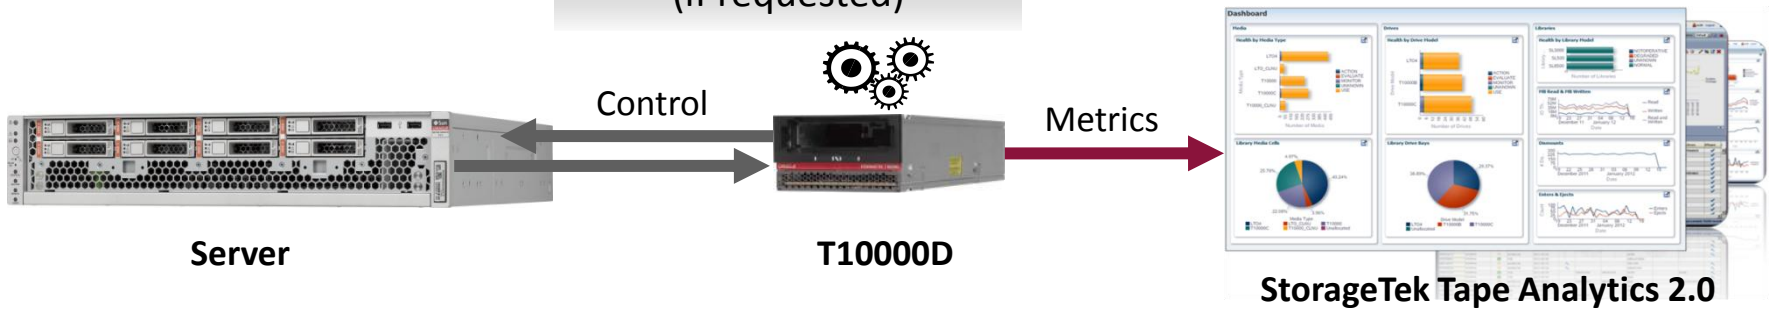

In-Drive Validation Saves Network Bandwidth and Server Resources

Media Validation with StorageTek Tape Analytics

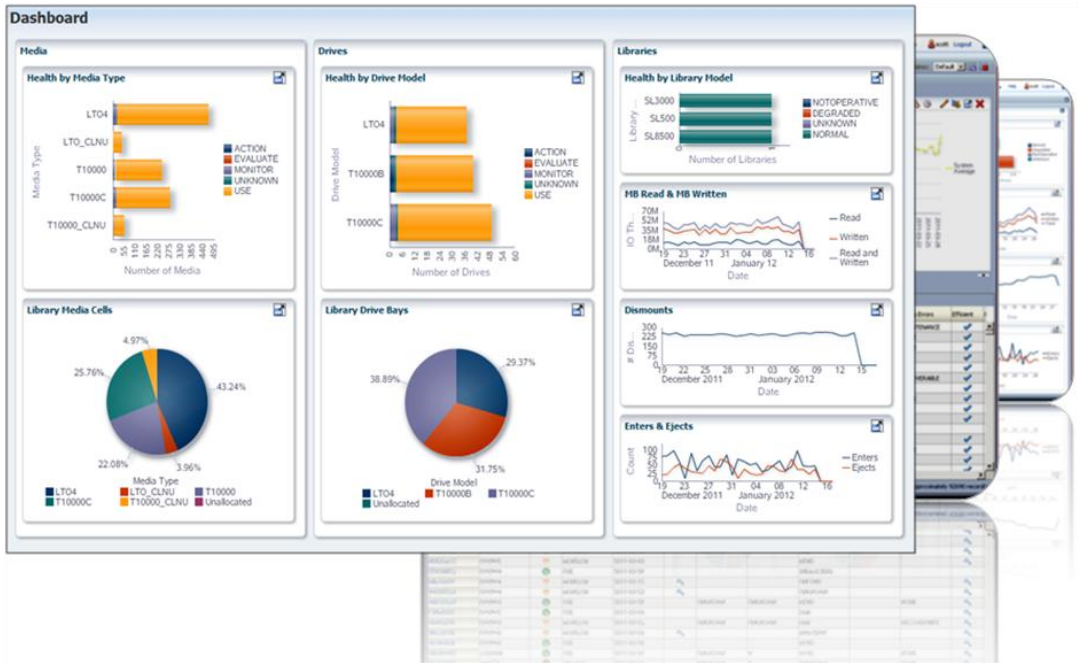

- ✓ **Savings: Saves Resources**
  Actively monitors all drives, media, and robots so you can focus your resources elsewhere

- ✓ **Smart: Leverages Intelligent Analytics**
  Proprietary algorithms provide proactive health indicators that can be trusted

- ✓ **Simple: Worry Free Deployment**
  Gathers performance data through the library without interrupting the live data path

- ✓ **Scalability: Grow with Peace of Mind**
  Scales to meet your needs and supports monitoring multiple globally dispersed libraries

Oracle's Tape Data Recovery Lab

Special Support for Enterprise Customer

- One-of-a-kind service included with standard support
- Occurrences very rare, but tape can break
- Oracle provides a dedicated, secure lab for recovering data
- Individual attention, best effort recovering specific data
- Special error recovery code creation
- Upon success, original and recovered copy returned

New! Read Files in Optimum Order for Faster Access

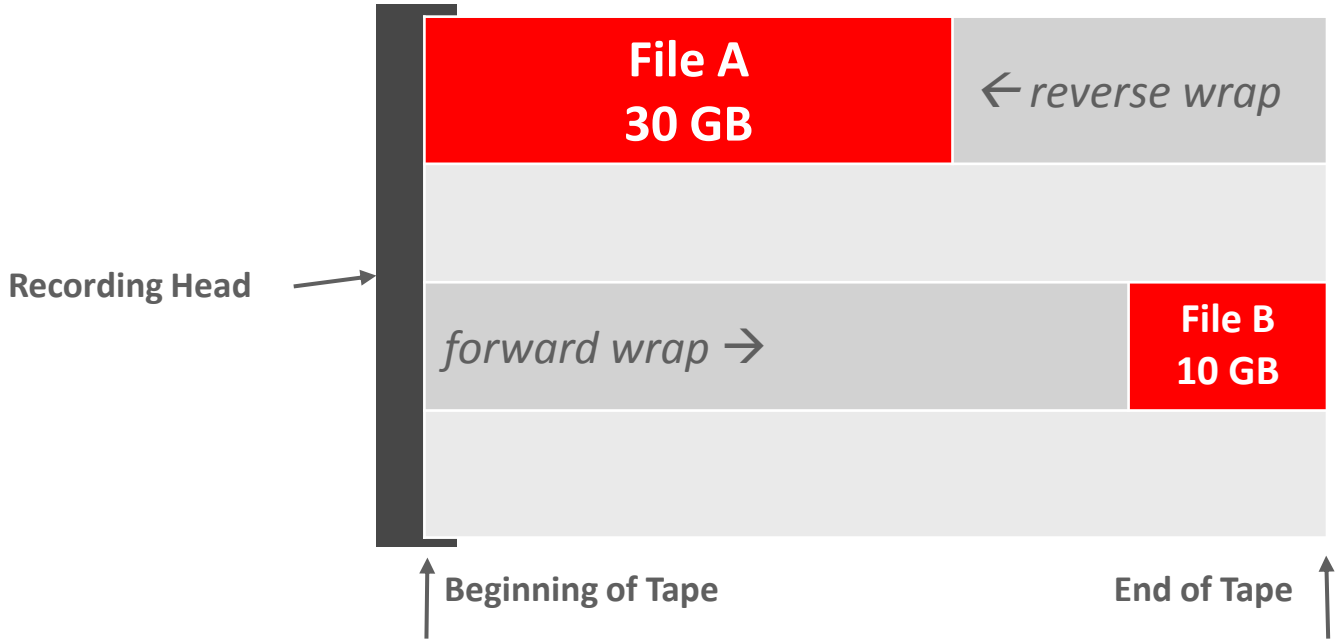

	Time
Reading A then B	00:06:40
Reading B then A	00:04:49

← Reading Alphanumerically

← Reading in Optimum Order

New! StorageTek File Access Accelerator

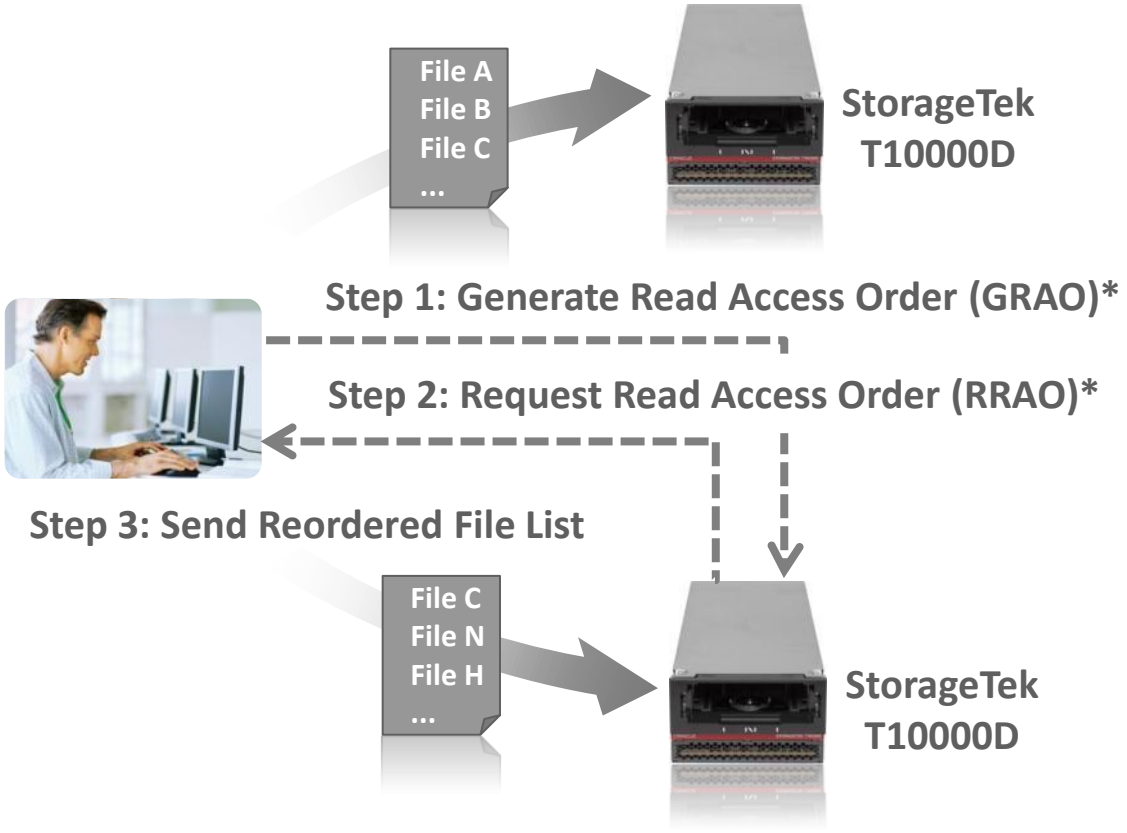

Read Files A thru U
without File Access Accelerator

00:33:21

00:21:47

Read Files A thru U
with File Access Accelerator

30%-60% time savings measured at Lawrence Livermore with HPSS

https://huf2015.scinet.utoronto.ca/abstracts/?task=download&attachmentID=34

***Supports T10 ANSI Standard 13-266r4**

New! StorageTek Direct Copy

Migration Challenges

- **Time**: months or years to complete
- **Cost**: CPU hungry with little business value
- **Complexity**: manage multiple formats

New! StorageTek Direct Copy

Migration Challenges

- **Time**: months or years to complete
- **Cost**: CPU hungry with little business value
- **Complexity**: manage multiple formats

Solution: Direct Copy

Copy **directly** from drive to drive

- Enable parallel migration
- Free CPU cycles & personnel
- Get to single format faster
- Full data protection

Migrate Data without Impacting Production

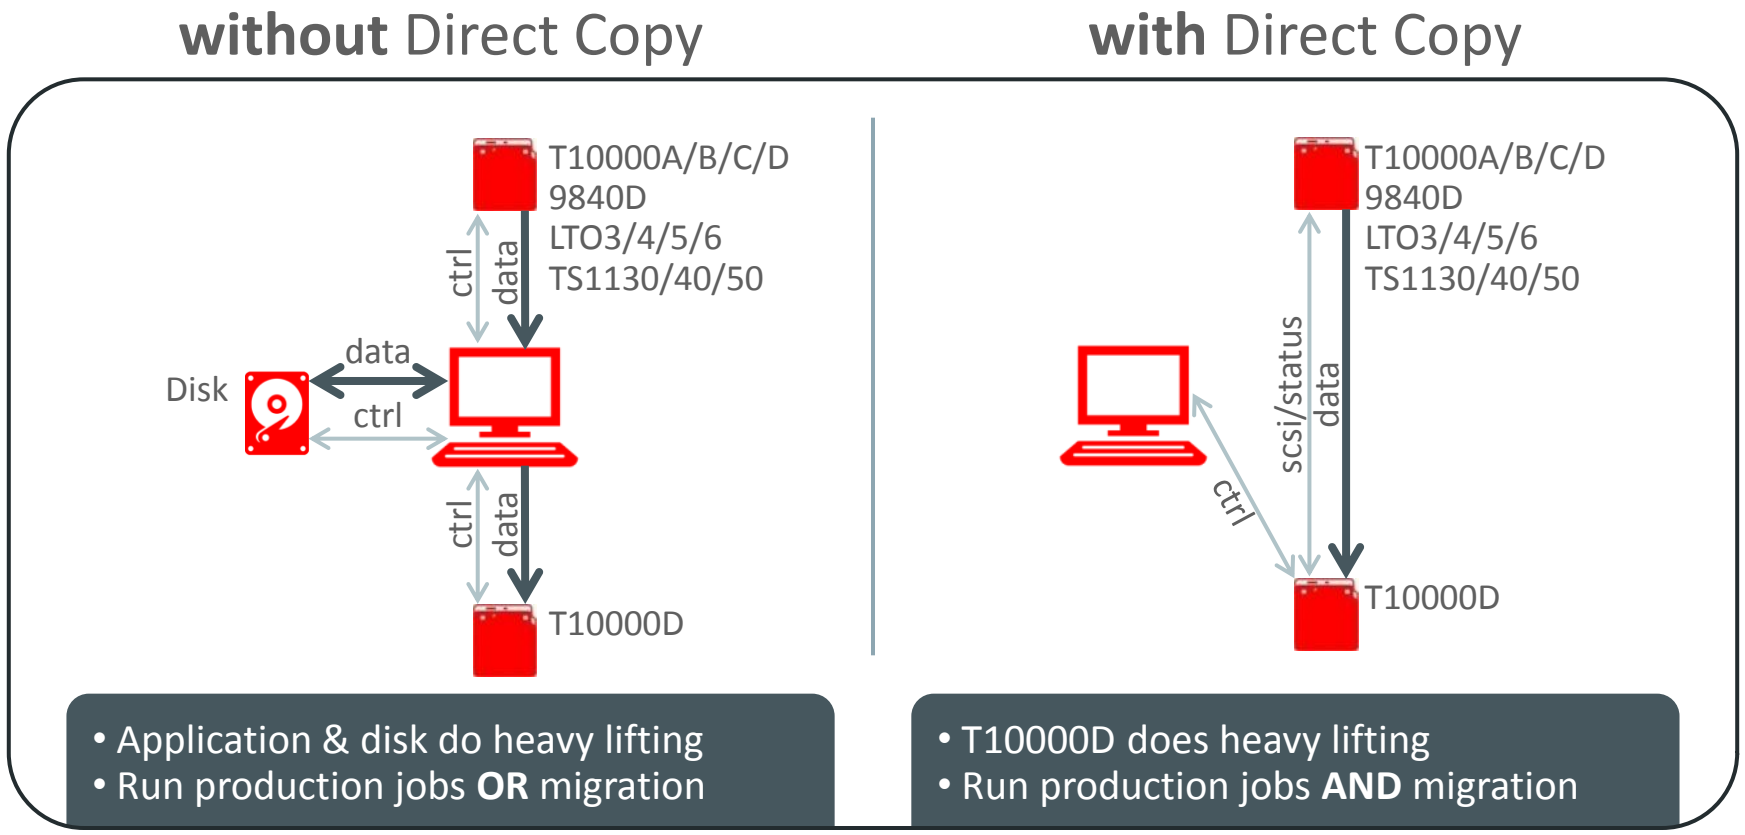

Migrate All Data or **Repack** a Portion of Data

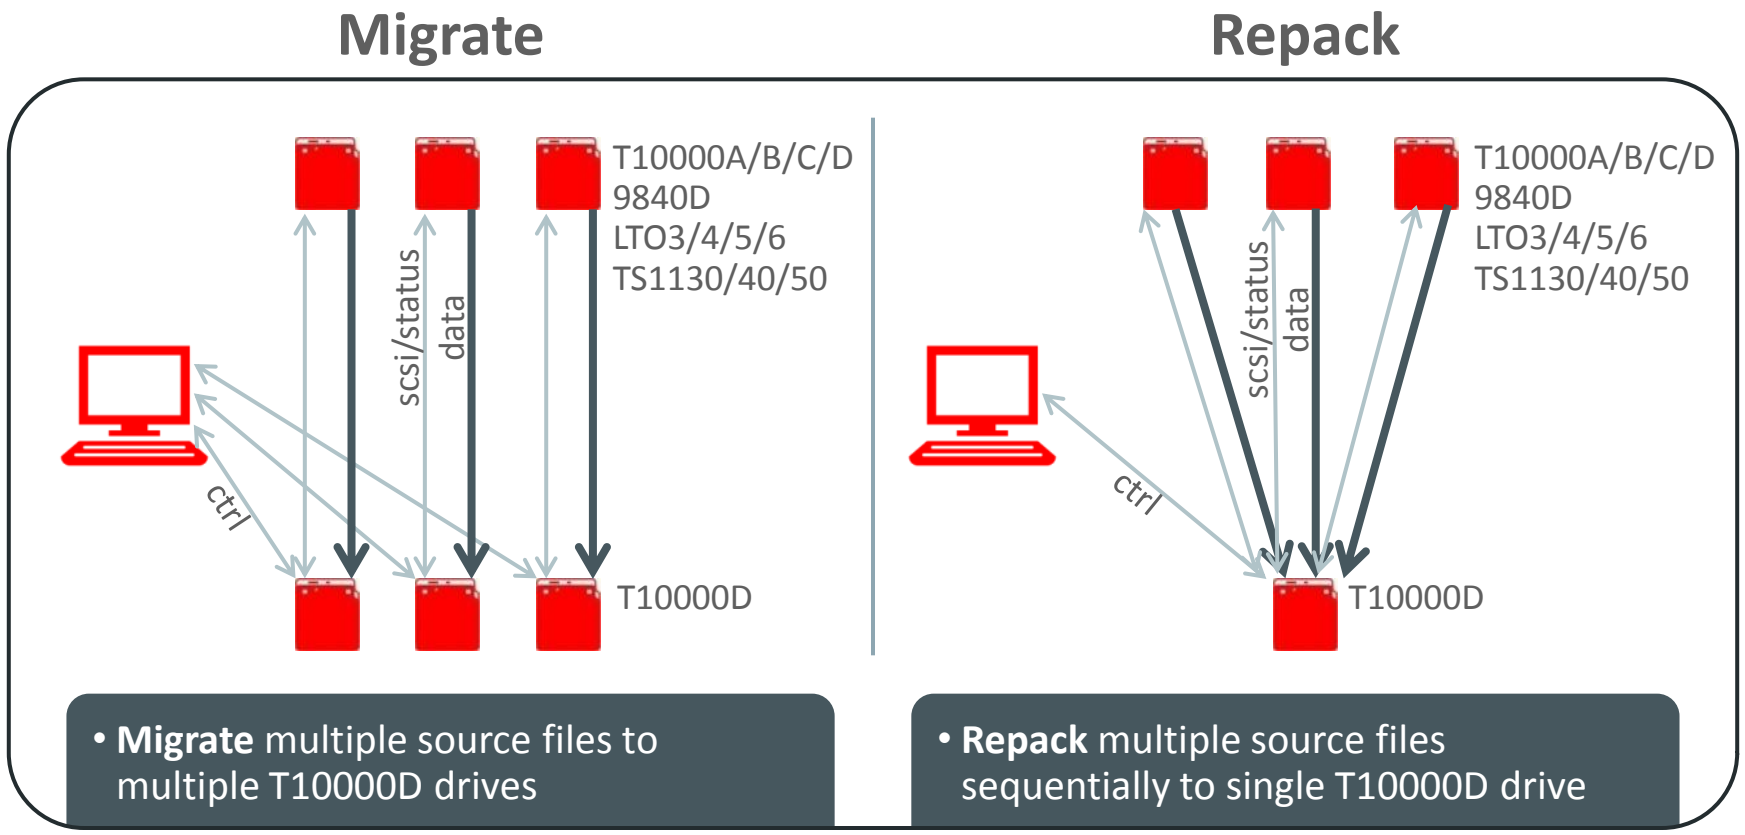

Migrate Example

- 6.3 PB from LTO-5 to T10000D
- 5x LTO-5 drives streaming volumes at 140MB/sec to 5x T10000D drives
- 5 streams x 140MB/sec = 2.5 TB/hour
- Migrate 6.3 PB in 15 weeks (24x7)
  - near-zero CPU overhead, run in parallel with production
  - no cache required
  - no HWM tuning

Reflect Multiple Copies or **Snap** Copies of Existing Data

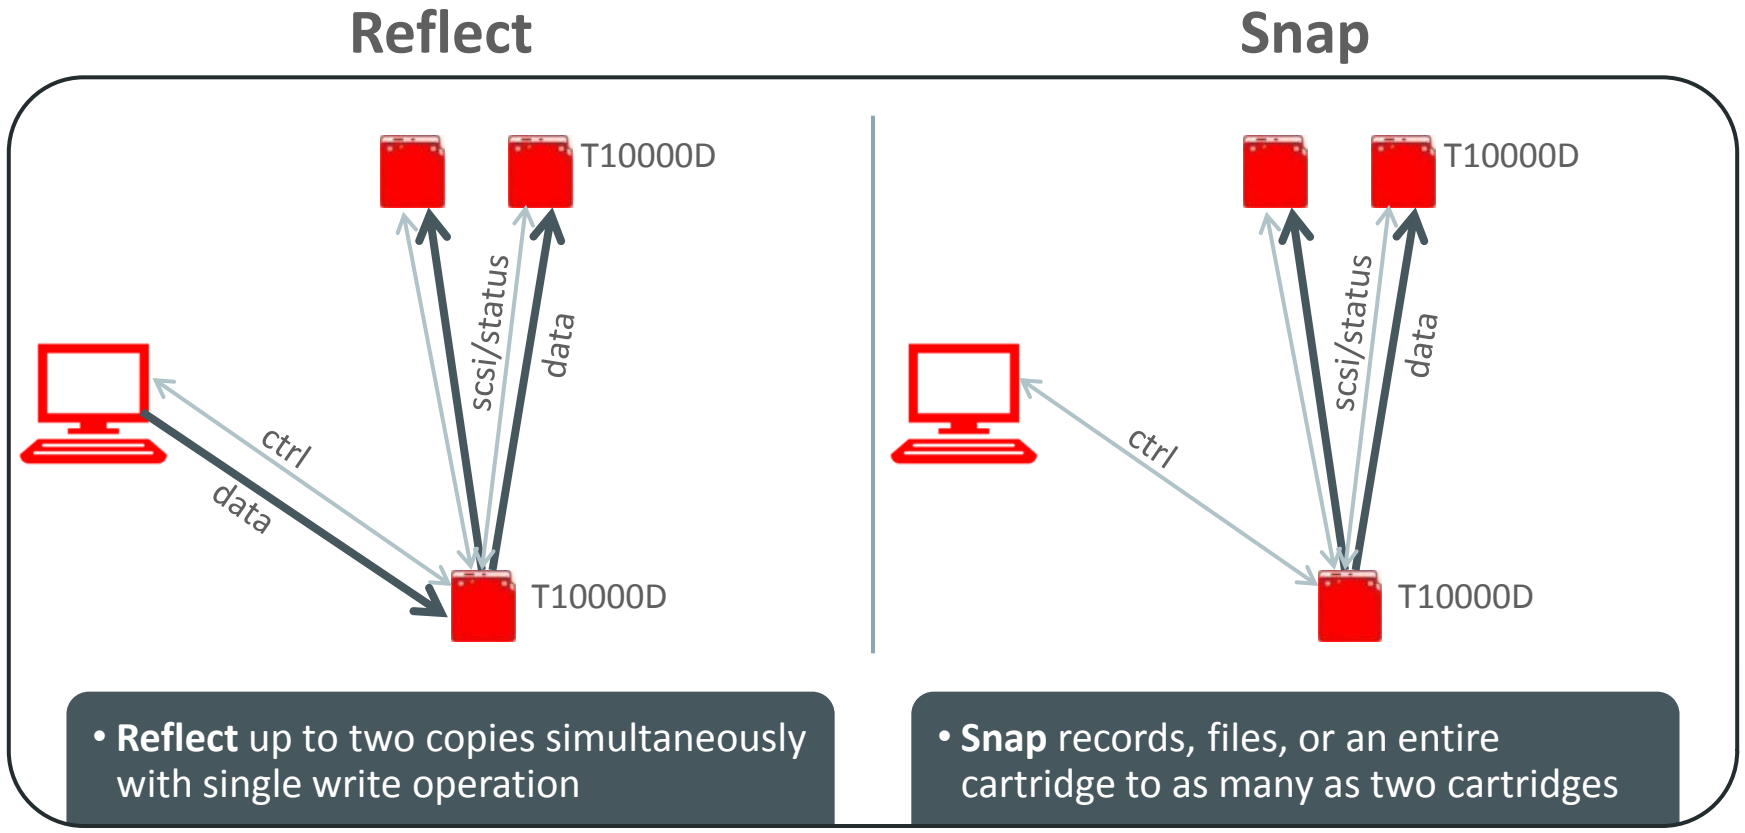

- **Reflect** up to two copies simultaneously with single write operation
- **Snap** records, files, or an entire cartridge to as many as two cartridges

New! RAIT Deployment

HPC Tape Storage Challenges

Time: store TB+ files to tape

Cost: dual copy of all data for high availability

Solution: RAIT with T10000D

Write files with **multiple** tape drives

- 4X or 8X faster performance
- Parity protection w/o dual copy
- High availability with single copy

Program Agenda

1. Oracle StorageTek Strategy
2. StorageTek Libraries Update
3. StorageTek T10000D Update
4. Future Technology Update

Areal Density Trends

Chart provided courtesy of the Information Storage Industry Consortium (INSIC)

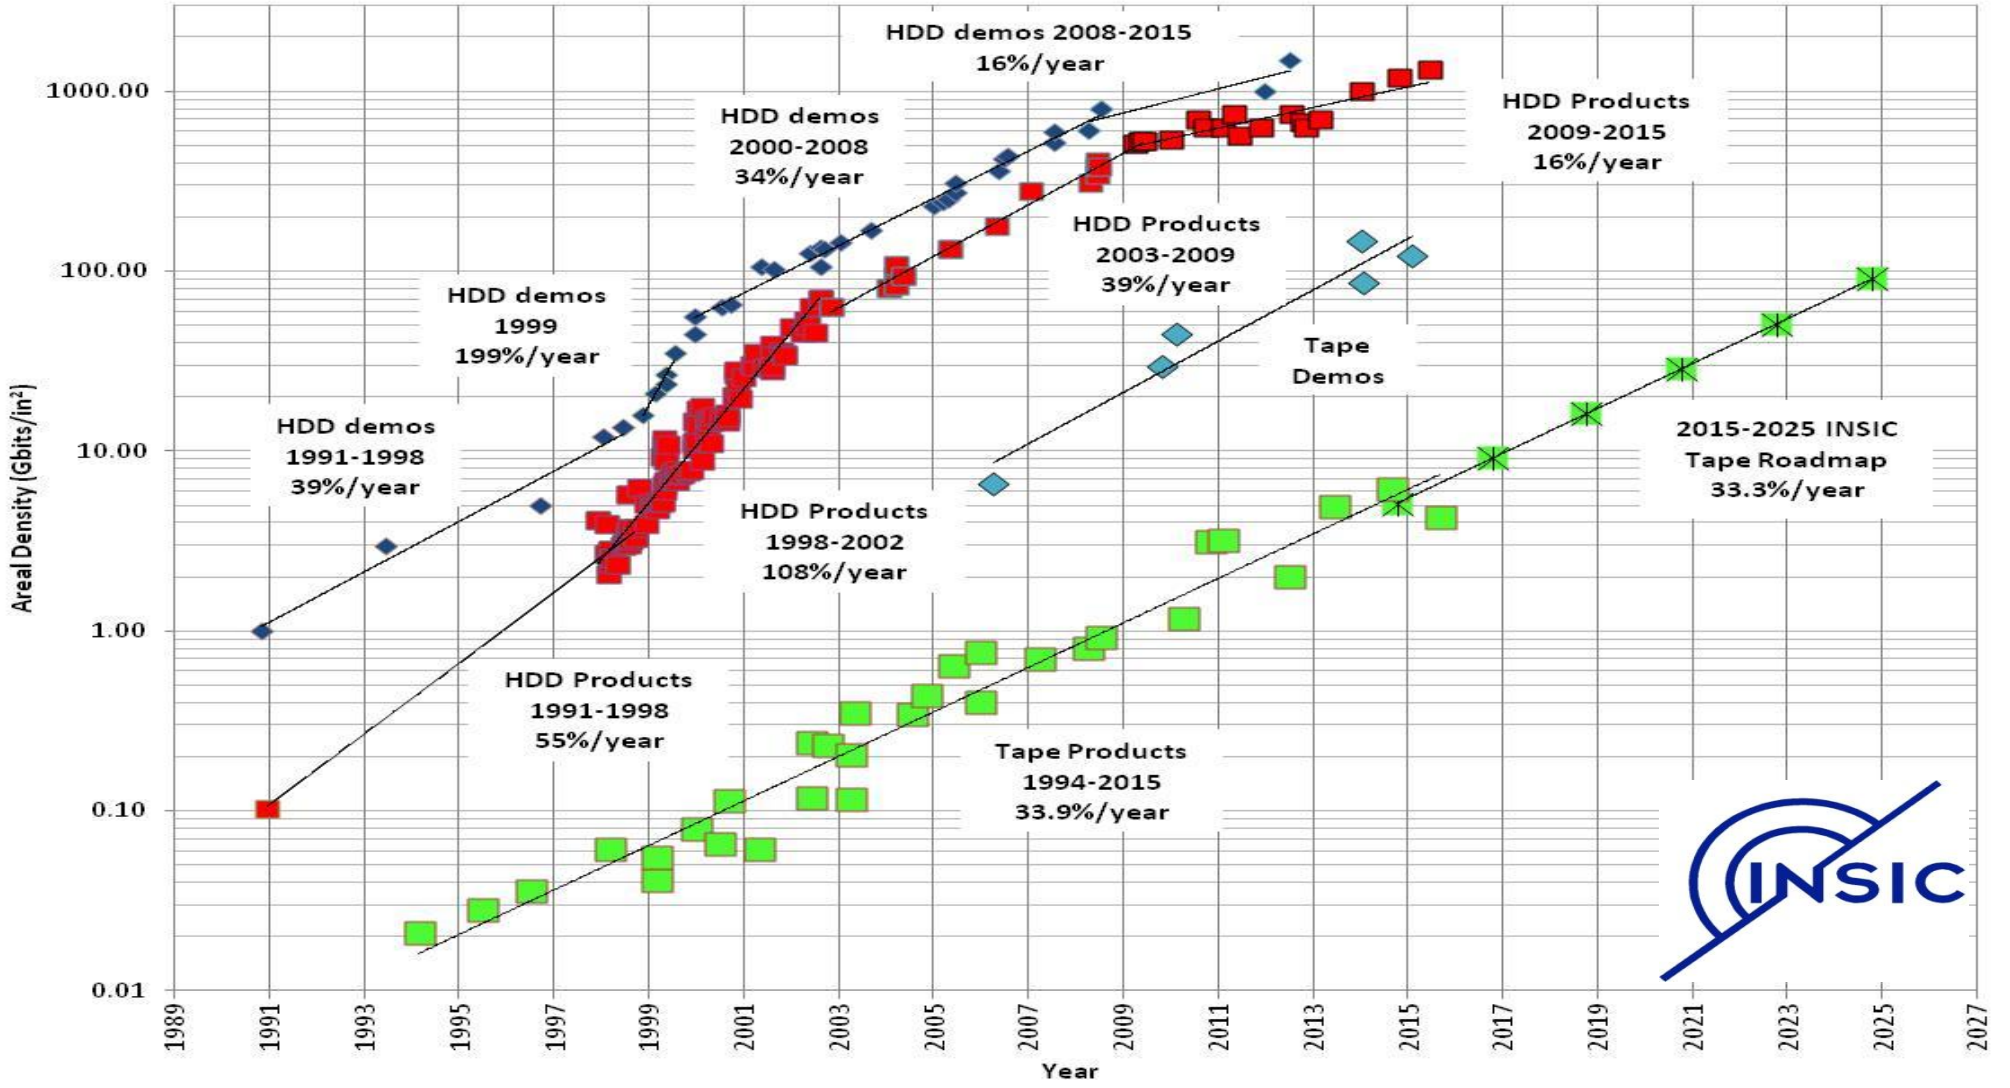

INSIC Shows Technology Path to 248 TB per Cartridge

Media Vendors Demonstrate Technology Needed to Achieve Roadmap Goals

Sputtered Media demo

Areal Density: 148 Gb/in^2

Cartridge Capacity: “185 TB”

Advanced BaFe demo

Areal Density: 123 Gb/in^2

Cartridge Capacity: “220 TB”

http://www.sony.net/SonyInfo/News/Press/201404/14-044E/

http://www.fujifilm.com/news/n150409_03.html

What Does it Take to **Increase Capacity**?

Tape Area	Linear Density	Track Density

All three variables impact cartridge capacity

What Does it Take to **Increase Capacity**?

Tape Area | Linear Density | Track Density

Cartridge Form Factor

Tape Thickness

Tape Length

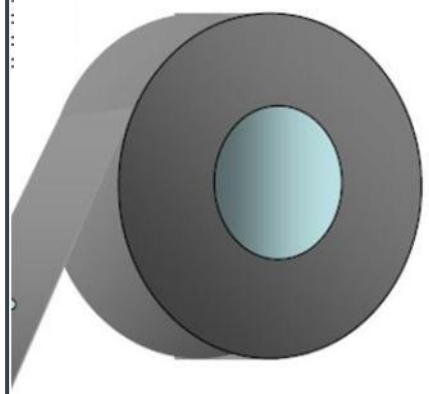

What Does it Take to **Increase Capacity**?

Tape Area | **Linear Density** | Track Density

Media

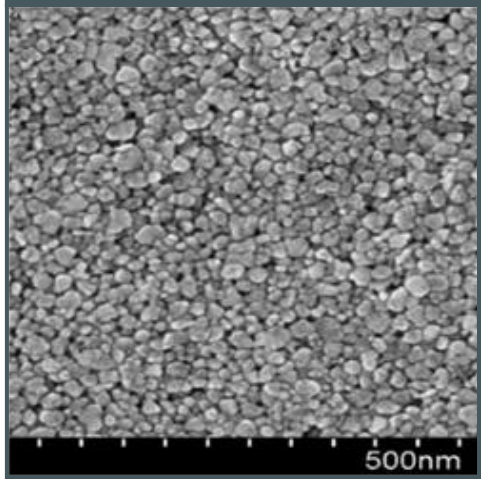

Read/Write Channel

Recording Head

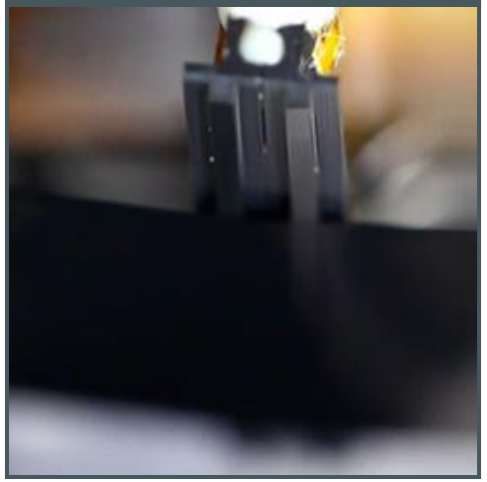

What Does it Take to **Increase Capacity**?

Tape Area | **Linear Density** | Track Density

Media | Read/Write Channel | Recording Head

- Smaller particles
- Improved magnetic properties
- Aligned particles (perpendicular orientation)
- Smoother, more uniform surface
- Higher signal to noise ratio

What Does it Take to **Increase Capacity**?

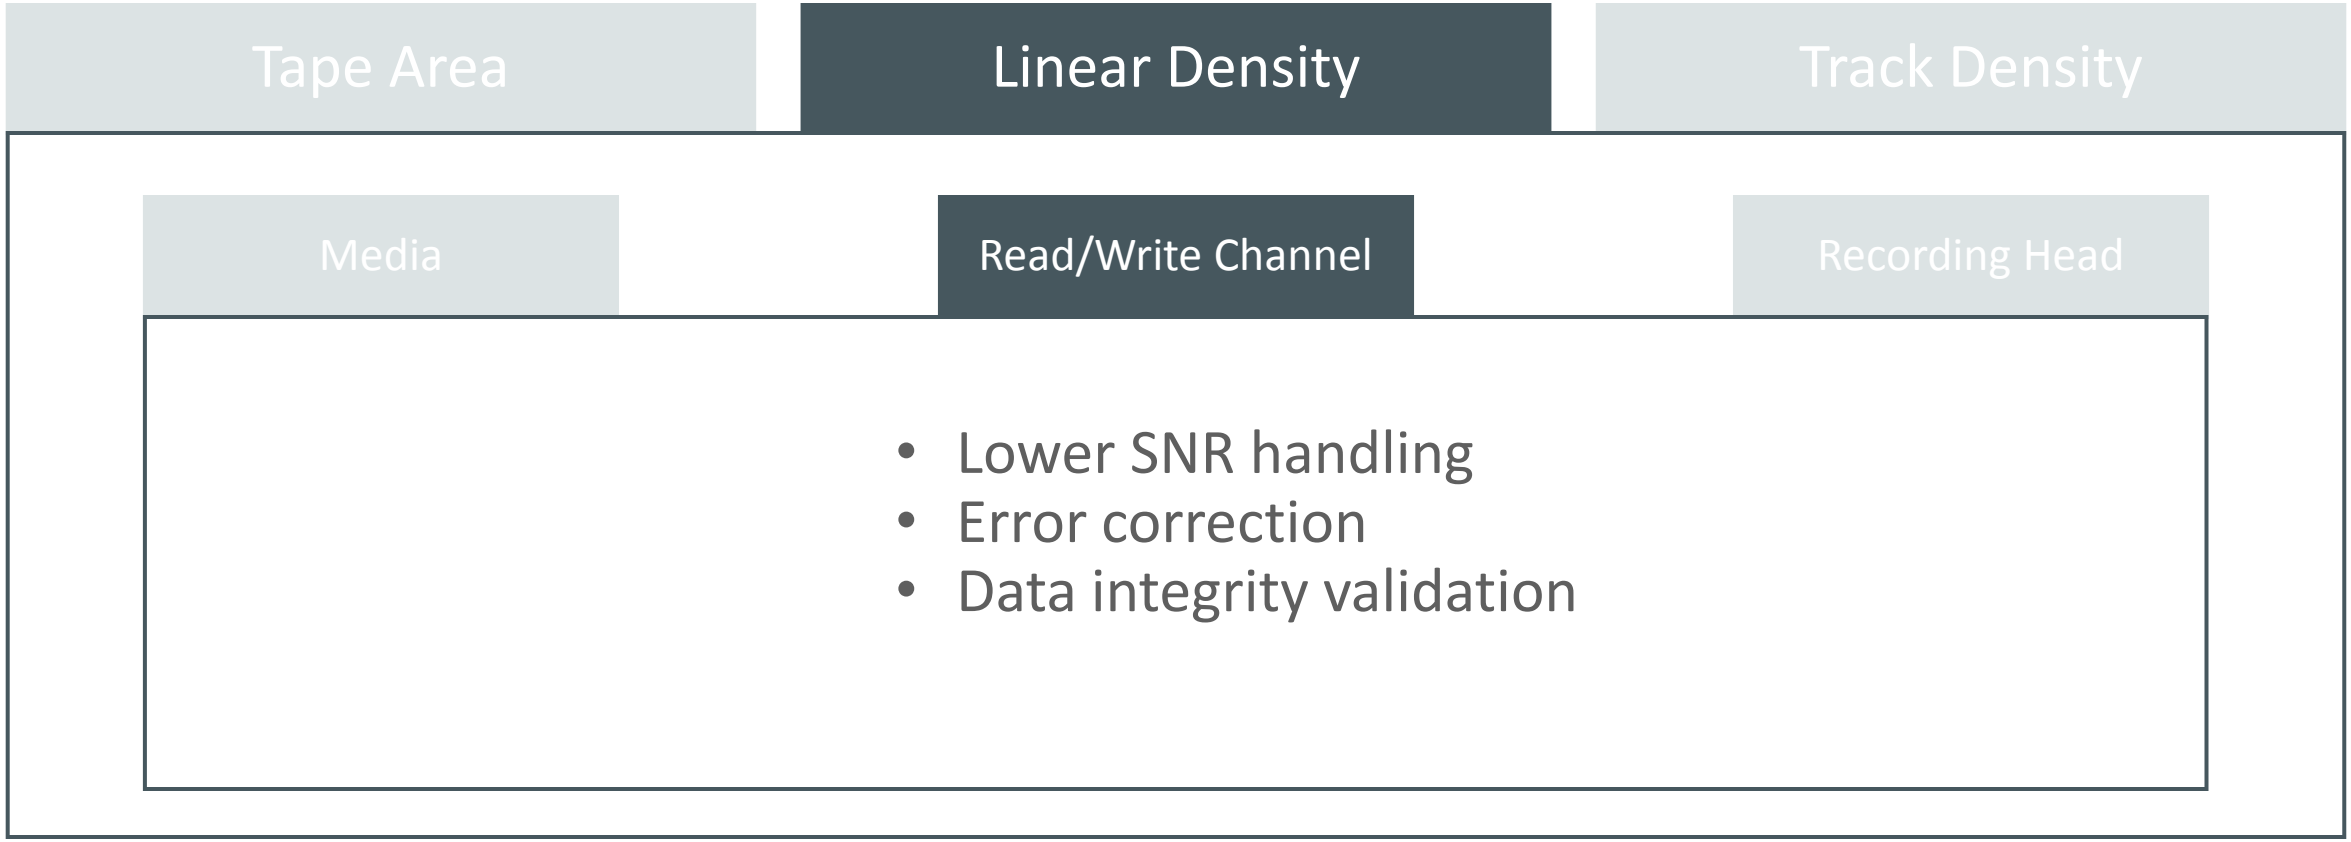

What Does it Take to **Increase Capacity**?

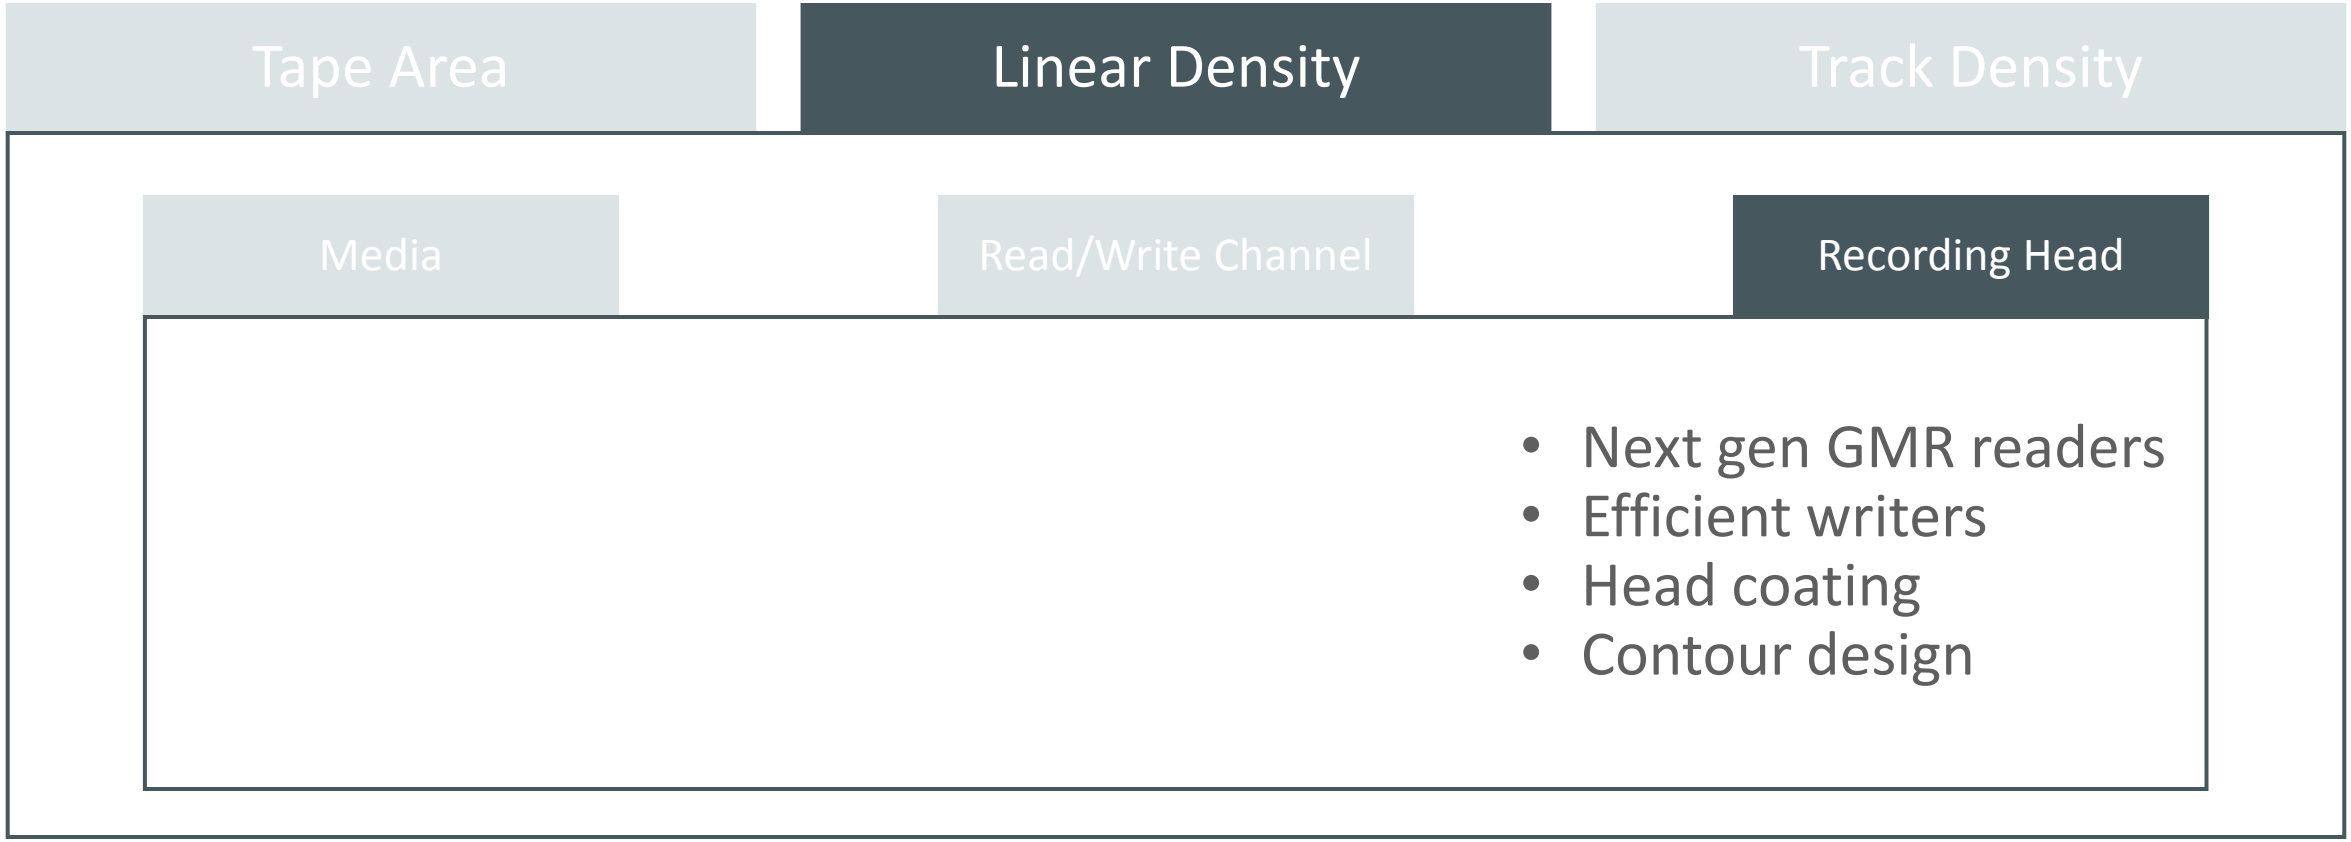

What Does it Take to **Increase Capacity**?

Tape Area | Linear Density | **Track Density**

Track Pitch

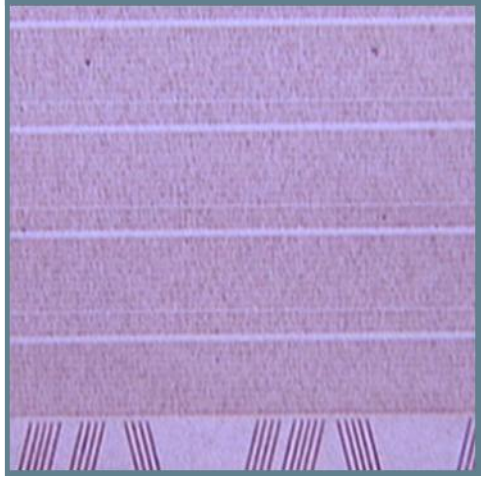

Track Following

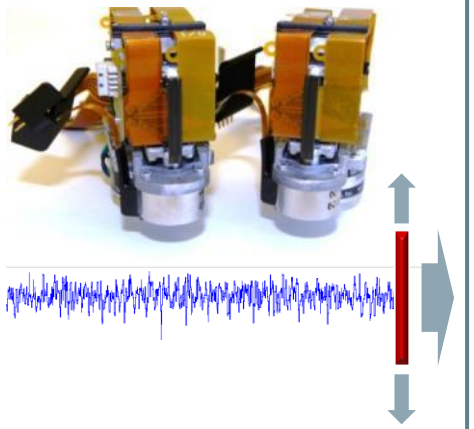

Basefilm Stability

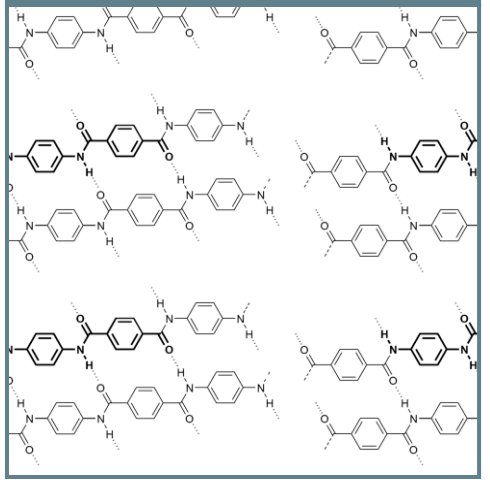

What Does it Take to **Increase Data Rate**?

Electronics	Data Format	Tape Velocity

All three variables impact data rate

What Does it Take to **Increase Data Rate**?

Electronics	Data Format	Tape Velocity

Integrated Circuits

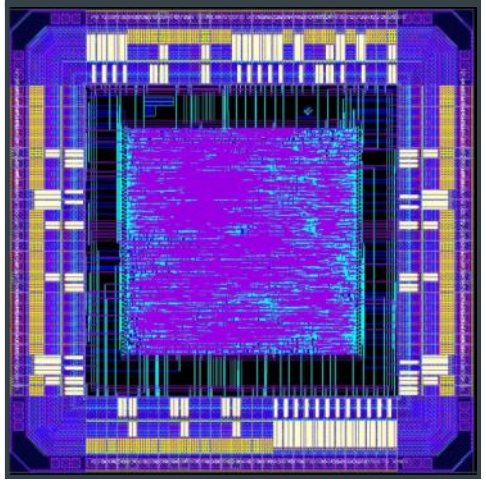

Printed Circuit Boards

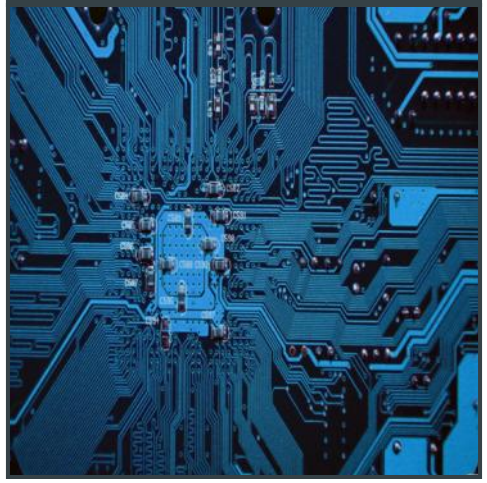

New digital architecture with ASICs integration

What Does it Take to **Increase Data Rate**?

Electronics | **Data Format** | Tape Velocity

Format Efficiency

Error Correction Code

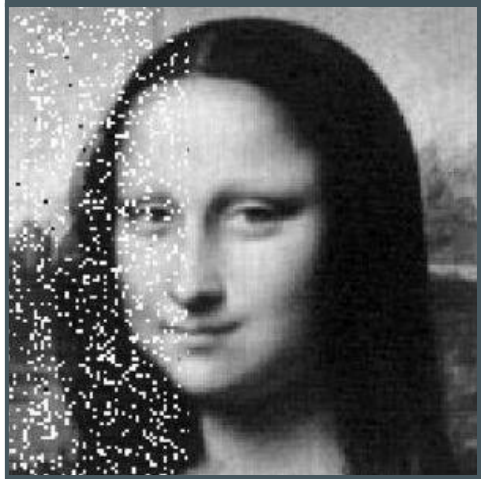

What Does it Take to **Increase Data Rate**?

Electronics	Data Format	Tape Velocity

Nanometer scale position control at high speeds

Servo Tracking

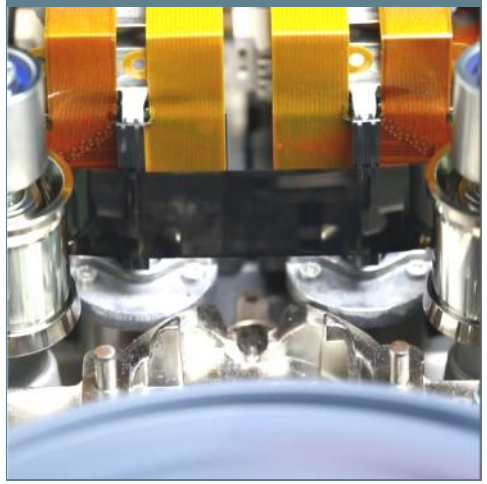

Tape Path

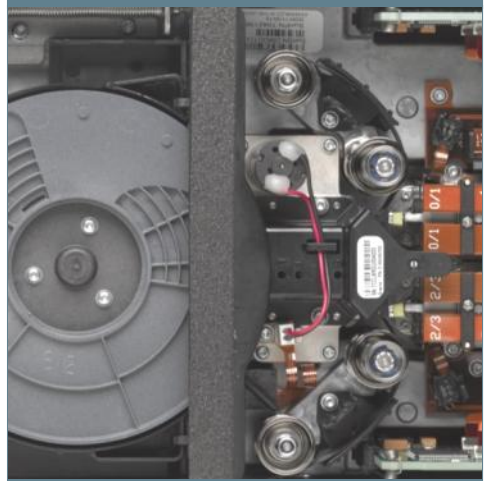

Focus Areas Moving Towards T10000E

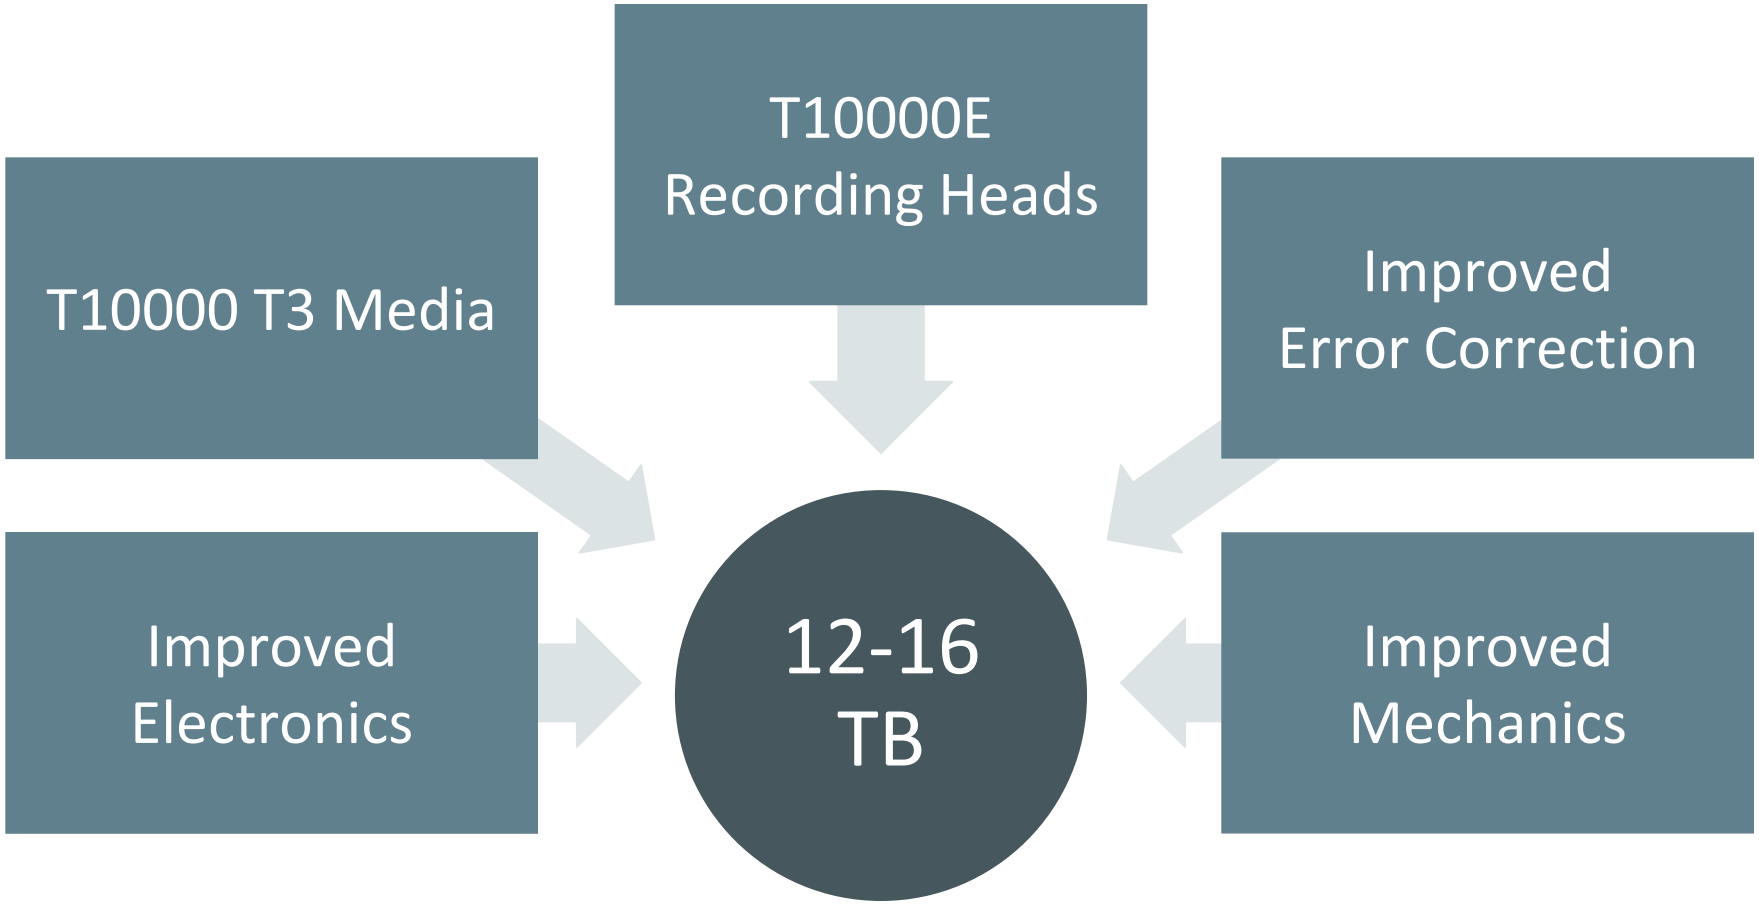

Need a StorageTek Roadmap Discussion?

Ask a member of the Oracle team to arrange a meeting.

Non-disclosure forms available, if needed.

Thank You!!

ORACLE®

STORAGETEK

ORACLE®

Integrated Cloud

Applications & Platform Services


ORACLE®