The Archive Solution Continuum

Chris Powers
Vice President, HPE Storage

The world is changing and accelerating

Big Data is no longer just a Buzzword – It's EVERYWHERE and growing …

Cloud

Big Data

Mobility

Volume

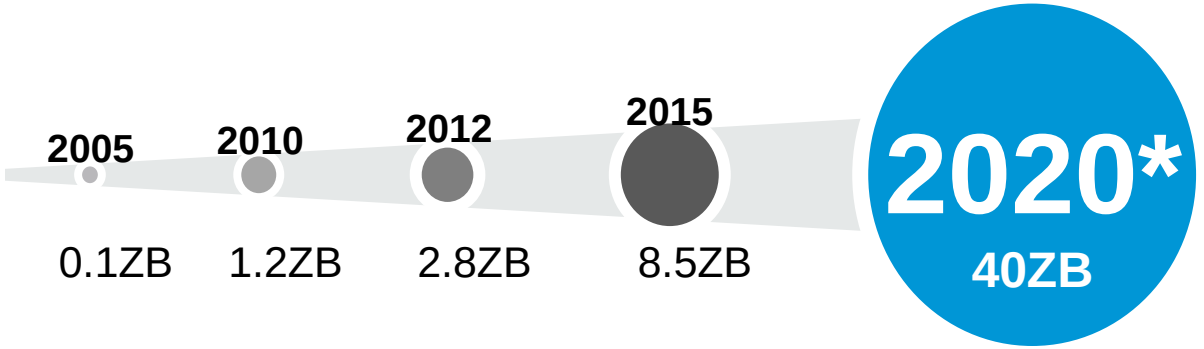

Variety

Velocity

IDC Estimates that by 2020, business transactions on the internet - business-to-business and business-to-consumer - will reach 450 billion per day.

Hewlett Packard Enterprise

*Source : IDC Digital Universe in 2020

Why archiving is essential

Data Growth

- Data is growing @ **40%** per year – yet IT budgets are relatively flat
- **90%** of the growth is **unstructured data** - growing faster than any other data type
- Putting pressure on backup window

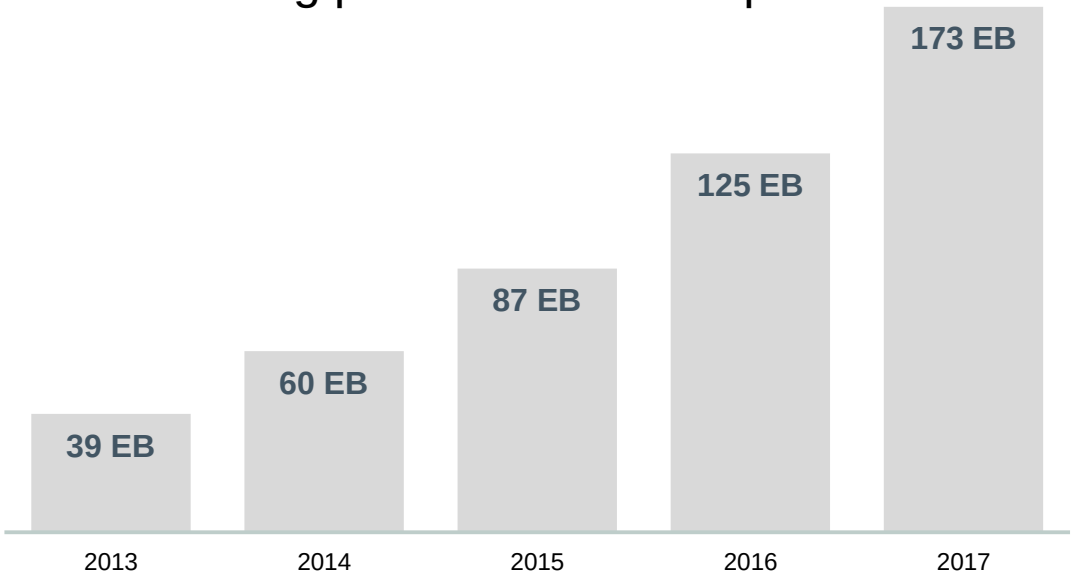

Source: Worldwide File- And Object-Based Storage 2013 – 2017 Forecast, IDC

Data Usage

- **80%** of unstructured data is rarely or never accessed after **90 days**
- But it still needs to be kept safe and available

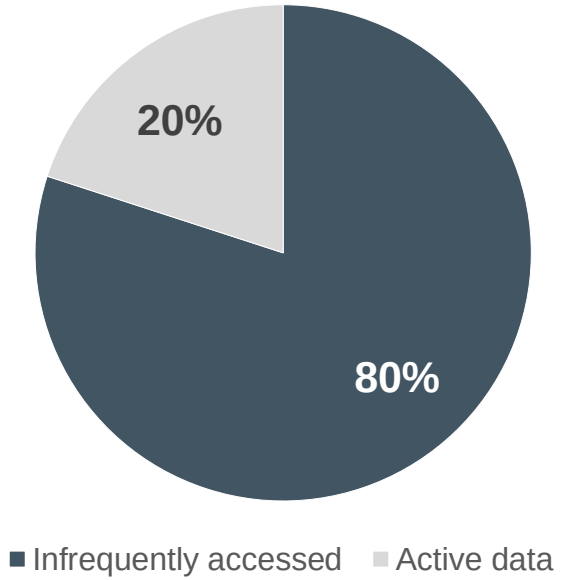

Source: National Energy Research Scientific Computing Center –

Hewlett Packard Enterprise

Why store and backup static, unstructured data on primary storage?

The benefits of archiving - Data center efficiencies

Lower Storage TCO

Reduce primary storage costs by removing unstructured data

Lower Backup TCO

Reduce backup costs/network load, meet backup window and improve SLAs by removing unstructured data from backup process

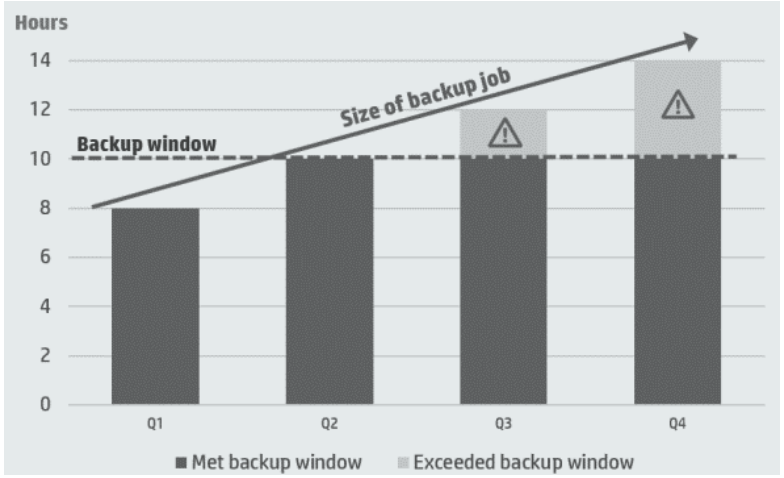

Increase Storage ROI

Retain, access and extract value from your data - longer for less

Archiving enables you to save more on your storage and do more with your data

Hewlett Packard Enterprise

What makes an effective archive solution?

Data access
Seamless access to data wherever it's stored – integrated with existing workflows – open standard

Scalability
Easily add capacity and scale to petabytes of storage

Lowest TCO
Reduce TCO by matching media type to SLA's by moving archival data to lower cost storage

Reliable long term retention
Reduce your risk with self protecting archive tier and workflow

Tiered Storage Alternatives

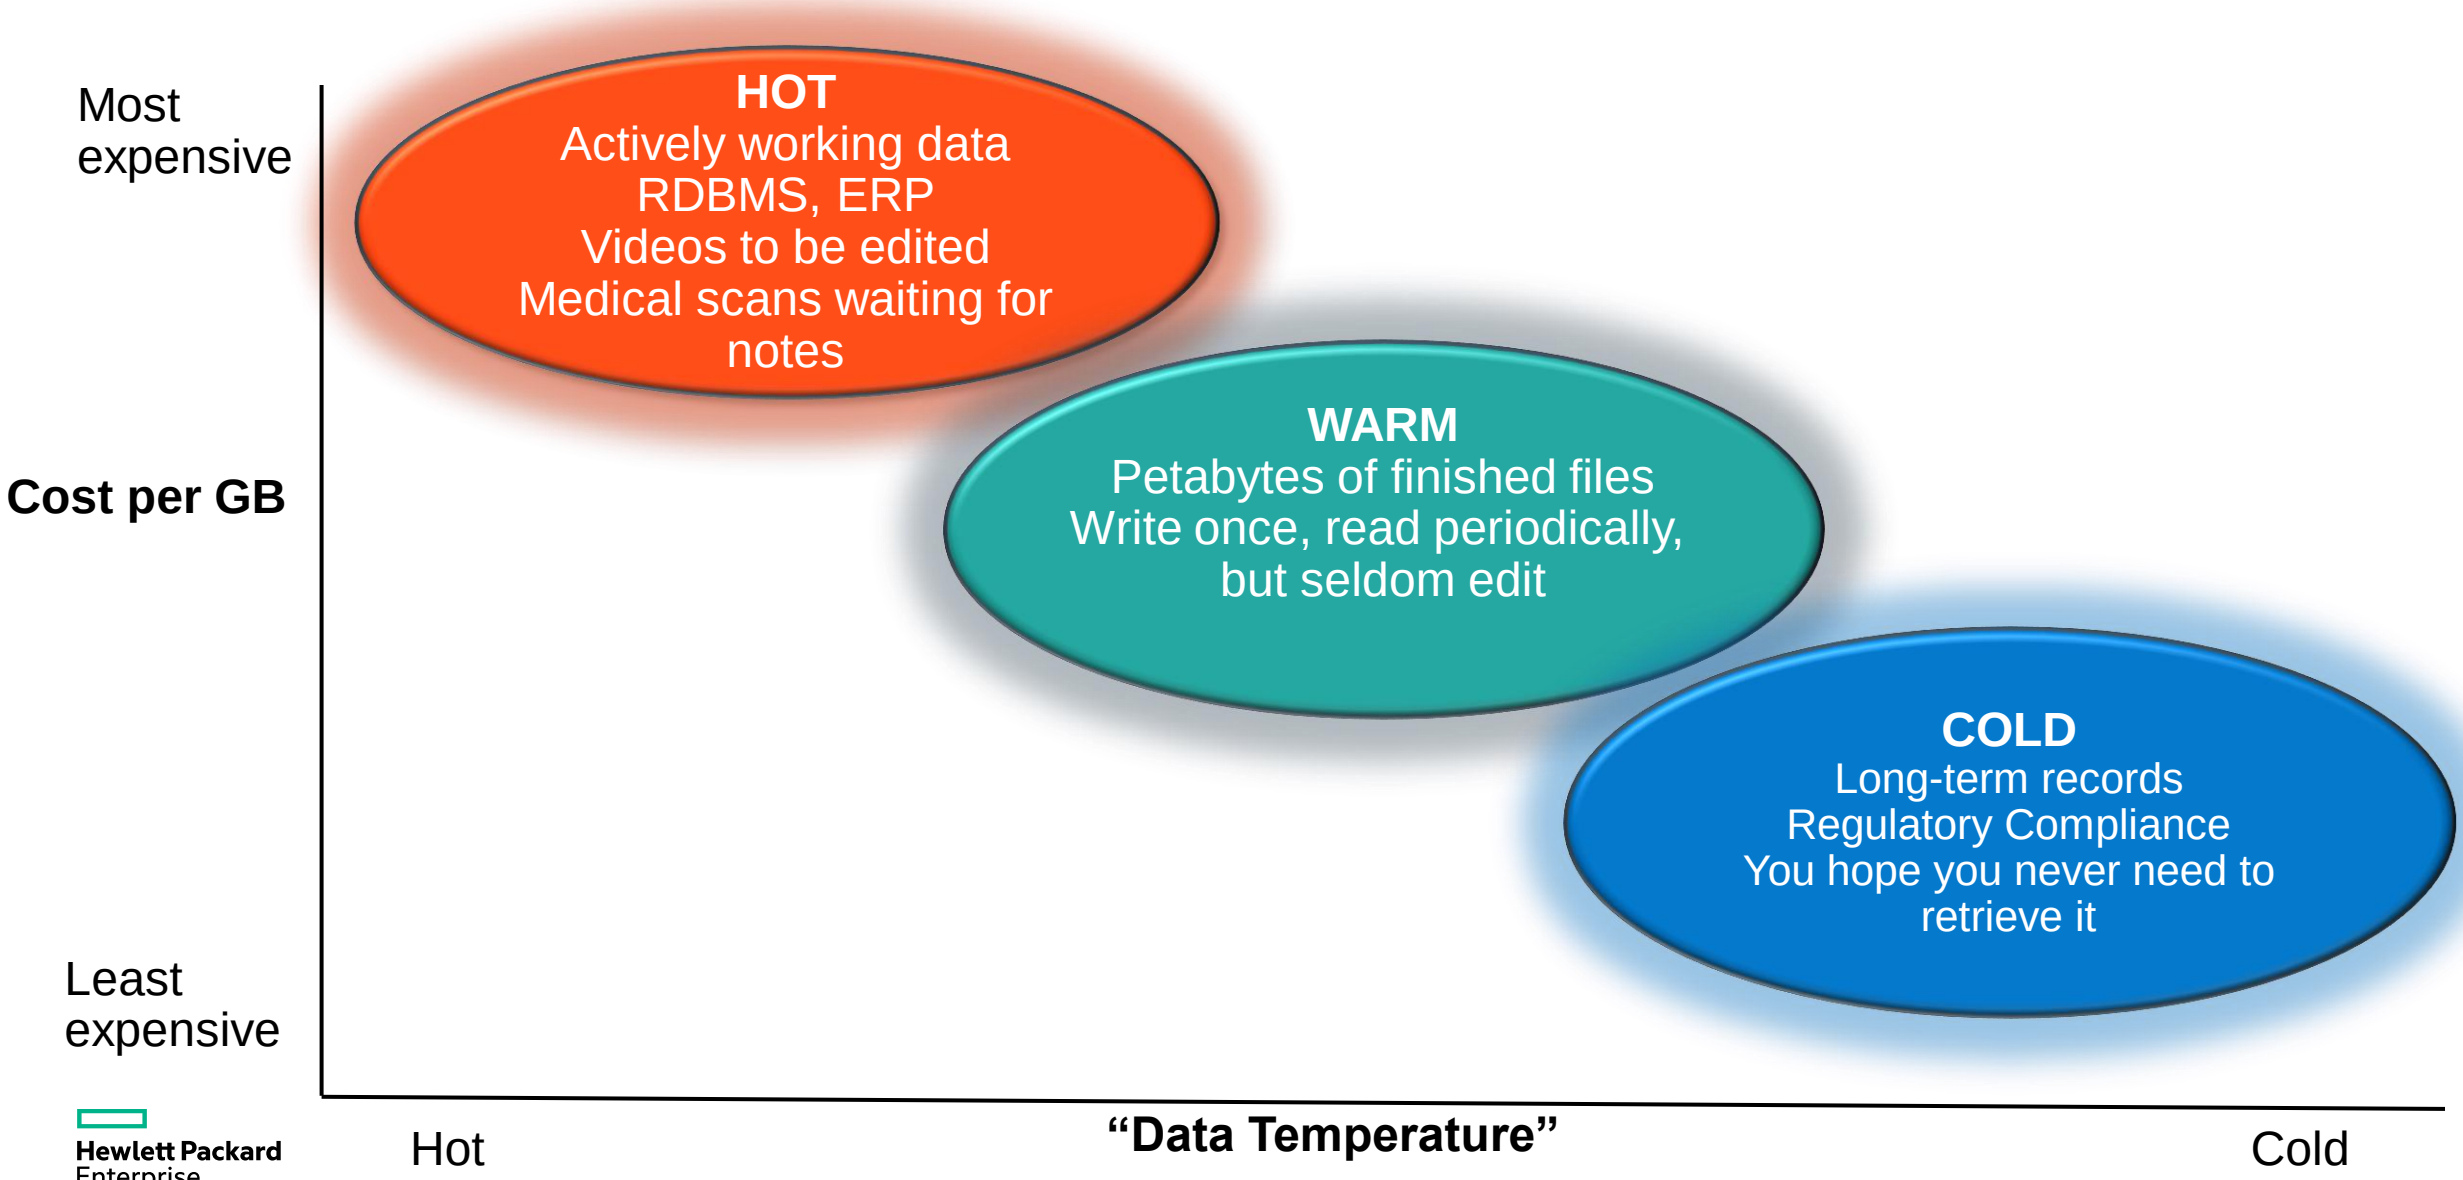

Hewlett Packard Enterprise

Warm: Object

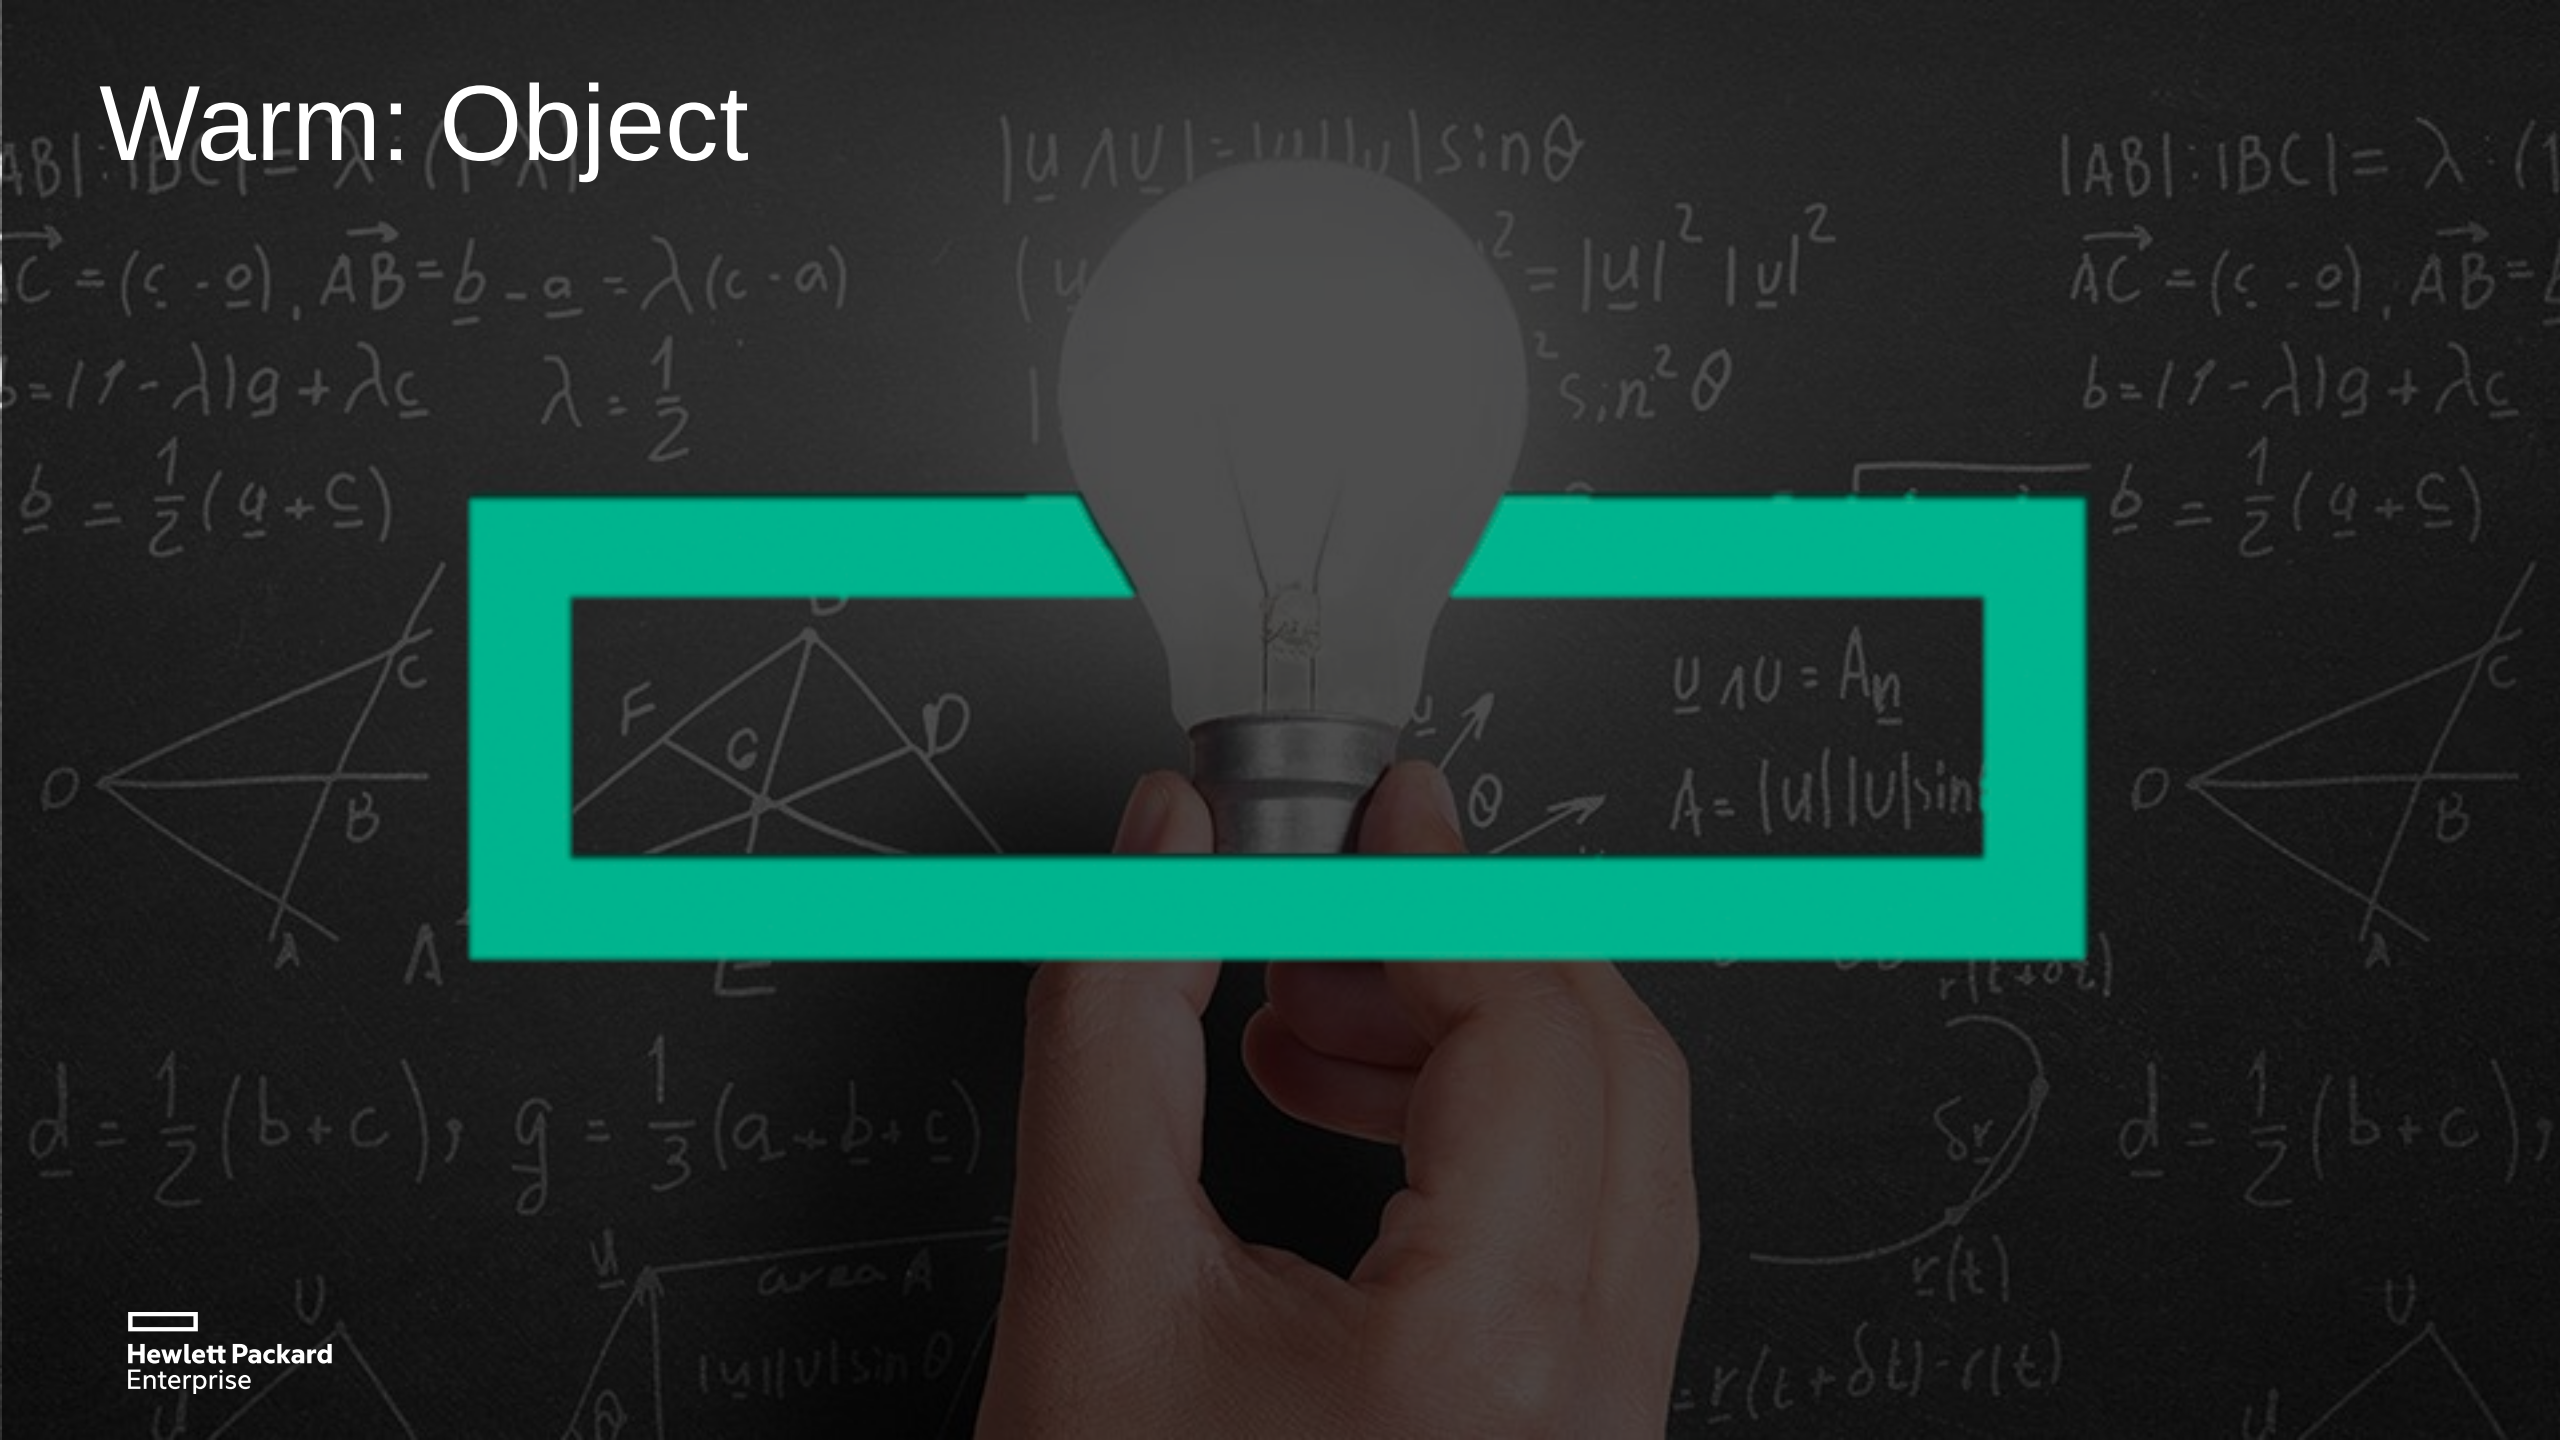

Hewlett Packard
Enterprise

Evolving feature requirements

Affordable. Hyperscale. Resilient. Autonomic

Affordable

- Storing huge quantities of data is only economic if you use the cheapest components available
- Operational costs must be minimized
- You can't afford to hold backups as well
- Preserve investment in existing file applications using a NAS to Object capability

Hyperscale

- Scale capacity from 500TB up to tens or even hundreds of PBs per solution
- Millions to Billions of objects per namespace
- Support high concurrency and throughput for access

Resilient

- Continue to operate despite multiple failed components
- Support multi data center configurations
- Allow swap out of failed components without service interruption

Autonomic

- Data durability of up to 11 x 9s
- Autonomic operation of data consistency and integrity
- Simplified configuration and provisioning

Demands a new type of storage

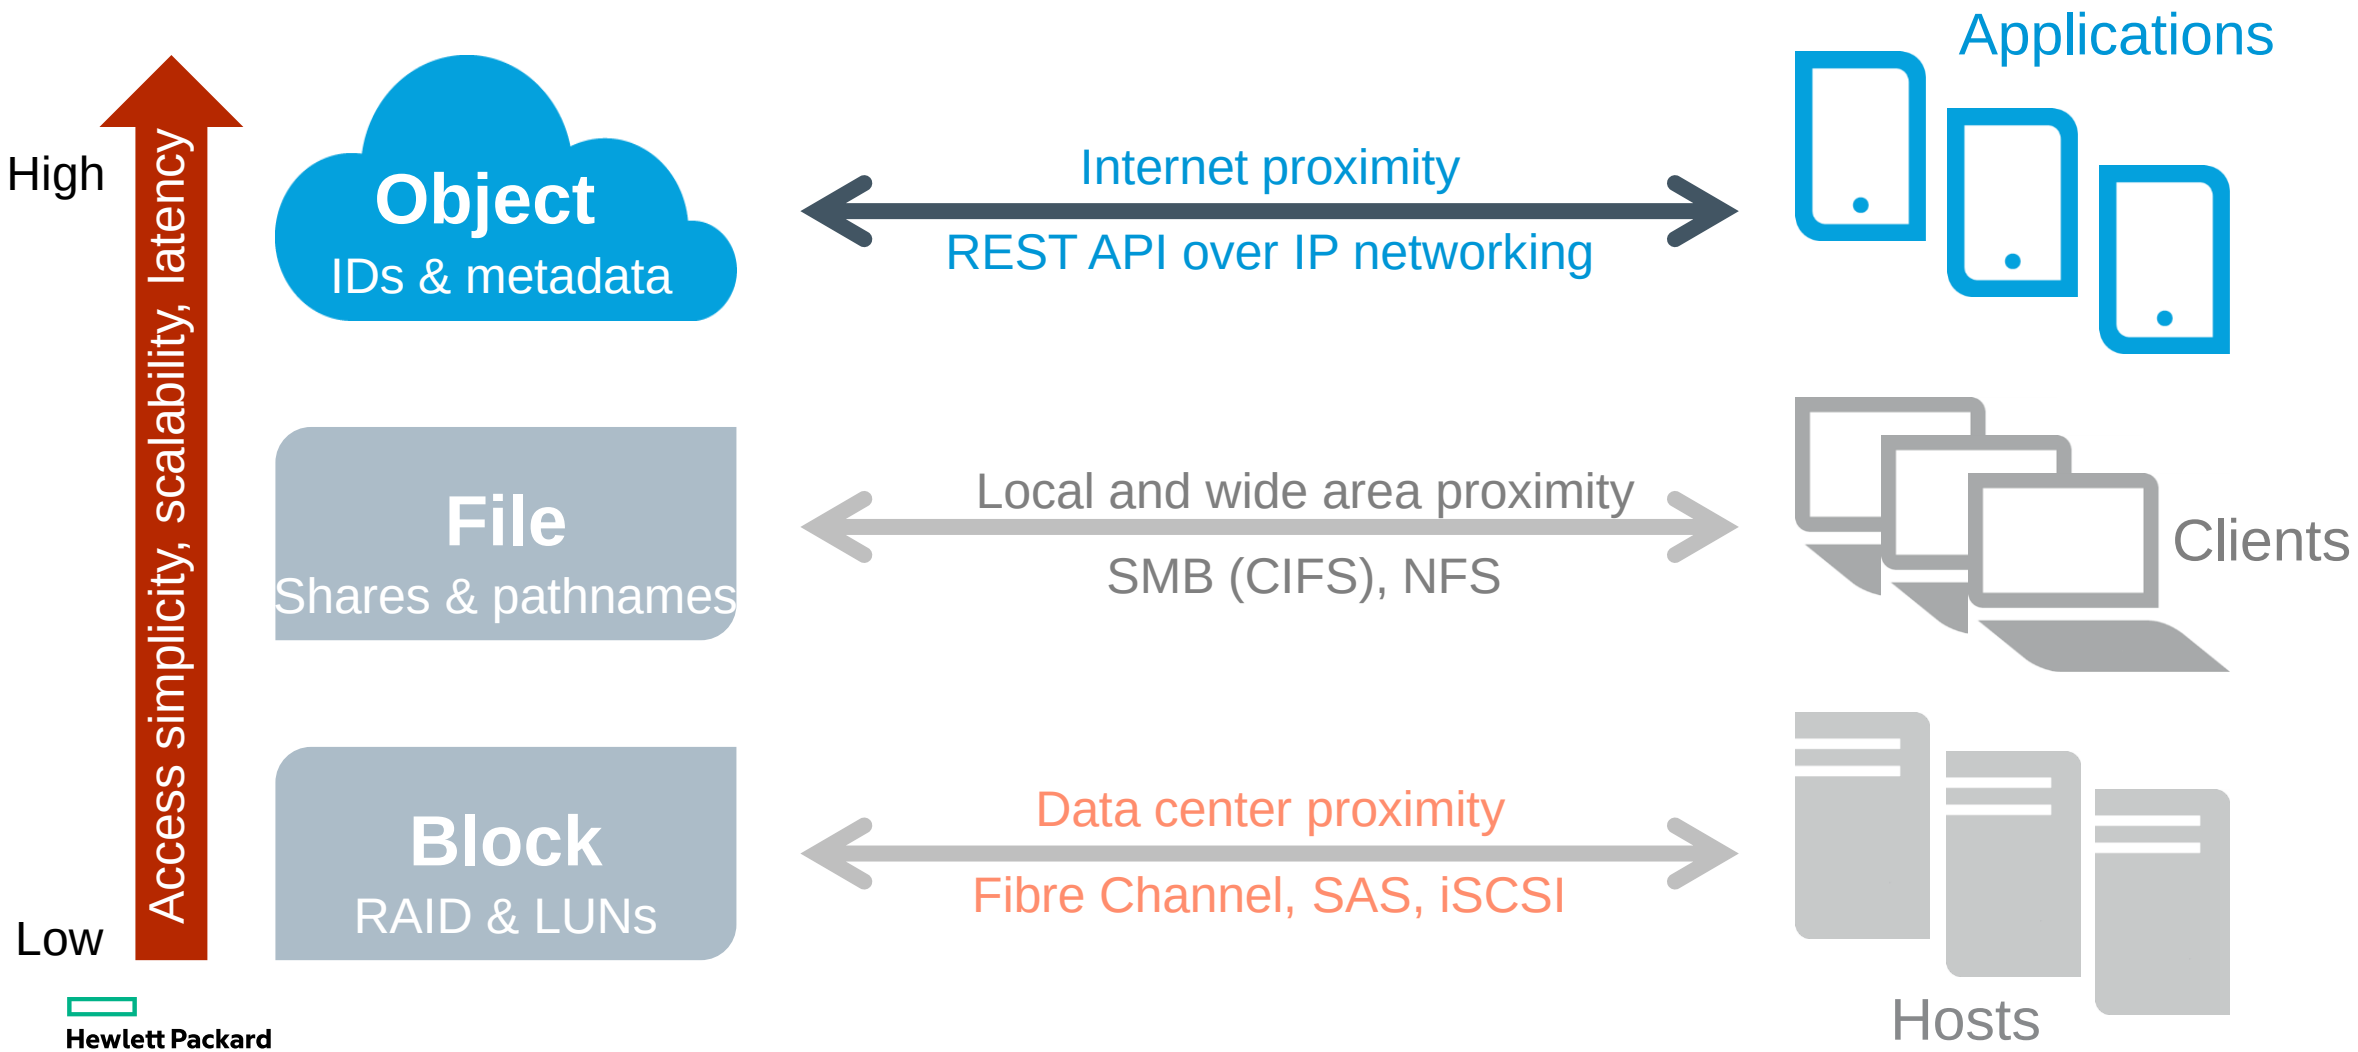

Hewlett Packard Enterprise

Is Object Storage the right choice?

Could you make more money and get more business insight if storing and retrieving your static data* were easy and inexpensive?	
Your current storage cannot handle the sheer quantity of modern Big Data	**That's what Object Storage was invented to address** Flat name-space, virtually unlimited scale, and built-in data protection, at reasonable costs.
You need a complimentary approach to long-term storage on tape in order to meet SLA requirements	Is the latency a key consideration for meeting your SLAs? What are your scaling consideration for the the data sizes you need to search and store?
You need answers from your static data *	Could you monetize that data, if it were more readily available? Do you need an 'active' archive solution? Is your traditional storage appliance causing performance issues at scale?
You need simpler storage for your customers or line-of-business	Would your customers pay for self-provisioned storage and on-line retrieval? Are your line-of-businesses using public clouds for file sync & share, for the convenience of quick provisioning?

* Static data is typically created once, then read frequently. Examples are video libraries, seismic data, or email and phone records. Object storage makes static data more accessible and less costly to store.

Object Storage Readiness

It is all about the Applications

- Applications access Object Storage via APIs
  - There are RESTful and Proprietary APIs available with various object storage environments
  - Verify applications support object storage APIs available with chosen object storage solution
  - Industry de facto standard RESTful APIs are Swift, S3 and simple HTTP
- Cloud Gateways for non-compliant applications
  - NFS/CIFS and block access to private and/or public object storage can be provided using cloud gateways
  - Cloud gateways can be physical appliances or software instances typically deployed as VMs
  - Gateways vary in application specialization
    - Global namespace for NFS/CIFS
    - Backup archive
    - Block storage access

Object Storage Solutions across major industries

Use Cases

- Professional video production, CDN
- Consumer/social video, images
- Security, surveillance, CCTV, body cams
- Satellite, Energy, Geological archives
- Enterprise backup / Enterprise archive
- Collaboration/document sync-n-share
- Big data analytics data tiering

Vertical Industries

Healthcare

- Genomics Data Hub
- Medical Imaging
- Medical Records
  - PACS
  - EMR – Electronic Medical Records
  - VNA – Vendor Neutral Achieve

Oil & Gas/ Manufacturing

- Seismic data archiving
- Production line, test & sensor data
- Autonomic car data

Service Providers

- Consumer/Business web services (Email, drop box, webapps)
- XaaS offerings

Financial Services

- Credit Card Transactional Data
- Check Imaging Backup

CME

- Content Delivery Networks
- Video Archiving
- CDR Archiving

Others

- Education
- Government

Hewlett Packard Enterprise

Cold: Tape

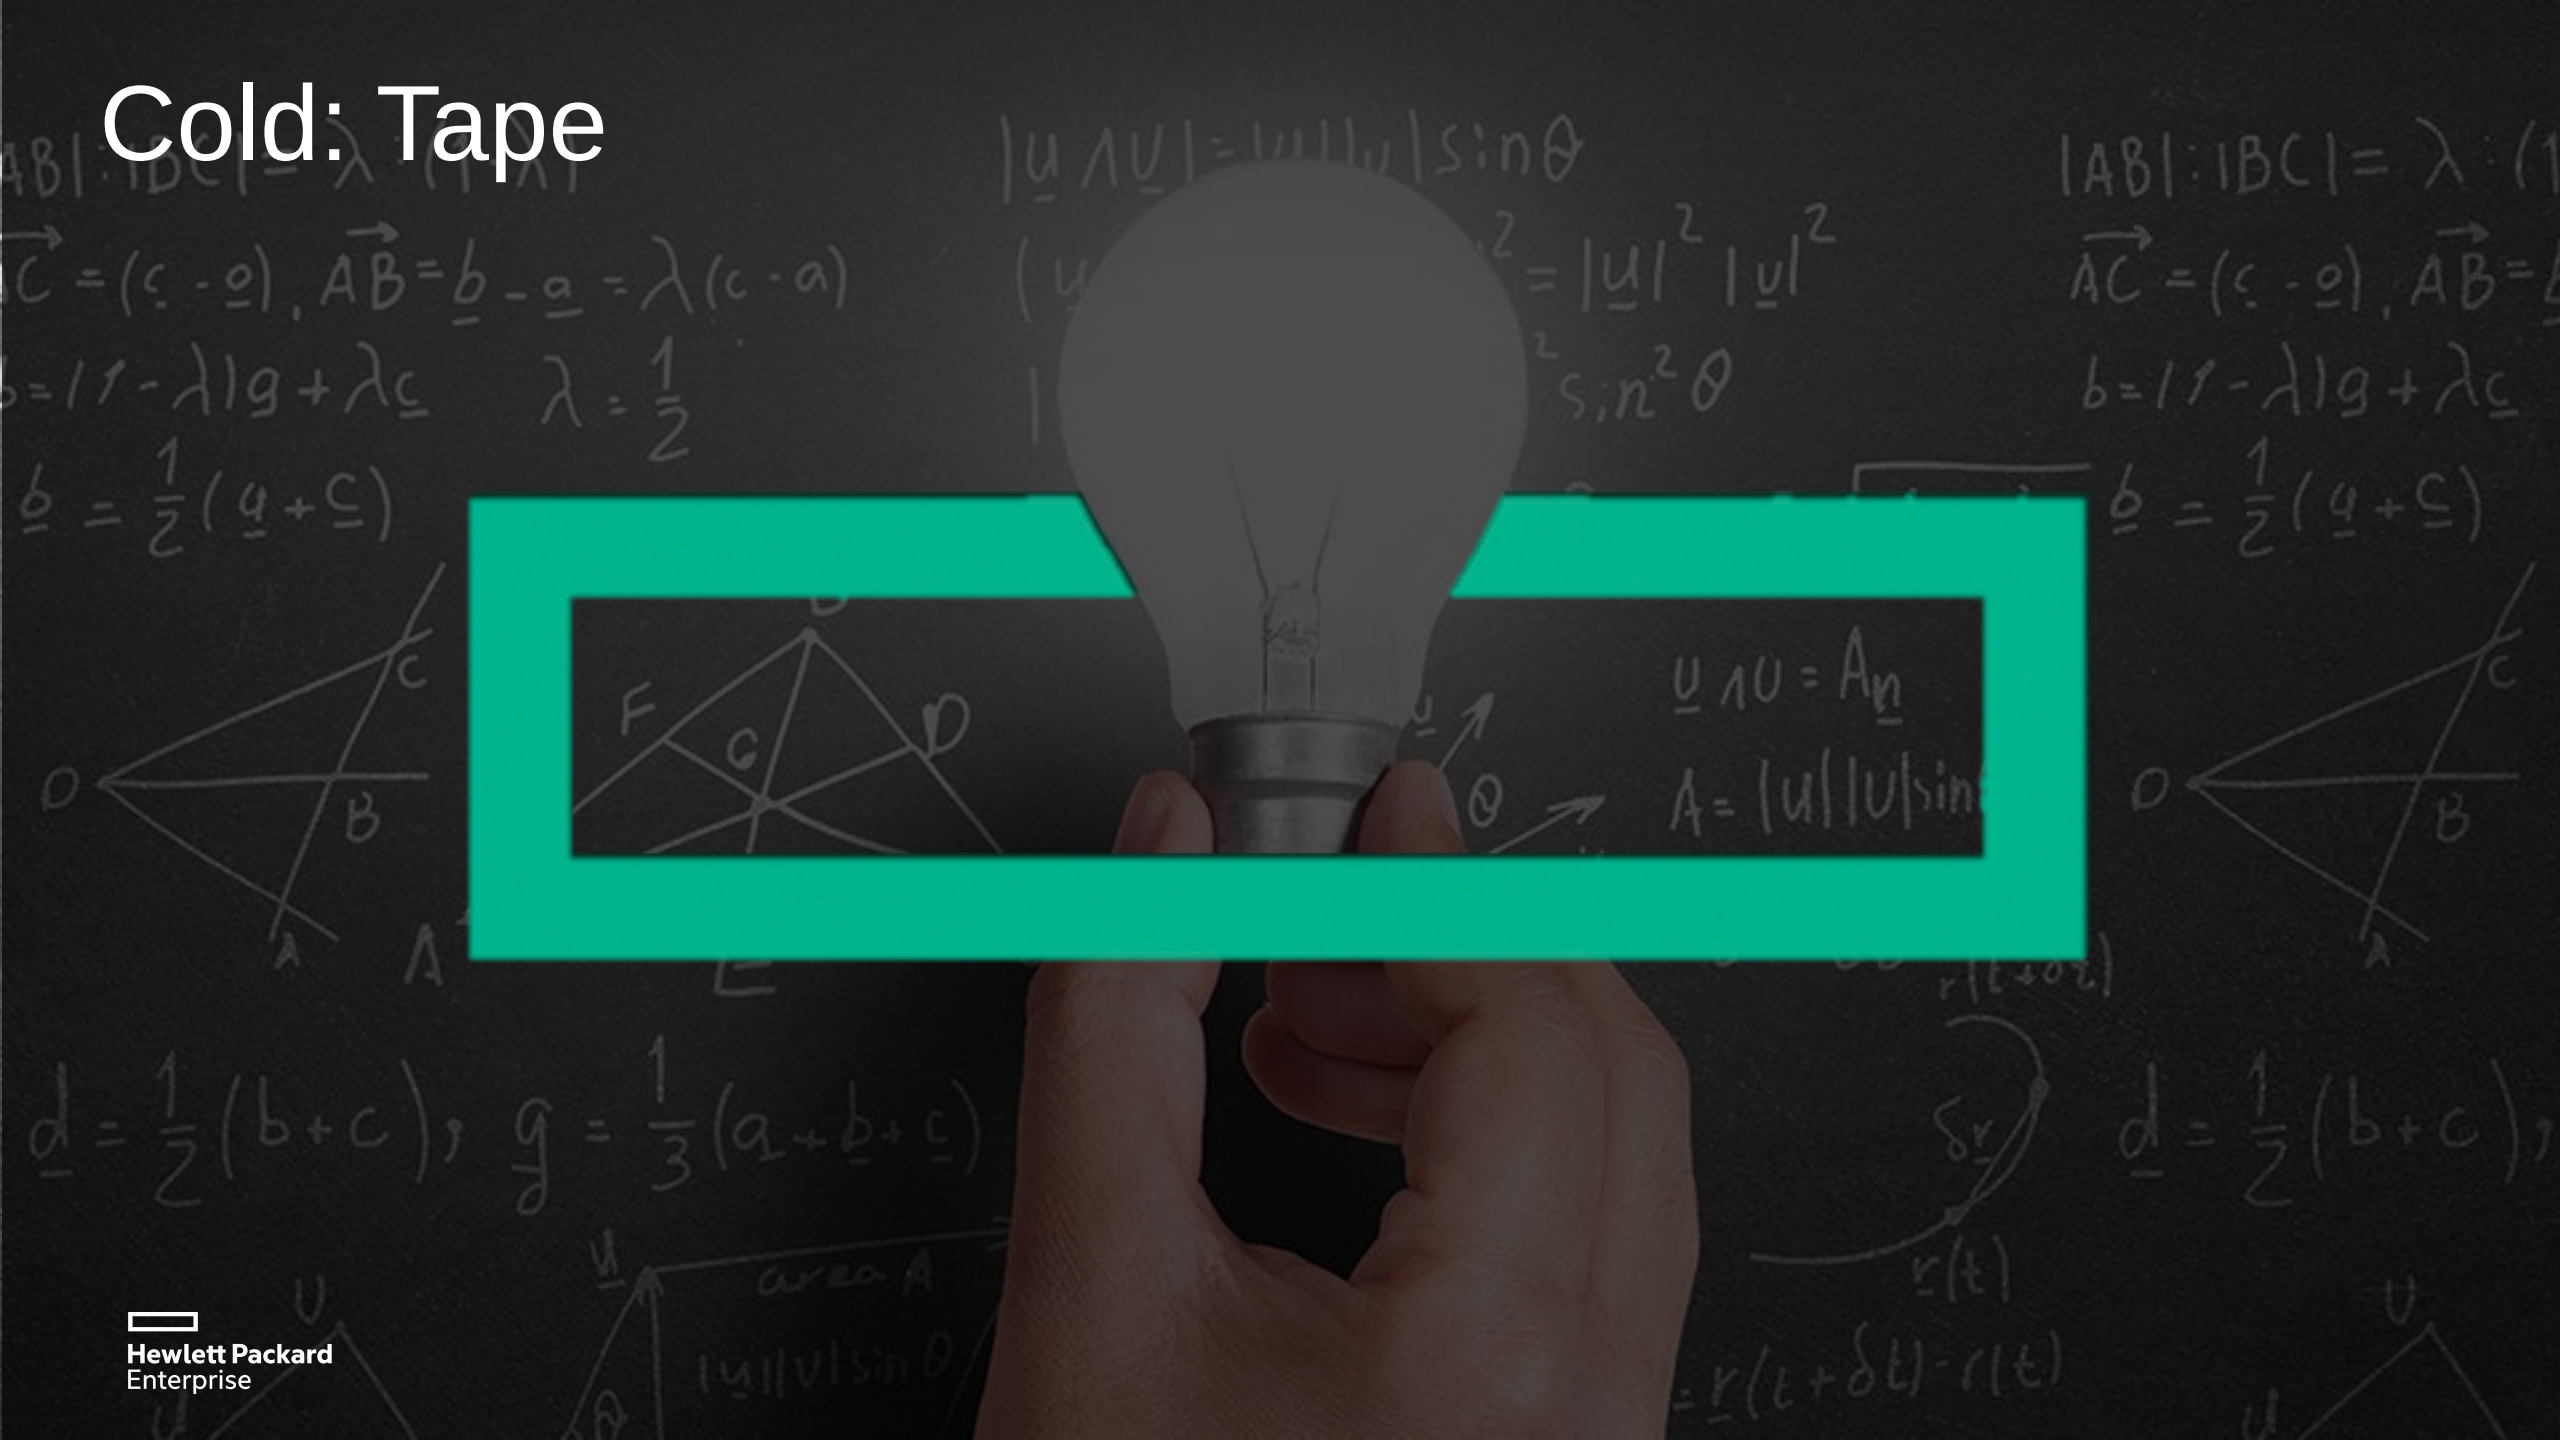

Hewlett Packard
Enterprise

Combining disk accessibility with economics and retention of tape

Efficient

Save time with direct and seamless access to archive data

Scalable

Manage data growth with non-disruptive, scale-out performance and capacity

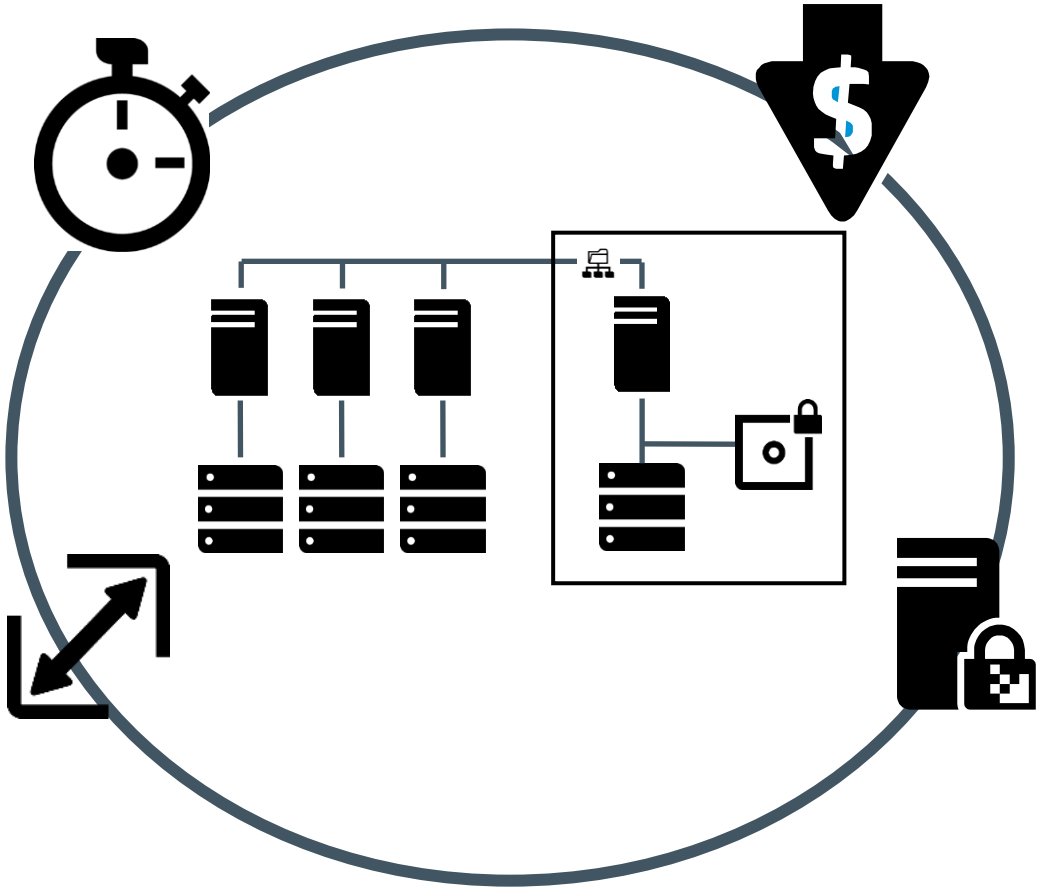

Economic

Reduce your costs with lowest TCO for long term data retention

Secure

Lower your risk with secure and dependable data availability

Hewlett Packard Enterprise

Save time with direct and seamless access to archive data

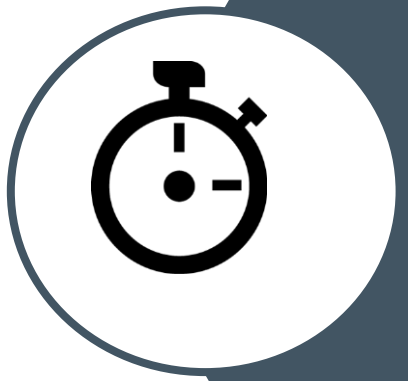

Easy file access

- Users and applications can directly access archived files
- No media management, media is self-allocated
- Data stored and accessed from standard subdirectories available through Explorer

Works with existing applications and file system tools

- Archive files are directly available to existing applications and file system utilities.
- Access to archive data and file searches no longer requires proprietary backup applications!

Leverages LTFS Standard

- Stores data in the Linear Tape File System (LTFS) open standard, supporting data transportability and long-term data accessibility.

Keep your data protected and available. Always.

Security

- Safeguard data with drive encryption, key management and WORM
- Automated retention management to meet governance and compliance requirements
- Off-line data protection provides last line of defence versus virus attacks, natural disasters or data corruption.

Reliability

- Media archival life of 30 years - Better bit error rate than disk!
- Integrated media verification software to analyze cartridges and take remedial action

Availability

- Automatically create additional media copies to be stored off-site for cost-effective archive disaster protection
- Optional real-time replication of data between two or more archive stores, locally or globally.
- Higher data availability - data always exists in two or more places!

Combining flash accessibility with economics & retention of tape

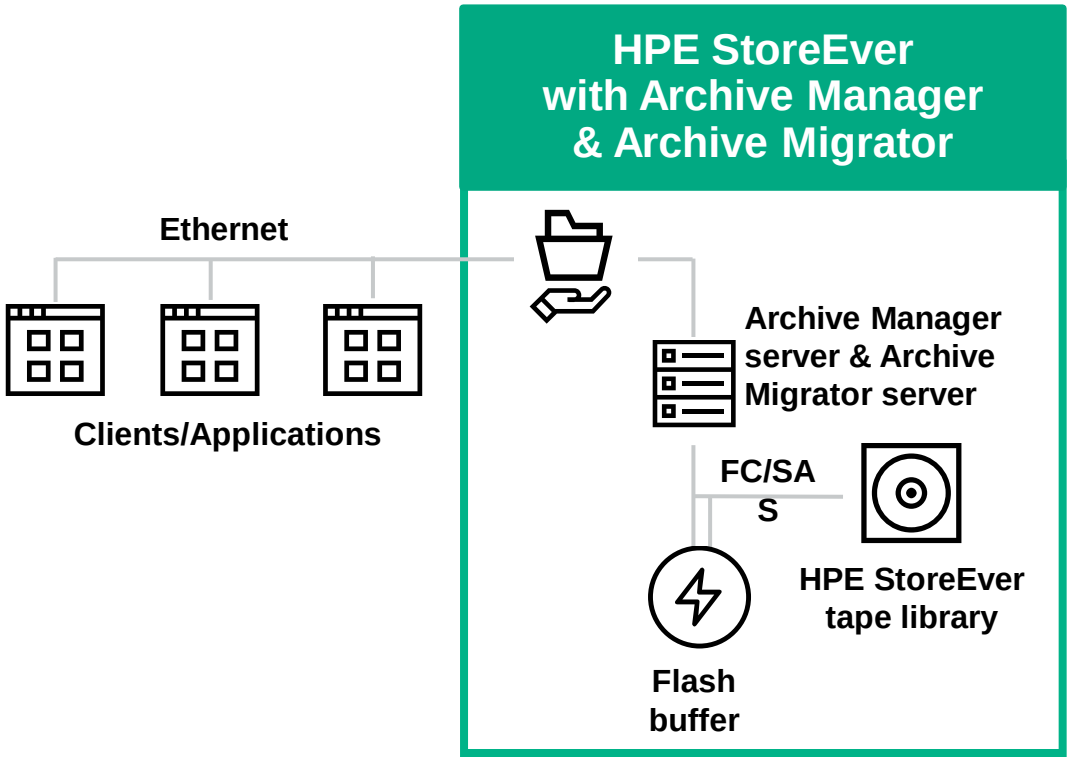

Data migrated from primary storage, either **manually** or **automatically**, by HPE StoreEver Archive Migrator

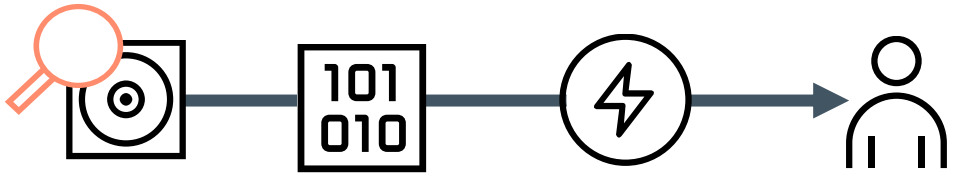

Archived files appear in searches on tape until users request access, then data is moved to StoreEver Archive Manager flash buffer, then to user

Reduce your costs with lowest TCO for long term data retention

Example: storing and protecting 200TB of static data

200 TB Static Content	Purchase more primary storage + backup storage	HPE StoreEver Archive Manager
Primary Storage	$200,000	
Backup Storage	$100,000	
Backup Software	$150,000	
DR Storage	$15,000	Included
Data Mover Software		$30,000
Archive Storage		$140,000
Total	$465,000	$130,000
Amount available for more storage		**$335,000**

Extending to cloud

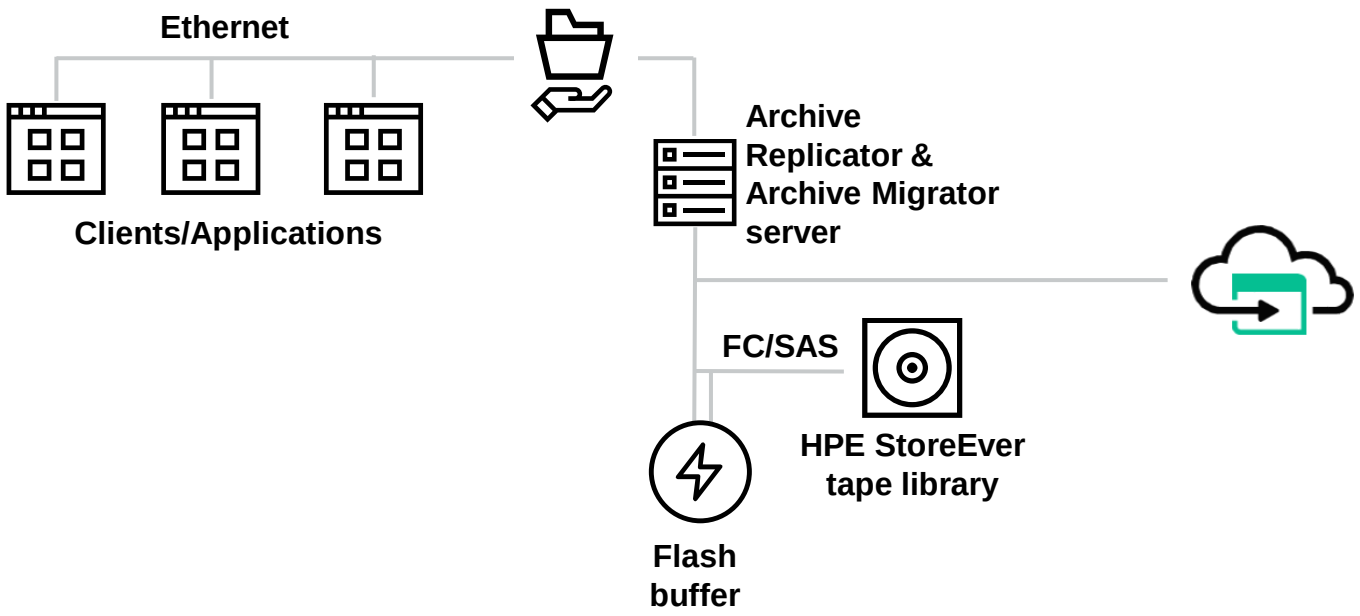

Data migrated from primary storage, either **manually** or **automatically**, replicated to remote cloud

Files asynchronously copied using appropriate API functions provided by respective clouds

File Metadata retained at primary site for simple searches

Extending to object

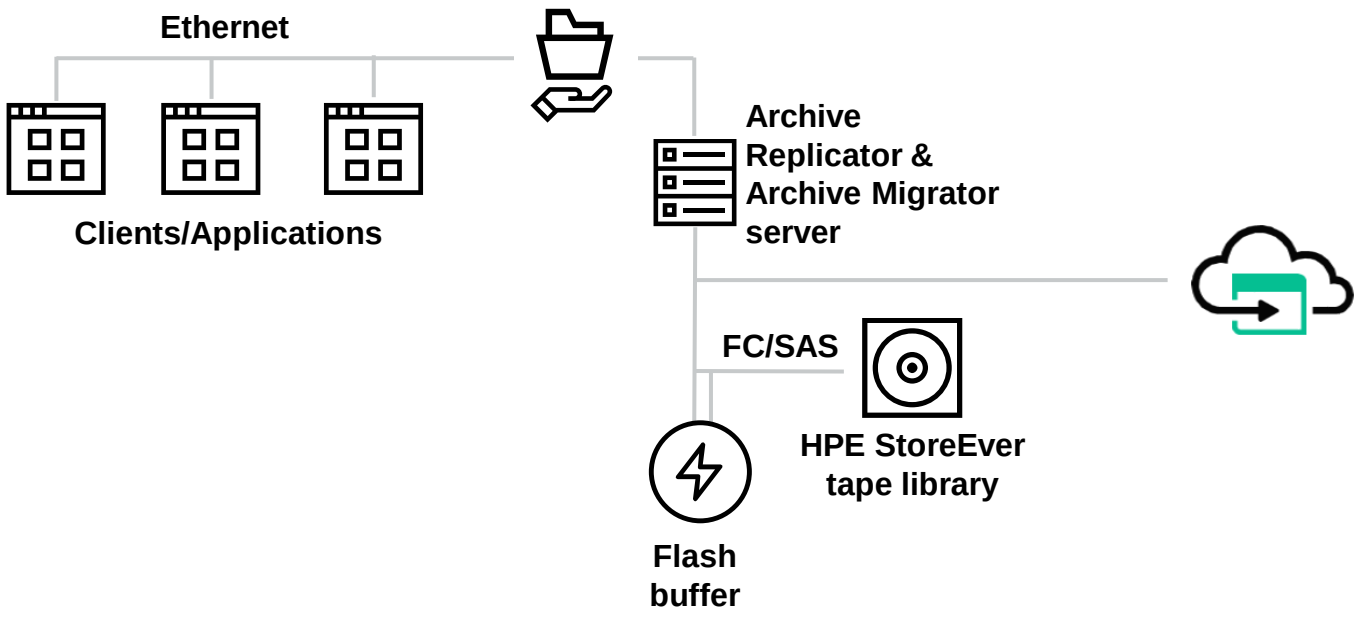

Data migrated from primary storage, either **manually** or **automatically**, replicated to object target

Files asynchronously copied using appropriate API functions provided by respective object targets

File Metadata retained at primary site for simple searches

Hewlett Packard
Enterprise

Thank You