

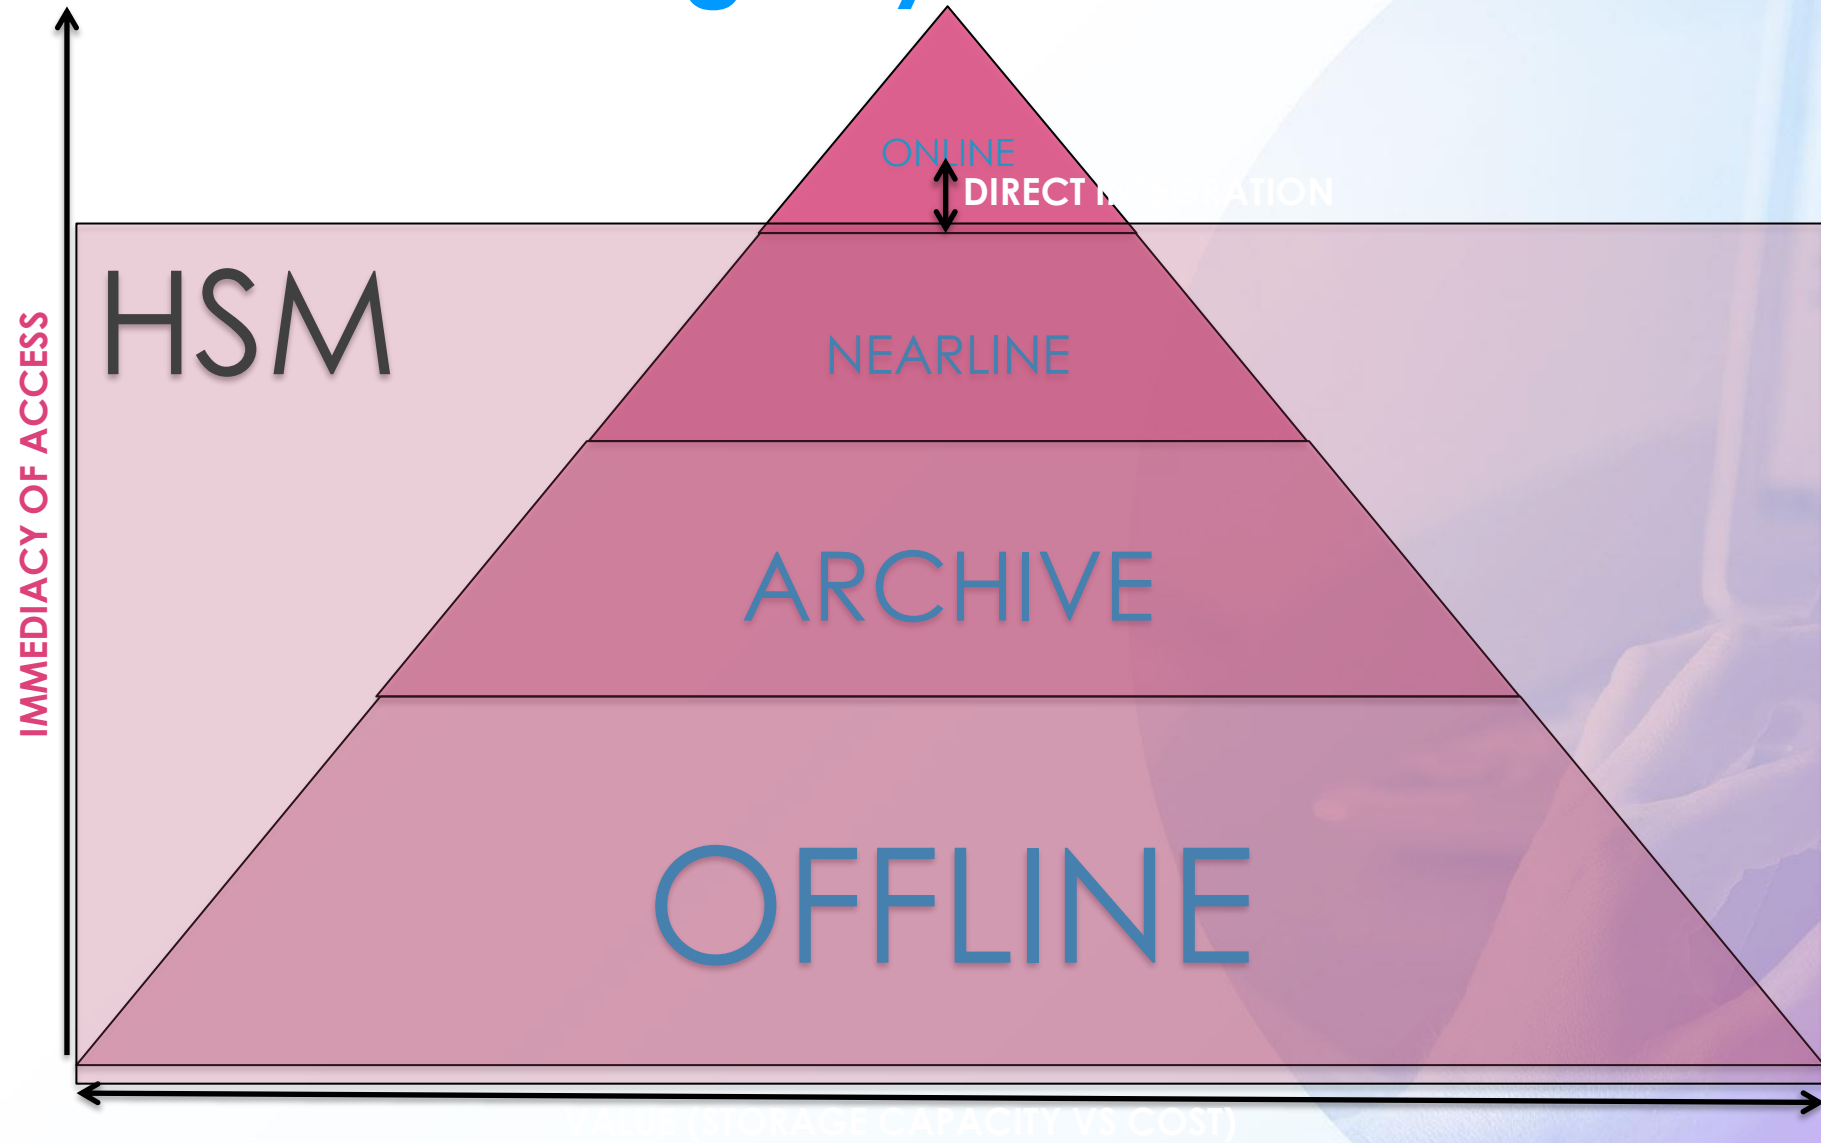
Building Media in the Cloud



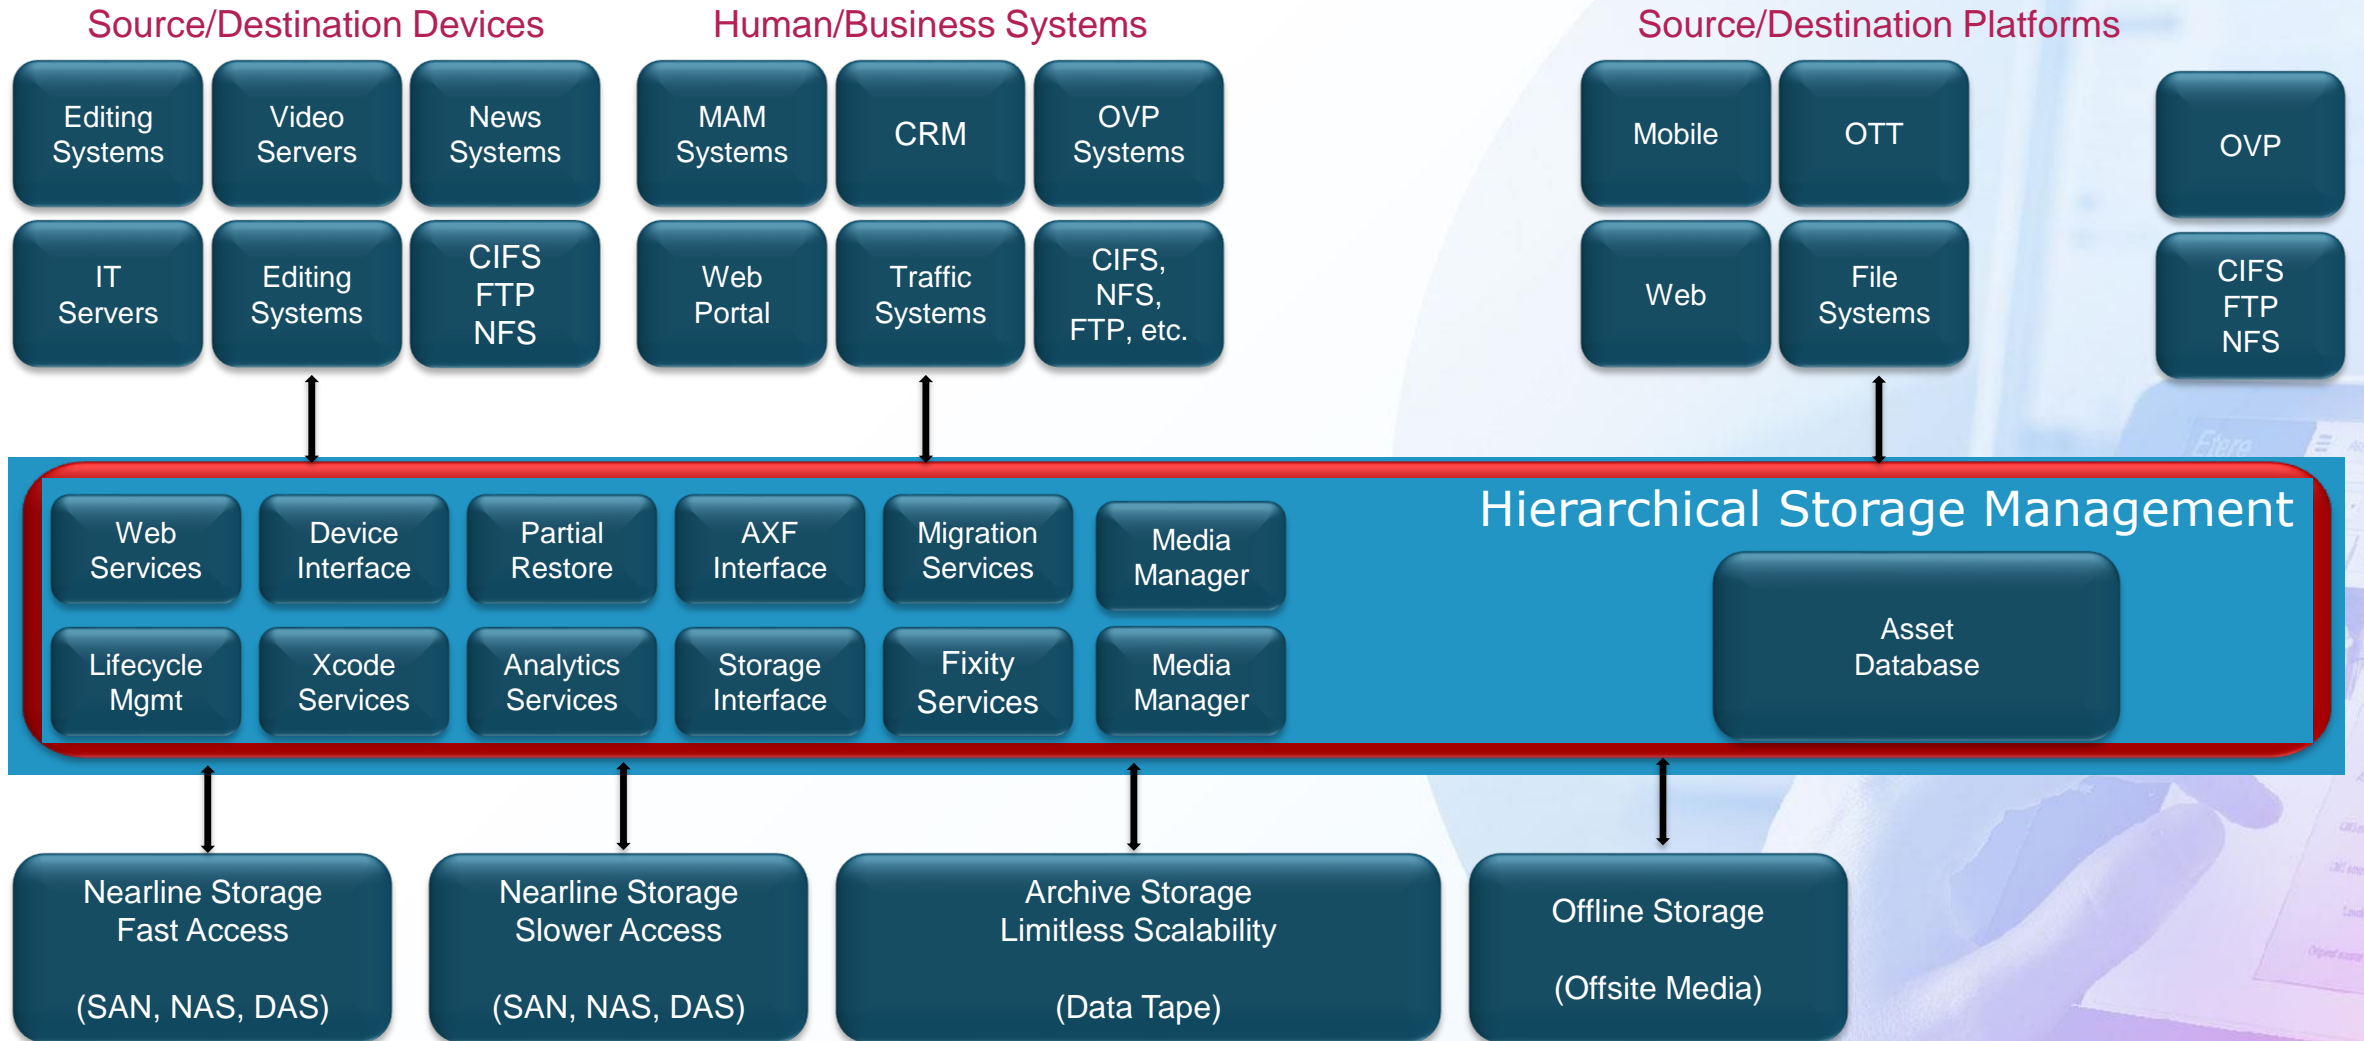
Hierarchical Storage Management (HSM)

- Etere HSM is a object aware, intelligent archive management solution
- HSM is the middleware which resides between any devices which produce or consume file-based content and commodity storage
- HSM is VERY different than disk migration solutions (SAMFS, StorNext, TSM, etc.) which do not handle content as “objects”
- Etere HSM in media applications offer “content aware” features such as transcoding, timecode partial restore, asset analytics, etc. but also extend to non-media applications as well!
- HSM offer direct integration with business systems to fundamentally link processes, resources and metadata with the content repository

Storage Pyramid



Hierarchical Storage Management



Multi-Tier Storage with Varied QoS

Storage Technology and Trends



Global Storage Trends

- Over the next five years the media and entertainment industry will see a ~10X increase in required digital storage capacity
- Currently 93% of the total storage capacity is used for content archiving and long term preservation
- This number is expected to grow to ~96% by 2015 due to more efficient conversion services, lower overall storage costs and a greater ROI on long tail content

Media Type Trends	2010	2015
Data Tape	86%	83%
Hard Drive	10%	14%
Optical	4%	2%
Flash	0.2%	0.6%


Storage Format Choices

	Hard Disk Storage	Flash Storage	Optical Storage	Data Tape Storage
Pros	<ul style="list-style-type: none">• Fast transfer speeds• Good storage density• Random access media• Fast mechanical times• Multiple R/W streams	<ul style="list-style-type: none">• Random access• Persistent storage• Low energy costs• Acquisition format• Fast access• Rugged	<ul style="list-style-type: none">• Random access• No contact read/write• Acquisition format• Emerging holographic• Low replication costs	<ul style="list-style-type: none">• Fast transfer speeds• Storage density• Rugged media• Cost per TB• Extremely portable• Expansion costs• Low replication costs
Cons	<ul style="list-style-type: none">• Tied to host chassis• Not portable• Cost per TB• Cost of ownership• Highly mechanical• Expansion costs• High replication costs	<ul style="list-style-type: none">• Limited read/write• Cost per TB• Proprietary formats• Low transfer speeds• Low storage density	<ul style="list-style-type: none">• Low storage density• Low transfer speeds• Questionable shelf life• Portability• Single R/W stream	<ul style="list-style-type: none">• Sequential data access• Head and tape wear• Slower access• Single R/W stream

Storage Technology Summary

	Type	Media Capacity (GB)	Media Capacity 50Mbps (Hours)	Drive Cost	Street Price (\$/TB)	Drive Speed (MB/s)
BLU RAY	Optical	50	2	\$\$	\$600	14
HOLOGRAPHIC	Optical	300	13	\$\$\$	\$600	20
SOLID STATE (SSD, P2,...)	Flash	16-256	1 - 10	HIGH	\$2,000	~250
SAIT2	Tape	800	35	\$\$\$	\$200	45
LTO4	Tape	800	35	\$\$	\$25	120
LTO5	Tape	1500	65	\$\$\$	\$55	140
TS1140	Tape	4000	175	\$\$\$\$	\$90	240
T10000C	Tape	5000	220	\$\$\$\$	\$70	240

Real World Performance (50Mbps)



Storage Technology Migration

- Data tape “shelf life” need not be a problem once assets have been digitized and placed under HSM control
- HSM systems automatically manage migration to newer storage technologies as a background process at much faster than real time rates
- Each storage generation brings better performance and higher storage density
- Checksums are used during all migration operations to ensure exact byte-for-byte copies are made


Data Tape is the Only Solution

- Unmatched storage density and value
- High performance media ensures no bottlenecks
- TCO is low as data tapes consume no power or cooling
- Low media costs allow automatic offsite asset replication
- Checksums confirm all subsequent file operations
- Automatic migration to new formats ensures longevity
- Initial capital costs are high but flatten out very quickly



a consistent system

What about the Cloud?



What about the Cloud?

- Amazon, Microsoft and many others offer scalable cloud based storage “services”
- All claim reduced capital costs and high availability via globally distributed data centres
- Unfortunately these solutions are not tuned for “big data”
- Amazon Simple Storage Service (S3) cloud storage costs:
 - \$0.125/GB per month to store data
 - No charge for archive transfers (uploads)
 - \$0.10/GB for restore transfers (downloads)
- Let's take a look at a real world Amazon S3 example...

What about the Cloud?

<u>Assumption</u> 10TB/Month Growth Over Five Years	Annualized WAN Fees	Annualized Cloud Storage Fees	Repository Size at the End of Each Year	Total Annualized Cloud Storage Cost
Year 1	\$30,000	\$97,500	120TB	\$127,500
Year 2	\$30,000	\$277,500	240TB	\$307,500
Year 3	\$30,000	\$457,500	360TB	\$487,500
Year 4	\$30,000	\$637,500	480TB	\$667,500
Year 5	\$30,000	\$817,500	600TB	\$847,500
Amazon S3 Five Year Total Cost				~ \$2.5M

Pricing based on Amazon S3 cloud services in \$USD

Results in an average cost of more than \$4,000/TB over five years

Annual storage fee of \$900,000 continues forever!!!

Can Cloud Storage Work?

- Commodity IT-centric cloud solutions such as Amazon cannot work in “big data” environments like those typical in the media industry
- However HSM-based solutions “in the cloud” are now available and do make real economic sense
- These solutions offer cost effective “big data” storage as well as the content-centric features of “HSM as a service”
- How do different cloud services topologies fit and can the transition be handled in a staged manner?

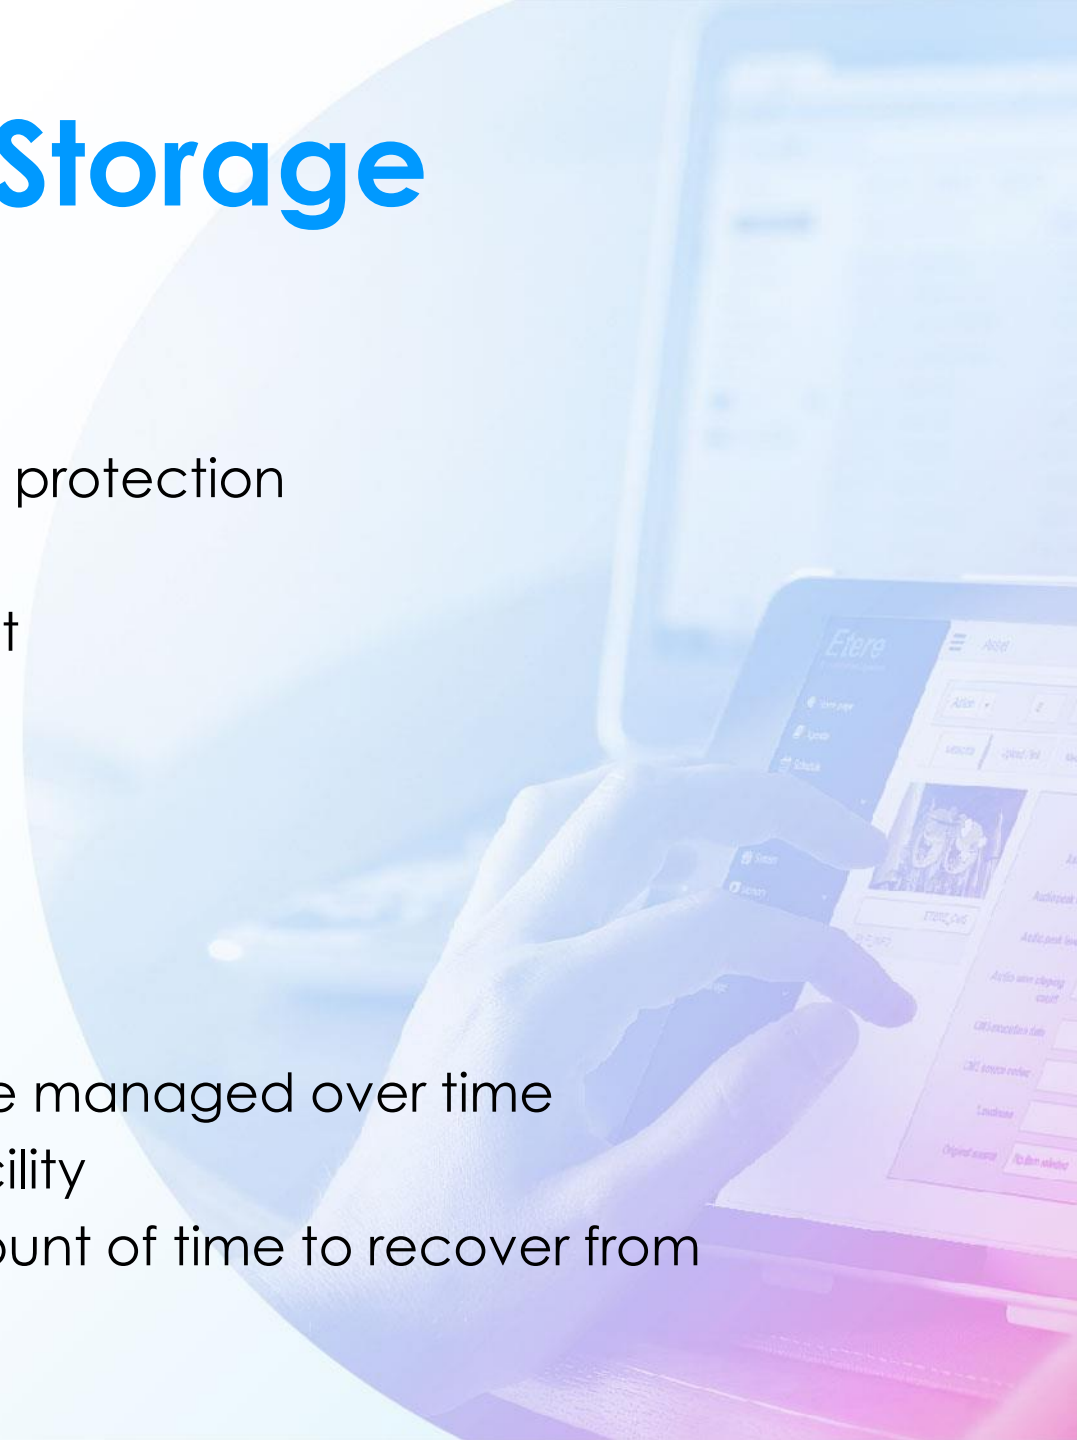
Offsite Media Storage

PROS

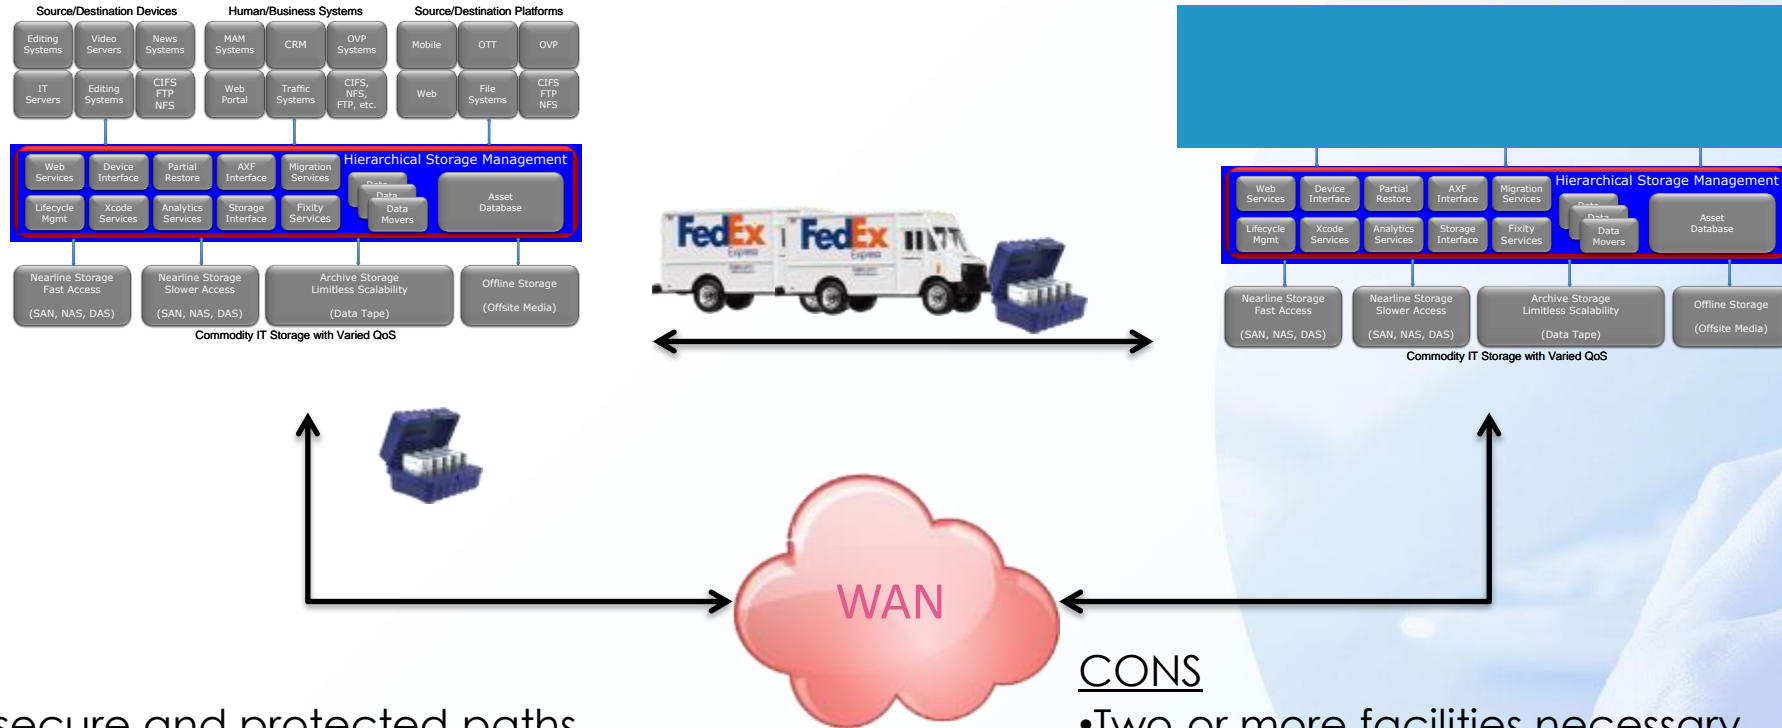
- Cost effective content protection
- Data tapes can be shipped daily offering rapid protection
- “FedEx” offers fast transfers!
- Scale to any size with little additional investment
- Capital costs are minimal
- Workflows are very well known...

CONS

- Workflows are human intensive
- Recall times can be excessive
- Tapes on the shelf are not “active” and must be managed over time
- Does not offer any redundancy to the main facility
- Main facility disaster can take a significant amount of time to recover from



The Private Cloud



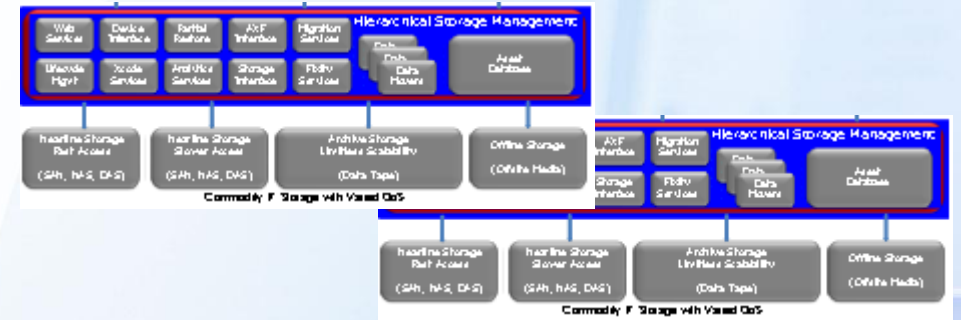
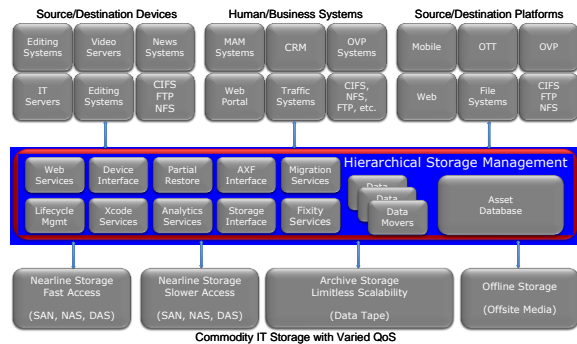
PROS

- Private, secure and protected paths
- Controlled and fairly predictable workflows
- Business continuance possible
- Shipping of high density data tape media is FAST and INEXPENSIVE

CONS

- Two or more facilities necessary
- Infrastructure, capital and staffing costs are simple multiples
- Technical challenges are multiplied by the number of facilities
- Heavy capital and operations costs

The Hybrid Cloud



HSM Cloud Service Provider

WAN

PROS

- Protected and encrypted paths
- Controlled and predictable managed workflows
- Business continuance as a service possible
- Shipping of high density data tape media is FAST & INEXPENSIVE
- Services can be “tuned” on demand

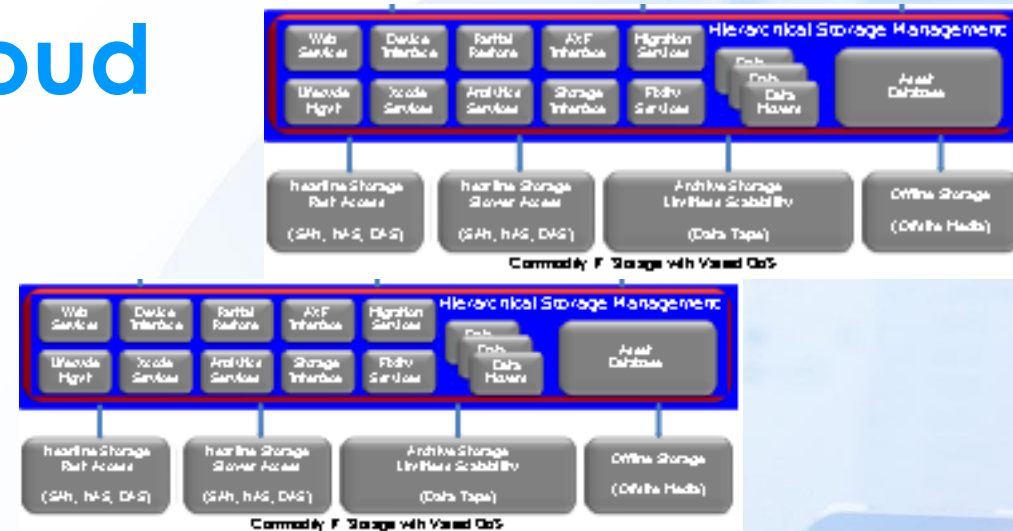
CONS

- Content must be stored in a “hosted” environment
- Security is a key concern which must be carefully managed

HSM in the Cloud



Customer Facilities



Cloud Service Facilities



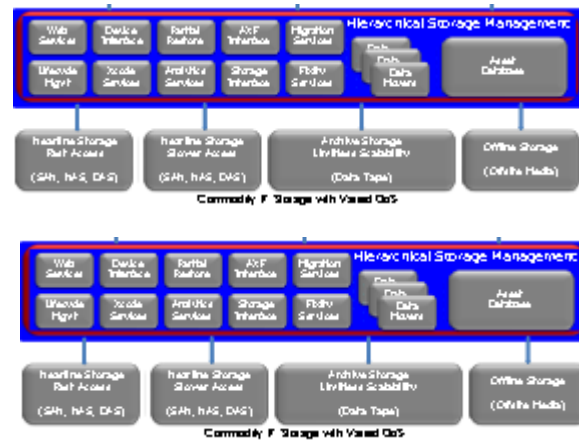
PROS

- HSM as a service
- Protected and encrypted paths
- Controlled and predictable managed workflows
- a as a service
- Services can be tuned on demand
- Global distribution and diversity
- Perfect for “elastic” applications

CONS

- Content stored in a “hosted” environment
- Network costs can be prohibitive
- High availability WAN must be part of the solution
- Local “caching” may be necessary to ensure rapid local access

Cloud Services Evolution



Disaster Recovery and Global Replication

Content Distribution Services

The Connected Enterprise

Business Continuity as a Service

Online Video Publishing as a Service

Elastic Transcoding Farm

Elastic Facility Services

Storage Migration Services

Global Media Asset Management (MAM)

Metadata Mining as a Service


Hub and Spoke Distributed Facility

How are my Valuable Assets Actually Stored?



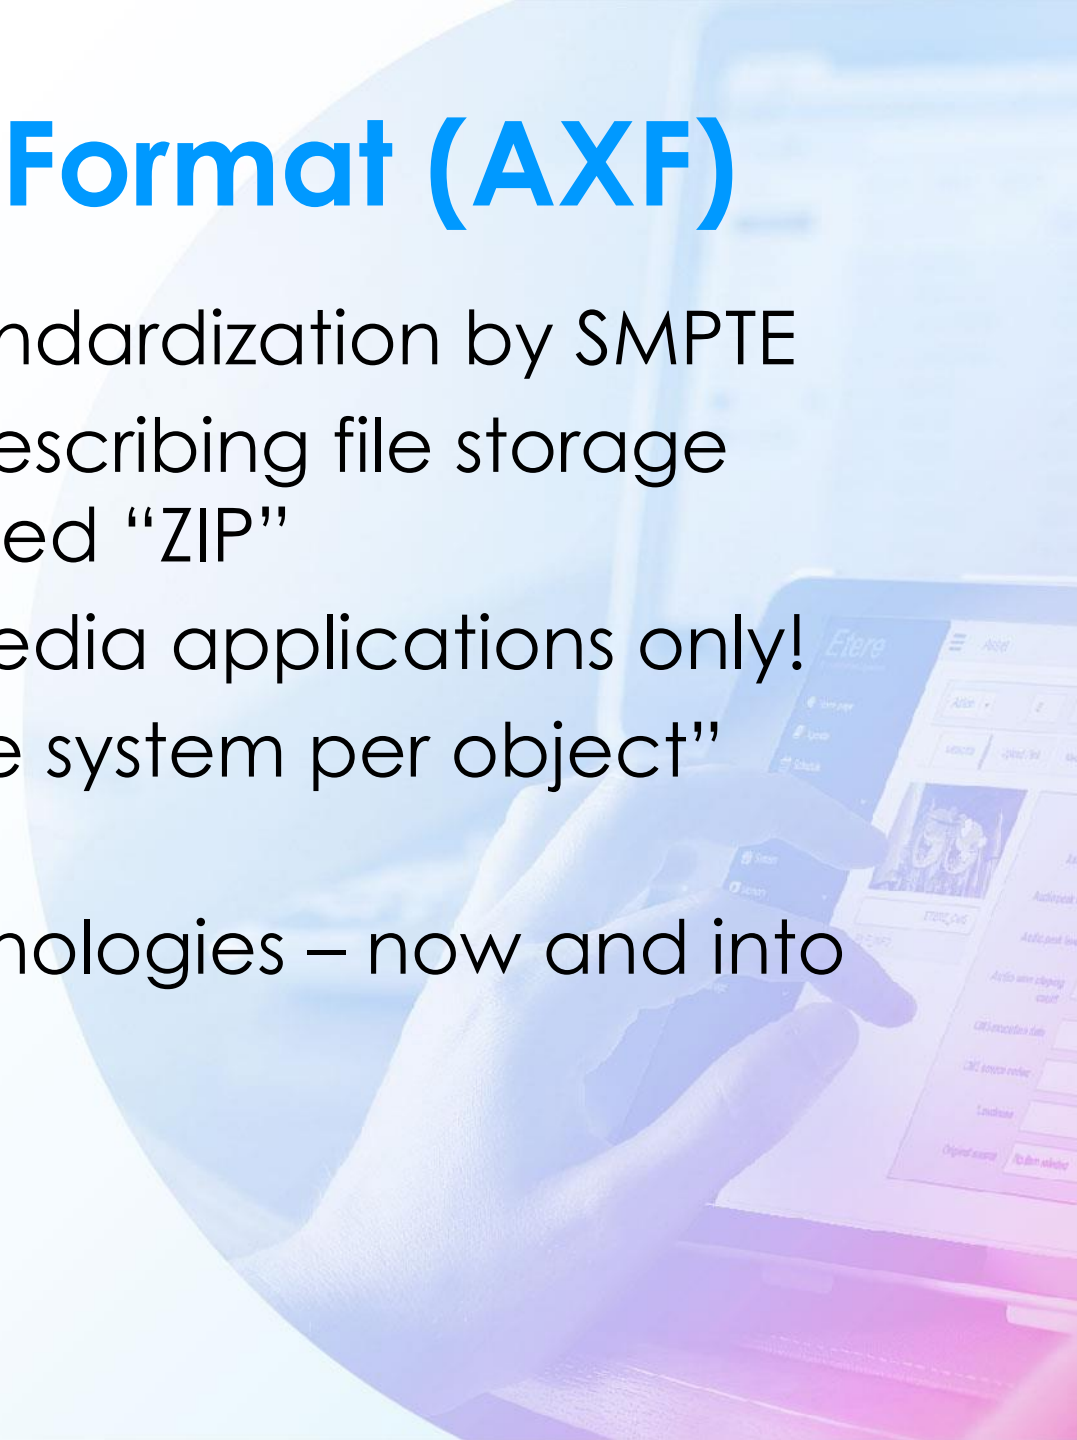
How Are My Valuable Assets Stored?

- We must take steps to ensure long term accessibility to valuable file-based assets
- Key goals of the “ideal storage format” include:
 - Ensure long term accessibility
 - Self descriptive assets and self descriptive storage media
 - Preservation (OAIS) features such as fixity, provenance, etc.
 - File encapsulation to wrap related metadata and files
 - Scalability for any number of elements of any size and type
 - Standardized regardless of storage media technology
 - Transportability and compatibility between systems
- What choices are there?

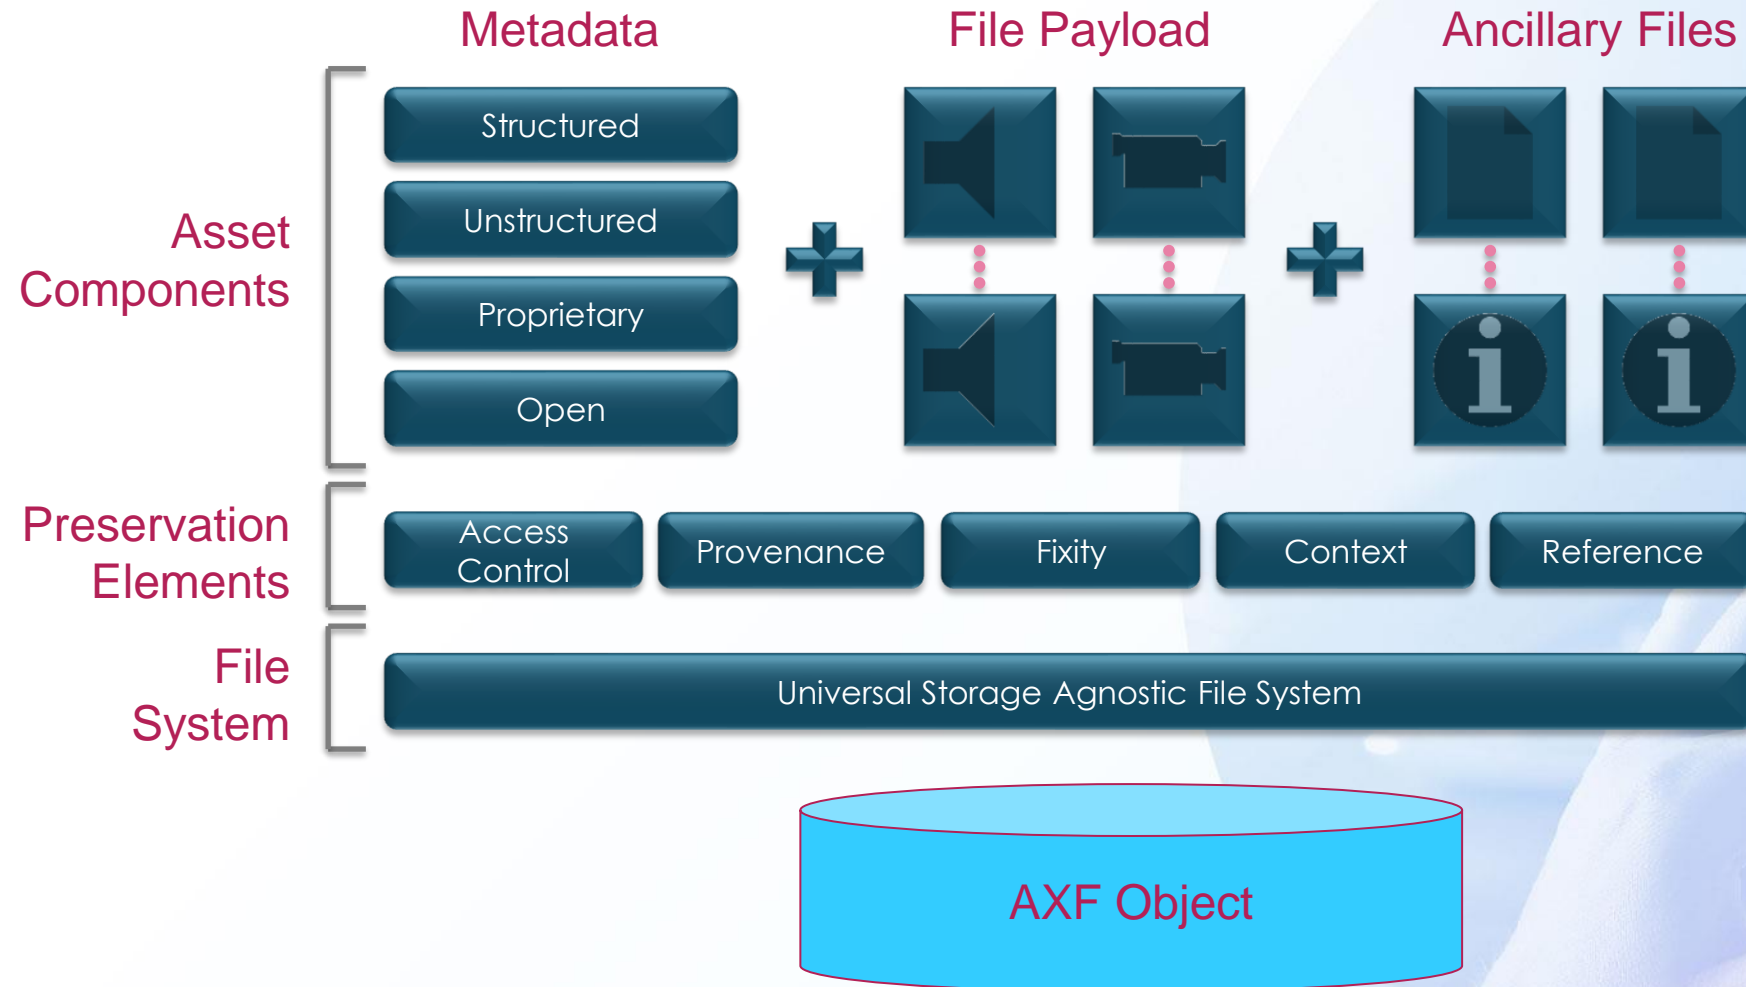


Archive eXchange Format (AXF)

- AXF was in the final process of standardization by SMPTE
- AXF is a fully self-contained, self-describing file storage "container" – think of it as advanced "ZIP"
- AXF is IT-centric and not tied to media applications only!
- AXF is based on an innovative "file system per object" approach
- AXF fully supports all storage technologies – now and into the future!

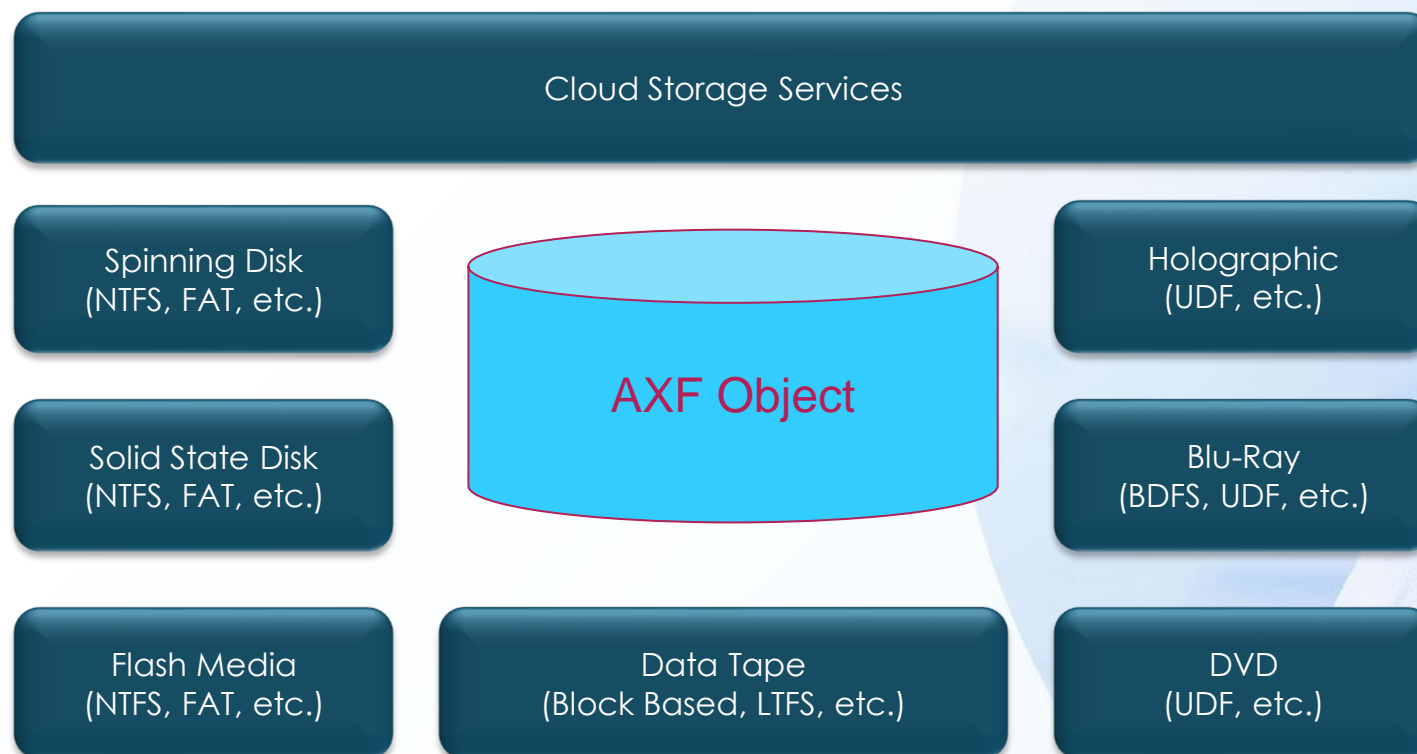


What is AXF?



AXF and Storage

AXF is storage technology, operating system and file system agnostic



AXF Benefits

- AXF Objects can scale to any size and encapsulate any number of files with full support for media spanning
- No need to upgrade existing storage infrastructures
- AXF guarantees long term compatibility and resiliency with self-describing features for both AXF Objects and AXF Media
- AXF overcomes all the technical, operational and functional limitations of other options such as TAR and LTFS
- AXF is an IT-centric implementation and is not limited to media files alone (documents, imaging data, etc.)
- Visit OpenAXF.org for more information