

Using Rosetta, StorageGRID, and New IBM Tape Solutions to Implement State-of-the-Art Digital Preservation

Gary T. Wright

wrightgt@ldschurch.org

ELUNA 2012

May 2012



image courtesy of IBM

Introducing the Church

- The Church of Jesus Christ of Latter-day Saints
- Global Christian church with 14 million members
- 700,000 students enrolled in religious training
- 3 universities, 1 college
- State-of-the-art audio-visual capabilities
- Scriptural mandate to keep and preserve records



photo by Henok Montoya

Church History Department

- Preserves records of enduring value from Church leaders, departments, universities, and affiliations (more than 35 organizations)
- Helps each organization develop a records management plan
- Typically, less than 10% of records are archived



Church History Library on Temple Square

Church Audiovisual Capabilities

- Audiovisual records will consume vast majority of archive capacity



Free Bible videos from biblevideos.lds.org

- 100+ PB in a decade for a single copy!



Mormon Tabernacle Choir and Orchestra




Conference Center on Temple Square

Architecting a Preservation System

- National Library of New Zealand business requirements

NATIONAL DIGITAL
HERITAGE ARCHIVE
PROGRAMME

- OAIS Reference Model



- Minimize cost of archival storage

- Scalability




image courtesy of NetApp

Minimizing Cost of Archival Storage

Total cost of storage ownership study

- Over ten years, ownership and operating costs of tape are **one third** of associated costs for disk arrays
- Power and cost per TB advantages of tape are expected to increase over time
- Conclusion—for now, tape is required to sustain a multi-PB digital archive
- But . . . tape presents some challenges



Scalability to Meet Church Needs

- After research, scalability concerns eliminated open source repositories
- Ex Libris Rosetta SPOC
- White paper at exlibrisgroup.com
- Conclusion—Rosetta can meet Church needs
- CHIPS—Church History Interim Preservation System



DRPS Architecture

Digital Records Preservation System

Fixity
Creation

DRPS Ingest Tools

Preservation
Functions

ExLibris Rosetta

Fixity
Bridge

Storage Extensions

Information
Lifecycle
Management

 StorageGRID
NetApp®

Tape
Interface

IBM
Tivoli Storage Manager

Ensuring Archive Data Integrity

- A critical requirement of digital preservation—*differentiates an archive from other tape farms*
- Fixity information (secure hash checksum) enables data integrity validation by checking it every time data is written, transferred, moved, or copied
- End-to-end fixity checking should be performed from file ingest to permanent storage to delivery
- Periodic validation of the *entire* archive should also be performed to detect bit flips (i.e., bit rot)
- DRPS uses a variety of integrity values for fixity

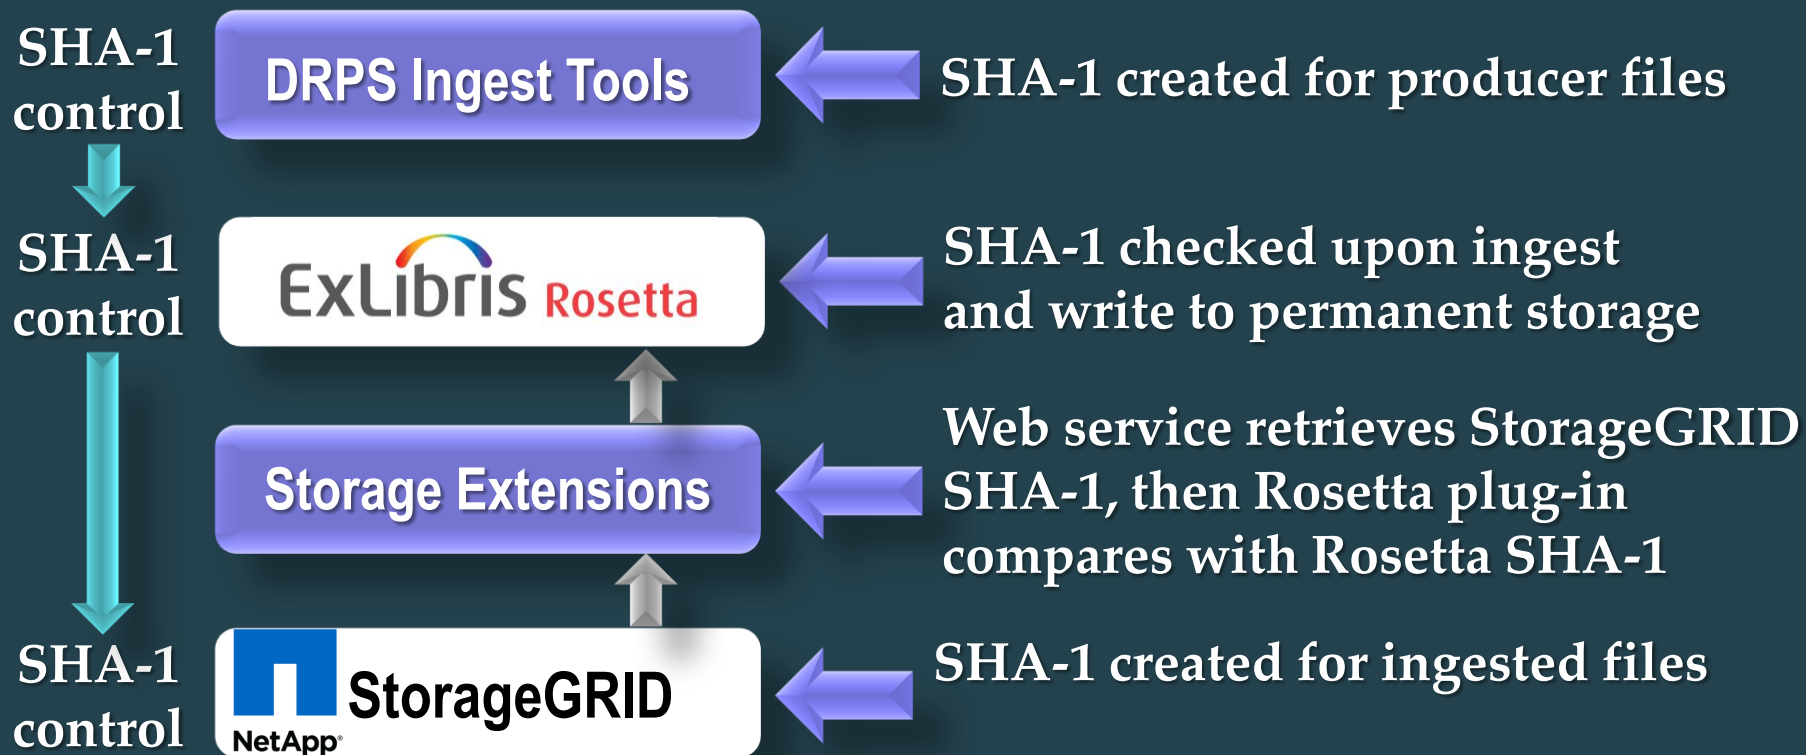
Sizing Bit Errors

- Recent validation of entire DRPS archive resulted in a 3.3×10^{-14} bit error rate
- USC Shoah Foundation Institute visit
- 8 PB tape archive of videotaped interviews of Holocaust survivors and other witnesses
- Experienced 1500 bit flips in 8 PB (2.3×10^{-14} bit error rate)
- Real life measurements that provide guidance for tape archives



100110110100001011

DRPS Data Integrity Validation



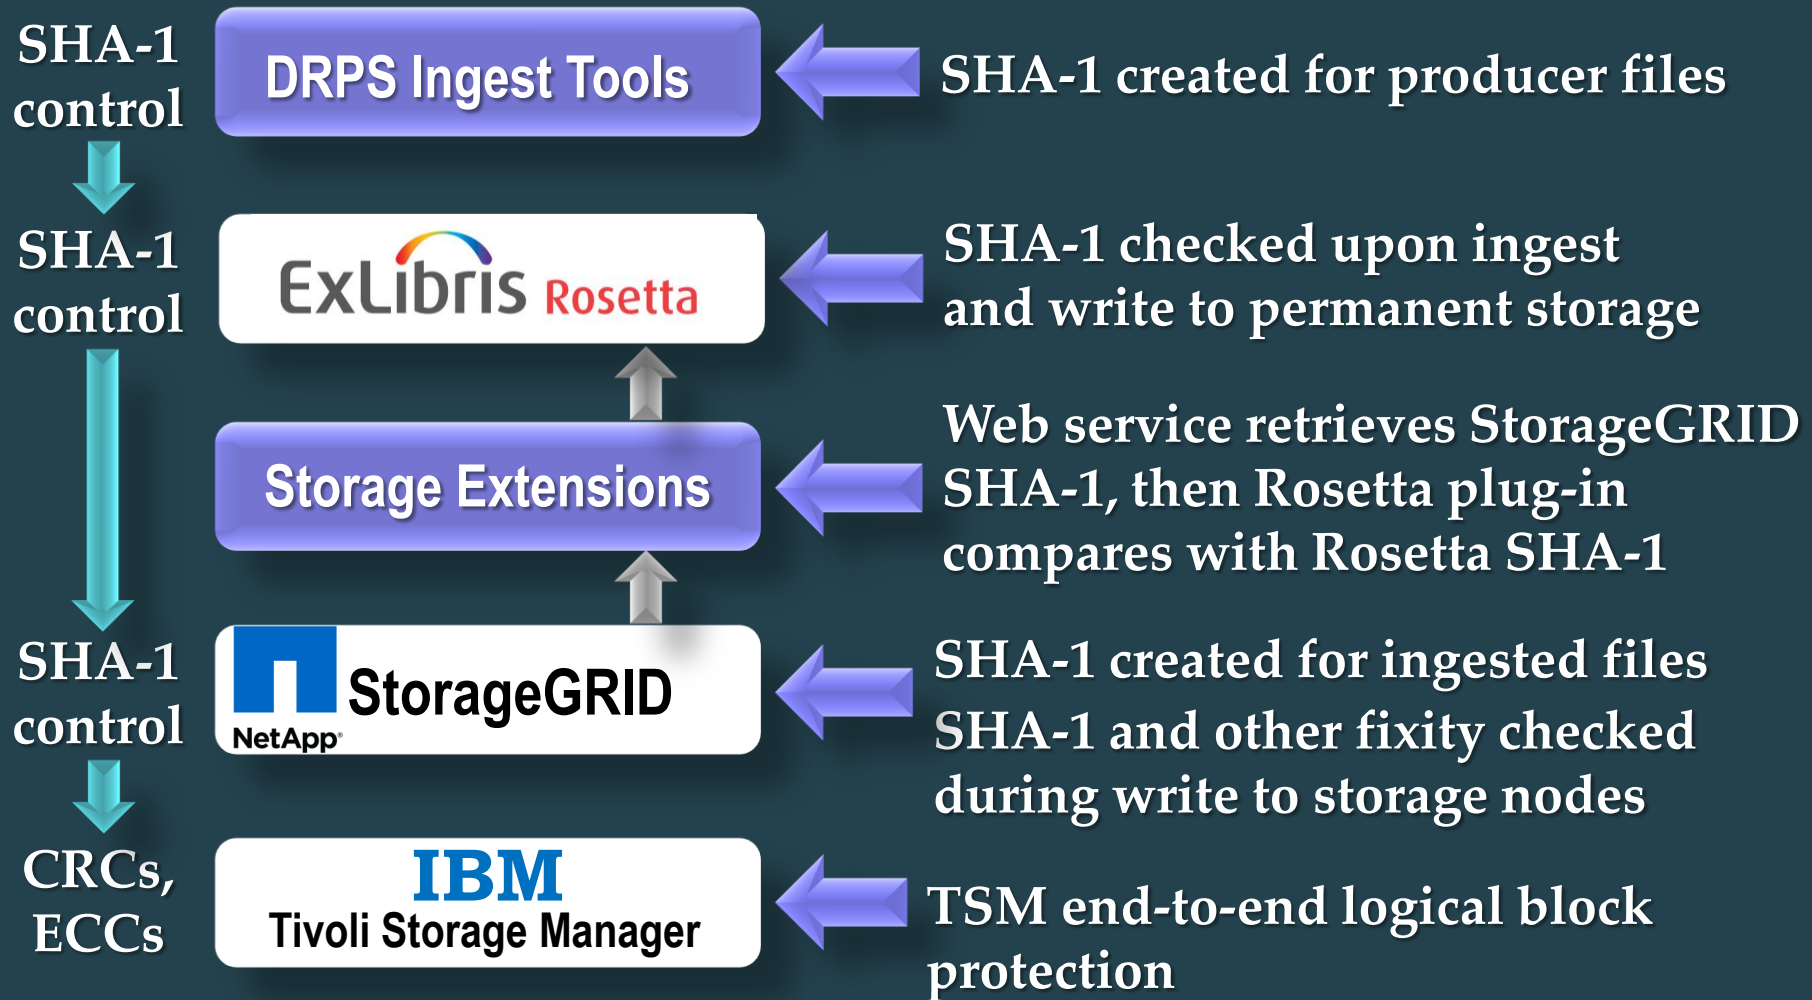
StorageGRID Fixity Checking

- StorageGRID is constructed around the concept of object storage



- Provides a layered/overlapping set of protection domains to guard against object data corruption
 1. SHA-1 object hash—checked on store and access
 2. Content hash—checked on access
 3. CRC checksum—checked with every operation
 4. Key-based hash value—checked on access

DRPS Data Integrity Validation




TSM End-to-End Logical Block Protection


- Supersedes SHA-1 fixity information with cyclic redundancy check values (CRCs) and error-correcting codes (ECCs)
- Enabled with new, state-of-the-art functionality of IBM LTO-5 and TS1140 tape drives
- Seamlessly extends validation of data integrity until AIPs are correctly written to tape




TSM End-to-End Logical Block Protection



1. TSM server calculates and appends “original data CRC” to AIP logical block






2. Tape drive computes its own CRC and compares to original data CRC




TSM End-to-End Logical Block Protection

3. As logical block is loaded into drive data buffer, on-the-fly verifier checks original data CRC





4. In parallel, a “C1 code” (Reed-Solomon ECC) is computed and appended

TSM End-to-End Logical Block Protection





5. An additional ECC, referred to as “C2 code,” is added to the logical block



6. More powerful than the original data CRC, the C1 code is checked every time data is read from the buffer

TSM End-to-End Logical Block Protection

- 7. Data written to tape at full line speed with read-while-write process
- 8. Just written data loaded to buffer and C1 code checked




Successful read-while-write operation assures no data corruption from TSM client to tape



TSM End-to-End Logical Block Protection

9. When tape is read, all codes (C1, C2, original data CRC) are checked by drive

10. Original data CRC appended to logical block



11. TSM server verifies original data CRC, completing TSM end-to-end logical block protection cycle

Ongoing Archive Data Integrity

- Must assume that bits will flip after being written correctly to tape

100110110100001011

- Actual validation results indicate a 10^{-14} bit error rate

- Therefore, *all* tapes must be read periodically to identify and correct bit errors



image courtesy of IBM

Ongoing Archive Data Integrity

- Staging AIPs so a server can verify integrity is resource intensive!
- IBM LTO-5 and TS1140 tape drives provide a more efficient solution
- During “Verify” operation, a tape is mounted, drive checks all codes (C1, C2, original data CRC) as data is being read (at full line speed)
- Only status is reported as these internal checks are completed



image courtesy of IBM


DRPS Data Integrity Summary

- Fixity information is the key to data integrity
- SHA-1 values ensure data integrity to StorageGRID
- TSM end-to-end logical block protection ensures data integrity to tape
- In-drive validation enables ongoing integrity checks for the entire archive

SHA-1
control



SHA-1
control



SHA-1
control



CRCs,
ECCs

DRPS Ingest Tools

ExLibris Rosetta

Storage Extensions

StorageGRID
NetApp

IBM
Tivoli Storage Manager

Image Preservation File Formats

- Lossless JPEG 2000 for still images (50% - 60% archive capacity savings with no loss of resolution)
- Motion JPEG 2000 wrapped in MXF containers for HD video files

Parish or Township of <i>St. Thomas</i>		Ecclesiastical District		City or Borough of <i>New Salem</i>		Town of		Village of	
No. of books bound	Name of Street, Place, or Road, and Name or No. of House	Name and Christian name of each Person who works in the house on the Night of the 30th March, 1851	Relation or Rank of Family	Condition	Age of Males	Age of Females	Rank, Profession, or Occupation	Where Born	Whether Free or Slave
72	<i>High Street</i> <i>number 6</i>	<i>Elizabeth Morris</i>	<i>Wife</i>	<i>U</i>		17	<i>Dressmaker</i>	<i>Wilton, Salisbury</i>	
		<i>Sarah Anne B</i>	<i>Daughter</i>	<i>U</i>		16		<i>B</i>	<i>B</i>
		<i>Emma Elizabeth B</i>	<i>Daughter</i>	<i>U</i>		13	<i>Scholar</i>	<i>B</i>	<i>B</i>
73	<i>High Street</i>	<i>George Firtlee</i>	<i>Head</i>	<i>M</i>	62		<i>Stone and Bone maker</i>	<i>B</i>	<i>B</i>
		<i>Mary Ann B</i>	<i>Wife</i>			63	<i>W & B Assistant</i>	<i>B</i>	<i>B</i>
		<i>Sarah B</i>	<i>Daughter</i>	<i>U</i>		17	<i>Apprentice to Milliner</i>	<i>B</i>	<i>B</i>
		<i>Emily B</i>	<i>Daughter</i>			13	<i>Netting</i>	<i>B</i>	<i>B</i>
74	<i>High Street</i>	<i>Benjamin Blake</i>	<i>Head</i>	<i>M</i>	36		<i>Upholsterer</i>	<i>Dorset, Blandford</i>	
		<i>Frederick B</i>	<i>Wife</i>	<i>M</i>		31		<i>Hants, Bishops Cleeve</i>	
		<i>John B</i>	<i>Son</i>			11	<i>Scholar</i>	<i>Wilton, Dorchester</i>	
		<i>Caroline B</i>	<i>Daughter</i>			6	<i>B</i>	<i>Dorset, Blandford</i>	
		<i>Elizabeth B</i>	<i>Daughter</i>			2		<i>Wilton, Salisbury</i>	
		<i>George B</i>	<i>Son</i>			1 year		<i>B</i>	<i>B</i>
		<i>Pauline Dean</i>	<i>Apprentice</i>			15	<i>Cabinet Maker's Apprentice</i>	<i>B</i>	<i>Farnham</i>
75	<i>High Street</i>	<i>James Crawford</i>	<i>Head</i>	<i>M</i>	36		<i>Bonemaker</i>	<i>Dorset, Exton</i>	
		<i>Abigail B</i>	<i>Wife</i>	<i>M</i>		35	<i>Washer</i>	<i>Hampshire, London N.E.</i>	
76	<i>High Street</i>	<i>Francis Dixon</i>	<i>Head</i>	<i>M</i>	30		<i>Brushmaker</i>	<i>Wilton, Salisbury</i>	
		<i>Mary Ann</i>	<i>Wife</i>	<i>M</i>		27		<i>B</i>	<i>B</i>
		<i>James B</i>	<i>Son</i>			6		<i>B</i>	<i>B</i>



Scene from biblevideos.lds.org

MXF Challenges

- General Conference MXF containers package HD video, ASL video, up to 96 repeating audio tracks
- No known tool can validate file formats and extract metadata within an MXF container
- MediaInfo tool only extracts from MXF container
- Rosetta does not yet support ingest of repeating tracks



Scene from biblevideos.lds.org

MXF Solutions

- Preservation team developed a Rosetta plug-in that utilizes MediaInfo to extract metadata
- MXF Extraction Tool concatenates data from repeating tracks in the metadata it extracts
- A modified MXF Extraction Tool given to Ex Libris (one video, one audio)
- *Available to all Rosetta users!*



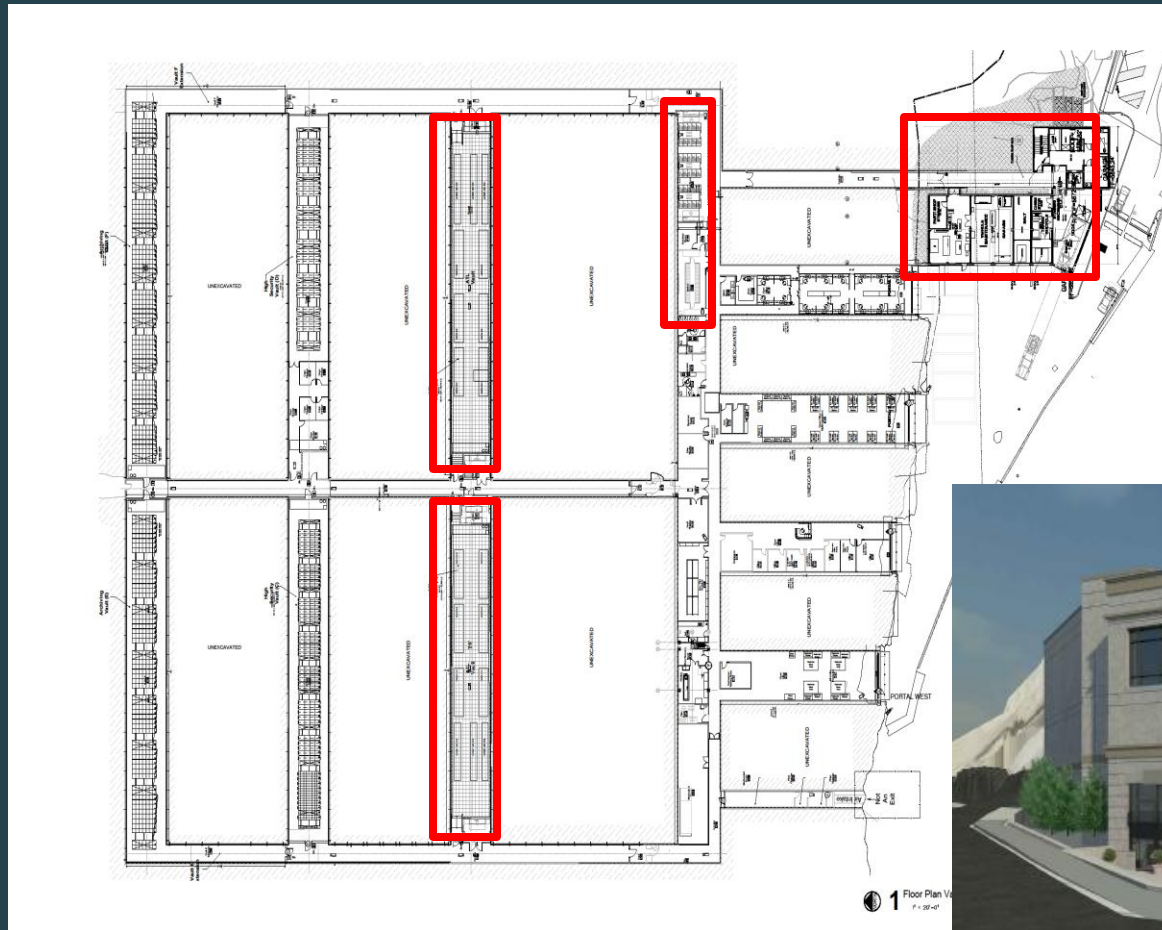
Scene from biblevideos.lds.org

Granite Mountain Records Vault



- Six tunnels bored into a solid granite mountain
- Stores FamilySearch microfilm collection and priceless artifacts
- Plans recently developed to renovate the facility for digital preservation


Granite Mountain Records Vault



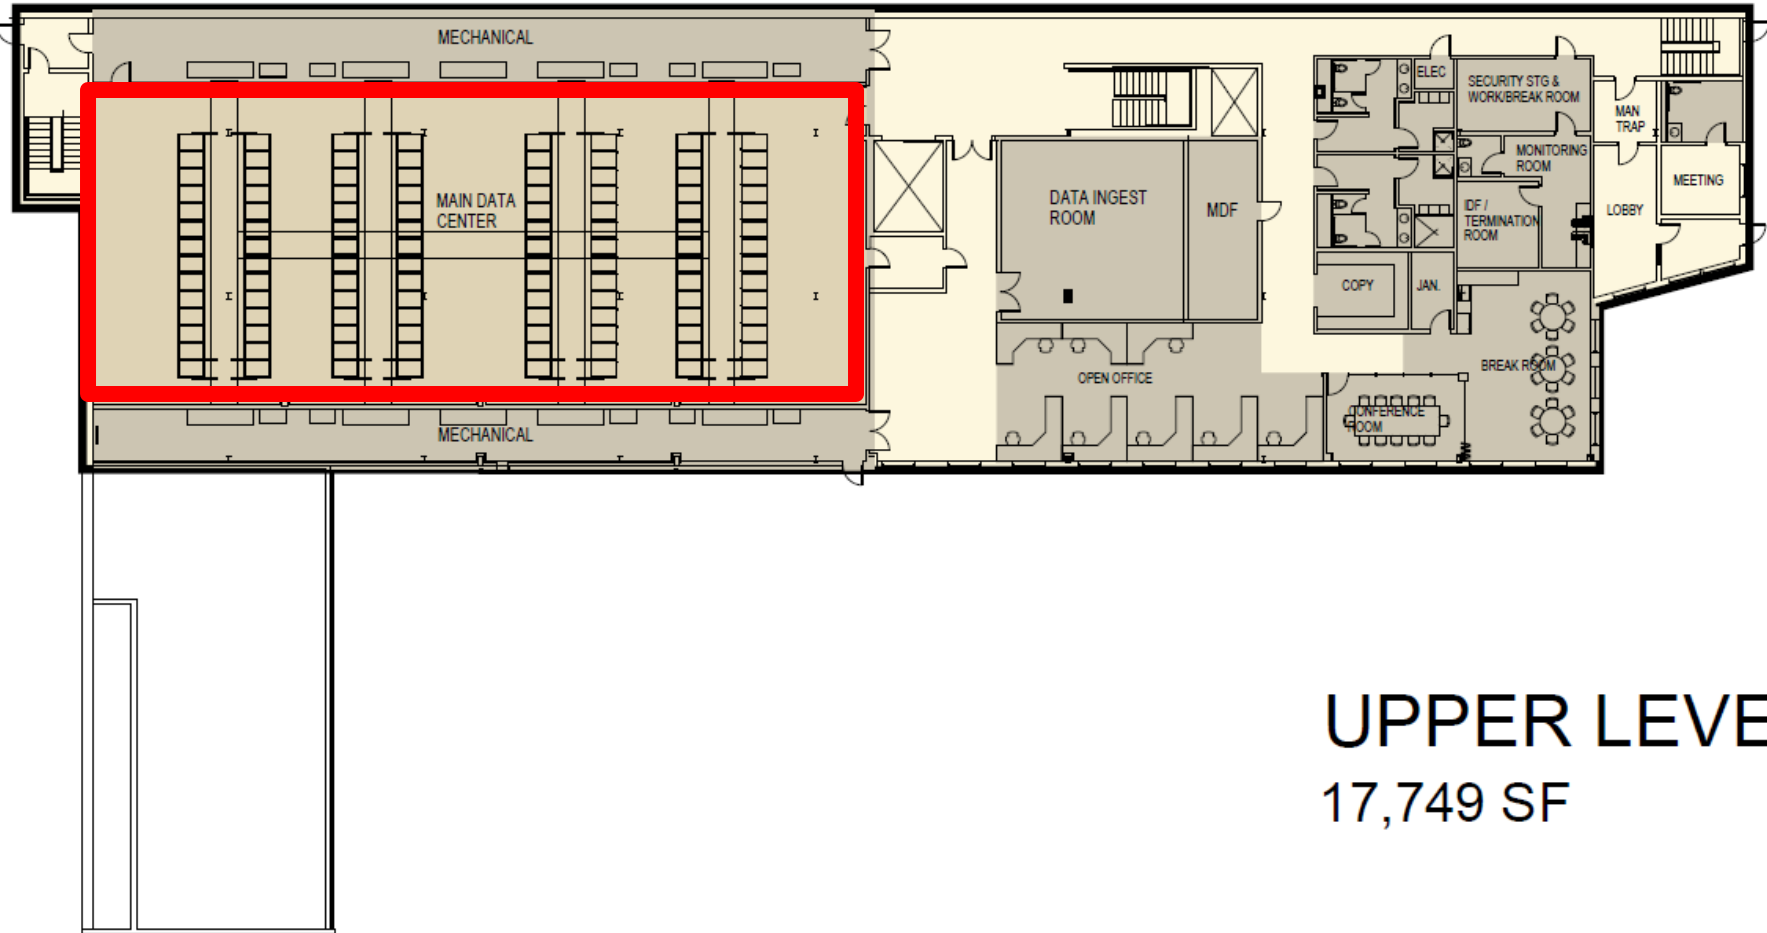
Remote Digital Preservation Facility



Remote Digital Preservation Facility



Remote Digital Preservation Facility



UPPER LEVEL
17,749 SF

Why Is the Church Doing All This?

- Build character
- Strengthen families
- *Foster personal and family happiness*



Thank you!

Questions?

Trademarks

The Ex Libris logo and Rosetta are trademarks of Ex Libris Group.

The NetApp logo and StorageGRID are trademarks of NetApp, Inc.

The IBM logo and Tivoli Storage Manager are trademarks of International Business Machines Corporation.