



The Ability to Preserve a Large Volume of Digital Assets

A Scaling Proof of Concept

CONFIDENTIAL INFORMATION

All of the information and material herein, inclusive of text, images, logos, and product names, is either the property of, or used with permission of Ex Libris Ltd. The information included in this document may not be distributed, modified, displayed, or reproduced, in whole or in part, without the prior written permission of Ex Libris Ltd.

TRADEMARKS

Ex Libris, the Ex Libris logo, Aleph, SFX, SFXIT, MetaLib, DigiTool, Verde, Primo, Voyager, MetaSearch, MetaIndex, and other Ex Libris products and services referenced herein are trademarks of Ex Libris and may be registered in certain jurisdictions. All other product names, company names, marks, and logos referenced herein may be trademarks of their respective owners.

DISCLAIMER

The information contained in this document is compiled from various sources and provided on an "as is" basis for general information purposes only without any representations, conditions, or warranties, whether express or implied, including any implied warranties of satisfactory quality, completeness, accuracy, or fitness for a particular purpose.

Ex Libris, its subsidiaries, and related corporations ("Ex Libris Group") disclaim any and all liability for all use of this information, including loss, damages, claims, or expenses any person may incur as a result of the use of this information, even if advised of the possibility of such loss or damage.

© Ex Libris Ltd., 2010

The Ability to Preserve a Large Volume of Digital Assets: A Scaling Proof of Concept

Contents

The Need for Digital Preservation	4
The Church of Jesus Christ of Latter-day Saints' Preservation Needs.....	5
Overall Objective.....	7
The Rosetta Digital Preservation Solution.....	8
Overview of Rosetta's Scalability Architecture	9
Scaling Proof of Concept Goals	10
Testing Methodology and Criteria	11
Test Results.....	12
Conclusions	13
Appendix: Test Configuration Technology	14
Glossary of Terms	15

The Need for Digital Preservation

International research has shown that the “digital universe” is growing at an unforeseen rate. Indeed, current forecasts expect the volume of digital content—from CCTV footage and personal photographs to academic research, research output in digital formats, Web sites, and corporate financial records—to double in size every 18 months.¹

As the options for creating and duplicating digital items multiply, comparable options for storing those items remain limited. Documents, images, and film produced in recent decades will be inaccessible to future generations as a result of continuing waves of innovation that render older storage systems obsolete.

As described in a white paper from Preservation and Long-term Access through Networked Services (Planets), “The more information we produce, and the more we can hold in a given space, the shorter the time we seem to be able to keep hold of it for.”²

Libraries, museums, and archives face particularly stark realities in this digital environment. Not only do they continue to store and provide access to the thousands of objects and resources in their care, but they must ensure that the huge quantities of digital items being acquired with every passing month are stored, preserved for the long term, and available for future generations.

With increasing numbers of resources that are born digital, rather than being digitized from a physical copy, the challenge for “memory” institutions (libraries, museums, and archives) is unique. Whereas the records of governments, companies, or agencies might be expected to become outdated and irrelevant

¹John F. Gantz & David Reinsel, “As the Economy Contracts, the Digital Universe Expands,” an IDC Multimedia White Paper sponsored by EMC, May 2009.

²Pauline Sinclair, “The Digital Divide: Assessing Organisations’ Preparations for Digital Preservation,” a Planets (Preservation and Long-term Access through Networked Services) white paper, March 2010.

over time, memory institutions are committed to preserving the records in their care in perpetuity.

A study by Planets (Preservation and Long-term Access through Networked Services)—a European Union-funded project to assess digital preservation needs—found that while 80% of organizations currently store documents and images, by 2019 over 70% will need to preserve databases, Web sites, audio files, and video files, as well.³ The emerging task for libraries is not only how to store the huge diversity of digital assets being generated but also how to equip institutions to scale their preservation efforts to meet the rapidly expanding volume of data.

Even experts have not accurately predicted the speed at which the digital universe is expanding. Research by the International Data Corporation (IDC) in 2010 underestimated the previous year's rate of growth by 3%, while their earlier paper in 2008 misjudged the global volume of digital items by 10%.⁴ As institutions look to a future of absorbing millions of digital objects each year, scalability is a fundamental requirement for their digital preservation architecture.

The Church of Jesus Christ of Latter-day Saints' Preservation Needs

The Church of Jesus Christ of Latter-day Saints ("the Church") is a worldwide organization of over 13 million members in over 28,000 congregations. The worldwide headquarters for the Church has been located in downtown Salt Lake City, Utah, since pioneers settled the community in 1847. Today, Church leaders oversee Church operations in some 160 countries from this location. However,

³ Pauline Sinclair, "The Digital Divide: Assessing Organisations' Preparations for Digital Preservation," a Planets (Preservation and Long-term Access through Networked Services) white paper, March 2010.

⁴ John F. Gantz et al., "The Diverse and Exploding Digital Universe: An Updated Forecast of Worldwide Information Growth through 2011," an IDC white paper sponsored by EMC, March 2008.

leadership is heavily decentralized throughout the world, with local, regional, and national levels of administration operating mainly through a lay clergy.

FamilySearch, a service provided by the Church, is the largest genealogy organization in the world. Millions of people use FamilySearch records, resources, and services to learn more about their family history. For over 100 years, FamilySearch has been actively gathering, preserving, and sharing genealogical records worldwide. FamilySearch manages the largest collection of genealogical collections in the world—2.5 million rolls of microfilm representing over 13 billion names and millions of additional digital images from over 100 countries worldwide. In 2007, the Church announced plans to begin digitizing and indexing its collection for broader, more economic online access—starting with popular collections like the U.S., Canada, and U.K. censuses. FamilySearch has created free online indexes to date for the 1850, 1860, 1870, 1880, and 1900 U.S. censuses.

In addition to the collection of genealogical records, the Church also maintains a huge library of other records and media in a digital format. For example, the Church broadcasts its semiannual worldwide conference in high-definition video, which is then translated into 92 languages and made available through the Church's Web site, www.lds.org.

With genealogical, audio-visual, and other records, the Church's need for digital preservation is growing at a tremendous pace. This growth and the sacred nature of the records have provided a need for the Church's Information and Communications Systems (ICS) department to investigate digital preservation strategies and particularly to evaluate the ability of the Ex Libris Rosetta system to scale up to meet the challenge. Figure 1 shows the projected growth of digital assets that need to be preserved by the Church.

Growth of Digital Preservation

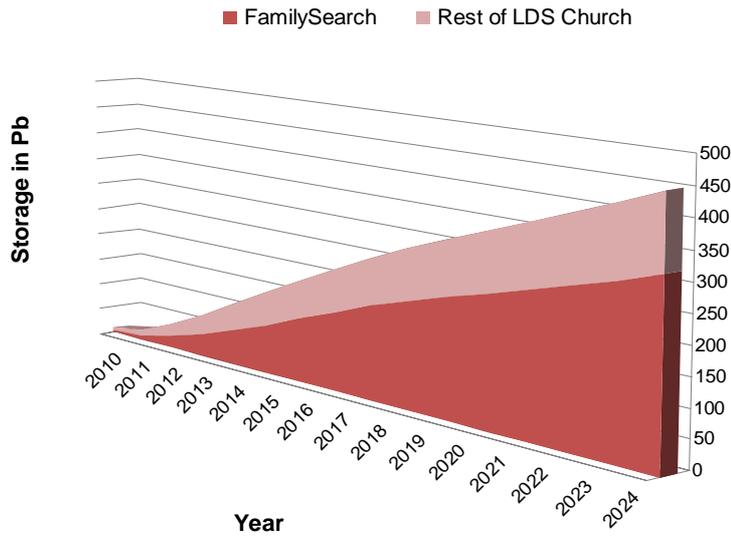


Figure 1. Growth of preservation storage

Overall Objective

Given the challenges described earlier, the Church desired to create a digital record preservation system (DRPS) that would securely and cost-effectively preserve, in perpetuity, all Church information of enduring value and permit straightforward accessibility for all Church departments and ecclesiastical leaders with the appropriate viewing permissions. The Church chose to evaluate Rosetta, a standards-based, commercial digital archive and preservation system developed by Ex Libris, and to explore the mission-critical nature of the system. The Church wanted to know whether Rosetta could ingest and store the number of records required to support the Church's preservation needs. A "scaling" proof of concept project was defined and carried out to answer this question.

The Rosetta Digital Preservation Solution

The Open Archival Information System (OAIS) reference model, an ISO-recognized standard, describes the characteristics of a digital preservation system.⁵ The model has become widely accepted among preservation bodies and experts worldwide and has been used as a guideline to evaluate current implementations of preservation and archiving initiatives.⁶

The OAIS model specifies six high-level functions that must be present in a digital preservation system:

- Ingest of digital objects
- Storage of digital objects
- Data management
- Administration
- Preservation planning
- Provision of access

These six functions are at the core of the Ex Libris Rosetta digital preservation system. Developed in partnership with the National Library of New Zealand and reviewed by a peer group of world-renowned preservation experts, Rosetta addresses libraries' and archives' need to collect, manage, and preserve a wide variety of digital objects in different formats and structures.

Released in 2009, Rosetta provides memory institutions with end-to-end management capabilities for digital entities—from submission to dissemination. A rule-based workflow engine and open architecture enable institutions using the system to develop unique plug-in tools and applications to enhance the system's ingest, management, preservation, and delivery processes.

⁵ <http://public.ccsds.org/publications/archive/650x0b1.pdf>

⁶ See, for example, "Assessment of UKDA and TNA compliance with OAIS and METS Standards," at http://www.jisc.ac.uk/uploaded_documents/oaismets.pdf.

In addition to conformance with the OAIS reference model, Rosetta supports international industry standards such as the Metadata Encoding and Transmission Standard (METS), Preservation Metadata: Implementation Strategies (PREMIS), Dublin Core, and the Open Archives Initiative Protocol for Metadata Harvesting (OAI-PMH).

The system supports the acquisition, validation, ingest, storage, management, preservation, and dissemination of different types of digital objects that are in various formats and originate from many sources. Consisting of separate yet interactive modules that correspond with the life cycle of a digital object, Rosetta keeps these digital objects secure while allowing institutions to implement multiple digital preservation policies and strategies.

Overview of Rosetta's Scalability

Architecture

Rosetta provides a scalable infrastructure to address the ever-growing need to preserve and manage digital materials. This scalable solution can be implemented on a robust distributed architecture that allows the deposit module, working area, delivery module, permanent repository, and database to be deployed on separate physical or virtual servers. Each module can be scaled with the provision of additional servers.

Crucially, the system can also be deployed on the Rosetta "all-in-one" scalable architecture, in which all modules are implemented on the same physical or virtual server. The all-in-one architecture can be deployed simultaneously on many servers to increase the system's computing power.

In addition, Rosetta enables institutions to add dedicated servers, called workers, to perform specific tasks, such as virus checks and [fixity checks](#) (Figure 2). This flexibility allows an institution to start with a small hardware configuration and expand Rosetta to meet the needs of the institution's growing collection.

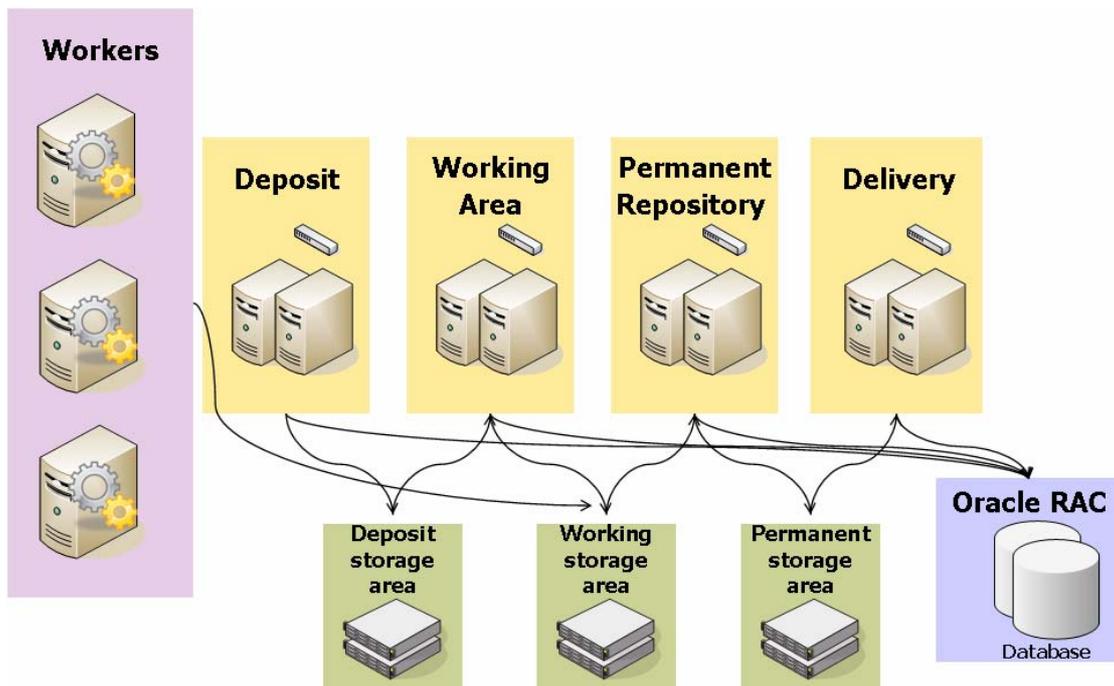


Figure 2. Rosetta's scalable architecture

Beyond the benefit of scalability, the architecture provides institutions with a system that offers redundancy and has no single point of failure. Moreover, a storage abstraction layer allows institutions to attach different storage hardware for each module of the system.

Scaling Proof of Concept Goals

The purpose of this project was to assess the ability of the Rosetta solution to meet the performance needs of an organization-wide preservation implementation. The scope of this effort was defined as an evaluation of the scalability of Rosetta with regard to ingest rate and the number of records a given instance will support. Because Rosetta supports partitioning through a [sharding](#) technique, the Church believed that understanding the performance capabilities of a given [shard](#) would help the organization carry out capacity planning for the larger system. Success of this proof of concept was a necessary precursor to a potential Church-wide solution.

The Church chose to explore Rosetta’s scalability in terms of ingest rate and shard capacity rather than, for example, testing the solution’s storage and search capabilities, in order to achieve the largest possible set of results within a limited time frame and budget.

The objectives of the proof of concept were as follows:

- To determine whether Rosetta would allow 200,000 files to be ingested in a 24-hour period—thus allowing the ingest of two [petabytes](#) (PB) in a year
- To determine whether a single instance (that is, a shard) of Rosetta would accommodate 50 million records—thus allowing one billion records in a 20-instance implementation

Testing Methodology and Criteria

The Church wanted to obtain some baseline benchmarks on an empty database and then benchmark a full database (50 million records) to test the scalability of database performance of the system. The Rosetta all-in-one approach, as specified in the [appendix](#), was used. The following tests were performed to establish whether the proof of concept was successful:

- Using synthetic data, 50 million files were loaded on one instance.
- 200,000 files of 10 KB each were all loaded on one server in a 24-hour period.
- Additional load tests were performed using 200,000 files of 100 KB, 10 MB, or 500 MB each.

The rationale for these criteria is that in order to meet organizational objectives of loading two petabytes of data within one year, the system must be able to load 200,000 files of 10 KB each within a 24-hour period. While the actual file size of the Church’s data varies from 10 KB to one terabyte, with the average closer to 10 MB, the amount of storage needed to perform the test with that much data was cost prohibitive. The Church therefore limited the tests to the objectives specified.

It should be noted that the following conditions existed and may have affected the test results:

- Database and file storage were on a [storage area network](#) (SAN) that was shared with the production environment.
- While the Rosetta application and database servers were dedicated, the SAN storage and [network attached storage](#) (NAS) were not.
- The network connection between the application server and the NAS was shared with the production environment.

Test Results

Over approximately six months, the Church and Ex Libris worked jointly to complete the scaling proof of concept project. The initial results of the tests met the Church's goal; however, it was agreed that the results were not optimal. The Church considered the performance to be significantly slower than in tests run at the Ex Libris lab with a somewhat similar setup, so the team endeavored to improve performance. Over the course of the project, the following technical modifications were made in order to enhance performance:

- Because of a known Oracle 11g performance problem caused by high insert and update rates and related to auto-extent management of tablespaces, the extent management was changed to manual, and free lists and free list groups were added to the Oracle database.
- Additional tuning of the Oracle database and Rosetta database included changing the log checkpoint interval and moving the LOB out of inline storage to its own tablespace.
- The various Rosetta shares were spread across multiple cabinet arrays to avoid read/write contention.
- The file storage, which was originally on SATA disks, was moved to [Fibre Channel](#).
- The Rosetta application server's JVM heap size was increased from the default setting.

The final results of the proof of concept exceeded the Church's expectations in both the ingest rate and the ability of Rosetta to accommodate large amounts of files in a single instance. Figure 3 illustrates the results of the scaling tests.

No. of submission information packages (SIP)	No. of intellectual entities (IE) per SIP	Total no. of files	File size	Total Processing Time	Comments
40	500	200,000	10 KB	4:46	10,000 baseline run 1
40	500	200,000	10 KB	4:05	10,000 baseline run 2
40	500	200,000	10 KB	5:40	10,000 baseline run 3
40	500	200,000	100 KB	5:08	100,000 Baseline run 1
40	500	200,000	100 KB	4:39	100,000 baseline run 2
40	500	200,000	100 KB	4:49	100,000 baseline run 3
40	125	50,000	10 MB	16:12	10 million baseline run 1
40	125	50,000	10 MB	16:16	10 million baseline run 2
10	10	1,000	500 MB	11:55	500 million baseline run 1

Figure 3. Selected test results

Conclusions

The tests conducted by the ICS department of the Church were successful in demonstrating that Ex Libris Rosetta can provide the targeted rate of ingest for files and can accommodate a large volume of data. Specifically, the tests demonstrated that:

- The ingest of 200,000 files of 10 KB each could be accomplished in far less than 24 hours
- One instance of the Rosetta system could easily accommodate 50 million records

Based on these results, the Church will conduct further tests to demonstrate search and retrieval times and the ability of the Rosetta system to scale over multiple instances or shards.

Scalability is a key facet of Rosetta 2.0, enabling institutions to preserve an ever-growing amount of digital material. The proof of concept carried out by the Church and Ex Libris demonstrates Rosetta's ability to manage very large collections and confirms Rosetta as the leader in complete digital preservation solutions.

Appendix: Test Configuration

Technology

The test configuration consisted of the following technology choices:

Hardware

- A single HP BladeSystem c7000 blade enclosure
- 16 HP Intel BL460c G6 CTO blades
- 16 HP BLc QLogic FC HBAs
- 2 Cisco MDS 9124e 24-port fabric switches
- 4 Cisco 3120 blade switches

Storage

- File storage: NetApp V series using EMC SAN Fibre Channel storage, 5 TB
- Database Storage: HBA attached EMC SAN Fibre Channel storage, 3.5 TB

Application Servers

- VMs running in an ESX VMware host environment
- 4 cores, 16 GB RAM per VM host
- Red Hat Enterprise Linux 5 64-bit OS
- JBoss 4.2.3

Database Servers

- 8 cores, 32 GB RAM per RAC node
- SLES 10
- Oracle 11g R1

In addition, the Church was required to modify the custom interim preservation system application used to create a submission information package (SIP) to handle the ingest rate specified. The Church used a custom Java tool from Ex Libris for this purpose. Figure 4 illustrates the infrastructure for the project.

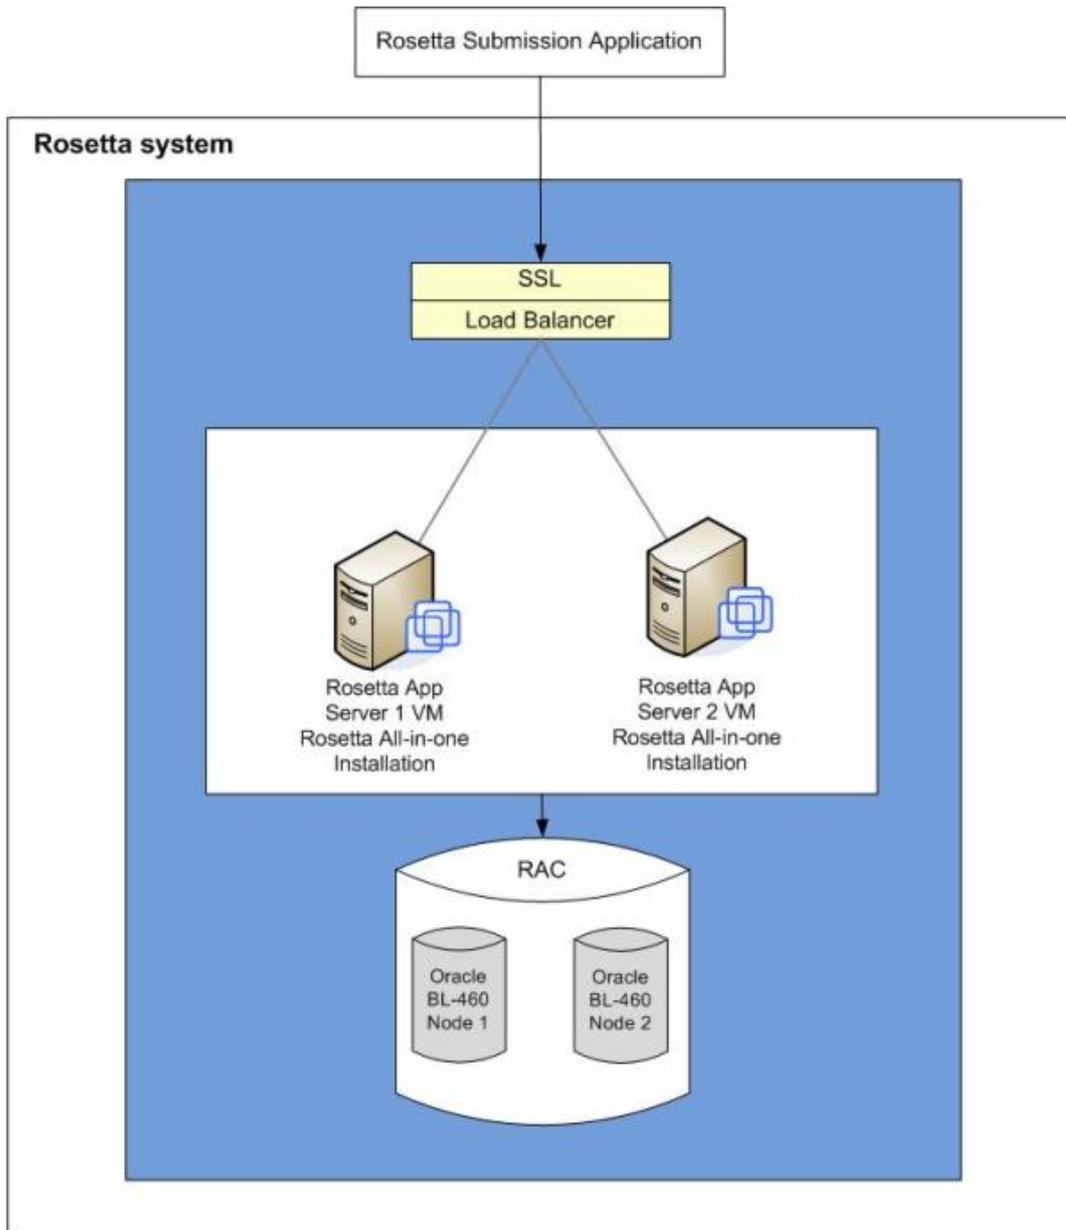


Figure 4. Project infrastructure

Glossary of Terms

Fibre Channel. A standard technology for physical layer networking

fixity check: A process that checks whether a digital object has changed in a given period of time

network attached storage (NAS). Computer data storage that incorporates a file system and is accessible on a computer network

petabyte. A unit of measure for digital information. One petabyte equals 1000 terabytes, or 1 million gigabytes.

sharding. Sharding is method for horizontal partitioning the database across several database instances. An advantage of sharding is that if one database reaches capacity, the user can add another one with a similar structure and use a single entry point or application to work with both of them. A shard is a single instance within that framework.

storage area network (SAN). A local network of multiple storage devices