Application Developers,

It is Time to Adopt XAM

February, 2010

Principal Authors:

Michael Peterson Chief Strategy Advocate for the SNIA's XAM Initiative, and President of Strategic Research Corporation and TechNexxus, LLC

Supporting Authors:

David Martin Chair of the XAM Initiative, and Product Strategy Manager for the Information Management Portfolio, Hewlett-Packard Corp.

Mark Carlson SNIA Technical Council member and Sr. Architect, Oracle

Produced and supported by the SNIA XAM Initiative

Executive Summary

XAM is a standard, open-source, application to storage interface for file sharing similar to CIFS[1] on Microsoft platforms and NFS on Unix-like systems. But, unlike file sharing protocols, XAM also gives applications and supporting storage systems an expanded metadata container, enabling powerful information and retention management capabilities. Expanded metadata is the key for improved information management in today's complex and litigious information-centric world. Today's datacenter will soon be an operating environment in which metadata needs to be controlled at the same level as the data. This means both have to be retained, deleted, protected, secured, shared, and preserved as a unit along with associated audit logs to verify authenticity of information and to mitigate risk. The need for these capabilities is amplified in domains that have to retain and preserve information assets for either legal purposes (as in eDiscovery, litigation support, or for litigation holds) or for extended periods of time (as in regulatory compliance or for archival purposes). Going forward, retention and preservation services in the datacenter will require the extended metadata that only a metadata container like XAM's accommodates. This means that XAM is strategic to applications because there is established customer need and a proven opportunity to empower the business' retention and information management practices from those applications. Application developers that integrate XAM and take advantage of the enhanced information management and data services that XAM enables will have a market advantage.

"The XAM API enables businesses to enhance custom applications and easily realize a new paradigm in information management and operations."

Figure 1

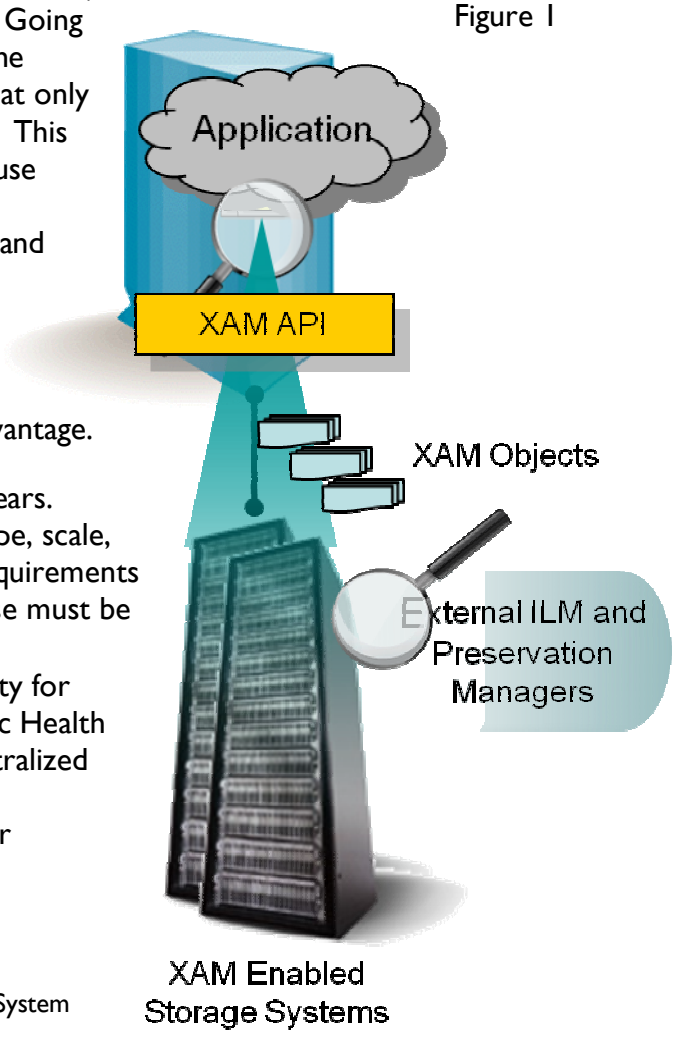

To illustrate this point, just look forward a few years. Today's metadata methods are inadequate in scope, scale, capabilities, support, and portability for future requirements and have to change. New paradigms such as these must be supported:

- The need for secure information portability for important programs such as the Electronic Health Records (EHR) system that envisions centralized and shared medical records accessible, confidential, and preserved for 15 years or more.

1 CIFS: Common Internet File System, NFS: Network File System

- The need for expanded metadata carrying business and application requirements that enable automated information and retention management practices across the datacenter and the organization.
- The need for protection of the authenticity of information throughout its lifecycle for compliance and litigation support. Metadata will become critical in verifying authenticity of evidence.
- The need for contextual search and discovery practices that include the metadata associated with specific data.

Figure 2

XAM and Metadata

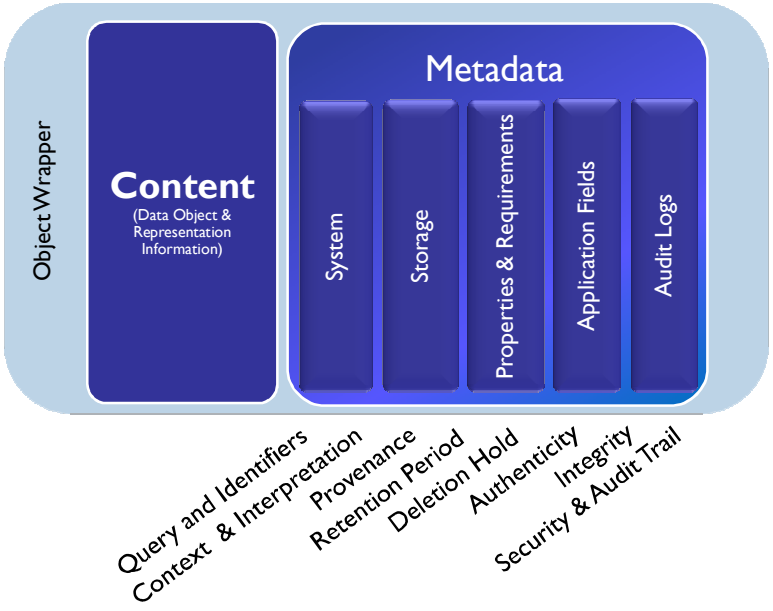

What is XAM?

For those of you not familiar with XAM, it is implemented at the application level as an API similar to block file system APIs. The analogy of a file sharing system similar to CIFS/NFS[2] is appropriate, but it is better described as an application to storage interface and access method. The two key places that XAM is different from a file system are these:

1. To aid the application, XAM abstracts the access method from the storage, a capability necessary to create location independence, policy-based automation, and portability. A side benefit is storage interoperability across storage systems that support XAM.
2. Physically, XAM wraps the data into what is called an "XSet," similar to an XML wrapper. The data or files as they are written to storage are placed into an XML-like information object with rich metadata fields incorporated into the object as illustrated in Figure 2. These expanded metadata fields are able to accommodate various metadata standards and application specific requirements useful for data and information management services in any environment.

"The value of XAM is that it solves the access method conundrum by replacing the many proprietary interfaces we have to develop today. The current storage interfaces are all very different for each vendor. Standardizing the interface is a big deal to us."

2 Thus the name, eXtensible Access Method

Why Adopt XAM?

As an application developer, XAM provides differentiation plus empowers internal operating cost reductions by reducing storage system support overhead. Even better it provides capabilities consistent with your customer's expanding information management requirements and needs. For example, through XAM's expanded metadata fields you can integrate value-add compliance, preservation, eDiscovery, risk management, or retention and information lifecycle management features into your application. The following examples expand these points:

XAM Value Propositions for Application Developers

- **Differentiation and Market Expansion:**
  - XAM allows applications and information management utilities or management consoles to differentiate their products and add valuable new services to address reducing the ever increasing burden of high overhead costs due to legal and compliance requirements and ensuing risk. Examples include retention and deletion management, preservation services, ILM-based practice automation, litigation support services, medical records, etc.
  - XAM provides the ability for application and management developers to define new storage services and features without coding specifically for each storage system they support. The result includes benefits such as storage interoperability and federation, portable information objects for secure information sharing and centralization, enhanced preservation services for long-term retention or litigation support, and better retention management including permanent deletion practices and enforcement of litigation holds.

> ***"As metadata becomes more critical in compliance and litigation support, so does XAM. As preservation practices become more important for litigation support, compliance, health care, or even protecting our digital history, so does XAM. Looking to the future, it is all about the metadata and our ability to use it effectively!"***

- **Enable New Functionality:**
  - Migration: XAM provides the basic functionality of a preservation object, providing the ability to control authenticity and integrity through logical and physical migration processes. These services are essential for long-term retention and preservation.
  - Regulatory Compliance: Improve, simplify, and cost reduce records retention practices due to integrated retention and disposition metadata.
  - Retention and Deletion Management: Each XAM object carries its own retention period and disposition metadata supporting automation.
  - eDiscovery: XAM allows data and metadata to be searched allowing objects to be better discovered for litigation review. Application-independent structured discovery avoids the impact of application obsolescence.
  - Litigation Support: Improved verification and control of authenticity, improved and cost reduced ability to preserve information and control litigation holds.

- Security and Legal Risk: Reduced exposure due to improved retention management including litigation holds and deletion practices supported by expanded metadata.
- Electronic Health Records: XAM provides secure information portability, supporting centralization and interoperability for the sharing and preservation of records over their extended lifecycles.
- Digital Libraries and Preservation Stores: XAM based objects meet many of the criteria needed for a preservation object and support typical preservation services that need expanded metadata.
- Cloud Services: XAM is cloud friendly and can be mapped over cloud interfaces such as SNIA's Cloud Data Management Interface (CDMI) to provide a full and rich metadata container that supports enhanced information portability and control. Retention periods, billing, secure information sharing, and verification of integrity and authenticity are good examples of types of services that need to map into the cloud.

- **Direct Application Developer Benefits:**
  - Reduced application development and maintenance costs, including savings in storage interface development, access to multiple storage interfaces, and storage interoperability
  - Improved storage system interoperability: Applications can work with any XAM conformant storage system allowing information to be migrated, relocated, and securely shared
  - Provide enhanced standards-based functionality that is of value to your customers

What Will Help XAM Adoption Move Ahead Faster?

As with any new and similar technology, broad market adoption requires time for demand to develop. XAM adoption requires that application developers see benefit in using XAM for differentiation and value-add information management features. Similarly, application developers need to help storage vendors see benefit in supporting XAM to become preferred partners with the application developers as well as for differentiation. And finally, IT and Legal need to begin asking their vendors for this functionality because they need it to better manage risk and control costs.

XAM is in early stages of market adoption. While XAM activity within SNIA is 5 years old now, the specification was just released in July of 2008 and is now in standardization through ANSI (expected completion in 2010). Standards like this typically take 6 to 8 years for market adoption to move through the early adopter phase. That said, there are many examples of XAM adoption picking up in the marketplace as the importance of the drivers and the strength of the value proposition elevate the strategic importance of XAM.

How to get started using XAM

Today, there is a growing community of experienced developers around XAM and a growing number of development resources available. Here is a listing of what is available to XAM developers:

- Download the current specification and the supporting technical documents from links at (http://www.snia.org/forums/xam/technology)
  - XAM Part 1: Architecture – The starting point and reference for storage vendors creating vendor interface modules (VIMs) and application developers writing APIs.
  - XAM Part 2: C API – a reference document for 'C-based' application development.
  - XAM Part 3: JAVA API - a reference document for 'JAVA-based' application development.
- Download the software developer's kit, SDK, and some of the sample applications. The SDK also contains many supporting documents. (http://www.snia.org/forums/xam/technology/software)
- Participate in the developer's community where you can get access to a number of SNIA experts for assistance as well as download resources that will help you – (http://groups.google.com/group/xam-developers-group)
- Participate in the XAM Plugfests to see how others are implementing XAM and to test your own code (http://www.snia.org/forums/xam/news/plugfest_reg/)
- Talk with your favorite storage vendors to evaluate their implementations, test yours, and to encourage others to support XAM

"XAM is a new kind of interface to storage that, for the first time, includes a standard for extensive metadata that can be interpreted by data services in the storage device. Prior to this, all vendor interfaces for this type of storage were proprietary, which limited the market for these types of devices."

Recommendations

SNIA is working to enable and promote XAM or XAM-like functionality to be adopted by the industry. Expanded metadata containers must come to market to meet the types of requirements identified in this paper. For example, today in the long-term retention and digital library community, a standards-based preservation object is something the industry really needs. In the Health Care Industry, centralized, secure, and accessible medical records are a vision that will not succeed without portable information objects supporting preservation services. Without XAM, the market only has proprietary storage access models and continued information silos, no interoperability, no information portability, no ability to scale discovery with context, no ability to automate information management, and no ability to broadly verify and control authenticity. If the future unfolds as expected, XAM's functionality is essential.

Of all the options on the table, XAM remains the best and only interoperable choice. It is our recommendation that you evaluate it and begin implementation in environments requiring these attributes and storage interoperability. The differentiation XAM enables can be strategic to your business.

About the XAM Initiative

The core mission of the XAM Initiative is to drive adoption of the eXensible Access Method (XAM) specification. The Initiative serves a XAM community that includes storage vendors, independent software vendors, and end users to ensure that the specification fulfills market needs for a reference information management interface standard. For more information, visit www.snia.org/forums/xam or participate in our open online community – (http://groups.google.com/group/xam-developers-group)

About the Storage Networking Industry Association:

The Storage Networking Industry Association (SNIA) is a not-for-profit global organization, made up of some 400 member companies spanning virtually the entire storage industry. SNIA's mission is to lead the storage industry worldwide in developing and promoting standards, technologies, and educational services to empower organizations in the management of information. To this end, the SNIA is uniquely committed to delivering standards, education, and services that will propel open storage networking solutions into the broader market. For additional information, visit the SNIA web site at www.snia.org