



The Power of XAM

July, 2006

**By: Christina Casten,
Chair XAM Committee, SNIA**



The Power of XAM
Christina Casten

In late 2005, the Storage Networking Industry Association (SNIA) voted to accept a contributed technical specification called the *eXtensible Access Method* (XAM) as a starting point for development of a standard interface and interoperability between applications and object storage devices characterized as CAS, Content Aware Storage systems. The XAM specification is profoundly important for the future of information-based management. Here's why.

- A standard application to storage interface allows interoperability between CAS systems with all the incumbent benefits
- Standard metadata semantics empower automation of Information Lifecycle Management, (ILM), based practices such as retention, archive, discovery, integrity, authentication, availability, protection, and secure-deletion. Standard metadata also resolves issues with Grid technology and even long-term archive
- Metadata is essential for the automation of ILM practices to reduce operating costs and improve the ability to conform to compliance, legal discovery, and reduce security risks

The alternative to life without XAM is not pretty. The ever growing chaos and complexity of owning and managing trillions of data objects is insurmountable with conventional approaches based on discrete, and often competing, management processes. Even the promise of ILM is incomplete without a robust metadata standard. When complete, XAM will have a positive impact on all types of information systems.

XAM – How it Works:

The XAM Interface specification defines a standard interface (access) method between “Consumers” (application and management software) and “Providers” (storage systems) to manage information. XAM annotates objects with metadata providing for the management of information at a semantic level. This coupling allows external ILM-based policy services to make intelligent decisions about the management of objects without referring back to the application and without impacting the application.

As an interface, XAM abstracts the access method from the storage. This supports the mobility of information independent from the storage to allow longevity, distribution, and management of information. The XAM Interface is intended to achieve interoperability, storage transparency, and automation for Information Lifecycle Management-based practices, long-term records retention, and information assurance (security).

In contrast, existing applications and CAS products use unique, incompatible access and management methodologies (e.g. naming, retention, expiration and placement of content) which require that each application's interface be modified, often extensively, to operate with any particular storage platform. No provisions are made for information sharing between different applications. No standards exist for movement of content between different CAS-type storage platforms. And, unlike file systems, XAM tags information with metadata and provides a

technology independent namespace, allowing software to interpret the content independent of the application.

To better understand the power of XAM-based standard metadata semantics consider the following illustration. What is the difference between specialized grocery stores with shelves stocked with unlabeled cans versus the second scenario of standardized labeling?.

Impact of XAM on ILM



Which is the corn? **Go to the corn store or Open the can.**
 What does it cost? **Ask the clerk.**
 How many calories? **Ask the vendor**
 How does the store automatically manage inventory? **You can't**

With Metadata Standards & Interoperability



Standardized labeling allows multiple vendors to consistently represent information to consumers

| Nutritional Facts | |
|--------------------------------|---------------|
| Serving Size 1/2 cup (130g) | |
| Servings per container about 3 | |
| Amount per serving | |
| Calories 130 | Fat Cal 5 |
| | % Daily Value |
| Total Fat 0.5g | 0% |
| Saturated Fat 0g | 0% |
| Cholesterol 0mg | 0% |
| Sodium 260mg | 11% |
| Total Carbohydrates 22g | 7% |
| Dietary Fiber 5g | 22% |
| Sugars 0g | |
| Protein 10g | 20% |
| Vitamin A 0% | Vitamin C 0% |
| Calcium 4% | Iron 10% |

Extended labeling for Lines of Business Users

One of the major areas of impact for XAM technology will be on the adoption and implementation of Information Lifecycle Management-based practices. For example, XAM provides an object-oriented location and storage technology-independent approach to information storage. By tying information to a globally unique name (a positive characteristic of object storage methods), ILM-based practices can efficiently manage the information without application concern for a specific

physical location of that information or technology it resides on. XAM also raises metadata – contextual information about the information being stored – to the same level of importance as the information itself. By bundling information and metadata together, applications can more effectively implement ILM and share information. And, XAM metadata allows other ILM-based policy services to make intelligent decisions about the management of objects without referring back to the application – the ultimate in automating ILM-practices.

ILM-based practices begin with the collaborative process of information classification then move to a step of assigning service requirements and setting policies called data classification. Business requirements get translated into IT infrastructure operating policies. An example is long term retention. The business requirements may be dictated by a regulatory constraint or a legal risk. In either case, the information of the business that has that requirement placed on it, demands it be managed to meet those requirements over its lifecycle in a cost efficient manner. Yet, how will you assure long term readability – longer than the life of the storage systems themselves? When implemented, XAM compliant applications, storage systems, and ILM practices will make adhering to business requirements much easier. By abstracting the physical aspects of storage and giving information objects metadata, the underlining storage infrastructure can independently age and run the course of its product lifecycles independent of the stored information. ILM-based practices will utilize the XAM-based architecture to instrument and automate migration and provide a consistent view of the managed information and all of its attributes. This capability is extremely important as the adaptive datacenter will also be able to age and retire physical assets while preserving the various services through other abstractions for the different types of resources.

XAM and Grids

An example of another important role of XAM is in the emerging concept of *Data Grids*. A *Data Grid* can be looked upon as a data resource location service that provides a unified view of many, if not all, information repositories or sources within a predefined domain. This can include databases, file-level and object-level network services, and reference information archives. XAM can contribute the metadata framework for this important new information architecture, enabling enterprise wide data mining and analysis, enterprise wide digital asset discovery, shared information repositories, and many other purposes.

Status of XAM Development

Recognizing the industry demand and importance of these benefits, more than 95 individual members from 34 companies representing storage vendors, application providers, end users and academic leaders are contributing to the XAM development effort within SNIA's *Fixed Content Aware Storage Technical Work Group* (FCAS TWG). The current plan is to publish the initial version of the specification in late 2006 and begin demonstrating XAM-based implementations at *Storage Networking World 2007*.

As customers, utilizing applications and storage systems, you will benefit immensely from XAM. You can help this process along by asking your vendors now for XAM support. Put pressure on

them to deliver. That will keep them engaged in the standards development process and in achieving interoperability.

By: Christina Casten
Chair of the SNIA's XAM Committee
July 2006
www.snia-dmf.org/xam

SNIA and the DMF

The Storage Networking Industry Association's (SNIA) Data Management Forum (DMF) is a cooperative initiative of IT professionals, integrators and vendors working to define, implement, qualify and teach improved and reliable methods for the protection, retention and management of electronic data and information. The Data Management Forum hosts three initiatives focused on data protection, information lifecycle management (ILM), and long term archive and compliance. Within the long-term archive initiative is a special interest group working to conduct market development for XAM. Overall, the DMF organization operates as a market development and educational service for the SNIA. It is currently spearheading two technical working groups within SNIA, hosts the "*Enterprise Information World*" Conference, participates world-wide in events and conferences speaking on these topics, produces and delivers tutorials and best practices training materials, and is the authority in the broad field of data and information lifecycle management. Participation is open to individuals and member companies at www.snia-dmf.org.