

The IEEE Storage System Standards Reference Model Current Status

John L. Cole

Army Research Laboratory
APG, Maryland 21005-5067
+1-410-278-9276, FAX +1-410-278-9276
cole@arl.mil

Abstract

The IEEE Reference Model for Open Storage Systems Interconnection (OSSI) approved September 1994[1] and Version 5 of the Mass Storage System Reference Model (MSSRM) introduced in April 1993[2] at the Twelfth IEEE Symposium on Mass Storage Systems are compared. The status of standards development for components of the Model is given. And a brief description of the Reference Model is presented.

Introduction

While storage devices and media continue a decade of dramatic improvements in capabilities and capacities, storage *systems* have just developed in recent years and are now beginning to advance[3] based on the Reference Model.

The OSSI Model previously known as the Mass Storage System Reference Model (MSSRM) Version 5, provides the conceptual and functional framework for development of standards describing application and user interfaces of open storage systems.

The Reference Model identifies the high-level abstractions underlying modern storage systems. It defines common terminology and concepts allowing the architectures of existing and future systems to be described and compared.

Success of the Reference Model is demonstrated by the influence of the MSSRM Version 5 on the architecture of the National Storage Laboratory's High Performance Storage System (HPSS) project.

The SSSWG, developer of the Model, first met in July 1990. During its first year or so, the SSSWG produced the initial version of the Model, which then evolved through five versions. In September 1994, the SSSWG finished modifications to the Model, renamed it the OSSI Model, and voted to approve Version 5 of the OSSI Model for public review. Seven model components and their respective standards are being developed, with plans to submit several standards to the IEEE balloting process in 1995. A Version 6 of the Model will be produced to incorporate public review, correct inconsistencies in Version 5, and to realign the model as needed with the standards developed.

Overview

It is important to begin with a distinction between the Model and its associated standards. No conformance to the Model is possible, because it will never be balloted as a standard. Standards associated with the Model are being developed, and conformance to these will be possible, especially since they will allow the greatest possible implementation variation.

The scope of the Reference Model is “open storage systems” and the environments in which such systems must function. A “storage system” is a set of mechanisms necessary to store and retrieve a stream of bits. And an “open storage system” is one which complies with the requirements of OSSI standards. The term “open storage system” is an abstraction used to describe storage system implementations. Such implementations must also comply with the OSSI standards in their interactions with clients in order to be “open”. The implementation of a storage system is achieved through a set of devices, associated software, operators, physical processes, etc., which serve to store and retrieve information. The domain of standards applicable to information storage services is defined by the Reference Model.

Openness

Through the mutual use of standards within this domain, two or more storage systems may inter-operate openly.

However, no particular system implementation, technology, or connectivity is required by the Reference Model or its standards. Nor does the Reference Model, by itself, specify the detailed and precise functioning or internals of an open storage system. The Reference Model does not serve as an implementation specification, the basis for appraising the conformance of implementations, or as a precise definition of the standards for services and protocols of the interconnection architecture.

The technology and application independence of the Reference Model accommodates descriptions of advanced technologies and expansion in user demands. This flexibility also supports the phased transition from existing implementations to OSSI standards. The Model’s conceptual and functional framework allows teams of experts to work productively and independently on the development of standards for OSSI.

Modeling Method

The SSSWG developed the Reference Model through a process which identified storage system services as partitioned, disjoint sets for which external behavior could be described and against which commercial products could reasonably be developed.

An abstract storage system was decomposed into Modules of services, Storage Objects, and an Application Environment.

By this method areas are exposed where new storage system standards are needed or where standards need improvement through development of the Reference Model. The Reference Model enables the coordination of storage system standards development and provides a common perspective for existing standards.

Application Environment

The variety of requirements of specific systems means that storage system implementations will operate in diverse environments. In an open system, the application environment provides common services such as those involving communication, services, location, and security. These common services rely on mechanisms governed by commercial availability, management policy and implementation-dependent architecture requirements. Site-specific constraints will also place distribution and scalability requirements on the open storage system implementation. The Reference Model permits an abstract view of the open storage system independent of these environmental, distribution and scalability contexts.

Moreover, the Application Environment permits a variety of commercial implementations to be conformant to standards through accommodation of vendor preference profiles.

Client Service Modules

All services in an open storage system are coupled to one or more objects, and these abstract the fundamental elements within storage system implementations. The Reference Model describes services which report or affect the state of storage objects and their relationships. The term “object” is used as an abstraction representing an individual and identifiable entity with a well-defined role in a storage system.

Modules of the Reference Model provide disjoint sets of services, and define the client’s view of the Model. Clients request services from modules through programmatic interfaces supported by services in the environment (e.g., communications and security). A module can act as a client and request services of another module.

Modules execute independently of their clients, responding to requests for services and optionally notifying clients of events associated with an object.

Clients and modules are implemented by a collection of hardware and software.

The client’s view of the service interface is effected through an abstract communications service which accepts requests and optionally provides return replies. A non-exclusive list of possible implementations includes subroutine call mechanisms with event notification, message passing, and remote procedure calls.

Key Concepts of the Reference Model

ABSTRACTION. The OSSI Model is based on abstraction. It must be recognized that concepts contained in the description of open storage systems are abstract, despite a similar appearance to elements of storage system implementations.

TRANSPARENCY. The OSSI Model allows many varieties of transparency including device, location and replication transparency.

- Device transparency supports device independent programmatic interfaces.
- Location transparency eliminates the need for clients to know the actual location of stored information or services.
- Replication transparency allows clients to remain unaware that replicas of their stored objects exist.

SEPARATION OF POLICY AND MECHANISM. The OSSI Model separates policy from supporting mechanism. For example, the Reference Model does not specify recovery policies, but it does provide mechanisms to define, associate and execute such policies.

The term “policy” is used as in standard English; a course of action, guiding principle, or procedure considered expedient, prudent, or advantageous.

SEPARATION OF CONTROL AND DATA FLOWS. The OSSI Model distinguishes control flows from data flows occurring between a client, a data source, and a data sink. Control flows carry requests, replies, and asynchronous notifications between a client and the data source or sink device. Control flows between the data source and data sink carry source-sink protocol information to manage the flow of data. Data flows pass only from a source to a sink. By logically separating control and data flows, the OSSI Model offers the possibility of optimizing each flow through separate implementation.

THIRD-PARTY TRANSFERS. The OSSI Model allows data to flow directly between independent sources and sinks, under the control of a third party, the initiating and controlling agent or client. Each entity separately performs operations such as data flow control, error reporting, or initiating and terminating the transfer.

LAYERED OBJECT NAMING. The OSSI Model presents a uniform name space for all objects. Clients may build on this uniform naming scheme to construct arbitrary naming schemes.

HIERARCHICAL STORAGE MANAGEMENT. The OSSI Model supports the optional creation and automated management of storage hierarchies, as well as the creation and association of policy to manage movement of data within the hierarchy. A storage hierarchy is an open storage system that provides multiple levels of service (performance, reliability, availability, etc.) and provides for movement of data among levels of service according to management policy. Examples of these policies include caching policies for movement of data to a higher-performance store, and policies for movement of data to a lower cost store.

SCALABILITY. The OSSI Model is intended to describe open storage systems of any size. Therefore, the OSSI Model does not specify any physical size or limit. Storage system implementations, based on the OSSI Model, can use different communication, location, and naming structures for differently sized implementations.

Within the OSSI Model's environment, storage system management supports system scalability with services to create storage object groups. These groups can be associated with distinct management policies allowing autonomous operation.

DISTRIBUTION. The OSSI Model does not assume or specify any specific framework for distribution or centralization. Therefore, the OSSI Model can describe open storage systems that are distributed as well as centralized.

SYSTEM MANAGEMENT. The OSSI Model prescribes a standard framework for extensible system management. This extensibility allows management of storage systems at a system level, as well as definition and imposition of policy within individual modules and storage objects.

EVENT NOTIFICATION. The OSSI Model provides event notification services to loosely couple storage system management to the storage modules and objects. Clients may register for notifications associated with storage system events, such as a storage object state change. Modules can post information (i.e., notifications) for registered clients when these events occur. Multiple clients may register for the same notifications.

Programmatic Interfaces And Environment

IEEE P1244 standards, which define open storage systems, consist of two basic components: programmatic interfaces for client access to open storage system services, and a detailed identification of the environment for the open storage system. The standardization efforts, within Project 1244, address this definition through the development of standard Application Programming Interfaces (API) for each module, as well as a complete definition of standard environmental profiles.

- Application Environment Profile (AEP) - specifies the environmental software interfaces required by open storage system services.
- Object Identifier (SOID, 1244.1) - defines globally unique, time/space immutable object identifiers within open storage systems, and the format and algorithms used to generate them.
- Physical Volume Library (PVL, 1244.2) - defines software interfaces for services that manage removable media cartridges and optimize drive use within a storage system.
- Physical Volume Repository (PVR, 1244.3) - defines human and software interfaces for services that stow removable media cartridges and selectively mount removable media cartridges onto drives.
- Data Mover (MVR, 1244.4) - defines the software interfaces for services that transfer data between two endpoints.
- Storage System Management (MGT , 1244.5) - defines a framework that permits the development of consistent and portable services to monitor and control IEEE P1244 storage system resources as motivated by site-specified storage management policies.
- Virtual Storage Service (VSS, 1244.6) - defines software interfaces to access and organize persistent storage presented as a single virtual storage image.

Functional Overview

As a framework for storage systems design, the Reference Model decomposes a complete storage system into storage objects, storage modules, and an overall computing environment (Figure 1). Storage objects and their respective modules are introduced:

- Devices are storage objects that copy data to and from storage media. The Mover is a storage module that presents the Model-defined interface to devices to manage the transfer of data between source and sink devices.
- Cartridges are storage objects that contain storage media. The Physical Volume Repository (PVR) is a storage module that stows cartridges and mounts these cartridges onto devices, employing either robotic or human transfer agents.

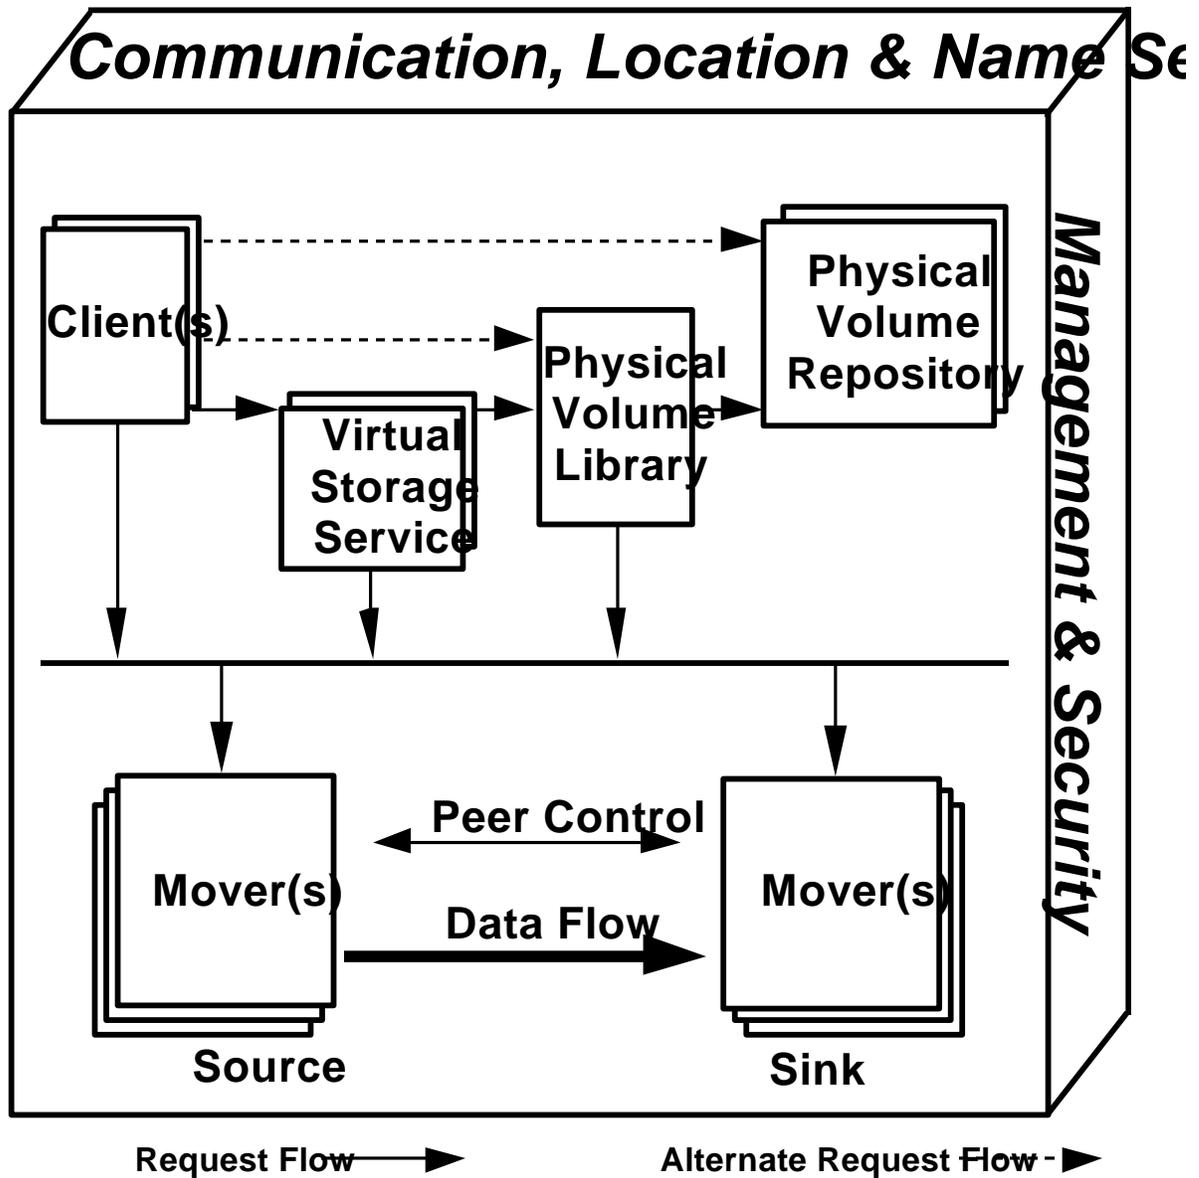


Figure 1. Reference Model for Open Storage Systems Interconnection

- Physical volumes are storage objects that provide an abstraction of data storage media. Physical volumes may have transient or fixed associations with devices. An association describes a group of links with common structure and semantics. A link is a physical or conceptual connection between object instances. These associations provide a foundation to construct device and media independent storage services consistent with the OSSSI Model. The Physical Volume Library (PVL) is a storage module that creates and utilizes these associations providing the capability to execute location independent cartridge mount/dismount requests.

- Stores are storage objects that provide addressing and transfer models for either physical volumes or other stores. A Virtual Storage Service (VSS) is a storage module that creates stores and builds associations between stores and physical volumes or other stores. The VSS uses these associations to translate access requests into sets of requests to the appropriate Mover(s) and, if the associated physical volume is not mounted, to the appropriate PVL.

The storage system environment encompasses the storage-system-wide management and security frameworks, and the communication, location and name services of the OSSI Model. The following paragraphs briefly describe the services provided within each of the frameworks:

- Storage system management is the collection of functions responsible for control, coordination, monitoring, performance and utilization of the storage system. The VSS, PVL, PVR and Mover use common storage system management mechanisms to monitor and control storage system resources in conformance with site-specified management policies.
- The security framework is the collection of functions associated with authentication of principals, authorization of principals to access storage objects and enforcement of access policies.
- The communication service connects all storage modules and their respective clients.
- The location service maps a Standard Object Identifier (SOID) to its abstract location in communication space.
- The name service associates arbitrary, structured storage object names with object identifiers for client convenience.

Status and Future of the SSSWG

The IEEE SSSWG held its first meeting in July 1990. Now in its fifth year, the SSSWG is producing working drafts of standards, and subjecting these to public review. Standards development is notoriously slow, but given that the SSSWG has had to work in the context of abstraction, the SSSWG has progressed rapidly in its work. Without pre-existing implementations to standardize, the work of the SSSWG has been unusually difficult. In contrast, standards traditionally have been developed to bring consistency and reliability to existing implementations.

In August 1994 the SSSWG chair passed to two individuals instead of a single chairperson. For those taking up the duties of the chair time was required and used in “learning the ropes” and in sharpening the focus of the SSSWG toward the goal of producing standards in 1995.

Plans are made for the SSSWG's work in refining the Model and developing standards in 1995 and 1996. After that, the sunset of SSSWG is contemplated.

SSSWG information is now available on the World Wide Web at URL <http://www.computer.org/SSSWG.html>. This information includes: SSSWG meeting agenda, minutes, announcements; points of contact within the SSSWG; working drafts of standards; the Reference Model; a history of the effort to standardize storage systems; pointers to other, relevant Web pages; and the electronic mail traffic of the SSSWG reflector.

SSSWG Issues, Discussions

At this writing several issues for discussion are before the SSSWG. One of these is the issue of self-consistency and style among the subcommittee (model component) standards. For example, should the standards settle on the use of structure passing versus function calls in their Application Program Interfaces (APIs)? Plug-and-Play (PnP) issues related to the standards are debated. And base data typing for APIs is a third discussion. These issues are being weighed through a mechanism which the WG has chosen to call a "lightweight RFP"¹.

SSSWG Online

The SSSWG has an electronic mail reflector, an anonymous FTP site, and will have a World Wide Web home page by May 1995. These may be accessed at the following addresses:

Reflector: ieee+mss@larc.nasa.gov
(to SUBSCRIBE, use ieee+mss-REQUEST@larc.nasa.gov)

FTP: [swedishchef.lerc.nasa.gov:/mass_store/ossi*](ftp://swedishchef.lerc.nasa.gov:/mass_store/ossi*)
(postscript files of the Model)

WWW: <http://www.computer.org/SSSWG.html>

Status and Future of the Reference Model

The Reference Model was approved for public review by internal SSSWG ballot on September 8, 1994.

¹ The term "RFP" means "request for proposals" and is normally used in procurement actions of the Federal Government. Here it is used to mean simply a request, made on the SSSWG mail reflector, for responses to position taken.

Far from being OSSIfied², it is expected that a sixth version will be produced to modify the Model's components in order to remain in agreement with the related standards, as well as to correct inherent errors and inconsistencies in the version 5 draft Reference Model. Public review of the Model will be captured and, to an extent, integrated with the Model. A sixth version would not represent a major departure from the current fifth version.

The Model will be retained, not as a standard, but as a guide.

Status of Model Components and Standards

Presented are the significant changes to the Model and to each of the seven Model components in the period between the introduction of the Model (April 1993) and the approved draft of the Model (September 1994). Additionally, plans for standards development related to each Model component are noted.

- Application Environment Profile (AEP)

This component does not have an IEEE Project Authorization Request number, and the component group is non-existent, lacking a chair or members devoted to the component development. As a consequence, there are no immediate plans to develop draft standards related to the AEP. This may result in functionality of the AEP being absorbed into one or more of the active component efforts. No significant changes to the AEP are reported for the period between the Model's introduction and approval.

- Storage Object Identifier (SOID)

The development of this component and its standard(s) is delayed for lack of a chairman and group. As with the AEP component, the functionality of the SOID component may be absorbed to some extent by one or more of the active component efforts. The feeling among members of the SSSWG is that this component does not need to be made standard. Since it remains an important area, though, a white paper position on the SOID may be published by the WG. No significant changes to the SOID component are reported for the period between the Model's introduction and approval.

- Physical Volume Library (PVL)

By September 1995, the PVL committee plans to draft and release the PVL Service API. The API will describe synchronous and asynchronous semantics, POSIX-style security,

² "OSSIfication" - A tendency toward or a state of being molded into a rigid, conventional, sterile, or unimaginative condition.

and an extensible architecture. In addition, the PVL committee is driving the base data types and consistency requirements proposals, which are two technical issues with implications for all of the 1244 subcomponents.

The most significant change to PVL reported for the period between the introduction of the Model and its approval has been to delete the media set object.

- Physical Volume Repository (PVR)

The PVL group plans to have a ballotable standard in 1995.

No significant changes to this component of the Model were reported for the period between introduction of the Model and the approval of the model.

- Data Mover (MVR)

The MVR committee has been working on and reviewing Mover requirements and the Mover API. The High Performance Storage System (HPSS) Mover API has been used in development of the SSSWG MVR API.

Note that there are actually two interfaces which could (and should) be standardized. First, is the client/server interface which the PVL and VSS use to submit requests. The second is the interface between Movers used to effect data transfer. The MVR group has focused on the client/server interface first, and a related draft standard will be produced in 1995. The second interface will be dealt with later. The scope and functionality of the Mover are little changed from Version 5 of the Reference Model.

- Storage System Management (MGT)

The acronym for this group has changed from "SSM" to "MGT".

The MGT committee succeeded in completing the Model component for MGT, and developed an outline for a standard before the committee efforts ceased.

There were no significant changes to the MGT component between introduction of the model and its approval.

At this time the MGT committee is not active. The other SSSWG committees have determined that it is appropriate to further develop their standards and then to restart the management standard effort to reflect and incorporate their progress. At some point in the future, the management committee will either be restaffed or its function will be absorbed by other committees.

- Virtual Storage Service (VSS)

Robert Baird, Hewlett-Packard Company, the VSS Chair, has published a “Virtual Storage Architecture Guide (VSAG)”[4], which expands on the virtual storage concepts of the Model.

No significant changes were reported for the VSS component in the period between the introduction and approval of the Model.

Summary

In its fifth year, the SSSWG has produced approved draft Version 5 of its Reference Model, and has embarked on development of draft standards for balloting. Although the model itself is subject to revision in order to track with standards developed, and the MGT module is the most susceptible to realignment with the other modules.

The challenge remains for the SSSWG to develop standards for the Reference Model components in time to influence commercial implementations of storage systems. Several of the “dot” groups (model component groups) plan to produce ballotable standards in the remainder of 1995.

Acknowledgments

The IEEE SSSWG membership represents a cross-section of computer hardware and software vendors, corporations, academic institutions, and government agencies involved with storage systems. The “Contributors” section of the OSSI Model recognizes all IEEE SSSWG members that, at some time, obtained voting rights. The “Executive Contributors” section identifies the organizations that supported the members’ involvement in the IEEE SSSWG.

References

[1] “Reference Model for Open Storage Systems Interconnection - Mass Storage System Reference Model Version 5”, developed by the IEEE Storage System Standards Working Group (Project 1244), September 8, 1994 (available on the World Wide Web at URL <http://www.computer.org/SSSWG.html>). Richard Garrison, Martin Marietta Corporation, Technical Editor.

[2] R.A. Coyne and H. Hulen, “An Introduction to the Mass Storage System Reference Model, Version 5”, Proc. 12th IEEE Symp. Mass Storage Systems, IEEE Computer Soc. Press, Los Alamitos, Calif., Apr. 1993.

[3] R.W.Watson and R.A.Coyne, "The National Storage Laboratory (NSL): Overview and Status", Proc. 13th IEEE Symp. Mass Storage Systems, IEEE Computer Soc. Press, Los Alamitos, Calif., Jun. 1994.

[4] R. Baird, "Virtual Storage Architecture Guide (VSAG)", 1312th IEEE Symp. Mass Storage Systems, IEEE Computer Soc. Press, Los Alamitos, Calif., September 1995.