The IEEE Storage Systems Standards Effort

John L. Cole, United States Army Research Laboratory, Aberdeen Proving Ground, Maryland USA

ABSTRACT: *A quick tutorial of the IEEE Reference Model for Open Storage Systems Interconnection (OSSI), better known as the Mass Storage System Reference Model (MSSRM) is presented, followed by a brief description of the IEEE Storage System Standards Working Group (SSSWG, Project 1244), which produced the Model. Current activities of the SSSWG and its plans are discussed. The best and current information may be viewed at URL http://www.arl.mil/IEEE, including copies of various documents relating to the Model.*

Introduction

The IEEE Reference Model for Open storage Systems Interconnection (OSSI) [1], better 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; on the design of the National Center for Atmospheric Research's Mass Storage system; in a wide variety of commercial products for components such as the PVR; and in references to the Model in a variety of system specifications and net attached storage studies.

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, added the OSSI Model name, and voted to approve Version 5 of the OSSI Model for public review. Of seven model components, two are being developed standards are being developed (PVR, PVL) for submission to the IEEE balloting process this year. 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. Versions 6 discussions have begun.

Overview

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.

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.

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.
- 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 OSSI 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

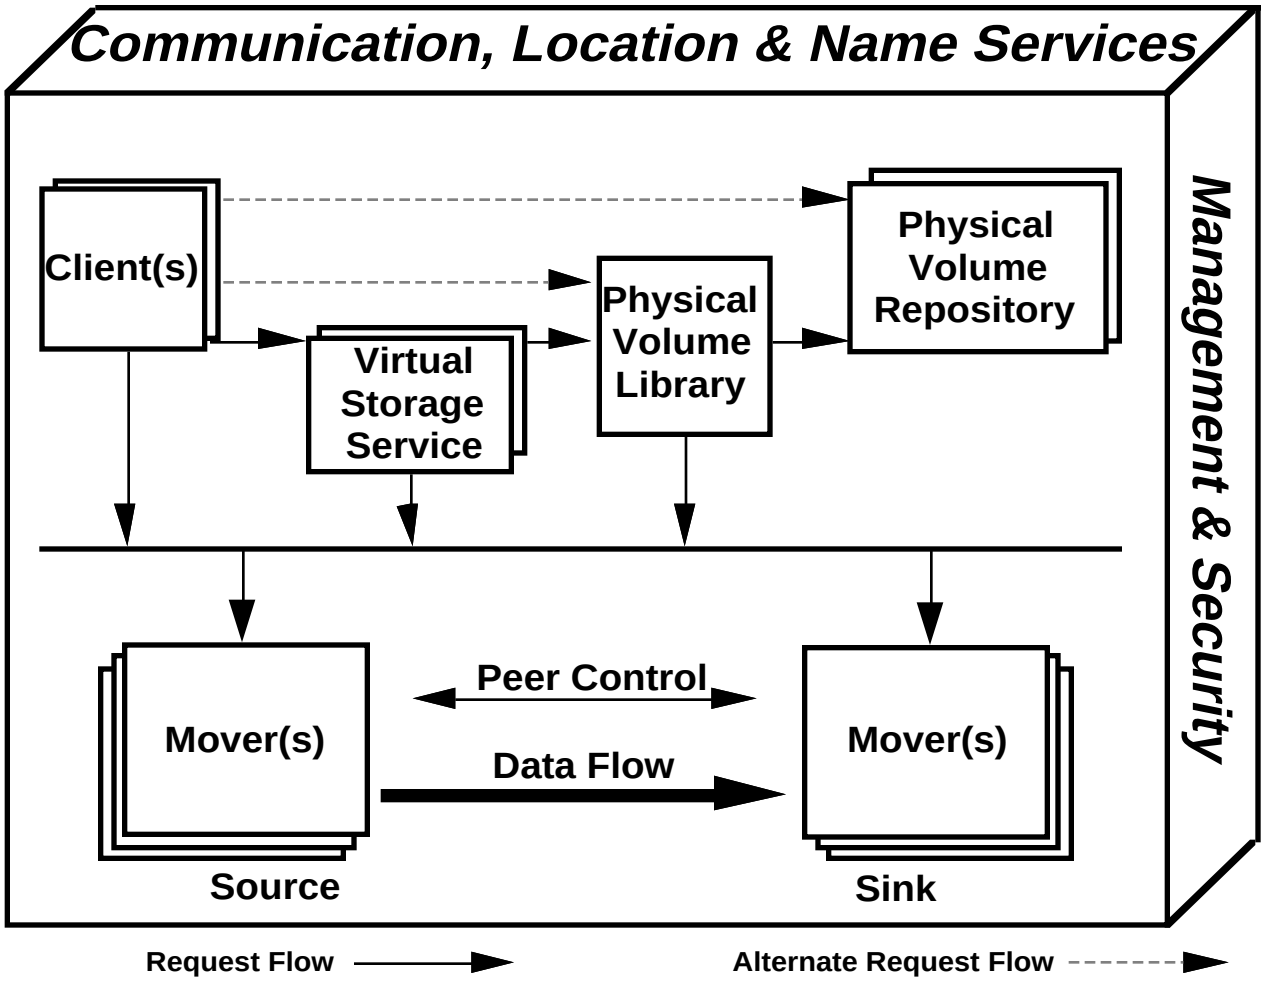

Figure 1: Reference Model for Open Storage Systems Interconnection

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.

The SSSWG and Development of Standards

Impact of Distributed Computing

The current interest and need to accommodate computing in an environment where users and resources are distributed has lead to a an even greater netcentric perspective for storage. The emphasis of the SSSWG Charter and the Model on distributed computing presaged this trend, and points up the value of a Model that, to a degree, leads the “real world”. References to various studies and projects relating to net attached storage may be seen at the SSSWG URL.

Metadata

While the IEEE effort in the Metadata arena has just begun, it has been suggested to the SSSWG that a recommend practices guide for a common subset of metadata may reduce the burden and difficulty of transitioning a pool of archived or migrated data from one application to another. Since this problem is so prevalent, this suggestion will be strongly considered.

Information Access

The focus of storage systems must not be "mass storage", but transparent, easy information access. The forces which drive developments in storage have only to do with the need to access information in a distributed, heterogeneous environment which continually makes such accses challenging.

Human Dynamics of "Standards" Development

The word "standards" was placed in quotes to point out the fact that the SSSWG does more than develop standards. In fact, the SSSWG has not yet produced a single standard. As well or instead, the effort provides guidance for implementors of storage systems and system components, and a common perspective for other groups developing related standards.

The human dynamics of "standards" development for storage systems devolves about several points. The first is that "storage" remains an invisible requirement behind computing systems, application software, networks, etc. In this regard, more "storage protagonists" are needed to raise the general awareness of the need for coordinated, systematic approaches to information access (was known as "mass storage"). Greater end user involvement is needed to help define the needs of storage systems, especially in regard to transparency, management, and minimized labor requirements. Finally, better coordination among the wide variety of groups developing standards for storage is needed. For example, at this writing several groups are discussing world wide identifiers for removable storage media!

SSSWG Pedigree and Charter

The SSSWG "pedigree" is that its efforts are part of the IEEE Computer Society, which traces its requirements and methods back through IEEE and ANSI to ISO. When the SSSWG began in 1990, its Charter was written so that the group would model generic systems, of any scale (not just HPC), and emphasize distributed systems created with an object oriented approach. This Charter is still effective.

Revision of the Model

The Reference Model was approved for public review by internal SSSWG ballot on September 8, 1994. It is expected that the sixth version will be produced before the end of 1997. This will be done in order to modify the Model's components to remain in agreement with related standards, as well as to correct inherent errors and inconsistencies in the version 5 draft Reference Model. It is unlikely that the sixth version will represent a major departure from the current fifth version.

Status of Model Components and Standards

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 relating to the AEP.

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 SSSWG will participate in discussions of other groups relating to object identifiers.

Physical Volume Library (PVL)

Working group participants have devoted a great deal of effort to the PVL and PVR standards, and drafts of these are viewable at the SSSWG URL. The PVR standard will move to IEEE ballot in a few months, and PVL will follow shortly thereafter.

Data Mover (MVR)

No work has been accomplished on the MVR component over the last year, but it is anticipated that quick progress will be made after the PVR and PVL standards are sent to ballot.

Storage System Management (MGT)

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

At this time the MGT committee is not active. Current members of the SSSWG have determined that some management functions belong with each Model component. In the two years since the SSSWG MGT committee was active, many commercial storage management products and a few standards efforts have arisen. To the extent possible, SSSWG will try to embrace these efforts, and standardize them as existing practices or works in progress.

Virtual Storage Service (VSS)

Robert Baird, Hewlett-Packard Company, is the chair of VSS, and several SSSWG members are partial members of the VSS group. Mr. Baird has published a "Virtual Storage Architecture Guide (VSAG)" [3], which expands on the virtual storage concepts of the Model. No recent changes have been made to this guide, and the SSWG overall views VSS as too complex to implement and to standardize at this time.

Medium Changer Service Standard

The SSSWG has submitted a Project Authorization Request (PAR) to the IEEE to begin a standard by this title (known simply as MCS). This standard will enable some plug and play capabilities for end users of automated tape libraries, reduce the time to market for new storage products, and be based on the SCSI-II command set. It is expected that this standard will be sent to ballot in a very short time.

Summary

In its sixth year, SSSWG has embarked on a revision effort for the Model, begun a new standard, and is about to ballot two standards for components of the Model. In addition, the SSSWG is attempting to broaden its participation in the activities of other

groups, as well as broadening participation by individuals in SSSWG through redefinition of its attendance requirement (see the SSSWG URL).

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.

Literature Citations

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

References

See http://www.arl.mil/IEEE

IEEE
Storage Systems
Standards Report

Jack Cole

Army Research Laboratory
Aberdeen Proving Ground, Maryland USA

http://www.arl.mil/IEEE

What Follows?

- **General Focal Points**
- **Netcentric Flavor of Storage**
- **Storage Standards Efforts**
- **IEEE SSSWG**
- **Quick Model Tutorial**

General Focal Points

- **Distributed Computing, Users, Resources**
- **Metadata**
- **Information Access, Not Mass Storage**
- **Storage Protagonists Needed**
- **End User Involvement Needed to Define Requirements**
- **Coordination of Standards, API Efforts**

Netcentric Flavor of Storage

not an exclusive list

- **VAXcluster of the '80s**
- **SSSWG Model - Separation of Control & Data**
- **HPSS**
- **Netstation Project/USC**
- **NCR Net Computing Architecture**
- **Lab Comp Sci & Engrg (LCSE), AHPCRC, UMN (SC'95 SGI PCA, Seagate, ...**
- **ARPA/ CMU/ NASD/ NSIC, CMU PDL "SAD", Taxonomy of Net Attached Storage Arch.**
- **EMC**

Storage Standards Efforts

- **ANSI**
  SCSI (world wide removable media ids)
- **IEEE**
  SSSC, PASC, Metadata, Embedded Sys
- **IETF**
  WG for Mass Storage Mgt
- **Vendor groups like NSIC**

SSSWG

Storage Systems Standards Working Group
IEEE Computer Society Project 1244

- **Pedigree**
- **Charter**
- **Activities**

SSSWG Pedigree

Storage Systems Standards Working Group
IEEE Computer Society Project 1244

- **ISO -> ANSI -> IEEE**
- **IEEE COMPUTER SOCIETY (CS)**
- **IEEE CS Stds Activity Board (SAB)**
- **SAB Storage Systems Standards Cmte (SSSC)**
- **SSSC Working Group (SSSWG)**

SSSWG CHARTER

- **To Model Generic Systems**
- **Create "Guide and Reference Model"**
- **Develop Recommended Practices Both for Standalone and Distributed Systems**
- **Emphasize Distributed Systems**
- **Without Favor Include Storage Systems of Every Scale in its Studies**
- **Use an Object Oriented Approach**

SSSWG Activities

- **Model Revision**
- **New PAR for MCS**
- **PVR Move to Ballot**
- **Non Attendee Participation (cf IETF)**
- **Coordination with THIC, NASD, PASC, and others**
- **Need Greater End User Participation**

Model Revision

- **IEEE Imposed Lifetimes**
- **Model Must Reflect "Real World" Which is Always Changing**
- **Model Should Lead Real World To A Degree**
- **Provide Common Perspective, Framework for Standards, Practices, Interoperability**

MCS API

- **Medium Changer Service Standard**
- **Changers: Automated Tape Libraries, Optical Jukeboxes, etc**
- **Will remove OS and Software App Dependence**
- **Embrace Existing Practices, Works in Progress**
- **SCSI-II Based, But is OS Independent (SCSI Medium Changer Interface only describes communication)**
- **Plug & Play, New Products Available More Quickly**

PVR

- **Physical Volume Repository (PVR)**
- **"Lower Level" Component of the Model**
- **Human And Software Interfaces That Stow And Selectively Mount Removable Media**
- **Protein PVR ;-)**
- **"Sees" Cartridges and Drive Mount Points**
- **Major Operation is to Mount Cartridges.**

Quick Tutorial of the Model

IEEE Mass Storage Systems Reference Model
MSSRM

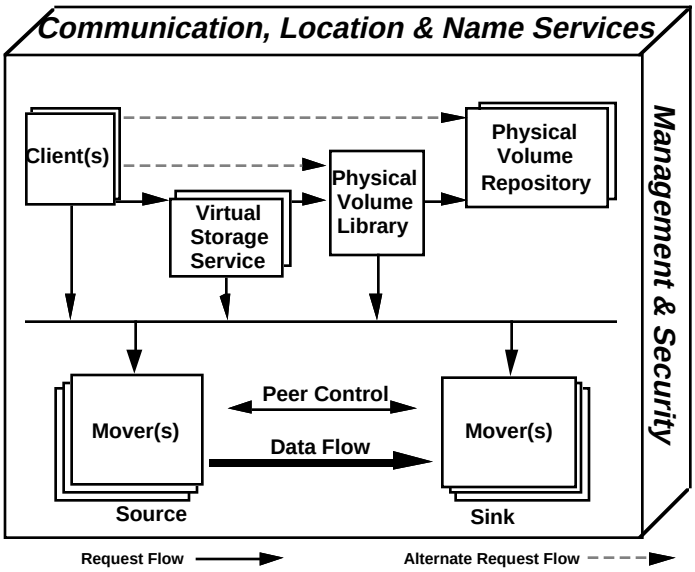

Components - Parfait Cake

AEP	**Application Environment Profile**
MGT	**Storage System Management**
VSS	**Virtual Storage Service**
PVL	**Physical Volume Library**
PVR	**Physical Volume Repository**
MVR	**Data Mover**

Some Key Concepts of Model

- **Abstraction**
- **Transparency (Device, Location, Replication)**
- **Separations of**
  - **Policy, Mechanisms to Effect Policies**
  - **Control and Data Flows**
    (implied net attached storage)
- **Unrestricted Scalability and Distributed Storage**

Some Characteristics of The Model

- **Vital, Influential At Numerous Sites**
- **Provides Common Perspective**
- **Defines Domain of Standards for Open Storage Systems**
- **Two Or More Systems Using OSSI Standards Will Be Able To Inter-Operate Openly**

Model Components

"Higher Level"
Programmatic Interfaces & Environment

Application Environment Profile (AEP)
Includes Implementor Declaration of Transfer Agent (e.g., DCE), Policy Modules, ...

Storage Object Identifiers & Generators (SOID)
Namespace Issues
World Wide Removable Media Identifier Debate

Model Components

"Higher Level"
Programmatic Interfaces & Environment

Application Environment Profile (AEP)
Includes Implementor Declaration of Transfer Agent (e.g., DCE), Policy Modules, ...

Storage Object Identifiers & Generators (SOID)
Namespace Issues
World Wide Removable Media Identifier Debate

Schedule of Upcoming Meetings

- **OSV Meeting on MCS**
  November 4th, Cupertino, CA Hotel
- **SSSWG Meeting November 5-6-7**
  HP Facilities, Cupertino, CA
- **January Meeting Seattle**
  Joint Meeting with THIC

Model Components

"Higher Level"

Storage System Management (MGT)
Services To Monitor And Control Resources As Dictated By Site-Specific Policies

Virtual Storage Service (VSS)
Access And Organization Of Persistent Storage Presented As Single Virtual Storage Image
Create and Manage Virtual Stores
Store to Volume Mapping

Comments About Standards

- **Openly Developed**
- **No Such Thing as "Vendor Lock"**
  Except that Others Do Not Participate
- **Revision of Model, For Example,**
  Is An Open Book Ready to be Written In

Model Components

"Lower Level"

Mover (MVR)

Transfer Of Data Between Two Endpoints

Manages Data Transfer
Designed for High Speed Data Transfer
Loads Media to Media Access Points,Transfer Data

http://www.arl.mil/IEEE