December 3-6, 2007, Santa Clara Marriott, Santa Clara, CA

Open Source Storage Management Aperi and SMI-S for Linux

Robert Wipfel
rawipfel@novell.com

Todd Singleton
toddsing@us.ibm.com

MANAGEMENT DEVELOPERS CONFERENCE

Agenda

- Open Source & Standards
- Open Storage Management
  - Introduction to Aperi
    - Storage Resource Management
  - SUSE Linux Enterprise 10
    - The “Well Managed” System
      - SMI-S Providers for Linux Servers
- Demo
- Conclusions

Open Source

o m c

open management with CIM

Open Standards

MANAGEMENT DEVELOPERS CONFERENCE

Collaboration Meets Innovation

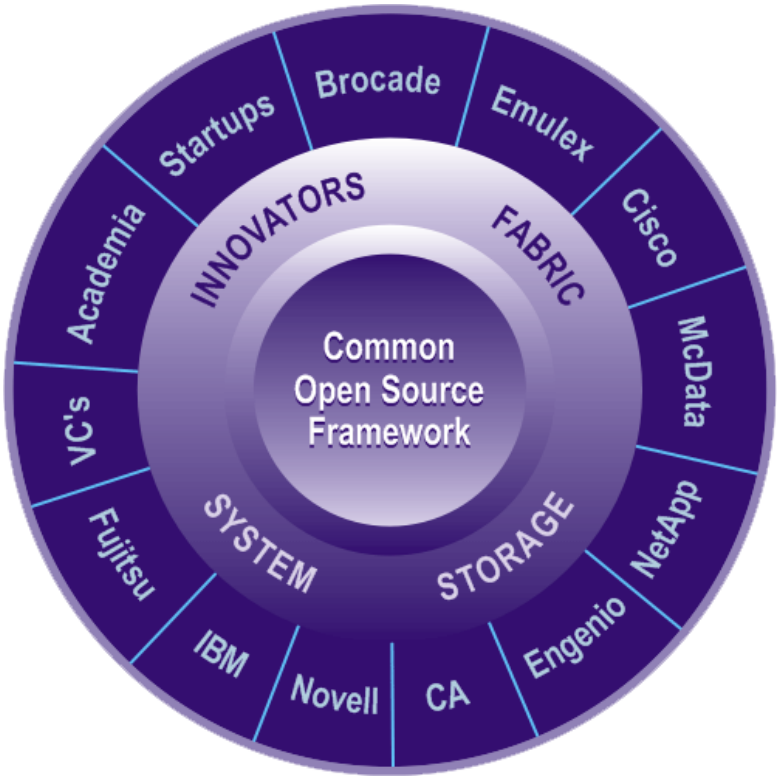

- 9 leading vendors founded an Eclipse project to accelerate storage standards adoption and spur innovation
- Vendor-neutral framework that includes an implementation of SNIA's SMI-S standard
- Initial IBM contribution of 1 million lines of code
- Follow-on contributions from community
- Aperi framework passed SMI-S interoperability conformance testing program (CTP)
- Latest contribution includes SAN simulator

""SNIA's planned relationship with Aperi will include interoperability programs for SMI-S, the use of SNIA facilities for Aperi interoperability programs, and advancing current and new storage standards. The IT industry will benefit from Aperi helping to drive SMI-S implementations, storage technologies and open standards." -- *Wayne Adams*, SNIA Chairman

Aperi's Mission

- Provide an open, extensible, standards-based storage management framework
- Give customers more flexibility and choice on how to manage their storage
- Simplify the infrastructure customers need to manage storage
- Drive adoption of industry standards

Currently in incubation phase of Eclipse project development

Framework Consists of Two Products: SRM

- Storage Resource Manager
  - Discover, monitor, control storage resources
  - Reporting, event management
  - Graphical CIM client with topology viewer
  - Supports several device and vendor types
    - Storage subsystem configuration
      - LUN assignment, and zoning
    - SAN fabric manager
    - Tape library discovery and reporting
    - File system capacity reporting
      - size, %used, %free only

Framework Consists of Two Products: SAN Simulator

- Storage Area Network Simulator
  - Simulates network of storage devices
    - Switches, tape, subsystems, etc
  - Reduces the need for expensive hardware resources
  - Extensible to support new vendor devices
  - Device setup done in two ways:
    - Manual configuration
    - Snapshot of real device
  - “Productizing” for v0.4 release

Leverages Eclipse Development Environment

- The Eclipse platform consists of runtimes, tooling, and components
- Eclipse provides a community oriented development environment
  - CVS, mailing lists, wikis, newsgroups
- Get up and running quickly within the Eclipse IDE
  - Download Aperi source code from Eclipse CVS
  - Download additional 3rd party libraries
  - Compile Aperi code
  - Run Aperi's components from within the IDE
    - With integrated source level debug environment

MANAGEMENT DEVELOPERS CONFERENCE

Aperi Architecture

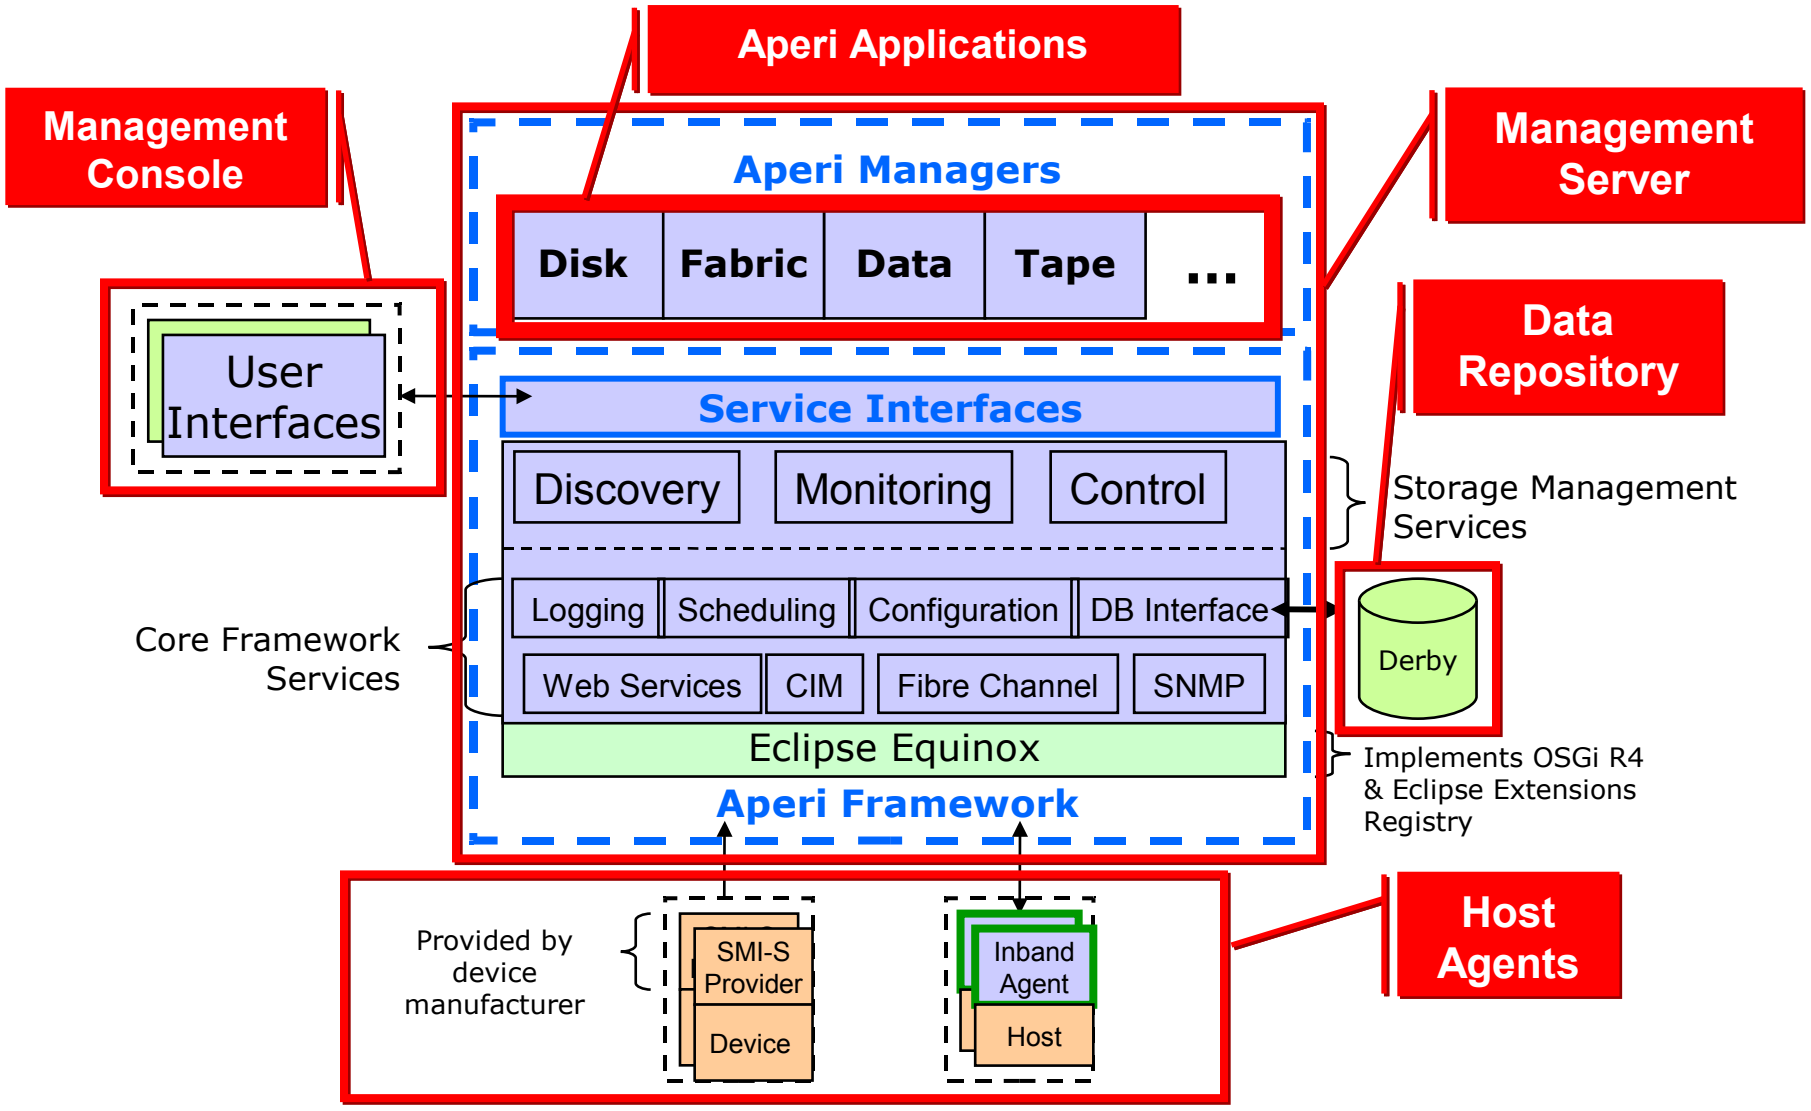

MANAGEMENT DEVELOPERS CONFERENCE

Recent Aperi Activities

- Novell donates support for XEN virtualization
- YottaYotta joins community
- Aperi Webcast (available online)
- Linux Technical Review article published
- Milestone 0.4 completed
- Brocade, LSI and NetApp join Eclipse Foundation
- Novell Brainshare futures demo includes Aperi and N_Port_ID virtualization for Xen VMs
- Eclipse Foundation Aperi booth at SNW
- SNIA and Eclipse-Aperi Alliance (2006)

MANAGEMENT DEVELOPERS CONFERENCE

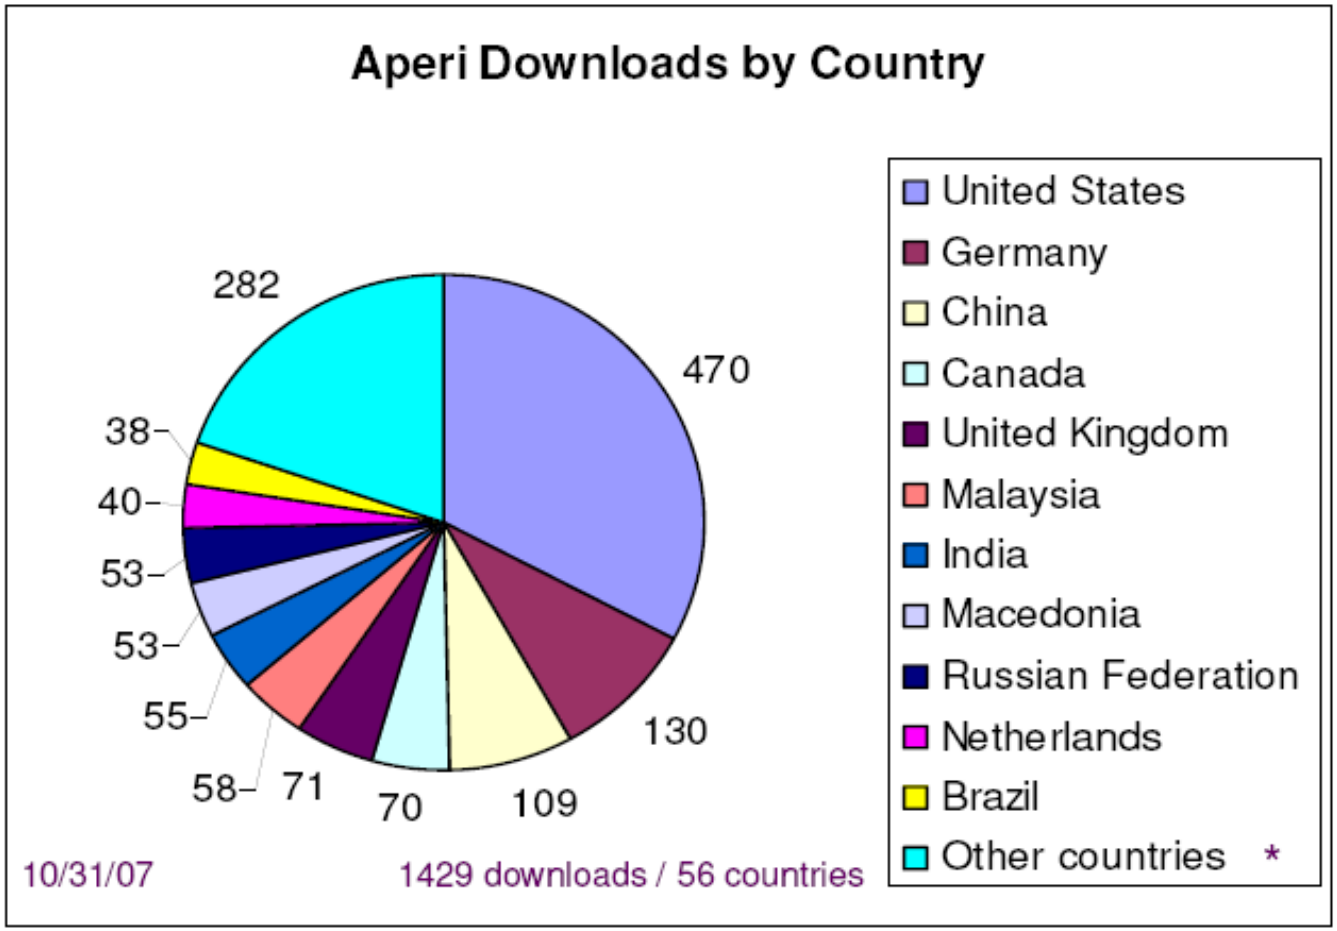

•45 Countries with 30 or less downloads: Israel, Republic of Korea, France, Japan, Italy, Indonesia, Taiwan, Turkey, Austria, Switzerland, Ireland, Hungary, Australia, Belgium, Hong Kong, Mexico, Tunisia, Sweden, Thailand, South Africa, Norway, Spain, Bulgaria, Denmark, Greece, Lithuania, Poland, Finland, Ukraine, Romania, Ghana, Czech Republic, Venezuela, Philippines, Albania, Pakistan, Colombia, Portugal, Saudi Arabia, New Zealand, Sri Lanka, Singapore, Cyprus, Vietnam

MANAGEMENT DEVELOPERS CONFERENCE

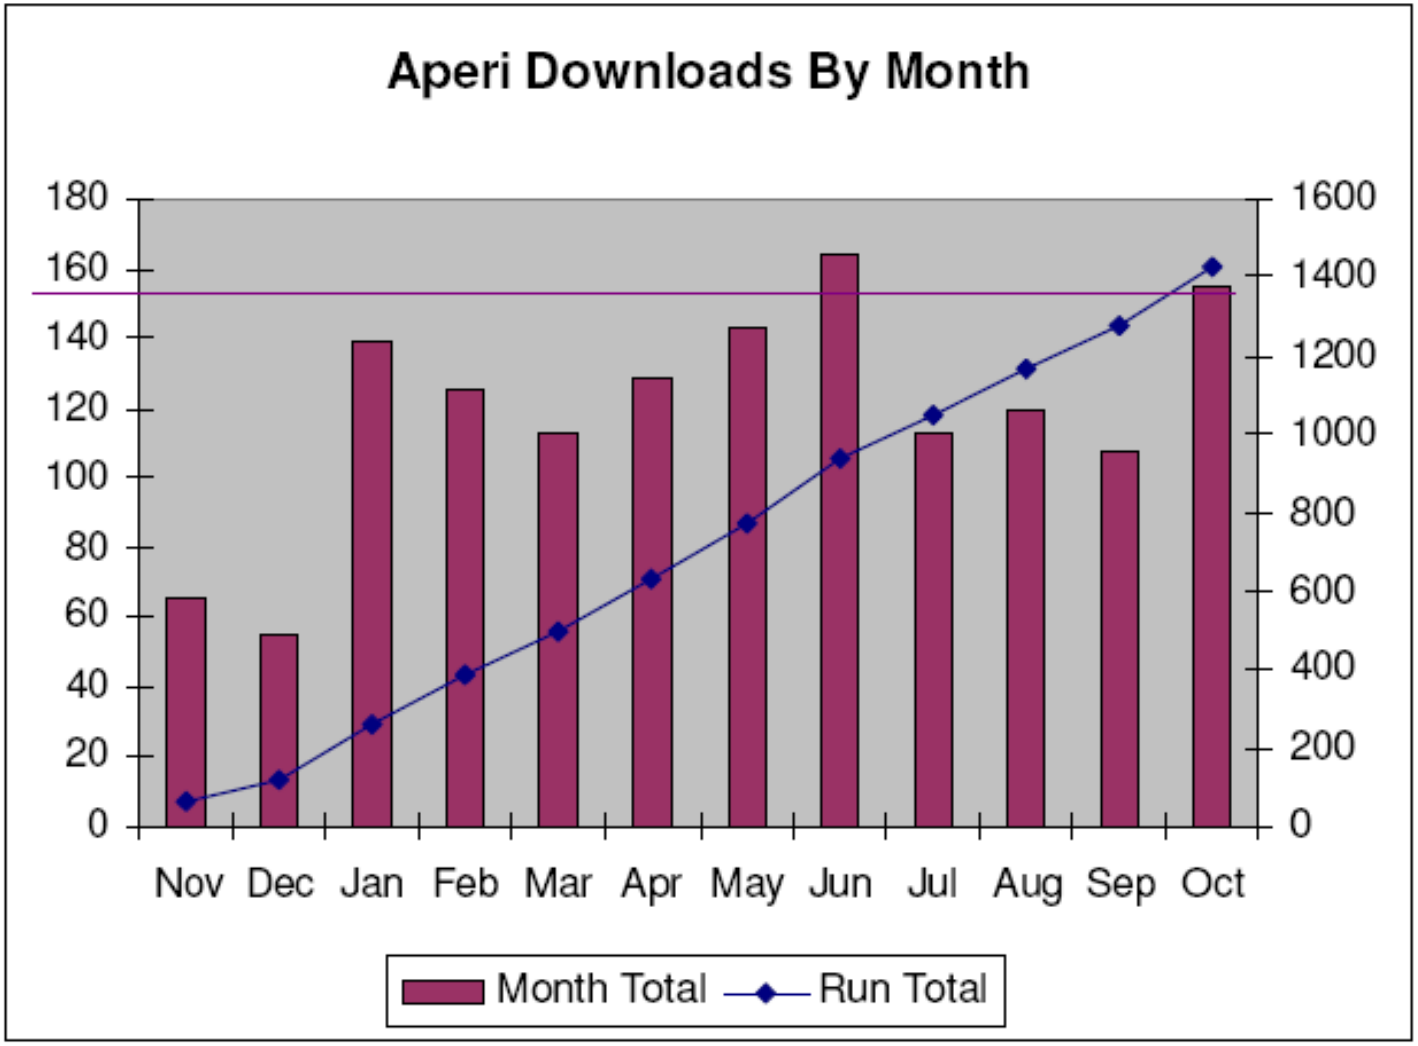

MANAGEMENT DEVELOPERS CONFERENCE

HOWTO - Get Involved

- Aperi Roadmap
  http://wiki.eclipse.org/images/f/f8/Aperi_Technical_Roadmap_v7.pdf
- Aperi Website & Demo
  http://www.eclipse.org/aperi/
- Aperi Project Collaboration
  http://dev.eclipse.org/mailmain/listinfo/aperi-dev
  http://wiki.eclipse.org/index.php/Aperi_Storage_Management_Project
  news://news.eclipse.org/eclipse.technology.aperi

MANAGEMENT DEVELOPERS CONFERENCE

SUSE® Linux Enterpise 10

- The platform for the open enterprise
  - Built-in certified application security
  - Integrated systems management
  - Virtualization and HA storage foundations
  - Supported by major IHV platforms

The Well-managed System

- The need for standardized management is driven by IT customers who want to manage all their systems – standalone, rack mount, blades and storage – using integrated tools. This requires a focus on the intersection of open management standards in the server, storage and virtualization areas.

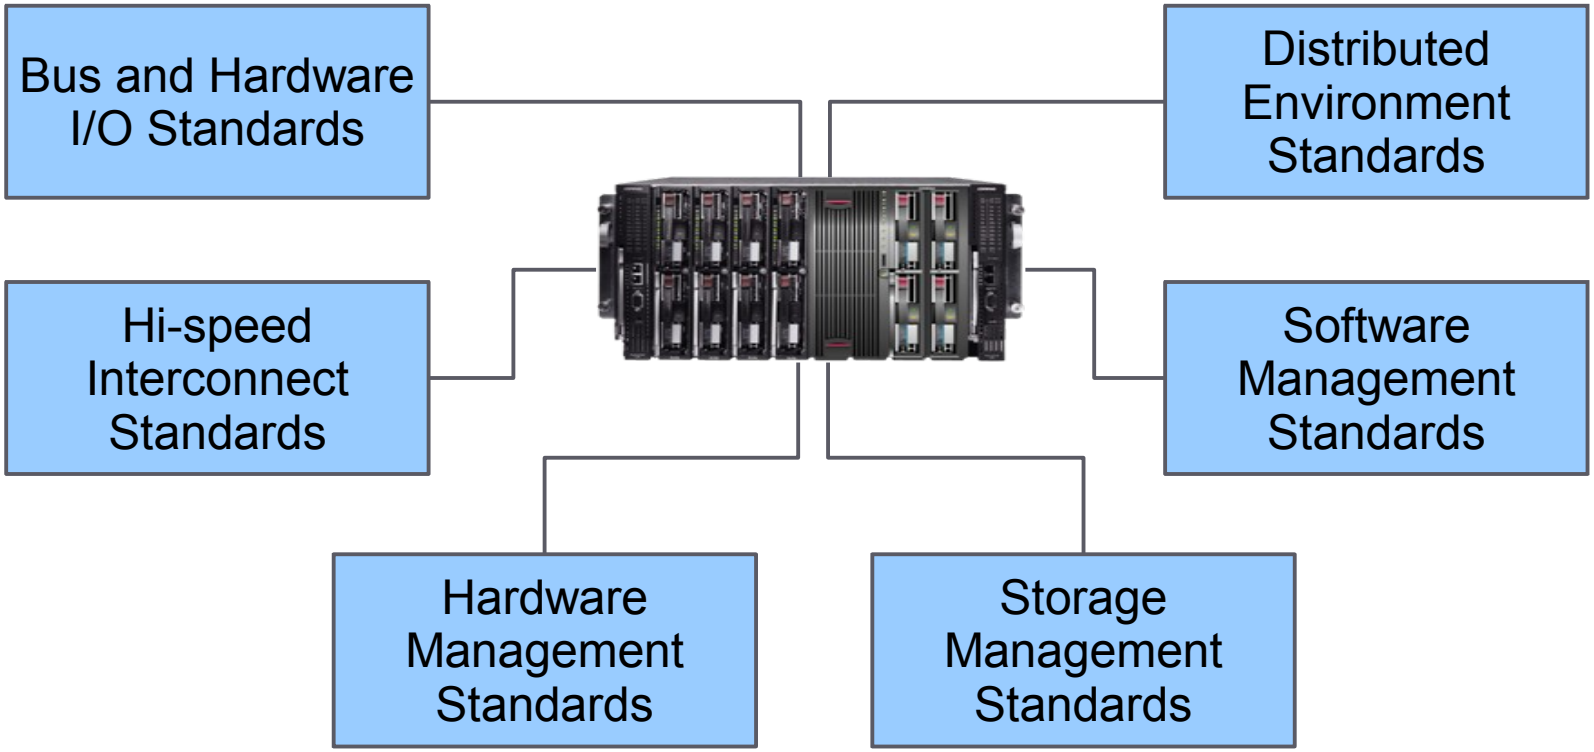

CIM based Standards

- Intelligent Platform Management Interface (IPMI)
- Systems Management Architecture for Server Hardware (DMTF SMASH)
- Storage Management Initiative Specification (SNIA SMI-S)
- System Virtualization, Partitioning and Clustering (DMTF SVPC-V & C)

Why Standards?

- Interop
  - Be a 1st class citizen in enterprise networks
- Model
    - Associates inter-dependent managed elements
      - Virtual machine -> cluster resource -> cluster node -> SAN path -> SAN LUN
    - Necessary for emerging enterprise-scale problems
      - Configuration management for virtualized data centers
- Policy
  - Monitoring plus Model = foundation for Policy
      - Auto-migrate virtual machine on multi-path failure indication

MANAGEMENT DEVELOPERS CONFERENCE

Reference Model

Representing Services as Graphs

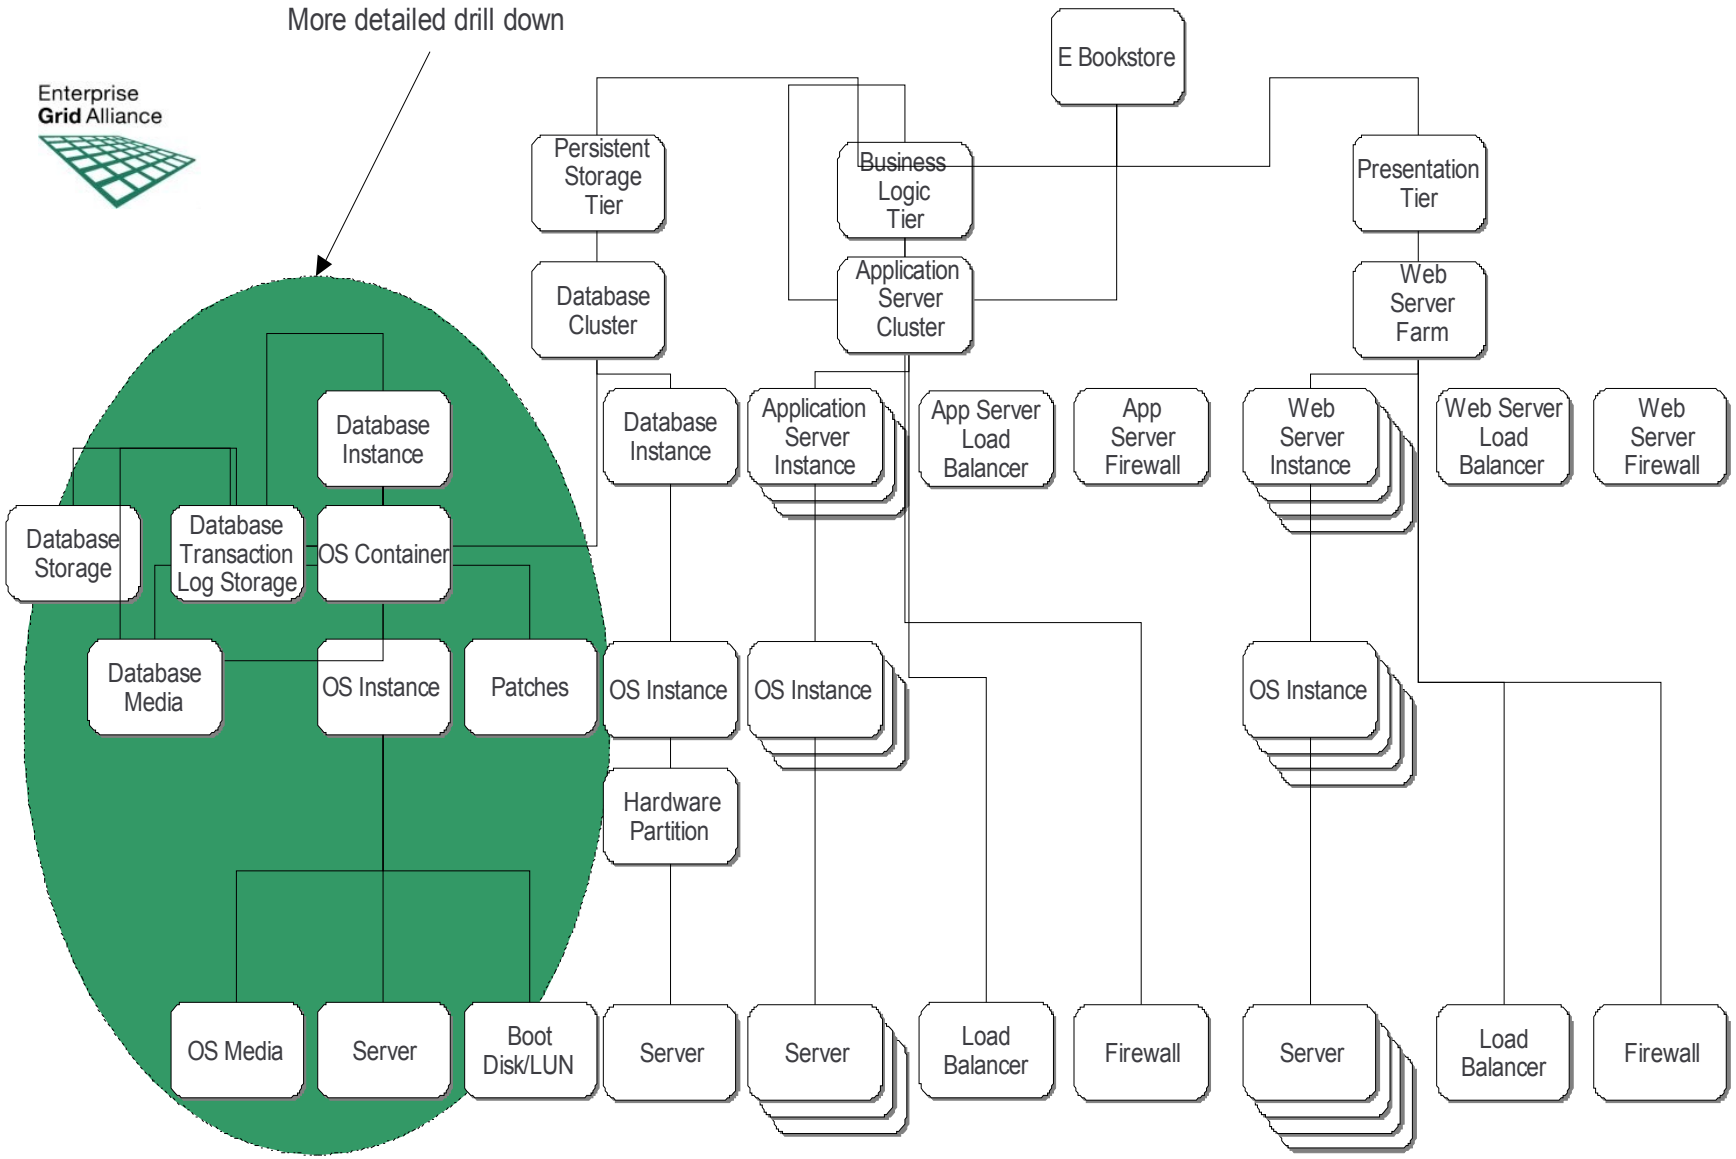

MANAGEMENT DEVELOPERS CONFERENCE

Virtual Machines

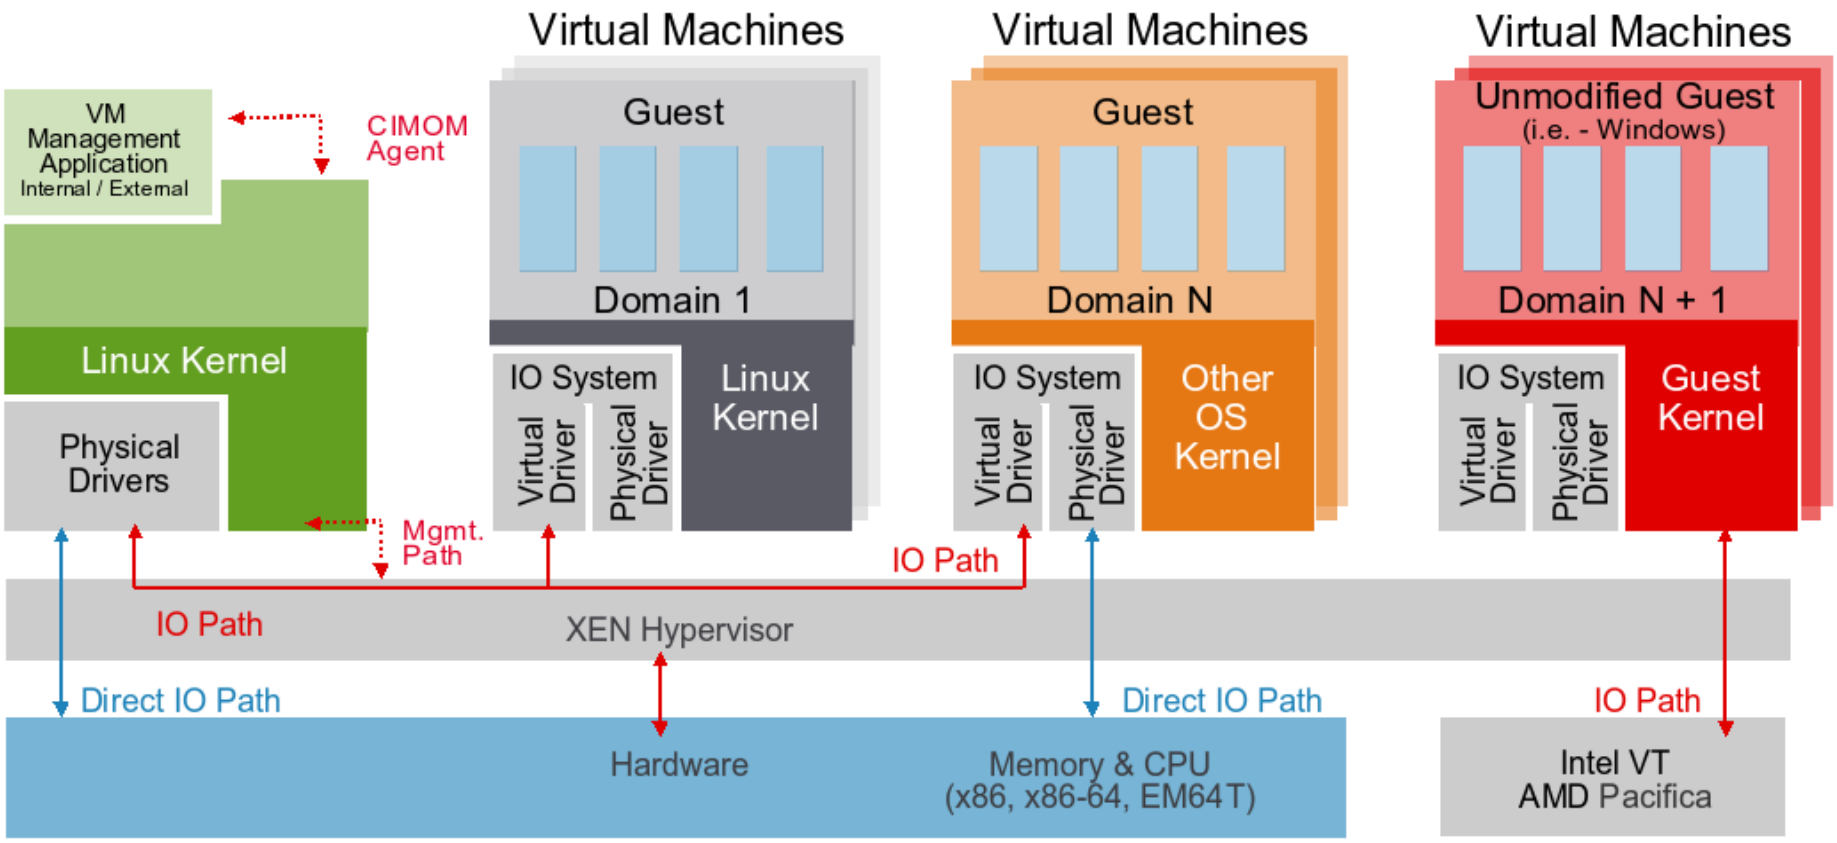

MANAGEMENT DEVELOPERS CONFERENCE

H/A Storage

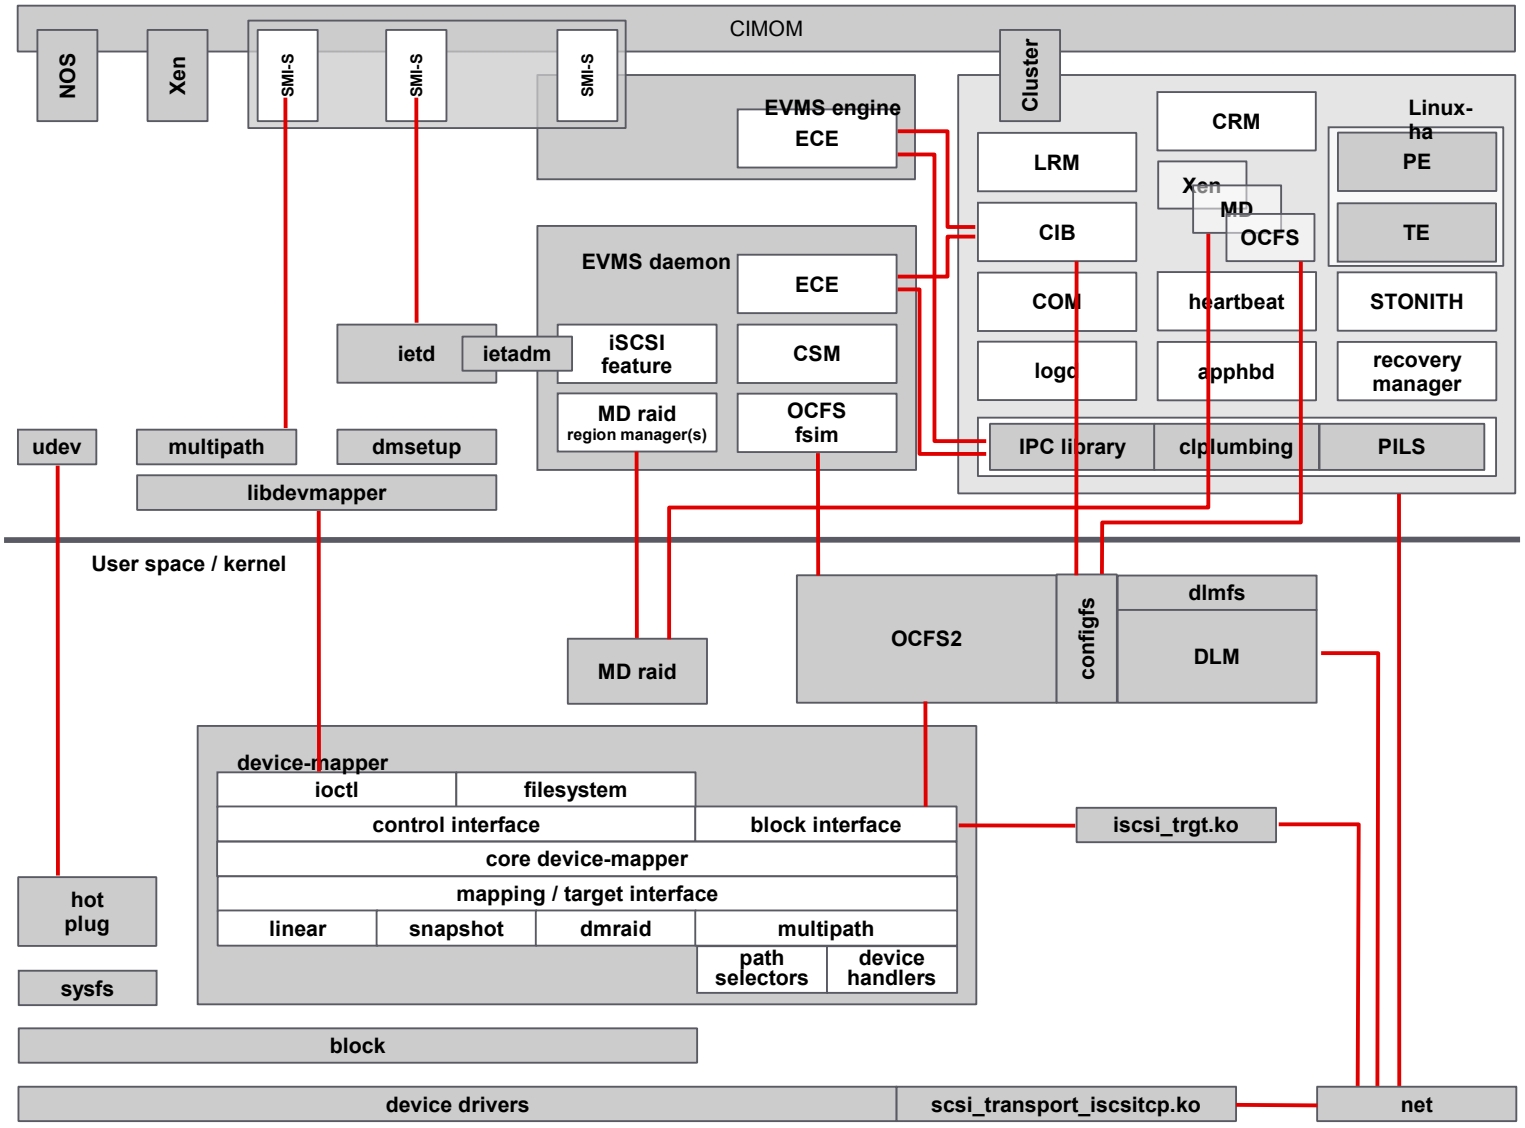

SNIA SMI-S Profiles

- Implemented (per CTP)
  - Server Profile
  - Volume Management Profile
    - Block services
    - Health packages
  - Extent Composition Subprofile
  - Indication Subprofile
- Work in Progress
  - Copy Services Subprofile
  - iSCSI and Multipath SCSI I/O

Storage Pool Manipulation

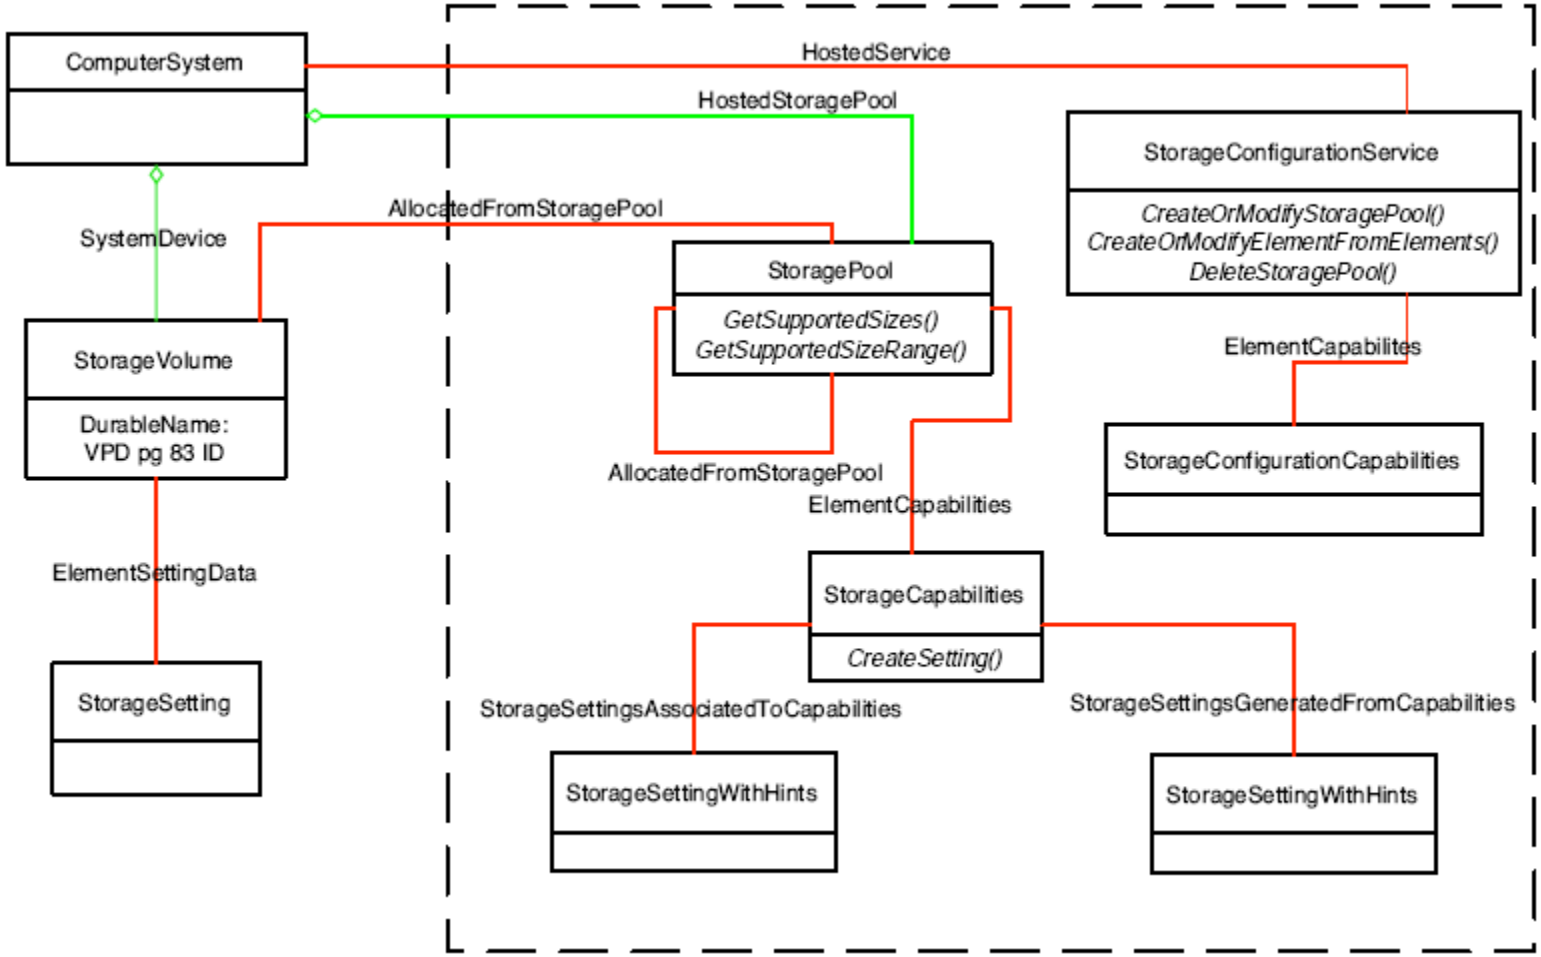

Provider Implementation

- CMPI
  - OpenPegasus or OpenWbem
- Developed using ECUTE Analyzer
- Dependencies
  - Open management with CIM
  - Enterprise Volume Management System
- EVMS abstracts MD and DM
  - Modeled via SMI-S StorageCapabilities
    - Multidisk (MD) for software RAID
    - DeviceMapper (DM) for LVM2 regions

MANAGEMENT DEVELOPERS CONFERENCE

Demo

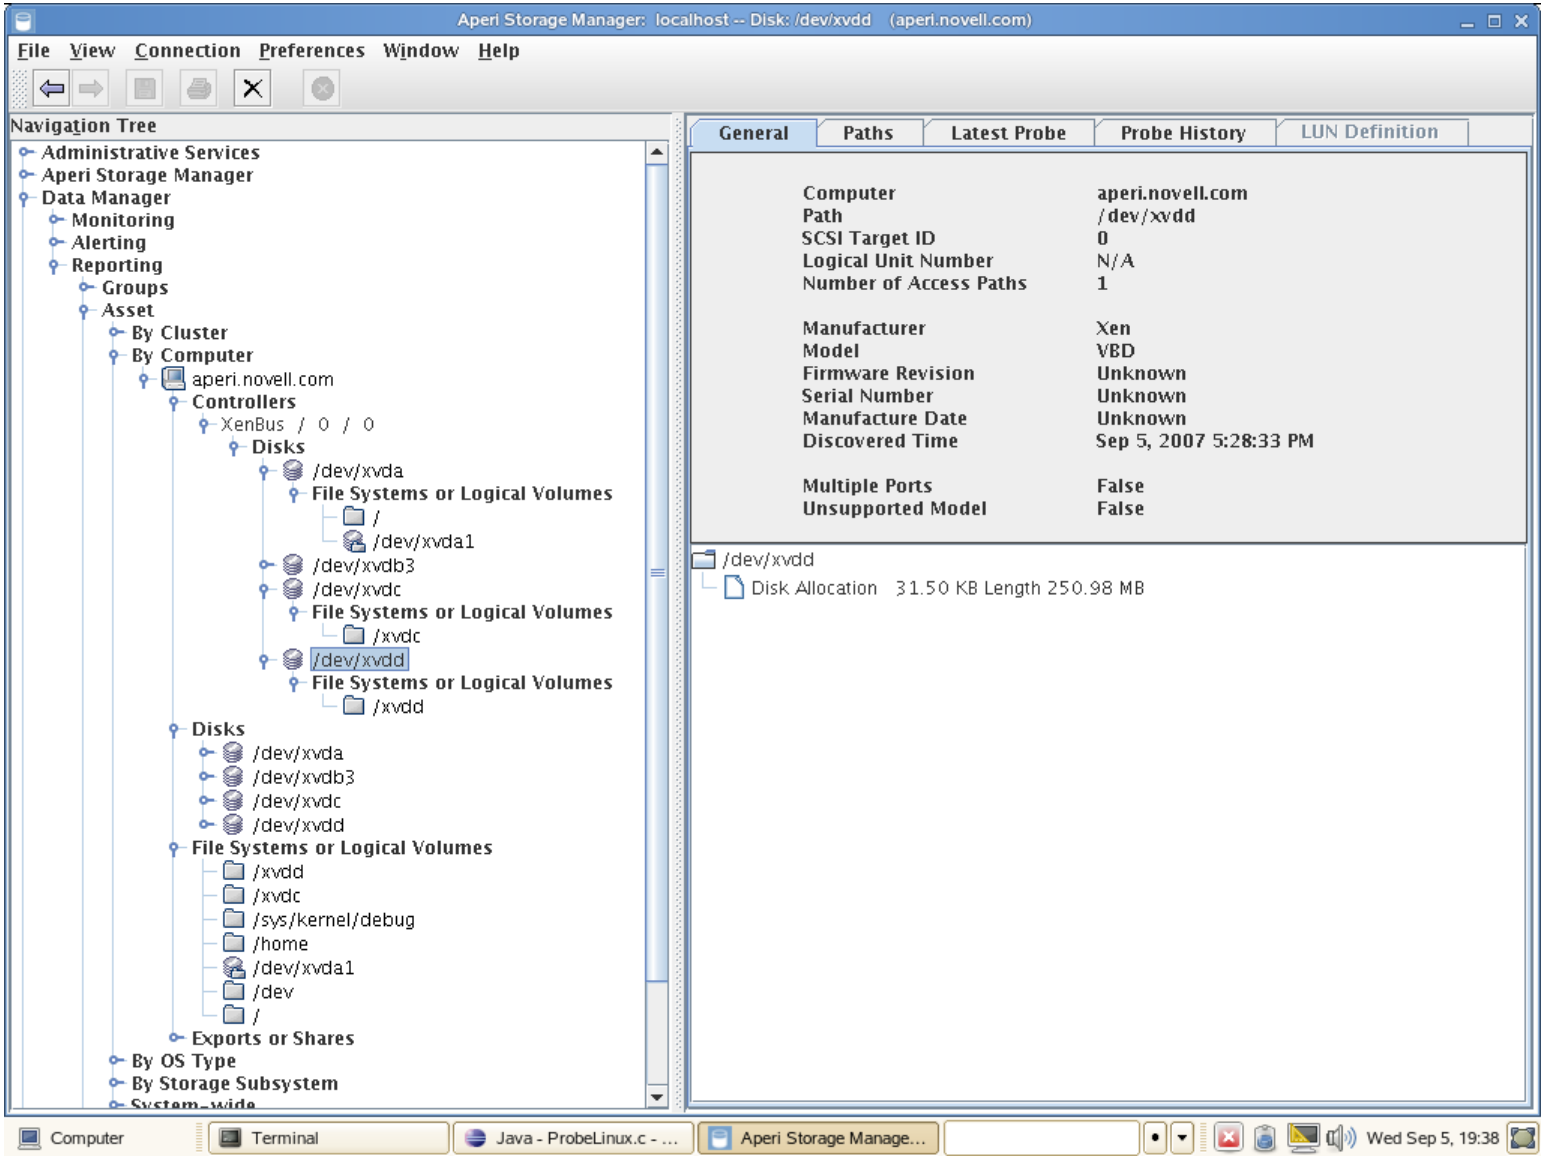

MANAGEMENT DEVELOPERS CONFERENCE

Example: Xen Cluster

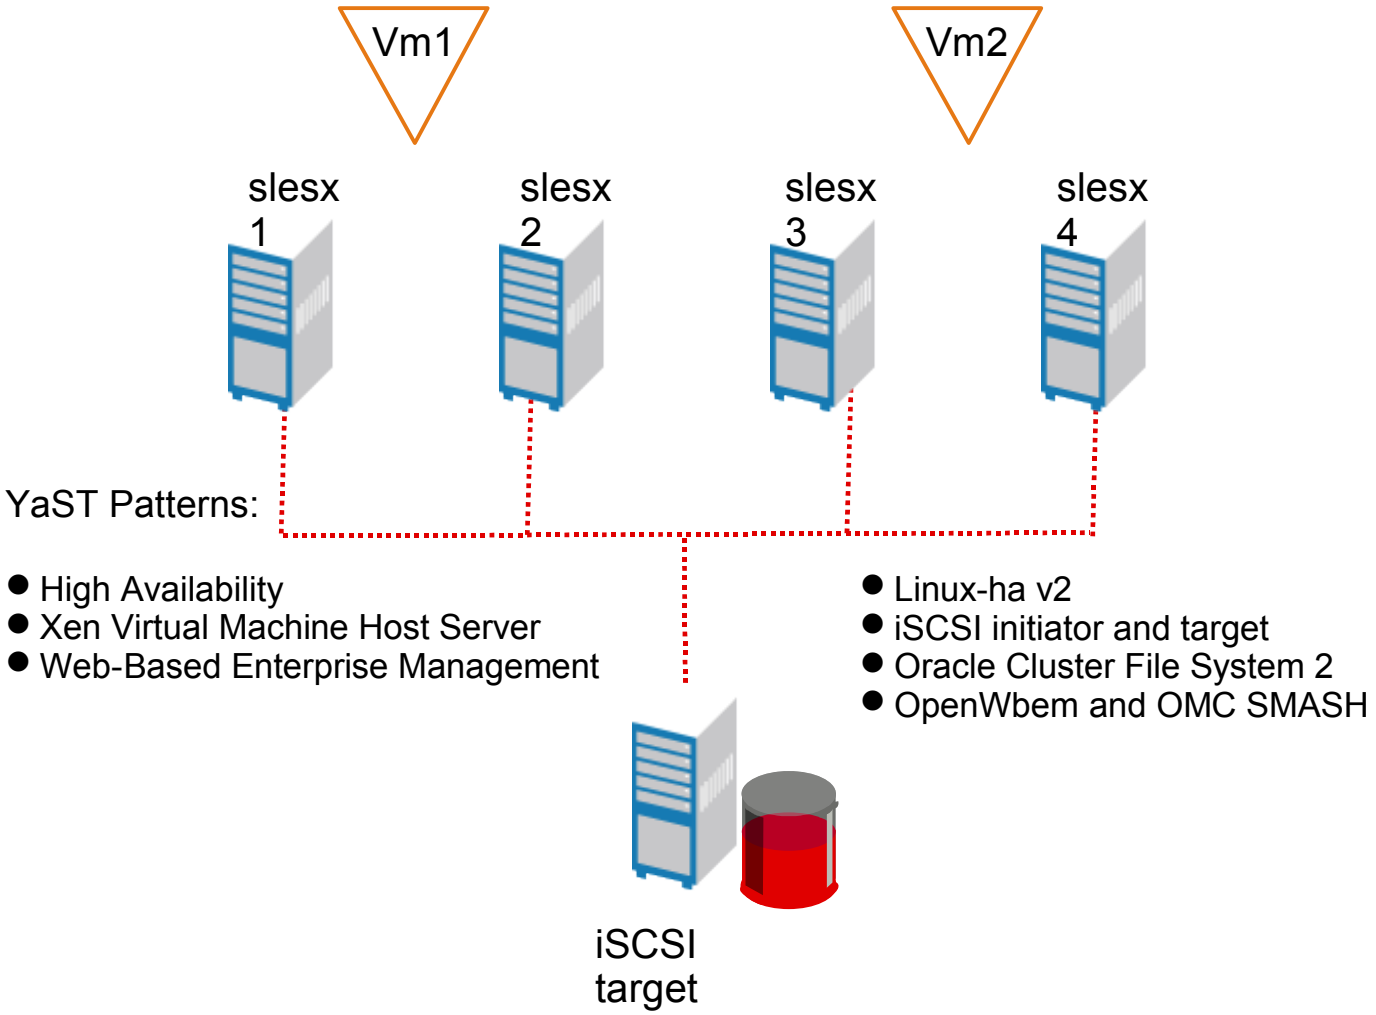

Conclusions

- Open Source
- Open Standards
- Model is Required
  - For Service Oriented Infrastructure
  - Relate Storage to other Resource Types
- Mixed Source Software
  - For the Next Generation Data Center

Legal Notices

- Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both
- Other company, product, or service names may be trademarks or service marks of others