



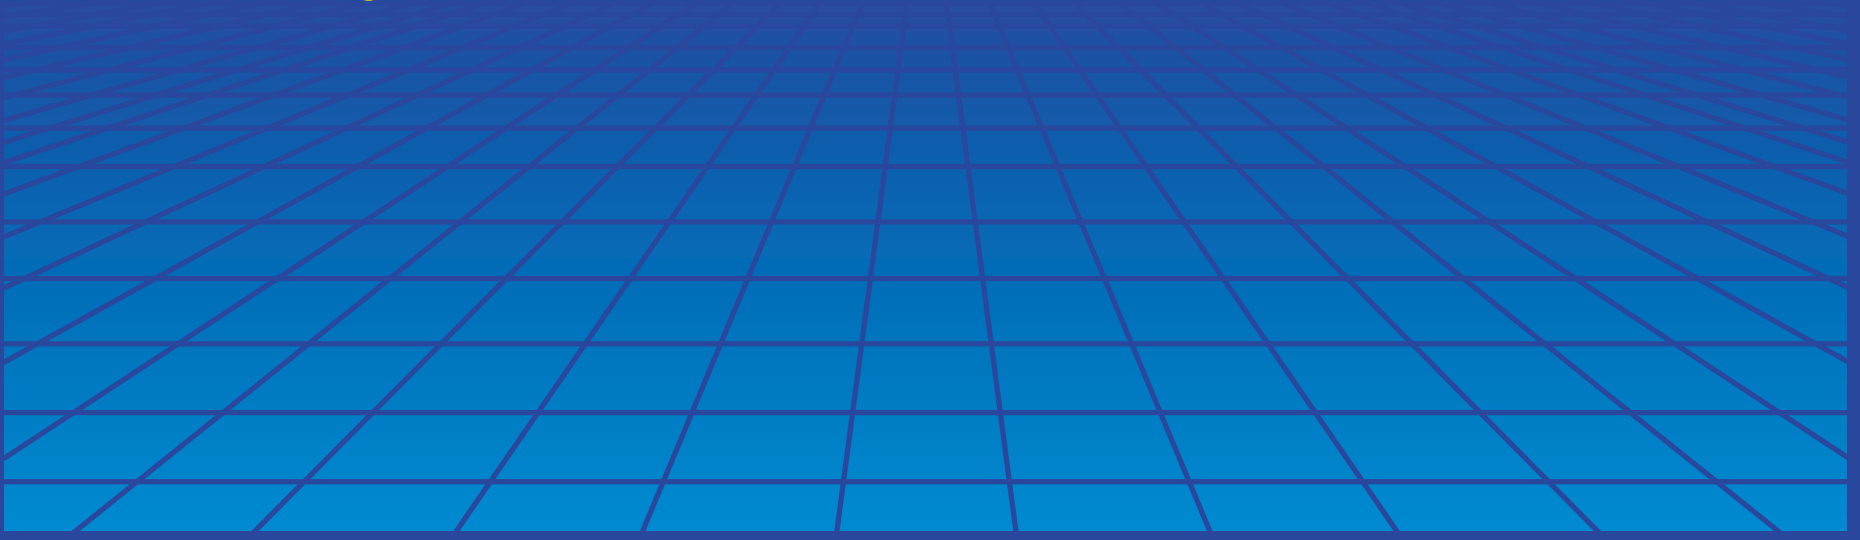
Enterprise Volume Management System Project

April 2002



Mission Statement

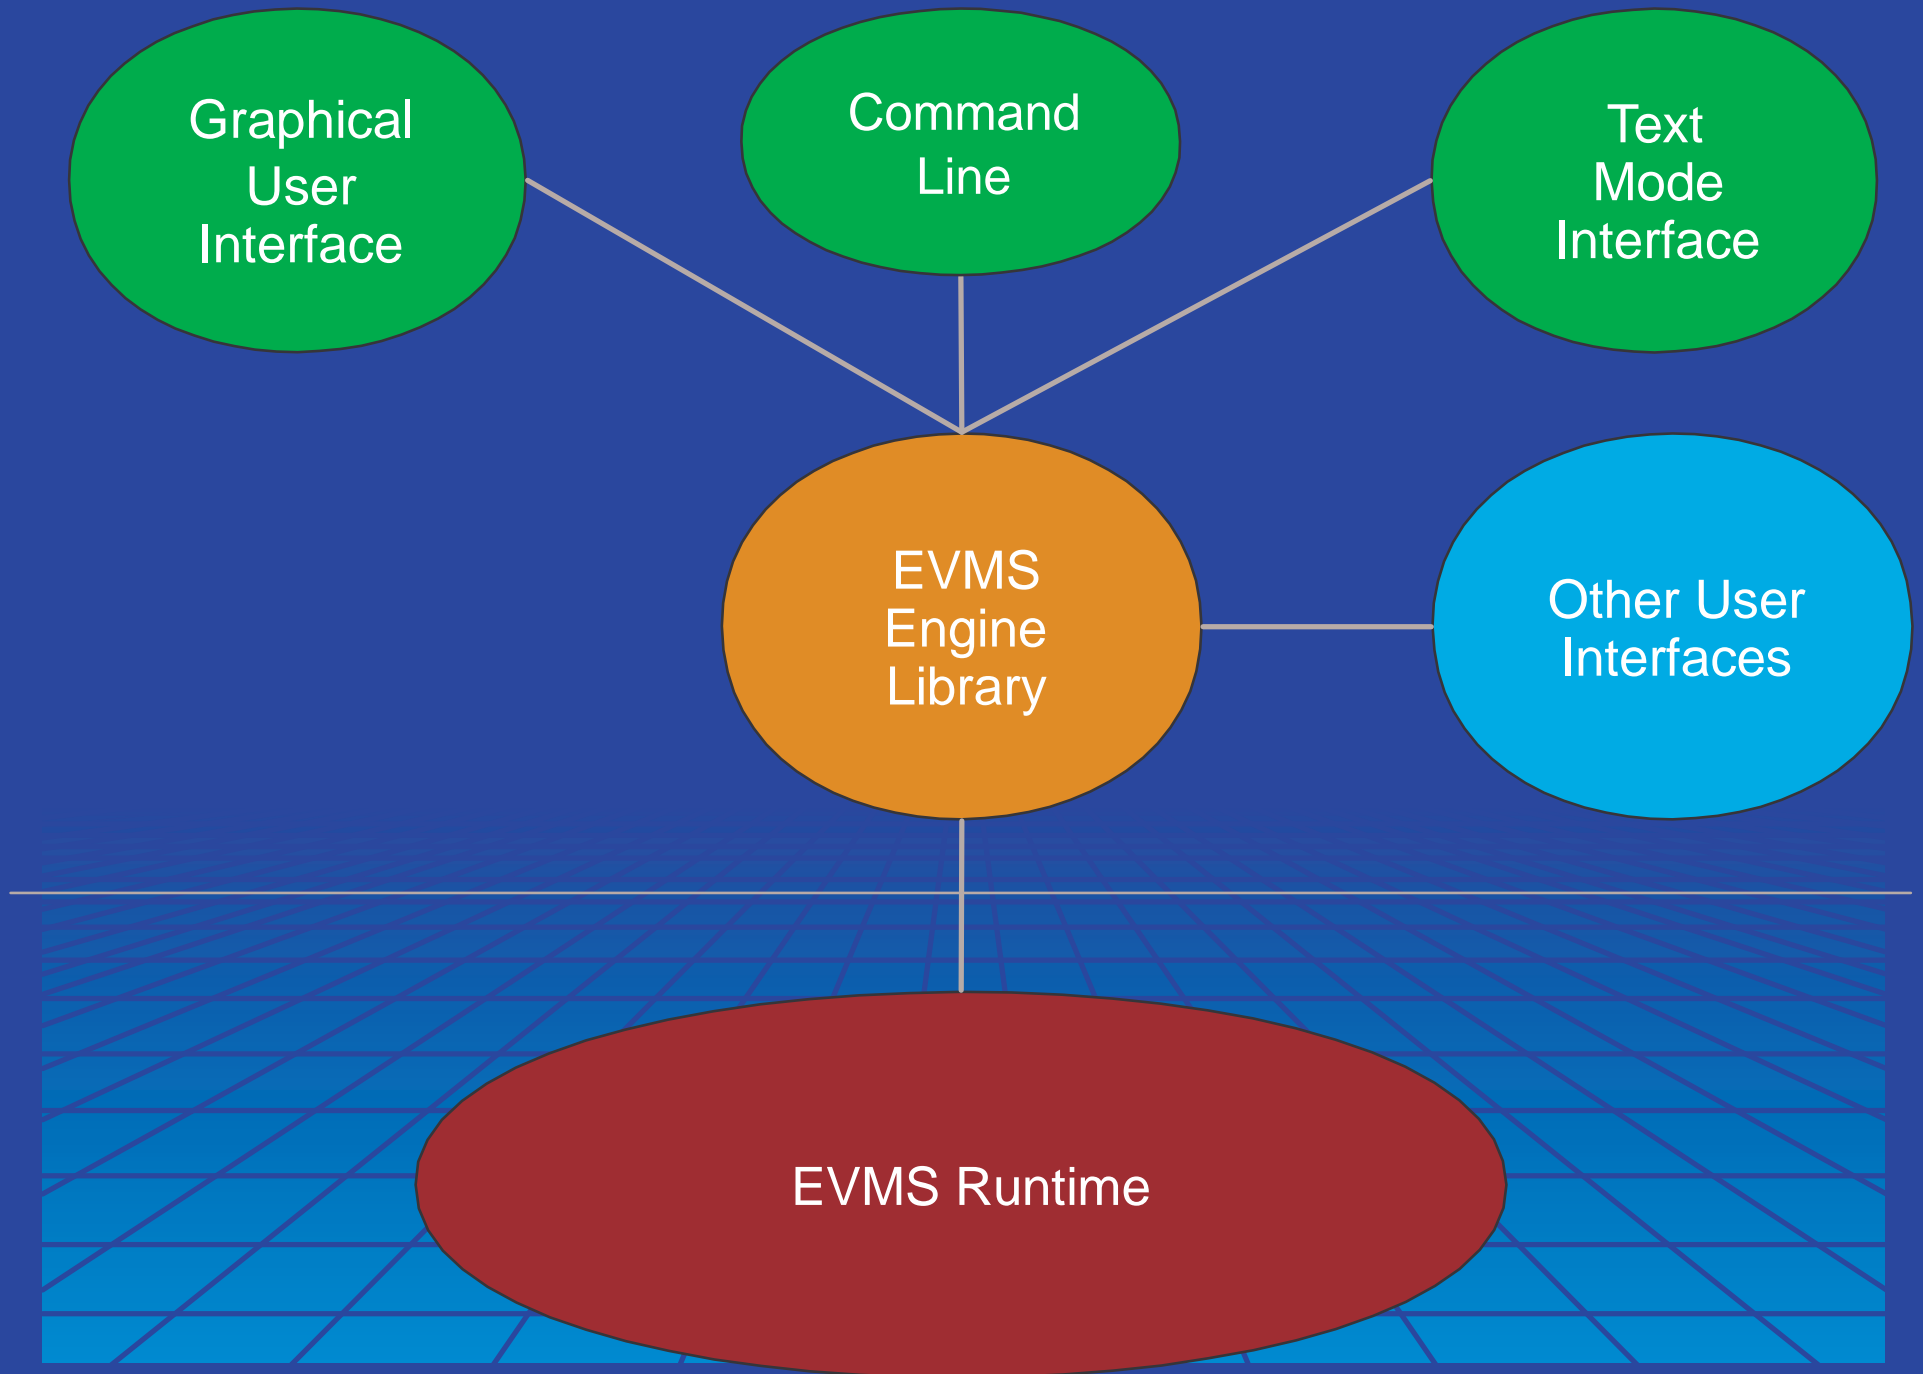
To create a state-of-the-art, enterprise level volume management system for Linux® which will also reduce the costs associated with migrating to Linux.



EVMS Architectural Overview

- ▶ Two components: EVMS Engine and EVMS Runtime
 - The EVMS Runtime resides in kernel space and provides access to volumes
 - The EVMS Engine provides API support for the creation, configuration, management, and deletion of Volumes, Volume Groups, Partitions, and Disks.
- ▶ The EVMS Engine and EVMS Runtime are highly modular, and both accept plug-in modules to expand their capabilities
- ▶ 7 classes of plug-ins are used:
 - Device Management Plug-ins
 - Partition Management Plug-ins
 - Region Management Plug-ins
 - Feature Plug-ins
 - Filesystem Interface Plug-ins
 - Cluster Management Plug-ins
 - Distributed Lock Management Plug-ins

User & Kernel space Diagram



Graphical
User
Interface

Command
Line

Text
Mode
Interface

EVMS
Engine
Library

Other User
Interfaces

EVMS Runtime

EVMS Plug-in Modules

- ▶ Device Management Plug-ins allow the EVMS to interface to local and remote devices. An example of a Device Management Plug-in would be a SAN plug-in.
- ▶ Partition Management Plug-ins allow the EVMS to recognize, use, and manipulate disks partitioned by other operating systems, such as the Macintosh®, Windows®, etc.
- ▶ Region Management Plug-ins - allow the EVMS to recognize, use, and manipulate the volumes and volume groups created by operating systems which employ a volume group based volume manager, such as AIX® or Linux with the Linux LVM.
- ▶ Feature Plug-ins provide the tools to build volumes, such as software RAID, encryption, n-way mirroring, drive linking, bad block relocation, etc.

EVMS Plug-in Modules (cont.)

- ▶ Filesystem Interface Plug-ins allow the EVMS to coordinate its activities on a volume with the filesystem present on that volume, such as when the EVMS is going to resize a volume or snapshot a volume.
- ▶ Cluster Management Plug-ins allow the EVMS to interface to the cluster management software controlling a cluster. EVMS is capable of operating in a cluster/shared storage environment. However, the EVMS built-in cluster support requires certain services which are normally provided by the cluster management software used to control the cluster. The Cluster Management Plug-ins allow EVMS to access these services.
- ▶ Distributed Lock Management Plug-ins allow the EVMS to make use of a Distributed Lock Manager (DLM), should one be available, when operating in a cluster/shared storage environment.

EVMS Runtime Diagram

Kernel IOCTLS

Kernel Buffer Cache

EVMS Runtime Driver

Generic Features

Generic Snapshot

Generic N-way Mirroring
(2002)

EVMS Features

Bad Block Relocation

Drive Linking

Encryption

Software RAID

EVMS Runtime Common Services

Region Managers

Linux LVM Volume Group Emulator

AIX LVM Volume Group Emulator

OS/2 Region Manager

EVMS Volume Group Emulator

Partition Managers

DOS Partition Manager

Linux LVM Partition Manager

AIX LVM Partition Manager

Device Managers

Local Device Manager

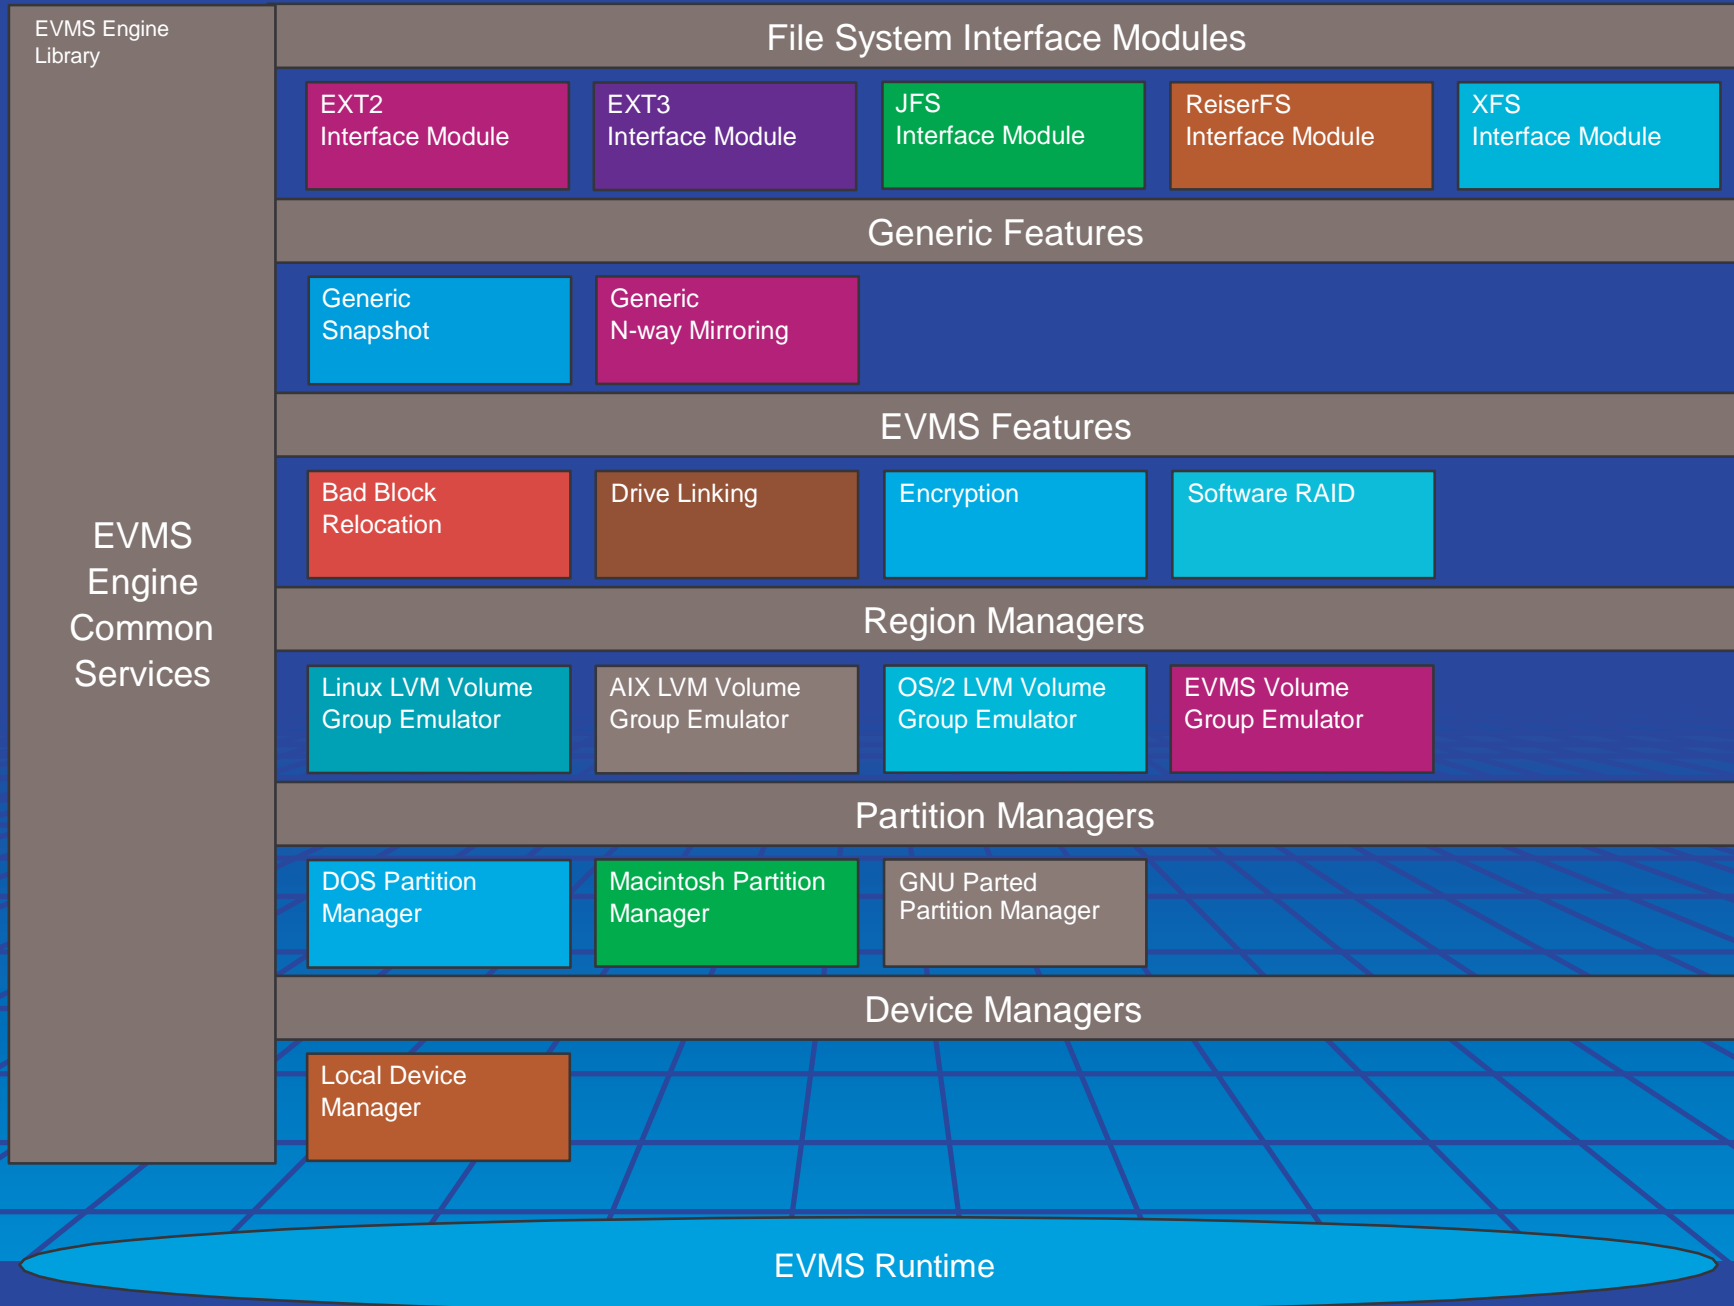
Kernel Device Drivers

EVMS Engine Diagram

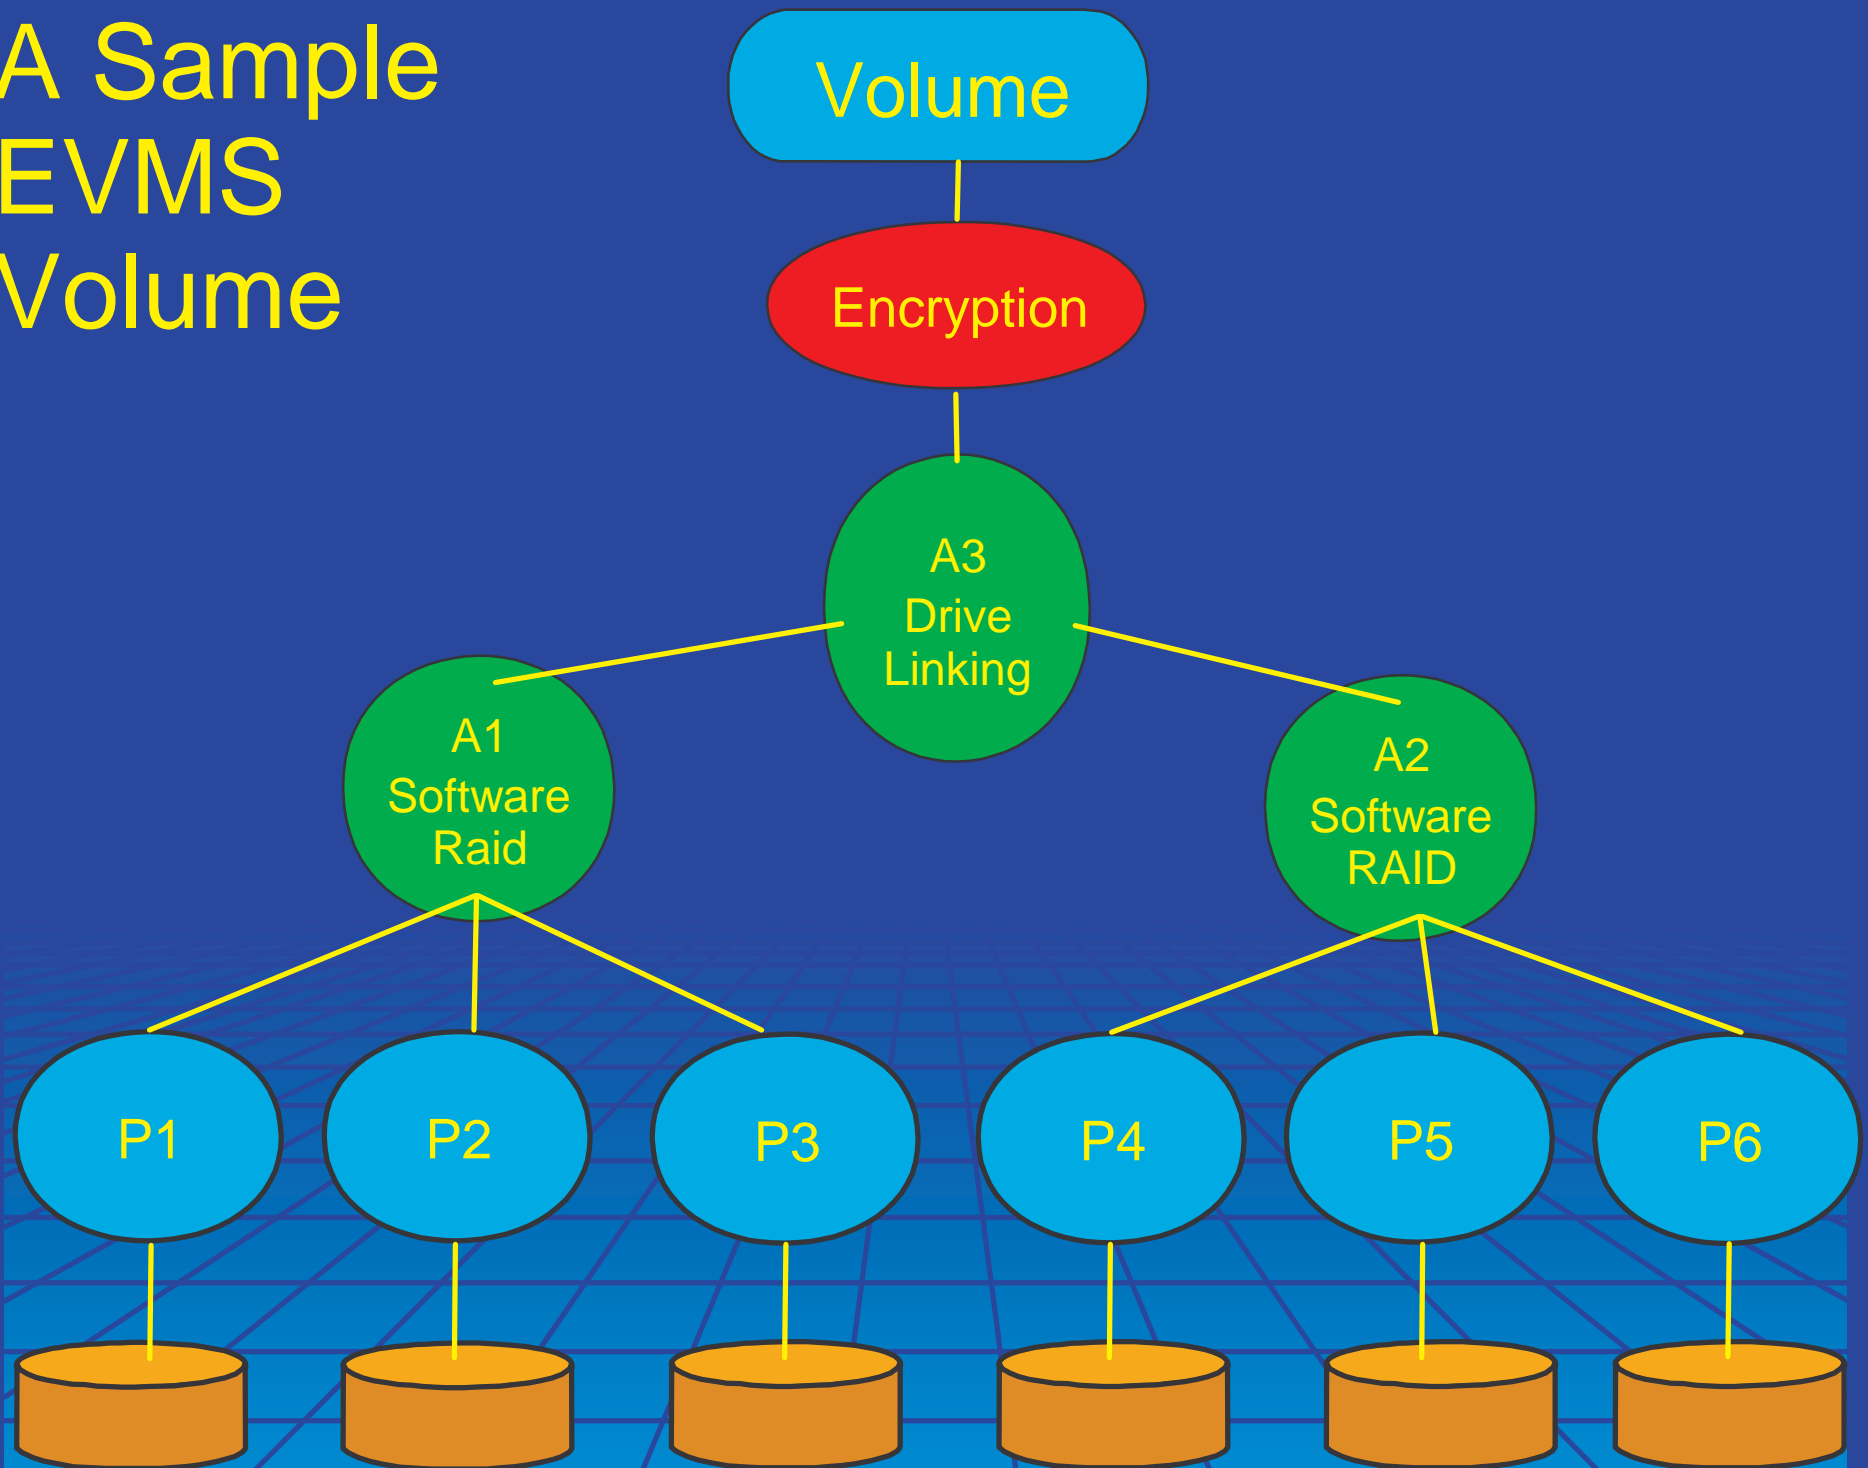
Command Line

Text UI

Graphical UI



A Sample EVMS Volume



Emulating other Volume Managers

- ▶ Since the interface to the EVMS is through the APIs provided by the EVMS Engine, multiple user interfaces can exist for the EVMS. This allows a user interface to be created which resembles the user interface of another volume manager.
- ▶ Plug-in modules can be created to allow the EVMS to understand the disk structure used by other volume managers.
- ▶ Thus, with the right user interface and plug-ins, it is possible for the EVMS to emulate another volume manager.

Summary of EVMS Features

- ▶ EVMS integrates all aspects of disk, partition, and volume management into a single coherent system
- ▶ EVMS has no design limit on the number of disks, partitions, or volumes that it can handle
- ▶ EVMS minimizes reboots due to disk, partition, or volume changes.
- ▶ Due to its highly modular design and support of plug-in modules, EVMS is very extensible
- ▶ EVMS is designed to be scalable
- ▶ EVMS is designed to support clusters
- ▶ With the proper plug-in modules, EVMS can read, write, and manipulate volumes created by other volume managers
- ▶ EVMS can reduce the costs associated with migrating to Linux

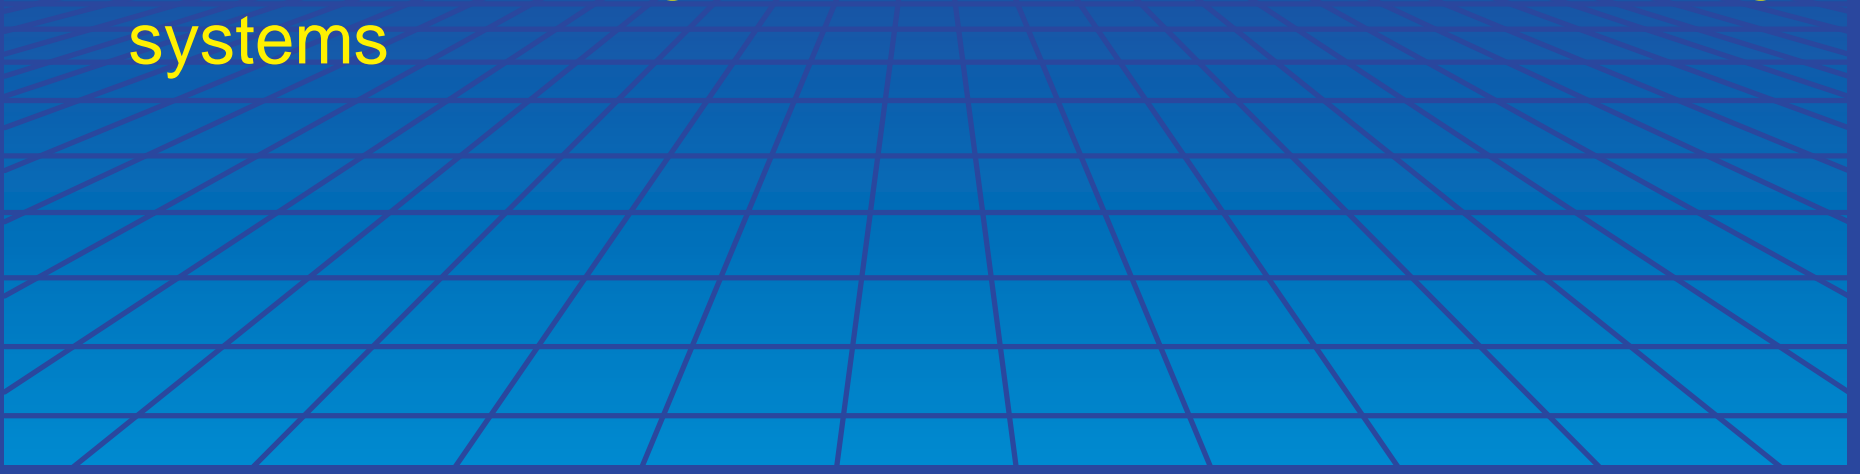
The EVMS Roadmap

- ▶ The EVMS project is hosted on SourceForge (<http://sourceforge.net/projects/evms>)
 - Live CVS tree
 - EVMS mailing list
 - Design documents available on-line
- ▶ Stable code drops approximately twice a month
- ▶ Conference calls with the community are held approximately twice a month
- ▶ EVMS V1.0 went GA on 3/28/2002 for all IBM® platforms
 - EVMS Kernel Code
 - EVMS Userspace Code
 - Linux MD emulation package
 - Linux LVM emulation package
 - Graphical User Interface
 - Command Line Interface
 - Text mode User Interface

The EVMS Roadmap (cont.)

- EVMS Generic Snapshot package
- EVMS Drive Linking package
- EVMS Bad Block Relocation package
- AIX read/write support
- OS/2® read/write support
- ▶ AIX Compatibility Package targeted for GA in the fourth quarter of 2002
- ▶ OS/2 Compatibility Package targeted for GA in the fourth quarter of 2002
- ▶ Additional planned enhancements for the EVMS:
 - EVMS Generic N-Way mirroring package
 - support for additional non-DOS partitioning schemes
 - Move and slide support
 - Cluster Support

Linux Community Involvement

- ▶ We are encouraging the Linux Community to take the lead in developing the following:
 - Encryption plug-in
 - Cluster Management plug-ins
 - Distributed Lock Management plug-ins
 - Partition Management plug-ins for non-DOS partitioning schemes
 - Emulation packages for the LVMs on other operating systems
- 

Notes:

The following terms are trademarks of International Business Machines Corporation in the United States, other countries, or both:

AIX

IBM

OS/2

Windows is either a registered trademark or trademark of Microsoft Corporation in the United States and/or other countries.

Other company, product, and service names may be trademarks or service marks of others.