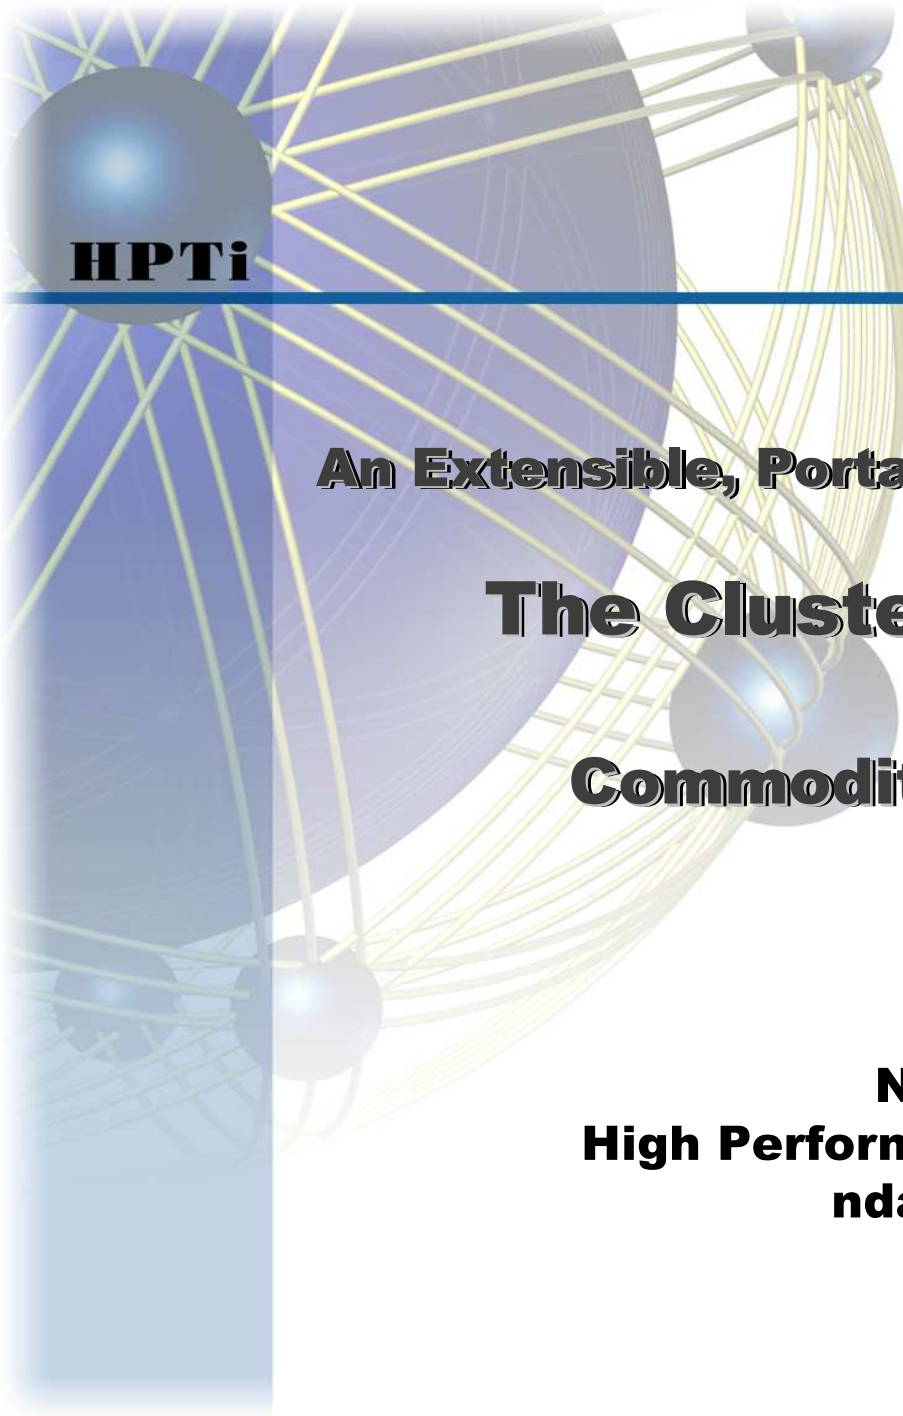

HPTi

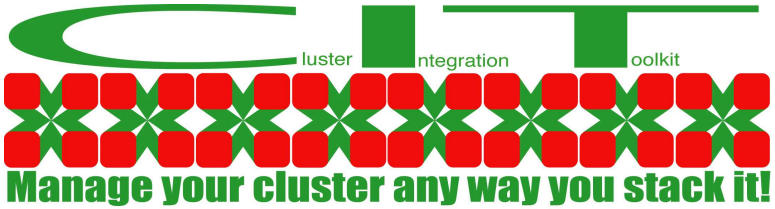

An Extensible, Portable, Scalable Cluster Management Solution

The Cluster Integration Toolkit

Commodity Cluster Symposium

July 22, 2003

Nathan Dauchy
High Performance Technologies, Inc.
ndauchy@hpti.com

Sandia National Laboratories

HPTi

What is the CIToolkit?

- Set of tools for configuring, testing, and managing a Linux cluster
- Software Architecture and Implementation
- Based on an easily extensible object oriented architecture
- Created to meet several goals, where other alternatives came up short

Sandia National Laboratories

Presentation Outline

HPTi

- Goals
- Key architecture points
- Key implementation details
- Tools in the Toolkit
- Cluster Integration Process
- Future work
- Q&A

Sandia National Laboratories

CIToolkit – Overall Goals

HPTi

- Provide the tools for an efficient and standardized cluster integration processes
- Automate as many install and implementation procedures as possible
- Leverage a single configuration for systems management, testing, performance management, etc.
- Increase reliability of the process to bring new clusters online
- Scale easily to 1000's of nodes, both in performance and manageability.

Sandia National Laboratories

HPTi

Additional Goals

- **Ease Cluster Management Tasks**
  - Supported by Unix system administrators not "cluster experts"
  - Same interface for every cluster, large or small
- **Target clusters in general rather than specific clusters used for specific purposes**
  - Usable as a platform to support various (and multiple) runtime environments
  - Support legacy, current, and future architectures (topologies and components)
  - ***Make as few assumptions as possible!***
- **No kernel mods required**

Sandia National Laboratories

More CIT Features

HPTi

- **Open Source (LGPL)**
- **Manage non-node devices**
  - Terminal servers, myrinet switches, etc.
- **Integrated, extensible, diagnostic suite**
- **Integrated hardware issue tracking database**
- **Software is as generic as possible**
  - Anything that can be considered site-specific is isolated. (eg. naming conventions)
  - Easy to find and modify if necessary
- **Easy distribution, installation, and configuration of 3rdparty management tools (open source, of course)**

Sandia National Laboratories

Key Architecture Points

HPTi

- **Device Class Hierarchy**
  - (device drivers)
- **Persistent Object Store**
  - (database)
- **Layered Utilities**
  - (tools)

Sandia National Laboratories

Architecture in Detail

HPTi

- *"An Extensible, Portable, Scalable Cluster Management Software Architecture"*
- **Published in the proceedings of the International Conference on Cluster Computing (Cluster 2002) October 2002.**

Sandia National Laboratories

Device Class Hierarchy

HPTi

- Hierarchical organization of devices and their capabilities
- “Device drivers” to provide the toolkit with consistent access to devices and their capabilities.
- Object Oriented provides Inheritance

Sandia National Laboratories

Device Class Hierarchy

HPTi

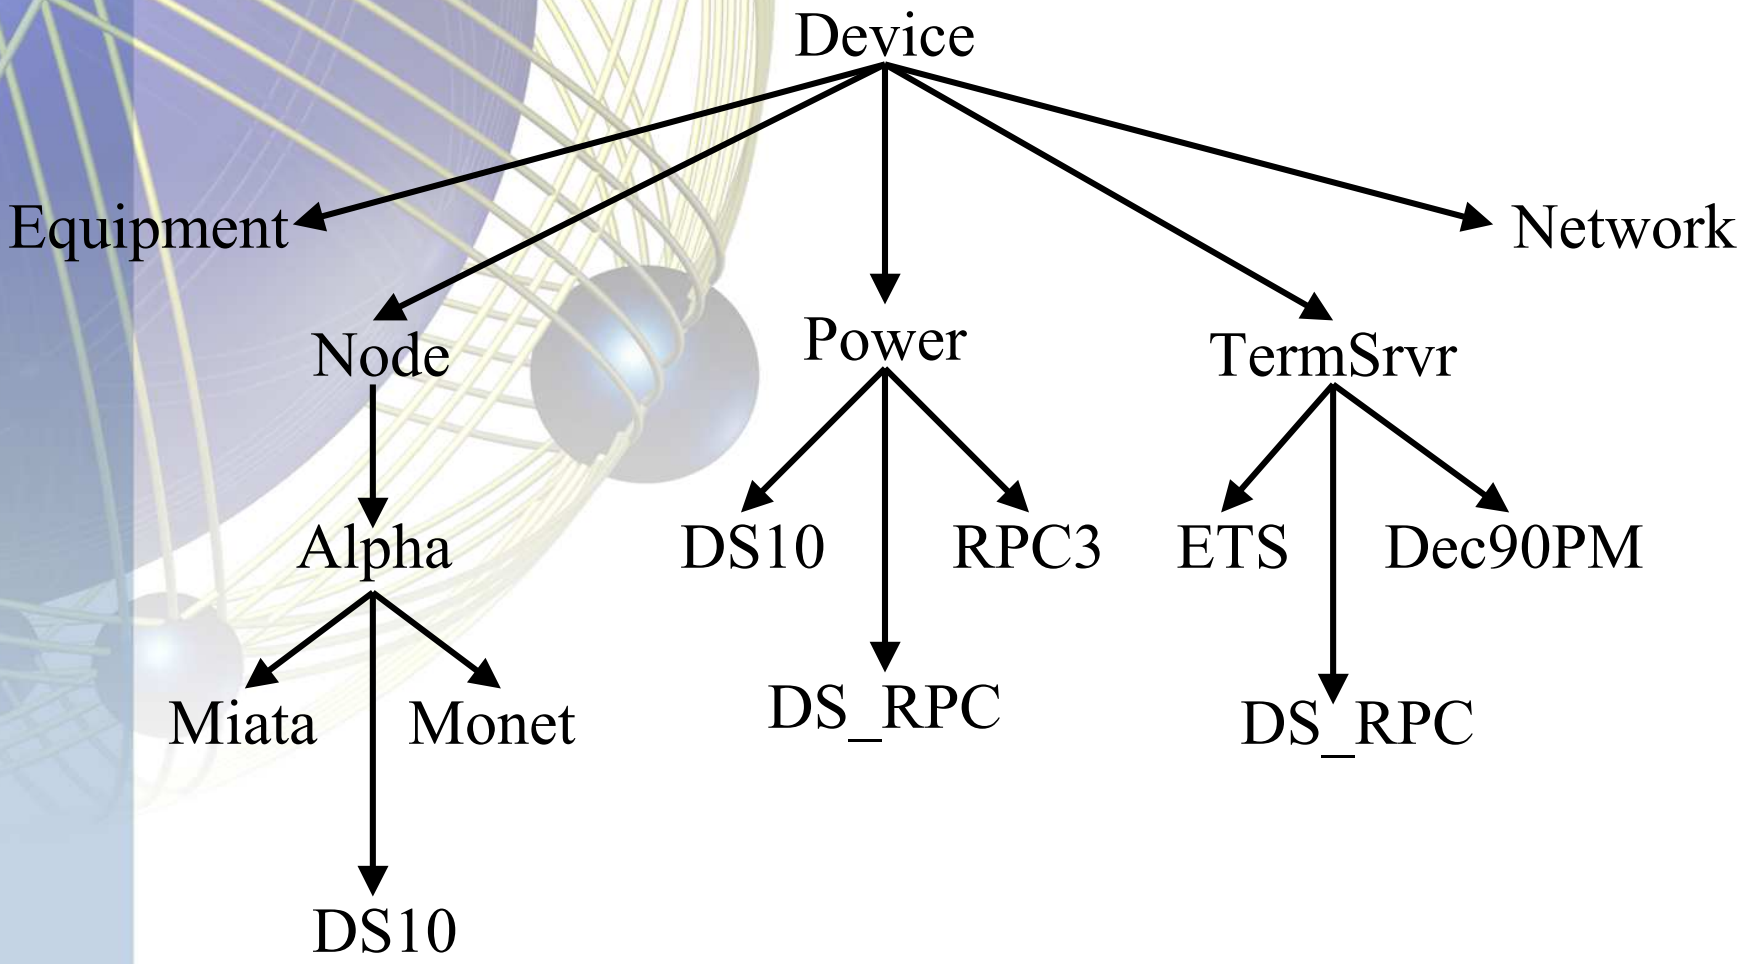

Sandia National Laboratories

Adding a Device Class

HPTi

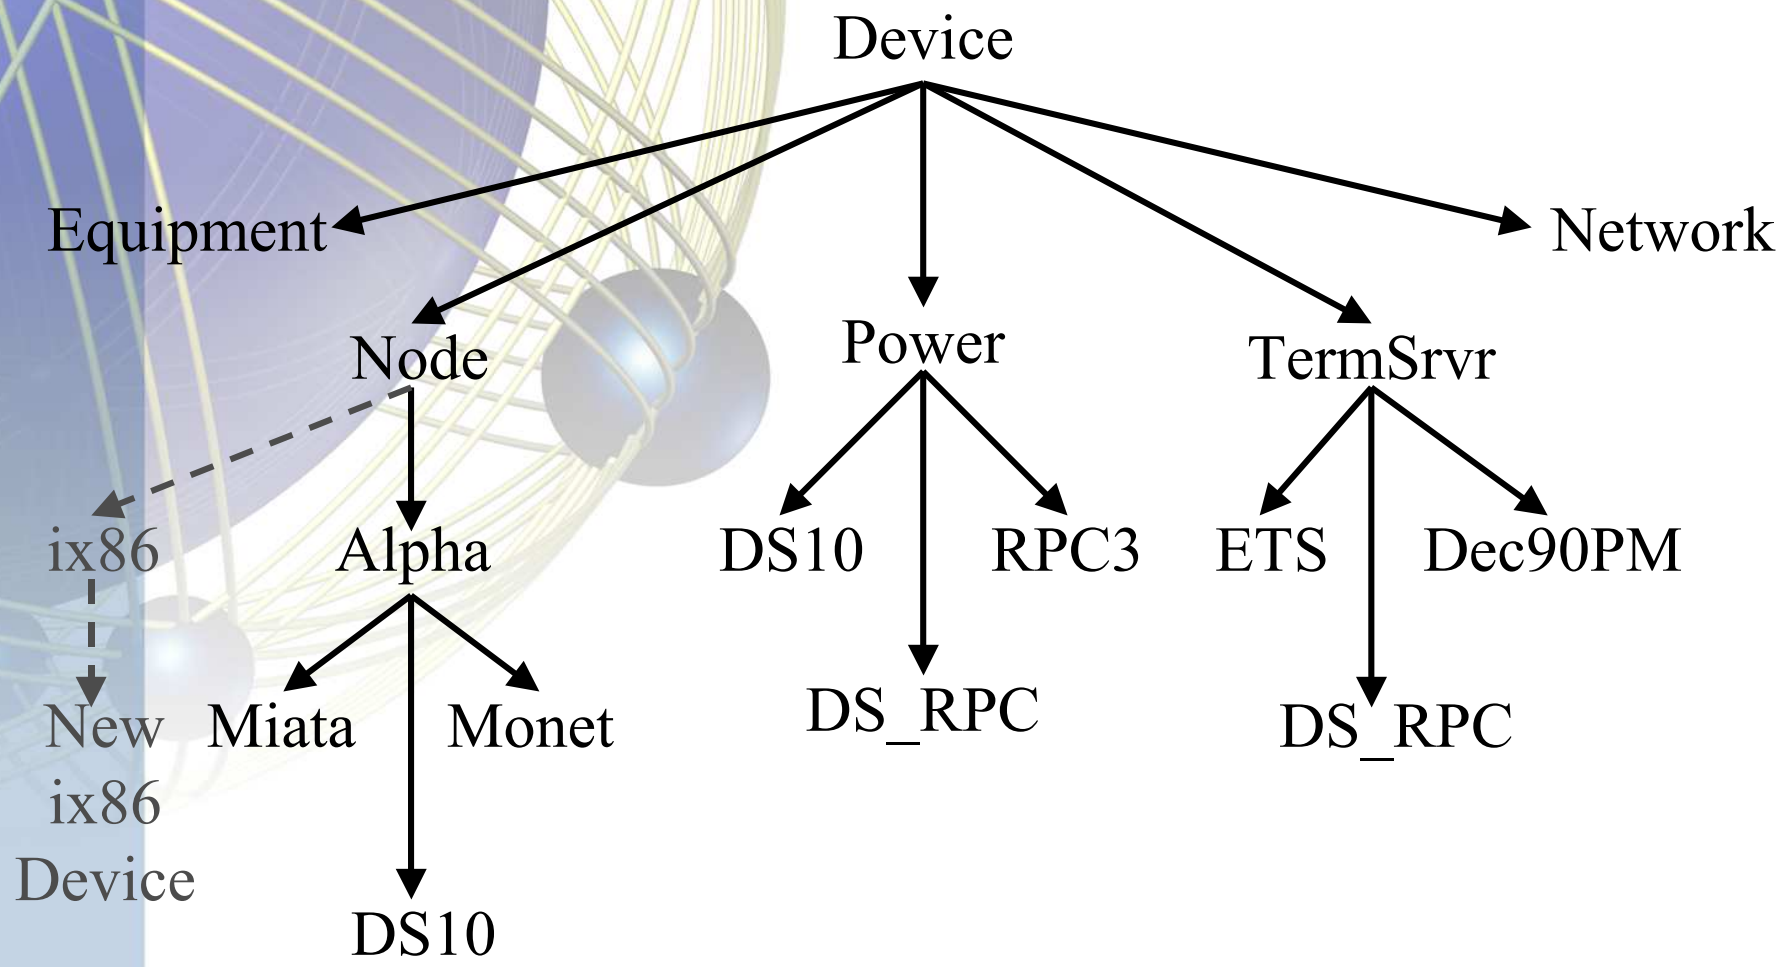

Sandia National Laboratories

Persistent Object Store

HPTi

- Representation of the Physical Cluster (hardware and topology)
- Instantiated Objects from Device Classes
- Linkage describes the specific topology of the cluster
- Provides foundation for tools
- **Database Interface Layer**
  - Supports GDBM, LDAP, SQL, etc

Sandia National Laboratories