GRID
Information Services

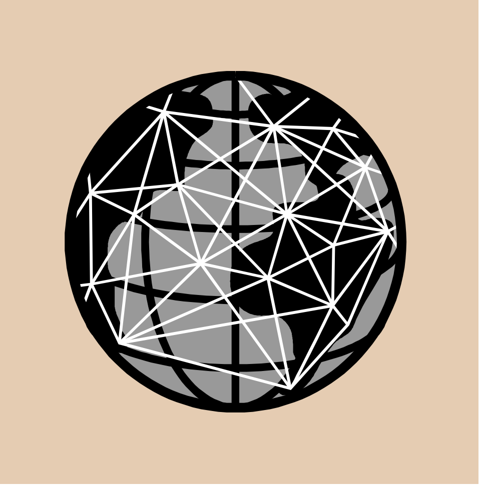

Panickos Neophytou

Department of Computer Science
University of Cyprus

4th CrossGrid meeting

09/10/2002

Outline of presentation

- The Grid and Grid architecture
- Information Services Definition
- LDAP (Lightweight Directory Application Protocol)
- Current Implementation of Information Services
- Description of Data Models in GRID

About the GRID in general

- Grid focuses on:
  - Large-scale resource sharing
  - Innovative applications
  - High-performance orientation
- “Grid problem”
  - Flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions and resources – virtual organizations.
  - Unique authentication, authorization, resource access, resource discovery and other challanges.

Some Definitions

- *Virtual Organization:* A group of people or institutions with some common purpose or interest that need to share (computer) resources to further their objectives.

- *Resource:*
  - An entity that is to be shared
  - Not always a physical entity
  - Defined in terms of interfaces, not devices.

Grid Architecture.

"Coordinating multiple resources": ubiquitous infrastructure services, app-specific distributed services

"Sharing single resources": negotiating access, controlling use

"Talking to things": communication (Internet protocols) & security

"Controlling things locally": Access to, & control of, resources

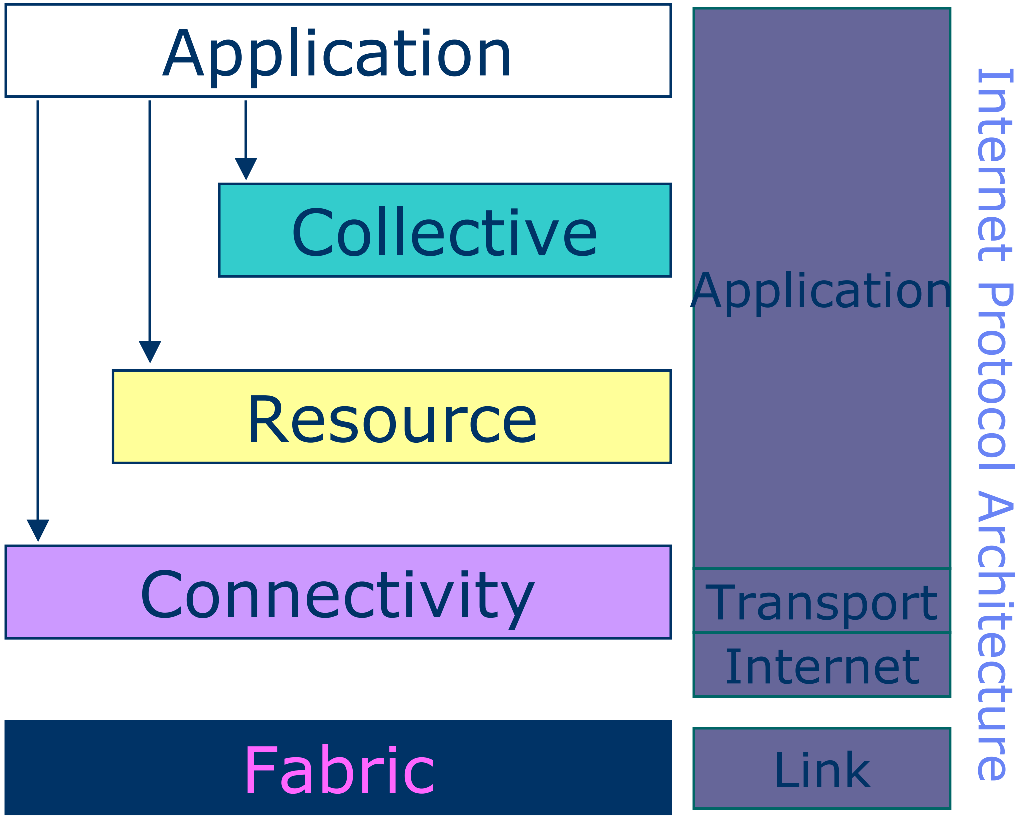

Information Services in Resource and Collective layers.

- IS in Resource layer
  - Grid Resource Information Service (GRIS)
    - Access to structure & state information
- IS in Collective layer
  - Index servers aka metadirectory services
    - Custom views on dynamic resource collections assembled by a community

Information Services

Information Services

- System information is critical to operation of the grid and construction of applications
  - What resources are available?
    - Resource discovery
  - What is the “state” of the grid?
    - Resource selection
  - How to optimize resource use
    - Application configuration and adaptation?
- We need a general information infrastructure to answer these questions

What are Information Services?

- *Information Services* are a vital part of any Grid software infrastructure, providing fundamental mechanisms for **discovery** and **monitoring** and thus for **planning** and **adapting application behavior.**
- *Information Services* have to be designed to support the **initial discovery** and **ongoing monitoring** of the existence and characteristics of **resources**, **services**, **computations**, and **other entities.**

Examples of applications

- *Service Discovery*
- *Superscheduler*
- *Replica selection service*
- *Application adaptation agent*
- *Troubleshooting*
- *Performance diagnosis tool*

Natural and other problems

- Information is always old
  - Time of flight, changing system state
  - Need to provide quality metrics
- Distributed state hard to obtain
  - Complexity of global snapshot
- Component will fail
- Scalability and overhead
- Many different usage scenarios
  - Heterogeneous policy, different information organizations, etc.

GIS Requirements

- 1. Distribution of information providers

Information sources are necessarily distributed and individual sources are subject to failure.

Any information delivered will be old

• Information producers should use timestamps and time-to-live metadata.

• Information service should transport information as rapidly and efficiently as possible.

GIS Requirements (cont.)

- 2. Managing Failures
  - Information service should behave robustly in the face of failure of any component.
  - Failure should not prevent users from gaining information, even if this is partial or inconsistent.

•Information services should be as distributed as possible.

•Information services components should be built under the assumption that failure is the rule not the exception

GIS Requirements (cont.)

- **3. Diversity in Information Systems**
  - A new VO may involve many entities and have unique requirements for discovery and monitoring

Define once, ahead of time, the discovery and enquiry mechanisms that must be supported by any Grid entity.

  - Information is often provided with restrictions

•We must have robust authentication and authorization mechanisms that information owners will trust.
•Robust providers may wish to assert policy over which VOs they are prepared to join.

Architecture Overview

- Two fundamental entities:
  - Information Providers
    - Common VO-neutral infrastructure
    - Providing access to detailed dynamic information about GRID entities
  - Specialized aggregated directory services
    - Specialized VO-specific views of federated resources services etc.

Architecture Overview (cont.)

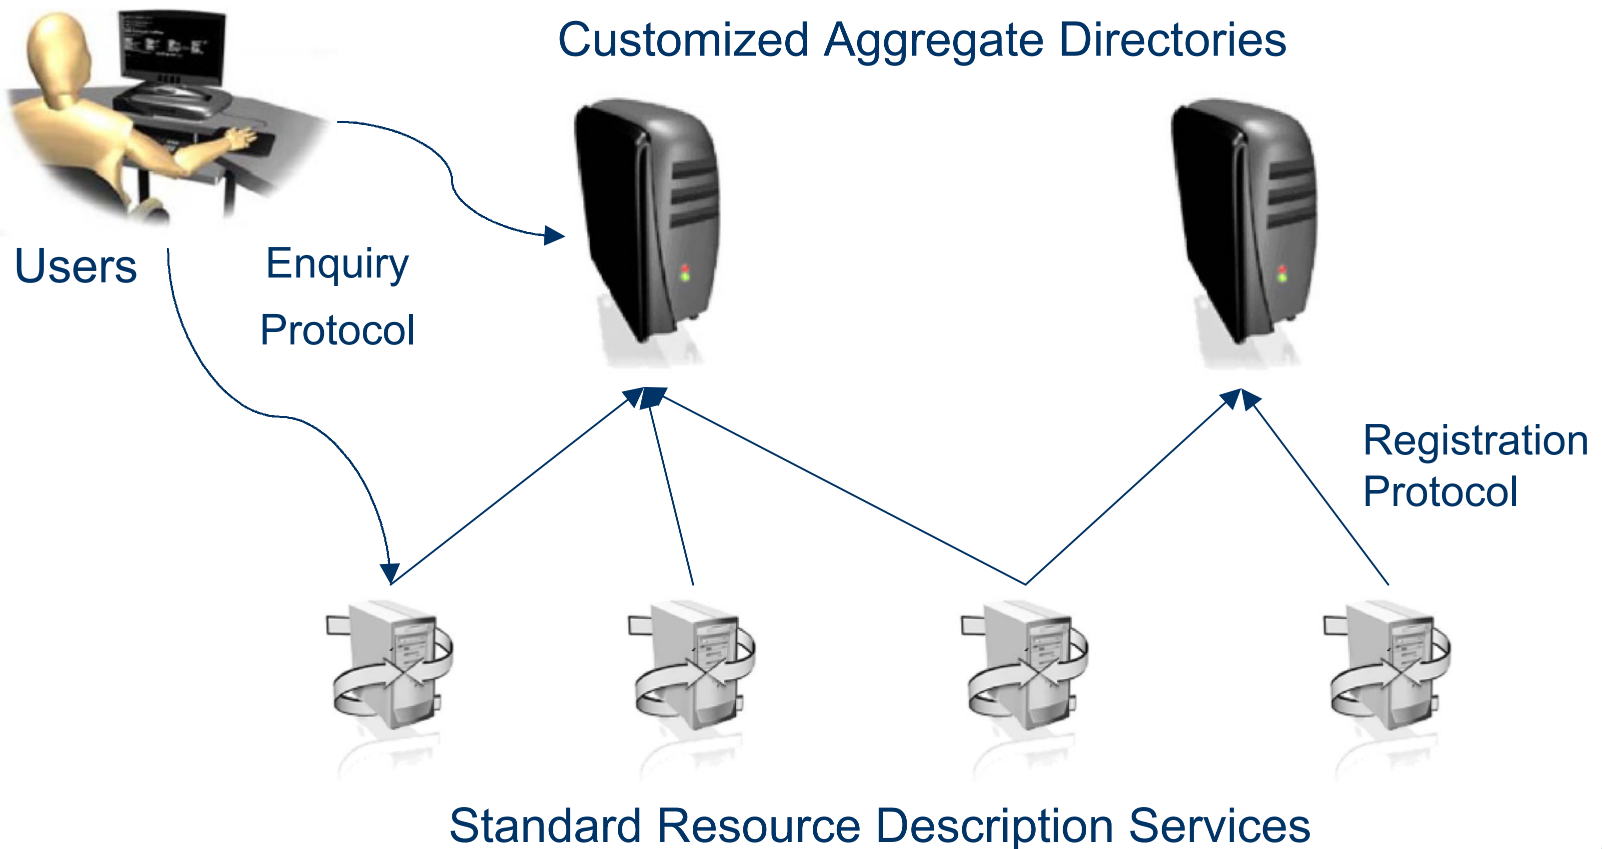

Security Issues

- The information providers and aggregate directory have the same data access policy and the providers trust the directory.
- The information providers limit the information that is available to an aggregate directory.
- The information provider makes no information known to other than its existence
- The information provider places no restriction on the information provided.

LDAP and the Grid

LDAP and the GRID

- Lightweight Directory Access Protocol
  - Stripped down version of X.500 DAP protocol
  - Supports distributed storage/access (referrals)
  - Supports authentication and access control
- Defines:
  - Network protocol for accessing directory contents
  - Information model defining form of information
  - Namespace defining how information is referenced and organized
- Example:
  - Domain Name Service (DNS)

LDAP and the GRID (cont.)

- The information model and namespace are based on entries.
- An entry is used to store attributes.
- An attribute is an associated type and can have one or more values.
- Each entry in a namespace has a distinguished name that allows easy identification.

LDAP and the GRID (cont.)

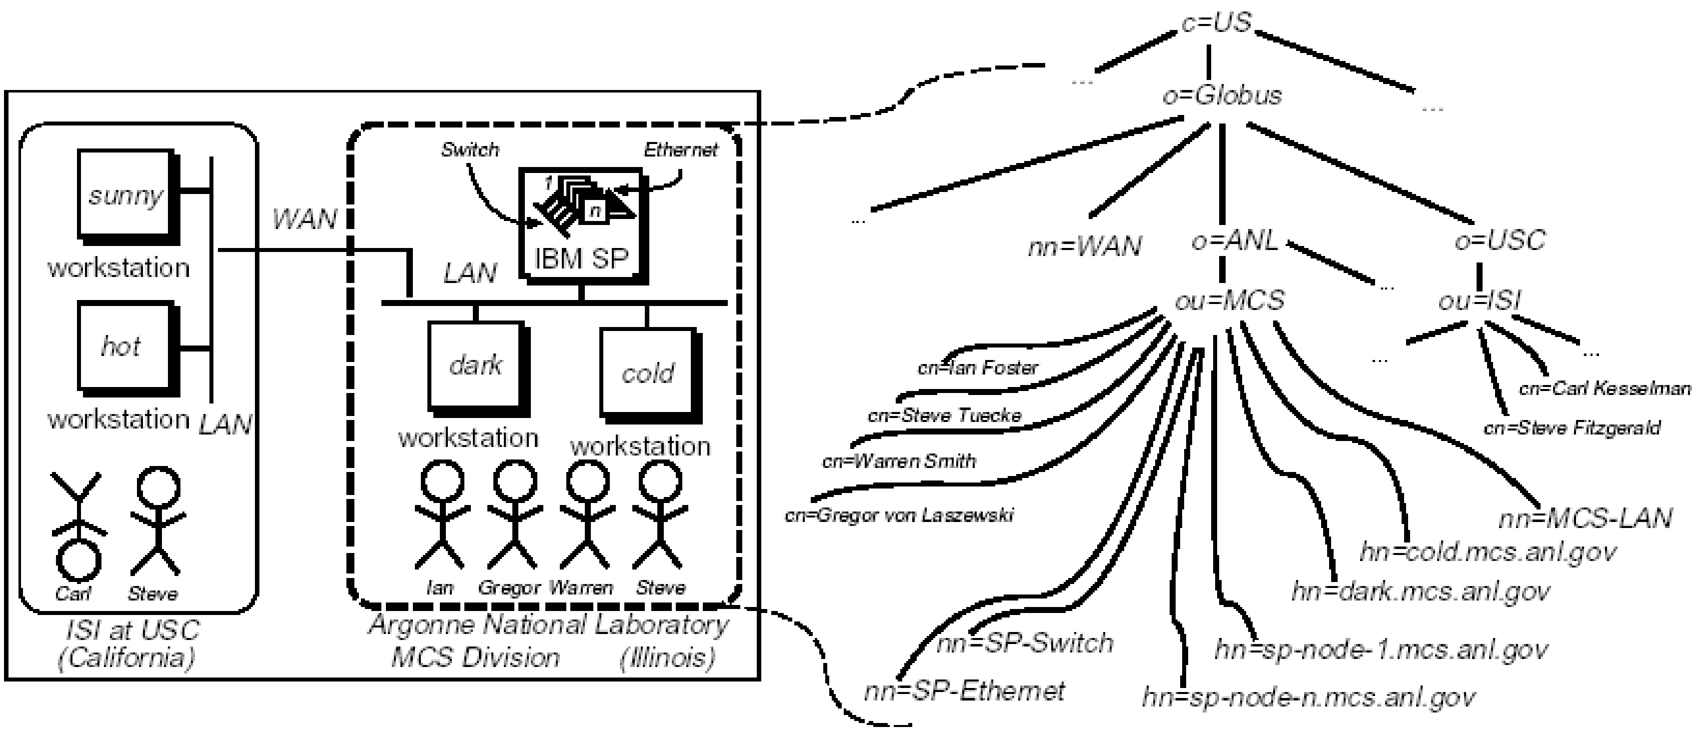

Current Implementation of Grid Information Services

the globus project
www.globus.org

Globus
MDS-2

Base protocols

- GRid Information Protocol – GRIP
  - Users, aggr. Directories and applications use GRIP to obtain information from an information provider about the entities the provider possesses information.
  - Discovery and enquiry are supported.

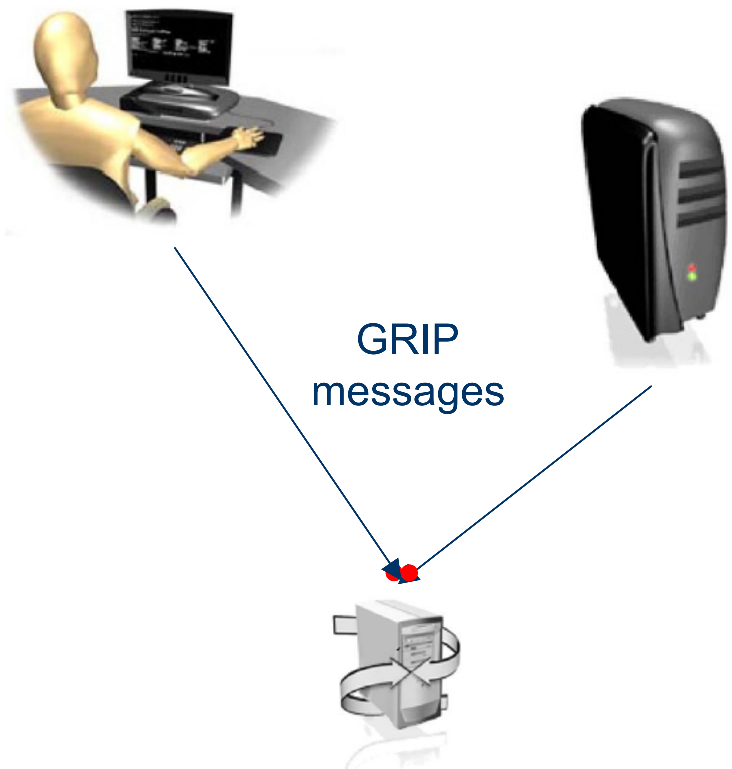

Base protocols (cont.)

- Implementation of GRIP:
  - Using OpenLDAP.
- Limitations of Information Protocol
  - LDAP query language cannot specify relational “joins”

Base protocols (cont.)

- Grid Registration Protocol – GRRP
  - Used by an information provider to notify an aggr. directory of its availability for indexing
  - Or by an aggr. directory to invite an information provider to join a VO.

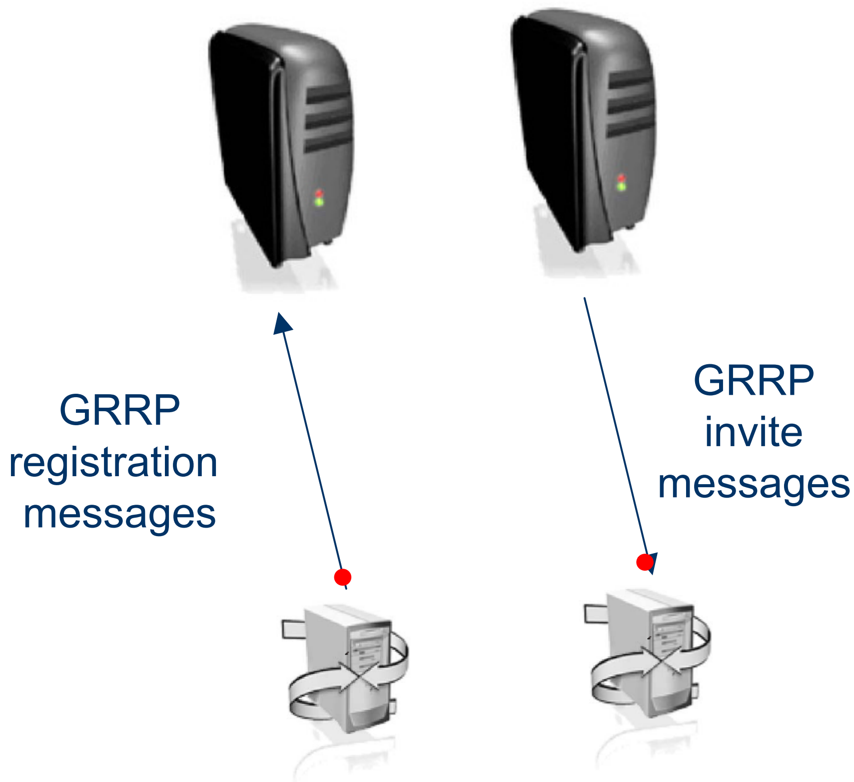

Base protocols (cont.)

- GRRP is a soft-state protocol.
  - Information can be discarded unless refreshed by a stream of subsequent notifications.
  - This way is resilient to failure.

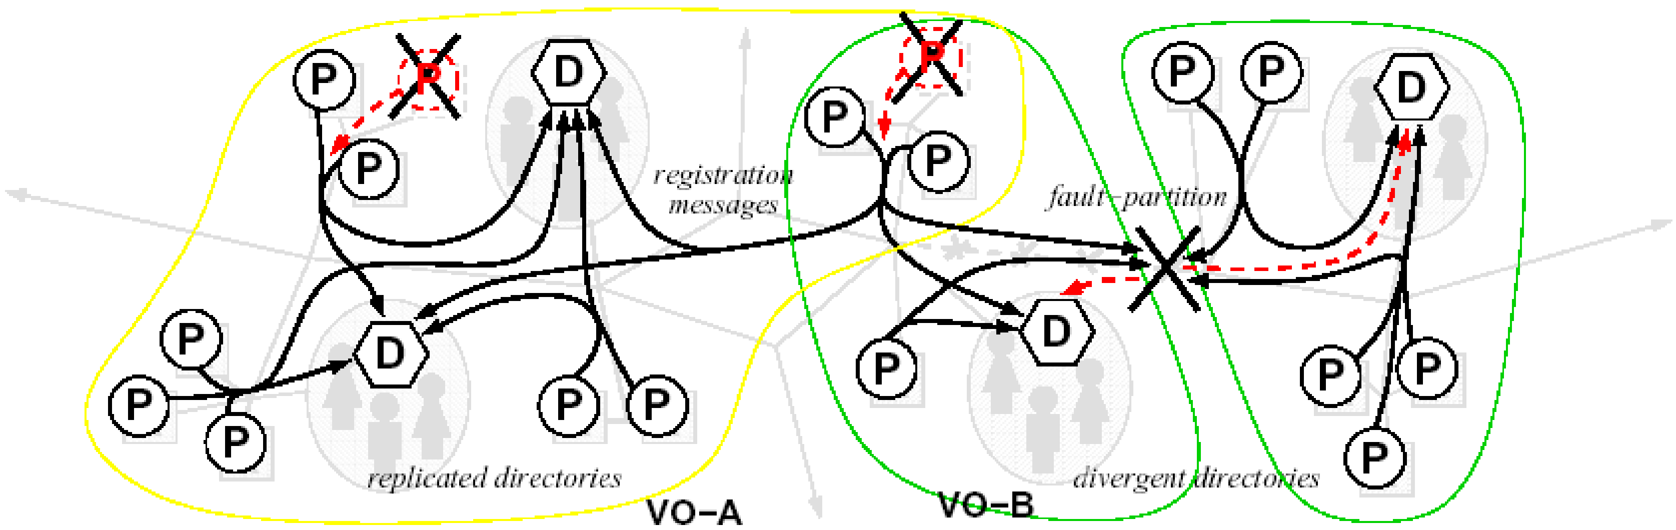

GRIS - Grid Resource Information Service

- Server which runs on each resource
  - Given the resource DNS name, you can find the GRIS server (implemented as an OpenLDAP server backend)
- Provides resource specific information by parsing each GRIP request
  - Much of this information may be dynamic
    - Load, process information, storage information, etc.
    - GRIS gathers this information on demand
- “White pages” lookup of resource information
  - How much memory does machine have?
- “Yellow pages” lookup of resource options
  - Which queues on machine allows large jobs?

GIIS - Grid Index Information Service

- GIIS describes a class of servers
  - Gathers information from multiple GRIS servers via GRRP messages
  - Each GIIS is optimized for particular queries
    - Ex1: Which Cygrid machines are >16 process SGIs?
    - Ex2: Which Cygrid storage servers have >100Mbps bandwidth to host X?
  - Relative to web search engines
- Organization GIIS
  - The Globus Toolkit ships with one GIIS
  - Caches GRIS info with long update frequency
    - Useful for queries across an organization that rely on relatively static information (Ex1 above)
- Can be merged into GRIS

Logical MDS Deployment

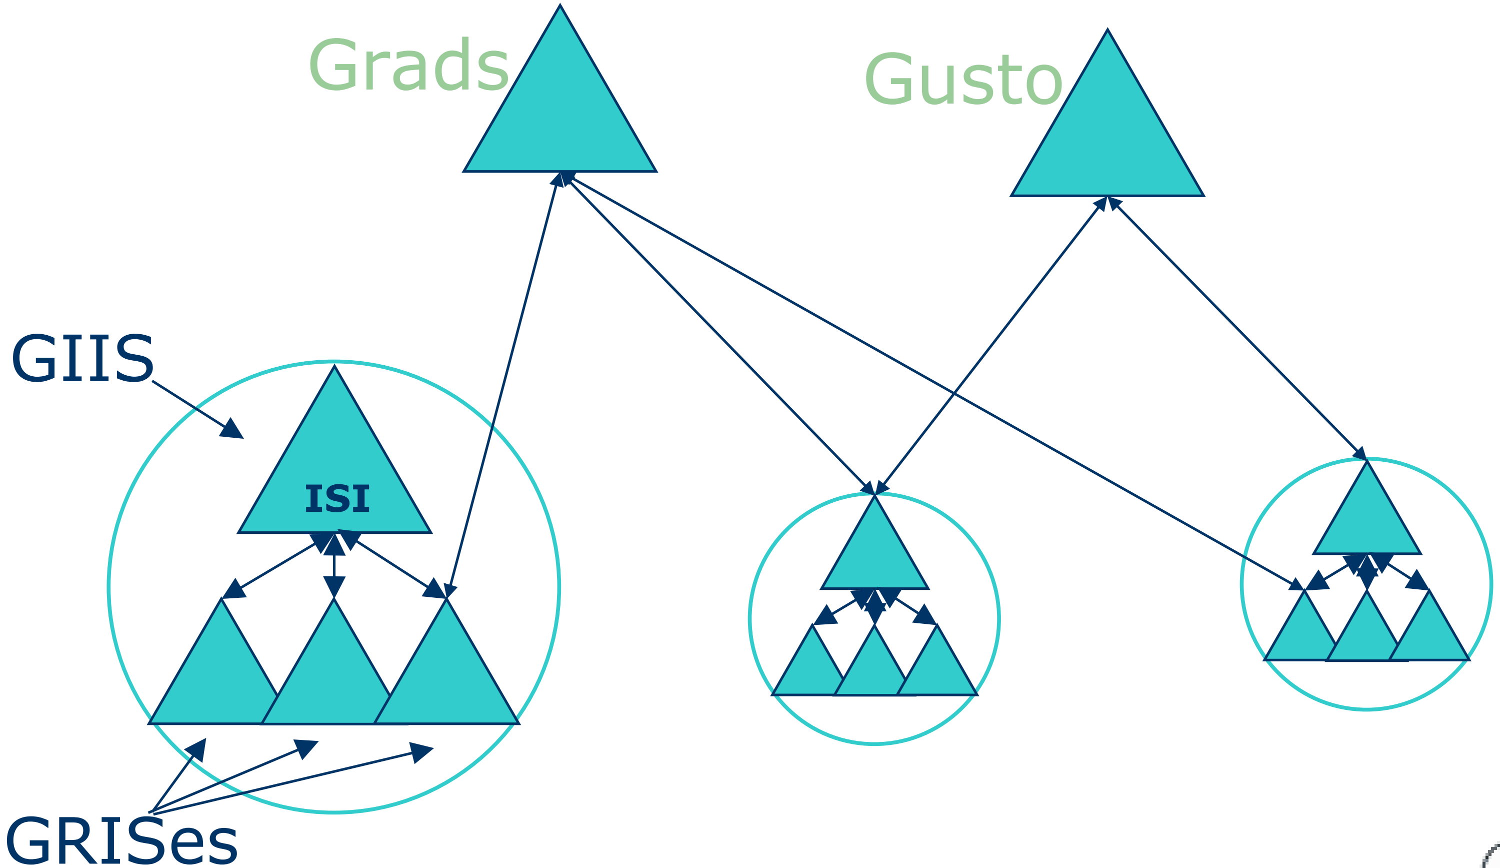

Description of current Data Modeling with LDAP

- Structural information
  - Resource hierarchy maps to objects
  - Named positions in LDAP DIT
- Merged information
  - Some parents “join” child data
  - Simplifies common query patterns
- Auxiliary information
  - Uniform representation of leaf/parent data
  - Uses LDAP auxiliary objectclasses

MDS 2.1 Information Model

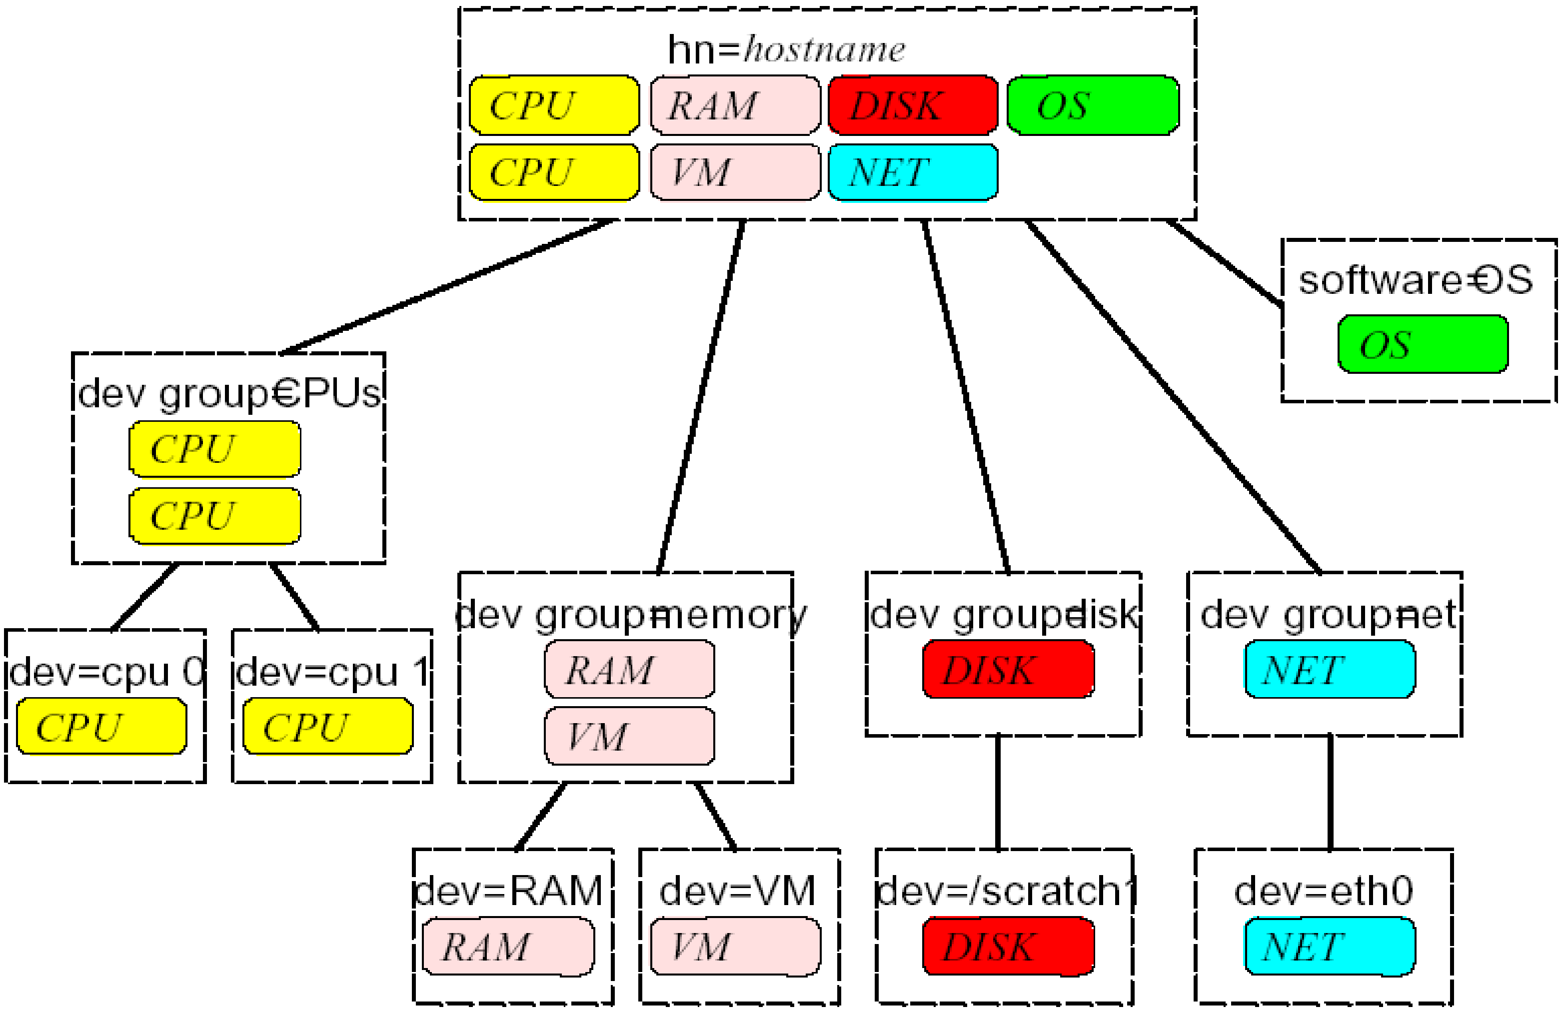

GRIS Object Hierarchy

Mds-Host-name=*hostname*

- Mds-Software-Deployment=operating system
- Mds-Device-Group-name=processors
  - Mds-Device-name=cpu 0
- Mds-Device-Group-name=memory
  - Mds-Device-name=physical memory
  - Mds-Device-name=virtual memory
- Mds-Device-Group-name=filesystems
  - Mds-Device-name=/scratch1
  - Mds-Device-name=/scratch2
- Mds-Device-Group-name=networks
  - Mds-Device-name=eth0

GRIS Structural Class Hierarchy

Mds

Attr: Mds-validfrom (like createtime)

Attr: Mds-validto (accuracy metadata)

Attr: Mds-keepto (discard metadata)

MdsHost

MdsDevice

MdsDeviceGroup

MdsSoftwareDeployment

- Every MDS object: name, time metadata

GRIS Auxiliary Class Examples

MdsCpu

- Attr: Mds-Cpu-vendor
- Attr: Mds-Cpu-model
- Attr: Mds-Cpu-speedMHz

MdsCpuCache

- Attr: Mds-Cpu-Cache-L1kB

MdsCpuSmp

- Attr: Mds-Cpu-Smp-size

MdsCpuTotal

- Attr: Mds-Cpu-Total-count

Conclusions

Conclusions

- The implementation of this architecture, MDS-2, has been widely deployed in a number of different configurations as part of the Globus.

Future work.

- Explore the construction of different and more specialized types of aggregate directories
- Investigate update versus freshness tradeoffs in directory implementation
- Explore applications in different settings and domains
- Develop flexible configuration tools to enable lightweight VO formation and
- Extend security models to incorporate capabilities and delegation to enable more sophisticated directory construction and caching of information provider values.

End of presentation

Closing with a flyby over the TeraGrid network

Argonne
I-WIRE
NCSA
NSF DTF Backbone
Multiple 10 GbE
Caltech
SDSC