Grid Computing—The Hype, the Truth and the progression of ideas

Sudharshan Vazhkudai

Network and Cluster Computing, CSMD

Oak Ridge National Laboratory

http://www.csm.ornl.gov/~vazhkuda

vazhkudaiss@ornl.gov

Acknowledgments: ORNL

Outline

- The Evolving Computing Landscape
- The Evolving Storage Landscape
- The Evolving Networking Landscape
- Grids

Evolving Computing Landscape

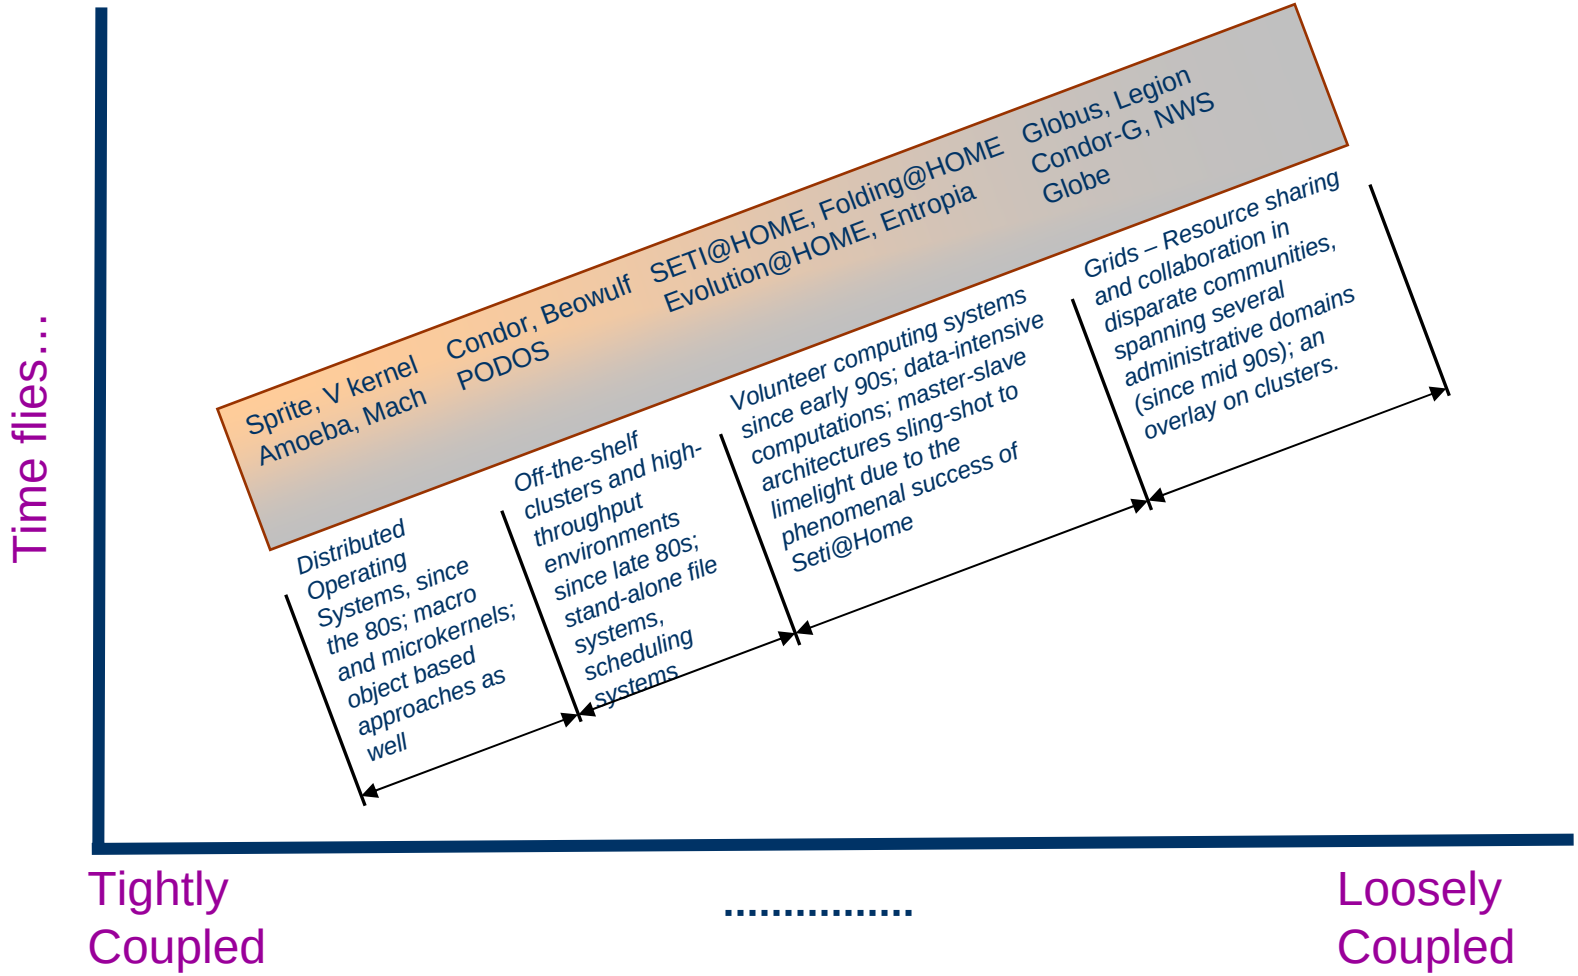

Distributed Operating Systems

- Highlights:
  - Transparency
  - Single system image
  - Performance
  - Resource Sharing
  - Availability
  - Fault Tolerance
- Illustrious Examples:
  - Sprite
  - Amoeba
- Drawbacks:
  - Commercial viability?

In Search for Clusters…

- Highlights:
  - COTS
  - Excellent price/performance ratio
  - Relaxed constraints
  - High sustained throughput
- Illustrious Examples:
  - Beowulf
  - Condor
- Drawbacks:
  - Single domain
  - High maintenance

Volunteer Computing Systems

- Highlights:
  - Harnessing idle compute cycles from Internet PCs
  - Amassing tremendous aggregated CPU power
  - Loosely coupled
  - Divide & Conquer
  - Client Server paradigm
- Illustrious Examples:
  - Seti@Home
  - Folding@Home
- Drawbacks:
  - Trust
  - Malicious users
  - Problem specific
  - User contribution proportional to mass appeal of project

Evolving Storage Landscape

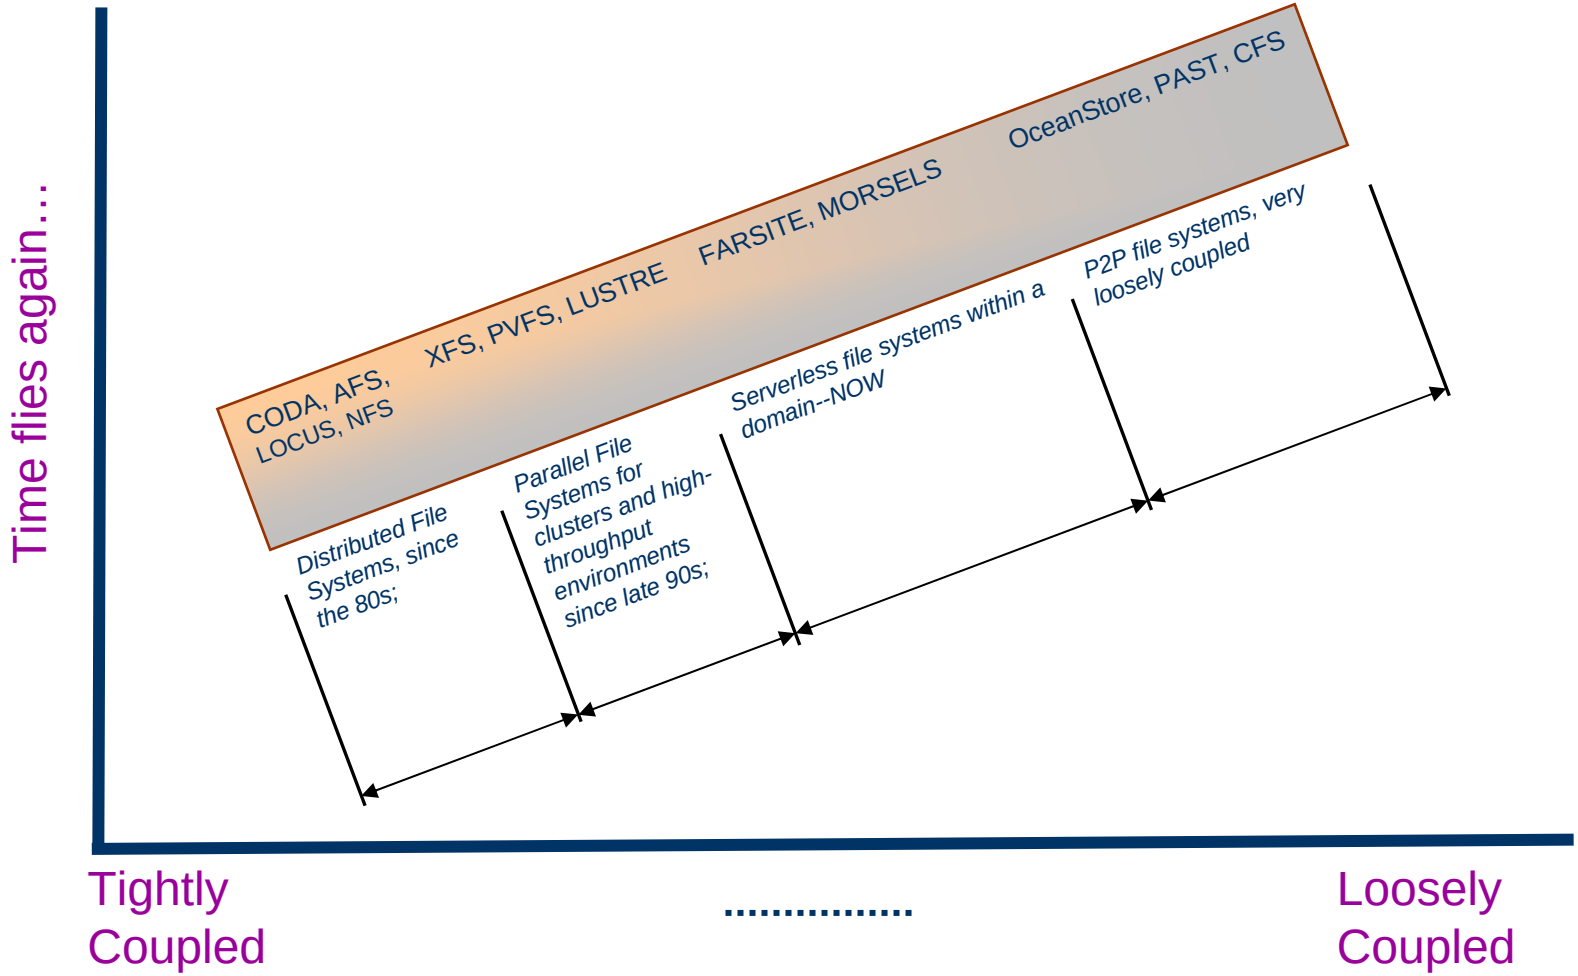

Family of File Systems

- Distributed File Systems:
  - Single system image (LOCUS)
  - Availability (Zebra)
  - Fault Tolerance (CODA)
- Networked File Systems:
  - Performance (NFS, PODOS)
- Parallel File Systems:
  - High performance for clusters
  - High-speed connectivity
  - Tightly coupled environments
- Serverless File Systems:
  - Resource sharing
  - Loosely coupled but within a domain (xFS, FARSITE, Morsels)
  - Small file sizes except in Morsels
- P2P File Systems
  - Very loosely coupled
  - Fault tolerant
  - Idle storage harnessing

Evolving Networks

- Simple LANs
- High-speed cluster interconnects: infiniband, myrinet, etc
- Optical Links
- TeraGrid

- UltraNet

Factoids

3 DVDs?

OC192	10Gbps
OC48	2.4Gbps
OC12	622Mbps
OC3	155.5Mbps
OC1	51.8Mbps
T3	44.7Mbps
T1	1.54Mbps

Meta-Message: *When multigigabit LANs and WANs become ubiquitous, and when message passing applications can tolerate high latency, the Grid becomes a Beowulf.* – Jim Gray

Why Grids…?

Challenge for the 21st Century

- All the kinds of problems we can solve using distributed data-intensive science
  - High Energy Physics (GriPhyN, EUDataGrid, PPDG)
  - Digital Sky Survey (SDSS)
  - Earthquake Simulations (NEESgrid)
  - SETI@Home, FightAIDS@Home, Folding@Home Evolution@Home
  - Bio Terrorism (Small Pox Research Grid)
  - BioInformatics, Geo Spatial Analysis
  - Medical data, digitizing patient records (Health Grid)

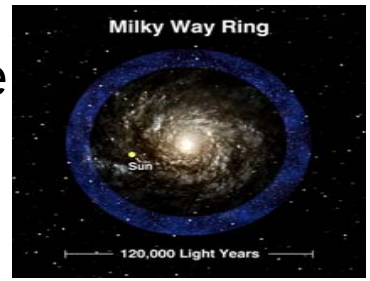

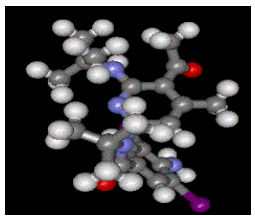

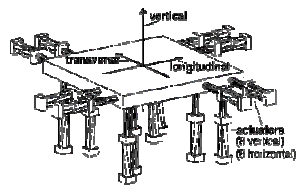

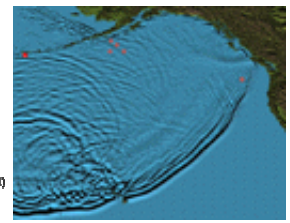

Technology Trends

- Moore's Law: Price/Performance doubles every 18 months
- Storage capacity beating Moore's Law – doubles every 12 months
- What does it mean…?
  - New Storage = sum of all old storage
- Gilder's Law: Networks are improving every 8 months
  - WANs are becoming faster than LANs!
- So, Distributed Data-Intensive Science is the Next Step in Evolution…
- E. coli doubles every 20 minutes!

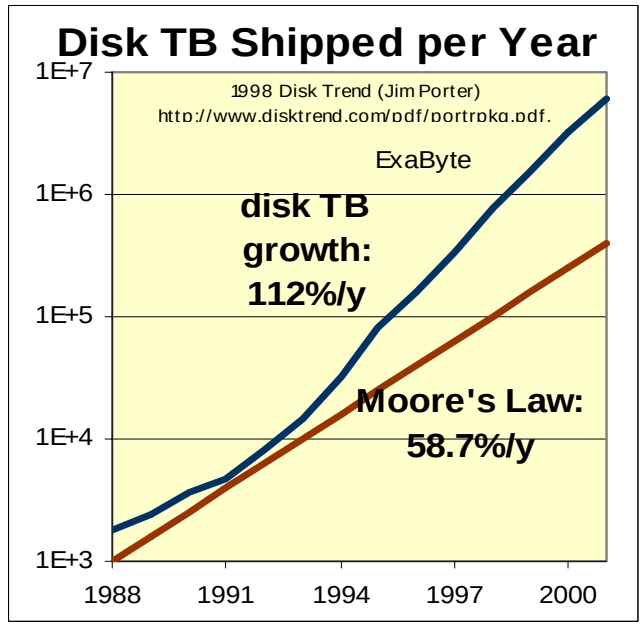

Meta-Message: Should we spend millions building petaflop supercomputers or spend it in building distributed software infrastructure?

So, What is the "Not-so-good" News?

- The Data Deluge…
  - 2000 ~0.5 Petabyte
  - 2005 ~10 Petabytes
  - 2010 ~100 Petabytes
  - 2015 ~1000 Petabytes?
- Problem:
  - Massive data collections
  - Computationally intensive analyses
  - Rapid access to large subsets
  - Data, computers, and users distributed across networks of varying capability
  - How to collect, manage, access and interpret such huge quantities of data?

Factoids

10^{24} Yotta

Zetta

Exa

10^{15} Peta

Tera

Giga

Mega

Kilo

Meta-Message: "Remember this: With great power comes great responsibility!!" – Spiderman

What are Grids anyway?

- Grids Galore:
  - Data Grids, Compute Grids, Commodity Grids, Science Grids, Terra Grids, PC Grids, Bio Grids, Rent-A-Grid?
- The Grid Problem:

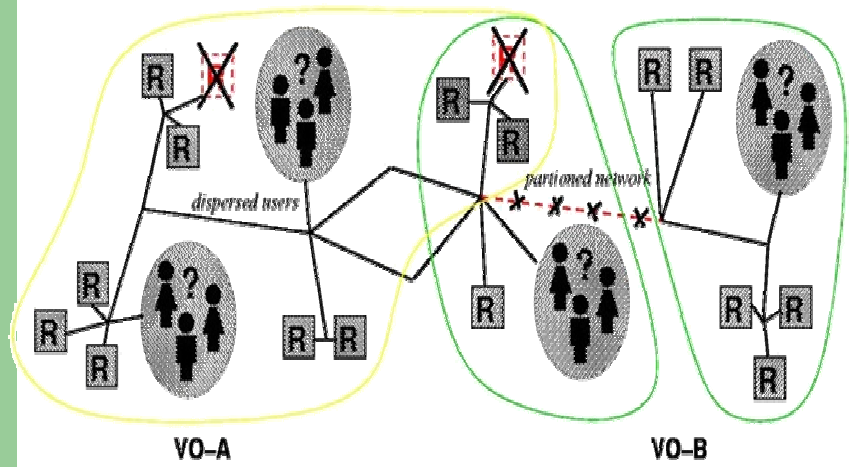

- Problem Elements:
  - Resource Sharing across massively disparate Virtual Organizations
  - No implicit trust
  - Coordinated problem solving (data, computation, networks)
- What constitutes a Grid then?
  - Middleware
  - No centralized control
  - The Web is NOT a Grid (not yet!)
  - Grid is NOT a distributed OS
  - "On-the-Internet" = Speak IP
  - "On-the-Grid" = Speak Intergrid Protocols

Meta-Message: Provide dependable, consistent, pervasive access to high-end computational capabilities – Carl & Ian

The Role of Grid Middleware



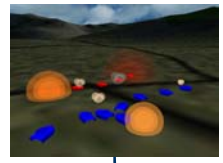

Collaboration Tools

Data Mgmt Tools

. . .

Distributed simulation

Information services

Resource mgmt

Fault detection

. . .

Remote access

Remote monitor


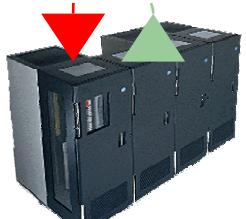
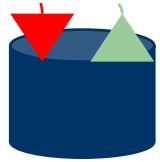

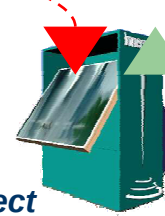

Image Courtesy: Globus Project

Problem Scope

- Problem:
  - Resource sharing across domains

2000 Physicists, 150 Institutes, 32 Countries

Globus Grid Architecture

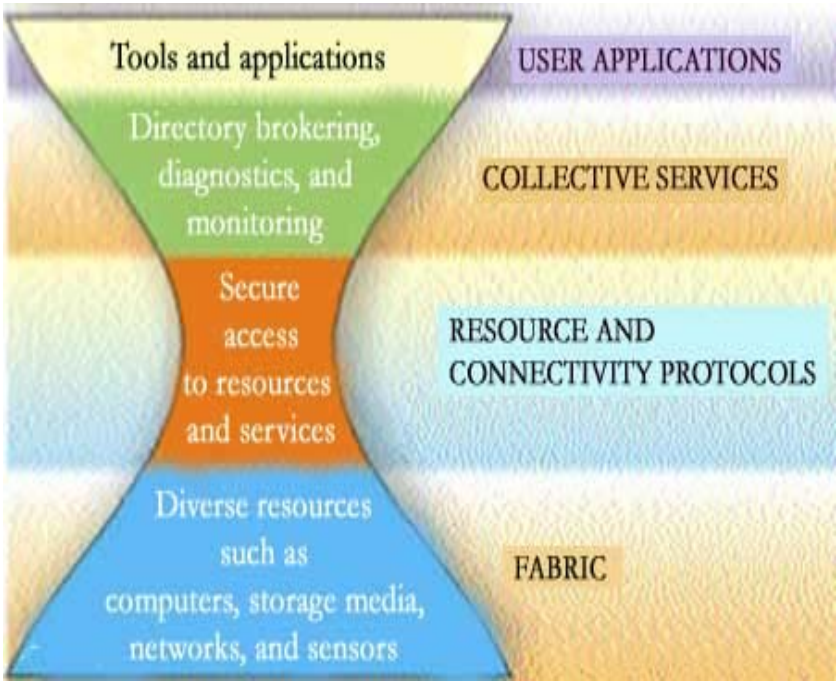

Constructing Layered Services to handle Data Management

How to manage/select/co-allocate replicas?

How to access/retrieve data? GridFTP, SRB, etc.

How to authenticate to these? Kerberos, GSI, etc.

How to create/delete files? DPSS, HPSS, UNIX FS, etc.

Talk Resource & Connectivity Protocols and you are on the Grid.

Meta-Message: Bottom-Up approach Works!!

Image Courtesy: Globus Project

Resource Management Architecture

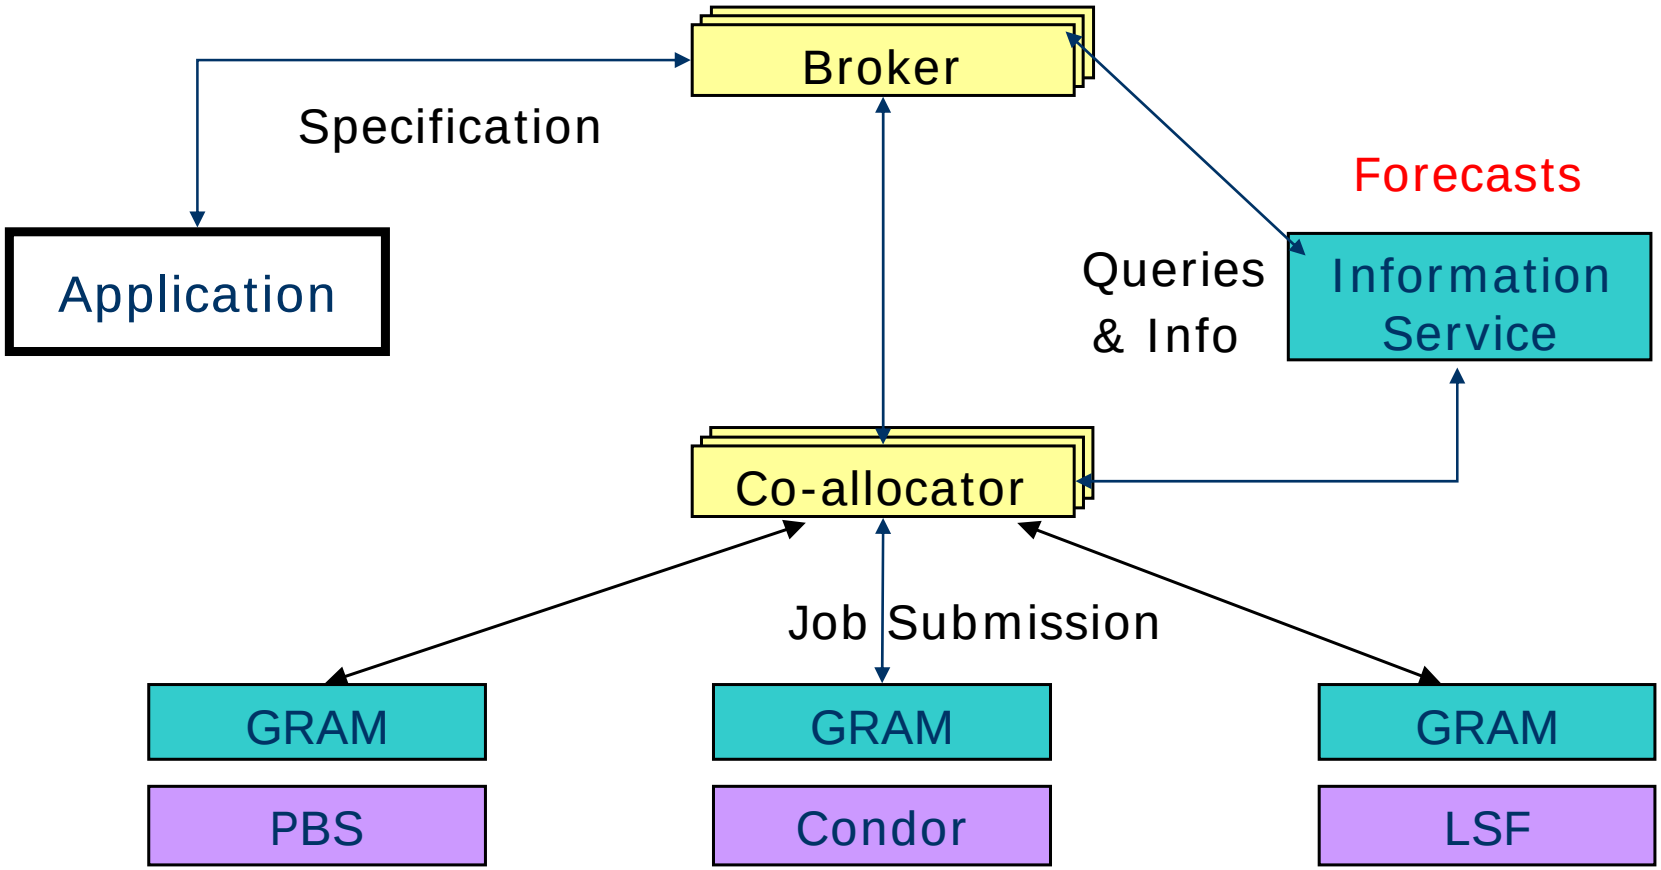

Image Courtesy: Globus Project

Grid Information Service

- Mechanisms to expose:
  - Identify, collect and publish server load details
  - Discovery in the context of an information service

- Publish/Query Infrastructure

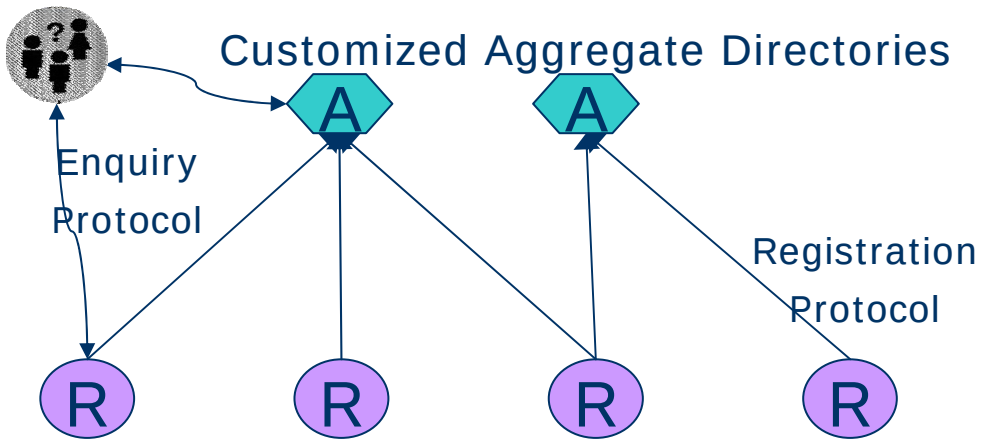

Standard Resource Description Services

Sample Storage Information Provider

dn:"140.221.65.69,
hostname=dpsslx04.lbl.gov,dc=lbl,dc=gov,o
=grid"
cn:"140.221.65.69"
gridftpurl:"gsiftp://dpsslx04.lbl.gov:61000"
minrdbandwidth:1462K
maxrdbandwidth:12800K
avgrdbandwidth:6062K
……
Policy: only these sites & at this time

Image Courtesy: Globus Project

Data Location and Scheduling Architecture

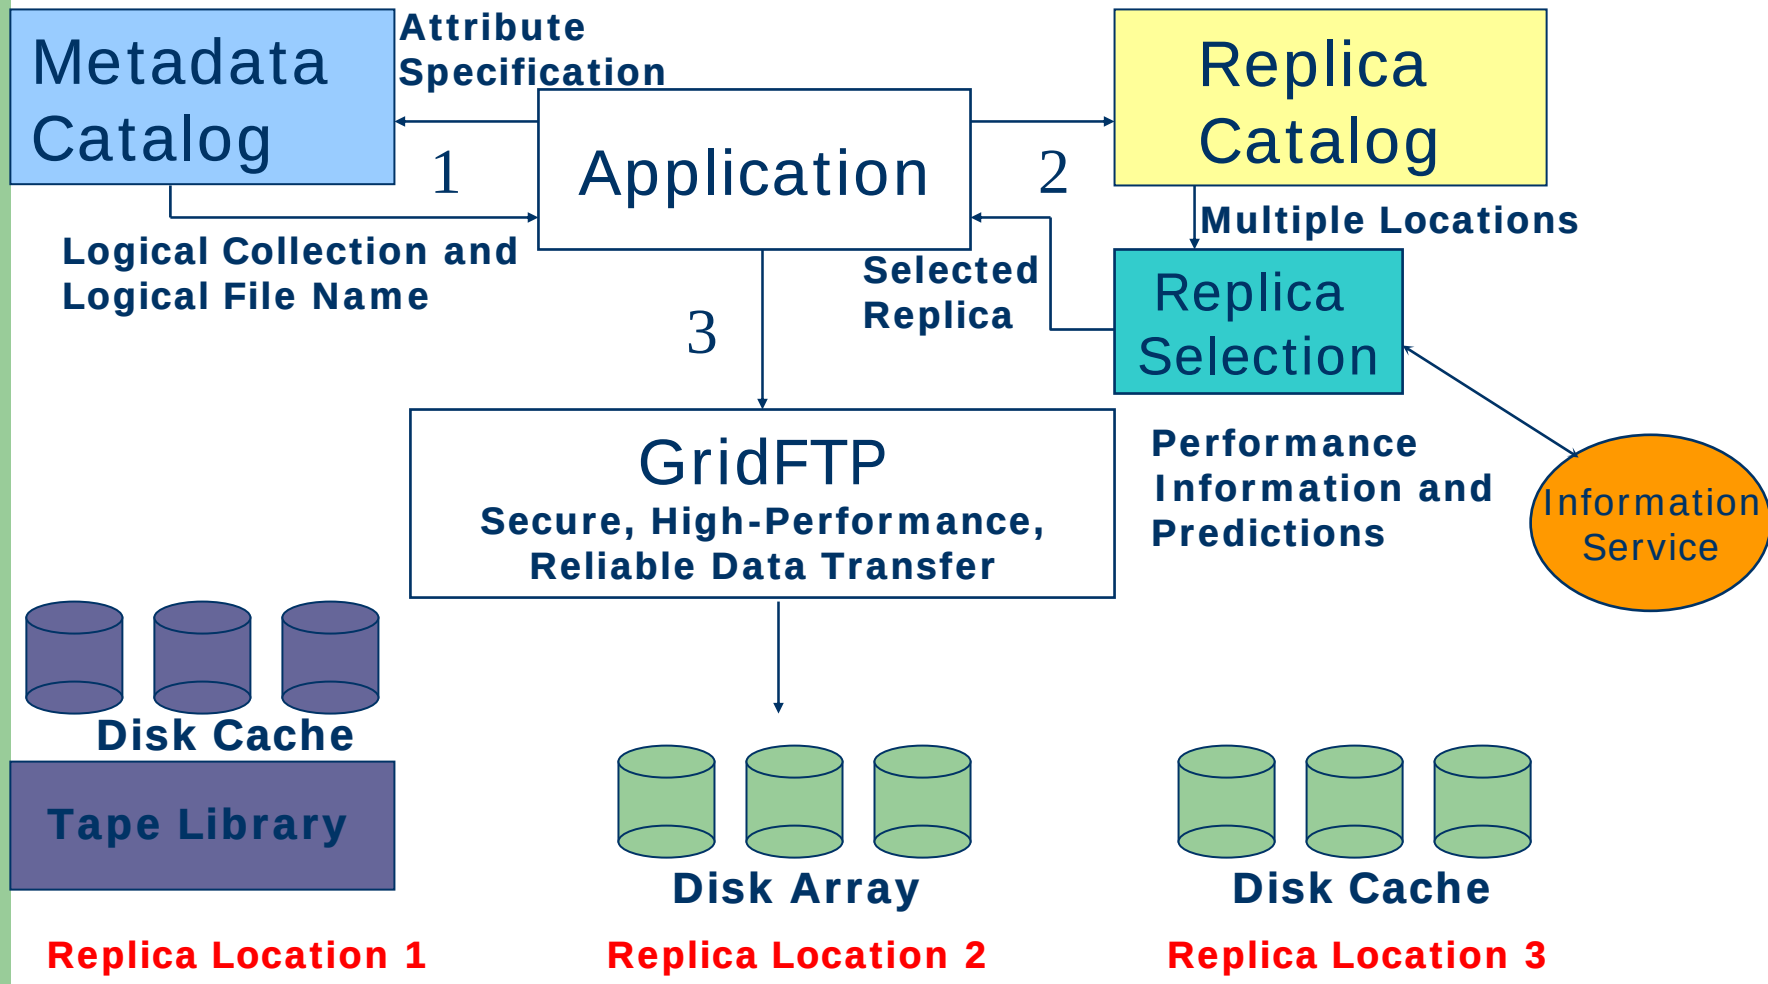

Image Courtesy: Globus Project

Tying it all together…

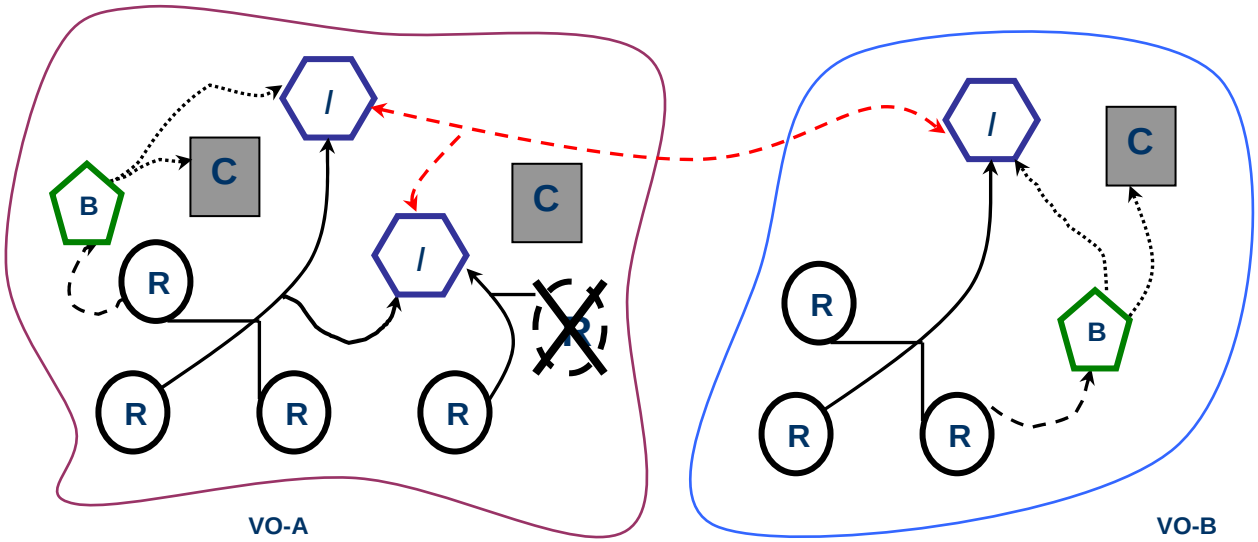

Further Information…

- Grid Forum: www.ggf.org
- My Website: www.csm.ornl.gov/~vazhkuda
- Email: vazhkudaiss@ornl.gov

Research sponsored by the Laboratory Directed Research and Development Program of Oak Ridge National Laboratory (ORNL), managed by UT-Battelle, LLC for the U. S. Department of Energy under Contract No. DE-AC05-00OR22725.