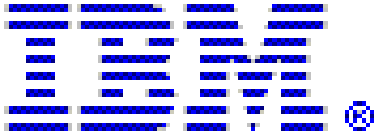

Grid Computing and Scaling Up the Internet

Brian E Carpenter

IBM and 6NET

Ian Foster

Argonne National Laboratory, University of Chicago, and Globus

May 2003

brian@hursley.ibm.com
foster@mcs.anl.gov

Topics

- General Introduction
- Brief introduction to Grid computing
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

General Introduction

Why discuss Grids and IPv6 together?

- Grid computing represents a fundamental shift in the approach to distributed computing, like the fundamental shift in information access introduced by the Web ten years ago.
- IPv6 represents a major step function in the Internet's ability to scale, like the introduction of IPv4 twenty years ago.
- Inevitably there is synergy between these two game changers.
- Let's share a common goal of reaching 10 billion Internet nodes.

Topics

- General Introduction
- Brief introduction to Grid computing
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

The Grid Is …

a) A collaboration & resource sharing infrastructure for scientific applications
b) A distributed service integration and management technology
c) A disruptive technology that enables a virtualized, collaborative, distributed world
d) An open source technology & community
e) A marketing slogan
f) All of the above

Grid Past, Present, Future

- Past
  - Origins and broad adoption in eScience, fueled by open source Globus Toolkit
- Present
  - Rapidly growing commercial adoption
  - Open Grid Services Architecture (OGSA)
- Future
  - Key enabler of new applications & industries based on resource virtualization and distributed service integration

Not Exactly a New Idea …

- “The time-sharing computer system can unite a group of investigators …. one can conceive of such a facility as an … intellectual public utility.”
  - Fernando Corbato and Robert Fano , 1966
- “We will perhaps see the spread of ‘computer utilities’, which, like present electric and telephone utilities, will service individual homes and offices across the country.”
  - Len Kleinrock, 1967

But, Things are Different Now

- Networks are far faster (and cheaper)
  - Faster than computer backplanes
- The Internet has already radically changed the practice of “Computing”
  - Our “computers” have already disintegrated
  - E-commerce increases size of demand peaks
  - Entirely new applications & social structures
- We’ve learned a few things about software

But Wait A Minute—Computing isn't Really Like Electricity!

- I import electricity but must export data
- "Computing" is not interchangeable but highly heterogeneous
  - Computers, data, sensors, services, …
- Ok, so the story is more complicated
- But more significantly, the sum can be greater than the parts
  - Real opportunity: Construct new capabilities dynamically from distributed services

  ⇒ Virtualization & distributed service mgmt

Virtualization & Distributed Service Management

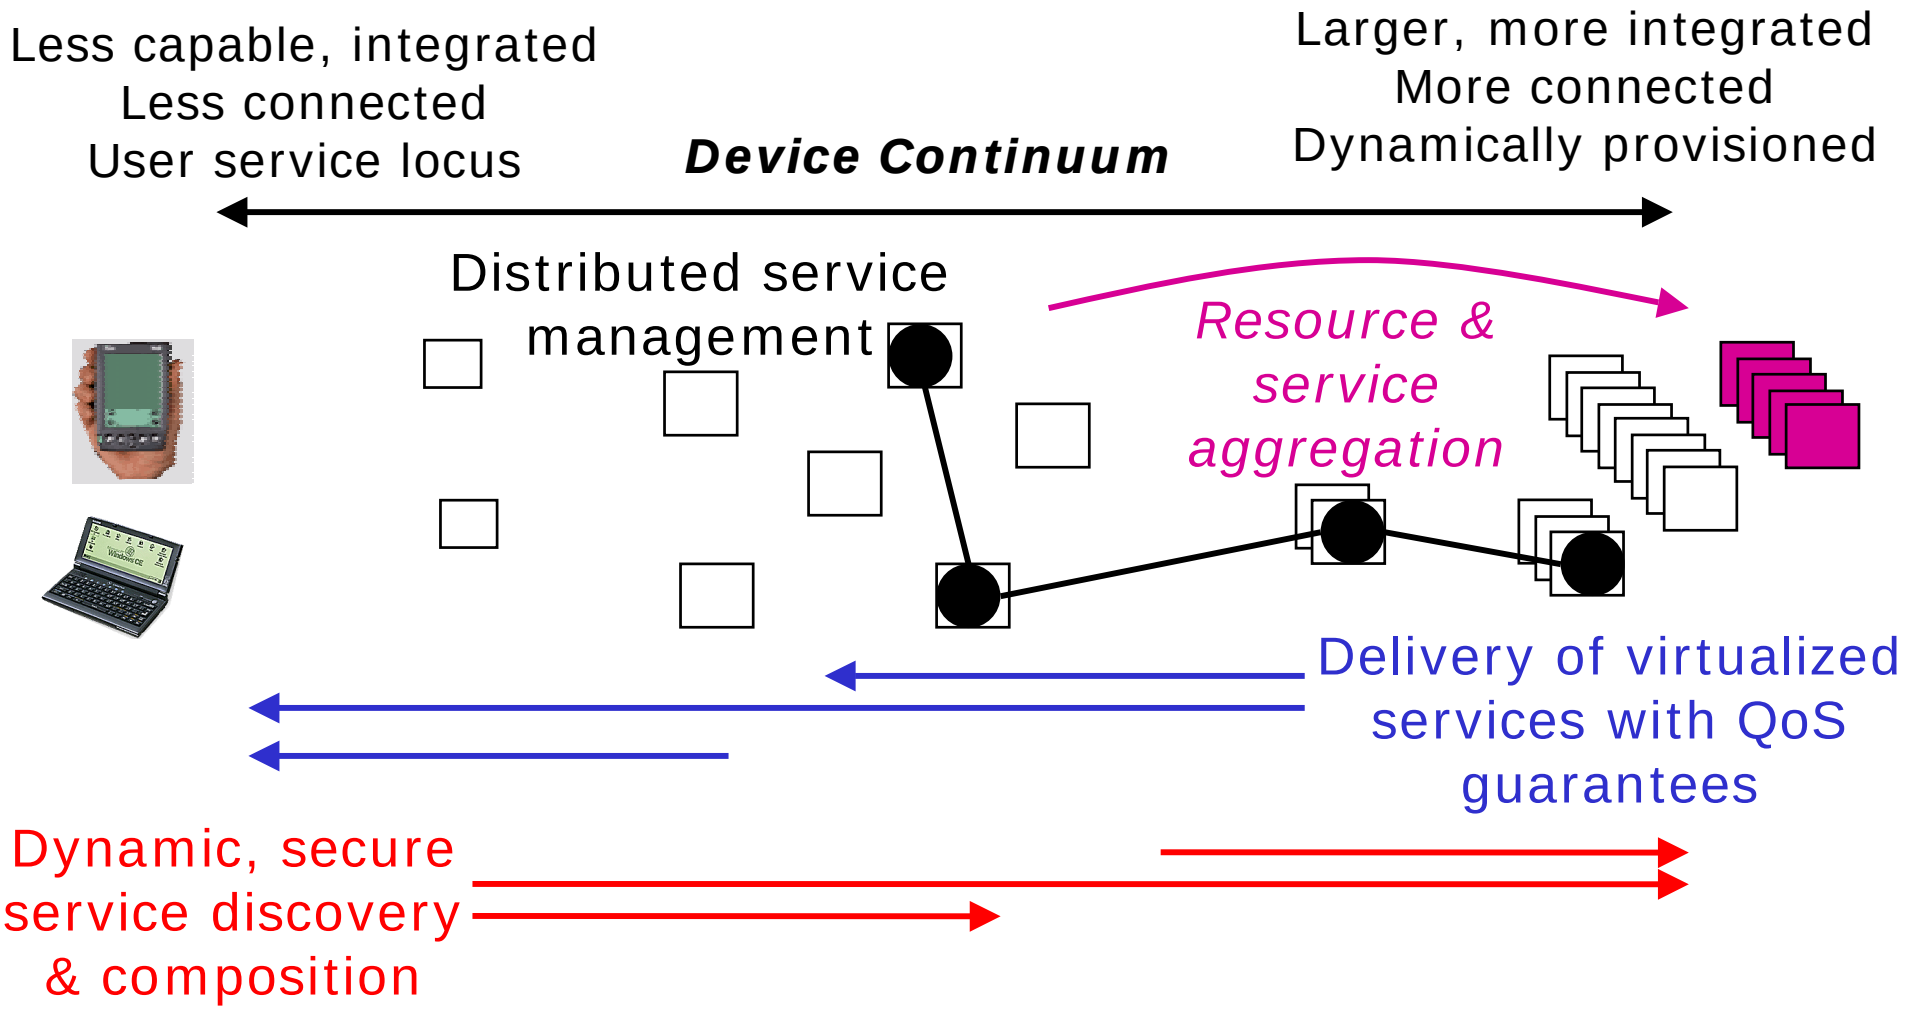

Why the Grid?
Origins: Revolution in Science

- Pre-Internet
  - Theorize &/or experiment, alone or in small teams; publish paper
- Post-Internet
  - Construct and mine large databases of observational or simulation data
  - Develop simulations & analyses
  - Access specialized devices remotely
  - Exchange information within distributed multidisciplinary teams

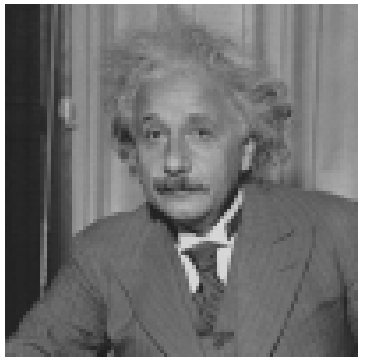

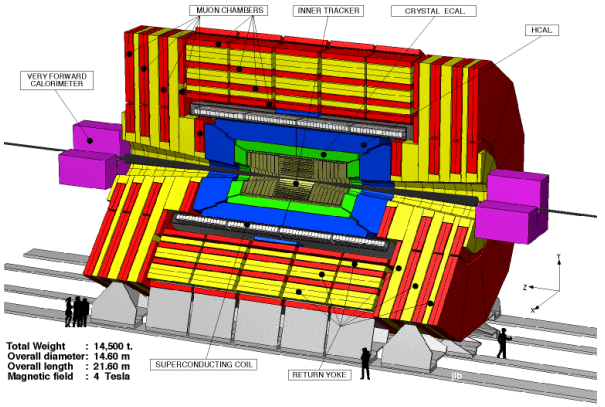

Example Science Grids

Building the National Virtual Collaboratory for Earthquake Engineering Research

NEESgrid

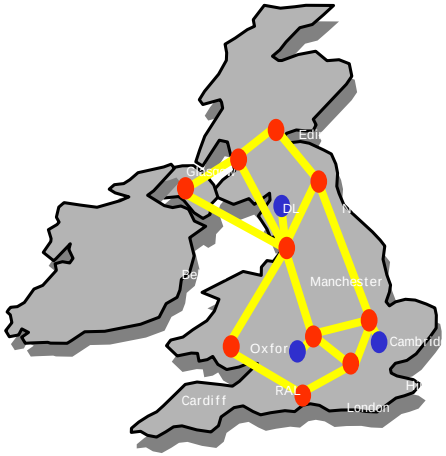

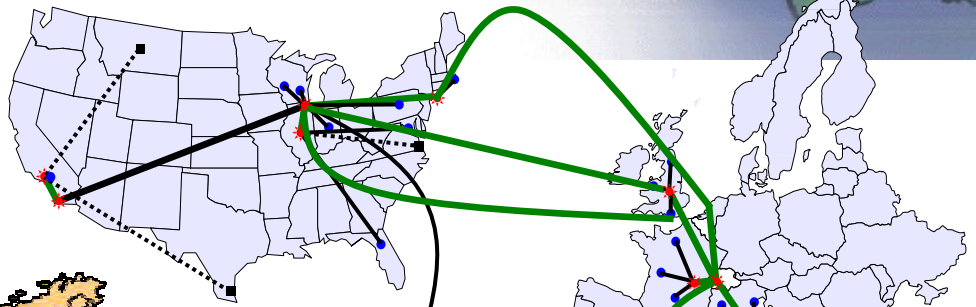

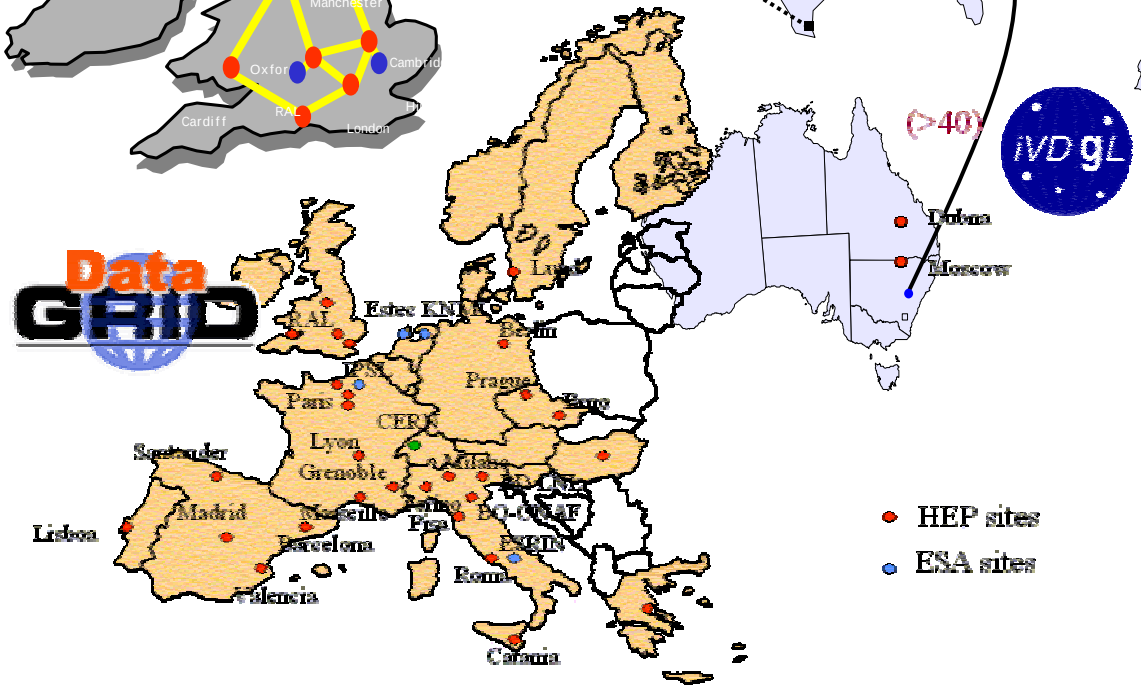

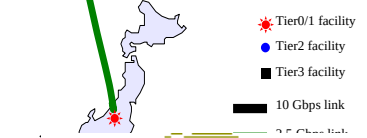

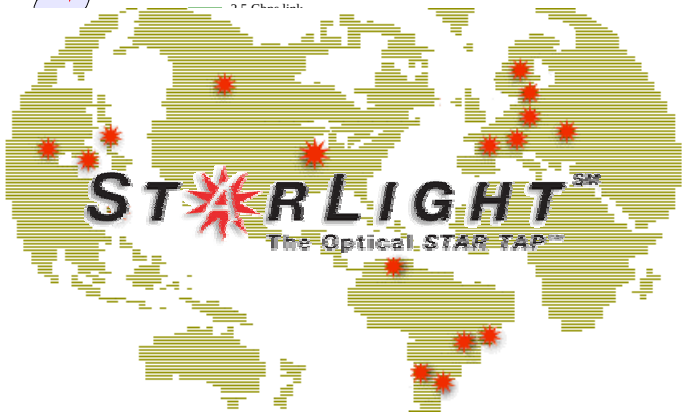

The Grid/eScience World: Status

- Dozens of major Grid projects in scientific & technical computing/research & education
  - Deployment, application, technology
  - www.mcs.anl.gov/~foster/grid-projects
- Globus Toolkit™ broadly adopted as de facto standard for major protocols & services
- Global Grid Forum a significant force for community building and standardization
  - GGF8: Tokyo, March 2003, 850 people
  - www.gridforum.org; 200+ organizations; Boeing, Merck, Ford, J&J, IBM, Platform, …

Why the Grid?
(2) Revolution in Business

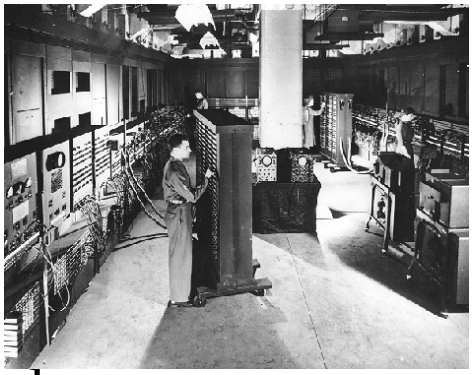

- Pre-Internet
  - Central data processing facility
- Post-Internet
  - Enterprise computing is highly distributed, heterogeneous, loosely coupled, inter-enterprise (B2B)
  - Business processes increasingly computing- & data-rich
  - Outsourcing becomes feasible => on-demand service providers of various sorts

Common eScience/eBusiness Vision

- Link dynamically acquired resources
  - From collaborators, customers, eUtilities, … (members of evolving “virtual organization”)
- Into a “virtual computing system”
  - Dynamic, multi-faceted system spanning institutions and industries
  - Loose coupling of heterogeneous systems
  - Configured ***on demand*** to meet instantaneous needs, for:
- Multi-faceted QoS for demanding workloads
  - Security, performance, reliability, …

Grids and Open Standards

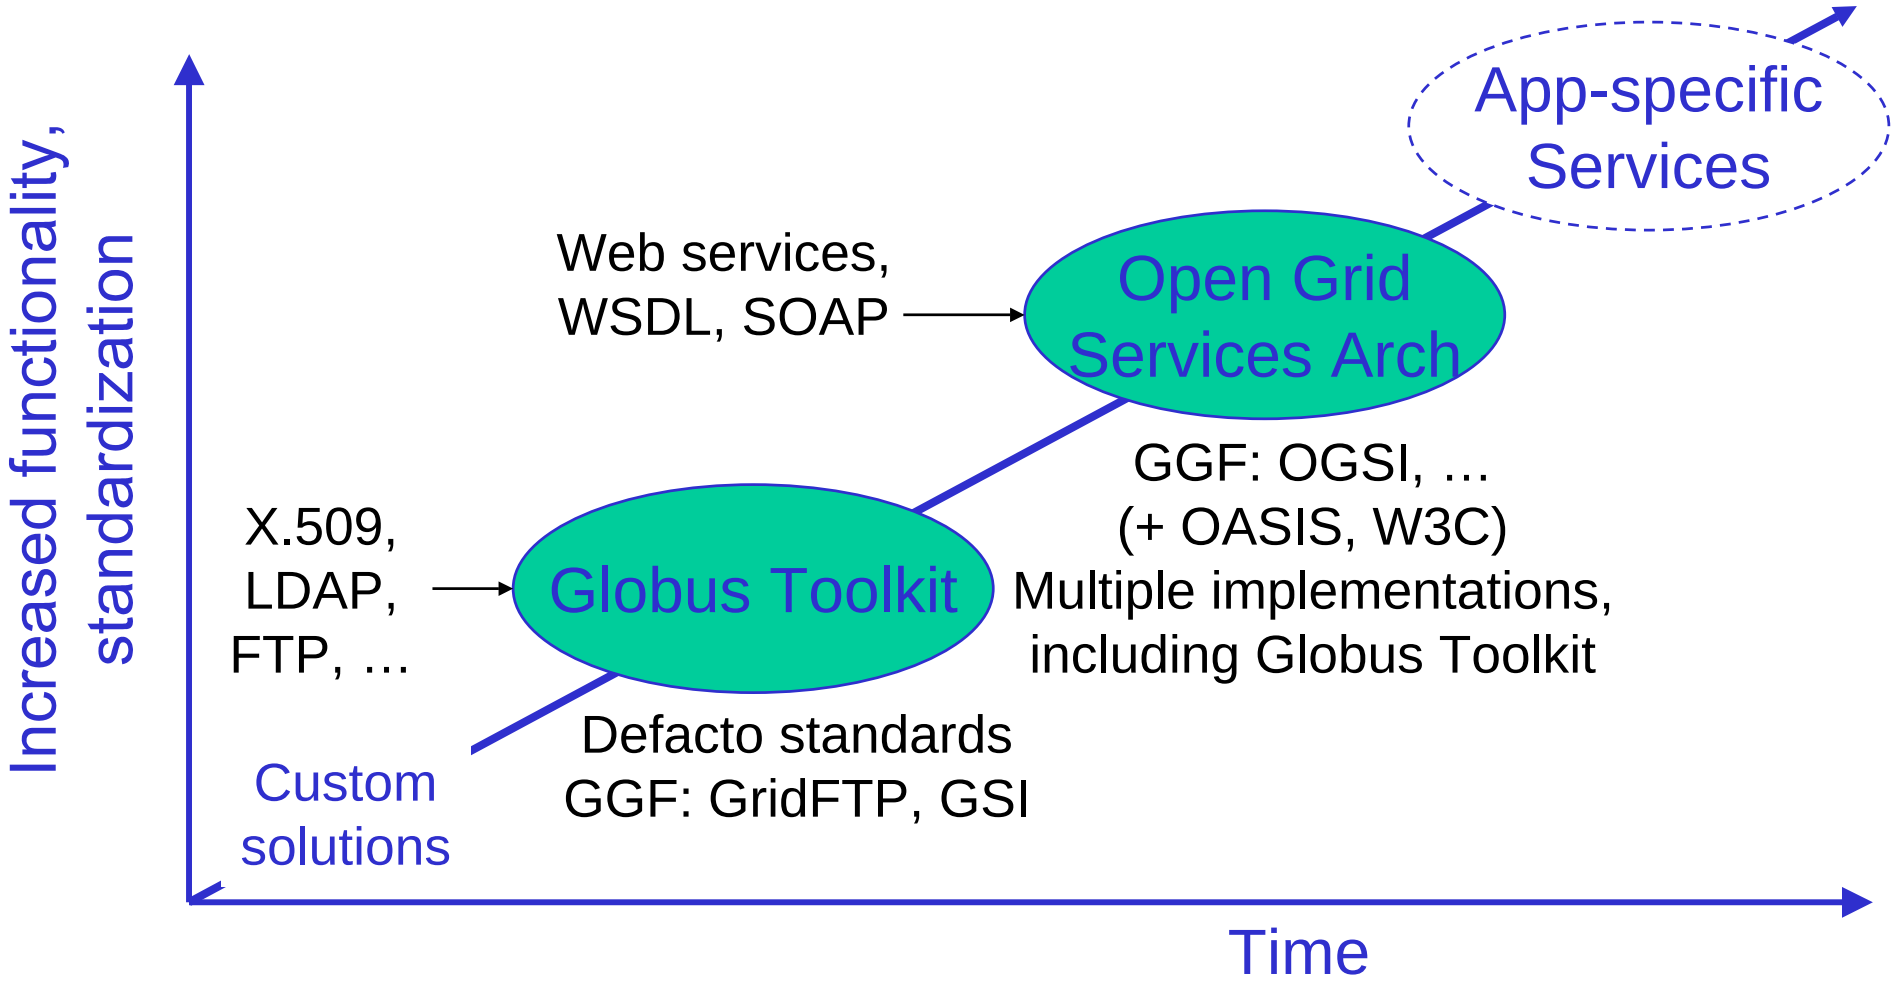

Open Grid Services Architecture

- Service-oriented architecture
  - Key to virtualization, discovery, composition, local-remote transparency
- Leverage industry standards
  - Internet, Web services
- Distributed service management
  - A “component model for Web services”
- A framework for the definition of composable, interoperable services

“The Physiology of the Grid: An Open Grid Services Architecture for Distributed Systems Integration”, Foster, Kesselman, Nick, Tuecke, 2002

Open Grid Services Infrastructure

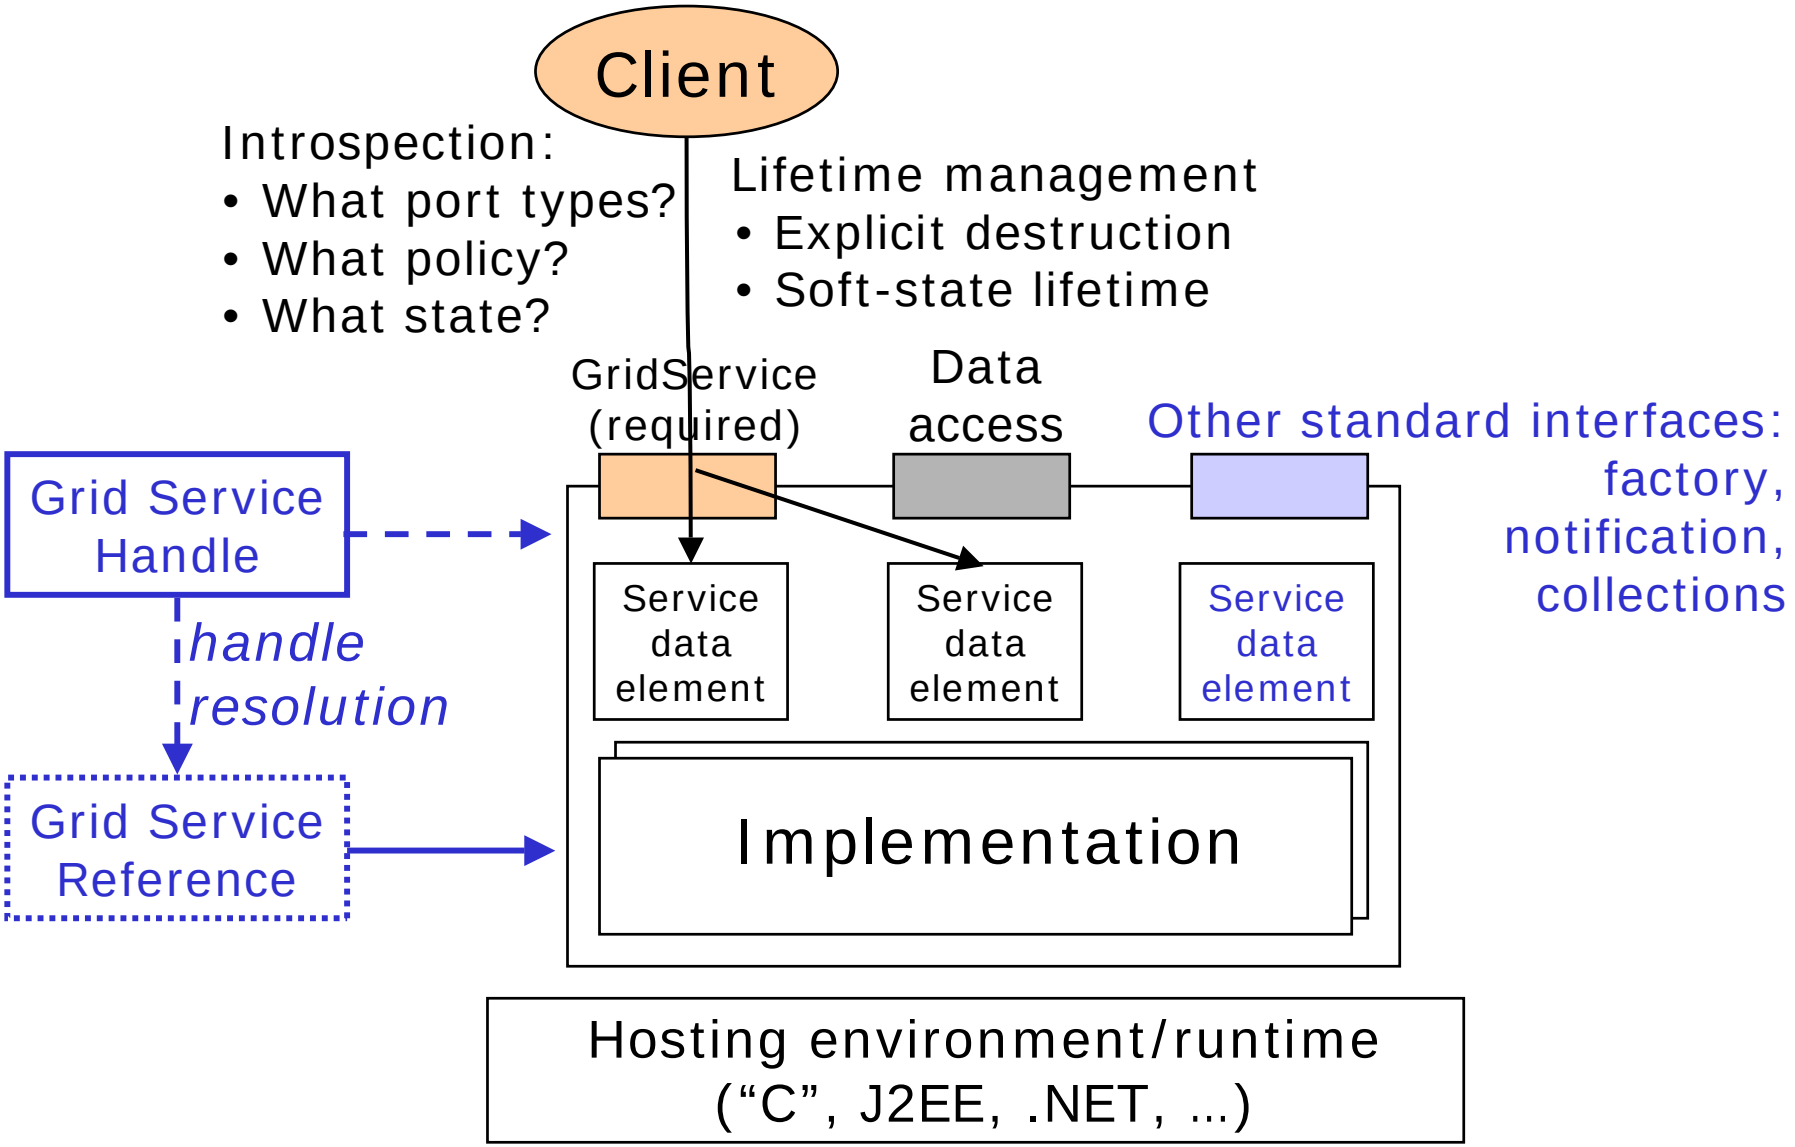

OGSA Interactions

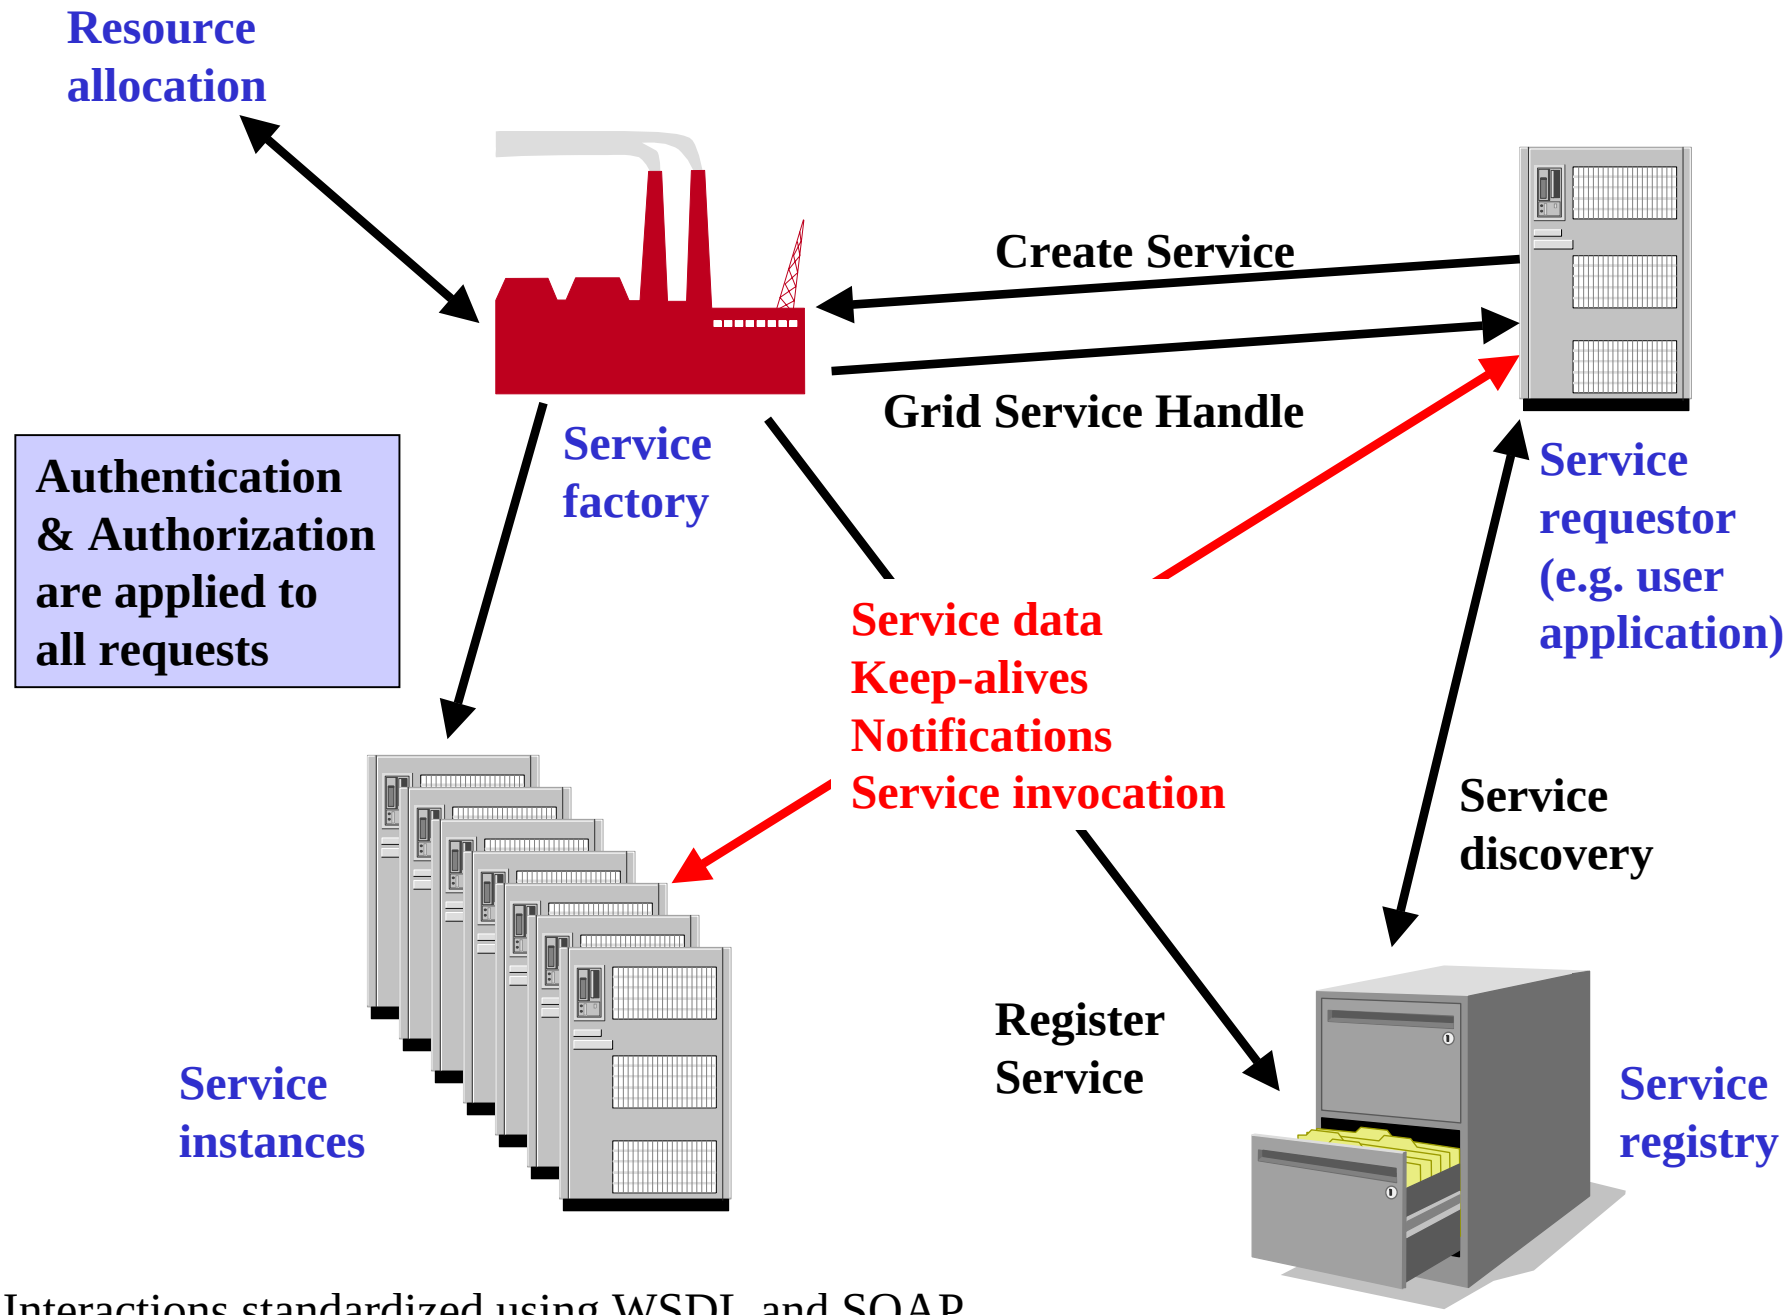

Interactions standardized using WSDL and SOAP

Globus Toolkit v3 (GT3)
Open Source OGSA Technology

- Implements OGSI interfaces
- Supports primary GT2 interfaces
  - High degree of backward compatibility
- Multiple platforms & hosting environments
  - J2EE, Java, C, .NET, Python
- New services
  - SLA negotiation, service registry, community authorization, data management, …
- Rapidly growing adoption and contributions: “Linux for the Grid”

Topics

- General Introduction
- Brief introduction to Grid computing
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

Grid Deployment Options

A function of business need, technology and organizational flexibility

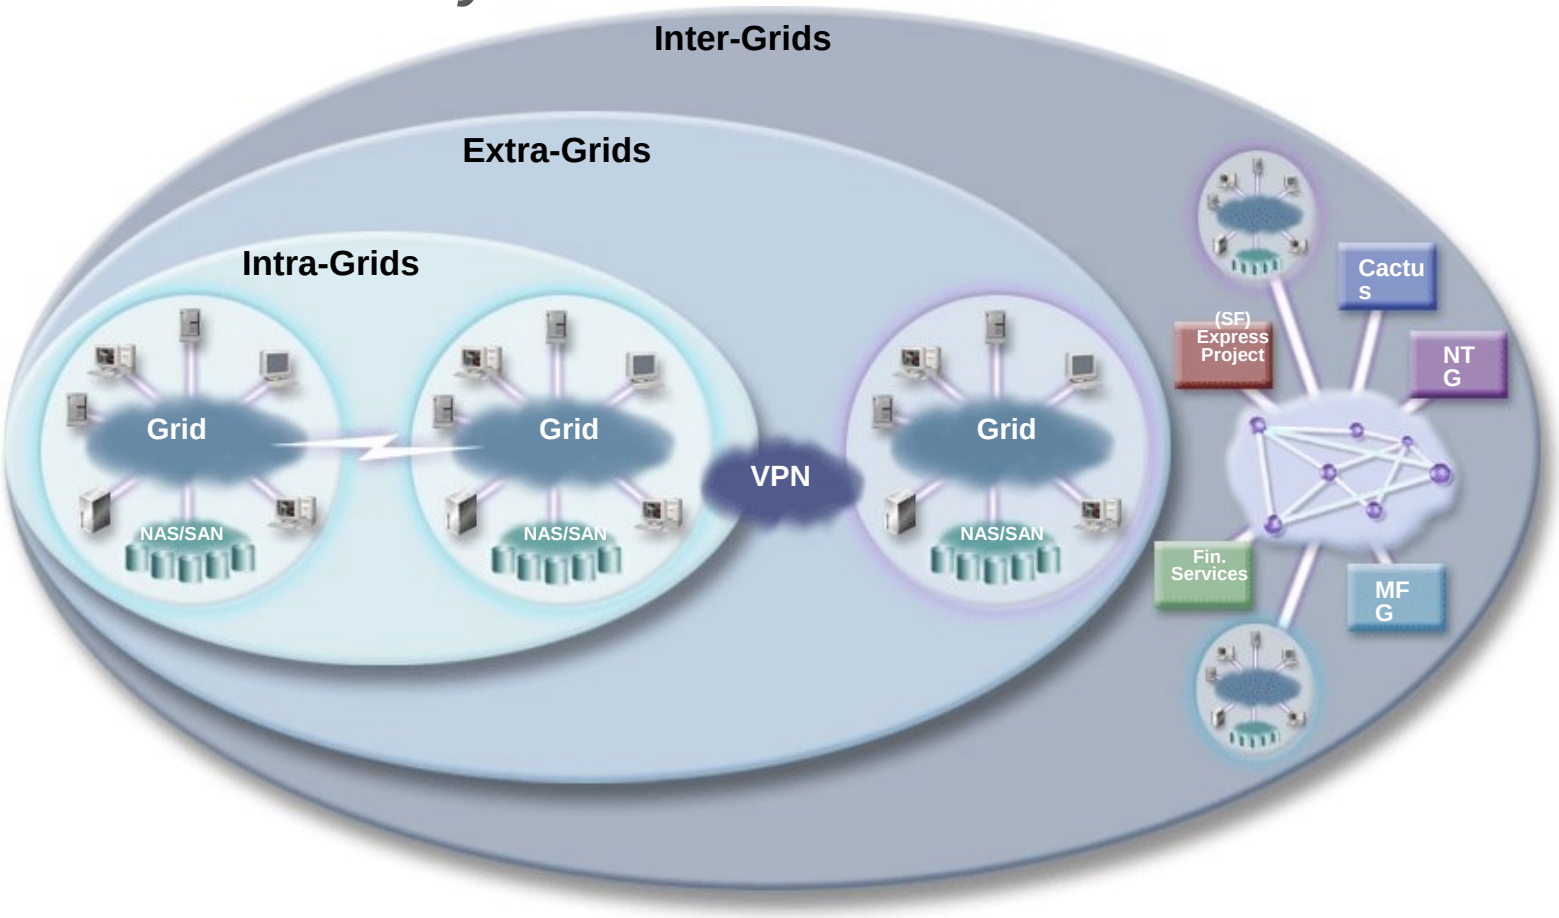

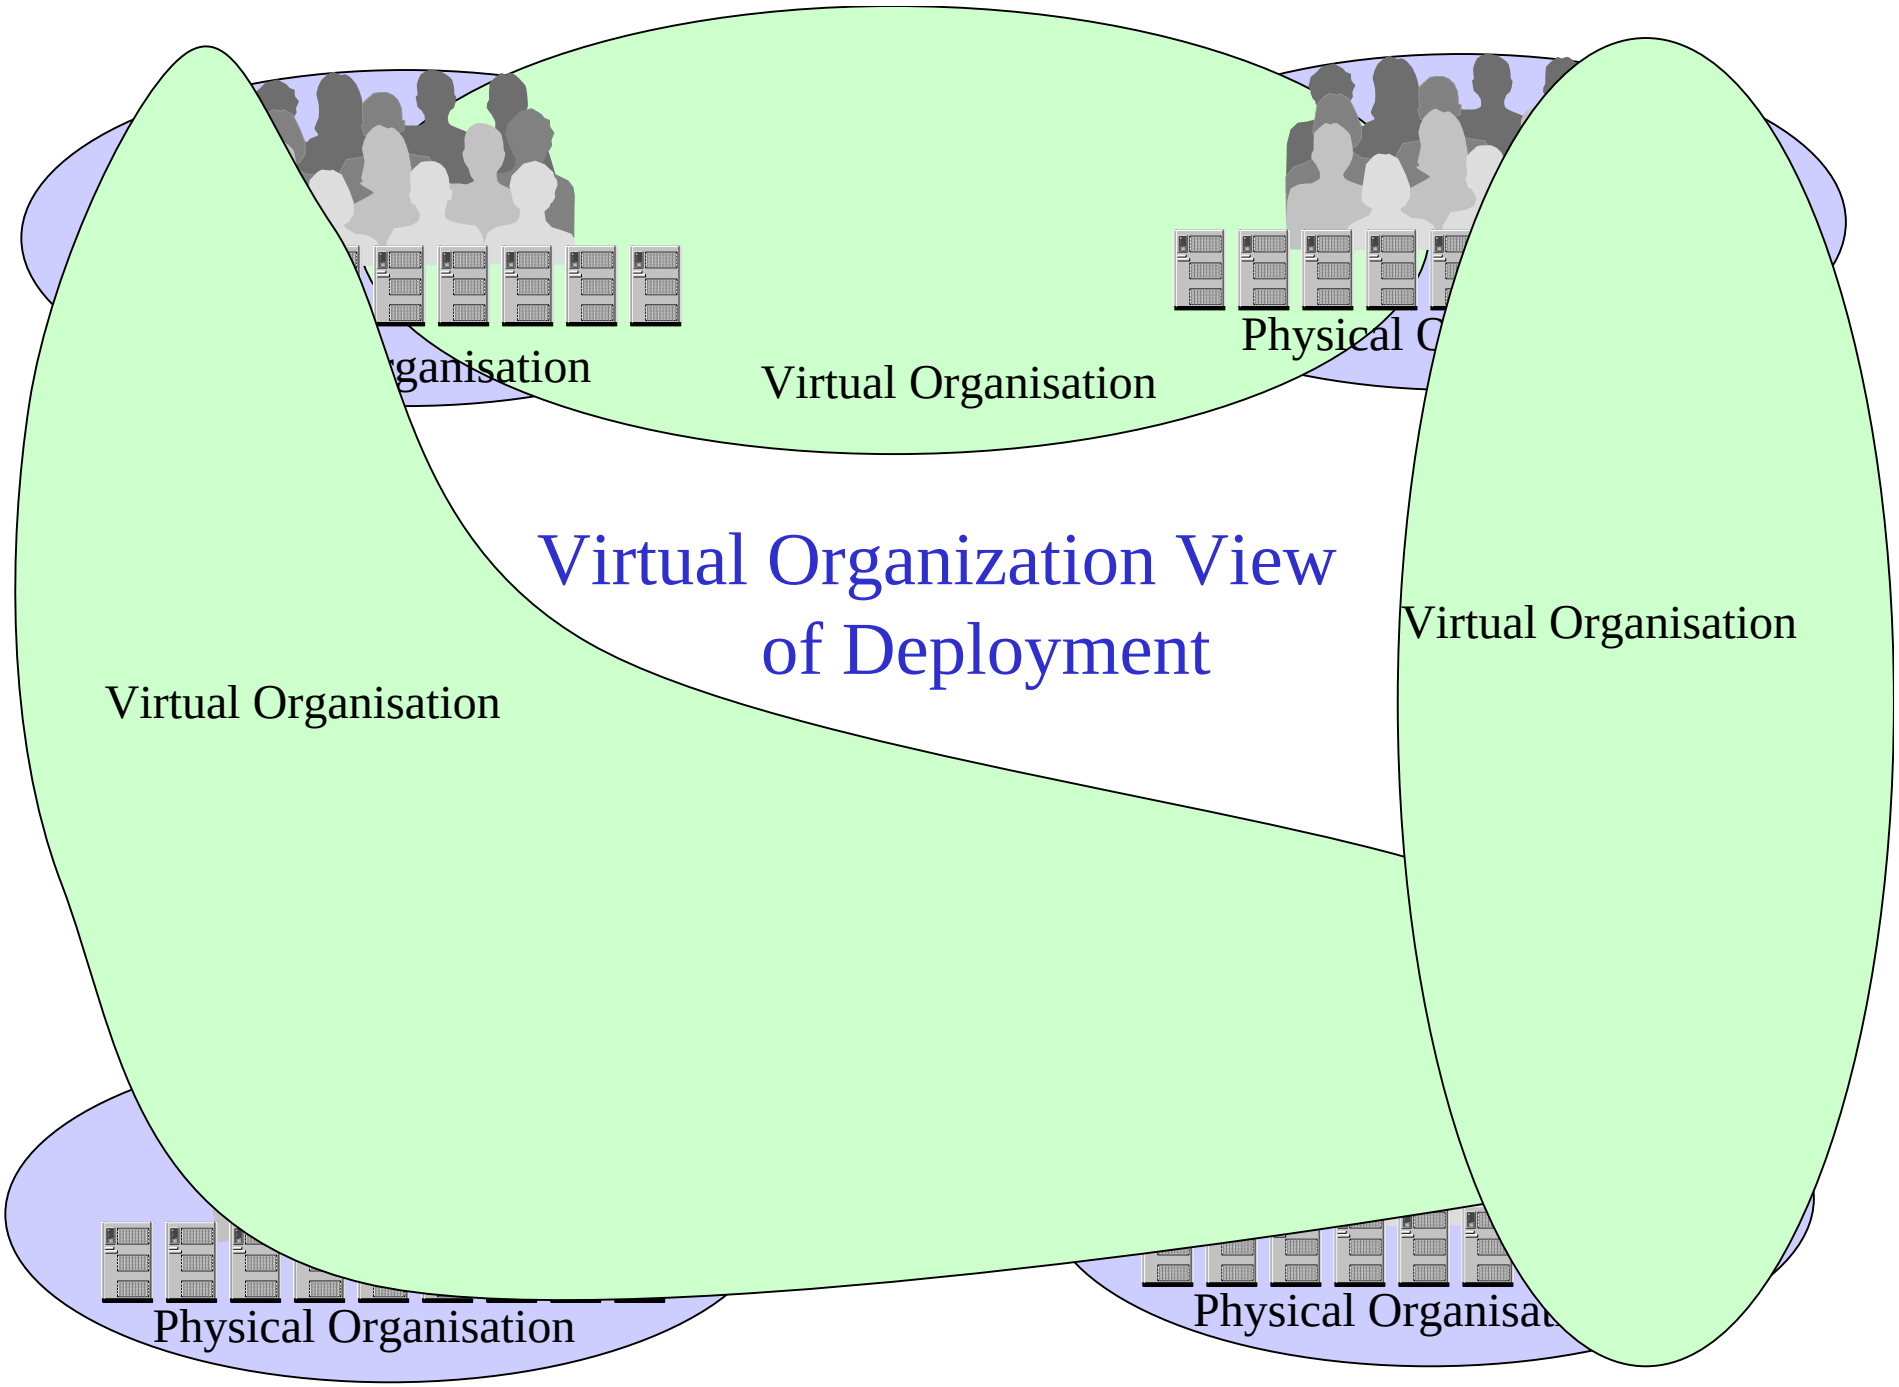
ganisation
Virtual Organisation
Physical O
Virtual Organization View
of Deployment
Virtual Organisation
Virtual Organisation
Physical Organisation
Physical Organisat

Virtual organizations look like dynamic mergers & acquisitions

- The effect of a Grid VO on networks is like a temporary partial merger of the organizations.
- Merging two networks is very painful today:
  - "private" IPv4 address space becomes ambiguous
  - worst case: forced to renumber both networks
- Temporary partial mergers of an arbitrary number of IPv4 networks is unthinkable.

➢ IPv4 based Grids are forced to rely on HTTP proxying between organizations: inefficient, and cannot exploit network level security.

Overlapping virtual organizations

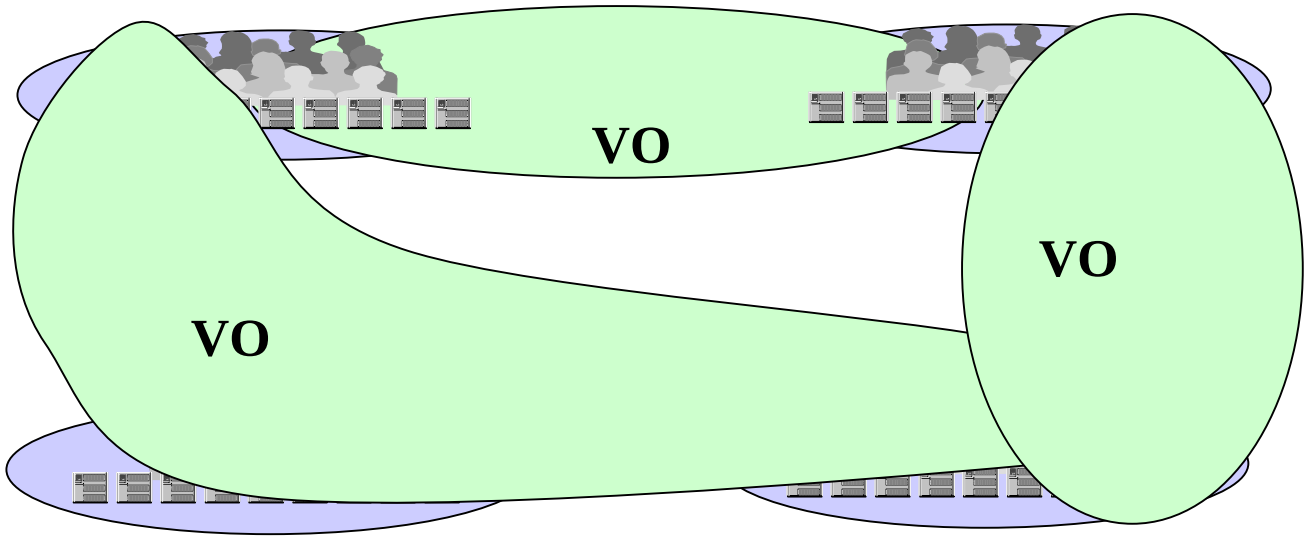

- Any system can be in any number of VOs with any number of other systems
  - needs uniform address space to avoid proxies & allow IPSEC
  - addressing ambiguities unacceptable
  - security boundaries ♑ organization boundaries
  - **can't meet these constraints at massive scale with IPv4**

Critical advantages of IPv6 for OGSA

- Potential for massive scaling
- Uniform global address space eliminates the problem of ambiguous “private” addresses and network address translation.
  - Wasteful proxies can be avoided
  - Network level security can be used
- Autoconfiguration is a big plus for infrastructure configuration

There's no such thing as an IPv6 killer app, but...

- It would be nice to find the killer app that only works on IPv6.
- OGSA won't be that, but there is a good chance that it will be the first major middleware suite to be IPv6-capable out of the box almost from Day One.
- The IPv6 community should make the most of it.

Topics

- General Introduction
- Brief introduction to Grid computing
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

Testing, testing

- 6NET is a three-year European Union funded project to demonstrate that continued growth of the Internet can be met using IPv6.
- It includes a work package for ***IPv6 Middleware and User Application Trials*** (led by IBM).
- Globus will be the subject of a trial (lead site: UCL)
  - Target is Globus Toolkit 3, i.e. OGSA
  - GT3 (OGSA) alpha code is now available and being tested on IPv6/Linux at UCL
  - *Credits: Sheng Jiang, Piers O'Hanlon, Peter Kirstein*

Status (evolving daily)

- Background: GT3 runs mainly over Java; Java2 Development Kit 1.4 supports IPv6, although GT3 alpha ships with JDK 1.3.
- In principle, all components of GT3 implemented in Java will "just work" with IPv6 by switching to JDK 1.4
  - In practice, it is not so simple.
  - Jim Bound of HP is also looking into this
- UCL has a GT3 testbed and has begun testing with JDK 1.4.
  - The GT3 Master Host Environment listens on both IPv6 and IPv4. An elementary test initiated from IPv6 completes OK, but some IPv4 packets are observed.
- Exact environment is GT3 alpha code with Java SDK 1.4.1. on Redhat Linux 7.3 and 8.0.

Status (update May 6th)

- Tested the postgresql IPv6 patch; found one bug there: after removing IPv6 items from the configuration file, the IPv6 address was still enabled.
- Working to deploy GT3 core on Apache Tomcat Java servlet container (IPv6 enabled).
- Waiting for IBM Websphere to do the same.
- Starting to port the OGSA stand-alone web container to be IPv6-enabled
  - *Sheng Jiang*

GridFTP (evolving daily)

- GridFTP (striped FTP) as shipped with GT3 alpha is C code that does not support IPv6 sockets.
- Globus is developing a new generic I/O module called XIO that does support IPv6 sockets.
- GridFTP is being rewritten, still in C, to exploit XIO.
- Any other components of GT3 that remain in C can also support IPv6 via XIO.

6net

Further plans (evolving daily)

- Plan is to make more extensive tests with successive GT3 alphas, with about 10 nodes
  - Issues with IPv6 will be reported into the Globus bug-tracking system
  - Good relations established between 6NET and Globus teams
- Also need to consider what is required to operate GT3 in the cases of:
  - IPv6 only
  - IPv6 and IPv4 coexistence
- Final goal is a realistic systematic trial between 6NET sites

Topics

- General Introduction
- Brief introduction to Grid computing
- Why Grids and IPv6 need each other
- Practical aspects of enabling GT3 for IPv6
- Future directions and summary

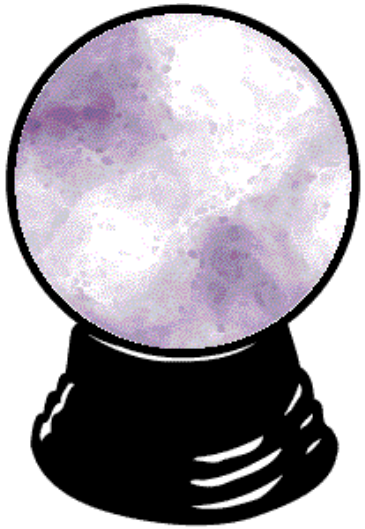

Future Directions

- GT3 will become fully functional for IPv6.
- Global Grid Forum needs to chase down any IPv4 dependencies in its standards.
- Grid computing will become a key enabler of new applications based on resource virtualization and loosely coupled distributed service integration.
- IPv6 will enable Grid Virtual Organizations to span existing network boundaries smoothly and securely.

Summary

- Grid computing is the new model for sharing networked IT resources efficiently & securely.
  - Transforming the Internet into a computing platform for e-business on demand
- The key toolkit is the Globus open source package based on the **Open Grid Services Architecture**.
- A marriage between OGSA and IPv6 is the key to massive scaling in a fully connected but secure network environment.

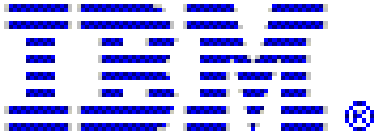

Pointers

www.globus.org

www.gridforum.org

www.ipv6forum.org

brian@hursley.ibm.com
foster@mcs.anl.gov