

Enabling Science: current and future trends in Grid Computing

Kate Keahey
Argonne National Laboratory
keahey@mcs.anl.gov

My Presentation

- Computational Grids 101
 - ◆ a perspective on forces driving the development of Grids
 - ◆ From yesterday to tomorrow
- Trends
 - ◆ Service orientation
 - Open standards and protocols
 - Emerging new directions: virtualization and QoS
 - ◆ Quality of service and policy-based execution
 - Agreement-based architecture
 - User execution environments
 - ◆ Grid community and its potential
- Conclusions

Computing has changed...

- It's easy to forget how different 2003 is from 1993
 - ◆ Ultra-high-speed networks: 10+ Gb/s
 - Global optical networks
 - ◆ Huge quantities of computing: 100+ Top/s
 - Moore's law
 - ◆ Increasingly better price/performance ratios
 - Networks, CPU, storage...
 - ◆ Ubiquitous Internet: 100+ million hosts
 - Collaboration & resource sharing the norm
 - ◆ Emerging Internet standards

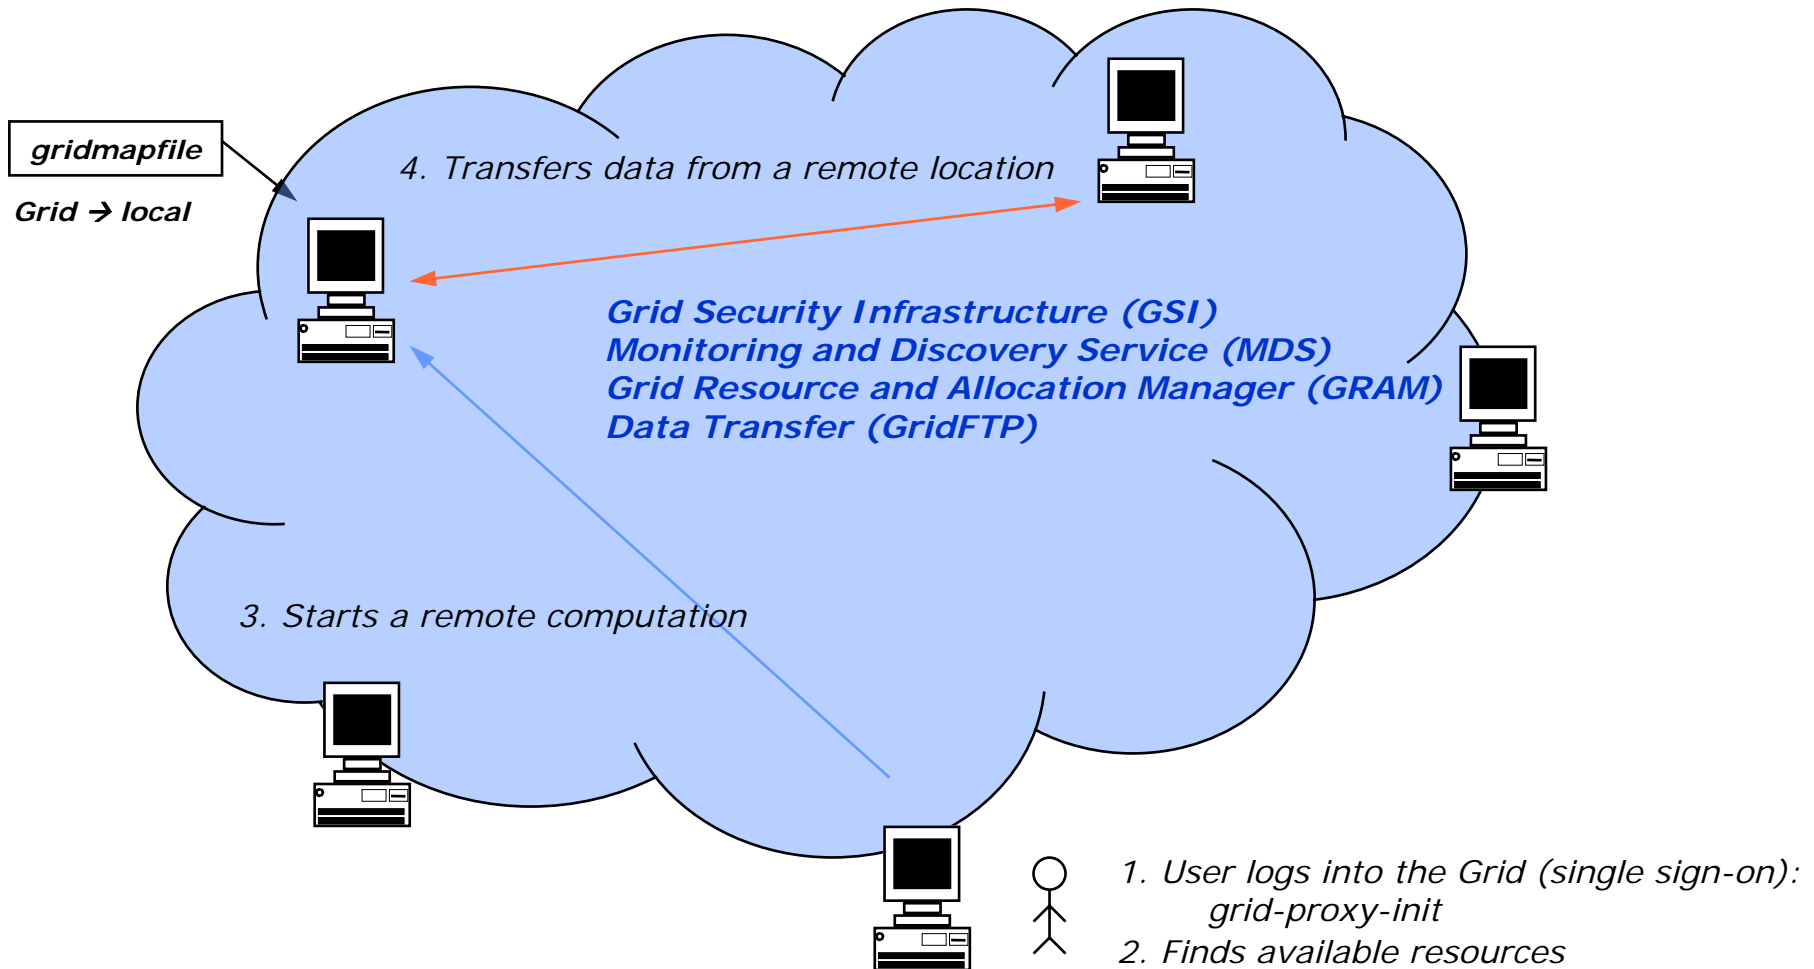
... and we have changed with it

- The Emergence of Grids
 - ◆ Teams organized around common goals
 - Communities: “Virtual organizations”
 - ◆ With diverse membership & capabilities
 - Heterogeneity is a strength not a weakness
 - ◆ And geographic and political distribution
 - No location/organization possesses all required skills and resources
 - ◆ Must adapt as a function of the situation
 - Adjust membership, reallocate responsibilities, renegotiate resources

Grid technologies

- Infrastructure (“middleware”) for establishing, managing, and evolving multi-organizational federations
 - ◆ Secure, coordinated sharing
 - ◆ Dynamic, autonomous, domain independent
 - ◆ On-demand, ubiquitous access to computing, data, and services
- Globus toolkit
 - ◆ An implementation of the most basic capabilities
 - ◆ A de facto implementation standard

Coordinated resource sharing

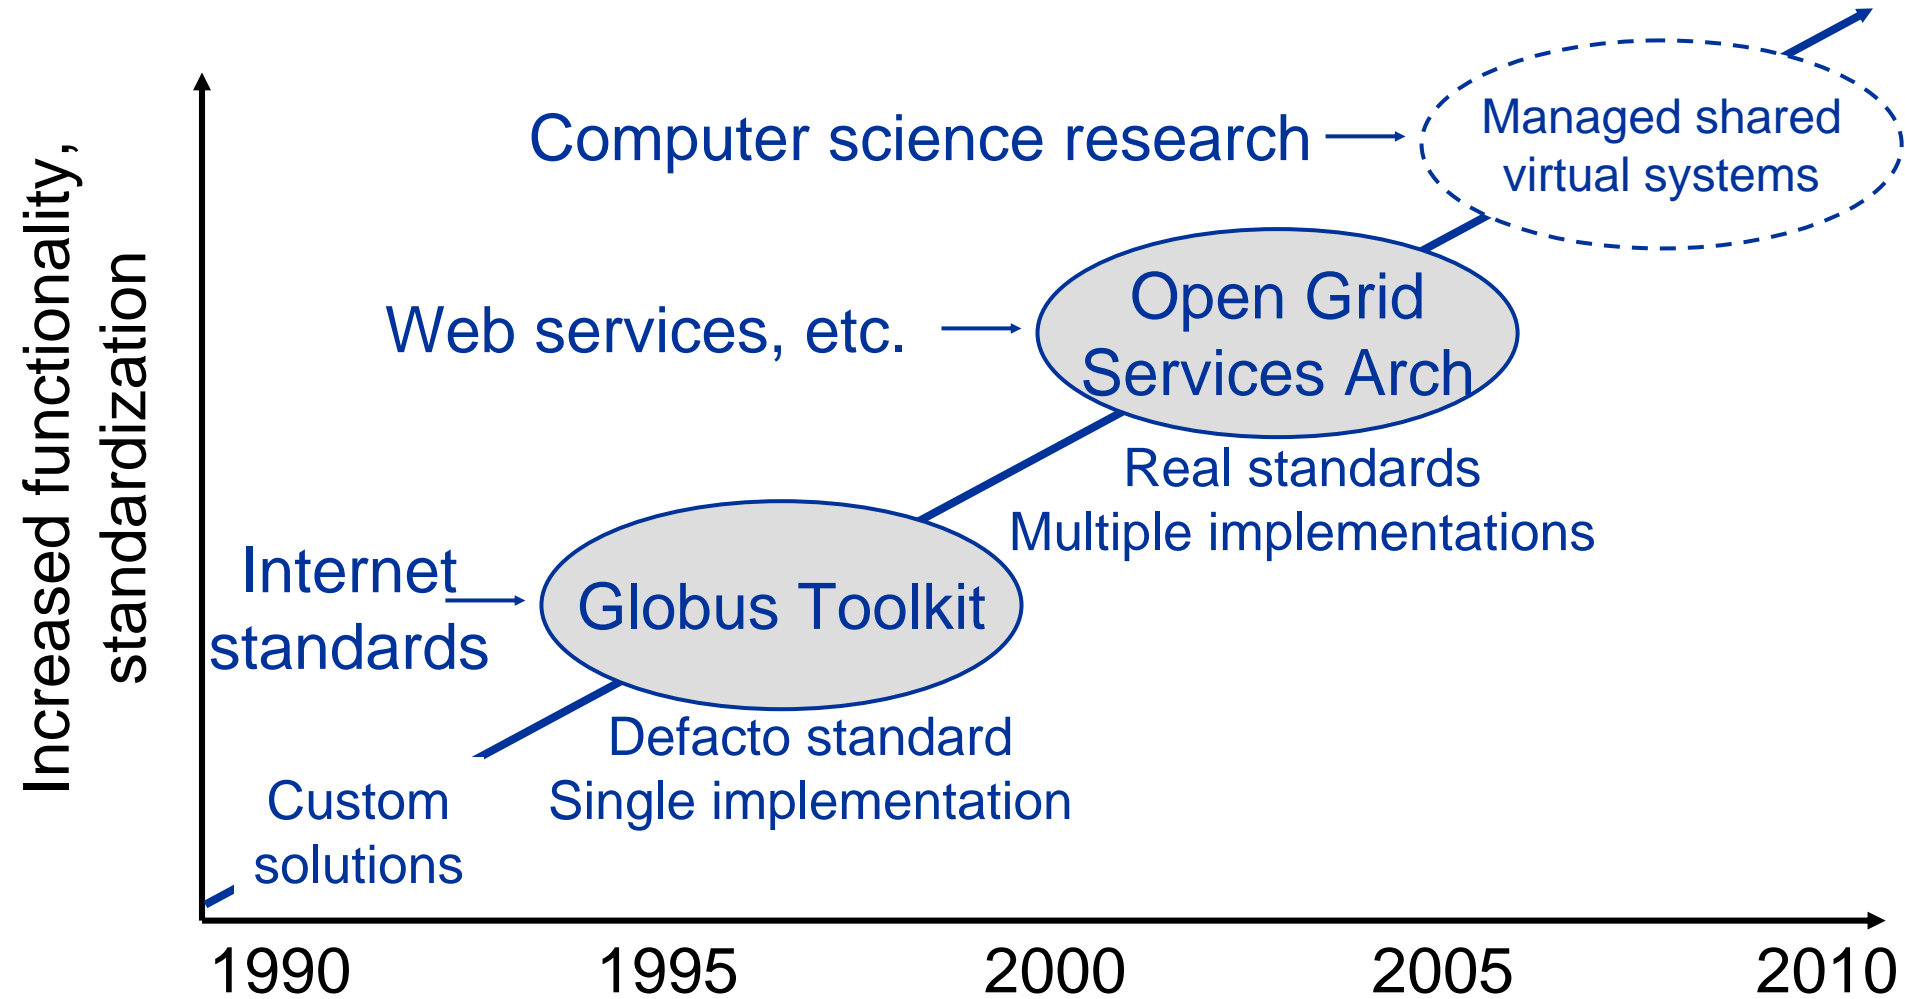


Globus Toolkit

- Pluses
 - ◆ Good technical solution for basic problems
 - Authentication and authorization (GSI)
 - Resource discovery and monitoring (MDS)
 - Reliable remote service invocation (GRAM)
 - High-performance remote data access (GrdiFTP, Data Grids)
 - ◆ Good quality reference implementation
 - ◆ Widespread acceptance by many communities
- Minuses
 - ◆ General ongoing “feature discovery” (firewalls, etc.)
 - ◆ Heterogeneous protocol basis
 - ◆ Little work on total system properties: dependability, QoS,
 - ◆ Need for a higher level of abstraction
- Direction: standardization efforts



The Emergence of Open Grid Standards



OGSA Structure

- Key abstraction: the Grid service
 - ◆ Standard interfaces and behaviors that address key distributed system issues: naming, service state, lifetime, notification
 - ◆ A Grid service is a Web service
 - ... supports standard service specifications
 - ◆ Agreement, data access & integration, workflow, security, policy, diagnostics, etc.
 - ◆ Target of current & planned GGF efforts
 - ... and arbitrary application-specific services based on these & other definitions
-

Open Grid Services Infrastructure

Client

Introspection:

- What port types?
- What policy?
- What state?

Lifetime management

- Explicit destruction
- Soft-state lifetime

Grid Service Handle

handle resolution

Grid Service Reference

GridService (required)

My service

Other standard interfaces:

factory,
notification,
collections

Service data element

Service data element

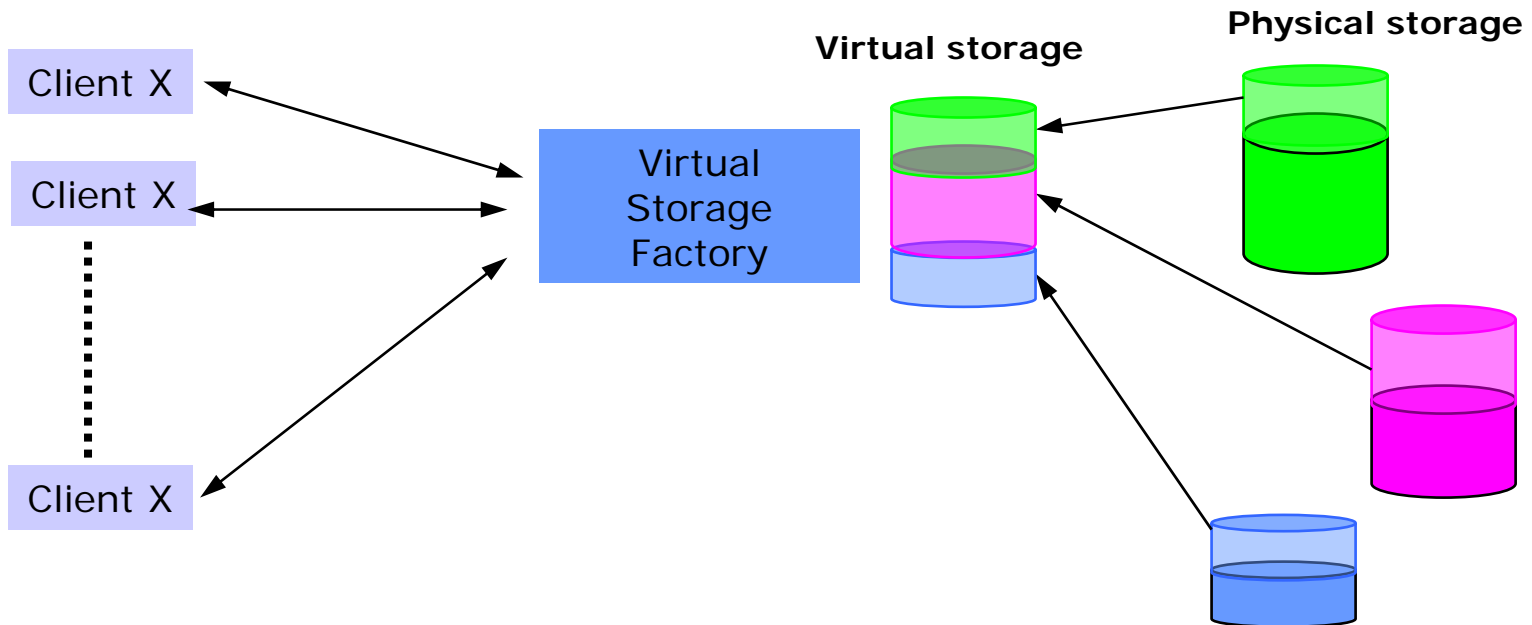
Service data element

Implementation

Hosting environment/runtime ("C", J2EE, .NET, ...)

Virtualization: Example

- Application: Virtual Storage
 - ◆ Garbage collecting unused space in an organization
 - ◆ Providing it to users as “virtual storage”



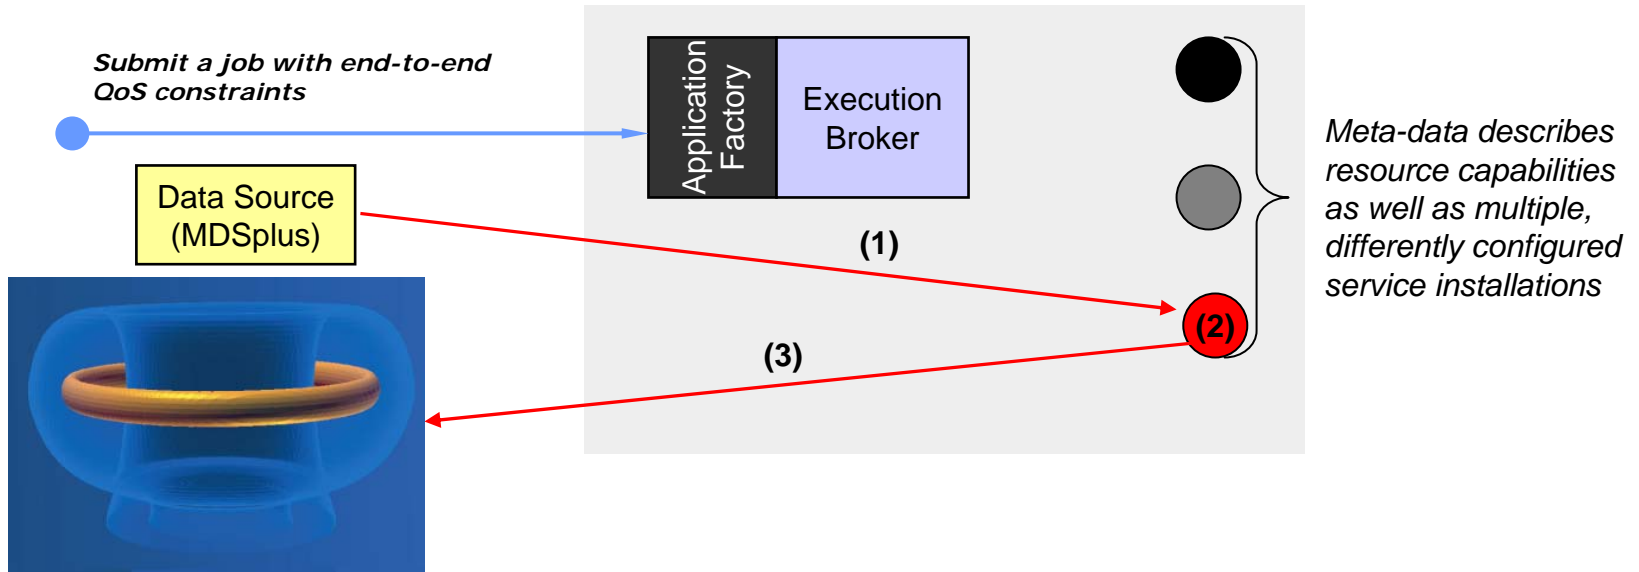
Quality of Service in the Grids

- If power grids were like computational grids...
 - ◆ Toaster analogy:
 - Where is/are your power station(s)?
 - How much power do they produce?
 - Does the power supply pattern vary through the day?
 - What if your power station goes down?
 - ◆ Unless you can handle these questions...
 - ◆ No toast!
- Solutions
 - ◆ Switch to serial ; -)
 - ◆ Keep trying
 - ◆ A different kind of breakfast

Resource Management (NFC)

Client's view:

Service Provider's view:

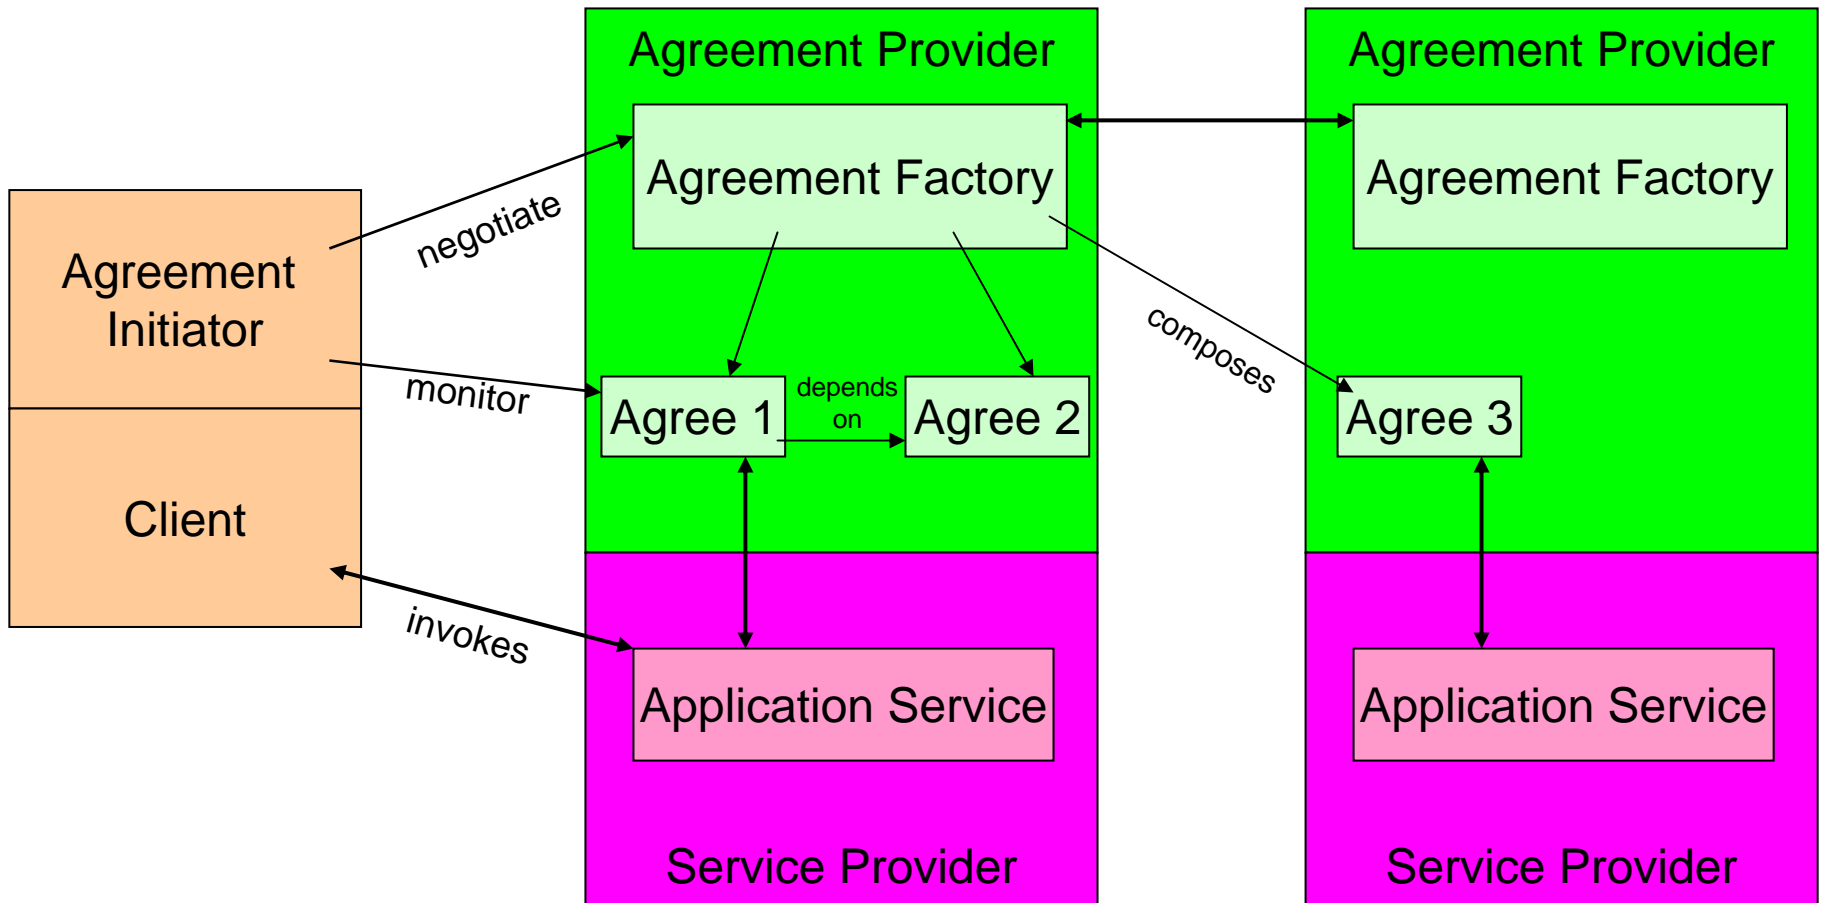


- Service Level Agreements (time critical)
- End-to-end QoS-driven execution
- Managing a workflow of multiple activities
- Adaptive capabilities to satisfy composite QoS

QoS Management Requirements

- Service terms: a way to express QoS
 - ◆ Execution time, responsiveness, etc...
 - ◆ Payment (money, barter, exchange)
- Capability to combine different service terms
 - ◆ Clients are interested in application-specific and resource-independent terms
 - ◆ Resource owners want to express their offering in resource-specific terms
- Capability to combine terms automatically
 - ◆ Dynamically changing Grid environment
 - ◆ Reliability
 - ◆ Adaptation
 - ◆ Autonomic computing
- Need an architecture expressing those needs

OGSI Agreement



OGSI Agreement

- Benefits
 - ◆ Client can compare different “offers” from different providers
 - ◆ Resource owner can plan and optimally manage resource allocation
- Current directions
 - ◆ Tools for fine-grain resource management on platforms
 - Flexible sharing of resources
 - Virtual machines
 - ◆ Adaptation
 - ◆ Protection, resource management, state management

User Execution Environment

- Dynamically created local execution environment
 - ◆ Obtained from OGSF factory after proper authorization
 - ◆ Protection of local execution environment
 - Unix accounts, VMs, sandboxes
 - ◆ Resource management and sandboxing
 - ◆ User state
 - ◆ Can be composed to form distributed execution environments
- Benefits
 - ◆ Alleviates administrative burden
 - ◆ Formalizing session management
 - Leverage different implementations with the same abstraction
 - ◆ Portable sessions
- Status: a GT3 prototype for NFC

A View of Grid Community

- Grids grew out of scientific community
 - ◆ The most immediate need for sharing
 - ◆ Access to resources important
 - ◆ Primarily academic involvement
- International community
 - ◆ Multiple international projects: UK e-science, Gridlab, European Data Grid, NorduGrid, others, collaborators on other continents
- Industry involvement
 - ◆ Widespread industry support: IBM, Microsoft, Sun, HP, Oracle, SGI, Cray, Platform, Entropia, Avaki, others...
 - ◆ Committed to use of GT3 & OGSA
- Open standards
 - ◆ Global Grid Forum (GGF)
- Combination of open standards, open source & industry involvement

Summary

- Grids are maturing
 - ◆ Grids have fulfilled their basic promise
 - ◆ The community is growing
 - More users using Grid for more things
 - More developers, support
 - Stronger outreach to community, standardization, etc.
 - ◆ The infrastructure is becoming more reliable
- The Grid is still a frontier town
 - ◆ Reaching out to new capabilities
 - ◆ Long way from cooking your toast ; -)
- Your requirements could be driving tomorrow's research
- More information
 - ◆ www.globus.org
 - ◆ Kate Keahey (keahey@mcs.anl.gov)