the globus project™
www.globus.org

The Globus Toolkit™: Current Status, Future Directions

Carl Kesselman

Director of the Center for Grid Technologies
Information Sciences Institute
University of Southern California

Chief Software Architect
National Partnership for Advanced computational Infrastructure

Partial Acknowledgements

- Open Grid Services Architecture design
  - Karl Czajkowski @ USC/ISI
  - Ian Foster, Steve Tuecke @ANL
  - Jeff Nick, Steve Graham, Jeff Frey @ IBM
- Grid services collaborators at ANL
  - Kate Keahey, Gregor von Laszewski
  - Thomas Sandholm, Jarek Gawor, John Bresnahan
- Globus Toolkit R&D also involves many fine scientists & engineers at ANL, USC/ISI, and elsewhere (see www.globus.org)
- Strong links with many EU, UK, US Grid projects
- Support from DOE, NASA, NSF, Microsoft

The Grid Problem

Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations

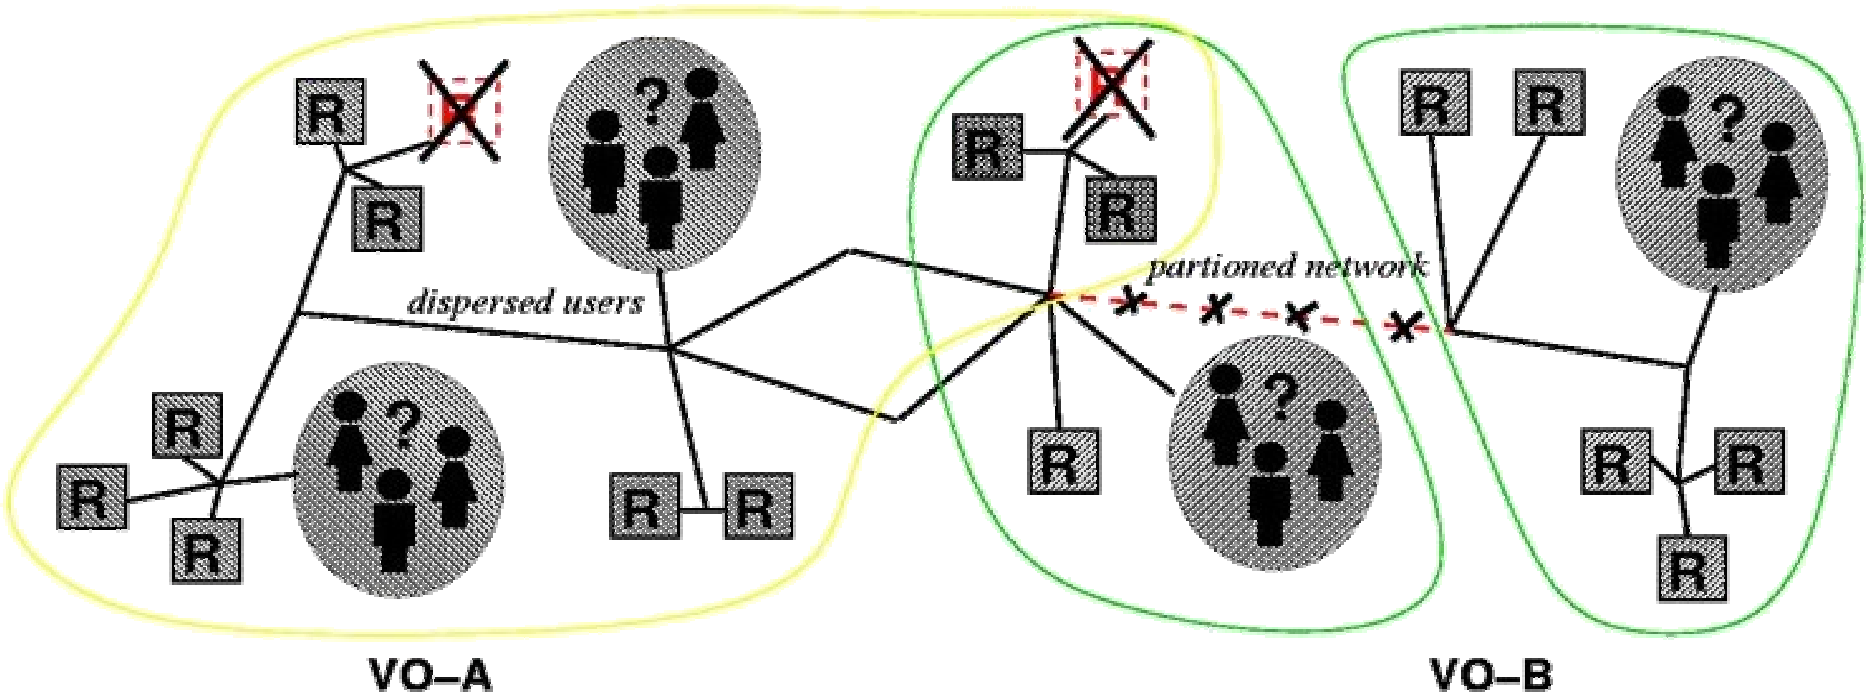

Grid Computing Concept

- New applications enabled by the coordinated use of geographically distributed resources
  - E.g., distributed collaboration, data access and analysis, distributed computing
- Persistent infrastructure for Grid computing
  - E.g., certificate authorities and policies, protocols for resource discovery/access
- Original motivation, and support, from high-end science and engineering; but has wide-ranging applicability

Broader Context

- "Grid Computing" has much in common with major industrial thrusts
  - Business-to-business, Peer-to-peer, Application Service Providers, Internet Computing, …
- Distinguished primarily by more sophisticated sharing modalities
  - E.g., "run program X at site Y subject to community policy P, providing access to data at Z according to policy Q"
  - Secondarily by unique demands of advanced & high-performance systems

Grid Communities & Applications: Data Grids for High Energy Physics

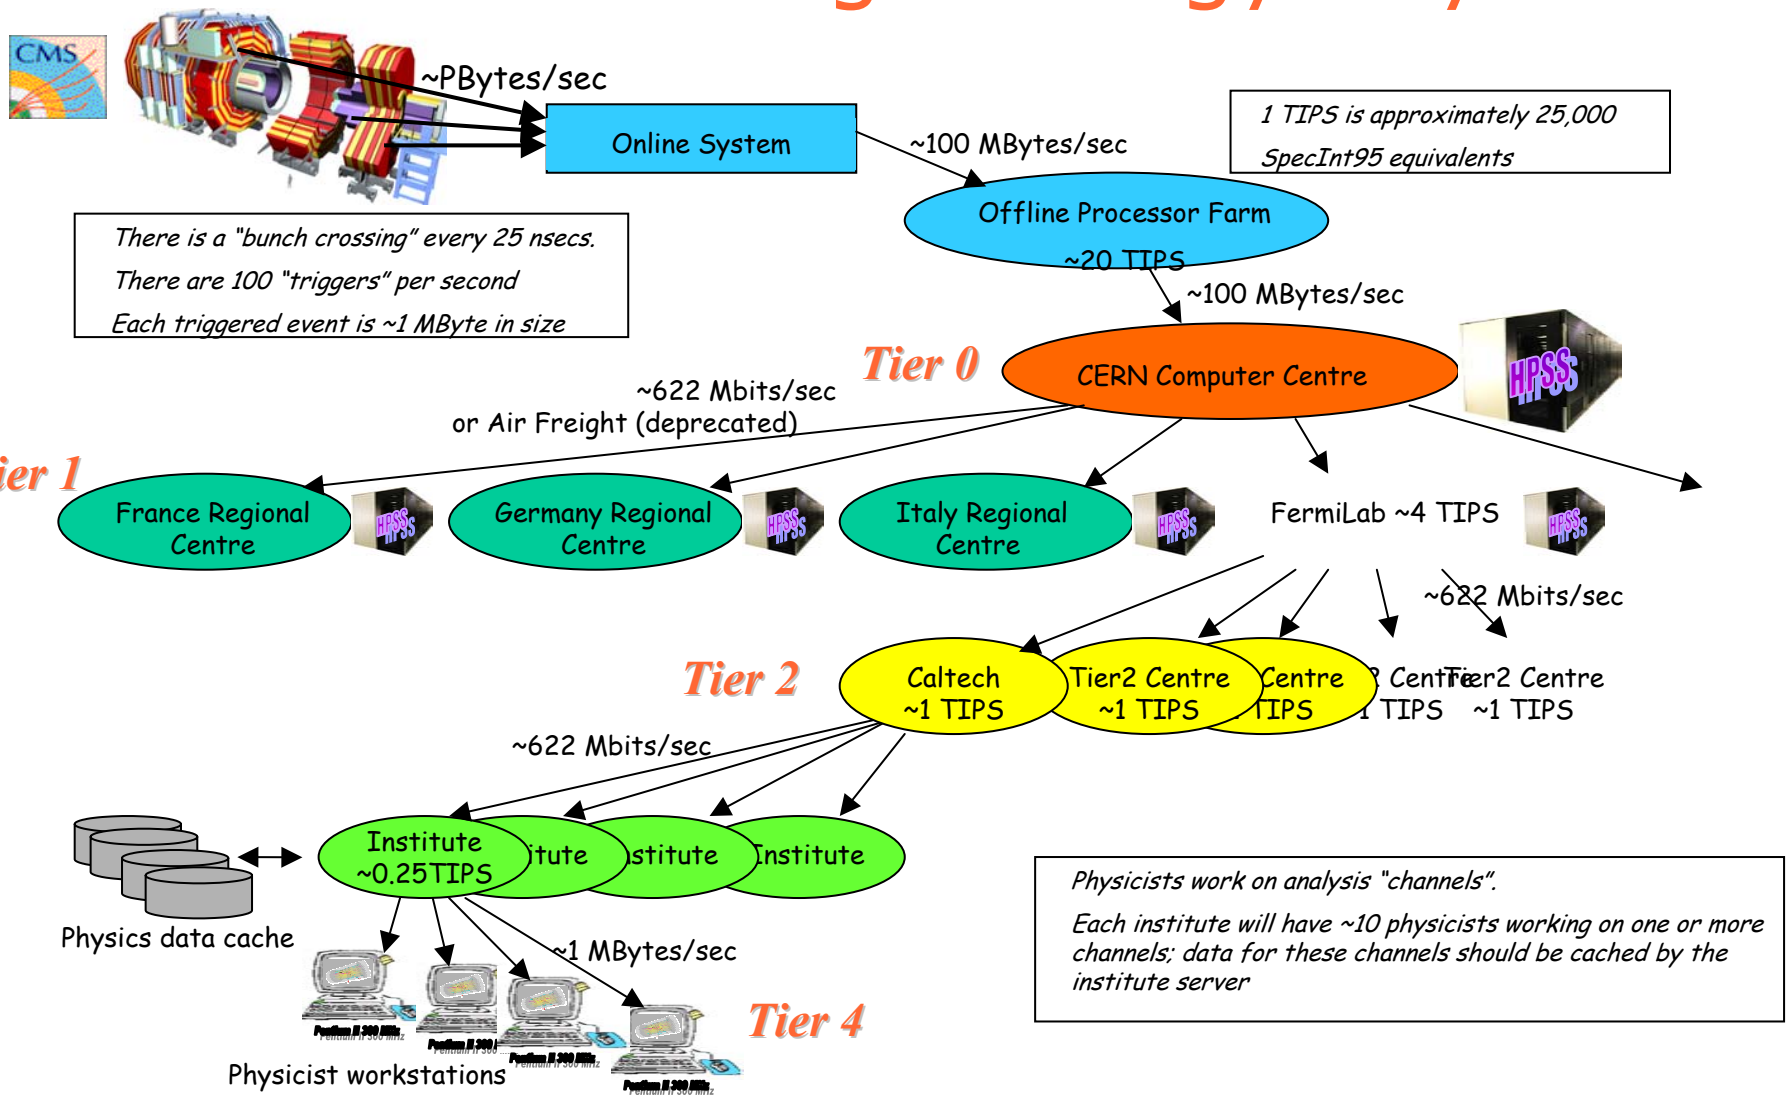

Grid Communities and Applications: Network for Earthquake Eng. Simulation

- NEESgrid: US national infrastructure to couple earthquake engineers with experimental facilities, databases, computers, & each other
- On-demand access to experiments, data streams, computing, archives, collaboration

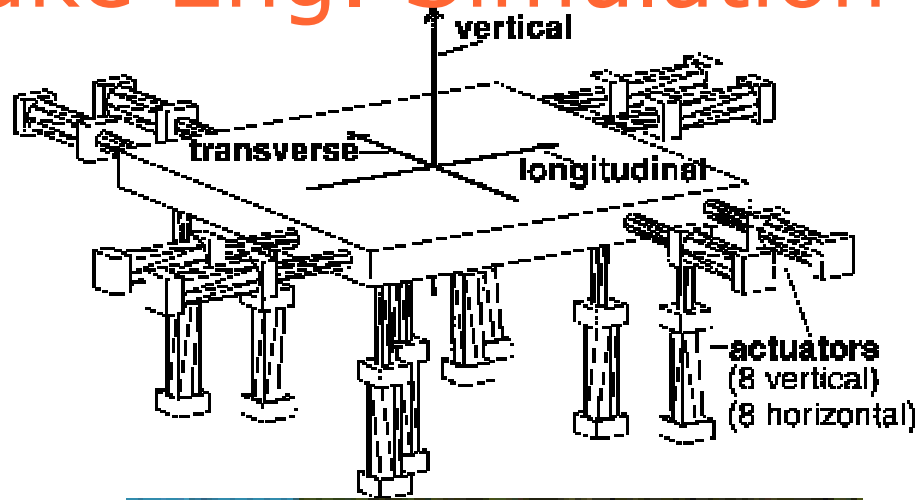

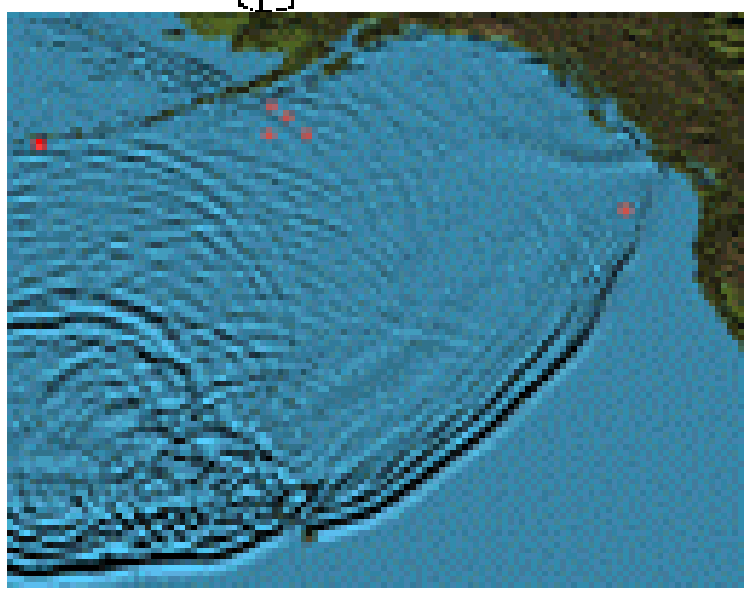

NEESgrid: Argonne, Michigan, NCSA, UIUC, USC www.neesgrid.org

Intelligent Infrastructure: Distributed Servers and Services

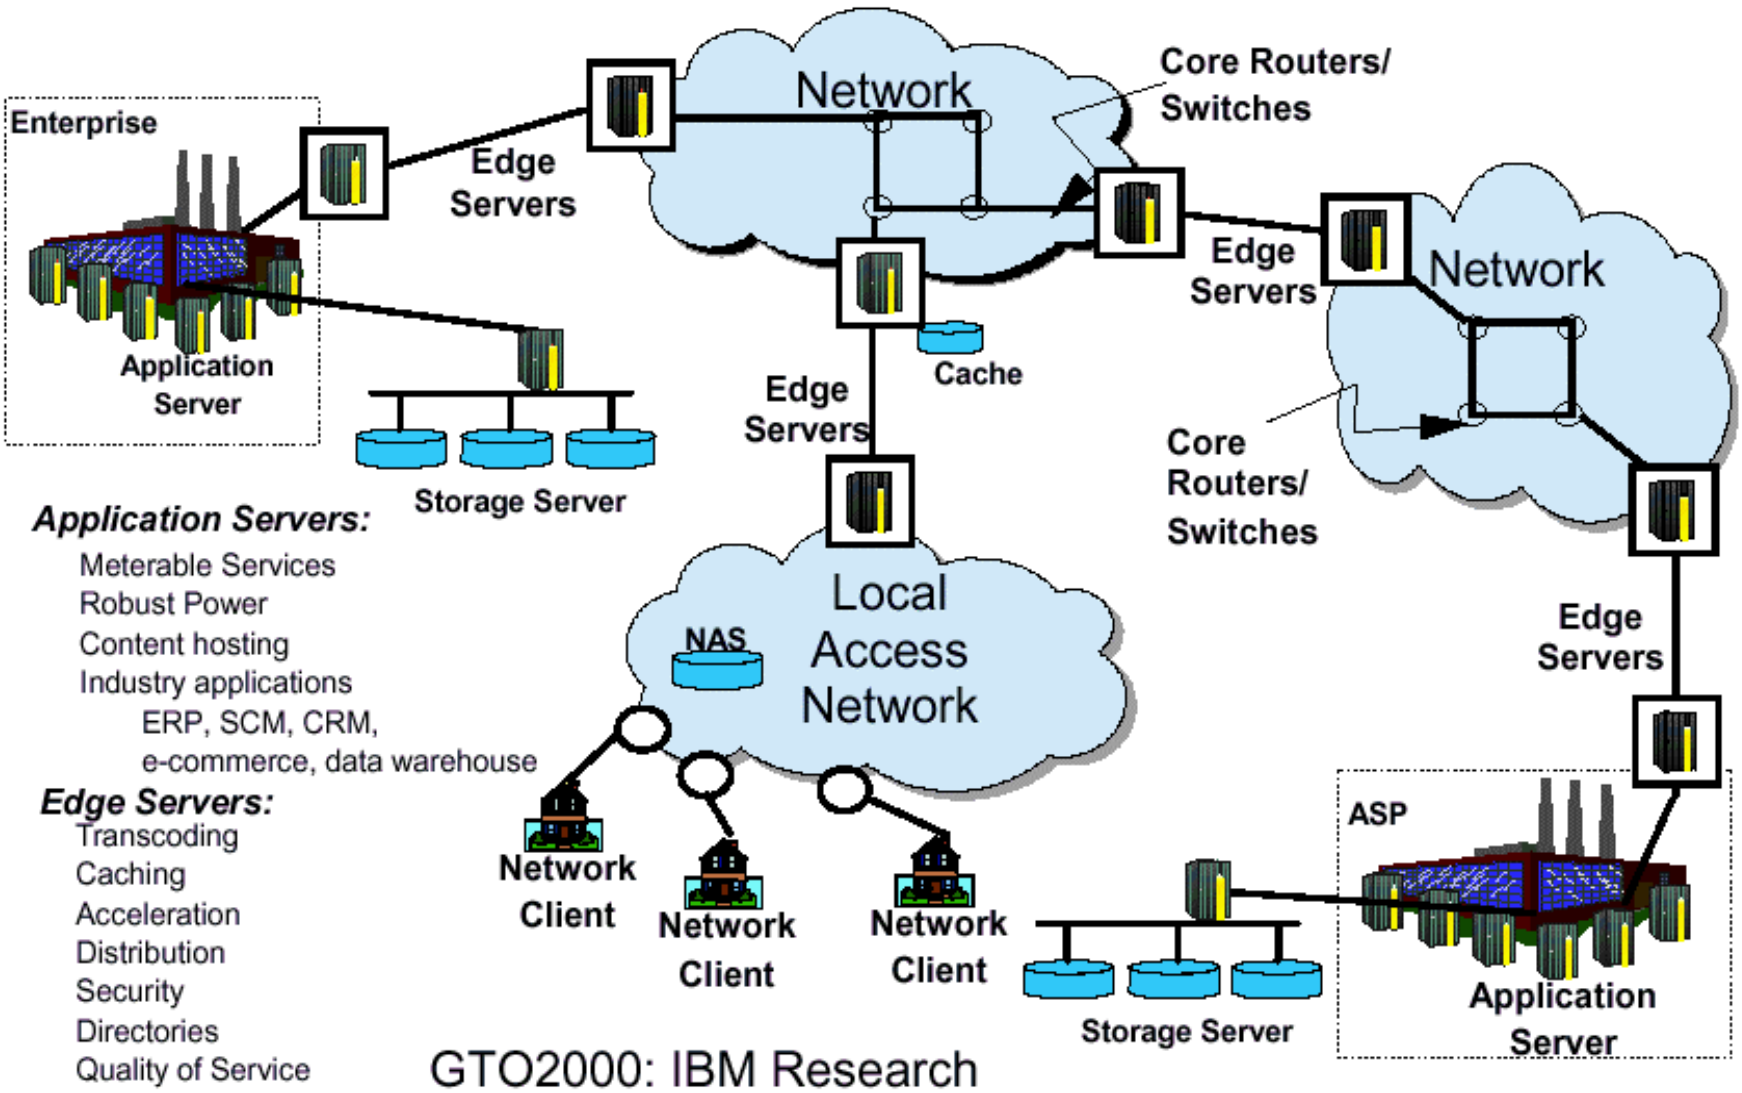

Application Servers:

- Meterable Services
- Robust Power
- Content hosting
- Industry applications
  - ERP, SCM, CRM, e-commerce, data warehouse

Edge Servers:

- Transcoding
- Caching
- Acceleration
- Distribution
- Security
- Directories
- Quality of Service

GTO2000: IBM Research

Grids: Why Now?

- Moore's law ⇒ highly functional end-systems
- Ubiquitous Internet ⇒ universal connectivity
- Network exponentials produce dramatic changes in geometry and geography
  - 9-month doubling: double Moore's law!
  - 1986-2001: x340,000; 2001-2010: x4000?
- New modes of working and problem solving emphasize teamwork, computation
- New business models and technologies facilitate outsourcing

Elements of the Problem

- Resource sharing
  - Computers, storage, sensors, networks, …
  - Heterogeneity of device, mechanism, policy
  - Sharing conditional: negotiation, payment, …
- Coordinated problem solving
  - Integration of distributed resources
  - Compound quality of service requirements
- Dynamic, multi-institutional virtual orgs
  - Dynamic overlays on classic org structures
  - Map to underlying control mechanisms

The Grid World: Current Status

- Dozens of major Grid projects in scientific & technical computing/research & education
  - Deployment, application, technology
- Considerable consensus on key concepts and technologies
  - Open source Globus Toolkit™ a de facto standard for major protocols & services
  - Far from complete or perfect, but out there, evolving rapidly, and large tool/user base
- Global Grid Forum a significant force
- Industrial interest emerging rapidly

Layered Grid Architecture (By Analogy to Internet 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

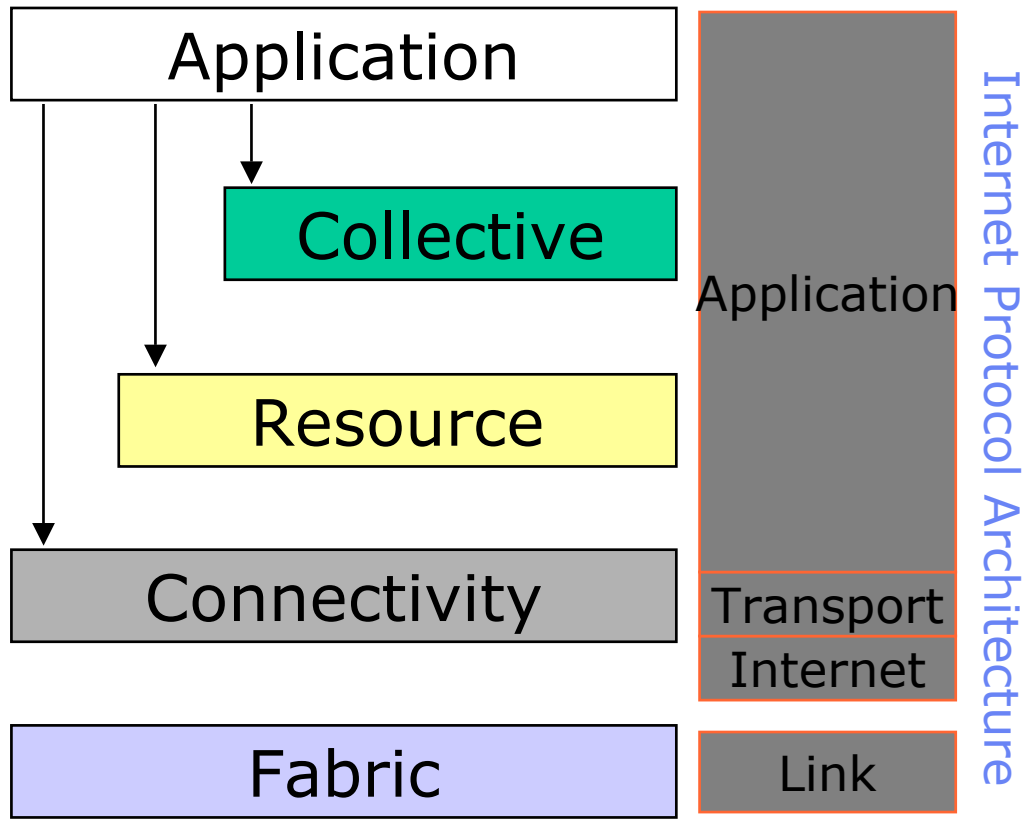

The Globus Toolkit in One Slide

- Grid protocols (GSI, GRAM, …) enable resource sharing within virtual orgs; toolkit provides reference implementation (= Globus Toolkit services)

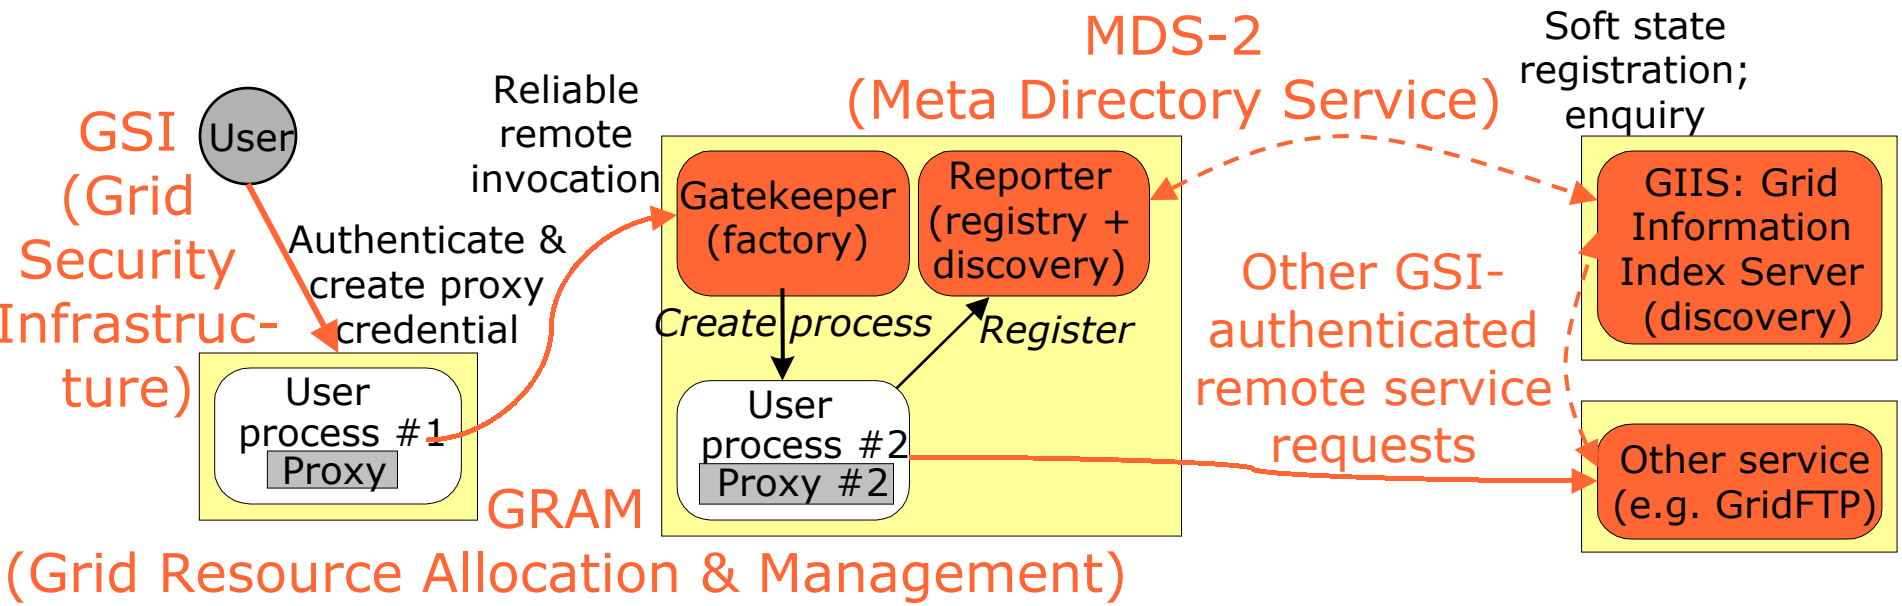

- Protocols (and APIs) enable other tools and services for membership, discovery, data mgmt, workflow, …

Globus Toolkit: Evaluation (+)

- Good technical solutions for key problems, e.g.
  - Authentication and authorization
  - Resource discovery and monitoring
  - Reliable remote service invocation
  - High-performance remote data access
- This & good engineering is enabling progress
  - Good quality reference implementation, multi-language support, interfaces to many systems, large user base, industrial support
  - Growing community code base built on tools

Globus Toolkit: Evaluation (-)

- Protocol deficiencies, e.g.
  - Heterogeneous basis: HTTP, LDAP, FTP
  - No standard means of invocation, notification, error propagation, authorization, termination, …
- Significant missing functionality, e.g.
  - Databases, sensors, instruments, workflow, …
  - Virtualization of end systems (hosting envs.)
- Little work on total system properties, e.g.
  - Dependability, end-to-end QoS, …
  - Reasoning about system properties

"Web Services"

- Increasingly popular standards-based framework for accessing network applications
  - W3C standardization; Microsoft, IBM, Sun, others
- WSDL: Web Services Description Language
  - Interface Definition Language for Web services
- SOAP: Simple Object Access Protocol
  - XML-based RPC protocol; common WSDL target
- WS-Inspection
  - Conventions for locating service descriptions
- UDDI: Universal Desc., Discovery, & Integration
  - Directory for Web services

Transient Service Instances

- "Web services" address discovery & invocation of persistent services
  - Interface to persistent state of entire enterprise
- In Grids, must also support transient service instances, created/destroyed dynamically
  - Interfaces to the states of distributed activities
  - E.g. workflow, video conf., dist. data analysis
- Significant implications for how services are managed, named, discovered, and used
  - In fact, much of our work is concerned with the management of service instances

the globus project™
www.globus.org

OGSA Design Principles

- Service orientation to virtualize resources
  - Everything is a service
- From Web services
  - Standard interface definition mechanisms: multiple protocol bindings, local/remote transparency
- From Grids
  - Service semantics, reliability and security models
  - Lifecycle management, discovery, other services
- Multiple "hosting environments"
  - **C**, J2EE, .NET, …

OGSA Service Model

- System comprises (a typically few) persistent services & (potentially many) transient services
  - Everything is a service
- OGSA defines basic behaviors of services: fundamental semantics, life-cycle, etc.
  - More than defining WSDL wrappers

Open Grid Services Architecture: Fundamental Structure

1) WSDL conventions and extensions for describing and structuring services
   - Useful independent of "Grid" computing
2) Standard WSDL interfaces & behaviors for core service activities
   - portTypes and operations => protocols

WSDL Conventions & Extensions

- portType (standard WSDL)
  - Define an interface: a set of related operations
- serviceType (extensibility element)
  - List of port types: enables aggregation
- serviceImplementation (extensibility element)
  - Represents actual code
- service (standard WSDL)
  - instanceOf extension: map descr.->instance
- compatibilityAssertion (extensibility element)
  - portType, serviceType, serviceImplementation

Use of Web Services

- A Grid service definition is a WSDL extension (serviceImplimentation) containing:
  - A serviceType definition
    - A list of portTypes
      - A set of operations
        - An exchange of messages
- A Grid service implementation is a service element containing
  - Implements declaration referencing a serviceImplimentation

Structure of a Grid Service

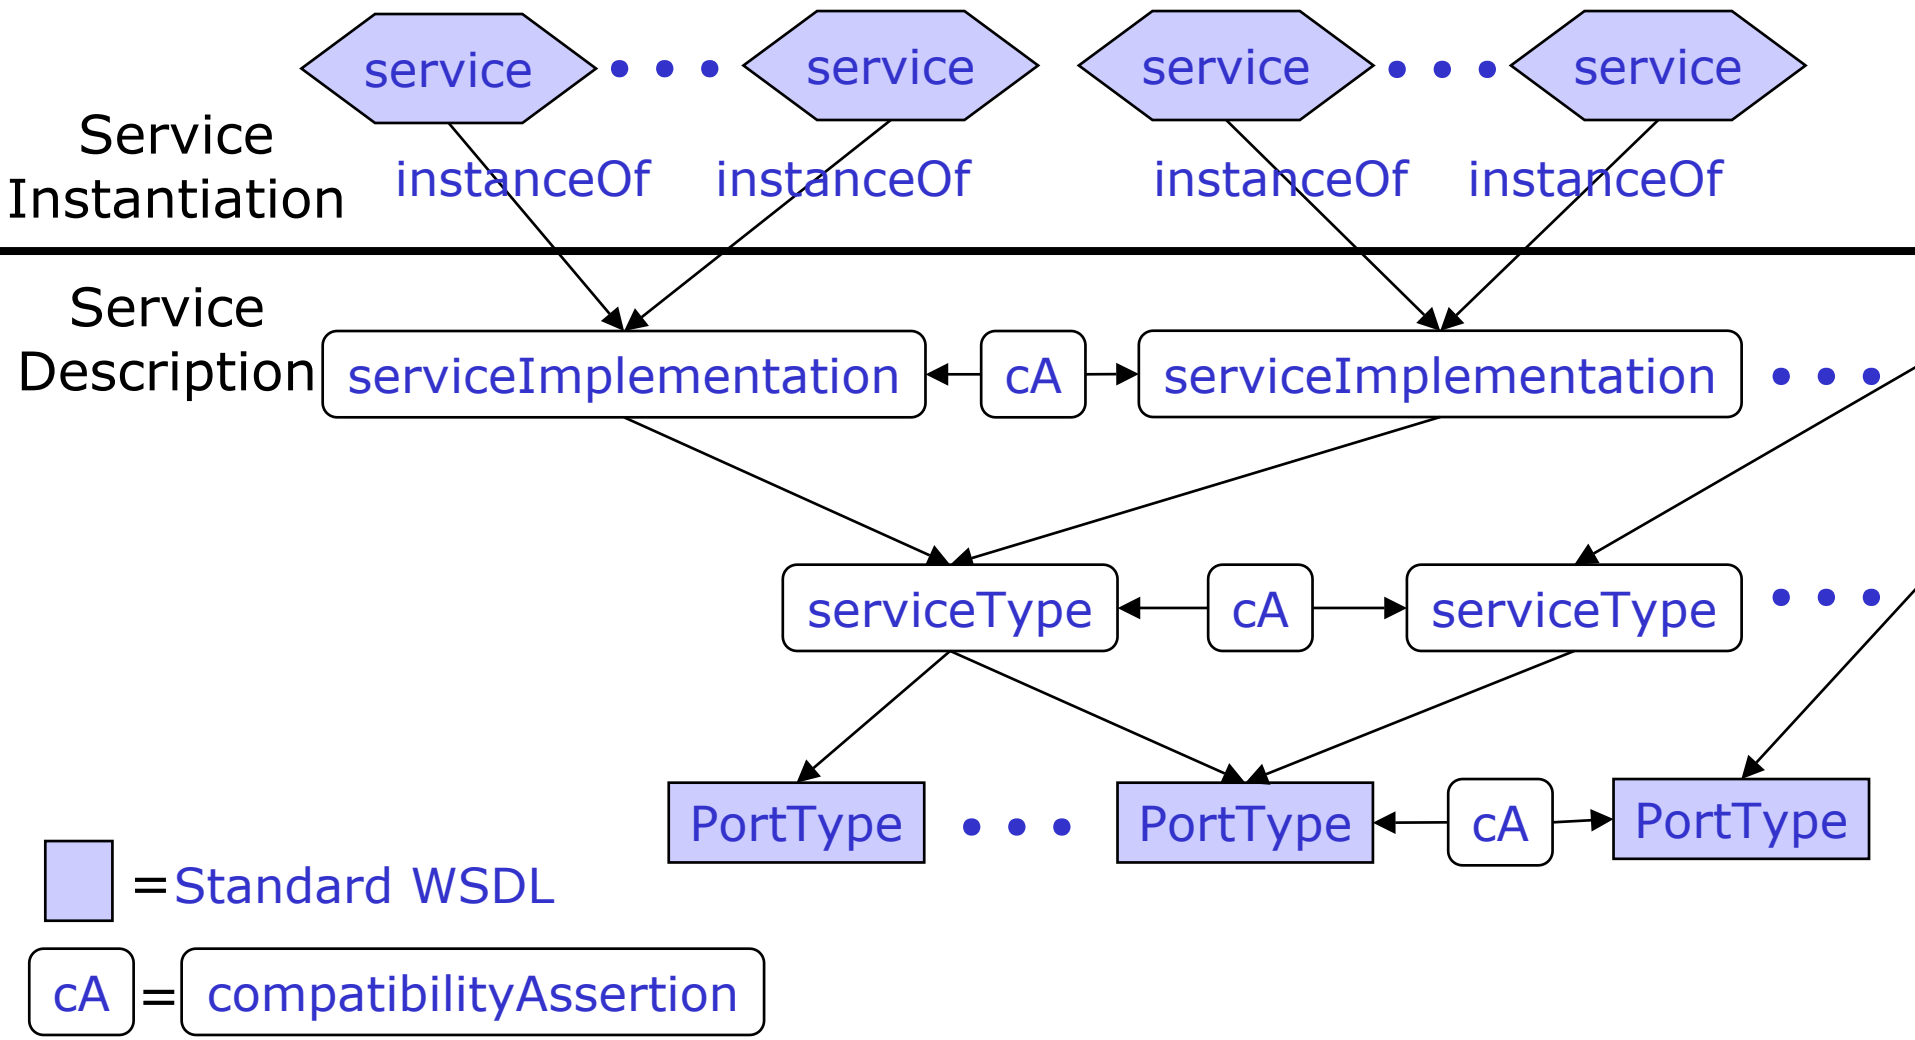

Standard Interfaces & Behaviors: Four Interrelated Concepts

- Naming and bindings
  - Every service instance has a unique name, from which can discover supported bindings
- Information model
  - Service data associated with Grid service instances, operations for accessing this info
- Lifecycle
  - Service instances created by factories
  - Destroyed explicitly or via soft state
- Notification
  - Interfaces for registering interest and delivering notifications

OGSA Interfaces and Operations Defined to Date

- GridService *Required*
  - FindServiceData
  - Destroy
  - SetTerminationTime
- NotificationSource
  - SubscribeToNotificationTopic
  - UnsubscribeToNotificationTopic
- NotificationSink
  - DeliverNotification
- Factory
  - CreateService
- PrimaryKey
  - FindByPrimaryKey
  - DestroyByPrimaryKey
- Registry
  - RegisterService
  - UnregisterService
- HandleMap
  - FindByHandle

Authentication, reliability are binding properties
Manageability, concurrency, etc., to be defined

OGSA and the Globus Toolkit

- Technically, OGSA enables
  - Refactoring of protocols (GRAM, MDS-2, etc.)—*while preserving all GT concepts/features*!
  - Integration with hosting environments: simplifying components, distribution, etc.
  - Greatly expanded standard service set
- Pragmatically, we are proceeding as follows
  - Develop open source OGSA implementation
    - > Globus Toolkit 3.0; supports Globus Toolkit 2.0 APIs
  - Partnerships for service development
  - Also expect commercial value-adds

Globus Toolkit Refactoring

- Grid Security Infrastructure (GSI)
  - Used in Grid service network protocol bindings
- Meta Directory Service 2 (MDS-2)
  - Native part of each Grid service:
    - > Discovery, Registry, RegistryManagement, Notification
- Grid Resource Allocation & Mngt (GRAM)
  - Gatekeeper -> Factory for job mgr instances
- GridFTP
  - Refactor control channel protocol
- Other services refactored to used Grid services

Summary: Evolution of Grid Technologies

- Initial exploration (1996-1999; Globus 1.0)
  - Extensive appln experiments; core protocols
- Data Grids (1999-??; Globus 2.0+)
  - Large-scale data management and analysis
- Open Grid Services Architecture (2001-??, Globus 3.0)
  - Integration w/ Web services, hosting environments, resource virtualization
  - Databases, higher-level services
- Radically scalable systems (2003-??)
  - Sensors, wireless, ubiquitous computing

Summary

- The Grid problem: Resource sharing & coordinated problem solving in dynamic, multi-institutional virtual organizations
- Grid architecture: Protocol, service definition for interoperability & resource sharing
- Globus Toolkit a source of protocol and API definitions—and reference implementations
  - And many projects applying Grid concepts (& Globus technologies) to important problems
- Open Grid Services Architecture represents (we hope!) next step in evolution

For More Information

- The Globus Project™
  - www.globus.org
- Grid architecture
  - www.globus.org/research/papers/anatomy.pdf
- Open Grid Services Architecture
  - www.globus.org/ogsa