The Grid Architecture

William E. Johnston

Distributed Systems Department
Computational Research Division
Lawrence Berkeley National Laboratory
dsd.lbl.gov

What are the Core Grid Middleware Services?

- A common security model for Grid services and Grid applications provides uniform and versatile authentication, authorization, and privacy
  - Basis of the cyber-trust that enables collaboration among the many organizations of a large science project
  - Preserves local autonomy of resource owners
  - Correctly used, provides pretty good security
- Standardized access to computing systems and data storage systems
- Tools and services supporting construction and management of collaborations (virtual organizations)

What are the Core Grid Middleware Services

- Services for dynamic construction of execution environments supporting complex distributed applications
  - locating and co-scheduling many resources to support, e.g., transient and complex, science and engineering experiments that require combinations of instruments, compute systems, data archives, and network bandwidth at multiple locations
- Management of dynamic pools of underlying resources
  - Automatic resource registration and de-registration
  - Resource discovery
- Evolving to a Web Services / object oriented model of core services

SensorNet Architectural Considerations

- Web Services approach is very likely to be dominate mechanism for distributed applications interfaces
- Combining Grids with Web Services brings dynamic resource management to Web Services (major industry push lead by IBM)
- Grid Services are Web services with a few additional behavoirs (ports) defined to do service initiation (factory) and lifecycle management of transient, stateful (e.g. Grid) services – i.e. dynamically created services
- Global Grid Forum (www.ggf.org) – international standards org. modeled on IETF (about 500 members from N. America, Europe, and Asia Pacific)

User Interfaces

Application Frameworks (e.g. XCAT, SciRun) **and Portal Toolkits** (e.g. XPortlets)

Applications (Simulations, Data Analysis, etc.)

Higher-level Services / OGSA (Data Grid Services, Workflow management, Visualization, Data Publication/Subscription, Brokering, Job Mg'mt, Fault Mg'mt, Grid System Admin., etc.)

Core Grid Services / OGSI: Uniform access to distributed resources

- Grid Information Service
- Uniform Computing Access
- Uniform Data Access
- Unix and OGSI hosting
- Co-Scheduling
- Global Event Services, Auditing, Monitoring
- Systems management and access
- Communication Services
- Authentication Authorization
- Security Services

Grid Managed Resources

ESNet

X.509 CA

NERSC Supercomputing & Large-Scale Storage

LBNL

PNNL

Asia-Pacific

Europe

ESnet

DOE Science Grid

ANL

ORNL

Funded by the U.S. Dept. of Energy, Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division

scientific instruments

Supernova Observatory

Synchrotron Light Source

NWGISS Top Level Architecture

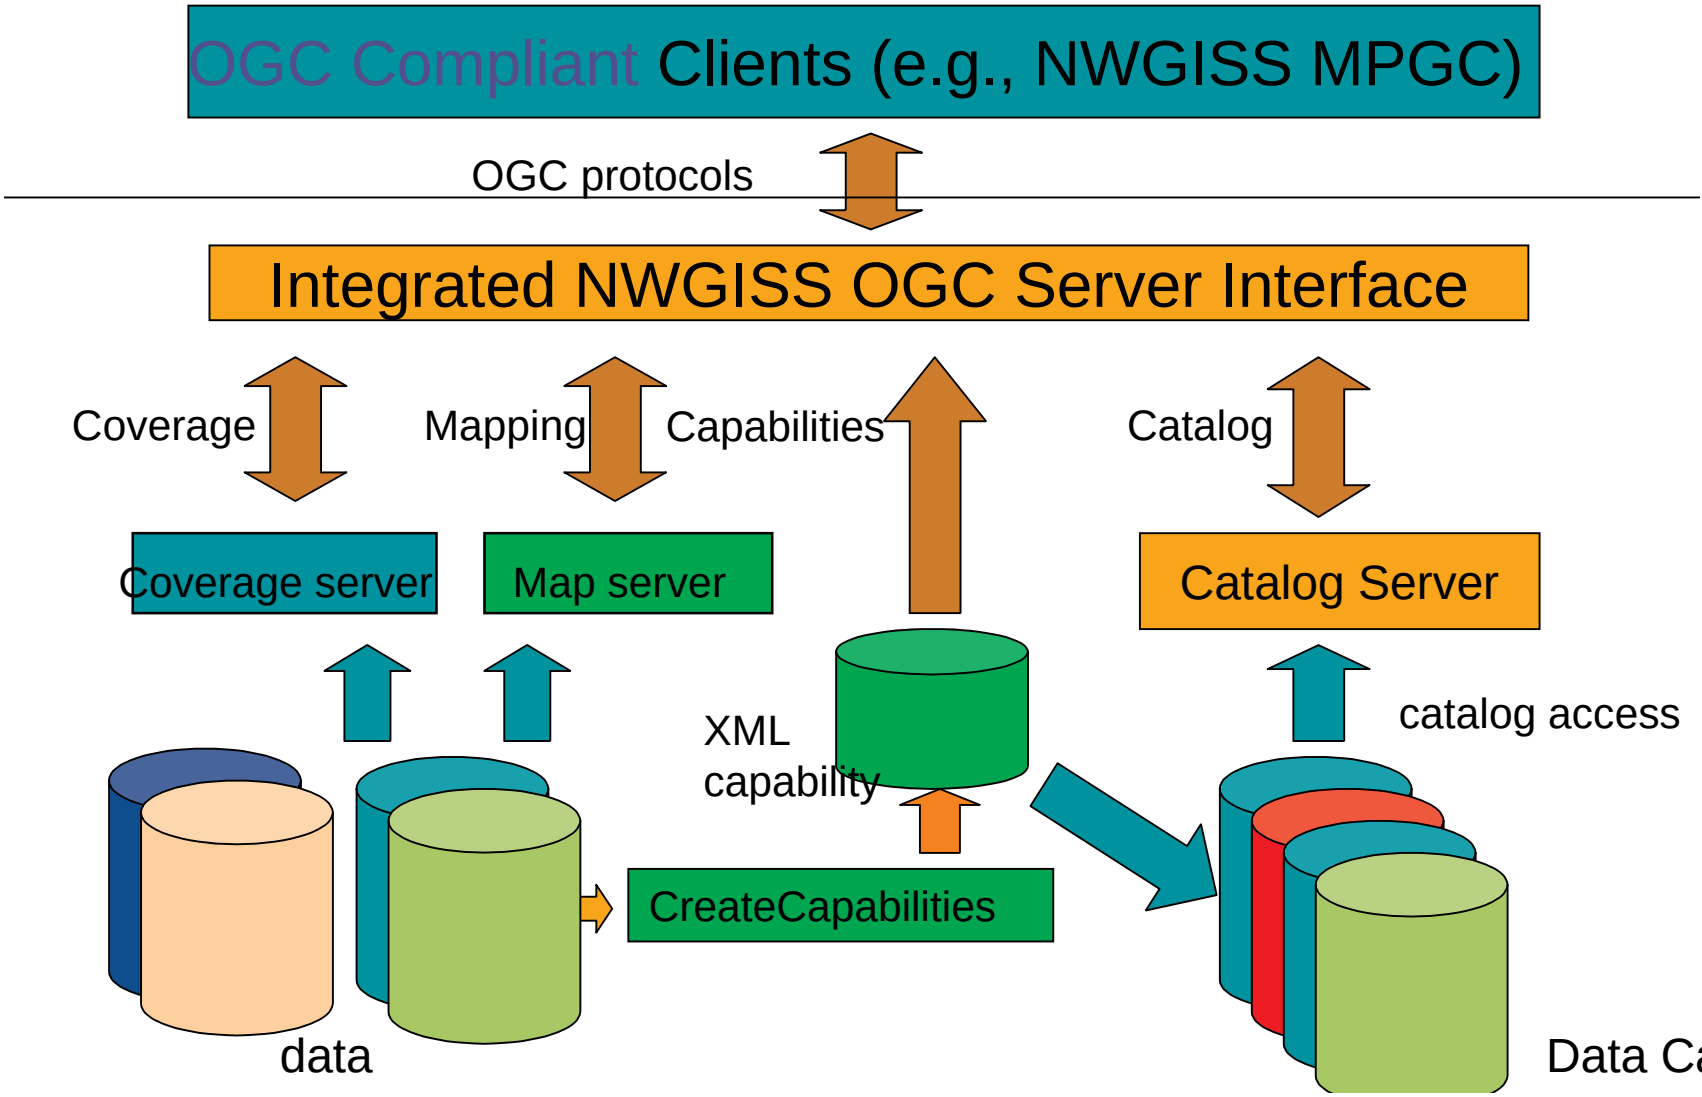

Portals

Higher Level Services

Frameworks

OGC Compliant Clients (e.g., NWGISS - NASA EOS DAAC access)

OGC protocols

Integrated NWGISS OGC Server Interface

Coverage | Mapping | Data generation prescriptions | Catalog

Coverage server | Map server | Catalog Server

Data managed by Grid Data Services

- Portal libraries
  - e.g. XPortlets (IU)
- Frameworks
  - e.g. XCAT (distributed CCA framework- IU)

Mostly Web Services Based

Toolkits and Collective services

- Workflow engine
  - WSFL
  - current state reporting
- Grid Data Services
  - version mg'mt
  - master dataset mg'mt
  - reliable file xfer
  - net caches
  - metadata cat'lg
- Replica Services
  - metadata
  - replica location
- Virtual Data Services
  - materialized data cat'lg
  - virtual data cat'lg
  - abstract planner
  - concrete planner
- DataBase access
  - e.g. OGSA DAIS

Core Grid Services / OGSI: Uniform access to distributed resources

Grid Information Service | Uniform Computing Access | Uniform Data Access | Unix and OGSI hosting | Co-Scheduling | Global Event Services, Auditing, Monitoring | Systems management and access | Communication Services | Authentication Authorization | Security Services

Grid Managed Resources

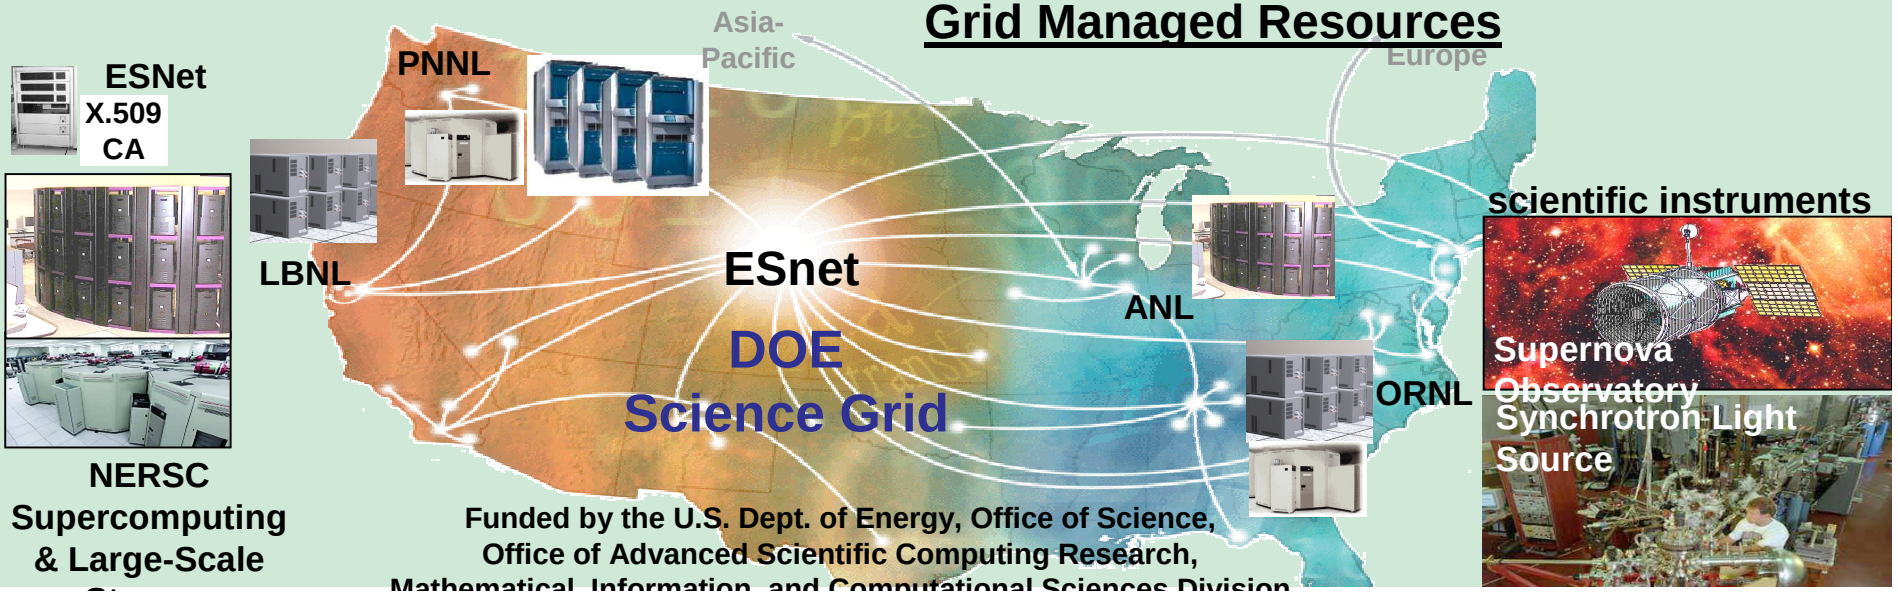

Funded by the U.S. Dept. of Energy, Office of Science, Office of Advanced Scientific Computing Research, Mathematical, Information, and Computational Sciences Division

Scalable Monitoring, Activation, and Publication Service

Brian Tierney

Dan Gunter

Distributed Systems Department

Computational Research Division

Lawrence Berkeley National Laboratory

Basic GGF, Grid Monitoring Architecture

(evolved from our early work with monitors – a direct descendent of MAGIC Testbed technology)

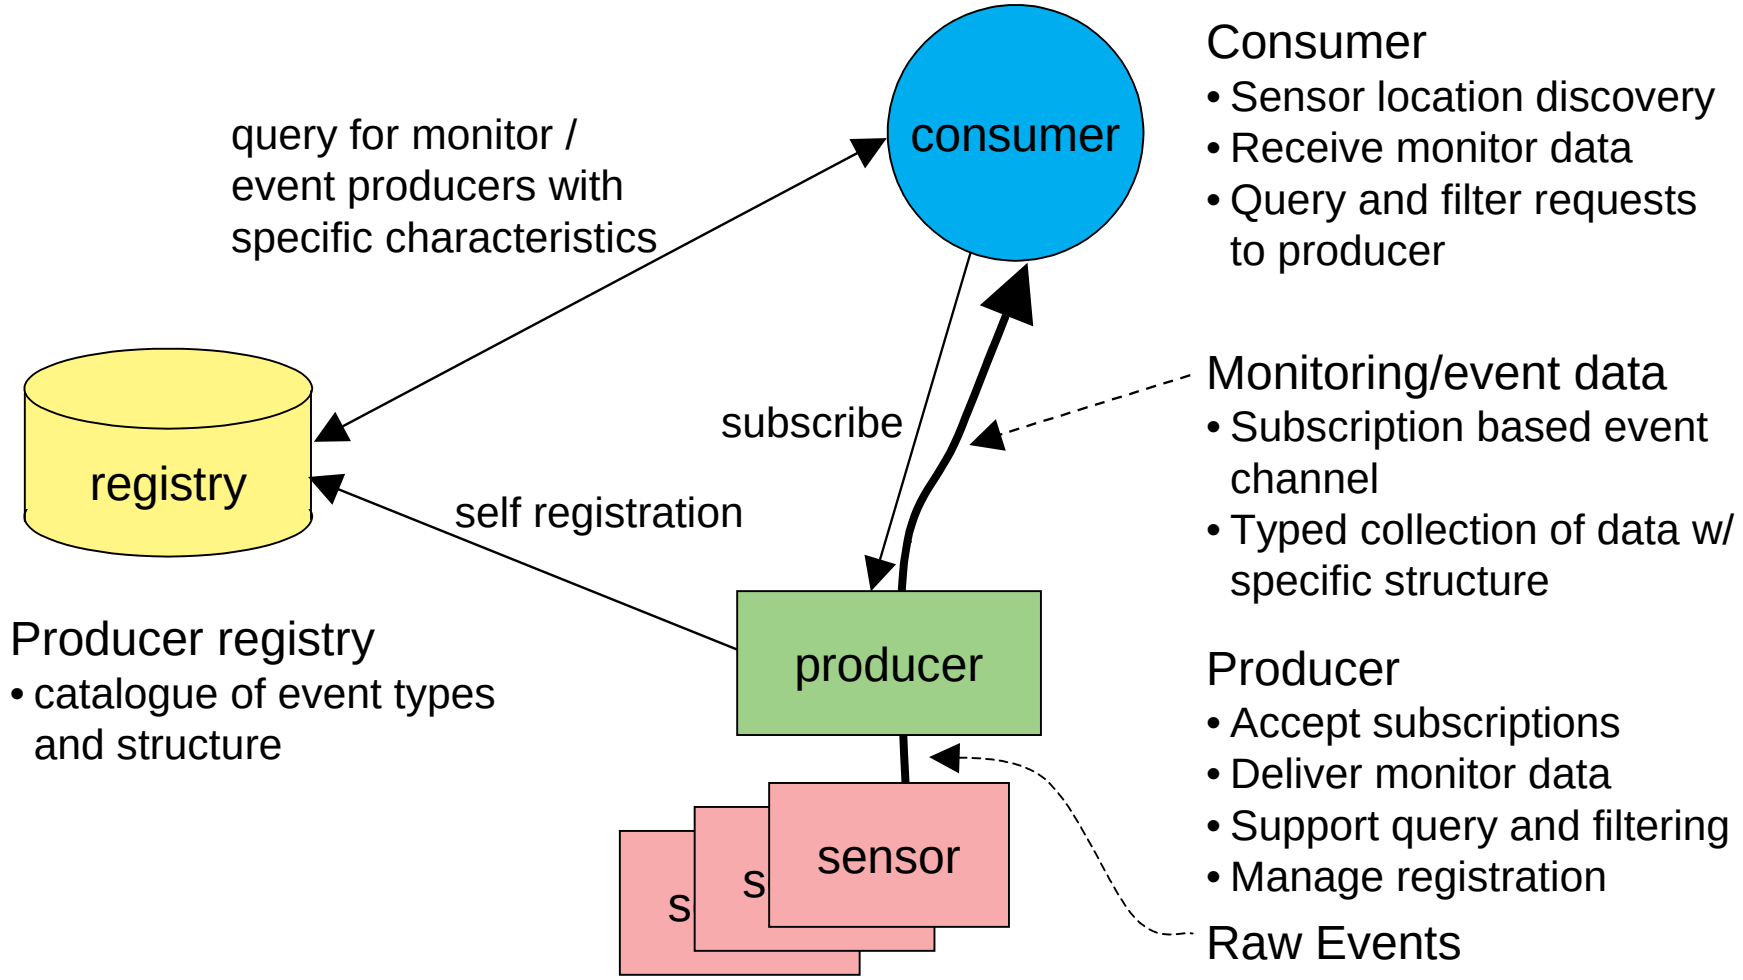

Consumer
- Sensor location discovery
- Receive monitor data
- Query and filter requests to producer

Monitoring/event data
- Subscription based event channel
- Typed collection of data w/ specific structure

Producer registry
- catalogue of event types and structure

Producer
- Accept subscriptions
- Deliver monitor data
- Support query and filtering
- Manage registration

Raw Events
- untyped data w/o specific structure

The GMA Producer and Consumer Interfaces are Well Defined – An Important Feature of the Architecture (Allows construction of standardized services)

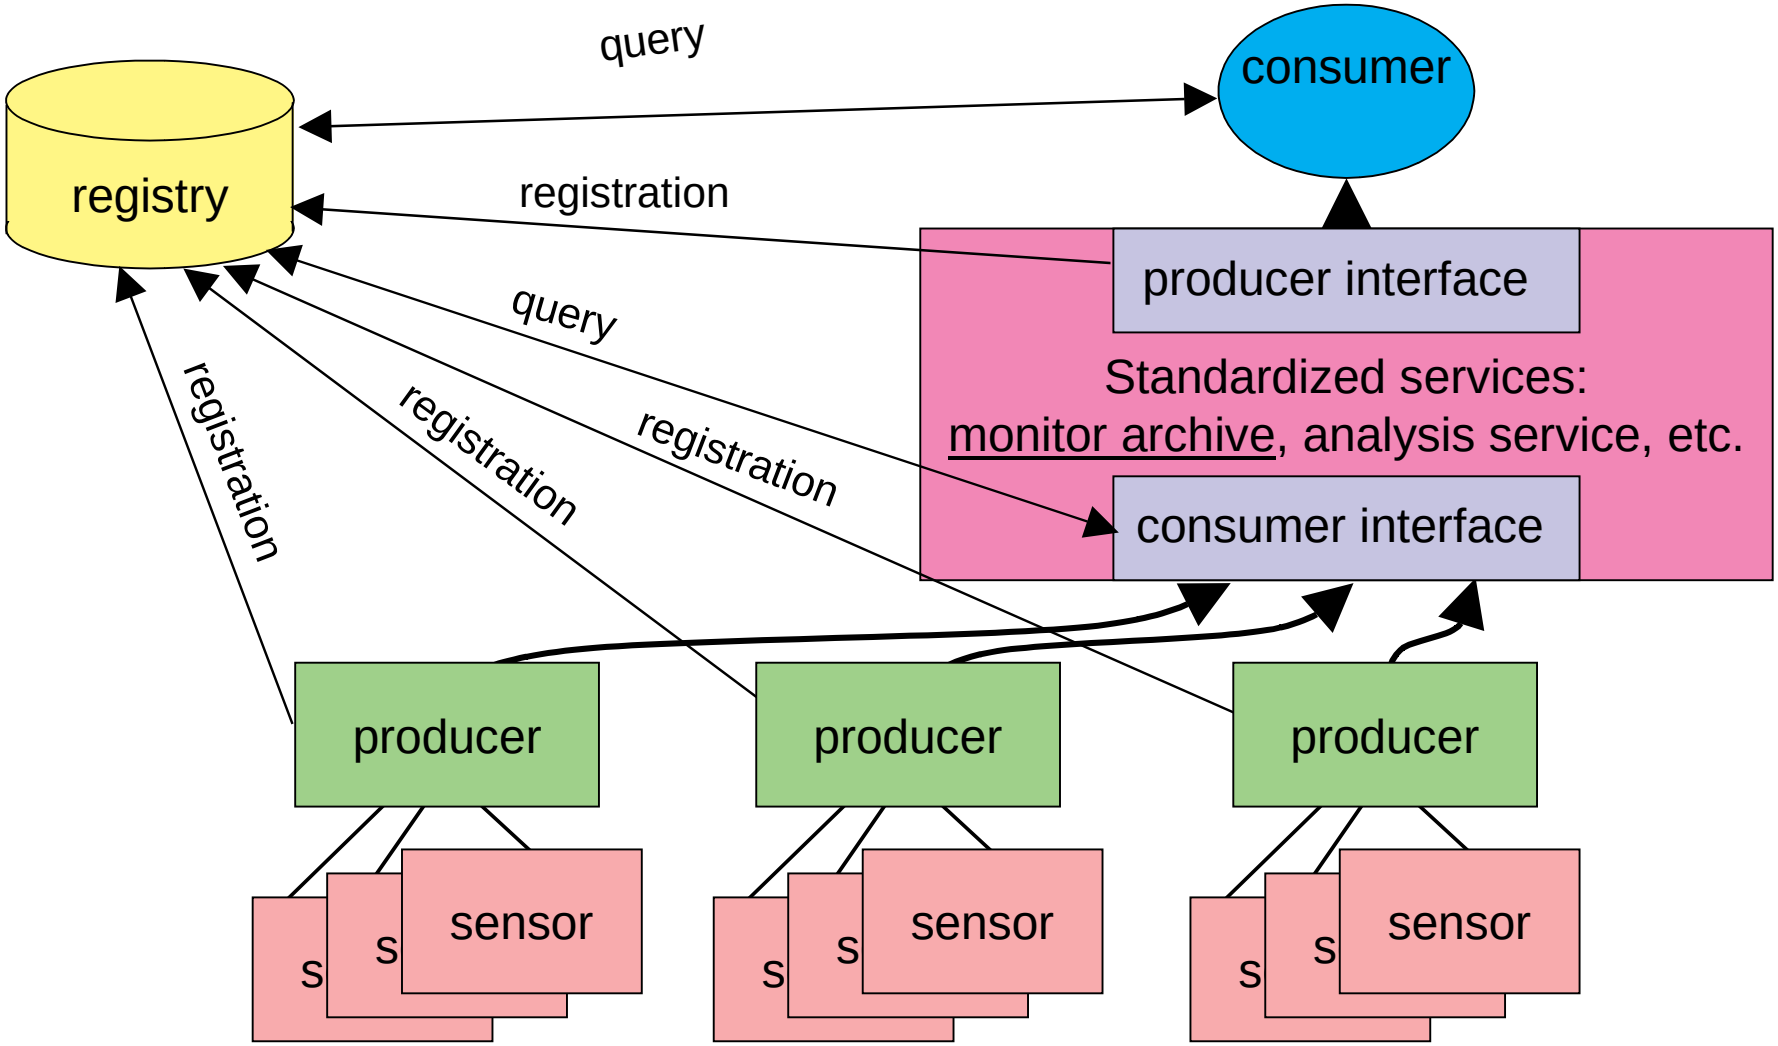

"A Grid Monitoring Architecture," B. Tierney, et al, Global Grid Forum Document.
http://www-didc.lbl.gov/GGF-PERF/GMA-WG/

Scalable Monitoring, Activation, and Publication Service (SMAPS)

- We are generalizing the basic Grid Monitoring Architecture in several dimensions in order to address various aspects of scalability
  - Lots of producers
  - Lots of consumers
  - Lots of sensors, high data rates
  - Self organizing, self healing collections of all of these
  - Complex producer subscriptions
    - e.g. generalized queries and downloadable event filters and analysis modules

D. Gunter, B. Tierney, C. E. Tull, V. Virmani, "On-Demand Grid Application Tuning and Debugging with the NetLogger Activation Service," 4th International Workshop on Grid Computing (Grid2003)

SMAPS: Separation of Command and Data and GMA Cache

(scalability for number and operational modes of sensors)

- Separation of command and data allows for standard cmd. protocol (SOAP) and API (WSDL defined) and efficient data transport
- Division of functionality using the standardized service interfaces allows, e.g.
  - load reduction for consumer
  - light-weight producers
    - § move heavy weight functions to a GMA cache system – e.g. filtering and compression for satellite up-link
    - § this gives light-weight producers, only managing sensors, providing standard messaging, subscription, etc. – more suitable for field / mobile units

Status: operational (skeletal registry)

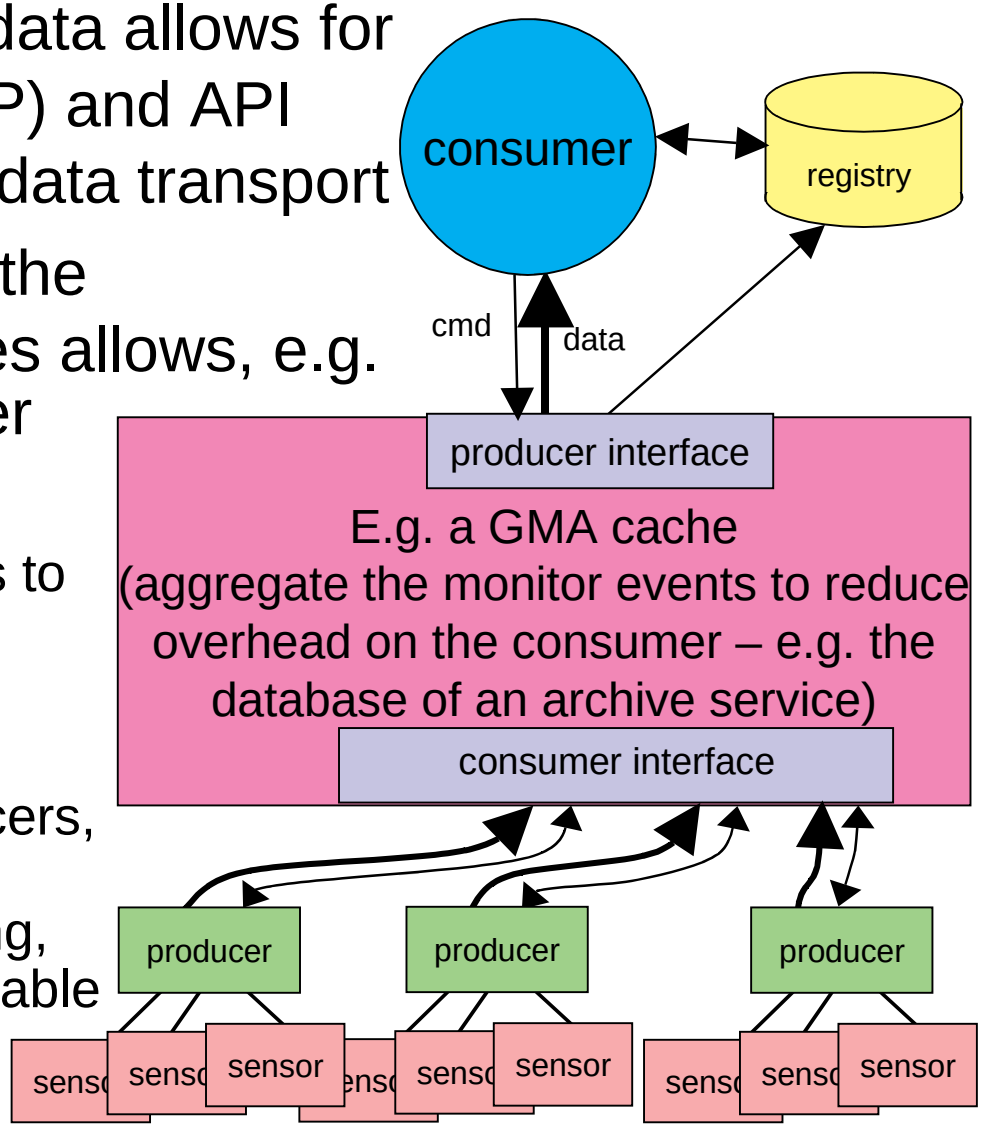

SMAPS: Sensor Activation
(scalable management of sensor data)

- Do not want all monitoring data collected all the time
  - Potentially too much data
  - Allow for adjusting the level of monitoring as needed for:
    - § operation
    - § debugging
    - § performance tuning
    - § error analysis
- Add a standard sensor control channel and “manager” function
  - manager may be combined with producer

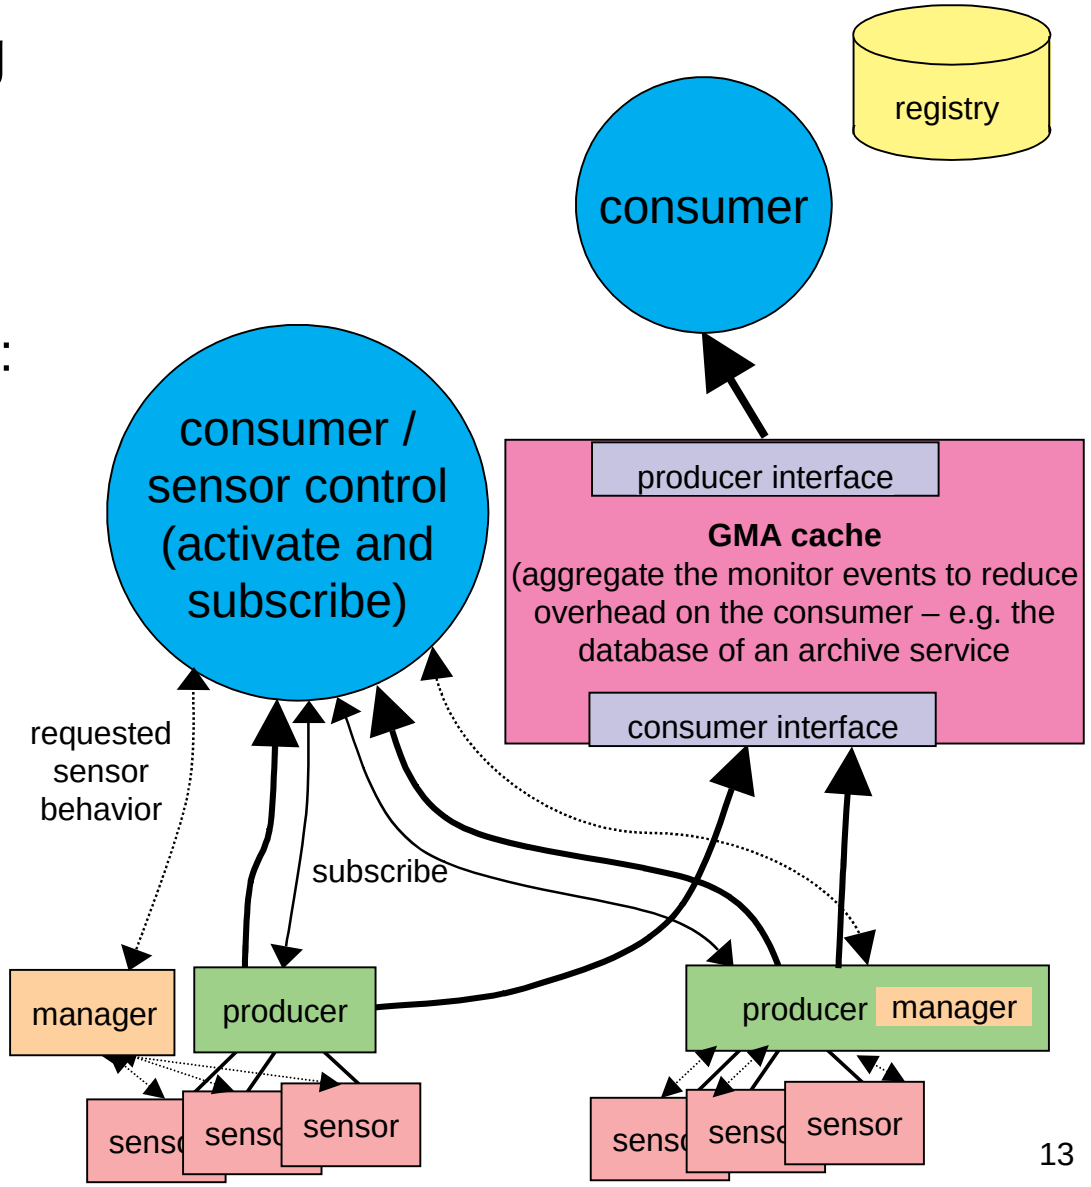

Status: operational

SMAPS: Security and Group Management

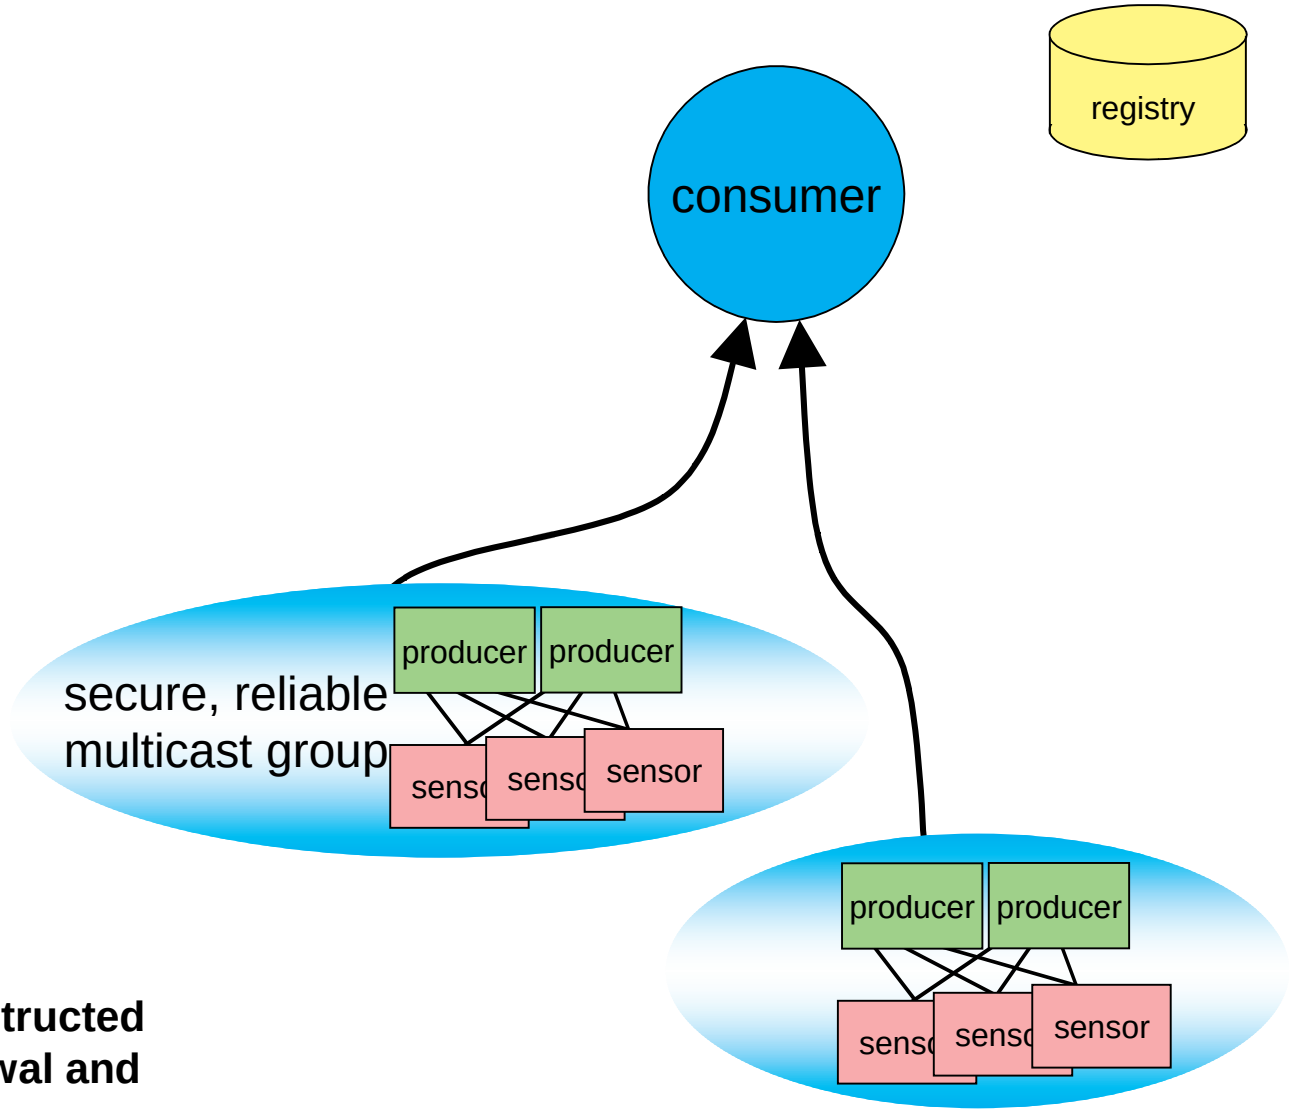

Status: prototype being constructed
Collaboration with Deb Agarwal and Olivier Chevassut

SMAPS: Autonomic (self-healing) and Self Organizing Sensor Nets

- Using Peer-to-Peer services provides scalable discovery of all available registries containing specific data types
- Combining P2P with reliable multicast provides:
  - P2P with security
  - Very rapid location of any available rendezvous point
    - § e.g. sensors can select among any available, compatible producers ⇒ self-organization (e.g. sensors that are isolated from original producer can rapidly locate any suitable producer and reconnect to it
    - § if all original producers are lost, then setting up an authorized producer on any system on the remaining network will reestablish a rendezvous point, and the sensors will automatically start reconnecting
    - § Since producers can serve as a light-weight registry
  - resource scalability – as servers get loaded, just add more – the rendezvous protocol could select a target for new sensors based on load

Status: experiments be conducted and design options considered. Collaboration with Wolfgang Hoschek and Karlo Berket