Saga Working Group

- **Goals**
  - Derive requirements for Grid APIs from Astro services
  - Develop a bulleted list of requirements (both in scope of API, and in general properties of APIs)
- **International Virtual Observatory Alliance**
  - Developing standard services for:
    - Accessing catalogs
    - Accessing image archives
    - Event notification
    - Registration of services
- **Observe that the interfaces are dealing with increasing levels of sophistication**
  - Not simple access to a file. Instead are accessing collections, managing assertions on collections, managing events.

Intellectual Property Policy

- I acknowledge that participation in GGF8 is subject to the GGF Intellectual Property Policy.
- Intellectual Property Notices Note Well: All statements related to the activities of the GGF and addressed to the GGF are subject to all provisions of Section 17 of GFD-C.1 (.pdf), which grants to the GGF and its participants certain licenses and rights in such statements. Such statements include verbal statements in GGF meetings, as well as written and electronic communications made at any time or place, which are addressed to: the GGF plenary session,
- any GGF working group or portion thereof,
- the GFSG, or any member thereof on behalf of the GFSG,
- the GFAC, or any member thereof on behalf of the GFAC,
- any GGF mailing list, including any working group or research group list, or any other list functioning under GGF auspices,
- the GFD Editor or the GWD process
- Statements made outside of a GGF meeting, mailing list or other function, that are clearly not intended to be input to an GGF activity, group or function, are not subject to these provisions.
- Excerpt from Section 17 of GFD-C.1 Where the GFSG knows of rights, or claimed rights, the GGF secretariat shall attempt to obtain from the claimant of such rights, a written assurance that upon approval by the GFSG of the relevant GGF document(s), any party will be able to obtain the right to implement, use and distribute the technology or works when implementing, using or distributing technology based upon the specific specification(s) under openly specified, reasonable, non-discriminatory terms. The working group or research group proposing the use of the technology with respect to which the proprietary rights are claimed may assist the GGF secretariat in this effort. The results of this procedure shall not affect advancement of document, except that the GFSG may defer approval where a delay may facilitate the obtaining of such assurances. The results will, however, be recorded by the GGF Secretariat, and made available. The GFSG may also direct that a summary of the results be included in any GFD published containing the specification.

  GGF Intellectual Property Policies are adapted from the IETF Intellectual Property Policies that support the Internet Standards Process.

Grid Challenge

- **Applications are moving from manipulation of a file to manipulation of**
  - Consistency constraints
    - Compare data on storage versus records in collection for consistency
    - Checksum validation of files on storage
    - Synchronization of files
  - Data structures and metadata structures
    - DFDL standard is still being developed
    - Space-Time-Coordinate ontology for interpreting events
      - Specific to Astrophysics
  - Transport properties
    - Bulk operations
    - Firewall interaction
    - Parallel I/O streams

Which Software Level is in Control?

Application	Application	Application	Application	Application
Storage system	Grid service	Grid service	Workflow	Workflow
	Storage System	Data Grid	Grid Service	Digital Library
		Storage System	Data Grid	Grid Service
			Storage System	Data Grid
				Storage System

At which level should grid services be standardized to provide:

- Robust operation
- Minimal latency
- Best performance

What operations are required to specify the state information required for control at each level?

Example - Image Cutout Service

- **From an image in a sky survey archive, cutout a specified sub-image**
- **Three approaches:**
  - Grab entire file, move to the client where subset is generated
    - Highly inefficient for large 2-Gbyte sized files. Minutes to execute
  - Have client issue I/O commands to do partial file reads
    - Improved efficiency, but time dominated by network latency for multiple I/O requests
  - Execute remote procedure at the storage system - fastest
    - Eliminates most network latency
    - Examples are HDF5 parsing library, Datacutter filters, OpenDAP server
    - Requires knowledge of structure of file, protocol for defining cutout area of interest, and support for executing a specified procedure
      - Hope is that DFDL will provide the data structure characterization
    - Moving knowledge from the client into the data grid

Example - Mosaic Service

- **Two approaches that rely on Montage (IPAC/Caltech)**
  - Workflow based mosaic generation
    - Generate a list of files that will be analyzed
    - Generate workflow that maps output from a procedure to the next procedure
    - Distribute workflow over multiple compute servers
    - Move each file to the server where the processing is done locally
      - All data accesses become local file system accesses
    - Store results into a digital library
  - Parallel application mosaic generation
    - Move entire collection to the processing environment (Teragrid)
    - Generate list of files to be analyzed
    - Do data decomposition, distribute tasks to nodes in cluster
      - All data accesses become local file system accesses
    - Use MPI for collective operations between nodes
    - Store results into a digital library

2MASS All-Sky Mosaic

- SDSC collaborated with R. Williams, J. Good, B. Berriman, and R. Cutri on the creation of standard plates for the 2MASS sky survey. The Montage mosaic code was used to generate the standard plates. L. Brieger used an MPI Python module, MYMPI, to manage the alternately sequential and parallel steps of the Montage process so that many mosaics were computed in a single parallel job. This allowed us to alternate between two models of parallelism in a single job: that of many, sequential steps executing simultaneously for independent mosaics and that of a single MPI parallel job executing on many CPUs for a single mosaic. This approach allowed us to parallelize all steps of the mosaic process.
- The processing of the 2MASS sky survey was a major exercise in use of Teragrid resources. The total number of images processed was 4,121,440 files from an input archive of 8 TB that had been replicated to the Teragrid. However, since many of the images contribute to more than one mosaic, they are included in the processing of more than one plate. If one counts the multiplicity of the processing, this results in actually 6,275,494 files, 14 TB worth, being processed for the construction of all plates. The amount of processing was 100,000 CPU hours.

Are there fundamental grid operations?

- **Challenges:**
  - Network devices - firewalls, load levelers, private virtual networks
    - Desired operation may not work through network device
  - Collections
    - For large numbers of files, operations are on the collection
  - Desired services are composites of independent operations
    - Skyquery - integrate records from multiple sky catalogs, return result of join across catalogs as a single file
      - Requires knowledge of semantics used in catalogs
      - For IVOA, have a standard Uniform Content Descriptor for catalog columns
      - Query on a catalog maps from local columns to standard UCDs
  - Migrate from operation on a single file using structure knowledge embedded in the application to operation on a collection, with knowledge encapsulated in the community service

Generic Astronomy Operations

- **Based on**
  - Standard data format - FITS
  - Standard semantics - UCD
  - Standard ontology - Space-Time-Coordinate
- **Bit level manipulation**
  - Efficient in grid environment when processing time dominates latency Open/Close/Read/Write/Sync/Stat/ …
  - Note that synchronization operations needed for manipulating large collections
- **Consistency assertions**
  - Checksum creation and validation
  - FITS header validation - requires comparing extracted star locations with catalog, adjusting FITS header metadata for coordinate system
  - Validation of compliance of event metadata with IVOA Space-Time-Coordinate ontology
- **Catalog operations**
  - Query across catalogs using UCDs to define semantics
  - Order queries to minimize metadata returned

IVOA Services

- **SIAP - Simple Image Access protocol returning pointers to the requested information encapsulated in a VOTable format.**
  - Requires extraction of a subset of an image from a remote collection
  - SAGA / Grid File System / DFDL / GridFTP / Byte IO / OGSA naming
  - Grid Remote Procedure Call
- **SSAP - Simple Spectral Access Protocol for one-dimensional spectra, time series, and Spectral Energy Distributions. Returns pointers to the information encapsulated in a VOTable format.**
  - OGSA DAIS / DFDL / GridFTP / OGSA naming
  - Grid Remote Procedure Call
- **Cone-Search – catalog access**
  - OGSA DAIS / DFDL / GridFTP / OGSA naming / Grid Information Retrieval
- **VOSpace – management of shared data collections**
  - OGSA Naming / OGSA data replication / Grid File Systems /

IVOA Services

- **VOEvent – event notification service**
  - Transaction management / Information Dissemination
- VOEvent defines the content and meaning of a standard information packet for representing, transmitting, publishing and archiving the discovery of a transient celestial event, with the implication that timely follow-up is being requested. The objective is to motivate the observation of targets-of-opportunity, to drive robotic telescopes, to trigger archive searches, and to alert the community. VOEvent is focused on the reporting of photon events, but events mediated by disparate phenomena such as neutrinos, gravitational waves, and solar or atmospheric particle bursts may also be reported.
- Structured data is used, rather than natural language, so that automated systems can effectively interpret VOEvent packets. Each packet may contain one or more of the "who, what, where, when & how" of a detected event, but in addition, may contain a hypothesis (a "why") regarding the nature of the underlying physical cause of the event. Citations to previous VOEvents may be used to place each event in its correct context. Proper curation is encouraged throughout each event's life cycle from discovery through successive follow-ups.
- VOEvent packets gain persistent identifiers and are typically stored in databases reached via registries. VOEvent packets may therefore reference other packets in various ways. Subscribers, human or machine, receive immediate notification of events, based on previously defined criteria. Packets are required to be small and to be processed quickly. This standard does not define a transport layer or the design of clients, repositories, publishers or brokers; it does not cover policy issues such as who can publish, who can build a registry of events, who can subscribe to a particular registry, nor the intellectual property issues.

GGF Services-Are their operations expressible in SAGA?

- **Infrastructure Standards Groups**
  - Ipv6
  - Network Measurement
  - Data Transport
  - Grid High-Performance Networking
  - Network Measurement for Applications
- **Data Standards Groups**
  - Data Access and Integration Services
  - Grid File Systems
  - Data Format Description Language
  - GridFTP
  - Grid Storage Management
  - Information Dissemination
  - OGSA Data Replication Services
  - Transaction Management
  - OGSA Data
  - Byte IO
- **3Compute Standards Groups**
  - Grid Resource Allocation Agreement Protocol
  - Job Submission Description Language
  - Grid Scheduling Architecture
  - OGSA Basic Execution Services

GGF Services

- **Architecture Standards Groups**
  - Open Grid Services Architecture
  - Grid Protocol Architecture
  - OGSA Naming
- **Applications Standards Groups**
  - Grid Remote Procedure Call
  - Grid Information Retrieval
  - Distributed Resource management Application API
  - Simple API for Grid Applications
  - Grid Checkpoint Recovery
- **Management Standards Groups**
  - Application Contents Service
  - Configuration Description, Deployment, and Lifecycle Management
  - Grid Economic Services Architecture
  - OGSA Resource Usage Service
  - Usage Record
- **Security Standards Groups**
  - Open Grid Service Architecture Authorization
  - OGSA-P2P-Security
  - Firewall Issues
  - **Trusted Computing**

Missing Ingredients

- **Management of state information**
  - Are there mechanisms to access the state information generated by each grid service?
- **Can consistency assertions be made across operations?**
  - Extension of concept of directing grid management (parallel, serial) to statements on completeness, validation of integrity (checksum)
  - Example might be guaranteed retry
- **Can discipline specific operations be executed - equivalent of actors in workflow systems?**
- **How does this compare to the e-Minerals project in the UK for supporting computation on distributed data?**