

# Status of the Information Power Grid

August 22, 2000

William E. Johnston<sup>1</sup> and Leigh Ann Tanner<sup>2</sup>

## 1 Introduction

IPG is a “Grid” (see [1]), and as such provides the software service for building large-scale, dynamically constructed problem solving environments. These services also provide uniform access and management for geographically and organizationally dispersed resources, including high performance computing and data handling systems. See Figure 1

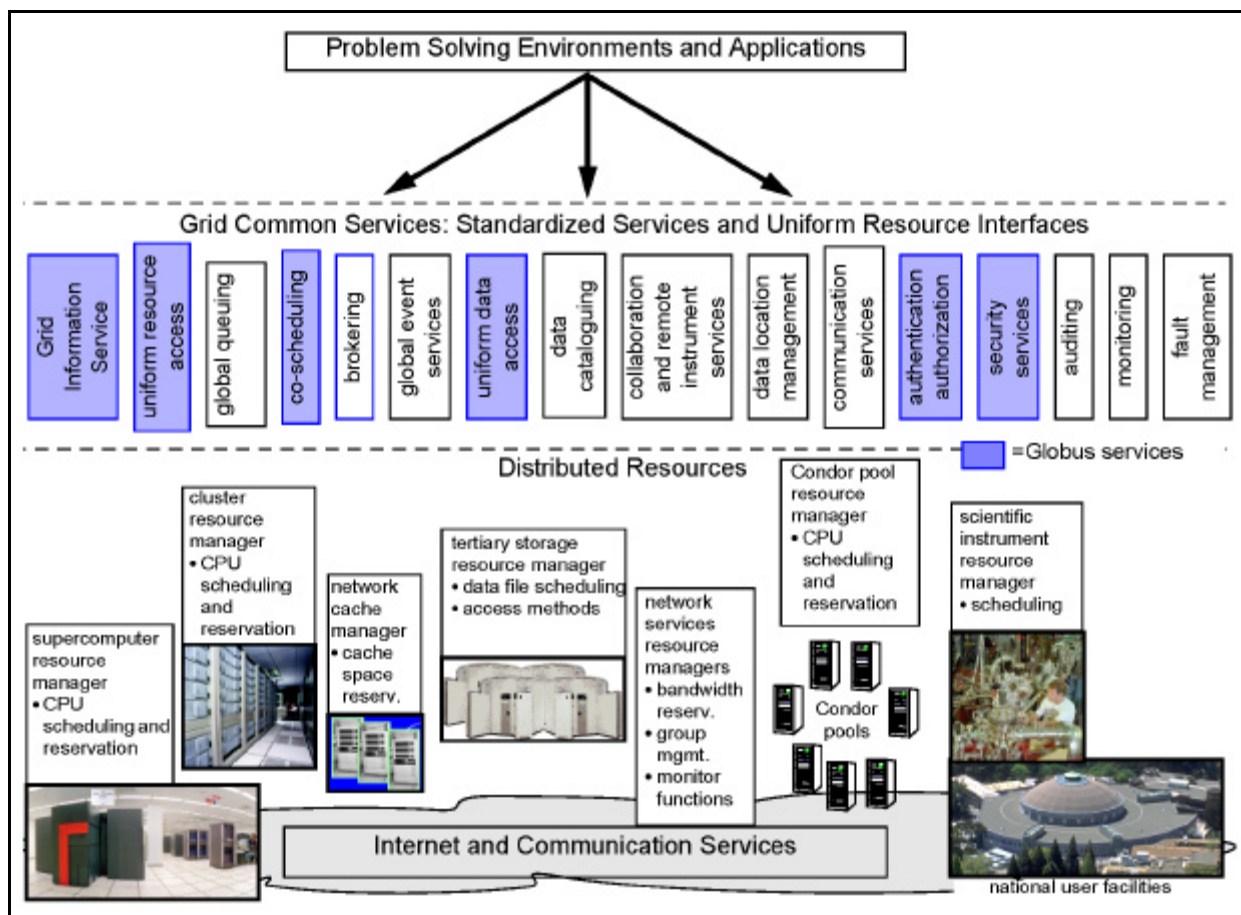


Figure 1 The Concept of a Grid

<sup>1</sup>Numerical Aerospace Simulation Division, NASA Ames Research Center and National Energy Research Scientific Computing Division, Lawrence Berkeley National Laboratory– wej@nas.nasa.gov

<sup>2</sup>Numerical Aerospace Simulation Division, NASA Ames Research Center – latanner@mail.arc.nasa.gov

The overall motivation for IPG is to enable the resource interactions that facilitate large-scale science and engineering such as aerospace systems design, astrophysics data analysis, climatology, large-scale remote instrument operation, etc.

The vision for IPG – a computing, data, and instrument Grid – is that it will provide significant new capabilities to scientists and engineers by facilitating *routine* construction of information based problem solving environments. That is, IPG will routinely – and easily, from the user’s point of view – facilitate applications such as:

- coupled, multidisciplinary simulations too large for single computing systems (e.g., multi-component turbomachine simulation – see [2])
- management of very large parameter space studies where thousands of low fidelity simulations explore, e.g., the aerodynamics of the next generation space shuttle in its many operating regimes (from Mach 27 at entry into the atmosphere to landing)
- use of widely distributed, federated data archives (e.g., simultaneous access to metrological, topological, aircraft performance, and flight path scheduling databases supporting a National Air Transportation Simulation system)
- coupling large-scale computing and data systems to scientific and engineering instruments so that real-time data analysis results can be used by the experimentalist in ways that allow direct interaction with the experiment (e.g. operating jet engines in test cells and aerodynamic studies of airframes in wind tunnels)
- single computational problems too large for any single system (e.g. extremely high resolution rotocraft aerodynamic calculations)

IPG development and deployment is addressing requirements obtained by analyzing a number of different application areas, in particular from the NASA Aero-Space Technology Enterprise. This analysis has focussed primarily on two types of users: The scientist / design engineer whose primary interest is problem solving (e.g., determining wing aerodynamic characteristics in many different operating environments), and whose primary interface to IPG will be through various sorts of problem solving frameworks. The second type of user is the tool designer: The computational scientists who convert physics and mathematics into code that can simulate the physical world. These two primary users of IPG have rather different requirements.

IPG will provide significant new capabilities to scientists and engineers by facilitating the solution of large-scale, complex, multi-institutional / multi-disciplinary, data and computational based problems using CPU, data storage, instrumentation, and human resources distributed across the NASA community. This entails technology goals of:

- independent, but consistent, tools and services that support various programming environments for building applications in widely distributed environments
- tools, services, and infrastructure for managing and aggregating dynamic, widely distributed collections of resources - CPUs, data storage / information systems, communications systems, real-time data sources and instruments, and human collaborators
- facilities for constructing collaborative, application oriented frameworks / problem solving environments across the NASA enterprise based on the IPG infrastructure and applications
- a common resource management approach that addresses, e.g., system management, user identification, resource allocations, accounting, security, etc.
- an operational Grid environment incorporating major computing and data resources at multiple NASA sites in order to provide an infrastructure capable of routinely addressing larger scale, more diverse, and more transient problems than is possible today

## 2 An Overall Model for Grids

Analysis of requirements of the work processes of some NASA user communities ([3]), as well as for remote instrument operation, and some anticipation of where the technology and problem solving needs are going in the future, leads to a characterization of the desired Grid functionality. This functionality may be represented as a hierarchically structured set of services and capabilities which are described below, and it's general structure is illustrated in Figure 2.

### 2.1 Problem Solving Environments, Supporting Toolkits, and High-Level Services

A number of services directly support building and using the Grid, e.g., by engineers or scientists. These include the toolkits for construction of application frameworks / problem solving environments (PSE) that integrate Grid services and applications into the “desktop” environment. For example, the graphical components (“widgets” / applets) for application user interfaces and control; the computer mediated, distributed human collaboration that support interface sharing and management; the tools that access the resource discovery and brokering services; tools for generalized workflow management services such as resource scheduling, and managing high throughput jobs, etc. E.g. SCIRun [4], ILab [5], and Condor-G [6].

An important interface for developers of Grid based applications is a “global shell,” which, in general, will support creating and managing widely distributed, rule-based workflows driven from a published / subscribed global event service. Data cataloguing and data archive access, security and access control are also essential components.

PSEs must also provide functionality for remote operation of laboratory / experiment / analytical instrument systems, remote visualization, and data-centric interfaces and tools that support multi-source data exploration, and services supporting human collaboration.

### 2.2 Programming Services

Tools and techniques are needed for building applications that run in Grid environments. These cover a wide spectrum of programming paradigms, and must operate in a multi-platform, heterogeneous computing environments. IPG, e.g., will require Globus support for Grid MPI [7] as well as Java bindings to Globus services. CORBA [8], Condor [9], Java/RMI [10], Legion [11], and perhaps DCOM [12] are all application oriented middleware systems that will have to interoperate with the Grid services in order to gain access to the resources managed by the Grid.

### 2.3 Grid Common Services

“Grid Common Services” refers to the basic services that provide uniform and location independent access and management of distributed resources. Much of the operational effort to run Grids is involved in maintaining these services.

Globus [13] has been chosen as the initial IPG runtime system and supplies the basic services to characterize and locate resources, initiate and monitor jobs, and provide secure authentication of users.

#### *Execution Management*

Several services are critical to managing the execution of application codes in the Grid. The first is resource discovery and brokering. By discovery we mean the ability to ask questions like: how

to find the set of objects (e.g. databases, CPUs, functional servers) with a given set of properties; how to select among many possible resources based on constraints such as allocation and scheduling; how to install a new object/service into the Grid; and how make new objects known as a Grid service. The second is execution queue management, which provides global views of CPU queues and their user-level management tools. The third category is distributed application management. The last category includes tools for generalized fault management mechanisms for applications, and for monitoring and supplying information to knowledge based recovery systems.

### ***Runtime***

Uniform naming and location transparent access must be provided to resources such as data objects, computations, instruments and networks. This, in turn requires uniform I/O mechanisms (e.g. read, write, seek) for all access protocols (e.g. http, ftp, nfs, Globus Access to Secondary Storage, etc.) and richer access and I/O mechanisms (e.g. “application level paging”) that are present in existing systems.

Data cataloguing and publishing services include the ability to automatically generate meta-data about data formats, and management of use conditions and access control. The ability to generate model-based abstractions for data access using extended XML and XMI [14] data models is also likely to be important in the complex and data rich environment of, e.g., aero-space design systems.

High-speed, wide area, access to tertiary storage systems will always be critical for the science and engineering applications that we are addressing, and we require data management services to provide global naming and uniform access. High-performance applications require high-speed access to data files, and the system must be able to stage, cache, and automatically manage the location of local, remote and cached copies of files. Management of data location with respect to, e.g., computing resources, will require large, dynamically managed, distributed “user-level” caches and “windows” on off-line data. Support for object-oriented data management systems will also be needed.

In IPG we are using SDSC’s Meta Data Catalogue / Storage Resource Broker (“MCAT/SRB”) [15] to provide widely distributed access to tertiary storage systems, independent of the nature of the underlying mass storage system implementation.

Services supporting collaboration and remote instrument control, such as secure, reliable group communication (“multicast”) are needed. In addition, application monitoring and application characterization, prediction, and analysis, will be important for both users and the managers of the Grid.

Other runtime services include checkpoint/restart mechanisms, access control, a global file system, and Grid communication libraries such as a network-aware MPI that supports security, reliable multicast and remote I/O. Monitoring services will include precision time event tagging for distributed, multi-component performance analysis, as well as generalized auditing of data file history and control flow tracking in distributed, multi-process simulations.

### ***Environment Management***

The key service that is used to manage the Grid environment is the “Grid Information Service.” This service – currently provided by Globus GIS (formerly MDS, see [16]) – maintains detailed

characteristics and state information about all resources, and will also need to provide access to dynamic performance information, information about current process state, user identities, allocations and accounting information.

Autonomous system management and fault management services provide another aspect of the environmental services.

## **2.4 Resource Management for Co-Scheduling and Reservation**

Grid services provide a uniform interface to resource management systems such as local batch queuing systems. For example, Globus' GRAM provides a uniform interface to local queuing systems such as LSF, NQE, PBS, etc.

One of the most challenging and well known Grid problems is that of scheduling scarce resources such as supercomputers and large instruments to work on a single problem. This is the issue of co-scheduling multiple resources. Any solution to this problem must have the agility to support transient problems based on systems built on-demand for limited periods of time, and in the case of remote instrument operation, at a scheduled time. CPU advance reservation scheduling and network bandwidth advance reservation are critical components to the co-scheduling services. In addition, tape marshaling in tertiary storage systems to support temporal reservations of tertiary storage system off line data and/or capacity is likely to be essential.

In several of these cases, most local resource managers do not currently provide the required functionality (e.g. advance reservation of CPUs) and new functionality must be added to the local resource managers.

## **2.5 Operations and System Administration**

To operate the Grid as a reliable, production environment is a challenging problem. Some of the identified issues include management tools for the Grid Information Service that provides global information about the configuration and state of the Grid; diagnostic tools so operations/systems staff can investigate remote problems, and; tools and common interfaces for system and user administration, accounting, auditing and job tracking. Verification suites, benchmarks, regression analysis tools for performance, reliability, and system sensitivity testing are essential parts of standard maintenance.

Implementing a persistent, managed Grid requires tools for deploying and managing the system software. In addition, tools for diagnostic analysis and distributed performance monitoring are required, as are accounting and auditing tools. Operational documentation and procedures are also essential to managing the Grid as a robust production service.

## **2.6 Access Control and Security**

Grids are inherently widely distributed, and frequently operate in open network environments. Security is an essential and integral part of the Grid environment.

The first requirement for establishing a workable authentication and security model for the Grid is to provide a single-sign-on authentication for all Grid resources based on cryptographic credentials that are maintained in the users desktop / PSE environment(s) or on one's person. In addition, end-to-end encrypted communication channels are needed in for many applications in

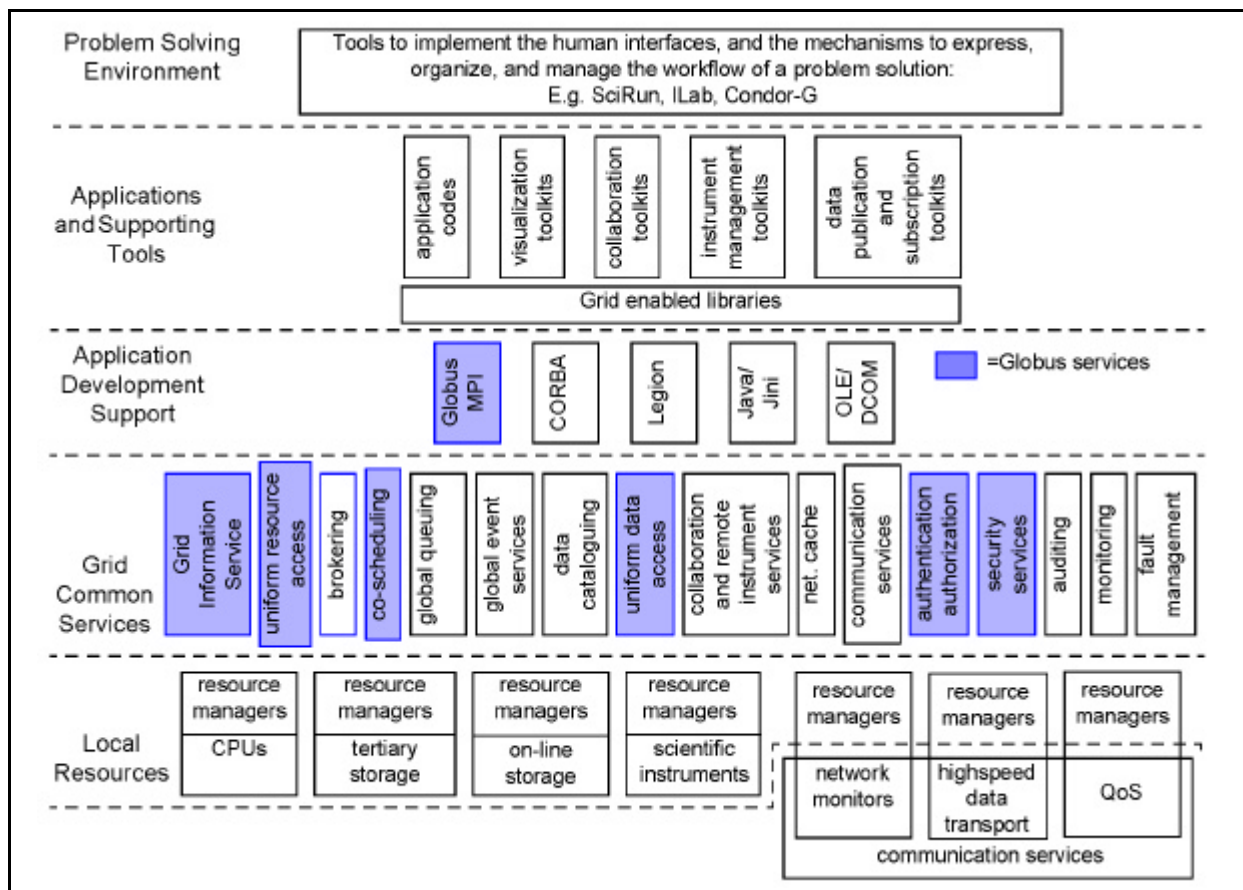
order to ensure data integrity and confidentiality. This is provided by X.509 identity certificates (see [17]) together with the Globus security services.

The second requirement is an authorization and access control model that provides for management of stakeholder rights (use-conditions) and trusted third parties to attest to corresponding user attributes. A policy-based access control mechanism that is based on use-conditions and user attributes is also a requirement. Several approaches are being investigated for providing these capabilities.

Security and infrastructure protection are, of course, essential requirements for the resource owners. This area includes things such as operating and system configuration standards for Grid platforms, design and coding standards for distributed applications, etc. Technologies such as IPSec and secure DNS to authenticate IP packet origin, secure router and switch management, etc. (see, e.g., [18]), and the plans are to deploy these in an IPG security testbed.

## 2.7 Grid Architecture: How do all these services fit together?

We envision the Grid as a layered set of services (see Figure 2) that manage the underlying resources, and middleware that supports different styles of usage (e.g. different programming paradigms and access methods).



**Figure 2 Software Architecture of a Grid**

However, the implementation is that of a continuum of hierarchically related, independent and interdependent services, each of which performs a specific function, and may rely on other Grid services to accomplish its function.

Further, the “layered” model should not obscure the fact that these “layers” are not just APIs, but usually a collection of functions and management systems that work in concert to provide the “service” at a given “layer.” The layering is not rigid, and “drill down” (e.g. code written for specific system architectures and capabilities) must be easily managed by the Grid services.

### **3 How is IPG being accomplished?**

Three main areas must be addressed in order to accomplish the goals of IPG:

- 1) new functionality and capability
- 2) an operational environment that encompasses significant resources
- 3) new services delivery model

The first area has already been discussed.

The second area, an operational system, is discussed below.

In the third area, Grids, such as IPG, effectively represents a new business model for operational organizations delivering large-scale computing and data resources. Grids require that these services be delivered in ways that allow them to be integrated with other widely distributed resources controlled, e.g., by the user community. This is a big change for, e.g., traditional supercomputer centers. Implementing this service delivery model requires two things: First, tools for production support, management, and maintenance, of integrated collections of widely distributed, multi-stakeholder resources must be identified, built, and provided to the systems and operations staffs. Second, organizational structures must be evolved that account for the fact that operating Grids is different than operating traditional supercomputer centers, and management and operation of this new shared responsibility service delivery environment must be explicitly addressed.

### **4 Progress Toward the Level 1 Milestones**

In this section we describe progress toward the first three IPG L1 milestones.

#### **4.1 Milestone 1 - 2QFY00**

*Milestone: “Distributed mass storage access. Description: Demonstrate software system to enable seamless access to catalogued archival data and information distributed throughout multiple NASA Centers and collaborator sites.”*

*Metric: Two archival data sites linked; 10 - 50 Mb/secs sustained transfer rate demonstrated.*

*Outcome: On-demand access to archived data and information; enhanced engineering and scientific collaboration*

One of the critical factors in making IPG a success will be its ability to manage distributed data. Data will live in different archives because the communities of experts that generate, analyze, and catalogue the data will typically maintain that data in a single location. Other groups that need to

use that data for secondary analysis or to drive simulations, will have to access the archives where the data is maintained.

Another aspect of data in the IPG environment is that simulations and analysis codes will execute wherever they find appropriate and available computing resources. This means that even after data is located at its primary archive, that data may have to be used (read for processing) at many other, typically remote locations. The effective location management of that data – caching, replication, local or remote access, etc. – will be a primary factor in determining the success or failure of the widely distributed systems built with the IPG services.

For both of these reasons, together with the fact that some relatively mature Grid data management software exists, the first milestone was designed to promote the integration of this software with IPG, and to demonstrate that it represents an effective tool for distributed data management.

The milestone was achieved by having a data mining application use IPG services to run on systems at NAS, NASA Ames, and access data at several remote sites. From [23]:

*This paper describes the development of a data mining system that is to operate on NASA's Information Power Grid (IPG), which is built using the Globus toolkit. The data mining system targets the mining of remotely sensed satellite data, which is characterized by its potentially large volume. A definition of data mining from a recent NASA data mining workshop states that "Data mining is the process by which information and knowledge are extracted from a potentially large volume of data using techniques that go beyond a simple search through the data." This paper represents a snap-shot into a project that is ongoing, presenting a scenario of grid-based mining, an architecture for a grid-based miner, and some initial experimental results.*







# Distributed Data Access in IPG



- Data access capabilities of IPG are demonstrated by parallel data mining
- 512 node SGI Origin at Ames uses IPG uniform interface data access tools to simultaneously mine hydrology data from four sites
  - SDSC
  - CalTech
  - GRC
  - Washington U.

Result from one agent



**Background:** A key function of Grids is to provide uniform access to widely distributed resources, including heterogeneous distributed archival data and information systems

**Objective:** Demonstrate software to enable seamless access to catalogued archival data and information that is distributed throughout multiple NASA Centers and collaborator sites

**Accomplishment:** Multiple data archive sites are accessed using a metadata catalogue and uniform access methods (SRB and GASS); high-speed remote data access is achieved

**Significance:** On-demand access to widely distributed archived data and information; enhanced engineering and scientific collaboration

**Future Plans:** Make this capability a permanent part of the IPG infrastructure and integrate with archival storage systems at ARC, GSFC, and JPL; incorporate the IPG/Globus security mechanisms to provide strong access control and secure remote data access




Figure 3

Highlights from IPG Milestone 1

## 4.2 Milestone 2 - 4QFY00

*Milestone: Distributed computing environment. Description: Prototype heterogeneous, distributed computing environment.*

*Metric: System tools and software provided; testbeds (2 or more classes of machines) at 3 NASA centers linked; application demonstration completed.*

*Outcome: Reduction in end-to-end turnaround time for aerospace simulation problems; peak performance, cost performance*

With this milestone the basic structure of IPG will be in place: An operational and persistent, “large-scale” prototype-production Information Power Grid providing access to computing, and data resources at NASA Ames, Glenn, and Langley.

This initial IPG consists of software services, operational services, and a base of hardware resources.

Current progress is reflected in the IPG Engineering Working Group tasks (see [3]): 30+ tasks have been identified as critical for the baseline system, and groups have been organized around the major task areas:

- ◆ IPG Prototype Startup Tasks
  - Globus deployed across Ames, GRC, and LaRC (Task 1.0)
  - IPG “common grid information base” (Task 2.0)
  - IPG X.509 Certification Authority and certificate server (Task 3.0)
  - “Global” queuing and user-level queue management capability on top of Globus (Task 4.0)
  - Computing resources for the initial IPG multi-center testbed (“CX”) (Task 5.0)
  - Networking for the IPG Testbed: Ames, GRC, LaRC (Task 6.0)
  - IPG Access for Archival and Published Data: SDSC’s Metadata Catalogue (MCAT) and the Storage Resource Broker (SRB) (Task 7.0)
  - Heterogeneity in the IPG testbed: Condor (Task 8.0)
  - Heterogeneity in the IPG testbed: High performance clusters (Task 9.0)
- ◆ IPG Operational Tasks
  - Security (Task 10.0)
  - IPG Information Base / MDS database maintenance (Task 11.0)
  - IPG/Globus system administration (Task 12.0)
  - Automatic monitoring of IPG components (Task 13.0)
  - Trouble ticket model (Task 14.0)
  - Condor support (Task 15.0)
  - CORBA support (Task 16.0)
  - Legion support (Task 17.0)
  - Documentation (Task 18.0)
  - User services (Task 19.0)
  - Account management (automated generation and maintenance mechanisms) (Task 20.0)
  - Globus with multiple MDS and PKI (Task 21.0)
  - Allocation Management and Accounting (Task 22.0)

- System testing: Verification suites, benchmarks, and reliability/sensitivity analysis for IPG (both static and dynamic) (Task 23.0)
- ◆ IPG Functionality Tasks
  - CORBA in the IPG environment. (Task 24.0)
  - Integration of Legion (Task 25.0)
  - CPU resource reservation (Task 26.0)
  - High Throughput Computing (Task 27.0)
  - Programming Services (Task 28.0)
  - Distributed debugging (Task 28.1)
  - Grid enabled visualization (Task 28.2)
- ◆ IPG Functionality Tasks
  - Network bandwidth reservation (Task 29.0)
  - High-Speed Network Testbed and Applications (Task 30.0)
- ◆ Characteristic Applications
  - OVERFLOW port & tune (Task 31.0)
  - NPSS port & tune (Task 32.0)
  - Parameter study (Task 33.0)
  - Heterogeneous testbed application: Condor (Task 34.0)
  - Heterogeneous testbed application: High performance clusters (Task 35.0)

### ***Software services tasks progress***

The basic software services provide for resource discovery, uniform computing and data resource access, job management, single sign-on, security, inter-process communication, and resource management.

Considering these services in terms of the organization given in Figure 2:

#### Problem Solving Environment

The ILab ([5]) high throughput job manager is functioning in IPG (task 27).


#### Applications and Supporting Tools

Several applications are being tested in the IPG environment. For example, the ILab manager was used to address the following problem (see the paper referenced at [5]):

*As an example, we chose the Overflow three-dimensional Navier-Stokes flow solver [11]. This modern CFD code implements the overset grid method (overlapping curvilinear grids exchange interpolated boundary information at each time-step).*

*The MPI parallel version of Overflow groups neighboring grids for solution onto individual processors. We applied the Overflow code to the solution of the flow field of the X38 Crew Return Vehicle (CRV), a NASA space vehicle designed as an astronaut escape pod potentially for the international space station.*

Figure 3 depicts the X38 CRV and several of the body-fitted curvilinear grids which define its surface. The geometry of the X38 CRV is defined by 13 curvilinear body-fitted grids and 115 off-body grids (which are strictly Cartesian in topology). This grid system contains



approximately 2.5 million points and since the Overflow code requires some 40 double-precision words of memory per grid point, this results in a total memory requirement per run of approximately 800 Megabytes per run.

[ILab created and ran] a 16 x 12 parameter study for two significant flow variables in a portion of the glide regime of the X38: Mach number (basically, the vehicle velocity) and Alpha (the “angle-of-attack”). This results in a two-dimensional parametric study consisting of 192 runs. Each run for the X38 vehicle requires four processors.

### Application Development Support

Grid MPI [7] and CORBA programming middleware systems are functioning in the IPG environment (tasks 24 and 28).

Distributed debugging tools have been developed and demonstrated, with the p2d2 portable, distributed debugger being ported to IPG (task 28.1).

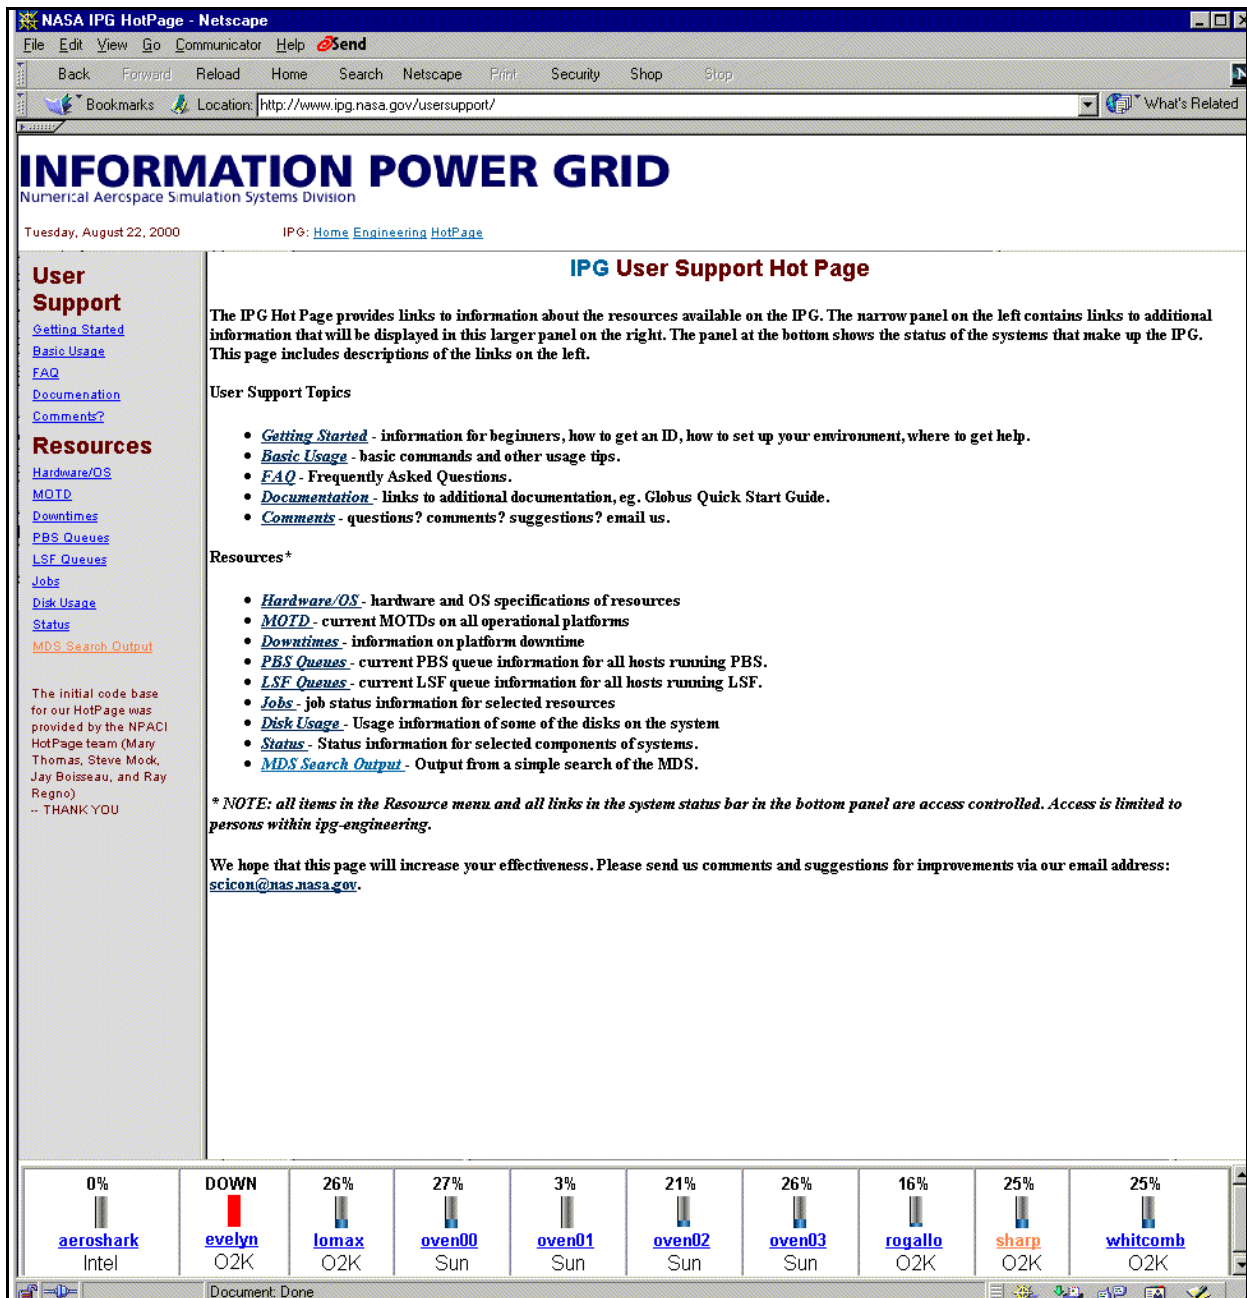
### Grid Common Services

Deployment of Globus ([13]) as the initial IPG runtime system is being accomplished in tasks 1, and 21. Globus 1.1.3 is up and operational at all sites.

Definition and implementation of a reliable, distributed Grid Information Service that characterizes all of the IPG resources – a critical service – is being accomplished in tasks 2, 11, and 21. The GIS namespace has been defined, and the LDAP servers that implement GIS are up and running at all sites.

Global management of CPU queues, and job tracking, and monitoring throughout IPG is being accomplished in tasks 4 and 13. Global queuing is under development and job tracking is implemented. See Figure 4 (“IPG User Support Hot Page” at <http://www.ipg.nasa.gov/usersupport>).

Public-key security infrastructure integration and deployment to support single sign-on using X.509 cryptographic identity certificates (see [17]) is being accomplished in tasks 3, 10, and 21.



**Figure 4 The IPG Hot Page**

The X.509 Certificate Authority is up and running and the Globus Security Infrastructure services are using the certificates in IPG.

Tertiary storage system metadata catalogue and uniform access system (based on MCAT/SRB - [15]) is being accomplished in task7, and was demonstrated as part of the first L1 milestone.

### Local Resource managers

The Portable Batch System is being extended to accommodate advance reservations (task 26).

Network bandwidth reservation is under development (task 29).

### ***Operational Services tasks progress***

Addressing the issue of a new service delivery model, the NAS Division at NASA Ames is identifying the new services that will be delivered by IPG, and is creating groups that will develop (as necessary), test, deploy, and support these services. In addition to new local organizational structure and local R&D, NAS is coordinating related activities at the NSF supercomputer centers [19], and at several universities, to provide various components of the new operational model.

Operational and system administration procedures for the distributed IPG are being addressed in tasks 12, 14, 15, 18, 19, 20, 22, and 23. Globus system administration procedures have been established and the operations staff is being trained. A distributed trouble ticket model has been defined and implemented. A Condor pool is installed and maintained. The Globus Quick Start Guide has been released. The User Services staff is trained and assisting IPG/Globus users. Account management is in the late stages of development. System verification suites have been developed and are being used.

### ***Hardware Resources***

Task 5 is accomplishing the identification and testing of heterogeneous computing and storage resources for inclusion in IPG. The baseline IPG prototype-production system (Figure 5) is in place and includes approximately 600 CPU nodes in half a dozen SGI Origin 2000s at four NASA sites, several workstation clusters, and 30-100 Terabytes (depending on the tape library configuration) of uniformly accessible mass storage.

A Condor pool of almost 200 workstations is functional and being used (task 15).

Task 6 is providing high speed, wide area network interconnects.

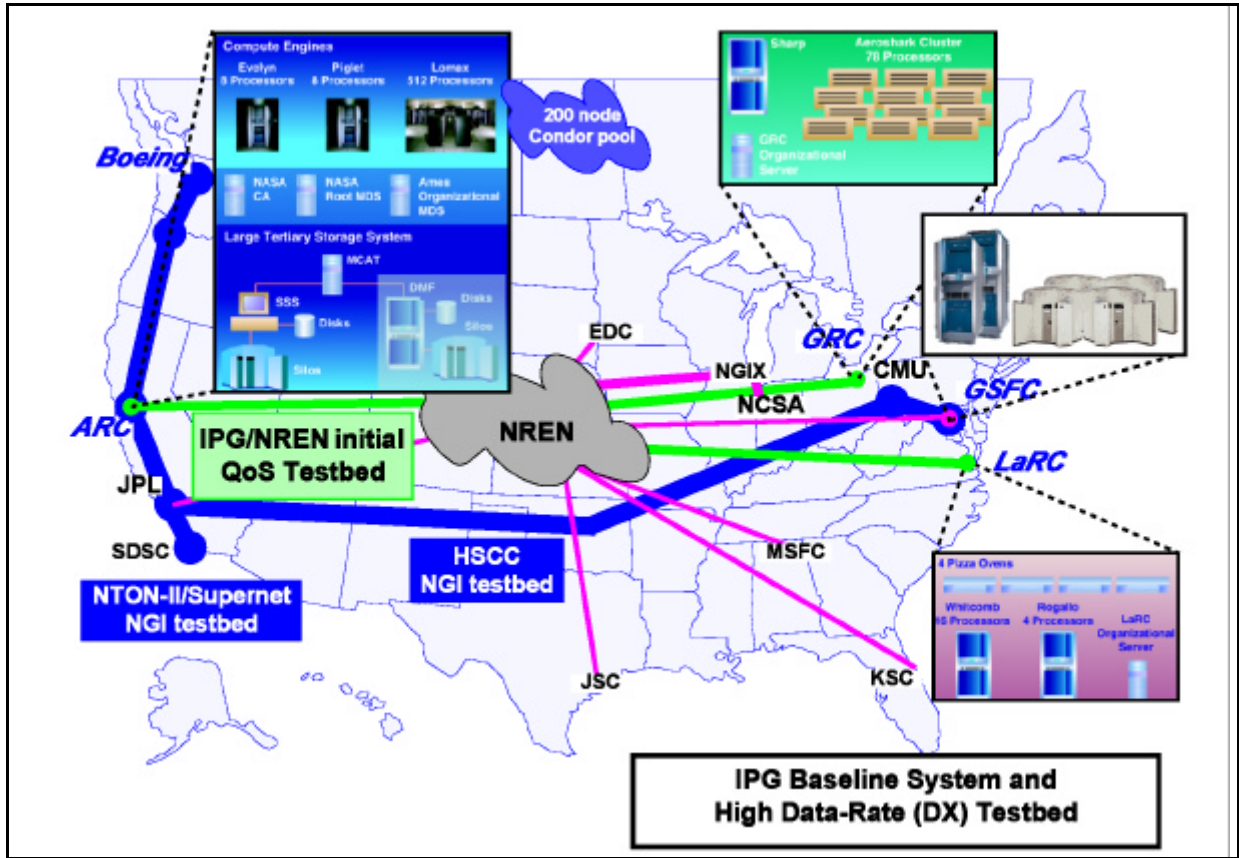


Figure 5

First Phase of NASA's Information Power Grid

### 4.3 Milestone 3 - 1QFY01

*Milestone: Integrate large-scale computing node into distributed environment. Description: Acquire and incorporate new large-scale computing systems and demonstrate seamless operations with heterogeneous distributed computing environment*

*Metric: Large-scale computing node installed and incorporated in computing environment; distributed resources capability to generate a solution data set for a complex aerospace research problem.*

*Outcome: Reduction in end-to-end turnaround time for aerospace design and simulation problems*

## 5 Acknowledgements

Almost everyone in the NAS division of the NASA Ames Research Center, numerous other people at the NASA Ames, Glenn, and Langley Research Centers, as well as many people involved with the NSF PACIs (especially Ian Foster, Argonne National Lab., Carl Kesselman, USC/ISI, Randy Butler, NCSA, and Reagan Moore, SDSC) have contributed to this work. Thanks also goes to Dennis Gannon (Indiana University), Bill Nitzberg (Veridian Systems, PBS Products Dept.), and Alex Woo.

The NASA Research and Education Network (NREN) has played a critical role in building the initial IPG.

We would also like to specifically thank Bill Feiereisen, NAS Division Chief, and while the NASA HPCC Program Manager the initiator of IPG. IPG is funded primarily by NASA's Aero-Space Enterprise, Information Technology (IT) program (<http://www.nas.nasa.gov/IT/overview.html>).

## 6 References

- [1] Foster, I., and C. Kesselman, eds., *The Grid: Blueprint for a New Computing Infrastructure*, edited by Ian Foster and Carl Kesselman. Morgan Kaufmann, Pub. August 1998. ISBN 1-55860-475-8. [http://www.mkp.com/books\\_catalog/1-55860-475-8.asp](http://www.mkp.com/books_catalog/1-55860-475-8.asp)
- [2] NPSS - see <http://hpcc.lerc.nasa.gov/grndchal.shtml>
- [3] "Information Power Grid." See [www.nas.nasa.gov/~wej/IPG](http://www.nas.nasa.gov/~wej/IPG) for project information, pointers, and the IPG implementation plan.
- [4] SCIRun is a scientific programming environment that allows the interactive construction, debugging and steering of large-scale scientific computations. SCIRun can be used for interactively:
  - Changing 2D and 3D geometry models (meshes).
  - Controlling and changing numerical simulation methods and parameters.
  - Performing scalar and vector field visualization.

SCIRun uses a visual programming dataflow system. SCIRun is extensible to a variety of applications and will work with third party modules written in Fortran, C, and C++. <http://www.cs.utah.edu/~sci/software/>



- [5] Yarrow, Maurice, Karen M. McCann, Rupak Biswas, "Advanced User Interface Approach for Complex Parameter Study Process Specification on the Information Power Grid." <http://www.nas.nasa.gov/Research/Reports/Techreports/2000/nas-00-009-abstract.html>
- [6] Livny, M., "Condor-G."
- [7] Foster, I., N. Karonis, "A Grid-Enabled MPI: Message Passing in Heterogeneous Distributed Computing Systems." Proc. 1998 SC Conference. Available at <http://www-fp.globus.org/documentation/papers.html>
- [8] Otte, R., P. Patrick, M. Roy, *Understanding CORBA*, Englewood Cliffs, NJ, Prentice Hall, 1996.
- [9] Livny, M, et al, "Condor." See <http://www.cs.wisc.edu/condor/>
- [10] Sun Microsystems, "Remote Method Invocation (RMI)." See <http://developer.java.sun.com/developer/technicalArticles//RMI/index.html> .
- [11] Grimshaw, A. S., W. A. Wulf, and the Legion team, "The Legion vision of a worldwide virtual computer", *Communications of the ACM*, 40(1):39-45, 1997.
- [12] Microsoft Corp., "DCOM Technical Overview." November 1996. See [http://msdn.microsoft.com/library/backgrnd/html/msdn\\_dcomtec.htm](http://msdn.microsoft.com/library/backgrnd/html/msdn_dcomtec.htm) .
- [13] Foster, I., C. Kesselman, Globus: A metacomputing infrastructure toolkit", *Int'l J. Supercomputing Applications*, 11(2);115-128, 1997. (Also see <http://www.globus.org>)
- [14] "XML Metadata Interchange" (XMI). See "XML News and Resources" <http://metalab.unc.edu/xml/>
- [15] Moore, R., et al, "Massive Data Analysis Systems," San Diego Supercomputer Center. See <http://www.sdsc.edu/MDAS>
- [16] Fitzgerald, S., I. Foster, C. Kesselman, G. von Laszewski, W. Smith, S. Tuecke, "A Directory Service for Configuring High-Performance Distributed Computations." Proc. 6th IEEE Symp. on High-Performance Distributed Computing, pg. 365-375, 1997. Available from <http://www.globus.org/documentation/papers.html> .
- [17] Public-Key certificate infrastructure ("PKI") provides the tools to create and manage digitally signed certificates. For identity authentication, a certification authority generates a certificate (most commonly an X.509 certificate) containing the name (usually X.500 distinguished name) of an entity (e.g. user) and that entity's public key. The CA then signs this "certificate" and publishes it (usually in an LDAP directory service). These are the basic components of PKI, and allow the entity to prove its identity, independent of location or system. For more information, see, e.g., RSA Lab's "Frequently Asked Questions About Today's Cryptography" <http://www.rsa.com/rsalabs/faq/>, *Computer Communications Security: Principles, Standards, Protocols, and Techniques*. W. Ford, Prentice-Hall, Englewood Cliffs, New Jersey, 07632, 1995, or *Applied Cryptography*, B. Schneier, John Wiley & Sons, 1996.
- [18] "Bridging the Gap from Networking Technologies to Applications." Workshop Co-sponsored by HPNAT & NRT (High Performance Network Applications Team & Networking Research Team of the Large Scale Networking (Next Generation Internet) Working Group). NASA Ames Research Center, Moffett Field, Mountain View CA. Moffett Training and Conference Center, August 10 - 11, 1999. To be published at [http://www.nren.nasa.gov/workshop\\_home.html](http://www.nren.nasa.gov/workshop_home.html) ("HPNAT/NRT Workshop")

- [19] The NSF PACIs are the Alliance/NCSA (<http://www.ncsa.uiuc.edu/>) and NPACI/SDSC (<http://www.npaci.edu/>).
- [20] “Globus Ubiquitous Supercomputing Testbed Organization” (GUSTO). At Supercomputing 1998, GUSTO linked around 40 sites, and provides over 2.5 TFLOPS of compute power, thereby representing one of the largest computational environments ever constructed at that time. See <http://www.globus.org/testbeds> .
- [21] “Distance Computing and Distributed Computing (DisCom2) Program.” See <http://www.cs.sandia.gov/discom> .
- [22] Grid Forum. The Grid Forum ([www.gridforum.org](http://www.gridforum.org)) is an informal consortium of institutions and individuals working on wide area computing and computational grids: the technologies that underlie such activities as the NCSA Alliance's National Technology Grid, NPACI's Metasystems efforts, NASA's Information Power Grid, DOE ASCI's DISCOM program, and other activities worldwide.
- [23] “Data Mining on NASA’s Information Power Grid,” Thomas H. Hinke and Jason Novotny. Available at [www.ipg.nasa.gov](http://www.ipg.nasa.gov)