Supporting Users of the NAS Facility at the NASA Ames Research Center

Chuck Niggley (niggley@nas.nasa.gov)
Ed Hook (hook@nas.nasa.gov)
Computer Sciences Corporation
NASA Advanced Supercom[puting Division
NASA Ames Research Center
Moffett field CA 94035-1000

The NAS Systems Division is part of the Information Sciences and Technology Directorate at NASA Ames Research Center, Moffett Field, California. The "NAS" in our name stands for NASA Advanced Supercomputing, and reflects our experience as pioneers in developing supercomputing technology and techniques to aid in the design of aerospace vehicles.

In 1996, NASA Administrator Daniel Goldin selected Ames as the NASA Center of Excellence in Information Technology. As the high-performance computing division of Ames, NAS is now involved in leading this technology sector for the entire agency.

Our participation in cutting edge projects such as the Information Power Grid (IPG) underscores our commitment to ensuring that the benefits of high performance computing as a critical technology are available to programs and projects throughout NASA.

As we enter the 21st century, we find ourselves at an exciting time in the development of high-speed computing. Recent advances in technology allow the realistic possibility of linking NASA's vast computing resources together to create an intelligent, large-scale, and adaptive computational model that is a great deal more than the sum of its parts.

As the high-speed computing component of the NASA Center of Excellence for Information Technology, NAS is leading the effort to develop, demonstrate, and implement this distributed heterogeneous computing capability for the Agency.

In the midst of this significant technological shift, we realize it is more important than ever to maintain our customer focus by continuing to work within the context of our two-fold mission:

1. To lead the country in the research and development of high performance computing for NASA's Programs and Missions by being the first to develop, implement, and integrate new high performance computing technologies into useful production systems.

2. To provide NASA and its customers with the most powerful, reliable, and usable production high performance computing systems in the country.

As you can tell, research and production are closely tied together in our mission statement. Either effort would be significantly weakened without the influence of the other. This ongoing exchange between the practical and the possible has been a key ingredient in our past successes and will be an essential element in the future development of the Information Power Grid.

At NAS, our work is driven by our mission, defined by project goals, and delivered on the completion of tasks. Most of the funding for the division's projects is provided through various NASA programs, offices, and projects. Here is a brief summary of current projects that shape our daily efforts.

The Information Power Grid is currently the largest project at NAS. The IPG represents an agency-wide effort to develop and deploy the next generation of computational capabilities to enable NASA missions. Nationally, the IPG will be an important instance of the emerging "grid" paradigm for high-speed computing applications. Primary funding of this project is provided by the Advanced Computing, Networking, and Storage element of the Information Technology Base R&T Program (IT/ACNS) and the Computational Aerosciences element of the High Performance Computing and Communication Program (HPCC/CAS).

Several aspects of the above mentioned funding sources are being reworked and will be combined into the new called Computing Information Communication Technology Program.

Consolidated Supercomputing Management Office (CoSMO). NAS is one of the primary organizations providing production computing cycles for NASA missions through CoSMO. We support the activities of about 400 CoSMO users on the CRAY C90 and SGI Origin systems.

Data Assimilation Office (DAO). As a key element of the NASA Earth Science Strategic Enterprise, the Data Assimilation Office at Goddard Space Flight Center is charged with the responsibility for processing huge amounts of atmospheric data accumulated from terrestrial satellites. Not surprisingly, the DAO has significant high-speed computing requirements in creating assimilated data sets from the raw data. NAS supports about 100 DAO users on several systems dedicated to this effort.

Internal Research and Development. Here, NAS works on high-risk, high-payoff research of interest to all four primary NASA Strategic Enterprises: Aero-Space Technology, Earth Science (formerly Mission to Planet Earth), Human Exploration and Development of Space and Space Science. Important examples of Internal R&D efforts are projects in the areas of astrobiology, nanotechnology, and device modeling. Research in these areas has the potential to revolutionize NASA missions in the 21st century.

NASA Research and Education Network (NREN). NREN is an experimental networking project whose goal is to provide both a high performance network application testbed and a network research testbed for the NASA community and its partners. NREN is the cornerstone of NASA's participation in the government-wide Next Generation Internet (NGI) project. Funding for this project is through the High Performance Computing and Communications (HPCC) Program.

Network Research for IT/ATED. The Analytical Tools and Environments for Design element of the IT Base R&T Program is involved in creating a next-generation advanced environment for the design of aerospace vehicles. The goals of this effort are aligned with the NASA Aero-Space Technology Enterprise strategic goal of reducing the design cycle time for aerospace vehicles by half.

Support and Infrastructure Activities. This project provides the infrastructure needed by all of the other projects at NAS, including: local area networks, workstations, desktop systems, distributed software, user support, publications and multimedia, and web work.

The NAS Systems Division, led by Division Chief Bill Feiereisen, is comprised of three branches: Applications, Engineering, and Research.

This basic organization reflects NAS's two-fold mission of researching and developing high performance computing technologies, and providing production high performance computing systems to NASA and its customers. Further, NAS has an underlying matrix organization that dovetails with the three branches and seven high-level projects encompassing all the division's work.

Each branch consists of smaller working groups that are structured around these projects, all funded and motivated by NASA missions related to space science, Earth science, aeronautics, and manned space exploration. Project managers work directly with each NASA customer to understand their requirements and determine our myriad division tasks.

The division's nearly 250 civil service and contractor employees support the Information Sciences and Technology Directorate at Ames Research Center, one of 11 NASA installations nationwide. Ames is designated as NASA's Center of Excellence for Information Technology.

When NAS was on the drawing boards in the late 1970s, the idea of supercomputing in a production environment was still very new. Supercomputers were rare, as was the expertise required to use them effectively When NAS opened its doors in the mid-'80s, one of its major goals was to make the use of supercomputers in aeronautics routine by making them readily available to scientists.

We accomplished this goal by pioneering many of the technologies and techniques that have become standards for integrating supercomputers into a production environment: networking to other computers with common operating systems, using scientific workstations to visualize datasets, and developing transparent methods of handling data transfer and storage.

By the early'90s, NAS had become the standard by which other supercomputer centers were evaluated. In fact many other successful supercomputer centers around the world adopted the architectural and operational paradigms that were
first implemented here.

A few highlights from our first 15 years:

First to put UNIX on supercomputers.

First to implement TCP/IP networking in a supercomputing environment.

First to link supercomputers and workstations together to distribute computation and visualization (what is now known as client/server)

.

Developed Aeronet, the first high-speed wide-area network (WAN) connecting supercomputing resources to remote customer sites.

Developed first batch queuing system for supercomputers, NQS, which became an industry standard.

Developed the first UNIX-based hierarchical mass storage system(MSS/NAStore).

Developed the NAS Parallel Benchmarks (NPB) which became the industry standard for objective evaluation of parallel computing architectures.

Developed PBS, the first batch queuing software for parallel and distributed systems.

Developed or contributed to the development of numerous visualization software applications which have become industry standards, including Plot3D and FAST (Flow Analysis Software Toolkit).

Co-created the NASA Metacenter, the first successful attempt to dynamically distribute real-user production workloads across Agency supercomputing resources at geographically distant locations.

In the few years since supercomputer use has become routine for researchers, the high-speed computing landscape has changed dramatically. In 1999, the power of the multimillion-dollar supercomputers of the 1980s is available in workstations costing just thousands of dollars. At the same time, advances in information processing, networking, and data storage have removed many of the physical and geographical barriers to sharing computing resources that existed just a short while ago.

Recognizing this world-wide change, we understand that the "supercomputer center" paradigm we once pioneered has fulfilled its role in the evolution of high-speed computing-and is rapidly becoming as impractical as the large vector computers it was once centered around. So, in our ongoing commitment to providing NASA missions and projects with unique and powerful computing resources, we are leading NASA's effort to create a new paradigm for high-speed computing in the next century-the Information Power Grid(IPG).

The mission of the IPG is to develop, demonstrate, and deliver innovative, distributed heterogeneous computing capabilities to enable NASA projects and missions

The Division has approximately 250 civil servants and contractors (employed by Advanced Management Technology Inc., Computer Sciences Corporation, and Raytheon). We support approximately 1100 scientific and engineering users across the country. The Division's main skills are in high performance computing and networking, distributed computing, parallel programming tools, nanotechnology and device modeling. Our main collaborators are Langley Research Center, Glenn Research Center, the National Science Foundation's NCSA Alliance and National Partnership for Advanced Computational Infrastructure, the Department of Defense High Performance Computing Modernization Program, Stanford University, SGI, and a host of others.

The key place where NAS system users can connect with helpful NAS technical support people, and get status on all our major systems is the “help desk”. Our User Services staff works 24 hours a day, 7 days a week to ensure system availability and assist researchers in using their allocations effectively. Our scientific consultants provide both production and development system users with specialized tools and techniques for optimizing code performance.

The primary function of the Control Room Analyst is to actively monitor the 23 super-computers. The computers are clustered, based on the user groups they support: CoSMO & HPCC, DAO and the experimental work conducted by TAG (Teraflop Application Group). In addition IPG users get accounts on several machines that are not located at Ames.

Additionally we monitor related peripheral systems such as firewall machines, web-servers, Large storage devices,fileservers and systems which provide administrative support for common applications that don’t necessarily run on the compute engines.

Shifts are meant to last 8 hours, and consist of from two to five personnel. We provide tertiary monitoring and support for four buildings, a special access area, the liquid and conventional cooling systems, electrical support in the case of power outages, and emergency notification/disaster recovery.

To monitor the systems we run up to 171 programs that allow us to monitor: internal and external network traffic flows, file system space requirements, the status of tapes in the 16 tape carousels (80,000+ tapes), and all the associated support processes.

When systems have problems analysts refer to on-line ISO.9002-approved guidelines that are agreed between the various second tier support agencies (there are 8 different ones of those) and the control room lead. If procedures call for coordination the Analysts then contact the POC and coordinate what actions can be taken (there are phones in the Control Room, as well as the floor, mobile phones and beepers, to ensure analysts are always in contact -- we monitor 2 buildings). The analysts must capture dump files or crash logs and mail them to SGI, or other appropriate agencies for further analysis and to alert them to potential hardware repair requirements.

We also serve as the first-line interface to users who have problems. We handle phone calls, emails, 4 different web-page-generated help requests, and those requests created using the Remedy trouble ticket tracking software. We use Remedy to track tickets, create metrics, and monitor computer up/down time. Our version of Remedy is the ancient client/server version, and we’re in the process of moving to the Web based version.

Once a ticket is created the Analysts must attempt to resolve the problem within four hours, or send it to second tier support (target is 90%). 75% of calls are resolved during the phone conversation. We use a series of questions to ascertain if the problem can be quickly fixed within our skill sets, or whether it involves steps that must be taken by the next level, and then create the Remedy ticket. Tier two support personnel are required to contact the user within 24 hours, if they are not able to resolve the problem in that period. Up-to-the-minute Information on the status of the systems, the amount of space in critical file systems, historical trends, and other key statistical data is provided to the users via the Web

An internal web-page exists for the Control Room analysts to track all current situational data concerning down systems, recent changes, issues, and other items that were previously kept in a written manual. This allows support personnel to peruse the logs from home or elsewhere in the building, search for trends in previous reports, and make comments or suggestions that other analysts can read and use, in a near real-time fashion (to include dropping in commands or chunks of code from other files or data they may have elsewhere).

The Help Desk personnel will open an average in excess of 800 tickets per month and about 1/3 of those are closed at the time of the call or soon afterwards by the control room staff. The other tickets are forwarded to one of the 15 other groups for resolution and closure.

The “scicon” group is staffed with folks that have many years of Fortran experience and have supported users on many different platforms and programming environments. We provide comprehensive customer support and deal with all questions concerning system usage, programming, data conversion, optimization, library usage, graphics packages, PBS, Message Passing techniques, system tools and utilities, etc. We perform extensive testing of any new software installation prior to its being made available for general use.

The System Administrators handle systems installation, provide operating system support and generally administer all of the computational and storage systems. Maintaining over 20 large systems with several versions of the operating systems, requires staff to be on-call 7 days a week 24 hours a day.

All of this infrastructure is needed to support a workload that typically sees more than 10000 jobs run each month. These jobs run the gamut from those needing 2 CPUs to those requesting 496 CPUs (which is currently the largest number that we can provide to a single job). These jobs can run for as little as a couple of seconds or they can grind away for as long as 35 hours. Jobs of the latter sort require special handling and we provide the needed mechanism on each of our large systems.

Further information about any of the topics discussed above can be found on the World Wide Web by starting at the URL *http://www.nas.nasa.gov.*