

NASA Advanced Supercomputing (NAS) User Services Group

John Pandori, Chris Hamilton and Charles Niggley, Computer Science Corporation,
Moffett Field, California, USA

ABSTRACT: *The primary function of the NASA Advanced Supercomputing (NAS) User Services Group is to continuously provide first level maintenance support to our 28 SGI Origin supercomputer systems and our new Cray SV1ex supercomputer. We also monitor more than 500 workstations, PCs & Macs, 24 file and web servers, 6 silo servers that manage 76 tape devices, and 25 RAID systems. In order to monitor these myriad systems, we make use of a variety of both commercially available tools and in-house products. These tools help us to set up and maintain user accounts across all of our systems, manage the scheduling of maintenance or special activities on individual systems and keep users informed of system status changes in a timely manner. In addition to first level maintenance support for the systems, we provide administrative help to approximately 1500 active users and staff, who serve in one of several different research groups, spread throughout the United States. This paper will describe the operations, tools and procedures that we use to accomplish all of these tasks and many others as well.*

Background

The NAS Systems Division is part of the Information Sciences and Technology Directorate at NASA Ames Research Center, Moffett Field, California. NAS has been one of the pioneers in developing supercomputing technology, and techniques to aid in the design of aerospace vehicles.

In 1996 Ames was selected as the NASA Center for 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. As the high-speed computing component of NASA our mission is to develop, demonstrate, and deliver innovative, distributed heterogeneous computing capabilities to enable NASA projects and missions.

The five primary programs supported by the User Services Group include: the Consolidated Supercomputing Management Office (CoSMO), Computing, Information, and Communications Technology (CICT), the Data Assimilation Office (DAO), the Earth Space Technology Office (ESTO), and NASA's Code-Y Division.

Operations - Monitoring Computers, Networks & the Environment

The first time you walk onto the main computer room floor at NAS it can be a fairly daunting experience. There are literally rows of super computers, RAID devices, power distribution units and air pushing units howling away. Everywhere you look there are little lights blinking at you on machines that all seem to have different name plates. How do we manage it? Where do we start? Who monitors what?

Our solution was to develop a chart that is posted in the workspace of each analyst, who we refer to as Control Room Analysts (CRA). The chart outlines which systems, programs and tasks a CRA is responsible for on each day of the week. We rotate the tasks each day so everyone has variety, and also to ensure that different eyes are looking at the myriad machines to discern potential problems. From a managers perspective the chart allows us to ensure all the important systems are included for coverage, and we can further refine the level of coverage that each machine receives; for example production systems are more important than test beds.

We provide 24 hour coverage utilizing a mixture of full and part time contractor personnel, a government worker and an intern. They serve in 8 1/2 hour shifts, which includes a meal break, and 1/2 overlap at shift change. The shift turnover is when CRAs are meant to pass on important, ongoing issues and problems. The shift maintains a minimum of two CRAs on the weekends, but can grow during the middle of the week to include as many as five analysts present. Because it is not always possible to answer phones we have an answering service, and we can take mobile phones when traveling between the four buildings that we support.

CRAs learn to perform their tasks through procedures documented in the Operations Manual that is stored in hard copy, as well as on-line in the NAS internal web pages. Our procedures and training are ISO 9001 compliant. Updates are made to both of these documents on a weekly basis to reflect current procedures promulgated in staff meetings, or as a result of coordination with other sections. We maintain the paper copy of the document to ensure it is available if all the systems go down, such as during a complete power outage. However the primary source for guidance is the online documentation. Many of the documented procedures include exact examples of frequently used

commands, and their locations, which are easily cut and pasted into active windows, thus reducing the possibility of committing an error.

Critical Processes

When critical processes fail on a system it ceases to function. Logically User Services put the bulk of our effort into monitoring these processes, to ensure the maximum uptime.

The most important processes involve network connectivity. If a computer cannot share files, pass messages, or transmit mail, it then cannot complete basic processes. Which will eventually cause the computer to bind up, and frequently cause it to crash.

Filesystems are also closely monitored by User Services personnel. Most NAS filesystems are NFS mounted (although we are moving toward SANS solutions), thus if a filesystem is not mounted correctly, computers cannot utilize the programs, storage space, or data located within them. Further as a filesystem begins to reach capacity system performance will begin to deteriorate. If the filesystem reaches it's capacity, it will fail, and all the processes dependent on that filesystem will be suspended in the process table until space is freed. Full system partitions such as *root (/)* or */var* will crash the computer system.

Next, User Services monitors system utilization. If there is an unusually high load level, then the system can appear sluggish. By noting and tracking this data one can develop reference points to ascertain the difference between normal operating loads and those associated with a true problem. Investigation may reveal user who is using resources outside the scope of prescribed usage, such as a user running an interactive job on a batch system. Recognizing how the system is being used and properly communicating to users the appropriate procedures for using each system is integral to maintaining a high rate of utilization.

Finally, we use a program originally developed at NAS to allow users to run their jobs called the Portable Batch System (PBS). This program also provides tools to track the number of CPU's utilized, jobs that are actually running, estimated runtime of each job, memory allocations, and various other parameters which are useful to the CRAs in assessing the overall ability of each compute engine. PBS uses three daemons. If one of the daemons for this program fails, then all of the jobs will go into a suspended state until PBS daemon(s) are successfully restarted.

Monitoring Tools

One of the primary tools we use at NAS is the Centralized Test Management System (CTMS). CTMS was developed by NAS programmers. CTMS consists of a series of scripts that run periodically on each of the super computers measuring: network connectivity, file

system availability, file system capacity, system load, and critical daemon processes. To monitor the results of CTMS processes the CRA launches a client window on their workstation. All CRAs, along with a good number of the technical staff run CTMS. The output from CTMS is gathered by a server and sent to all analysts who are running the client. Everyone gets all CTMS messages, so each CRA needs to filter through the messages and ascertain which apply to the systems for which they are responsible. Whenever a critical event occurs, such as a system crashing, a special alert window is launched onto the monitor of everyone running CTMS -- this ensures all the CRAs are informed of critical events immediately. Each messages received from CTMS requires a response acknowledging the event. The CRA responsible for the affected system is also responsible for acknowledging to the message and taking required actions. CTMS, along with reporting the event, also provides possible corrective actions to be taken.

Recently CTMS was augmented with a web based monitoring system called *status* (<http://www.nasa.gov/cgi-bin/nas/status>). *Status* provides information in both color graphics and textual windows, on subjects such as: whether a system is currently available on the network, system utilization, whether PBS is operating properly, how many jobs are running or awaiting run time (by job name), file system utilization, CPU utilization by machine, the Message of The Day (MOTD), scheduled outages, and the amount of compute hours each group has utilized. More poignantly, this information is available to the users on both the internal and external web pages, so users anywhere can easily gain vital data regarding system's status without having to run commands or contacting User Services.

CRA's also use scripts and utilities, which can be launched from their workstation, or from a compute engine. For instance, they can issue the *oper* command on a super computer, which will launch a utility displaying system messages, or they can run the *qstat* command for viewing the status of PBS jobs running on the super computer. CRAs have a variety of scripts that monitor system logs for critical errors, testing whether daemons are running, or actually submitting a job in PBS, as a means of testing whether it is working correctly. Most of these are fairly simple, and are redundant examples of the tests that have been automated in CTMS or status. However, they provide immediate feedback, and can be run from a CRA or user terminal, so as to troubleshoot or replicate problems that a user is observing.

To monitor the status of the LAN networks we use a stand-alone system running third party software called "What's Up Gold". This software package runs on an internal server and it can be configured to capture throughput information on selected circuits. We primarily use it determine whether communications are working between the super computers, the mass storage devices, the external long haul circuits, the internet, and

to a selected number of the web/file system servers. As we support over 500 PCs and workstations, we do not continuously monitor the connectivity of these generic devices, however, we can add systems if users complain of problems, in order to gather data. Users are the primary contact for PC or workstation connectivity problems. PC and workstation connectivity problems are also discovered daily at 1800 hours, when *nastruck* (an automated workstations backup program), begins to run on these systems. Through the *nastruck* backup program and various security monitoring programs, messages are sent regarding affected systems. Finally, we rely on the most important indicator, the user, to tell us when they experience connectivity trouble with external systems. Wide area connectivity to other NASA facilities is handled by NASA's Integrated Services Network (NISN). When wide area problems exist a CRA will contact the NISN support desk located at the Marshall Space Flight Center, and provide data which confirm connectivity problems exist (normally *traceroute* and *ping* results).

To monitor the environmental status of the computer floor and various peripheral computer rooms throughout the facility, a stand alone computer system called Data Acquisition System (DAS) is used. DAS is managed by a separate contractor that provides all of the various building engineering services throughout the installation. Our system monitors the amount of electricity flowing to the building, the presence of smoke, temperature levels on the computer floors, and then it reports this back to the base central operations cell. To augment this system, we make a visual inspection of the computer floor every two hours, and we check for errors on the system's consoles. Any error messages on the power distribution units (PDUs), indications of inordinately high temperatures on the air handlers, or alarms, are noted in *logit* (a recording tool which will be explained below) after hours, and then during regular work hours this is passed for action to the repair teams. Whenever we note items that might cause the systems to fail we immediately contact the service provider and the appropriate team is dispatched. We are also responsible for noting any safety issues after hours, and serve as the central command point in the event of a disaster we have a limited UPS capability.

The NAS facility, like many computer centers, is not open to the general public. We use a keypad and card operation to allow access to most important rooms, and the card grants holders access to the building(s) proper. The Card Access System (CAS) computer, allows us to monitor the status of all the important access points, including the primary entry points. Whenever these are open for more than 90 seconds we have to respond to determine the reason, and we have to check those personnel who don't have a permanent badge into the building.

Finally, on Silicon Graphics Incorporated (SGI) systems, the field engineers run the Embedded Support Partner (ESP) which is a proprietary software suite similar

to CTMS that can be configured to monitor important system processes, and then send emails to administrators based on preset criteria.

Recording Tools

The recording and tracking of system issues is handled through the *Remedy* ticket system, a commercial product of the Peregrin Corporation. NAS currently runs the client-server version, but we will soon field a web-based version. *Remedy* serves as the backbone of NAS reporting. Every problem and request is logged and tracked using one of several *Remedy* schemas or automated pages for entering data. There are several different schemas in use that allow the CRAs to accomplish a variety of tasks.

The OPERATIONS schema is used to track system events -- all problems affecting the operational status of NAS systems are recorded in this schema. Entries in this schema are used to create inputs to a Sybase database. Subsequently this data is then used to compute system operational metrics in a program called Down Log or (DLOG). It also maintains a database of system problems and resolutions that analysts can review.

The PAGER schema allows CRAs to send pager alerts and emails to other analysts regarding problems that require a rapid response. This ensures timely action on critical issues, and as the response time data is stored, it can be used to measure efficiency.

There is a schema to allow for system SCHEDULING within the *Remedy* package. This schema is used to schedule pre-planned system down times. The scheduling schema reports scheduled maintenance to users via the MOTD (message of the day) upon user login, and through a program called *schedule*. Scheduling can also be found through status web pages. It also interfaces with PBS, the batch job *scheduler*, on hosts to ensure that jobs are not started that would be killed by the dedicated time, thus preventing the loss of compute time.

In addition to *Remedy* we have developed a web based logging system called *logit*. *Logit* allows us to record much of the same data we would put in *Remedy* about systems being up/down, but in a sequential listing based on time and for all systems. This active file, and all the previously produced files, can then be shared out via the internal web to all of the analysts within NAS to review in near-real-time (to include those working from home). Further, it automatically sends emails to affected users when systems are down or up, along with informational statements regarding system status. *Logit* also has the capacity to send emails and pages to selected on-call analysts, thus providing a backup to *Remedy*. It is a more abbreviated process than *Remedy* which makes system information readily available to staff. It creates flat files which can be queried by staff without contacting User Services, thus providing historical or trend data. Previously, we used to produce this product as a simple

paper log, and one had to walk into the control room to review it.

To escalate problems to second level support, NAS maintains a web based listing of all sections' on-call personnel through Web Action Groups Section (WAGS). The CRA launches WAGS in the background. Whenever there is a reason to contact members of another group we determine the means each group uses to contact the on-call person (this allows each on-call person to specify the means they should be contacted; phone, cell phone or pager).

Major Problem Support

When a failure of a super computer occurs there are a variety of required actions the CRAs must accomplish. All of these are outlined in an online spread sheet called the System Failure Notification Checklist. It specifies which staff members must be contacted, the medium for contacting them, what hours the machine is supported and by which section, which users to inform, and how to create proper documentation of the incident for vendor support, metrics and trend analysis. When a system crashes the CRA refers to this document and utilizes several of the tools discussed previously.

The first task the CRA accomplishes, after a problem is identified, is determining the state of the computer. Frequently the computer has simply initiated an auto-reboot, and will be in the process of restarting. If the system is hung or there are some significant processes not working correctly, then WAGS is employed to determine who the second level support Point of Contact (POC) is, and an attempt to contact that individual is initiated.

As soon as there is reasonable certainty about what caused the computer problem, the CRAs will make an entry into *logit*. When we make the entry that a machine is down an automated message is sent to users. It is also published on the status page, further the icon for the system will appear to be down. If an estimated time for the system to recover is available, it is put in *logit* and it goes to the users as part of the message. When the system is back up and the critical daemons are all running correctly another entry will be made in *logit*, thus generating a system up message to the same user group, and it will appear in *status* that the computer is again working.

While the CRA is awaiting a response from the POC the CRA will attempt a select series of procedures outlined in the Operations Manual to collect information on what caused the system problem and restart the computer. If the CRA is successful in restarting the computer, the CRA will send an informational message to the POC (freeing up the POC to work on other problems). Next the CRA will create a problem ticket in the *Remedy* system.

If these steps fail to restart the computer, then the CRA and the rest of the staff continue to attempt to

contact the POC, who is responsible for either providing the CRA with additional instructions, or coming in and getting the machine back up. If there are hardware issues, and the system is supported by a vendor for either software or hardware, then the CRA will fill out the appropriate forms (either online, through an on-call help desk, or in writing in the field technicians on-site log books). The vendor's trouble ticket number will then be noted in both *logit*, and in the *Remedy* ticket created for this outage. This facilitates the ability to contact the vendor and cross reference their actions with those noted in the *Remedy* ticket. If hardware is the problem, then the CRAs will, in some cases, contact the vendor's on-call technician.

Additionally, a ticket is created in the *Remedy* operations *DLOG*. This log is used to record all planned and unplanned outages, for whatever reason. *DLOG* is used by metrics programs to determine the total up time, and the reason for all down times. This data is reviewed weekly by the High Speed Processing (HSP) Group manager. It allows examination of all outages with the various vendors, contractors and support staff, ensuring down time is kept to a minimum and the various component organizations work in a well coordinated fashion. The *DLOG* ticket isn't closed until the status of the machine is clarified. If a machine is running in a degraded mode (with say less than its full compliment of CPUs), this will be noted until it is functioning on all processors. *DLOG* tickets are also created when the computers are down for scheduled upgrades in software/hardware, baseline testing, or because of problems in the building (such as power or air conditioning). We created 4011 tickets in *DLOG* in 2001.

To summarize, whenever a super computer fails, the CRA will: create a *Remedy* ticket, contact selected members of the staff, email affected users, create comments in *logit* concerning when the system was down and then up, as well as comments concerning the CRAs actions, crash data if available will be stored and sent to the appropriate vendor (if the system is supported), and the event will be recorded in *DLOG*.

Storage System Problems

NAS has very large storage requirements to meet the needs of its customers. For long-term storage NAS has three mass storage systems which utilize the Data Migration Facility (DMF) The systems are attached to 16 large silos which contain upwards of 80,000 tapes and on a tape up to 60 Gb of uncompressed capacity. Additionally, half the silos and tapes are mirrored in a separate building for disaster recovery. The local file systems used by the super computers are actually very large RAID systems, which would be used for permanent storage on most systems.

When monitoring the mass storage systems we look at critical daemons (data migration & tape

management), and utilize various utilities and commands specific to these systems. *Oper* displays tape system messages, such as the status of tape mounts, import requests, and device connectivity errors. *Tmstat* is a tape status tool which shows whether a device is up, down, idle or in use, and whether tapes are mounted at all. *Tmgstat* displays time of the mount and session process number. We use *tmfrls* to release failed or stagnate mounts. Daemon logs display tape activities, and tape silo servers display hardware status of drives.

Using the tools outlined above we can configure devices up or down when they've failed, such as when the logs indicate that a particular device shows excessive error rates (indicating that service may be needed). We can release stagnate devices or mounts when a tape is no longer readable. When the system requests imports of tapes no longer contained in the silos (older data) CRAs must go to the tape vault, find the needed tape and load the tape into the appropriate tape silo.

During periodic floor checks we also discover problems with the RAID drives. We maintain drives for replacing failed drives (most of our RAID systems are RAID-5, which allow for one hot spare -- don't want to loose that second one...).

In every problem case identified above the CRA will note the events in logit, create a Remedy ticket, note the number of a drive/tape, changed/removed/replaced, location, serial number, create a ticket to the vendor (and all storage systems are vendor supported), and DLOG the event.

In those cases where the CRA is not able to resolve a problem rapidly it can become a serious impediment to operations of other systems. Therefore, the CRA will page the POC and/or the vendor. The vendor support contracts for storage are very comprehensive, because you can't run large jobs if you can't store the output. It is very important to get problems with storage system hardware identified early, so that replacement parts can be procured and on-site.

User Support Services

The control room serves as the first line of support to the users and the staff, providing services 24 hours a day. We receive problem calls or requests via phone, email, walk-ins, and fax. We utilize two distinct entities to provide support to users.

The primary service entity is the help desk, which handles the same sort of questions most help desks receive: delete, and archive accounts, restore data, change file permissions, share data, reset passwords, why won't this run, where is my email, how do I get an account, the printers broken, etc. To help CRAs answer questions and provide support to the users, we maintain an Operations Manual web page with sections on each important system at NAS. These pages are written by the systems administrators, and they're to be reviewed and updated quarterly. Most NAS sections also provide guides or

tutorials for users on how to accomplish their tasks on the NAS external web page. On the internal web pages NAS sections provide instructions for section personnel on how to administrate the machines, all of which is available to the help desk personnel.

The second entity is the User Interface Coordinator (UIC). The UIC creates all user accounts, all project groups identification numbers, and interfaces with all the NASA program managers to determine how many compute hours are distributed to each group. Because of the tremendous importance of ensuring the hours get correctly allocated, tracked and reported, and that the right people get into each group, we have assigned this function to an individual, rather than distributing it. Every year each of the 5 major programs go through a New Operational Period (NOP). During NOP all users must revalidate their accounts, or obtain one if they've never had one. Research Program Managers promulgate information about NOP and how to obtain compute hours throughout NASA and the various research agencies in government, higher education, and industry, by using the NAS web page. Once a project has been approved and an account request has been received, new project groups are created, users are added to them, and hours are then allocated to the groups. When a user starts a job in PBS it will cross check with our accounting program to determine how many hours the user's group is allocated, and how many they have collectively used. If a job exceeds project allocations the job will not be permitted to run. Ergo, tracking the hours used by one of the 600+ projects, spread across the 5 programs, and myriad users within each project is a bit sticky at times, and frequently contentious. The key is ensuring the proper account and project paperwork has been received and approved, then entered in the correct programs. Keeping managers well informed about the correct procedures, and the hours currently allocated is also done via the web.

All actions on questions or problems are recorded in Remedy, which was discussed previously in the recording tools section. When we create a Remedy ticket the system queries the CRA for the user's identification, and then it creates their complete name, phone number, and other important details. The CRA is meant to ask which machine the user is signed into, which machine or software they are trying to use or access, and specific questions regarding the nature of the problem they are encountering. This is then recorded in the ticket. Tickets generally fall into the categories of simple or complex. There are two schemas within the *Remedy* system used to record users' problems CONTROL ROOM, and HELP DESK. The CONTROL ROOM schema allows for the quick entry of simple requests. Many of the fields requiring analyst input are predefined. The HELP DESK schema is used for the more complex questions or problems. *Remedy* also includes a problem prioritization feature assisting analysts in determining problem severity or request urgency. For

example when the print spooler fails it will normally have a higher precedence than a single printer.

Each phone call, email, discussion in the hallway or fax, produces a separate ticket, which allows us to track it, send it between sections, and annotate all actions associated with its resolution. To electronically create a Remedy ticket, staff can: access the software directly, users send an email to one of several aliases, or they can click on a create button at a variety of locations throughout the NAS internal and external web pages, which will then develop an email that goes to Remedy. Most tickets are created by emails, sent either directly from the user, or by using the web pages, to create a ticket electronically. In addition, members of the other NAS sections can create tickets based on inquiries or emails they receive.

We handled 12802 tickets in 2001. NAS maintains internal web pages that report on efficiency at resolving questions and problems by time, and section, and this is part of the monthly reports to management. Various sections have different goals for resolving problems by category, and whether they have responded to the user, within preset guidelines. Creating and monitoring metrics on problem resolution is part of the ISO 9001 procedures.

Those questions or problems that we define as simple can be resolved while the user is on the phone. Examples include; resetting passwords, updating email addresses and answering simply user inquiries. Most of the time the CRA will attempt to resolve all requests this way, and then close the ticket. Handling a problem during the initial discussion reduces the need for call backs to conduct testing or gather additional information. However, as a CRA becomes more seasoned, they can gage those issues which they can sort out quickly, from those which require the user to capture error data for inclusion in a ticket.

When the problem is too complex to resolve during a phone call, such as a restore, the CRA will attempt to gather all of the data required to accomplish the task, and record it using the HELP DESK schema. If the problem is one the Control Room normally handles, CRAs will work the issue, in conjunction with the rest of their duties, noting what work they have accomplished or coordination they have done regarding the request. Often, CRAs on one shift get the question electronically, but work at different hours than the user who asked the question. Thus, follow on shifts can continue to work on the issue, and when necessary open tickets will be added in *logit* for shift turnover reporting. Once the problem is resolved, the user should receive a phone call and/or email notification to confirm that the problem has been resolved to the users satisfaction.

Whether we receive complex problems through emails or phone calls, the CRAs will frequently attempt to replicate the problem so as to isolate it. Problems generally occur; because the user may be using poor command syntax, user's workstation isn't functioning

correctly, the super computer isn't functioning correctly, a software product is not running successfully, or the users' account, or file permissions aren't set correctly.

When the CRA has determined the fault, we then have to determine whether we have the authority to fix it. For example security reserves the right to change certain files, and workstations personnel prefer to avoid remotely rebooting their systems without actually checking the hardware to ensure cables haven't become loose.

The HELP DESK schema in Remedy has a listing of all the other sections that are responsible for performing hardware/software maintenance, the systems they are responsible for, and frequent problem types. CRA's take the ticket, include that information they have learned from attempting to re-create the problem, their estimate of what is wrong, and then forward it to the appropriate section. When that section is done resolving the problem, they annotate the actions they took, coordinate with the user and close the ticket. If additional help from other sections is required then they simply forward the ticket onward. Some tickets have been used to identify hardware and software shortfalls that take months to resolve, such as bug reports or software upgrades. So using Remedy for the ticket system is a fairly robust tool.

Other User Support Tools

The creation of accounts by the UIC is accomplished using software called Login Account Maintenance System (LAMS). LAMS was developed by Boeing personnel working at NAS, and it allows for the secure creation and distribution of user accounts, group accounts, and passwords from a server within the NAS domain. The server runs on a separate computer and communicates to client software that is loaded on all systems administered by NAS. LAMS can be used to simultaneously create accounts on multiple computer systems, update passwords on multiple systems, change the groups that a user belongs to, change the shell a user uses, their email address and personal details.

Message Of The Day, MOTD, this is a banner page that is displayed upon login to all users. It contains usage instructions. Information on software updates and new features are also displayed on the MOTD. Planned system maintenance as well as file system status are shown on many NAS MOTDs.

Accounting tools include the *account_ytd* and *acct_query* commands. These tools will tell a user about how many hours their group has utilized during either a selected amount of time or during the whole Year-to-Date (YTD). In addition there is a GUI version of these commands which can be run by issuing the *aqua* command. *Aqua* is an older client/server program which utilizes *mosaic*, however, it is much more intuitive than simply issuing the commands. Finally, as previously mentioned, the *status* program now offers myriad ways to

run both of these commands and get a spread sheet style output, via the web.

Conclusion

There are certainly many ways to operate a user services organization, and we don't claim that we've cornered the market on ideas. Our software is an eclectic mixture that runs from curses to xml. We must work as a team with personnel from the other sections, so we're constantly holding informal discussions to assess our effectiveness and maintain the levels of commitment. Like all help desks we get a certain percentage of highly frustrated users, and yes our people do make mistakes. At the end of the day though, we've managed to make it work, and the vast group of users at the other end of the telephone line are happy.

About the Authors

John Pandori, works as a member of the 1st in the Control Room, and can be contacted at pandori@nas.nasa.gov

Chris Hamilton, serves as the Control Room lead, and can be contacted at hamilc@nas.nasa.gov

Charles Niggley, serves as the manage of the User Services Group, managing both the Control Room and the Scientific Consultants Team, he can be reached at niggley@nas.nasa.gov