NCRC Grid Accounting and Allocation

Frank Indiviglio
Geophysical Fluid Dynamics Laboratory
National Oceanographic and Atmospheric Administration
Princeton, NJ
frank.indiviglio@noaa.gov

Ron Bewtra
Geophysical Fluid Dynamics Laboratory
National Oceanographic and Atmospheric Administration
Princeton, NJ
ron.bewtra@noaa.gov

Abstract—**NOAA has deployed an accounting system for the purpose of coordinating HPC system usage between NOAA user centers and the NCRC located at Oak Ridge National Laboratory. This solution provides NOAA with a centralized location for reporting and management of allocations on all production resources located at the NCRC and at NOAA Laboratories. This paper describes the design, deployment, and details of the first year of production using this system. We shall also discuss the future plans for extending its deployment to other NOAA sites.**

The NCRC hosts Gaea, a leadership class system that is part of NOAA's Research and Development High Performance Computing System. In addition, NOAA has an HPC system that is now entering acceptance at their new NESCC facility in Fairmont, WV, a climate post-processing and archive solution in Princeton, NJ, and a system to support Hurricane Forecast Improvement Program in Boulder, CO. NOAA requires an allocation system to provide holistic management and centralized reporting for its users and stakeholders. For this deployment Moab and Gold from Adaptive Computing was chosen. This was used to create a multi-site scheduling and allocation architecture to support NOAA's scientific workflow.

This paper describes the successes and issues that the engineering teams encountered as they deployed the initial solution between NOAA's resources in Princeton and the NCRC, located at ORNL. Moab and Gold are being utilized to schedule, allocate, manage, and to facilitate data transfers. Challenges ranged from generating reports, client and server interactions with the scheduler, and administration issues with the allocation software. Also addressed in this paper is the next phase of deployment where the current system will be integrated with the remaining NOAA R&D sites.

Finally, comments about the general usability of the architecture will be provided from both a workflow developer and a user perspective. Issues in this area cover a range of items from general usability of reporting tools, validating the solution as new systems become available within the NCRC, implementing the allocation architecture, and integration issues of the scientific workflow into the scheduling architecture

Keywords: Gold, Moab

I. INTRODUCTION

The National Climate-Computing Research Center (NCRC) is a joint computing center between the National Oceanographic and Atmospheric Administration and Oak Ridge National Laboratory. The NCRC supports NOAA's Research and Development Climate and weather modeling initiatives across its entire R&D community. This community is spread out across the entire country, with concentrations at GFDL in Princeton, NJ, NCEP in Camp Springs, MD, and ESRL in Boulder, CO.

The NCRC hosts Gaea, which is comprised of a Cray 30912 core XT6, 79,812 core XE6. In addition to the XT and XE systems there are 8 external login nodes, 16 local data transfer nodes and 8 remote data transfer nodes. The external nodes are also schedulable resources that contribute to the execution of the site-wide workflow.

The NCRC is the first NOAA remote computing center to support multi-site distributed workflows for climate and weather modeling. This remote computing paradigm has been extended with the establishment of the National Environmental Security Computing Center (NESCC) in Fairmount, WV. The NESCC hosts Zeus a 27,360 processor SGI ICE cluster, with 8 login nodes and 2 data transfer nodes. Zeus, which began its production life in March of 2012, joins Gaea in supporting the climate and weather modeling research for NOAA's users across the country.

Together with computing assets located at each of the user labs, NOAA's R&D computing has been extended beyond the user's centers and across the country and created a distributed HPC infrastructure that allows for multi-site workflows.

A critical part of the distributed system was the scheduling and allocation system that supports the science

groups in running and prioritizing the experimentation streams across the shared R&D platforms. This scheduling and allocation infrastructure needed to provide the following features to support the vision of the shard R&D platform:

- support centralized reporting
- support monthly allocations
- easy transfer of time between projects
- provide the ability to adjust the priority of the workload based on allocation weight
- support standing transfers of time between groups
- support multiple systems
- support for multiple allocation groups
- flexible and accurate reporting

To support these requirements the NCRC selected Moab and Gold to supply the scheduling and allocation infrastructure. This infrastructure originally supported the grid environment between GFDL and NCRC, and will be latered expanded to the NESCC. In the future NOAA plans to centralize allocation reporting and perhaps the allocation process across the R&D HPC enterprise.

II. COMPUTING ENVOIRNMENT

The NCRC computing environment is a distributed environment that incorporates several centers across the country. At the center of this environment is Gaea. Gaea is housed at the NCRC at Oak Ridge National Laboratory. Gaea is comprised of two large compute partitions. The first of which is the c1ms, which went into production in January of 2011. The c1ms is 30,912 cores and is accompanied by 4 external login nodes, 8 ldtn(loacal data transfer) nodes and 4 rdtns(remote data transfer nodes). The c1ms is an XT6 and supported the bulk of NOAA's remote computing until 2012. In January of 2012, the c2 was added to the Gaea environment. The c2 is an XE6 and includes 79212 processors, and an additional 4 login nodes, 8 ldtns, and 4 rdtns. The systems are currently being treated as two separate machines due to their different architectures.

Due to the separation of these systems, they required separate allocations and scheduler instances. This was mainly done to allow the c2 to be brought into production and verified, while the c1ms remained the production system of the NCRC. Later in 2012, the c1ms will be removed from service and upgraded to match the c2. Once this upgrade is complete, the two systems will be joined into a single scheduler instance.

Initially to deal with these two systems, two allocation groups were created, so each science group where provide separate allocations for Gaea, one for c1ms and one for c2.

In addition to Gaea, the NCRC also utilizes post-production and analysis resources located at GFDL in Princeton, NJ. NCRC and GFDL instances of Moab are linked into a grid to facilitate job flow from the production resources at the NCRC to the post-production, archive, and analysis facilities at GFDL.

Post-processing usage is not allocated, however tracking it's use will be useful, and in the future tighter control on usage of Pan may be needed as special projects or laboratory's strategic goals change over time. In the initial installation, only jobs dispatched from Gaea where allocated, and local job on post-processing was not counted as an allocated job. A future goal for the post-processing system is to provide general usage and accounting statistics in addition to providing the current allocations on Gaea.

III. R&D ALLOCATIONS

To support the new initiative of remote shared computing, and to support the development of computational climate and weather science within NOAA, the Office of Atmospheric Research put in a place an allocation board to determine to use of the shared computational resources. The allocation board assigns resources for the major projects and scientific initiatives within the R&D labs and programs. As these allocations are distributed to the labs and programs, they may be further broken down into smaller groups within the allocation as projects are distributed amongst the lab's scientific groups. In all cases these minor allocations are further subdivided into working group or per scientist basis depending on how a science group is organized. While the allocation system only provides management of the major allocations, it's reporting infrastructure needed to provide a framework for the individual groups to have enough information to correctly manage their usage. Allocations are provided on a monthly basis, and may change depending on need or scientific priority. Allocations are further altered by trading between the labs and sciences groups, this trading is mostly dependent on scientific priority and needs of the laboratory.

Each laboratory handles allocations in a different manner. Some labs spread their allocations across small very specific tasks, while other have traditionally only provided major allocations to a science group and then allow the science group to self police there allocation on a monthly basis. In addition each lab has a unique understanding of what the system needed to provide in terms of allocation management. Several labs with in the R&D science organization expect that all users will need to select a project in each job, leaving the preferred allocation choice up to the user. They also expect the system to reject jobs if a user chooses an incorrect allocation group. Other groups expect the system to provide their jobs with the correct allocation at job submission. This cultural mixing within the new infrastructure provided an interesting challenge for allocation management to find an acceptable middle ground for all users on the NCRC system. This also provides us with a constant need for refinement and improvement of the allocation process. Given that allocation period is one

month, we can provide a flexible and dynamic platform to support the science needs of NOAA.

It is also important to note that not all of the R&D HPC infrastructure is an allocated resource. While Gaea and Zeus are certainly allocated resources, major supporting systems remain localized to a lab or program and are not managed via the allocation system. Providing the flexibility to include these systems in the future, or provide basic accounting to track usage statistics was also desirable. An example of this type of supporting system is the post processing and analysis system located at GFDL. This system is generally unallocated, and only restricts usage by job limits. However during times of heavy usage, or specialized usage, or finer grained control maybe needed to allow prioritized work to be completed.

IV. THE NCRC ALLOCAITON SYSTEM

The allocation system was installed as a service to support the shared infrastructure of the R&D HPC environment. This service was accessible to all systems within the NCRC workflow. All of the resources within the workflow are not allocated resources. This way any resource could have it's usage tracked even if an allocation had not been assigned to it.

This was to allow for allocation management for large-scale resources, like Gaea, and basic accounting if desired for smaller resources such as post processing and analysis systems.

The allocation system was designed to manage each of the sites from a central location. This design called for a high-availability implementation and strong tie in with grid scheduler elements of the system. The ultimate goal of the allocation system is to manage large groups of the computing resources from a central point. This was desired for several reasons, most desired was the ability to report on allocation usage from a single point in the infrastructure. This would allow administration and management a whole picture of system usage across NOAA R&D HPCS.

Initially all allocations and accounting where set up against the whole system. For Gaea, this meant that all partitions shared the same allocation. Each group was given their percentage based on usage across the compute pool, external services, initial post-processing, and data movement. Each user was assigned a default allocation group, which corresponded, to their primary science group. They did not need to specify an allocation for their jobs, as the allocation system selected their project for them. This allowed the users and framework developers to rely on the system to choose a default project, and still gives them the ability to select a secondary group if needed.

In early 2012 Gaea was upgraded with a second compute partition (c2). This required that allocations be created for each compute partition, as they were treated like separate systems. This presented an additional burden on the user base, as they would now have to select an allocation project for production runs carried out on each system.

Once the c1ms is upgraded these allocations would merge again so that user would not have to specify an allocation tag. Originally the upgrade was scheduled for early in 2012. However the schedule was extended that now places the upgrade in summer of 2012. To allow for an easier transition for the users, Adaptive suggested that each project be given several accounts to allow for the default project to remain the same as it was on c1ms. These subaccounts would be tied to specific hardware and charge against the hardware to which they where submitted. This design provided a convenient means to allow the user base to no longer worry about project specification. During this conversion and re-design a bug was discovered and prevented this from moving forward. The issue encountered was a confusion of the internal scheduler construct within each Moab instance. The internal scheduler construct is present in all instances of Moab as a reserved name for internal scheduling operations. However the bug we uncovered was attempting to charge the jobs with the internal instances name. When the allocation system received the information, it rejected the job, since there is no system named internal. This rejected jobs caused gaps in the reporting, and also caused a number of deferred jobs within the partition queues. A patch for this bug is awaiting installation; once the patch is installed we will move forward with creation of subaccounts tied to hardware.

Future improvements on the accounting and allocation system at NCRC will be continued through out this year. The largest of these will be the change from a strict debit system to a mixed debit/credit system. Currently all cpu hours are charged, even for support systems such as data movement. This was initially configured in this fashion as a means for monitoring use of the supporting systems and to keep their use along the same lines as the compute cluster. However over time, working with Adaptive, we have devised a plan to convert the external systems to a credit based tracking system, and the compute clusters to a debit based system.

This will allows flexibility to track usage of the external nodes and uphold allocated percentages on the system, while focusing the bulk of the allocations on the compute cluster. The way this will be configured is all supporting systems will start and 0 and work to a negative total, and all compute systems will start the month with a positive number of hours which reflects the groups percentage of the system and be charged downwards to zero as their allocation is used throughout the month.

Originally the allocation system was designed to collect information from the NCRC and GFDL, with the goal to collect information from other sites to provide NOAA management with a single reporting platform with the R&D HPCS. This plan called for a centralized instance of the allocation system to collect information for the whole enterprise. Using the NCRC as an initial grid configuration encompassing two sites, we implemented a single allocation system to capture data from the two compute clusters and supporting nodes, and to track the initial start of post-processing as the job stream moved from the assets located in Oak Ridge to the post-processing assets located at GFDL. Based on the initial results, Adaptive has suggested that the design be altered to include an allocation system for new

systems. To provide the needed reporting infrastructure an additional reporting instance of Gold would need to be installed to collect data from all of the sites and provide combined reporting for all of the HPC sites. This collection instance will automate collection via the underlying database software. This instance is currently planned to be installed later this year to provide this extension of the current reporting infrastructure. Currently reports are complied from each site and combined manually to provide an overall utilization report if the whole enterprise.

A second HPC system entered production at the NESCC in Fairmont, West Virginia. Currently it uses a stand-alone allocation system, and the long-term goal is to integrate it into the overall reporting infrastructure.

V. REPORTING INFRASTURUCTURE

A. Current Implmentation

Over the course of Gaea's lifetime, the reporting infrastructure has undergone several changes. From using built in gold tools, to several iterations of custom reports. Our experience with built in reporting tools within gold, prompted us to work with Adaptive to develop custom reports for our users.

The built-in tools for reporting provided a good format for our users and were generally well received. Initially reports were generated once every four hours. However over the course of the first six months we found that we needed to reduce the number of times the reporting ran per day. Slowly the report generation time went from minutes to over an hour. Over time using the built in tools, the query times while reporting was running slowed to a point that jobs trying to request allocation charges and reservations timed out. This, in turn, caused intermittent job deferment and slower start and end times of the jobs across all partitions. This trend led us, with Adaptive, to develop custom reports with the gold shell tools. This report provided basic information that the users required and ran in less time than the built-in tools.

Currently the reporting is run automatically once every six hours on the gold server. The results are copied to a central location where a simple script parses the results based on the users groups to display the correct group usage. Simple options are provided to display addition group information. Users have further built on these simple parsers to further parse general group information in to sub project reporting. This mechanism seems to provide the user base the ability to monitor their group's usage and provide additional information based on their individual needs. It also provides an acceptable interval for improved performance of the allocation system.

The current report generation script was developed by Adaptive Computing and provides a method of direct query of the allocation data. This allows us to speed reporting mechanism down from minutes to seconds. This has allowed us to improve job start times and reduce timeouts across the scheduling infrastructure.

Currently the report provides information about the usage over the month by user; we are currently reworking the report to add additional information to include user percentages and additional allocations used by each user. These additional allocations include windfall usage, which is usage over the guaranteed availability rate of the system, which is the usage of any additional non-allocated time on the system. This additional time is not always available however scientific groups what to be able to see how much of this time is being used by their groups when it is available.

Once a report is generated it is transferred to each site and is viewed from the parser script. This reduces the load on the allocation system and allows users to run reports as needed.

B. Future plans

Future reporting will include a more centralized instance of the allocation system that will provide the centralized reporting functionality requested by NOAA management. This centralization will provide another layer of reporting that will convert site localisms to enterprise terms. At each site project names have evolved organically to reflect original Unix groups and scientific group names. At a higher-level NOAA names each major project to a scientific initiative. The goal is to allow the centers to retain the group names that they have become accustomed to while allowing the reporting to reflect the NOAA major project names. Currently we plan to make this conversion will occur as data is imported into the database and allow for a revised report that provides allocation management from a high-level scientific project perspective. An example of this conversion can be seen from a group at GFDL. The Modeling Services Group is named the "f" group based on historical precedence. In moving from a GFDL centric computing environment it was requested that they keep this nomenclature and received the project gfdl_f. NOAA management refers to the f group's allocation as Performance and Portability of NOAA Coupled Models. Currently the conversion from the center's group name to the NOAA project name is done manually in the creation of the R&D HPCS wide report. Future reporting is planned to convert the site nomenclature to enterprise nomenclature, so that NOAA management will be able to provide spot reports as needed throughout the month.

C. Dual Running

When systems are brought into production, or major changes are introduced into the system GFDL scientists will run each job twice to ensure that bit-wise reproducibility is maintained across the experimentation. For the first 30 days of production each job run on the system is run twice and the resulting data is compared for reproducibility. The same process is used for code changes or major changes on the system. An example of a major change would be the change of a default complier or OS upgrade. To allow for tracking of these runs each allocation is cut in half and each group is assigned two allocations. For new system introduction the assignment of these allocations is straightforward.

The primary and secondary allocations are provided fro the 30 day period and the secondary allocation is removed after the dual run period ends. However there is a requirement to continue dual running for a longer duration for test jobs and Modeling Services purposes to test code changes and other software changes. Our initial dual running using the split allocations worked fairly well, but as we moved to a period of continued period of smaller percentage dual running, splitting the allocations did not seem to fit the process. To allow for a lower percentage of dual running and still retain the ability to track these dual runs as repeated runs a rework was required.

To allow for this we added a quality of service to the moab configuration and then reported on this qos to track the dual runs. This presented an issue with the current reporting infrastructure as gold tracks the qos information in a separate table structure than the job information that is currently used for reporting. This required the addition of a second report to allow users and group leads to view dual runs on the system. This report was implemented in a similar fashion as the normal allocation report, it is produced every six hours and then the data is transferred to each site and parsed by a custom script. This was put into production on May 1 2012. This new system of tracking dual runs has provided a more manageable way of tracking dual runs, although the addition of an extra report has lead to some user complaints and requests for the centralization of reports to reflect this new non-standard allocation method.

We will be incorporating this information in a future revision of the allocation report, which unfortunately will require a complete rewrite of the custom reporting now in place. Work is currently underway to allow for this reporting functionality and is desired across the R&D HPC enterprise.

This method of implementing qos for tagging allocation conditions has been so popular in the first week of production management is currently evaluating ways to further use the qos tagging for obtaining priority usage information, and possible expansion for use in tracking sub projects.

VI. Issues encountered

During our use of Gold, several large issues were encountered that took significant effort to overcome to provide better usage of the product. Most of these issues were encountered in the interaction of Moab and Gold. Initially these issues caused mis-charging of jobs and confusion of actual usage of the system.

A. Mis-Charging

Several instances of mis-charging have occurred within the NCRC. These issues have resulted in charging of queue time of canceled jobs and charging the incorrect cluster. This was attributed to a mis-communication of moab and gold. These types of issues are mainly attributed to mis-communications between Gold and Moab and required patching of Moab. While these types of issues have been few, there effects within the user community are felt far beyond the actual impact of the issue. We found that this particular type of issue caused the users to lose confidence in the product. Regaining confidence was a challenge, but over time is returning.

B. Reporting

Providing reporting has been an issue over the life of the system. This is due to several factors, but mostly can be attributed to issues with mis-charging, and difficulties we have had with developing custom tools and initial experiences with built-in tools. With NOAA's requirement of monthly allocations, we have worked on steering the reporting to better facilitate this. During this process we have had to reset deposit numbers at the beginning of the month, and compensate for end of month transactions. These have mostly been worked through but working through the reporting development took the bulk of the development time of the allocation system.

VII. Conclusion

Gold has provided NOAA with an adequate platform to report the allocation and utilization of the NCRC and other R&D systems. It however does not provide a robust reporting infrastructure without significant customization and development on the part of both Adaptive and NOAA. During its usage in the first years of production at the NCRC, it was grown to handle several large compute platforms.

We found that extension of standard reporting options required complete reworks of our custom reporting and the organization of the data within the gold product did not easily allow for additions of new data into the reporting infrastructure.

Despite the downsides of the experiences we went through, Gold has kept up with increased utilization and continues to evolve towards NOAA's needs.