



Fermi National Accelerator Laboratory

FERMILAB-Conf-87/139

[PN-320]

NULLJOB Product*

Nancy Hughart and David Ritchie
Fermilab Computing Department
Fermi National Accelerator Laboratory
P.O. Box 500, Batavia, Illinois 60510

May 1987

*Presented at the Fifth Conference on Real-Time Computer Applications in Nuclear, Particle, and Plasma Physics (IEEE), San Francisco, California, May 11-15, 1987.



Operated by Universities Research Association Inc. under contract with the United States Department of Energy

NULLJOB PRODUCT

Nancy Hughart, David Ritchie

Fermilab Computing Department
P.O. Box 500
Batavia, IL 60510

May 1987

This paper was presented at the 1987 IEEE
Conference on Real-Time Computer Applications in
Nuclear, Particle and Plasma Physics.

NULLJOB PRODUCT

Nancy Hughart, David Ritchie
Fermi National Accelerator Laboratory
P.O. Box 500
Batavia, Illinois 60510
(312) 840-4186, 840-3940

ABSTRACT

The ever increasing demand for more CPU cycles for data analysis on our Central VAX Cluster led us to investigate new ways to utilize more fully the resources that were available. A review of the experiment and software development VAX systems on site revealed many unused computing cycles. Furthermore, these systems were all connected by DECnet which would allow easy file transfer and remote batch job submission. A product was developed to allow jobs to be submitted on the Central VAX Cluster but actually to be run on one of the remote systems. The processing of the jobs was arranged, to the greatest extent possible, to be transparent to the user and to have minimal impact on both the Central VAX Cluster and remote systems.

IMPLEMENTATION

Overview

This product was implemented using command procedures and the facilities of the default batch queues on both the Central VAX Cluster and the remote systems. The actual user job is run on one of the remote systems in a batch queue having a CPU priority of one, the lowest on the system, so that all other activities of the system take precedence. A detached process, running on the Central VAX Cluster, maintains a database of information on the jobs as well as the state of the remote systems. It copies the user files to the remote systems, initiates processing on the remote systems, and retrieves user files from the remote systems to the Central VAX Cluster. Proxy logins between the systems allow transparent file transfer and remote batch job submission.

User Interface

The user interface consists of commands which allow the user to send a job to the NULLJOB product, monitor the status of his/her jobs only or all jobs using the NULLJOB product, and receive the completed job. A command is also available to abort the job while it is execution on a remote system. The status commands use the public domain product, FINGER, to return job status information. A FINGER.PLN file, if created and updated by the user's job while executing, may be used to display job specific status information.

Central VAX Cluster Manager and Job Flow

NULLMGR, a detached process running on the Central VAX Cluster, serves as the manager of the NULLJOB product. When the user issues a command (NULLJOB_SEND) to send a job to the NULLJOB product, information about the job is placed in a work area. When the NULLMGR awakes, it determines whether it has room for another job in its "outbound" holding area. If it does, it backs up the job's files to a save set for storage in the "outbound"

holding area until a remote system becomes free for a new job. It then determines the availability of the remote systems by examining remote system status database maintained and updated by the remote systems on the Central VAX Cluster.

When the NULLMGR finds a remote system available, it copies the user save set to the remote system and issues a "SUBMIT /REMOTE" command to submit a monitoring job on the remote system. The remote system executes this via its default batch queue. The monitoring job's first task is to restore the user save set to the NULLJOB_USR area on the remote system. At completion of this restore, the monitoring job submits the user job (now residing in the NULLJOB_USR area) to the NULLJOB_BATCH batch queue on the remote system. The monitoring job waits until this user job completes. It then makes a save set of the NULLJOB_USR area after the user job's execution and updates the status data on the manager VAX with the information that the job is done. The NULLMGR copies the save set back to the Central VAX Cluster where it is placed in the "inbound" holding area. When the user decides to receive the job (via NULLJOB_RECEIVE), the save set in the "inbound" holding area is restored to a sub-directory specified by the user.

Remote VAX

The remote system has a very ordinary, nonprivileged account (NULLJOB_USR) along with some procedures used by the NULLMGR on the Central VAX Cluster to process the user job on the remote system. The local system manager is provided with tools to stop and start the NULLJOB queue on his system according to the needs of that system on both a temporary and long term basis. Security is provided to the remote system by allowing only the NULLMGR on Central VAX Cluster proxy access to the nonprivileged account on the remote system. The NULLMGR can submit a job remotely which starts the user job after it has been copied to the remote system. However, users from the Central VAX Cluster do not have (and do not need) direct access to the remote system.

USER GUIDELINES

Appropriate Jobs

From a computational point of view, the type of job which is most appropriate for use with this product is one whose calculations result in few numerical results (due to limited disk space on the remote systems and tape access not being supported) and whose computation time is long compared to the transmission time to process the job to and from remote systems.

A limit of ten thousand blocks of disk space and two megabytes of virtual memory for the job are imposed by the resource limitations on the remote

VAX's. This is not a fundamental restriction on the product. It is essentially an estimate of the maximum resource requirements likely to be needed by the majority of NULLJOB users but which at the same time would not be so large as to impact the remote systems.

The job should require at most FORTRAN, DECnet, and VMS. In our particular case, all remote sites are provided with this software. The job should not be sensitive to specific versions of this software. Due to the large number of sites and their individual mission requirements, exactly the same versions of the software can not be guaranteed. Generally, versions are within one or two minor revisions of one another and the variation does not cause a problem. A possible approach for activities finding this restriction awkward is to do part of the processing on the Manager VAX and to submit only the less version-sensitive executable to the NULLJOB product.

CPU Time Available

This product is useful for CPU-bound jobs which do not have time critical project deadlines. The user is not guaranteed that his/her job will run to completion or that the job will receive a certain amount of CPU time during a specified amount of elapsed time. In particular, a remote job may be aborted or suspended by the local system manager of the remote system as required. Since all job files are returned to the user regardless of how the job terminated, useful computation may occur, provided the job checkpoints itself frequently.

All user jobs are submitted to the NULLJOB_BATCH queue with the "/RESTART" qualifier so that the job will continue after a shutdown or crash. The user job may check the \$RESTART symbol to determine if the job has been restarted and take appropriate action, such as skipping once-only initialization, etc.

DECnet Access

Jobs should not be dependent upon the network links being up. Although no guarantees of long-term network accessibility are made, as with any VAX VMS job, files can be accessed over DECnet. Such access should be relatively short and the job should make provision for the possibility of the network being down.

SYSTEM RESOURCES

Central VAX Cluster

The Central VAX Cluster resources which are required to support NULLJOB are a process slot for the NULLMGR process and sufficient disk space for the backup save sets of user jobs. Currently, user jobs are limited to ten thousand blocks. Ten remote systems run the NULLJOB product. Thus one hundred thousand blocks of disk space has been allocated for NULLJOB. Also, the presence of a default batch queue on the Manager VAX is assumed by the product since the NULLMGR submits batch jobs during the processing of the user job. In particular, this is done for copying files to and from the remote systems and for backing up and restoring user files.

Remote Systems

A NULLJOB_BATCH queue, where the NULLJOB user job runs, is needed on the remote system. It has CPU priority of one, the lowest on the system, so that it does not compete with interactive or normal batch jobs. A default batch queue, SYS\$BATCH, is assumed since this is the queue which is supported across DECnet by the "SUBMIT/REMOTE" command used by the Central VAX Cluster NULLMGR. This queue is used for control jobs such as restoring save sets, backing up completed files to a save set, and cleaning up after jobs have finished. Twenty thousand blocks of disk space (ten thousand blocks for the save set which is copied to/from the remote system and ten thousand blocks for the restored files) and two megabytes of memory are required on the remote system. At the discretion of the remote system's manager, the impact of the NULLJOB product on the remote system may be reduced by stopping the NULLJOB_BATCH queue during the busy hours. The batch job is then swapped out of main memory if the system needs that memory for normal use. When the busy hours are over, the queue may again be started. Batch jobs to perform this start/stop function automatically at particular times are provided as part of the NULLJOB command files on the remote systems.

RESTRICTIONS AND LIMITATIONS

DECnet

A reliable DECnet link is required. However, some effort has been made to insure that the product continues despite network errors and outage conditions.

Failure Recovery and Feedback

An important part of this product was the provision of adequate failure recovery. Effort was made to insure that the systems would recover from various sorts of failures. For example, all user files are returned so that an aborted job can be salvaged if the user has checkpointed the job. The user batch job is submitted with the "/RESTART" qualifier so that the job will continue after a shutdown or crash. Mutex files are used to insure the integrity of the database on the Central VAX Cluster and checking is done to insure the success of DECnet file transfers.

Another important part of this product was the provision of sufficient job status feedback so that the user could be sure where his/her jobs were in the several systems. The status commands indicates whether the submission is waiting for backup, waiting for a free remote system, in transmission to/from a remote system, in execution, or completed.

EXPERIENCE

History

A Design Specification for the NULLJOB product was written in April of 1986. The Design Specification estimated that two months would be required for the implementation of the product. An early version was working by July of 1986. However, additional failure recovery provisions and added functionality in areas such as user monitoring of jobs were needed. In November of 1986, three remote VAX systems were made available to a few users on a test basis. The

three systems were put into a production mode in December of 1986. By February of 1987, ten systems were being used by the NULLJOB product. These systems include two 11/785's, seven 11/780's, and one MicroVAX II. The number of user jobs being processed per month has gone from nine to forty one and from a few hours per month to over two thousand CPU hours per month on the ten systems. (This latter figure is approximately ten percent of the time available for similar batch jobs on the Central VAX Cluster.)

Operationally, once the initial bugs were sorted out (e.g. experience in the post-January time frame), the product has been very maintenance free. The Central VAX Cluster, the remote systems, and DECnet have all been brought up and down according to their usual maintenance and unusual crash cycles. Little consideration has been given to the NULLJOB product and the jobs have survived. One job in fact went through some thirteen crashes and automatic restarts before it finished its task and returned its results. Clearly this was a credit to user checkpointing of intermediate results.

User Job Profile

User jobs have varied in size from a few hundred blocks to seven thousand blocks and have required from a few hours to as many as thirteen CPU days for completion. Since the NULLJOB_BATCH queue is operational when the remote system is not normally busy, jobs generally obtain as much as ninety five percent of the CPU. Depending on the Central VAX Cluster batch load, this can be considerably more effective. For example, in one case on a remote VAX, a job obtained thirteen hours of CPU in thirteen clock hours, while on the Central VAX Cluster, the thirteen hours required forty seven clock hours. When coupled with the ability to submit jobs to the several CPU's, a high degree of "macro" parallel processing results.

The principal users of the NULLJOB product to date have been individuals doing Lattice Calculations and GEANT Monte Carlo calculations.

Remote System Profile

Generally during the week, jobs are processed only during the 8 p.m. to 8 a.m. interval. On weekends, jobs are processed twenty-four hours per day. During some periods of light remote system load, remote VAX system managers have run the product continuously.

The remote VAX's consist of software development systems, experiment data-taking systems, and analysis systems. The software development and analysis systems have provided a consistent and steady amount of time for the NULLJOB user. The experimental systems have, of course, provided substantial time when data-taking was not in progress. Tools are provided that allow for the easy removal of a system from the known list of available systems when the experiment begins taking data.

FUTURE PLANS

Accounting

Accounting is currently done on a limited basis using the log created by the batch job on the remote system for CPU statistics. Accounting information on jobs which do not finish execution due to a system shutdown is not easily associated with a specific job. This is due to the inability to distinguish one job from another because of the generic nature of the jobs as seen by the remote system and also the nonprivileged nature of the job running on the remote system. However, this information is available. We expect that a future release of the NULLJOB product will obtain more complete CPU statistics on each job. It should be noted that no accounting information is retained by VMS on batch jobs which are executing during a system crash.

Remote Cluster Capability

Remote systems which are part of a cluster are not supported at this time. However, the introduction of Local Area Clusters by DEC increases the desire for running NULLJOB on all such remote systems. The product currently requires a default batch queue that is specific to a remote systems. Further, it assumes that the queue in which the user job executes is always called NULLJOB_BATCH. These two aspects do not work on a cluster where generally the default batch queue is cluster wide and where only one NULLJOB_BATCH queue may exist. A possible solution is to make the command procedures more intelligent and to incorporate the node names into the queue names used by the user job.

Specific System Configurations

The variety of system configurations and software uses as well as different user requirements suggests that some effort be made to allow users to specify parameters which could be matched to the system resources available. For example, a feature which allowed users to send their jobs only to systems with large virtual address space capability would be a useful addition.

CONCLUSION

The NULLJOB product extends the computing resources available to users of the Central VAX Cluster by utilizing the null time on remote systems. For those jobs appropriate to the use of the NULLJOB product, CPU time not otherwise available may be found for execution. As more remote systems are brought online, computing resources for Central VAX Cluster users may be increased further through the use of the NULLJOB product.

ACKNOWLEDGMENTS

We would like to acknowledge the valuable aid in testing and development which our early NULLJOB users provides. In particular we would like to acknowledge the efforts of Hank Thacker and James Sexton of the Fermilab Theory Department, the cooperation of the many remote system managers, and the contributions of the Fermilab Data Acquisition Group. We should also mention that the multiple node software development formalism reported elsewhere at this conference was essential to the success of this product.