Preprint
UCRL-JC-147996

# SLURM: Simple Linux Utility for Resource Management

*Morris A. Jette, Andy B. Yoo, and Mark Grondona*

This article was submitted to
Workshop on Job Scheduling Strategies for Parallel Processing,
Seattle, WA, June 24, 2003

**April 3, 2003**




Approved for public release; further dissemination unlimited

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

This is a preprint of a paper intended for publication in a journal or proceedings. Since changes may be made before publication, this preprint is made available with the understanding that it will not be cited or reproduced without the permission of the author.

This report has been reproduced directly from the best available copy.

Available electronically at http://www.doc.gov/bridge

Available for a processing fee to U.S. Department of Energy
And its contractors in paper from
U.S. Department of Energy
Office of Scientific and Technical Information
P.O. Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Facsimile: (865) 576-5728
E-mail: reports@adonis.osti.gov

Available for the sale to the public from
U.S. Department of Commerce
National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Telephone: (800) 553-6847
Facsimile: (703) 605-6900
E-mail: orders@ntis.fedworld.gov
Online ordering: http://www.ntis.gov/ordering.htm

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
http://www.llnl.gov/tid/Library.html

# SLURM: Simple Linux Utility for Resource Management*

Morris A. Jette Andy B. Yoo Mark Grondona

Lawrence Livermore National Laboratory
Livermore, CA 94551
{jette1 | yoo2 | mgrondona}@llnl.gov

**Abstract**

A new cluster resource management system called Simple Linux Utility Resource Management (SLURM) is developed and presented in this paper. SLURM, initially developed for large Linux clusters at the Lawrence Livermore National Laboratory (LLNL), is a simple cluster manager that can scale to thousands of processors. SLURM is designed to be flexible and fault-tolerant and can be ported to other clusters of different size and architecture with minimal effort. We are certain that SLURM will benefit both users and system architects by providing them with a simple, robust, and highly scalable parallel job execution environment for their cluster system.

## 1 Introduction

Linux clusters, often constructed by using commodity off-the-shelf (COTS) componnets, have become increasingly populuar as a computing platform for parallel computation in recent years, mainly due to their ability to deliver high perfomance-cost ratio. Researchers have built and used small to medium size clusters for developing and debugging applicationsd [3, 15]. Now, continuous decrease in the price of the COTS parts in conjunction with the good scalability of the cluster architecture has made it feasible to economically build large-scale clusters with thousands of processors [17, 18].

An essential component that is needed to run user applications on a cluster is cluster management system. A cluster management system (or cluster manager) performs such crucial tasks as scheduling user jobs, monitoring machine and job ststus, launching user

*This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes. This work was performed under the auspices of the U. S. Department of Energy by the University of California, Lawrence Livermore National Laboratory under Contract No. W-7405-Eng-48. Document UCRL-TBD.

applications, and managing machine configuration, An ideal cluster manager should be simple, efficient, scalable, and fault-tolerant.

Unfortunately, there are no open-source or proprietary cluster management systems currently available, which satisfy these requirements while being immediately available for clusters of any architecture. A survey [10] has revealed that the existing cluster managers have poor scalability and hence they are not suitable for large clusters of thousands of processors [12, 9]. Many of these cluster managers also cannot be easily ported to other clusters since they are tied to specific components of the cluster that they are targeted to [20, 12, 9]. Furthermore, their cluster management functionality is usually provided as a part of a specific job scheduling system package. This mandates the use of the given scheduler just to manage a cluster, even though the scheduler does not necessarily meet the need of organization that hosts the cluster. Clear separation of the cluster management functionality from scheduling policy is desired.

This observation led us to set out to design a simple, highly scalable, and portable cluster management system that performs the core of cluster resource management functions. The result of this effort is Simple Linux Utility Resource Management (SLURM[1]). SLURM was developed with the following design goals:

- *Simplicity*: SLURM is simple enough to allow motivated end-users to understand its source code and add functionality. The authors will avoid the temptation to add features unless they are of general appeal.

- *Open Source*: SLURM is available to everyone and will remain free. Its source code is distributed under the GNU General Public License [8].

- *Portability*: SLURM is written in the C language, with a GNU *autoconf* configuration engine. While initially written for Linux, other UNIX-like operating systems should be easy porting targets. SLURM also supports a general purpose *plugin* mechanism, which permits a variety of different infrastructures to be easily supported. The SLURM configuration file specifies which set of plugin modules should be used.

- *Interconnect independence*: SLURM supports UDP/IP based communication as well as the Quadrics Elan3 and Myrinet interconnects. Adding support for other interconnects is straightforward and utilizes the plugin mechanism described above.

- *Scalability*: SLURM is designed for scalability to clusters of thousands of nodes. Jobs may specify their resource requirements in a variety of ways including requirements options and ranges, potentially permitting faster initiation than otherwise possible.

- *Robustness*: SLURM can handle a variety of failure modes without terminating workloads, including crashes of the node running the SLURM controller. User jobs may be configured to continue execution despite the failure of one or more nodes on which they are executing. Nodes allocated to a job are available for reuse as soon as the job(s) allocated to that node terminate. If some nodes fail to complete job termination in a timely fashion due to hardware of software problems, only the scheduling of those tardy nodes will be effected.

[1] A tip of the hat to Matt Groening and creators of *Futurama*, where Slurm is the most popular carbonated beverage in the universe.

- *Secure*: SLURM employs crypto technology to authenticate users to services and services to each other with a variety of options available through the plugin mechanism. SLURM does not assume that its networks are physically secure, but does assume that the entire cluster is within a single administrative domain with a common user base across the entire cluster.

- *System administrator friendly*: SLURM is configured a simple configuration file and minimizes distributed state. Its configuration may be changed at any time without impacting running jobs. Heterogeneous nodes within a cluster may be easily managed. SLURM interfaces are usable by scripts and its behavior is highly deterministic.

The main contribution of our work is that we have provided a readily available and inexpensive tool that anybody can use to efficiently manage clusters of different size and architecture. The SLURM is highly scalable[2]. The SLURM can be easily ported to any cluster system with minimal effort with its plug-in capability and can be used with any meta-batch scheduler or even a Grid resource broker [7] with its well-defined interfaces.

The rest of the paper is organized as follows. Section 2 describes the architecture of SLURM in detail. Section 3 discusses the services provided by SLURM followed by performance study of SLURM in Section 4. Brief survey of existing cluster management systems is presented in Section 5. Concluding remarks and future development plan of SLURM is given in Section 6.

## 2 SLURM Architecture

As a cluster resource manager, SLURM has three key functions. First, it allocates exclusive and/or non-exclusive access to resources to users for some duration of time so they can perform work. Second, it provides a framework for starting, executing, and monitoring work on the set of allocated nodes. Finally, it arbitrates conflicting requests for resources by managing a queue of pending work. Users and system administrators interact with SLURM using simple commands.

Figure 1 depicts the key components of SLURM. As shown in Figure 1, SLURM consists of a `slurmd` daemon running on each compute node, a central `slurmctld` daemon running on a management node (with optional fail-over twin), and five command line utilities, which can run anywhere in the cluster.

The entities managed by these SLURM daemons include nodes, the compute resource in SLURM, partitions, which group nodes into logical disjoint sets, jobs, or allocations of resources assigned to a user for a specified amount of time, and job steps, which are sets of tasks within a job. Each job is allocated nodes within a single partition. Once a job is assigned a set of nodes, the user is able to initiate parallel work in the form of job steps in any configuration within the allocation. For instance a single job step may be started which utilizes all nodes allocated to the job, or several job steps may independently use a portion of the allocation.

[2]It was observed that it took less than two seconds for SLURM to launch a thousand-task job on a large cluster currently being built for Lawrence Livermore National Laboratory.

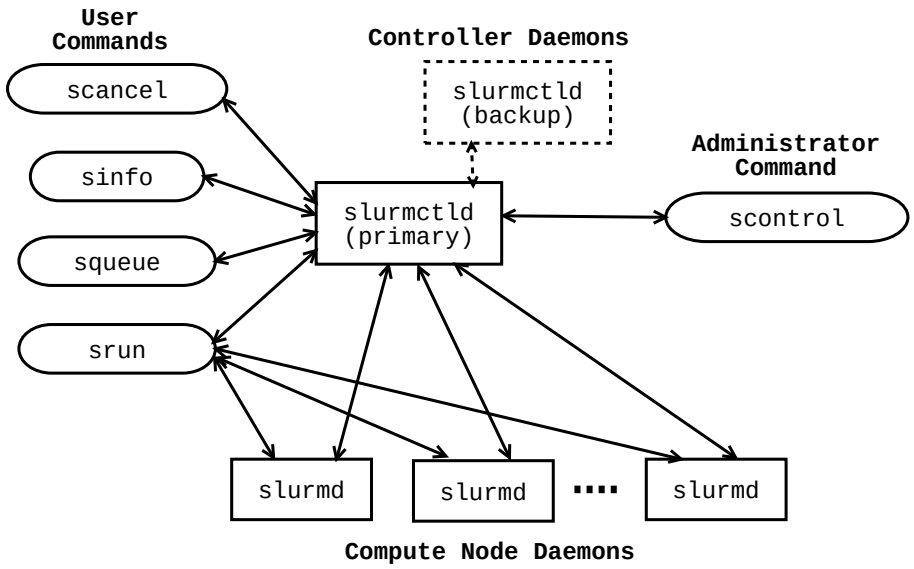



Figure 1: SLURM Architecture

Figure 2 exposes the subsystems that are implemented within the `slurmd` and `slurmctld` daemons. These subsystems are explained in more detail below.

## 2.1 SLURM Local Daemon (Slurmd)

The `slurmd` is a multi-threaded daemon running on each compute node. It reads the common SLURM configuration file, notifies the controller that it is active, waits for work, executes the work, returns status,then waits for more work. Since it initiates jobs for other users, it must run with root privilege. It also asynchronously exchanges node and job status information with `slurmctld`. The only job information it has at any given time pertains to its currently executing jobs. The `slurmd` performs five major tasks.

- *Machine and Job Status Services*: Respond to controller requests for machine and job state information, and send asynchronous reports of some state changes (e.g. `slurmd` startup) to the controller.
- *Remote Execution*: Start, monitor, and clean up after a set of processes (typically belonging to a parallel job) as dictated by the `slurmctld` daemon or an `srun` or `scancel` command. Starting a process may include executing a prolog program, setting process limits, setting real and effective user id, establishing environment variables, setting working directory, allocating interconnect resources, setting core file paths, initializing the Stream Copy Service, and managing process groups. Terminating a process may include terminating all members of a process group and executing an epilog program.
- *Stream Copy Service*: Allow handling of stderr, stdout, and stdin of remote tasks. Job input may be redirected from a file or files, a `srun` process, or /dev/null. Job output may be saved into local files or sent back to the `srun` command. Regardless of the

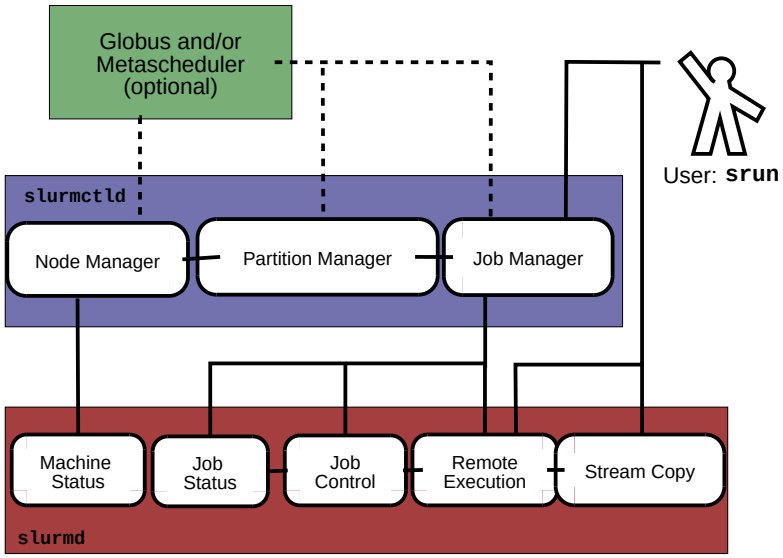



Figure 2: SLURM Architecture - Subsystems

location of stdout or stderr, all job output is locally buffered to avoid blocking local tasks.

- *Job Control*: Allow asynchronous interaction with the Remote Execution environment by propagating signals or explicit job termination requests to any set of locally managed processes.

## 2.2 SLURM Central Daemon (Slurmctld)

Most SLURM state information is maintained by the controller, `slurmctld`. The `slurmctld` is multi-threaded with independent read and write locks for the various data structures to enhance scalability. When `slurmctld` starts, it reads the SLURM configuration file. It also can read additional state information from a checkpoint file generated by a previous execution of `slurmctld`. Full controller state information is written to disk periodically with incremental changes written to disk immediately for fault-tolerance. The `slurmctld` runs in either master or standby mode, depending on the state of its fail-over twin, if any. The
slurmctld need not execute with root privilege. The `slurmctld` consists of three major components:

- *Node Manager*: Monitors the state of each node in the cluster. It polls `slurmd`'s for status periodically and receives state change notifications from `slurmd` daemons asynchronously. It ensures that nodes have the prescribed configuration before being considered available for use.

- *Partition Manager*: Groups nodes into non-overlapping sets called *partitions*. Each partition can have associated with it various job limits and access controls. The partition manager also allocates nodes to jobs based upon node and partition states and

configurations. Requests to initiate jobs come from the Job Manager. The `scontrol` may be used to administratively alter node and partition configurations.

- *Job Manager*: Accepts user job requests and places pending jobs in a priority ordered queue. The Job Manager is awakened on a periodic basis and whenever there is a change in state that might permit a job to begin running, such as job completion, job submission, partition-up transition, node-up transition, etc. The Job Manager then makes a pass through the priority-ordered job queue. The highest priority jobs for each partition are allocated resources as possible. As soon as an allocation failure occurs for any partition, no lower-priority jobs for that partition are considered for initiation. After completing the scheduling cycle, the Job Manager's scheduling thread sleeps. Once a job has been allocated resources, the Job Manager transfers necessary state information to those nodes, permitting it to commence execution. When the Job Manager detects that all nodes associated with a job have completed their work, it initiates clean-up and performs another scheduling cycle as described above.

# 3 SLURM Operation and Services

## 3.1 Command Line Utilities

The command line utilities are the user interface to SLURM functionality. They offer users access to remote execution and job control. They also permit administrators to dynamically change the system configuration. These commands all use SLURM APIs which are directly available for more sophisticated applications.

- `scancel`: Cancel a running or a pending job or job step, subject to authentication and authorization. This command can also be used to send an arbitrary signal to all processes on all nodes associated with a job or job step.

- `scontrol`: Perform privileged administrative commands such as draining a node or partition in preparation for maintenance. Many `scontrol` functions can only be executed by privileged users.

- `sinfo`: Display a summary of partition and node information. A assortment of filtering and output format options are available.

- `squeue`: Display the queue of running and waiting jobs and/or job steps. A wide assortment of filtering, sorting, and output format options are available.

- `srun`: Allocate resources, submit jobs to the SLURM queue, and initiate parallel tasks (job steps). Every set of executing parallel tasks has an associated `srun` which initiated it and, if the `srun` persists, managing it. Jobs may be submitted for batch execution, in which case `srun` terminates after job submission. Jobs may also be submitted for interactive execution, where `srun` keeps running to shepherd the running job. In this case, `srun` negotiates connections with remote `slurmd`'s for job initiation and to get stdout and stderr, forward stdin, and respond to signals from the user. The `srun` may also be instructed to allocate a set of resources and spawn a shell with access to those resources.

### 3.2 Plugins

In order to make the use of different infrastructures possible, SLURM uses a general purpose plugin mechanism. A SLURM plugin is a dynamically linked code object which is loaded explicitly at run time by the SLURM libraries. A plugin provides a customized implemenation of a well-defined API connected to tasks such as authentication, interconnect fabric, task scheduling. A set of functions is defined for use by all of the different infrastructures of a particular variety. For example, the authentication plugin must define functions such as: `slurm_auth_activate` to create a credential, `slurm_auth_verify` to verify a credential to approve or deny authentication, `slurm_auth_get_uid` to get the user ID associated with a specific credential. It also must define the data structure used, a plugin type, a plugin version number. The available plugins are defined in the configuration file.

### 3.3 Communications Layer

SLURM presently uses Berkeley sockets for communications. However, we anticipate using the plugin mechanism to easily permit use of other communications layers. At LLNL we are using an Ethernet for SLURM communications and the Quadrics Elan switch exclusively for user applications. The SLURM configuration file permits the identification of each node's hostname as well as its name to be used for communications.

While SLURM is able to manage 1000 nodes without difficulty using sockets and Ethernet, we are reviewing other communication mechanisms which may offer improved scalability. One possible alternative is STORM[**?**]. STORM uses the cluster interconnect and Network Interface Cards to provide high-speed communications including a broadcast capability. STORM only supports the Quadrics Elan interconnnect at present, but does offer the promise of improved performance and scalability.

### 3.4 Security

SLURM has a simple security model: Any user of the cluster may submit parallel jobs to execute and cancel his own jobs. Any user may view SLURM configuration and state information. Only privileged users may modify the SLURM configuration, cancel any jobs, or perform other restricted activities. Privileged users in SLURM include the users *root* and `SlurmUser` (as defined in the SLURM configuration file). If permission to modify SLURM configuration is required by others, set-uid programs may be used to grant specific permissions to specific users.

We presently support three authentication mechanisms via plugins: `authd`[**?**], `munged` and `none`. A plugin can easily be developed for Kerberos or authentication mechanisms as desired. The `munged` implementation is described below. A `munged` daemon running as user *root* on each node confirms the identify of the user making the request using the `getpeername` function and generates a credential. The credential contains a user ID, group ID, timestamp, lifetime, some pseudo-random information, and any user supplied information. The `munged` uses a private key to generate a Message Authentication Code (MAC) for the credential. The `munged` then uses a public key to symmetrically encrypt the credential including the MAC. SLURM daemons and programs transmit this encrypted credential with communications. The SLURM daemon receiving the message sends the credential to `munged`

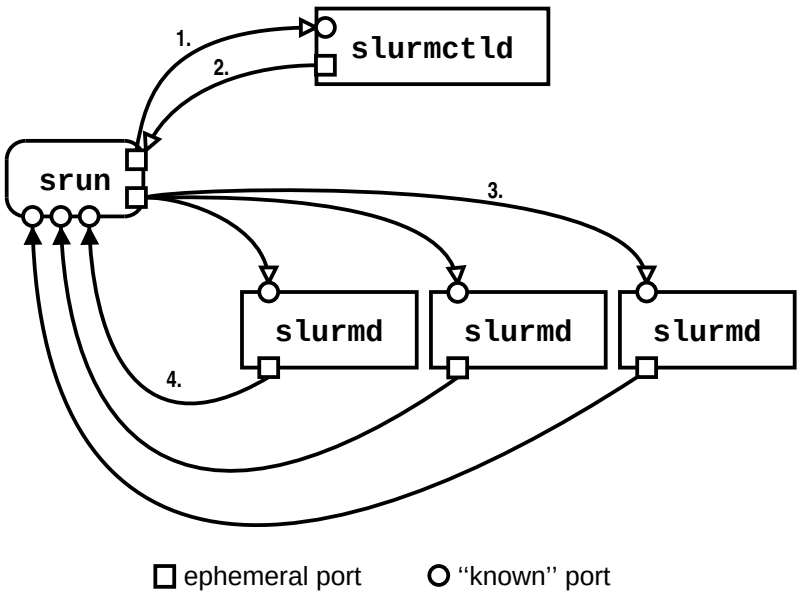



Figure 3: Job initiation connections overview. 1. The `srun` connects to `slurmctld` requesting resources. 2. `slurmctld` issues a response, with list of nodes and job credential. 3. The `srun` opens a listen port for every task in the job step, then sends a run job step request to `slurmd`. 4. `slurmd`'s initiate job step and connect back to `srun` for stdout/err.

on that node. The `munged` decrypts the credential using its private key, validates it and returns the user ID and group ID of the user originating the credential. The `munged` prevents replay of a credential on any single node by recording credentials that have already been authenticated. In SLURM's case, the user supplied information includes node identification information to prevent a credential from being used on nodes it is not destined for.

When resources are allocated to a user by the controller, a *job step credential* is generated by combining the user ID, job ID, step ID, the list of resources allocated (nodes), and the credential lifetime. This job step credential is encrypted with a `slurmctld` private key. This credential is returned to the requesting agent (`srun`) along with the allocation response, and must be forwarded to the remote `slurmd`'s upon job step initiation. `slurmd` decrypts this credential with the `slurmctld`'s public key to verify that the user may access resources on the local node. `slurmd` also uses this job step credential to authenticate standard input, output, and error communication streams.

### 3.5 Job Initiation

There are three modes in which jobs may be run by users under SLURM. The first and most simple is *interactive* mode, in which stdout and stderr are displayed on the user's terminal in real time, and stdin and signals may be forwarded from the terminal transparently to the remote tasks. The second is *batch* mode, in which the job is queued until the request for resources can be satisfied, at which time the job is run by SLURM as the submitting user. In *allocate* mode, a job is allocated to the requesting user, under which the user may manually run job steps via a script or in a sub-shell spawned by `srun`.

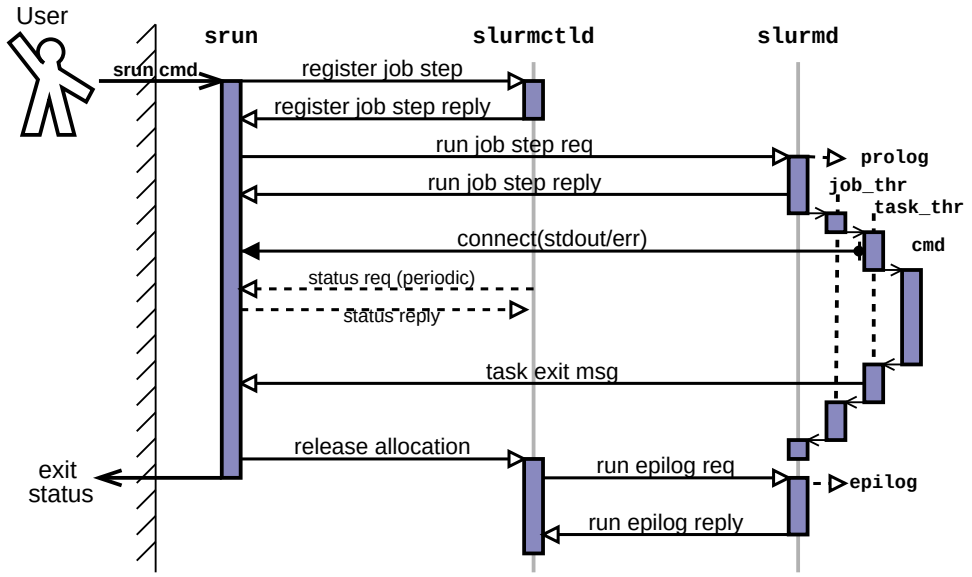



Figure 4: Interactive job initiation. `srun` simultaneously allocates nodes and a job step from `slurmctld` then sends a run request to all `slurmd`'s in job. Dashed arrows indicate a periodic request that may or may not occur during the lifetime of the job.

Figure 3 gives a high-level depiction of the connections that occur between SLURM components during a general interactive job startup. The `srun` requests a resource allocation and job step initiation from the `slurmctld`, which responds with the job ID, list of allocated nodes, job credential. if the request is granted. The `srun` then initializes listen ports for each task and sends a message to the `slurmd`'s on the allocated nodes requesting that the remote processes be initiated. The `slurmd`'s begin execution of the tasks and connect back to `srun` for stdout and stderr. This process and the other initiation modes are described in more detail below.

### 3.5.1 Interactive mode initiation

Interactive job initiation is illustrated in Figure 4. The process begins with a user invoking `srun` in interactive mode. In Figure 4, the user has requested an interactive run of the executable "`cmd`" in the default partition.

After processing command line options, `srun` sends a message to `slurmctld` requesting a resource allocation and a job step initiation. This message simultaneously requests an allocation (or job) and a job step. The `srun` waits for a reply from `slurmctld`, which may not come instantly if the user has requested that `srun` block until resources are available. When resources are available for the user's job, `slurmctld` replies with a job step credential, list of nodes that were allocated, cpus per node, and so on. The `srun` then sends a message each `slurmd` on the allocated nodes requesting that a job step be initiated. The `slurmd`'s verify that the job is valid using the forwarded job step credential and then respond to `srun`.

Each `slurmd` invokes a job thread to handle the request, which in turn invokes a task

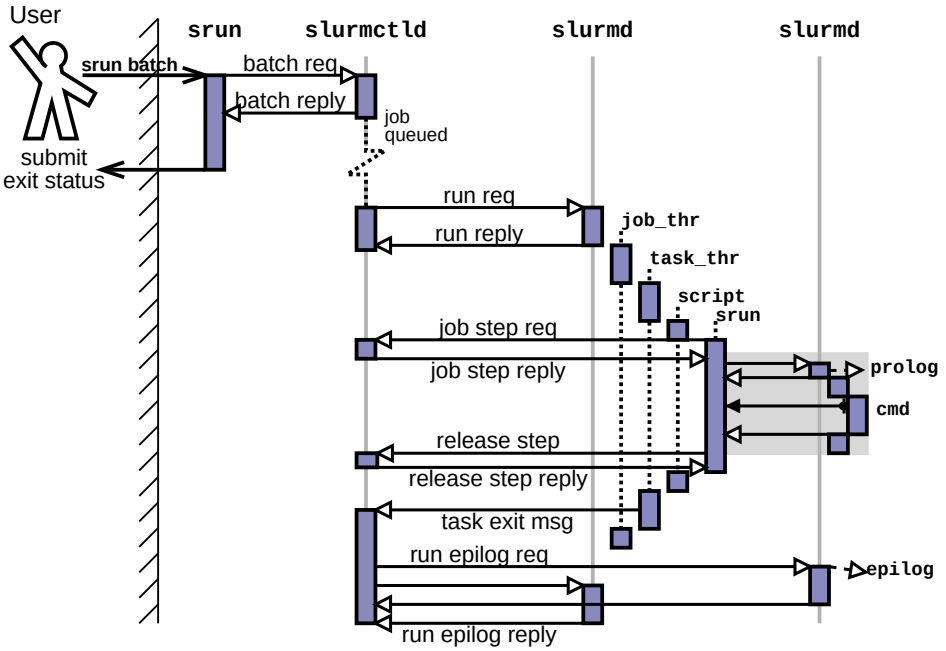



Figure 5: Queued job initiation. `slurmctld` initiates the user's job as a batch script on one node. Batch script contains an srun call which initiates parallel tasks after instantiating job step with controller. The shaded region is a compressed representation and is illustrated in more detail in the interactive diagram (Figure 4).

thread for each requested task. The task thread connects back to a port opened by `srun` for stdout and stderr. The host and port for this connection is contained in the run request message sent to this machine by `srun`. Once stdout and stderr have successfully been connected, the task thread takes the necessary steps to initiate the user's executable on the node, initializing environment, current working directory, and interconnect resources if needed.

Once the user process exits, the task thread records the exit status and sends a task exit message back to `srun`. When all local processes terminate, the job thread exits. The `srun` process either waits for all tasks to exit, or attempt to clean up the remaining processes some time after the first task exits. Regardless, once all tasks are finished, `srun` sends a message to the `slurmctld` releasing the allocated nodes, then exits with an appropriate exit status.

When the `slurmctld` receives notification that `srun` no longer needs the allocated nodes, it issues a request for the epilog to be run on each of the `slurmd`'s in the allocation. As `slurmd`'s report that the epilog ran successfully, the nodes are returned to the partition.

### 3.5.2 Batch mode initiation

Figure 5 illustrates the initiation of a batch job in SLURM. Once a batch job is submitted, `srun` sends a batch job request to `slurmctld` that contains the input/output location for the job, current working directory, environment, requested number of nodes. The `slurmctld` queues the request in its priority ordered queue.

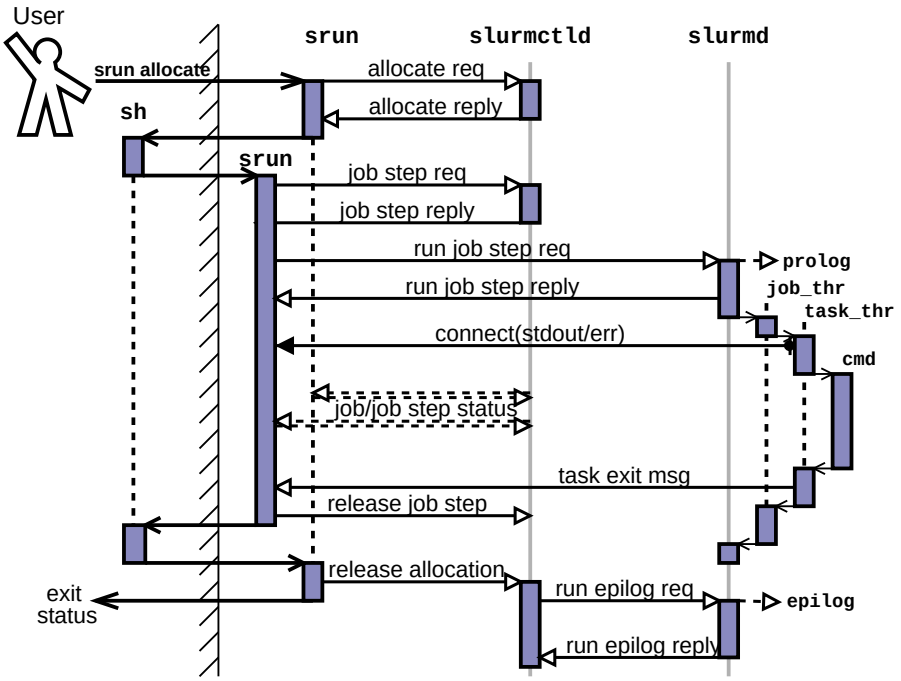



Figure 6: Job initiation in allocate mode. Resources are allocated and `srun` spawns a shell with access to the resources. When user runs an `srun` from within the shell, the a job step is initiated under the allocation.

Once the resources are available and the job has a high enough priority, `slurmctld` allocates the resources to the job and contacts the first node of the allocation requesting that the user job be started. In this case, the job may either be another invocation of `srun` or a *job script* which may have multiple invocations of `srun` within it. The `slurmd` on the remote node responds to the run request, initiating the job thread, task thread, and user script. An `srun` executed from within the script detects that it has access to an allocation and initiates a job step on some or all of the nodes within the job.

Once the job step is complete, the `srun` in the job script notifies the `slurmctld` and terminates. The job script continues executing and may initiate further job steps. Once the job script completes, the task thread running the job script collects the exit status and sends a task exit message to the `slurmctld`. The `slurmctld` notes that the job is complete and requests that the job epilog be run on all nodes that were allocated. As the `slurmd`'s respond with successful completion of the epilog, the nodes are returned to the partition.

### 3.5.3 Allocate mode initiation

In allocate mode, the user wishes to allocate a job and interactively run job steps under that allocation. The process of initiation in this mode is illustrated in Figure 6. The invoked `srun` sends an allocate request to `slurmctld`, which, if resources are available, responds with a list of nodes allocated, job id, etc. The `srun` process spawns a shell on the user's terminal with access to the allocation, then waits for the shell to exit at which time the job is considered complete.

An `srun` initiated within the allocate sub-shell recognizes that it is running under an allocation and therefore already within a job. Provided with no other arguments, `srun` started in this manner initiates a job step on all nodes within the current job. However, the user may select a subset of these nodes implicitly.

An `srun` executed from the sub-shell reads the environment and user options, then notify the controller that it is starting a job step under the current job. The `slurmctld` registers the job step and responds with a job credential. The `srun` then initiates the job step using the same general method as described in the section on interactive job initiation.

When the user exits the allocate sub-shell, the original `srun` receives exit status, notifies `slurmctld` that the job is complete, and exits. The controller runs the epilog on each of the allocated nodes, returning nodes to the partition as they complete the epilog.

# 4 Related Work

### Portable Batch System (PBS)

The Portable Batch System (PBS) [19] is a flexible batch queuing and workload management system originally developed by Veridian Systems for NASA. It operates on networked, multi-platform UNIX environments, including heterogeneous clusters of workstations, supercomputers, and massively parallel systems. PBS was developed as a replacement for NQS (Network Queuing System) by many of the same people.

PBS supports sophisticated scheduling logic (via the Maui Scheduler[3]). PBS spawn's daemons on each machine to shepherd the job's tasks. It provides an interface for administrators to easily interface their own scheduling modules. PBS can support long delays in file staging with retry. Host authentication is provided by checking port numbers (low ports numbers are only accessible to user root). Credential service is used for user authentication. It has the job prolog and epilog feature. PBS Supports high priority queue for smaller "interactive" jobs. Signal to daemons causes current log file to be closed, renamed with time-stamp, and a new log file created.

Although the PBS is portable and has broad user base, it has following drawbacks. First, PBS is implementted using single thread and hence exhibits poor performance especially when a compute node in the system dies in which case the PBS tries to contact down nodes, while other activities must wait. Second, PBS has weak mechanism for starting and cleaning up parallel jobs. Finally, the PBS has very poor scalability and is not suitable for large clusters.

### Maui Scheduler

Maui Scheduler [16] is an advance reservation HPC batch scheduler for use with SP, O2K, and UNIX/Linux clusters. It is widely used to extend the functionality of PBS and LoadLeveler

[3] http://superclustergroup.org/maui

### Distributed Production Control System (DPCS)

The Distributed Production Control System (DPCS) [6] is a resource manager developed at Lawrence Livermore National Laboratory (LLNL). The DPCS provides basic data collection and reporting mechanisms for prject-level, near real-time accounting and resource allocation to customers with established limits per customers' organization budgets, In addition, the DPCS evenly distributes workload across available computers and supports dynamic reconfiguration and graceful degradation of service to prevent overuse of a computer where not authorized.

While DPCS does have these attractive characteristics, it supports only a limited number of computer systems: IBM RS/6000 and SP, Linux with RMS, Sun Solaris, and Compaq Alpha. DPCS also lacks commercial support.

### LoadLeveler

LoadLeveler citeLoadLevelerManual,LoadLevelerWeb is a proprietary batch system and parallel job manager by IBM. LoadLeveler supports few non-IBM systems. Very primitive scheduling software exists and other software is required for reasonable performance such as Maui Scheduler and DPCS. The LoadLeveler is simple and very flexible queue and job class structure is available operating in "matrix" fashion. The biggest problem of the LoadLeveler is its extremely poor scalability. It was observed that it took more than 20 minutes to start-up a 512-task job using LoadLeveler on one of the IBM SP machines at LLNL. In addition, all jobs must be initiated through the LoadLeveler, and a special version of MPI is requested to run a parallel job.

### Load Sharing Facility (LSF)

LSF [13] is a proprietary batch system and parallel job manager by Platform Computing. Widely deployed on a wide variety of computer architectures, it has sophisticated scheduling software including fair-share, backfill, consumable resources, an job preemption and very flexible queue structure. It also provides good status information on nodes and LSF daemons. The LSF share many of its shortcomings with the LoadLeveler: job initiation only through LSF, requirement of a spwcial MPI library, etc.

### Condor

Condor [5, **?**, 1] is a batch system and parallel job manager developed by the University of Wisconsin. Condor was the basis for IBM's LoadLeveler and both share very similar underlying infrastructure. Condor has a very sophisticated checkpoint/restart service that does not rely upon kernel changes, but a variety of library changes (which prevent it from being completely general). The Condor checkpoint/restart service has been integrated into LSF, Codine, and DPCS. Condor is designed to operate across a heterogeneous environment, mostly to harness the compute resources of workstations and PCs. It has an interesting "advertising" service. Servers advertise their available resources and consumers advertise their requirements for a broker to perform matches. The checkpoint mechanism is used to relocate work on demand (when the "owner" of a desktop machine wants to resume work).

### Memory Channel

Memory Channel is a high-speed interconnect developed by Digital/Compaq with related software for parallel job execution. Special version of MPI required. The application spawns tasks on other nodes. These tasks connect themselves to the high speed interconnect. No system level tool to spawns the tasks, allocates interconnect resources, or otherwise manages the parallel job (Note: This is sometimes a problem when jobs fail, requiring system administrators to release interconnect resources. There are also performance problems related to resource sharing).

### Linux PAGG Process Aggregates

PAGG [11] consists of modifications to the linux kernel that allows developers to implement Process AGGregates as loadable kernel modules. A process aggregate is defined as a collection of processes that are all members of the same set. A set would be implemented as a container for the member processes. For instance, process sessions and groups could have been implemented as process aggregates.

### Beowulf Distributed Process Space (BPROC)

The Beowulf Distributed Process Space (BPROC) is set of kernel modifications, utilities and libraries which allow a user to start processes on other machines in a Beowulf-style cluster [2]. Remote processes started with this mechanism appear in the process table of the front end machine in a cluster. This allows remote process management using the normal UNIX process control facilities. Signals are transparently forwarded to remote processes and exit status is received using the usual wait() mechanisms.

### LAM / MPI

Local Area Multicomputer (LAM) [14] is an MPI programming environment and development system for heterogeneous computers on a network. With LAM, a dedicated cluster or an existing network computing infrastructure can act as one parallel computer solving one problem. LAM features extensive debugging support in the application development cycle and peak performance for production applications. LAM features a full implementation of the MPI communication standard.

## 5 Performance Study

We were able to perform some SLURM tests on a 1000 node cluster at LLNL. Some development was still underway at that time and tuning had not been performed. The results for executing simple 'hostname' program on two tasks per node and various node counts is show in Figure 7. We found SLURM performance to be comparable to the Quadrics Resource Management System (RMS) [20] for all job sizes and about 80 times faster than IBM LoadLeveler [12, 9] at tested job sizes.

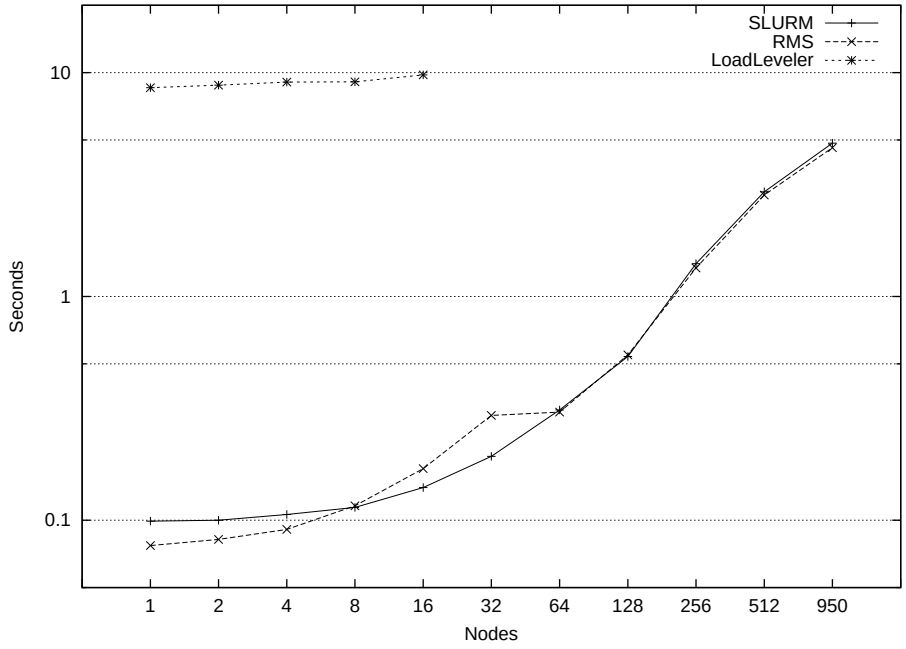



Figure 7: Time to execute /bin/hostname with various node counts

# 6 Conclusion and Future Plans

We have developed and presented SLURM, a simple, highly scalable, robust, and portable cluster resource management system, in this paper. The contribution of this work is that we have provided a immediately-available and inexpensive tool that virtually anybody can use to efficiently manage clusters of different size and architecture. We expect SLURM to begin production use on LLNL Linux clusters starting in March 2003 and be available for distribution shortly thereafter.

Looking ahead, we anticipate adding support for additional operating systems. We anticipate adding a job preempt/resume capability, which will provide an external scheduler the infrastructure required to perform gang scheduling, and a checkpoint/restart capability. We also plan to use the SLURM for IBM's Blue Gene/L platform [4] by incorporating a capability to manage jobs on a three-dimensional torus machine into the SLURM.

# Acknowledgments

Additional programmers are responsible for the development of SLURM include: Chris Dunlap, Joey Ekstrom, Jim Garlick, Kevin Tew and Jay Windley.

## References


[1] J. Basney, M. Livny, and T. Tannenbaum. High Throughput Computing with Condor. *HPCU news*, 1(2), June 1997.

[2] Beowulf Distributed Process Space. `http://bproc.sourceforge.net`.

[3] Beowulf Project. `http://www.beowulf.org`.

[4] Blue Gene/L. `http://cmg-rr.llnl.gov/asci/platforms/bluegenel`.

[5] Condor. `http://www.cs.wisc.edu/condor`.

[6] Distributed Production Control System. `http://www.llnl.gov/icc/lc/dpcs_overview.html`.

[7] I. Foster and C. Kesselman. *The GRID: Blueprint for a New Computing Onfrastructure*. Morgan Kaufmann Publishers, Inc., 1999.

[8] GNU General Public License. `http://www.gnu.org/licenses/gpl.html`.

[9] IBM Corporation. *LoadLeveler's User Guide, Release 2.1.*

[10] M. Jette, C. Dunlap, J. Garlick, and M. Grondona. Survey of Batch/Resource Management-Related System Software. Technical report, Lawrence Livermore National Laboratory, 2202.

[11] Linux PAGG Process Aggregates. `oss.sgi.com/projects/pagg`.

[12] Load Leveler. `http://www-1.ibm.com/servers/eservers/pseries/library/sp_books/loadleveler.htr`

[13] Load Sharing Facility. `http://www.platform.com`.

[14] Local Area Multicomputer. `http://www.lam-mpi.org`.

[15] Loki – Commodity Parallel Processing. `http://loki-www.lanl.org`.

[16] Maui Scheduler. `mauischeduler.sourceforge.net`.

[17] Multiprogrammatic Capability Cluster. `http://www.llnl.gov/linux/mcr`.

[18] Parallel Capacity Resource. `http://www.llnl.gov/linux/pcr`.

[19] Portable Batch System. `http://www.openpbs.org`.

[20] Quadircs Resource Management System. `http://www.quadrics.com/website/pdf/rms.pdf`.