


IIT DELHI POWERS INDIA'S RESEARCH COMMUNITY

The Indian Institute of Technology Delhi (IIT Delhi), established in 1961, is one of the seven oldest Indian Institutes of Technology. Its mission is to contribute to India and to the world through excellence in scientific and technical education and research, and its 320-acre campus hosts its own high-performance computing (HPC) resources including the "PADUM" hybrid supercomputer. In 2020 IIT Delhi allocated dedicated computational resources to COVID-19 research projects including drug repurposing, molecular modeling, and pathogen movement.

Workload Management on PADUM

In 2016, the PADUM system at IIT Delhi ranked 217th on the Top500 list, fourth in India, and 28th on the Green500 list. PADUM had a total of 438 nodes (10,384 cores), including 238 conventional CPU nodes, 200 accelerated nodes including NVIDIA K40-based nodes, and 8 login nodes. IIT Delhi has since added 184 more Skylake nodes to the cluster which comprises 144 conventional CPU nodes and 40 GPU nodes with NVIDIA V100 cards.

To support this centralized HPC system, the team at IIT Delhi needed a commercial workload manager. They also needed support and customization to implement scheduling requirements, queue policies, and system fine-tuning, as well as a simple but fully functional user interface.



Support for critical research with global impact



Computational resources to power discovery
Images are representational only

They brought Altair six workload management requirements:

- Consistent and easy-to-use interface
- Job priority sorting
- Restricted access to system resources and device isolation
- User and group limits and allocation
- Node health checking
- Utilization reporting

The Altair® PBS Professional® workload manager satisfied each of those requirements, and PBS Professional was implemented with functionality including usage-based charging via Allocation Manager and reporting by PBS Analytics.

Consistent and Easy-to-use Interface

IIT Delhi users needed the ability to easily choose between high-memory or regular-memory nodes, top/high/standard priority queues, number of GPUs, number of CPU cores, jobs requiring exclusive access, and projects with accounts under which jobs are submitted, along with the ability for a user to belong to multiple projects.

The team at IIT Delhi also needed several checks in place:

- Number of CPU cores rounded up to the number of GPUs requested
- High memory defaults to exclusive access; exclusivity at the per-node level
- In exclusive mode, the number of cores/GPUs is rounded up to nodes, or users can directly request nodes
- Guided interface to verify job requests at the time of submission for unrealistic, unauthorized resource values; prompts users for probable mistakes and provides suggestions
- Entire cluster can be used for maximum one day per user (more time is allowed by reducing the resources requested); multiplication of duration by core numbers must be constant

Altair's PBS Professional solution utilizes the PBS Professional plugin framework (a.k.a. hooks). By using the queuejob hook event, Altair satisfied all these requirements for IIT Delhi. The hook checks each resource specified by a user in a job script, and if a user requests conflicting resources not in line with site policies, it rejects the job with an explanatory error message.

The hook can be configured to make suggestions to users, and the site administrator can adjust messages by simply updating the hook. The hook also checks the wall time of each job to make sure it's within specified constraints, and it confirms that the user-specified project is valid and the user is authorized to submit jobs in it.

Job Priority Sorting

Because high system utilization was critical to IIT Delhi, we needed to ensure that jobs requesting GPUs or high-memory resources were given higher priority than jobs requesting conventional CPUs. This priority evaluation would allow important jobs to be scheduled on these special resources before conventional jobs. Job priority is also influenced by how long a job has been waiting for resources and what queue it's in.

Because no single scheduling algorithm can include all these parameters in the required weightage, PBS Professional has a versatile feature known as `job_sort_formula` that allows users to embed their own job priority intelligence in the system and have the scheduler sort jobs based on the resulting values. The `job_sort_formula` can use parentheses, exponents, unary + and - operators, and the ternary operator (which must be Python). All operators use standard mathematical precedence. The formula can use standard Python mathematical operators and those in the Python math module. To satisfy the job priority requirement, a sophisticated `job_sort_formula` was plugged in with several weighted parameters.

PBS Professional's `backfill_depth` capability boosted system utilization by allowing smaller, conventional jobs to run on the remaining available resources.

Restricted Access to System Resources and Device Isolation

IIT Delhi needed to ensure that what the user requested for the job is exactly what the job was allocated — no more. For example, a user can ask for fewer resources but unknowingly consume more, causing other jobs on the same node to suffer from poor performance, or worse, fail because of insufficient resources. The same is true for GPU devices. The team needed for a user requesting one GPU on a node to see only one GPU on that node, regardless of the number of GPUs actually available.


Altair's integration with Linux cgroups, which restricts access to system resources and accounts for resource usage, satisfied this requirement. The PBS Professional cgroups hook creates a cgroup (and other cgroups subsystems) for each job and every node allocated to the job. When a job is started, the hook creates a set of directories for the configured subsystems based on the resource requirements of the job, then places the job process within the cgroup. The kernel then enforces resource restrictions based on the newly created cgroup settings while the job is running. When the job finishes, the hook gets final resource usage measurements and removes the cgroup directories it created when the job was initialized.



High-performance computing for leading-edge science

The cgroups hook automatically detects and configures GPUs and will restrict jobs to specific allocated GPUs, providing GPU isolation so a job can access only as many GPU devices as requested.

The cgroups hook works for multi-node jobs if MPI integration is in place, and we achieved that by customizing the SSH configuration on the nodes. Doing separate integrations for separate MPIs isn't effective in an academic environment because users can employ many different MPIs by installing them in non-standard paths like home directories.



Leading-edge research in a broad range of disciplines

User and Group Limits and Allocation

Another requirement for IIT Delhi was the ability to control the cluster's usage based on groups, users within those groups, and funds available in various projects to which they have access. The site has three queues, and each queue has defined controls for users, groups, projects, and budgets.

- No-limit (low) queue – low priority, infinite budget, all users have access
- Default (standard) queue – limited budget, all users have access
- High-priority queue – high priority, limited budget, accessible only by privileged users and groups
- Submission to the high-priority queues requires specific permissions
- Non-default limit allowed at a group/user level (a user can belong to multiple groups)
- Even for group-based access, usage log and remaining compute-minutes are per user

Altair satisfied these requirements with existing capabilities of PBS Professional and customization of Altair's Allocation Manager.

PBS Professional allows usage tracking based on a virtual entity like "project." All the groups to which a user belongs are created as separate projects, and a mapping of users versus projects is created. A user can select a project at the time of job submission using the -P option and can submit jobs in valid projects for which they are authorized.

Altair's Allocation Manager allows users to track compute-minutes in terms of money. It also allows them to configure separate charging units like Cost and LowQueueCost. Cost is for high-priority runs with limited money and quotas. LowQueueCost is charged only for low-priority queue jobs and is unlimited. Per-unit charges for cores and GPUs can be defined separately. Allocation Manager provides a simple user interface to query budget status and usage.

Node Health Checking

An additional requirement at IIT Delhi was the ability to detect faulty or misconfigured nodes in order to avoid job failures. A cluster manager continuously monitors cluster health, and there are events during a job's lifecycle that need to be checked and double-checked to ensure successful job execution. Upon detecting a faulty configuration, IIT Delhi required that the site administrator be notified and the faulty node(s) be excluded from scheduling eligibility until the site administrator could troubleshoot.

Altair's solution took advantage of the PBS Professional plugin framework (a.k.a. hooks). An execjob_begin hook was developed to check the file systems on all nodes assigned to a job before launching the job on them. It also checks to see that LDAP service is running and checks /scratch file system directory write permission. If any of these checks fail, the node is taken offline and a comment about the failure cause is tagged on the node. Admin is notified of any faulty nodes.

The hook also provided the IIT Delhi team peace of mind that if a failure was detected the job was automatically requeued, saving the user from the needing to resubmit the same job. If a job requeues by more than a specified count, both the user and the site admin are notified.

Utilization Reporting

Finally, IIT Delhi required workload and cluster utilization reporting so its team could do advance quarterly capacity planning and budget allocation.

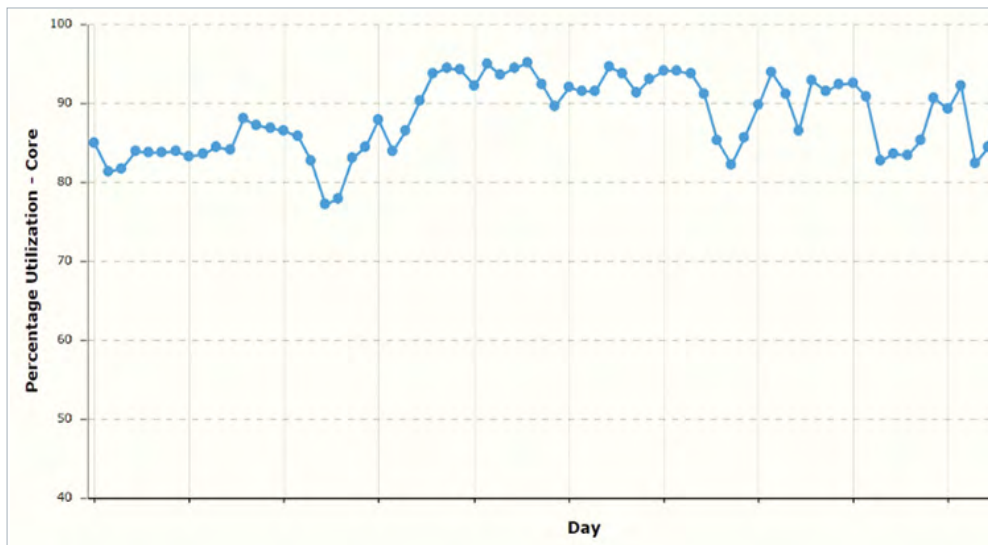
Altair's PBS Analytics provides a global view of IIT Delhi's HPC usage for chargebacks, capacity planning, troubleshooting, and project management. PBS Analytics ingests the PBS Professional accounting logs and delivers detailed reporting via the "auto report generation" customization feature. It places charts and job csv files in a predefined location to be used as a basis for budget population for each project in the Allocation Manager system. Reports can be sent automatically to stakeholders.



One of seven original Indian Institutes of Technology

IIT Delhi Powering Science and Research

The PADUM system at the Indian Institute of Technology Delhi continues to power discovery in India. With its HPC resources running efficiently and supporting more science and research than ever before, PADUM is a critical piece of the country's computing infrastructure and an important resource for the academic community.



The compute-hours it has dedicated to COVID-19 research are likely to make significant contributions to the global knowledge base for the disease's behavior, treatment, and prevention.