

What's New Platform LSF 8

An Overview of New Features for Platform LSF Administrators



Written with Platform LSF administrators in mind, this brief provides a short explanation of significant changes in Platform's latest release of Platform LSF, with a specific emphasis on scheduling and workload management features.

About Platform LSF 8

Platform LSF is the most powerful workload manager for demanding, distributed high performance computing environments. It provides a complete set of workload management capabilities, all designed to work together to reduce cycle times and maximize productivity in mission-critical environments.

This latest Platform LSF release delivers improvements in performance and scalability while introducing new features that simplify administration and boost user productivity. This includes:

- **Guaranteed resources** - Aligns business SLA's with infrastructure configuration for simplified administration and configuration
- **Live reconfiguration** - Provides simplified administration and enables agility
- **Delegation of administrative rights** - Empowers line of business owners to take control of their own projects
- **Fairshare & pre-emptive scheduling enhancements** - Fine tunes key production policies

Platform LSF 8 Features

Guaranteed Resources Ensure Deadlines are Met

In Platform LSF 8, resource-based scheduling has been extended to guarantee resource availability to groups of jobs. Resources can be slots, entire hosts or user-defined shared resources such as software licenses. As an example, a business unit might guarantee that it has access to specific types of resources within ten minutes of a job being submitted,

even while sharing resources between departments. This facility ensures that lower priority jobs using the needed resources can be pre-empted in order to meet the SLAs of higher priority jobs.

Because jobs can be automatically attached to an SLA class via access controls, administrators can enable these guarantees without requiring that end-users change their job submission procedures, making it easy to implement this capability in existing environments.

Live Cluster Reconfiguration

Platform LSF 8 incorporates a new live reconfiguration capability, allowing changes to be made to clusters¹ without the need to re-start LSF daemons. This is useful to customers who need to add hosts, adjust sharing policies or re-assign users between groups "on the fly", without impacting cluster availability or running jobs.

Changes to the cluster configuration can be made via the `bconf` command line utility, or via new API calls. This functionality can also be integrated via a web-based interface using Platform Application Center. All configuration modifications are logged for a complete audit history, and changes are propagated almost instantaneously. The majority of reconfiguration operations are completed in under half a second.

With Live Reconfiguration, down-time is reduced, and administrators are free to make needed adjustments quickly rather than wait for scheduled maintenance periods or non-peak hours. In cases where users are members of multiple groups, controls can be put in place so that a group administrator can only control jobs associated with their designated group rather than impacting jobs related to another group submitted by the same user.

¹This capability pertains to both the LSF "base" and LSF "batch" configurations

Delegation of Administrative Rights

With Platform LSF 8, the concept of group administrators has been extended to enable project managers and line of business managers to dynamically modify group membership and fairshare resource allocation policies within their group. The ability to make these changes dynamically to a running cluster is made possible by the Live Reconfiguration feature.

These capabilities can be delegated selectively depending on the group and site policy. Different group administrators can manage jobs, control sharing policies or adjust group membership.

More Flexible Fairshare Scheduling Policies

To enable better resource sharing flexibility with Platform LSF 8, the algorithms used to tune dynamically calculated user priorities can be adjusted at the queue level. These algorithms can vary based on department, application or project team preferences. The Fairshare parameters `ENABLE_HIST_RUNTIME` and `HIST_HOURS` enable administrators to control the degree to which LSF considers prior resource usage when determining user priority.

The flexibility of Platform LSF 8 has also been improved by allowing a similar “decay rate” to apply to currently running jobs (`RUN_TIME_DECAY`), either system-wide or at the queue level. This is most useful for customers with long-running jobs, where setting this parameter results in a more accurate view of real resource use for the fairshare scheduling to consider.

Performance & Scalability Enhancements

Platform LSF has been extended to support an unparalleled scale of up to 100,000 cores and 1.5 million queued jobs for very high throughput EDA workloads. Even higher scalability is possible for more traditional HPC workloads.

Specific areas of improvement include the time required to start the master-batch daemon (MBD), `bjobs` query performance, job submission and job dispatching as well as impressive performance gains resulting from the new Bulk

Job Submission feature. In addition, on very large clusters with large numbers of user groups employing fairshare scheduling, the memory footprint of the master batch scheduler in LSF has been reduced by approximately 70% and scheduler cycle time has been reduced by 25%, resulting in better performance and scalability.

More Sophisticated Host-based Resource Usage for Parallel Jobs

Platform LSF 8 provides several improvements to how resource use is tracked and reported with parallel jobs. Accurate tracking of how parallel jobs use resources such as CPUs, memory and swap, is important for ease of management, optimal scheduling and accurate reporting and workload analysis. With Platform LSF 8 administrators can track resource usage on a per-host basis and an aggregated basis (across all hosts), ensuring that resource use is reported accurately. Additional details such as running PIDs and PGIDs for distributed parallel jobs, manual cleanup (if necessary) and the development of scripts for managing parallel jobs are simplified. These improvements in resource usage reporting are reflected in LSF commands including `bjobs`, `bhist` and `bacct`.

Improved Ease of Administration for Mixed Windows and Linux Clusters

The `lspasswd` command in Platform LSF enables Windows LSF users to advise LSF of changes to their Windows level passwords. With Platform LSF 8, password synchronization between environments has become much easier to manage because the Windows passwords can now be adjusted directly from Linux hosts using the `lspasswd` command. This allows Linux users to conveniently synchronize passwords on Windows hosts without needing to explicitly login into the host.

Bulk Job Submission

When submitting large numbers of jobs with different resource requirements or job level settings, Bulk Job Submission allows for jobs to be submitted in bulk by referencing a single file containing job details.

Platform Computing is the leader in cluster, grid and cloud management software - serving more than 2,000 of the world's most demanding organizations for over 18 years. Our workload and resource management solutions deliver IT responsiveness and lower costs for enterprise and HPC applications. Platform has strategic relationships with Cray, Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat® and SAS®. Visit www.platform.com.

World Headquarters

Platform Computing Corporation
3760 14th Avenue
Markham, Ontario
Canada L3R 3T7
Tel: +1 905 948 8448
Fax: +1 905 948 9975
Toll-free Tel: 1 877 528 3676
info@platform.com

Sales - Headquarters

Toll-free Tel: 1 877 710 4477
Tel: +1 905 948 8448

North America

New York: +1 212 888 6270
San Jose: +1 408 392 4900

Europe

Bramley: +44 (0) 1256 883756
London: +44 (0) 20 3206 1470
Paris: +33 (0) 1 41 10 09 20
Düsseldorf: +49 2102 61039 0
info-europe@platform.com

Asia-Pacific

Beijing: +86 10 82276000
Xi'an: +86 029 87607400
asia@platform.com
Tokyo: +81(0)3 6302 2901
info-japan@platform.com
Singapore: +65 6307 6590
wliaw@platform.com

