

Platform LSF 7 Update 6

An Overview of New Features for Platform LSF Administrators



With Platform LSF 7 update 6 (LSF 7u6), Platform has introduced several new features and enhancements, extending Platform's leadership in workload management. Written with Platform LSF administrators in mind, this brief provides a short explanation of what's new in Platform's latest release of Platform LSF.

About Platform LSF 7 Update 6

The emphasis of Platform LSF 7u6 is on ease of use features, improving user productivity, and increasing the operational efficiency of the cluster to drive higher throughput and reduce cost.

This release lays the foundation for the new Platform HPC portal, a flexible, intuitive interface for both cluster users and administrators that can optionally be added to your Platform LSF 7 cluster. The HPC Portal includes an easy to use Application Interface builder with integrated job and cluster management features.

In this release Platform continues to enhance Platform LSF to meet the needs of sophisticated HPC environments. As customers grow in their sophistication using Platform LSF to obtain competitive advantage and reduce internal costs they often find that their needs have been anticipated with features already present in the Platform LSF family. This degree of refinement, made possible by Platform's sustained engineering investment, is one of the key reasons that sophisticated customers choose Platform LSF to optimize the use of their IT infrastructure.

Platform LSF 7 Update 6 Features

Flexible Job Descriptions

Platform LSF customers often use the job name switch (-J) when submitting jobs to provide meaningful names for their workloads. In large environments, with many related jobs however, it can become difficult to keep track of which job is which.

With Platform LSF 7u6, detailed job descriptions can be provided, and several commands have been enhanced to allow users to search for jobs based on keywords in the job description. Up to four kilobytes of descriptive text can now be associated with any job using a new "-Jd" command line option. The new capability affects a number of Platform LSF commands including *bsub*, *bjobs*, *bmod* and *bhist* and job descriptions are written into the Platform LSF accounting data. This ensures that detailed job descriptions are accessible for reporting purposes.

Particularly useful is the ability to search the text in the job descriptions for keywords. As an example, a command like *bmod -Jd "*stress fracture*"* would find and act on all jobs with the string "stress fracture" embedded in the job description. This makes it easier to manage clusters with large volumes of jobs.

Improved Exception Reporting

Error messages in Platform LSF 7u6 have been made easier to understand. As an example, when jobs on a cluster terminate unexpectedly, cluster users would often use the command "*bjobs -l <job-id>*" to diagnose what happened. The *bjobs* output provided only an exit error code that was difficult to interpret. With Platform LSF 7u6, a full description of the termination reason is provided in the *bjobs* output.

Platform LSF MultiCluster Enhancements

Scheduling policies for sharing work between clusters has become significantly more flexible in Platform LSF 7u6. Prior to this release, when considering where to schedule jobs in a MultiCluster environment, the scheduling algorithms in Platform LSF MultiCluster considered only the slots on each remote cluster that were available. It did not take into account considerations like the number of slots that might be made available on a cluster if workloads were pre-empted (as an example).

With Platform LSF 7u6, these policies in the Platform LSF MultiCluster job forwarding model have been made more flexible. Several factors can now be considered in making a scheduling decision including the status of remote queues (including running slots and priorities), preempted available slots, and pending slots on preemptable queues on the remote cluster.

These enhancements support several different use cases for how jobs can be distributed among multiple clusters, making Platform LSF MultiCluster more effective for a wider range of requirements.

Multi-Phase Resource Reservation Enhancements

In some production environments, with long-running jobs, the resource requirements for a job can vary in known and predictable ways. For example, consider a job that runs for three hours and requires a significant amount of memory for the first 20 minutes (1 GB of RAM perhaps) and then requires less memory after that. Allocating 1 GB to the job via an *rusage* directive in a resource requirements string for the full three hour jobs is wasteful of resources since the inflated memory requirement can prevent other jobs from running even though there may be ample memory resources actually available.

Platform LSF 7u6 provides new flexibility in how resources are reserved supporting reservations in different job phases. Each phase can have a different pre-set duration. Granular controls also exist from phase to phase to control whether the resource requirement grows or decays linearly between phases, or whether it changes in a step-wise fashion.

While this might seem like a minor enhancement to some, this new flexibility in Platform LSF 7u6 translates directly into tangible financial benefits. By intelligent selection of Multi-Phase Resource Reservation policies, more jobs can be “packed” onto a particular host and completed in a set period of time. This improved use of resource allows sites to realize higher levels of throughput without the need to purchase additional hardware resources, often avoiding considerable cost.

Simplified Management through Resource Reservation Limit Enhancements

With prior versions of Platform LSF, queue level resource requirement settings (**RES_REQ**) acted as a ceiling that job level and application profile level settings could not exceed. Also, jobs submitted without a resource requirement specification would inherit the queue level requirement, meaning that jobs often reserved more resources than required leading to sub-optimal resource utilization.

Platform LSF 7u6 introduces a new **RESRSV_LIMIT** expression that allows a range to be configured for multiple consumable resource requirements at the queue level. By supporting a range within which the **RES_REQ** must fall for each consumable resource, a single queue definition can be used for a wider variety of jobs and applications. This allows administrators to significantly simplify the configuration of their environments by avoiding the need to configure multiple queues with different queue level resource requirements.

Support for X11 Tunneling via SSH

Many customers wish to submit X-Windows client applications as jobs and have them display back to their local desktop. With desktop PCs running X-Windows terminal emulators, a common scenario is that users will log into a job submissions host and submit their job from that intermediate host. Platform LSF will then dispatch this job to an execution host, and ideally assuming the job dispatched is an X-Windows client, we would like that job to open a secure *ssh* connection back to our desktop PC rather than the intermediate host that the job was submitted from.

Platform LSF 7u6 specifically addresses this requirement with internal communication mechanisms that set the display variable on the execution host such that the X-Windows client application can display back to the desktop through an X11. This process is transparent to Platform LSF users, but involves a call back to the *res* process on the execution host when a job is submitted with the **-IXF** arguments and *nios* launching an *xagent* process via *ssh* to setup an X11 forwarding tunnel and set the display variable appropriately.

Further Information

Please contact:

Bill McMillan
Director, Product Management
bcmillan@platform.com
Tel: +44 7775 657721 / +44 1256 883041