



Platform LSF[®]

The HPC workload management standard

Add-ons to Platform LSF[®]:

- Platform Adaptive Cluster
- Platform Analytics
- Platform Application Center
- Platform License Scheduler
- Platform Make
- Platform MultiCluster
- Platform Process Manager
- Platform RTM
- Platform Session Scheduler

Other complementary products:

- Platform Cluster Manager
- Platform MPI

The most powerful HPC workload management platform

Platform LSF is the most powerful workload management platform for demanding, distributed and mission-critical high performance computing environments. It provides a comprehensive set of intelligent, policy-driven scheduling features, enabling you to fully utilize all of your compute infrastructure resources and ensure optimal application performance. A highly scalable and available architecture allows you to schedule complex workloads, and manage up to petaFLOP-scale resources.

Customer Benefits	
Maximum Flexibility	<ul style="list-style-type: none">• Heterogeneous platform support• Policy-driven• CLI, web services, APIs
Designed for Extensibility	<ul style="list-style-type: none">• Powerful workload scheduler• Broad application support
Extreme Scalability	<ul style="list-style-type: none">• 100K+ cores, 1.5M+ pending jobs
Industry-Leading Support	<ul style="list-style-type: none">• Large, worldwide development & support team• Extensive partner ecosystem• Nearly two decades of HPC experience

Better, faster, smarter computing

With the Platform LSF product family, you can manage and accelerate workload processing for compute- or data-intensive applications, resulting in faster time to solution. It enables you to intelligently schedule and guarantee the completion of workloads across a distributed, virtualized IT environment and fully utilize all IT infrastructure resources regardless of operating system or architecture.

Most power; best TCO

The Platform LSF product family helps you ensure that all available resources are fully utilized. By delivering maximum computing power, it accelerates the performance of the most demanding, compute- or data intensive applications.

Reduce operational & infrastructure costs

Platform LSF helps reduce total cost of ownership (TCO), by providing optimal SLA management and greater flexibility, visibility and control of job scheduling. This enables IT to improve the service provided to stakeholders. By ensuring optimal utilization of existing IT infrastructures, more work is done with fewer resources, reducing additional hardware and administration costs.

Improve user productivity

Platform LSF enables you to fully utilize hardware resources, whether they are just down the hall or half-way around the globe. By improving utilization, resources are more readily available to users, enabling them to get more work done in a shorter amount of time.

Leverage investments in existing resources

Platform LSF pools computational resources and manages application workloads across highly distributed environments – from single and local departmental clusters to a globally dispersed, multi-cluster infrastructure. It enables you to distribute workload to any mix of hardware systems including desktops, servers and supercomputers.

Deep partner network delivers integrated solutions

Platform Computing has partnered and integrated with more than 100 key industry-leading application vendors including ANSYS Inc., Mentor Graphics, MSC Software, and Schlumberger Information Solutions. Platform Computing's strategic partnerships with Cray, Dell™, Fujitsu, HP, IBM®, Intel®, Microsoft®, Red Hat® and SAS® round out the company's ability to deliver well-integrated, comprehensive HPC management solutions.

Most complete management solution

Platform LSF is the most comprehensive workload management solution for complex, distributed HPC environments. It provides a complete set of workload management capabilities; all designed to work together to address your high performance computing needs.



Platform LSF
Intelligent workload management

Platform LSF provides the capabilities to manage and accelerate workload processing across heterogeneous distributed compute environments. It is comprised of a comprehensive set of intelligent scheduling policies to ensure that the right resources are automatically allocated to the right jobs, for maximum application performance and efficiency.

Platform Application Center
Web portal for end users and administrators

Platform Application Center simplifies HPC by making it easier for users to run applications without having to understand application code. Scripting guidelines and application templates simplify job submission, reduce setup time and minimize operation errors. The web-based interface enables remote job monitoring, easy access to job related data, and the capability to easily perform basic operations like stopping, suspending, resuming or re-queuing jobs through a web browser.

Platform Process Manager
Workflow automation and scheduling

By reducing or removing the need for operator intervention to trigger computational workflows, Platform Process Manager compresses end-to-end cycle time. It also boosts the productivity and efficiency of your HPC infrastructure with tools that automate lengthy, complex, or repetitive tasks that are prone to human error. The result is a more cost-effective, logical, self-documenting solution for workflow automation and scheduling.

Platform Analytics
Comprehensive business decision analysis tool

Platform Analytics is a business decision solution for Platform LSF environments that employs OLAP techniques to analyze long-term historical data from HPC clusters and grids for a variety of purposes useful to the business. It provides timely, complete, and trustworthy analyses to support capacity planning decisions.

Platform RTM
Flexible operational dashboard

Platform RTM is an operational management environment for Platform LSF. Dashboards provide comprehensive reports to support the day-to-day administrative tasks associated with managing single and multiple Platform LSF cluster environments. It provides timely information on the current status of your HPC environment to help improve decision-making, reduce costs and increase service levels.

Platform MultiCluster
Extend resource sharing across many clusters

Platform MultiCluster extends an organization’s reach to share resources beyond a single cluster, to span multiple geographic locations. Local ownership and control is maintained, ensuring priority access to local resources while providing global access across an enterprise’s entire HPC infrastructure. Using Platform MultiCluster you can complete workload processing faster with greater productivity and faster time to results.

Platform License Scheduler
Policy-driven application license optimization

By using Platform License Scheduler, you can manage and optimize application license usage between sites and projects by allocating licenses based on an established distribution policy. You can also optimize performance and sharing where licenses are primarily shared between clusters, and then between projects within clusters. This simplifies license sharing, increasing overall access to license resources.

Platform Session Scheduler
High throughput, low latency scheduling

By creating a virtual private cluster for just your workload, Platform Session Scheduler provides an innovative approach to scheduling large numbers of short duration jobs. Hundreds or even thousands of users can process jobs that are individually comprised of over 50,000 tasks per user with exceptionally low-latency. Whether you have short or long jobs, by using Platform Session Scheduler you can realize the benefits of high-throughput, low-latency task submission without re-architecting applications.

Platform Adaptive Cluster
HPC private cloud environments

Platform Adaptive Cluster turns static clusters and grids into dynamic, shared HPC private cloud computing environments. It allocates resources dynamically based on Platform LSF workload demands. By being able to quickly repurpose servers and application environments based on flexible, workload-aware policies, you can enjoy tangible business benefits. These include eliminating cluster and queue silos and reducing job completion times and power consumption.

Platform Make
Reduce application build times

Platform Make enables developers to distribute and execute makefiles in parallel across multiple processors and hosts. As a result it empowers developers to harness the power of Platform LSF to reduce build and compile times for complex software applications.

Complementary products

Platform MPI

Platform MPI is a fully integrated message passage interface (MPI) library that enables users to build high performance applications while delivering faster performance through optimizing MPI communication and core affinity. It also enables application developers to support a wider range of customer environments while lowering the number of binaries required.

Platform Cluster Manager

Platform Cluster Manager includes all the tools required to deploy, run and manage clusters with unprecedented ease. Leveraging Platform’s leadership in HPC management software, Platform Cluster Manager is built for flexibility and productivity, while offering management, support and real-time software update capabilities.

Challenges and Solutions

Challenge	Platform Solution
Accelerate time-to-market and results	<ul style="list-style-type: none"> High performance architecture improves workload throughput Clustering distributed IT resources into a single virtual supercomputer delivers processing power as needed
Reduce operational & infrastructure costs	<ul style="list-style-type: none"> Maximum utilization rates and system scalability for more power with fewer resources Leverage commodity-based hardware and open industry standards Optimize utilization of application licenses Detailed analytics and reporting to support fact-based project planning decisions
Minimize business risk by increasing system reliability	<ul style="list-style-type: none"> Fault-tolerance and fail-over capabilities ensure always-on computing Compliant with open industry standards
Improve quality of results	<ul style="list-style-type: none"> Failed jobs are automatically re-run Guaranteed service levels and execution



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