

Platform LSF HPC

Intelligent, policy-driven high performance computing (HPC) workload processing

Platform LSF Family of Products At A Glance

- Platform LSF®
- Platform LSF HPC

Complementary Products

- Platform Enterprise Grid Orchestrator™
- Platform VM Orchestrator™
- Platform LSF MultiCluster
- Platform LSF License Scheduler
- Platform Process Manager

Our Solution

Platform LSF HPC is software for managing and accelerating High Performance Computing (HPC) mission-critical workloads. With Platform LSF HPC you can intelligently schedule parallel and serial workloads providing the capability to solve large, complex compute tasks while utilizing the available computing resources at maximum capacity and capability. It enables you to take full advantage of high performance network interconnects available on clustered systems and supercomputers. Platform LSF HPC features patent-pending; topology-based scheduling that enables maximum application performance for industry leading interconnects.

Product Highlights

Topology-Based Scheduling – Consistent, Optimal Performance

Platform LSF HPC is integrated with system vendor and third-party high performance interconnects enabling workloads to be scheduled based on the optimal performance. Topology-based scheduling determines which processors are closest to each other based on the unique characteristics of each high performance interconnect. This guarantees the best application performance and consistency every time an application is run.

IT Resource Virtualization - Self-Managing & Guaranteed Service

With Platform LSF HPC, built on EGO, a virtualized pool of computing resources is created, which orchestrates on-demand and policy-based allocation of resources to create in-effect a large virtual supercomputer that performs optimal resources utilization, maximizing existing hardware investments. The self-adaptive capabilities of Platform LSF HPC reduce administration and associated costs by dynamically allocating resources to adapt to a changing IT infrastructure.

Technically Superior – Robust, Proven & Standards-Based

Core to Platform LSF HPC is Platform Enterprise Grid Orchestrator™ (EGO), the first and only grid platform that will deliver the power of virtualization, automation and the sharing of all IT resources to every application type. Platform EGO provides a flexible and modular way to orchestrate all enterprise application types into a single, cohesive, efficient system.

By de-coupling resource management from workload management, Platform EGO can effectively allocate, prioritize and manage the supply of resources with business policies across the enterprise. This functionality provides organizations the ability to scale up and scale out, while improving application performance, resource utilization and achieving better overall Service Level Agreements (SLAs).

Platform EGO is the by-product of Platform Computing's 14+ years of experience in large-scale distributed HPC computing environments. Platform's experience has been in supporting enterprise-wide large-scale production grids, managing multiple different types of applications on top of an infrastructure shared by multiple design teams, projects or user communities comprised of tens of thousands of CPUs. Given Platform's experience in large-scale distributed and multi-site grids, it is uniquely qualified to generalize this experience to all types of workloads and applications, hence enabling a true enterprise grid.

Platform Enterprise Grid Orchestrator (EGO) enables improved performance and organizational efficiency to:

- Integrate virtualization and business-critical applications into a single, cohesive, efficient system
- Operate a single grid environment across geographically dispersed sites
- Prioritize resource allocation with business policies across all enterprise resources
- Respond to changes in customer and supply chain business requirements in real-time
- Grow and expand by taking a modular approach for application integration

Highlights

- Solve large, complex problems by optimizing High Performance Computing environments
- Complete high priority projects on time by aligning IT resources and business priorities
- More cost effective compute power for mission-critical workloads

Features	Benefits
Optimized Application, System and Hardware Performance	<ul style="list-style-type: none"> Guaranteed, consistent and reliable parallel workload processing with high-performance interconnect support Out-of-the-box parallel integration support Maximized application performance with topology-aware scheduling Finer grain control at the CPU level for parallel tasks with task geometry support Ensured application runtime consistency by automatically allocating similar processors Site configurable scheduling algorithm and dynamic priority formula providing optimal scheduling for parallel and serial workloads with a job-weight scheduler plug-in
Enhanced Accounting, Auditing and Control	<ul style="list-style-type: none"> Track Return on Investments (ROI) and effectiveness of IT investments Full parallel job control and accounting across multiple nodes Determine job termination reasons Out-of-the-box integration with vendor supported accounting packages Fine grain audit information for charge-back accounting
Advanced Self-Management	<ul style="list-style-type: none"> Adaptive to dynamically allocate resources to changing IT infrastructure Resilient workload execution for superior fault tolerance Guaranteed zero downtime with automated fail-over to the nth node
Commercial Grade System Scalability and Reliability	<ul style="list-style-type: none"> Proven, risk-free investment with over 2,000 customers including deployments in over 100+ clusters in multi-cluster environment, 200,000 CPUs and 500,000 active jobs per cluster Running on several of the top 10 supercomputers in the world Directly launch, manage & control 8,192 way parallel jobs IPv6 support including mixed IPv4/IPv6 Full Kerberos Support
Extensive Hardware Support	<ul style="list-style-type: none"> Supports industry leading vendors, supercomputing hardware and clusters – such as Cray, IBM, HP, SGI and Dell Fully integrated with Linux HPC cluster solutions Support for 32-bit and 64-bit operating systems Support for Microsoft Windows, 32 and 64 bit including CCS

Platform LSF HPC Supported Systems

Vendor	Platform	Operating System	Vendor	Platform	Operating System	
Apple	Apple Intel	Macintosh OS X 10.x	Sun	SPARC	Solaris 7, 8, 9 (64-bit)	
Cray	XD1	SLES 9				Solaris 10 (64-bit)
	XT3	SLES 9 Front end		AMD64	Solaris 10 (64-bit)	
HP	HP 9000 Servers (PA-RISC)	HP-UX 11i (64-bit)		Intel X86	Solaris 10 (32-bit)	
	HP Integrity Servers (Itanium2)	HP-UX 11i (IA64)	HP Alpha	kernel 2.4, glibc2.2		
	XC clusters	3.0	Intel X86	RHEL 3, RHEL 4, RHEL5, SLES 9, SLES10, RHEL5 RHEL 2.1, SLES 8		
IBM	IBM Power Series	AIX 5.x (32-bit)	Linux	AMD64, EM64T	RHEL 4, SLES 9, SLES10, RHEL5 RHEL 3 RHEL 2.1, SLES 8 RHEL 4, SLES 9 with BPROC 4 RHEL 4, SLES 9 with RMS 2.82/3.00	
		AIX 5.x (64-bit)				
	IBM Power5	RHEL4, RHEL5, SLES9, SLES10				
IBM BlueGene/L	SLES9 Frontend					
Microsoft	Intel X86	Windows 2000 (SP1 — SP4)		IA64	RHEL 4, SLES 9, SLES10, RHEL5 RHEL 3 RHEL 2.1, SLES 8	
	Intel X86, IA64, EMT64	Windows 2003				
	Intel X86, IA64, EMT64	Windows XP (SP1/SP2)				
SGI	Altix	ProPack4 (SLES 9)				
		ProPack5 (SLES 10)				
	Origin	IRIX 6.5.24 and up				

Supported High Performance Interconnects

<ul style="list-style-type: none"> IBM® High Performance Switch and SP Switch (Federation and Colony Switch) 	<ul style="list-style-type: none"> Quadrics QsNet
<ul style="list-style-type: none"> HP Crossbar Switch 	<ul style="list-style-type: none"> SGI® NumaFlex
<ul style="list-style-type: none"> Myrinet (Including IntelMPI, LAMMPI, MPICH-GM, MPICH-P4, MPICH-SHMEM and ScalMPI support) 	<ul style="list-style-type: none"> SGI® NumaLink

For the latest list of resources supported by Platform LSF please visit:
<http://www.platform.com/Products/Platform.LSF.Family/Platform.LSF/Product.Information/System.Requirements/>