

# Platform MultiCluster 8

The Power of Unified Clusters



## Key Benefits

- Simplifies Platform LSF cluster consolidation
- Improves visibility to global resource use
- Boosts user productivity
- Optimizes utilization
- Reduces complexity

## Key Features

- Distributed HPC resource consolidation
- Automatically load balances workload
- Full visibility to remote compute resources
- Multiple modes of operation
- Flexible compute resource sharing policies
- Resource reservations across clusters
- Preserves Platform LSF job semantics

## Overview

Platform MultiCluster consolidates compute resources from multiple, distributed Platform LSF clusters so that workload can be distributed more efficiently across the datacenter or the globe. Just as Platform LSF enables more effective sharing of resources within a single cluster, Platform MultiCluster extends this concept to enable flexible hierarchical or peer-to-peer workload distributions between multiple clusters.

By ensuring that local administrators retain control of their own assets while still allowing remote clusters to tap idle capacity, Platform MultiCluster promotes efficiency by removing the problem of underutilized compute resources. Unlike competing approaches that require sites to implement an expensive, customized scheduling layer to share resources between clusters, Platform MultiCluster transparently extends existing Platform LSF capabilities, making Platform MultiCluster exceptionally fast and cost-efficient to deploy

## Flexible and configurable compute resource sharing

Platform MultiCluster is the only solution that transparently links multiple Platform LSF clusters into a single shared compute resource while preserving compatibility with existing Platform LSF commands, scheduling options and configuration files. Because it is built specifically for Platform LSF clusters, users can use already familiar commands to submit and manage jobs on any cluster. Configurable sharing policies allow administrators to strike the right balance between sharing compute resources enterprise-wide while respecting resource ownership and the autonomy of local clusters and business units.

## Product Capabilities

### Full Visibility to Global Compute Resources

Platform MultiCluster provides full visibility to remote compute resources and jobs as if they were local. With full visibility to compute resources enterprise wide, organizations can access the systems best able to run their applications, no matter where they reside. This promotes efficiency, saving cost and delivering a better quality of service to all users.

### Multiple Resource Sharing Models

When it comes to compute resource sharing, organizations have diverse needs. To provide flexibility, Platform MultiCluster provides two distinct resource-sharing models: a job-forwarding model and a resource-leasing mode. The job-forwarding model automatically forwards jobs to remote clusters when suitable resources are not available locally. The resource-leasing model enables administrators to selectively share remote resources as if they were local. With granular configuration options and the ability to deploy these models in combination, Platform MultiCluster provides administrators with tremendous configuration flexibility, allowing them to define their own resource sharing policies.

## Promote Sharing while Preserving Local Control

Concerns about control, asset ownership, and preserving local service levels can be major barriers to compute resource sharing. Platform MultiCluster eliminates these concerns by enabling administrators to selectively control the amount of resource they are willing to share. Facilities such as resource leasing and advanced reservations provide access to resources on a pre-negotiated basis. In a job forwarding arrangement, clusters receiving jobs can override submission options, ensuring that administrators can share their idle capacity while still providing ample resource for critical projects on their own local cluster.

## Automated, Secure Job Forwarding and Data Transfer

Platform MultiCluster takes the complexity out of submitting jobs to remote clusters with automated file transfers, user-account mapping and configurable network security. User account-level mapping transparently handles cases where user name space may not be global, allowing jobs to execute

remotely under a guest user ID. Configurable authentication schemes ensure that jobs are accepted only from users that are actually authorized to use cluster resources and that the global environment remains secure.

## Seamless Integration with Platform LSF

Global workload sharing with Platform MultiCluster works just as users familiar with Platform LSF would expect. Resources in remote clusters can be made transparently available to local users without any special knowledge while administrators can continue to use the same commands they always have. Because the standard command vocabulary in Platform LSF is already Platform MultiCluster aware, Platform MultiCluster can be deployed quickly with minimal training requirements. In a Platform MultiCluster environment, administrators have better visibility to remote clusters and more configuration flexibility, making them more productive and freeing up their valuable time to work on other projects.

## Platform MultiCluster 8 Supported Environments

Operating System Support	<ul style="list-style-type: none"><li>• Linux® on x64 architectures including RHEL 2.1, 3, 4, 5.x, 6, SUSE Linux Enterprise Server including SLES 8, 9, 10, 11 and generic Linux distributions using 2.6 or greater kernels with glibc 2.3 (Debian, CentOS, Ubuntu, Scientific Linux and others).</li><li>• Linux on ia64 systems including RHEL 4,5, SLES 9,10,11 and generic Linux distributions using 2.6 kernels and glibc 2.3 and later.</li><li>• Microsoft Windows® on x32 &amp; x64 platforms including Windows 7, Windows Vista, Windows XP, Windows Server 2003 &amp; 2008 standard &amp; enterprise editions, Windows HPC server 2008</li><li>• MacOS X 10.4.x, 10.5.x on Apple hardware</li><li>• HP-UX 11i1, 11i2 &amp; 11i3 on HP hardware</li><li>• IBM AIX 5.3,6 &amp; 7 on IBM hardware</li><li>• Sun/Oracle Solaris 7,8,9 on SPARC 32 &amp; 64 bit, Solaris 10 on SPARC 64 bit and x86_64 systems</li><li>• Cray Unicos/Ic 2.x</li></ul>
Master Host Requirements	<ul style="list-style-type: none"><li>• Minimum 2 GB of physical memory (RAM) recommended</li><li>• Available SWAP space twice physical memory</li><li>• Minimum one high-speed network interface</li><li>• Secondary master host recommended in large clusters</li></ul>
Compute Host Requirements	<ul style="list-style-type: none"><li>• 1 GB of physical memory (RAM) recommended</li><li>• 40 GB of free disk space</li><li>• Minimum one high-speed network interface</li></ul>

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](http://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](mailto: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](mailto:info-europe@platform.com)

### Asia-Pacific

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

