

Platform Adaptive Cluster 7.0.6

Transform static, low utilization clusters into dynamic, high performing resources



Features

- Workload driven dynamic node re-provisioning
- Automated job checkpoints and migration
- Flexible policy controls
- Smart performance controls
- Dynamically switch nodes between physical machines and hypervisors

Benefits

- Optimize utilization, maximize throughput
- Eliminate costly, inflexible silos
- Increase reliability of critical workload
- Maintain maximum performance
- Improve user and administrator productivity
- Increase automation, decrease manual effort

Overview

Platform Adaptive Cluster turns static Platform LSF clusters into dynamic, shared cloud infrastructure. By automatically changing the composition of clusters to meet ever-changing workload demands, more work can be done while improving service levels to give your organizations greater productivity without added cost.

Unlike other less capable solutions, Platform Adaptive Cluster provides the flexibility to automatically provision mixed physical and virtual environments. Your existing investments in hypervisors, management tools and virtual machine templates are leveraged to create a dynamic private cloud environment. With smart policies and numerous features that ensure maximum efficiency, sites running Platform Adaptive Cluster realize improved utilization, better reliability and increased productivity, while reducing administrator workload.

The most flexible HPC cloud software

Platform Adaptive Cluster uniquely supports mixed physical and virtual environments, which are provisioned dynamically based on smart policy controls and workload demand. Physical nodes can be static or repurposed automatically to run native operating systems. By leveraging the templating capabilities of your chosen hypervisor, Platform Adaptive Cluster can quickly and easily provision a practically unlimited set of operating system and application environments, providing the flexibility to quickly and easily deploy multiple different types of application workloads.

Maximum efficiency in HPC cloud management

The efficient scheduling and provisioning of resources are essential to realizing the productivity gains and savings possible with a cloud implementation. Platform Adaptive Cluster incorporates unique features that ensure efficient workload placement, taking into account application resource requirements and the state of the cluster. The ability to intelligently reuse virtual machines for batch jobs with similar resource requirements avoids excessive stopping and starting of VMs, keeping service levels as high as possible. In environments with high demand for different types of resources, configuration settings can control the rate of automated node switching, ensuring that unnecessary thrashing does not inhibit performance and utilization. These and other Platform Adaptive Cluster features combine to maximize throughput and resource utilization, resulting in the highest possible levels of productivity and efficiency.

Product Capabilities

Dynamic node provisioning

Platform Adaptive Cluster can dynamically change the software profiles of compute nodes to meet the requirements of diverse and changing application workloads. Clusters can be comprised of static hosts and dynamic hosts that can be

re-purposed to run your preferred hypervisor or native operating system. As jobs are submitted to the cluster, resource requirements are examined, and the scheduler finds the host best able to run the application. If a suitable resource is not available to run a specific job, Platform Adaptive Cluster provisions an appropriate virtual or physical machine and automatically integrates it into the cluster. Therefore, when using a Platform Adaptive Cluster environment, users are no longer constrained to a specific set of static resources to run their jobs. The cluster can simply flex to meet application demand, reducing user wait-time, maximizing productivity and ensuring more efficient use of the cluster.

Flexible policy controls

Platform Adaptive Cluster extends the rich set of scheduling and sharing policies in Platform's workload managers. You can specify the degree to which resource allocations can flex during periods of high demand, ensuring that minimum resource amounts are always available maintain service levels. Configurable thresholds serve to avoid scenarios where hosts are continuously changing by ensuring that no more than a threshold number of machines are switched per configurable scheduling cycle. These policy controls ensure that service levels are maximized and that cloud resources are used as efficiently as possible.

Automated job checkpoints with VM Save/Restore

Platform Adaptive Cluster eliminates the need to kill lower priority jobs when they are preempted. In addition, suspended jobs will no longer consume resource on the host running higher priority workload. By making jobs on dynamic hosts transparently "checkpointable" using a virtual machine manager's native suspend/resume and migration functionality. This means that pre-empted jobs will not continually need to re-start from scratch and can resume execution where they left off, saving time and valuable resources.

Dynamically switch nodes between physical machines and hypervisors

With Platform Adaptive Cluster nodes running a hypervisor such as Red Hat KVM to be dynamically re-provisioned as a physical machine with a native OS to run particularly resource and I/O intensive jobs. These machines can then be changed back automatically when demand is more efficiency met by virtual resource rather than physical hosts. This is particularly advantageous when running multi-threaded or I/O intensive applications that perform optimally when jobs are run natively on large servers, rather than within the confines of a virtual machine. This ability to switch between physical hosts and hypervisor hosts based on workload profiles and user demand provides yet another dimension of flexibility, ensuring that organizations can maximize their productivity while minimizing unnecessary infrastructure investments.

Fully certified and supported by Platform Computing

Platform Adaptive Cluster comes complete with full certification and support from Platform Computing, the world leader in HPC management solutions.

System Requirements

Operating System Support

- Linux® on x64 architectures including RHEL 4, 5.x, 6, SUSE Linux Enterprise Server including SLES 9, 10, 11 and generic Linux distributions using 2.6 or greater kernels with glibc 2.3 (Debian, CentOS, Ubuntu, Scientific Linux and others)

Platform LSF Compatibility

- Platform LSF 7.0.6 or higher

Hypervisor Compatibility

- RedHat Enterprise Linux AS 5.5, 5.6 with KVM

Platform Computing is the leader in cluster, grid and cloud management software - serving more than 2,000 of the world's most demanding organizations since 1992. 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®, Fujitsu 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

Asia-Pacific

Beijing: +86 10 82276000
Xi'an: +86 029 87607400
Tokyo: +81(0)3 6302 2901
Singapore: +65 6307 6590

