

Librato® Load Manager™

Key Benefits

- Deliver predictable application performance and quality of service to guarantee application SLAs
- Consolidate multiple applications onto fewer servers
- Maximize server utilization
- Set and enforce policies for system resource utilization based on business priorities
- Reduce management costs with fewer OS instances and servers
- Enable application resource usage analysis for:
 - Consolidation planning
 - Capacity planning
 - Chargeback billing

The Challenge

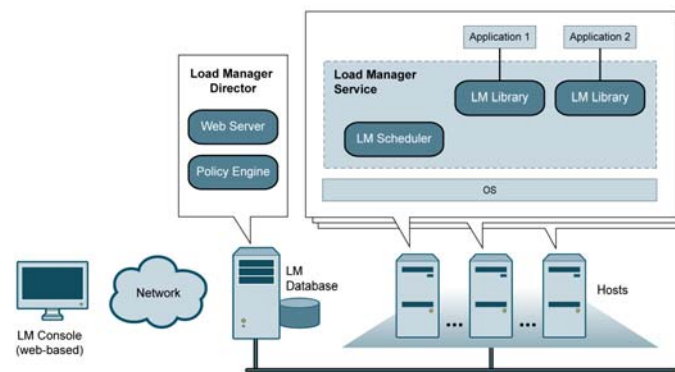
IT managers often face the challenge of having to guarantee a predictable level of service to end users while dealing with ever increasing application workloads. The traditional “one application per server” approach helps meet SLAs for enterprise applications but it usually increases administration and operating costs as resources are provisioned for peak demand, resulting in low utilization rates.

Consolidating multiple applications onto fewer servers helps control costs and increases server utilization. However, this can result in unpredictable performance levels. While advances in virtual machines (VMs) have made server consolidation a reality, VMs are not suited for all application workloads such as I/O intensive applications¹.

The Librato Solution: Load Manager

Librato Load Manager is a dynamic workload management solution that enables fine-grained monitoring and control of major system resources such as CPU, memory, network I/O and storage I/O. As a policy-based fair share scheduler, Load Manager ensures enterprise applications are allocated the system resources required for optimal performance and predictable service levels. By dynamically assigning unused capacity, Load Manager dramatically increases server utilization, enabling IT organizations to get the most out of existing resources.

As a groundbreaking, user space technology, Load Manager requires no modifications to the underlying operating system (OS) or the applications under management. Load Manager is also very lightweight with typically less than 1% overhead.



Load Manager Architecture

Fine-grained Monitoring and Control

Load Manager monitors and regulates major system resource consumption along 4 axes – CPU, memory, network I/O and storage I/O – up to 100 times per second. Using detailed monitoring data – both real time data and historical trends – administrators can understand how applications use system resources and set optimal resource allocation policies, enabling precise control over resource scheduling.

¹ Gartner, "Server Workloads: What Not to Virtualize," G00156214, March 26, 2008: 1

USE CASES

Policy-driven Fair Share Scheduler

Load Manager enables administrators to explicitly allocate shares or weights of major system resources to specific applications. These allocation policies are automatically enforced, making systems easier to manage. More importantly, these policies enable IT organizations to link business priorities and SLAs to resource utilization, ensuring higher priority applications get larger shares of system resources. Load Manager delivers predictable application performance levels and eliminates unexpected resource contention.

Dynamic Resource Rebalancing

Any spare or unused capacity is constantly monitored, managed and automatically redistributed to applications that need more capacity. This assures efficient use of system resources and enables IT organizations to maximize server utilization.

Transparent, Lightweight and Stateless Architecture: Easy and Safe to Deploy

Load Manager is transparent to both the OS and applications, requiring no modifications or re-linking of the applications. Load Manager's lightweight architecture has negligible performance overhead, typically less than 1%.

In addition, the architecture is stateless for fault tolerance such that in the event of failure of any Load Manager software component, applications are unaffected and will continue to run through completion. These features ensure Load Manager can be easily and safely deployed, and seamlessly integrated into existing IT infrastructure.

Centralized Management

In addition to a rich command line interface, Load Manager provides a web-based, easy-to-use, centralized management GUI for configuring allocation policies, creating logical groups of applications and monitoring current and historic workloads with a dashboard of resource usage graphs. This gives administrators a unified view of the data center and a single point of control.



Load Manager Console

Flexible Access to Detailed Monitoring Data

Load Manager collects and stores detailed load monitoring statistics in a database, enabling analysis of system resource consumption for capacity planning, chargeback billing and application profiling. This data can easily be accessed and imported into 3rd party ODBC compliant applications such as Crystal Reports.

Platform Support

Load Manager supports x86 or x86_64 servers running Red Hat Enterprise Linux 3, 4 and 5, SUSE Linux Enterprise Server 9 and 10, and Microsoft Windows Server 2003 and 2008.

Server Consolidation

With policy-driven resource allocation and dynamic rebalancing capabilities, Load Manager allows multiple enterprise applications to safely co-exist on the same servers by guaranteeing the amount of system resources available to each application. This increases server utilization without sacrificing quality of service of these applications while minimizing the number of OS instances, thereby simplifying management of server infrastructure. It also helps alleviate server sprawl.

“Sponge” Unused Capacity

Load Manager ensures no resource goes unused, making 100% server utilization possible. Load Manager allows allocating 100% of system resources to business critical applications as foreground applications. When the foreground applications are not using all of their allocations, Load Manager enables background or sponge applications to safely absorb unclaimed system resources. Load Manager then dynamically reassigns system resources the instant they are needed by the foreground applications. This guarantees SLAs for business critical applications while allowing administrators to run other background applications, thereby maximizing server utilization.



2900 Lakeside Drive
Suite 229
Santa Clara, CA 95054

Tel: 408.588.1716
Fax: 408.588.1776

Support: 866.921.4743
Information: info@librato.com
Sales: sales@librato.com