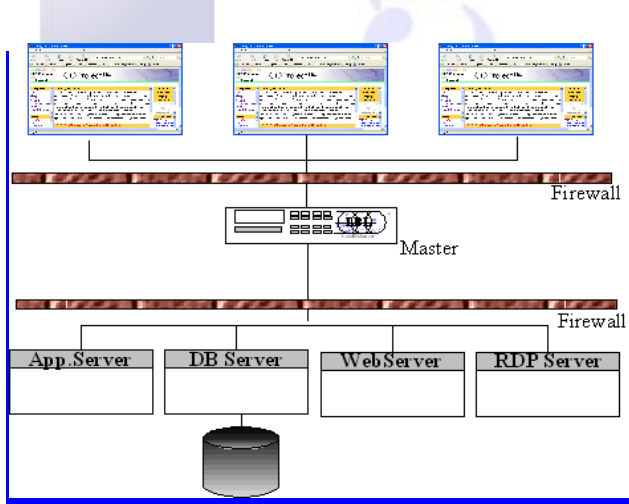
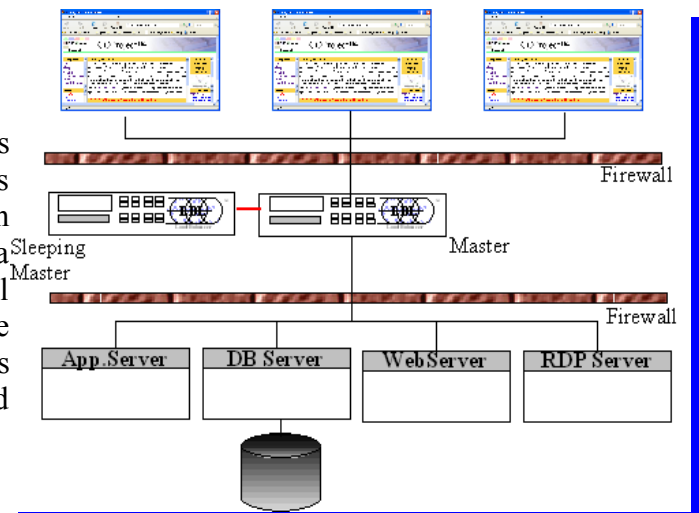


**LBL® Application Availability Infrastructure** is a suite of tools designed to increase application services availability. Heart of the solution is **LBL® LoadBalancer** a software tool for balancing traffic data at layer 4 OSI (port forwarding) and application layer 7 OSI (HTTP/S) with characteristics of session affinity (aka sticky session and load balancers managed session) in capable of providing a highly scalable on modern multiprocessor systems and on systems with multithreaded processors. **LBL® Application Availability Infrastructure** comes from a long experience in many mission-critical projects that contributed the product to achieve typical features of simplicity and reliability in this area. **LBL® Application Availability Infrastructure** includes several products released in commercial distributions: **LBL® Monitor**, **LBL® LoadBalancer Platform**, **LBL® LoadBalancer Standard HA**, **LBL® LoadBalancer Enterprise HA**, **LBL® DNSManager**, **LBL® Surface Cluster Work Flow & Decision Engine**.




**LBL® LoadBalancer Platform** Edition distinguishes itself thanks to its ease of installation and implementation. Its main targets are environments in which there is a high mutability of conditions or in asymmetric cluster environments. This version is particularly suitable for software development in distributed environments because it provides a valuable tool for testing multi-instance servers and the ability to fine-tracking data traffic.

**LBL® LoadBalancer Standard HA** is designed for high availability environments thanks to its “master / sleeping master” architecture in which each balancer instance can keep track of sessions routing data on both nodes. This version allows a theoretical availability index of 99.99% \*\*. These features are realized through the management of virtual addresses (VIP) that allow nodes to take control of failed addresses (IP-TakeOver) in case of a master failure.



**LBL® LoadBalancer Enterprise HA** implements an advanced system for managing parallel service requests that lets you manage multiple nodes both simultaneously and symmetrically (true active-active architecture). This solution is particularly useful in environments where fault tolerance and uptime values of 99.999%\*\* are mandatory. This exclusive feature, unique of its kind, is particularly valid for service and content provisioning, even geographically distributed and mediated thanks to either an upstream round robin DNS or firewall\*\*\*. While guaranteeing Fail-Over capability to the whole system, it preserves global services continuity – Fault-Tolerance - by maintaining overall routing sessions consistency. This



version is extremely scalable and exceeds by far normal routing limits imposed by master / sleeping-master balancer configurations. To achieve the highest performance in managing dynamic name <> address associations, **LBL® LoadBalancer Enterprise HA** can be matched to cooperate with **LBL® DNSManager**, specifically developed to address this kind of requirement.

\*\* Results not typical. Availability depends on several factors, including the hardware architecture, software applications, mission-critical processes and professional services  
\*\*\* The "Round Robin DNS or firewall roundrobin" not part of TCOProject®

<i>Main features</i>		Platform	Standard	Enterprise
<b>End-Points Grouping</b>	Group of services in conjunction with a listener (Layer 4 and 7 - name / port address)	●	●	●
<b>Domain virtualization</b>	Grouping of services based on their belonging to a specific domain pattern (Only Layer 7 HTTP / S - es.dominio: www.papaia.fruit www.ananas.fruit)	●	●	●
<b>URI Path Grouping</b>	Group of services determined by the request URI path (Only Layer 7 HTTP / S - es.URI Path: / mango or / cherry)	●	●	●
<b>Content Rewriting</b>	Feature that allows rewriting HTTP contents: HEADER and BODY	●	●	●
<b>Balancing algorithms</b>	RoundRobin, Adaptive, FailOver	●	●	●
<b>Geographical Balancing</b>	Geographical information of each managed service is specifiable in order to achieve least cost routing (LCR) to preserve bandwidth.	●	●	●
<b>“Out-Of-Order” &amp; “Ready-Again” Service detection</b>	LBL® LoadBalancer can handle offline backend services by removing them from the active services list. As soon as services are back online, LBL® LoadBalancer automatically reinstates them within the list.	●	●	●
<b>Session Affinity &amp; LoadBalancer Managed Session</b>	Session management can be handled in different ways depending on whether you are using Layer 4 TCP or Layer 7 HTTP / S. By balancing Layer 4 traffic, sessions are handled according to TCP client addresses, through a parameter on the listener. If you are handling Layer 7 HTTP / S traffic you can route application sessions in two different ways: • Application-managed sessions (inspection) • Sessions fully generated and managed by LBL® LoadBalancer Layer 7 HTTP / S sessions can be managed by means of cookie, URI parameters and query strings	●	●	●
<b>URL Redirection</b>	In response to a service request URL redirection service is available when backend application results offline. LBL® LoadBalancer is capable of determining the service state by carrying out health check requests to the backend service	●	●	●
<b>HA feature implementation for other services</b>	LBL® LoadBalancer can be implemented for the sole purpose of making virtual addresses available to applications that deliver services on the same host. Services that can be placed in high availability could be DNS, firewall, proxy services, and in general any stateless service. Firewall or DNS system configuration is very straight-forward and can provide low maintenance costs over time.		●	
<b>Balancer virtualization</b>	Balancer virtualization technology permits to run multiple instances of balancers on the same LBL® LoadBalancer Process. This technology facilitates management in complex scenarios where different zones are diversified according to the service offered, or whenever mutual take-over balancing services are required.	●	●	●
<b>Mutual Take-Over</b>	Technology that achieves high availability (HA) on multiple nodes simultaneously active on different services		●	
<b>Logging &amp; Log-rotation</b>	LBL® LoadBalancer is designed to be implemented in unattended environments without service interruption. For this reason the events are historicized per day in text files that are deleted periodically to avoid host resources saturation.	●	●	●
<b>SSL session termination</b>	LBL® LoadBalancer can be used as a SSL session end-point by coupling the listener with digital certificates both at layer 4 (TCP Port Forwarding) and layer 7 HTTP / S. At Layer 7 HTTP / S seamless URL rewriting is provided by LBL® LoadBalancer towards backend services.	●	●	●
<b>SSL Tunneling</b>	At layer 4 (TCP Port Forwarding) SSL tunneling is available.	●	●	●
<b>SSL Re-encryption</b>	Where required for privacy or security reasons, LBL® LoadBalancer Reencryption offers encrypted channels to/from backend services	●	●	●
<b>Protocol correction (Layer 7 HTTP/S)</b>	In some circumstances -especially in presence of 'URL Rewriting' layers- HTTP content header correction is required to assign the correct mime type	●	●	●
<b>Dynamic Address Listen and Network Adapter Translation</b>	Listening and monitoring systems can dynamically manage change of addresses or network cards on which you accept connections	●	●	●
<b>Network Federation &amp; Automatic Role Assumption (Hot Plug Nodes)</b>	Standard and Enterprise versions are able to hot-add additional nodes to preserve service continuity. LBL® LoadBalancer Automatic Role Assumption services (lookup discovery) will permit new nodes to spontaneously federate. This feature allows each LBL® LoadBalancer instance to configure itself in the role of either Master, Sleeping Master or Joint Active by joining the community of existing nodes. In case of failure of one node, LBL® LoadBalancer will make arrangements for the best operating condition by means of its Automatic Role Assumption feature		●	●
<b>Multiplatform availability</b>	LBL® LoadBalancer is a software balancer and is currently certified by TCOProject® on IBM AIX®, Sun® Solaris®, MS Windows®, Linux® (refer to the latest white papers available or the constantly updated "compatibility matrix" available online)	●	●	●
<b>Master/Sleeping-Masters High Availability</b>	This feature is expected to achieve high availability through a master-sleeping master configuration by means of creation and management of virtual addresses		●	
<b>Active-Active High Availability</b>	Balancing configuration where all nodes simultaneously contribute to the service erogation through the mediation of an upstream "DNS roundrobin" ***			●
<b>Web Console</b>	Remote administration console with secure communication via HTTPS (SSL) and user authentication	●	●	●
<b>External event communication</b>	Run-time interaction for each service, service groups or overall balancing feature	●	●	●
<b>Custom process management plugin</b>	Custom process life-cycle management capable of interacting with the overall system balancing structure	●	●	●
<b>Cache control</b>	This feature allows clients and proxies to exploit RFC2616 to control systems cache responses where required (RFC2616: Cache-control Mechanisms, Cache-Control and in particular cache-response-directive).	●	●	●
<b>Historicized stats on relational DBs</b>	The historicization of statistics are stored in a relational DB. LBL® LoadBalancer can work in the following scenarios: Embedded DB (default), Networked DB, Centralized Networked DB (Standard and Enterprise). If installed on a Networked DB, operators can inquire the Relational DB by means of the SQL query language within the entire database. The maximum temporal depth of logs data is controlled by user-definable parameters. LBL® LoadBalancer WebConsole provides common queries that cover most of the needs in daily traffic and system usage monitoring.	●	●	●
<b>Event notification via e-mail &amp; HTTP</b>	Facing certain events, LBL® LoadBalancer can notify administrators via email boxes or HTTP Web services. This feature permits straight-forward integration in existing monitoring systems within the datacenter	●	●	●
<b>Courtesy messages</b>	Courtesy messages can be set up to be shown to operators when services are placed in maintenance or temporarily out of order. This feature is very useful in enterprise organizations to clearly warn that the service is temporarily under maintenance and thus preventing helpdesk calls congestion.	●	●	●
<b>Business Continuity &amp; Disaster Recovery</b>	LBL® LoadBalancer is designed to distribute traffic data loads and to achieve high availability even in scenarios that include one or more sites for disaster recovery. Due to its hierarchical node management, LBL® LoadBalancer can be configured to handle both manual and automatic DR site take-over. This feature is already included in the product and requires only the additional licenses for the extra nodes located on the disaster recovery sites. LBL® LoadBalancer together with LBL DNSManager® allows you to control traffic on a geographic scale as well by making services available even on sites not having preferential communication channels between them.		●	●

\*In the pictures, LBL®LoadBalancer is depicted by a physical system for presentation simplicity. LBL®LoadBalancer is a software product.

\*\*Theoretical results. Availability is determined by various factors, among which hardware architectures, backend software applications, mission-critical processes and professional services

\*\*\*Round Robin DNS & FireWall are not part of the TCO Project® solution

Sun & Java are registered trademarks of Sun Microsystems, inc. MS Windows, Linux SUSE & RedHat are registered trademarks of their respective owners.

LBL & TCOProject are registered trademarks, all rights reserved. Contents within this document can change without notice.