

Local Load Balancing

Open Enterprise Systems
Internet Technologies Group

Introduction

This document will discuss a technique, which Open Enterprise Systems (OES) employs, to utilize a redundant web server cluster on an LAN to deliver one seamless service. This technique fulfills the following usage requirements:

1. **High Throughput.** The cluster is able to service requests at a higher rate than one server alone. The throughput of the cluster should scale linearly with the number of servers in the group.
2. **High Availability.** If a server experiences a problem (for example, a software or hardware crash), the machine will be removed from the availability pool, and all further requests for service will be routed to the other servers.
3. **Flexibility.** This system must work with all major TCP/IP services, including web services (including HTTP and HTTPS), email (including SMTP, POP3, and IMAP), and databasing (including LDAP, ODBC, and any databases that support replication, such as Oracle, PostgreSQL, and MySQL).
4. **Transparency.** Some services, such as web shopping with cookies, require one server to handle a user transaction persistently. For example, shopping carts using cookies and SSL have this requirement. The service must be able to persistently deliver a user request to the same server if necessary.

The most effective method to meet these criteria is using a smart router to dynamically route new incoming requests to the available servers in the pool.

Load-Balancing Routers

Load-balancing routers can be implemented various ways, and at various levels of the networking protocol stack. OES suggests using the Linux Virtual Server (<http://www.linuxvirtualserver.org/whatis.html>) (LVS) technology, which routes data at the TCP level.

Using LVS satisfies all of the stated requirements:

1. **High Throughput.** LVS clusters are limited by the available LAN bandwidth (usually 100 Mbps), which is fast enough for even the highest loads seen in the industry.
2. **High Availability.** Open source technologies used with LVS are capable of monitoring and smartly handling problem events which happen to the server cluster. For more information, please see the LVS web page about high availability (<http://www.linuxvirtualserver.org/HighAvailability.html>).
3. **Flexibility.** Since LVS operates at the TCP/IP level of the networking stack, all TCP/IP-based services will be compatible, including web services (HTTP and HTTPS), email (SMTP, POP3, and IMAP), and databasing (LDAP, ODBC, Oracle, PostgreSQL, and MySQL).
4. **Transparency.** Persistent connections are a design feature of the LVS system. LVS has well-established support for persistent connections for protocols such as FTP and HTTP. For more information, please see the LVS documentation about persistence (<http://www.linuxvirtualserver.org/docs/persistence.html>).

Furthermore, LVS provides the following advantages over alternative products or methods:

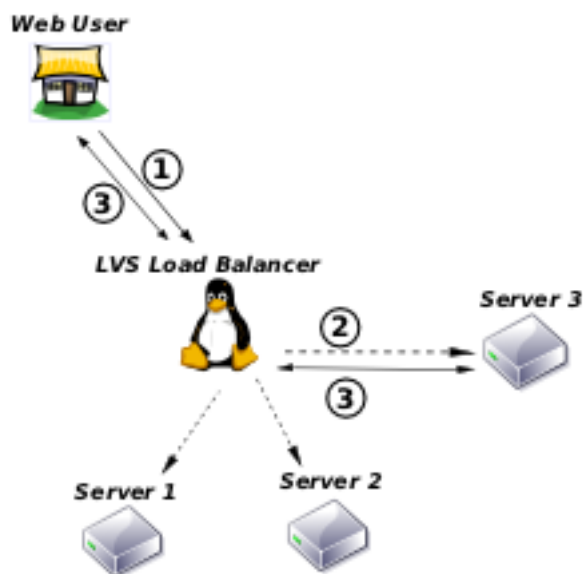
- **Compatibility.** LVS routers are compatible with servers using any operating system or software based on TCP/IP networking. This includes every major Internet software product, such as Linux, Unix, and Windows operating systems and their associated applications.
- **Freedom.** LVS is free software. As part of the Linux operating system kernel, you can be sure that your technology will always remain under your control. Furthermore, since the LVS source code is available to you, you may customize LVS to any degree to fit your needs.
- **Cost.** LVS is available at no cost. However, the LVS technology is industry-proven. LVS is used in TurboLinux's Cluster Server, and Red Hat's Linux Clustering Solution; and LVS is used to deploy several of the Internet's high-profile websites (please see the deployment page (<http://www.linuxvirtualserver.org/deployment.html>)).

Method

OES implements local load balancing using LVS plus additional monitoring tools to deliver a full management suite.

Example

Figure 1. Hypothetical Service Transaction



Hypothetical Service Transaction

1. The first time a user accesses the web site, she will access the load balancing LVS router.¹
2. The LVS router will direct the connection to the least busy available server.

3. The server will perceive the connection as a standard Internet connection and it will service the connection as usual. In this example, all communications are transparently routed through the load balancer.

Redundant LVM Routers: In practice, multiple LVS servers are used in a failover mode for redundancy; however, the conceptual protocol is the same.

OES Service

Open Enterprise Systems provides this service to customers worldwide. We offer two approaches for our customers to use this system.

Remote Collaboration

OES will work with your organization to establish a proper hardware specification for all servers used. When this is complete, international customers may purchase this hardware, either independently or through OES.

After the hardware phase is complete, OES will regularly deliver an installation CD-ROM image with all software. Since we have agreed on a hardware platform, booting from the CD-ROM will initiate a hands-free process which installs all necessary software. In addition, this CD-ROM may be used as an emergency recovery disc.

Remote Training and Assistance

Alternatively, OES is prepared to deliver all of our training materials to your organization. These materials are distilled from our experiences deploying load balancing solutions in a variety of environments. They include our recommendations for various common server and network designs.

Additionally, OES will make itself available for regular VoIP voice collaboration sessions, wherein we examine the current status of the project and work together toward its completion.

Notes

1. This is the initial TCP SYN packet requesting service.