

Load Balancing Hardware

Deal With Heavy Loads On Web Servers, Application Servers & Network Traffic

by Paul Ferrill

WHILE SOME COMPANIES would love to have the problem of heavy traffic on their Web site, those that do have to either deal with it or lose business. Nothing turns Internet shoppers away faster than a slow Web site. The same goes for a company-wide application that takes forever to return the information you're looking for.

Load balancing is a generic term used to describe a combination of hardware and software designed to distribute network or application traffic evenly across the available resources. Choosing the right hardware depends a lot on what problem you're trying to fix. While most of the vendors will sell you a system that will help with a pure network or Web traffic problem, they all take a slightly different approach.

Understanding the true nature of the problem will help to determine a solution that will work. For example, the symptom may be slow response time to a product search on the company Web site. In reality the root problem could be in either a database bottleneck or an undersized application server.

Most high-volume Internet sites utilize some type of clustered server architecture for reliability, performance, and scalability reasons. A load balancing hardware solution will sit between the Internet and the server cluster in order to evenly distribute the traffic to the different Web servers. In most cases it will also route traffic

around a failed cluster node to provide that extra level of reliability.

Different Problems, Similar Solutions

While these network-related problems may seem different, the solution is often a similar approach. For pure capacity issues such as an undersized server or poor performing database, you can often fix the problem with better



hardware. Load balancing can help make better use of available resources by intelligently choosing where to handle the processing.

The OSI seven-layer model has been used for years to describe how networks interact at different layers. The seven layers are:

1. Physical
2. Data Link
3. Network
4. Transport
5. Session
6. Presentation
7. Application

A network load balancing solution would operate at layers 2, 3, and 4 while a purely application balancing approach would operate at layer 7.

Cisco Systems

Cisco Systems has made its name on building networking infrastructure hardware that keeps the most demanding networks up and running. The Catalyst 6500

system is Cisco's highest selling product. It was designed to deal directly with problems at layers 2, 3, and 4. Built on a modular chassis, the Catalyst 6500 is expandable to handle everything from switching to firewall services to SSL and IPSec services.

The Cisco CSS 11000 is a hardware-based load balancer that uses a similar modular approach to the Catalyst. You can start out with a basic configuration and add additional modules as the need arises. With other modules you can add hardware SSL offload and acceler-



SNAPSHOT

Cisco

\$9,995 to \$39,995

Cisco's CSS11000 family of products, as well as an add-on module to the Catalyst 6500 data center switch, directly addresses server load balancing

www.cisco.com

ation along with other application acceleration capabilities.

Cisco's expertise in the network infrastructure game makes it uniquely qualified to build products that understand the flow of traffic, be it TCP or UDP. Its products are architected for growth and flexibility. The key to buying from Cisco is that it is a one-stop shop for all your networking infrastructure needs.

F5 Networks

The BIG-IP line of products from F5 addresses both the application acceleration problem (or layer 7) and the network traffic management problem. In addition, BIG-IP helps address the ever present threat of Web attacks such as denial of service and

application attacks. F5 pioneered the concept of application traffic management with the BIG-IP product line. Version 9 of the BIG-IP systems adds a suite of application acceleration and optimization features. F5's Traffic Management Operating System architecture is designed to increase network effi-



SNAPSHOT

Coyote Point

E350si: \$5,995

Delivers layer 7 functionality to differentiate among users, letting you direct traffic to the right server and handle the most demanding ecommerce application; the Equalizer also functions as a failure monitor and will automatically switch traffic to a different server when it detects a problem

www.Processor.com/CoyotePoint2

ciency and control while improving end-user performance. Support for a broad range of load balancing algorithms, HTTP compression, and the ability to avoid unavailable servers help make the BIG-IP family a market leader.

Coyote Point

Coyote Point believes there is demand for this type of equipment for the small to medium-sized enterprise. As more and more companies rely on the Internet for



SNAPSHOT

XINCOM

XC-DPG503: \$299

Cost-effective load balancing solution targeted at the small to medium-sized enterprise or for a branch office in a larger organization; leveraging commercially available broadband connections makes this product a real cost saver for small IT budgets

www.Processor.com/Xincom

both outward facing customer-centric Web sites and inbound employee and business partner traffic, it's easy to see where they might benefit from the same hardware solutions as the big players.

Coyote Point offers the Equalizer line of products focused at both the SME and higher-end markets. The E350si is its hottest selling product, delivering a number of features for an affordable price. The Equalizer's layer 7 functionality makes it possible to differentiate among users to direct them to the appropriate server. Other supported functionality includes URL parsing, persistent cookies, and a virtual cluster technology that makes it possible to assign pools of servers to the most heavily used applications and services.

Coyote Point's Envoy product is a geographic load balancing add-on for the Equalizer. Envoy allows any Equalizer to cooperate with its peers to enable intelligent request distribution across a geographically distributed server cluster. Traffic is routed to the appropriate region based on the DNS address of the requestor.

XINCOM

The XINCOM XC-DPG503 is targeted at small businesses that

are looking to get the most out of any combination of two broadband connections, including cable, DSL, or T1. It would also appeal to companies with multiple branch offices seeking to establish a

secure connection using available network connectivity. Using two XC-DPG503 boxes connected over the Internet, you can establish a dual-redundant VPN to securely link the two networks together.

From a load balancing perspective, the XC-DPG503 will let you use both WAN ports simultaneously for increased bandwidth. XiNCOM uses a session-based load balancing approach, meaning it operates at layer 5 of the OSI model. An auto crossover feature makes it simple to connect to an existing network. Other features include advanced Network Address Translation configurations, a built-in DHCP server, stateful packet inspection firewall, and more.

Remote monitoring and management makes it possible to control the box from anywhere on the Internet as long as the box has connectivity. The Web browser-based management console is easy to use and even provides a firmware update option. The XC-DPG503 also supports dynamic DNS in the

Buying Tips

Here are some things to consider when you're purchasing load balancing hardware.

- Try to identify the problem you want to address as specifically as possible, even down to the individual user or application.
- Gather the necessary data to justify your investment and make it as detailed as possible. This could include employee labor hours saved or an increased peak Web traffic capacity.
- Plan for a pilot test project. Some vendors will even provide hardware for you to try out on a loaner basis.
- Understand your budget realities. Don't propose a solution that you know will never make it past the first manager who sees it.

case where your ISP doesn't provide a fixed IP address.

The Right Fix

Load balancing hardware can help improve many performance-related issues. The key is in understanding the nature of the problem and then applying the right fix. Many of the vendors mentioned here also provide customer support engineers to help identify the right product to meet your needs. All you have to do is ask. ■

To see Processor Showcase advertisements for products in this category, look in the Solution Index on page 2 and find the Load Balancing ads under Communications/Content Networking.

