

Deploying Load-Balanced Server Clusters with Cobalt RaQ and Coyote Point Equalizer

Web hosting customers expect their sites to be up and running 24x7 with rapid response. They also expect scalable servers that transparently meet growing user demand. Deploying ever-larger web servers to meet these customer expectations is an expensive, short-term approach, doomed to eventual failure. Top-tier ISPs are finding that combining load balancing with clusters of low-cost, high-density web servers is a far more cost-effective, flexible, and reliable strategy. Cobalt Networks and Coyote Point Systems have joined forces to meet ISP needs by marketing an integrated solution for deploying load-balanced server clusters.

What are Server Clusters?

A cluster is a group of physical servers with identical content, networked together to form a single *virtual server*. Requests directed to the virtual server can be handled by any physical server in the cluster. Clustering enables transparent growth: physical servers can be added without externally-visible network changes. Clustering improves fault-tolerance: a physical server can be taken down for maintenance or repair without resulting in a service outage. A server cluster can exhibit high-availability and throughput characteristics far beyond those achievable by even the largest single server.

The most cost-effective way to build scalable server clusters is to deploy low-cost, high-density physical servers. The Cobalt RaQ™ series of Internet server appliances are an ideal choice. The RaQ offers turn-key virtual hosting in a rack-mountable enclosure designed for high-density installation: up to 80 can be accommodated in a standard equipment rack. Starting under \$1000, each RaQ can handle millions of web page requests per day at a fraction of the cost associated with traditional Unix or Windows NT web servers. The RaQ arrives pre-configured with Linux, the Apache web server, CGI and Perl scripting support, Front Page server extensions, Sendmail, DNS, FTP, and SSL support for full-featured web hosting. Simple browser-based administration enables "while-you-wait" provisioning of new customer accounts. Dedicated servers can be deployed for high-end customers with premium service demands, while shared servers can support up to 200 customer sites using virtual domains. High performance options like dual 10/100 Ethernet and SCSI interfaces further increase the RaQ's versatility.




Figure 1. Cobalt RaQ Server Cluster

Figure 1 illustrates a cluster of RaQ servers that have been deployed by an ISP to meet growing user demand and transaction volume at a reasonable cost.

What is Load Balancing?

DNS can be used to redirect requests to physical servers in a round-robin fashion. But simple round-robin DNS cannot exploit more powerful servers, react to current conditions, or avoid unavailable servers. For a server cluster to achieve its high-performance and high-availability potential, *load balancing* is required. Load balancing optimizes request distribution based on factors like capacity, availability, response time, current load, historical performance, and administrative weights. A load balancer sits between the Internet and a physical server cluster, acting as a virtual server. As each request arrives, the load balancer makes near-instantaneous intelligent decisions about the physical server best able to satisfy each incoming request. A well-tuned adaptive load balancer ensures that customer sites are available 24x7 with the best possible response time and resource utilization.

Coyote Point's Equalizer is a premier high-volume, low-cost load balancing solution. Starting at less than \$4000, the Equalizer 250 balances up to 64 clusters of 8 servers each, supporting 64,000 simultaneous connections at T1 rates. The Equalizer 450 can handle an unlimited number of 64-server clusters and 4 million connections, with an aggregate bandwidth of 100 Mbps. All Equalizers are rack-mountable and can be deployed in a hot-backup configuration for maximum reliability. Designed to meet the extreme demands imposed by heavily-loaded, mission-critical web sites, the Equalizer can handle over 130,000 HTTP GET operations per minute and includes security features that resist denial-of-service attacks. In addition to HTTP, the Equalizer can also balance email, news, and FTP traffic, and supports "sticky connections" required to efficiently handle Active Server Pages and SSL. Active content verification uses a configurable end-to-end request to ensure that target applications are fully operational on each active server, thereby avoiding failures that might go undetected by other load balancers.

Another essential ingredient in a load balancer is manageability. Administrators can use the Equalizer's browser-based GUI to define and monitor virtual server clusters. Clusters are configured with a wide variety of built-in load balancing algorithms. Cluster summaries provide an instant view of server status. Hit rate, service time, and computed load statistics are readily available at both cluster and server level, presented in an easily-understood graphic format. Custom event handling can even be configured to react automatically to critical events like physical server failure.

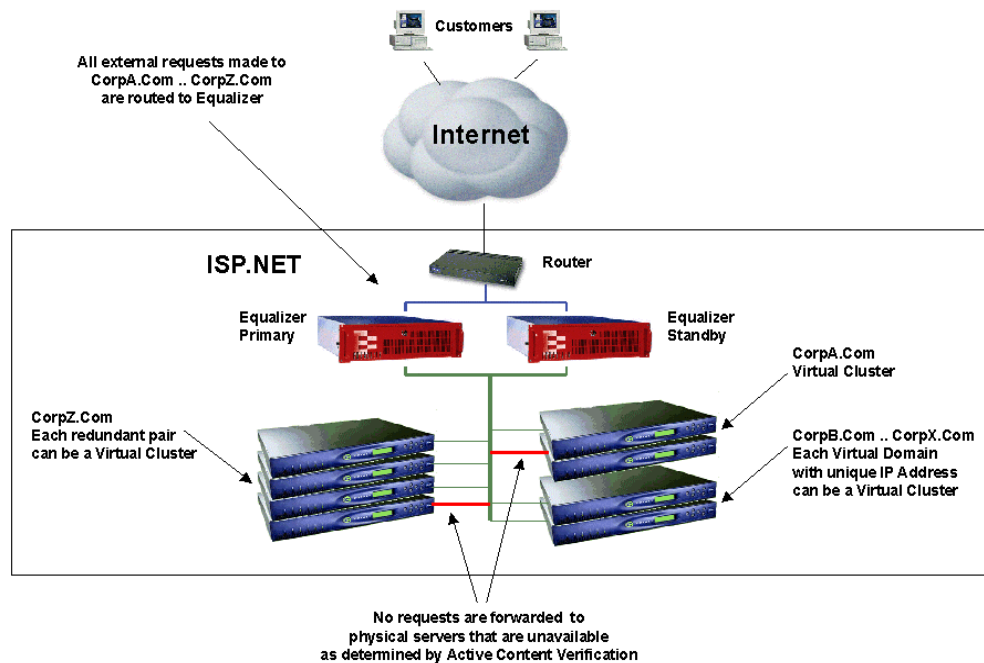


Figure 2. Coyote Point Equalizer balancing traffic for Cobalt RaQ Server Clusters

Figure 2 illustrates a redundant pair of Equalizer load balancers deployed in front of a RaQ server farm to support very high availability, very high volume traffic. Note that each high-end customer web site can be managed and monitored as a separate virtual server cluster. Each load-balanced cluster is also well-positioned for transparent growth as traffic volume increases.

When to Use Load-Balanced Server Clusters

A load-balanced server cluster is a cost-effective strategy for any enterprise or ISP that requires:

- scalable growth beyond more than a handful of stand-alone single servers, or
- robust, reliable "zero down time" web hosting performance.

A product like the Cobalt RaQ is less expensive than even a low-end UNIX or Windows NT server, both in start-up and administration costs. Increasing the throughput of a customer site by wholesale replacement of a small NT server with a larger NT server is time-consuming and expensive. Increasing throughput by adding another RaQ is, by comparison, trivial. If a load balancing infrastructure is already in place, all RaQ resources can be utilized to their fullest potential, further increasing return-on-investment. For example, four RaQ's clustered behind an Equalizer 250 could easily support 64 high-end customers at an entry cost of less than \$8,000 -- that's just \$125 per customer.

The value of high-availability is more difficult to quantify, but is no less compelling. High-end customers expect – and pay for – 100% availability, 24x7, with zero down time as perceived by the end user. Each "server down" response may well represent a lost sale. ISPs who do not offer robust, reliable hosting services ultimately experience loss of revenue through inability to support enterprise-class, mission-critical web sites. A load balancing product like the Equalizer offers high-availability at an entry cost comparable to a single well-equipped NT server. Furthermore, the Equalizer includes top-notch performance monitoring and 60-day 5x8 technical support at no additional cost. Redundant Equalizer 250s can be installed for less than \$8,000, while competitor products start at \$35,000, including installation fees.

How to Deploy Load-Balanced Server Clusters

When deploying a load-balanced cluster of servers, the following steps must be taken.

Network Planning: A load balancer like the Equalizer typically operates as a dual Ethernet device that routes between two subnetworks: the external network and the server cluster network. It is also possible to use the Equalizer with one Ethernet in a single subnetwork. Designing your network, identifying the external and internal subnets, and defining the IP addresses to be used as virtual cluster addresses is the first step in deployment. You may also need to modify default gateways and routes so that traffic to and from your server cluster is routed through the Equalizer.

Web Server Deployment: As compared to a traditional NT or UNIX server, the RaQ is relatively simple to install. All server software is pre-configured; simply use the RaQ's intuitive, browser-based interface to define a virtual domain for each customer site. Virtual domains can share a single IP address, or a single RaQ can support several IP addresses (one per virtual domain). The ISP or customer administrator can create or modify user accounts and services for each virtual domain. Dedicate an entire RaQ to a high-end customer or support up to 200 customers with a shared RaQ: the choice is yours.

Server Content Replication: It is imperative that all redundant servers in the same cluster host identical content. This can be accomplished by using NFS to remotely mount a shared file system on redundant servers, or by using rdist / rsync to synchronize replicated data stored in separate file systems. Another method of content replication is to stage new content on a single server, then "push" content to redundant servers at the same time, using custom distribution scripts or commercial web publishing software.

Load Balancer Deployment: Install the Equalizer by connecting it to your external and internal subnetworks, in accordance with your network plan. If deploying redundant Equalizers, simply connect both – the Equalizers will communicate with each other and automate failover from primary to standby as needed. Coyote Point will preconfigure the Equalizer with IP addresses for your network at no extra cost. Thereafter, administration involves using a browser to create and monitor virtual clusters.

Virtual Cluster Administration: Each virtual cluster is identified by its external IP address and port (the address associated with the customer's virtual domain name.) Add physical servers to the cluster by internal (RaQ) IP address and port, assigning a static weight to each. All servers in a homogenous cluster should have the same weight, but more powerful servers should be assigned a higher weight. Configure load balancing options for each virtual cluster, choosing between round robin, static weight, and adaptive. When using adaptive load balancing (the default), performance indicators include server response time, active connection count, and server agent response. Load balancing responsiveness, optimization threshold, and sticky seconds may also be specified to fine-tune the Equalizer's behavior.

Tune Load Balancing Algorithms: Once this product combination has been deployed, performance should be monitored. Adjust static weights and algorithms in order to achieve optimum utilization of all resources. Increasing capacity is as simple as adding another RaQ. Traffic may be skewed towards larger RaQs or away from smaller RaQs by adjusting static weights. Results are viewed through the Equalizer's powerful monitoring interface.

Adding Geographic Load Balancing: In large installations, geographic load balancing may be introduced by adding Coyote Point's Envoy to the Equalizer. Customers who require web presence on a global level may benefit by deploying geographically-distributed server clusters to improve response time, reduce WAN bandwidth consumption, and serve up regionalized content. Envoy costs just \$2,495 per site.

The Bottom Line

Combining adaptive load balancing with low-cost, high-density servers is a cost-effective, flexible strategy to provide high-availability, high-performance web hosting. The Cobalt RaQ and Coyote Point Equalizer together represent an excellent turn-key solution for ISPs seeking to implement this strategy.

Prepared by:
Lisa Phifer
Core Competence, Inc.
lisa@corecom.com