Transparent Web Cache Redirect

Transparent Web Cache Redirect

ABSTRACT

"Increase the efficiency of data networks, and the world will beat a path to your door." Redundant requests for web content are a fact of life. This makes the caching of web page content an important element of network design. This paper specifies the advantages of using Transparent Web Cache Redirect on a Riverstone router to improve the efficiency of networks that carry world wide web traffic.

WEB CACHE BASICS

How does web caching work? The Web Cache Server is a key part of any caching solution. It may reside directly in the path to the Internet, monitoring all traffic, or to the side with only HTTP traffic redirected to it (see Figure 1). When the Web Cache Server sees the first request for a page, it sends the request to the intended site, and caches the reply. The next time the Web Cache Server sees a request for that page, it traps the request and responds to the client with the cached page, keeping the traffic local.

There are many different Web Cache Server vendors, and many ways to cache web traffic. Most Web Cache Servers set time limits on how long an object is stored locally. Regardless of which method is used, Riverstone Routers enhance the network's Web Cache Server.

When the Web Cache Server sits on the route to the Internet, it sees all traffic going to and from the Internet. This solution is transparent to the clients on the network. There is no need to configure the clients to send their traffic to the Web Cache Server. However, this solution processes other forms of traffic in addition to HTTP, which may cause performance degradation. It also introduces an unnecessary point of failure. For these reasons, this configuration is not recommended.

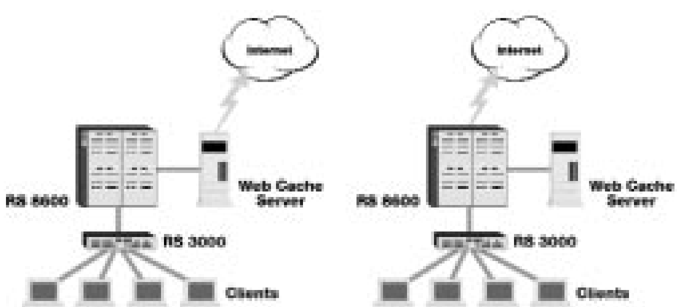

Figure 1: Direct or Side Configuration for a Web Cache Server

A better alternative is to have the Web Cache Server on the periphery of the network, with only HTTP traffic directed to it. This removes the bottleneck caused by other network traffic. However, this approach means that all clients on the network must be configured to send their HTTP requests to the Web Cache Server. It is complex to configure every client on the network to send HTTP requests to the Web Cache Server, and all other Internet traffic through normal paths. Aside from the onerous task of configuring these systems, some of these Web Cache Servers can only support a limited number of simultaneous connections. Once the limit is hit, new sessions are dropped. If the server goes down, HTTP traffic stops; obviously not ideal.

The solution is to use Transparent Web Cache Redirect.

TRANSPARENT WEB CACHE REDIRECT

Riverstone Routers address these issues with Transparent Web Cache Redirect. Instead of reconfiguring all the clients, the network administrator needs only to configure the RS Router. With a few commands, the administrator can transparently redirect all (or selected) HTTP traffic to a Web Cache Server.

Riverstone Routers can be configured to keep track of the number of open connections and compare that connection count to the maximum number of connections that the Web Cache Server can support. Once the maximum is reached, new traffic is sent directly to the Internet, bypassing the Web Cache Server until the number of sessions falls below the threshold.

ADVANCED WEB CACHE REDIRECT FEATURES

When some Web Cache Servers receive a request for a page that isn't cached yet, they make a request to the website using the Web Cache Server's IP address as the source address. Some websites use the source IP address for accounting purposes or for access control. Since some Web Cache Servers modify the source IP address, the website may see this as an invalid address, and respond with an error message. The Web Cache Server would then incorrectly cache the error message for future requests. Riverstone routers can be configured to route traffic directly to those sites, not through the Web Cache Server.

There may also be instances when a proxy server is used for HTTP requests. It may be configured to receive requests on a port other than port 80 for HTTP traffic, such as port 8080, but it would send its HTTP requests to the Internet using port 80. In this situation, the RS Switch Router can be configured to redirect HTTP traffic on ports other than port 80. The RS Router sits between the proxy and the clients. As traffic is sent to the proxy, the RS Router redirects it to the Web Cache Server, which then processes the packet normally.

Some sites may be concerned that one Web Cache Server is handling all the traffic. Riverstone routers solve this problem. The system administrator can organize a group of Web Cache Servers into a pool, and the RS routers balance the HTTP traffic over the pool of Web Cache Servers.

SAMPLE CONFIGURATIONS

The Web Cache Servers and the RS Routers can be located either at the enterprise or the service provider. They can also be at both locations. The machines at the service provider would cache service provider traffic, while the machines at the enterprise would cache traffic specific to that enterprise.

The enterprise network administrator frees up bandwidth by reducing web traffic. This optimization reduces the need for additional bandwidth. The service provider has less redundant traffic from enterprise customers. In both cases, there is an immediate improvement in the available bandwidth and performance.

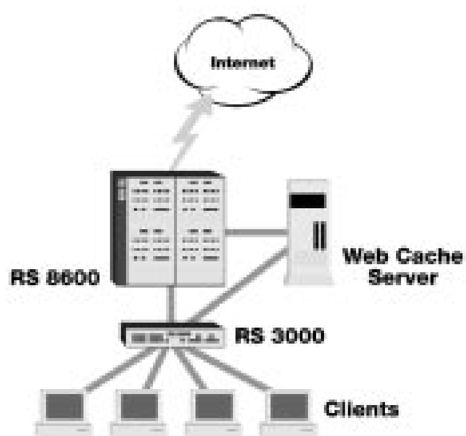

Figure 2: Configuring an RS Switch Router and a Web Cache Server

This is best illustrated through examples. Suppose an enterprise manager has purchased an RS 8600 and a Web Cache Server. The Web Cache Server has an IP address of 186.89.10.51, and access to the Internet is on interface ip1 (see Figure 2).

The selected Web Cache Server supports 10,000 active connections. There is no HTTP proxy server and only one Web Cache Server installed. Some of the users access www.riverstonenet.com at 128.121.123.251. The site checks the source address to confirm that the client is allowed access. While in the configuration mode on the RS Router, the following commands would be entered:

web-cache c1 create server-list s1 list 186.89.10.51

web-cache c1 set-maximum-connection s1 10000

web-cache c1 create bypass-list list 128.121.123.251

web-cache c1 apply interface ip1

These commands create a Web Cache Redirect with the name c1 and a server list associated with it by the name s1. The server list s1 has only one Web Cache Server entry:

186.89.10.51

The set-maximum-connection command sets the total number of active connections.

The bypass list allows traffic destined to address 128.121.123.251 to be sent directly to the Internet, and not redirected to the Web Cache Server. It is also possible to apply an ACL profile on the traffic, instead of a specific address (or list/range of addresses). You do this by replacing list 128.121.123.251with profile p1, where p1 is the name of the profile to apply. Load balancing is achieved by putting a list of the IP addresses of the Web Cache Servers within quotation marks. The example below shows how to configure the Riverstone router if there were two Web Cache Servers in the s1 list and traffic with port 81 needed redirection as well. The commands are entered when the RS Router is in configuration mode.

web-cache c1 create server-list s1 list 186.89.10.51 186.89.10.52

web-cache c1 set-maximum-connection s1 10000

web-cache c1 set http-port 81

web-cache c1 create bypass-list list 128.121.123.251

web-cache c1 apply interface ip1

CONCLUSION

These simple commands provide dramatic performance improvements on the network.

The RS Router enhances the power of a Web Cache Server, and delivers the improved performance without making you reconfigure every client on the network. The transparent redirect feature of the RS Router Family simplifies the implementation of a Web Cache Server. By integrating redundancy and override features, the redirection feature alleviates concerns over Web Cache Servers becoming single points of failure on the network.

A Web Cache Server also improves response time by bringing data closer to clients. This increases available WAN bandwidth by reducing redundant traffic. When this is added to other features found in the RS Product Family (such as VRRP for redundancy, QoS for traffic shaping, IP policy routing for security, and traffic management at wire speed), the administrator sees a tremendous return on investment. The Riverstone product family (RS 1000, RS 3000, RS 8000/8600, RS 16000 and RS 38000) shares the features that provide this heightened return.

Acronyms

ACL	*Access Control List*
ANSI	*American National Standards Institute*
ASIC	*Application-Specific Integrated Circuit*
ASP	*Application Service Provider*
ATM	*Asynchronous Transfer Mode*
CBR	*Constant Bit Rate*
CWDM	*Coarse Wave Division Multiplexing*
DS1/DS3	*Digital Signal, Level 1 (1.54 Mbps) or 3 (44.7 Mbps)*
DSL	*Digital Subscriber Line*
DWDM	*Dense Wave Division Multiplexing*
DVMRP	*Distance Vector Multicast Protocol*
E1/E2	*European Trunk 1/2 (2 Mbps/34.3 Mbps)*
ERP	*Enterprise Resource Planning*
HSSI	*High Speed Serial Interface*
ISP	*Internet Service Provider*
ITU	*International Telecommunications Union*
LAN	*Local Area Network*
LEC	*Local Exchange Carrier*
MAC	*Media Access Control*
MAN	*Metropolitan Area Network*
MDU	*Multiple Dwelling Unit*
MLPPP	*Multi Layer Point-to-Point Protocol*
MPLS	*Multiple Protocol Label Switching.* *See "MPLS in Metro IP Networks,"* *http://www.riverstonenet.com/technology/mpls.shtml*
MTU	*Multiple Tenant Unit*
OC-3/OC-12	*Optical Carrier 3/12 (155 Mbps/622 Mbps)*
PDH	*Plesiochronous Digital Hierarchy*
PIM	*Protocol Independent Multicast*
POS	*Packet over SONET*
PPP	*Point-to-Point Protocol*
PVC	*Private Virtual Circuit*
QoS	*Quality of Service*
RED	*Random Early Discard*
SONET	*Synchronous Optical NETwork* *See http://www.techguide.com/comm/sec_html/sonet.shtml*
SLA	*Service Level Agreement*
SPE	*Synchronous Payload Envelope*
SRP	*Spatial Reuse Protocol* *See RFC 2892*
T1	*Trunk 1 (1.544 Mbps)*
TCP/IP	*Transport Control Protocol/Internet Protocol*
TDM	*Time Division Multiplexing*
UBR	*Undefined Bit Rate*
VBR	*Variable Bit Rate*
VLAN	*Virtual LAN*
VoD	*Video on Demand*
WAN	*Wide Area Network*
WDM	*Wave Division Multiplexing*
WRED	*Weighted Random Early Discard*

Riverstone Networks, Inc.
5200 Great America Parkway, Santa Clara, CA 95054 USA

877 / 778-9595 or **408 / 878-6500** or **www.riverstonenet.com**

Printed in the USA

v 1.4 12/01