



# SENTINEL RIDGE CONSULTING GROUP INC.

**“YOUR COMPUTER SECURITY EXPERT”**

## **All About Proxy Servers**

Ian Parker, GCUX, GCFW, CCSA, CSSE  
President  
Sentinel Ridge Consulting Group, Inc.

# 1 Introduction

As explained briefly in my article “*A Basic Primer On Firewalls*”, available for download at [www.sentinelridgeconsulting.com/resources](http://www.sentinelridgeconsulting.com/resources), the concept of the proxy server emerged from the earlier proxy firewall, a security device that is seldom used anymore, having lost out in the marketplace to the stateful inspection firewall.

A proxy server retains the proxy firewall's ability to proxy client connections but leaves most access control decisions to a separate device.

This article takes an in depth look at how proxy servers work and the benefits they provide in a typical security infrastructure.

## 2 Functions of a Proxy Server

The basic function of a proxy server is to act as an intermediary on behalf of clients who wish to access resources on a server. A common example is a workstation on the corporate network that requires access to a Web site on the Internet. If this was the only function that proxy servers performed, this would be a very brief article. Fortunately, they also perform a number of other functions either directly or indirectly that, in combination, make proxy servers a very important component of an organization's security infrastructure.

A proxy server within a corporate environment can perform one or more of the following functions:

- caching of frequently accessed resources
- user authentication
- enforcement of a corporate acceptable use policy for Internet access
- Internet usage reporting
- scanning of inbound content for malware
- scanning of outbound content for leakage of classified information

In addition to their many legitimate uses, proxy servers can also be used for more nefarious purposes. For example, so called “anonymizing” proxy servers are prevalent on the Internet. The purpose of such devices is to bypass security controls enforced by a company by hiding the true identity of the client. The use of these servers is not discussed further in this article. However, corporate customers need to be aware that they can pose a significant threat to a company's security.

## 3 Proxy Methods and Terminology

The world of proxy servers contains a number of confusing terms that mean different things to different people. For example, you will frequently hear the term *transparent* used to describe a proxy server that operates transparently to the user and intercepts all client connections without any required browser configuration. However, the correct term for this is actually an

*intercepting* proxy. By contrast, RFC<sup>1</sup> 2616 defines a transparent proxy as one that does not modify the request or response. By the same token, intercepting proxies are sometimes known as *forced* proxies. However, the latter term can sometimes be used to describe a *non*-intercepting proxy, if the person using the term is referring to the fact that the user is forced to configure their browser for proxy operation in order to access the Internet.

Despite the ambiguous terminology, from the user's perspective, the key point is that the security infrastructure will be configured in one of two ways to allow employees to access the Internet:

- configuration of the browser with the host name or IP address of the corporate proxy server or
- routing of all client traffic bound for the Internet through the proxy server without the necessity of changing the browser configuration

If the first option is used, the proxy configuration may be performed manually or may occur automatically through the use of a proxy auto-configuration file. In either case, the browser settings are often locked to prevent users from bypassing the proxy server. However, the Internet firewall would normally block such direct connectivity in any case.

## 4 Caching

Proxy caching operates in a manner similar to caching within a browser. However, whereas a browser cache is specific to a single user, a proxy server can cache content and then serve that content to an entire network of users. This capability can have a tremendous impact both on latency, i.e. the speed at which Web pages are delivered to users, and network bandwidth. For example, an organization with limited bandwidth for its Internet connection can still accommodate a large number of users through judicious use of proxy caching.

## 5 User Authentication

Some companies may have a policy whereby only certain employees are allowed to obtain Internet access. In such instances, the proxy server needs to be able to authenticate the user prior to completing the connection. Even in cases where all employees are allowed Internet access, it is useful to be able to record user names in the event that an incident investigation is required in the future, otherwise only the IP address of the client is available.

Proxy servers generally offer the same kind of authentication options as Web servers. The best choice for a particular company will depend on the supporting infrastructure. For example, some proxy servers can use Windows domain controllers to transparently authenticate the user using the credentials with which they logged onto the network. Companies with an existing Windows Active Directory environment would find this to be a much more attractive option than forcing users to manually authenticate every time they wish to browse the Internet.

---

<sup>1</sup> RFC stands for Request for Comments. RFC documents represent a formal mechanism used to describe communications standards for the Internet.

## **6 Policy Enforcement**

One of the most useful value-added services that proxy servers provide is that of policy enforcement concerning acceptable use of Internet services. Companies are legitimately concerned about the potential for time wasting if employees are provided with unlimited access to the Internet. There are also legal ramifications if employees download such things as illegal software or pornography.

Content filtering capabilities may be built into the proxy server software or the capability may be provided by a third party. In the latter case, some proxy servers make use of a special protocol called the Internet Content Adaptation Protocol (ICAP) to integrate with the device that performs the content filtering.

Content filtering may be as simple as a static list of allowed or blocked URL's. More commonly, the customer will subscribe to a service that places Internet Web sites into various categories and then configure the proxy server or third party device to allow or block access to each category. More advanced proxy servers allow the user to implement very fine grained filtering rules.

## **7 Internet Usage Reporting**

Proper logging configuration is important for all security devices. However, it is especially important for a proxy server, as these logs often form the centrepiece of an incident investigation, especially one related to violation of the acceptable use policy.

A large number of Internet users can produce voluminous proxy logs, so companies need to give careful consideration to two issues, namely, how the logs will be read and how they will be archived.

Some proxy servers produce very detailed logs but provide little in the way of user friendly tools to read them. Presumably this is based on the assumption that the customer will be sending the logs to another centralized logging device, such as a Syslog server or security information and event management product. If these facilities do not currently exist, customers should be aware that the cost of a project involving the installation of a new proxy server may exceed their initial expectations.

In a similar vein, local storage capacity on the proxy server may be limited and additional external storage may need to be allocated for archival of proxy logs.

## **8 Malware Scanning**

Scanning of malware in HTTP or FTP traffic is very similar to anti-virus scanning on a workstation or server. As is the case for URL filtering discussed earlier, many proxy servers offload this functionality to a third party device and use the ICAP protocol for communication back and forth.

## **9 Data Loss Prevention**

Until recently, companies have been more concerned about content that may be entering the

corporate network than that which may be leaving it. That is now changing, as there are growing concerns over the ease with which classified information can leave a company, inadvertently or otherwise. Data Loss Prevention (DLP) products attempt to address this issue by inspecting outbound traffic for sensitive content and applying the appropriate DLP policy. Again, these devices often work in conjunction with a proxy server and make use of the ICAP protocol.

## **10 Proxy Server Deployment**

Proxy servers are usually located within the corporate network, on a subnet that can be reached by all internal clients. The Internet firewall must be configured appropriately to allow the proxy server to access the Internet. As the proxy server will normally not have an Internet-routable IP address, network address translation must also be configured on the firewall.

## **11 Reverse Proxy Servers**

When people speak of proxy servers, they are usually referring to an outbound proxy server, one that proxies connections between an internal client and a server on the Internet. A reverse proxy server, on the other hand, proxies connections from clients on the Internet to an internal server that hosts publicly-accessible resources. Reverse proxy servers are also known as surrogates. Typically, the target server is located in a demilitarized zone (DMZ), which is a segregated network segment behind the firewall. The use of DMZ's is beyond the scope of this article. However, the basic purpose of a DMZ is to allow external access to certain internal resources without posing a risk to the remainder of the corporate network.

The main functions that reverse proxy servers perform are the following:

- prevention of direct connectivity to internal servers
- load balancing
- caching of frequently used resources
- compression of content
- spoon feeding to reduce resource usage
- extranet publishing for servers that are truly internal as opposed to residing on a DMZ
- termination and acceleration of Secure Sockets Layer (SSL) connections

The last point bears some explanation. If a company hosts a number of publicly-accessible servers, all of which require SSL connectivity, a saving in certificate costs can be achieved by using a single SSL certificate on the proxy server, as opposed to a certificate on each Web server. Some reverse proxy servers also utilize special hardware to speed up encryption and decryption operations, which can otherwise consume a lot of CPU resources.

## **12 Proxy Server Maintenance**

Once a proxy server has been deployed, its network configuration is not likely to change. The same, however, cannot be said for any content filtering policy that may be applied. Depending

on the complexity of the policy, it may require tweaking on a daily basis. There are several reasons for this. As discussed earlier, many customers subscribe to a service that attempts to place Internet Web sites into different categories e.g. Business/Economy, Government/Legal, Health etc. Customers then configure the proxy server to allow or block access to entire categories. While such a service is extremely valuable, the service providers do make mistakes and occasionally place a site into the wrong category. Usually, the problem can quickly be corrected by notifying the vendor. However, in the meantime, it will be necessary to create a policy exception (if the proxy server contains this feature) to allow access to the site in question. Also, on an ongoing basis, the user community will come up with creative reasons why a certain site is really business-related, even though it is in a blocked category, and will request that they be allowed access. A procedure should be developed such that management sign-off is required before such requests are granted.

The proxy policy should be tested periodically (ideally every time a change is made) to confirm that it is doing what it is supposed to do and that sites which should be blocked are in fact blocked. Users are unlikely to complain of having too much access, so you probably won't be alerted to the fact that all of a sudden all sites that were previously blocked are now accessible.

*Ian Parker is the president of Sentinel Ridge Consulting Group, a firm specializing in the security of corporate computer networks. You can reach him at [sentinelridgeconsulting.com](http://sentinelridgeconsulting.com).*