



Web Request Routing and Redirection

What's the best option for your web security deployment?

Choosing the right method for redirecting traffic to your secure web gateway is absolutely essential to maximize protection for users who are accessing the Internet either on site within the corporate firewall or on the go. How you do it depends on the available configuration options for your web security solution, your unique IT infrastructure, and your chosen method of deployment. Never before have organizations of all sizes and infrastructure configurations had so many deployment choices for multilayered web security. Flexible implementation options include on-premises appliances and blade servers, Software-as-a-Service (SaaS), or hybrid combinations. Secure web gateways in all of these form factors aim to protect users, devices, and networks against increasingly malicious and complex web threats through inbound and outbound web filtering, advanced antivirus and anti-malware, and policy enforcement.

Location. Location. Location.

Where are your users and devices when they access the Internet? This is a key consideration in determining the best methods for redirecting web traffic to a secure web gateway solution. Whichever methods you choose, you must take into account user experience in terms of latency (performance degradation), location of the user, the device in use, ease of deployment and management, and above all, level of protection.

Behind the corporate firewall

When a user is on a corporate network behind the firewall, latency is generally not an issue, and redirection or routing of web traffic to a secure web gateway is typically transparent. On the other hand, the minute the user walks out of the office, the picture changes. Redirection could impact the protection level of the user and connection performance, depending on how and where the user goes online.

Virtual private network (VPN) and captive portal connections

When a user connects to a VPN, there is often a significant slow-down in performance, but there can be a higher degree of security when the user's exploration of the Internet is protected by the secure web gateway solution within the corporate network. But what happens when the user is on a business trip and attempts to access the Internet at a hotel? Typically, guests are asked to open a browser and then pass through a captive portal that requires them to log in on a gated web page

Web Request Routing and Redirection

Key routing methods

- Browser plug-ins.
- Explicit proxy.
- Client proxy.
- Proxy auto-configuration (PAC).
- Generic routing encapsulation (GRE).

Key supporting technologies

- Web proxy auto-discovery (WPAD).
- VPNs, firewalls, routers, and perimeter devices.
- Mobile device management.
- IP anycast proxy URLs.

Technical Brief

before they are able to access the internet. This is commonplace at Wi-Fi hotspots, such as cafés, and at wired locations as well, such as business centers and apartment buildings. In some cases, captive portals are poorly secured. Packet sniffers can easily spoof the IP or media access control (MAC) address of the authenticated target—the user's computer—and follow the path to the network gateway. However, with the use of a client-side proxy, web access through a captive portal can still be directed to a secure web gateway to stop these threats.

Off network, on a mobile device

And last, but not least, there are mobile devices—laptops, smartphones, and tablets. What redirection methods can you use to prevent Internet threats originating on these devices from entering the corporate network? What if you want to enforce certain policies and monitor Internet usage, just as you would with any other corporate workstation? Regardless of location or device, the choice of how you route users to your secure web gateway can have a direct impact on the security posture of your business and employee productivity.

Redirection Methods

Let's try to address this challenge by exploring some of the most widely used mechanisms for routing web traffic to a secure web gateway solution by taking a look at their capabilities, advantages, and limitations.

Browser plug-ins

Browser plug-ins are applications or programs that extend the capabilities of your browser. Browser plug-ins supplied by some vendors are a simple method of redirecting web traffic to gateway solutions.

However, there are a number of user experience and security issues that arise from their implementation:

- Some of these plugins are visible to users, which means that they may be easy to bypass or uninstall by those who are technology savvy. In most cases, the bypass is as simple as running an unsupported browser from the device itself or even from a USB drive.
- Certain browsers, like Firefox, are ideal for handling add-ons, but with others, like Internet Explorer, plug-in implementation may be more challenging and time-consuming for the administrator.
- Browser plug-ins have full access to users' browsers, clearly seeing content that is sent over HTTPS. Attackers, of course, are well aware of this and could easily deceive users into downloading malicious programs in the guise of legitimate web security plug-ins that can gain access to user surfing data, logins, and passwords sent over unencrypted channels.
- Skype and other applications that connect to the Internet directly are not covered by many plug-ins, completely bypassing the security solution.

While browser plug-ins are a simple and convenient method of redirection, plug-in software running in the browser could inherently slow down the performance of the browser by adding extra processing load, degrading aspects of the browsing experience.

Explicit proxy

Using an explicit proxy, devices are configured so that they direct web traffic to a secure web gateway, which is usually deployed behind a firewall. Setting up an explicit proxy involves configuring each client browser to point to a specific proxy, like a secure web gateway solution. Explicit proxy deployments are typically done manually.

A major disadvantage of using explicit proxies is that users can disable the settings and bypass the proxy. The process of setting up an explicit proxy is time-consuming for IT, so it may not be the best option for far-reaching, complex corporate networks, but it might do just fine for a small business.

Client proxy

A new development in the security industry is the client proxy, which solves some of the major issues mentioned thus far. Client proxies use agents to redirect web traffic from laptops, regardless of where they are located. Client proxies can be location-aware; for example, they recognize when a laptop is inside the corporate network, connected via VPN, or outside the network.

The client proxy should work transparently—so there's no user interaction. When the user and device are inside the corporate network, the client proxy will stand down and not re-route the web traffic to something other than what has been configured on the network and will allow traffic to route through an on-premises secure web gateway. When the user is outside the network, the client-proxy can utilize encrypted metadata to identify the user without requiring the user to authenticate. Then it will intelligently route to the cloud-based secure web gateway, if such a solution has been deployed. This solution may also allow you to define certain policies, such as bypass lists of safe domains, addresses, and ports that users can connect to directly without redirection, and block lists, such as unauthorized browsers, URLs, or applications. This type of solution also works well in a captive portal environment once the user logs in and establishes a connection.

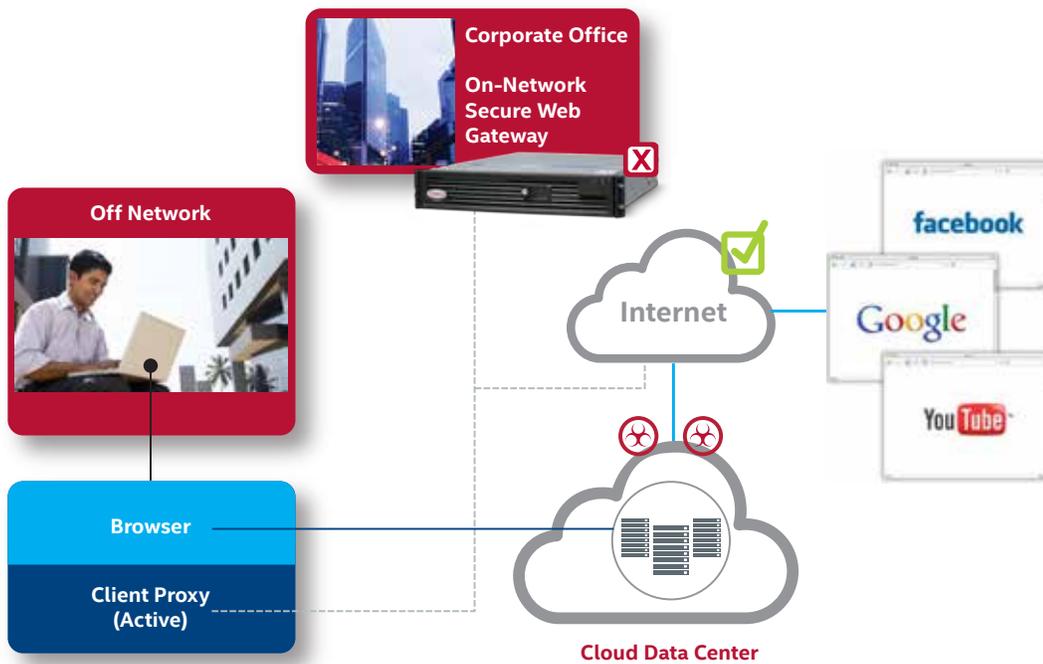


Figure 1. Typical off-network user routing with a client proxy.

Administrators typically deploy this type of solution with a centralized security management system. And since the user cannot disable it easily, it's relatively tamperproof. Another advantage of the client proxy is that it works with multiple layers of network traffic—including applications, such as Skype, Hulu, or IM clients—not just browser traffic.

Proxy auto-configuration

A method that simplifies administration and reduces the burden on IT, proxy auto-configuration (PAC) files can be hosted on the network or live on the client to redirect traffic to the secure web gateway. A built-in JavaScript function in the PAC file defines settings for finding the appropriate proxy for particular web addresses—a valuable feature in a complex corporate network with multiple proxies or gateway appliances.

The great benefit of PAC files is that they can be easily pushed out over the network, and browser settings can be controlled centrally by IT administrators. PAC files also feature automatic detection, which when turned on, automatically configures the browser to direct web traffic to the assigned proxy without the need to customize each and every browser for each and every user. For example, if a user downloads an updated version of a browser from the Internet, PAC files automatically configure the user's browser. PAC files also allow IT administrators to configure browsers to bypass redirection for URLs that do not need to be accessed through a proxy, such as a corporate intranet. PAC implementations are used for on-premises, SaaS, or hybrid secure web gateway configurations.

Generic routing encapsulation

Generic routing encapsulation (GRE) is a protocol that acts as an 'envelope,' encapsulating packets of network traffic and carrying them over an IP network to their destinations. The packets within the envelope are not parsed by IP routers—only the outer packet or 'envelope' is parsed as it moves along to its endpoint. When the endpoint is reached, the GRE envelope is removed and the payload is sent to its final destination.

With GRE, the connection must be built at the main router and the endpoint. This endpoint knows about the path to the router and its IP address. The only level of security to protect the actual data packets is the password/validation word used to authenticate the tunnel. Packets travelling inside a GRE tunnel are not encrypted, as GRE does not encrypt the tunnel but only encapsulates it with a GRE header (the 'envelope').

If you require data protection at your organization, Internet Protocol Security (IPSec) must be configured to provide data confidentiality—and then a GRE tunnel can be transformed into a secure VPN GRE tunnel. While one might think that a GRE IPSec tunnel between two routers is similar to a site-to-site IPSec VPN (crypto), it is not. The key difference is that GRE tunnels allow multicast packets to traverse the tunnel, whereas IPSec VPN does not support multicast packets. In large networks where routing protocols such as OSPF and EIGRP are necessary, GRE tunnels are a legitimate option. In addition, GRE tunnels are relatively easy to configure, so most administrators prefer to use GRE rather than IPSec VPN.

Supporting Technologies

Routing web traffic to a secure web gateway can be a complex task. This next section explores additional mechanisms, which, while not required, can help optimize your configuration.

Web Proxy Autodiscovery (WPAD) protocol

The WPAD protocol used widely in enterprise networks, is designed to help browsers or streaming media applications on Microsoft Windows (or any system that supports proxy servers) discover which proxy server to use for HTTP(S) web traffic. WPAD enhances the user experience, significantly speeding up the process of establishing a connection with URLs because it caches popular web pages and stores them on a proxy so that clients don't have to travel through the network to access a particular page.

WPAD uses a variety of different methods to connect to the appropriate proxy, regardless of whether the user is on or off network. It starts with Dynamic Host Configuration Protocol (DHCP). If that doesn't work, it uses Service Location Protocol (SLP) to find proxies within an enterprise network. And if that proves unsuccessful, it searches through domain name system (DNS) records. When the proxy is discovered, WPAD automatically connects to that location for web requests.

While WPAD reduces latency, especially in a corporate setting, it has its vulnerabilities. Improper configuration of WPAD can open the door to man-in-the-middle attacks, which can easily exploit WPAD by using legitimate penetration testing tools to redirect web traffic to malicious proxies.

Technical Brief

Let's take a closer look at a WPAD man-in-the-middle attack from a network security monitoring perspective.



Figure 2. Man-in-the-middle attack exploiting WPAD.

The attacker and the victim above are connected to the same LAN, which is a typical scenario when users connect to the Internet at airports, conference centers, or hotels. All traffic from the local network is also captured by a sniffer via a monitor/SPAN port. The attacker machine is running Backtrack Linux, which contains Metasploit, a legitimate penetration testing tool often used by hackers.

These are the steps cybercriminals follow to mount the attack:

1. Metasploit is updated to the latest version, which contains the WPAD module.
2. Metasploit's command line tool "msfconsole" is initiated.
3. NetBIOS Name Service (NBNS) responses for WPAD are spoofed.
4. The WPAD module is set up to fool clients into using the attacker machine as web proxy.

Clients on the local network with WPAD configured will now try to use the attacker's machine as proxy for HTTP and HTTPS traffic. The attacker will then proxy all outgoing web traffic via TCP port.

Administrators need to exercise strict control over WPAD domains. If there is no WPAD configuration for an organization, users who request web pages and content will be redirected to external locations that have the next WPAD site in the domain hierarchy and use that for its configuration. This could have potentially extreme negative consequences. For example, hackers could register for a WPAD subdomain for a particular country and perform man-in-the-middle attacks on much of that country's Internet traffic by setting themselves as a proxy for all traffic. The other downside of WPAD, is that it executes JavaScript files—even potentially malicious ones—on all users' browsers, even if JavaScript has been disabled.

VPNs, firewalls, routers, and perimeter devices

Enterprise VPNs are used to redirect Internet traffic to a secure web gateway solution through an IPSec tunnel. A big advantage is that the enterprise VPN can be configured to forward traffic from all ports so that it can be analyzed for malware. This type of scenario is particularly effective in a shared infrastructure environment with multiple IP addresses redirecting to the secure web gateway.

The major drawback with this type of routing is that when failover occurs and processes have to be switched to a redundant or alternative server or network, some latency can be introduced. In addition, in environments like branch offices that may have a variety of routers from different vendors, support may be an issue, as each implementation may be different.

Mobile device management

Smartphones have been embraced widely by business and educational institutions for everyday productivity, computing, and communication. Different technologies are needed to secure Internet usage and ensure proper redirection on corporate-issued or other managed devices such as school-owned smartphones than on laptops and desktops.

Mobile device management solutions, a recent development in the security industry, help organizations apply policies and guarantee that Internet connectivity for smartphones has corporate policy enforced on the device, including redirection to a secure web gateway solution, regardless of the location of the user. Many of these work for on-premises, SaaS, and hybrid deployments. PAC, a browser-based proxy, or a VPN profile needs to be downloaded to smartphones that connect to a corporate network.

Mobile device management available from some security vendors can easily deploy and update the software transparently.

IP Anycast Proxy URL

More of an optimization than a routing mechanism in itself, an IP Anycast Proxy URL can be viewed as a supplemental technology to minimize latency for users who are already set up with one of the routing mechanisms already discussed. This technology, which is used primarily in SaaS deployments, sends Internet requests to the nearest node or nearest data center in the network topology.

Anycast has no concept of latency, so that does not factor into its determination of where to route traffic—it simply uses the fewest router hops to determine the closest data center. This is advantageous in a scenario where servers are distributed over multiple geographic locations. Additionally, Anycast has a built-in failover. It continually monitors the health of access points, so if connecting to the target fails, traffic is redirected to the next healthy target within the cluster. Proxy failover uses routing to redirect the traffic from a failed cluster to an active cluster.

Solutions with IP Anycast enabled are optimal for global deployments of a secure web gateway solution.

What's the Right Routing Choice for You?

Your choice of optimum redirection or routing methods ultimately comes down to the unique requirements of your organization based on your existing infrastructure, location of workers, types of devices to be secured, and available vendor offerings.

When evaluating the most appropriate option, you'll need to take into consideration the following:

Your overall infrastructure

- The size of your network, the number of clients it serves, and where clients and offices are located.
- Plans for deployment of secure web gateway solutions at branch offices or global locations.

Devices

- The devices that need to be protected by your web gateway security solution—desktops, laptops, smartphones.
- Need for a mobile device management solution.
- User location: on site, mobile, or a mix of both.

Technical Brief

Policies and management

- The size of your IT department.
- Requirement for centralized management and administration of redirection methods.
- Corporate web access policies.
- Extent of user interaction you are comfortable with.

Performance/user experience

- Importance of connectivity performance and your organization's tolerance for latency (for example, whether your organization is particularly concerned about decreasing latency while keeping protection levels as high as possible).
- Routing methods your vendor of choice supports and recommends as being the most efficient and offering the best possible protection.
- Preferred form factors of your secure web gateway deployment—on-premises appliances, virtual machines, SaaS, or a hybrid solution.

The information provided here will help you make more informed decisions when choosing methods of routing web traffic to your secure web gateway, and which solution is right for you. Each deployment will have unique requirements, and not every solution will fit your business needs or your infrastructure configuration. Industry-leading solutions will provide you with the access, security, and flexibility you need to run your business, unhindered by undesirable web activity or malicious threats.

More Information

For more information on Intel® Security products including McAfee® web security solutions—both on-premises and Software-as-a-Service (SaaS), please visit: [McAfee.com/websecurity](https://www.mcafee.com/websecurity).

