

Deploying with Websense Content Gateway

Websense® Content Gateway is a high-performance Web proxy that improves network efficiency and performance by caching frequently accessed information at the edge of the network. Websense Content Gateway is deployed as an add-on module with Websense Web Security or Websense Web Filter. Websense Content Gateway can also be combined with Websense Web Security and Websense Active Security Module to create the Websense Web Security Gateway.

The Websense Content Gateway module offers:

- ◆ Proxy and caching capabilities
- ◆ Automatic categorization of dynamic Web 2.0 sites
- ◆ Automatic categorization of new, unclassified sites
- ◆ HTTPS decryption/encryption

User requests for Web content pass through Websense Content Gateway on the way to the destination Web server (origin server). If the Websense Content Gateway cache contains the requested content, it serves the content directly. If the cache does not have the requested content or the content is not recent enough, Websense Content Gateway fetches the content from the origin server, while keeping a copy to satisfy future requests.

This document provides system requirements and a brief overview of deployment considerations for Websense Content Gateway with Websense Web filtering software. Deploying Websense Content Gateway as a highly available proxy or in a proxy chain is also covered.

For more information about deploying Web filtering software, see the Websense Web Security and Websense Web Filter *Deployment Guide*, and the appropriate deployment guide supplement for your network size.

For more information on Websense Content Gateway operation, see the Websense Content Gateway *Installation Guide* and Websense Content Manager Online Help.

Deployment issues

A plan to deploy Websense Content Gateway as a proxy in your network should include at least the following considerations:

- ◆ Site requirements for hardware
- ◆ Websense Content Gateway system requirements
- ◆ Advantages and disadvantages of various proxy network configuration options
- ◆ Authentication considerations
- ◆ HTTPS decryption/encryption requirements
- ◆ Mitigation plan for proxy/client issues

Physical requirements

Physical requirements can include such issues as plant size, the power and cooling requirements for the hardware, available rack space, and network connectivity. This issue should be considered in conjunction with Websense Content Gateway system requirements, covered in the next section.

Websense Content Gateway system requirements

Hardware requirements for a basic Websense Content Gateway deployment appear below, including CPU, memory and disk requirements. See the Websense Web Security and Websense Web Filter *Deployment Guide* for a list of the requirements for other Websense system components.

CPU	Quad-core running at 2.8 GHz or faster
Memory	4 GB
Disk space	2 disks: <ul style="list-style-type: none">◆ 100 GB for the operating system, Websense Content Gateway, and temporary data.◆ 147 GB for caching If caching will not be used, the disk is not required.◆ The caching disk:<ul style="list-style-type: none">– Must be a raw disk (not a mounted file system)– Must be dedicated– Must <i>not</i> be part of a software RAID– For best performance, use a 10K RPM SAS disk on a controller that has at least 64 MB of write-through cache.
Network interfaces	2

Transparent proxy deployment includes either a Web Cache Control Protocol (WCCP)-enabled router or Layer 4 switch to redirect traffic, as described below:

Router	<p>WCCPv1 routers support redirection of HTTP only. If your deployment requires additional protocols, such as HTTPS, your router must support WCCPv2.</p> <p>A Cisco router must run IOS 12.2 or later.</p> <p>The clients, the destination Web server, and Websense Content Gateway must reside on different subnets.</p>
—or—	
Layer 4 switch	<p>You may use a Layer 4 switch rather than a router.</p> <p>To support WCCP, a Cisco switch requires the EMI or IP services image of the 12.2SE IOS release (or later).</p> <ul style="list-style-type: none"> • Websense Content Gateway must be Layer 2 adjacent to the switch. • The switch must be able to rewrite the destination MAC address of frames traversing the switch. • The switch must be able to match traffic based on the Layer 4 protocol port (i.e., TCP port 80).

Websense Content Gateway runs on a Red Hat Enterprise Linux Advanced Server Release 4, Update 5, kernel 2.6.9-55

Advantages and disadvantages of proxy deployment options

Network configuration options should be investigated. Websense Content Gateway proxy is used in either an explicit or transparent proxy deployment. With an explicit proxy deployment, client software is configured to send a request directly to Websense Content Gateway. Transparent proxy deployment means that a client request for Web content is intercepted (usually by a router) and sent to the proxy. The client is unaware that it is communicating with a proxy.

Both deployments have advantages and disadvantages that should be considered for Websense Content Gateway deployment. See [Explicit and transparent proxy deployments](#) for more information.

Authentication Considerations

The issue of how to identify and authenticate users needs to be decided. Authentication is the process of verifying a user via a username and password. The use of Websense transparent identification (XID) agents to identify users is recommended over proxy-based authentication. When XID agents are not an option, Websense Content Gateway supports the following proxy authentication methods:

- ◆ NTLM (NT LAN Manager)
- ◆ LDAP (Lightweight Directory Access Protocol)
- ◆ Radius

The issue of authentication is particularly important in a deployment in which multiple proxies are chained. See *In a proxy chain* for more information.

See the Websense Content Manager Online Help and the Websense Web Security/Websense Web Filter *Deployment Guide* for details about authentication methods.

HTTPS decryption/encryption

SSL Manager is an optional feature for transmitting secure data over the Internet. When you use Websense Content Gateway with the SSL Manager enabled, HTTPS data can be decrypted, inspected for policy, and then re-encrypted as it travels from the client to the origin server. The SSL Manager includes a complete set of certificate-handling capabilities. See Websense Content Manager Online Help for information on using the SSL protocol.

Mitigation plan

Preparation for Websense Content Gateway deployment also needs to include some plan for handling Web site requests that should bypass the proxy. The plan should also deal with situations in which key fobs or tokens are used to access the network and for cases of highly coupled client/server Web applications. The type of proxy deployment determines how these situations are handled.

Explicit and transparent proxy deployments

Websense Content Gateway supports installation on a Red Hat Enterprise Linux Advanced Server Release 4, Update 5 (kernel 2.6.9-55). Websense Web filtering software and its reporting components can be installed on either Windows or Linux machines. Running reporting on a Windows system provides more reporting features.

Websense Content Gateway provides the following proxy deployment options:

- ◆ *Explicit proxy deployment*, where the user's client software must be configured to send requests directly to Websense Content Gateway
- ◆ *Transparent proxy deployment*, where user requests are automatically redirected to a Websense Content Gateway cache, typically by a Layer 4 switch or WCCP-enabled router, on the way to their eventual destination

Explicit proxy deployment

Use of Websense Content Gateway in an explicit proxy deployment is an easy way to handle Web requests from users. This type of deployment is recommended for small networks with only a few users. Explicit proxy is also used effectively when proxy settings can be applied by group policy. It requires minimal network configuration, which is an advantage during troubleshooting efforts.

Individual client browsers may be manually configured to send requests directly to the proxy. They may also be configured via proxy configuration instructions downloaded from a Proxy Auto-Configuration (PAC) file or by using Web Proxy Auto-Discovery (WPAD) to download configuration instructions from a WPAD server. See Websense Content Manager Online Help (*Explicit Proxy Caching*) for more information about these options.

A group policy that points to a PAC file for easy configuration changes is a best practice for explicit proxy deployments.

Exception handling instructions can be included in the PAC file or WPAD instructions. For example, requests for trusted sites can be allowed to bypass the proxy.

Websense Content Gateway can scale in an explicit proxy deployment from a single node into multiple nodes that form a cluster, improving system performance and reliability. With management clustering, the nodes in a cluster share configuration information. A configuration change on one node is automatically made in all other nodes. See Websense Content Manager Online Help (*Clusters*) for more details about proxy clusters.

Disadvantages of explicit proxy deployment include a user's ability to alter the client configuration and bypass the proxy. This type of configuration is difficult to maintain for a large user base because of the lack of centralized management.

Transparent proxy deployment

With Websense Content Gateway as a transparent proxy, the use of a Layer 4 switch or WCCP-enabled router to redirect traffic can provide redundancy and load balancing features for the network. Because system management is centralized, users cannot bypass the proxy so easily.

Users request Internet content as usual, without any special browser configuration, and the proxy serves their requests. The Adaptive Redirection Module (ARM) component of Websense Content Gateway processes requests from the network device and redirects user requests to the proxy engine. The user's client software (typically a browser) is unaware that it is communicating with a proxy. See Websense Content Manager Online Help (*Transparent Proxy Caching and ARM*) for more details.

This type of deployment requires the implementation of a network device that is not required in the explicit proxy deployment. The overall system is more complex and therefore requires more network expertise to construct and maintain.

Figure 1 shows some basic components of Websense Content Gateway in a transparent proxy deployment.

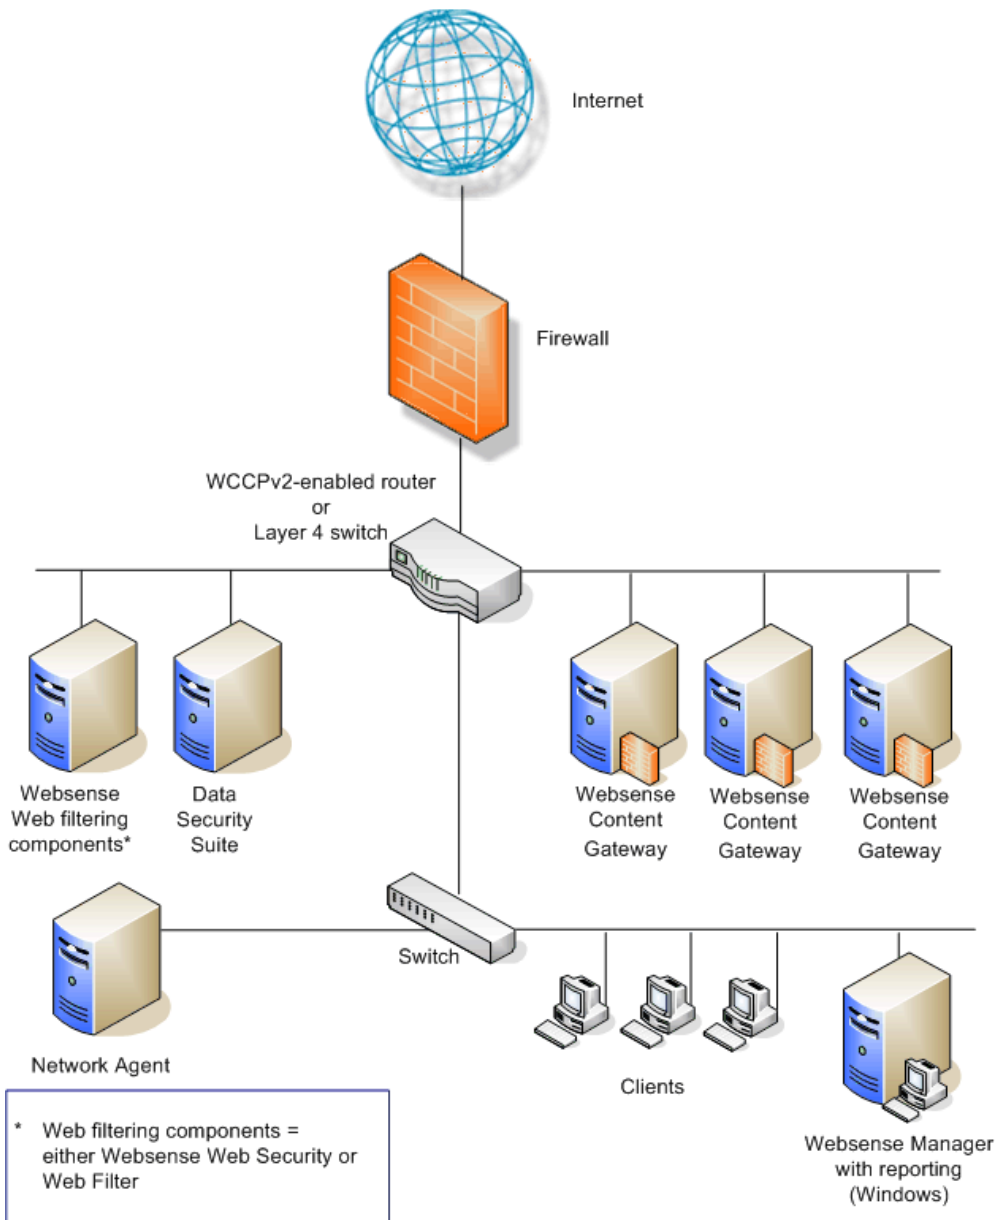


Figure 1 Sample transparent proxy deployment

Websense Content Gateway can also be deployed with a load balancer to distribute the processing of Internet requests. A load balancer not only routes traffic intelligently among all available servers, but can also detect whether a proxy is nonfunctional. In that case, the load balancer re-routes traffic to other, available proxies.

A comparison of how some activities are handled in explicit and transparent proxy deployments appears in the following table:

Activity	Explicit Proxy Deployment	Transparent Proxy Deployment	Proxy Chain
Client HTTP request	Direct connection to proxy by browser to port 8080 (default)	Redirected to proxy by network device using GRE encapsulation or by rewriting the L2 destination MAC address to the proxy's address	Direct connection to parent proxy from child proxy
Exception management	Exclude site, CIDR, etc., using browser configuration settings and PAC file settings.	Static or dynamic bypass rules	Child/parent proxy configuration rules
Proxy authentication	Standard proxy challenge using 407 Proxy Authentication Required code	Nonstandard challenge using server-based authentication scheme (client is not aware of proxy)	Proxies in a chain may share credential information, or a single proxy in the chain can perform authentication (nonstandard solutions).
Redundancy proxy management	Proxy virtual IP pool shared across multiple proxies (management clustering)	WCCP pool with multiple proxies (management clustering)	Parent/child configuration points to proxy virtual IP addresses.
Load balancers	Supported	Supported	Parent/child configuration points to a virtual IP owned by load balancer.

Special Deployment Scenarios

Websense Content Gateway can be deployed in proxy clusters with failover features that contribute to high availability. The proxy can also be deployed in a chain, either with other Websense Content Gateways or third-party proxies. This section describes these deployment scenarios.

Highly available Web proxy

The goal of high availability for Web proxies is continuous, reliable system operation. Minimizing system downtime increases user access and productivity.

High availability may be accomplished via a proxy cluster that uses various failover contingencies. Such deployments may involve either an explicit or transparent proxy configuration, load balancing, virtual IP addresses, and a variety of switching options.

This section summarizes some possibilities for high availability Web proxy deployments.

Using explicit proxy

As previously mentioned for the explicit proxy deployment, clients must be specifically configured to send requests to the proxy cache. The configuration can be accomplished manually, or via a PAC file or a WPAD server.

An explicit proxy deployment for high availability can benefit from the use of *virtual IP failover*. IP addresses may be assigned dynamically in a proxy cluster, so that one proxy can assume traffic-handling capabilities when another proxy fails. Websense Content Gateway maintains a pool of virtual IP addresses that it distributes across the nodes of a cluster. If Websense Content Gateway detects a hard node failure (such as a power supply or CPU failure), it reassigns IP addresses of the failed node to the operational nodes.

Active/Standby

In the simple case of an active/standby configuration with 2 proxies, a single virtual IP address is assigned to the virtual IP address “pool.” The virtual IP address is assigned to one proxy, which handles the network traffic that is explicitly routed to it. A second proxy, the standby, assumes the virtual IP address and handles network traffic only if the first proxy fails.

This deployment assumes the proxy machines are clustered in the same subnet, and management-only mode is used to configure the cluster (that is, both proxies have the same configuration). Figure 2 illustrates this deployment.




Figure 2 Active/standby explicit proxy deployment

Active/Active

In an active/active configuration with 2 proxies, more than one virtual IP address is assigned to the virtual IP address pool. At any point in time, one proxy handles the network traffic that is explicitly directed to it. This deployment is scalable for larger numbers of proxies.

DNS round robin capabilities can support traffic distribution on a rotating basis for the 2 proxy servers. The first client request is served by the first server IP address in the list, and that server IP address moves to the bottom of the list. The second client request is served by the second server IP address, and so on. Management-only clustering mode is assumed.

An increase in the number of proxy machines makes the use of a PAC file for specifying client configuration instructions convenient. A PAC file may be modified to adjust for proxy overloads, in a form of load balancing, and to specify Web site requests that can bypass the proxy.

As with the active/standby configuration, an available proxy can assume a failed proxy's load. Figure 3 illustrates the active/active explicit proxy configuration.




Figure 3 Active/active explicit proxy deployment

With load balancing

A load balancer is a network device that not only distributes specific client traffic to specific servers, but also periodically checks the status of a proxy to ensure it is operating properly and not overloaded. This monitoring activity is different from simple load distribution, which routes traffic but does not account for the actual traffic load on the proxy.

A load balancer can detect a proxy failure and automatically reroute that proxy's traffic to another, available proxy. The load balancer also handles virtual IP address assignments. Figure 4 shows a load balancer added to an explicit proxy configuration.

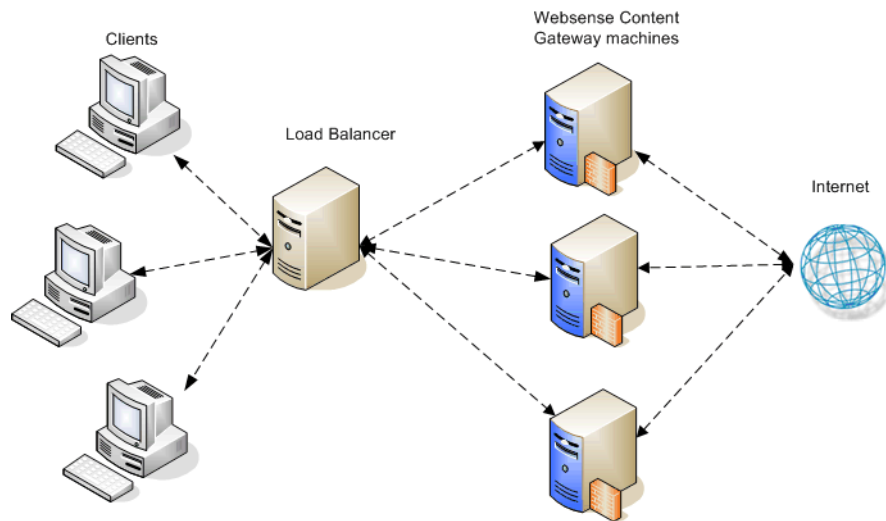


Figure 4 Explicit proxy configuration with load balancing

Using transparent proxy

In a transparent proxy deployment for high availability, traffic forwarding may be accomplished using a Layer 4 switch or a WCCP enabled router. Routers or switches can redirect traffic to the proxy, detect a failed proxy machine and redirect its traffic to other proxies, and perform load balancing.

With a Layer 4 switch

In one simple form of transparent proxy, a hard-coded rule is used to write a proxy's Media Access Control (MAC) address as the destination address in IP packets in order to forward traffic to that proxy. Traffic that does not include the specified proxy address for forwarding is passed directly to its destination. Figure 5 illustrates this deployment.

As described for the explicit proxy, virtual IP addresses can be used in this scenario to enhance availability in case a proxy machine fails.




Figure 5 Layer 4 switch with transparent proxy

Policy-based routing can also be used to allow a Layer 4 switch to change the destination IP address to be that of the Websense Content Gateway proxy. The following tasks need to be performed for this scenario:

1. Create an access list of IP addresses for Web traffic requests that should be redirected.
2. Create a route map that defines how Web traffic request packets are modified for redirection.
3. Apply a “redirect to proxy” policy to the router interface.

Using WCCP

WCCP is a service that is advertised to a properly configured router, allowing that router to automatically direct network traffic to a specific proxy. In this scenario, WCCP distributes client requests based on the proxy server’s IP address, routing traffic to the proxy most likely to contain the requested information.

A proxy and a router communicate via a set of WCCP “Here I am” and “I see you” messages. A proxy that does not send a “Here I am” message for 30 seconds is removed from service by the router, and client requests that would have been directed to that proxy are sent to another proxy.

In a proxy chain

Websense Content Gateway can be deployed in a network that contains multiple proxy machines, including one or more third-party proxies. A *proxy chain* deployment can involve different scenarios, depending on where Websense Content Gateway is located in relation to the client. The proxy that is closer to the client is called the *downstream* proxy. The other proxy is the *upstream* proxy.

In Figure 6, one simple example of proxy chaining shows Websense Content Gateway as the downstream proxy, and the other shows Websense Content Gateway as the upstream proxy.




Figure 6 Simple proxy chain deployments

Websense Content Gateway can participate in flexible cache hierarchies, where Internet requests not fulfilled in one cache can be routed to other regional caches, taking advantage of their contents and proximity. For example, a cache hierarchy can be created as a small set of caches for a company department or a group of company workers in a specific geographic area.

A cache hierarchy consists of levels of caches that communicate with each other. In a hierarchy of proxy servers, Websense Content Gateway can act either as a parent or child cache, either to other Websense Content Gateway systems or to other caching products. A cache hierarchy with multiple parent caches is an example of *parent failover*, in which a parent cache can take over if another parent has stopped communicating.

Web content returned from the origin server by a parent cache is copied by each parent/child cache as it travels back to the client/requestor. See Websense Content Manager Online Help (*Hierarchical Caching*) for more information.

Basic deployment options for proxy chaining using Websense Content Gateway and a third-party proxy are described in the following sections. Specific instructions for using Blue Coat ProxySG or Microsoft ISA server as the downstream proxy are also included.

Websense Content Gateway is downstream

A basic deployment has Websense Content Gateway as the downstream proxy, closer to the client. In this scenario, Websense Content Gateway security features are well positioned for maximum protection and network performance.

This scenario uses the Websense Content Gateway authentication feature to validate client credentials. You must disable authentication on the third-party proxy.

If the upstream third-party proxy requires authentication, you must disable authentication on Websense Content Gateway and enable the pass-through authentication feature via an entry in the **records.config** file, for example:

```
CONFIG proxy.config.http.forward.proxy_auth_to_parent INT 1
```

You can then use an XID agent (for example, the Logon Agent) to facilitate client identification. Websense Content Gateway can additionally send the client IP address to the upstream third-party proxy using the X-Forwarded-For HTTP header via an entry in **records.config**. To enable this, the following entry would be made:

```
CONFIG proxy.config.http.insert_squid_x_forwarded_for INT 1
```

The X-Forwarded-For HTTP header is the *de facto* standard for identifying the originating IP address of a client connecting through an HTTP proxy. Some proxies do not utilize the X-Forwarded-For header.

For information about installing and deploying XID agents, see the Websense Web Security/Websense Web Filter installation and deployment guides.

Websense Content Gateway is upstream

When Websense Content Gateway is the upstream proxy, the downstream third-party proxy can perform authentication and send client IP and username information in the HTTP request headers. Websense Content Gateway authentication must be disabled.

In this scenario, caching must not be enabled on the third-party proxy. Allowing the third-party proxy to cache Web content effectively bypasses Websense Content Gateway's filtering and inspection capabilities.

For an upstream Websense Content Gateway to identify users:

- ◆ Enable authentication on the third-party proxy.
- ◆ Designate Websense Content Gateway as the parent proxy in the third-party proxy's configuration.
- ◆ Set the **Read authentication from child proxy** option in the Websense Content Gateway Configure pane (**Configure > My Proxy > Basic > Authentication**). This option allows Websense Content Gateway to read the X-Forwarded-For and X-Authenticated-User HTTP headers. The downstream third-party proxy passes the client IP address via the X-Forwarded-For header and the user domain and username in the X-Authenticated-User header.

If the third-party proxy can send the X-Forwarded-For header but not the X-Authenticated-User header, the following step is also required:

- ◆ Deploy an XID agent to facilitate client identification by Websense Content Gateway. For information about installing and deploying XID agents, see the Websense Web Security/Websense Web Filter installation and deployment guides.

Websense Content Gateway can be configured to read authentication from the following proxies in the downstream position:

Blue Coat ProxySG	210 series
Microsoft Internet Security and Acceleration (ISA) server	2004, 2006

With Blue Coat ProxySG

You can configure the Blue Coat proxy to send X-Forwarded-For and X-Authenticated-User headers for Websense Content Gateway to read either by manually editing a policy text file or defining the policy in a Blue Coat graphical interface called Visual Policy Manager. Both methods are described here.

Note that for Blue Coat to service HTTPS requests properly with the following setup, you must have a Blue Coat SSL license and hardware card.

Editing the local policy file

In the Blue Coat Management Console Configuration tab, click **Policy** in the left column and select **Policy Files**. Include the following code in the current policy text file, using one of the Install Policy options:

```
<Proxy>
action.Add[header name for authenticated user](yes)

define action dd[header name for authenticated user]
set(request.x_header.X-Authenticated-User,
"WinNT://$(user.domain)/$(user.name)")
end action Add[header name for authenticated user]

action.Add[header name for client IP](yes)

define action dd[header name for client IP]
set(request.x_header.X-Forwarded-For,$(x-client-address))
end action Add[header name for client IP]
```

Using the Blue Coat graphical Visual Policy Manager

Before you configure the Blue Coat header policy, ensure that NTLM authentication is specified in the Blue Coat Visual Policy Manager (**Authentication > Windows SSO**). Set Websense Content Gateway as the forwarding host (in the Blue Coat Management Console Configuration tab, **Forwarding > Forwarding Hosts**).

In the Blue Coat Management Console Configuration tab, click **Policy** and select **Visual Policy Manager**. Click **Launch** and configure the header policy as follows:

1. In the Policy menu, select **Add Web Access Layer** and enter an appropriate policy name in the Add New Layer dialog box.

2. Select the **Web Access Layer** tab that is created.
3. The Source, Destination, Service, and Time column entries should be **Any** (the default).
4. Right-click the area in the Action column, and select **Set**.
5. Click **New** in the Set Action Object dialog box and select **Control Request Header** from the menu.
6. In the Add Control Request Header Object dialog box, enter a name for the client IP Action object in the Name entry field.
7. Enter **X-Forwarded-For** in the Header Name entry field.
8. Select the **Set value** radio button and enter the following value:
`$(x-client-address)`
9. Click **OK**.
10. Click **New** and select **Control Request Header** again.
11. In the Add Control Request Header Object dialog box, enter a name for the authenticated user information Action object in the Name entry field.
12. Enter **X-Authenticated-User** in the Header Name entry field.
13. Select the **Set value** radio button and enter the following value:
`WinNT://$(user.domain)/$(user.name)`
14. Click **OK**.
15. Click **New** and select **Combined Action Object** from the menu.
16. In the Add Combined Action Object dialog box, enter a name for a proxy chain header in the Name entry field.
17. In the left pane, select the previously created control request headers and click **Add**.
18. Select the combined action item in the Set Action Object dialog box and click **OK**.
19. Click **Install Policy** in the Blue Coat Visual Policy Manager.

With Microsoft Internet Security and Acceleration (ISA) server

A Microsoft ISA server can be deployed as a downstream proxy from Websense Content Gateway via a plug-in from Websense, Inc. This plug-in allows Websense Content Gateway to read the X-Forwarded-For and X-Authenticated-User headers sent by the downstream ISA server. The plug-in and ReadMe file are available for download at

<http://installers.websense.com/Downloads/files/v7.1/full/ISAPIPlugin/Websense-AuthForward.ISAPI.zip>

Install the plug-in as follows:

1. Copy the **Websense-AuthForward.dll** file to the Microsoft ISA installation directory. The default directory is **C:\Program Files\Microsoft ISA Server**.
2. Open a Windows command prompt and change directory to the Microsoft ISA installation directory.

3. From the command prompt, type
`regsvr32 Websense-AuthForward.dll`
4. Verify the plug-in was registered in the ISA management user interface (**Start > Programs > Microsoft ISA Server > ISA Server Management**). In the Configuration section, select **Add-ins**, then click the Web-filter tab. The WsAuthForward plug-in should be listed.

To uninstall the plug-in, run the following command in a Windows command prompt from the ISA installation directory.

```
regsvr32 /u Websense-AuthForward.dll
```

