Deploying in a Distributed Environment

Distributed enterprise networks have many remote locations, ranging from dozens to thousands of small offices. Typically, between five and 50 employees work at each remote office. Many of these offices have Internet access, but no dedicated IT staff.

Some organizations use a decentralized network topology that provides each remote office with its own Internet connection. The challenge is to apply consistent, cost effective filtering of Internet requests across the entire organization.

Websense filtering components can be deployed regionally and communicate over the Internet to apply uniform filtering policies to hundreds of remote offices.

Network topology

To reduce network infrastructure costs, each remote office firewall in a decentralized network is connected directly to the Internet, rather than to a corporate WAN.

A small office/home office (SOHO) firewall is connected to an ISDN, DSL/cable, or T1 connection. Except for corporate application data that may use a virtual private network (VPN) connection, each outbound Internet request from a remote office is sent through a local Internet service provider (ISP) to the Internet.

Figure 1 shows the network topology of this type of remote office.

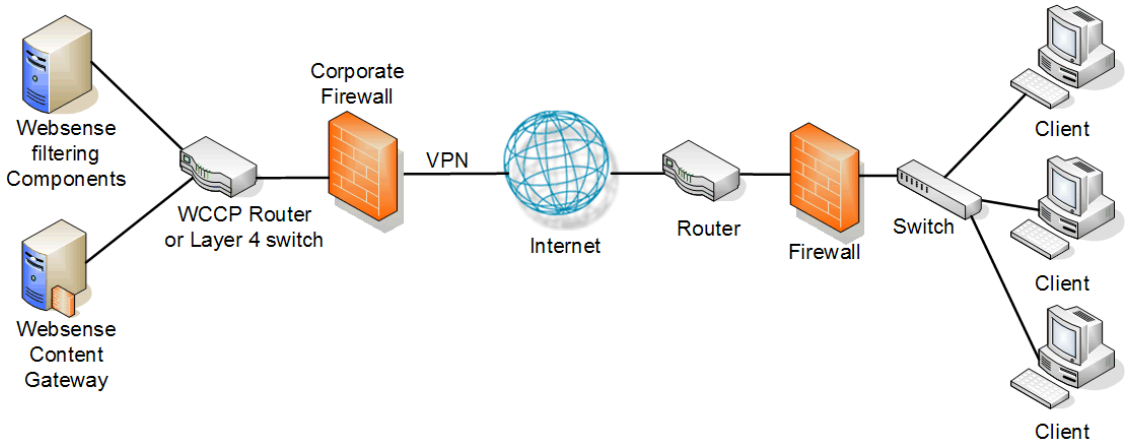

Figure 1 Remote office topology in a decentralized network

A large distributed enterprise, as shown in Figure 2, can have hundreds, or even thousands, of remote offices connected to the corporate network via the Internet.

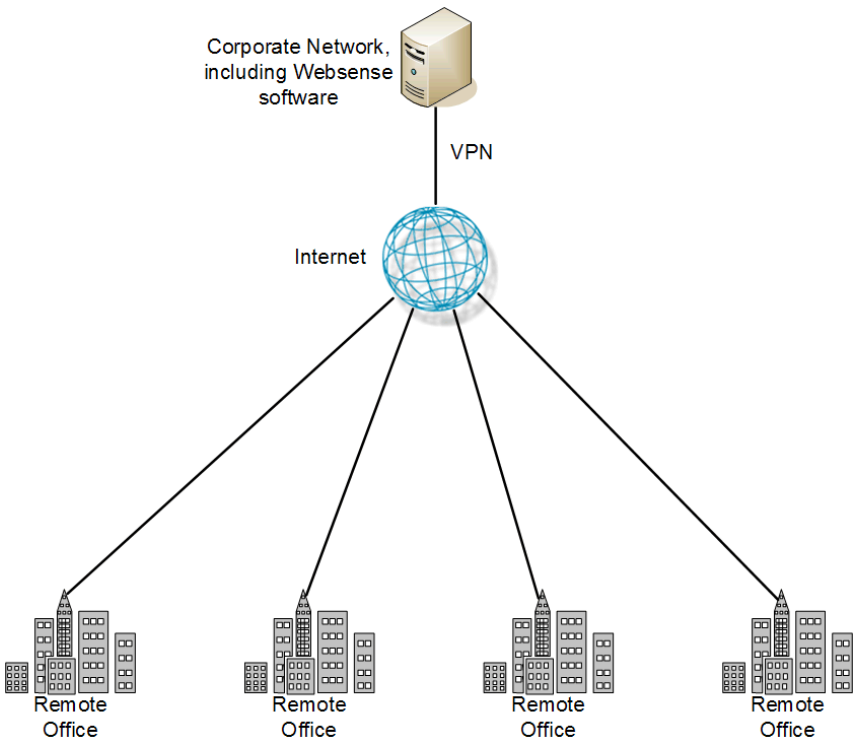

Figure 2 Distributed enterprise

Distributed enterprises with remote Internet connectivity have a complex set of filtering considerations. Remote Filtering and Citrix integration are discussed in *General Deployment Recommendations* in the *Deployment Guide*.

- Remote offices must have Internet access.
- Internet access is provided by independent Internet service providers, often using low to medium-bandwidth connections.
- Web page requests are sent directly to the Internet and are not first routed through a central corporate network.
- Internet access must be filtered to permit only appropriate content.
- Cost considerations prohibit deploying a dedicated filtering server at each office.
- Given the relative low speed of each office's Internet connection, a slightly slower response from the filtering product is acceptable.
- All remote offices can be filtered using to the same policies.

Deploying Websense software in a distributed enterprise

In centralized organizations that route all outbound Internet requests through a single large Internet connection, the server running Websense software is normally placed physically close to the firewall, proxy server, or network appliance.

Remote offices in a distributed enterprise have a direct local connection to the Internet, and no centralized point of control.

Rather than deploying Websense software at each remote office firewall, Websense components can be deployed in a geographically central location. Since Websense software is accessible from the Internet, the Websense components should be protected by a firewall that allows URL lookup requests to pass through.

A SOHO firewall at each remote office is configured to communicate with the centralized Websense components. The firewall does not distinguish between accessing Websense software over the Internet and accessing it through a LAN connection at a remote office.

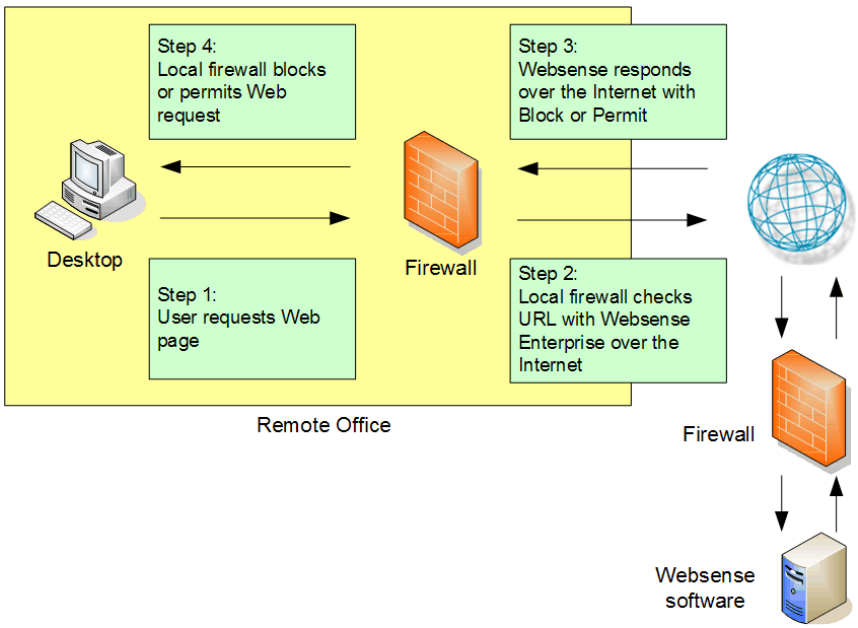

Figure 3 Remote office communication strategy

Figure 4 shows a distributed enterprise with multiple remote offices.

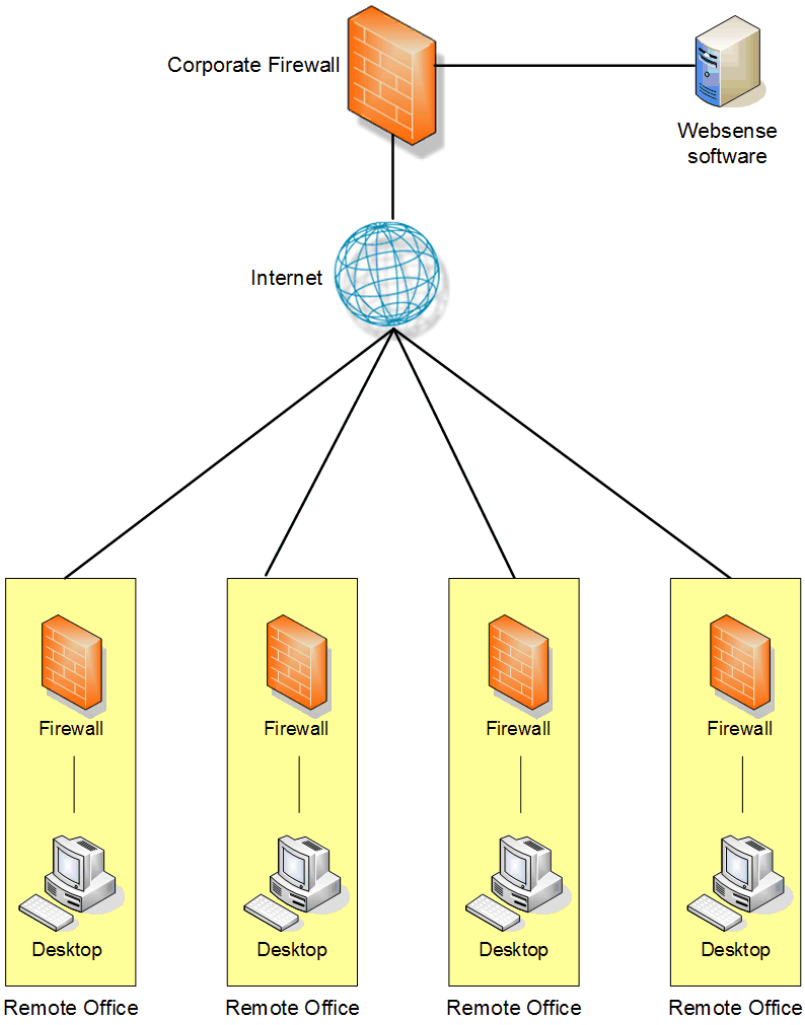

Figure 4 Distributed enterprise communicating with Websense software

Websense has tested this configuration in cooperation with several of its integration partners. The same deployment methodology described here can be used with any network security product integrated with Websense software. A full list of supported integration products can be found at:

www.websense.com/global/en/Partners/TAPartners/SecurityEcosystem/

Centralized filtering:

- Provides distributed enterprises with Websense filtering for each remote office.
- Eliminates the need for a separate Websense installation at each location.
- Provides uniform filtering policies at each remote office.

- Eliminates the cost of additional hardware to provide filtering servers at each remote office.
- Allows the enterprise to centrally configure, administer, and maintain a limited number of Websense filtering machines.

Deployment models

Deployment scenarios vary with different enterprise configurations. For example, an organization with 50 remote offices, all located in the same general region, deploys Websense software differently than a company with remote offices spread throughout the world. This section discusses 3 of the general deployment models available for distributed enterprises:

- Remote offices located within one region
- Remote offices located within one region, with a growing number of employees or offices (or both)
- Remote offices located nationally or globally

Regional offices

The simplest Websense deployment for a distributed enterprise is a network with remote offices in a single region, such as San Diego County. Most organizations with this configuration can use a single Websense deployment, centrally located within that region, to provide filtering for all clients. See Figure 4, page 4.

Expanding regional offices

Some organizations deploy Websense software within a given region and later decide to increase the number of remote offices in that area. To compensate for the additional offices and employees, the organization can:

- **Improve the performance of the machines running Websense components**. Increasing the RAM and CPU, and installing faster hard drives on the Websense machines allows Websense software to respond to an increased number of requests without additional latency. This type of upgrade can help with a moderate increase in head count, or the addition of a few more offices.
- **Deploy additional machines to run Websense components.** If a significant number of new users or offices is added, the deployment of additional instances of certain Websense components, such as Filtering Service and Network Agent, distributes the load and provides optimum performance for each remote office. See Figure 5, page 6.

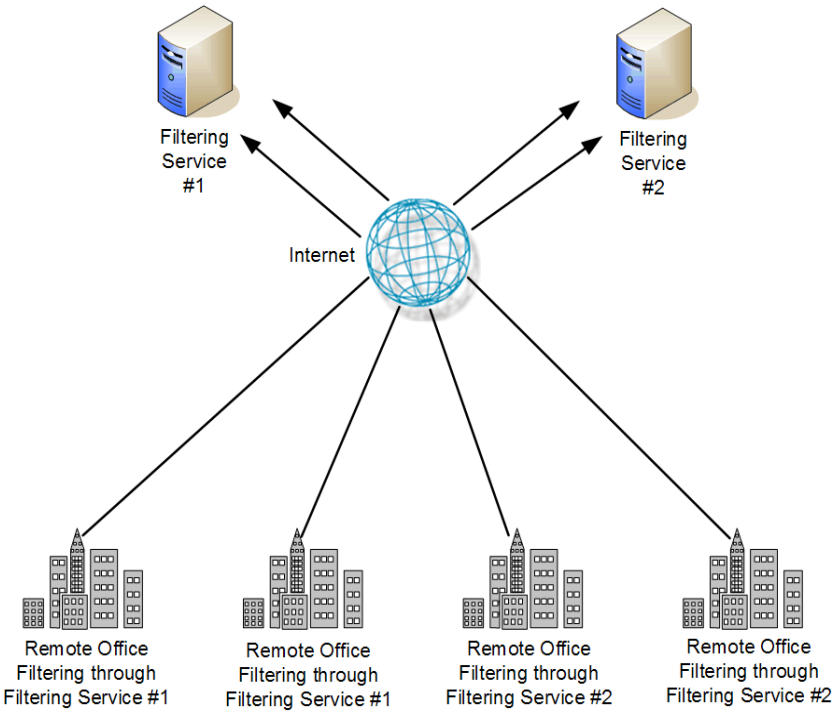

Figure 5 Adding Websense filtering

Additional instances of Websense components can be deployed within the region as the number of offices continues to grow.

National or worldwide offices

Some organizations have hundreds of remote offices spread through a country or even around the world. In such cases, one or two Websense installations are not enough because:

- Each remote office is geographically distant from the Websense components. Request lookups would to have to travel further over the Internet to reach Websense software. This distance increases the total latency of the response and may lead to slower Internet access for end users.
- Large numbers of employees generate more Internet requests than recommended for one or two Websense machines, leading to delays in returning Web pages to requesting clients.

These organizations should divide their offices into logical regions and deploy Websense software in each region. For example, a distributed enterprise might group their United States offices into a western region, a central region, and an eastern region. Websense software is deployed at a central office in each region.

The logical division of offices into regions depends on the location and grouping of remote offices and the total number of employees at each office. For example, a company with a large number of remote offices in a concentrated area, such as New York City, may need to deploy multiple machines running Websense software within that area. Or, an enterprise may only have three offices in California with 100 to 250

employees at each office. In this case, a single Websense installation might be deployed for all three offices. This enterprise also can deploy Websense software locally at each office (rather than using a distributed approach), particularly if an IT staff is present at each location.You may consider installing instances of Filtering Service, Network Agent, and possibly Policy Server and Websense Content Gateway to improve response time for filtering.

Given the significant number of variables, large organizations should contact Websense Sales Engineering to plan a rollout strategy before deployment.

Secure VPN connections

For URL lookup requests and replies, some firewalls allow administrators to set up a secure VPN connection between the remote office firewalls and Websense software. Permitted requests then are fulfilled directly from the Internet, providing an optimum combination of speed and security. See the firewall documentation to determine if the firewall supports this capability.

If a RADIUS server is being used with the VPN service, a Websense RADIUS Agent can be used for transparent user identification. See *General Deployment Recommendations* in the *Deployment Guide* for information about deploying RADIUS Agent. See the Websense Web Security and Websense Web Filter *Installation Guide* for more information about installing the RADIUS Agent.

Calculating TCP connections

In a distributed enterprise, Internet requests may be sent to multiple Filtering Services from hundreds of remote-office firewalls that are configured for persistent TCP connections. In this type of deployment, the number of TCP connections available for a Filtering Service may be exceeded. By default, each Filtering Service is configured to accept a maximum of 500 connections. When the remote offices' firewalls exceed the maximum allowed number of connections that can be accepted by Websense, the firewalls either block all subsequent requests or permit all requests, depending upon how those firewalls are configured.

This section provides the instructions for calculating the number of connections required for a Websense deployment and the number of Filtering Service instances needed under different traffic loads.

Calculating connections

The number of TCP connections opened by different integration products varies widely. In a distributed environment of remote offices, 1-3 connections should be sufficient for each remote office firewall, depending upon the load (number of requests per second) from each office. Contact the manufacturer of the integration product to determine how to limit TCP connections. If the integration product cannot be reconfigured to open fewer connections, additional Filtering Services may be needed to handle the extra connections requested by the remote offices' firewalls.

> **Note**
> Switching connections from TCP to UDP in a distributed enterprise may solve a connection problem. Consult the integration product documentation to determine if the integration can be configured for UDP connections.

To calculate the number of Websense connections required to filter Internet requests from remote offices, multiply:

(*number of integration machines*) x (*number of connections opened by each integration machine*)

To calculate the number of Filtering Service instances an enterprise needs to filter the traffic from remote offices, divide:

(*number of Websense connections required*) / (*number of connections each Filtering Service is configured to accept*)

System requirements for a high performance machine running Filtering Service:

- Quad-Core Intel Xeon processor, 2.5 GHz or greater
- 2 GB RAM (including 1 MB of memory for each connection)

Sizing information

Websense Web filtering performance is dependent upon the machine's processor speed and available memory under a given load (requests per second). An increased load requires more CPU time and supports fewer connections. If fewer connections are supported, additional Filtering Service instances are required to filter the requests from the remote offices.

The following tables display sizing information for remote offices with differing numbers of users.

- Estimates are based on the system requirements for a high performance machine as described in the previous section.
- The number of connections from the integration has been set at one for this example but may need to be higher as the load increases.
- A remote location could have one or multiple firewalls, depending on the network configuration and user location. See the number of firewalls in the following tables.
- As the number of users increases, the required number of Filtering Service instances increases to meet the need to filter a greater number of Internet requests.

Table 1 10 Users per Firewall

Number of users	Connections from integration	Connections allowed by Websense software	Number of Filtering Service instances
1000	1	1000	1
2000	1	1000	2
5000	1	1250	4

Table 2 25 Users per Firewall

Number of users	Connections from integration	Connections allowed by Websense software	Number of Filtering Service instances
1000	1	1000	1
2000	1	1000	2
5000	1	1250	4

Table 3 50 Users per Firewall

Number of users	Connections from integration	Connections allowed by Websense software	Number of Filtering Service instances
1000	1	1000	1-2
2000	1	1000	2-3
5000	1	1250	7-8

Filtering Service faces an increased demand as more users are added behind a firewall. Due to the increased traffic, each installation of Filtering Service is able to handle fewer connections, as seen in the table below.

Table 4 100 Users per Firewall

Number of users	Connections from integration	Connections allowed by Websense software	Number of Filtering Service instances
1000	1	500	2
2000	1	500	4
5000	1	500	10-11

Configuring Websense connections

The number of connections that Websense accepts can be increased.

> **Note**
> Contact Websense Technical Support for assistance with this procedure.

1. Stop the Websense Filtering Service.
2. Navigate to the Websense **bin** directory (C:\Program Files\Websense\bin or /opt/Websense/bin, by default) and locate the **eimserver.ini** file.
3. Make a backup copy of the **eimserver.ini** file in another directory.
4. Open the **eimserver.ini** file with a text editor.
5. Add this line to the file:

   `MaxWISPConnections=<number>`

   Where <*number*> is a value between 501 and 1500.
6. Save and close the file.
7. Restart Filtering Service.

For instructions on stopping and starting Filtering Service, see the *Installation Guide* or Websense Manager Help.

Optimizing network performance

Websense software introduces minimal latency when deployed on a server physically close to a firewall, proxy server, or caching appliance. Websense has also tested the distributed deployment approach discussed in this document to ensure a similarly low level of delay. Latency is the time a network packet needs to reach its destination. Even though outbound Web requests from remote offices must travel over the Internet to the Websense installation, in most situations end users at remote offices are not aware of the filtering process unless they are blocked from a Web site. Total latency depends on these factors:

- Speed (bandwidth) of the Internet connection at each remote office.
- Distance from the remote office to the machine running Websense filtering.
- Number of users and connections to the machine running Websense filtering.
- Speed of the Websense machine.

Internet Connection Speed

Overall filtering performance is dependent upon the speed of the Internet connection at each remote office, which is determined when the parent corporation sets up the office. A DSL, cable, or T1 line is appropriate for an office of 5-25 employees and is fast enough to provide responsive URL lookups through Websense software. A 56K dial-up modem is not recommended because of the additional time needed to retrieve Websense responses.

Important

Match an appropriate class of firewall to the number of employees at each remote site. For example, a remote office with ten employees can use a SOHO-class firewall, while a remote office of 100 employees should use a firewall with greater capacity.

Distance from the Websense machine

The Internet is a large collection of servers and routers that pass data from point to point until it reaches its destination. The more points (hops) a Websense lookup request has to travel, the longer it takes the remote office to receive a reply and fulfill the end user's request. The number of hops required to reach the Websense machine, and the time required for each hop, is generally tied to the geographical distance between machine running Websense software and the source of the request. The closer the server is to a remote office, the faster the Websense lookup and overall performance improves for the end user.

Websense recommends that distributed enterprises deploy the Websense machine no more than 20 hops from each remote office. Similarly, the total trip for an ICMP

(Internet Control Message Protocol) ping from each remote office to Websense machine should take no more than 100 ms to provide satisfactory browsing speeds.

Trip time for a ping and the number of hops can be determined through the use of commands.

From a DOS prompt (Windows):

- `ping`—test network connection and discover the total trip time.
- `tracert`—traces the route to the remote host.
- `pathping`—combines the ping and tracert functions.

From Linux:

- `ping`—test network connection and discover the total trip time.
- `traceroute`—prints the route traffic takes to the remote host. This command requires super-user or administration privileges and has many options.
- `tracepath`—traces the route to the network host. This command has fewer options than traceroute, but can be used by all users.
- `netstat`—prints the network connections, routing tables, and other network data depending on the options that are entered

For more information on these commands and their options, see the Linux man pages.

Hardware performance

The number of requests per second coming in from remote offices can be quite high, as can the number of connections being opened to the Websense machine. The Websense machine must be capable of handling the anticipated traffic load without adding to the latency of the system.

The speed of the SOHO (small office/home office) firewall is also an important consideration. A slower firewall requires additional time to contact Websense software, resulting in slower overall Web page responses. A faster firewall at each remote office processes the Websense response in less time and provides faster overall performance.

Caching

Certain Websense partners have included a filtering enhancement that significantly improves performance with Websense software for distributed enterprises. Juniper Networks NetScreen, Check Point, and Cisco firewalls cache the responses received from Websense software. This cache keeps track of common Web requests and the Websense response (Permit/Block) so the firewall does not have to check with Websense software for every requested Web page. The Websense Content Gateway also offers this feature.

For example, if all employees of an organization are allowed to visit www.cnn.com, these firewalls allow the request to be fulfilled by the destination Web server without

first checking with Websense software (after the first request has been verified). This use of caching can dramatically improve performance.

Websense recommends using a firewall from one of these three vendors when configuring remote offices for filtering. For information on determining if caching is appropriate for the environment, see the *Websense Installation Guide Supplement* for the integration.

Best practices for distributed enterprises

Enterprises with multiple remote offices often find it impractical to deploy Websense software at each location. Other companies do not have the network infrastructure in place to feed all outbound Internet requests through a single, central control point. Using the guidelines and deployment methodologies outlined in the *Deployment Guide* and its supplements, together with careful planning, distributed enterprises can deploy Websense software in an efficient, high-performance, and cost-effective manner.

The main considerations in deploying Websense software in a distributed enterprise are:

- **Response caching**—deploy Websense with a firewall that supports Websense response caching (such as from Juniper Networks NetScreen, Check Point, and Cisco). Other network security products integrated with Websense software may also be used, but end-user performance may be higher with firewalls from mentioned vendors.
- **Distance to the Websense filtering machine**—deploy Websense software no more than 20 hops and 100 ms from remote offices. Organizations with offices spread over a wider area should deploy one or more FIltering Services in each geographical region. Each server should conform to the CPU and RAM requirements described in *Deployment Guide* supplements
- **Configuring connections**—be sure to configure an adequate number of persistent TCP connections for all remote office firewalls. Increase the number of connections that Websense software accepts to accommodate the number of connections opened by the remote firewalls. Provide enough Filtering Services for the anticipated traffic.
- **Internet connection speed**—remote offices should use the fastest Internet connection possible. Filtering is virtually undetectable when using a fast Internet connection. Cable or DSL connections are the minimum requirement for use with Websense software in distributed enterprises.
- **Server speed**—Websense machines must be capable of handling the anticipated traffic load and the number of connections opened by the remote office firewalls. Deploy high performance machines.
- **Filtering policy when Websense is unavailable**—some firewalls and cache appliances give administrators the option of allowing Web requests out to the Internet—unfiltered by Websense—if they receive more Internet requests than they can handle. If this feature is enabled, and a performance problem with the

Internet causes Websense lookups to take longer than normal, users at each remote office can still access the Internet. Filtering is enabled again as soon as Internet performance returns to normal. Websense recommends that administrators enable this option, if available.

- **VPN connection**—use a VPN connection between the remote office firewall and the Websense machine for maximum security (if supported by the firewall).