Deployment Guide

Websense® Web Security Solutions

v7.5

10240 Sorrento Valley Rd., San Diego, CA 92121, USA

Published 2010

Printed in the United States of America and Ireland

Trademarks

Contents

List of Figures

1 Introduction

Use this guide to plan your Websense software deployment before installation. The guide provides an overview of how Websense software can be deployed in a network, as well as operating system and hardware requirements.

This guide applies to version 7.5 of Websense Web Security Gateway Anywhere, Web Security Gateway, Websense Web Security, and Websense Web Filter. The term *Websense software* is used to refer to all or any of these solutions. When information or instructions apply to particular solutions, they are referred to individually by name.

> **Note**
> The technical papers and other documents mentioned in this guide are available from the Documentation > Planning, Installation, and Upgrade folder in the Websense Knowledge Base (www.websense.com/docs).

Websense software consists of components that work together to monitor Internet requests, log activity, apply Internet usage filters, and report on activity. Websense software is highly-distributable, providing the flexibility to scale a deployment to suit your needs. Components can be installed together on one machine for smaller organizations; or they can be distributed across mutliple machines, and multiple sites, to create a high-performing deployment for larger organizations. The appropriate deployment is determined by network size and configuration, Internet request volume, hardware performance, and filtering needs.

This manual provides system recommendations to optimize Websense component performance. Performance can also be improved by using more powerful machines for resource-intensive components.

This chapter introduces Websense filtering, reporting, and interoperability components. See also:

- *Chapter 2: General Deployment Recommendations*—operating system requirements for running Websense components, component limits, tips for maximizing performance, plus recommendations for deploying transparent identification agents, Remote Filtering, and Websense software as a stand-alone installation. Version requirements are also included for various integrations.

- *Chapter 3: Deploying Network Agent*—information for deploying across single and multiple segment networks. Also provides Network Agent placement details, settings, and relationship to hubs, switches, and gateways.
- *Chapter 4: Web Security Gateway Anywhere Deployments*—description of modules in addition to Websense Web Security, including Websense Content Gateway, Websense Data Security Management Server, Websense Sync Service, and Websense Directory Agent.
- *Chapter 5: Integration Deployment*—overview of deploying Websense software with firewalls, proxy servers, caching applications, network appliances, or other integration products or devices.

For Websense Content Gateway deployment information see the *Deploying with Websense Content Gateway* supplement. The gateway provides Web and proxy caching, dynamic classification of Web sites, Web 2.0 categorization, and an optional SSL manager. See the Websense Content Gateway documentation for more information on this product.

Note

Please contact Websense Sales Engineering for assistance in designing your Websense software deployment. A Sales Engineer can help you optimize Websense component deployment and understand the associated hardware needs.

Websense Components

Table 1 provides a brief description of the Websense components. This table groups the components into *core* (included in a standard deployment), *reporting*, *optional*, and *interoperability* (allowing communication and interaction between Web security components, data security components, and the hybrid service in a Websense Web Security Gateway Anywhere deployment).

Review these descriptions to better understand the interaction between components. See Table 2, on page 16, and Table 3, on page 20, for information on the operating system versions needed to run these components.

> **NOTE**
>
> Certain integrations include Websense filtering plug-ins. These are discussed in Table 7, on page 40.

Table 1 Websense Components

Component	Definition
Core Components	
Policy Database	Stores global Websense software settings (configured in TRITON - Web Security) and policy information (including clients, filters, and filter components). • Policy Database is installed in the background with Policy Broker. • Policy Database stores policy-related data; configuration data is stored separately, by Policy Server. In multiple Policy Server environments, a single Policy Database holds policy and general configuration data for multiple Policy Servers.
Policy Broker	Manages requests from Websense components for policy and general configuration information stored in the Policy Database. A deployment can have only one Policy Broker, which is bundled with Policy Database.

Table 1 Websense Components

Component	Definition
Policy Server	• Identifies and tracks the location and status of other Websense components in a deployment. • Logs event messages for Websense components. • Stores configuration information specific to a single Policy Server instance. • Communicates configuration data to Filtering Service for use in filtering Internet requests. Policy and most configuration settings are shared between Policy Servers that share a Policy Database. Policy Server is typically installed on the same machine as Filtering Service. Large or distributed environments can include multiple Policy Servers. Each Policy Server may communicate with up to 10 Filtering Services (see *Filtering Services per Policy Server*, page 25).
Filtering Service	Works with Network Agent or an integration product to provide Internet filtering. When a user requests a site, Filtering Service receives the request and determines which policy applies. • Filtering Service must be running for Internet requests to be filtered and logged. • Each Filtering Service instance downloads its own copy of the Websense Master Database. Filtering Service is typically installed on the same machine as Policy Server. Large or distributed environments may include multiple Filtering Service instances, up to 10 per Policy Server (see *Filtering Services per Policy Server*, page 25).
Network Agent	Works with Filtering Service to enable protocol management, bandwidth-based filtering, and reporting on bytes transferred. • In a stand-alone deployment, enables HTTP and non-HTTP filtering • In an integrated deployment, enables filtering for protocols not managed by your integration product and provides enhanced logging information A deployment may have up to 4 Network Agents per Filtering Service (see *Network Agents per Filtering Service*, page 24).
Master Database	• Includes millions of Web sites, sorted into more than 90 categories and subcategories • Contains more than 100 protocol definitions for use in filtering protocols Download the Websense Master Database to activate Internet filtering, and make sure that the database is kept up-to-date. If the Master Database is more than 2 weeks old, no filtering can occur. A copy of the Master Database is downloaded by each Filtering Service instance.

Table 1 Websense Components

Component	Definition
TRITON - Web Security	Serves as the configuration and management interface to Websense Web Security software. A deployment can have only one instance of TRITON - Web Security per Policy Broker. (Additional, non-reporting instances are possible, but only in a specific configuration, see *Multiple TRITON - Web Security instances*, page 23.) Use TRITON - Web Security to define and customize Internet access policies, add or remove filtering clients, configure Websense software components, generate reports, and more.
Usage Monitor	Enables alerting based on Internet usage. Usage Monitor tracks URL category and protocol access, and generates alert messages according to the alerting behavior you have configured. Alerts can be sent via email or on-screen display, or an SNMP alert can be sent to an SNMP Trap Server. A deployment can have only one Usage Monitor per Policy Server.
User Service	Communicates with an LDAP or NTLM-based directory service to apply filtering policies based on users, groups, domains, and organizational units. A deployment can have only one User Service per Policy Server. The directory service is not a Websense product or component.
Control Service	Tracks the installation, configuration, and removal of Websense components and services on a particular machine. Control Service is automatically installed upon installation of any Websense component.
Reporting Components	
Log Database (*requires a supported database engine*)	Stores Internet request data collected by Log Server for use by Websense reporting tools. The database is created when Log Server is installed. TRITON - Web Security communicates with Log Database to generate Investigative and Presentation Reports, and populate Today and History charts. Reporting components require either Microsoft® SQL Server or MSDE. (MSDE can be installed from the Websense Web site.) Note that SQL Server Express is not supported, and that MSDE is not supported on Windows 2008.

Table 1 Websense Components

Component	Definition
Log Server	Logs Internet request data, including: • The request source • The category or protocol associated with the request • Whether the request was permitted or blocked • Whether keyword blocking, file type blocking, quota allocations, bandwidth levels, or password protection were applied Log Server can log to only one Log Database at a time, and only one Log Server can be installed for each Policy Server. Environments with a high volume of Internet activity should place Log Server on a separate machine. Log Server processing can consume considerable system resources.
Optional Components	
DC Agent[1]	• Offers transparent user identification for users in a Windows-based directory service. • Polls domain controllers in the network to transparently identify users. • Communicates with User Service to provide up-to-date user logon session information to Websense software for use in filtering.
eDirectory Agent [1, 2]	• Works with Novell® eDirectory™ to transparently identify users. • Gathers user logon session information from Novell eDirectory, which authenticates users logging on to the network. • Associates each authenticated user with an IP address, and then works with User Service to supply the information to Filtering Service.
Logon Agent[1]	• Provides unsurpassed accuracy in transparent user identification in Linux and Windows networks. • Does not rely on a directory service or other intermediary when capturing user logon sessions. • Detects user logon sessions as they occur. Logon Agent communicates with Logon Application on client machines to ensure that individual user logon sessions are captured and processed directly by Websense software.
Logon Application	Works with Logon Agent. Runs from a logon script on a domain controller to capture logon sessions as users log on to, or log off of, Windows domains in the network. Logon Application, `LogonApp.exe`, runs as a process on client Windows machines. Upon log on, Logon Application identifies the user and sends the information to Logon Agent.
RADIUS Agent[1]	Enables transparent identification of users who use a dial-up, Virtual Private Network (VPN), Digital Subscriber Line (DSL), or other remote connection to access the network.

Table 1 Websense Components

Component	Definition
Remote Filtering Client	• Resides on client machines outside the network firewall. • Identifies the machines as clients to be filtered. • Communicates with Remote Filtering Server, installed inside the organization's firewall.
Remote Filtering Server	• Allows filtering of clients outside a network firewall. • Acts as a proxy that accepts requests from Remote Filtering Client and submits them for filtering. • Communicates with Filtering Service to provide Internet access management of remote machines.
Interoperability Components	
Directory Agent	In Websense Web Security Gateway Anywhere deployments, collects user and group information from a supported directory service for use in filtering by the hybrid service. See *Chapter 4: Web Security Gateway Anywhere Deployments*.
Filtering Plug-In	Websense software can integrate with a third-party firewall, proxy, cache, or similar product (referred to as an *integration product*). For certain integration products, a Websense filtering plug-in may be required to enable communication between Filtering Service and the integration product.
Linking Service	In Websense Web Security Gateway Anywhere deployments, or in environments that combine Websense data and Web security solutions: • Gives data security software access to Master Database categorization information and user and group information collected by User Service. • Enables shared administrative access to TRITON - Web Security and TRITON - Data Security. See *Chapter 4: Web Security Gateway Anywhere Deployments*.
Sync Service	In Websense Web Security Gateway Anywhere deployments: • Sends policy updates and user and group information to the hybrid service. • Receives reporting data from the hybrid service. See *Chapter 4: Web Security Gateway Anywhere Deployments*.

1. Certain combinations of transparent identification agents are supported within the same network, or on the same machine. For more information, see *Deploying transparent identification agents*, page 27.
2. Running eDirectory Agent and DC Agent in the same deployment is not currently supported.

2 General Deployment Recommendations

Before deploying Websense software, ensure that your hardware and network configuration meet the recommendations provided in this document. This chapter focuses on:

- *Operating system requirements*
- *Component limits*
- *Component ratios*
- *Required external resources*
- *Deploying transparent identification agents*
- *Maximizing system performance*
- *Stand-alone deployments*
- *Remote Filtering Server and Client*
- *Supported integrations*

See *Websense Components*, page 9, for descriptions of Websense components. Note that Websense filtering is based on protocols (like HTTP and FTP), not on the operating system of the computer being filtered.

Note

Websense software supports only TCP/IP-based networks. If your network uses both TCP/IP- and non-IP-based network protocols, only users in the TCP/IP portion of the network are filtered.

Operating system requirements

The tables in this section list the operating systems supported by Websense components. They also list any applications required by Websense components.

Table 2 lists each component and its supported operating systems, along with other software required to run the component. Table 3, on page 20, organizes the requirements by operating system.

Table 7, on page 40, lists the supported integration versions.

Table 2 Components and Required Software

Component	Supported Operating Systems	Other Required Software
*Note: Unless otherwise stated, only **32-bit** versions of operating systems are supported.*		
DC Agent	• Windows Server 2008 SP2 • Windows Server 2003 SP2 • Windows Server 2003 R2 SP2	One of these directory services: • Windows Active Directory® • Windows NT Directory
Directory Agent	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	One of these directory services: • Windows Active Directory® 2003 or 2008 • Novell® eDirectory 8.51 or later
eDirectory Agent	• Windows Server 2008 SP2 Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	• Novell eDirectory 8.51 or later • NMAS authentication is supported. • Recommend Novell Client v4.83 or v4.9 (v4.81 and later are supported)
Filtering Service	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	If Network Agent is used for protocol filtering *and* User Service is installed on a Linux machine, Samba client (v2.2.8a or later) is required on the User Service machine to allow Windows clients to display protocol block messages.
Linking Service (Windows only)	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2	

Table 2 Components and Required Software

Component	Supported Operating Systems	Other Required Software
Log Database (Windows only)	• The Log Database is dependent on the database engine (see *Other Required Software* column), and not the operating system version.	One of these database engines: • Microsoft SQL Server 2008 • Microsoft SQL Server 2005 SP2 or SP3 (Workgroup, Standard, Enterprise, or 64-bit* edition) • MSDE 2000 SP4 * If using 64-bit SQL Server, Log Server must be installed on a machine separate from SQL Server because it does not support 64-bit operating systems.
Log Server (Windows only)	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2	• Log Server creates and maintains the Log Database. One of the database engines listed for Log Database above, must be installed.
Logon Agent	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	Can be used with: • Windows NT Directory (NTLMv1 and v2, only v1 if domain controller runs on Windows Server 2008) • Windows Active Directory (native or mixed mode) • Other LDAP-based directory services
Logon Application (Windows only)	• Any Windows	Logon Application is used with Logon Agent to identify users as they log onto their Windows machines. See *Logon Application*, page 12.
Network Agent	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	Samba client (v2.2.8a or later) is required on the machine running User Service to enable Windows clients to display protocol block messages, if Network Agent is used for protocol filtering and User Service is installed on a Linux machine.

Table 2 Components and Required Software

Component	Supported Operating Systems	Other Required Software
Policy Broker	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	
Policy Server	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	
RADIUS Agent	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	Most standard RADIUS servers are supported. The following servers have been tested: • Microsoft IAS • Merit AAA • Livingston (Lucent) 2.x • Cistron RADIUS server • NMAS authentication
Remote Filtering Client	• Windows 7 • Windows XP SP3 • Windows Vista SP2	
Remote Filtering Server	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	
Sync Service	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	

Table 2 Components and Required Software

Component	Supported Operating Systems	Other Required Software
Usage Monitor	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	
User Service	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3 • Red Hat Enterprise Linux 4, update 7	Supports: • NTLM-based directory services • Active Directory 2003 or 2008 • Sun Java System Directory Server, 4.2 and 5.2 • Novell eDirectory, 8.51 and later Samba client (v2.2.8a or later) is required to enable Windows clients to display protocol block messages, if Network Agent is used for protocol filtering and User Service is installed on a Linux machine.
TRITON - Web Security	• Windows Server 2008 SP2 • Windows Server 2003 R2 SP2 • Windows Server 2003 SP2 • Red Hat Enterprise Linux 5, update 3	• Internet Explorer 7 or 8 Firefox 3.0.x - 3.5.x* • Adobe Flash Player 8 or later *On Linux, Firefox 3.5.x is required for access to all reporting features. For more information, see the TRITON - Web Security Help.

Table 3 lists the operating systems on which the Websense components run. Note: this is the same information contained in Table 2, but organized to show components

grouped by common operating systems. The same component may appear multiple times because it supports multiple operating systems.

Table 3 Operating Systems

Operating System	Component
Microsoft Windows (32-bit versions supported only)	
Windows Server 2008 SP2 Windows Server 2003 R2 SP2 Windows Server 2003 SP2	All Websense components: • Log Database (Microsoft SQL Server or MSDE database engine; note: MSDE is not supported on Windows Server 2008) • DC Agent • Directory Agent • eDirectory Agent • Explorer • Filtering Plug-in • Filtering Service • Linking Service • Log Server • Logon Agent • Logon Application • Network Agent • Policy Server • RADIUS Agent • Remote Filtering Client • Remote Filtering Server • Sync Service • Usage Monitor • User Service • TRITON - Web Security
Windows 7 Windows XP SP3 Windows Vista SP2	• Remote Filtering Client
Any Windows	• Logon Application

Table 3 Operating Systems

Operating System	Component
Linux (32-bit versions supported only)	
Red Hat Enterprise Linux 5, update 3 Red Hat Enterprise Linux 4, update 7	• Directory Agent • eDirectory Agent • Filtering Plug-in • Filtering Service • Logon Agent • Network Agent • Policy Broker • Policy Server • RADIUS Agent • Remote Filtering Server • Sync Service • Usage Monitor • User Service • TRITON - Web Security (supports only Red Hat Enterprise Linux 5, update 3)

Network considerations

To ensure effective filtering, Websense software must be installed so that:

- Filtering Service can receive HTTP requests from Websense Content Gateway; an integrated firewall, proxy server, caching application; or Network Agent.

  In a multi-segmented network, Filtering Service must be installed in a location where it can both receive and manage Internet requests from the integration product and communicate with Network Agent.

- Network Agent:
  - Must be deployed where it can see all internal Internet traffic for the machines that it is assigned to monitor.
  - Can be installed on a dedicated machine to increase overall throughput.
  - Must have bidirectional visibility into Internet traffic to filter non-HTTP requests (such as instant messaging, chat, streaming media, and other Internet applications and protocols).
  - Multiple instances of Network Agent may be required in larger or distributed networks. Each Network Agent monitors a specific IP address range or network segment.

    Using multiple Network Agents ensures that all network traffic is monitored, and prevents server overload. The required number of Network Agents depends on network size and Internet request volume.

  For more information, see *Chapter 3: Deploying Network Agent*.

- As a network grows and the number of Internet requests increases, components can be deployed to additional, non-dedicated machines to improve processing performance on the dedicated machines.
  - You can deploy multiple Filtering Service instances, connected to one Policy Server. This is useful for remote or isolated sub-networks.
  - Because a maximum of 10 Filtering Service instances per Policy Server is recommended (see *Filtering Services per Policy Server*, page 25) multiple Policy Servers may be needed.

> **Note**
> Network Agent can be deployed with the filtering components or on a separate machine. Network Agent should **not** be deployed on the same machine as response-critical components. For more information, see *Chapter 3: Deploying Network Agent*.

> **IMPORTANT**
> **Do not** install Websense components on a firewall machine.

Component limits

In the case of certain components, there can be only one instance of particular dependent components. When deploying Websense software, the following restrictions must be considered.

> **NOTE**
> Even when the number of dependent components is not limited to one, there are best practice component-to-dependent-component ratios. See *Component ratios*, page 24.

Per entire deployment:

- One Policy Broker
- One Sync Service (Websense Web Security Gateway Anywhere deployments)

Per Policy Broker:

- One TRITON - Web Security (only one instance for reporting, other administration-only instances may be deployed; see *Multiple TRITON - Web Security instances* below)

Per Policy Server:

- One Log Server
- One User Service
- One Usage Monitor
- One Directory Agent (Websense Web Security Gateway Anywhere deployments; see *Multiple Directory Agent instances* below for additional information)

Per Filtering Service:

- One primary Remote Filtering Server

Multiple TRITON - Web Security instances

There can be only one instance of TRITON - Web Security that generates and schedules reports. Typically, only one instance is needed in a deployment.

It is possible to install additional instances of TRITON - Web Security in a deployment. However, these must be used as configuration- and administration-only instances (referred to as administration-only instances). They cannot be used to generate reports.

Each administration-only instance of TRITON - Web Security must be associated with a separate Policy Server instance that is not associated with a Log Server. Because the administration-only instances are not associated with a Log Server, they will not display Today and History charts. Also, reporting options will not be available. Only configuration and administration functions will be available.

Multiple Directory Agent instances

Typically, only one Directory Agent instance is required in a deployment. Multiple instances may be deployed if necessary. However, specific configuration of the additional Directory Agent instances is required. See the TRITON - Web Security Help for more information and configuration instructions.

Important

In a V-Series-Appliance-based deployment of Websense Web Security Gateway Anywhere, be aware that Directory Agent is already installed on the appliance. Additional instances of Directory Agent are not typically necessary. If you need to deploy additional instances, see the TRITON - Web Security Help for important configuration instructions.

Component ratios

This section includes best practice component deployment ratios. The optimal deployment may vary based on network configuration and Internet traffic volume.

Larger systems (more than 2500 users) may require a more distributed deployment for load balancing and support of multiple languages.

- Multiple Network Agent instances may be required, for example, to detect outbound traffic on individual network segments.
- It may be appropriate to install multiple Filtering Service instances for load balancing. Some load balancing configurations allow the same user to be filtered by different Filtering Service instances, depending on the current load.

For limits on transparent identification agents, see *Deploying transparent identification agents*, page 27.

For more information about the interaction of Websense components, see the TRITON - Web Security Help.

Network Agents per Filtering Service

As a best practice, no more than 4 Network Agent instances should be deployed per Filtering Service. One Filtering Service instance may be able to handle more than four Network Agents, depending on the number of Internet requests, but if Filtering Service or Network Agent capacities are exceeded, filtering and logging inconsistencies may occur.

Network Agent can typically monitor 50 Mbps of traffic per second, or about 800 requests per second. The number of users that Network Agent can monitor depends on the volume of Internet requests from each user, the configuration of the network, and the location of Network Agent in relation to the computers it is assigned to monitor. Network Agent functions best when it is close to those computers.

Contact your Websense software provider for technical assistance with specific Network Agent sizing guidelines.

Filtering Services per Policy Server

As a best practice, no more than 10 Filtering Service instances should be deployed per Policy Server. A Policy Server instance may be able to handle more, depending on the load. If the number of Filtering Service instances exceeds the Policy Server's capacity, however, responses to Internet requests may be slow.

Multiple Filtering Service instances are useful to manage remote or isolated sub-networks.

The appropriate number of Filtering Service instances for a Policy Server depends on:

- The number of users per Filtering Service
- The configuration of the Policy Server and Filtering Service machines
- The volume of Internet requests
- The quality of the network connection between the components

  If a ping command sent from one machine to another receives a response in fewer than **30 milliseconds (ms)**, the connection is considered high-quality. See *Testing the connection*, page 25 for more information.

If Filtering Service and Policy Server become disconnected, all Internet requests are either blocked or permitted, depending on which option you have chosen in TRITON - Web Security. For more information, see the *Getting Started* topic in the TRITON - Web Security Help.

Filtering Service machines running behind firewalls or remotely (at a great topological distance communicating through a series of routers) may need their own Policy Server instance. In a multiple Policy Server environment, a single Websense Policy Database holds the policy settings for all Policy Server instances. See the TRITON - Web Security Help for more information.

Testing the connection

Run a **ping** test to check the response time and connection between the Policy Server and Filtering Service machines. A response time of fewer than 30 milliseconds is recommended.

1. Open a command prompt (Windows) or terminal session (Linux) on the Policy Server machine.
2. Enter the following command:

   ```
   ping <IP address or hostname>
   ```

   Here, <*IP address or hostname*> identifies the Filtering Service machine.

Interpreting your results

When you run the **ping** command on a Windows machines, the results resemble the following:

```
C:\>ping 11.22.33.254
Pinging 11.22.33.254 with 32 bytes of data:
Reply from 11.22.33.254: bytes=32 time=14ms TTL=63
Reply from 11.22.33.254: bytes=32 time=15ms TTL=63
Reply from 11.22.33.254: bytes=32 time=14ms TTL=63
Reply from 11.22.33.254: bytes=32 time=15ms TTL=63
Ping statistics for 11.22.33.254:
Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 14ms, Maximum = 15ms, Average = 14ms
```

In a Linux environment, the results look like this:

```
[root@localhost root]# ping 11.22.33.254
PING 11.22.33.254 (11.22.33.254) 56(84) bytes of data.
64 bytes from 11.22.33.254: icmp_seq=2 ttl=127 time=0.417 ms
64 bytes from 11.22.33.254: icmp_seq=3 ttl=127 time=0.465 ms
64 bytes from 11.22.33.254: icmp_seq=4 ttl=127 time=0.447 ms
64 bytes from 11.22.33.254: icmp_seq=1 ttl=127 time=0.854 ms
```

Ensure that **Maximum** round trip time or the value of **time=x.xxx ms** is fewer than 30 ms. If the time is greater than 30 ms, move one of the components to a different network location and run the ping test again. If the result is still greater than 30 ms, locate and eliminate the source of the slow response.

Required external resources

Websense software relies on certain external resources and network characteristics to function properly in your network.

- **TCP/IP**: Websense software provides filtering in TCP/IP-based networks only. If your network uses both TCP/IP and non-TCP protocols, only those users in the TCP/IP portion of your network are filtered.
- **DNS server**: A DNS server is used to resolve requested URLs to an IP address. Websense software or your integration product requires efficient DNS performance. DNS servers should be fast enough to support Websense filtering without becoming overloaded.
- **Directory services**: If Websense software is configured to apply user- and group-based policies, User Service queries the directory service for user information. Although these users and group relationships are cached by Websense software, directory service machines must have the resources to respond rapidly if Websense software requests user information. See *Supported directory services*, page 27.

- **Network efficiency**: The ability to connect to resources such as the DNS server and directory services is critical to Websense software. Network latency must be minimized if Filtering Service is to perform efficiently. Excessive delays under high load circumstances can impact the performance of Filtering Service and may cause lapses in filtering.

Supported directory services

If your environment includes a directory service, you can configure Websense software to filter Internet requests based on policies assigned to users, groups, and domains (organizational units).

Websense software can work with the following directory services:

- Windows NT Directory and Windows Active Directory (Mixed Mode)
- Windows Active Directory (Native Mode)
- Sun Java System Directory Server
- Novell eDirectory

For information on configuring Websense software to communicate with a supported directory service, see the TRITON - Web Security Help. Websense software does not need to run on the same operating system as the directory service.

Deploying transparent identification agents

If you are using Websense software as a stand-alone deployment, or if your integration product does not send user information to Websense software, use Websense transparent identification agents to identify users without prompting them for a user name and password.

There are 4 optional transparent identification agents:

- DC Agent
- eDirectory Agent
- Logon Agent
- RADIUS Agent

> **Note**
> DC Agent must have domain administrator privileges to retrieve user information from the domain controller.

If you have deployed Websense software in a single network location, it is a best practice to have a single transparent identification agent instance.

In deployments that cover multiple locations, you can install an agent instance in multiple domains.

For example:

- One **DC Agent** instance can handle multiple trusted domains. Add additional instances based on:
  - The load placed on DC Agent
  - Whether a DC Agent instance can see all the domains on the network, including remote offices

  Load results from the number of user logon requests. With a large number of users (10,000+ users, 30+ domains), having multiple DC Agent instances allows for faster identification of users.

  If multiple Filtering Services are installed, each Filtering Service instance must be able to communicate with all DC Agent instances.
- One **eDirectory Agent** is required for each eDirectory Server.
- One **Logon Agent** is required for each Filtering Service instance.
- One **RADIUS Agent** instance is required for each RADIUS server.

  It is a best practice to install and run RADIUS Agent and the RADIUS server on separate machines. (The agent and server cannot have the same IP address, and must use different ports.)

In some environments, a combination of transparent identification agents may be appropriate within the same network, or on the same machine. See *Combining transparent identification agents*, page 28.

Refer to the Websense Web Security and Websense Web Filter *Installation Guide* for transparent identification agent installation instructions. See the TRITON - Web Security Help for detailed configuration information. More information is also available in the *Transparent Identification of Users* technical white paper.

Combining transparent identification agents

Websense software can work with multiple transparent identification agents. If your environment requires multiple agents, it is best to install them on separate machines.

- eDirectory or RADIUS Agent can be installed on the same machine as Filtering Service, or on a separate server on the same network.
- Do not run eDirectory Agent and DC Agent in the same deployment.

Table 4 lists supported combinations of transparent identification agents.

Table 4 Deploying Multiple Transparent ID Agents

Combination	Same machine?	Same network?	Configuration required
Multiple DC Agents	No	Yes	Ensure that all instances of DC Agent can communicate with Filtering Service, and that the individual DC Agents are not monitoring the same domain controllers.

Table 4 Deploying Multiple Transparent ID Agents

Combination	Same machine?	Same network?	Configuration required
Multiple RADIUS Agents	No	Yes	Configure each agent to communicate with Filtering Service. Multiple instances of the RADIUS Agent cannot be installed on the same machine.
Multiple eDirectory Agents	No	Yes	Configure each instance to communicate with Filtering Service.
Multiple Logon Agents	No	Yes	Configure each instance to communicate with Filtering Service.
DC Agent + RADIUS Agent	Yes	Yes	Configure each agent to use a unique port for communication with Filtering Service.
DC Agent + eDirectory Agent	No	No	Communication with both a Windows directory service and Novel eDirectory is not supported in the same deployment. However, both agents can be installed, with only one agent active.
DC Agent + Logon Agent	Yes	Yes	Configure both agents to communicate with Filtering Service. By default, each agent uses a unique port, so port conflicts are not an issue unless these ports are changed.
RADIUS Agent + Logon Agent	Yes	Yes	Configure all agents to communicate with Filtering Service.
eDirectory Agent + Logon Agent	No	No	Communication with both Novell eDirectory and a Windows- or LDAP-based directory service in the same deployment is not supported. However, both agents can be installed, with only one agent active.
RADIUS Agent + eDirectory Agent	Yes	Yes	Configure all agents to communicate with Filtering Service. When adding agents to TRITON - Web Security, use an IP address to identify one, and a machine name to identify the other. See the *Transparent Identification of Users* white paper for details.
DC Agent + Logon Agent + RADIUS Agent	Yes	Yes	This combination is rarely required. Configure each agent to use a unique port for communication with Filtering Service.

Maximizing system performance

Adjust Websense components to improve filtering and logging response time, system throughput, and CPU performance. This section discusses some ways to optimize performance.

Network Agent

Network Agent can be installed on the same machine as other Websense components, or on a separate machine.

As the number of users grows, or if Network Agent does not block Internet requests as accurately as needed, place Network Agent on a different machine from Filtering Service and Policy Server. You can also add a second Network Agent and divide the network monitoring between the 2 agents.

If Websense software is running in a high-load environment, or with a high capacity Internet connection, you can increase throughput and implement load balancing by installing multiple Network Agent instances. Install each agent on a different machine, and configure each agent to monitor a different portion of the network.

Important

Network Agent must have bidirectional visibility into the network or network segment that it monitors.

If multiple Network Agents are installed, each agent must monitor a different network segment (IP address range).

If a Network Agent machine connects to a switch, the monitor NIC must plug into a port that mirrors, monitors, or spans the traffic of all other ports. *Multiple segment network*, page 47, and *Network Agent location*, page 45, discuss locating Network Agent in more detail.

HTTP reporting

You can use Network Agent or an integration product to track HTTP requests and pass the information to Websense software, which uses the data to filter and log requests.

Network Agent and some integration products also track bandwidth activity (bytes sent and received), and the duration of each permitted Internet request. This data is also passed to Websense software for logging.

When both Network Agent and the integration partner provide logging data, the amount of processor time required by Filtering Service increases.

If you are using both Network Agent and an integration product, you can avoid extra processing by configuring Websense software to use Network Agent to log HTTP

requests (enhanced logging). When this feature is enabled, Websense software does not log HTTP request data sent by the integration product. Only the log data provided by Network Agent is recorded.

Consult the TRITON - Web Security Help for configuration instructions.

Database Engine

The Websense Log Database can be created and maintained by any of the following database engines:

- Microsoft SQL Server 2008
- Microsoft SQL Server 2005
- Microsoft Database Engine (MSDE) 2000

> **Note**
> See the software requirements for Log Database in Table 2, on page 16 for specific SQL Server or MSDE version and service pack requirements.

Log Server logs Internet activity information to only one Log Database at a time.

Microsoft SQL Server

Microsoft SQL Server works best for larger networks, or networks with a high volume of Internet activity, because of its capacity for storing large amounts of data over longer periods of time (several weeks or months).

Under high load, Microsoft SQL Server operations are resource intensive, and can be a performance bottleneck for Websense software reporting. You can tune the database to improve performance, and maximize the hardware on which the database runs:

- If Log Server is installed on the database-engine machine, alleviate resource conflicts between Log Server and Microsoft SQL Server by increasing the CPU speed and/or the number of CPUs.
- Provide adequate disk space to accommodate the growth of the Log Database. Microsoft SQL Client Tools can be used to check database size.
- Use a disk array controller with multiple drives to increase I/O bandwidth.
- Increase the RAM on the Microsoft SQL Server machine to reduce time-consuming disk I/O operations.

Note

Consult Microsoft documentation for detailed information about optimizing Microsoft SQL Server performance.

MSDE

Microsoft Database Engine (MSDE) is a free database engine best-suited to smaller networks, organizations with a low volume of Internet activity, or organizations planning to generate reports on only short periods of time (for example, daily or weekly archived reports, rather than historical reports over longer time periods). MSDE cannot be optimized.

With MSDE, the maximum size of the Log Database is approximately 1.5 GB. When the existing database reaches this limit, it is saved (rolled over), and a new Log Database is created. Use the ODBC Data Source Administrator (accessed via Windows Control Panel) to see information about databases that have been saved.

If the Log Database is rolling over frequently, consider upgrading to Microsoft SQL Server.

Note

Consult the Websense Web Security and Websense Web Filter *Installation Guide* for detailed information about selecting the appropriate database engine for the deployment.

When using MSDE, make sure that the latest service packs have been applied. See the software requirements for Log Database in Table 2, on page 16 for specific MSDE version and service pack requirements. Microsoft SQL Server service packs can be applied to MSDE 2000. The service pack updates only those files relevant to MSDE.

Log Database disk space recommendations

Log Database requirements vary, based on the size of the network and the volume of Internet activity. The following baseline assumptions are made to provide general recommendations:

- An average user requests 100 Web pages per day.
- The Log Database creates a record for each visit.
- Each log record is approximately 500 bytes.
- Each Web page request requires roughly 10 HTTP GETS (*hits*).

If the Log Database is configured to write a record for each hit, the size of the database may increase by a factor of 10. During installation, you are provided options for minimizing the size of the Log Database.

After installation, additional configuration options, including selective category logging, are available to help manage the size of the Log Database. A small Log Database means report generation is faster, but data consolidation may also negatively affect report accuracy. See the TRITON - Web Security Help for details.

Logging visits (default settings)

When you log visits, one log record is created for each Web page requested rather than each separate file included in the Web page request. This creates a smaller database and allows faster reporting.

If the Log Database is configured to record **visits** (the default), you can calculate the disk space required for the database as follows:

(# of Web pages) x (# of bytes) x (# of users)

If an average user generates 50 KB per day (100 pages x 500 bytes), and is logged on for 20 work days per month, that user consumes 1 MB in the Log Database each month (20 days x 50 KB/day). Extrapolating to 500 users, the database would use 500 MB per month to record visits. In this case, the database would use 6 GB for one year’s worth of data.

Logging hits

When you log hits, a separate log record is generated for each HTTP request to display any element of a Web page, including graphics and ads. This type of logging results in a larger and more detailed database than the logging visits option.

If the Log Database is configured to record each **hit**, you can calculate the disk space required for the database as follows:

[(avg. # of Web pages) x (avg. # of hits) x (# of bytes)] x (# of users)

If an average user generates 250 KB per day (100 pages x 5 gets per page x 500 bytes), and is logged on for 20 work days per month, that user consumes 5 MB in the Log Database each month (20 days x 250 KB/day). Extrapolating to 500 users, the database would use 2.5 GB per month.

In this example, the Log Database would requires 30 GB of disk space for one year's worth of data (500 users at 500 hits per day).

Due to the large amount of disk space required, and due to the performance impact on reporting, it is a best practice to not keep live data from large networks for a year. When you break up the database into smaller pieces, you can generate reports much more quickly.

Logging full URLs

Enabling full URL logging creates a larger database than with logging hits, and also provides the most detailed reports. Log records include the domain name and the full path to specific pages requested. Use this option if you want reports of real-time scanning activity.

If the Log Database is configured to log the full URLs, each URL recorded can be up to 1000 characters, or 1 KB, in length. When full URL logging is turned off, a log entry requires only 256 bytes per URL.

If the Log Database is growing too quickly, you can turn off full logging to decrease the size of each entry and slow growth by a factor of 4.

Configure URL logging options in TRITON - Web Security. See the TRITON - Web Security Help for details.

Consolidation

Consolidation helps to reduce the size of the database by combining Internet requests that share the same value for all of the following elements, within a certain interval of time (1 minute, by default):

- Domain name (for example: www.websense.com)
- Category
- Keyword
- Action (for example: Category Blocked)
- User

For example, the user visits *www.cnn.com* and receives multiple pop-ups during the session. The visit is logged as a record.

- If consolidation is turned off (the default), and the user returns to the site later, a second visit is logged.
- If consolidation is turned on, additional visits to the site within a specified period are logged as a single record, with a hits (i.e., visits) count indicating the number of times the site was visited in that period.

Protocol logging

If your deployment includes Network Agent, you have the option to log non-HTTP protocol traffic (for example, instant messaging or streaming media traffic) in addition to HTTP and HTTPS traffic.

The more protocols you choose to log, the greater the impact on the size of the Log Database. See the TRITON - Web Security Help for information about filtering and logging non-HTTP protocols.

Log Database strategy

Using the hits and visits calculations provided under *Logging hits*, page 33, even without logging full URLs, storing data for 1 year could require (assuming 500 users):

- 30 GB for hits
- 6 GB for visits

Generating reports against such large amounts of data can significantly slow report processing.

Use database partitions to limit the scope of the data use to generate reports.

- A database rollover is triggered by a time or size limit.
- New data is collected in a new partition.
- Older data is preserved in other partitions.
- You configure which partition you want to use to generate reports.

Adjust the partition or rollover limits to maximize reporting performance and ease the management of data. See the TRITON - Web Security Help for details.

Stand-alone deployments

When Websense Web Security or Websense Web Filter is deployed as a stand-alone product, Network Agent rather than a third-party integration product (i.e., firewall, proxy, or gateway product or device) monitors network traffic and enables filtering of all protocols, including HTTP, HTTPS, and FTP.

In a stand-alone deployment, Network Agent:

- Detects all TCP/IP Internet requests (HTTP and non-HTTP)
- Communicates with Filtering Service to see if each request should be blocked
- Calculates the number of bytes transferred
- Sends a request to Filtering Service to log Internet activity

For more information, see the Websense Web Security and Websense Web Filter *Installation Guide* or the TRITON - Web Security Help.

Stand-alone installations of Websense software run on the operating systems listed earlier in this chapter (see Table 2, on page 16, and Table 3, on page 20). Components may need to be distributed over multiple machines for load balancing and improved performance in larger networks.

Table 5, on page 36, provides system recommendations for stand-alone deployments of Websense software, based on network size. System needs vary, depending on the volume of Internet traffic.

The following baseline is used to create the recommendations:

- 1 - 500 users = 1 - 100 requests/second
- 500 - 2,500 users = 100 - 500 requests/second
- 2,500 - 10,000 users = 500 - 2,250 requests/second

If your network traffic exceeds these estimates, more powerful systems or greater distribution of components may be required.

Important

- Do not install Websense components on a firewall machine. Firewall and Websense software function and performance may be affected.
- Each Network Agent machine must be positioned to see all Internet requests for the machines that it is assigned to monitor.
- eDirectory Agent or RADIUS Agent can be installed on the same machine as Filtering Service, or on a separate machine in the same network, but not on the same machine as Log Server.

Table 5 System Recommendations for Stand-Alone Deployment

Network Size	Filtering Components	Reporting (Windows)
1 - 500 users	**Windows** or **Linux** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 4 GB RAM • 10 GB free disk space (Free space must equal at least 20% of total disk space.)	**Windows** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 4 GB RAM • 100 GB free disk space • Microsoft SQL Server 2008 or 2005, or MSDE 2000*
500 - 2,500 users	**Windows** or **Linux** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 4 GB RAM • 10 GB free disk space (Free space must equal at least 20% of total disk space.)	**Windows** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 4 GB RAM • 100 GB free disk space • Microsoft SQL Server 2008 or 2005, or MSDE 2000*

Table 5 System Recommendations for Stand-Alone Deployment

Network Size	Filtering Components	Reporting (Windows)
2,500 - 10,000 users	**Windows** or **Linux** • Load balancing required • Quad-Core Intel Xeon 5450 or better processor, 3.0 GHz or greater • 4 GB RAM • 10 GB free disk space (Free space must equal at least 20% of total disk space.) • See the *Important* note below.	**Windows** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 4 GB RAM • 200 GB free disk space with a disk array (The Log Database requires a disk array to increase I/O reliability and performance.) • High-speed disk access • Microsoft SQL Server 2008 or 2005*

* See the software requirements for Log Database in Table 2, on page 16 for SQL Server or MSDE service pack requirements.

Important

Two Network Agent instances, running on separate machines, are required for 2500-10000 user networks. The machines should have:

- Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater
- At least 1 GB of RAM

Multiple Filtering Service machines may also be needed. Machine requirements depend on the number of users being monitored and filtered.

To run both filtering and reporting on the same machine in the two smaller network sizes, increase the RAM to 6 GB (if supported by your operating system), and consider using a faster processor and hard drive to compensate for the increased load.

Remote Filtering Server and Client

For all Websense Web Security solutions, you can monitor computers outside your network using remote filtering components. The **Remote Filtering Client** must be installed on each remote machine.

The remote clients communicate with a **Remote Filtering Server**, which acts as a proxy to Filtering Service. This communication is authenticated and encrypted.

> **NOTE**
>
> If you have Websense Web Security Gateway Anywhere, you can also use the hybrid service to monitor users outside your network. This does not require software installation on remote machines. See *Chapter 4: Web Security Gateway Anywhere Deployments*.

When installing remote filtering components:

- The Remote Filtering Server should be installed on a dedicated machine that can communicate with the Filtering Service machine. See Table 6, on page 38.
- Do **not** install Remote Filtering Server on the same machine as Filtering Service or Network Agent.
- Each Filtering Service instance can have only one primary Remote Filtering Server.
- As a best practice, the Remote Filtering Server should be installed inside the outermost firewall, in the DMZ outside the firewall protecting the rest of the corporate network. This is strongly recommended.
- See Table 2, on page 16, for operating system requirements for Remote Filtering Server.

Table 6 Remote Filtering Server System Recommendations

Network Size	Hardware Recommendations
1-500 clients	**Windows** or **Linux** • Quad-Core Intel Xeon 5420 or better processor, 2.5 GHz or greater • 2 GB RAM • 20 GB free disk space
500+ clients	**Windows** or **Linux** • Quad-Core Intel Xeon 5450 or better processor, 3.2 GHz or greater • 4 GB RAM • 20 GB free disk space

Remote Filtering Client system recommendations:

- 32-bit Windows only (see Table 2, *Remote Filtering Client*, page 18)
- Pentium 4 or greater
- Free disk space: 25 MB for installation; 15 MB to run the application
- 512 MB RAM

Figure 1 provides an example of a Remote Filtering deployment. The illustration does not include all Websense components. For more information, see the Websense *Remote Filtering Software* technical paper.

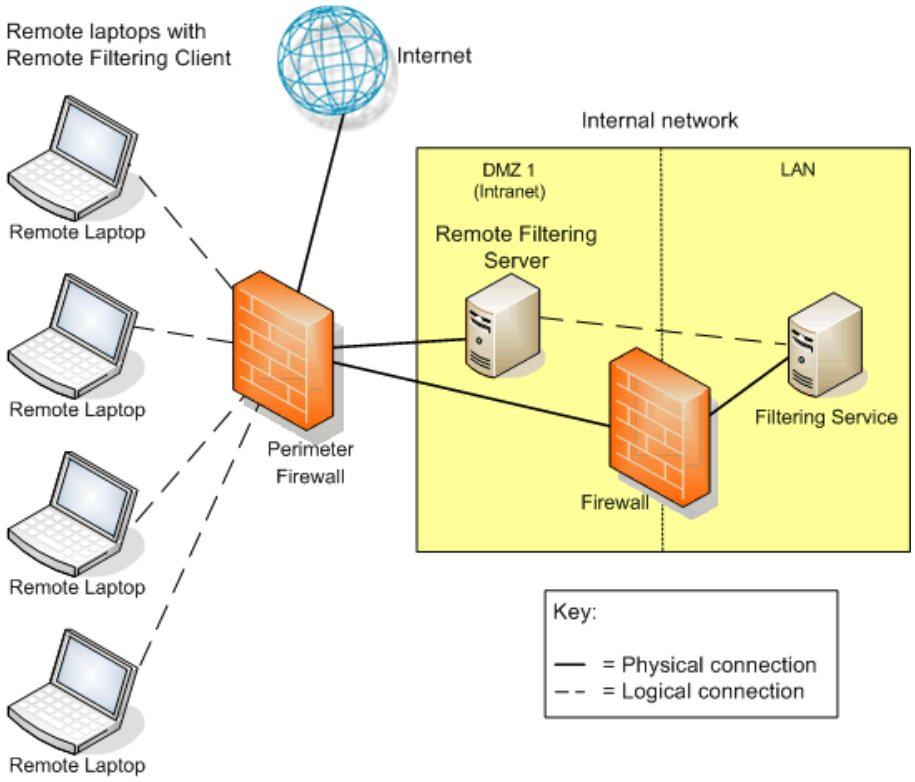

Figure 1 Example of Remote Filtering Deployment

Supported integrations

Websense software can be integrated with the following firewalls, routers, proxy servers, and caching applications (collectively referred to as **integration products**) to provide Internet filtering.

Table 7 Supported Integrations

Integration	Version Supported	Comments
Cisco®	• Cisco PIX Firewall Software v5.3 or greater • Cisco Adaptive Security Appliances (ASA) Software v7.0 or greater • Cisco Content Engine ACNS v5.4, v5.5, and v5.6 • Cisco Routers with Cisco IOS Software Release 12.3 or greater	
Check Point®	• FireWall-1 FP1 or greater • FireWall-1 NG AI • FireWall-1 NGX • Check Point Edge • Check Point UTM-1™ Edge	Contact Check Point for assistance in determining which FireWall-1 version is running.
Citrix® – MetaFrame® Presentation Server – Presentation Server™ – XenApp	• MetaFrame Presentation Server 3.0 • Citrix Presentation Server 4.0 • Citrix Presentation Server 4.5 • Citrix XenApp 5.0	**Requires Websense Plug-in**: The Websense Citrix Integration Service supports Presentation Servers 3.0, 4.0, and 4.5 on the following versions of Windows only: • Microsoft Windows 2000 Server (32-bit x86) • Microsoft Windows Server 2003 (32-bit x86) The Websense Citrix Integration Service supports XenApp 5.0 on the following versions of Windows only: • Microsoft Windows Server 2003 (32-bit x86) • Microsoft Windows Server 2008 SP2 (32-bit x86)

Table 7 Supported Integrations

Integration	Version Supported	Comments
Microsoft® Internet Security and Acceleration (ISA) Server	Integration products: • Microsoft ISA Server 2004, Standard Edition and Enterprise Edition • Microsoft ISA Server 2006, Standard Edition and Enterprise Edition Clients: • Firewall • SecureNAT • Web proxy	**Requires Websense Plug-in**: The Websense ISAPI Filter for Microsoft ISA Server supports the Windows versions supported by ISA Server 2004 and 2006, except Windows Server 2000.
Microsoft Forefront® Threat Management Gateway (TMG)	Integration product: • Microsoft Forefront TMG 2010 Clients: • Forefront TMG (formerly Firewall) • SecureNAT • Web proxy	**Requires Websense Plug-in**: The Websense ISAPI Filter for Microsoft Forefront TMG supports Windows Server 2008 (x64) only.
Squid Web Proxy Cache	• Squid STABLE v2.5 • Squid STABLE v2.6	**Requires Websense Plug-in**: The Websense redirector for the Squid Web Proxy Cache is supported only on 32-bit Red Hat Linux release 4, update 7 and release 5, update 3.

3 Deploying Network Agent

When your Websense software deployment includes Network Agent, the positioning of Network Agent and other Websense filter components depends on the composition of your network.

For the most part, Ethernet networks are built of **segments** (very simple networks are an exception). A segment is a sort of neighborhood for a group of machines, which are connected to the rest of the network via a central connection point (router, bridge, switch, or smart hub). Most of these devices keep local traffic within a segment, while passing traffic intended for machines on other segments. This architecture reduces network congestion by keeping unnecessary traffic from passing to the whole network.

A very simple network may require only a single Network Agent. A segmented network may require (or benefit from) a separate Network Agent instance for each segment. Network Agent functions best when it is closest to the computers that it is assigned to monitor.

This chapter provides configuration information and sample deployment diagrams to help you position Network Agent in your deployment. It contains the following sections:

Network Agent

Network Agent manages Internet protocols (including HTTP, HTTPS, and FTP), by examining network packets and identifying the protocol.

As with third-party integration products (like firewalls, routers, proxies, or network appliances), Network Agent can be configured to monitor HTTP requests and query Filtering Service to determine whether to allow or block a request, and then logs the results of the query. Network Agent can be configured to do the same for non-HTTP requests as well.

Network Agent must be installed on the **internal** side of the corporate firewall, in a location where it can see all Internet requests for the machines it is assigned to monitor. The agent then monitors HTTP and non-HTTP requests from those machines, and the response that they receive.

Network Agent only monitors and manages traffic that passes through the network device (switch or hub) to which it is attached. Multiple Network Agent instances may be needed, depending on the size, volume of Internet requests and the network configuration.

The Network Agent machine can connect to the network via a switch or a hub. See *Hub configuration*, page 50, and *Switched networks with a single Network Agent*, page 51.

Network Agent can be installed on the same machine as some integration products. See *Gateway configuration*, page 55.

Warning

Do **not** install Network Agent on a machine running a firewall or Remote Filtering Server. On a firewall, Network Agent's packet-capturing may conflict with the firewall software. On a Remote Filtering Server, machine resources may be too heavily taxed.

Network Agent settings

Configure Network Agent global (applying to all agent instances) and local (specific to a single agent instance) settings in TRITON - Web Security. These settings tell Network Agent which machines to monitor and which to ignore.

- Global settings:
  - Specify which machines are part of your network.
  - Identify any machines in your network that Network Agent should monitor for **incoming** requests (for example, internal Web servers).
  - Specify bandwidth calculation and protocol logging behavior.

- Local settings:
  - Specify which Filtering Service is associated with each Network Agent.
  - Identify proxies and caches used by the machines that this Network Agent monitors.
  - Determine which network interface card (NIC) the Network Agent instance uses to monitor requests and which it uses to send block pages.

    Configuration settings for the NIC used to monitor requests determine which segment of the network the agent instance monitors.

Network Agent location

Network Agent must be able to see all outgoing and incoming Internet traffic on the network segment that it is assigned to monitor. Multiple instances of Network Agent may be needed to monitor an entire network.

- Multiple Network Agents may be needed for larger or high-traffic organizations.
- A Network Agent instance can be placed in each internal network segment. Each instance should monitor its own segment without overlapping any other agent's segment.

The Network Agent machine may be:

- Connected to a **switch**.
  - Configure the device to use a mirror or span port, and connect Network Agent to this port, to allow the agent to see Internet requests from all monitored machines. (On most switches, you can change a port mode to spanning, mirroring, or monitoring mode. The term varies by manufacturer; the function is the same.)

    Note

    Not all switches support port spanning or mirroring. Contact the switch vendor to verify that spanning or mirroring is available, and for configuration instructions.

  - It is a best practice to use a switch that supports bidirectional spanning. This allows Network Agent to use a single network interface card (NIC) to both monitor traffic and send block pages.

    If the switch does not support bidirectional spanning, the Network Agent machine must have at least 2 NICs: one for monitoring and one for blocking. See *Using multiple NICs*, page 57.
- On a dedicated machine, connected to an **unmanaged, unswitched hub** located between an external router and the network.

To ensure that Network Agent is able to monitor the expected traffic, you must position the Network Agent machine appropriately and configure Network Agent

settings in TRITON - Web Security. See the TRITON - Web Security Help for instructions.

The following sections illustrate possible single- and multiple-Network Agent configurations.

Single segment network

A single segment network is a series of logically connected nodes (computers, printers, and so on) operating in the same portion of the network. In a single segment network, Filtering Service and Network Agent must be positioned to monitor Internet traffic across the entire network.

Figure 2 shows the filtering components in a stand-alone deployment of Websense software, installed in a central location to see both HTTP and non-HTTP traffic.

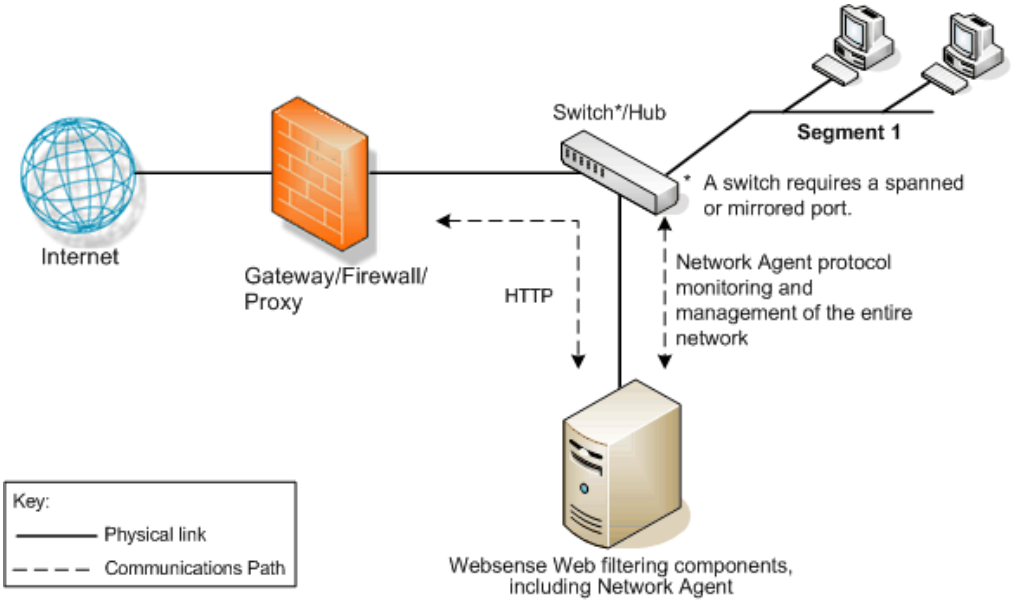

Figure 2 Websense software in a single-segment network

To learn more about installing Network Agent in a network:

- With a hub, see *Hub configuration*, page 50.
- With a switch, see *Switched networks with a single Network Agent*, page 51.
- With a gateway, see *Gateway configuration*, page 55.

Multiple segment network

Depending on the device used to connect network segments, some traffic may not be sent to all segments. A router, bridge, or smart hub serves as traffic control, preventing unneeded traffic from being sent to a segment. In this environment, the Websense filtering components must be deployed to see all network traffic.

- Filtering Service must be installed where it can receive and manage Internet requests from the integration product, if any, and communicate with Network Agent.
- Each Network Agent instance must be able to see all Internet requests on the segment or segments that it is configured to monitor.

Deploying multiple Network Agents

Multiple Network Agent instances may be needed in a multiple segment network to capture all Internet requests. A Network Agent can be installed on each segment to monitor the Internet requests from that segment.

> **Note**
> A limit of 4 Network Agents is suggested for each Filtering Service. It may be possible to use more agent instances, depending on system and network configuration and the volume of Internet requests. See *Network Agents per Filtering Service*, page 24.

If multiple Network Agent instances are installed:

- Ensure that the instances are deployed such that they, together, monitor the entire network. Partial deployment results in incomplete filtering and loss of log data in network segments not watched by Network Agent.
- Network Agent instances must not be configured to monitor overlapping IP address ranges. An overlap can result in inaccurate logging and network bandwidth measurements, and improper bandwidth-based filtering.

  The network segment or IP address range monitored by each Network Agent is determined by the NIC settings for the agent configured in TRITON - Web Security. See the TRITON - Web Security Help for instructions.
- Avoid deploying Network Agent across different LANs. If you install Network Agent on a machine in the 10.22.x.x network, and configure it to communicate with a Filtering Service machine in the 10.30.x.x network, communication may be slow enough to prevent Network Agent from blocking an Internet request before the site is returned to the user.

For examples of central and distributed Network Agent placement, see:

- *Hub configuration*, page 50
- *Switched networks with a single Network Agent*, page 51.
- *Gateway configuration*, page 55

Central Network Agent placement

A network with multiple segments can be filtered from a single location. Install Filtering Service where it can receive Internet requests from both the integration product, if any, and each Network Agent.

If the network contains multiple switches, Network Agent instances are inserted into the network at the last switch in the series. This switch must be connected to the gateway that goes out to the Internet.

In Figure 3:

- One Network Agent instance is installed with Filtering Service on Machine A. This machine is connected to the network via a switch that is configured to mirror or span the traffic of network Segment 1.
- A second Network Agent is installed on Machine B, which is connected to the same switch as Machine A. Machine B is connected to a different port that is configured to mirror the traffic of Segments 2 and 3.
- Both Network Agents are positioned to see all traffic for the network segments they monitor, and to communicate with other Websense components.
- The switch is connected to the gateway, allowing the Network Agent instances to monitor network traffic for all network segments.

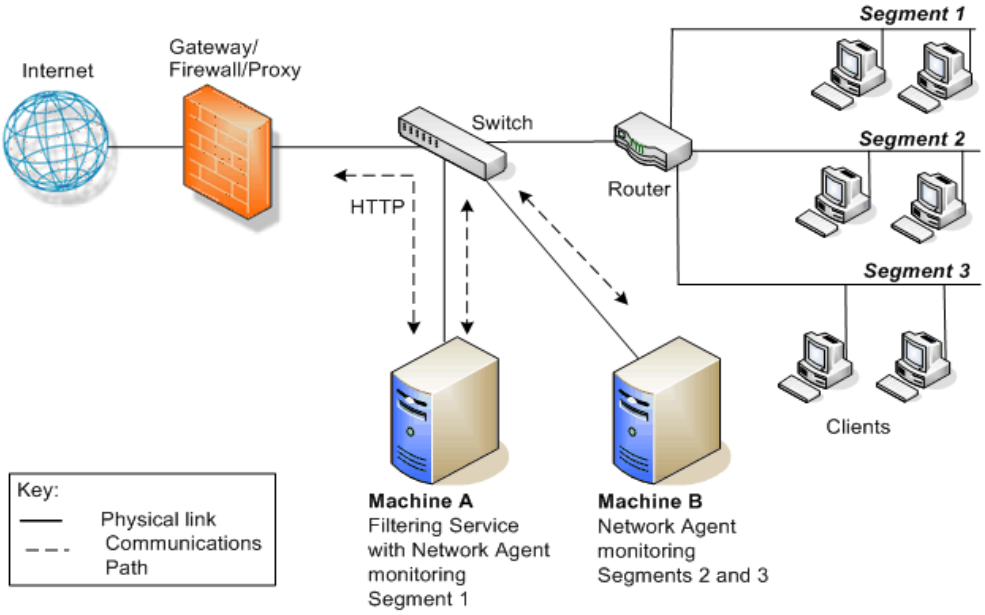

Figure 3 Websense software in a multiple-segment network

Distributed Network Agent placement

The network diagram below shows a single Filtering Service with 3 Network Agents, one for each network segment. A deployment like this might be useful in organizations with satellite offices, for example.

- Filtering Service (Machine C) must be installed where it is able to receive and manage Internet requests from both the integration product (if any) and each of the Network Agent instances in all network segments.
- Each Network Agent (machines A, B and C) is connected to the network segment it monitors via the span or mirror port of a switch.

See *Deploying multiple Network Agents*, page 47, for more information.

In Figure 4, the switches are not connected in a series. However, each switch is connected to the router, which is connected to the gateway.

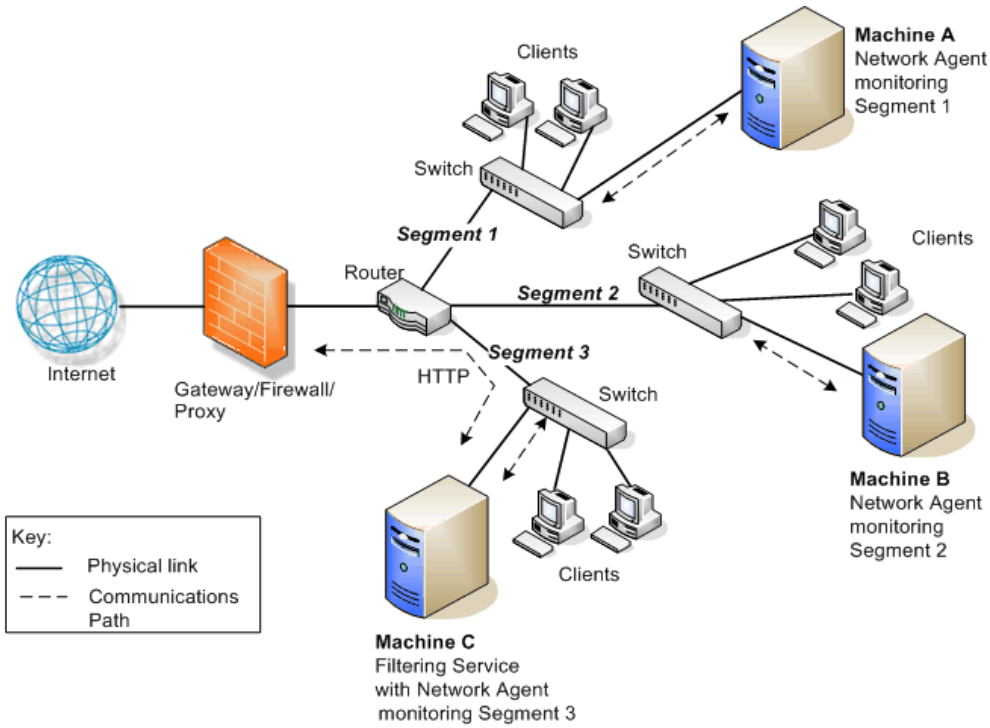

Figure 4 Multiple Network Agents in a multiple-segment network

Hub configuration

At the simplest level, a network hub provides a central connection point for the segments in a network and the devices in those segments. The port to which the Network Agent machine connects is dependent on the type of hub. Some hubs broadcast traffic to all of their ports, while others do not.

Network Agent must be able to see the traffic for the network segments it is assigned to monitor.

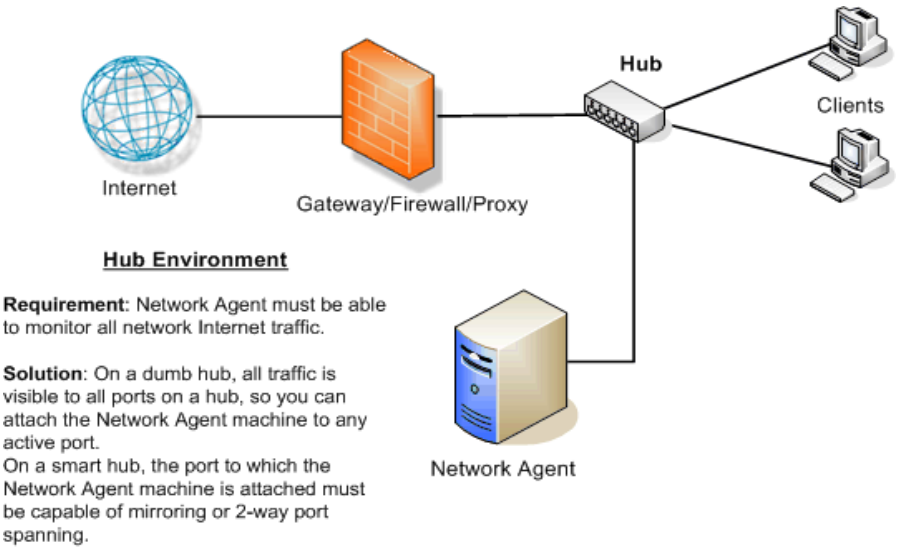

Figure 5 Network Agent connected to a hub

Switched networks with a single Network Agent

A switch is a bridge that routes traffic between network segments. It prevents all traffic from going to all segments, reducing network congestion. Since not all traffic going through a switch is visible to all devices on the network, the machine running Network Agent must be connected at a point where it can monitor all Internet traffic for the network.

Connect the Network Agent machine to the port on the switch that mirrors, monitors, or spans the traffic on the gateway or firewall port. The span or mirror port sees all the traffic that leaves each network segment.

Note

Not all switches support bidirectional port spanning or mirroring. Contact the switch vendor to verify that spanning or mirroring is available, and for configuration instructions.

If bidirectional communication is not available, at least 2 network interface cards (NICs) are needed to monitor traffic and communicate with other Websense components.

If port spanning is not available, Network Agent cannot properly monitor the network.

Figure 6 shows a network with a single switch. The Network Agent machine is attached to the port that mirrors all traffic from connected clients. Subsequent illustrations show multiple switch and multiple subnetwork configurations.

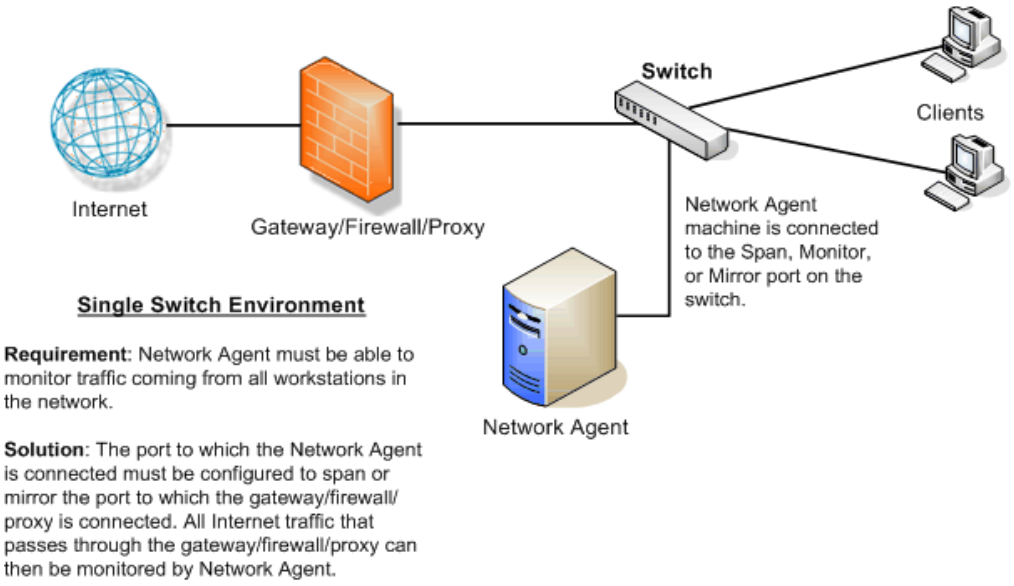

Figure 6 Simple deployment in a switched environment

Figure 7 shows the use of additional switches to create 2 network segments. All Internet traffic from these network segments must pass through Switch #3, to which Network Agent is attached. In a multiple-switch environment, failure to enable port spanning or mirroring could result in missed filtering and inaccurate reports.

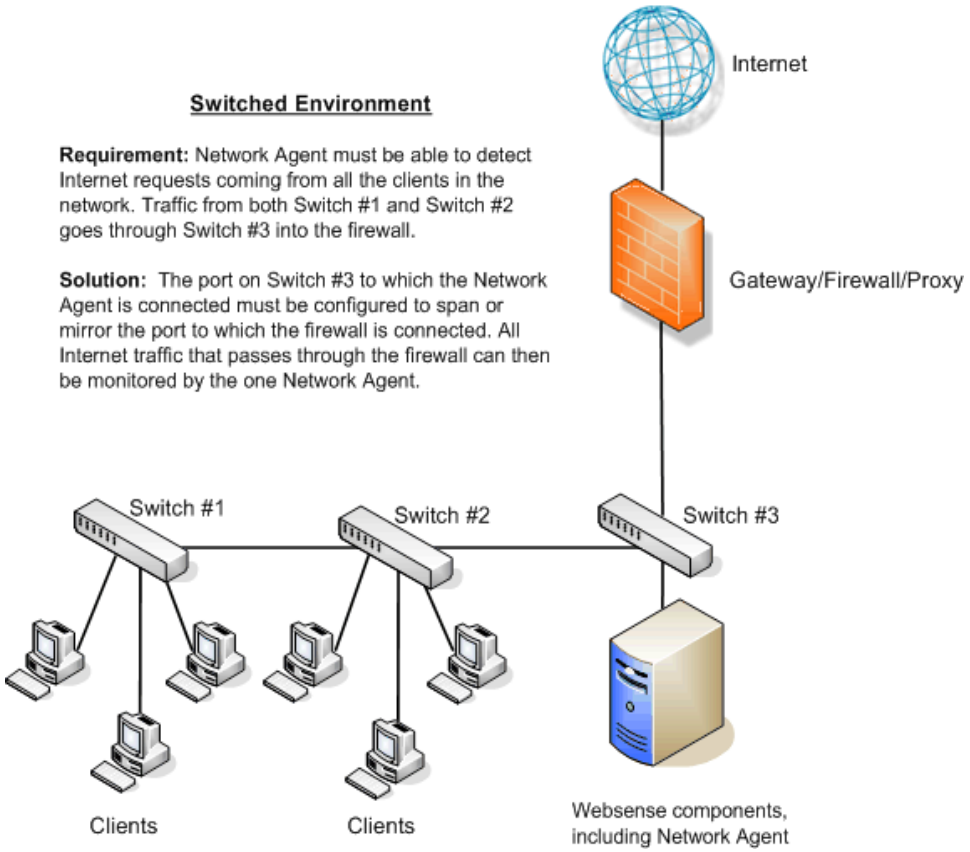

Figure 7 Multiple segments in a switched environment

Figure 8 also contains multiple network segments. A remote office is filtered by installing another instance of Network Agent and configuring it to communicate with the Filtering Service at the main office.

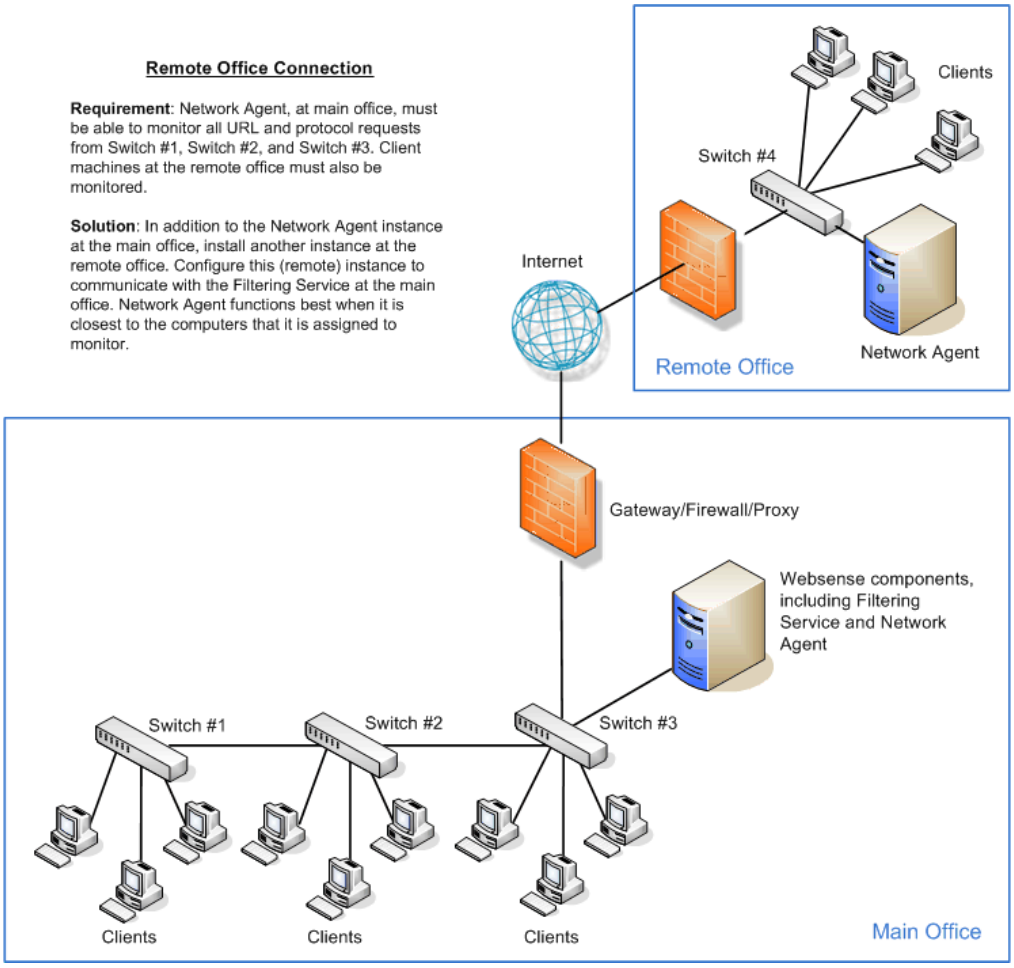

Figure 8 Switched environment with a remote office connection

To improve performance, Network Agent can be deployed on its own, dedicated machine. Network Agent can also be positioned closer to the clients, as shown in Figure 9, page 54.

Switched networks with multiple Network Agents

A busy network may need multiple Network Agents to monitor different network segments or IP address ranges. Network Agent operates best when it is close to the computers it is assigned to monitor. Figure 9 shows a network in which multiple Network Agent instances monitor separate network segments.

See *Deploying multiple Network Agents*, page 47, for more information.

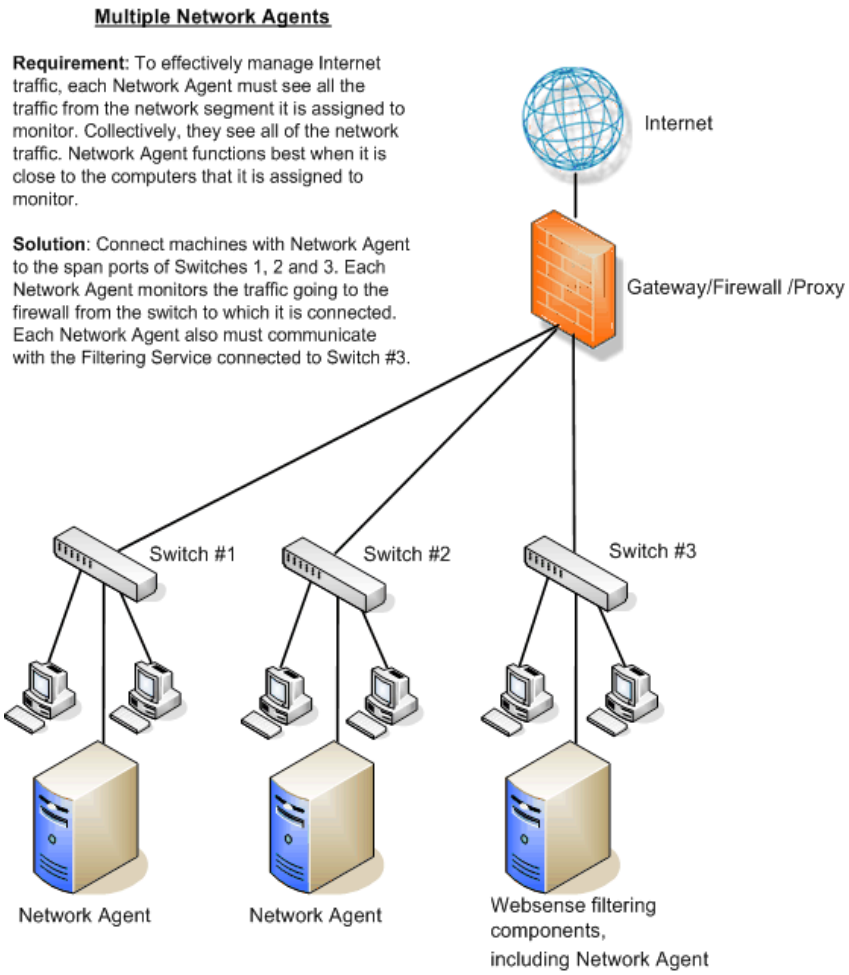

Figure 9 Multiple Network Agents in a switched environment

Gateway configuration

A gateway provides a connection between two networks. The networks do not need to use the same network communication protocol. The gateway can also connect a network to the Internet.

Network Agent can be installed on the gateway machine, allowing Network Agent to manage and monitor all Internet traffic. The gateway can either be a third-party proxy server or a network appliance. Do **not** install Network Agent on a firewall.

Important

The gateway configuration shown here is best used in small to medium networks.

In larger networks, performance can suffer as a result of resource competition between the gateway software and Network Agent.

Figure 10 shows Network Agent monitoring the Internet traffic at the proxy gateway or caching appliance directly attached to the firewall.

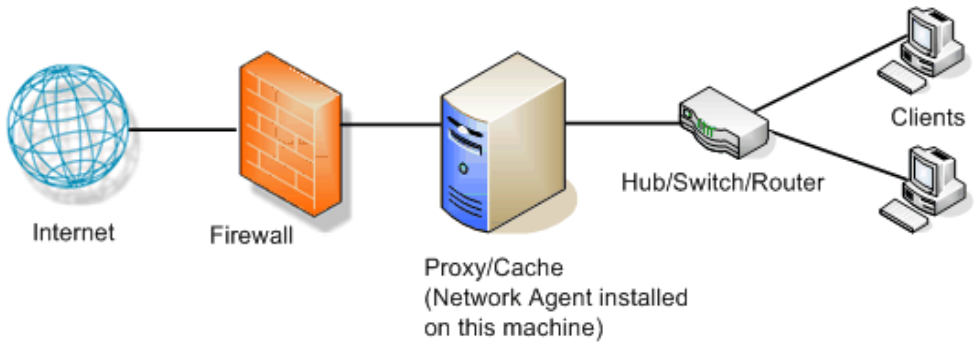

Figure 10 Network Agent installed on the gateway

Figure 11 shows Network Agent deployed in a network that includes Websense Content Gateway. Do not install Network Agent on the Websense Content Gateway machine.

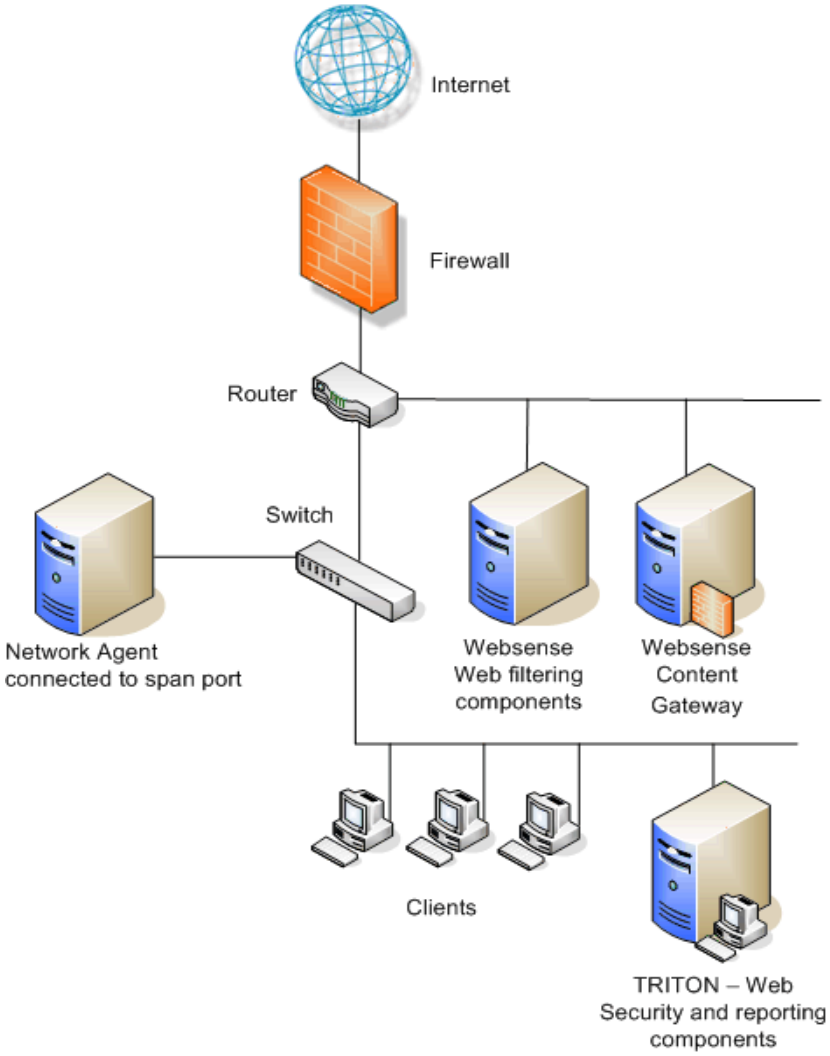

Figure 11 Network Agent deployed with Websense Content Gateway

Using multiple NICs

Network Agent is capable of using more than one network interface card (NIC).

- Best practices suggest a maximum of 5 NICs
- The NICs can be connected to ports on the same network device (switch or router), or to different network devices.

If the machine running Network Agent has multiple NICs:

- Only one instance of Network Agent can be installed on the machine.
- The blocking or inject NIC (used to serve block pages) must have an IP address.
- Each NIC can be configured to monitor or block Internet requests, or both.
- Each NIC can be configured to monitor a different network segment.
- At least one NIC must be configured for blocking.

When you configure separate network cards to monitor traffic and send block messages (shown in Figure 12, page 58):

- The monitoring and blocking NIC do not have to be assigned to the same network segment.
- The monitoring NIC must be able to see all Internet traffic in the network segment that it is assigned to monitor.
- Multiple monitoring NICs can use the same blocking NIC.
- The blocking NIC must be able to send block messages to all machines assigned to the monitoring NICs, even if the machines are on another network segment.
- A monitoring NIC can be set for **stealth mode**. For information on configuring stealth mode, see the Websense Web Security and Websense Web Filter *Installation Guide*.
- The blocking NIC **must** have an IP address (cannot be set to stealth mode).

During installation, you specify which NIC is used by Websense software for communication and which NIC or NICs are used by Network Agent. For more information, see the Websense Web Security and Websense Web Filter *Installation Guide*.

For information on configuring multiple NICs, see the TRITON - Web Security Help.

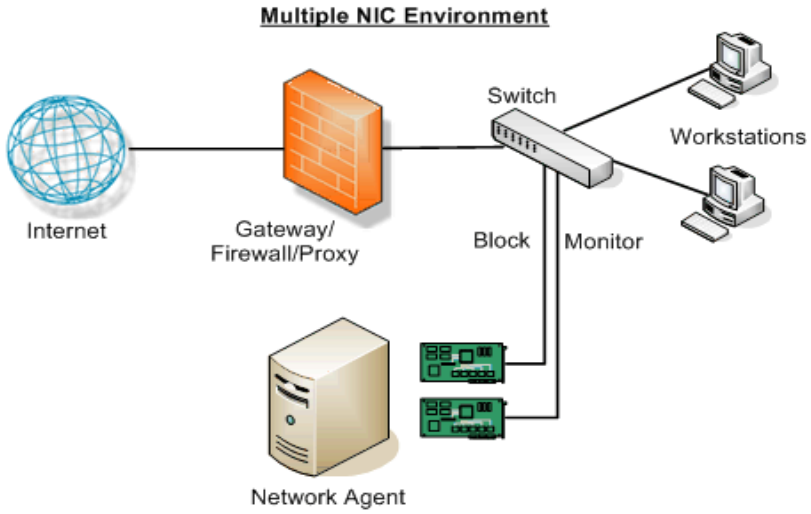

Figure 12 Dual NIC configuration

NAT and Network Agent deployment

If you use Network Address Translation (NAT) on internal routers, Network Agent may be unable to identify the source IP address of client machines. When Network Agent detects traffic after being passed through such a router, the agent sees the IP address of the router's external interface as the source of the request, rather than the IP address of the client machine.

To address this issue, either disable NAT, or install Network Agent on a machine located **between** the NAT router and the monitored clients.

4 Web Security Gateway Anywhere Deployments

Websense Web Security Gateway Anywhere is a Web security solution that consists of several modules in addition to Websense Web Security. These include:

- Websense Content Gateway
- Websense Data Security Management Server
- Websense Sync Service
- Websense Directory Agent

This chapter discusses where to place these components in your network. It contains the following sections:

For instructions on installing and configuring Web Security Gateway Anywhere, refer to the Websense Web Security Gateway Anywhere *Getting Started Guide*.

Web Security Gateway Anywhere

Websense Web Security Gateway Anywhere is a Web security solution designed for distributed enterprises with one or more branch offices and multiple remote users.

Web Security Gateway Anywhere offers an alternative to pure service- or appliance-based solutions. Rather than choosing between an in-the-cloud or on-premises Web filtering solution for your entire enterprise, you can deploy a blended solution that encompasses the best of both worlds, and you can manage it from a single user interface—the TRITON Unified Security Center.

You can decide which method to use for which users, but typically you use our robust on-premises Web filtering for your corporate office (business) or main campus

(education), and filter your regional offices or satellite locations through our hybrid service.

In addition, Web Security Gateway Anywhere protects you from data loss over the Web, providing security for outbound content as well. You identify sensitive data and define whether you want to audit or block attempts to post it to HTTP, HTTPS, FTP, or FTP-over-HTTP channels.

And finally, Web Security Gateway Anywhere provides flexible solutions for users who travel or work from a location outside of your network, such as a home office. You can install a Web filtering client on remote users' machines, or you can monitor remote activity using our hybrid filtering service.

Web Security Gateway Anywhere is available on the V-Series appliance or as software. Which version you purchase affects how you deploy your system.

Appliance configuration

If you purchase an appliance-based Web Security Gateway Anywhere solution, the following components are pre-loaded on the appliance:

- Websense Web Security core components, including:
  - Policy Database
  - Policy Broker
  - Policy Server
  - Filtering Service
  - User Service
  - Usage Monitor
  - Control Service
  - Directory Agent
  - TRITON - Web Security (optional; can be run on separate Windows server)
    - Investigative Reports Scheduler
    - Manager Web Server
    - Reporting Web Server
    - Reports Information Service
- Websense Content Gateway
  - Content Gateway Manager
  - Content Cop
  - Download Service
- Network Agent (optional)

Larger enterprises might use 2 or more Websense appliances, with one designated as the *policy source* machine (the only machine to run Policy Broker and Policy Database, along with other components). All other appliances point to the *policy*

source machine for policy updates. Alternatively, you can add a Windows or Linux server and designate it as the *policy source*.

In all cases, Network Agent and Websense Content Gateway run as separate modules on each appliance, if they are enabled.

Regardless of how many appliances you have, the following Websense Web Security components must be installed separately. Most are Windows-only components.

- Log Server
- Sync Service
- Linking Service
- (optional) Transparent identification agents
  - DC Agent
  - Logon Agent
  - eDirectory Agent
  - RADIUS Agent

Note that TRITON - Web Security can be installed on one or more machines in addition to the appliance. TRITON - Web Security on the appliance is enabled by default. For production use, Websense recommends running it off the appliance. (You configure which manager to use in the Appliance Manager.)

In addition, the Websense Data Security Management Server must be installed on a Windows server. This includes:

- Policy Engine
- Crawler
- PreciseID Fingerprint Repository
- Forensics Repository

The off-box Web and data security components can be installed on the same Windows machine using VMWare, if desired. (See the Websense Web Security Gateway Anywhere *Getting Started Guide* for more information.)

Finally, you are required to have a Windows database server running Microsoft SQL Server. This is where the Log Database is built. It does not matter whether SQL Server is installed in a virtualization environment or not, or whether it’s on the same hardware as the Web and data security components.

Software configuration

If you purchase Web Security Gateway Anywhere as software, you are required to install the various components as described below:

- Websense Web Security filtering components on a Windows or Linux machine.
- TRITON - Web Security and Data Security Management Server on a Windows machine in 2 virtual machines. TRITON - Web Security can be installed on a Windows or Linux machine.
- Log Server, Sync Service, Directory Agent, and Linking Service on a Windows machine. It is a best practice to install Sync Service on the same machine as Log Server. Note that, while Log Server and Linking Service run on Windows only, Sync Service and Directory Agent can run on either Windows or Linux (these components could be installed on a separate Linux machine if you choose).
- Log Database on your Microsoft SQL Server database engine.
- Websense Content Gateway on a Linux machine (includes Content Gateway Manager).

Optionally, you may install transparent identification agents as needed.

Websense Content Gateway requirements

Hardware

CPU	Quad-core running at 2.8 GHz or faster
Memory	4 GB
Disk space	2 disks: • 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: – Must be a raw disk (not a mounted file system) – Must be dedicated – Must *not* be part of a software RAID – For best performance, use a 10k RPM SAS disk on a controller that has at least 64MB of write-through cache.
Network Interfaces	2

To support transparent proxy deployments:

Router	WCCP v1 routers support redirection of HTTP only. If your deployment requires additional protocols, such as HTTPS, your router must support WCCP v2. A Cisco router must run IOS 12.2 or later. The clients, the destination Web server, and Websense Content Gateway must reside on different subnets.
—*or*—	
Layer 4 switch	You may use a Layer 4 switch rather than a router. To support WCCP, a Cisco switch requires the EMI or IP services image of the 12.2SE IOS release (or later). 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).

Software

Linux operating system:

- Red Hat Enterprise Linux 5, update 3 or later, base or Advanced Platform (32-bit only)
  - Only kernels shipped with the above Linux versions are supported by Websense Content Gateway. Visit www.redhat.com for kernel information. To display the kernel version installed on your system, enter the command:

    ```
    /bin/uname -r
    ```

Important

If SELinux is enabled, disable it before installing Websense Content Gateway.

- PAE (Physical Address Extension)-enabled kernel required
  - By default, Red Hat Enterprise Linux 5, update 3 and later has PAE enabled. If you are running the non-PAE kernel, reboot with the PAE-enabled kernel before installing Websense Content Gateway.
- RPM compat-libstdc++-33-3.2.3-47.3.i386.rpm (or higher version of this package)
  - To display a list of RPMs installed on your system with the string "compat-libstdc" in their name, enter the command:

    ```
    rpm -qa |grep compat-libstdc
    ```
- GNU C library (glibc) version 2.5-42
  - Note that Red Hat Enterprise Linux 5, update 3 ships with glibc version 2.5-34. Be sure to update it to version 2.5-42.
  - Example command to update this library (running as `root`): `yum update glibc`.

Supported browsers:

- Websense Content Gateway is configured and maintained with a Web-based user interface called the Content Manager. Supported browsers include:
  - Internet Explorer 7 or 8
  - Mozilla Firefox 3.0.x - 3.5.x

Data Security Management Server requirements

Hardware

	Minimum Requirements	Recommended
CPU	2 Dual-core Intel Xeon processors (2.0 GHz) or AMD equivalent	2 Quad-core Intel Xeon processors (2.0 GHz) or AMD equivalent Note: The Management Server can **not** have more than 8 cores.
Memory	2 GB	4 GB
Hard drives	4 - 72 GB	4 - 146 GB
Disk space	144 GB	292 GB
Hardware RAID	1 + 0	1 + 0
NICs	1	2

Software

- Windows 2003 R2 Standard Edition with the latest SP.

  For optimized performance of Websense Web Security Solutions, verify that the operating system's file cluster is set to 4096B. For more information, see the Websense knowledge-base article: "File System Performance Optimization."
- Windows installation requirements:
  - Set the partition to 1 NTFS Partition. For more information, see the Websense knowledge-base article: "File System Performance Optimization."
  - Regional Settings: should be set according to the primary location. If necessary, add supplemental language support and adjust the default language for non-Unicode programs.
  - Configure the network connection to have a static IP address.
  - The Data Security Manager computer name must not include an underscore sign. Internet Explorer does not support such URLs.
  - Short names cannot be enabled.
  - Create a local administrator to be used as a service account.
  - It's necessary to set the system time accurately on the server onto which you install the Data Security Server.

- Application Server
  - ASP.NET
  - Create a local administrator to be used as a service account.
  - It's necessary to set the system time accurately on the server onto which you install the Data Security Server.

In addition, the Data Security Management Server requires the following to support the Data Security Manager user interface:

- Adobe Flash Player v8 or beyond

  This is required for the Data Security Today and System Health dashboards. All the other functions of the manager interface can operate without Flash. If users do not already have Flash Player, they are prompted to install it when they log on, and a link is supplied.
- One of the following Web browsers:
  - Internet Explorer 7
  - Internet Explorer 8
  - Firefox 3.0.x - 3.5.x

  If you have another browser or version, the user interface may behave in unexpected ways or report an error.

Network diagram - appliance

This is what a Web Security Gateway Anywhere deployment might look like when deployed on a V-Series appliance in a small network. Larger networks might have 2 Windows machines.This diagram depicts a sample deployment. Alternative deployments are possible. For example, the components on Windows Server 1 could be distributed across several machines. And some of those components could be installed on Linux machines. See *General Deployment Recommendations*, page 15 for supported operating systems and software requirements for each component.

Figure 13 Websense Web Security Gateway Anywhere on appliance

Network diagram - software

This is what Web Security Gateway Anywhere deployment looks like when software is installed in a small to medium network. In larger networks, filtering components may be distributed across multiple machines.This is just an example of how the modules might fit together. You can arrange them in numerous ways to fit your needs. For example, Sync Service and Directory Agent (on Server 3) could be installed on a Linux machine. See *General Deployment Recommendations*, page 15 for supported operating systems and software requirements for each component.

Figure 14 Websense Web Security Gateway Anywhere as software

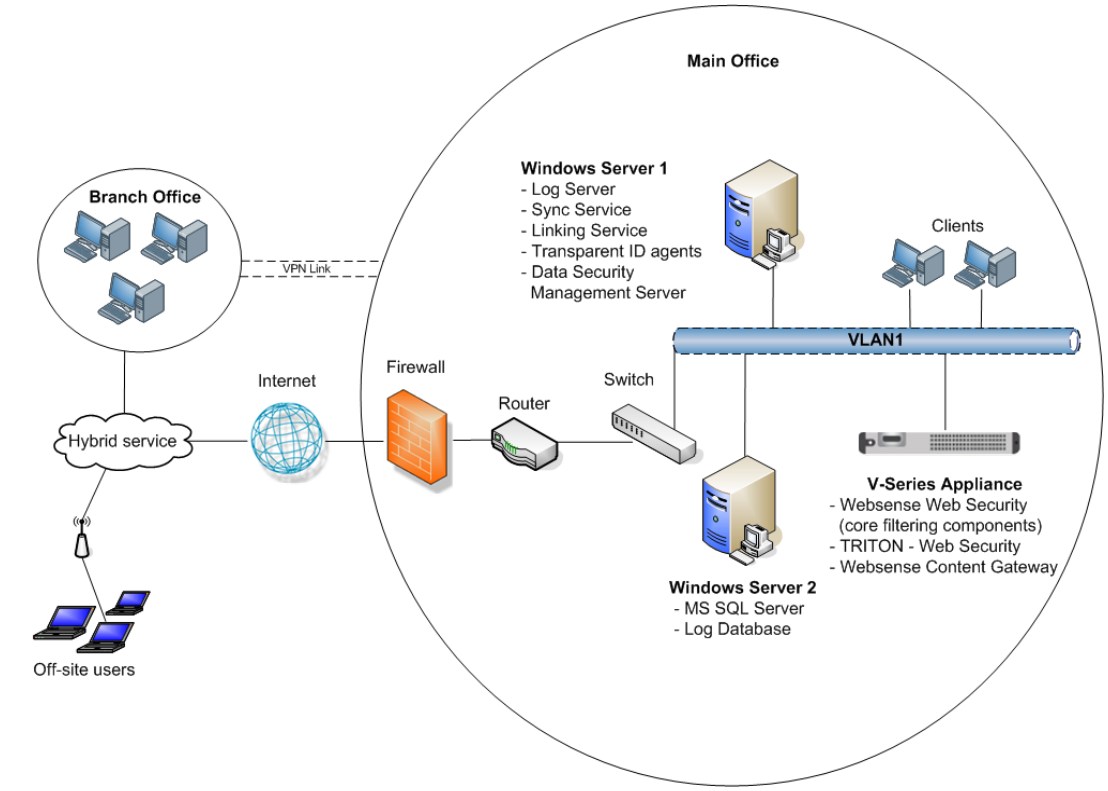
Main Office
Windows Server 1
- Log Server
- Sync Service
- Linking Service
- Transparent ID agents
- Data Security
Management Server
Clients
VLAN1
Branch Office
VPN Link
Internet
Firewall
Router
Switch
Hybrid service
V-Series Appliance
- Websense Web Security
(core filtering components)
- TRITON - Web Security
- Websense Content Gateway
Windows Server 2
- MS SQL Server
- Log Database
Off-site users

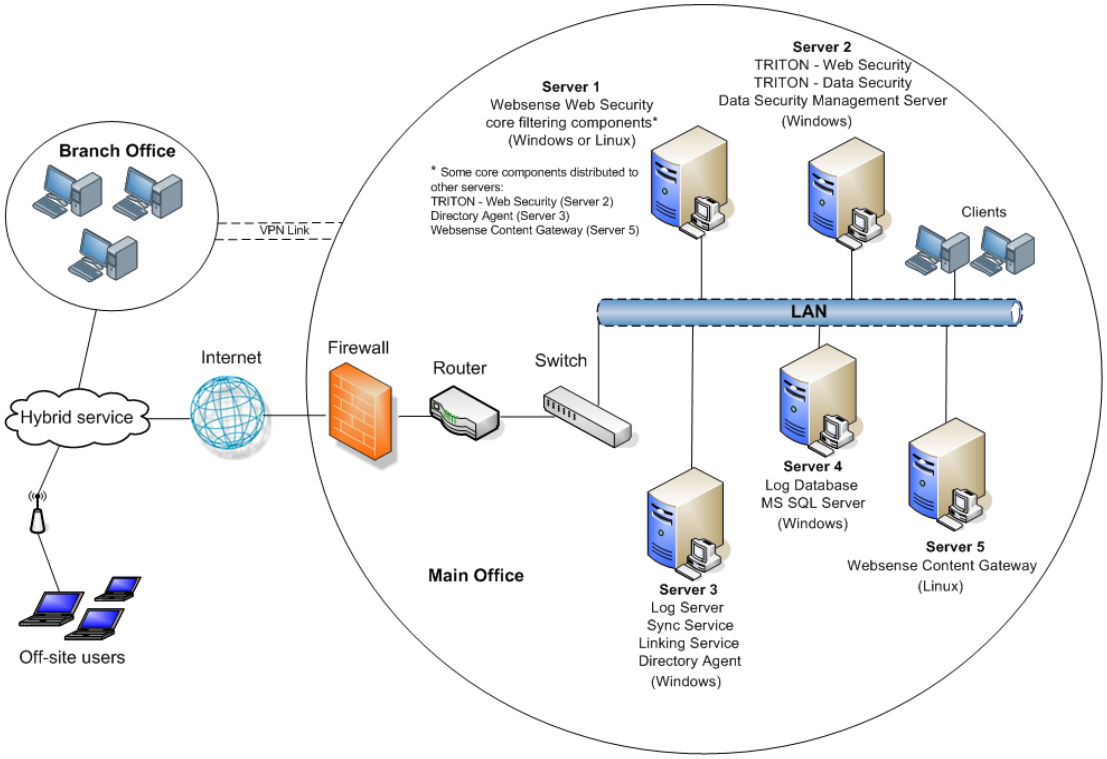
Server 1
Websense Web Security
core filtering components*
(Windows or Linux)
Server 2
TRITON - Web Security
TRITON - Data Security
Data Security Management Server
(Windows)
* Some core components distributed to other servers:
TRITON - Web Security (Server 2)
Directory Agent (Server 3)
Websense Content Gateway (Server 5)
Branch Office
VPN Link
Clients
LAN
Internet
Firewall
Router
Switch
Hybrid service
Server 4
Log Database
MS SQL Server
(Windows)
Server 5
Websense Content Gateway
(Linux)
Main Office
Server 3
Log Server
Sync Service
Linking Service
Directory Agent
(Windows)
Off-site users

5 Integration Deployment

This chapter addresses considerations for deploying Websense components with an integration product (such as a firewall, proxy, or caching application). It contains the following sections:

Most of the network diagrams included in this chapter show a typical small network installation (500 users or fewer). The diagrams show the recommended location of your integration product relative to Websense components.

- The diagrams are intended to provide a general overview, and do not show all Websense components.
- Larger networks require Websense components to be distributed across several dedicated machines.

> **Note**
> DC Agent is listed as the transparent identification agent in many of the diagrams in this chapter. Logon Agent can also be used.

Websense Content Gateway

Websense Content Gateway™ is a central gateway for controlling Web content that offers:

- The advantages of a proxy cache, improving bandwidth usage and network performance by storing requested Web pages and, if the page is still considered "fresh," serving the Web page to the requesting client.
- Real-time content categorization. This feature examines the content of uncategorized sites and sites that include rapidly changing content, and then returns a recommended category to Filtering Service.

Websense Content Gateway can run in explicit or transparent proxy mode.

- In explicit proxy mode, client browsers must be configured to point to Content Gateway.
- In transparent proxy mode, the client request is intercepted and redirected to the proxy. Traffic is redirected through a router or a Layer 4 switch and the ARM (Adaptive Redirection Module) feature of Content Gateway.

Websense Content Gateway can participate in flexible cache hierarchies, where Internet requests not fulfilled in one cache can be routed to other regional caches. Content Gateway also scales from a single node into multiple nodes to form a cluster, improving system performance and reliability.

Websense Content Gateway is installed on a Linux machine, separate from other Websense components. See the Websense Content Gateway and Websense Web Security Gateway *Installation Guide* for more information.

Figure 15 shows Websense Content Gateway and Websense Data Security deployed with Websense Web filtering components (including Policy Broker, Policy Server, Filtering Service, User Service, and a transparent identification agent).

- The Websense Data Security, Websense Content Gateway, and Websense filtering component machines access network traffic through a router.
- Network Agent is installed on a separate machine, attached to the span port on a switch.

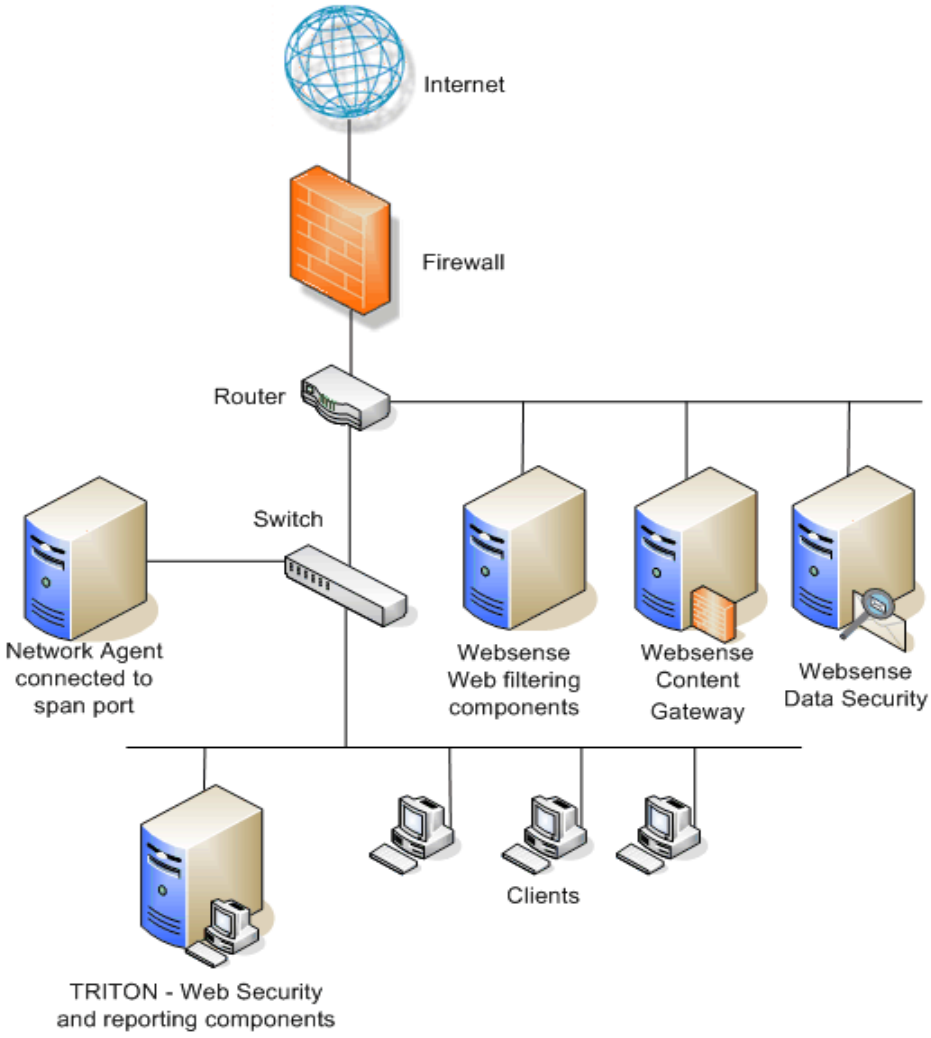

Figure 15 Integration with Websense Content Gateway

Microsoft ISA Server or Forefront TMG

Note

In this section, “ISA/TMG” refers to ISA Server and Forefront TMG collectively. When information differ for the two products, they are referred to specifically as “ISA Server” or “Forefront TMG”.

When you integrate Websense software with Microsoft ISA/TMG, the Websense ISAPI Filter must be installed on the ISA/TMG machine. The Websense ISAPI Filter allows ISA/TMG to communicate with Filtering Service, and must be installed on every ISA/TMG machine that communicates with Websense software.

You can install Policy Broker, Policy Server, Filtering Service, and User Service on the same machine as Microsoft ISA Server.

IMPORTANT

No Websense components, other than the ISAPI Filter and Control Service, can be installed on a Forefront TMG machine. Control Service is automatically installed when installing the plug-in.

If your environment includes an array of ISA/TMG machines, install Websense software on a machine outside the array.

See the *Installation Guide Supplement for use with Microsoft ISA Server or Forefront TMG* for instructions and more information.

Single Microsoft ISA Server configuration

Figure 16 shows all Websense components, including the Websense ISAPI Filter, running on the same machine as Microsoft ISA Server. Unless the Internet traffic volume is light, this configuration requires a powerful machine.

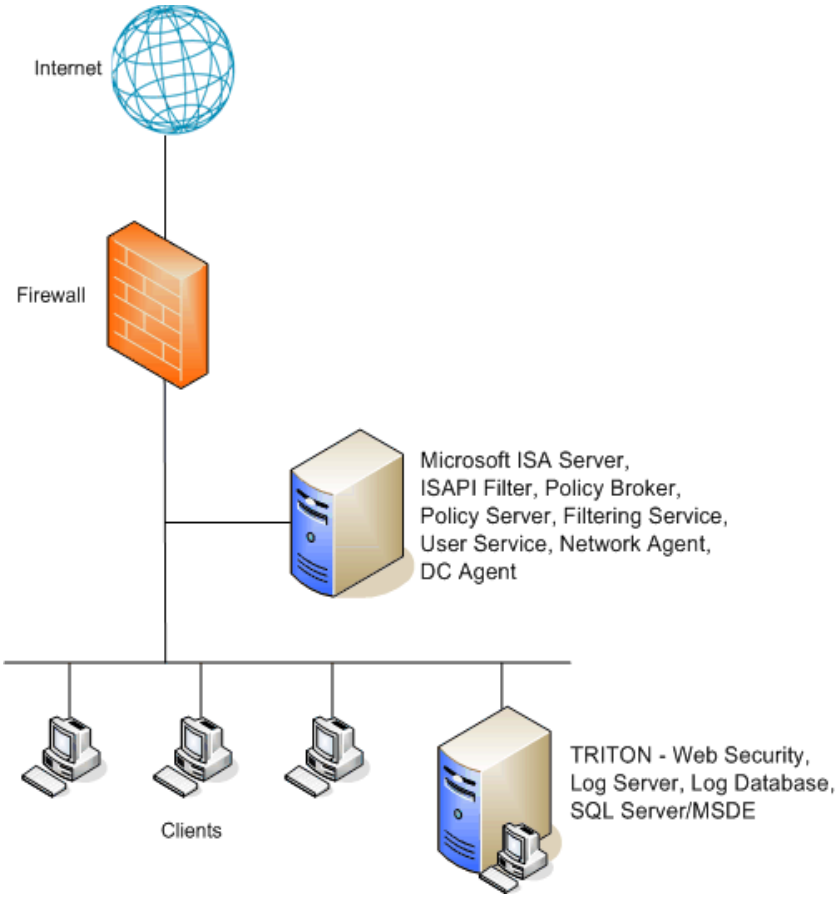

Figure 16 Filtering components installed with Microsoft ISA Server

> **NOTE**
>
> This configuration—Websense components on the same machine as ISA Server—is not possible with Forefront TMG. Websense components cannot run on the same machine as Forefront TMG because they do not support 64-bit Windows Server 2008 (x64).

An alternative setup, Figure 17, places Websense filtering components on a Windows machine separate from the ISA/TMG machine. This configuration eases the load on the ISA/TMG machine.

- The ISAPI Filter must be installed on the ISA/TMG machine so that Internet activity information can be communicated to Filtering Service.
- The Filtering Service and ISA/TMG machines must be able to communicate over the network.

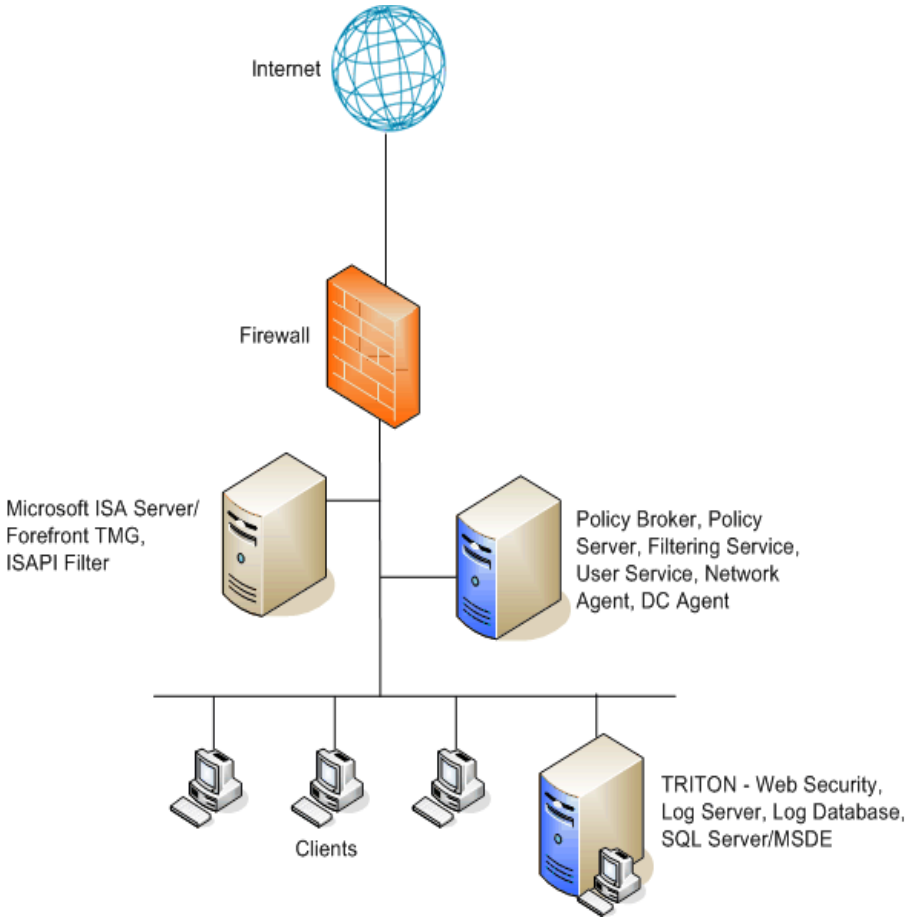

Figure 17 Filtering components installed separately from Microsoft ISA Server/ Forefront TMG

Array configuration

Websense software is compatible with most array configurations, including Cache Array Routing Protocol (CARP) arrays. It is a best practice to install Websense software outside an array of ISA/TMG machines. Install the Websense ISAPI Filter on each member of the array. See Figure 18.

> **NOTE**
> With an array of Forefront TMG machines, all Websense components (other than the ISAPI Filter and Control Service) must be installed outside the array. Forefront TMG runs on 64-bit Windows Server 2008 (x64), which is not supported by Websense components other than the ISAPI Filter and Control Service.

When Websense software is deployed in this configuration, all array members send Internet requests to Filtering Service outside the array.

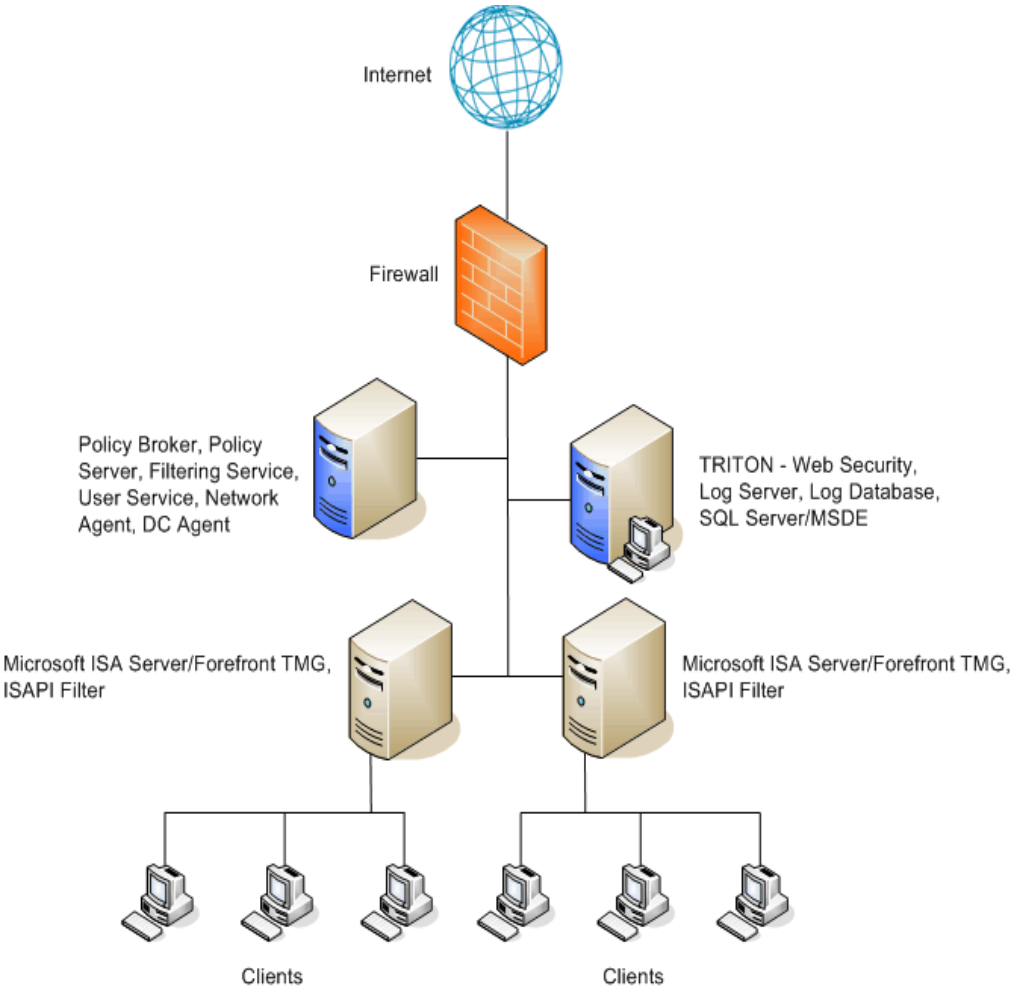

Figure 18 Microsoft ISA Server/Forefront TMG array configuration

Other configurations are possible. See your Microsoft ISA/TMG documentation for information about ISA/TMG configurations.

Cisco deployment

Cisco PIX/ASA

A simple and common network topology places Websense filtering components on a single machine, or group of dedicated machines, communicating with a Cisco PIX Firewall or Cisco Adaptive Security Appliance (ASA) via TCP/IP.

- TRITON - Web Security and reporting components are installed on a separate machine.
- If you install Network Agent, it must be positioned to see all traffic on the internal network.

See the *Installation Guide Supplement for use with Cisco Integrated Products* for configuration instructions.

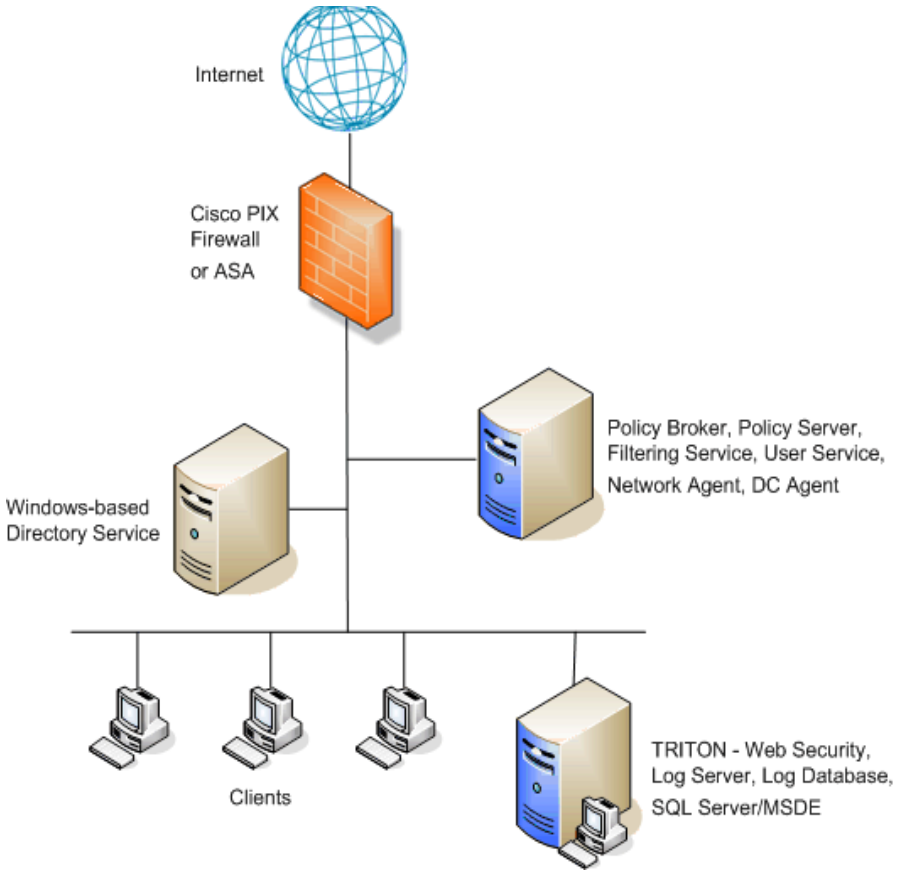

Figure 19 Common Windows Network Configuration for Cisco PIX Firewall or ASA

Other configurations are possible. See your Cisco PIX Firewall or ASA documentation and the information in this section to determine the best configuration for your network.

Cisco Content Engine

In this common configuration, Websense filtering components are installed on a single machine, communicating with the Cisco Content Engine through TCP/IP.

- TRITON - Web Security and reporting components are installed on a separate machine.
- If you install Network Agent, it must be positioned to see all traffic on the internal network.

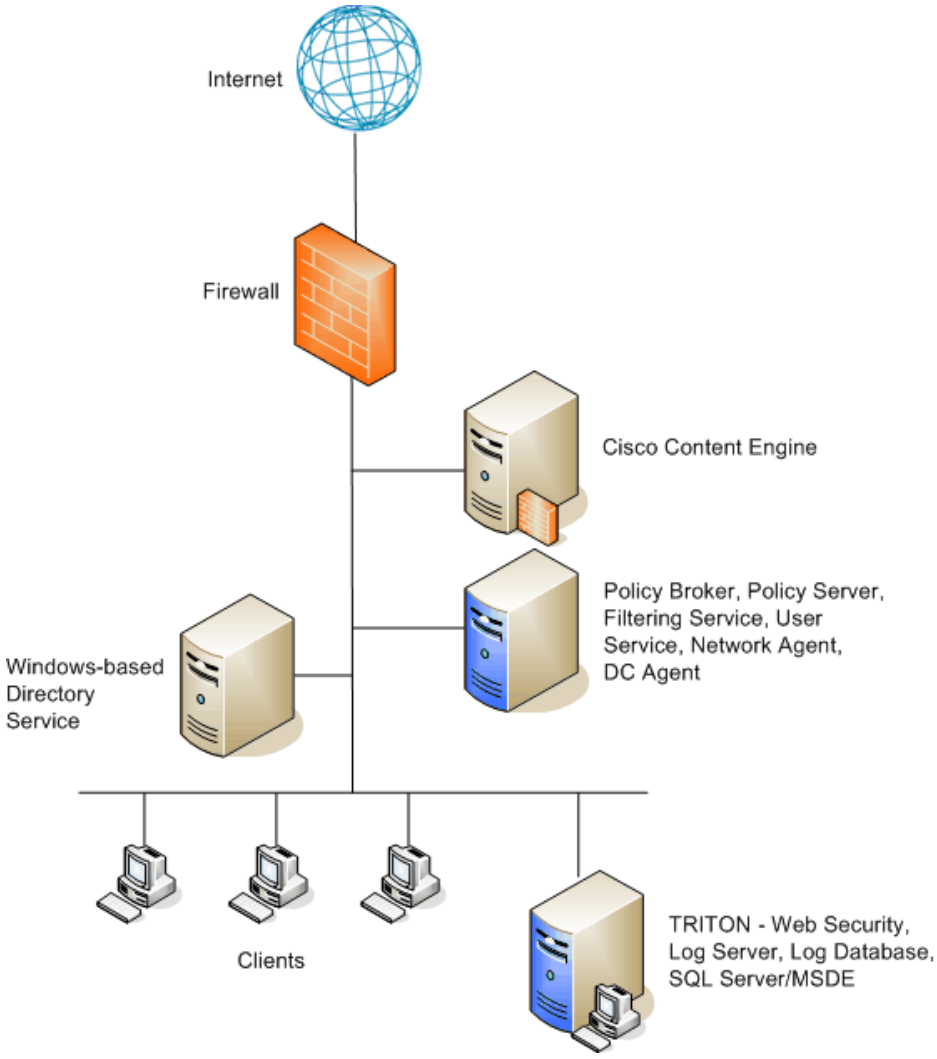

Figure 20 Common Windows network configuration for Cisco Content Engine

Other configurations are possible. See your Content Engine documentation and the information in this chapter to determine the best configuration for your network.

Cisco IOS Routers

In this common configuration, Websense filtering components are installed on a single machine, communicating with the Cisco IOS Router.

- TRITON - Web Security and reporting components are installed on a separate machine.
- If you install Network Agent, it must be positioned to see all traffic on the internal network.

The router has firewall functionality and can be used with or without an accompanying firewall.

If the Cisco IOS Router is used with a separate firewall, ensure that all Internet traffic is configured to pass through the router and is not set to bypass the router and go directly to the firewall. Traffic filtered through the separate firewall cannot be filtered by the Websense software.

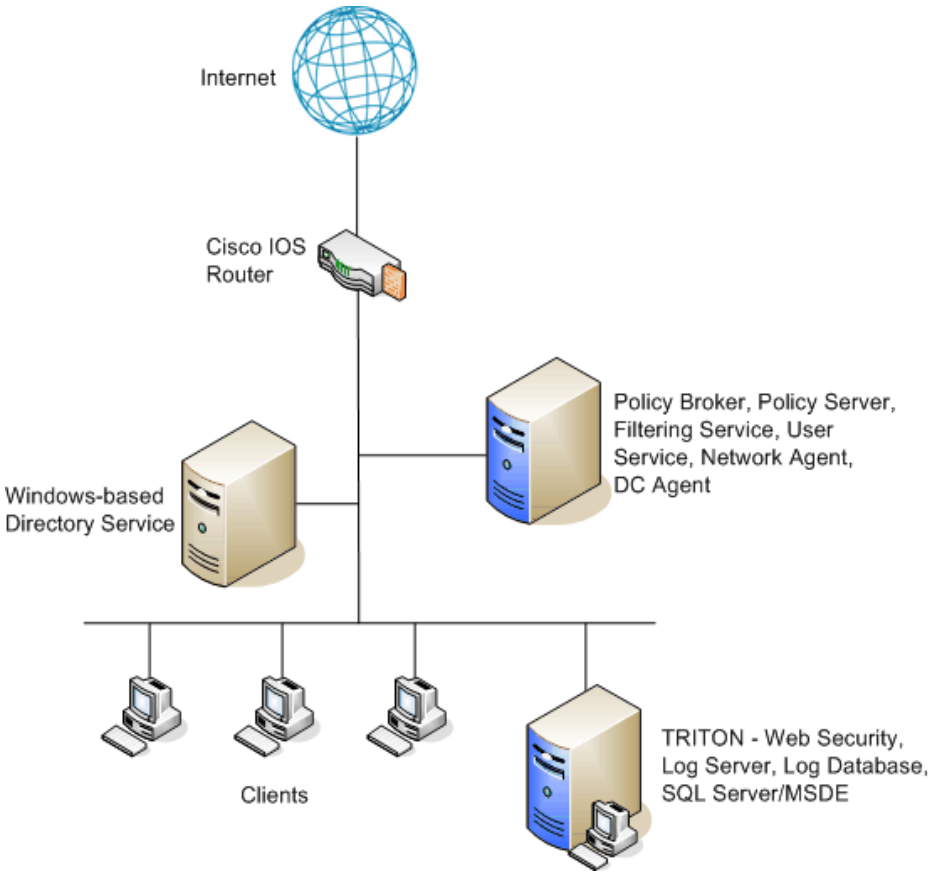

Figure 21 Common Windows network configuration for Cisco IOS Routers

Other configurations are possible. See your Cisco Router documentation and the information in this chapter to determine the best configuration for your network.

Check Point

This section includes a general discussion of 2 common Check Point integration deployment options: simple deployment with unified components, and distributed deployment. See the *Installation Guide Supplement for use with Check Point Integrated Products* for configuration instructions.

Simple

In the simplest and most common network topology, an organization has one firewall that resides on a dedicated server. All Websense filtering components are installed on a separate machine on the internal network.

- TRITON - Web Security and reporting components are installed on a separate machine.
- If you install Network Agent, it must be positioned to see all traffic on the internal network. HTTP requests are handled by the Check Point appliance, and non-HTTP traffic is managed by Network Agent, which is positioned to detect all outbound traffic.

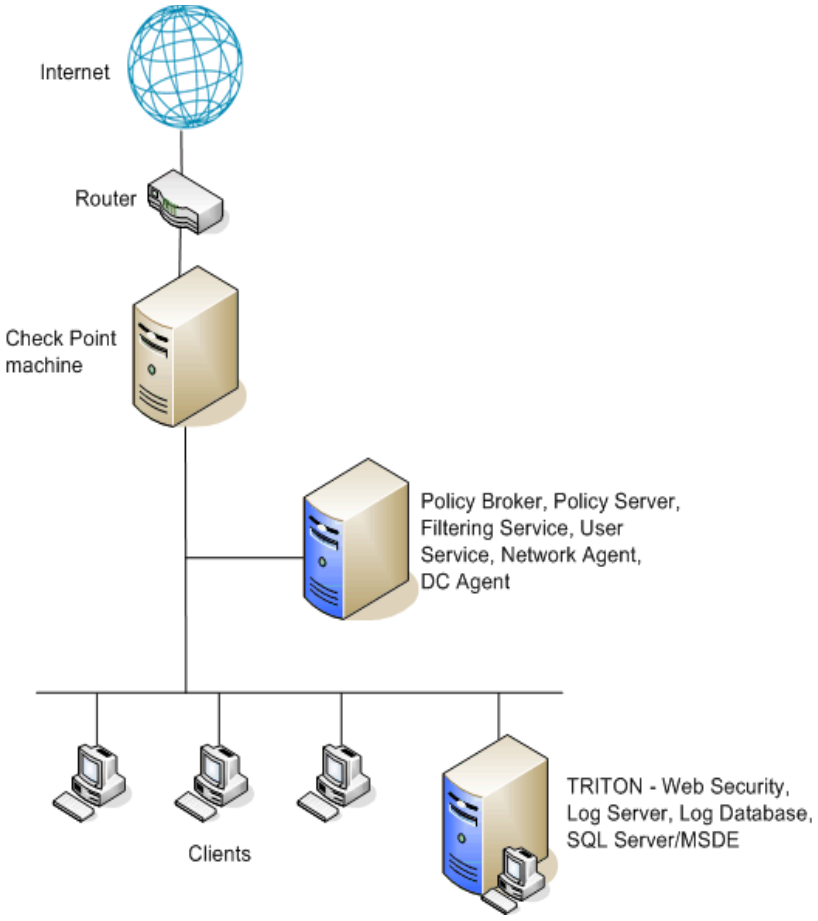

Figure 22 Simple network configuration

Distributed

In Figure 23, Websense filtering software is installed on a single machine in a central location where it can manage both protocol and HTTP traffic. HTTP requests are handled by the Check Point appliance, and the non-HTTP traffic is managed by Network Agent, which is positioned to detect all outbound traffic.

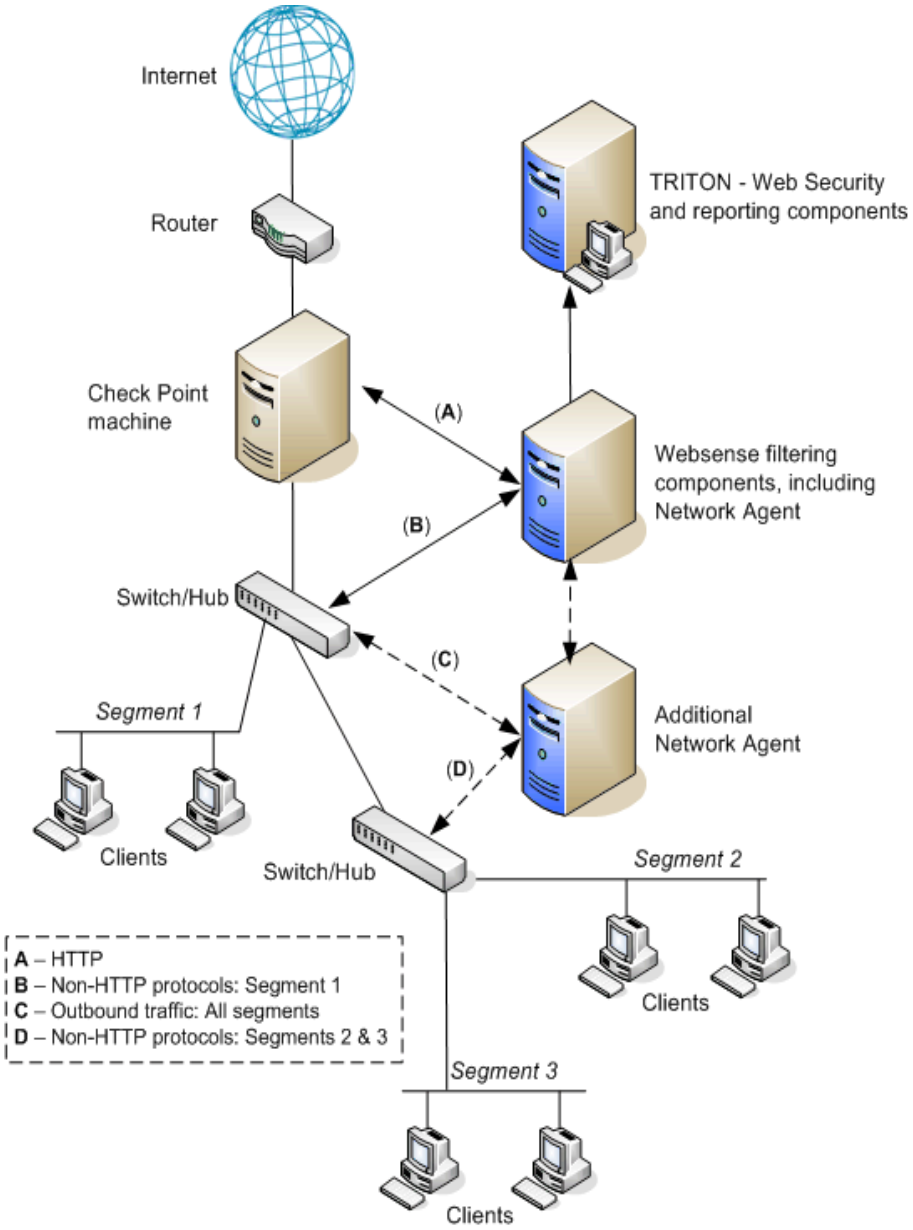

Figure 23 Multi-segment network configuration

To avoid performance and security issues, do **not** install Websense components on a machine running Check Point software. Network Agent will not function correctly if installed on the Check Point machine.

Warning

Websense, Inc., and Check Point do not recommend installing Websense software and Check Point on the same machine. Do **not** install Network Agent on the same machine as Check Point software.

Squid Web Proxy Cache deployment

Websense filtering components can be installed on the same machine as Squid Web Proxy Cache, on a separate machine, or on multiple machines.

Squid Web Proxy Cache machines may be deployed in an array to share the load in a larger network.

A Websense Squid redirector plug-in must be installed on each machine running Squid Web Proxy Cache.

The diagrams in this section assume a Linux installation.

Single Squid Web Proxy Cache configuration

Figure 24 shows the Websense filtering components, the Websense redirector plug-in, and Squid Web Proxy Cache running on the same machine. You can also install a

Websense transparent identification agent on the same machine, or on a separate machine.

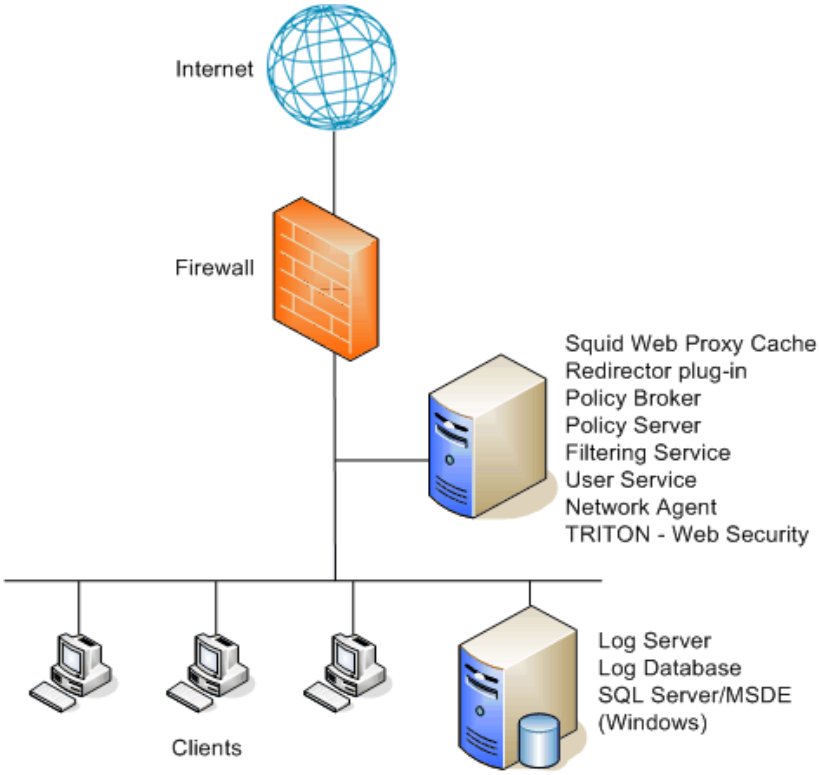

Figure 24 Filtering components installed with Squid Web Proxy Cache

An alternative deployment places all Websense filtering components on a separate machine from Squid Web Proxy Cache. This configuration eases the load on the Squid Web Proxy Cache machine.

- The Websense redirector plug-in must be installed on the Squid Web Proxy Cache machine to enable communication with Filtering Service.
- The Filtering Service and Squid Web Proxy machines must be able to communicate over the network

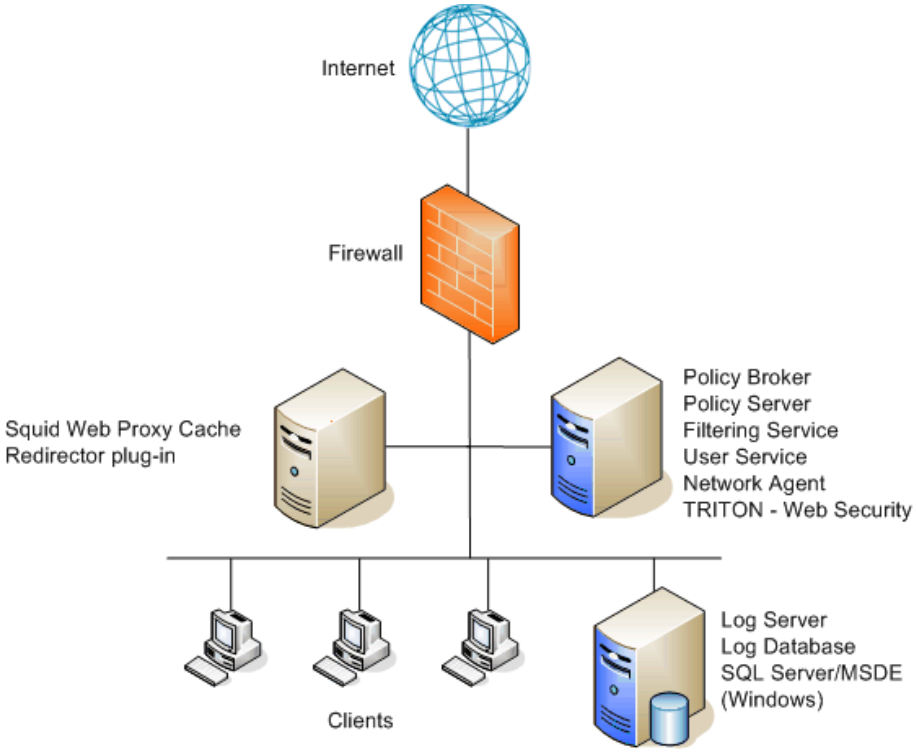

Figure 25 Filtering components and Squid Web Proxy Cache on separate machines

Array configuration

Websense software is compatible with most array configurations, including Cache Array Routing Protocol (CARP) arrays. If the Squid Web Proxy Cache machines in an array can run Websense software without a performance impact, install the main Websense filtering components on one of the array machines.

Figure 26 shows the Websense filtering components running on a Squid Web Proxy Cache machine, with Websense reporting components on a separate machine.

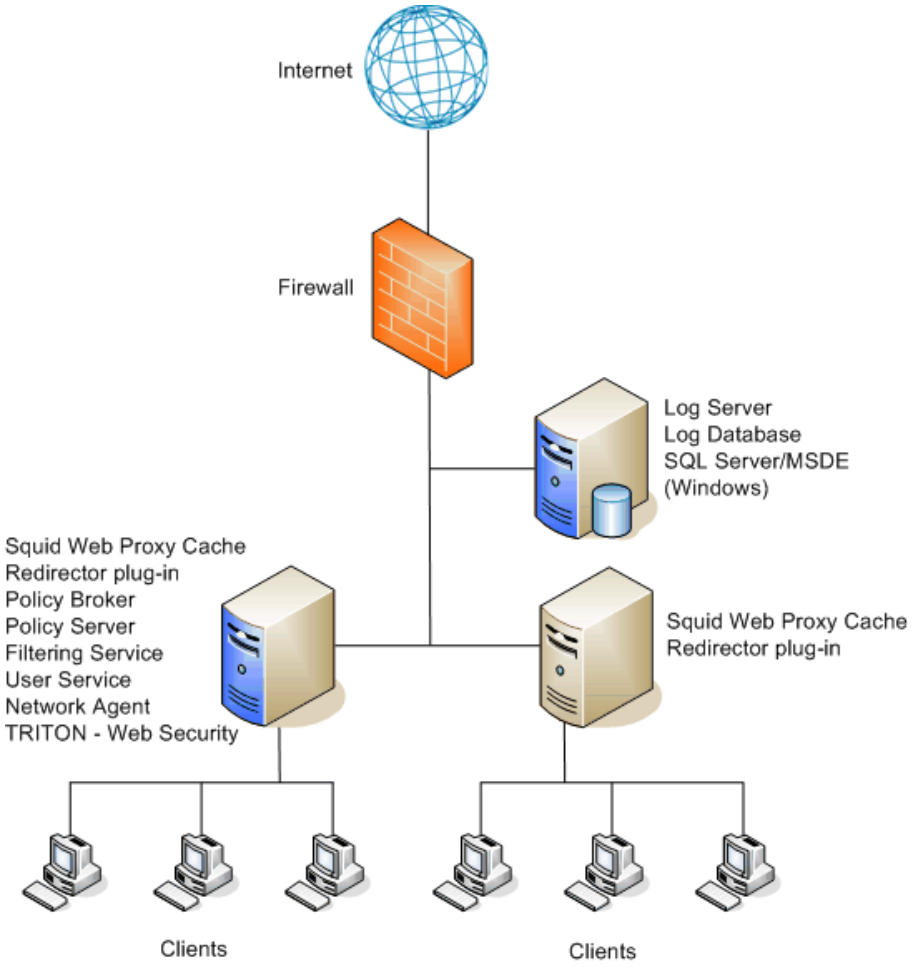

Figure 26 Squid Web Proxy Cache array configuration #1

If installing the Websense filtering components on the Squid Web Proxy Cache machine is likely to have a performance impact, install Websense software on a separate machine outside the array, and then install the Websense redirector plug-in on each member of the array. See Figure 27.

When Websense software is installed in this configuration, all array members send Internet requests to Filtering Service outside the array.

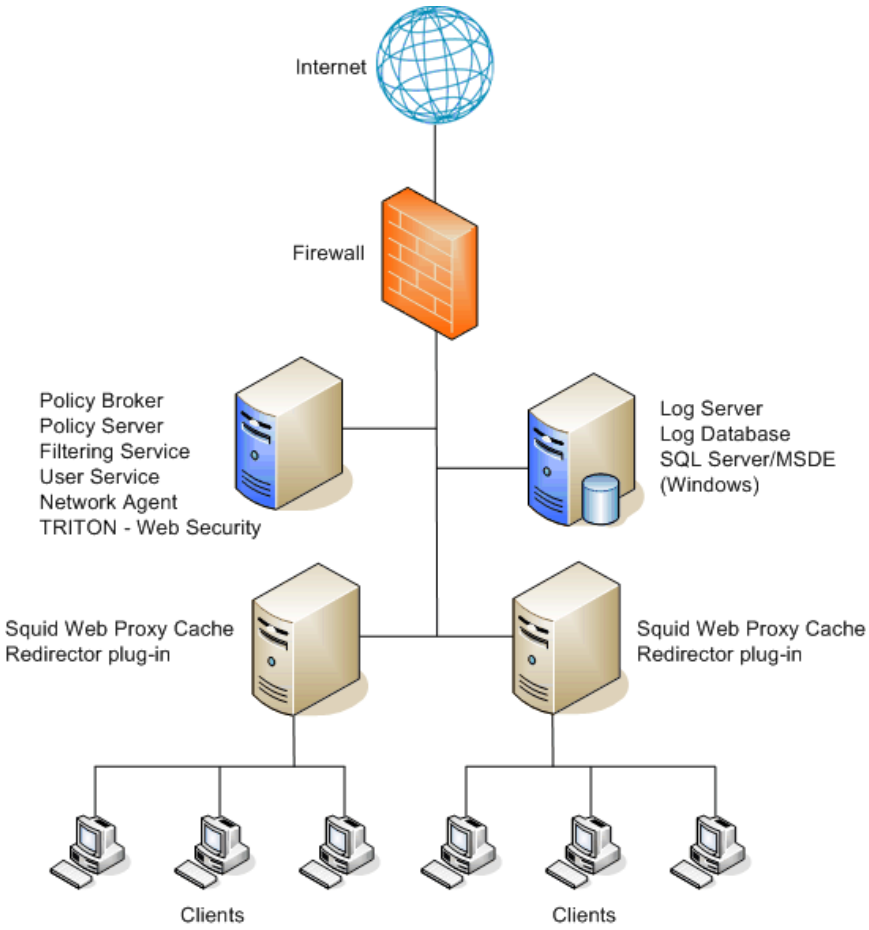

Figure 27 Squid Web Proxy Cache array configuration #2

Other configurations are possible. See your Squid Web Proxy Cache documentation for information about array configurations. See the *Installation Guide Supplement for use with Squid Web Proxy Cache* for Websense software configuration instructions.

Citrix

Websense software integrated with a Citrix server can monitor HTTP, FTP, and SSL requests from individual Citrix users. Network Agent can be used to filter other protocols, if needed.

> **Note**
> "Citrix server" refers to Citrix Presentation Server or XenApp. For the versions are supported by Websense software, see *Supported integrations*, page 40.

Figure 28 shows a typical deployment used to filter both users who access the Internet through a Citrix server and users who access the Internet locally.

- The Websense filtering components are installed on a dedicated machine that can filter Citrix server clients (non-Citrix clients are filtered by a separate integration product or Network Agent; see the *Installation Guide supplement for use with Integrated Citrix Servers*).
- The Websense Citrix Integration Service must be installed on each Citrix server to allow it to communicate with Filtering Service.
- No other Websense components can be installed on a Citrix server.

Separate Network Agent instances are needed for the Citrix and non-Citrix users.

To simplify the diagram, not all individual Websense components are shown.

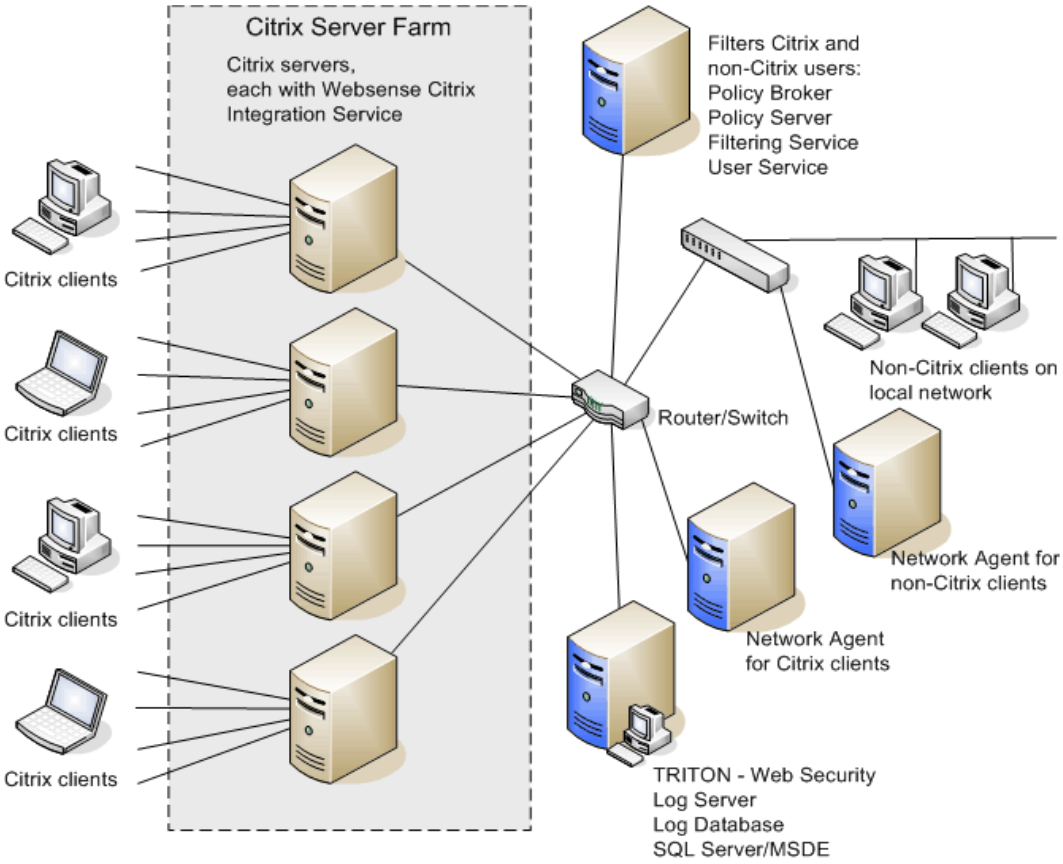

Figure 28 Citrix integration

Other integrations also can be used in the non-Citrix portion of the network. See the *Installation Guide Supplement for use with Integrated Citrix Servers* for Websense software configuration instructions.

Universal integration

If your firewall, proxy server, caching application, or network appliance is not one of the products listed in this chapter, you may still be able to integrate it with Websense software. Check the list of Websense Technology Partners at www.websense.com/global/en/Partners/TAPartners/SecurityEcosystem/ to see if Websense software can be integrated with the product. If your integration product is listed, that product has been specifically enhanced to integrate with Websense software.

Typical configurations include networks with a single firewall, proxy server, or caching application, and networks with an array of firewalls, proxy servers, or caching appliances. A Websense transparent identification agent (DC Agent, Logon Agent, eDirectory Agent, or RADIUS Agent) can be installed on the Filtering Service machine or on a separate machine.

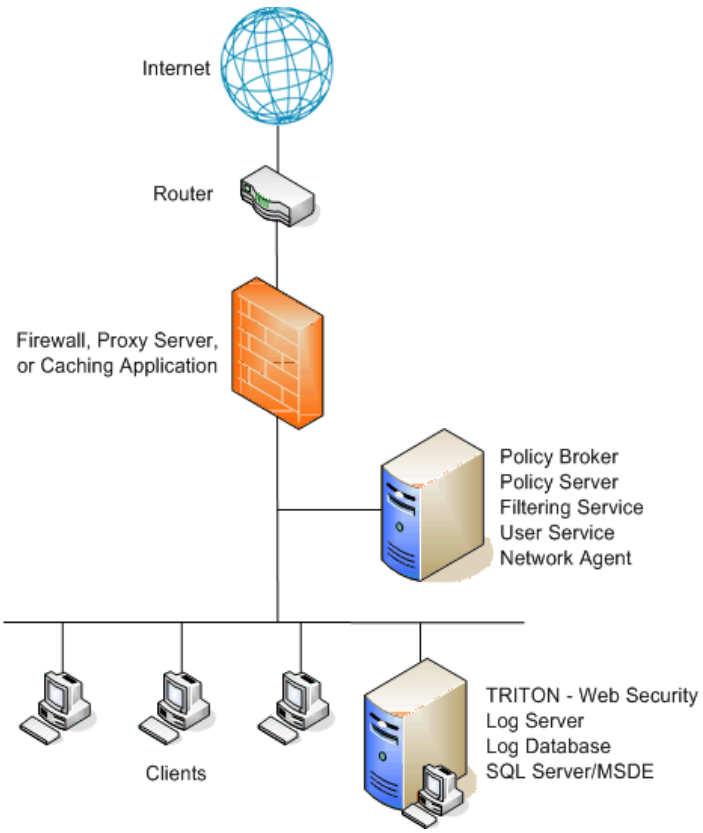

Figure 29 Common network configuration

Other configurations are possible. See your integration product's documentation for other recommendations. See the *Installation Guide Supplement for use with Universal Integrations* for Websense software configuration instructions.

6 Distributed Enterprise Deployments

Distributed enterprise networks have many remote locations, ranging from dozens to thousands of small sites. Typically, between 5 and 50 employees work at each remote site. Many of these sites have Internet access, but no dedicated IT staff.

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

Distributed enterprises with remote Internet connectivity have a complex set of filtering considerations:

- Remote sites 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 site.
- Given the relative low speed of each office's Internet connection, a slightly slower response from the filtering product is acceptable.
- All remote sites can be filtered using the same policies.

Websense Web Security and Websense Web Security Gateway are on-premises solutions in which Websense filtering components can be deployed regionally and communicate over the Internet to apply uniform filtering policies across all offices.

NOTE

The information in this supplement about Websense Web Security generally applies to Websense Web Filter as well. For the purposes of this supplement, *Websense Web Security* should be taken to refer to both solutions collectively, unless otherwise stated.

Websense Web Security Gateway Anywhere is a hybrid solution, allowing a combination of on-premises and in-the-cloud filtering. Additionally, off-site users outside the network of any site (e.g., telecommuters or traveling employees) can be filtered through the hybrid service.

Basic Network Topology

Websense Web Security and Web Security Gateway

To reduce network infrastructure costs, each remote-site 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 site is sent through a local Internet service provider (ISP) to the Internet.

Figure 31 shows the network topology of this type of remote site for Websense Web Security.

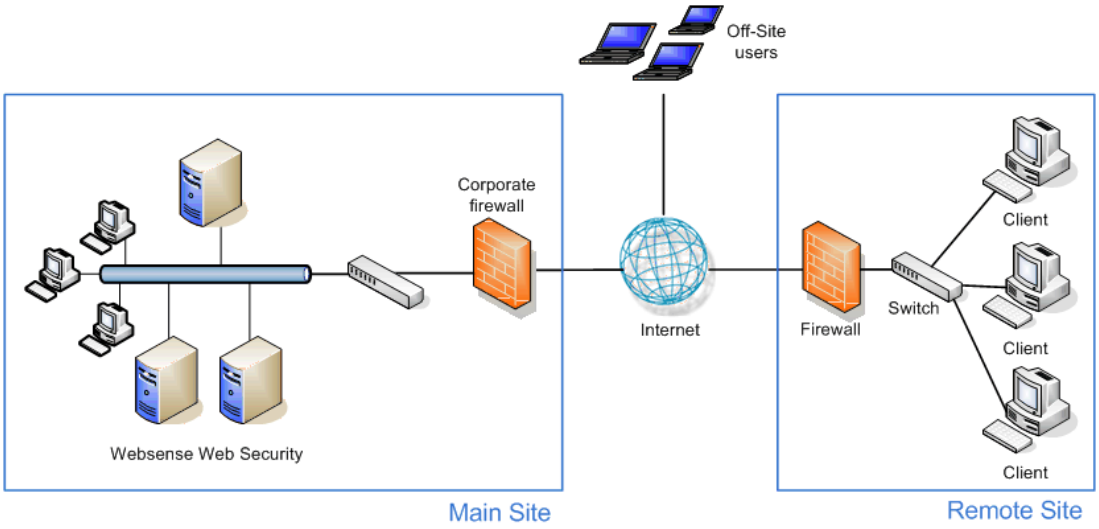

Figure 30 Remote site topology in a decentralized network (Websense Web Security)

Websense Web Security Gateway adds Websense Content Gateway to the deployment, as shown in Figure 31.

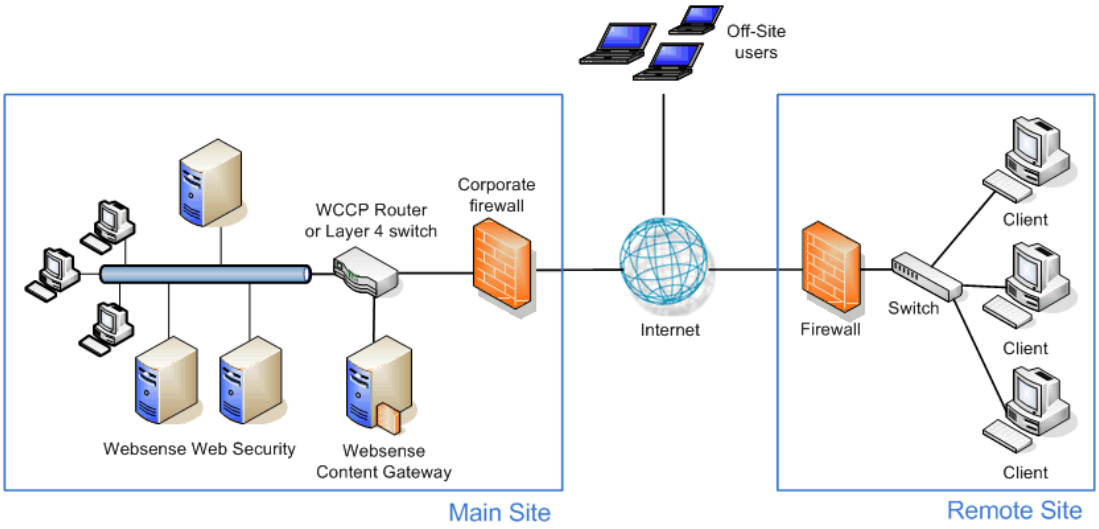

Figure 31 Remote site topology in a decentralized network (Websense Web Security Gateway)

Off-site users (i.e., remote users outside the corporate or remote-site network) can be filtered using Websense Remote Filtering. Note that the Remote Filtering Server must be deployed in the main site network (not depicted in the above illustrations) and Remote Filtering Client must be installed on each off-site machine. See the Websense *Remote Filtering Software* technical paper.

Figure 30 and Figure 31 depict a high-level scheme only. Details about how Websense filtering components might be distributed across separate machines, Websense Content Gateway deployment, Network Agent placement, use of an integration product, and so forth are not addressed here. See the *Deployment Guide* for Websense Web Security Solutions and its *Deploying in an Enterprise Network* supplement.

Websense Web Security Gateway Anywhere

In a basic Websense Web Security Gateway Anywhere deployment, the remote site and off-site users can be filtered through the hybrid service rather than by the filtering software at the main site. See Figure 32.

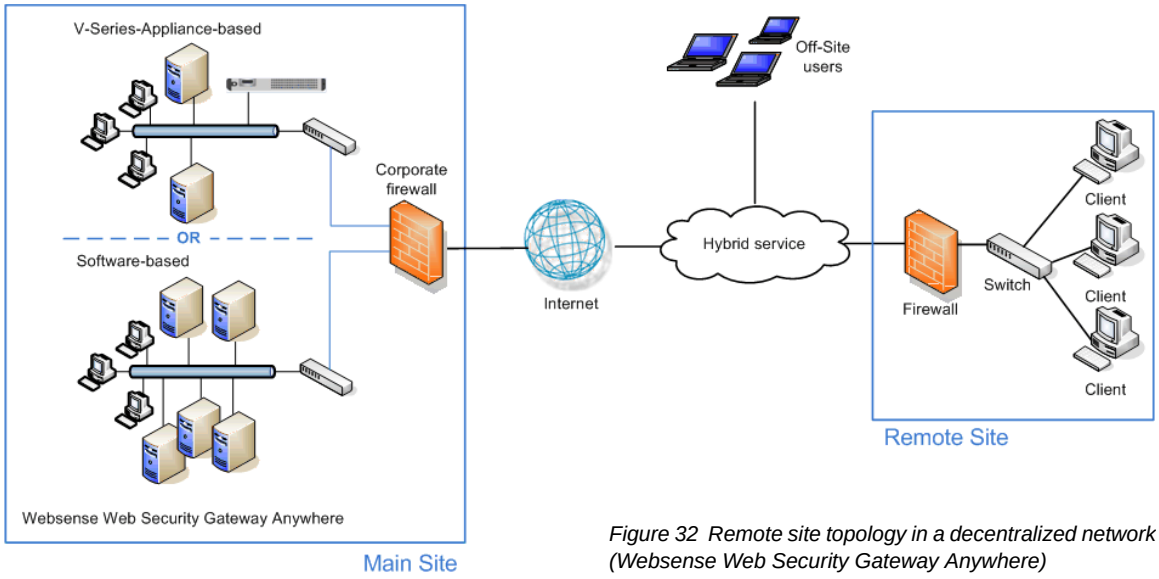

Figure 32 Remote site topology in a decentralized network (Websense Web Security Gateway Anywhere)

A V-Series-Appliance-based or software-based deployment of Websense Web Security Gateway Anywhere is installed at the main site. A V-Series-Appliance-based deployment consists of a Websense V-Series Appliance running core filtering components, plus additional servers running reporting and interoperability components (allowing communication between Web and data security components and also between on-premises components and the hybrid service). A software-based deployment consists of all the same components, distributed across a number of servers. See the *Deployment Guide* for Websense Web Security Solutions and the Websense Web Security Anywhere *Getting Started Guide* for details about the components deployed at the main site.

No additional software is required at remote sites or on off-site machines to be filtered through the hybrid service (some configuration at the main site and deployment of a PAC file to client machines is required; see the TRITON - Web Security Help and the Websense Web Security Anywhere *Getting Started Guide*).

Filtering Remote Sites

Websense Web Security or Web Security Gateway

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

Remote sites 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-site 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.

Filtering is performed by the Websense components at the main site. Remote sites must be equipped with a firewall that can be integrated with Websense software (i.e., configured to check with Websense software to permit or block Web requests) or an instance of Websense Network Agent must be deployed at the remote site. *Firewall* is used here as a generic term to refer to a firewall, gateway, or proxy.

Websense, Inc. 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 site.
- Eliminates the need for a separate Websense software installation at each location.
- Provides uniform filtering policies at each remote site.
- Eliminates the cost of additional hardware to provide filtering servers at each remote site.
- Allows the enterprise to centrally configure, administer, and maintain a limited number of Websense filtering machines.

Figure 33 shows a typical sequence of events in filtering a client machine at a remote site.

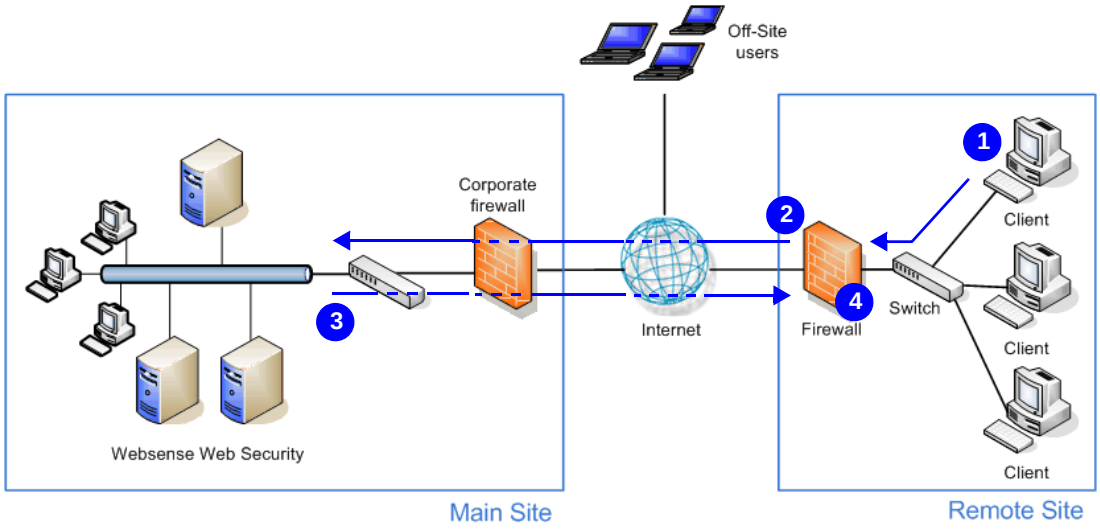

Figure 33 Filtering a remote-site client machine (Websense Web Security and Web Security Gateway)

1. User requests a Web page.
2. Local firewall checks the URL of the requested page with Websense Web Security/Web Security Gateway over the Internet.
3. Websense Web Security/Web Security Gateway responds over the Internet, indicating whether the request should be permitted or blocked.
4. Local firewall permits or blocks the request as directed.

Note that Figure 33 is a simplified diagram showing the main conceptual sequence of events. Details of Websense component distribution and placement in the corporate network, network routing and internal firewall usage, segmentation of networks, and so forth are not addressed here. For Websense component deployment details, see the *Deployment Guide* for Websense Web Security Solutions and its *Deploying in an Enterprise Network* supplement.

In the case of multiple remote sites, each remote site communicates with Websense components at the main site in the same manner shown in Figure 33.

Off-site user machines are filtered by deploying Websense Remote Filtering Server at the main site. Websense Remote Filtering Client is installed on each off-site machine to be filtered. See the Websense *Remote Filtering Software* technical paper for details.

Websense Web Security Gateway Anywhere

In a Websense Web Security Gateway Anywhere deployment, remote sites can be filtered by the hybrid service rather than the Websense software or appliance at the main site.

Network latency issues are addressed by the fact that a remote site and off-site users are filtered by the nearest Websense hybrid service cluster and not the main site. Without the hybrid service, geographically distant remote sites could experience slowing of the Web browsing experience that might be unacceptable to some end users. In such cases, without the hybrid service, additional Websense filtering components would have to be installed at the remote site or geographically close regional site (incurring additional hardware costs for servers to run Websense software locally).

Figure 34 shows how remote-site filtering works in Web Security Gateway Anywhere. Remote site client machines are filtered by the hybrid service directly rather than instructing the local firewall to permit or block a request. A user's request for a Web page is directed to the hybrid service, which permits or blocks the request based on the applicable policy.

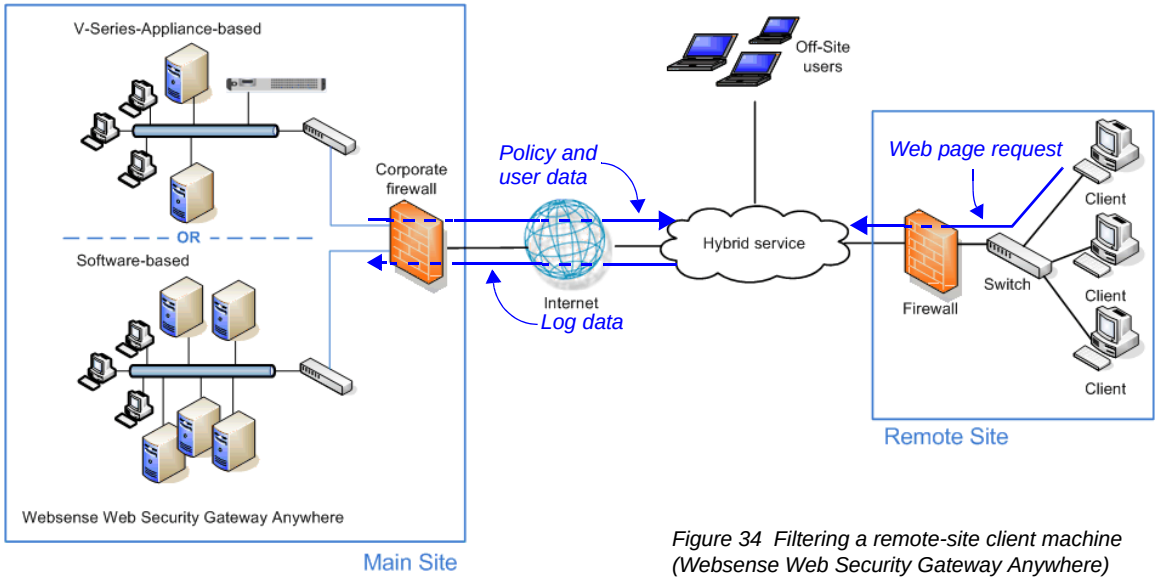

Figure 34 Filtering a remote-site client machine (Websense Web Security Gateway Anywhere)

Policy settings are defined at the main site and uploaded automatically to the hybrid service at preset intervals. User information, for user- or group-based filtering, is also uploaded.

Log data for reporting is downloaded from the hybrid service to the main site automatically and is incorporated into the Websense Log Database (at the main site). Hence, reports can cover users at all offices.

Off-site users are filtered by the hybrid service as well. Alternatively, off-site users can be filtered using Websense Remote Filtering Server (deployed at the main site). In that case, Websense Remote Filtering Client must be installed on each off-site user's machine. See the Websense *Remote Filtering Software* technical paper for details.

Deployment models

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

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

Sites in a region

The simplest Websense deployment for a distributed enterprise is a network with remote sites in a single region, such as San Diego County, California, U.S.A. Most organizations with sites like this can use a single Websense Web Security or Web Security Gateway deployment, centrally located within that region, to provide filtering for all clients. See Figure 35.

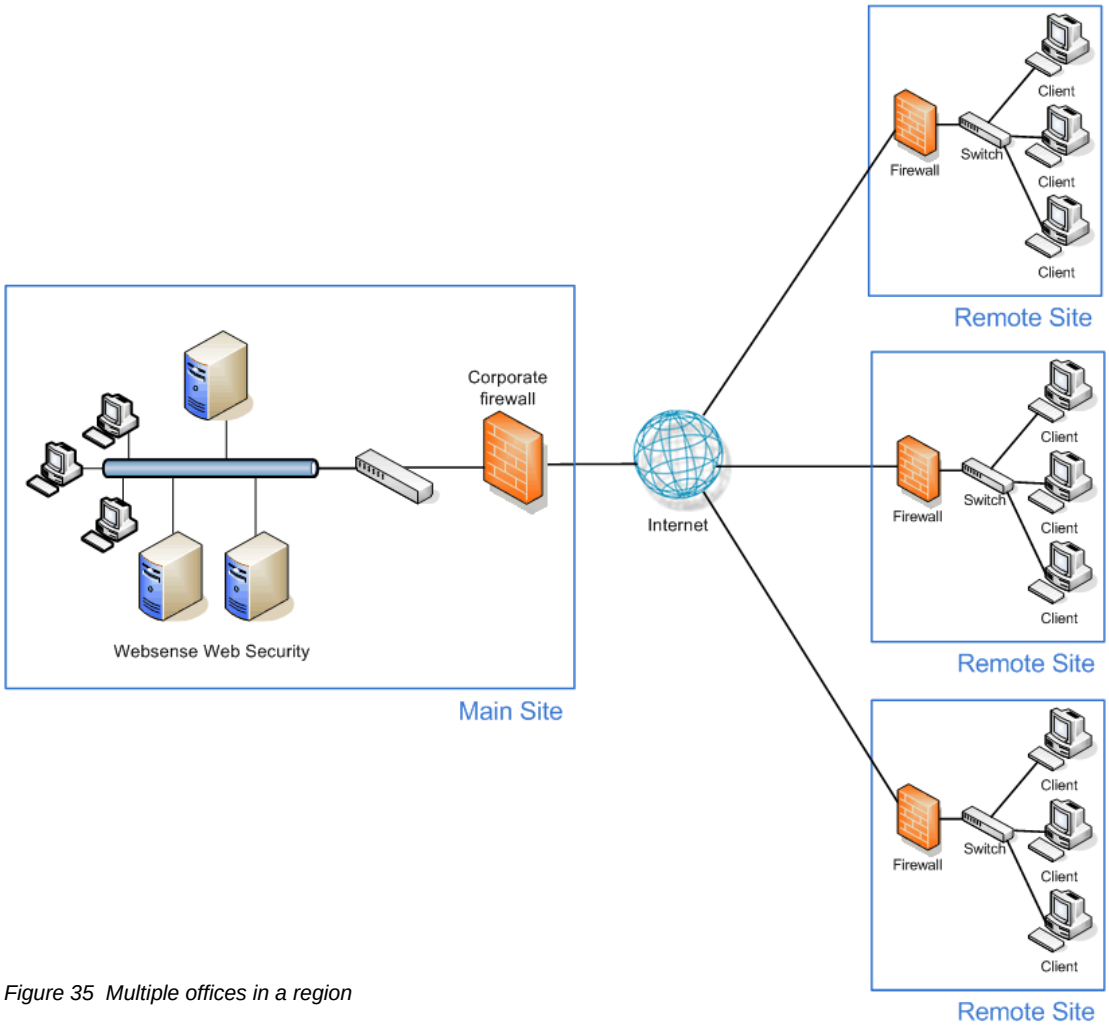

Figure 35 Multiple offices in a region

Each remote site would be filtered as shown in Figure 33, page 92. The site in which Websense software is deployed is represented as the "main site", but need not be truly a main site in your organization. It is whichever one houses Websense software.

Off-site users, not shown in Figure 35, can be filtered using Websense Remote Filtering Server (deployed at the main site). Websense Remote Filtering Client must be installed on each off-site user's machine. See the Websense *Remote Filtering Software* technical paper for details.

Expanding sites in a region

Some organizations deploy Websense Web Security or Web Security Gateway within a given region and later decide to increase the number of remote sites in that area.

To compensate for the additional sites 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 sites 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 site. See Figure 36.

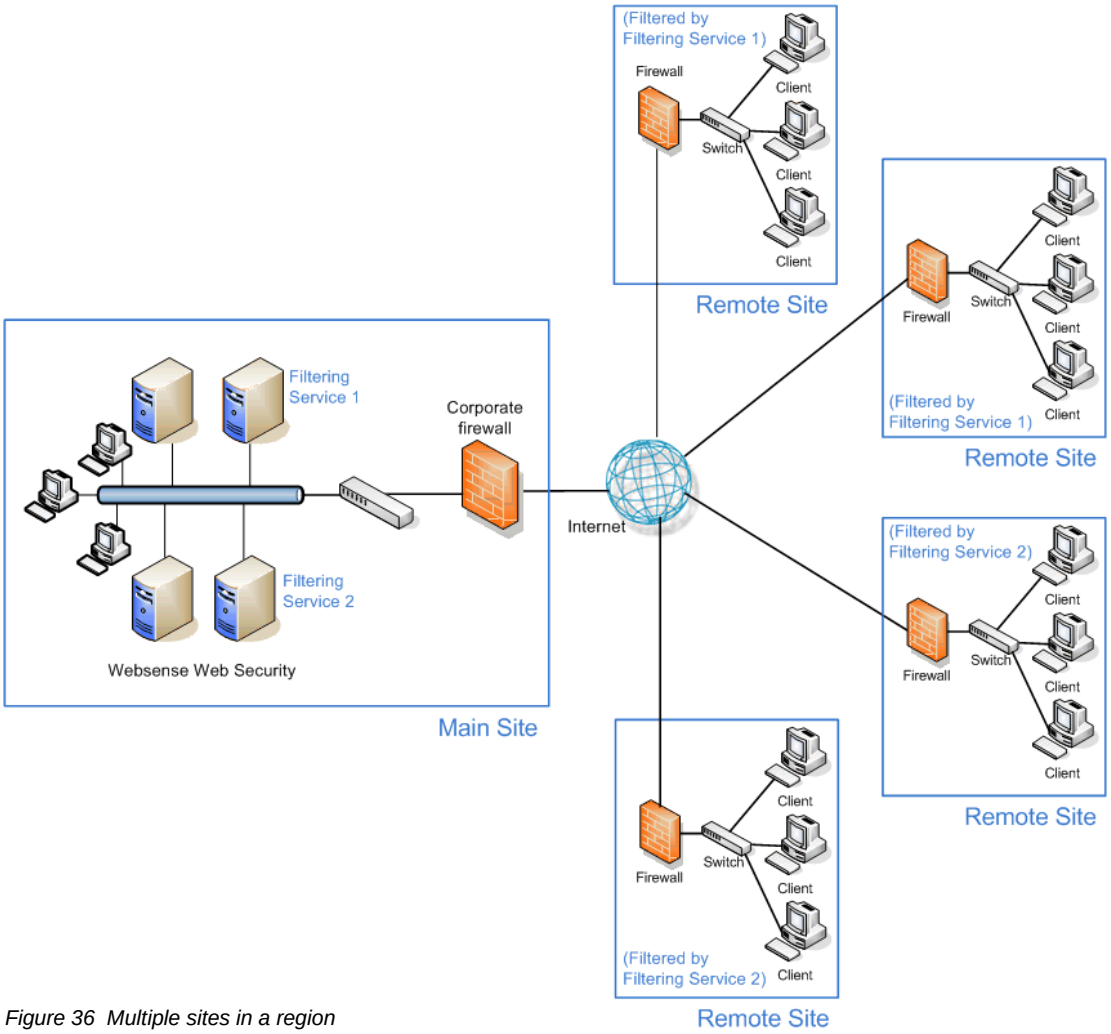

Figure 36 Multiple sites in a region

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

Off-site users, not shown in Figure 36, can be filtered by Websense Remote Filtering Server (deployed at the main site). Websense Remote Filtering Client must be installed on each off-site user's machine. See the Websense *Remote Filtering Software* technical paper for details.

National or worldwide offices

Websense Web Security or Web Security Gateway

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

- Each remote site is geographically distant from the Websense components. Request lookups would 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 sites into logical regions and deploy Websense software in each region. For example, a distributed enterprise might group their United States sites into a western region, a central region, and an eastern region. Websense software is deployed at a central site in each region.

The logical division of sites into regions depends on the location and grouping of remote sites and the total number of employees at each site. For example, a company with a large number of remote sites 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 sites in California with 100 to 250 employees each. In this case, a single Websense software installation might be deployed for all three sites. This enterprise also can deploy Websense software locally at each site (rather than using a distributed approach), particularly if 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.

Websense Web Security Gateway Anywhere

Websense Web Security Gateway Anywhere is particularly well-suited for organizations with sites distributed nationally or worldwide.

Single main site

An organization with one main site (e.g., headquarters office or main campus) and multiple, geographically dispersed remote or branch sites can deploy Websense software at the main site (with main-site users filtered by the on-premises components) and have all remote sites filtered through the hybrid service. See Figure 37.

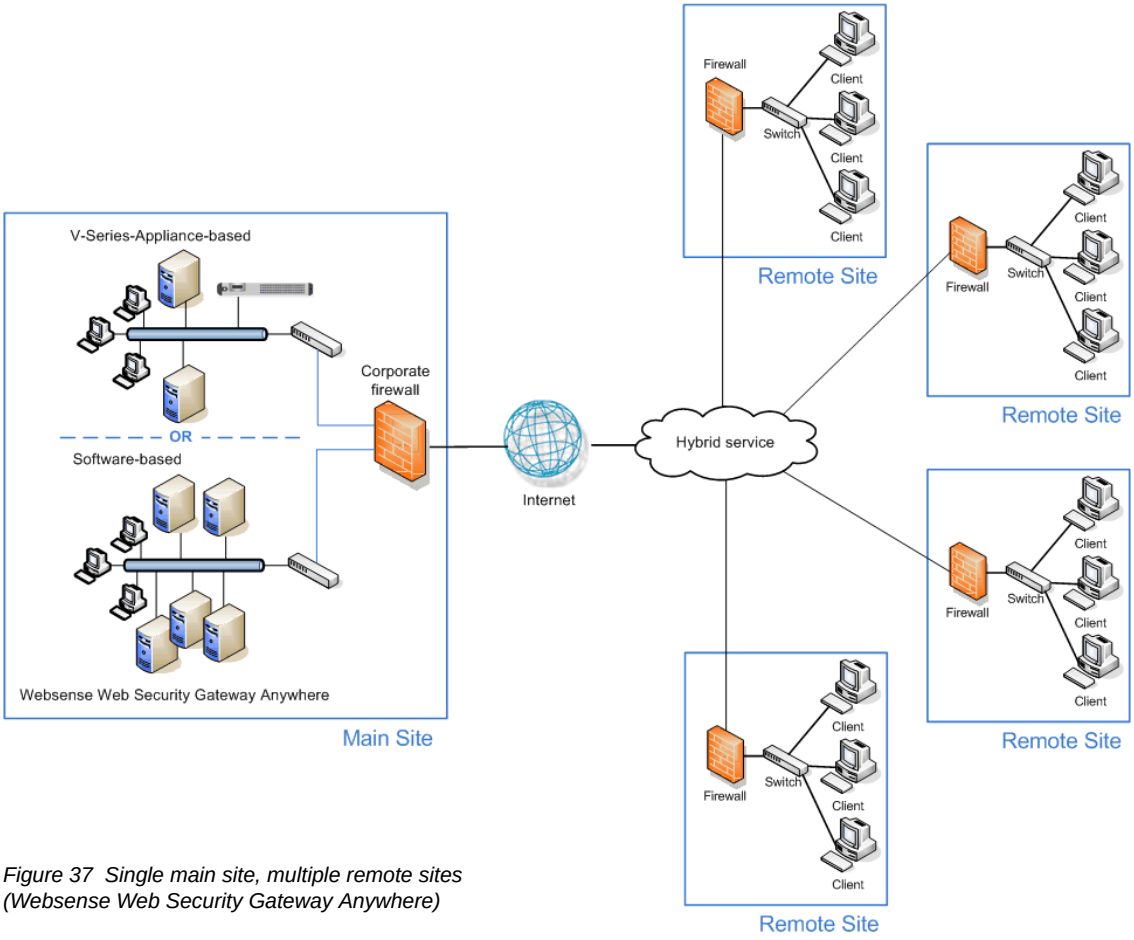

Figure 37 Single main site, multiple remote sites (Websense Web Security Gateway Anywhere)

Off-site users, not shown in Figure 37, are filtered through the hybrid service. Alternatively, they could be filtered by Websense Remote Filtering Server (deployed at the main site). In that case, Websense Remote Filtering Client must be installed on each off-site user's machine. See the Websense *Remote Filtering Software* technical paper for details.

Multiple large sites

Organizations with multiple large sites (e.g. main headquarters and regional headquarters) can deploy on-premises filtering at the larger sites while filtering small, remote sites through the hybrid service. See Figure 38.

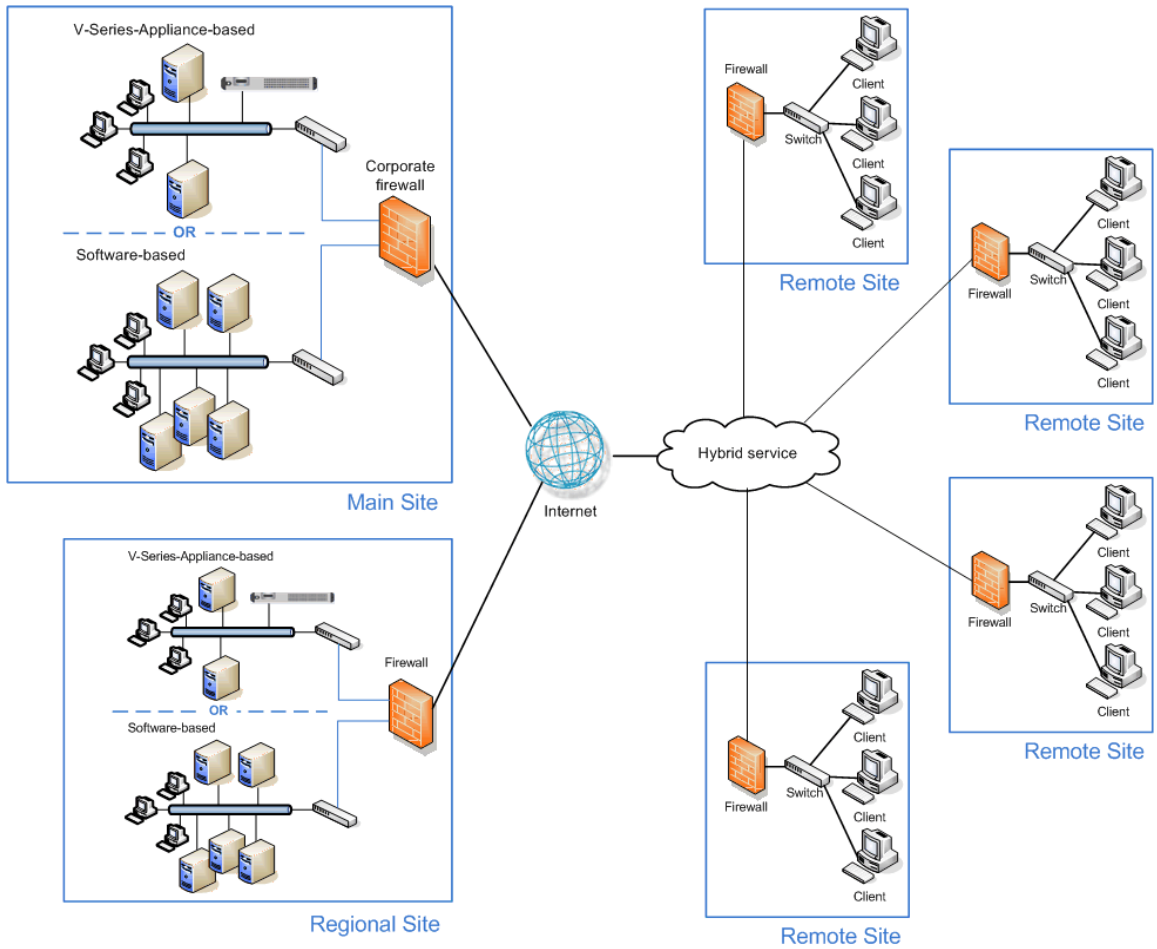

Figure 38 Multiple large sites (Websense Web Security Gateway Anywhere)

When there are multiple on-premises deployments of Websense Web Security Gateway Anywhere components, there are important things to note:

- There must be only one Policy Broker and one Sync Service in the entire deployment (at the main site). See the TRITON - Web Security Help and the *Deployment Guide* for Websense Web Security Solutions
- For unified configuration and policy-application, V-Series Appliances deployed at regional sites should be configured to use the appliance at the main site as the *policy source*. See the *Getting Started Guide* for Websense V-Series Appliance and the Websense Appliance Manager Help.
- All Log Server instances should be configured to send data to the main Log Database at the main site. See the TRITON - Web Security Help and the *Deployment Guide* for Websense Web Security Solutions

Off-site users, not shown in Figure 38, are filtered through the hybrid service. Alternatively, they could be filtered by Websense Remote Filtering Server (deployed at the main site or a regional site). In that case, Websense Remote Filtering Client must be installed on each off-site user's machine. See the Websense *Remote Filtering Software* technical paper for details.

Secure VPN connections

This information applies to Websense Web Security and Web Security Gateway. For URL lookup requests and replies, some firewalls allow administrators to set up a secure VPN connection between the remote-site 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, Websense RADIUS Agent can be used for transparent user identification. For information about deploying RADIUS Agent, see *General Deployment Recommendations* in the *Deployment Guide* for Websense Web Security Solutions. For more information about installing RADIUS Agent, see the Websense Web Security and Websense Web Filter *Installation Guide*.

Calculating TCP connections

The information in this section applies to Websense Web Security and Web Security Gateway. For a Websense Web Security Gateway Anywhere deployment, this section applies only to those cases where a remote site is filtered by on-premises components at a main site instead of by the hybrid service. In such cases, remote-site filtering operates just like it would in a Websense Web Security or Web Security Gateway deployment. The hybrid service is administered and maintained by Websense, Inc. Scaling of the hybrid service to fulfill requests from your remote sites is handled by Websense, Inc.

In a distributed enterprise, Internet requests may be sent to multiple Websense Filtering Service instances from hundreds of remote-site 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 sites' firewalls exceed the maximum allowed number of connections that can be accepted by Websense software, the firewalls either block all subsequent requests or permit all requests, depending upon how those firewalls are configured.

This section provides 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 sites, 1-3 connections should be sufficient for each remote-site firewall, depending upon the load (number of requests per second) from each site. 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 Service instances may be needed to handle the extra connections requested by the remote sites' 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 sites, 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 sites, 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 sites.

The following tables display sizing information for remote sites 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 8. 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 9. 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 10. 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 11. 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 Filtering Service 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 Websense Web Security and Websense Web Filter *Installation Guide* or the TRITON - Web Security 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 for Websense Web Security and Web Security Gateway 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 sites must travel over the Internet to the Websense installation, in most situations end users at remote sites 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 site.
- Distance from the remote site to the machine running Websense filtering.
- Number of users and connections to the machine running Websense filtering.
- Speed of the Websense machine.

In a Websense Web Security Gateway Anywhere deployment, latency due to filtering is minimized when remote sites are filtered by the hybrid service. Each remote site is filtered by the geographically closest hybrid service cluster.

Internet Connection Speed

Overall filtering performance is dependent upon the speed of the Internet connection at each remote site, which is determined when the your organization sets up the site. A DSL, cable, or T1 line is appropriate for a site 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 software responses.

Important

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

Distance from the Websense filtering machine

When remote sites are filtered by on-premises Websense software and not the hybrid service (i.e., Websense Web Security and Web Security Gateway; and Websense Web Security Gateway Anywhere deployments in which remote sites are not filtered through the hybrid service but, instead, are filtered through the on-premises deployment at a main site) the more points (hops) across the Internet a Websense filtering 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 filtering machine, and the time required for each hop, is generally tied to the geographic distance between the machine running Websense software and the source of the request. The closer the server is to a remote site, the faster the Websense filtering lookup and overall performance improves for the end user.

It is a best practice to have the Websense filtering machine no more than 20 hops from each remote site. Similarly, the total trip for an ICMP (Internet Control Message Protocol) ping from each remote site to Websense filtering 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 the following commands.

From a Windows command prompt:

- **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

When remote sites are filtered by on-premises Websense software and not the hybrid service, the number of requests per second coming in from remote sites can be quite high, as can the number of connections being opened to the Websense filtering machine. The 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 site processes the Websense software response in less time and provides faster overall performance.

Index

S

T

U

V

W

X