

Netsweeper Inc.
Corporate Headquarters
104 Dawson Road
Suite 100
Guelph, ON, Canada
N1H 1A7
CANADA
T: +1 (519) 826-5222
F: +1 (519) 826-5228

Netsweeper Inc. India
Apt. No.: 9J, Block 2
Ceebros Shyamala Gardens
136, Arcot Road, Saligramam
Chennai – 600 093
INDIA
T: +91 44 426 426 25
F: +91 44 426 426 35

Netsweeper Inc. Europe
41 Marlowes
Hemel Hempstead
Hertfordshire
HP1 1EP
UNITED KINGDOM
T: +44 (0) 1442 355 160
F: +44 (0) 1442 355 001

Netsweeper Inc.
Australia/New Zealand
13 Bareena Drive
Mt. Eliza, Victoria
3930
AUSTRALIA
T: +61 (0) 3 9787 2284
F: +61 (0) 3 9787 0965

Netsweeper Whitepaper

Deploying Netsweeper
Internet Content Filtering
Solutions

Document Date: 2010

www.netsweeper.com

©1999-2010 Netsweeper Inc.
All rights reserved.

Every effort has been made to ensure the accuracy of this document. However, Netsweeper Inc. makes no warranties with respect to this documentation and disclaims any implied warranties of merchantability and fitness for a particular purpose. Netsweeper Inc. shall not be liable for any error or for incidental or consequential damages in connection with the furnishing, performance, or use of this document or the examples herein. The information in this documentation is subject to change without notice.

Netsweeper and Netsweeper Inc. are trademarks or registered trademarks of Netsweeper Incorporated in Canada and/or in other countries. Other product names mentioned in this document may be trademarks or registered trademarks of their respective companies and are the sole property of their respective manufacturers.

Table of Contents

Deploying Netsweeper Internet Content Filtering Solutions	4
How Netsweeper Works	5
User to Integration Level	6
Integration to Distribution Level	7
Distribution to Categorization Level	7
In Practice.....	8
Considerations for Deploying the Netsweeper Enterprise Filter	9
Enterprise Filter.....	9
Policy Server.....	10
Reporting Server	10
Web Server and Administrator	10
Estimating Server Requirements	11
Enterprise Filter.....	11
Policy Servers	11
Reporting Servers.....	11
Failover and Load Balancing Requirements.....	11
Deployment Examples	13
High Demand Network.....	13
Modest Demand Network	14
Conclusion.....	15
About Netsweeper	15

Deploying Netsweeper Internet Content Filtering Solutions

In a very short period of time, the Internet has firmly established itself as an essential research and communication tool in virtually any business or institution around the world. Every organization and individual that is connected to the Internet is also exposed to the threats the Internet brings to data, productivity, financial safety, and moral sensibilities. By its global reach, the Internet regularly defies laws, policies, and regulations established by governments and lawmakers.

Adopting filtering services available over Internet protocol (IP), businesses, organizations, and users can avoid offensive and often intrusive websites and the spyware, adware, and malware that lurk outside every network Internet connection.

On considering it, no one doubts the case for filtering services over IP in their business or institution to protect themselves from Internet threats. The question is, which of the many filtering (and security) tools will provide the necessary control without requiring complex and/or expensive solutions that can make deployment a nightmare, daily operation an exercise in frustration, and maintenance seem hopeless? According to IDC (International Data Corporation), “a key challenge for IT managers is to maximize their return on investment by seamlessly integrating security solutions into their existing environment.”

Netsweeper, Inc. offers an advanced enterprise-calibre filtering system for services over IP. With a methodology that responds to actual Internet traffic and a simple deployment methodology that scales easily with network expansion, Netsweeper’s filtering solution warrants serious consideration for maximizing any organization’s return on its IT security investment.

This paper describes the typical Netsweeper Enterprise Filter deployment and operation.

How Netsweeper Works

Netsweeper's unique architecture provides effective, flexible services-over-IP filtering through a series of Internet-connected servers that access one of the largest URL databases of any IP filtering provider. Netsweeper houses most of the filtering technology in secure and redundant locations, so an organization needs only to set up a Netsweeper Policy Server and an Enterprise Filter to handle its unique network use and traffic flow requirements. *(Netsweeper clients can now also opt to use Netsweeper's hosted environment or a Netsweeper NS PROx Web Filter Appliance to accomplish the same.)*

The Internet is a constantly changing matrix of web sites and services. Netsweeper was designed to respond immediately to surfing patterns and new sites. By design, the most commonly requested sites are already categorized and available in a cache as near the user as possible.

If so little is required to successfully deploy a Netsweeper filtering solution in an enterprise network environment, how is it that Netsweeper actually accomplishes such responsive, comprehensive filtering? Figure 1: *URL Flow through Netsweeper Architecture* shows a simplified version of what happens when an outgoing URL request is made through a Netsweeper Enterprise Filter solution.

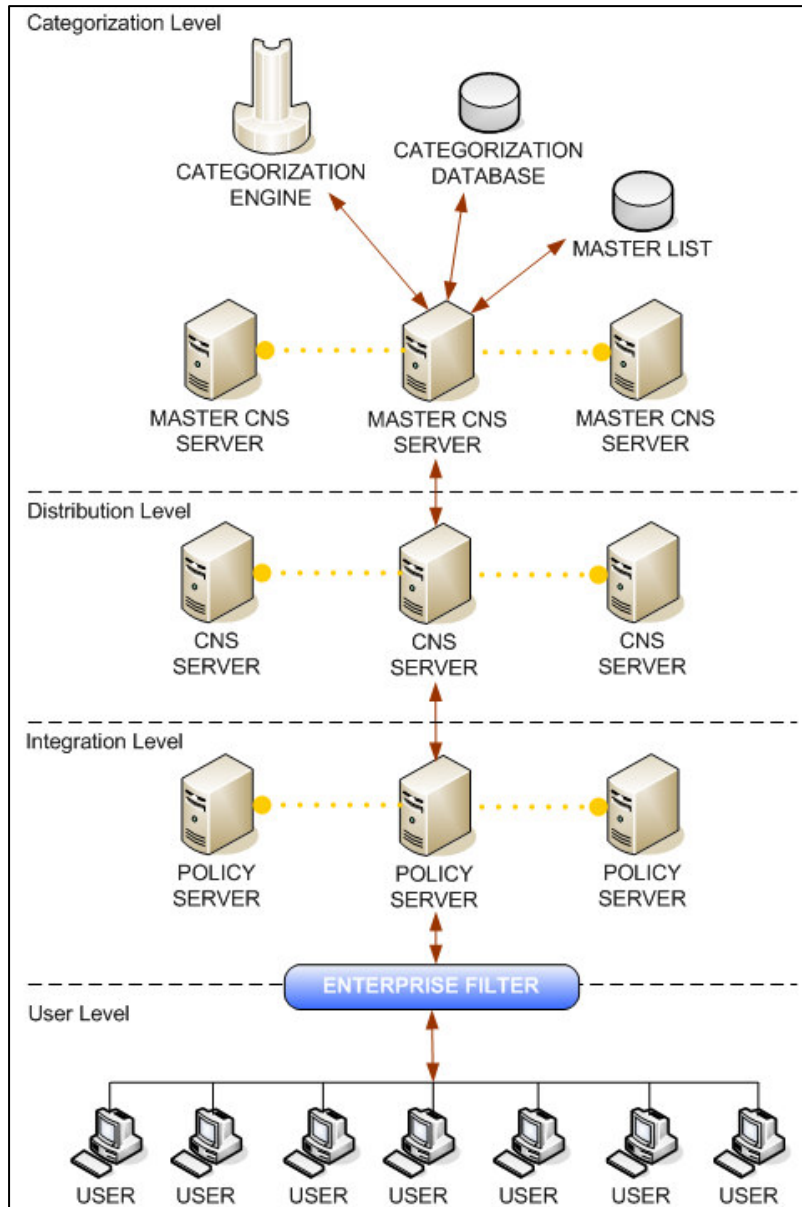


Figure 1: URL Flow through Netsweeper Architecture

User to Integration Level

When a user makes an outgoing request to the Internet, the Netsweeper Enterprise Filter intercepts the request and asks the Policy Server for a ruling - whether to allow or deny the connection. The Policy Server must first categorize the outgoing request: Is it a protocol request or an HTTP request? For non-HTTP requests (such as messaging or file sharing), the Policy Server is always able to make the categorization itself. If it is an HTTP request, the Policy Server checks its own cache for the URL. If the URL is there, the Policy Server categorizes the request.

Once categorized, to process the outgoing request and to respond to the Enterprise Filter, the Policy Server looks up the group policy associated with the user who made the outgoing

request. Policies can be defined as blanket policies covering all users, groups of users, or an individual. (It is also possible to define different policies for different times of the day.) If the specific policy allows the outgoing request, the Enterprise Filter is told to process the request. If the specific policy does not allow the category of the outgoing request, the Enterprise Filter is instructed to return a deny page to the user.

Integration to Distribution Level

If the Policy Server cannot locally categorize an HTTP request, it sends the URL to the Netsweeper Category Name Server (CNS) asking for a category ruling. Like the Policy Server, the Category Name Server maintains a local cache of recently requested URLs and first looks here to assign a category to the URL. If the URL is in its cache, the Category Name Server returns the category for the URL to the Policy Server. If the Category Name Server does not have the requested URL's category in its cache, the Category Name Server requests a category ruling for the URL from the Netsweeper Master Category Name Server (Master CNS) and allows the request from the Policy Server to time out (default setting of time out is one second).

Normally, the Enterprise Filter and Policy Server are located within the client's network. The Category Name Server is hosted on the Internet by Netsweeper. In certain circumstances, a Category Name Server can be dedicated to a particular client or group of clients and may contain its own local URL list – for example, static allow/deny lists. On the request time out, the Policy Server proceeds to process the initial request from the Enterprise Filter using "New URL" as the category. Now having a category for the URL, the Policy Server looks up the ruling and responds to the Enterprise Filter to allow or deny. The Policy Server stores the URL in its cache with the category of No Category.

Distribution to Categorization Level

Continuing upstream, if the Master Category Name Server does not have the URL in its own cache, it allows the Category Name Server request to time out, which results in New URL being stored in the Category Name Server cache. The Master Category Name Server then requests a category ruling for the URL from the Categorization Database. If the URL is not in the Categorization Database, the Categorization Service sends the URL to the Categorization Engine for categorization and sets the category for the URL in its own cache to New URL.

The Categorization Engine is made up of a number of daemons/servers running over 800 processes; each processing URL categorization requests. Through this dedicated categorization process, the Categorization Engine reviews the Web page content from a request, and within milliseconds, assigns a category to it.

When the Categorization Engine receives a request, it retrieves the URL, parses the data, reports any found links to the Master Category Name Server for their own category ruling, and proceeds to determine a category for the original URL request. Once it determines a category for the URL, it passes the data to the Master Category Name Server which updates the Categorization Database.

The Categorization Database is made up of several SQL database servers that balance the URL request load.

In Practice...

New URL is one of several special system categories. The administrator can set the filtering policy to allow or deny URLs with the New URL category (or other system categories) to tailor the overall response. For New URL categorizations, the servers (Policy, Categorization Name, Master Categorization Name) know to request a refresh the category for the URL (since the Categorization Engine will have properly categorized the URL at this point and updated the Categorization Database).

The entire Netsweeper categorization process—from initial outgoing Internet request for a URL never seen by the system before (worldwide) to Categorization Engine categorization and storage in the database—takes as little as one second and at most about five seconds, depending on the global location of the network user and the speed of connection to the requested URL web server.

Users and administrators are able to request a human review of URLs—either to add a URL to a category, remove a URL from a category, or add a URL to multiple categories. All sites reviewed manually are immediately updated in the Categorization Database and are available to the Master Category Name Server. These sites/updates are also downloaded nightly to the Category Name Server and Policy Server caches.

Considerations for Deploying the Netsweeper Enterprise Filter

The Netsweeper Enterprise Filter solution consists of several components, most of which can be run concurrently on the same server hardware or, as scaling requires, separately on independent/load balanced server hardware.

The two major components are the Enterprise Filter (which intercepts outbound Internet traffic and ultimately allows or denies that traffic) and the Policy Server (which makes the categorization decision and, based on the categorization decision, makes the allow or deny decision). Other components are the Reporter Server and the Web Server and Administrator.

Enterprise Filter

Deploying the Netsweeper Enterprise Filter can be done in three different ways:

1. Default Gateway Router (inline solution) – Following this deployment method, the Enterprise Filter will monitor and filter traffic as it travels from one sub-net to another within a local network.
2. Transparent Network Bridge (inline solution) – Installing the Enterprise Filter using this method will require all workstations on a network to have their default gateway configured to send all traffic to the Netsweeper Enterprise Filter software. Policy decisions will be made for each request and if allowed, forward the request on to its default gateway.
3. Pass-by filtering (not an inline solution) – Using a switch to that is capable of copying and forwarding packets (also known as an IDS or Port Mirroring switch), packets will be copied and sent to the Enterprise Filter simultaneously for identification. Should the Policy Server determine that the request is to be blocked, the Enterprise Filter will inform the switch to cancel the request and serve up a deny screen.

Regardless of the deployment method deployed, the following types of outgoing Internet requests are recognized and processed:

- HTTP
- FTP
- Text messaging (also known as instant messaging, or IM)
- Peer-to-peer file sharing (P2P)
- Mail
- Other UDP and TCP based protocols.

After intercepting outgoing requests, the Enterprise Filter sends them to a Netsweeper Policy Server. Based on the reply from the Policy Server, the Enterprise Filter then blocks the request or forwards it to the Internet.

The Enterprise Filter is an OSI model-based, Layer 7 protocol analyzer that can handle 30 Mbps of Internet traffic per hardware server. It does not need inbound packets to be returned the same way they were sent, making it an ideal solution for asymmetric routing environments: the Enterprise Filter checks outgoing requests only. This also introduces bandwidth savings as the request is never sent to the remote web server if the content is deemed inappropriate.

Policy Server

The Netsweeper Policy Server is the core Netsweeper component. It receives requests regarding outgoing Internet requests from the Netsweeper Enterprise Filter, categorizes the request, maps the requests to a policy, and determines whether the request should be allowed or blocked.

If the Policy Server is unable to make a categorization decision locally (using its own cache and rules), it communicates with upstream Netsweeper devices to assign a category for the requested URL.

The Policy Server is not in-line with the Internet traffic. It can be hosted locally, within the enterprise or remotely at a central location that is accessible.

It is the Policy Server that records the request result in the report log, not the Enterprise Filter.

In its smallest deployment, the Netsweeper Policy Server is a single hardware server that is running the web server for the administrative functions and the Policy and the Reporter services. In an ultra-small deployment, the Enterprise Filter can also be run on the same hardware server as the Policy Server.

In its largest deployment, the Netsweeper Policy Server consists of multiple policy servers, a separate web server, and a separate reporter server, plus load balancing appliances.

Reporting Server

The Reporting Server receives and stores log files that are transferred from the Policy Server in real time as outgoing requests are being processed. Through a web interface on the Policy Server, network administrators can use the log files on the Reporting Server as a source for generating Internet activity reports for all network clients and for each network workstation.

The Reporting Server can export reports to standard programs, including Crystal Reports and Microsoft Excel.

Web Server and Administrator

The Policy Server is controlled and administered through a web interface. The web server and system administrator allows complete remote administration of the filtering, reporting, and configuration.

Estimating Server Requirements

To define a custom Netsweeper deployment strategy, the following network variables can help determine the estimated server requirements for an organization's unique network needs:

- For Netsweeper Enterprise Filters, the average number of Mbps of network traffic.
- For Netsweeper Policy Servers, the average number of concurrent network connections.
- For Reporting Servers:
 - The total number of connected networks
 - The length of time for storing logs and reports.

Enterprise Filter

The number of filters required for a Netsweeper deployment is directly related to the average number of Mbps of network traffic. In general, the following formula determines how many filters are required:

- 30 Mbps of traffic = 1 Enterprise Filter and/or
- 100,000 of concurrent TCP/UDP connections = 1 Enterprise Filter

Note: Some ISPs may choose to use a transparent or explicit proxy server with a Netsweeper Policy Server instead of opting for a Netsweeper Enterprise Filter. Although these proxy servers can cache requested URLs and DNS queries, they generally can only handle 15 Mbps of Internet traffic and do not offer filtering for text messaging (IM), peer-to-peer file sharing (P2P), Mail, and other UDP and TCP based protocols.

Policy Servers

The number of Policy Servers required for a Netsweeper deployment is directly related to the average number of concurrent connections that a network needs to support. In general, the following formula determines how many Policy Servers are required:

- 8,000 concurrent connections = 1 Policy Server

If necessary, organizations can split the Policy Server functions into subcomponents over multiple servers to accommodate Internet traffic load balancing and system failover.

Reporting Servers

For Reporting Server storage requirements, consider:

- The total number of connected networks to determine the effects on processing power
- The length of time that you want to archive logs and reports to determine hard disk space (100GB minimum is recommended)

In general, having a separate server for reporting can save processing power for the Netsweeper Enterprise Filters and Policy Servers. However, on a simple network, the Reporting Server can be located on a Policy Server.

Failover and Load Balancing Requirements

An organization's service level agreement may dictate further environment modifications to allow for failover and load balancing. To comply, the Netsweeper deployment can include

multiple Policy Servers, Enterprise Filters, Reporting Servers with RAID disk arrays, and load-balancing OSI Layer 4/7 devices.

Note: Some models of OSI Layer 4/7 switch do not support both failover and load balancing. If both are required, the device performing the load balancing and failover may need to be upgraded to comply with these requirements.

Deployment Examples

The following examples represent only two of the many possibilities of Netsweeper Enterprise Filter deployment strategies that address the unique needs of two sample network environments.

High Demand Network

In a typical, high demand network Netsweeper deployment, multiple Policy Servers and Enterprise Filters are installed to accommodate a high volume of concurrent connections and outgoing Internet traffic, and to provide failover support. Inbound traffic does not travel through the Enterprise Filter.

The OSI Layer 4 switch manages load balancing by routing or forwarding URL requests to available Policy Servers and Enterprise Filters. In addition, a standalone Reporting Server is set up to provide maximum processing power for request reviews and filtering on the Policy Servers and Enterprise Filters. The administrator web server is generally put on one of the Policy Servers.

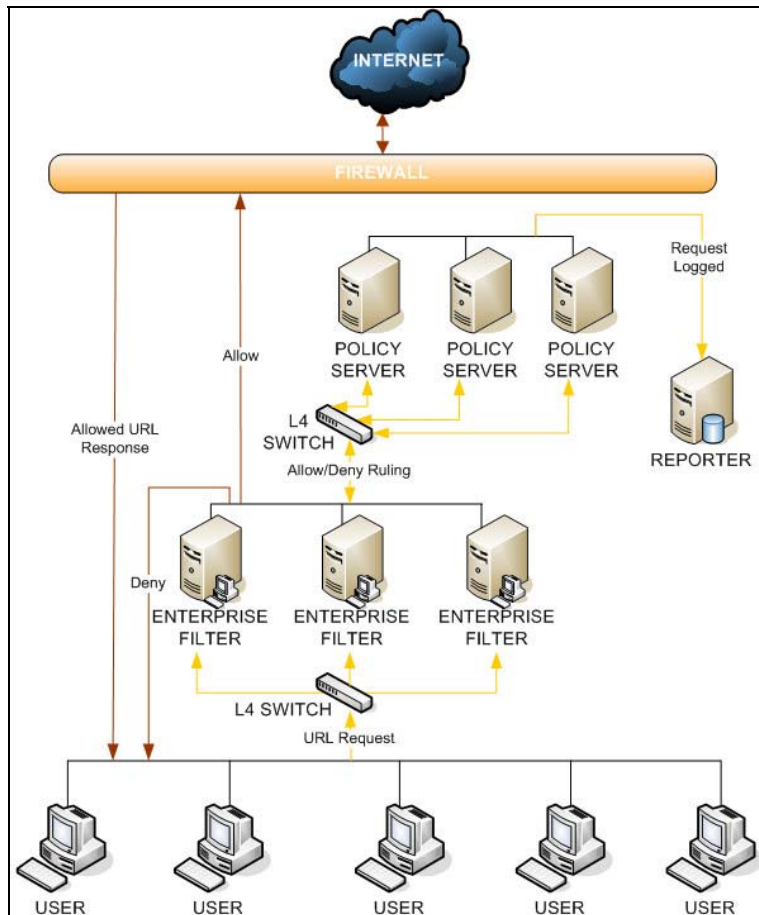


Figure 2: Large, high demand network deployment

Modest Demand Network

In a modest demand network Netsweeper deployment, with a low volume of concurrent connections and outgoing Internet traffic, it's possible to have the Policy Server (and all of its components) and the Enterprise Filter all located on one hardware server. If no failover or load balancing support is needed, a OSI Layer 4/7 switch is not needed.

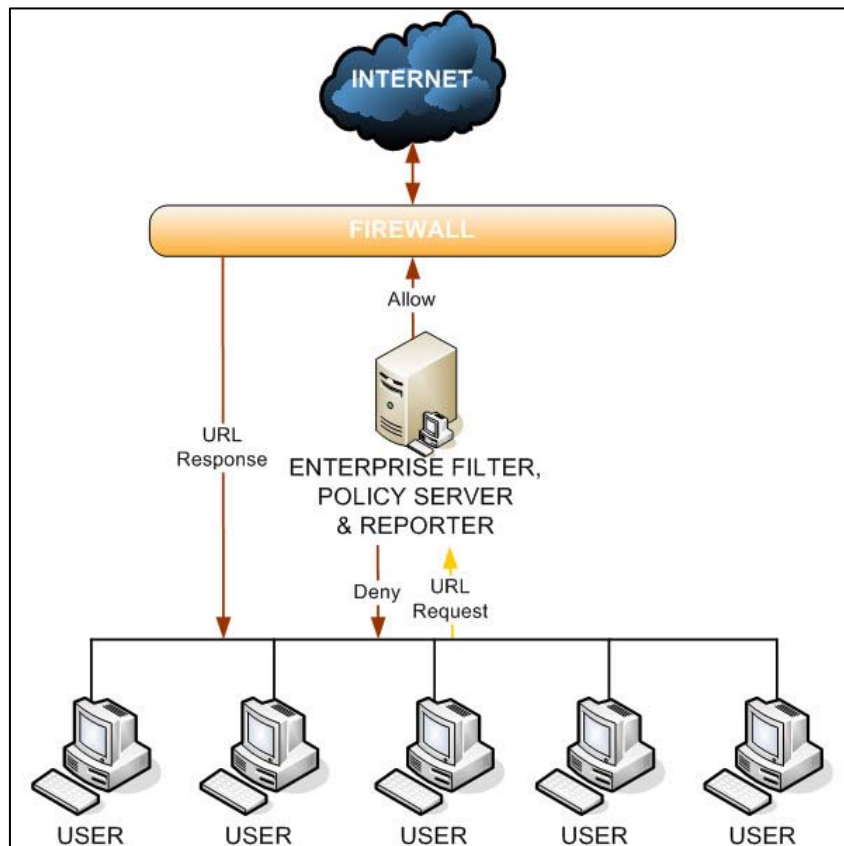


Figure 3: Small, modest demand network deployment

Conclusion

There's no doubt that services-over-IP filtering has become essential in an Internet-connected world. With every network connected through the Internet, it's a two-way street with abundant access to information, communication, and products and services offset by a vulnerability to performance loss, network complexity, and ethical, and even criminal intrusion. The best way for an organization to realise the benefits of the Internet, and maximize productivity and network management is to deploy an effective, tailor-made IP-services filtering system.

Netsweeper offers maximum filtering along with scalability, robust functionality, and best of all, a simple deployment that conforms to each organization's unique IT infrastructure. From a single server that houses the complete filtering, caching, and reporting solution to multiple servers that manage, filter, balance, and report on high volumes of outgoing Internet requests, Netsweeper provides the flexibility to meet any organization's IP-services filtering needs.

About Netsweeper

Netsweeper, Inc. specialises in content filtering software solutions and holds possibly the industry's most advanced proprietary global filtering system for corporations, Internet service providers, educational institutions and government organizations.

Netsweeper's content filtering products operate on a model that categorizes new sites on demand, makes that categorization available to all Netsweeper users worldwide, stores the categorization for fast retrieval and periodic reclassification, and effectively uses local caches to reflect the nature of the local Internet users. With over 1 billion pages currently logged and constant updates occurring daily, Netsweeper's filtering matrix system evolves to offer the organizations and individuals that deploy its software the most protected and secure Internet experience available on the market.

Netsweeper's flexible and customizable technology enables deployment on a wide variety of networks. Netsweeper clients are located on every continent and in every industry vertical.

The company is headquartered in Guelph, Ontario, Canada with offices in India and the UK and distribution channels situated around the world.