



Global Information Assurance Certification Paper

Copyright SANS Institute
Author Retains Full Rights

This paper is taken from the GIAC directory of certified professionals. Reposting is not permitted without express written permission.

Interested in learning more?

Check out the list of upcoming events offering
"Security Essentials: Network, Endpoint, and Cloud (Security 401)"
at <http://www.giac.org/registration/gsec>

Internet Filtering and Websense

David Mangini

April 1, 2002

Abstract

This practical will focus on the management, legal and human resource issues of Internet access provided by businesses for their employees. Included will be a discussion of the importance of an Internet Access Policy (IAP) and Internet access filtering along with a review of the methods available to accomplish Internet filtering. The focus will then change to the implementation of the Websense filtering product on a Windows 2000 platform in conjunction with a Cisco PIX firewall. The conclusion will include a troubleshooting section and how to avoid potential operational difficulties to insure a more successful implementation.

Internet Filtering

When you first address the issue of whether or not you need an Internet filter the answer may seem obvious. A corporation that is providing Internet access to its employees may view this as just another tool, a resource intended to increase productivity by allowing access to almost unlimited research possibilities. After all, businesses have been providing access to external research material for years through traditional sources such as newspapers and trade magazine subscriptions. It's not the access to business related topics that is the issue with Internet access. It's the ability to surf to millions of Internet sites which are totally unrelated to any business function. If you think that employees will not be tempted to check last night's scores or do a little Internet shopping while waiting for that monthly report to arrive, you are wrong. Actually, your organization may not care if the Internet is used for personal business or recreation. How many businesses look the other way concerning employees telephone usage? Large corporations can certainly track phone usage statistics while smaller businesses may not care, as long as it is not interfering with the daily workload. While personal telephone usage may be officially discouraged, occasional use may be acceptable at some organizations.

Internet access, like telephone usage, may also enjoy the same leniency. But is this warranted? If you believe that Internet access provided by employers is not being abused, here are some statistics that may change your mind. [Websense.com]

- Nearly one-third of employees' Internet use at work is recreational.
USAToday, November 1999
- U.S. Department of Labor estimates that wasted time on the Internet costs corporations up to \$3 million a year per every 1,000 employees

Industry Standard, January 2000

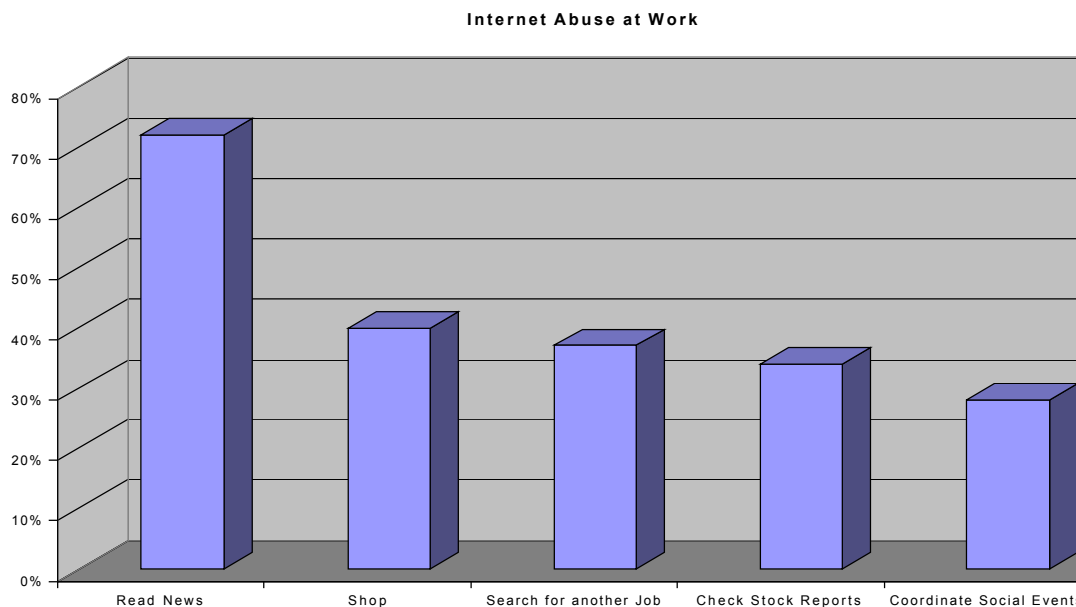
- Much to the chagrin of many managers and supervisors, people are spending more time surfing the Internet at work than they are at home, mainly because home Web connection speeds pale in comparison to the faster connections that companies give their employees.

ZDNet Interactive Investor, 02/18/00

- For the week ending April 2, the top site of the week for female surfers at-work was LTDCCommodities.com, an online shopping catalog and the top site of the week for male surfers at-work was TheOnion.com, a satirical weekly newspaper.
Nielsen/NetRatings, April 11, 2000

- Wasting time online accounts for 30% to 40% of all lost productivity, according to International Data Corporation.

Some reports have stated that as many as 90% of employees have used the Internet for personal reasons while on the job. The following chart gives an indication of what employees are doing when engaged in personal surfing.



[Trudeau and Wynn]

In addition to these activities you can also include Internet chat sessions, e-gambling,

pornography, online banking, music downloads, checking sports scores.... You can be sure, if it's not already on the Internet, it soon will be and your employees will have access to it.

Are you still not convinced that the unfiltered Internet access can be detrimental to your business? There is one way to put this into more concrete terms, impact on the bottom line. If you need help translating lost productivity by employees with unfiltered Internet access into dollars, please visit this site.

<http://www.group1iam.com/cbcalc.html>

This is a cost/benefit calculator that will compute your potential cost of unfiltered Internet employee access and relate it to an investment in an Internet filter. While some organizations may not be willing to expend financial resources for vague Human Resources and Legal concerns, lost productivity when converted into dollars will get some attention.

But let's not gloss over the human resource and legal issues too quickly. These are the issues that are the most complex and the most difficult problems to solve. You started out by providing your employees with Internet access for them to perform their job function most effectively. You now realize that some of them, I mean most of them, are using the Internet for personal reasons. And like the telephone usage, you are willing to look the other way for the most part and accept the financial impact of lost productivity. How can someone surfing at work effect the organization that provided the access?

There are legal issues that you may not have considered. Internet copyright violations are no different from the classic definition just because they are done via a download. Almost everyone knows about the Napster music "sharing" website and software. It allows digital music files to be transmitted over the Internet, a very popular activity at colleges. So popular in fact that Indiana State University's network began to overload with Napster traffic. But that wasn't the worst of it. Indiana State University was named in a copyright infringement lawsuit by two record labels and a heavy metal rock band for allowing its students to participate in Napster transfers and providing the infrastructure to accomplish it. All of this has been resolved, but situations like are better off avoided. [Shuchman]

These violations include not only artistic property but literary works as well. Once again, although your company did not participate in the download which caused the copyright violation, you may find that you are defending your company against a contributory infringement claim. This is what happened to NetCom when it was sued by the copyright holders of works of L. Ron Hubbard, late founder of Church of Scientology; Religious Technology Center v. Netcom On-Line Communications Services, Inc., 907 F. Supp.

1361 (N.D. Calif.) [Diotalevi]

Most employees may not even consider the legal implications of their action. How many companies provide their employees with training concerning the proper use of copyrighted material? If your company is actively providing this type of training, you are to be congratulated. Everyone else may want to consider, at the very least, adding a section to your Internet usage policy.

As serious as the last two examples of copyright infringement may be, as well as the impact that they could have on your company, artistic and literary copyright infringements are dwarfed by software piracy. A survey commissioned by the Software & Information Industry Association (SIIA) and the Business Software Alliance (BSA), found that losses from pirated software topped 59 billion dollars in the 5 years leading up to 1999.

[Beruk] Curious as to whether your company could be liable for this type of copyright violation? The following quote is from the FAQ section of the Software & Information Industry Association website and is in reference to illegally installed software.

“Under "vicarious liability" of the US Copyright Act, an employer is liable for acts committed by its employees when those acts are within the scope of their employment duties. Another theory of liability is the doctrine of contributory copyright infringement, whereby a party who does not do an infringing act but who aids or encourages it is liable for the infringement.” (SIIA)

Your Human Resource department has a vested interest in the way you manage Internet access. Once again you are providing access to the Internet to your employees to do research. Vault.com states that you can expect 1 out of every 25 employees to visit a pornography site while at work. With the loss of productivity aside, you may think that your organization hasn't been affected by this event. With the introduction of Title VII, this situation can very quickly become a Human Resources issue. Should someone else happen to view these images, whether accidentally or intentionally shown to them, your company can be charged with sexual harassment.

You may not understand the connection immediately. However, the definition of sexual harassment includes the existence of a hostile work environment. A hostile work environment can include things beyond the obvious sexual contact to include the distribution of explicit jokes, personal comments directed towards an individual and displaying of sexual objects or photos within the work environment. [Grossman p.3]

Two 1998 supreme court decisions in the Burlington Industries, Inc. v. Ellerth and Faragher v. City of Boca Raton expanded the rights of individuals to sue employers for sexual harassment. The employee no longer had to prove that they experienced any adverse action from the event, they were not fired, demoted, etc. to sue their employer for

sexual harassment. The fact that they were subject to a hostile work environment by being exposed to sexually explicit material was enough to initiate the charge of sexual harassment. [Towns p. 3]

Since some surveys have estimated that 5% of employees will visit pornography sites while at work, isn't it reasonable that in company of 20 employees or more that someone is doing something that is putting your company at risk? Some international corporations and government organizations have taken definite action to insure that they are not allowing a hostile environment to exist.

- Compaq Computer investigates the distribution of pornography via its email system and terminates 20 employees found to be involved.
- The New York brokerage firm, Morgan Stanley was sued for 30 million dollars and six employees were disciplined for transmitting racist jokes over the company's network
- Dow Chemical Company terminated about four dozen employees and disciplined hundreds more because of inappropriate use of company email.
- New Jersey Transit terminated employees for computer misuse.
- The U.S. Department of Defense terminated more than 100 workers for circulating sexual content over its computer system.
[Trudeau and Wynn]

As an employer and provider of Internet access to your employees, how do you defend your company against potential sexual harassment lawsuits and the financial consequences? There are two basic ways to defend yourself. Did the employee take advantage of the policies and procedures that your company has in place to prevent such events? This one can be very difficult to prove and should not be your only avenue of defense. The second concerns what measures did your institute to prevent this from happening in the first place? The standard here is reasonable care. [Towns p. 3]

Your reasonable care defense must have a firm foundation from which to build. This is why your employee Internet Access Policy (IAP) is so important. Employees must have a clear understanding of what is and isn't allowed when using company provided Internet access. You may want to preface the policy with an introduction such as the following from Surfcontrol.

“Use of the Internet by Company employees is permitted and encouraged where such use is suitable for business purposes and supports the goals and objectives of the Company and its business units. The Internet is to be used in a manner that is

consistent with the Company's standards of business conduct and as part of the normal execution of an employee's job responsibilities.” [Surfcontrol]

Your IAP should also contain the following points

- State to your users that their Internet access is to be used for business purposes only and that personal use is prohibited.
- Make it clear in your policy that their usage will be occasionally monitored. (employees are much more likely to agree with the employer’s right to monitoring if they are fully aware of the policy)
- Instruct them in the proper use of copyrighted material.
- Any use of the employees Internet access for illegal purposes is strictly prohibited
- Make it clear that Internet access is provided by the employer, is owned by the employer and that the employee should not expect their activity to be private.
- Violations of the policy will result in discipline or termination.
[Surfcontrol.com]

Once your policy is written, reviewed and signed by Management, Human Resources and Legal, make sure it is properly distributed to the employees. This is not something that you want hidden away. It is suggested that the policy be distributed to employees and that they sign off on a statement indicating that they have read and understand the policy. This may seem like an unnecessary step, but it will be invaluable should you be faced with a legal challenge to the policy.

Reasonable care should also include an Internet filtering product as part of your defense. Just like your IAP can not guarantee that you will never be subjected to a lawsuit as a result of an employee’s wayward Internet surfing, Internet filtering products are not perfect. It’s the combination of an appropriate Internet Access Policy and the application of an Internet Filter that is going to provide the proper level of protection. These two implementations will allow you to meet the reasonable care standard.

There are various methods utilized to accomplish the goal of preventing a user from reaching undesirable web pages. Currently you can choose from software or hardware; keyword, host, protocol, site or content rated blocking lists; centralized or local implementations and various combinations of these functions.

Your choice of an Internet filtering product will depend on a number of factors. You may

want to answer these questions before you begin to consider which Internet filtering method is right for you.

- How many users will be filtered?
- What is the user's proximity to the administrative resources?
- Does the product need to be periodically updated and how is this accomplished?
- If the product incorporates predefined blocking lists, how up to date are they?
- Will everyone be filtered to the same degree?
- How much will be budgeted for the purchase and continued administration?
- Do you have a need for logging and reporting of user activity?
- Does the product have an automatic or even manual failover capability?
- Do you need users authentication or can the product use existing mechanisms?

The One Notable Exception

Hopefully I have convinced you of the need to filter your Internet access. However there is one exception, Internet Access provided by publicly funded libraries. There are various court battles contending that the filtering of Internet access provided by publicly funded libraries is unconstitutional and is a violation of an individual's freedom of speech. This is a very contentious issue especially since the Children's Internet Protection Act was passed by Congress in December 2000 and made Internet filtering a requirement for certain federal funding.

On March 25, 2002 this issue was brought before a three judge panel in Philadelphia's Third Circuit Court of Appeals. The intent of the CIPA was to protect children from undesirable content on the Internet. Various library associations and the ACLU contend that Internet filtering is ineffective and puts individuals that can not afford private Internet access at a disadvantage. Numerous solutions have been proposed including separate library sections for filtered and unfiltered access and parental consent forms for children. The problem actually completes a full circle when you consider that the library patrons may require unfiltered Internet access. But what about the library workers? Don't they deserve to be able to work without being subjected to a hostile work environment? This issue will not be solved quickly. Any appeal to the Philadelphia's Third Circuit Court of Appeals will go directly to the Supreme Court. [Pruitt]

Software vs Hardware or Appliances

When you discuss hardware in the context of Internet access you are really talking about blocking not filtering. Hardware, such as routers, have been used for years to prevent users from access resources. If you didn't want user to access the Internet, block the IP address to the external gateway. If you didn't want your Internet users to chat, block port 194 for Internet relay chat. At this point two things should become clear, hardware blocks

user access, not filters it. And hardware used to block Internet access can be very labor intensive. Hardware can't block http traffic to specific sites unless you specify the IP address or range. Doing that manually, considering the proliferation of objectionable websites, is almost impossible.

The Software implementation of Internet filtering is much more flexible and comes in two types, client and server. Client software involves the actual installation of a filtering product on each client PC. This includes such traditional products as Cyber Patrol, Cybersitter, Net Nanny and Surfwatch. Because of the administrative burden of the local installation and maintenance of the software, these are best used for single PC environments such as home use. Although recently, server side implementations of these products have become available. [Impact magazine]

If your user base is large enough to warrant the expense, you will have to at least consider a software server solution. The advantages of server Internet filtering are

- Centralized administration
- Scalability
- No workstation overhead
- No workstation Installation
- Harder to bypass than client software
- Centralized reporting

The benefit of centralized administration can not be overstated, especially in a geographically disparate network. The advantage of making global or even individual user changes to Internet access from one location is a great saving in both time and money. It does create another problem that needs to be addressed, single point of failure. If your budget allows, most products allow for redundant server installations. You lose a server and Internet filtered access is still available. If providing redundancy on your Internet Filter is not the next item on your wish list, you could always accept the risk of being without Internet access for a short time or allow unfiltered access. That is your decision. However, you should be aware of the risk and make that decision before hand, not after the server is off-line.

Software Internet filtering is accomplished by a number of methods. Like the hardware implementations, software usually allows for specific protocol and address blocking. Blocking, as opposed to filtering is generally more reliable due to its objective nature. It is when software filtering utilizes its other methods of filtering that inaccuracies become evident. Some of these methods are as follows:

Protocol Blocking: This is not an actual filtering technique but more of a method to block entire communication ports to disallow specific types of traffic. This has the same result

as the hardware version previously mentioned.

Site blocking: One of the earliest forms of blocking technology employed by filtering products. Sites are actually blocked, not filtered based on their URL inclusion on a predefined list. This method has progressed to the point that filtering products provide periodic automatic updates to the site blocking list.

Keyword blocking: Sites are disallowed based on the existence of certain words contain on predefine word list. This is one of the first attempts at automating the filtering process. The major problem with this method is the misinterpretation of the word when taken out of context. For example, blocking sites related to criminal activity could cause law enforcement site to be blocked as well, due to their reference to criminal activity.

Rating Systems: Sites are rated according to various categories to determine if the site should be blocked and under which category. Digitally signed labels are embedded into the web pages which are read by the requester's web browser. Access to the pages is then controlled by the individual browser settings. [Beuselinck]

One browser that utilizes a rating system is Microsoft's Internet Explorer. You can find this option in Microsoft's Internet Explorer 6.0 by selecting Tools, Internet Options, Content and Enable the Content Advisor. To effectively block undesirable content you will want to deselect the option that allows the user to view site that are not rated. If this is not selected you are opening a large hole in you filtering process. In conjunction to this option there is an override feature that allows a "supervisor" to type a password to allow access to restricted content.

Once you have enabled this feature, you will have the choice of five levels of access for four different categories of filtering. The categories are language, sex, nudity and violence.

In the language category for example the levels are

Level 0: Inoffensive slang
Inoffensive slang; no profanity.

Level 1: Mild expletives
Mild expletives or mild terms for body functions.

Level 2: Moderate expletives
Expletives; non-sexual anatomical references.

Level 3: Obscene gestures
Strong, vulgar language; obscene gestures. Use of epithets.

Level 4: Explicit or crude language

Extreme hate speech or crude language. Explicit sexual references

Source: Microsoft Internet Explorer 6.0

If you would like more information on the organization that compiles these ratings, please visit: <http://www.rsac.org/ratingsv01.html> .


Current Internet filtering products are actually using a combination of all of these methods to at least some degree. The biggest change in this technology has been the way sites are reviewed. Websense, the product that will be mentioned next in this paper, develops a digital fingerprint of the site based on fifty factors. It is this digital fingerprint when applied to the categories that you select that determines if a site is blocked.

WebSense

Websense provides passthrough filtering to provide Internet access. Passthrough filtering means that all Internet traffic must first be passed to a control point such as a Proxy server, firewall or some type of caching device. For purposes of this discussion, Websense is installed on a Windows 2000 server in conjunction with a PIX firewall in a Windows NT 4.0 domain environment. Websense can also be installed on Windows NT 4.0, Sun Solaris and Linux Red Hat. The control point can be any one of at least 20 Proxy server, firewall and cache products. Almost all of the major players in this field are represented.

Websense currently contains over 75 categories from which you can select to customize your Internet filtering. For a complete listing of the categories and a description of their meanings, please visit <http://www.websense.com/products/about/database/categories.cfm>. Websense also has two premium categories that are available for an extra fee. Premium Group 1 is described as Productivity Management and includes Advertisements, Freeware/Software Download, Instant Messaging, Message Boards & Clubs, Online Brokerage & Trading and Pay-to-Surf sites. Premium Group 2 is described as Bandwidth Management and includes Internet Radio & TV, Streaming Media, Peer-to-Peer File Sharing, Personal Network Storage/Backup and Internet Telephony. In addition, Websense has the ability to dynamically add new categories as needed. Categories can be over ridden by specifically allowing or denying sites in a custom URL listing.

Categories are selected and access is grant or denied based on various filtering options. Besides the expected permit and block, access can be limited to time quotas or deferred to after work hours and are depicted as follows:



[Websense]

The collection of categories and their filtering options is called a category set. The category sets are combined with date and time restriction to form policies. It is the policies that are then applied to the users to grant Internet access.

One of the most attractive features of Websense is the ability to use authentication from other systems to allow access. LDAP, Active Directory, Windows NT or another Websense server can be selected for authentication purposes. When Windows NT is selected you need to install a service on the primary and all backup domain controller to handle the authentication passthrough to Websense. This is the DCAgent. The function of the DCAgent is to monitor logins at the controller and pass the IP address and user name information on to Websense, using default port 30600. It is this IP/user information that Websense will use to compare to applied policy to determine if the web traffic is allowed or blocked. The reason for this is that Internet traffic sent to the PIX firewall and then to Websense will not contain the user information, only the Source IP address. Websense must have that relationship beforehand to know what policy is to be applied.

All Domain controllers must be identified to Websense by inclusion of their IP addresses into the Websense server configuration section of the interface. Also on this screen is the option to allow users to be presented with a Windows NT authentication prompt if they are unknown to Websense at that time. This is convenient should the user be attempting Internet access without having had a previous successful domain authentication.

Users and groups can be added directly into Websense from the NT domain. You will have more control over the user population if they are added directly to the Websense server than if you decide to add global NT groups to Websense. This is true in organization where there is a select population of Internet users or users with varying degrees of filtering. It could be possible for an NT user to be copied to create a new user and the new user would be granted Websense access based on a copied NT group. This is nothing new to NT administration, just be aware that it could happen across the two

platforms.

Websense access can also be granted by individual IP addresses or IP ranges under the Network heading. Although this is not the way to go in a DHCP environment with users granted varying degrees of access, you will want to use this feature for any servers that require internet access. An example of this would be a server responsible for anti-virus updates from the Internet. If the update function was running as a service, meaning that no user need to be logged in for it to run, Websense would block the traffic. Inputting the static server IP address would allow the service Internet access.

By default, Websense contacts the domain controller every three hours to obtain updates to the user and group information stored on the NT Security Account Manager (SAM) database. You may find it necessary to change this setting to a more frequent interval. An example of this is: a new user has just been granted Internet access by inclusion into a NT global group that is defined on Websense. Based on the default setting of three hours between SAM updates, the access will not be in effect immediately. If you determine that this lag is not acceptable, you will need to change the default setting. This is accomplished on the Websense server via the websense.ini file. Open the file and find the section labeled websenseserver and add the following line.

PolicyCacheTimeout=20

The value of 20 is in minutes. Websense Server service will need to be restarted before the changes to apply.

Another setting that you may want to change from the default is how long user IP combination are stored in the DCAGENT tables on the domain controllers. The default is 24 hours. This means that the DCAGENT will only resolve a user IP combination if the current entry is more than 24 hours old. The DCAGENT has a thread that will remove the older entries and allow for the new combination to be established. You may not see the importance of this until a laptop user connects into the network and obtains a DHCP address that someone else had less than 24 hours ago. If you have users with varying degrees of access or no access, that laptop user may receive the dreaded Websense blocked screen:

[Websense]