



The Next Generation in Web Filtering

7.0

Protecting Your Business with Internet Filtering

A White Paper for Today's Businesses



TABLE OF CONTENTS

Introduction..... 3

Why Filter Internet Access..... 4

 Enhance Employee Productivity..... 4

 Minimize Legal Liabilities..... 5

 Increase Security 6

 Maximize Bandwidth 6

How Internet Filtering Works 7

 Interrupt versus Prevent 9

 Content Evaluation 12

How the Technology is Evolving 13

Conclusions..... 15

Works Cited..... 16

INTRODUCTION

When allowing or requiring employees to use the Internet in the workplace, organizations today are faced with the same two issues as they were yesterday—decline in work productivity and threats of possible litigation.

This white paper examines reasons for deploying Internet filtering solutions and reviews how filtering and monitoring tools are used to enhance the management of network resources.

Internet filtering and monitoring technology has developed very rapidly. The change has been from basic tools that simply logged visited IP addresses to sophisticated management solutions that show a much more complete picture of who is using the network, when the network is being used, and how the network is being used.

These tools have undergone an evolution that has closely mirrored the rapidly changing demands of managers and Internet users. Bombarded with unmanaged content, they want to more effectively use the Internet as a business resource without the associated problems and legal repercussions associated with unmanaged Internet access.

As this technology continues to develop, it is certain that the business needs that drive its deployment will also continue to effect changes in the function of the solutions as well as the way in which those solutions are used.

WHY FILTER INTERNET ACCESS

*...six in every ten (10) workers
admit to wasting time at
work...*

It is important to understand the business requirements that were driving the development of Internet filtering. When Internet filtering first emerged, many organizations only had basic Internet connectivity, usually consisting of a small number of employees with e-mail accounts. As e-mail became a more prevalent form of business communications, more people within the organization began demanding e-mail access. This rapidly built the case for more widespread Internet access throughout the organization.

At around the same time, Web development had reached a point where the Internet was becoming a cost effective mechanism for the delivery of information. Nevertheless, the quantity of non-business related content on the Web remained, and still remains unimaginably large. New Web sites are added to the Web each day, some of which contain malware, viruses, and spyware that could potentially attack important business information (Web of Trust). With the ever growing need for employees to access work-related information on the Internet, the distraction of accessing non-business material, and the threat of breached security, has become a more serious issue.

Enhance Employee Productivity

This kind of access entails a number of very significant business problems. Perhaps the most obvious is that of employee productivity. Studies have shown that a considerable amount of non-work browsing takes place during the average workday—one online survey found about six in every 10 workers admit to wasting time at work with the average employee wasting 1.7 hours of a typical 8.5 hour working day (Reuters).

With social networking sites, such as Facebook, growing at 5 million new users per week (CNN Money), more distractions are entering the workplace which can potentially lower productivity and company efficiency. Many employees prefer to do their personal browsing during the workday for several reasons: boredom, the workday is too long, being underpaid, and a lack of challenging work (Reuters). Thus, employees can easily spend up to 8.5 hours per week browsing the Internet while at work (Reuters).

One can see that Internet browsing at work negatively impacts the organization's bottom line. For an employee who makes \$20 per hour, the total amount of wasted resources per week is \$170, and \$8,840 per year. For a small business of 50 employees, that is a waste of \$442,000 per year which could be spent on needed capital expenditures.

Eleven (11) new pornographic sites are added to the Web everyday.

...companies are required to prevent, and not simply react to a hostile workplace.

Minimize Legal Liabilities

Every organization has a responsibility, recognized under law, to protect its employees. Should an employee be forced to work in an environment where he or she is made to feel threatened or discriminated against because of race, age, religious belief or gender, then the employer is held responsible.

An employee can sue their employer for sexual harassment if a colleague views pornography on a computer screen in the workplace. Title VII of the Civil Rights Act states that, “Unwelcome sexual advances...and other verbal or physical conduct of a sexual nature constitute sexual harassment. This conduct can explicitly or implicitly affect an individual’s employment, unreasonably interfere with an individual’s work performance, or create an intimidating, hostile, or offensive work environment” (Sexual Harassment).

The proliferation of pornographic material on the Internet has made the problem of hostile workplace lawsuits a particular issue among organizations that provide Internet access to their employees. According to Web of Trust, 34% of Internet users received unwanted exposure to porn in 2008. Also, research done by Web of Trust found that eleven new pornographic sites are added to the Web everyday.

A 1998 U.S. Supreme Court decision (Faragher v. City of Boca Raton) ruled that companies are required to prevent, and not simply react to, a hostile workplace. Should someone be exposed to such material while at work, the employer is potentially liable for damages, and this kind of lawsuit can be highly expensive and potentially embarrassing (oyez.org).

Web content itself may also be illegal. For example, numerous Web sites exist that offer free downloading of hacked or cracked copies of commercial products, in direct violation of international copyright law.

Again, if an employee downloads illegal copies of commercial products onto an organization’s network, it is the employer, not the employee that may be held responsible.

Increase Security

Anytime a browser connects to a Web page, there is an inherent vulnerability that can be exploited by malicious Web sites to perform a wide variety of potentially damaging acts.

According to PC World, in December 2008 Microsoft was forced to release a number of emergency patches and security bulletins to address problems with Web sites exploiting vulnerabilities in Internet Explorer, infecting over 2 million machines. Malware enabled personal information, such as passwords, to be stolen from at least 10,000 sites which Microsoft then needed to address and prevent.

Although an Internet filtering solution can't eliminate these problems, it can at least provide a degree of protection by preventing access to Web sites likely to host these kinds of files, often the very same Web sites that are also providing illegal copies of applications.


Maximize Bandwidth

Finally, it should always be remembered that Internet access makes demands on a finite resource, that of network bandwidth. As increasingly demanding content became available to people within organizations, such as large file downloads of music, streaming on-line radio, and now even view-on-demand movies, so the amount of that limited and expensive network capacity became consumed by non-business activities.

In an attempt to maximize the capacity of their existing infrastructure, as well as the previously mentioned issues, organizations began to look for a way to manage their Internet access.

HOW INTERNET FILTERING WORKS

Most Internet filtering and monitoring solutions have a number of elements and design assumptions in common. The basic technique is to place a monitor/filter between the client browser and the outside world. The filter evaluates a request for Web content with a set of pre-defined rules or “policies.” If there is a violation of those policies, the request is either blocked from establishing the connection, or the filtering software terminates the existing connection. Figure 1 illustrates how this process works.



Filtering Steps

- ① User requests Web site.
- ② Internet monitor/filter retrieves categorization of requested Web site from database and compares to policy.
- ③ Policy determines whether to allow or block the request.
- ④ a. Connection is blocked/denied.
b. Connection is allowed, Web page is fetched and sent to user system.
- ⑤ Activity is recorded in the activity log.

Figure 1
Basic monitoring and filtering processes

Rules are usually constructed to enforce an organization's Internet Access Policy (IAP). The majority of policies **block** access to Web sites that are deemed inappropriate, such as Adult or Gambling sites. Other policies may take the opposite approach, and **allow** access to certain Web sites under specific conditions. For example, they may allow access to shopping and travel-related Web sites, but only during non-working hours.

Rules are usually applied at three levels of increasing security:

- Level 1 Rules that apply generally to all entities within the organization. These prohibit access to Web sites regarded as potentially damaging, such as adult material, on-line gambling, hacking sites, pirated software sites, instant messaging and so on.
- Level 2 Rules that are applied at the *departmental or work group level* that are business specific, such as allowing certain groups to have access to Web sites generally not widely available or possibly available to others only at certain times of the day.
- Level 3 Rules that apply only to *certain individuals*. For example, Systems Administrators may need access to all Web sites, even those that would normally be blocked for other users, and rules are often constructed to allow them to bypass other access restrictions.

Normally, Internet filtering solutions also provide some degree of reporting capability, such that managers can review the type of access taking place within their organization and see which groups or individuals are most active on the network.

In order to enforce such rules, an Internet filtering solution must also have the capability to decide if any given request for content violates a previously defined rule. As a result, it must understand something of the nature of the Web site that is being visited, or conversely, evaluate the content as it is being transmitted.


To do so, there must be a further set of rules that define how to evaluate content, or there must be a knowledge base of Web sites against which the solution can evaluate the request.

Interrupt versus Prevent

In order to achieve these goals, there are a number of different strategies adopted by the various Internet filtering solutions. The first requirement, interrupt or prevent a given Web request, can be addressed in one of two ways.

Pre-connection Evaluation

Pre-connection evaluation does not allow Web access until the request has been fully evaluated against all the necessary rules, providing access only to safe Web sites. This approach effectively blocks all access until it has been allowed, and is something akin to an assumption of guilty until proven innocent.



Filtering Steps

- ① User requests Web site.
- ② Internet monitor/filter retrieves categorization of requested web site from database and compares to policy.
- ③ Policy determines whether to allow or block the request.
- ④ a. If the policy denies the request, the user is sent a block message.
 b. If the policy allows the request, the connection is made to retrieve the Web site from the Web server and display it on the user's system.
- ⑤ Activity is recorded in the activity log.

Figure 2

Pre-connection evaluation processes


This approach certainly provides a high degree of confidence that all Web requests will be evaluated; however, it does not provide any greater degree of assurance that a given Web request will be evaluated correctly.

The problem with this technique is that it essentially provides a single point of failure for the Internet filtering solution. Furthermore, because of the nature of the evaluation bottleneck, it presents a number of issues when the solution must be scaled across a large, wide area network. Additionally, those sites blocked by filtering software can be subjective according to the person or country where the list was compiled (Australia CISRO).

Post-connection Evaluation

Post-connection evaluation, or blacklists, allows all requests to proceed uninterrupted, but then evaluates the browser connection in parallel, and disconnects the client should a rule violation be determined. The advantage of post-connection evaluation is access to the entire Internet. The disadvantage of this approach is that an incorrectly scaled solution may become overtaxed in a very heavily loaded network, and allow unwanted sites to escape the filter. Figure 3 illustrates the post-connection approach.

However, any solution that is sufficiently flexible in the way it positions network monitors will not experience this problem and will be able to address the Internet filtering bottleneck problems associated with the former approach.



Filtering Steps

- ① a. User requests Web site and the connection is made to fetch the Web page.
 b. At the same time, the Internet monitor/filter sees the Web site request as it watches the network.
- ② Internet monitor/filter retrieves categorization of requested Web site from database and compares to policy.
- ③ Policy determines whether to allow or block the request.
- ④ a. If the policy denies the request, the user is sent a block message and receipt of the Web page is blocked.
 b. If the policy allows the request, the connection is made to retrieve the Web site from the Web server and display it on the user's system.
- ⑤ Activity is recorded in the activity log.

Figure 3

Post-connection evaluation processes

Content Evaluation

Another area where Internet filtering solutions differ is in their approach to evaluating the content of the Web site itself. As previously mentioned, there are effectively two approaches.

Real-time Classification

This approach attempts to evaluate content in real-time while the connection is taking place. This on-the-fly method would at first glance appear to be the most scalable. However, the current state of natural language processing is simply not up to the task of categorizing, with any degree of accuracy, the content of Web sites that may be anything from car maintenance tips to movie reviews.

Not only the content itself needs to be considered, but also the context of that content, both within the Web site and in relation to links to other sites. To attempt to evaluate all of this in real time introduces a great deal of unnecessary processing, slowing down Web access.

Also, without the option to include some degree of human intelligence in the decision-making process, categorizations are prone to the old problems of over or under blocking.

Database Classification

This approach compares the destination Web site to a pre-categorized list and evaluates the connection on that basis. This approach has some limitations.

The rapidly changing nature of the Web is such that using a database of pre-categorized Web sites requires constant work to update the content and maintain the accuracy of the database. Nevertheless, this approach still provides the most reliable and accurate way to categorize Web sites, and accuracy of categorization is of paramount importance. There can be little confidence that the underlying rules themselves are being correctly applied if the accuracy of the Web categorization methodology is questionable.

Also, poor categorization techniques dramatically reduce the quality and usefulness of reports produced by the Internet filtering system. It would be unwise to trust a report showing that no one is attempting to access adult material on the Web, if the basis of that report is incomplete or inaccurate Web site categorization data.

Accuracy and scalability have been proven time and again to be major considerations in any organization's Internet filtering strategy.

HOW THE TECHNOLOGY IS EVOLVING

The process of evolution has been spurred by the fluid nature of the demands placed upon it. The changing nature of the demands themselves has come from the growing understanding among organizations of the ways in which they can use the Internet as a business resource.

The early days of Internet filtering were beset by solutions that proved highly unreliable in their ability to accurately filter, whether inclusion or exclusion filtering. Those solutions are still available today, a fact that consistently provides ammunition for those who argue that Internet filtering solutions are prone to over blocking, under blocking, or subjectivity.

Nevertheless, the new generation of Internet filtering solutions has been built to address, more closely, the underlying business needs driving the adoption of this technology. Accuracy and scalability have been proven time and again to be major considerations in any organization's Internet filtering strategy.

In much the same way that the gradual mapping of the human genome has slowly begun to reveal the nature of complex interactions within the human body, the mapping of the Web by Internet filtering-solution providers has enabled organizations to begin to more fully appreciate the way in which their employees are using the Internet (cost reduction, increased speed, limited liability).

As more scalable and accurate Internet filtering solutions are now available, it is also possible for organizations to provide a more granular approach to their Internet access strategies, targeting specific kinds of access for particular groups of users in varying locations. Also, the ability to apply rules consistently to employees regardless of how or where they enter the network is essential in a true enterprise-wide solution.

...the same technology used to provide classification for content across the Web can be used to classify and monitor content within the organization's network.

A further enhancement of Internet filtering technology is the capability to work in concert with other products to provide an enhanced level of management. This synergistic effect is particularly important, as organizations are working to more fully understand how the business resources of employee time and network bandwidth are being utilized.

Internet filtering solutions are now working closely with e-mail filtering solutions to provide consistent enforcement of Internet access policy for organizations. Similarly, the same technology used to provide classification for content across the Web can be used to classify and monitor content within the organization's network. This synergistic cooperation between all levels of filtering provides a much higher degree of protection as well as a consistent level of security.

Internet filtering solutions are becoming a necessity for an organization's efficiency with company resources and the safety of their employees and customers. The ideal solution is one that provides both a high degree of accuracy and the ability to operate in concert with other informational gathering tools to filter and monitor across the entire spectrum of Internet activity.

CONCLUSIONS

As discussed, true enterprise-based Internet filtering solutions are going to provide two interrelated areas of benefit to the organization.

The first and most immediate benefit is to provide a method to protect both employees and employer from the dangers of unmanaged access.

These dangers, such as loss in productivity, legal liability, and bandwidth waste, are very real and cost organizations enormous sums of money each year. Unmanaged Internet usage also exposes the organization to costly data corruption or loss, and resource damage due to downloaded software viruses or other malicious code usage.

As stated before, an organization with only 50 employees can potentially lose up to \$442,000 in lost wages due to Internet browsing. This does not include other potentially wasted resources such as: overhead, bandwidth, protection, etc. Ultimately any non-business process that wastes funds in this way is a problem even in the best of economies.

The second benefit of Internet filtering is offering a far better understanding of how their Internet resource is being used, and how to maximize their investment.

Providing Internet access is usually a significant investment for any organization regardless of size. The actual physical connection and software associated with it are usually a small proportion of the total costs.

Additional Internet hardware, security costs, management costs and employee time all quickly add up. Without a method of understanding how this is being utilized, organizations have no methodology for measuring their return on this investment.

Businesses need to build better methods to access information in order to better manage and address real business needs. This requires, at the very least, a method to accurately and consistently monitor, filter and report on Internet activity.

The ideal solution is inevitably one that provides both a high degree of accuracy and the ability to operate in concert with other informational gathering tools to filter and monitor across the entire spectrum of Internet activity.

The need for businesses to access the Internet isn't going to go away. The quantity and the range of content available on the Web is going to continue to grow, challenging organizations to remain efficient in an ever changing global environment.

The challenge for Internet filtering solutions is to continue to build technology that meets the need for safe, managed access to essential business resources, while providing enhanced understanding of how the Internet is being used and at the same time eliminating an ever increasing range of harmful or disruptive content.

WORKS CITED

- “About Filtering.” OpenNet Initiative. Oxford, Cambridge, Harvard, Toronto. 23 Feb. 2009 <<http://opennet.net/about-filtering>>.
- Australia. CSIRO: Mathematical and Informational Sciences. ACMA. By Paul Greenfield, Peter Rickwood, and Huu Cuong Tran. Sept. 2001. 30 Feb. 2009 <<http://www.acma.gov.au/webwr/aba/about/recruitment/filtereffectiveness.pdf>>.
- Brenner, Bill. “Does Microsoft’s Patch Tuesday Need Fixing?” PC World 15 Feb. 2009. 23 Feb. 2009 <http://www.pcworld.com/businesscenter/article/159563/does_microsofts_patch_tuesday_need_fixing.html>.
- Content Watch. Working paper. 25 Feb. 2009 <<http://www.contentwatch.com/images/documents/papers/appliance.pdf>>.
- Helsinki, Finland. “WOT Finds Increased Security Threats in the Internet’s Red Light District: Adult Sites Cause the Most Damage to Internet Users.” Web of Trust. 18 June 2008. 23 Feb. 2009 <<http://www.mywot.com/en/press/wot-study-internets-red-light-district>>.
- Hempel, Jessi. “How Facebook is Taking Over Our Lives.” CNN Money 11 Mar. 2009. 12 Mar. 2009 <http://money.cnn.com/2009/02/16/technology/hempel_facebook.fortune/>.
- Sparrer, Curtis. “PC Tools Issues Warning About Looking for Love Online.” PC Tools 13 Feb. 2009. 23 Feb. 2009 <http://www.marketwatch.com/news/story/pc-tools-issues-warning-about/story.aspx?guid=24EEA405-3857-444A-92E3-0B8D4CC727B4&dist=msr_1>.
- The Oyez Project, *Faragher v. City of Boca Raton*, 524 U.S. 775 (1998), 25 February 2009. <http://www.oyez.org/cases/1990-1999/1997/1997_97_282>
- United States of America. The U.S. Equal Employment Opportunity Commission. Sexual Harassment. 11 Mar. 2009. 13 Mar. 2009 <http://www.eeoc.gov/types/sexual_harassment.html>.
- “Wasting Time at Work? You’re Not Alone: Survey.” Reuters 26 July 2007. 23 Feb. 2009 <<http://www.reuters.com/article/lifestyleMolt/idUSN2541395620070726>>.