



Avoiding Legal Traps on the Web: Protecting Your Organization with Internet Filtering

by Bret Hill, IIS MVP

As a security specialist, I find that companies pay plenty of attention to hardening their servers and networks but pay little attention to how uncontrolled Internet access from within an organization can represent a significant legal and security risk. For example, a university in California was warned it might be held liable for permitting downloads of a copyrighted movie¹. And a company in Arizona was fined \$1 million for file sharing activity hosted on the company's computers². These and other events add up to an increase in exposure to liability as a result of a worker's illegal activity on the Internet while using company equipment.

Additionally, users who browse a malicious Web site can become infected with a Trojan or other malware without their knowledge as a result of vulnerabilities in Microsoft Internet Explorer (IE) and the difficulty of keeping computer systems updated. Internet filtering technology, which is typically associated with increasing work efficiency by reducing time spent on non-essential Internet sites, is a key player in mitigating these threats. This white paper discusses the various methods available for Internet filtering and how to use them to increase security and decrease legal exposure.

¹ <http://www.timesonline.co.uk/article/0,,2-625793,00.html>

² <http://news.com.com/2100-1023-879308.html>

→ Contents

Reasons to Control Internet Access	1
Holding Organizations Accountable for Illegal Internet use	1
Providing Logs for Use in Forensic Studies	2
Meeting Legal Requirements	3
Preventing Access to Hacker or Cracker Sites	3
Towards a Solution: A Review of Available Technologies	3
Filtering Methods.....	3
Filtering Implementation Methods	4
Requirements of an Effective Solution..	5
iPrism Quick Installation	6
iPrism Integration	6
Integration with Network Architecture	6
Integration with Existing DNS	7
Integration with Existing SMTP	7
Integration with Existing Web Proxy.....	8
Integration with Web Cache Coordination Protocol Support	8
Managing Internet Access	8
iPrism Filtering and How it Works....	8
ACLs and Profiles: What They Are and How They Work	8
Users and Groups	10
iPrism Security.....	10
Reporting	10
Summary	11
References	11
About the Author.....	11



Copyright 2004 St. Bernard. All rights reserved

Avoiding Legal Traps on the Web: Protecting Your Organization with Internet Filtering

Reasons to Control Internet Access

Organizations have several incentives for wanting to control Internet access among their users, including

- holding organizations accountable for illegal Internet use
- providing logs for use in forensic studies
- meeting legal requirements
- preventing access to hacker or cracker sites

Holding Organizations Accountable for Illegal Internet Use

The Internet is obviously a valuable resource for many organizations. However, in my work with companies of all sizes, I find that many are exposed to security and legal liability concerns because they fail to control Internet access across their organizational assets. Often, IT departments simply put in place patch management and virus scanning software to secure the internal network, then consider the job complete. However, such measures don't protect an organization from criminal Internet use, intended or otherwise.

Additionally, most companies don't take formal steps to protect their systems against Trojans, viruses, listeners, and other malware that lure users to official-looking Web sites where users give approval to download updates. For example, such malware might lure a user to a site with the promise of free music downloads if the user simply installs a special "player." The player is actually a Server service that, once installed, is used to distribute illegally pirated DVDs. Such activity can bring down a network due to overwhelming traffic and result in lawsuits from the copyright holders. In some cases, the material might be ethically objectionable, resulting in distasteful materials being distributed from your organization's systems. In such situations, you might be required to keep such distribution on line while criminal investigations are occurring with Federal agencies.

As CNET's News.com reported³, Nancy Flynn, executive director of the ePolicy Institute, summarized a survey published by ePolicy with the American Management Association on surveillance in the workplace. "Productivity is a concern; loss of confidential information is still a concern; security breaches are a concern. But...the No. 1 concern is liability. Employers are afraid of being sued," she said. In my opinion, they have reason to be concerned.

The International Federation of the Phonographic Industry (IFPI) represents the interests of the recording industry worldwide. IFPI made headlines in the United States by suing more than 5700 individual file sharers for copyright infringement, not including 750 suits filed during October 2004. The same organization has carried out similar actions against individuals in the UK, Austria, Denmark, Germany, and Italy⁴.

Although these particular actions are against individuals and not organizations, it's just a matter of time until organizations without reasonable precautions in place to prevent illegal Internet use with company assets will be charged with liability. This next step was made clear in March 2004 when Sony, EMI, and Universal took legal action against the University of Melbourne, Sydney and Tasmania over students illegally downloading music⁵.

In fact, in *Copyright Use and Security Guide for Companies and Government*⁶, Jason Berman, Chairman and CEO of IFPI, states

Unfortunately, employees of companies and government bodies sometimes engage in unauthorised copying of music on the organisation's computer systems. This activity not only wastes the organisation's time and system resources, it is illegal. Such activities on your systems can put your organisation at risk of legal prosecution [emphasis added], tarnish your organisation's reputation and increase security risks for your computer systems.

IFPI sent a letter to every university in Britain reminding them of the implications of unlicensed Internet copying. IFPI went on to state

The legal risks include injunctions, damages, costs and possible criminal sanctions against the institutions and their heads where systems are used for copyright theft.

In the United States, the same article³ states

One student at a California college tripped an "electronic alarm" at Warner Bros when he downloaded a Clint Eastwood film. The company threatened to prosecute the college and the student had to write a letter of apology to Warner Bros as punishment.

In one of the most aggressive actions in the United States, News.com reported that the Recording Industry and Association of America (RIAA) filed a suit against Integrated Information Systems of Arizona because employees were using the company's resources to distribute copyrighted music; the suit was settled for \$1 million⁷. In perhaps the clearest statement to date of an organizations exposure to liability, Matt Oppenheim, RIAA senior vice president of business and legal affairs, said in a statement of the settlement, "This sends a clear message that there are consequences if companies allow their resources to further copyright infringement."

Providing Logs for Use in Forensic Studies

An increasing number of cases exist in which electronic documents are at the heart of criminal studies. Often, copies of email, Web server, or firewall logs are used as evidence of a crime. However, in many cases, no logs exist that identify how the Internet was used from workstations within the organization. Clearly, a detailed log of what sites where accessed, by whom, and for how long would be a tremendous value in identifying those individuals involved who might be using the Internet for illegal purposes,

3 http://news.com.com/Mind+those+IMs—your+cubicles+walls+have+eyes/2100-1014_3-5423220.html

4 <http://www.ifpi.org/site-content/press/20041007.html>

5 <http://www.timesonline.co.uk/article/0,,2-625793,00.html>

6 <http://www.ifpi.org/site-content/library/copyright-use-and-security-guide-english.pdf>

7 <http://news.com.com/2100-1023-879308.html>

gambling, downloading obscene material, trading stocks, or participating in other activities that violate policy.

Additionally, if company resources were used for illegal activity, such logs would provide detailed records of the activity, which would be invaluable in not only prosecuting the guilty parties but in shielding the organization from liability to the degree that policies in place to prevent such activity are actively circumvented or ignored.

Meeting Legal Requirements

Your organization might be legally required to install filtering to receive federal funding or meet security guidelines. For example, legislation that dictates such requirements includes the Children's Internet Protection Act (CIPA), which was enacted by Congress at the end of 2000 as part of house appropriations bill H.R. 4577. Under this law, K-12 schools and libraries are required to adopt an "Internet Safety Policy" and install filtering technology to receive certain types of federal funding. This law applies to all schools and libraries that receive discounted rates for the purchase of equipment and services used to access the Internet under the E-Rate program, through the Library Services and Technology Act (LSTA), or Title III of the Elementary and Secondary Education Act (ESEA).

Preventing Access to Hacker or Cracker Sites

When I teach users how to secure a server, I recommend disabling IE browsing by blocking port 80 outbound using either IP Security (IPSec) or Internet Connection Firewall (ICF). This is sound advice because of the continuing stream of discoveries regarding weaknesses in IE and other Windows components, including the buffer overrun condition described in Microsoft Security Bulletin MS04-028: "Buffer Overrun in JPEG Processing (GDI+) Could Allow Code Execution (833987)".⁸ This very serious flaw lets malicious code turn on a system simply by downloading a JPEG image in IE, which is a very common occurrence. Microsoft has addressed some serious problems with IE, but new vulnerabilities continue to surface, making IE an unacceptable risk to use on a highly secure sever.

Your options are far more limited on a client system. Even if you switch to a non-Microsoft Web browser, risks associated with contacting malicious Web sites still exist. Without leveraging any technical weaknesses, users can be fooled into downloading updates or shareware that contain malware. Once installed, the client's workstation can become a spam relay, a host for illegal material, a station that participates in illegal hacking activities unknown to the user, or a console for an attacker to further penetrate your internal network. Monitoring activity related to such sites

and restricting access for most users is clearly in an organization's best interest.

Clearly, the trend toward holding companies accountable for illegal network activity is a matter of increasing concern. Developing a written Acceptable Use Policy is an important first step but without the ability to enforce a policy, you can still be liable for the poor judgment of workers. Many organizations aren't waiting for this problem to become front-page news before taking action to implement technologies that enforce an Acceptable Use Policy. As I've indicated, these organizations will realize tangible benefits above and beyond shielding them from potential legal action. So what are the technologies available for implementing a filtering system and of these, which technology is best? Let's take a closer look.

Towards a Solution: A Review of Available Technologies

In general, the technology of limiting access to the Internet is collectively referred to as a "filtering" technology. Numerous ways exist to put a filter in place that will inspect outgoing traffic and enforce rules about what is and is not permitted. In this section, we'll look at the methods for filtering and the technologies by which they can be implemented.

Filtering Methods

Three primary ways exist for inspecting a request to see whether it meets Acceptable Use Policies. Each has its own pros and cons.

Keyword filtering. The simplest method of Web filtering is to block by keyword. Keyword filtering scans for specific words within the text of a page as it's downloaded. The filtering application blocks the page if it detects any one of the listed words.

The major downfall of keyword blocking is that it can't take context into account; as a result, keyword blocking often blocks acceptable content. For instance, keyword filtering on the word "breast" may inadvertently block sites that contain valuable research information about "breast cancer."

Additionally, because keyword filtering is sensitive to only text on a Web site, it's of little value for blocking pages that contain nothing but images (such as those found on pornography sites). In addition, keyword blocking can be slow, which impacts user response time while accessing the Internet.

⁸ <http://www.microsoft.com/technet/security/bulletin/MS04-028.msp>

IP filtering. More sophisticated filtering methods employ a database of URL or IP address information to block access to specific, predetermined sites. Blocking all traffic from a specific IP address has the advantage of being simple and fast. However, in today's environment where several different sites might be virtually hosted at a shared IP address, the IP filtering method blocks either *all* the virtually hosted sites or *none* of them. In this scenario, blocking one inappropriate site on a hosted server might preclude access to dozens of other hosted sites with acceptable content.

URL-based filtering. Unlike IP filtering, URL blocking can provide blocking down to specific pages within a Web site. Using this approach, a generic photo archive site containing mostly acceptable content, with only a few pages dedicated to nudity, could be made accessible with only the objectionable pages blocked.

Today's more sophisticated Internet access management techniques, such as those found in St. Bernard's iPrism appliance generally combine extensive URL filtering with the ability to handle direct entry of IP addresses.

Filtering Implementation Methods

After you've determined what filtering system (i.e., keyword, IP, URL-based) you want to use, you have to select where to place the system. Should you implement software on each PC? Should you dedicate a server for the filtering software? Should you use a stand-alone appliance? Each implementation method has far-reaching effects on your deployment, administration, and in the end, satisfaction.

Client-side filtering software. Client-side filtering software blocks Internet content based on the current listings in a local database, which must be updated periodically at each PC. Popular with consumer-oriented Web filtering, it's cost effective for very small numbers of users. However, this approach has several shortcomings when used in large organizations:

- **Deployment and maintenance**—The coordination and effort involved with updating many different desktop-resident databases can present a challenge in large environments, resulting in outdated and inconsistent filtering. In addition, when updates are required to the software itself, you must deploy hotfixes and updates to all systems.
- **Decentralized administration**—Each user has a lot of control over his or her computer and the software that it runs. Users can potentially reconfigure the filtering system, adjust the filtering database, or otherwise compromise the installation. Additionally, often no

simple way exists to enforce configuration standards from a centralized server, such as using Active Directory (AD) Group Policy or another policy enforcement technology.

- **Lack of flexibility**—Because the filtering software is tied to each machine, it's difficult to provide special permissions for roving individuals (e.g., teachers) or different permissions for multiple people who share one PC.

Server-based Web filtering software. Server-based Web filtering systems address many of the limitations of client-based filtering solutions but present their own unique set of problems. Instead of installing the filtering software on each local PC, you typically install the filtering engine on an existing application server (e.g., Internet Security and Acceleration—ISA—Server, Windows Server 2003, Windows 2000 Server, Windows NT Server 4.0, UNIX) configured either as a proxy server or as part of a firewall.

The server-based filtering engine automatically intercepts all packets requesting external Web-based content. It screens and filters all inappropriate content for the entire network from a unified central location using one database.

This approach is less prone to tampering by individual users. It's also simpler to administer because you only need to update one database. Existing server hardware *may* be able to be used, so this solution might have lower costs upfront, as compared with other solutions such as turnkey server solutions and appliances (discussed later in this section).

However, a server-based filtering solution has hidden costs and limitations, including:

- **Potential security risks**—Maintenance of the host OS becomes crucial. An intruder who exploits a vulnerability on the host server can compromise the filtering software. As such, you must take steps to protect such systems, which require a high degree of vigilance to harden, maintain, and monitor any mission-critical server.
- **Single point of failure**—If the server should fail due to hardware or software issues, this can result in an immediate halt to outbound Internet Web site traffic. Rerouting clients to a new filtering server can be complicated and take time. Such delays can be costly, especially for those businesses that are involved in time-sensitive transactions or that are dependent on access to mission-critical Web sites.
- **High administration and support costs**—Server-based filtering solutions typically require a dedicated server and an additional OS license. Additionally, such a

solution can require a significant amount of IT administration time to install, test, configure, and maintain the server, especially as new versions or patches of the add-on software or server OS are released.

Add-ons to firewalls. Filtering services can be added to firewalls or other network devices. However, firewalls are typically not optimized for large Web site database lookups, which often results in performance and scalability bottlenecks. Firewall filtering software has also, in certain situations, compromised the firewall security.

Dedicated filtering appliances. Dedicated filtering appliances are easier to implement than other filtering solutions because you can install them into an existing network with little effort or impact on the rest of the network's performance. Filtering appliances are optimized for Internet filtering, so they tend to offer the highest performance and scalability. Another benefit is a single point of contact for service, centralized administration, and centralized logging. Overall, dedicated filtering appliances with fault-tolerant capabilities offer the best solution with the lowest total cost of ownership (TCO) when compared with other filtering solutions. Let's look at how a dedicated filtering appliance coupled with intelligent URL filtering addresses weaknesses of other implementations.

Requirements of an Effective Solution

In reviewing many of the current technologies for Web filtering solutions, I've enumerated the benefits and shortcoming of each. With this list, and some additional requirements, we can identify the characteristics of an effective solution and apply them to the solutions listed above and see what comes out on top. For the purposes of this white paper, the Web filtering requirements that I've identified must include

1. URL-based blocking—Any of the technologies that I've mentioned above can implement URL blocking, but not all do.
2. Human review of Web sites—Thousands of new Web sites are added to the Internet every day. As a result, one of the core criteria in selecting a Web filtering solution is how accurate and current the filtering database is. No matter what method of blocking is used, any filtering scheme must include mechanisms for continuous updates and maintenance of the database used to make blocking decisions. Even the most comprehensive database will rapidly become
3. Centralized administration—Client-based filtering doesn't meet this criteria so it can be effectively eliminated as a solution.
4. Security—Any Web filtering solution must have a secure means for authentication to the administration console and reporting system. The solution should also provide for a real-time and messaging-based means to override rules that prevent users from accessing sites they legitimately need to access. Also, the system should offer protection against virus and worm attacks. Finally, downloads to a centralized database from a vendor should occur over a secure channel to avoid interception and modification.
5. Integration/ease of use—As a key criterion to any Web filtering solution, your Web filtering system should know about your existing users but be transparent to the network. Ideally, you would install and start running the solution with just a few modifications, and it would automatically know about the users on your network, intercept HTTP traffic, and start enforcing rules immediately with minimal setup requirements.
6. Reporting—Reports should allow for various degrees of detail and have facilities for being automatically delivered to department heads. Custom reports should be possible.
7. Scalability/availability—If the system should fail, traffic should either pass through or shut down depending on your requirements. If traffic is passed through, in the event of a failure, the system should ensure that Internet access is not impeded. The solution should also scale to support a large network and provide failover to other identical filtering systems.
8. Functionality beyond filtering Web-based content—Access to Internet services such as ICQ and IRC chat, FTP downloads, RealAudio broadcasts, and MP3 music can expose your organization to criminal lawsuits,

outdated without constant review and updating. In addition, structuring the database into logical categories with mechanisms for adding customer-defined categories increases the value of the filtering product because these functions enable organizations to tailor their filtering policies to meet their specific requirements.

The importance of human review in the database maintenance process cannot be overstated. Using automated tools to identify new Web sites that *might* contain unacceptable content is a great place to start. However, relying solely on those tools for making the final determination can lead to many inaccuracies that end up annoying users.

consume significant bandwidth, and degrade employee productivity. Web filtering products that address these types of content and services add value to access control.

When reviewing these requirements as whole, it becomes clear that an Internet appliance is the solution of choice, providing that the vendor's implementation meets your requirements. One such appliance is St. Bernard's iPrism. This device stands out from other Web filtering appliances for many reasons, including allowing for human review of Web sites, offering ease of integration into the existing network, and providing a very intuitive administrative interface.

iPrism Quick Installation

Very few steps are required to get iPrism up and running. The default settings provide immediate benefits, but you'll want to configure iPrism to implement your Acceptable Use Policies to suit your requirements. The key here is that you can install and start using iPrism very quickly, with little or no modifications to the client systems or firewalls.

You can install iPrism "in line" so that it simply and automatically monitors and acts on HTTP and other traffic. In this configuration, iPrism acts as a network bridge; as a result, it's transparent to the network. Clients require no additional configuration because their traffic will automatically pass through iPrism.

Configuring iPrism for immediate filtering is simply a matter of running the configuration wizard from a client workstation. The Java-based iPrism administration tool lets you run the initial configuration wizard, which asks for basic network-related information. Figure 1 shows the first screen of the wizard, which is where you specify the IP settings.

After you enter this information, you must complete a few other simple forms, after which you can access the iPrism administration system from your Web browser or the




Figure 2: Placing iPrism in line with the existing network hardware

client tools, whichever you prefer. The wizard will also step you through enabling the default filtering rules that filter objectionable content immediately.

With just these few steps, you've configured the iPrism device for administration and to filter Web access for objectionable sites according to predefined rules in the default policy. The database of sites that iPrism uses will automatically be updated over port 80, which most firewalled environments permit. Of course, you'll want to fine-tune your filtering settings—this is where iPrism really shines. First, let's look at some features that let you integrate iPrism into an existing network infrastructure.

iPrism Integration

You wouldn't want to reconfigure your network or add new servers to support your filtering solution. Ideally, you want an appliance that can integrate with your existing network infrastructure and leverage your existing services. iPrism has the flexibility to be placed into a network as a bridge or router and use existing DNS, SMTP, and other network services.

Integration with Network Architecture

iPrism satisfies the requirement that our Web filtering solution must integrate easily with an existing network. As I mentioned, iPrism acts a bridge and you can simply place it in line to the router so that all traffic to be monitored automatically

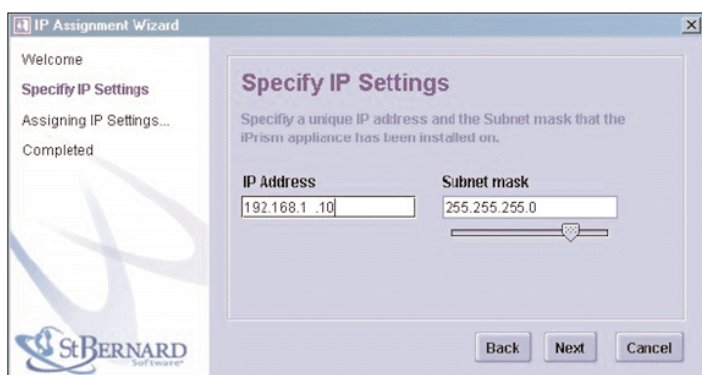


Figure 1: iPrism Initial Configuration Wizard




Figure 3: Installing iPrism in a complex network environment

passes through the iPrism device, as Figure 2 shows.

Although this configuration is the most common, you can also configure iPrism to route network traffic as required. You can also install iPrism as a bridge in more complex network environments, as Figure 3 shows.

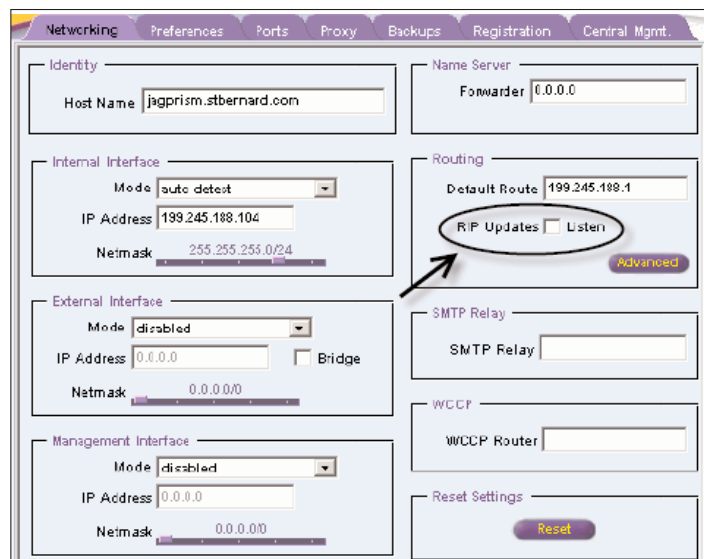


Figure 5: iPrism support for RIP

In the above example, iPrism is placed in line in front of the internal firewall. iPrism is generally placed on the inside of the trusted network to filter all outbound traffic without the restrictions of the firewall in place.

If you need to route traffic from other networks to and from iPrism, you can configure iPrism to use network traffic static route tables, as Figure 4 shows.

Alternately, you can easily enable iPrism to automatically create and dynamically adjust routing tables by listening for inter-router communications via Routing Information Protocol (RIP). iPrism supports RIP 1 and RIP 2, as Figure 5 shows.

Integration with Existing DNS

iPrism constantly resolves Internet host names to their IP address, as well as reverse-map IP addresses to their host names. If iPrism's installed environment allows direct Internet access, it will (by default) use its built-in name resolver to perform all DNS tasks. However, some installations require that iPrism defer all DNS lookups to another name server, called the "fowarder" name server. iPrism lets you designate multiple forwarding name servers or use its internal name server, as you prefer.

Integration with Existing SMTP

iPrism uses SMTP to perform the following types of communications:

- Reports/log files exports
- Email alerts
- System notifications (upgrades, filter list problems, registration)
- Access requests

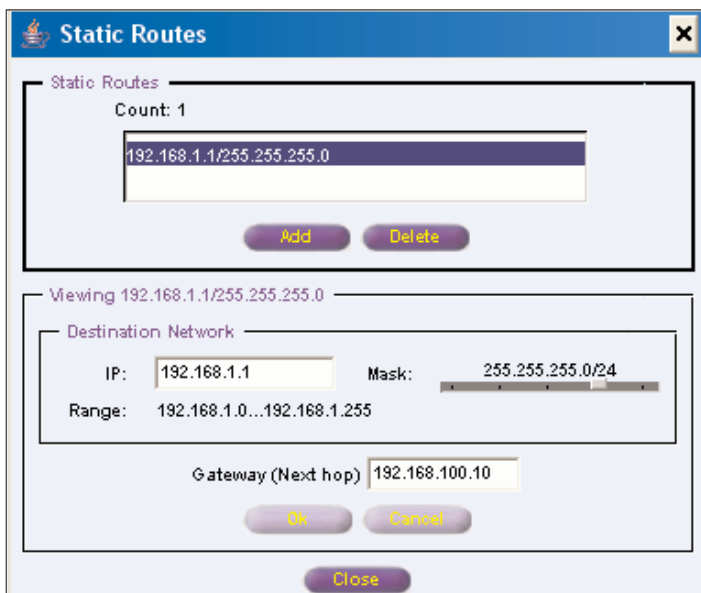


Figure 4: Configuring iPrism to use static route tables

By default, iPrism performs a DNS (MX record) lookup to deliver these emails. If you install iPrism in a network where a DNS server isn't available and an SMTP Smarthost is used (for efficiency), you can configure the Smarthost's IP address in the SMTP Relay field in the iPrism interface. If an SMTP Relay is specified, iPrism will delegate the delivery of the email to the relay and not attempt to directly contact the recipient's mail server.

Integration with Existing Web Proxy

Under typical circumstances, iPrism contacts remote sites (that aren't blocked) and retrieves Web pages on behalf of the original requester. However, if you have a local Web-caching server, you can configure iPrism to make requests to the caching server, rather than to the Internet. This configuration lets you maximize your existing investment in Web caching while simultaneously introducing intelligent Web filtering. Additionally, you can configure iPrism to obtain its database updates through the Web caching device. In this way, all iPrism traffic is sent to or from the Web caching device, thereby optimizing integration and eliminating the need to create special rules on the firewall for the iPrism device.

Integration with Web Cache Coordination Protocol Support

iPrism supports the Web Cache Coordination Protocol (WCCP)⁹ v1, which provides fault tolerance by automatically re-routing traffic directed to Web sites in the event that iPrism is turned off, disconnected, or a system failure occurs. When the client workstation generates traffic outbound to Web servers on the Internet, a WCCP-enabled router detects the HTTP traffic and diverts that traffic to iPrism using a Generic Routing Encapsulation (GRE) tunnel. iPrism then makes the request to the server on behalf of the client and responds directly to the client. However, from a client's perspective, the response appears to come directly from the server the client requested, so the client doesn't know it's communicating with iPrism.

Managing Internet Access

You can see that you can configure iPrism in numerous ways to integrate into your network. But what iPrism does best is that it lets you configure Internet filtering to suit your requirements. Let's look at how iPrism accomplishes this task.

iPrism Filtering and How it Works

iPrism can filter access to Web sites by checking each client's Web request against an extensive URL database.

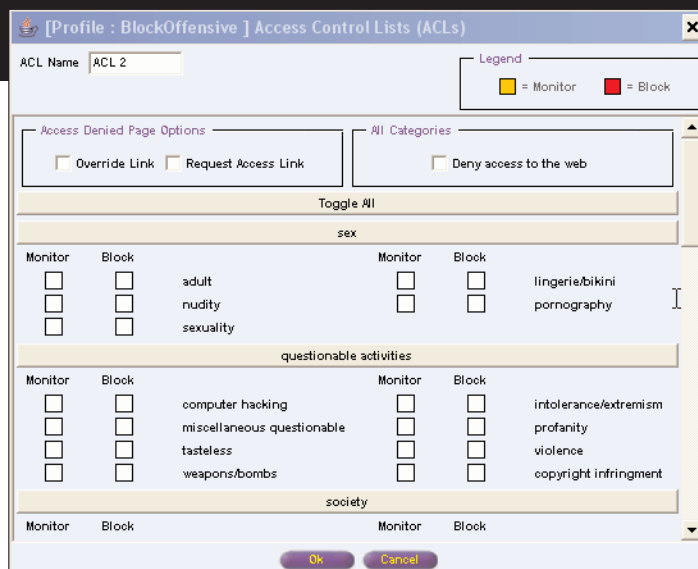


Figure 6: Selecting categories of Web sites to block or monitor

This database classifies each listed Web site according to one or more categories (e.g. adult, pornography, violence, drugs, sports, entertainment). If the requested Web site falls within one of the categories to which the iPrism administrator has chosen to block access, then the client receives a Denied Access page instead of gaining access to the Web site. This customizable page notifies the client that the Web page the user tried to access belongs to a category that's currently being blocked.

An iPrism "profile" determines whether a specific user can access a specific site, and this profile is either active for that user or for that user's network. The profile consists of one or more ACLs assigned to the profile. ACLs are time sensitive in that you can enforce them or disable them, depending on the time of day.

In essence, the process for determining whether iPrism will take any action depends on iPrism making a determination for each request that the user is identified as an individual, member of a group, or working on a network assigned to an active profile, and that the requested Web site is associated with an active ACL set to either monitor or block the action.

ACLs and Profiles: What They Are and How They Work

ACLs are the building blocks of profiles—they alone determine which types of Web sites get blocked. An ACL can also designate that a Web site is permitted, but should be monitored rather than blocked.

ACLs are built or modified by simply selecting the Block or Monitor check box next to a preconfigured, but customizable, list of categories, as Figure 6 shows.

iPrism contains many more categories out of the box than the figure shows. For a complete list of automatically

⁹ <http://www.cse.ohio-state.edu/cgi-bin/rfc/rfc3040.html>

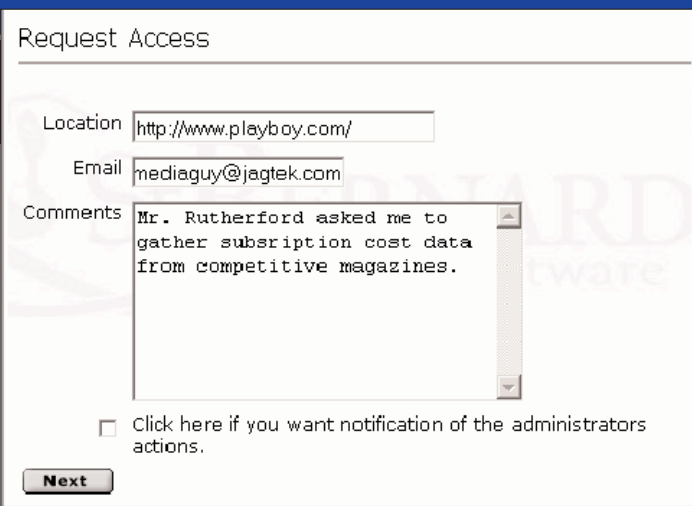


Figure 7: Sending a Request Access email

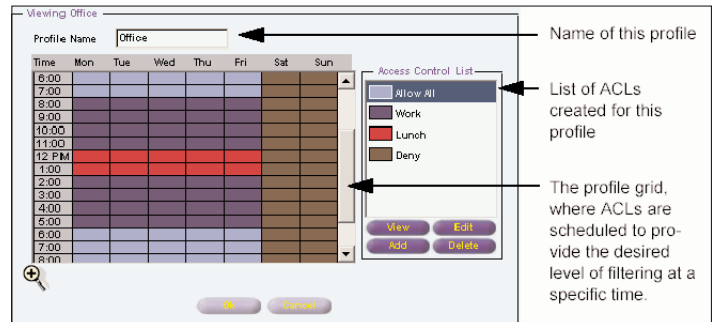


Figure 9: Assigning characteristics to a profile

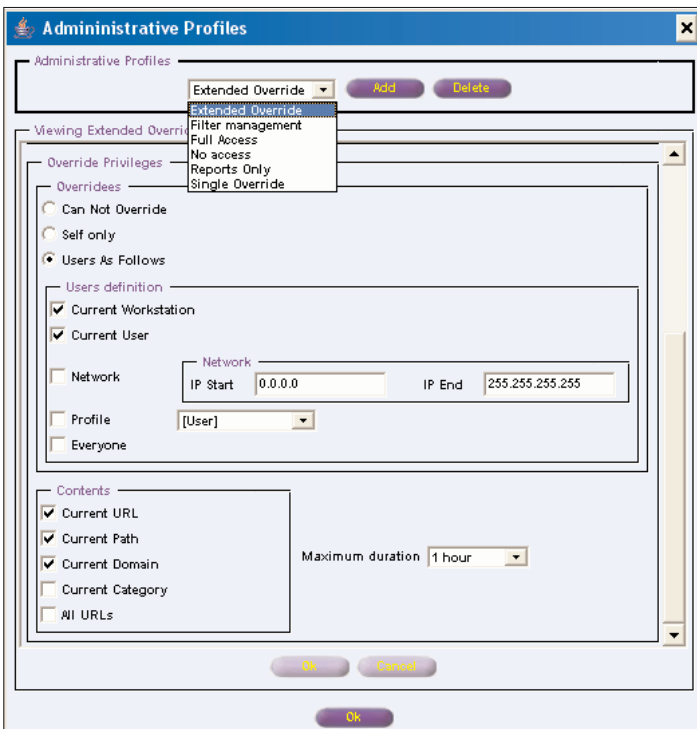


Figure 8: Configuring administrative rights for a profile

maintained categories, visit the iPrism Web site¹⁰. St. Bernard employs a staff who continually review new and existing content for proper classification. Nothing can substitute using a continuously updated, automatically deployed database of classifications for sites that's 100 percent human-reviewed. Even so, every business is unique, so you can add your own sites and customize the categories, as you prefer. Categories you create are called *Local categories*.

The box titled Access Denied Page Options in Figure 6 controls the specific action iPrism takes when access to a Web site is not permitted. When you select Override Link, the customizable Denied Access page will include an Override button. Then, if the user isn't authorized to view the Web page, iPrism prompts the user for authorized credentials before denying access to the Web site. If you select

Request Access Link, the Denied Access page will include a Request Access button, which lets the user send an email to the administrator requesting an override for the block, as Figure 7 shows.

One feature that I find particularly well designed is the ability to configure override privileges for individual users. For example, if a user is doing market research about Internet gambling, you can grant permission for just that user's workstation to override any Web site blocks that you have in place. These overrides are time sensitive, so the next time the user visits a blocked site, he or she must again re-authorize to gain access. The Administrative Profile form, which Figure 8 shows, lets you determine the profile for "Administrative" actions associated with the iPrism for each user or network.

Similarly, you can give the department head the ability to override a block for the entire network or a group of users as identified by a profile. This kind of flexibility in real time is one of the key areas that makes or breaks a filtering solution.

With iPrism, you can provide a descriptive name, assign ACLs, and configure the time of day for your profiles, as Figure 9 shows.

You can assign up to eight ACLs to a profile, which permits a high degree of customization per profile. Finally, you can implement profiles, such as the Office profile shown above, in a several ways. For instance,

- you can apply a profile to the entire network, a subnet, or a single workstation.
- if you're using authentication, you can assign a profile directly to individual users, or any combination of users and networks. When you assign a profile to a user, that profile has precedence over any profile that you might apply to the network. Profiles assigned to a user are always applied to that user, regardless of which workstation the user logs into.

10 http://www.stbernard.com/products/iprism/products_iprism-cats.asp

Users and Groups

iPrism doesn't require you to authenticate users to filter or monitor Internet activity. Instead, you can simply implement generic rules that are in place for all users and create accounts for administration purposes that are local to iPrism. Many organization, however, have domain controllers (DCs) in place and prefer to integrate their existing users and groups with the iPrism profiles.

When iPrism works with authenticated users, you gain enhanced reporting, blocking override, and granularity with ACL and profile design. iPrism has three different ways for dealing with users and groups:

1. Integration with Windows systems via Windows NT LAN Manager (NTLM)—The NTLM authentication feature lets iPrism access users directly from Windows DCs. iPrism joins the domain, which allows seamless authentication of Windows Server 2003 and Windows 2000 Server users and the ability to obtain group assignments. With this information, iPrism can control and monitor Windows domain user access to the Web. iPrism provides a simple mapping scheme in which Windows groups are associated with iPrism access profiles and administrative privileges. This functionality allows for extremely detailed control of externally authenticated users.
2. Obtaining user and group information using Lightweight Directory Access Protocol (LDAP)—iPrism can authenticate users and optionally obtain access information (i.e., an iPrism access profile name) for those users from a remote LDAP server. Although this approach isn't as full featured as NTLM authentication, it can be a useful method for obtaining user and group information.
3. Locally defined users (on iPrism)—If you're not using an external authentication method but you still want to reap the full benefits of running iPrism in authenticated mode, you can create user accounts directly in iPrism. You can create these user accounts individually or import them from a text file. Groups can't be defined using this method.

iPrism Security

Any time you add a network device that communicates with an untrusted network, security is a key consideration. St. Bernard designed iPrism to be secure in multiple layers:

- The OS running on the iPrism device isn't subject to attacks directed at Microsoft OSs.
- iPrism contains a built-in administration Web interface that can't be accessed from the Internet side.
- Administrative communication to and from iPrism is

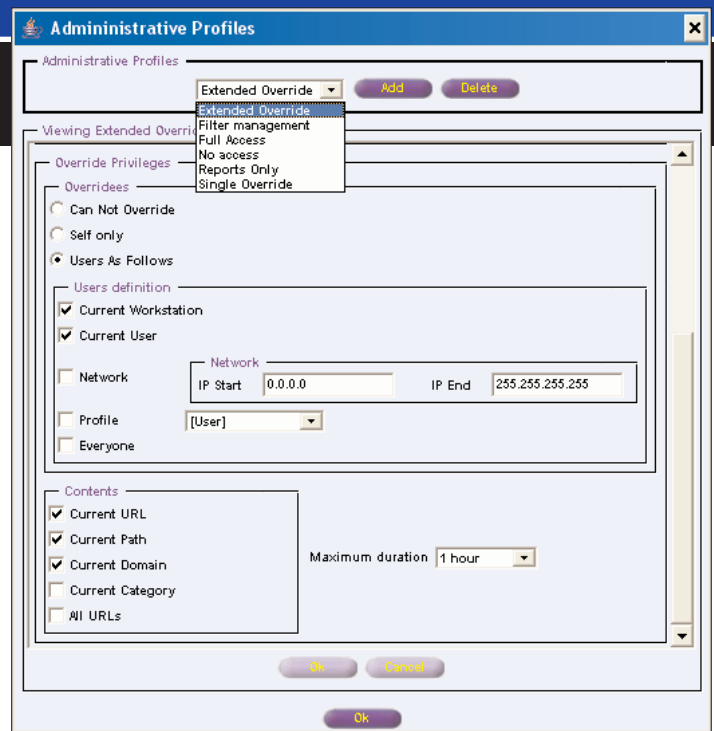


Figure 10: iPrism's predefined reports

- encrypted via Java or Secure Sockets Layer (SSL).
- Authentication to iPrism uses NTLM on the private network to authenticate users when required.
- ACLs let you create blocks based on file extensions to prevent users from downloading potentially harmful content.
- Automatic updates to iPrism occur over an authenticated and encrypted channel.
- iPrism runs no custom user code.
- iPrism is a dedicated device that you can load balance with other slave iPrism devices to help mitigate Denial of Service (DoS) attacks.
- iPrism filters protocols other than HTTP to control outbound non-HTTP traffic. This filtering prevents downloading of illegal or harmful content via FTP or Instant Messaging (IM) applications.

Reporting

Reporting is another area in which iPrism excels. The predefined reports, as Figure 10 shows, are well designed such that in many cases, these will be all that your require.

If you want additional depth or focus, you can customize the reports to suit your needs. You can access iPrism's reports on request using the administrative interface, via email, or publish them to an FTP server. A quick look at the Most Accessed Sites report, as Figure 11 shows, will tell you whether a lot of stock trading or other non-business related activity is occurring on your network.

One of the most useful features built into the reporting system is the ability to designate how frequently you want a report to run and to whom the report should be emailed. This functionality lets you easily schedule delivery of

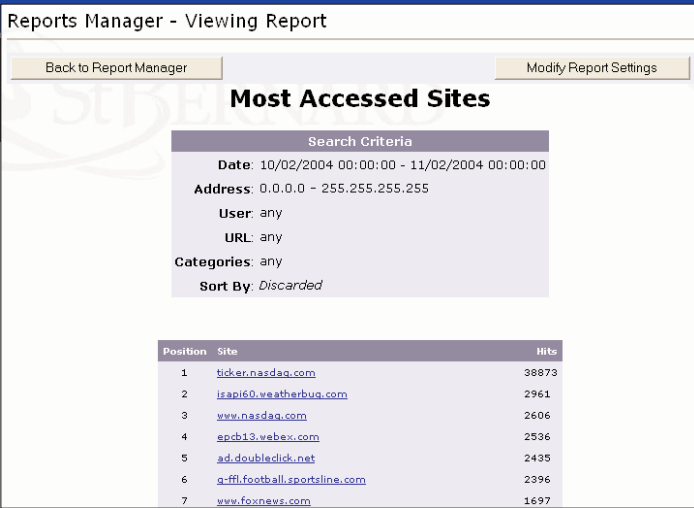


Figure 11: Viewing the Most Accessed Sites report

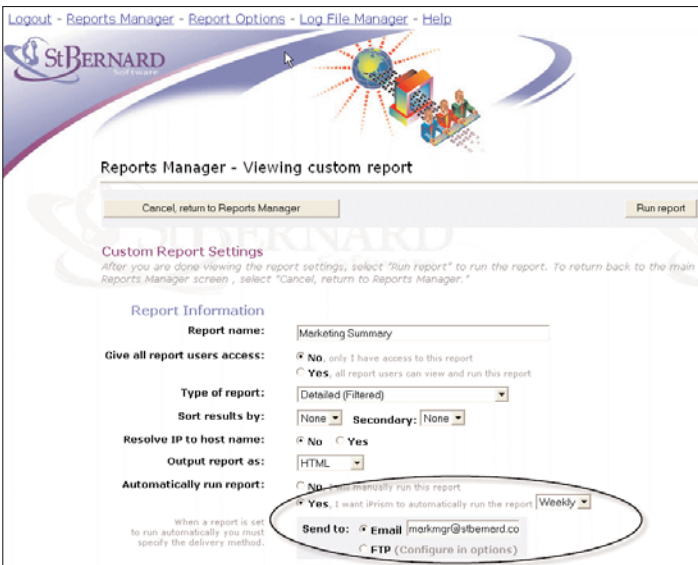


Figure 12: Scheduling delivery of reports in Reports Manager

reports related to a specific department to the department head. You can easily configure these settings in the Reports Manager, as Figure 12 shows.

In this way, you can deliver a summary of activity to managers so they get a snapshot of important events on an automated basis. This automatic report delivery keeps them informed on a regular basis so they can intercede wherever problems exist or adjust to new trends.

Summary

We started by looking at the problems of increased liability to organizations resulting from illegal use of their network resources. In a growing trend, more and more organizations are being held responsible for securing their networks from activities associated with copyright infringement. I believe it's just a matter of time until a company is sued because of hacking or other illegal activity that occurs using their computers. This risk has caused many IT departments to consider the means by which they can control Internet access on their networks.

In reviewing available filtering technologies, a dedicated appliance is the clear winner in terms of ease of integration, security, scalability, and administration. St. Bernard's iPrism is a clear example of such a device and sets a high standard in this market. In particular, iPrism's ease of use gets high marks as well as its flexibility to be as simple or complex as you require, integrate into your existing network infrastructure, and design (and deliver) custom reports. In my opinion, this is the box to beat.

References

iPrism Product Page

<http://www.stbernard.com/iprism>

iPrism Administrator's Guide

http://www.stbernard.com/products/docs/ip34_adminguide/

iPrism Installation Guide

http://www.stbernard.com/products/support/iprism/support_iprism-iguide.asp

iPrism Technical Support FAQs

http://www.stbernard.com/products/support/iprism/support_iprism-faqs.asp

Thinking Inside the Box: Benefits of an Internet Filtering Appliance

<http://www.stbernard.com/products/docs/ThinkingInsideTheBox.pdf>

iPrism Tech Notes

http://www.stbernard.com/products/support/iprism/support_iprism-tnotes.asp

Includes:

- iPrism in a Proxy Server Environment
- iPrism Auto-Login in Transparent Mode
- iPrism Auto-Login in Proxy Mode
- External Authentication in iPrism Using NTLM
- Windows 2000/2003 LDAP Authentication
- iPrism and LDAP for Novell 5.x and 6.x
- iPrism Central Management
- F5 Load Balancing
- iPrism Using WCCP Router
- How Does iPrism Work?
- Log File Format and Export
- WebTrends Integration
- iPrism Bypass Mode
- iPrism SSL Certificate Management (for Windows Networks & Internet Explorer)
- iPrism with the Cisco CSS 11000 Load Balancer
- Citrix/Terminal Server with NTLM/Auto-Login
- iPrism on Networks with Multiple Subnets
- iPrism Custom Filters

Brett Hill (brett@iisanswers.com) is a contributing editor for *Windows IT Pro* and operates <http://www.iistraining.com>. He is an IIS MVP, a Microsoft IIS consultant, an author, and a technical trainer offering IIS classes.