Comparing Architectures For Internet Content Filtering Solutions

Innovative website caching and rating architecture delivers an affordable, enterprise-class filtering solution for businesses, schools, and libraries

CONTENTS

***Abstract**: Businesses, schools, and libraries with Internet connections need the ability to control access to objectionable or inappropriate content. Without that control, businesses risk productivity loss, erosion of available bandwidth, and legal liability. The stakes are just as high for schools and libraries, which stand to lose federal funding if they do not provide the content filtering mandated by the Children's Internet Protection Act (CIPA) of 2000.*

Content filtering solutions generally conform to one of three architectures: client, server, or gateway. An organization's choice of architecture determines the solution's effectiveness, cost, and manageability. The SonicWALL Content Filtering Service (CFS) is a next-generation, gateway-based solution. SonicWALL CFS leverages an innovative website caching and rating architecture to deliver the scalability and flexibility of an enterprise-class filtering solution, at a breakthrough price point. The service provides content ratings in over 50 categories for more than 4 million sites, with more sites added daily. The management interface, accessed via a Web browser, provides centralized, flexible control with very low administrative costs.

This white paper makes the business case for content filtering, explains the pros and cons of different solution architectures, and describes how the SonicWALL solution enables optimum protection and productivity for small and mid-sized businesses, schools, and libraries.

Content filtering: its rising importance

As Internet use grows in business and education, so too do the risks of uncontrolled access. When workers inadvertently or deliberately access sites containing inappropriate, illegal, or dangerous content, organizations lose productivity, expose themselves to legal liability, and in some cases experience degraded network performance.

For businesses, installing a content filtering solution can eliminate many of these problems. A good filtering solution can help by:

- Improving employee productivity - Restricting Web access to appropriate sites can enable companies to prevent excessive non-productive Web surfing.
- Minimizing liability exposure - Businesses that prevent access to objectionable content shield themselves from potential legal liability that can arise if an employee repeatedly sees offensive material on a co-worker's compute.
- Preserving network bandwidth - Preventing downloads of large video files and pictures unrelated to business helps preserve bandwidth and application performance.

- Preventing hacker attacks and protecting privacy - Blocking automatically downloaded files such as Java applets and ActiveX scripts helps protect the network from viruses and hacker attacks, while rejecting cookies helps ensure privacy is not compromised.

Education institutions and libraries benefit from content filtering for many of the same reasons as businesses. In addition, content filtering addresses certain other requirements unique to schools and libraries:

- Keeping students focused - Shielding students from distraction helps keep them focused on the curriculum.
- Protecting federal funding - Schools and libraries must adhere to certain conditions to receive discounted rates for Internet access under the Federal E-rate program. They are entitled to federal assistance for Internet access only if they install software that blocks obscene or pornographic images and prevents minors from accessing harmful materials. The requirement began applying to libraries as well as schools after a June 2003 decision by the United States Supreme Court. Content filtering is also a mandate for libraries that apply for grants under the Library Services and Technology Act (LSTA).

Architecture options for content filtering solutions

Content filtering solutions conform to one of three architectures: client, server, or gateway. These architectures vary in terms of their effectiveness, cost, and manageability.

Client architecture

Client-based content filtering solutions are software products installed on individual desktops. The software includes a management interface and a database of blocked websites; the user downloads database updates via the Internet. Leading vendors of client solutions include Zone Labs, Net Nanny, and Computer Associates. In addition, some Internet Service Providers (ISPs), such as Microsoft MSN and AOL, integrate client-based filtering into their standard broadband service offerings.

The advantage of the client-based approach is the low initial cost of software-only solutions, making them popular with home users. The low initial cost, however, is counterbalanced by higher costs later on. One reason for the high total cost of ownership is that database growth is constrained by the amount of available storage on each desktop, limiting scalability. Another reason is a lack of manageability. The organization cannot enforce regular, frequent database updates. Rather, the administrator must visit every desktop to install software upgrades or change filtering policies.

Client-based solutions can also provide less effective content filtering than other architectures. For example, because users download database updates at their own convenience, inappropriate content might be accessible for some time before it is blocked. What's more, savvy workers or students can sometimes uninstall or disable the filtering function entirely.

Server architecture

The server architecture consists of two components: a dedicated database server platform and a separate gateway or firewall (see figure 1). The database contains URLs, their content categories (such as nudity or violence), and specifications on which categories are allowed or blocked. The gateway or firewall enforces the content filtering policies on the server. Leading server solutions include Websense®, N2H2™, and SurfControl®.

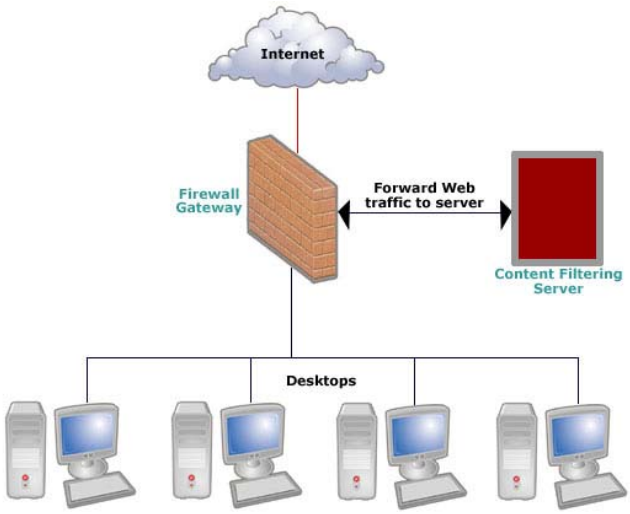

Figure 1.

Server-based content filtering solutions are more manageable than client-based solutions. The network administrator can create a filtering policy once at the gateway, and then apply it across all desktops. This eliminates the need to update every desktop individually and ensures that policy is applied consistently.

But the simplified manageability comes at a high cost. The business, educational institution, or library must purchase two hardware devices—the dedicated server and its companion gateway or firewall—as well as the content filtering software. With two devices to manage, organizations double their equipment cost and management burden, and may even be forced to dedicate resources to manage this solution. For distributed organizations, deploying an effective content filtering solution at each site can be prohibitively expensive.

While the server architecture is more manageable than the client architecture, it does not address a chief drawback of client-based solutions: limited scalability. Hundreds of potentially objectionable sites are launched each week. Therefore, as the database of websites grows, the business, educational institution, or library eventually will be forced to purchase additional storage for its database server.

Gateway Architecture

The gateway architecture improves on the server architecture by consolidating management and processing in a single device, thereby reducing capital and operational expenses. Some gateway-based solutions provide central policy management and enforcement, ensuring timely policy updates, consistent enforcement, and ultimately a lower total cost of ownership. Key vendors of gateway-based content filtering solutions include Fortinet and Symantec.

Unfortunately, traditional gateway solutions cost too much for budget-conscious small and mid-size businesses, schools, and libraries. Moreover, like server solutions, most of today's gateway solutions fail to address the need for scalability. Rather, the growth of the database is limited by appliance memory and hard disk space. Cost-conscious organizations need a solution they can count on for years to come without the looming need for hardware upgrades.

Criteria for the ideal gateway-based content filtering solution

Of the three architectures—client, server, and gateway—the gateway architecture comes closest to meeting the needs of small and mid-size businesses, educational institutes, and libraries. And yet traditional gateway solutions fail to satisfy these customers' scalability and affordability requirements.

The ideal gateway-based solution would meet the following criteria:

- Scalable - Support any size database without requiring hardware or memory upgrades—as would be possible if the database were provided by a hosted service
- Cost-effective - Eliminate the need to purchase, maintain, and administer multiple hardware devices, while maintaining a consistent cost model, even as the user base grows
- Manageable - Ensure policies are updated in a timely fashion and enforced consistently on all desktops, for all users
- Comprehensive - Store millions of URLs, IP addresses, and domains, while distinguishing among rating categories such as pornography and sex education
- Flexible - Allow administrators to create different policies and apply them to groups of users or individual departments customizing policy enforcement. Allow administrators to create custom lists of blocked and allowed sites, unblock specific sites or switch off filtering entirely.
- Secure - Support multiple authentication databases and provide User Level Access (ULA) control

SonicWALL Content Filtering Service

A next-generation gateway solution, SonicWALL Content Filtering Service (CFS) delivers the traditional advantages of gateway architectures as well as breakthrough scalability, low cost, and manageability. The solution enforces protection and productivity policies by using an innovative content rating and caching architecture that blocks objectionable and inappropriate Web content such as pornography, nudity, and violence. CFS also blocks access to unproductive sites such as sports, online shopping, gaming and streaming audio/video sites. SonicWALL CFS requires no additional hardware or software, enabling small and mid-size businesses, educational institution, and libraries to obtain enterprise-class features at an affordable price.

SonicWALL Content Filtering Service is available in three different versions:

	SonicWALL CFS	SonicWALL CFS Premium Business Edition	SonicWALL CFS Premium Gov/Ed Edition
Categories	12	50+	50+
Dynamic Rating	No	Yes	Yes
Multiple Group Policies	No	Yes	Yes
User Level Authentication	Yes	Yes	Yes
Bandwidth Management	Yes	Yes	Yes
Reporting Software	Optional	Optional	Included

Architectural components

SonicWALL CFS consists of two key architectural components. One is a comprehensive database of more than 4 million URLs, domains, and IP addresses and their ratings, hosted remotely and provided as a subscription-based service. The other component, the local SonicWALL appliance, caches acceptable URLs and their ratings, enforces the content filtering policy, and provides the management interface (see figure 2).

The SonicWALL CFS solution features support for reporting applications—SonicWALL Global Management System (GMS) software and SonicWALL ViewPoint™ reporting software. Both provide a comprehensive view of filtering that includes sites blocked, bandwidth usage, and top sites visited.

Dynamic database and rating system

At the heart of SonicWALL Content Filtering Service is a dynamic database and rating system that blocks objectionable and inappropriate Web content SonicWALL CFS provides ratings for over 50 categories, including pornography, drugs, gambling, chat/instant messaging and streaming media/MP3. The database resides in managed co-location facilities around the world, ensuring redundancy and optimum performance. It rates over 4 million URLs, IP addresses, and domains. Hundreds of sites are added daily, their ratings determined by human as well as artificial intelligence. Human rating enables the distinction between, say, a pornography and a sex education site.

For requested sites not already contained in the database, a Dynamic Rating Engine is used to categorize the site based on content in the Web page and the context of phrases. These dynamic ratings are then placed in the database for future reference by subsequent requests. A third-party service provider maintains the database so that scalability is not an issue for businesses, educational institutions, and libraries that subscribe to the service.

Because the database lists actual website URLs in addition to IP addresses and domains, users are protected from sites that utilize constantly changing dynamic IP addresses to evade blocking by traditional content filtering solutions.

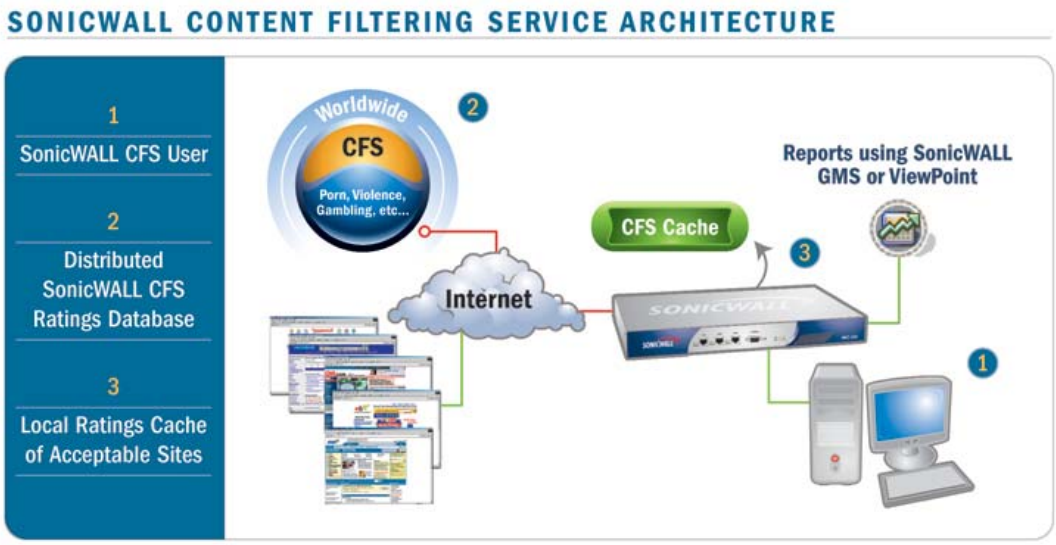

Figure 2.

SonicWALL appliance and caching architecture

The SonicWALL appliance, which resides at the network edge, provides two key solution functions. One is to cache acceptable URLs requested in the last 24 hours, as well as their ratings. That is, the first time that any user requests a website, the appliance requests a rating from the dynamic database. It caches acceptable URLs and their ratings for 24 hours. During this period, subsequent requests for the same website are processed locally on the SonicWALL appliance, enabling instantaneous response times and reducing off-net traffic. After 24 hours, the URL and its rating are removed from the cache. The next request for that URL is once again checked against the dynamic database, helping to ensure that the cached ratings remain synchronized with the dynamic database.

The other function of the SonicWALL appliance is to provide an interface for flexible policy management. Using the Content Filtering Service interface, the network administrator can flexibly define policy, selecting from numerous management options:

- Automatically block websites based on 50+ pre-defined categories. The network administrator can block any combination of categories, changing them on-demand as organizational policies change. When the administrator changes the policy, SonicWALL CFS immediately begins comparing the ratings in the cache against the new policy.
- Create multiple customized policies. SonicWALL CFS offers the ability to create multiple polices representing different filtering levels. This gives administrators the flexibility to enforce custom policies for groups of users on the network. For example, a school could create a filtering policy for teachers that contrasts with the student's policy. Similarly a network administrator could create different filtering policies for various different departments within a company.
- Create a customized list of blocked and allowed sites. SonicWALL CFS provides organizations with complete control over filtering by allowing network administrators to override policy for specific sites. An administrator can provide access to an individual site whose rating is disallowed, simply by categorizing it as an "allowed domain." For example, a high school social studies teacher might request access to sites bearing the "Hate/Racism" rating during a unit on the Civil Rights Era. Similarly, to block a site that does not fall into one of the 50+ categories, the administrator can identify it as a "blocked domain." For instance, a business might add a sports site to the list of blocked domains in the SonicWALL CFS cache to maintain employee productivity during a popular sports event.
- Assign filter-bypass privileges. Administrators can allow certain users and guests to bypass the filter policy. If an adult library patron asks for unfiltered Web access, for example, the

librarian can assign the patron a pre-defined username/password combination with bypass privileges or instantly create a custom account.

- Control Internet access by individuals through User Level Authentication (ULA). Administrators can support organizational goals for control and protection by specifying the users who will be granted Internet access, and their priority. Through ULA, the network administrator can require individuals to log on to the network with their username and password. ULA works with existing authentication databases such as RADIUS, LDAP (Lightweight Directory Access Protocol), and Active Directory.
- Block by time of day. Network administrators can choose the hours when content filtering applies. A school might filter certain content categories during school hours, for instance, and remove that filter after school. Similarly, a business might allow unfiltered Internet access to employees during lunch hours or after work. Time of Day filtering applies not only to websites, but also to blocking of cookies and keywords.
- Create custom acceptable use policy. Businesses, educational institutions, and libraries can create acceptable use policies for Internet access to ensure all network users are aware of their Internet access policies, strengthening their protection from legal liability. The administrator can check a box to display the acceptable use policy when users log on, and optionally include the same text in HR manuals.
- Block files that are automatically downloaded, such as Java applets, ActiveX scripts, and cookies. Because it blocks most automatically downloaded files, SonicWALL CFS provides more security than most firewalls. Java and ActiveX are often used for hacker attacks. Cookies create privacy concerns because they direct a remote Web browser to save small amounts of data on the local disk, and can be used to store preference information and track Web usage history.
- Manage Internet bandwidth. SonicWALL CFS gives network administrators the ability to dedicate bandwidth to high-priority traffic. This increases productivity by ensuring that critical traffic does not have to compete with high volumes of traffic with lesser priority.

INTEGRATED MANAGEMENT AND REPORTING

SonicWALL CFS can forward data directly into SonicWALL's award-winning Global Management System (GMS) or the SonicWALL ViewPoint™ reporting package to generate detailed reports on Internet usage and content filtering.

- SonicWALL GMS is the best choice for companies, educational institutions, and libraries with multiple SonicWALL appliances. GMS provides centralized reporting for multiple sites, as well as the tools to manage all SonicWALL appliances from one location.
- SonicWALL ViewPoint is a more cost-effective solution for organizations that have only one or two SonicWALL appliances. ViewPoint provides the same reports as GMS, without its management capabilities. (see figure 3).

SonicWALL GMS and SonicWALL ViewPoint can generate reports on:

- Bandwidth usage
- Top sites visited
- Top users
- Top sites per user
- Reporting by filter category
- Summary of firewall attacks and exceptions

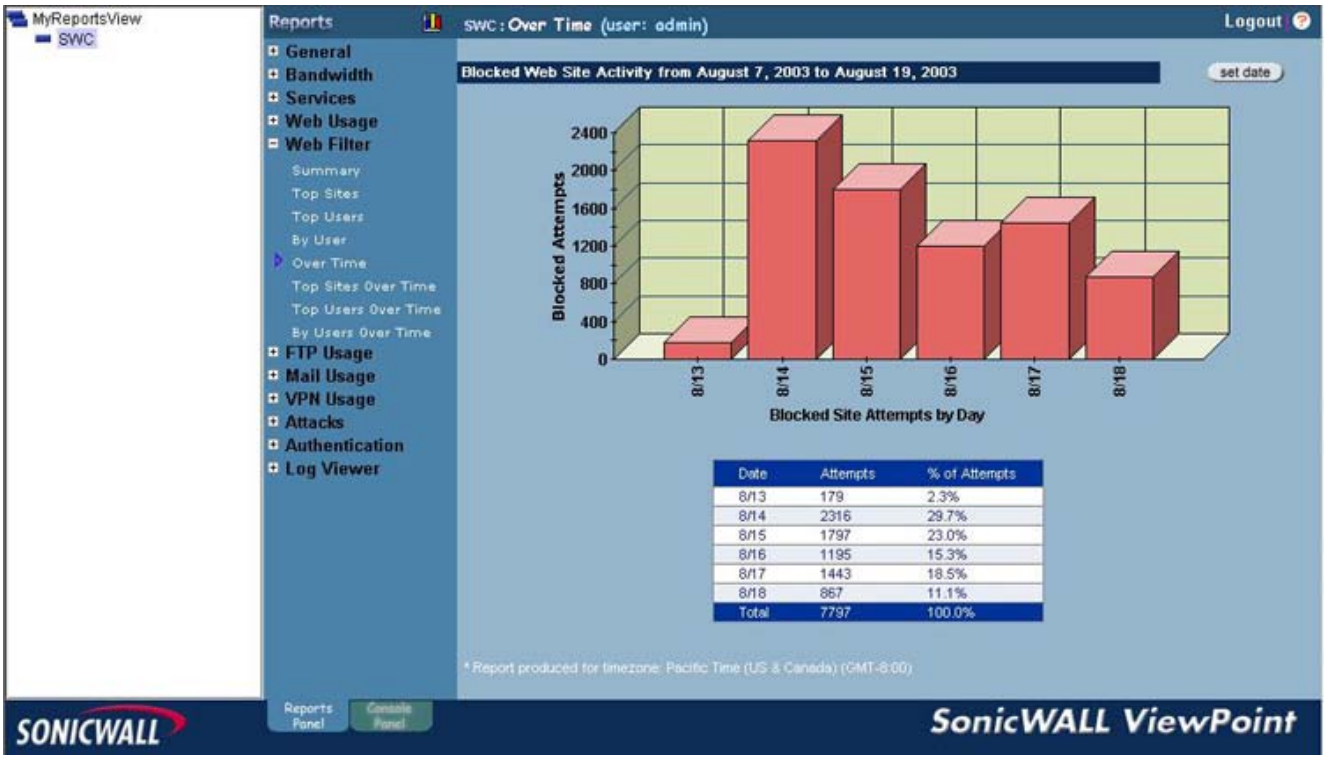

Figure 3.

Benefits of the SonicWALL solution

The following table summarizes the benefits of the SonicWALL CFS solution for businesses, educational institutions and libraries.

Requirement	SonicWALL CFS solution
Scalable	▪ Supports virtually any size database without requiring additional hardware or memory upgrades
Cost Effective	▪ Eliminates the costs of a separate database server, such as hardware, operating system, and ongoing maintenance costs ▪ Offers a fixed subscription rate for 50 or more users
Manageable	▪ Enforces content filtering policy on all desktops in the network from one central location, via the SonicWALL appliance ▪ Provides a simple Web-based interface that network administrators can access locally or remotely
Comprehensive	▪ Filters not only on URLs, but also domains and IP addresses ▪ Provides a database of over 4 million entries, updated continually ▪ Uses a combination of artificial and human intelligence to determine which websites belong in each category ▪ Uses a dynamic rating engine to categorize sites not already in the database
Flexible	▪ Allows administrators to block up to 50 different categories ▪ Allows administrators to create custom policies for groups of users or individual departments ▪ Enables administrators to override blocked and allowed sites to create a customized list of sites that can be visited ▪ Allows administrators to assign “bypass filter” privileges ▪ Supports filtering by time of day
Secure	▪ Enforces User Level Authentication (ULA), requiring users to log on to the network with a username and password ▪ Blocks automatically downloaded files such as Java, ActiveX, and cookies
Efficient	▪ Caches acceptable URL ratings locally on the SonicWALL appliance, eliminating potential delays when users request frequently-visited sites ▪ Allows administrators to prioritize traffic by user and to dedicate bandwidth to higher priority Web traffic

Usage scenario

When an organization begins using SonicWALL CFS, users are unaware that it is operating in the background. Following are the steps that occur behind the scenes.

1. User requests URL or IP address.

2. SonicWALL appliance checks its local cache for the site and associated rating.
   - If the URL/IP address is present and its rating complies with the policy, the request is processed. If the address is present and does not comply with the policy, the user is sent a message that the site is not allowed. The administrator can use the management interface to select the default message or write a custom message.
   - If the URL or IP address is not in the cache, the appliance queries the closest of multiple databases located around the world.

3. The ratings database looks up the rating associated with the URL/IP address. Ratings are associated with specific categories of content, such as Alcohol/Tobacco or Violence. Some sites can have multiple ratings, such as Nudism and Pornography.

4. The database returns the rating to the SonicWALL appliance, where it is compared against the current filtering policy.
   - If the rating complies with policy, then the page is displayed and the rating is cached on the SonicWALL appliance.
   - If the rating does not comply with policy, the user is informed that the content may be objectionable and therefore is not allowed.

Conclusion

Content filtering is essential for businesses seeking to improve productivity and avoid legal liability. Educational institutions and libraries have both a fiduciary and legal responsibility to install filtering solutions in order to protect their students and young patrons from objectionable websites. Those that neglect to comply with regulations risk losing federal funding for technology programs.

The three architectures for content filtering solutions—client, server, and gateway—vary in their effectiveness, manageability, and affordability. Gateway solutions best meet the needs of educational institutions, libraries, and small businesses, but until now have been prohibitively expensive for cost-conscious organizations.

SonicWALL CFS defines a new type of gateway solution. With its unique caching and website rating architecture, the SonicWALL CFS solution delivers unsurpassed scalability, control, and flexibility. And because it requires only one local device and is available as a subscription-based service, SonicWALL CFS makes comprehensive content filtering available at a breakthrough price point.

To learn more about SonicWALL Content Filtering Service (CFS), visit: http://www.sonicwall.com/products/cfs.html.

To view a Flash demo of the SonicWALL CFS solution, visit http://www.sonicwall.com/img/cfs-flash/cfs.html.