



# Comparing URL Filtering Databases

**A Broadband-Testing Report**

First published June 2005 (V3.0)

Published by Broadband-Testing  
La Calade, 11700 Moux, Aude, France

Tel : +33 (0)4 68 43 99 70  
Fax : +33 (0)4 68 43 99 71  
E-mail : [info@broadband-testing.co.uk](mailto:info@broadband-testing.co.uk)  
Internet : <http://www.broadband-testing.co.uk>

©2005 Broadband-Testing

All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the authors.

Please note that access to or use of this Report is conditioned on the following:

1. The information in this Report is subject to change by Broadband-Testing without notice.
2. The information in this Report, at publication date, is believed by Broadband-Testing to be accurate and reliable, but is not guaranteed. All use of and reliance on this Report are at your sole risk. Broadband-Testing is not liable or responsible for any damages, losses or expenses arising from any error or omission in this Report.
3. *NO WARRANTIES, EXPRESS OR IMPLIED ARE GIVEN BY Broadband-Testing. ALL IMPLIED WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT ARE DISCLAIMED AND EXCLUDED BY Broadband-Testing. IN NO EVENT SHALL Broadband-Testing BE LIABLE FOR ANY CONSEQUENTIAL, INCIDENTAL OR INDIRECT DAMAGES, OR FOR ANY LOSS OF PROFIT, REVENUE, DATA, COMPUTER PROGRAMS, OR OTHER ASSETS, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.*
4. This Report does not constitute an endorsement, recommendation or guarantee of any of the products (hardware or software) tested or the hardware and software used in testing the products. The testing does not guarantee that there are no errors or defects in the products, or that the products will meet your expectations, requirements, needs or specifications, or that they will operate without interruption.
5. This Report does not imply any endorsement, sponsorship, affiliation or verification by or with any companies mentioned in this report.
6. All trademarks, service marks, and trade names used in this Report are the trademarks, service marks, and trade names of their respective owners, and no endorsement of, sponsorship of, affiliation with, or involvement in, any of the testing, this Report or Broadband-Testing is implied, nor should it be inferred.

# TABLE OF CONTENTS

---

<b>EXECUTIVE SUMMARY .....</b>	<b>1</b>
<b>URL FILTERING DATABASES – WHAT ARE THEY? .....</b>	<b>2</b>
How to Compare and Contrast URL Filtering Databases .....	2
<b>THE URL FILTER DATABASE TESTS .....</b>	<b>5</b>
Sourcing the Test URL List .....	5
The Difficulties in Accurately Comparing Categories.....	6
Our Testbed Defined.....	7
<b>THE URL FILTERING DATABASE COVERAGE RESULTS.....</b>	<b>9</b>
<b>SUMMARY AND CONCLUSION.....</b>	<b>12</b>
<b>APPENDIX: THE TEST EQUIPMENT DETAILS .....</b>	<b>13</b>

## TABLE OF FIGURES

Figure 1 – Win Web Crawler.....	6
Figure 2 – Directly Comparable Categories Example .....	7
Figure 3 – Our Test bed Simulating Real-World Use.....	8
Figure 4 – The Blue Coat ProxySG Management Screen – Setting URL Database Parameters .....	8
Figure 5 – URL Filter Database Coverage Results: Adult Oriented .....	9
Figure 6 – URL Filter Database Coverage Results: Leisure and Hobbies .....	10
Figure 7 – URL Filter Database Coverage Results: Other .....	10
Figure 8 – URL Filter Database Coverage Results: Overall Coverage .....	11
Figure 9 – URL Filter Database Coverage Results: Overall & Number of Firsts in Category .....	12
Figure 10 – Spirent Avalanche 2500.....	13
Figure 11 – Creating An Avalanche 2500 Test.....	14
Figure 12 – Creating An Avalanche 2500 URL Filter Database Test.....	15

## Broadband-Testing

---

Broadband-Testing is Europe's foremost independent network testing facility and consultancy organisation for broadband and network infrastructure products.

Based in the south of France, Broadband-Testing offers extensive labs, demo and conference facilities. From this base, Broadband-Testing provides a range of specialist IT, networking and development services to vendors and end-user organisations throughout Europe, SEAP and the United States.

Broadband-Testing is an associate of the following:

- *NSS Network Testing Laboratories (specialising in security product testing)*
- *Broadband Vantage (broadband consultancy group)*
- *Limbo Creatives (bespoke software development)*

**Broadband-Testing Laboratories** are available to vendors and end-users for fully independent testing of networking, communications and security hardware and software.

**Broadband-Testing Laboratories** operates an **Approval** scheme which enables products to be short-listed for purchase by end-users, based on their successful approval.

Output from the labs, including detailed research reports, articles and white papers on the latest network-related technologies, are made available free of charge on our web site at <http://www.broadband-testing.co.uk>

The conference centre in Moux in the south of France is the ideal location for sales training, general seminars and product launches, and Broadband-Testing can also provide technical writing services for sales, marketing and technical documentation, as well as documentation and test-house facilities for product development.

**Broadband-Testing Consultancy Services** offers a range of network consultancy services including network design, strategy planning, Internet connectivity and product development assistance.



## EXECUTIVE SUMMARY

---

Following a number of tests of security appliances featuring URL filtering capabilities – among these Blue Coat, Equinet, Symantec, SonicWall and ServGate – it occurred to us at Broadband-Testing that the URL filtering databases themselves should be compared to see which offers the best range of categories, URL coverage – and most relevant URLs.

No URL list can include every web page on the Internet, so customers need to investigate the categories offered and the likely coverage each list will offer them. This report can help in that quest.

For this test, we focused on assessing the usefulness of the market-leading URL filtering databases for URL filtering devices, including Websense, Blue Coat Web Filter (BCWF), SmartFilter from Secure Computing, SurfControl, and WebWasher.

We found that not all URL lists are equal. The basic concept of URL lists is always the same – source the URLs to be blacklisted, categorize them, and update them. However, the net results, in terms of which URLs make the database, how they are categorized, and how new entries are treated, can differ markedly.

Our results show that the cost or market position of the URL filtering database/supplier is not a significant guide to its relative ability to categorize successfully. There were some significant showings from some of the lesser-known databases, ahead of their more prestigious rivals' products.

The differences can largely be attributed to the way each database supplier goes about the process of finding and updating URLs in their list. Some clearly are more efficient than others...

### **Aims of the Test**

This test is not designed to investigate the merits of URL filtering itself, but to take one security application platform that can support multiple URL filtering databases and compare the databases themselves.

Likewise, “accuracy” per se, has to be taken as read – that suppliers will be accurate with the list content. More important for us is to see suppliers coverage – how they categorize that content and how they deal with additions and updates.

The results are not intended to be a “one size fits all” answer to every user’s requirements – we appreciate that different types of companies will be interested in very different types of web content. However, we have made the tests broad enough so that the results are indicative of the overall strengths of each database supplier, to act as a guide to which databases you should consider ahead of others.

## URL FILTERING DATABASES – WHAT ARE THEY?

---

In order to prevent unwanted access to specific areas of the Internet – such as pornographic or gambling sites – many security appliance vendors have introduced URL filtering features, allowing these kinds of sites to be blacklisted and made unavailable to network users.

These appliances make use of URL filtering databases – regularly updated lists of categorized URLs. There are many different companies who sell URL filtering databases to help inform and control users' access to the Internet. These databases can be loaded onto proxy servers, firewalls and other network infrastructure devices or provided as a service across the Internet. The databases are split into multiple categories, such as sex and sport but – as we'll see – both the number of categories and the precise definition of each varies from supplier to supplier. There are no standards currently in place for this type of filtering.

Database providers deliver differing levels of sophistication depending on the platform the database is loaded onto. This might range from simple blocking of all access to a particular category at one extreme, to a more elaborate and flexible mix of access rules, based on user, group, time of day and other attributes. The best examples provide a rule-based setup, whereby you can configure a series of nested rules using And/Not/Or type qualifiers, so that URL filtering can be fine-tuned down to the individual user. An example would be allowing access to sports sites during specific hours, but denying access to streaming content hosted on those sports sites at the same time.

URL filtering databases are provided by a number of specialist suppliers, the current market leader being Websense, and may be licensed either directly or via the security appliance vendor. The latter may offer a range of different database supplier options, or specialise in just one – depending on the database and/or security appliance provider itself.

### How to Compare and Contrast URL Filtering Databases

An important point to make immediately here is that no two URL filtering databases are the same, despite the suppliers' claims of offering a nearly identical service. Not only do they each gather new sites in different ways and have different teams trying to categorize them, the categories themselves are not always comparable.

Each company might claim a total number of sites in their database, but – as a user - just going for the highest total as the “winning” feature may not be the best way to evaluate a database. Bear in mind that the Internet is constantly growing and changing and if that total is boosted by many old sites that are no longer active, or by sites that are in categories that are of little interest, then they are of no use. So the overall total is not necessarily the best criterion to judge by.

Therefore, as a prospective customer, you might want to look at the categories that the database provider delivers, in order to decide which are most useful to your company. Make sure the database lists the categories that need to be controlled, also checking that the list provides the depth and flexibility required. For example, if you want to block access to illegal drugs

but one vendor puts both illegal and legal drugs in the same category, you need to check whether that is too broad a definition for you.

Another set of issues to check is how far down within a URL the categorization is done. Does the database simply list top-level URLs, or does it create further sub-categories? What about a media site that has specific areas for news, sports, entertainment and such? Do you want to allow one area to be accessed but not another *and* does the database correctly identify each area to allow you to do this?

All databases are updated *daily* with new categorizations. Another factor to consider: How is the list populated? Is all categorization performed by individuals, or is there a computer-driven initial categorization? Perhaps this doesn't matter as long as the data is updated, but bear in mind that people **do** cost more than computers. A possible trade-off here is accuracy – perhaps using computers for simple categorization and using humans to categorize where the computer-driven program is unsure is the most appropriate combination.

Where are new URLs sourced? Is the database company constantly searching for new URLs, doing searches or watching spam links to harvest or does it bring in uncategorized URLs from its customers? There's little point in buying a database if that database is full of sites that are irrelevant to you, the customer. Systems that automatically send all uncategorized sites for addition will quickly become more effective once installed.

It is also worth considering whether the list has *too many* categories. If the reason for buying a list is to remove access to "bad" content, do you really need a list that includes various business categorizations such as real-estate, computing, and government, since the vendor is taking time (and in so doing spending your money) evaluating, and adding to, these categories?

Websense further complicates the decision by splitting its categories into four groups – the main database and then three further databases. As a customer, you have to buy each set you need separately. At the same time, it can be argued that what (to some) is over-elaboration is ideal for others. For example, Websense is also the only database to offer separate pro and anti-abortion categories though, at the same time it doesn't offer children's site categories whereas others do.

There are also certain categories offered by some database suppliers that are – at face value – overly broad; for example, 'image', 'content delivery networks,' and 'dynamic content'. These are likely to cover so many different types of content that setting a rule based on them is pretty much ineffective, or will simply give the wrong results. Some might claim that these categorizations are included primarily to reduce the number of URLs listed as un-categorized.

Each of the companies also makes claims on language and country-coverage. Some years ago, many of these lists were mainly populated with content from their country-of-origin (Websense and SmartFilter from the USA, WebWasher from Germany etc.) and ALSI is still mainly a Japanese database, though they are now doing their best to cover all countries.

Again, the customer needs to decide what countries are important (and not-important) for their users before deciding on what to implement. In the end, the best way to make a decision is to install multiple systems for a test

period simply to log user access and see which list provides the best categorization of the traffic.

This report, however, can provide some short cuts to the decision-making process, in so much as this is precisely what we've done here – to line up the different databases on a single security appliance and compare them directly.

### Coverage – a Discussion

A list made up of content you are not interested in is of no benefit to you or your company.

As an example, if someone provided a list of 10,000 sporting URLs all from Germany – that would be of use to German companies or those with a lot of German speaking users – but of little interest to a company who is wholly-based in Spain. If you take this approach further, a list of sites that is not accurately kept up to date is equally of little use.

Many pieces of spam are linking to new web sites all the time. Of course, that's exactly what happens with phishing attacks – an email is sent to millions of people and the phisher makes money on the first few victims who click on their site and pass over their credentials. The site doesn't need to be up for very long, perhaps less than an hour, to catch enough victims for the phishers to make their money. So, a list that is updated two days later to block the site is of no help at all. Bearing that in mind, if a URL filtering database supplier is surfing the web looking for new sites to add, how do you know that supplier is looking at the sites your users are accessing? Each supplier has a way for users to check categorization of individual sites and recommend new sites to add. Some of them also provide a service where uncategorized sites can be automatically sent for possible inclusion, while others make this optional; some systems do this by default. For example, the BCWF list automatically receives new URLs each day from its customers (receiving details of more than 50 million surfed URLs a day); the uncategorized URLs among these are prioritised by popularity so that the most popular are rated and added into the database first.

Therefore, a customer is most likely to receive ratings on sites that its users are attempting to visit, rather than random sites that have been found by the vendor itself. This is an approach we would like to see adopted by all the URL list suppliers.

## THE URL FILTER DATABASE TESTS

---

So just how do you test a URL filtering database list?

This is not simply a case of “take test device, put it in between client and server, test traffic generators, and monitor the results”, but rather requires real-world use, closely monitored. Even then, can you test for accuracy? If a site is mis-categorized, a user may get access to content that the organization wants blocked or not get access to a site needed for legitimate purposes. Sadly, testing for accuracy for a statistically significant number of sites is impossible. Further, every vendor has a public web site where customers can check on and recommend categorizations – so we have to accept that all these lists aim to be as accurate as possible.

However, if we can provide a large enough list – we can test for coverage. If we run a test and one database has twice as many URLs in a category as another list, then we can assume that the first list is likely to have categorized more of this type of content. Hence, this is precisely what we did – created a list that couldn't possibly form the basis of any single database supplier's own list, and tested accordingly.

### Sourcing the Test URL List

There are many public sources of URL lists for testing and – not surprisingly – all have their benefits and drawbacks.

On the Internet are various sources of URL lists, such as Squidguard. The problem is that these lists may already be used by one vendor and not by another and by their public-domain nature could easily be full of old, dead sites – a test of databases against sites that are no longer in use would penalise those databases that keep cleaning their database and help lists that are not so up to date.

A second option was to use browsing logs obtained from ISPs or enterprises. The problems here are two-fold. Firstly, there may have already been a filtering system deployed, thereby skewing the results and, secondly, any ISP or enterprise log is going to be filled with URLs that are popular with its own users and naturally skewed by their language, interests, etc., which may not be broad enough to satisfy our test requirements here.

A third option was to try to generate our own lists, in order to guarantee the quantity and depth of URLs required to make this a fair test. We chose to use the application Win Web Crawler to generate our lists and produced a list of over 2,000,000 URLs (both top-level and subdirectories). See [www.winwebcrawler.com](http://www.winwebcrawler.com) for more details. Win Web Crawler asks the user for a search term and then uses that term to search across multiple search engines. It produces a list of URLs in a file. Win Web Crawler searches through 148 search engines from over 30 different countries, is able to follow links from one page to another and use keyword matching within a URL. In addition, it can be set not to follow more than a specified number of pages on a single-site to guarantee as many unique results as possible.

When compiling our list, we chose keywords in multiple languages, common words, single letters, words likely to be in categorized sites and common surnames in various countries to try to broaden the scope of the list as much as possible.

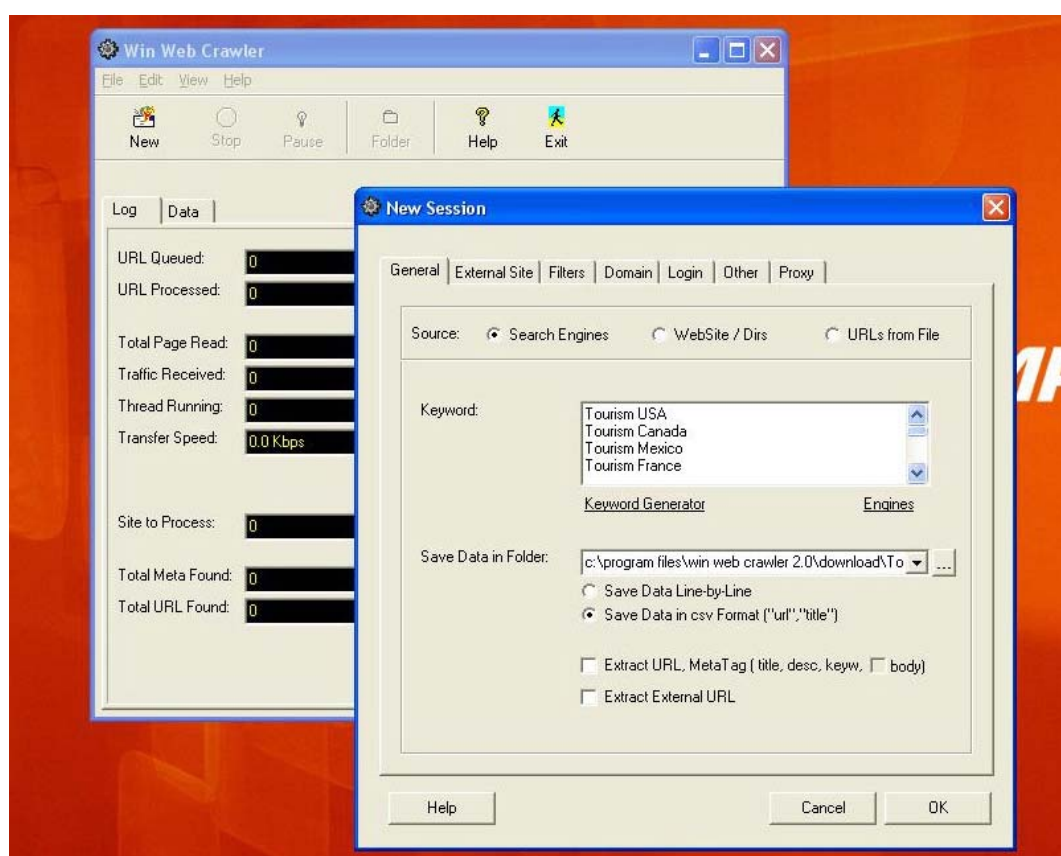


Figure 1 – Win Web Crawler

However, any list is naturally biased by the initial terms chosen and this list is angled against sites that do not use the Roman alphabet.

Here are some examples of terms we used:

1, 2, 888, a, aa, b, blank, cannabis, cart, date, drink, free, frei, Garcia, gun, Johansson, Kim, Lee, libero, libre, livre, Montag, me, Nagy, Nguyen, Novak, now, poker, set, shoot, Silva, Suzuki, shop, teen, tienda, Wang, xxx, Yang, you, thumb.

### The Difficulties in Accurately Comparing Categories

One of the issues we faced in compiling results was comparing categories between different database providers. The key issue is terminology – what one supplier labels “Adult/Mature Content”, another calls “Pornography” and yet another calls “Adult/Sexually Explicit”.

Our goal was to ensure that we were able to perform direct comparisons between the different database/category combinations to get accurate results. To that end, we attempted to group similar categories together.

For example, the table below shows how we found some directly comparable categories:

SurfControl	SmartFilter	Blue Coat Web Filter	WebWasher
Adult/Sexually Explicit	Pornography Nudity	Adult/Mature Content Pornography Nudity	Pornography Erotic / Sex
Spyware	Spyware Malicious Sites	Spyware/Malware Sources Spyware Effects	Computer Crime / Spyware
Glamour & Intimate Apparel	Provocative Attire	Intimate Apparel/Swimsuit	Swimwear / Lingerie / Nudity
Gambling	Gambling	Gambling	Gambling / Lottery
Violence Intolerance & Hate	Violence	Violence/Hate/Racism	Hate / Discrimination
Weapons	Weapons	Weapons	Weapons
Games	Games	Games	Computer Games
Sports	Sports	Sports/Recreation/Hobbies	Sports
Personals & Dating	Dating/Social	Personals/Dating	Dating / Relationships
Travel	Travel	Travel	Travel
Chat	Chat	Chat/Instant Messaging	Chat
Alcohol & Tobacco	Alcohol Tobacco	Alcohol/Tobacco	Alcohol Tobacco
Streaming Media	Streaming Media Media Downloads	Streaming Media/MP3	Music / Webradio
News	General News	News/Media	News / Magazines
Shopping	Shopping/Merchandizing Auction	Shopping Auctions	Shopping Auctions / Classified Ads
Web-based E-mail	Web Mail	Email	Web Mail
Religion	Religion and Ideology	Religion Cult/Occult	Religion Sects
Job Search & Career Development	Job Search	Job Search/Careers	Human Resources

*Figure 2 – Directly Comparable Categories Example*

## Our Testbed Defined

Having sourced our URL list using Win Web Crawler, we then created a test bed which enabled us to pass this list through each URL filtering database under test, under exactly the same test conditions.

In order to do this, we used a Blue Coat ProxySG 400 appliance as our base test device, enabling just its URL filtering and reporting capabilities. Key to this device is its ability to support several different databases, so it was the ideal test platform for us.

To generate the URL requests we used two different tools. The first was a Spirent Avalanche web traffic generator. With this we were able to feed in the entire list of URLs, point the Avalanche at the Blue Coat ProxySG, which filtered the URLs, while forwarding any DNS resolution requests to the Internet, via an ADSL connection. Here, however, we experienced a problem with the Websense database, since it requested DNS resolutions for every URL and this caused time-out problems.

## Broadband-Testing – URL Filter Database Test

To resolve this we used WGET – a public domain, command-line based utility which GETs a list of URLs directly or via a proxy; in this case, via the Blue Coat ProxySG. However, running from an Intel Pentium M-based laptop, it took several days to complete the URL run with the Websense database; the others completed in a fraction of that time. This should be borne in mind for real-world use.

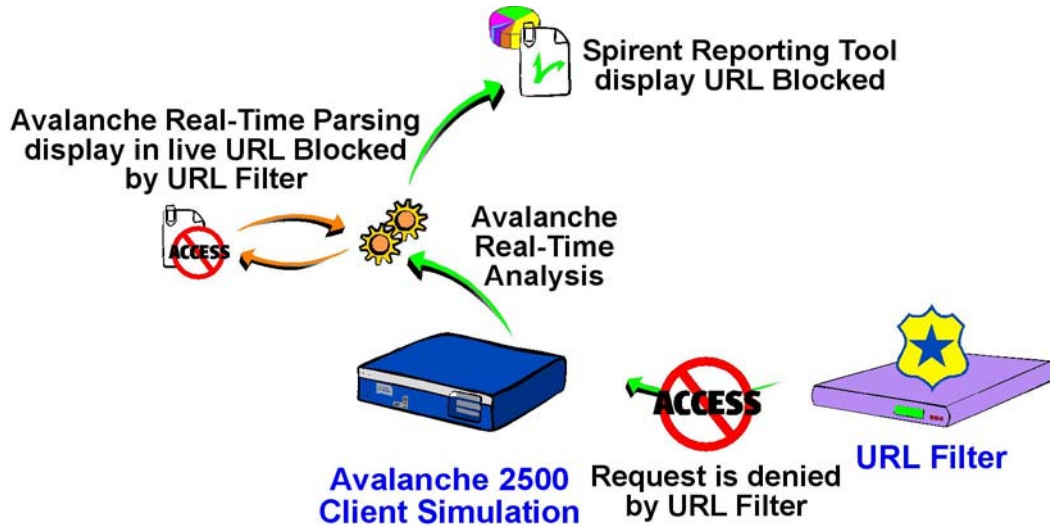


Figure 3 – Our Test bed Simulating Real-World Use

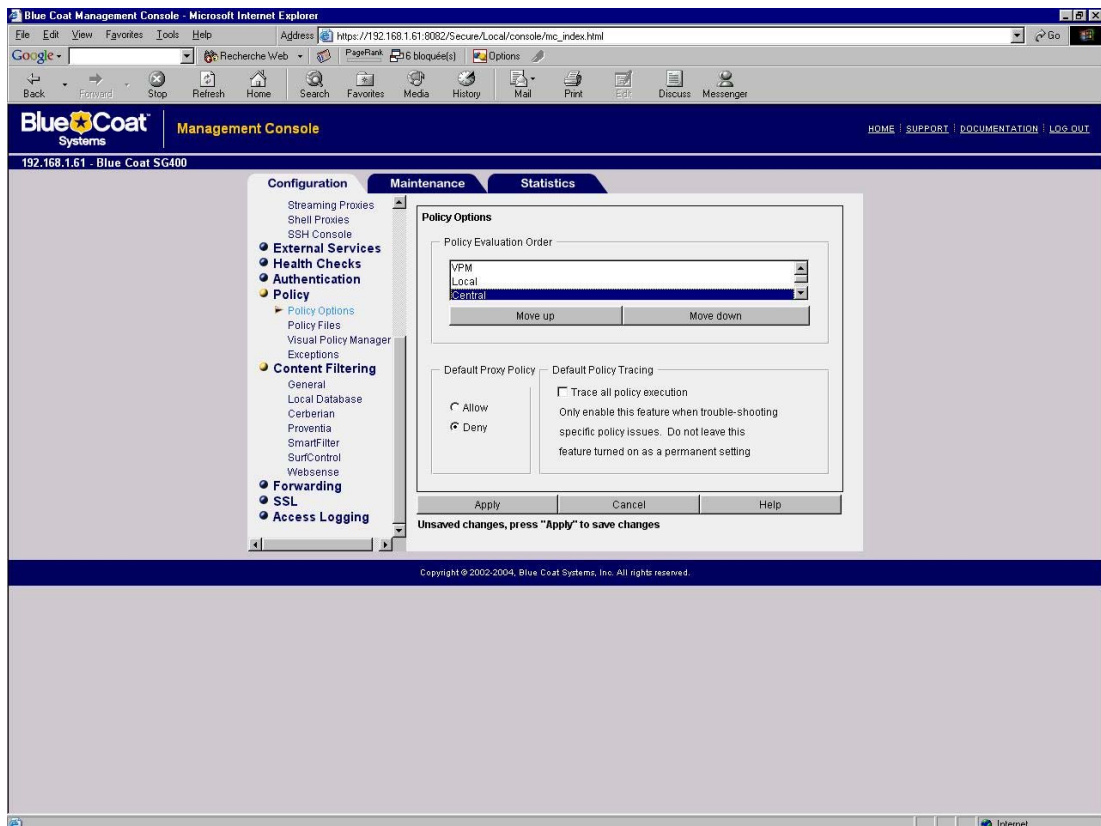


Figure 4 – The Blue Coat ProxySG Management Screen – Setting URL Database Parameters

## THE URL FILTERING DATABASE COVERAGE RESULTS

Our filter database coverage compares database results from Websense, BCWF (Blue Coat Web Filter), BCWF with Dynamic Real-Time Rating™ (DRTR – see below for explanation), SmartFilter from Secure Computing, SurfControl and WebWasher.

Taking our directly comparable categories, here are the results, starting with a separation – for ease of viewing – into three sub-categories, that we could loosely call adult-oriented, leisure-hobbies and other. In this first batch of categories, we see that BCWF and especially BCWF with DRTR databases score particularly well on Adult and Spyware categories while SmartFilter scores highly on Glamour and Intolerance. WebWasher excels on Weapons but disappoints in some of the other categories. SurfControl and Websense generally show mediocre results.

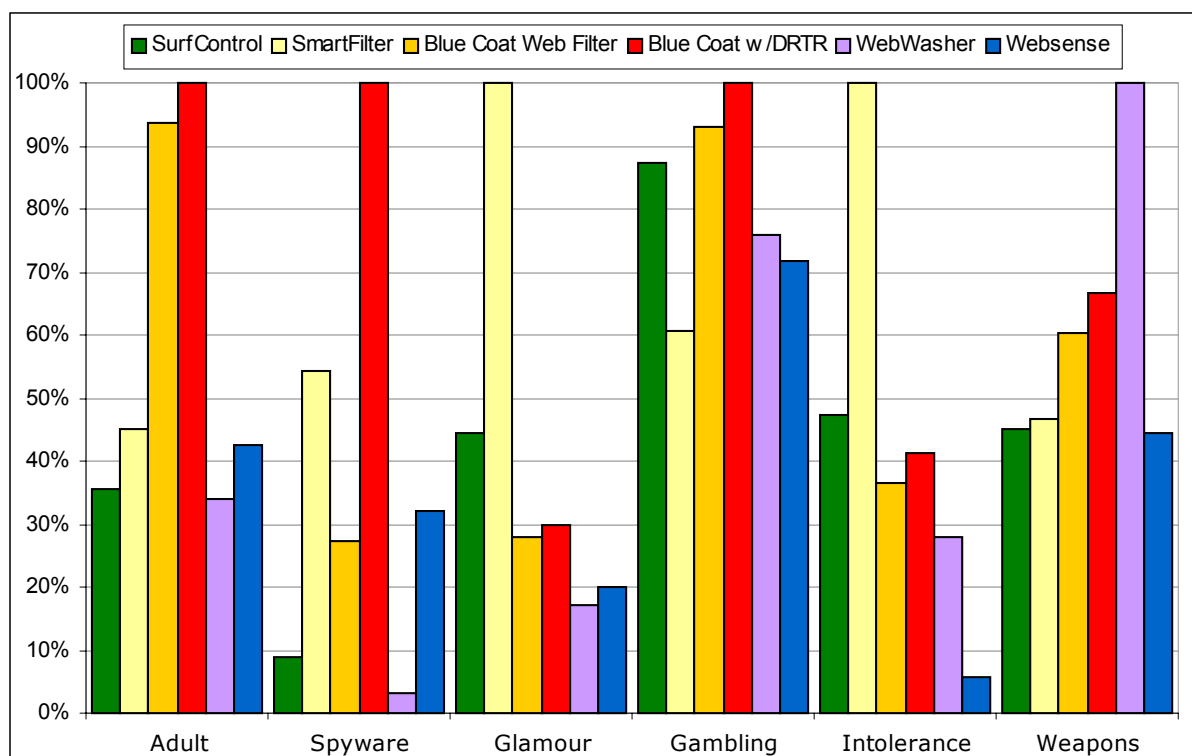


Figure 5 – URL Filter Database Coverage Results: Adult Oriented

Both SurfControl and Websense fare better in some of the leisure and hobbies categories, the latter notably with travel, but overall – again – BCWF with DRTR shows the best performance. SmartFilter scores highly on the Chat category and also shows well in Games. WebWasher gets an outstanding score in Dating. WebWasher also scores well in many of the “Other” categories, but still BCWF with DRTR comes out top more often than not (see after graphs for explanation of how Blue Coat’s DRTR technology works).

**Broadband-Testing – URL Filter Database Test**

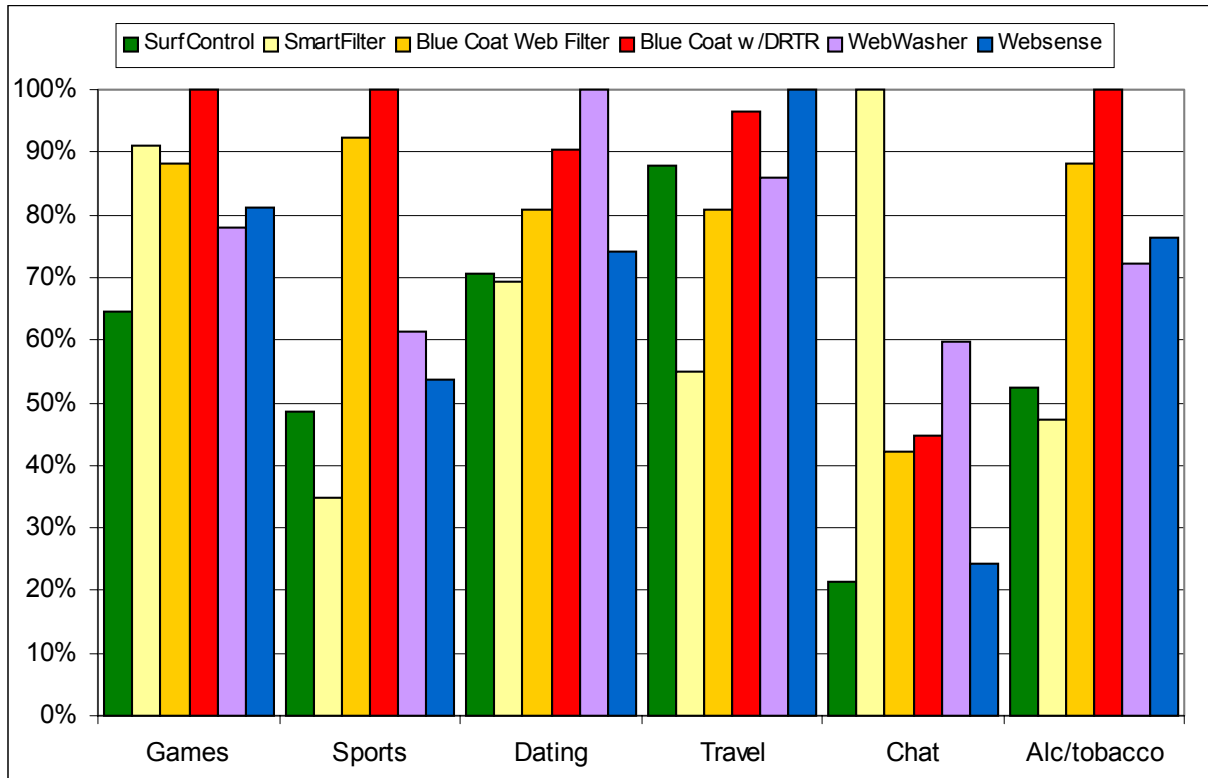


Figure 6 – URL Filter Database Coverage Results: Leisure and Hobbies

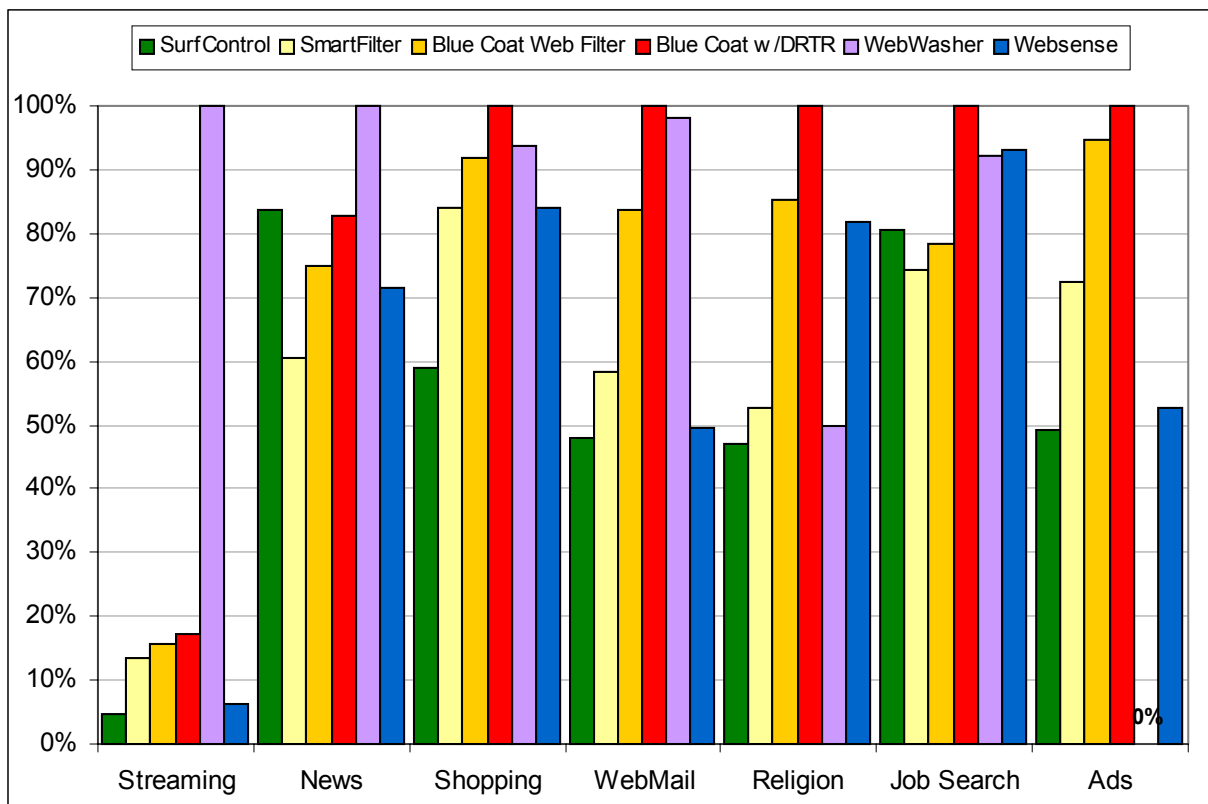


Figure 7 – URL Filter Database Coverage Results: Other

**Blue Coat’s Dynamic Real-Time Rating (DRTR) Explained Fully**

Since the Internet changes constantly, no ratings service can ever hope to categorize every Web page.

So when users request a new URL that has not been rated in the BCWF database, BCWF optionally uses what it claims to be a unique Dynamic Real-Time Rating (DRTR) technology to retrieve the page from its host server and analyse its content.

The DRTR service first determines the language of the web page, and then assesses its categorization. This categorization is based on a variety of information, including the words on the page, the context for each word, and the links on the page. New web sites for which DRTR is able to determine a rating with high confidence are categorized in real-time and immediately blocked or allowed, based on the user’s access policy. In those instances where DRTR cannot categorize the requested page with high confidence, it immediately returns an “unrated” response. All web pages, independent of the results from real-time categorization, are further processed by the BCWF background rating processes and added to the BCWF ratings database based on frequency of request.

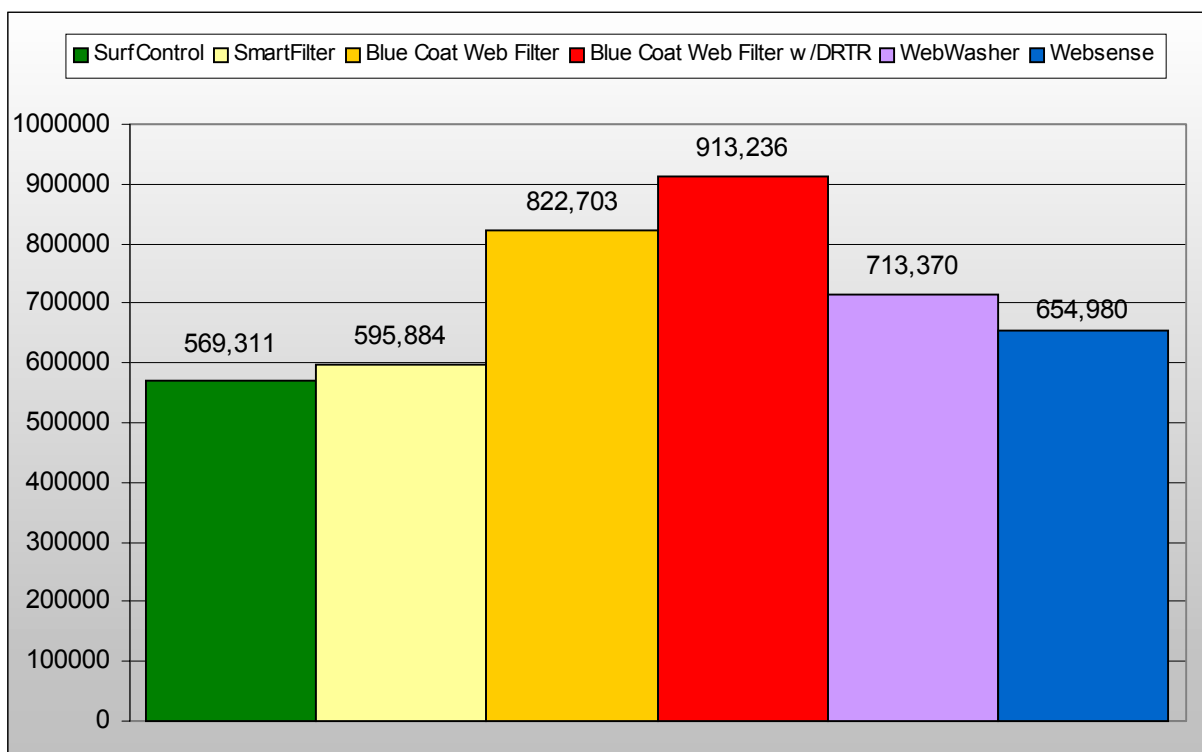


Figure 8 – URL Filter Database Coverage Results: Overall Coverage

Overall then, the BCWF database, notably with DRTR, scored highest in our study, followed by WebWasher, Websense, SmartFilter, and SurfControl, respectively. This was due to BCWF’s consistently high scoring across the broad range of categories selected for comparison. This result is reflected in the graph below, BCWF with DRTR gaining the highest

average mark, 5.4/6 and recording by far the highest number of firsts in category, 11, of any of the databases in the test. Even if DRTR is turned off (and it's optional), BCWF scored the highest number of firsts; 6.

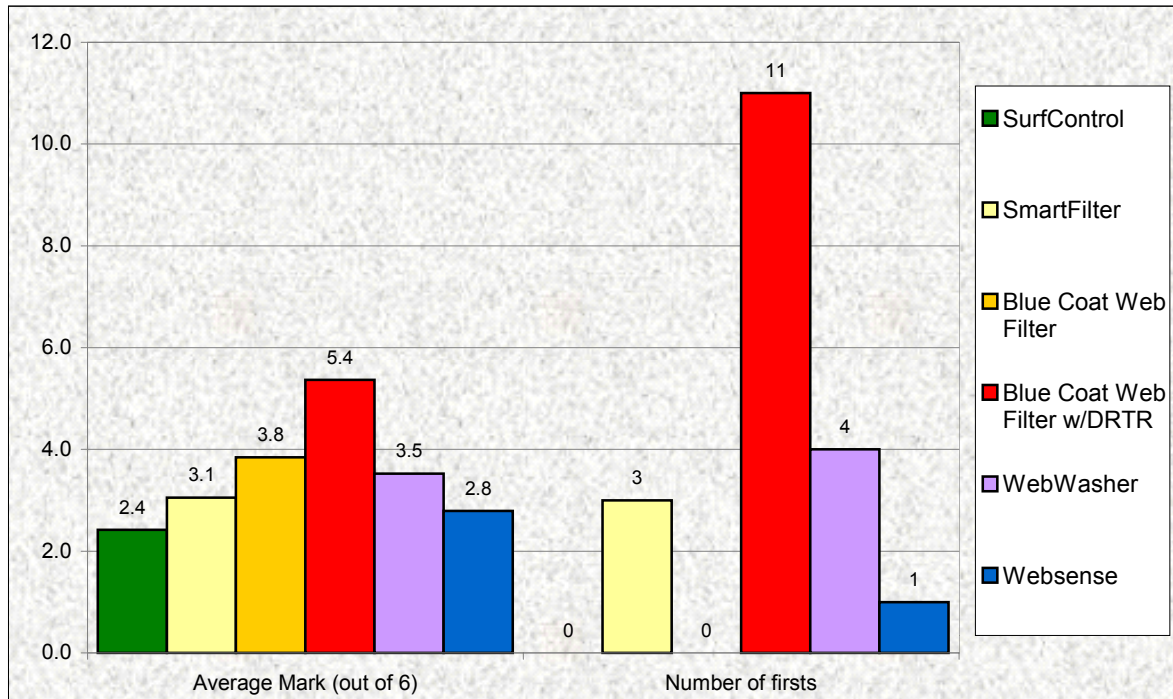


Figure 9 – URL Filter Database Coverage Results: Overall & Number of Firsts in Category

## **SUMMARY AND CONCLUSION**

---

While comparing web URL filter database coverage is not a purely scientific, squeaky-clean method for testing, we believe here that we have achieved a true reflection of each database's coverage capabilities, relative to each of the others.

As we see from the overall results position, Blue Coat Web Filter (BCWF), especially with DRTR enabled, came out the clear winner, notching up 11 category firsts along the way. SmartFilter showed well with three category firsts, and WebWasher slightly better still, with four category firsts and an overall third place behind the two Blue Coat Web Filter databases. Websense, somewhat disappointingly for a market leader, came out next to last, with just one category first – Travel – while SurfControl brings up the rear with no category firsts at all.

So, our Approved award goes to Blue Coat Web Filter with DRTR, for an outstanding overall performance in our coverage tests.



## APPENDIX: THE TEST EQUIPMENT DETAILS

Internet architectures are becoming increasingly complex.

Whether you're building network equipment or providing a service, you must deliver consistent performance under all conditions. Until now, capacity assessment at high-loads has been a costly and complex process. For this reason, Spirent Communications introduced the Avalanche 2500 and Reflector 2500 appliances to assist with the challenge. At Broadband-Testing we have taken these web application simulation and planning products and integrated them into our test-bed simulating real-life Internet conditions; those that the average user experiences daily.



Figure 10 – Spirent Avalanche 2500

Avalanche 2500 is described by Spirent as a capacity assessment product that challenges any computing infrastructure or network device to stand up to the real world load and complexity of the Internet or intranets. The system determines the architectural effectiveness, points of failure, and the performance capabilities of a network or system. Using Avalanche 2500 to generate Internet user traffic and Reflector 2500 to emulate large clusters of data servers, you can simulate even the world's largest customer environments. The system provides invaluable information about a site's architectural effectiveness, points of failure, modes of performance degradation, robustness under critical load, and potential performance bottlenecks. It is able to set up, transfer data over, and tear down connections at very high rates - all while handling cookies, IP masquerading for large numbers of addresses, and traversing tens of thousands of URLs.

Avalanche 2500 initiates and maintains more than a million concurrent connections, each appearing to come from a different IP address. This allows realistic and accurate capacity assessment of routers, firewalls, load-balancing switches, and Web, application, and database servers. It helps identify potential bottlenecks from the router connection all the way to the database. This accuracy is especially critical for gauging Layer 4-7 performance. The ability to additionally simulate error conditions such as HTTP aborts, packet loss, and TCP/IP stack idiosyncrasies can help

anticipate-and avoid-significant and previously unknown impacts on performance.

To enable more accurate load simulations across multi-tiered Web site architectures, the system also supports extremely realistic user modelling behaviours such as think times, click stream, and HTTP aborts that cause Web servers to terminate connections while back-end application servers continue to process requests. Configuring in this way is simple as both Avalanche 2500 and Reflector 2500 directly from a desktop browser to set up tests, review feedback in real time, and easily reconfigure test parameters.

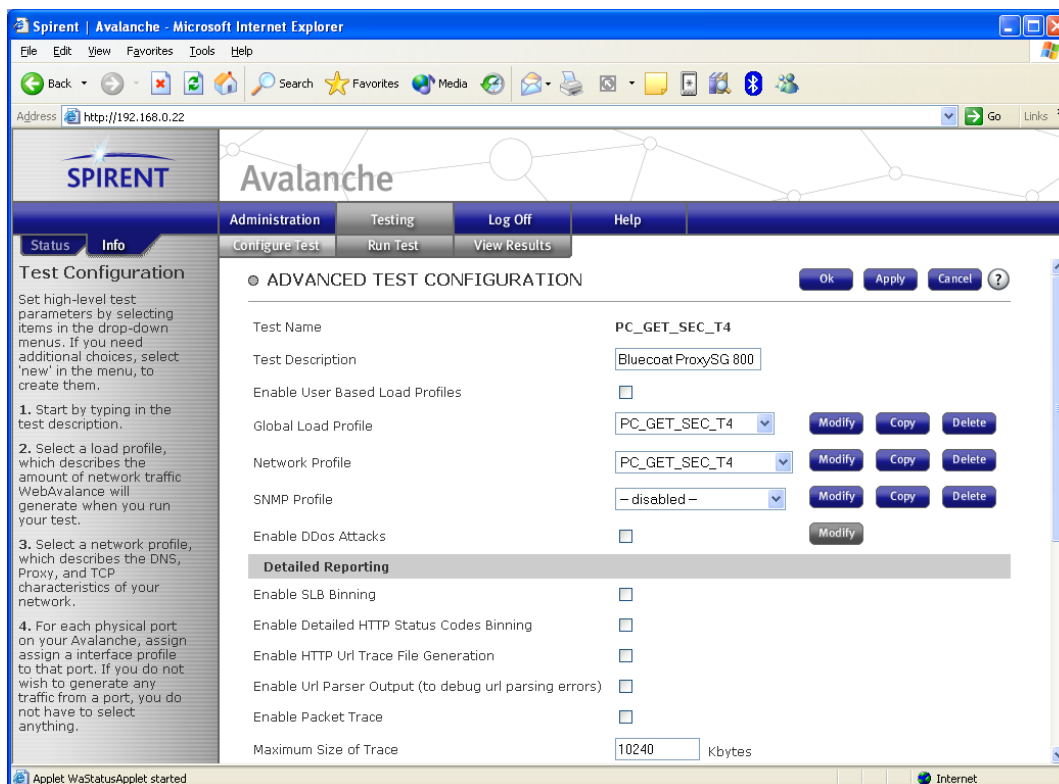


Figure 11 – Creating An Avalanche 2500 Test

The Avalanche 2500 also supports browser cookies, html forms, HTTP posts, and SSL-encrypted traffic. The system therefore gives you the flexibility to specify data sources and mix and match data sets to recreate accurate user behaviour at very high performance levels.

It also simulates SSL loads that can stress the world's most sophisticated secure e-commerce platforms. It also includes configurable cipher suites that enable you to emulate different types of browsers. Avalanche 2500 includes a high-accuracy delay factor that mimics latencies in users' connections by simulating the long-lived connections that tie up networking resources. Long-lived, slow links can have a far more detrimental effect on performance than a large number of short-lived connections, so this approach delivers more realistic test results.

While Avalanche 2500 focuses on the client activity, Reflector 2500 realistically simulates the behaviour of large Web, application, and data

server environments. Combined with Avalanche 2500 it therefore provides a total solution for recreating the world's largest server environments.

By generating accurate and consistent HTTP responses to Avalanche 2500's high volume of realistic Internet user requests, Reflector 2500 tests to capacity any equipment or network you connect between the two systems. Its protocol-level accuracy helps you assure the stability and performance of switches, routers, load balancers, firewalls, caches, and other Layer 4-7 devices. The system is ideal for helping infrastructure service providers validate, enforce, and maintain service level agreements (SLAs) as well as the kind of specific application used for in this report – namely URL filter database testing.

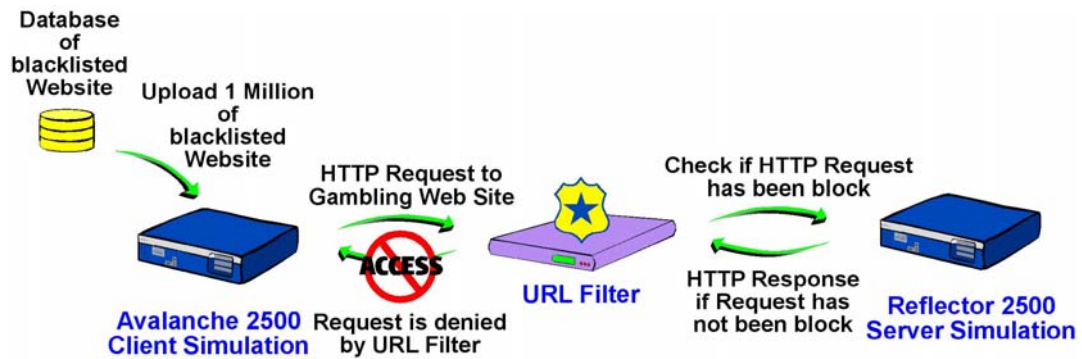


Figure 12 – Creating An Avalanche 2500 URL Filter Database Test

