



Internet Service Provider (ISP) Content Filtering Pilot Report

October 2009

Table of Contents

Executive Summary	1
Introduction.....	6
Background	6
Objective	6
Methodology.....	7
Pilot participants	7
Filtering Technologies	8
Accuracy Testing	10
Objective	10
Methodology.....	10
Network Performance Testing	15
Objective	15
Methodology.....	15
Web page download test.....	16
Streaming download test.....	16
Upload test.....	16
Latency test.....	16
Results	21
Summary of Findings.....	23
Circumvention Testing	25
Objective	25
Methodology.....	25
Results	25
Customer Feedback – Additional Content Filtering	27
Objective	27
Methodology.....	27
Results	27
Costs Associated with Internet Content Filtering	29
Appendix 1: Performance Graphs	33
Appendix 2: Customer Feedback Survey Questions.....	87
Glossary.....	91
Enex TestLab.....	94

Executive Summary

This Internet Service Provider (ISP) filtering pilot was undertaken to determine key information about different ISP-level content filtering technologies when these are applied to:

- Filtering a defined list of URLs, such as the current Australian Communications and Media Authority (ACMA) blacklist (around 1000 URLs at the time of testing)¹,
- In addition to the above blacklist, filtering a wider range and volume of material to provide some level of protection to children using the internet.

Testing was undertaken within an ISP's 'live' network. The different filtering solutions were assessed against a number of factors, including accuracy, effectiveness, impact on network speeds (performance), the relative ease of circumvention and the costs to implement. The filtering technologies that were tested included pass-by filtering; Deep Packet Inspection; pass-through filtering; and proxy filtering (see 'Introduction' for an explanation of these technologies).

Nine ISPs participated in the pilot—two large ISPs, one medium ISP and six small ISPs.

The customers/subscribers of ISPs involved in the pilot were offered a choice of participating. In some cases customers/subscribers had already chosen to receive a filtered service from their ISP.

Filtering of the ACMA blacklist

Three of the ISPs filtered the ACMA blacklist only. Two ISPs used Deep Packet Inspection (DPI) pass-through monitoring technology. The third ISP used a proxy filtering technology.

¹ Internet content is currently regulated through the *Broadcasting Services Act 1992* (Broadcasting Services Act). The existing ACMA blacklist is a list of internet web pages which are defined as 'prohibited' under the Broadcasting Services Act. The list is compiled in response to complaints from the public. Online content is assessed in accordance with the National Classification Scheme for classifying films, computer games and publications.

Accuracy

All participants in the pilot were successful in blocking 100 percent of the ACMA blacklist. This was a requirement of the pilot.

Initially, all filters had issues with loading the ACMA blacklist. These issues were resolved by the vendors, but this indicates a need for routine checking to ensure the blacklist is filtered correctly with each update.

Performance

Testing revealed that the three ISPs filtering only the ACMA blacklist had no noticeable performance degradation that could be attributed to the filter itself.

Circumvention

A technically competent user could, if they wished, circumvent the filtering technology. Testing showed that the filters used for the ACMA blacklist *only* were more easily circumvented than other more complex filters used to cover a wider range and volume of material.

Some ISPs proposed that further consideration be given to greater security and automation of how the ACMA blacklist is distributed and updates installed.

Other Categories of Content

In addition to filtering the ACMA blacklist, six ISPs filtered additional categories of content. The ISPs used DPI pass-through monitoring technology and pass-by hybrid technologies.

Accuracy

All six ISPs achieved 100 percent accuracy in blocking the ACMA blacklist. This was a requirement of the pilot.

In blocking additional categories of content all six ISPs achieved 78 percent to 84 percent accuracy when assessed against the test list of URLs compiled by Enex TestLab (Enex). These results represent an improved level of performance when compared to previous Enex testing, and suggest commercially available filtering products are increasingly effective at including additional categories of content on their filtering lists. One hundred percent accuracy using these commercial lists is unlikely to be achieved as the content on different commercial lists varies and there is a high rate at which new content is created on the internet.

Testing was also undertaken against a list of content, prepared by Enex, considered to be innocuous and which should not be blocked by a filter. All participants experienced some level of over-blocking in this test (i.e. blocking of some legitimate URLs). All filters blocked less than 3.4 percent of such content.

Performance

The majority of these filtering technologies, when correctly installed, enable the filtering of additional content with minimal or no performance impact. One technology, however, displayed a noticeable performance impact. This finding was similar to levels recorded by Enex in previous trials.

Circumvention

Filtering of additional categories of content enabled ISPs to implement measures which made some common circumvention techniques difficult. For example, a third party website which hides the origin of the requested content (proxy site) can be included in a wider list of URLs to be blocked.

As a general rule, there appears to be a relationship between measures to counter deliberate circumvention and impact on internet performance (i.e. stronger circumvention prevention measures can result in greater degradation of internet performance).

Customer feedback

Customers who received a filtered service which included additional categories of content beyond the ACMA blacklist were provided with the opportunity to complete a survey on their experiences during the pilot.

A small number of customers indicated they experienced some over-blocking and/or under-blocking of content during the pilot. These events were considered relatively minor and occurred only once or twice. A small number of customers also reported slower network speeds as a result of the service which filtered additional categories of content.

Overall the service offered by the ISPs was considered effective by customers, with around two-thirds of customers participating in the survey indicating that they would either probably or definitely continue using additional content filtering services.

Customers expressed the view that it was important for there to be mechanisms for self-management of the filter settings and improved visibility of the filter in action.

Costs

The cost to implement filtering of the ACMA blacklist and additional lists of content is influenced by a number of factors including:

- the nature of the filtering solution adopted;
- cost of the filtering hardware/software;
- the extent to which updates of the blacklist are automated;
- ISPs implementing ACMA blacklist-only filtering can expect minimal, if any, customer service costs; and
- fees associated with commercially provided lists for maintaining up-to-date URL lists and handling customer enquiries on filter settings.

Other related findings

Telstra

While not a formal participant in the pilot, Telstra undertook its own testing of ISP filtering of a blacklist of up to 10,000 URLs using a 'domain name server plus proxy server' filtering solution. No customers were involved in the Telstra trial and testing was conducted using Telstra's test environment (which is a replication of its network and used by Telstra for testing its products prior to release).

Telstra found that its filtering solution was 100 percent accurate at blocking a blacklist of 10,000 URLs. Telstra also found there was no discernible performance degradation.

Telstra did not test circumvention, because it considers that filtering can be circumvented by a technically competent user.

Telstra found its filtering solution was not effective in the case of non-web based protocols such as instant messaging, peer-to-peer or chat rooms. Enex confirms that this is also the case for all filters presented in the pilot. Telstra reported that heavy traffic sites could overload its trial filtering solution if included in the filtering blacklist. This is also the case for all filters presented in the pilot.

New Zealand

During the pilot, Enex TestLab engineers, in conjunction with the Department of Broadband, Communications and the Digital Economy (DBCDE) staff, met with staff in the New Zealand Department of Internal Affairs and key participants in the ISP filtering trials undertaken in New Zealand using a blacklist of around 7000 URLs. The Department of Internal Affairs advised that the outcome of the New Zealand trial broadly support the findings from the Australian pilot in terms of accuracy and performance. Circumvention testing was not conducted by New Zealand authorities when performing their filtering trial.

Variables affecting internet services

Testing internet performance in a live environment requires a number of uncontrollable variables to be considered. Although these variables are unrelated to the filter, they may affect the speed of an individual internet connection. For example, distance from exchange/tower, condition and configuration of Customer Premises Equipment and customer computer system(s), and the number of subscribers connected to a service, versus the capacity of the upstream connection that the ISP maintains, commonly referred to as the “contention ratio”, will all affect performance, and can do so at 40 percent performance degradation over theoretical maximum line-rate, or more in some cases. The performance results in the pilot are assessed in the context of these variables.

Filter Setup

The testing showed that while the configuration of the blacklist for a filter may present some initial difficulties, these can be resolved. Filter vendors were advised of these findings and responded to these in current versions/models. Some participants proposed that further consideration be given to greater security and automation of how the ACMA blacklist is distributed and updates installed. Enex considers that machine-to-machine transfer of the list is the most desirable distribution method.

Introduction

Background

Cyber-safety Plan

In May 2008, the Australian Government committed \$125.8 million over four years to support a range of cyber-safety measures, including education, law enforcement and internet filtering.

Live Pilot

On 10 November 2008, an Expression of Interest was issued, inviting interested ISPs to participate in a live ISP-level filtering pilot for a minimum six-week period.

The pilot aimed to test a range of content filtering solutions in a real-world environment, with the cooperation of ISPs and their customers.

A limited amount of funding was made available to assist ISPs with the costs associated with participation in the pilot. Companies that wished to apply for this funding were asked to submit their applications to the Department of Broadband, Communications and the Digital Economy (the Department) by 8 December 2008.

Sixteen applications to participate in the pilot were received, with representation from a cross-section of the industry and filtering solutions. The applications were assessed by the Department and its independent technical advisor, Enex TestLab.

Six ISPs withdrew their applications and one ISP was excluded for technical reasons.

Agreement was reached with nine ISPs to participate in the pilot.

Objective

The objective of the ISP filtering pilot was to test a range of filtering solutions implemented within a range of ISP network environments, and across a broad range of technical platforms.

For each type of filtering solution, the pilot aimed to investigate the:

- Accuracy of filtering, including over-blocking and under-blocking of content.
- Impact on network speed (performance) from the perspective of the user and ISP.
- Relative ease of circumventing the filtering.
- Ease of use from an end-user perspective for filtering solutions that require end-user involvement.
- Costs associated with introducing ISP filtering.
- Scalability of the filtering.
- Effectiveness of any additional functionalities of the filter product.

Methodology

The pilot was conducted on behalf of the Department by Enex TestLab (Enex) in cooperation with the participating ISPs and their customers. Enex is an established provider of information and communications technology testing and benchmarking services.

Enex provided technical assessments of ISP-level filtering, including reporting on the results of these assessments.

To reduce the impost on participants, testing was undertaken via a remote connection to the ISPs' networks and did not require a physical presence at the various ISP locations.

The participating ISPs provided Enex with access to both a filtered and an unfiltered service, enabling accuracy and circumvention testing as well as concurrent performance testing of the filtering technologies. Switching the filtered and unfiltered services during the pilot enabled Enex to determine whether it was the filter or other factors affecting network performance.

One participating ISP (hereby called Participant1) was unable to perform the switch over of the filtering between services at this mid-point in the performance test, due to the configuration and location of their filter and test services. All other service providers complied with this requirement.

The purpose of the pilot was to test a variety of metrics across participating filtering technologies in a live ISP environment but not to select or rate the filtering technologies. Therefore, for the purpose of anonymity, ISP names and filtering solution names are not provided in this report. ISPs are referred to as Participant1 to Participant9 and filtering solutions as SolutionA to SolutionI.

Pilot participants

Nine ISPs participated in the pilot.

Participants could perform filtering of the Australian Communications and Media Authority (ACMA) blacklist *only*, or filtering additional material based on the choice of the customer. Many filter vendors have a number of categories that permit administrators and customers to select from many subjects. Some vendors have over 100 categories with hundreds of thousands of sites listed. For the purposes of the pilot, filtering of additional content was optional for participating ISPs and their customers.

Participants were tested for accuracy in blocking the ACMA blacklist only and all nine participants achieved 100 percent accuracy - a base requirement of the pilot.

For accuracy, performance, and circumvention testing, the following ISPs provided filtering of the ACMA blacklist only: Participant1, Participant2 (SolutionF), Participant3 and Participant4.

For accuracy, performance and circumvention testing the following ISPs provided filtering of the ACMA blacklist *and* additional content simultaneously: Participant2 (SolutionA), Participant5, Participant6, Participant7, Participant8 and Participant9.

Filtering Technologies

The filtering solutions and types of filters tested during the pilot are listed below.

<i>ISP</i>	<i>Filtering solution/s</i>	<i>Type</i>
Participant1	SolutionD	DPI – Pass-through
Participant2	SolutionA SolutionF	Pass-by Pass-by
Participant3	SolutionD	DPI – Pass-through
Participant4	SolutionC	Proxy
Participant5	SolutionE	DPI – Pass-through
Participant6	SolutionB	Pass-by
Participant7	SolutionG	Pass-by
Participant8	SolutionH	Pass-by
Participant9	SolutionI	Pass-by

It is noted that the same filtering technology was used by a number of participating ISPs, but the solutions were licensed and provided under different vendor/product names.

Three key technologies were selected by pilot participants, pass-by, pass-through and proxy.

Pass-by filters

A pass-by filter does not require all traffic to pass through the filter. Pass-by filters comprise two types: hybrid and port mirroring. Only hybrid pass-by technology was used in this trial.

Hybrid pass-by filters consist of two-stages. The most common first stage populates routers with IP addresses relating to the URLs to be blocked (generally using Border Gateway Protocol). The routers then divert any traffic to those IP addresses to a second stage. The second stage examines the diverted traffic to see if the actual URL request matches that on the block list. Most commonly the action by a filter is a block-page being returned to the end-user, otherwise the traffic is allowed to pass on and the site returned to the user.

If the request is for a site with an IP address that is not on the filter's translated list, then the request is processed as a normal transaction with no further involvement by the filter. The premise of hybrid technology is that the majority of traffic on a network is not to sites sharing an IP address with a site to be blocked, and therefore most traffic can pass-by without examination (except for a cursory check against an IP list on the router).

Hybrid pass-by filters should not be confused with IP-blocking. IP-blocking occurs only at the router level, using a list of URLs translated to IP addresses to re-route those addresses to be blocked. With the growth in technologies such as virtual hosting and load balancing, however, a single IP address can host a number of different URLs, so simple IP blocking is not accurate. In tests previously conducted by Enex, simple IP blocking results in approximately 20 percent over-blocking of sites.

Port mirroring pass-by filters utilise a standard feature of routers to send a duplicate of traffic to a monitoring device. The filter is connected to the router using this port mirroring function, and examines a copy of all URLs flowing through the router. If the filter detects traffic to a blocked URL, it signals the router to interrupt, or reset that connection. Unlike the hybrid type, the filter must examine all the URLs to see if they are on the filter list. Vendors claim that because the normal traffic flow does not need to pass through the filter, there is potentially less performance impact.

Pass-through filters

A pass-through filter operates in-line with the traffic, essentially examining every address passing through. There are a number of pass-through technologies available, most of these are Deep Packet Inspection (DPI). DPI is not only used in content filtering devices but also to detect network attacks, malware and other security functions.

Proxy filters

A proxy server is a server that caches or stores requests from users for a variety of purposes—traditionally to increase network performance to downstream end-users if the upstream connection is limited. Today, proxy servers can also be used for content filtering by comparing a list of blocked URLs against the incoming requests and, where necessary, re-directing to a block page. Proxies are commonly used as the second stage in conjunction with hybrid filters (described above).

Accuracy Testing

Objective

A key aim of the pilot is to test the ability of filtering technologies to accurately block internet content, in particular where that content is a defined list of URLs.

Enex tested accuracy of the technologies to block a defined list of URLs (i.e. the ACMA blacklist).

Enex also tested additional category filtering and the extent to which these technologies under-block (i.e. allow access to content they are configured to block) and over-block (i.e. block access to content they should allow).

Methodology

Three lists of URLs were tested during the pilot. The first is the ACMA blacklist. The second test list was created by Enex and contained content 'inappropriate' for children. The third test list, also created by Enex, contained URLs of content that was 'appropriate' for children (innocuous).

The first list of URLs was used to test the extent to which the filter products blocked access to content contained on the ACMA blacklist.

The Online Content Scheme (the Scheme) introduced in 2000 under the *Broadcasting Services Act 1992 (Cth)* (the BSA), regulates content on the internet. The Scheme is contained chiefly in Schedules 5 and 7 of the BSA. Under the BSA, 'prohibited content' and 'potential prohibited content' include content that has been classified or is likely to be classified Refused Classification under the National Classification Scheme.

Where content is hosted in Australia and is found by ACMA to be prohibited, the ACMA has the authority to issue a take-down notice requiring the relevant content service provider to remove the content from their service.

Where content is not hosted in Australia and is prohibited, ACMA will notify the content to the suppliers of approved filters so that access to the content using such filters can be blocked. The compilation of URLs referred to filter providers is known as the ACMA blacklist. The blacklist is available to a small number of PC and local network based filter providers in Australia, and access is strictly controlled.

All nine participating ISPs tested filtering of the ACMA blacklist of URLs.

The second URL test list contained content considered inappropriate for children and included classifications in the range MA 15+ to X18+.

Material on the third test list contained URLs considered innocuous and were generally below MA15+.

The second and third lists were not known to the vendors. They therefore also served as a check of the filter's ability to identify and block/pass relevant additional content.

Six ISPs participated in testing the URLs on the second and third lists.

Control tests

Prior to performing the testing during the pilot using the three lists, each site on each list was tested by Enex to ensure that it was still live. Because each participant was tested at a different date/time, checking was performed for each site just prior to the formal accuracy testing being conducted.

Blocking of the ACMA blacklist

All nine participants tested filtering of the ACMA blacklist of URLs.

The first URL index tested the extent to which the filter products in the pilot block a defined list of URLs i.e. the ACMA blacklist of prohibited URLs. For the purposes of the pilot it was necessary for each participating ISP to achieve 100 percent blocking of these sites before any further testing was conducted.

Initially, several participants experienced difficulty loading and blocking the complete ACMA blacklist. Some of the filters needed adjustments to be made so that they could recognise URLs that were long and complex and included spaces. Others included colons, question marks and percentages. Some URLs were associated with more than one IP address and some URLs redirected the user to a second URL.

Following consultations with the product vendors, all issues experienced with loading URLs contained on the ACMA blacklist were resolved.

The result was that all participating ISPs achieved 100 percent blocking of the ACMA blacklist.

Results for accuracy tests – ACMA blacklist only

ISPs filtering the ACMA blacklist only

Service	Participant1	Participant2 (SolutionF)*	Participant3	Participant4
ACMA Blacklist	100%	100%	100%	100%

*Participant2 using SolutionF was involved in testing a simulated ‘hosted’ filtering solution, i.e. a filtering solution hosted remotely of the ISP and potentially offered as a third party service to a number of ISPs. The SolutionF was hosted in Enex’s data centre rather than at Participant2’s premises.

ISPs filtering the ACMA blacklist and additional content

Service	Participant2 (SolutionA)	Participant5	Participant6	Participant7	Participant8	Participant9
ACMA Blacklist	100%	100%	100%	100%	100%	100%

Blocking additional content (based on user choice)

The second list of URLs tested the extent to which the filter products blocked additional content, i.e. content that may be regarded as harmful or inappropriate for children, based on the choice of the user.

This list was drawn from an existing database of URLs held by Enex. The content on this list would likely be classified as MA15+, R18+ and X18+. A proportion of the content considered to be strong M was regarded as being close to the MA15+ classification, and was also included on the test list.

Inappropriate for children test list

Content on the inappropriate for children test list included:

Gambling	Adult
Lingerie/Swimsuit	Drug-advocacy
Nudism	Gross-content
Profanity	Racism/Hate
Sex	Terrorism/Crime

Blocking of innocuous content (over-blocking)

The third URL index test measured the extent to which the selected filter products over-block (i.e. block access to content that is regarded as innocuous). This third index was drawn from an existing database of URLs held by Enex. The content on this list is likely to be rated in the range G to M.

The content on this list, while innocuous, is also designed to potentially lead some filters into recording a false positive, for example, references to sperm whales and robin red breast for example.

Innocuous test list

Content in the innocuous test list included:

Animals	Plants
Cooking	Computers
Search/Reference	Government
Science	Children
News/Weather/Sport	Art/Literature
People/Travel	

Results for accuracy tests – additional content

Accuracy testing results are as follows:

- All filters participating in additional content filtering in the pilot blocked between 78.80 percent and 84.65 percent of inappropriate material.
- All filters participating in additional content filtering in the pilot blocked less than 3.37 percent of innocuous content.

ISPs participating in filtering additional content

Service	Participant2 (SolutionA)	Participant5	Participant6	Participant7	Participant8	Participant9
Inappropriate	80.72%	84.65%	80.97%	80.72%	78.80%	82.03%
Innocuous	2.76%	3.17%	2.78%	2.44%	2.87%	3.37%

Enex considers it unlikely that any filter vendor would achieve 100 percent blocking of the URLs inappropriate for children without significant over-blocking of the innocuous URLs because the content on different commercial lists varies and there is a high rate at which new content is created on the internet. Enex has also noted, through previous testing, that the higher the accuracy the higher the over-blocking.

Enex considers any score above 70 percent for the second list would be reasonable by industry standards, providing there is zero to very low over-blocking of the third list.

In terms of over-blocking the results of this trial show that, while an improvement on previous testing levels, this is still considered high.

A 'successful outcome', from a filter vendor perspective, of the second and third URL list tests would be a high percentage of blocking of URLs in list two, matched by a zero to low blocking of URLs in list three. Blocking rates of below 2 percent would be considered low.

Where additional category filtering is provided, reporting mechanisms to notify the filter vendor of potential over-blocking would be beneficial. It would also be beneficial to provide the user a visual check/confirmation that the filter is active and operational.

Network Performance Testing

Objective

The objective of the performance testing is to identify the impact, if any, on the end-user connection due to the introduction of a filtering solution.

Methodology

Each participating ISP provided Enex with two sample internet services: filtered and unfiltered. Each service provider was given two identical, pre-configured laptop computers to connect to their sample filtered and unfiltered services. Enex test engineers then remotely accessed these laptops and conducted the technical testing components of the filter pilot.

Mid-way through the performance testing period participants were required to switch the filtering to each sample service, swapping the sample filtered service with the sample unfiltered service. This enabled Enex to identify and record performance disparity in the live environment that may have been caused by factors other than the filter itself.

Using Enex's eMetric network performance toolset, the following tests were conducted on the services provided by each participating ISPs:

- Webpage download test;
- Streaming download test;
- Upload test; and
- Latency test.

All tests were performed to/from a single Melbourne data centre with sufficiently high bandwidth connectivity ensuring consistency in performance results across all participants. This also ensured that no bottlenecks in the transmission performance or latency were due to the test environment or tool.

At one participant's site, engineers installed a network load generator to simulate an artificial load on the filter and network under test to ascertain what, if any, network performance impact was introduced by the filter while under above average and sustained load. Load generators are used by test laboratories and network engineers to gauge the maximum network system throughput or identify network bottlenecks by simulating larger than normal traffic volumes.

Using a load generator in this participant's network ensured results were consistent with measuring the performance scalability of the participant's filtering solution and could be compared to the non-loaded test results.

Web page download test

The web page download test was performed using both filtered and unfiltered internet connections. The test involved downloading a set of predefined websites and recording the time taken to complete the download of each website. A combination of dynamic and static websites was used. The test was repeated at regular intervals over a period of time. The web pages' combined size was approximately 8MB. (A half size sample was used when assessing dial-up services due to the significantly slower speed - less than 56 kilobits per second.)

Frequency: Every 15 minutes over a 7-14 day period.

Streaming download test

The streaming download test was performed using both filtered and unfiltered internet connections. The test involved recording the time taken to download a single file (approximately 2MB in size). The file was encrypted so that it could not be compressed and was randomly generated at the time of testing to prevent caching. This test assessed how efficiently media files (i.e. music or video files) were downloaded.

Frequency: Every 15 minutes over a 7-14 day period. (Frequency for dialup testing was decreased)

Upload test

The upload test was performed using both filtered and unfiltered internet connections. The test involved recording the time taken to upload a single encrypted file (approximately 2MB in size) from a sample service to a designated server.

Frequency: Every 15 minutes over a 7-14 day period. (Frequency for dialup testing was decreased.)

Latency test

Latency is the time that it takes for a packet of data (that is a certain size) to travel from one location to another, and is generally measured in milliseconds (ms). An example is real-time audio such as a telephone call, where a delay of 300ms or more is noticeable to the human ear. In networking terms, engineers aim to have a round-trip time of less than 300ms. This ensures Quality of Service for real-time dependent traffic such as Voice over Internet Protocol (VoIP) services. For the purposes of the pilot, latency was measured to see if there was any noticeable impact when filtering was introduced.

The latency test was performed using both filtered and unfiltered internet connections. Latency was measured by determining the total round-trip time from the sample service to a designated server.

Frequency: Every 15 minutes over a 7-14 day period.

Testing a filtered and unfiltered service

Testing of both services was conducted in parallel (where possible) over a 7-14 day period. Enex ran the same tests for a further 7-14 days (the exception being Participant1—discussed later) switching the services for the test from filtered to unfiltered and vice versa. This is important because the test needed to identify any performance disparity in the live environment that might have been caused by factors other than the filter itself – for example, factors such as distance from exchange/tower, condition of cabling, transmission differential and modem or ISP equipment variances.

It is commonly recognised that internet service performance is variable due to these (and other) factors. This was taken into consideration when comparing the filtered and unfiltered services.

Scalability testing using a load generator

As the pilot was conducted in a live environment, Enex evaluated the filtering technology in a simulated heavy-load environment.

For the pilot, pass-through filtering was load tested. A load generator was installed in the ISP's network to simulate a user load of 1.6Gbps. Enex then conducted performance testing to assess scalability and whether the filter adversely affected internet performance under higher loads. During the pilot and testing in a live ISP environment it was not technically possible to introduce artificial loads across all the participating ISPs and the filtering technologies.

Conversion of URL lists to IP address lists

Recognition of all IP addresses associated with a given URL is a possible concern for filter vendors who rely on converting URL lists to IP address lists for initial interpretation at the router level. Many large-scale websites and hosts commonly implement a number of servers in a 'server farm' or 'cluster' for load balancing (to maintain performance). If a blacklisted URL resides on such a system, the filter vendor would need to identify all potential IP addresses associated with that URL.

An increasingly common technique used to negate filters is known as 'fast fluxing' and involves the rapid and automated change in IP addresses. This technique can be effectively employed against a number of filters that block using the IP address. Importantly, it should be noted that this technique is employed by providers of content rather than end-users.

The ACMA blacklist does not simply include top level URLs. It is, in fact, very granular and may specify detail right down to a particular target within a site (e.g. the actual page listed on a site). Other jurisdictions' blacklists (such as New Zealand's) contain the top level URL only, so everything hosted on that site is considered blacklisted.

Machine to machine transfer of the blacklist (from ACMA to vendors' filters or ISPs' filters) is the most desirable distribution method. However, such a system may need to include some checks and balances to ensure that URLs are not inadvertently modified when certain special characters are interpreted. A set of standards would be advisable, ensuring uniform handling of known and unknown URL formats. Enex considers the way in which standards-compliant browsers handle URLs would be a reasonable place to start.

Number of URLs handled by filters

The question has been raised about how many URLs could be handled by a filter before experiencing a load/performance degradation. It has been suggested by some stakeholders that 10,000 URLs may be a tipping point. During the pilot the ACMA blacklist comprised approximately 1000 URLs and it was not possible to perform a live test of more than this sample. However, it is known that the number of URLs on the filter vendor lists tested during the pilot (i.e. ISPs that tested filtering of additional content), ranged from 100,000 to millions of URLs and this did not have a discernible effect on network performance.

Telstra - Domain Name Server plus Proxy Server

While not a formal participant in this pilot, Telstra undertook its own testing of ISP filtering of a blacklist of 10,000 URLs using a 'domain name server plus proxy server' filtering solution. Customers were not involved in the Telstra trial.

Telstra found that there was no discernible impact on end-user experiences with a list size of 10,000 URLs. Telstra describes the impact as equivalent to one seventieth of the blink of an eye.

Pass-by technologies

Pass-by filters often encompass extremely large list sizes, for example some filters in the pilot had millions of URLs on their lists of content that could potentially be filtered based on consumer choice.

Enex found that performance is not dependent on the total number of URLs included in the blacklist. In the pilot, pass-by technologies were tested for the ACMA blacklist only, as well as more substantial lists, and there appeared to be no (or negligible) differences in performance impacts between the two. In Enex's view any variations in performance results were not due to the size of the lists.

Pass-through technologies

Pass-through technologies check all traffic that passes through the filter. For one of the pass-through filters in the pilot, the increase in list size did not result in any additional performance impact. For another of these filters, however, noticeable performance impact was identified. Enex could not ascertain whether this was due to the size of the list used by the participant/vendor or simply the technology itself.

The vendor for one of the pass-through technologies stated that its product could handle 100,000 URLs natively/statically, and more with the addition/integration of an external policy server. This filter displayed no noticeable performance impact with two participant ISPs in the live trial.

Proxy technologies

Proxy systems are designed to handle significant traffic flows inherently; therefore any increase in size of the URL list is unlikely to have any noticeable performance impact. One proxy filter vendor in the pilot claims to have licence agreements with many list providers and the capability of running multiple lists simultaneously. These lists could amount to hundreds of thousands, if not millions, of URLs which could be filtered subject to end-user choice.

Capacity of filters to handle high traffic loads/sites

In a pass-by filtering solution the actual traffic load placed through the filters is very low because only a small percentage of end-users would be attempting to access sites on the blacklist at any one time.

However, in situations where there is a potential for very high traffic sites, such as YouTube, to have pages on the filtering list, this could result in significantly higher traffic rates passing through the filter, even though the specific pages being accessed are not those on the blacklist. This could cause additional load on the filtering infrastructure and subsequent performance bottlenecks.

To support peaks in traffic, vendors recommend allowing additional network capacity of approximately four times the estimated traffic at the filter.

Variables that may affect Internet speeds

Conducting real world testing in a live environment has many uncontrollable variables. It is well documented that there are a number of factors that affect the speed of any individual internet connection. When testing live internet performance it is important to establish some margins to accommodate these variables - a conservative +/-10 percent, without impact on a filter.

A major factor affecting internet speed is the contention ratio maintained by the ISP. The contention ratio is the number of people connected to a service versus the upstream connection that the service provider maintains at the point of aggregation. The more users accessing the internet simultaneously, the smaller the amount of available upstream bandwidth will be, and the greater the impact on an end-user's speed.

Commonly cited explanations for actual degradation of performance over the theoretical performance include:

- configuration of the Customer Premises Equipment;
- distance from the exchange and condition of the line (in fixed-wire internet solutions);
- distance from the tower and physical/environmental conditions (in wireless internet solutions).

In some cases as much as 40 percent of an internet service performance could be lost to these factors. And in theoretically higher speed services, such as ADSL2, significantly greater performance losses are commonly accepted.

As anticipated, during testing all live services were seen to vary in performance from time to time – often fluctuating to as low as half of the maximum performance. This highlights why it was important to run performance testing over an extended period and average the results. Sometimes, technical problems unrelated to the filter may have prevented the tests being run for an adequate time, or cause results to vary erratically.

Baseline performance for the pilot

The baseline performance for the pilot was set by Enex as follows:

- Performance impact is considered negligible if it is less than +/-10 percent;
- Performance impact is considered minimal if between 10 and 20 percent;
- Performance impact is considered noticeable if more than +/-20 percent.

To put this in layman's terms, say it takes five seconds to download a particular webpage then anything more/less than an extra half a second increase/decrease would be considered minimal, but more than one second would be considered noticeable to the user.

It is noteworthy that sometimes a filtered internet service scored better performance results than an unfiltered service. There are two main reasons for this.

The first is that normal variations in internet speed may have affected the area or equipment serving a filtered computer less than it affected an unfiltered computer. This is because a slightly different path might be needed to connect a computer to the internet through a filter.

The second reason is that some proxy based internet content filters actually store frequently used information. By storing this information a filter can take a short cut and send data straight to a computer without having to collect it from somewhere else on the internet first. This second issue should only affect the 'Webpage download test' in this study. The data files used in the eMetric tests include random data which varies from test to test therefore invalidating any attempts to cache the data.

Results

Filtering the ACMA blacklist only

ISP	Streaming Download test	Upload-test	Latency test	Web page Download test
Participant1	-17.32%	-16.95%	-5.95%	0.46%
Participant2 with SolutionF	11.00%	27.20%	-4.90%	5.24%
Participant3	-1.42%	4.63%	1.41%	6.43%
Participant3 with load generator	2.05%	4.70%	1.84%	-7.16%
Participant4	1.65%	0.04%	-5.14%	2.98%

Filtering ACMA blacklist and additional content

ISP	Streaming Download test	Upload-test	Latency test	Web page Download test
Participant2 with SolutionA	44.15%	1.11%	3.71%	-0.10%
Participant5	35.80%	20.56%	-0.97%	36.45%
Participant6	-4.20%	-5.52%	3.06%	-4.07%
Participant7	-0.23%	-0.39%	0.46%	-4.68%
Participant8 ADSL2	6.92%	1.69%	-7.02%	3.02%
Participant8 Ethernet	0.30%	0.72%	4.90%	9.49%
Participant8 dial-up	8.85%	22.49%	3.77%	-0.53%
Participant9	-1.62%	-0.90%	-5.95%	10.02%

The following criteria were applied by Enex to measure the possible impact of the filters on network performance.

Negligible impact on network performance

10% and below: individual service performance impact is negligible to the end-user. It would be difficult for the test to distinguish the impact of the filter from any other factor potentially affecting network performance.

Minimal impact on network performance

10%-20%: individual service performance is impacted by the filter, but it is potentially minimal given the myriad of other factors that can contribute to the performance degradation of the service.

Noticeable impact on network performance

20% and above: is noticeable and shows that the filter affects the individual internet service performance.

Minus 10% and below: is considered an irregular/incorrect result, highly anomalous with reasonable expectations.

Note that a negative score in the tables indicates that the performance measurement was not degradation, but an increase in performance between the non-filtered and filtered services test. In relation to Participant1, this service was not provided to Enex with complete experimental control as the filtering could not be switched between the two sample services mid-way through the performance testing.

Summary of Findings

ISPs filtering the ACMA blacklist only

Participant1 – Results show that the *filtered* sample internet service performed considerably faster than the *unfiltered* sample internet service. The test data shows that the unfiltered connection performance fluctuated throughout the duration of the testing, whereas the filtered connection performance was relatively stable. Given the variability of the unfiltered service, Enex concludes that the performance of the sample unfiltered service was affected negatively by factors other than filtering. The difference in performance between the two services, therefore, cannot be attributed to the filter.

Participant2 (with SolutionF filter) – Minimal impact on file downloads, noticeable impact on file uploads and negligible impact on web page downloads were recorded.

Participant3 – Negligible impact on file downloads, file uploads and web page downloads were recorded.

Participant3 (with load generator) - With simulation of additional network traffic at 80 percent of capacity, the filter had a negligible effect on file downloads, file uploads and web page downloads during the pilot.

Participant4 – Negligible impact on file downloads, file uploads and web page downloads were recorded.

ISPs filtering ACMA blacklist and additional content

Participant2 (with SolutionA filter) – Noticeable performance impact was recorded for the filtered service when performing file downloads. Enex considers that it was unclear whether this was attributable to the filter or other non-related factors. The filter had a negligible impact on file uploads and web page downloads.

Participant5 – Noticeable performance impact was recorded for the file downloads, file uploads and web page downloads. Enex has found in previous testing that factors that may have affected performance were related to using a pass-through filter. The pass-through filter assesses all traffic that passes through the network. It can therefore have an affect on performance if not planned and implemented correctly and of sufficient scale and capacity to handle the flow of traffic. Furthermore it was determined, following consultation with the filter vendor, post-pilot, that the content list was maintained overseas, therefore it was suggested that additional latency and timeouts were generated for filtering against the vendor's list.

It was recommended that the purchase of an additional management appliance would enable the vendor lists to be downloaded for local reference and potentially solve this problem. This claim, however, was not tested by Enex as a component of the trial.

This service blocked more circumvention attempts than any other filter in the pilot.

Participant6 – Negligible impact on file downloads, file uploads and web page downloads were recorded.

Participant7 – Negligible impact on file downloads, file uploads and web page downloads were recorded.

Participant8 – Filters on the ADSL2 and Ethernet services had a negligible impact on file downloads, file uploads and web page downloads.

Filters on the dial-up service had a negligible impact on file downloads and web page downloads and a noticeable impact on file uploads.

Some technical issues were experienced with Participant8's dial-up service which resulted in customers contending for bandwidth during the testing period. Access to the dial-up network was provisioned by a separate provider with Participant8 acting as a reseller of this service. It is reasonable, however, to attribute the noticeable impacts on file upload performance to the services technical difficulties rather than the filter; particularly when compared to the vendor's filter performance over ADSL2 and Ethernet.

Participant9 – Negligible impact on file downloads, file uploads and web page downloads were recorded.

Circumvention Testing

Objective

Circumvention testing was undertaken to determine the ability of the filtering technologies to counter attempts by a user to circumvent the filtering implemented in the ISP network.

Methodology

Prior to the commencement of circumvention testing, participating ISPs were encouraged to configure the filtering solution in a way that would minimise the potential circumvention of internet filtering.

A number of circumvention techniques commonly used to bypass filtering were tested by Enex. The methods tested were non-intrusive - that is they did not attempt to penetrate the ISPs' networks or disable the filtering implemented.

Nine separate circumvention techniques were tested by Enex, which involved 37 individual circumvention tests being undertaken for each ISP.

The circumvention techniques tested will not be described in this report for public interest reasons.

Results

The number and percentage of circumvention attempts successfully blocked (out of 37 tests for each ISP) are outlined in the tables below.

ISPs filtering the ACMA blacklist only

ISPs	Number of circumvention attempts successfully blocked out of 37 tests	Percentage of circumvention attempts successfully blocked
Participant1	3/37	8.1%
Participant2 (using SolutionF)	4/37	10.8%
Participant3	6/37	16.2%
Participant4	5/37	13.5%

Circumvention attempts for ISPs filtering the ACMA blacklist and additional content

ISPs	Number of circumvention attempts successfully blocked out of 37 tests	Percentage of circumvention attempts successfully blocked
Participant2 (using SolutionA)	14/37	37.8%
Participant5	35/37	94.5%
Participant6	16/37	43.2%
Participant7	15/37	40.5%
Participant8	28/37	75.6%
Participant9	28/37	75.6%

Participant5 successfully blocked 35 out of 37 circumvention attempts resulting in the highest result in the pilot of 94.5 percent. It is noted during the pilot, however, that noticeable performance degradation was observed for the filtered service which was utilising a pass-through technology.

Participant8 and Participant9 both blocked 28 out of 37 circumvention attempts (75.6 percent). These results were better than other ISPs using the same filtering technology. This would seem to suggest that ISPs have some flexibility in configuring the filtering solutions according to their individual requirements. For example, countering of circumvention by ISPs may be adjusted in order to improve performance.

The testing for circumvention generally indicates that filtering of additional categories of content, enabled ISPs to implement measures which made some circumvention techniques more difficult to use. For example many commercial lists have a “proxy” category; proxies are a common form of filter circumvention.

Customer Feedback – Additional Content Filtering

Objective

An online feedback survey was undertaken with customers of the six ISPs that agreed to participate in the filtering of additional categories of content (i.e. where ISPs and customers chose to have content filtered beyond the ACMA blacklist).

No customer feedback survey was issued to the customers of ISPs that filtered the ACMA blacklist only.

Methodology

The six ISPs in the pilot participating in the feedback survey were Participant2, Participant5, Participant6, Participant7, Participant8 and Participant9. These ISPs and their customers had agreed to participate in the filtering of additional categories of content. ISPs were consulted on the survey prior to its distribution to customers.

Completion of the feedback survey was voluntary. Information about customers and ISPs remains anonymous and no personal information such as names, addresses or websites visited by customers was sought in the survey.

The feedback survey was designed to ensure minimum impact on ISPs and their customers. A link to the online survey was emailed to ISP customers and it took approximately five minutes to complete.

The questions in the online customer feedback survey are provided at Appendix 2.

Results

The customers completing the online survey were either parents or those responsible for the internet service.

Home/Business Users

Eighty two percent of respondents to the survey were home users. Use of the internet both for home and personal use accounted for over 50 percent and business use was 18 percent. The outcomes of the survey, therefore, are most relevant in the context of personal use.

Thirty seven percent of respondents had children 15 years or younger using an internet service that was being filtered.

Under-blocking

Participants were asked if they had experienced content that they felt should have been blocked during the pilot (e.g. to identify perceptions of under-blocking), and 80 percent of respondents indicated that they did not access such content during the pilot. Those who did access such content described the incidences as relatively minor and occurring infrequently.

Over-blocking

Participants were asked if access was blocked to sites that they thought they should have access to (e.g. to identify perceptions of over-blocking). Only a couple of respondents indicated that they were frequently blocked access to sites they should have had access to, while most respondents reported that they were blocked “possibly once or twice”. A number of users reported they were “unsure” as to whether they had been blocked unnecessarily and some customers commented on the benefit of seeing some “output” of the filters action.

Features of the service

70 percent of respondents indicated there was no feature of the service that was “least liked”. A small number of customers commented that they perceived slower network speeds.

Customers were asked if they thought the service should have other features, and 74 percent of respondents considered that no other features seemed necessary. Some respondents suggested more user controls in adjusting the filter settings to cater for different users within the household.

Overall satisfaction

Overall the service offered by the ISPs was considered effective by customers, with 80 percent of respondents indicating that the filtered service either entirely or generally met their needs. 72 percent of the sample stated that they would either probably or definitely continue to use the service.

Summary

Maintaining the balance between under- and over-blocking of additional content, and allowing customers to easily change the filter settings to suit their particular contexts, appears to be important to the customers in the pilot. A small number of customers believe they experienced some speed degradation. While the value of ISP based filters is clear to some, mechanisms for self-management of filter settings that provide additional categories of content and better visibility of the filter in action appear to be important.

Costs Associated with Internet Content Filtering

Introduction

This chapter outlines the nature of the costs that may be incurred by ISPs providing internet content filtering. These costs are associated with the procurement, implementation and ongoing operation and management of an internet content filtering solution.

For commercial reasons, a number of ISPs and filter vendors were not prepared to release detailed information on the costs of filtering.

Cost Elements

The actual costs incurred by individual ISPs will vary depending on:

- The size of the ISP and how they deliver their services (e.g. network size and architecture).
- The filtering solution selected by the ISP (e.g. pass-by technology, pass-through technology or proxy based solutions).

Initial Costs

The cost components associated with the implementation of ISP filtering of the ACMA blacklist and additional categories of content may include the following:

- Purchase costs associated with the initial acquisition of any hardware and software associated with the filtering solution (if not already owned by the ISP) or specific adaptation costs required for existing systems.
- Installation and configuration costs associated with the implementation of the filtering solution.
- Network reconfiguration – depending on the individual solution and the ISP's network arrangements, there may also be costs associated with the re-arrangement of their network to accommodate the filtering solution (e.g. redirecting user traffic flows).
- In addition, where customers choose to have filtering of additional content, modifications may be necessary to customer management applications to provide customers with the ability to self-select the type of services to be filtered.

Ongoing Costs

After the filtering is implemented, and depending on the solution chosen, the ISP will incur a range of costs associated with the ongoing operation, maintenance and management of the filtering solution.

These ongoing costs may include:

- Hardware and software maintenance costs --generally provided by solution vendor.
- Subscription fees – costs associated with ongoing subscriptions, mainly for additional content filtering lists.
- Service provider support and management costs – for ongoing support and management of the system specifically by the service provider staff.
- Additional customer help desk support costs arising from queries associated with content filtering.
- Software maintenance – costs associated with ongoing maintenance of the customer management system.
- Hardware refresh costs – costs associated with the growth in demand as well as the upgrade and replacement of the equipment once it reaches its end-of-life (typically around five years).

Third Party Hosted Filtering Services

For small ISPs third party hosted filtering services may be an option.

The costs associated with content filtering for smaller ISPs (i.e. 1000 to 5000 customers) can be significant on a per user basis in comparison to the larger ISPs (100,000 customers and above). This arises from economies of scale and discounting arrangements that would tend to favour the larger ISPs.

One mechanism for addressing costs to smaller ISPs may be through a third party filtering service provider solution that provides services to a number of small ISPs. The additional costs for the users of the third party service would be:

- The cost of a link to the third party host provider.
- Service fees and/or margins imposed by the third party host provider.

The pass-by filtering approach would be particularly suitable for this type of hosting arrangement.

Summary

An analysis of the costs associated with internet content filtering appears to suggest that the cost per user reduces with ISP size.

Other factors affecting cost include the filtering technology implemented by the ISP, the size of the ISP's customer base, the ISP's network configuration and other services supported by the filtering technology (e.g. additional filtering capability, traffic management capabilities, etc.).

