Network & Technology
Service Delivery & Applications

Blacklist Blocking Trial

Final Trial Report

Document Status:	*Final*
Document Issue:	*01.10.00*
Issue date:	*03/06/2009*
Issued by:	*Barrie Hall, General Manger, BigPond and SD&A Networks, Telstra Corporation Limited*

TABLE OF CONTENTS

1.1. Introduction

This trial evaluated hybrid DNS / Proxy Server based blacklisting blocking technology in a closed network environment.

No customers were involved in this trial.

The “closed environment” consists of an accurate replication of Telstra’s residential broadband Internet service (BigPond). This environment is used on a day-to-day basis for testing BigPond network and system changes before deployment into production.

The actual environment was Telstra’s BigPond test environment. This environment was shared with other projects during the trial period.

Telstra engaged Nominum to assist with testing. Nominum’s Vantio MDR and Centris product was used for this trial.

Evaluation of this blocking technology was executed against a defined set of test criteria and a defined set of URL’s. These URL’s bear no relationship to the actual ACMA blacklist (they are just a random list of 10,000 URL’s).

1.2. Executive Summary

The Blacklist blocking trial was conducted during April, 2009 following acquisition and installation of software during March. A Nominum representative was present on-site in Sydney to assist with the configuration of Nominum's software.

The solution trialled was a "DNS plus Proxy" Hybrid blocking solution. This type of solution uses a "poisoned" DNS server, loaded with a Black List, to redirect users to a proxy server if the user requests access to a domain listed on the Black List.

The trial progressed well and produced results which showed that the solution chosen was sound and operated correctly.

Key Points:

- Filtering was 100% accurate (no under or over blocking was seen).
- DNS response times were mildly degraded with a list size of 10,000 entries. However, this degradation was found to be imperceptible compared to overall all page load times. This figure of 10,000 was chosen based on expert opinion
- Customer experience was not noticeably degraded.
- It was noted that it is possible for the solution to **fail** if pages from a heavily trafficked site are added to the blacklist. This is due to volume limitations of a typical proxy server. These sites serve video content to end users. The volume of video traffic would be likely to overwhelm a proxy server.
- The solution keeps the Blacklist in encrypted form at all times. It was shown that the list could be managed and automatically updated by a third party, such as ACMA without the ISP or its staff ever having to see the list.

The trial did not test the solution's ability to prevent circumvention as these are seen as being in plain sight and well documented.

Specifically, the solution will fail to block content in these two situations:

- User configures their web browser to use a proxy server external to the ISP
- User configures any type of VPN/Tunnelling software on their PC

In any planned deployment Telstra would prevent customer access to non-Telstra DNS, thus eliminating an alternate DNS circumvention.

1.3. Trial Environment

The evaluation environment (known as the BigPond UAT facility) is located at 151 Clarence St Sydney.

This environment allows the end-to-end testing of each BigPond access method (ADSL, BigPond Wireless, BigPond Cable and BigPond Dial). We are not able to test the BigPond Satellite service in this environment.

Given the access-agnostic nature of the BigPond network, trials will be conducted using test ADSL services and Ethernet services. (We would fully expect the test results to be identical if the test user was using a Cable service or a Wireless service).

BigPond Pre-Production Test Environment

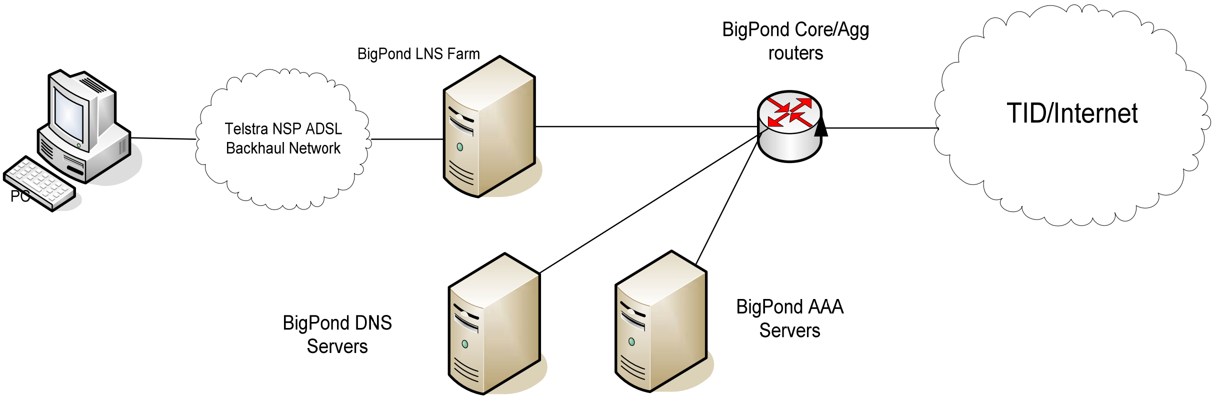

The above environment was augmented with a proxy server and a Blacklist manager server for the purposes of the trial as shown in the diagram below.

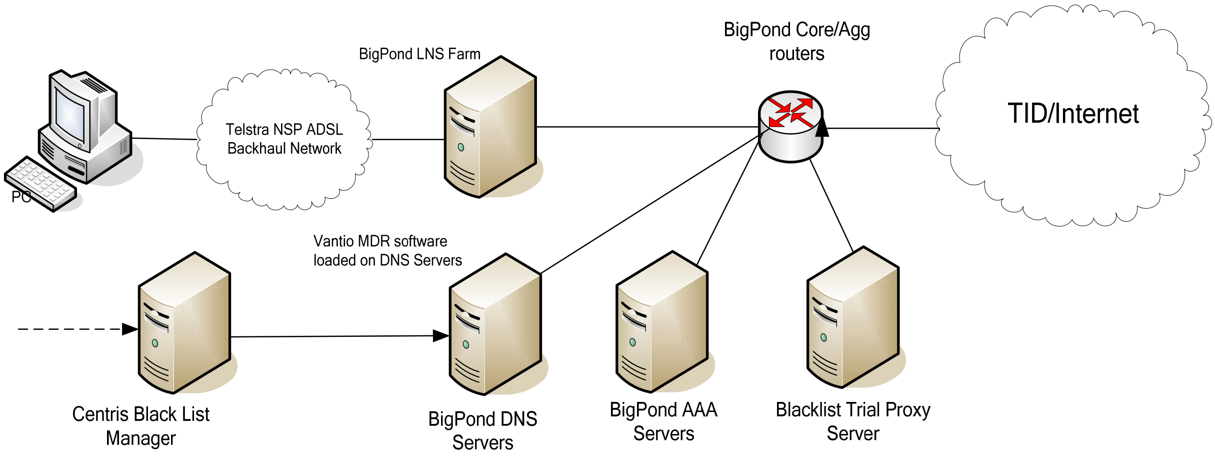

BigPond Pre-Production Test Environment configured for Black List Trial

BigPond LNS Farm
BigPond Core/Agg routers
TID/Internet
Telstra NSP ADSL Backhaul Network
PC
Vantio MDR software loaded on DNS Servers
Centris Black List Manager
BigPond DNS Servers
BigPond AAA Servers
Blacklist Trial Proxy Server

Note that all equipment types, model numbers and software versions present in this test environment are identical to those currently in production.

1.4. Evaluation Methodology

Telstra had originally planned to evaluate two blacklist blocking technologies in the environment described above. Both solutions are ‘hybrid’ in nature and both are designed to address concerns regarding “overblocking” without significantly degrading network performance or customer experience.

After some review of the available industry data it was decided that testing an IP Plus Proxy solution wouldn’t yield value due to a recent rising phenomenon known as “Fast Fluxing”.

> **Fast Fluxing** is the term used to describe the rapid movement of Internet content from one IP address to another in order to avoid IP based blacklisting. This technique is now widely used by individuals who are attempting to propagate objectionable material, especially over the last 6 months. This has become a popular method of distributing objectionable content because the content is actually hosted on PCs which have been hijacked. The IP address of a hijacked PC typically changes every time the (unsuspecting) Internet user connects to the Internet (which makes IP blocking ineffective).

DNS+Proxy based filtering solutions are totally immune to Fast Fluxing, hence the decision to pursue this solution.

The solution trialled was forecast to have engineering limits in terms of the number of URL’s which are able to be blocked. This limit is estimated to be 10,000 URL’s. The trial used this figure as the upper limit for the list.

The actual ACMA Blacklist was not used for testing purposes. Instead, a list of 10,000 URL’s provided by the vendor was used. The actual nature of the URL is immaterial in measuring its ability to be blocked.

1.5. Test Case Outline

In evaluating the proposed solution, the primary aim was to determine how the solution performs in terms of the following criteria:

- Overblocking
- Underblocking
- Network performance degradation

Given that the circumventions to this solution well known and in plain sight, **circumvention testing was not performed**.

Telstra did not perform any specific performance testing on the proxy server used. This element’s capacity and performance limitations are already well understood.

Given the nature of the solution trialled, testing focused on the following on the following specific areas:

- **DNS Query times**, with and without Blacklist blocking enabled
- **Overall user experience**, with and without Blocking enabled

1.6. Detailed Solution Description

DNS Plus Proxy Blacklist Blocking

This technique achieves a blocking action via the following mechanism:

1. The user's browser requests an IP address from a poisoned DNS server by passing the domain portion of the URL to the server.
2. The Poisoned DNS Server examines the domain. If the domain appears in the Blacklist, the DNS server returns the IP address of the ***proxy server***. If the domain ***does not*** appear on the Blacklist, the DNS query proceeds as usual and the user's PC is allowed to fetch the URL from the Internet.
3. The user's browser sends an http GET request to the IP address returned by the DNS server above. If the domain portion of the requested URL is in Blacklist, this request will be forwarded to the proxy server.
4. The proxy server accepts the http GET request from the user's browser and then performs a high speed database lookup (Centris Component), comparing the requested URL with a list of up to 10,000 URL's (The Blacklist).
5. If the URL ***does not*** appear in the proxy server database, the proxy server retrieves the requested web content from the Internet and returns it to the user's PC.
6. If the requested URL does appear in the Proxy Server's Database, the user's request is discarded (Blocked). The user's experience will as if the URL doesn't exist (no indication to the user that the page has been specifically blocked).

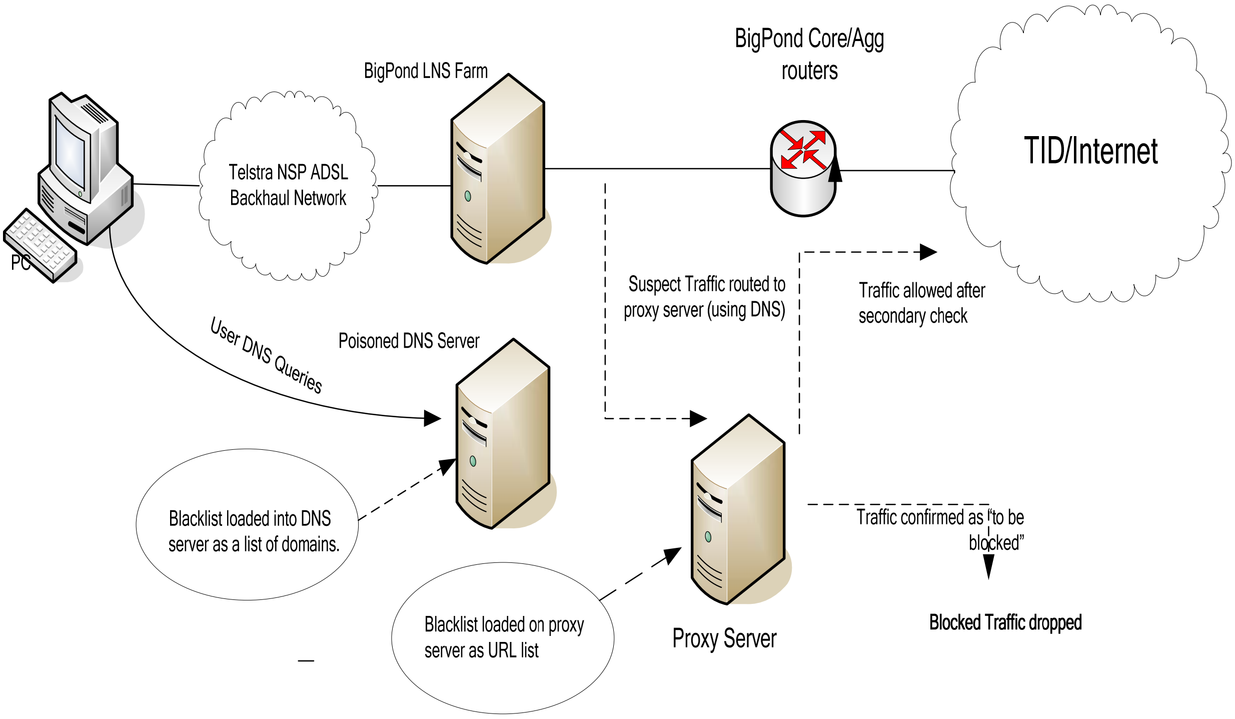
BigPond Core/Agg routers
BigPond LNS Farm
TID/Internet
Telstra NSP ADSL Backhaul Network
PC
Suspect Traffic routed to proxy server (using DNS)
Traffic allowed after secondary check
User DNS Queries
Poisoned DNS Server
Blacklist loaded into DNS server as a list of domains.
Traffic confirmed as "to be blocked"
Blocked Traffic dropped
Blacklist loaded on proxy server as URL list
Proxy Server

1.6.1.1. Evaluation Infrastructure and software

The relevant infrastructure for this solution consists of the following:

I. ***One SUN V240 Solaris servers running Nominum's Vantio MDR Poisoned DNS software and Nominum's Centris management software.***
II. ***One SUN V240 Solaris servers running Squid V3 proxy server software.***
III. ***Cisco Catalyst 6509 switch chassis***

This equipment is closely aligned with current BigPond production environments.

1.7. Test Results

Test results were as follows:

1.7.1. Blocking accuracy

Initial tests were performed to assess the accuracy of blocking. After environment configuration, testing showed that **blocking accuracy was 100%**. This result was expected based on the deterministic nature of the software. (The software and solution worked as designed).

Given this result, further testing focused on performance, specifically DNS performance under load.

1.7.2. DNS Performance

DNS Query performance is the key to the performance of the solution trialled. The results are tabulated below:

“Cold Cache” is a term which refers to the state of a DNS server which has been started from cold and has no cached entries, meaning that the DNS server has to make further DNS queries to answer a query itself.

“Warm Cache” refers to a DNS server which has seen significant volumes of diverse queries over time and is generally able to respond to a query without referral to other DNS servers.

Maximum Queries per Second

	Empty Blacklist	10,000 Entries
Cold Cache	12,900 QPS	12,750 QPS
Warm Cache	25,000 QPS	20,000 QPS

DNS Query Latency

	Empty Blacklist	10,000 Entries
Cold Cache	5ms	5.5ms
Warm Cache	1.7ms	2.2ms

QPS Degradation	20% when warm
Latency change	28% when warm

These results show that DNS performance does degrade with a Blacklist of 10,000 entries, however, the degradation does not impact end user experience due to extremely low cache latency figure.

With a Blacklist of 10,000 entries, we are able to see that response to a DNS query increases from 1.7ms to 2.2ms.

Allowing for a typical figure of 10 DNS lookups on a moderately busy web page, user experience (for a page not on the blacklist) would be an additional load time of 5ms. This would not be noticeable to the end user. In addition, these queries are often issued in parallel, further reducing impact to page load times.

1.7.3. End User Experience

This testing was performed using a modern medium specification Laptop computer connected to the test environment using a standard 1.5Mb/s ADSL service and, at times, via a direct Ethernet connection to the test environment.

The outcome of this testing showed **no measurable difference** between browsing via the Blacklist blocking solution and browsing using normal DNS infrastructure.

This is as expected when considering typical “busy page” load times are of the order of 5-10 seconds and DNS overheads imposed by the solution are of the order of an additional 5ms.

1.8. Analysis and Discussion

Analysis of the results revolves around the following areas:

- Accuracy of blocking action
- Performance and capacity
- User experience

1.8.1. Accuracy

The test results show that blocking action was 100% accurate with the overload caveats noted below. This is as expected, based on the solution design and software specification. No software defects were noted during the trial.

1.8.2. Performance and Capacity

Two key performance and capacity metrics are relevant to the solution trialled. These are sensitivities specific to the solution trialled. (Alternate solutions have different sensitivities).

- DNS Query performance
- Proxy server performance

DNS Performance

It would be expected to see some difference in DNS Query performance between a DNS server running standard software and a DNS server running blocking software. This is because the blocking function incurs a database lookup on the DNS server.

It was noted that DNS performance suffered some degradation, however, as noted in the results section, the degradation would add < 6ms to page load times for pages not on the blacklist.

Proxy Server Performance

Proxy Server performance is the most “sensitive” component in the trialled solution. It is obvious by inspection that the proposed proxy server hardware would **not be able to continue functioning** if the Blocking software were to divert a substantial percentage of Internet traffic via the proxy server.

Typical proxy server hardware (SUN Solaris Servers) has 1GB/s Ethernet interfaces while our core networks operate at up to 40Gb/s. In addition, the processing capacity of these servers would be unlikely to be able to process traffic at greater than 1Gb/s.

During the trial, we were not able to overload the proxy server used for testing but as noted above, there are clear upper limits to capability of this part of the solution.

There are two key determinants to how much traffic is directed via the proxy server:

1. Exact content of the Blacklist

2. Internet traffic patterns

The content of any blacklist is assumed to be determined by relevant government authorities (ACMA). Internet traffic patterns are largely predictable.

If the proxy server fails due to overload, the blocking system fails as follows:

URL's which have a domain portion which is on the blacklist but a RHS which isn't would be inadvertently blocked because the proxy server would be unable to process traffic.

Example:

If the URL somesite.com/really-bad-stuff appears in the blacklist, all somesite traffic is directed via the proxy server.

Video clips from high traffic sites are very popular with typical Internet users, accounting for up to 10% of traffic. If any content from sites distributing these video clips were to appear on the blacklist the blocking solution would fail because 10% of 40Gb/s of traffic is greater than the 1Gb/s capacity of a proxy server.

This mode of failure is triggered by large volumes of data being directed via the proxy server. Popular video content sites (both now and future) are seen as the major problem here.

A mitigation to this problem is provided by Nominum's software in the form of a "white list". This feature is able to prevent an operator from adding "highly trafficked" domains to the blacklist.

Of course, this feature is optional and assumes a policy of not needing to block URL's existing on highly trafficked domains.

The alternative to this approach is to block the entire domain and accept the end user consequences of over-blocking.

It should also be noted http requests which are routed via the proxy server all appear to originate from the same IP address. This can cause problems with some popular websites which place limitations on the number of concurrent sessions originating from the same IP address.

1.9. Summary

Telstra trialled a DNS+Proxy Blacklist blocking system in a model environment representative of Telstra's BigPond network PoP's.

A Blacklist with 10,000 entries was tested, and it was noted that the blocking action was 100% accurate.

The solution trialled utilised a proxy server to prevent overblocking. The proxy server retrieves content which would otherwise be overblocked in a DNS only Blocking Solution.

Available proxy server's have upper performance limits, typically 1-2Gb/s. This limitation becomes significant if the Blacklist contains pages from very popular web sites.

Telstra believe that the solution trialled would be satisfactory and fit for purpose in Telstra's production environment with the following caveats:

- The size of the Blacklist doesn't exceed 10,000 entries
- The Blacklist doesn't contains pages from "heavily trafficked" websites

The solution trialled maintained the content of the list in encrypted form at all times. Given the sensitivity of any Blacklist, this is seen as an important requirement in any Blacklist blocking system.