

# Ad and Tracker Blocking with Domains

**Peter Lowe**

[pgl@yoyo.org](mailto:pgl@yoyo.org) | <https://pgl.yoyo.org/adservers/>

November 2024

# Introduction: Me

---

## Who I am

- DNS Abuse Ambassador for [FIRST](#)  
Co-chair of the [DNS Abuse SIG](#)
- Accidental DNS person
- Cybersecurity enthusiast
- Closet privacy amateur
- No current employer



# Introduction: Blocking Ads and Trackers with Domains

---

- Control the DNS server
- Return fake results for ad and tracker domains
- ... profit
- But there's a bit more to it

# Introduction: How I Started My List

---

- Heard about the `/etc/hosts` trick
- Couldn't be bothered to maintain a huge list of individual hostnames
- Realised local nameservers could be authoritative for any domain

# Introduction: How I Started My List

---

- Threw up a text file on my website
- And the rest is history
- Now have a horrible pile of spaghetti code and a few tables in MySQL

# Blocking

# Blocking With Domains: Methods

---

- Via DNS
- Browser extension
- ProxyPAC for the browser
- hosts file
- Windows registry entries
- Little Snitch, Squid, or some other proxy

# Blocking With Domains: Downsides

---

- It's a hammer, and everything is a nail
- Can block legitimate content
- Have to be careful about the policy of the list you're using
- Need an easy way to exclude entries
- Can't do URL-level blocking

# Blocking With Domains: Upsides

---

- Can do network-level blocking
- Useful for all devices on the network
- Don't have to maintain detailed lists of URLs
- DNS blocking can be done various ways
  - Invalid IPs
  - NXDOMAINs
  - Localhost
  - "Block server"
- Can have combination of all of these
- Logging can give insight into failed requests

# My List

# My List

---

- Started back in 1997 (apparently)
- ~3.5k domains, 9K IPs (7k valid)
- 33 formats supported
- Various [format options](#) available - change IP, skip hostnames, date ranges, plaintext, etc.
- Weekly cleanup script keeps things tidy

# October 2001

INTERNET ARCHIVE <http://yoyo.org/~pgl/adserver/> Go SEP OCT JUN 06 2000 2001 2002 About this capture

Wayback Machine 97 captures  
6 Oct 2001 - 27 May 2004

[back home](#)

last updated: Fri, 5 Oct 2001 16:05:17 +0100

I run a local nameserver and regularly add the hostname of any ad server I come across as a master zone, so ad images can't load when I visit web pages. I thought I'd put up the list of ad servers just in case anyone else wanted to use it to block ads. if you know of any other lists like this, or any ad servers I don't have listed, please [mail me](#), or type it in here:

 

thanks.

I use this list with a local nameserver, but [you can use this list with windows to block ads](#). if you want to do this, it will probably be easier if you view this list in [hosts file](#) format.

for those that don't want to bother reading the above link, basically, all you do is copy the list (in [hosts file format](#)) to your windows directory as a file called "hosts" that's pretty much it (you might have to restart, actually). if you already have a hosts file, just add this list to the end of it. what happens is that every time your machine sees a hostname, it checks the hosts file to see what the ip address of that hostname is. because ad servers are listed as being your local machine (127.0.0.1 is always the local machine), the ads don't load (because your browser tries to load them from your machine).

**update:** I found [another site](#) listing ad servers. it has a good explanation of how to use the list to block ads using the hosts file in windows and by using the proxy server settings in netscape & ie. I've also stolen some of the hosts that weren't in this list before, after weeding out the nonexistent ones. the site hasn't been updated since 20th feb (at the time of writing), so looks like it might be dead or something.

since it looks like there's a slight different between that list and mine, I thought I'd explain how I choose what hosts / domains go in my named.conf, basically, for any ad company that I know of, like [doubleclick](#), I put in their whole domain, so I don't have to bother adding all the hosts as they pop up. also, any stats services go in there as well, because I don't particularly want my browsing habits being recorded and analysed by people I don't know. other than that, if I see an ad I'll find out what server it's coming from, and add the server name to the list if it's obviously hosted on a dedicated ad server (eg, the hostname is "ads.example.com").

a script checks my named configuration file every ten minutes and updates the list this page uses if new zones have been added. another script runs every night checking to see if the hosts still exist and mails me if they don't.

**update:** [another one!](#) woah, this one has shedloads of hostnames. they could probably make it easier to maintain if they just used second level domains instead of listing each ad server, but hey, not my site. probably not worth maintaining this page any more, but it's automatic so I might as well leave it in place. I found this site whilst browsing [dmoz's list of ad filtering pages](#).

**update:** [they just keep on coming](#).

here's the list (245 total). you can also view this list as a [hosts file](#).

```
247media.com
a.r.tv.com
a1576.g.akamai.net
abcnews-footprint.net
actionsplash.com
ad-flow.com
ad-up.com
ad.37.com
ad.eurosport.com
ad.infoseek.com
ad.iwin.com
ad.kvantum.cz
ad.linkexchange.com
ad.preferences.com
ad.seznam.cz
ad2.atlas.cz
ad2.linx.cz
ad2.seznam.cz
ad4.atlas.cz
adbot.theonion.com
adbureau.net
adbutler.com
adclick.com
adclick.gamespy.com
adclub.net
adcontent.gamespy.com
adcontroller.unicast.com
adcreative.tribuneinteractive.com
adfu.blockstackers.com
adi.mainichi.co.jp
adimages.been.com
adimg.egroups.com
```

# My List

---

- Probably around 50m users? (Directly and indirectly)
- Very rough guess
- Random day in July had about 57m downloads
- Used by (or available for) uBlock Origin, Opera, AdBlock Plus, Pi-Hole, AdGuard, HostsMan, ...
- Most new entries are trackers
- Slowly turning into a privacy list more than anything

# My List

---

- Started life as just a text file
- “Last updated” was just a stat call on the file
- Excited at first, but then found others doing the same
- Still excited
- Then uBlock Origin included my list a default option
- Weekly cleanup script is a big reason it has a good reputation
- Haven’t changed the main page for about 15 years

Check of ALL ad server hostnames

=====

Stats:

-----

Time script started: Fri, 25 Oct 2024 11:59:01

Time script finished: Fri, 25 Oct 2024 12:05:15

Total running time: 6 mins and 14 secs

Total servers: 3,561

Total IP addresses: 9,141

Valid hostnames found: 3,542

Valid IP addresses found: 7,770

Invalid hostnames found: 19

Invalid IP addresses found: 1,371

Hostnames deleted: 3

IP addresses deleted: 0

Weekly block list report: 3,561 entries; 3,542 valid / 19 invalid; 3 deleted #ads #adblocking #blocklist #trackers  
#pglblocklistreport

List URL: <https://pgl.yoyo.org/as/serverlist.php>

Output:

-----

1-1ads.com: good - updating IPs

101com.com: good - updating IPs

100white.de: good - updating IPs

# My List: Policy

---

- Highest level domains only (i.e., doubleclick.net, not ad.doubleclick.net)
- Not only ad servers - spyware servers, hit tracking servers, and servers that enable these technologies
- Other nasty hostnames are included such as some malware domains etc.
- Checked by hand before being added when submitted by a user
- Automatically checked weekly to see if they're still valid

# My List: Options

---

- Show only trackers, or ad servers
- Different formats for different applications
- Change output MIME type
- Add custom domains
- Change date range
- Skip entries
- Append or prepend strings to entries

# My List: Backend

---

- Worst codebase I've ever created... basically the same from PHP 5 days
- I will never share the code, it's far too embarrassing
- Little admin area for managing entries and news etc.
- Very basic DB structure for domains, and reverse IPs
- Custom "search" in browser to add domains

## Ugh

```
$useip = isset($useip)
? htmlentities(stripslashes(preg_replace('/[^\0-9.]\/', '', $useip)))
: '127.0.0.1';

// for iptables + ipchains
$device = isset($_GET['device']) ? htmlentities(stripslashes($_GET['device'])) : 'eth0';

if (empty($zonefilename)) {
    $zonefilename = '';
}

if (!$zonefilename = stripslashes($zonefilename))
    $zonefilename = 'null.zone.file';

header('Last-Modified: ' . date('D, d M Y H:i:s', ADSEVERERS_LASTUPDATED) . ' GMT');
header('Cache-Control: max-age=3600'); // 1 hour
header('Vary: Accept-Encoding'); // so non-gzip clients don't get gzipped versions

if ($hostformat == 'glimmerblocker')
    $hostformat = 'glimmerblocker.xml';

if (empty($mimetype)) {
}

if (in_array($hostformat, array('msie-nocookies-xml', 'pgl-xml', 'nohtml', 'ms-isa-server-2004-xml', 'ms-isa-serv
) {
    unset($showintro);

    $mimetype = 'plaintext';

    header('Content-Type: text/plain');
    header('Last-Modified: ' . adservers_lastupdated($timestampfile, true));
}
```

## Basic DB structure

Field	Type	Null	Key	Default	Extra
hostname	varchar(250)	YES	UNI	NULL	
valid	tinyint(3) unsigned	YES	MUL	1	
date_added	int(10) unsigned	YES	MUL	NULL	
ip	varchar(50)	YES		NULL	
last_checked	int(10) unsigned	YES	MUL	NULL	
failed_checks	int(10) unsigned	YES		0	
forced	tinyint(3) unsigned	YES		NULL	
hidden	tinyint(1)	YES		0	
comment	text	YES		NULL	

**9 rows in set (0.003 sec)**

# My List: Issues

---

- Constantly been a problem for my gracious server admins (Many, many thanks to Matt and Neil!)
- Load issues (forcing implementation of caching)
- ... although Matt said it keeps them busy!
- Various security issues (shout out to **iampytest** for excellent reporting!)
- Problems with PHP versions

# My List: Other Issues

---

- People get annoyed
- Being forced to remove entries
- Controversial entries (`t.co`, `go.icann.org`, ...)
- People don't realise I just maintain the list

# My List: People

---

- Most new entries are submissions now
- Most emails I get are delisting requests
- Most of the time people are understanding after I explain
- Still find it hard to do bulk additions

# Angry People

*"Why do I have to see your name in my computer. All I want us some sanity in my life with no bloody adverts pissing me off, my wife and I are both pensioners and all we want is to be left alone. So get yourself off of my computer."*

## Angry People

*“i dont understand this shit. i just know i cand go to gamestop.com unless i use a different browser.”*

## Angry People

*“Please get this off my  
computer!”*

# Understanding People

*“ Thank you Peter for hearing me out. One of the information is lawtrack.me on the email from my psychic. The other ones are very confusing. But I do thank you for helping. God Bless and have a terrific day.Aloha”*

# Understanding People

*“Thank you very much for the courtesy of your reply. I totally misunderstood the situation. I thought that you were the proprietor of the system I signed onto that resulted in the limination of a great number, but not all, of advertisements from my screen when I went online. That is clearly not the case and I apologize for taking up your time. It does sound as if you are involved with some project for the betterment of society and for that I thank you.”*

# My List: The Future

---

- Detect trackers by analysing passive DNS data (but need access to some passive DNS data first...)
- LLM trained on DNS usage that can recognise abnormal behaviour? (Bit of a daydream)
- Easier way to do bulk additions
- GitHub?

# Questions?

---

## Peter Lowe

- [pgl@yoyo.org](mailto:pgl@yoyo.org)
- <https://infosec.exchange/@pgl>, <https://twitter.com/pgl>
- <https://linkedin.com/in/peterlowe>

## Resources

- <https://pgl.yoyo.org/adservers/> - main page
- <https://pgl.yoyo.org/as/serverlist.php> - list of domains
- <https://pgl.yoyo.org/as/iplist.php> - list of IPs
- <https://pgl.yoyo.org/as/formats.php> - formats