Blocking Web Sites

Blocking Web Sites

There are several options for blocking websites with pfSense® software, some of which are described on this article. This is not an exact science, but these solutions typically function well enough for a majority of use cases.

See also

The pfBlockerNG package (pfBlocker-NG Package) offers mechanisms which can be useful in this area, such as DNSBL, geographic IP address blocking, and automation of AS lookups.

Using DNS

If the built in DNS Resolver or Forwarder are active an override can be entered there to resolve the unwanted website to an invalid IP address such as `127.0.0.1`.

Warning

Do not use DNS override functionality as the only means of blocking access to sites.

Blocking via DNS requires that local clients utilize the firewall as their only DNS source. See Redirecting Client DNS Requests and Blocking External Client DNS Queries for suggestions on ensuring clients get their DNS responses from the firewall. It will stop non-technical users, but it is easy to circumvent for those with more technical aptitude.

With the DNS Resolver, additional methods are possible via custom options.

This first example will prevent any host under the given zone from being resolved by clients:

```
server:
local-zone: "movie.edu" static
```

When the firewall enforces DNS resolution in this way, the firewall must also force clients to resolve DNS using the firewall. Otherwise, clients could bypass the restrictions by using alternate DNS

servers. See Redirecting Client DNS Requests for details.

This can be limited in scope using custom views. This example is similar to the above, but only blocks access for `10.6.0.100`:

```
server:
access-control-view: 10.6.0.100/32 blocksites

view:
name: "blocksites"
local-zone: "movie.edu" static
```

Using Firewall Rules

If a website rarely changes IP addresses, then it can be blocked by an alias. Create an alias containing its IP addresses and then use this alias in firewall rules.

Warning

This is not a feasible solution for sites that return low TTLs and spread the load across many servers and/or datacenters, such as Google and similar large sites. Most small to mid sized websites can be effectively blocked using this method as they rarely change IP addresses.

A hostname can also be inside a network alias. The firewall will resolve the hostname periodically and update the alias as needed. This is more effective than manually looking up the IP addresses, but will still fall short if the site returns DNS records in a way that changes rapidly or randomizes results from a pool of servers on each query, which is common for large sites.

Another option is finding all of the IP subnet allocations for a site. Create an alias with those networks and block traffic to those destinations. This is especially useful with sites such as Facebook that spread large amounts of IP space, but are constrained within a few net blocks. Using regional registry sites such as ARIN can help track down those networks. For example, all of the networks used by Facebook in the region covered by ARIN can be found at http://whois.arin.net/rest/org/THEFA-3.html under “Related Networks”. Companies may have other addresses in different regions, so check other regional sites as well, such as RIPE, APNIC, etc.

As an alternative to looking up the IP blocks manually, locate the BGP Autonomous System (AS) number for the target company by doing a `whois` lookup on one of their IP addresses. For example, the AS number for Facebook is `AS32934` and the following command will locate all of their allocations:

```
# whois -h whois.radb.net -- '-i origin AS32934' | awk '/^route:/ {print $2;}' | sort | uniq
```

Copy the results of that command into a new alias and it will cover all of their currently allocated networks. Check the results periodically for updates.

Using a Proxy

In modern environments a client proxy is not effective. HTTPS can sometimes be filtered via peek/splice to inspect SNI and similar aspects of connections, but even that fails with modern security practices like encrypted SNI. Using proxies for these tasks is no longer a recommended practice.

Prevent Bypassing Restrictions

With any of the above methods, there are many ways to get around the defined blocks. The easiest and likely most prevalent is using any number of proxy websites. Finding and blocking all of these individually and keeping the list up to date is impossible. The best way to ensure these sites are not accessible is using an external proxy or content filtering capable of blocking by category.

To further maintain control, use a restrictive egress ruleset and only allow traffic out to specific services and/or hosts. For example, only allow DNS access to the firewall or the DNS servers specifically used for LAN clients (Redirecting Client DNS Requests). Also, if a proxy is in use on the network, make sure to disallow direct access to HTTP and HTTPS through the firewall and only allow traffic to and/or from the proxy server.

Revision #1
Created 29 December 2023 03:34:32 by ColtM
Updated 7 August 2024 23:24:39 by ColtM