Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Domain Name System security for domain resolvers

First published: January 2021
Last updated: October 2021

Introduction

This publication provides information on Domain Name System (DNS) security for recursive resolution servers, as well as mitigation strategies to reduce the risk of DNS resolver subversion or compromise. Organisations should implement the mitigation strategies in this publication to improve the security of their DNS infrastructure.

DNS is a hierarchical naming system built on a distributed database for resources connected to the internet. DNS maps human-readable domain names to their associated Internet Protocol (IP) addresses.

A DNS resolver (recursive resolver) is a server that discovers a host name by querying the DNS server hierarchy to match and provide an IP address for a web browser to connect to.

Additional information on securing domains and authoritative name servers is available in the *Domain Name System security for domain owners* publication.

Background

A DNS resolver is a trusted agent between the client and the DNS hierarchy for locating an IP address. Compromising a DNS resolver can allow malicious actors to redirect client connections to malicious websites.

The common attacks involving DNS resolvers are:

- **DNS resolver hijacking:** Takeover of a DNS resolver by malicious actors.
- **DNS spoofing (or cache poisoning):** Subverting DNS processes to redirect users to malicious websites.
- **DNS reflection and amplification attacks:** Using DNS resolvers to perform denial-of-service attacks.
- **Surveillance of DNS requests:** Recording DNS requests for the purposes of intelligence gathering.
- **DNS as a malware covert channel:** Covert methods used by malicious actors to maintain command and control of malware infections and exfiltrate data.
- **Misuse of encrypted DNS:** Defeating the advantage of DNS monitoring in detecting malicious activity.

Recursive resolution

A DNS resolver (also known as a recursive resolver or recursive name server) searches for requested domains by querying the DNS hierarchy. As an example, when the DNS resolver receives a request for ‘example.com.au.’, it starts

by asking the root server, then moves through the hierarchy (i.e. the Top Level Domain ('.au') then second level ('.com'), until it reaches the authoritative server for its request ('example')). The authoritative server then provides the IP address and port information for the requested service.

DNS caching

A complete end-to-end DNS resolution can be slow. For efficiency, DNS resolvers store results in a cache (caching) for a nominated period (time to live). However, this methodology is vulnerable to cache poisoning attacks (see *DNS spoofing (or cache poisoning)*).

DNS forwarder

A DNS forwarder is a server configured to forward requests to a DNS resolver for resolution. Both consumer and commercial routers often perform this role. DNS forwarders support organisations by providing a point to perform DNS logging, caching of outbound DNS requests and filtering attempts to reach unauthorised domains.

DNS resolution server attacks and mitigations

DNS resolver hijacking

A DNS hijacking attack occurs when malicious actors take over the DNS resolver or redirects a client to a malicious DNS resolver. The hijacked DNS resolver subverts the normal DNS resolution process by causing the DNS resolver to return an incorrect address.

DNS resolvers are prominent on internal and external networks. Organisations should take steps to ensure their own DNS resolvers are secure against trusted insider or external attacks, for example, by:

- **Keeping DNS software up to date:** Most historical threats to DNS have existing mitigations in the current DNS software releases. Keep DNS services up to date and review the developer's configuration recommendations regularly, as they are likely to change over time as new threats emerge.
- **Using multi-factor authentication:** Where possible, administration should be done through methods that make use of MFA.
- **Restricting administration privileges:** As DNS is a core network service, access should be limited to administrators who have a reason to access DNS services and are trained to operate them.
- **Maintaining the server's operating system:** Update the server's operating system to the latest version. While most updates for systems in internal networks can be scheduled, the exposure of DNS resolvers merits special priority for operating system patching.
- **Hardening the server's operating system:** Consider the use of hardened operating system versions that put additional limitations on software and include specialised mitigations. The Center for Internet Security and the National Institute for Standards and Technology offer guidance for operating system hardening.

DNS spoofing

A DNS spoofing attack causes a DNS resolver to redirect traffic to servers controlled by malicious actors. This process includes cache poisoning whereby:

- an attacking system requests an IP address from a DNS resolver
- the DNS resolver does not have the answer, so sends a query for an authoritative answer
- the attacking system floods the DNS resolver with its own seemingly authoritative responses which the DNS resolver might accept as a legitimate response
- if accepted, future requests to the DNS resolver will return the cached IP address of the malicious server.

While newer DNS resolvers now have safeguards to reduce the likelihood of cache poisoning occurring, DNS Security Extensions (DNSSEC) query response authentication is a more effective prevention.

Transaction identifiers and randomised source ports

To restrict DNS spoofing, current versions of DNS software use randomised transaction identifiers and replace the default UDP/TCP source port 53 with a randomised port. As a result, the difficulty of tricking a DNS resolver is increased as an attacking system now needs to guess both the source port and transaction identifier pair through packet flooding.

DNS resolver source port randomisation poses an additional challenge for gateway design as gateway devices will, by default, apply Network Address Translation to DNS traffic. This modifies the source ports, disrupting DNS resolver source port randomisation and exposing the DNS resolver to spoofing. To address this, organisations should place their DNS resolver in a demilitarized zone with a public IP address, or configure gateway devices to assign a public IP address to the DNS resolver.

Authenticating DNS addresses through DNSSEC

DNSSEC is an extension for DNS which provides cryptographic integrity and a certified chain of trust. This allows name servers to prove that they are the authoritative server for the zone and that their responses have not been tampered with.

DNSSEC provides two extra records in each DNS response, a cryptographic signature to verify the validity of the DNS record and a second cryptographic signature to validate the DNS server. The second signature is validated by the DNS servers above it in the DNS hierarchy, which in turn has a signature validated by a higher DNS server. The root DNS zone's public key information is verified through a formal key signing ceremony.

This process means that when a client requests the address of a web server, they receive a response they can independently verify. Provided a client system is configured to use a DNS resolver with DNSSEC validation enabled, DNSSEC can prevent impersonation and cache poisoning attacks. Organisations should configure their DNS resolvers to validate DNSSEC where possible.

This process is similar to how Hypertext Transport Protocol Secure (HTTPS) is validated with Certification Authorities and the root signing keys used by web browsers. DNSSEC requires additional DNS requests to validate, but these responses are cached in the same way as DNS queries to keep DNSSEC's overheads to a minimum.

DNS reflection and amplification attacks

A DNS reflection and amplification attack is a variation of a denial-of-service attack using a large volume of DNS resolver responses to make a target inaccessible. Reflection attacks send a request to DNS resolvers with responses directed to the target's IP address. Amplification attacks rely on sending small queries that result in large responses that overwhelm a target server.

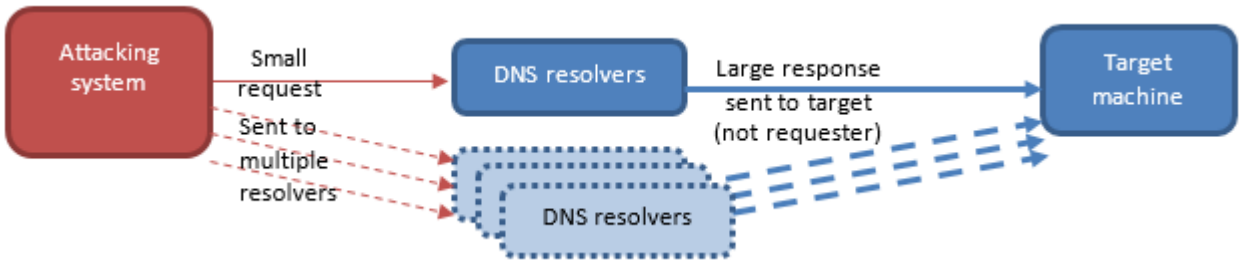

The process of reflection and amplification as seen above:

- malicious actors create a DNS packet with a spoofed origin indicating the target as the requestor (the packet specifically requests a reply of the maximum size allowed by DNS (i.e. 3876 bytes of data as per the Internet Engineering Task Force [IETF]'s *Request for Comments [RFC] 6894*)
- the DNS resolver treats it as a normal request and returns a response of the requested size (significantly larger than the request size) to the target
- this request can be sent to multiple DNS resolvers that will respond similarly
- the response packets are sent to the target exceeding the available bandwidth, exhausting the target's capacity to process or respond to incoming traffic.

Current versions of DNS servers restrict the use of DNS reflection and amplification attacks by limiting the size of the DNS response based on the size of the request. This limits the amplification that is possible through DNS reflection, but does not remove the potential for it to be used.

Surveillance of DNS requests

Surveillance of DNS request occurs when a third party intercepts DNS requests to monitor an organisation's web connections and uses these to infer what the organisation may be planning or doing.

DNS queries are transmitted unencrypted and can be read by any intermediary between a device and its DNS resolver. This may be necessary to monitor DNS usage for inappropriate activity on internal networks, but DNS traffic will eventually reach external networks where DNS surveillance is possible.

DNS surveillance permits malicious actors with the ability to monitor networks to read DNS requests and gradually build a record of what systems an organisation regularly interacts with. This might allow malicious actors to discern the organisation's strategic goals, indicators of internal difficulty or what might currently be under development.

DNS Query Name (QNAME) minimisation

QNAME minimisation (IETF RFC 9516) limits the query length to what is necessary for each step in the recursive resolution process. For example, instead of sending 'mx.example.com.au' to a root server, QNAME minimisation means:

- the root server is queried for .au
- the .au name server is queried for .com
- the .com name server is queried for .example

- only the authoritative name server for example.gov.au knows the mx server is being queried.

The interception of a single query no longer provides malicious actors with complete knowledge of the intended connection. This is useful in the event that malicious actors have control over a DNS server contacted during the resolution process.

DNS over TLS / DNS over HTTPS

If an organisation has opted to use an external DNS provider, they can provide cryptographic confidentiality by implementing DNS over TLS (DoT) or DNS over HTTPS (DoH).

DoT encapsulates DNS within Transport Layer Security (TLS). This is similar to how the Hypertext Transfer Protocol can be encapsulated within TLS to make HTTPS. In networks where DoT is unsupported, DNS requests can be sent over HTTPS to provide a comparable level of security.

DoT currently has limited adoption, while DoH adoption has begun to climb with the implementation of DoH by web browsers, including Mozilla Firefox and Google Chrome. While an improvement to security for individual users, organisations should ensure that web browsers are configured to disable DoH within their own networks for monitoring purposes, and implement DNS request encryption for DNS queries that traverse external networks.

Organisations that utilise external DNS resolvers should consider using a DoT capable DNS forwarder to provide DoT for all outbound DNS resolver queries. Details on how to provision this service can be found on the *DNS Privacy Project* wiki.

DNS as a malware covert channel

Malicious actors may use an organisation's DNS resolver to direct communications out of an organisation's network and connect malware infections back to their command and control systems.

Monitor DNS requests for malicious activity

Malicious actors' malware can potentially use DNS to contact their infrastructure. Malware infections typically beacon to announce their existence and can use DNS resolution to locate its command and control server. Malicious actors can use this to create resilient infrastructure that can be moved between DNS or IP addresses to evade takedowns or blocking.

Monitoring an internal DNS resolver or DNS forwarder for DNS requests from beaconing malware is a useful way to locate undiscovered malicious activity on internal networks. Potentially, this monitoring could be used locate and isolate malware.

Where possible, organisations should ensure DNS queries only use a monitored DNS service and block attempts to use unauthorised DNS resolvers. This can disable malware reliant on malicious actors' control of DNS. With limited service paths it is easier for organisations to inspect DNS requests and monitor for attempts to contact suspect domains.

Data communications via DNS tunnelling

A common technique used by malicious actors that have infiltrated a network is to encapsulate their communications inside DNS queries. DNS is often not monitored, readily accessible throughout the network and by design the protocol will route the request to malicious actors' DNS infrastructure.

Organisations logging all DNS queries through a DNS resolver can potentially use pattern recognition to detect suspicious payloads or a large number of queries heading to a suspicious domain.

Misuse of encrypted DNS

Privacy-enabling technologies can introduce challenges for an organisation's security monitoring activities. While recommended for use in public networks, DNS encryption can prevent detection of DNS exfiltration and indications of malware infection.

Organisations should test whether they are still able to monitor their system's activity when DoT and DoH are used. DoT capable DNS forwarders and TLS termination proxies at an organisation's network boundary may aid controlling and scrutinising DoT and DoH transactions. Some organisations currently use TLS termination proxies to inspect HTTPS traffic. All web browsers on an organisation's network will require configuration to retain organisational oversight.

Use of DNS encryption by users

Incomplete DNS monitoring poses a potential risk of allowing users to bypass DNS-based security filters and content management. This can occur when users request information from external DNS resolvers that provide a DoH interface. Without a TLS termination proxy, it is impossible to monitor DoH DNS requests.

It is recommended to disable DoH via Group Policy in enterprise environments, however, rapid adoption is expected in Bring Your Own Device environments and on misconfigured or un-configured web browsers.

Adoption of DoH may result in the bypassing of content restrictions, including for offensive material and inappropriate use of organisational resources. This can also potentially evade attempts to block downloading of malware through DNS filtering.

Use of DNS encryption by malicious actors

DoH is being adopted by malware authors and is starting to feature in malware infections. As previously shown in the *DNS as a malware covert channel* section, DNS has the potential to be used in the command and control of malware infections and as an exfiltration vector. Monitoring DNS has been a useful tool for discovering existing infections, but this may become difficult to detect with the adoption of DoH.

Unmonitored DNS traffic increases the risk borne by organisations. Organisations should not permit end-to-end use of DoH and DoT without inspection at their gateway. Both DoH and DoT can be inspected in transit, through interception by TLS termination proxy or via a DNS forwarder respectively. Proxy filtering and firewall rules should be adjusted to ensure that traffic is either monitored or blocked.

Further information

Further information on DNS is available from the following sources:

- Cisco's *DNS Best Practices, Network Protections, and Attack Identification* website
- Cloudflare's *DNS security* website
- Digital Ocean's *An Introduction to Managing DNS* website
- Internet Corporation for Assigned Names and Numbers' *DNSSEC – What Is It and Why Is It Important?* website
- National Institute for Standards and Technology's Special Publication 800-81-2, *Secure Domain Name System (DNS) Deployment Guide*

- Asia Pacific Network Information Centre's *Do you have DNSSEC validation enabled?* and *DNSSEC validation revisited* blog posts
- Dan Kaminsky's *Black Ops 2018: It's the End of the Cache As We Know It* presentation
- Drew Hjelm's *A New Needle and Haystack: Detecting DNS over HTTPS Usage* whitepaper
- Paul Wouter's *Defending your DNS in a post-Kaminsky world* presentation.

Contact details

If you have any questions regarding this guidance you can write to us or call us on 1300 CYBER1 (1300 292 371).

Disclaimer

The material in this guide is of a general nature and should not be regarded as legal advice or relied on for assistance in any particular circumstance or emergency situation. In any important matter, you should seek appropriate independent professional advice in relation to your own circumstances.

The Commonwealth accepts no responsibility or liability for any damage, loss or expense incurred as a result of the reliance on information contained in this guide.

Copyright

© Commonwealth of Australia 2021.

With the exception of the Coat of Arms, the Australian Signals Directorate logo and where otherwise stated, all material presented in this publication is provided under a Creative Commons Attribution 4.0 International licence (www.creativecommons.org/licenses).

For the avoidance of doubt, this means this licence only applies to material as set out in this document.

The details of the relevant licence conditions are available on the Creative Commons website as is the full legal code for the CC BY 4.0 licence (www.creativecommons.org/licenses).

Use of the Coat of Arms

The terms under which the Coat of Arms can be used are detailed on the Department of the Prime Minister and Cabinet website (www.pmc.gov.au/resources/commonwealth-coat-arms-information-and-guidelines).

For more information, or to report a cyber security incident, contact us:

cyber.gov.au | 1300 CYBER1 (1300 292 371)

Australian Government

Australian Signals Directorate