Australian Government
Australian Signals Directorate

ASD AUSTRALIAN SIGNALS DIRECTORATE
ACSC Australian Cyber Security Centre

Domain Name System security for domain owners

First published: January 2021
Last updated: October 2021

Introduction

This publication provides information on Domain Name System (DNS) security for domain owners, as well as mitigation strategies to reduce the risk of misuse of domains and associated resources. Organisations are recommended to 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.

Additional information on securing DNS resolvers is available in the *Domain Name System security for domain resolvers* publication.

Background

DNS is made up of root servers, recursive resolvers and a hierarchy of authoritative name servers that form a distributed database for mapping IP addresses to easily searchable names. Domain names are not owned by the registrant. Domain names are licensed from a domain name registrar for a set and renewable time period.

Name servers, through their responses, direct the flow of information on the internet. Malicious actors can target these servers as a means to disrupt, intercept or misdirect these information flows.

Common attacks against domain names are:

- **Typo-squatting:** Misleading domain names based on common typographic errors of real websites.
- **Domain registration hijacking:** Takeover of a domain registrar account to redirect traffic to a malicious website.
- **DNS hijacking:** Redirecting traffic to a website controlled by malicious actors.
- **Denial of service:** Disruption of traffic intended to reach an online service.
- **DNS spoofing:** Misleading responses to subvert DNS queries.
- **Unauthorised zone transfers:** Scanning services via unauthorised zone transfer requests.

Authoritative name server

An authoritative name server is the source of truth for a particular DNS zone. It is able to respond to queries such as 'what is the IP address of www.example.com' from its own data, without needing to reference another source.

Every domain requires at least one authoritative name server. All resources on the domain will become unreachable if there are no functioning authoritative name servers, thus best practice is to maintain two or more geographically dispersed authoritative name servers.

DNS resolvers

A DNS resolver (also known as a recursive resolver or recursive name server) searches for requested domains by querying the DNS hierarchy. DNS resolvers start by querying the root servers, then moving through the directory tree of high-level name servers until it reaches an authoritative name server that can return an IP address for a domain. DNS resolvers are covered in more detail in the *Domain Name System security for domain resolvers* publication.

Domain and authoritative name server attacks and mitigations

Typo-squatting

Typo-squatting (or Uniform Resource Locator hijacking) is where malicious actors register a domain with a name similar to another, with the intention of fooling users into visiting their decoy website. For example, with the domain 'example.com', malicious actors could register 'example-com.net' to fool inattentive users into thinking the decoy was the website they were intending to use. Malicious actors would then steal confidential information users believed they were providing to the real website. Such theft could include user names, passwords, browser login sessions, personal details or multi-factor tokens.

The DNS registry in charge of the .au namespace, auDA, discourages typo-squatting through their *Prohibition on Misspellings Policy* which has been in force since 2008.

Mitigations

Domain owners should consider the benefits of registering additional domain names that could be reached through common typing errors or with similar names. These should be redirected at the DNS level so users do finally connect with the legitimate website. For example, if 'example.com' was the domain name, consider registering:

- examples.com (plural)
- example.net and example.org (different top-level domains)
- examp1e.org (simple typograph attack [i.e. '1' replaces 'l'])
- exampel.com (common mistype).

As the variations of the original name increases, the number of potential typo-squatting domains will become costly to secure. Organisations should balance this cost against the value of their domain and the likelihood of typo-squatting.

Domain owners should check for attempts to imitate their website. Tools and monitoring services are available for detecting typo-squatting, and are generally sold as 'brand protection' or 'domain protection' and provide updates when new DNS entries are created that are similar to existing ones. The cost of the service should be weighed against the popularity of a website and the likelihood of imitation.

If an owner becomes aware of a deceptive website, efforts to takedown or block access should be initiated as soon as possible (see *Blocking access and takedowns*).

If a takedown is successful, domain owners may consider acquiring the typo-squatting domain to redirect further traffic to their legitimate website. Doing so may protect users targeted by active phishing campaigns and should also include a warning that the link was previously malicious. This scenario also provides an opportunity for user education.

Blocking access and takedowns

Domain owners can reduce the effectiveness of deceptive domains by adding them to the safe browsing block-list, and requesting a takedown of the domain or website.

- **The safe browsing block-list:** The safe browsing block-list is available for public users to identify addresses of websites hosting malicious software or phishing infrastructure. Following a submission, a review is conducted, and if validated, the malicious website is added to a block-list used by major web browsers including Google Chrome, Mozilla Firefox and Safari.

- **Domain takedown:** Domain takedowns use the abuse reporting facility of the offending domain's registrar, the registrar being a retail agent for a DNS registry. The DNS registry can also be asked to assist in taking down a domain if the registrar that sold the domain is unresponsive to a request. To request a domain takedown, a complainant needs to establish that they own the domain that is being impersonated, the impersonating domain does not have a legitimate interest or right to the domain, and the impersonating domain is being used in bad faith or for malicious purposes. For example, domain owners could provide screen-shot comparisons for each website to support their request, however, this should be done using a sandboxed system to minimise the risk of a malware infection. Note, the takedown process can be slow, but if successful, owners can consider acquiring the typo-squatting domain to redirect further traffic to their legitimate website or simply disable it. If the deceptive domain is within the .au namespace, refer to .auDA's *Prohibition on Misspellings Policy*.

- **Website takedown:** Website takedowns can be performed using the abuse reporting facility of the deceptive website's webhost or cloud provider, where available. The providers of most online services can be identified by performing a reverse DNS lookup to identify the owner of the IP address. When a provider has been identified, domain owners should attempt to make contact through the provider's abuse reporting mechanism. Although not guaranteed, a provider is expected to take down a deceptive website when presented with sufficient evidence of malicious activity.

Advice for .gov.au domain owners

Under the *Australian Government Domain Name Policy*, non-corporate Commonwealth entities must use a gov.au domain to support their website(s), and not use non-gov.au domain names (unless granted an exemption).

If you use a .gov.au domain and encounter a typo-squatting domain impersonating yours, reach out to the Australian Government's registrar at the Digital Transformation Agency.

Domain registration hijacking

Domain registration hijacking is the act of taking over a domain registration by accessing the registrar's administration system using compromised credentials or impersonating the owner to modify registration information (including transferring the domain to another registrar or by adding malicious services to existing records).

Registrar protections

Registrars and DNS registries typically provide policy controls to limit unauthorised transfers. These policies require that procedural steps must be taken by both the domain owner and registrar before any transfer can be performed. Domain owners should use these controls where possible to mitigate the risk of unauthorised domain name transfer hijacking. The follow are controls that may be implemented by registrars and registries:

- **Registrar/client locking:** This indicates a domain name should not be transferred to another registrant or registrar until the domain owner unlocks it. Depending on the parent registry's policy, locking can stop unauthorised and unsolicited transfers to another registrar. This service is often sold as a premium protection plan by registrars at an additional cost. This helps reduce the risk of malicious actors impersonating an organisation and requesting the transfer of a domain to another registrar, but does not prevent transfers through credential theft.

- **Registry locking:** This improves registrar locking by adding additional policy controls to DNS registry records that show a domain owner intends to freeze the DNS record against modification or updates. Implementation, and costs, can vary between providers. To unlock a record, the registry must carry out identity verification for the domain owner. Registry locking is typically obtained and managed through the domain's registrar.

- **Retrieving a hijacked domain:** In the event of a hijacked domain, domain owners should contact the registrar as soon as possible to report the incident and request a review. If the registrar is unresponsive, the next contact should be the registry responsible for the domain, and finally the Internet Corporation for Assigned Names and Numbers can also assist through their *Unauthorised Transfer Complaint* or *Transfer Dispute Resolution Policy*.

All the above controls are typically implemented via policy and may not be limited by technical controls. As such these controls may not provide protection if a registrar acts improperly or suffers a cyber security compromise. For high-value domains it may be worth implementing a DNS health monitoring service.

Protecting your domain administration account

In order to protect your domain administration account, the following controls are recommended:

- use multi-factor authentication (MFA) if available

- set strong credentials for logging on to the registrar's system (e.g. use a complex password or a passphrase, do not reuse credentials between services, and protect the credentials)

- review account access logs and report any suspicious activity to the registrar

- keep registration details, such as email addresses, up-to-date to ensure receipt of any unexpected registrar account change alerts.

DNS hijacking

In DNS hijacking, malicious actors compromise an organisation's name server to redirect traffic to a malicious service.

Protecting DNS name servers

Protecting DNS servers from compromise requires a strong commitment to cyber security, 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.
- **Hardening DNS software configurations:** Regularly review the configuration of DNS zone files to limit public visibility to only the services intended for public access and to provide security verification data for those services. Configure logging to provide monitoring for the performance and security of the server.
- **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 public exposure of name servers merits special priority for operating system patching. Consider rolling updates gradually across secondary name servers as updates are performed to mitigate against patch issues and downtime.
- **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.

Denial of service

During a targeted DNS denial-of-service attack, malicious actors flood an organisation's name server with a high volume of requests. This reduces the targeted service's capacity to respond to legitimate DNS traffic, and in doing so can prevent external users from finding or reaching the service on the internet.

Choosing a resilient provider

Having a service that can resist sudden bursts of traffic or sustained query load helps make the service more resilient to denial-of-service attacks. Strategies could include load balancing between secondary name servers as well as adding or removing secondary name servers to meet demand. Organisations should consider using specialised DNS providers to either host secondary name servers for load balancing or to host the full DNS service.

Commercial providers are able to provide a DNS service at scale. Organisations should choose a provider offering:

- **Wide geographic distribution:** The provider should have redundant name servers distributed across multiple locations so that if one location experiences connectivity issues, users can still reach services.
- **Anycasting and load-balancing:** Anycasting allows providers to route queries to a single IP address across multiple servers. Providers can adjust routing to balance traffic flow demand, and may offer sinkhole services to isolate denial-of-service attacks.
- **A service-level agreement:** To manage risk, organisations should negotiate a service standard enforced by penalties. Also consider avoiding limits on the number of requests made to the service as sustained denial-of-service attacks can potentially pass these limits and incur a cost to the organisation.

- **Secure administration:** A provider should offer secure and resilient administration options including:
  - **Secure authentication:** The provider's administration console should use MFA for logons.
  - **Secure encryption:** The administration console must use contemporary versions of Transport Layer Security or Secure Shell-based encryption, as specified within the *Guidelines for cryptography*.
- **24/7 helpdesk:** Attacks on name servers can occur at any time, making it important to have available support staff to respond to cyber security incidents.
- **Reporting and logging:** Detecting and investigating cyber security incidents relies on the existence of sufficient logs that are analysed in real-time or forensically. Automated log data summaries may also add value for DNS services.
- **DNS health monitoring:** Degraded performance and availability of name servers can significantly impact user experience. Health monitoring is encouraged for high value domains.

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.

Authenticating DNS addresses through DNSSEC

DNSSEC is an extension of DNS that 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.

This process is similar to how Hypertext Transfer Protocol Secure 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.

Organisations should implement DNSSEC where possible on their principal domains, and where practical, on any secondary domains. Few commercial DNS providers offer support for DNSSEC, most however will offer sufficient control of DNS records to permit manual implementation. If implementing DNSSEC manually, take care to test the implementation as widely as possible on a secondary domain before implementing on a principal domain.

Unauthorised zone transfers

An authoritative name server will provide responses for all the domains recorded in its zone file, including sub-domains and delegations to other name servers.

The primary authoritative name server can be configured to share query load and zone file data with secondary name servers for improved DNS availability. These secondary servers will regularly refresh their zone file from the primary authoritative name server via a zone transfer request.

Malicious actors probing for weaknesses will often attempt zone file transfers as a first step to discover what services are available on a network. Zone transfers can become a greater risk if an organisation's name server is responsible for both external and internal addresses. When internal addresses are resolvable, this information can be provided via zone transfer and give malicious actors knowledge of the contents of an internal network.

Zone transfer safeguards

Zone transfers may not appear to constitute a threat to most organisations, as most information available via a zone transfer is already public. However, taking steps to limit available information is a useful action for security conscious organisations, as every step that impedes or slows malicious actors increases the cost of attacking.

An organisation should limit bulk access of publicly available information by:

- **Using an internal name server:** To ensure internal network information on services will not be exposed in external zone file transfer requests. This is also known as split horizon or split brain DNS.
- **Using out-of-band transfers:** To disable zone transfers and replace them with another mechanism to limit potential exposure of external network addresses.

Further information

Further information on DNS is available from the following sources:

- National Institute for Standards and Technology's Special Publication 800-81-2, *Secure Domain Name System (DNS) Deployment Guide* publication
- Cloudflare's *DNS security* website
- Verisign's *Registry Lock* website
- Dan Kaminsky's *Black Ops 2018: It's the End of the Cache As We Know It* presentation
- 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