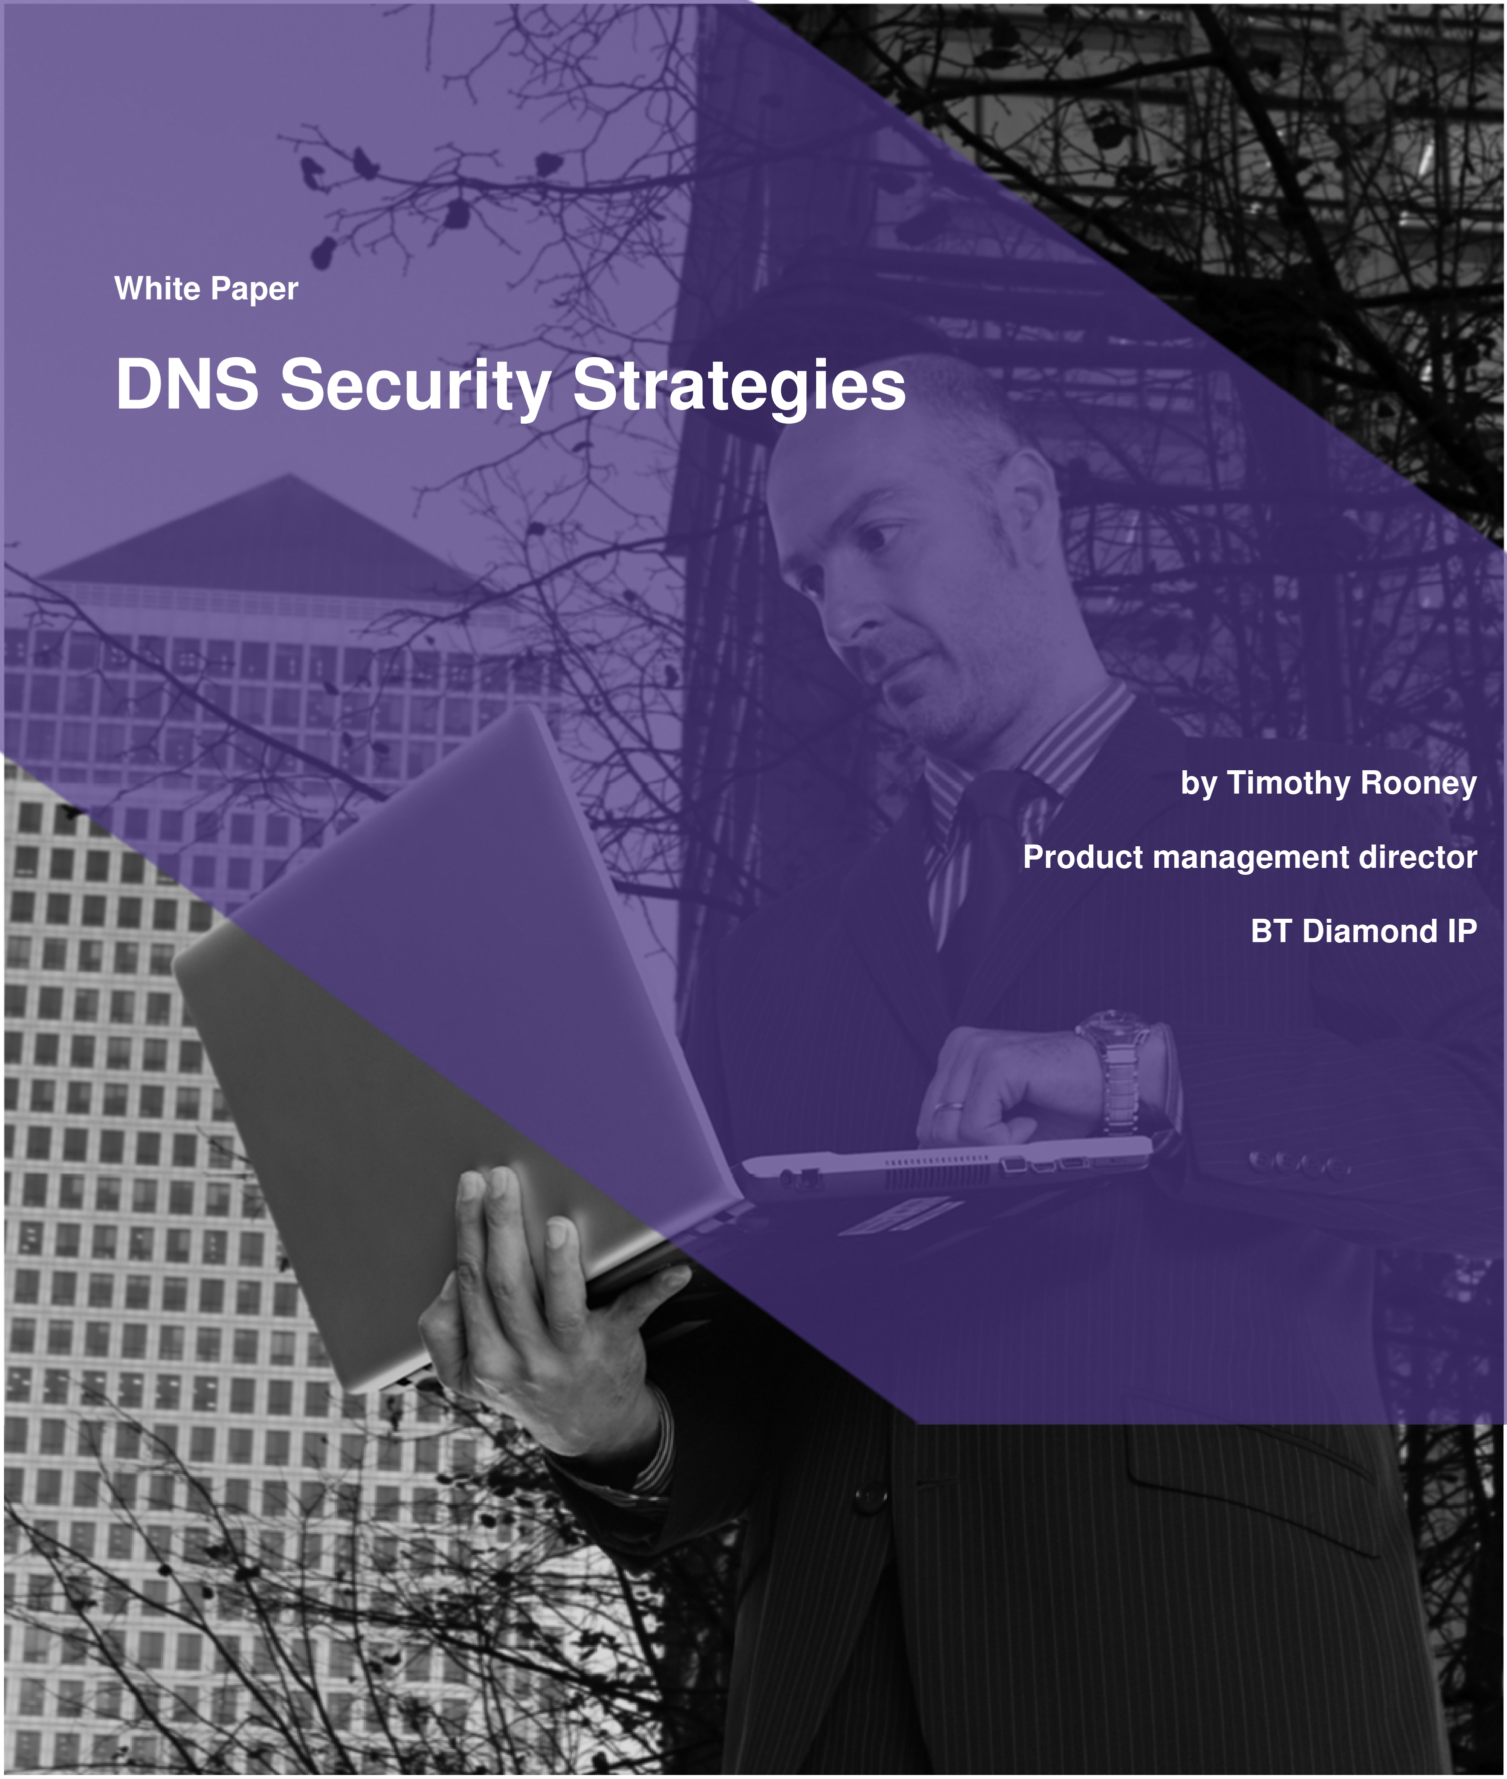

White Paper

DNS Security Strategies

by Timothy Rooney

Product management director

BT Diamond IP

BT Diamond IP

Introduction

The domain name system (DNS) is fundamental to the proper operation of nearly every IP network application, from web browsing, email, to multi-media applications and more. DNS provides the lookup and translation services from name to IP addresses that are used by computers to communicate. An attack that renders the DNS service unavailable or which manipulates the integrity of the data contained within DNS can effectively bring a network down. As a side effect of its necessity, DNS traffic is generally permitted to flow freely through networks, exposing networks to attacks that leverage this freedom of communications.

By its very nature, the global Internet DNS system serves as a distributed data repository containing host names (e.g., website and other addresses) and corresponding IP address information. The distributed nature of DNS applies not only to the global geographic distribution of DNS servers, but to the distribution of administration of the information published within respective domains of this repository. DNS has proven extremely effective and scalable in practice and most people take DNS for granted given this and its proven reliability. However, its essential function and decentralized architecture serve to attract attackers seeking to exploit the architecture and rich data store for sinister activities.

This white paper describes various forms of enterprise DNS attacks and strategies you can employ to mitigate these attacks. We'll start with a basic overview of DNS to establish a level set and present potential vulnerabilities, which we'll discuss next, followed by mitigation strategies that can be deployed to reduce exposure to various attack types.

DNS basic operation

Figure 1 illustrates the basic flow of a DNS query, along with the various data stores for DNS data and corresponding data sources. The DNS resolution process generally starts when a device's user seeks to connect to a website or other IP service. Upon entry of the desired destination by name, software running on the device called a *DNS resolver* issues a query to its configured *recursive DNS server*, starting on the left of the figure, unless the query had recently been resolved, where the answer may be returned immediately from the resolver cache if it exists.

The recursive DNS server's role is to resolve the query on behalf of the client by using its cache of previously resolved queries or by querying public DNS servers on the Internet. The latter process typically involves several queries to multiple DNS servers to first locate a DNS server that is *authoritative* for the domain for which the query relates and then to query an authoritative server itself to obtain an answer that can be passed back to the client, thereby completing the resolution process. The recursive server also caches the resolution information in order to respond more quickly to a similar query without having to re-seek the answer on the Internet.

Multiple authoritative DNS servers are deployed to provide services continuity in the event of a server outage. Generally, an administrator configures a *master* server that then replicates or transfers its domain information to one or more *slave* servers. Some DNS solutions enable alternative forms of data replication using LDAP-like replication for example.

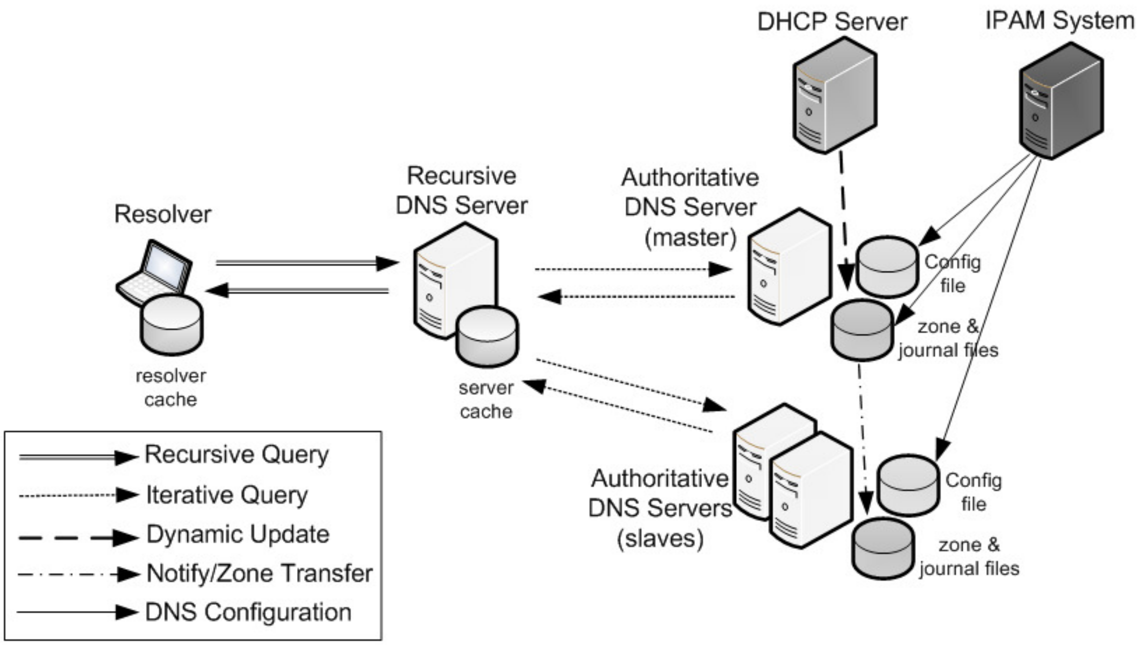

Figure 1 DNS Query Flow and Data Sources

DNS Trust Model

The DNS trust model refers here to how DNS information flows among these components of the DNS system. In general information received by other components in the system is trusted though various forms of validation and authentication can improve trustworthiness as we shall discuss later.

From the client resolver perspective, it trusts its resolver cache and the recursive server to provide answers to DNS queries. Should either trusted data source be corrupted, the resolver could inadvertently redirect the user application to an inappropriate destination.

The recursive server trusts its cache and the various DNS servers it queries internally on the enterprise network or externally on the Internet. It relies not only on accurate responses from authoritative DNS servers but on other domain servers which provide referrals to locate DNS servers authoritative for the domain in question. Referral answers are generally provided by the Internet root servers as well as top level domain (TLD) servers, but referrals may also be provided by other servers operated internally or by DNS hosting providers. These referral DNS servers are not shown in Figure 1 but may be queried by the recursive DNS server to walk down the hierarchical domain tree to locate the authoritative DNS servers.

Authoritative DNS servers are so called given that they are purportedly operated by or on behalf of the owner of a given DNS domain who is responsible for the information published on these servers. Resolvers attempting to resolve hostnames within the domain of the authoritative server trust the server to respond with accurate information, where *accurate* means *as published by the domain administrator*. Information published on authoritative DNS servers originates from a variety of sources including manually edited text files, inter-server transfers and updates, and/or use of IP address management (IPAM) solutions. Inter-server transfers refer to master-slave replication, while updates may originate from other DNS servers, DHCP servers, other systems, or even end user devices if permitted by administrators.

DNS Information Sources

The DNS trust model provides us one perspective on information flow among components and inherent trust relationships within the resolution process. Examining the sources of each component's information can help us understand how each is potentially configured both in terms of base configuration, such as for what domains a particular server is authoritative, whether it should answer queries from given addresses, etc. and actual DNS resolution data communicated via the DNS protocol.

The following table lists potential sources of this information respectively. By "referral" DNS servers, we include root servers, TLD servers and other authoritative servers within the domain tree that refer recursive DNS servers down the tree ultimately to the authoritative servers. These servers are administered by external organizations with IANA managing the Internet root zone, and TLD administrators managing corresponding DNS servers to publish proper mapping of child domains to corresponding authoritative or referral DNS servers.

Component	Configuration data sources	DNS data sources
Client resolver	• Configuration or properties file • DHCP/PPP server provided parameters	• Resolver cache (if equipped) • Recursive DNS server
Recursive DNS server	• Configuration or properties file • DNS control channel • Hints file of root DNS servers	• Server cache • Forwarder server if configured • Referral DNS servers • Authoritative DNS servers
"Referral" DNS servers	• Domain level DNS administrators and tools	• Child domain administrators provision of their DNS servers' names/IP addresses
Authoritative DNS server	• Configuration file or properties • Zone file(s) • DNS control channel	• DNS administrator via text editor, DNS GUI or IPAM system • DDNS updates to zone files • Zone transfers

DNS as an Attack Target

Attackers may target DNS services in and of themselves in order to stifle communications or to steer unwitting end users to imposter web servers or other destinations. Alternatively, DNS may serve as a facilitator for use with the scope of a broader network attack. Just as DNS enables users to connect to websites by resolving text-based destinations to IP addresses, it enables attacker malware to locate command and control centers or to tunnel information through firewalls. DNS by its nature also openly publishes potentially useful information about networks, host names and IP addresses for would-be attackers.

We'll examine first those attacks on DNS infrastructure, consisting of DNS servers within your organization and those on the Internet used in the process of name resolution. Then we will discuss more broadly targeted attacks that leverage the DNS. As we shall see given this broad and diverse set of attacks, no single mitigation technology can effectively combat them all; a comprehensive DNS security strategy is necessary to defend against them collectively. We will wrap up the paper with a discussion of various components of such a strategy.

Attacks on DNS infrastructure

By DNS infrastructure, we refer to the DNS servers themselves, the DNS resolvers which query for address lookups on behalf of user applications on end user devices, the integrity of DNS information, and the collective DNS service of resolving hostnames on behalf of resolvers.

DNS Service Denial

The familiar Denial of Service (DOS) or distributed DOS (DDOS) attack is invoked by an attacker to flood the DNS server with bogus DNS requests, overwhelming its ability to process legitimate DNS queries. From the DNS server's perspective, it merely attempts to process each query as it is received. As the volume of bogus queries is intensified beyond the query response rate supported by the server, the proportion of legitimate queries lessens and DNS resolution services capacity drops precipitously to only that small proportion that is processed.

- Denial of service – DNS, like other network services, is vulnerable to denial of service (DOS) attacks, which features an attacker sending thousands of packets to a server in hopes of overloading the server, causing it to crash or become otherwise unavailable to other queriers. The service is rendered unavailable and thus denied to others.
- Distributed Denial of Service - A variant of this type of attack is the use of multiple distributed attack points and is referred to as a distributed denial of service (DDOS) attack. The intent is the same, though the scale is larger, potentially impacting several servers.
- Wild Goose Chase – This attack attempts service denial through the flooding of a recursive server with queries for bogus domain names. This causes the server to initiate a "wild goose chase" and utilize resources to futilely locate the authoritative server within the domain tree.

D/DOS attacks can take the form of the issuance of a high volume of DNS queries to deny DNS service but may also comprise issuance of TCP SYN attacks or UDP floods targeted at DNS servers.

Cache Poisoning Style Attacks

DNS resolvers and recursive caching servers maintain a cache of resolved resource records to improve resolution performance as described earlier. If an attacker succeeds in corrupting a recursive server's cache, the corrupted information may be provided to several users requesting the same or similar domain name information. Corrupting the cache requires an attacker to provide a seemingly legitimate query answer albeit with falsified information in part or in total.

This can result in hijacking resolvers and hence applications to incorrect destinations, e.g., web sites.

The following forms of cache poisoning attacks have been identified:

- Packet interception or spoofing – Like other client/server applications, DNS is susceptible to “man-in-the-middle” attacks where an attacker responds to a DNS query with false or misleading additional information. The attacker spoofs the legitimate DNS server response, leading the server to resolve and cache this information.
- ID Guessing or Query Prediction – Another form of malicious resolution is ID guessing. The ID field of the DNS packet header is 16 bits in length, as is the UDP packet header ID. If an attacker can provide a response to the query with the correct ID field and UDP port number, the resolver will accept the response, whatever it contains. This enables the attacker to provide falsified results, assuming the query type, class and queried name are known or guessed by the attacker. Guessing a 2^{32} number is relatively easy even with brute force methods.
- Name Chaining – This attack features the provision of supplemental resolution information usually within the Additional or even Authority section of the DNS response packet, thereby poisoning the cache with malicious resolution information. This may for example attempt to falsify information for a popular website such as bt.com, google.com or the like, so when such a query is requested, the resolver will rely on this falsified cached information to essentially redirect the client to the attacker’s intended destination.
- The Kaminski DNS vulnerability – Dan Kaminsky identified a vulnerability which simplified instigation of cache poisoning attacks. An attacker could create a web page with hundreds or thousands of tags with URLs such as img tags, causing the web browser to attempt to resolve each URL using DNS. The attacker knows what queries will be asked, which reduces the challenge to properly guess the transaction ID and UDP port, which as mentioned above, is relatively easy. The query answers will all contain additional information which the server will cache, perhaps the falsified IP addresses of popular websites.

Authoritative Poisoning

Cache poisoning attacks attempt to corrupt DNS information cached within resolvers and recursive servers. Other forms of attack attempt to corrupt DNS information published within authoritative DNS servers. Unlike cached information which eventually times out, corruption of authoritative information could persist for lengthy time periods until detected.

- Dynamic updates – An attacker may attempt to inject or modify data in a DNS zone by attempting to issue a DNS Update message to the DNS server. This type of attack could manipulate resolution data, redirecting resolutions from clients for the intended destination to an attacker-specified destination.
- Server configuration – an attacker may attempt to gain access to the physical server running the DNS service. Configuring appropriate credentials and hardening the DNS server can help mitigate this attack vector as can securing communications to the server whether over SSH or via an IPAM system.

Beyond being able to manipulate configuration and zone information, an attack of this type could enable the use of the server as a stepping stone to other targets, especially if this server is trusted internally.

- Configuration errors – while typically not malicious (though most attacks are initiated from internal sources), misconfiguring the DNS service and/or zone information may lead to improper resolution or server behavior.

Server/OS Attacks

As with all network servers, vulnerabilities within the server operating system may be exploited by attackers in order to severely hamper or crash the server. These attacks can be of the following forms:

- Operating System attacks – An attacker may attempt to gain access to the server by overflowing the code execution stack or buffer. Such an attack may exploit a known vulnerability of the operating system or version of DNS software running on the server.
- DNS service attacks – An attacker may attempt to exploit a known vulnerability for a given version of DNS server software running on the victim server to shut it down or otherwise disrupt service.
- Control channel attack –The DNS server control channel provided in most implementations provides a convenient mechanism to remotely control the DNS server, such as stopping/halting the server's DNS software (e.g., 'named'), reloading a zone, and more. Such power may entice an attacker to attempt to access the control channel to perform nefarious functions such as stopping the DNS service thereby denying DNS service to querying servers and resolvers.

Resolver Attacks

The resolver on the client device must be initialized with at least one DNS server IP address to which DNS queries can be issued. This IP address is the destination address for all DNS queries originating from the client. Other resolver configuration information such as domains suffixes may also be defined. The resolver configuration may be performed manually by hard-coding the DNS server IP address in the TCP/IP stack, or automatically via DHCP or PPP.

- Corruption through DHCP/PPP – This type of attack seeks to redirect the resolver from the legitimate recursive DNS server to an attacker's DNS server to poison the resolver with malicious DNS query answers. Manipulation of client configuration obtained through DHCP or PPP would entail provision of a rogue DHCP or Radius server on the part of the attacker.
- Device infiltration – An attack to gain access to a device could provide the ability to edit the resolver configuration among other host information.

Broader attacks that leverage DNS

While several attack types target the DNS infrastructure itself, several broader network attacks leverage the DNS to inflict damage on other network components or to exfiltrate sensitive information outside the network.

Network Reconnaissance

DNS by design contains a repository of hostname-to-IP address mapping among other things. If an attacker desired to glean information about particular hosts that may be more attractive to attack than others, he/she may start with DNS.

- Name guessing - One brute force approach to such reconnaissance consists of guessing hostnames of interest (e.g. "payroll") and issuing standard DNS queries to obtain corresponding IP addresses if they exist.
- Zone transfers – Impersonating a DNS slave server and attempting to perform a zone transfer from a master is a form of attack that attempts to map or footprint the zone. That is, by identifying host to IP address mappings, as well as other resource records, the attacker attempts to identify targets for direct attacks.

Reflector Style Attacks

This form of attack attempts to use one or more DNS servers to send massive amounts of data at a particular target, thereby denying service for the target machine.

- Reflector attack - The attacker issues numerous queries to one or more DNS servers using the target machine's IP address as the source IP address in each DNS query.
- Amplification – Using the reflector approach while querying for resource record types with large quantities of data such as NAPTR, and DNSSEC signed answers amplifies this attack. Each responding server responds with the data to the "requestor" at the spoofed IP address to inundate this target with a large data flow.

Data Exfiltration

Data exfiltration refers to the transmission of data originating from within one security domain, e.g., an enterprise network to another entity or organization. There are two basic forms of data exfiltration using DNS:

- DNS as data protocol (tunneling) - DNS tunneling entails the use of the DNS protocol as a data communications protocol. This technique enables a user or device within the network to communicate with an external destination, easily traversing firewalls (DNS is generally permitted through firewalls).
- DNS as resource locator – an attacker may attempt to install malware on devices via phishing attacks that bait users into opening executable email attachments or installing software from an attacker website. Whether a device is attacked while inside the enterprise network or a user device is physically brought onto the network, if they are

trusted within the confines of an enterprise network they may have access to sensitive information. The malware may perform data collection, locating internal resources using DNS reconnaissance. In addition, DNS could be used to identify the current IP address of the attacker's external destination for exfiltration of the information.

Advanced Persistent Threats

Advanced Persistent Threats (APTs) are organized, stealthy forms of network intrusion where an attacker attains access within a target network to steal data, disrupt communications, or otherwise infiltrate network components. APTs are persistent in that the intent is to retain access to the network for a lengthy time frame, if not indefinitely, so they require continual evasion techniques to avoid detection.

Attackers may gain access within the network via phishing attacks to deploy malware, social engineering, or other methods. Once within the network, the attacker's malware typically attempts to communicate to a "command and control" (C&C) center, from which the attacker can instigate attacks, update malware code, or collect information. Many times, this "phoning home" process involves DNS queries to identify the current IP address of the C&C center. Use of DNS enables the attacker to change IP addresses quickly to avoid detection. In addition, C&C center domain names can be algorithmically generated to support a moving target on the query name as well.

DNS Mitigation Approaches

The diversity of DNS threats presents a serious challenge to IT administrators responsible for managing the operation and integrity of DNS services within their networks. With such a variety of attacks, no single mitigation technology or approach will suffice in defending against them all. A mitigation portfolio is necessary to provision defenses for each major attack threat. This section summarizes the major mitigation techniques available to IT administrators in assembling such a mitigation strategy.

Denial of Service Mitigation

Effective denial of service mitigation requires limiting the level of server processing of attack packets such that legitimate packets are still handled properly. Methods to reduce the impact of D/DOS attacks include:

- Inbound ACLs at the server and DNS service level regarding from which IP networks or hosts the server will process queries (recursive servers)
- Inbound rate limiting of packets from specific sources and/or by type (TCP, UDP, DNS, etc.) (recursive or authoritative servers)
- DNS anycast deployment where multiple DNS servers use a common IP address. This approach was proven effective against a DDOS attack against Internet DNS root servers in February, 2007 (authoritative servers)
- To protect against bogus query attacks, limit the number of outstanding queries per client.

Cache Poisoning Mitigation

Attackers attempt to steer clients to falsified websites for the purposes of stealing personal or other data or for other malicious purposes. The various forms of this style of attack basically entail the attacker answering a valid DNS query with falsified resolution data prior to the arrival of the legitimate answer. The only definitive mitigation to cache poisoning is DNSSEC (DNS security extensions).

DNSSEC eliminates the possibility of an attacker successfully poisoning the recursive server cache provided that the resolution data is DNSSEC-signed and the recursive server is configured to validate DNSSEC responses. DNSSEC provides origin authentication such that only the domain publisher can be authenticated as well as data integrity checking to verify no data manipulation occurred in transit between the authoritative server and the recursive server. Authenticated denial of existence of resource record information is also provided by DNSSEC.

Authoritative Poisoning Mitigation

Poisoning of authoritative DNS server information enables an attacker to modify resolution data directly on the server resolving DNS queries. Authoritative DNS servers may be provisioned with resource record data by the following means, each of which may be defended using the corresponding approaches.

Zone [file] editing requires either physical access to the server or the ability to transfer a file to the server using a file transfer protocol (FTP, TFTP, SCP, etc.) or via DNS zone transfer. If an administrator uses a separate configuration tool such as a DNS GUI or IPAM system, entry errors or uncaught errors could corrupt a zone's information. Zone information may also be modified using the DNS Update aka Dynamic DNS (DDNS) message. DDNS messages are directed to the master server and updates are replicated to slave servers typically using Notify messages in order to trigger an incremental zone transfer.

Mitigation of these forms of zone information corruption consists of the following techniques:

- DNS server host access controls
- DNS server operating system hardening to prevent access via unauthorized protocols
- DNS files permissions setting
- Use of a robust DNS GUI or IPAM system or use of DNS configuration or zone checking utilities to maximize server information integrity
- ACLs on hosts/networks from which the server will accept DDNS (and Notify) messages

DNS Server Attack Mitigation

Compromise of the DNS server or the DNS service running on a server can enable an attacker to reduce the availability of DNS service or to modify DNS configuration. Compromise may come in the form of host access, operating system or DNS service disruption through exploitation of a known vulnerability, or access to the DNS control channel if equipped, which may enable an attacker to stop the DNS service, freeze dynamic updates, or other disruptive activities.

The following approaches may be employed to defend against DNS server attacks:

- DNS server host access controls
- DNS server operating system hardening to prevent access via unauthorized protocols
- Monitor security advisories (e.g., CERT) for operating system or DNS service vulnerabilities and keep systems updated to prevent exploitation
- Inhibit responses to "version" queries
- Protect the DNS service control channel using available controls such as ACLs and authentication keys.

Resolver Attack Defenses

Attackers desiring to steer resolvers to illicit recursive servers may attempt to corrupt the "DNS servers" setting (among others) of the resolver accordingly. This setting may be configured by editing this information directly on the device or via parameters provided in initialization protocols, such as DHCP or PPP. Protecting against this class of attack necessitates the following actions:

- Device access controls
- Verify proper provisioning of DHCP/PPP parameter settings
- Monitor for rogue DHCP servers which may be setting improper parameters

Network Reconnaissance

The intent of DNS is to publish address information about hosts on the network. However, this information may be gathered and analyzed by an attacker to facilitate target identification for further attacks. Naming hosts intuitively certainly simplifies user accessibility though "attractive" names may tempt attackers, so this is a trade-off. However, limiting who can query for such information can help constrain the scope of access to this information for recursive or internal authoritative servers. Host information published in external DNS servers (where constraining query sources makes less sense) should be limited to those accessible via the Internet.

In summary, the following defense mechanisms may be put in place to protect against overt network reconnaissance:

- Implement ACLs limiting zone transfers to only other authorized authoritative DNS servers
- Implement ACLs limiting the scope of hosts that are permitted to query the server

Reflector Attack Mitigation

Reflector and amplification attacks entail an attacker issuing a high volume of DNS queries using the source IP address of the intended target; hence the target receives a high volume of DNS responses which may deny its ability to perform its intended function. The reflector attack may be amplified by issuing queries from multiple sources, all spoofing the target's IP address and/or by issuing queries for resource records containing large payloads. To protect against this style of attack, the following techniques may be employed:

- Provision ingress filtering on routers to minimize the ability to spoof addresses
- Implement response rate limiting to control the flow to a reasonable maximum

Data exfiltration

Stealing data from sources within a network and transmitting it to an attacker's system externally may prove to be a very attractive attack vector. Use of DNS to identify external system domain names or to serve as the transmission protocol itself facilitates this style of attack. Steps you can take to mitigate these include:

- Implement a DNS firewall to prevent resolution of known "bad domains" to reduce resolution possibilities for external attacker systems
- Monitor DNS transactions to identify potential tunnels

Advanced Persistent Threats

Advanced Persistent Threats (APTs) feature stealthy infiltration of a network typically for the purpose of data exfiltration or other malicious activities. Stealthiness is achieved by the use of rapidly changing code compilation and in other techniques to avert detection. Use of a DNS firewall can help to disable DNS resolution of APT malware bots that use DNS to resolve external command and control centers for instructions.

Summary of Major DNS Threats and Mitigation Approaches

As we've seen, several varieties of DNS attacks are possible to disrupt DNS or network communications in general or to leverage DNS' intended purpose to identify targets or attack systems with malicious intent. No single mitigation approach can eliminate vulnerabilities to all threats; thus, a multi-pronged mitigation strategy is required to reduce attack exposure. The following table summarizes the threats to DNS and corresponding mitigation approaches for each. Mitigation approaches are summarized in the section following.

Threat		Threat Summary	Mitigation Approaches
Denial of Service	Denial of service	Attacker transmits a high volume of TCP, UDP, DNS or other packets to the DNS server to inundate its resources	• Inbound rate limiting • Anycast deployment
	Distributed denial of service	Attacker transmits a high volume of TCP, UDP, DNS or other packets from multiple sources to the DNS server to inundate its resources	• Inbound rate limiting • Anycast deployment
	Bogus queries	Attacker transmits a high volume of bogus queries, causing the recursive server to futilely locate authoritative servers	• Limit the number of outstanding queries per client

Cache Poisoning	Packet Interception/ Spoofing	Attacker transmits a DNS response to a recursive DNS server in order to poison its cache, affecting DNS resolution integrity for	• DNSSEC validation on recursive servers • Source port and XID randomization • Qname case manipulation and verification on response
	ID Guessing/Query Prediction	Attacker transmits a DNS response(s) to a predicted query using a predicted or variety of XID values.	• DNSSEC validation on recursive servers • Source port and XID randomization • Qname case manipulation and verification on response
	Kaminsky Attack/Name Chaining	Attacker transmits a DNS response(s) with falsified answers in the DNS message Additional section. The Kaminsky attack produces deterministic queries to facilitate the attack.	• DNSSEC validation on recursive servers
Authoritative Poisoning	Illicit dynamic update	Attacker transmits a DNS Update message(s) to a master DNS server to add, modify or delete a resource record in the target zone	• Use ACLs on allow-update, allow-notify, notify-source. • ACLs may also be defined as requiring transaction signatures for added origin authentication
	Server attack/hijack	Attacker hacks into the DNS server which enables manipulation of DNS data among other server capabilities	• Implement host access controls • Use hidden masters to inhibit detection of the zone master • Harden server operating system and keep up to date • Limit port or console access
	DNS service misconfiguration	Vulnerability to configuration errors exposes the DNS service to improper configuration	• Use checkzone and checkconf or similar utilities • Use an IPAM system with error checking • Keep fresh backups for reload if needed

Server/OS Attack	Buffer overflows and OS level attacks	Attacker exploits server operating system vulnerability	• Harden operating system and keep updated
	Control channel attack	Attacker accesses the DNS service control channel to disrupt DNS service	• Control channel ACLs • Control channel keyed authentication
	DNS service vulnerabilities	Attacker exploits DNS service vulnerability	• Monitor CERT advisories and update DNS service • Do not expose DNS service version to version queries
Resolver/host attacks	Recursive DNS redirection	Attacker misconfigures resolver to point to illicit recursive DNS server	• Configure DNS servers via DHCP • Monitor for rogue DHCP servers • Periodically audit each client for misconfigurations or anomalies
	Resolver Configuration Attack	Attacker hacks into the device which enables manipulation of resolver configuration among other device capabilities	• Implement host access controls
Network Reconnaissance	Name guessing	Attacker issues legitimate DNS queries for names that, if resolved could serve as further attack target	• Avoid naming hosts with overly “attractive” names
	Illicit zone transfer	Attacker initiates a zone transfer request to an authoritative DNS server to obtain zone resource records to identify potential attack targets	• Use ACLs with TSIG on allow-transfer; and use transfer-source IP address and port to use a non-standard port for zone transfers
Reflector style attacks	Reflector attacks	Attacker spoofs the target’s IP address and issues numerous queries to one or more authoritative DNS servers to inundate the target	• Implement ingress filtering on routers to mitigate spoofing • Use DNS response rate limiting
	Amplification attacks	Attacker amplifies reflector attack by querying for “large” resource records to increase data flow to target per query	• Implement ingress filtering on routers to mitigate spoofing • Use DNS response rate limiting

Data exfiltration	DNS tunneling	Attacker transmits data through firewalls using DNS as the transport protocol	• Monitor DNS queries for frequent queries between a given client and server especially with large query and response payload
	Resource locator	Attacker infects internal device which uses DNS to locate command and control center	• DNS firewall
APT	Advanced Persistent Threats	Attacker deploys adaptable malware within a network to perform nefarious functions to disrupt communications and/or steal information	• DNS firewall

Summary

Available, responsive, and accurate DNS service is an absolute must in today's IP networks. But its critical role, not to mention its openness and flexibility render DNS an attractive target for attackers. A multi-faceted approach is required to combat attackers and provide a secure defense. BT Diamond IP products and services enable customers to implement and deploy these defensive measures to proactively protect their DNS infrastructure against attack and against its illicit use for nefarious purposes.

About BT Diamond IP

BT Diamond IP is a leading provider of software and appliance products and services that help customers effectively manage complex IP networks. Our next-generation IP management solutions help businesses more efficiently manage IP address space across mid-to-very large sized enterprise and service provider networks. These products include IPControl™ for comprehensive IP address management and Sapphire hardware and virtual appliances for DNS/DHCP services deployment. Our cable firmware management product, ImageControl™, helps broadband cable operators automate and simplify the process of upgrading and maintaining firmware on DOCSIS devices in the field. Our customers include regional, national and global service providers and enterprises in all major industries.

For more information, please contact us directly at +1-610-321-9000 worldwide, email to btdiamondip-sales@bt.com or consult www.btdiamondip.com.

IPControl is a trademark of BT Americas, Inc.