cira

BUILDING A BETTER
ONLINE CANADA

OVERVIEW OF THE DNS AND GLOSSARY OF TERMS

Protect your business

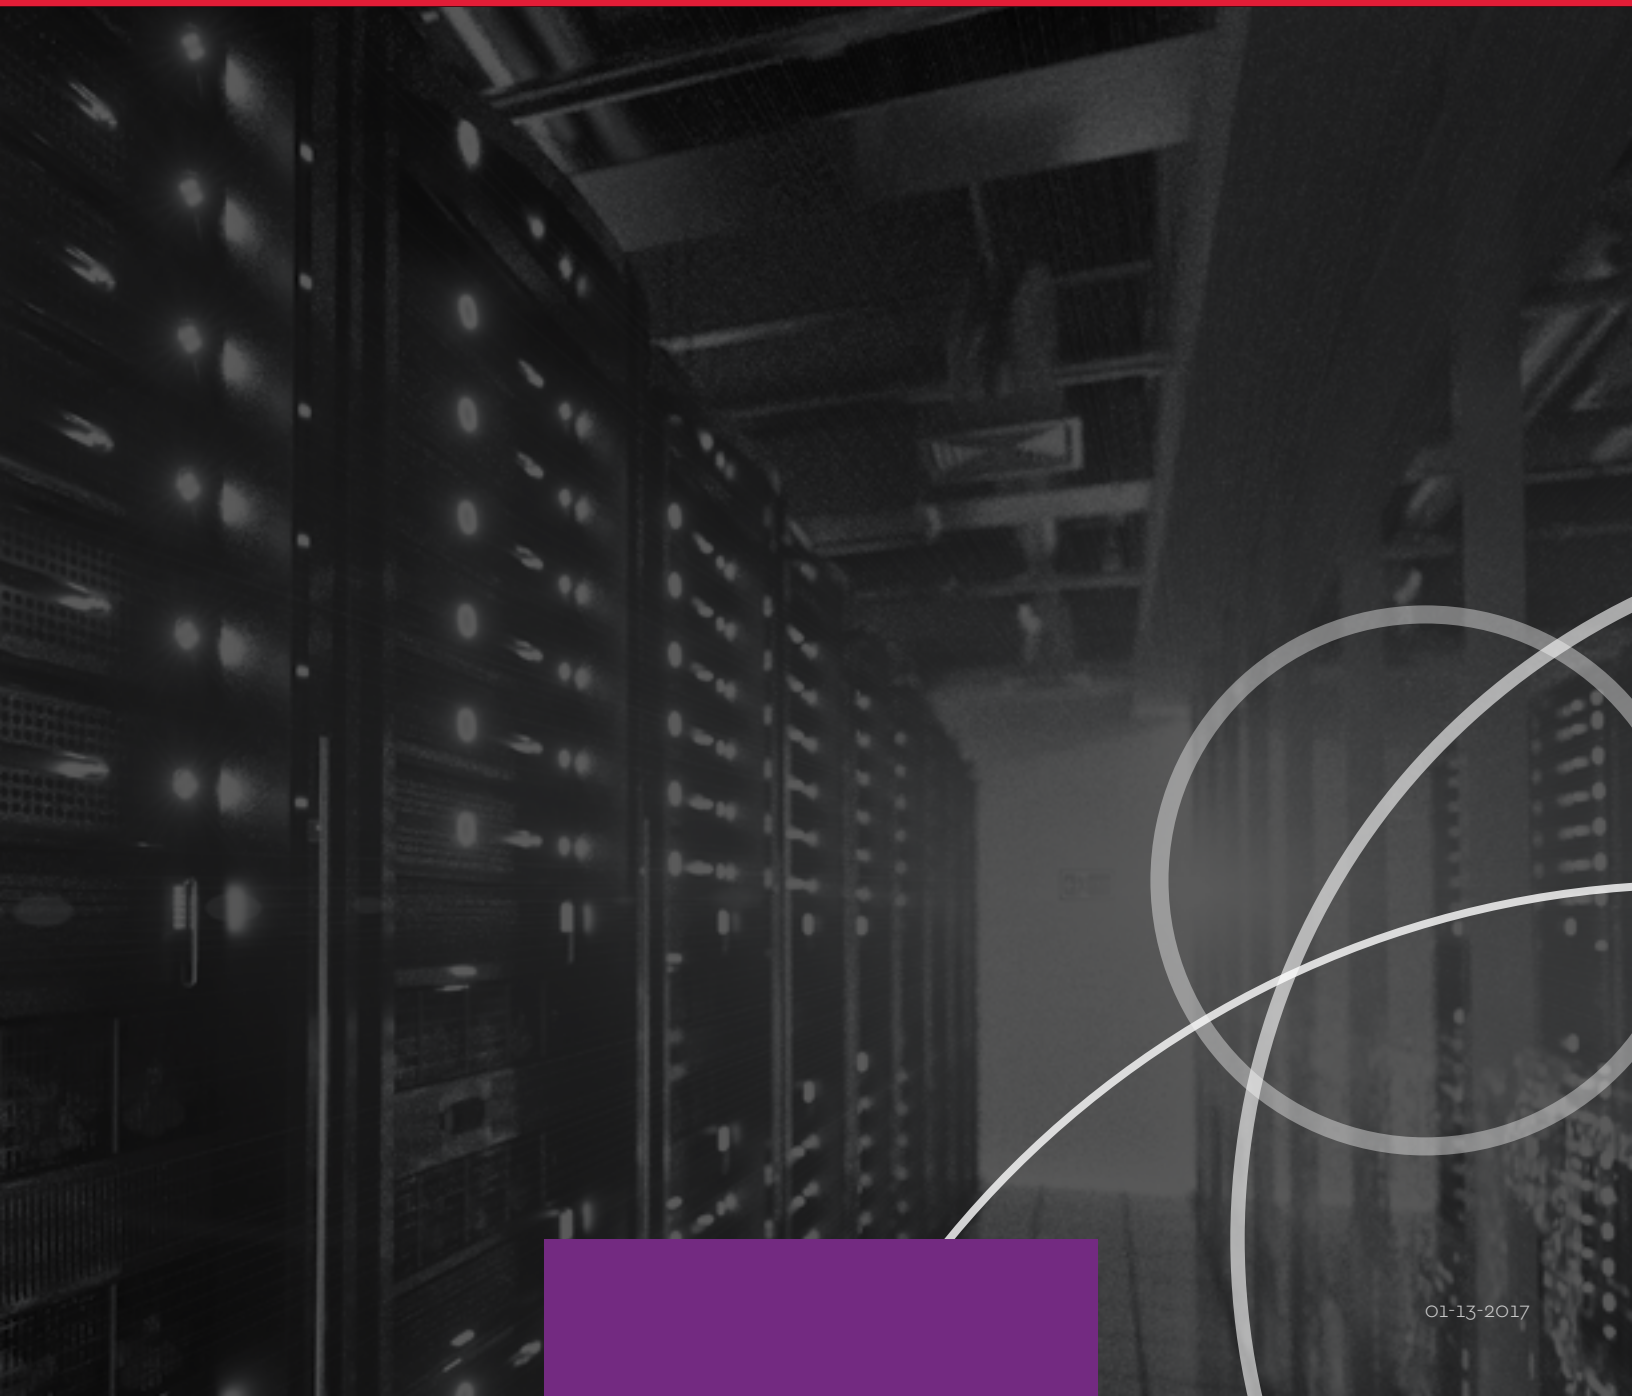

01-13-2017

OVERVIEW OF THE DNS AND GLOSSARY OF TERMS

The DNS is a technology that most IT managers don't think much about; it works well and usually does not require much attention to support organizational objectives. As businesses increasingly rely on their web strategy, the DNS infrastructure deserves attention for many reasons, including:

1. Increasing trend towards DDoS attacks against the DNS has raised awareness of weaknesses in traditional legacy DNS architecture that can easily be exploited.

2. Increasing complexity of websites and web applications, along with sourcing of content from multiple different sites, has resulted in longer page load times. DNS latency, although typically low, is part of the total sum and optimizing your DNS infrastructure will enhance your users' web experience.

3. Understanding that, for many organizations, there are benefits to keeping local Canadian traffic within Canadian borders. This helps to improve information security, mitigate geographically sourced malicious attacks, and increase speed. DNS architectures can be configured to optimize regional traffic while serving global traffic.

To help demystify the DNS, this short primer provides information to help IT managers or technical decision makers re-familiarize themselves with the technology.

HOW DOES THE DOMAIN NAME SYSTEM (DNS) WORK?

The Domain Name System (DNS) provides the core backbone of the Internet by providing the map between easily-readable hostnames (i.e. www.cira.ca) and IP addresses (192.228.29.1) by way of resource records. It is essential to the operation of the Internet by enabling the use of logical, human-readable names for locations rather than complex IPv4 or IPv6 addresses. It additionally provides mappings to things like mail servers, SIP servers, redirects, digital signatures, and more.

The DNS is a distributed database organized as a tree of interconnected nodes (server or server clusters) where each node is a partition of the database. Nodes are delegated to designated authorities and there can be only one authority for a node or group of nodes.

HOW ARE NODES DELEGATED?

The DNS is a top-down hierarchical distributed database with the Internet Corporation for Assigned Names and Numbers (ICANN) acting as "the root" at the top of the database. Top-level domains (TLDs) are delegated to both countries and commercial entities, who likewise delegate second-level domains to Registrants. For example, ICANN/IANA delegates the authority for .CA to the Canadian Internet Registration Authority (CIRA), who delegates the authority for gc.ca to the Canadian government, and the Canadian government may (and does) delegate authority for subdomains under gc.ca to its departments.

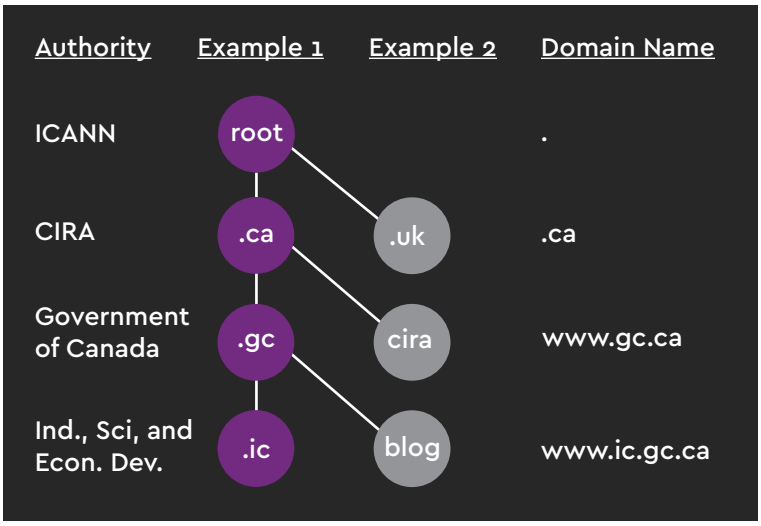

Example – Delegation of authority for ic.gc.ca showing the root level, top level, second level and third level. Grey circles represent alternative examples for each level.

WHAT IS A ZONE FILE?

A zone file is a list of DNS resource records for the zone. The responsibility for maintaining this zone file rests with the organization or individual with delegated authority for the zone.

The responsible organization or individual is also obligated to ensure that there are DNS nameservers available to respond to DNS queries for resource records in the zone. For some organizations, this function is offloaded to another company for management, and this is particularly so for small companies using hosting packages from a Registrar or other hosting company. A typical SOHO (small office/home office) would tend to not have the expertise needed to manage their own DNS nameservers. With medium and larger organizations within Canada, CIRA research shows that DNS authorities are using a mix of insourced DNS infrastructure and outsourced DNS service providers.

```
$ORIGIN example.com.     ; designates the start of this zone file in the namespace
$TTL 1h                  ; default expiration time of all resource records without their own TTL value
example.com.  IN  SOA   ns.example.com. username.example.com. ( 2007120710 1d 2h 4w 1h )
example.com.  IN  NS    ns                    ; ns.example.com is a nameserver for example.com
example.com.  IN  NS    ns.somewhere.example. ; ns.somewhere.example is a backup nameserver for example.com
example.com.  IN  MX    10 mail.example.com.  ; mail.example.com is the mailserver for example.com
@             IN  MX    20 mail2.example.com. ; equivalent to above line, "@" represents zone origin
@             IN  MX    50 mail3              ; equivalent to above line, but using a relative host name
example.com.  IN  A     192.0.2.1             ; IPv4 address for example.com
              IN  AAAA  2001:db8:10::1        ; IPv6 address for example.com
ns            IN  A     192.0.2.2             ; IPv4 address for ns.example.com
              IN  AAAA  2001:db8:10::2        ; IPv6 address for ns.example.com
www           IN  CNAME example.com.          ; www.example.com is an alias for example.com
wwwtest       IN  CNAME www                   ; wwwtest.example.com is another alias for www.example.com
mail          IN  A     192.0.2.3             ; IPv4 address for mail.example.com
mail2         IN  A     192.0.2.4             ; IPv4 address for mail2.example.com
mail3         IN  A     192.0.2.5             ; IPv4 address for mail3.example.com
```

A typical zone file (source: wikipedia)

WHAT ARE THE TYPICAL STEPS INVOLVED IN A DNS QUERY

Step 1 – Network connected devices (resolvers) send a request (or query) to a recursive nameserver. If the recursive server does not have the answer cached it moves to step two.

Step 2 – The recursive nameserver sends a query to the root nameservers to resolve the address for the top-level domain. The root nameservers for the top-level domain then return a referral to the recursive nameserver.

Step 3 – The recursive nameserver then sends a query to the top level nameservers (in this example, .CA) which then returns a referral to the second-level nameservers.

Step 4 – The recursive nameserver then sends a query to the second-level nameservers (in this example, gc.ca) which then returns a referral to the third-level nameservers.

Step 5 – The recursive name server then sends a query to the third-level nameservers (in this example, ic.gc.ca) which then returns an authoritative answer to the recursive nameservers.

Steps in a typical DNS query

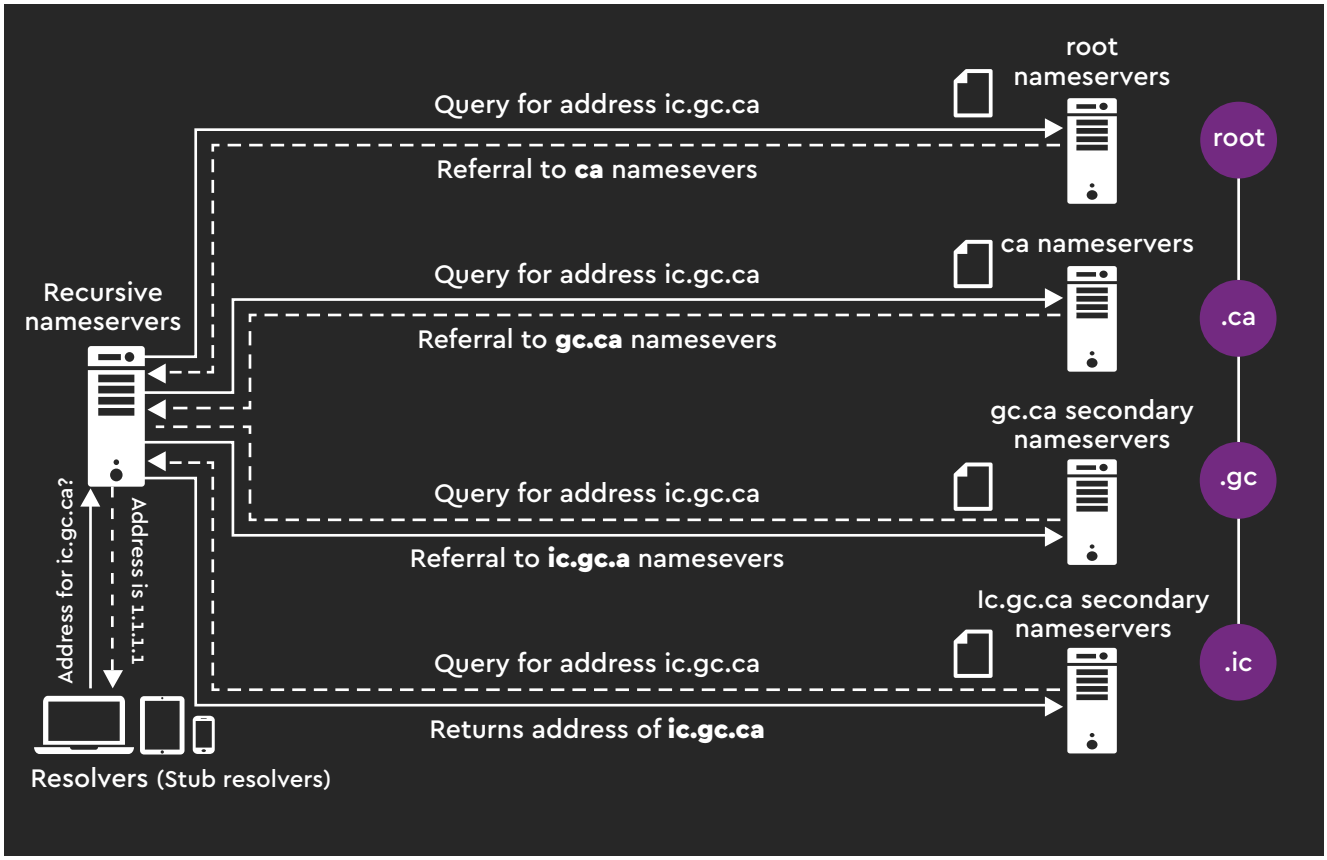

NAMESERVER ARCHITECTURE

The best practice for architecting nameservers is to use a hidden master (primary) nameserver and secondary nameservers. This has the benefit of keeping a hidden primary/master server within the corporate infrastructure for administration of the zones. With this structure in place, management of the server and/or downtime at this server will not impact access to the organization's web properties. For security, the firewall should be configured to only allow communication to and from the secondary servers.

Once this architecture is in place, the choices and scope of the secondary infrastructure is based on organizational needs and risk tolerance.

DNS NAMESERVER ADDRESSING METHODOLOGIES

Like any server infrastructure, it is best practice to have redundancy built into the DNS infrastructure. There are two methods of addressing DNS nameservers:

Unicast

Broadly applied, unicast is the communication from a single sender to a single receiver on a network. As it applies to the DNS it is a one-to-one association between a network address and the end-point. In other words, if the unicast DNS has 2 records (ns1 and ns2) then each corresponds to exactly one server. This does not preclude building in redundancy at unicast nodes or having more than one node online to answer queries, but it does not afford the full suite of benefits of an anycast solution for external DNS resolution.

Anycast

Anycast is one of several routing schemes that can be used to control the flow of traffic across a network, typically using the Border Gateway Protocol (BGP). With anycast DNS, multiple nodes that are copies of each other are geographically dispersed. In this way, multiple servers sit behind identical IP addresses and answers to queries are always done by the geographically closest node. Benefits of an anycast network include: lower latency, increased redundancy, and higher resilience to DNS DDoS attacks.

ANYCAST FOR THE SECONDARY DNS SERVICE

Much like other organizational decisions, the provisioning of DNS secondary services comes down to a, "build, buy, or both" decision. Whatever is chosen, adding an anycast secondary DNS service improves speed and resiliance for the organization's web properties and services.

If a secondary DNS service architecture is already in place, whether in house or from a service provider, this does not preclude adding additional secondary DNS services. By incrementally adding new services to their DNS an IT manager can improve speed and resilience for their web properties and can get additional benefits from geo-focused services. For instance, CIRA runs the D-Zone Anycast DNS Service, which offers both global and local secondary nodes, where the local nodes are set-up to serve Canadian traffic and to protect the Canadian traffic from external DDoS attacks. In business parlance, the D-Zone solution will help protect a company's Canadian audience and related objectives while also improving global reach through nodes located in Internet hubs around the world.

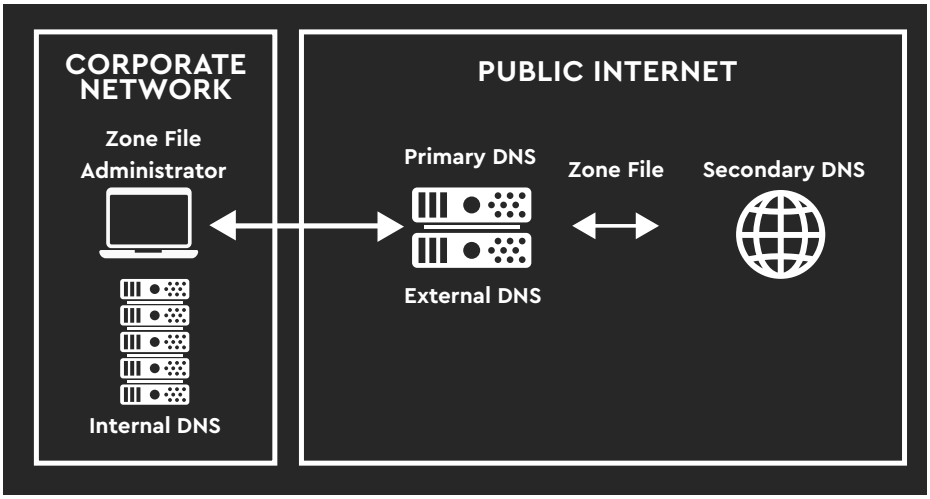

DNS architecture showing the secure corporate network, the use of a hidden primary DNS nameserver and the secondary service serving queries via an anycast cloud.

DDoS attacks on the DNS are a top operations threat with "one-third of DNS operators having experienced a customer-impacting attack"

- Arbor 2016 Worldwide Infrastructure Security Report

BENEFITS OF THE .CA D-ZONE ANYCAST DNS SERVICE

High Reliability
- ✓ 23 anycast servers are online
- ✓ Managed to provide a 100 per cent up-time service level agreement (SLA)

High Performance
- ✓ Queries are routed to the geographically closest node
- ✓ Seven redundant and load-balanced sites across Canada and close to Canadian customers
- ✓ Peered through Console Network Solutions (formerly IXReach) to additional datacenters in Europe and America

DDoS Protection
- ✓ Global footprint is 100 times over-provisioned to soak attacks
- ✓ Local and global nodes are designed to insulate Canadian traffic from DDoS attacks that occur off-shore

Easy to Implement
- ✓ Provision, monitor, and manage via web management console
- ✓ Integrate into existing systems using the REST API

Enterprise Support
- ✓ System is monitored and supported 24×7 by CIRA's networking and DNS experts
- ✓ 24×7 technical support

DNS Reporting
- ✓ Hourly, daily, and monthly reports for analyzing your DNS traffic
- ✓ Customized alerts based on changes in DNS activity inform the organization about potential problems

Secure
- ✓ Support for DNSSEC
- ✓ All zone transfers are done using TSIG

LEARN MORE

For information on ordering, finding a reseller, or becoming a reseller please contact info@d-zone.ca or visit cira.ca/d-zone.

ABOUT CIRA

The Canadian Internet Registration Authority (CIRA) manages Canada's .CA domain name registry as a 100 per cent up time service for Canadians and Canadian organisations. In addition to stewardship over .CA, CIRA develops and implements policies that support Canada's Internet community and the .CA registry internationally.

GLOSSARY OF TERMS

Anycast
Anycast refers to an addressing or routing method where queries are routed to the nearest node from several options. With respect to the DNS, an anycast network refers to answering DNS queries from the geographically closest node where all nodes share the same IP address. This has the benefit of protecting the DNS at any given node from malicious activity or failures at other nodes. Anycast nodes can be configured as global or local nodes. Local nodes provide lower latency, improve reliability and keep service local at wide-area links. Global nodes are distributed across the entire Internet.

Authoritative Nameserver
The authoritative nameserver refers to a server in the DNS that responds to questions about names in a zone. It is distinguished from a recursive DNS server in that recursive nameservers ask questions of authoritative name servers. Authoritative nameservers only provide answers about zones that are locally configured as authoritative zones. Hybrid nameservers are configured to act both authoritatively and recursively concurrently but are no longer recommended. Different "views" should be used to logically separate recursive from authoritative traffic.

BIND
BIND (Berkeley Internet Name Domain) is the most widely used nameserver software on the Internet. It originated at the University of California at Berkeley in the early 1980s as open-source software that implements the DNS protocols for the Internet. It is an enterprise-class component of the software stack with respect to query volume and stability. The software has three parts: (1) a DNS server, (2) a DNS resolver library, (3) software tools for testing.

CNAME (Canonical Name)
A resource record in the DNS that specifies a domain name is an alias for another domain name and not for an IP address. For instance, [blog.cira.ca CNAME dog.cira.ca]. It allows the running of multiple services (i.e. web server and FTP server) on different ports but sharing the same IP address.

BGP (Border Gateway Protocol)
The BGP is a system routing protocol for routing information on the Internet (versus internal networks). It is a highly scalable and robust protocol that uses route parameters to define routing policies and maintain a stable routing environment. This protocol allows gateway hosts in a network of autonomous systems to exchange routing information. The BGP has a few high-level features that make it important for the functionality of the Internet, including sending updated information only when a change is detected and a local preference attribute to reduce latency. It is the latter that can be used in an anycast network for increasing performance and security.

(DDNS) Dynamic DNS
On the public Internet, DDNS is used to refer to a method of automatically (and in real-time) updating a nameserver in the DNS using TSIG without manual editing. It is most often used to provide a mechanism to propagate the DNS for dynamic IP addresses.

DDoS (on DNS)
A Distributed Denial of Service describes a scenario where servers get overwhelmed by multiple queries

from distributed queriers, generally from bots, with the intention of overwhelming the service and making it unavailable to answer legitimate queries. Like all services on the Internet, the DNS is susceptible to attacks aimed at saturating the authoritative server's Internet connectivity with bad data. If an attack is of large enough scale, this resource exhaustion means that valid queries are never received, and thus answers can never be provided. While malicious activity targeted against the DNS are less common than other attacks on the application layer, DDoS against the DNS are still a large and growing problem.

Delegation
Delegation describes the action of delegating a server to be the authoritative name server for a domain name. This can be delegating responsibility for name resolution to a server owned by a DNS supplier or one of your own servers. For example, the .CA ccTLD delegates its subdomains, such as "companyexample.ca," to other servers.

Domain (or Domain Name)
A domain name is registered with, and delegated from the authoritative parent. For example, gc.ca is a domain name delegated to the Government of Canada (GoC) by CIRA. The GoC leases the right to the domain name gc.ca, but in the context of GoC's own DNS server administration, "gc.ca" is a zone that they must configure on their DNS servers. Further, "parl.gc.ca" and "health.gc.ca" may be part of the gc.ca zone, or they may be delegated to their subordinate organizations. If the latter were true, "parl.gc.ca" and "health.gc.ca" would also be referred to as "zones."

Domain Registration
Domains are not owned by organizations, but are registered for a period of time. Domains are registered through registrars who act like resellers for the top level domain registries such as .CA (CIRA) and .com (Verisign).

DNS Query
Since an IP address is the underlying method by which computers/devices can talk to each other, and since human readable words are the method by which web addresses are navigated, a DNS query is used to ask what human readable address corresponds to what
IP address.

DNS Resolver (or Recursive DNS Server)
Computers that respond to queries to resolve a domain name into an IP address.

DNS Spoofing
DNS spoofing is the practice of fooling an organization's recursive nameservers into accepting a bogus answer to a DNS question and having it accepted and stored in its cache. This is generally accomplished by spoofing (or forging) large numbers of illegitimate answers towards the recursive nameservers before the real authoritative servers have time to send the real answer.

DNSSEC
The Internet was designed to be open and trustworthy and the DNS protocol, as it was originally designed, met these objectives. As the Internet has grown it has remained open – but its trustworthiness is being challenged. DNS spoofing is the practice of assuming the DNS name of another system by compromising a DNS server for a valid domain. DNSSEC enables DNS records to be signed cryptographically allowing a server to validate the the response it receives is genuine. This is analogous to SSL, but for the DNS. It requires additional administration by the host organization, but helps to protect their customers from man-in-the-middle type attacks.

Forward Lookup
Forward lookup describes using a domain name to find an IP address. In practice, the URL address that someone types into their browser gets sent to the DNS to deliver the IP address. By contrast, a reverse lookup uses an IP address to find a domain name.

FQDN (Fully Qualified Domain Name)
A FQDN is a domain name that specifies its exact location in the tree hierarchy of the DNS (for example, "targetpage.example.ca"). It specifies all domain levels including the top-level and the root. Many DNS resolvers process a name without a dot by automatically appending the systems default.

GSLB (Global Server Load Balancing)
GSLBs are geographically distributed servers with authoritative nameservers running at each site where each domain is a sub-domain (ns1, ns2, ns3, etc). Load balancing is used to manage the traffic across the servers. Balancing can be simply a round-robin between the servers or via more intelligent protocols to manage traffic for reduced latency.

Global Node
A global node in the DNS is available to answer queries from anywhere on the Internet. Their existence on the Internet is advertised so that other hosts can announce them to their peers. In effect a global node is accessible from anywhere on the Internet.

IPv4
Internet Protocol V4 is the system that routes most traffic on the Internet. It is a connectionless protocol for packet switched networks based on best-effort delivery. It uses 32 bit addresses typically expressed in a human-friendly dotted decimal notation (i.e. 255.255.255.1). IPv4 numbers were "exhausted" in 2011 with the large numbers of devices in the world with IPv6 set-up as the next generation protocol, although at present most global traffic remains IPv4.

IPv6
IPv6 is the latest version of the Internet Protocol that provides the location system for devices on the Internet. It uses a 128-bit address to allow a virtually limitless number of addresses (2^128) when compared to what is currently in use with IPv4. IPv6 also has technical advantages that limit the expansion of routing tables, enable multicast addressing, and assist with security. IPv6 addresses are eight groups of four hexadecimal digits (2001:0db8:85a3:0000:0000:8a 2e:0370:7334).

Iterative Query
The DNS server will not get the complete answer to a query, but gives back a referral to the server that may have the complete answer. It will not query the root server on behalf of the original query. In this way the original requester, or DNS client, will be responsible for making a query to the next DNS server until it locates a DNS server that is authoritative, or until it times out.

Latency (of the DNS)
Latency describes the time between the end-user requesting DNS resolution and the response from the server. Although DNS latency it is not reflected in organizations' web server logs, it adds to the total load time of your website to the end user.

Load Balancing
Load balancing describes how traffic is managed between multiple servers located in a server cluster or node. Similar to any redundant infrastructure it can be managed with a simple round-robin approach or an intelligent approach to sending traffic to the least busy server in the node.

Local Node
A local node in the DNS is announced with the no-export BGP so that hosts do not announce them to their peers. They are typically located in local IXPs that are peered within the local community (such as a specific country). This makes the local nodes primarily accessible to local queries and mitigates the risk of malicious activity that is not peered to the local node.

Nameserver
A nameserver is a server on the Internet that answers DNS queries. A nameserver may be authoritative (providing answers) or recursive (asked questions on behalf of a third-party). A domain may be delegated to authoritative DNS servers that are subordinates to that domain (ie: NS01.CIRA.CA is a DNS server for CIRA.CA), or to nameservers outside of that domain (NS1.D-ZONE.CA is a DNS server for CIRA.CA).

Node
A DNS Node is a server or cluster of servers that answers DNS queries.

Primary DNS (server)
The Primary DNS is used to describe the server that is the primary source of valid zone files for the DNS records. This can be deployed to answer queries on the Internet or it can be kept as a hidden master for zone administration and communication only to secondary servers that are set-up to answer queries.

Recursive Query
The DNS server that receives your query will do the job fetching the answer and, if needed, query other DNS servers on the Internet to get the answer. This is done if it does not have the answer to the DNS request in a zone file or in its cache.

Recursive Server (also called recursive resolver)
Recursive server receives user-generated queries, checks its cache, and if not present, spawns its own set of queries to the respective authorities for each level (DNS Root, TLD, Second-level, etc.) and provides an answer back to the initial querier.

Registry
A registry is the organization responsible for the management of the top-level domain name such as .CA or .com. The registry is mandated by the Internet Assigned Numbers Authority (IANA) as a department of ICANN, to manage the domain name based on a set of guidelines set out in their mandate. This includes generic top level domains (gTLD), such as .net, .org, .com, the new gTLDs, and a domain for every country code (ccTLD) such as .CA for Canada, and .uk for the United Kingdom. These domains are under the responsibility of the registry to manage.

Secondary DNS (server)
Secondary DNS is a term most often applied to a backup to the primary DNS server and describes a redundant name service on a separate network to prevent downtime. In an anycast cloud it is used to describe the servers that are set-up to answer DNS queries and it receives zone files from the primary server.

Reverse Lookup
A reverse DNS lookup is the system of looking up a domain name when you have the IP address.

Root Servers (or Root)
The DNS is organized as a hierarchy and at the top of it is the root domain. The root domain contains all the top-level domain (TLD) names such as .CA and .com and can be envisioned as an empty string that occurs after the TLD. In the DNS, the authoritative nameservers that serve the root zone are called root servers. They are a network of servers throughout the world. Recursive resolvers need to configure a root hints file that contains the names and IP address of root servers so they can bootstrap DNS resolution.

RTTM (Real Time Traffic Management)
RTTM describes the managing of DNS traffic globally that routes traffic either to a geographically close node or a global node based on what servers deliver the lowest latency and fastest speed. In this scenario, servers are constantly being monitored rather than simply relying on geography or round-robin type techniques for traffic management.

SoA Record (Start of Authority)
The NS (Name Server) resource record identifies the name servers, not the SoA. The SoA contains the "source host" (generally, RFC-specified, but not always true), but that could be ANY identifier string, which is often but not limited to the name of the "hidden master" used to create the zone.

TLD (Top Level Domain)
The root is the highest level in the domain hierarchy and the root zone contains the delegations of all of the world's TLDs. TLDs are delegated to specific countries and organizations by ICANN. Within this grouping there are broadly two types of TLDs. Country Code TLDs (ccTLD) are assigned to specific countries such as .CA, .UK, and .SE. Generic TLDs (gTLD) are not country specific and include some of the larger TLDs by volume such as .com and .net, in addition to a large number of new gTLDs.

TSIG (Transaction Signature)
TSIG is the mechanism for sending zone updates securely between nodes/servers. It is the networking protocol used by the DNS to ensure that the information from a certain server is actually from that server by using a form of key-based infrastructure defined in RFC 2845. Because the DNS works in a question-answer model, TSIG is essential to ensure that the answer is sent based on more than just the IP address it originated from.

TTL (Time to Live)
In order to facilitate updating the DNS servers across the Internet the zone file in the authoritative name server specifies a TTL, which is to say how long a given recursive server should keep the DNS information in cache. In this way servers don't need to reach back to the authoritative server every time they need to respond to a query. It is not advantageous to set an arbitrarily low TTL because it results in recursive servers needing to ask questions about a given domain more often, which could be perceived as latency by the querier. It is also not advantageous to set an arbitrarily high value, because it can reduce the domain owner's ability to work around network issues or changes, as the "old data" may be cached on any number of recursive servers across the world until this counter ceases.

Unicast
Broadly applied, unicast is the communication from a single sender to a single receiver on a network. As it applies to the DNS, it is a one-to-one association between a network address and the endpoint. In other words, if the unicast DNS has two records (ns1 and ns2) then each of them correspond to exactly one server. This does not preclude building redundancy at unicast nodes or having more than one node online to answer queries, but it does not afford the full suite of benefits of an anycast solution for external DNS resolution.

Zone
A DNS zone references the domain name space where administration is delegated to a single manager. It is implemented in the domain nameserver. It is organized into zones to allow delegation of responsibility over sub-domains to the relevant authorities. Top-level domains, such as .CA, manage a zone in the DNS where the sub domains live. In this way a zone always has a domain boundary in which it operates.