

White Paper

DNS Anycast for High Availability and Performance

by Timothy Rooney

Product management director

BT Diamond IP

DNS Anycast for High Availability and Performance

By Tim Rooney, Director, Product Management

Introduction

We've all heard the mantra, "do more with less," – for some of us to the point of numbness. But when it comes to mission-critical DNS name resolution services, any suggestion of skimping is likely to evoke defiance from those responsible for maintaining its availability and performance. One strategy, however, that can make the deployment of DNS services more efficient is the use of anycast addressing.

This white paper shows how DNS servers can be deployed with anycast to achieve a number of benefits, including higher availability and resilience.

Defining Anycast

The term *anycast* refers to an IP address assigned to a set of interfaces (usually belonging to different nodes), any one of which can be the intended recipient. Anycast addresses are assigned from the same address space from which unicast addresses have been allocated. Thus, unlike private address space, one cannot visually differentiate a unicast address from an anycast address.

An IP packet destined for an anycast address is routed to the nearest interface (according to routing table metrics) configured with the anycast address. The concept is that the sender doesn't necessarily care which particular host, or in this case DNS server, receives the packet, just as long as one of those sharing the anycast address receives it.

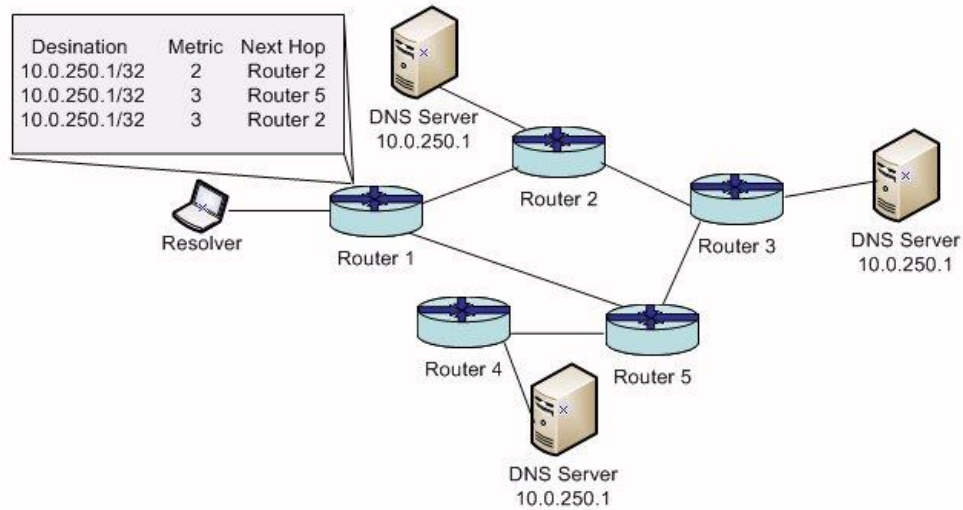
From the perspective of the DNS resolver, the operating system service that performs DNS lookups for the client, it can issue a query addressed to the anycast address, and the network will route the request to the nearest available DNS server configured with the anycast address. This enables the routing infrastructure to route the query to the DNS server closest to the resolver client (according to routing metrics) wherever the resolver happens to be. If the DNS server also participates in dynamic routing with the routing infrastructure, the routing infrastructure can also detect reachability and automatically route queries elsewhere if a given DNS server is down or unreachable. Thus, use of anycast addresses with dynamic routing on DNS servers can provide higher resolution performance as well as DNS server redundancy.

Deploying DNS Servers with Anycast Addresses

Configuring DNS servers with anycast addresses enables them to utilize a common IP address. Anycast allows a resolver to reach any one of the anycast addressable hosts without regard to which host is reached. The routing infrastructure handles routing metric updates to track

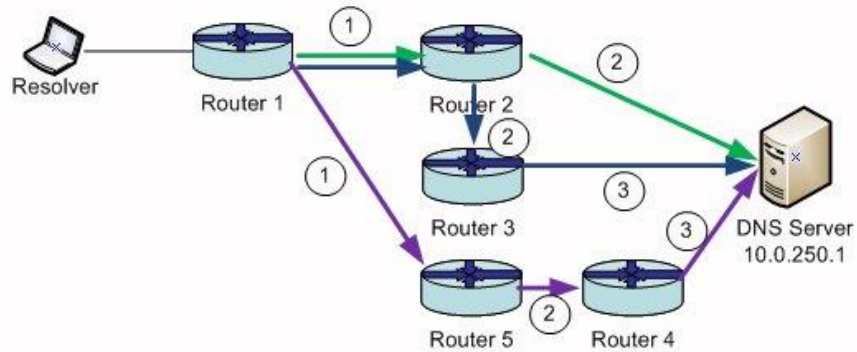
reachability and routing to the nearest host configured with the destination anycast address. Figure 1 illustrates three DNS servers configured with anycast address 10.0.250.1.

Figure 1: Anycast Routing Table Example



As depicted in Figure 1, Router 1 has three routes to anycast address 10.0.250.1/32, corresponding to the three servers. The closest server is the one homed on Router 2 and is two hops from Router 1. The next closest servers are both three hops away, routing to Router 3 via Router 2 (or Router 5) or to Router 4 via Router 5. The logical view from Router 1's perspective is illustrated in Figure 2, where the anycast IP address is considered a single destination, reachable via multiple paths.

Figure 2: Logical Routing Perspective from Router 1 Showing Hop Counts



The routers can be configured with static routes to the anycast address 10.0.250.1 in this example. The configuration statically defines the route from one router to the next in order to reach a given DNS server configured with the anycast address. Referring back to Figure 1, a static route would be defined in Router 1 to route all packets destined for IP address 10.0.250.1/32 to Router 2. As such, Router 1 will always direct DNS queries destined for the anycast address to Router 2.

Static routing enables configurable routing of DNS queries across the routing infrastructure according to the router administrator's choosing, which enables a resolver client to include the anycast address among its configured DNS server addresses. However, if the DNS server homed on Router 2 fails, static routing provides no automated alternate routing. Implementation of dynamic routing on the DNS server enables advertisement of reachability to the anycast address. The DNS server homed on Router 2 can provide reachability updates as part of a standard routing protocol such as Open Shortest Path First (OSPF). With dynamic routing, if the DNS server homed on Router 2 fails and stops transmitting routing updates to Router 2, Router 2 will update its routing table to reach 10.0.250.1/32 via Router 3 instead. This approach adds high availability through dynamic routing to the high performance inherent in shortest path routing.

Anycast Benefits

Deploying anycast provides a number of benefits:

- Simplified resolver configuration
- Improved resolution performance
- High availability DNS services
- Resilience from DNS denial-of-service attacks

As we've seen, resolvers configured with the DNS servers' anycast address would have their queries routed to the nearest DNS server configured with that anycast address. Thus, regardless of where the resolver host connects to the network, the same anycast IP address may be used by the resolver to locate a DNS server. This homogenized resolver configuration also helps provide improved performance of the resolution process. A query to a DNS anycast address should be routed to the closest DNS server, thereby reducing the round trip delay portion of the overall query process.

The outage of a DNS server can be communicated (by absence of communication) to the routing infrastructure in order to update routing tables accordingly. This requires the DNS server to run a routing daemon using the routing protocol of choice to communicate reachability to the local router. Participation in routing protocol updates enables the local router to update its routing table with an appropriate metric, and to pass this on to other routers via the routing protocol. Depending on the deployment (internal or external) of the DNS server, a corresponding interior or exterior routing protocol would need to be running on the DNS server. The server simply needs to communicate that its anycast address is reachable. It would be more useful if this routing update was linked to the status of the DNS daemon or service on the server, though application status is not generally considered when communicating IP address reachability.

Deploying anycast also provides mitigation against denial-of-service attacks as evidenced by the distributed-denial-of-service (DDoS) attack on multiple root servers on February 6, 2007¹. Of the six root servers targeted, the two most severely affected had not yet implemented anycast. The other four root servers, having deployed anycast, enabled the spreading of the attack across more physical servers. Thus a DDoS attack on the I-root server, which did not have anycast configured, severely impacted the ability of the server to respond to legitimate queries. Meanwhile, the attack on the F-root server, which was deployed with over 40 DNS servers sharing the F-root anycast address, was minimized by distributing the impact of the attack across these servers. This form of load sharing enabled the F-root server(s) to continue processing legitimate queries while suffering a barrage of artificial requests.

Anycast Caveats

While deploying anycast provides many benefits, there are constraints and caveats to be considered. Because resolvers may query any DNS server configured with the anycast address at a given time, it's important that the resolution information configured on the server be consistent. For example, the implementation on Internet root servers consists of a set of master servers with static information. These root servers do not accept dynamic updates. If anycast is desired for dynamic zones, then each server must have a unicast address in addition to its anycast address². This enables updates to be directed to the master's unicast address, which may in turn Notify its slaves via their respective unicast addresses. A hidden master configuration can be used with the slaves configured with anycast addresses; that is, only the slaves would be configured with the anycast address, not the master. The master is not to be queried by resolvers in this case.

Another consideration is the requirement to run a routing daemon on DNS servers configured with anycast addresses. While routing of packets to anycast addresses is primarily a routing function, the unreachability of a DNS server host may result in lost query attempts. Such would be the case if static routes are used to configure routers with fixed metrics for the DNS servers

¹ March 1, 2007, ICANN Factsheet: Root Server Attack on 6 February 2007, <http://www.icann.org/announcements/factsheet-dns-attack-08mar07.pdf>

² Every anycast server will require a unicast address for administration, but to support dynamic zones, an additional unicast address is required to provide an interface for updates, notify's and zone transfers.

configured with a common anycast address. Should a server become unavailable, the serving router has no way to detect this and would not re-route packets destined for the anycast address. Therefore, incorporating a routing daemon on the DNS server improves overall robustness. Should a server fail, the local router will determine that it is no longer reachable and will update its routing table and those of other routers via routing protocol updates. Internet root serves support BGP, given their deployment on the global Internet, though deployment within organizations will likely require support for OSPF, IGRP, or the interior routing protocol of choice.

IPControl™ Sapphire Anycast Support

IPControl Sapphire x-series DNS servers support anycast addressing for high availability and performance. Sapphire also incorporates OSPF routing services for dynamic routing. In addition, reachability is not based solely on the server availability, but also on the DNS service availability. If the server is up but the DNS service is not running, Sapphire will not announce reachability. This level of granularity communicates DNS service reachability, not simply the appliance itself.

IPControl Sapphire appliances are configurable using the IPControl centralized management system or via local or network reachable (SSH) console access. Using IPControl, anycast and the many other configurable elements and parameters of DNS can be configured centrally, using a holistic IP address management perspective. This approach eliminates the need to manually translate the addressing and domain plans from a spreadsheet or IP database into requisite components of a given DNS server configuration file. The intuitive web interface automates this process and thereby saves time and reduces entry errors.

Conclusion

Anycast addressing is an innovative technique for providing high performance and high availability DNS services for your constituents. Deployment of IPControl Sapphire appliances enables you to reap the benefits of DNS anycast addressing, while enjoying the many additional benefits of configuring DNS services within the scope of your overall IP address and domain plan.

IPControl Sapphire provides many market-leading differentiators, including the following key features:

- Highly scalable IP address management
- Four network interfaces, port ACLs, port bonding, and lights-out management³ using the industry-standard Integrated Platform Management Interface (IPMI)
- Centralized appliance dashboard enables monitoring, controlling and upgrading of deployed Sapphire appliances
- First integrated IPv4 and IPv6 address management system
- First all-appliance IPAM and DHCP/DNS solution

³ Sapphire 20 provides IPMI as standard and Sapphire 10 provides IPMI as an optional feature.

- Supports multivendor DHCP/DNS services configuration for native Microsoft, Internet Systems Consortium (ISC/BIND), and BT Diamond IP Sapphire.
- Backed by the global presence and support of BT, a worldwide leader in technology and services

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 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 btiamondip-sales@bt.com or consult www.diamondipam.com.

IPControl and ImageControl are trademarks of BT Americas, Inc.

Copyright © 2017, BT Americas, Inc.

This is an unpublished work protected under the copyright laws.
All trademarks and registered trademarks are properties of their respective holders.
All rights reserved.